Freie Universität Berlin

Department of Mathematics and Computer Science
Institute of Computer Science

Doctoral Dissertation

Measuring and Implementing Internet Backbone Security: Current Challenges, Upcoming Deployment, and Future Trends

Dissertation zur Erlangung des akademischen Grades eines
Doktors der Naturwissenschaften (Dr. rer. nat.) am
Fachbereich Mathematik und Informatik der Freien Universität Berlin

vorgelegt von

Dipl.-Inform. Matthias Wählisch
m.waehlisch@fu-berlin.de

Berlin 2015

Tag der Disputation: 20. Januar 2016

Examiners

Prof. Dr.-Ing. Jochen Schiller	**Prof. Dr.-Ing. Tanja Zseby**
Freie Universität Berlin	***Technische Universität Wien***

Abstract

In this thesis, we start from the observation that the Internet is a critical infrastructure, which needs severe protection. We take a practical view on Internet security, considering the whole ecosystem including the network, end devices, and services, protected and threatened by current and future Internet protocols. We contribute tools, methodologies, and measurement results to improve the current state of art as well as operational practice.

Acknowledgments

Successful research is linked to passion that is shared with the community. It is not based on working in the ivory tower but on open-minded discussions with friends and colleagues who share thoughts and *seriously* argue with you. Before and during my PhD thesis I was lucky to meet several of those people.

During my high school time I found my way to the *RZ*. This was a unique experience over several years. I am very grateful to worked with *rzstaff* in Berlin Karlshorst, and that they just let me play early with real networks.

The next big step was not only to start studying at FU Berlin but get involved in IETF work. I would like to thank Rajeev Koodli. He was our first research group chair, while we worked on mobile multicast. He made very clear to me that exploring the problem space precisely, is the (significant) step to identify the solution.

At FU Berlin, Georg Wittenburg was kindly (or brave?) enough to accept my Diploma thesis proposal on overlay multicast and thus gave me an entry to the CST group. It was also fun to share an office with him during the first phase of my PhD.

I thank Jochen Schiller for showing me that an illustrative presentation of problems and solutions is important. More importantly, I am very grateful for his trust, which allowed me to autonomously follow my project ideas. I am very much looking forward to continuing the collaboration.

I am also grateful to Tanja Zseby for taking the role of the second examiner, and for her incredibly detailed feedback on a first version of this thesis.

Many thanks to Alex Band, Tim Bruijnzeels, and Carlos Martinez to give feedback on our RPKI work from an RIR perspective, as well as Thorleif Wiik, Thomas King, and Arnold Nipper for discussions on the intricate topic of global peering from an IXP point of view.

Thanks to the SIDR working group for sharing insights into the RPKI; in particular, Rob Austein, Randy Bush, Sandy Murphy, Ruediger Volk, and Andy Newton.

Lixia Zhang, Jennifer Rexford, Steve Uhlig, Markus de Brün, and Thomas Häberlen are gratefully acknowledged for early and open discussions on the topic of nation-state routing; Christian Keil, Jochen Schönfelder, and André Vorbach for discussions on honeypots; and Sebastian Trapp for starting the work on interpersonal trust in ad hoc networks.

Special thanks to Eric Osterweil for controversial arguments on ICN, RPKI, and the overall system. I am looking forward to a joint paper, sooner or later.

Constrained (sic!) thanks to the Franco-German IoT gang, in particular Emmanuel Baccelli, Oleg Hahm, Hauke Petersen. Nordic thanks go to the INET research team in Hamburg.

I would like to thank Steve Uhlig for playing(?) the devil's advocate, Olaf Maennel and Johann Schlamp for discussions on routing security, and Gareth Tyson for bringing bizarre topics to the research community.

It was a pleasant experience to work with motivated BSc and MSc students. Particular thanks to Samir Al-Sheikh, Fabian Holler, Marcin Nawrocki, Andreas Reuter, Fabrice Ryba, Robert Schmidt, and Markus Vahlenkamp.

I am thankful to the organizers of the Dagstuhl Seminars *12361: Information-centric networking – Ready for the real world?*, *14291: Information-Centric Networking 3*, and *15102: Secure Routing for Future Communication Networks* for inviting me. Special thanks to Walter Willinger, Steve Uhlig, Georg Carle, Jochen Schiller, and Thomas Schmidt for the joint organization of *13322: The Critical Internet Infrastructure*. Dagstuhl is a great chance to discuss fresh ideas in a relaxed atmosphere.

Last but not least, I would like to express my gratitude to Thomas Schmidt for continuous discussions on Internet communication. Beyond this thesis, I would like to thank him for ten + seven years of joyful collaboration. Thanks, for being inspiring. And sorry that it took so long ;).

This work has been supported by the German Bundesministerium für Bildung und Forschung within the projects H∀Mcast, SKIMS, Peeroskop, and SAFEST. This funding offered the opportunity to work on the different subjects without a thought on my subsistence and travel budget—acknowledgments to the ministry for doing the right thing, again.

Bibliographical Notes

The author of this dissertation was the principle investigator and first author of the following research papers, which are the basis for this thesis:

Chapter 2 is based on

[144] 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 the AS Level," in *Proc. of the 13th Passive and Active Measurement Conference (PAM)*, ser. Lecture Notes in Computer Science, N. Taft and F. Ricciato, Eds., vol. 7192. Heidelberg: Springer, 2012, pp. 200–210.

Chapter 3 is based on

[141] M. Wählisch, F. Holler, T. C. Schmidt, and J. H. Schiller, "RTRlib: An Open-Source Library in C for RPKI-based Prefix Origin Validation," in *Proc. of USENIX Security Workshop CSET'13*. Berkeley, CA, USA: USENIX Assoc., 2013.

Chapter 4 is based on

[143] M. Wählisch, R. Schmidt, T. C. Schmidt, O. Maennel, S. Uhlig, and G. Tyson, "RiPKI: The Tragic Story of RPKI Deployment in the Web Ecosystem," in *Proc. of Fourteenth ACM Workshop on Hot Topics in Networks (HotNets)*. New York: ACM, 2015. First version arXiv:1408.0391, August 2014.

Chapter 5 is based on

[149] M. Wählisch, A. Vorbach, C. Keil, J. Schönfelder, T. C. Schmidt, and J. H. Schiller, "Design, Implementation, and Operation of a Mobile Honeypot," Open Archive: arXiv.org, Technical Report arXiv:1301.7257, Jan 2013. [Online]. Available: `http://arxiv.org/abs/1301.7257`

Chapter 6 is based on

[147] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, "Backscatter from the Data Plane – Threats to Stability and Security in Information-Centric Network Infrastructure," *Computer Networks*, vol. 57, no. 16, pp. 3192–3206, Nov. 2013. [Online]. Available: `http://dx.doi.org/10.1016/j.comnet.2013.07.009`

Contents

Acronyms

AFRINIC African Network Information Centre.

APNIC Asia-Pacific Network Information Centre.

ARIN American Registry for Internet Numbers.

AS Autonomous System.

ASN Autonomous System Number.

BGP Border Gateway Protocol.

CA Certificate Authority.

CCN Content Centric Networking.

CDN Content Delivery Network.

CNAME Canonical Name.

DDoS Distributed Denial-of-Service.

DNS Domain Name System.

DoS Denial-of-Service.

FIB Forwarding Information Base.

IANA Internet Assigned Numbers Authority.

ICN Information-Centric Networking.

IETF Internet Engineering Task Force.

IRTF Internet Research Task Force.

ISP Internet Service Provider.

IXP Internet Exchange Point.

LACNIC Latin America and Caribbean Network Information Centre.

LIR Local Internet Registry.

MOA Multiple Origin AS.

NDN Named-Data Networking.

ORDNS Open Recursive DNS.

PIT Pending Interest Table.

RIB Routing Information Base.

RIPE Réseaux IP Européens.

RIR Regional Internet Registry.

ROA Route Origination Authorization.

RPF Reverse Path Forwarding.

RPKI Resource Public Key Infrastructure.

RPKI/RTR RPKI to Router Protocol.

RTT Round-Trip Time.

SIDR Secure-Inter Domain Routing.

Chapter 1

Introduction

The Internet is a globally distributed interconnection of networks, also called autonomous systems—and it is one of the most complex and challenging networks that was built. The fundamental service that each network provides is transit, either for other networks or (ultimately) for end devices. It is the purpose of the Internet to enable content exchange between end systems attached to different networks using different services. This ecosystem of networks, end devices, and services, changes continuously.

Openness is a key design decision for the success of the Internet, as it supports this continuous change. The openness is implemented in several dimensions, technically as well as economically. In principle, any autonomous system (AS) is open to establish peering relationships with any other AS or buy upstream connectivity as long as there is a mutual agreement between these two parties. Any operator is open to deploy any technology within the AS as long as there is a common protocol available between ASes. Such an open and decentralized system challenges reliability and security.

The Need for Security

Services on top of the Internet have taken a critical role in our daily life. To illustrate this, we just need to imagine that the Internet, and thus the interconnection of these services, does not exist. Online business would not be possible, and people could not coordinate via social media, for example. The demand for working interconnects will get even more severe in the future with an increasing deployment of the Internet of Things. Maliciously broken connections have been demonstrated several times, either by attacking the backbone or the edge. Security mechanism should help to improve this situation.

Approaches to improve security and reliability are available on several layers. However, protection at upper layers has only limited effect if the lower layers remain vulnerable. For example, traffic redirection can break transport layer security. End devices with a local firewall can still be attacked with a denial of service. Many approaches have been proposed but only a few evolved to practically applicable solutions. Nevertheless, even well designed standards lack deployment. In this thesis, we argue for a better understanding of pitfalls that are made when deploying new security mechanisms, as well as for analysing the implications of architectural changes.

The Need for Nation-Centric View

At the beginning, the Internet structure followed a hierarchical model, where a limited number of large Internet Service Providers (ISP) were responsible for interconnecting major parts of the Internet, larger regional and local access providers connected smaller customers. With the increasing demands of web applications and the decreasing costs for interconnection, dedicated content delivery networks (CDNs) emerged. Those CDNs (or *Hyper Giants*) create their own transnational backbones and foster local peering [91]. This trend towards local interconnections is widely visible [11], and makes control even more challenging. In practice, the openness of the Internet is limited by governmental regulations in some countries, such as in China where local peering relations are predetermined.

Giving the Internet of today, a nation-centric perspective—without limiting the openness—may help to improve the robustness. Critical infrastructures of a country should operate independently of other countries, for example. However, more and more infrastructures rely on the Internet, without a clear understanding how traffic flows, and thus depend on external support. Identifying possible vulnerabilities caused by international connections is thus crucial.

In this thesis, we (i) discuss nation-centric aspects on the Internet, (ii) analyze current mechanisms to improve the security of the Internet backbone, (iii) and study the limits of a future, information-centric Internet architecture.

1.1 A Primer on Current and Future Internet Communication

Internet communication is threatened in two ways, either by attacking the delivery infrastructure or by attacking directly the end device.

1.1.1 Core: The Border Gateway Protocol

Since more than twenty years the Border Gateway Protocol (BGP) [119] is used to exchange network layer reachability information between autonomous systems. Any BGP speaker informs its neighboring ASes about its own IP prefixes or paths towards prefixes of other networks. With respect to security, the major problem of BGP is lack of mechanisms to verify the exchanged information. Consequently, an autonomous system may incorrectly claim to own an IP prefix or change path information. This may lead to traffic redirection, resulting in unreachability of addresses or traffic interception. It is worth noting that such incidents happen often. They may be due to misconfiguration or malicious behavior.

Securing BGP has been discussed since more than one decade in the research community [31]. Current efforts of the Secure-Inter Domain (SIDR) working group within the IETF lie in the standardization of a set of protocols to enhance the security of BGP. They focus on solving two problems: enable routers (a) to verify that a BGP update did originate at an authorized AS and (b) to verify that the AS path within the BGP update corresponds to the route traversed. Even though the latter is far from global deployment, first steps have been performed to establish route origin validation.

A fundamental part for securing BGP is the Resource Public Key Infrastructure (RPKI) [77], which consists of a distributed public key infrastructure responsible for Internet resources, i.e., AS numbers and IP prefixes. An RPKI repository stores certificates and Route Origin Authorization (ROAs) objects. A ROA provides a secure binding between one or multiple IP prefixes and an AS that is allowed to originate that prefix.

Using ROA data, an RPKI-enabled router is able to verify the BGP updates it receives. The prefix information within the BGP update might be *valid* (i.e., the origin AS is allowed to announce this prefix), *invalid* (i.e., the origin AS is incorrect or the announced prefix is too specific), or *not found* (i.e., the announced prefix is not covered by the RPKI). Rejecting an invalid route helps to successfully suppress an incorrectly announced prefix, which finally secures network layer reachability of services assigned with an IP address of this prefix.

As soon as correct routing information has been established between autonomous systems, end devices should be able to reach each other.

1.1.2 Edge: Remote Attacks

The original Internet design follows the end-to-end principle [126]. This means application endpoints are located at end hosts instead of intermediary nodes. More importantly, the network itself is not aware of services running on specific hosts. The network only delivers data based on end point addresses. Even though this simplifies the core network design, it makes end devices more vulnerable as it allows for unsolicited traffic.

A serious problem for end devices in the current Internet are remote attacks. In this scenario, an attacker (a) either tries to take over a local system by exploiting software that offers external connections; or (b) attempts to make a host unavailable by issuing unsolicited traffic. Two observations are worth noting. First, those attacks do not only harm the local system but may also lead to larger network outages. A prominent example are amplification attacks, which may disable both, an end device as well as points of presence of network operators due to unexpected, large (amplified) traffic. Second, even a small amount of packets arriving at the victim may pose a threat. Mobile devices have limited resources in terms of processing power and energy. Handling of malicious packets can drain those resources.

To overcome the problem of remote attacks, traffic filters are usually deployed. Ideally, these filters are applied near the source of the malicious traffic, i.e., within the network to protect end devices. However, in the current Internet this is a challenging task. On the one hand, as the network core is not aware of edge services, network operators would need to negotiate filter policies with end customers, which is not scalable. On the other hand, deploying a priori strict filter rules conflicts with the openness of the Internet.

Another, more drastic approach would abandon end point identifiers and thus inherently prevent unsolicited traffic towards the edge.

1.1.3 From Edge to Core: Information-Centric Networking

To improve both content delivery and security, Information-Centric Networks (ICN) [14] have been designed. The basic idea is to implement content awareness within the network. They are inspired by content distribution networks (CDNs), which facilitate an efficient, wide-area replication of static data for selected content providers. ICN gives up the end-to-end design of TCP/IP, to allow for asynchronous, global in-network caching of popular content. Communication to an end device then follows *only* content requests initiated by this device. Consequently, current denial of service attacks towards a dedicated host are impossible by design—as long as the consumer did not subscribe to the content, the end host will not receive any (pushed) data.

Named-Data Networking (NDN) is a prominent representative of ICN and implemented in Content-Centric Networking (CCN). NDN performs content retrieval by routing on names. A content consumer request content by sending Interest messages, which are distributed within the network towards the content source. Data itself is then forwarded along reverse paths (RPF), either by using IP as a lower layer, or without IP but by dedicated RPF states. As soon as the Interest pass a content store that can satisfy the request, data is delivered along the RPF path.

1.2 Challenges

1.2.1 Exposing a Nation-Centric View on the Internet

The Internet is a critical infrastructure for almost all countries. However, in contrast to common infrastructures such as power grids or rail transportation, the Internet is (a) a decentralized system, (b) even parts cannot obviously be assigned to countries, (c) it is a multi-service infrastructure.

For example, the Deutsche Telekom is clearly a German Internet Service Provider in the sense that this company is a spin-off of the German Post and that the German government still holds ≈15% of the company stocks directly. Deutsche Telekom provides Internet connectivity for German business customers (e.g., email providers, video distribution, web hosting) and is responsible for nearly 42% of the broadband access in Germany [2]. That these customers can connect to *www.ard.de*, the first public-service broadcaster in Germany, without involving international operators is not obvious. The content of *www.ard.de* is hosted within a content distribution network (CDN), which is operated by Akamai, a US-based company. On the other hand, an Internet company that has its headquarter in Germany may have major peering hubs outside of Germany, thus relying on additional infrastructure in foreign countries and might even be regulated under foreign law.

From this perspective, it is challenging to define which parts of the Internet are essential for a specific country. The reasons for this lie in the decentralisation and openness of the Internet. There is neither global (or national) knowledge about the geographical deployment of ISP infrastructure, nor is there knowledge about which services run on top of these networks. This implies several research questions: (a) How can we define a nation-centric view on the Internet? (b) How are different business sectors of a country interconnected? We look into these questions in more detail in Chapter 2.

Even after we have a clear understanding of a nation-specific notion of the Internet infrastructure, Internet connectivity would be still surprisingly fragile.

1.2.2 Protecting the Current Internet Backbone

When studying new protocols the main challenges are related to tools, data collection, and data interpretation. Realistic measurements and comprehensive monitoring in practice require a full-fledged implementation to analyze the feasibility and the deployment of the protocol. Such an implementation is rarely available at the beginning of the protocol evolvement. In the context of RPKI, router implementations are either not tailored to real-world deployment; or contributed by router vendors. System-related research questions are then restricted to closed platforms with a limited set of capabilities to measure system performance in detail.

Regarding data analysis, one of the most interesting questions is the identification of malicious incidents. The RPKI clearly distinguishes between valid and invalid BGP updates. An invalid update, though, is not necessarily an attack. Incorrectly configured ROAs and BGP updates respectively may lead to an invalid outcome as well.

Deploying backbone protection based on RPKI leads thus to the following research question: (a) How can we efficiently implement origin validation on routers? (b) What are the detailed reasons for invalid prefix announcements? (c) Does the RPKI introduce to new local threats on RPKI-enabled routers? (d) How does RPKI deployment relate to the availability of higher layer Internet services, such as the Web? We present tools, methodologies, and results to shed light on these questions in Chapter 3 and Chapter 4.

1.2.3 Disclosing Internet Attacks on Mobile Devices

The number of mobile end devices significantly increased over the last years. Today, smartphones and tablets are also connected with the Internet. In more than ten countries, such as the US, Internet access of mobile-only users actually exceeds Internet access via desktops [51]. In the near future, the Internet of Things will increase the number of wireless and wired devices even further. To effectively design device-specific countermeasures (e.g., network filters), a clear understanding not only of the attack strategy but also of the attack development is necessary.

Honeypots are a common tool to measure remote attacks. This type of software serves as a trap to better understand unsolicited connections.

For a scalable measurement environment we should simplify the deployment and maintenance of honeypots. Deploying honeypot software on any different mobile platform would be complex and time-consuming as mobile-specific apps need to be implemented, which is different to desktop environments. Abstraction may help to consolidate target platforms.

Crucial questions to better understand attacks and countermeasures for mobiles are: (a) Do attack signatures differ between network types? (b) What is the design space for a mobile honeypot to capture attacks? We discuss these questions in Chapter 5.

1.2.4 Identifying Potentials and Limits of ICN to Protect a Future Internet

The role of the backbone infrastructure changed in information-centric networks. The backbone does not only deliver data, it operates with content-awareness. To implement in-network content distribution within ICN, the network infrastructure needs to maintain content-specific states. As content is provided by end users and volatile, these states are data-driven states controlled by end users. Virtually, control plane and data plane are merged. This finally leads to an opening of the control plane (backbone) to continuous modification by the data plane (end user).

The ICN concept is in very strong contrast to the current Internet—BGP states are not controlled by end hosts. The conceptual change introduces several research questions with respect to the vulnerability of the ICN infrastructure: (a) Which (inherent) threats exist for an ICN-based Internet backbone? (b) Which drawbacks arise from merging control plane and data plane? (c) How does ICN state maintenance affect router performance? We analyse these aspects in Chapter 6.

Figure 1.1 illustrates the threat landscape for the current and future Internet.

1.3 Contributions and Outline

A nation-centric view on the Internet backbone. In Chapter 2, we report on a methodology and tool chain for identifying and classifying a 'national Internet', and evaluate detailed results for the example of Germany. Our contribution (i) identifies the ASes that are important for the country, (ii) classifies these autonomous systems (ASes) into functional sectors, (iii) constructs the AS routing graph of a country as well as subgraphs of specific sectors, and (iv) analyzes structural dependencies between key players. Our methods indicate the importance of examining individual IP-blocks held by individual organizations, as this reveals 25% more stakeholders compared to only looking at prefixes. We quantify the centrality of ASes with respect to specific sectors and the robustness of communication communities. Our results show that members of sectoral groups tend to avoid direct peering, but inter-connect via a small set

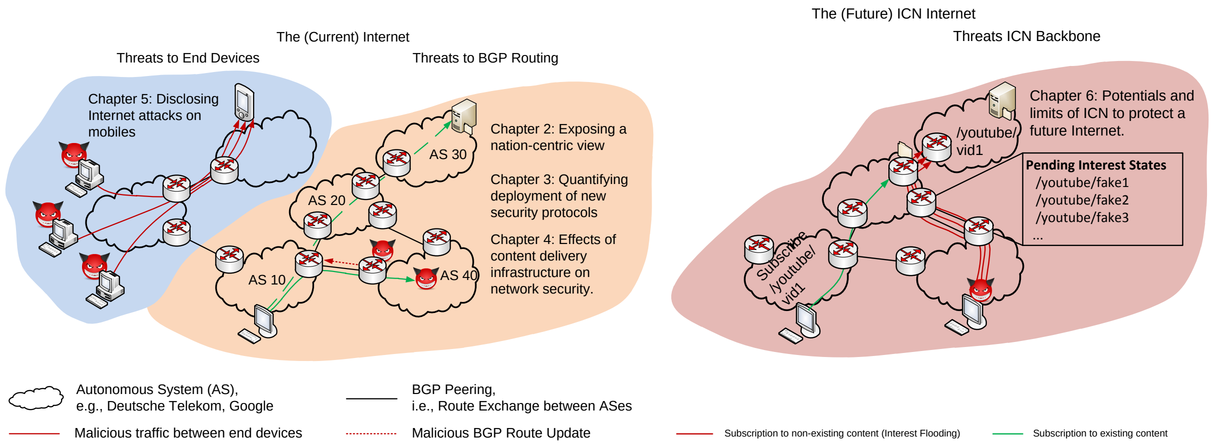
The (Current) Internet
Threats to End Devices
Threats to BGP Routing
Chapter 5: Disclosing Internet attacks on mobiles
Chapter 2: Exposing a nation-centric view
Chapter 3: Quantifying deployment of new security protocols
Chapter 4: Effects of content delivery infrastructure on network security.
AS 30
AS 20
AS 10
AS 40
Autonomous System (AS), e.g., Deutsche Telekom, Google
BGP Peering, i.e., Route Exchange between ASes
Malicious traffic between end devices
Malicious BGP Route Update
The (Future) ICN Internet
Threats ICN Backbone
Chapter 6: Potentials and limits of ICN to protect a future Internet.
/youtube/ vid1
Pending Interest States
/youtube/fake1
/youtube/fake2
/youtube/fake3
...
Subscribe /youtube/ vid1
Subscription to non-existing content (Interest Flooding)
Subscription to existing content

(a) Attacks in the current Internet

(b) Harming ICN infrastructure

Figure 1.1: Threat landscape

of common ISPs. Even though applied for Germany here, all methods are designed general enough to work for most countries, as well.

Analysis of RPKI deployment on BGP routers. In Chapter 3, we give first insights into the additional system load introduced by RPKI at BGP routers and present a new threat model for RPKI-enabled routers. Our observations are experimentally analysed using a real-time compliant, highly efficient reference C implementation of the RPKI router part, called RTRlib. RTRlib is integrated into existing BGP daemons. In particular, this is the only openly available tool for monitoring RPKI validation activities in real-time. After understanding the overhead and potentials of new attacks on the local system, we take a look into the real world by performing a long-term measurement using live BGP streams that evaluates the current impact of RPKI-based prefix origin validation on BGP routers. We observe that most of the invalid prefixes are probably the result of misconfiguration, while leading sources of deployment errors change over the measurement period. We measure a relatively small overhead of origin validation on commodity hardware (5% more RAM than required for full BGP table support, 0.41% load in case of ≈ 92,000 prefix updates per minute), which meets real-world requirements of today.

Analysis of the protection of web server infrastructure by RPKI. In Chapter 4, we present a first quantitative analysis of the protection of web servers by RPKI. We introduce an initial methodology that accounts for distributed content deployment and shall enable the content owners to estimate and improve the security of the web ecosystem. For a current snapshot, we find that less popular websites are more likely to be secured than the prominent sites. Popular websites rely on CDNs in larger parts, which did not start to secure their IP prefixes. Whenever CDN-content is protected by RPKI, it is located within third party ISP networks. This hesitant deployment is the likely cause why popular content experiences reduced security.

Analysis of context-specific attacks on mobiles. In Chapter 5, we argue for a simple long-term measurement infrastructure that allows for (i) the analysis of unsolicited traffic to and from mobile devices and (ii) fair comparison with wired Internet access. We introduce the design and implementation of a mobile honeypot, which is deployed on standard hardware for more than 1.5 years. Two independent groups developed the same concept for the system. We also present preliminary measurement results to gain insights whether attacks depend on the network access.

Analysis of threats to stability and security in Information-Centric Network infrastructure. In Chapter 6, we analyze threats to the stability and security of the content distribution system in (i) theory, (ii) simulations, and (iii) practical experiments. We derive relations between state resources and the performance of routers, and demonstrate how this coupling can be misused in practice. We further show how state-based forwarding tends to degrade reliability by decorrelating resources. We identify intrinsic attack vectors present in current content-centric routing, as well as possibilities and limitations to mitigate them. Our overall findings suggest that major architectural refinements are required prior to global ICN deployment in the real world.

We summarize our research questions and key contributions in Table 1.1 and Table 1.2.

1.4 How to Read This Thesis

This thesis covers several dimensions of security for today's and the future Internet. In this chapter, we presented basic background. Instead of presenting all further details in a single separate chapter dedicated to the background, we consider it beneficial to the reader to find the required background near to the related sub-topics. Following this concept, we present all sub-sections in a self-consistent way.

	Key Question	Methodology				Key Observations
		S	E	M	Data	
	Chapter 2: Exposing a nation-centric view on the distributed Internet for country-wise infrastructure protection.					
Internet Core	Which Internet resources comprise a nation-centric view of the Internet infrastructure?	✘	✘	✓	RIR DB, Maxmind	Using IP blocks instead of prefixes, we identified 25% more ASes. IP block to country mapping exhibits 0.2% false positives and false negatives compared with Maxmind.
	How robust is routing when viewed w.r.t. critical business sectors?	✘	✘	✓	RIR DB, BGP dumps	Members of the same public or business sector tend to not peer with each other, but interconnect via some selected national and also international ASes. Diversity of the first hop upstream depends on the sector.
	Chapter 3: Quantifying the challenges of deploying new security mechanisms in the Internet backbone.					
	Does RPKI protection slow down router performance?	✘	✓	✓	BGP dumps, RPKI data	A full RPKI validation table would lead to 5% increase of RAM. 10 million ROA entries can be loaded in less than one minute. Validating BGP live data (max. ≈ 92k per minutes) results in less than 0.5% CPU load.
↓	Does enhanced RPKI deployment increase router performance?	✘	✓	✘	Artificial BGP, RPKI data	Mixture of valid, invalid, and not found prefixes may change CPU load up to one order of magnitude.
	Are invalid BGP updates mainly hijacks?	✘	✘	✓	BGP dumps, RPKI data	In the beginning of RPKI deployment, up to 90% invalid updates are due to ROA misconfigurations. If the origin mismatches, in more than 90% of the cases, the ROA is only one hop away from BGP origin.
	Can we improve the number of erroneous invalids by changing a small number of ROAs?	✘	✘	✓	BGP dumps, RPKI data	30%-40% of invalid updates are caused by the top five ROA entries.
	Chapter 4: Analysing the effects of content delivery infrastructure on network security.					
Core / Edge Services	Do name prefixes influence service location?	✘	✘	✓	Active DNS data	For the first 100k Alexa domains more than 76% of the IP prefixes equal for www and w/o www domain names, for the remaining domains 94% exhibit the same prefix.
	How far does RPKI route security cover web server infrastructure?	✘	✘	✓	BGP dumps, active DNS	6% of the 1M Alexa domains are covered by the RPKI, where ≈ 0.09% of the BGP updates are invalid.
↓	Does website popularity correlate with route security?	✘	✘	✓	BGP dumps, RPKI, DNS	More popular domains are less secured. Only 4% of the 100k most popular domains are secured whereas 5.5% of the last 100k popular domains are secured.
	Why are more popular domains less secured?	✘	✘	✓	BGP dumps, RPKI, DNS	Popular web content is commonly served by CDNs, which are hesitant in deploying RPKI.

Table 1.1: Overview: Research questions, experiments, and key results. (S=Simulation, E=Emulation, M=Measurements)

	Key Question	Methodology				Key Observations
		S	E	M	Data	
	Chapter 5: Disclosing Internet attacks on mobile devices.					
Internet Edge	Does the amount of unsolicited remote connections depend on the network access type?	✘	✘	✓	Honeypot data	DSL, UMTS, and darknet nodes experience connections in the same order of magnitude.
	Do remote attacks specifically targeting mobile devices?	✘	✘	✓	Honeypot data	Intruders rarely conduct mobile-specific attacks. Rather, attackers frequently navigate through file system following common Linux structure.
↓	Does the location of attackers differ for mobile and stationary targets?	✘	✘	✓	BGP, Maxmind, honeypot data	Most of the attacks are initiated from the same set of ASes. The top-5 ASes are based in China and Russia.
Internet Edge / Core	***Chapter 5: Identifying potentials and limits of ICN to protect a future Internet.***					
	What is the attack surface of ICN infrastructure?	✘	✓	✘	CCNx	State maintenance may overload routers leading to dropping of data packets.
	How does current Internet delay space affect ICN performance?	✘	✓	✘	Analytical model, PingER data, CCNx	Internet delay is hardly predictable and highly heterogeneous. Network delays that exceed expiration timers lead to increased state maintenance. This requires significant over provisioning of router hardware.
	How does heterogeneous router hardware affect ICN network design?	✘	✓	✘	CCNx	The upstream of the weakest node needs 50%-500% more memory than any other node. Proper ICN delivery requires homogeneous hardware deployment.
	Is low network performance only a side effect of local, insufficient system resources?	✓	✘	✘	ndnSim, Sprintlink topology (Rocketfuel)	No. Our simulations show that ICN routing tends to not efficiently exploit transmission resources in realistic topologies because of statewise uncoordinated hop-by-hop forwarding.

Table 1.2: Overview (continued): Research questions, experiments, and key results. (S=Simulation, E=Emulation, M=Measurements)

Chapter 2

Exposing a Nation-Centric View on the Internet

2.1 Introduction

The Internet was originally shaped to offer open transmission services on a global scale, but has now turned into a mission-critical infrastructure of local relevance for most countries and dedicated players. The coherence of the Internet is defined by peering relations between ASes. Analyzing mutual impact, vulnerability and efficiency of the backbone requires the identification of ASes and corresponding transits between them.

Today, the global Internet is composed of more than 50,000 ASes with significantly more links, which challenge a clear picture on dependencies. Similar to traditional infrastructures, a country, its population as well as organizations share an obvious interest that the internal data exchange does not rely on weak third parties (cf., [136, 243 ff.]). In more detail, political regulations may prohibit that Internet traffic crosses specific countries. This might due to censorship reasons but also (positively motivated) to achieve political autonomy. Another example are traditional critical infrastructures such as power grids, which increasingly rely on a working Internet [103]. Having a clear picture about the international dependencies with respect to communication flows helps to improve resilience strategies for such traditional infrastructures.

In Internet terms, AS transits connecting key players of a country should be part of an apprehensible Internet ecosystem. However, the Internet is a globally distributed network without boundaries, which makes the identification of nationally relevant subparts hard and leads to the following questions: Which Autonomous Systems are important for a reliable interconnection of the Internet infrastructure of a country? How do sectors of a country communicate? Rigorous insight into the country-wise nature of the Internet thus carries fundamental importance and it is somewhat surprising that only recently the inter-network structures of nations attracted attention [84, 123].

This chapter contributes with the following first steps to answer these questions:

1. A promising methodology to derive a country-centric view on the Internet structure.

This is exemplarily verified for our home country Germany. The approach starts with IP-blocks. Compared to pure prefix-based techniques, we are able to identify approximately 25 % more members.

2. A novel, non-hierarchical AS taxonomy, as well as a heuristic sectoral classification technique. Both allows us to identify ASes with national relevance. By adding routing information, we are able to generate, visualize, and analyze the structure of communication flows between relevant public and business sectors. This has not been evaluated before.

3. The evaluation of the German inter-AS structure based on common and new graph metrics. This reveals for example that most eyeball providers peer dependent on the target AS, whereas the financial sector operates on static paths.

4. Finally, extracted and visualized data will be provided to the community for subsequent research and analysis.

The remainder of this chapter is structured as follows: Section 2.2 discusses the current state of Internet backbone analysis in the context of nation-state routing. Our methodology and corresponding toolchain, which allows for a nation-centric view on the Internet, is described and evaluated for Germany in Section 2.3. Section 2.4 presents a visual approach to get an intuitive view on the derived structure. Section 2.5 analyses AS (sub-)graphs from Germany in detail based on graph metrics. Finally, Section 2.6 concludes this chapter.

2.2 Related Work

Research on the Internet AS structure continues to attract significant interest since more than one decade. This includes the analysis of structural properties of the inter-AS graph [53], the (mainly hierarchical) classification of ASes [67], and the inference of the relationship between them [58]. Active measurements within dedicated countries [161] reveal geographic reachability of ASes and thus follow a direction distinct from our work. Until now, there is only little work on a nation-state understanding of the Internet backbone routing, as well as on a horizontal classification of ASes that aims to identify key players of relevance for the Internet services of a country.

Dimitropoulus *et al.* [50] introduce a broader AS taxonomy (large/small ISPs, customer ASes, universities, IXPs, and network centers), but this does not include a detailed decomposition of customer ASes into dedicated sectors. The proposed inference algorithm analyzes the description value of the Internet Routing Registry [79] and follows a text classification technique. The authors focus on a complete mapping of ASes to classes. This differs from our perspective, as we do not intend to classify *all* ASes, but concentrate on *important* players of a country, viewed in further differentiated sectors. Cai *et al.* [32] introduce the interesting idea of an Internet AS ecosystem, which is based on a novel AS to organization map. The authors normalize contact records of the RIR WHOIS data to cluster AS numbers that belong to the same organization. Although this work could be applied to peering analysis and planning as well as threat analysis, our work is orthogonal as we classify ASes according to roles in a country.

The first paper that proposes a nation-state view on the Internet routing measures the impact of countries on the global data forwarding [84]. Karlin *et al.* start from IP prefixes, which they map to ASes. Routing paths are derived from an approximation of active traceroute measurements. We will show that IP prefixes are too coarse-grained to obtain an in-depth picture of a country and miss 25 % of German ASes. It is important to note that the base set of the national classification should be as complete as possible to judge on relevance.

Roberts and Larochelle [123] present a mapping project that visualizes and quantifies the relevance of ASes for countries. The AS to country mapping is based on the external service Team Cymru [5], which starts from IP prefixes. The authors introduce network maps of countries. Each map abstracts connections to the outside by a single Autonomous System that subsumes all foreign ASes. The relevance of an AS increases with the number of prefixes reachable via this AS. ASes with multiple upstream peers share routes equally among parents. This model oversimplifies common practice in Internet backbone routing. Normally, countries do not have a single entry point apart from China with the exclusive entry China Telecom. Furthermore, multilateral peering allows for path selection depending on the target AS. This diversity is not reflected and causes a weight distortion of AS importance.

To the best of our knowledge, current approaches do not provide sufficient mechanisms for identifying country-specific ASes, categorize ASes in business sectors, nor analyze importance of inter-AS communication between network domains for a country including international inter-connects. We will address these topics below.

2.3 Methodology

Deriving Nation-Centric Subsets of the Internet

We want to identify all Autonomous Systems of the global Internet that host organisations from a specific country. Many organizations are normal ISP customers and do not own a prefix or AS. Thus we argue that the appropriate granularity must be IP blocks. An IP block is a subset of a prefix and will be assigned internally by the prefix owner to departments or customers. An organization and thus its hosting AS is coined to a country if the organization or its administrative contact person of an IP block is located therein. Consequently, we include also ASes in our view with primary base outside of the investigated country, as well as national organizations with IT infrastructure outsourced to foreign countries.

To demonstrate and validate our approach, we choose our home country Germany (*DE*). The introduced methodology, however, can be applied to other countries, as well. Its implementation is easily extendable and thus provides a good base for the community and subsequent work. In

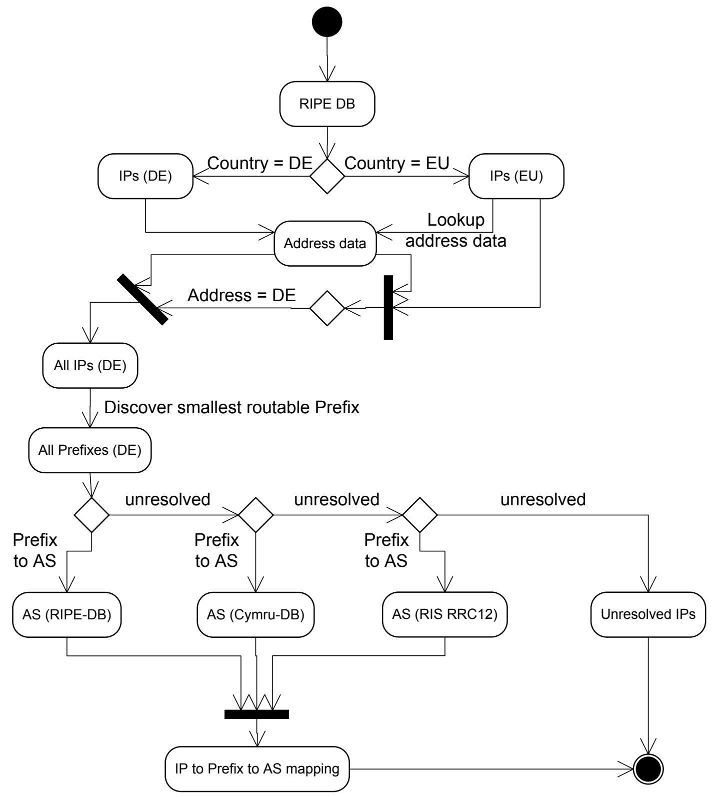

Figure 2.1: Tool chain for identifying national ASes.

this section, we present the data sources, our inference and classification algorithm to derive a nation-state view, and the the fully automated construction of existing interconnects. These results enable us to analyse the composition of the nationally relevant part of the Internet in detail.

2.3.1 From Internet Members to ASes

Regional Internet Registries (RIR) maintain network and contact details of their region. IP addresses and AS information related to Germany are registered at the RIPE database (DB). We start by extracting all `inetnum` records, which represent IP-blocks, from the RIPE DB that carry the mandatory `country` attribute of either *DE* or *EU*. Additionally, we collect address data for the associated `admin-c` and `org` objects. Based on the latter, we created keyword lists of synonymic country codes (e.g., Germany, DE), local city names, and international dialing codes (e.g., 0049, +49) representative for Germany. Applying the keywords on the contact record allows us to further resolve *EU* IP-blocks to *DE* and to verify the *DE* classification of IP-blocks. The result is a list of all IP-blocks allocated by organizations in Germany.

Next, we determine the longest covering IP-prefix for each IP-block. Prefix lengths are

Approach	**DE**	**EU**	**other**
IP-Block	6,278	–	–
Prefix (RIPE DB)	5,243	–	1,035
Prefix (Team Cymru)	4,395	947	936

Table 2.1: Number of identified prefixes

subject to aggregation and thus depend on the point of observation. Using passively measured BGP data from distant route collectors would be too coarse grained and yield less specific prefixes. Assuming that RIRs provide the most detailed prefix mapping, it is reasonable to query the RIPE DB. The inter-AS route is specified in the `route` record, which is referred by the `inetnum` object.

Finally, we map the prefixes to origin Autonomous System Numbers (ASNs) by the `route` object. However, in this step using the RIPE DB alone would lead to several unresolved mappings. Therefore, we also consider data of Team Cymru [5] and the route collector RRC12 of the RIPE RIS [4]. The latter peers at the largest German Internet Exchange Point (DE-CIX) and thus provides localized data. We apply the different data sources in the following order to maximize the number of resolvable ASNs: (1) RIPE DB, (2) Team Cymru, and (3) RIS RRC12. In cases of Multiple Origin ASes, we keep all discovered ASNs. The resulting list contains the ASes that compose the nation-centric part of the Internet for the example of Germany.

Our fully automated tool chain was applied in Oct. 2010 and yielded 246,861 German IP-blocks. Thereof 240,237 are embedded in 6,278 IP-prefixes that belong to 1,471 ASes, ≈ 2 % could not be resolved to a prefix. To estimate errors of our IP-block-to-country mapping, we checked back with the well-known MaxMind GeoLite Country service [3]. Deviations were found below 0.2% for both false positives and false negatives.

Our method of starting from IP-blocks rather than IP-prefixes identifies significantly more prefixes that carry relevance for the country Germany (cf., Table 2.1). When considering prefixes alone, only ≈ 84% can be identified as ’German’ using RIPE-DB, while Team Cymru yields ≈ 70% *DE* prefixes. Thus a significant fraction of prefixes that route traffic relevant for Germany is not directly associated to country or address values from the this country.

Providers from outside Germany are also selected by our scheme. The corresponding 301 ASes (≈ 20 %) are classified relevant for nation-state routing and internationally distributed as follows. More than a third (110) ASes originate from direct geographical neighbors, another third (107) from the remaining Europe, thereof 57 British, and 18 % (54) are North American ASes. These classifications are again based on RIR databases and Team Cymru with an estimated error of about 15 %.

2.3.2 Tier and Sector Classification of Autonomous Systems

Having categorized the nationality of the stakeholders, we add two further classifications to the selected ASes. First, we harvest the topological hierarchy (tier1, large/small ISP, and stub) from [158]. Additionally, we investigate the role of ASes within those public or business sectors that are operationally relevant for the country. As there is no AS classification available that describes the professional role of an organization in relation to the global BGP routing, we introduce a sectoral categorization. This extends the taxonomy of critical infrastructure published by the Federal Office for Information Security (BSI), Germany. We determine sectoral classification by applying an optimized and manually verified keyword spotting to names, descriptions, and address fields of the previously derived AS data. Our approach uses general terms such as “bank”, but also specific company names (e.g., Siemens, Daimler) as keywords associated with classes to identify important ASes. Keywords are correlated to enhance the identification. Thereof we obtain an additional list of the ’relevant national ASes’ including branch tags such as *financial services* (cf., Table 2.3).

99 % of the classified ASes belong to exactly one sector, five ASes are assigned to two sectors. Companies may attain multiple roles. For example, AS 31438 is a municipal utility responsible for waste water and DSL access in the City of Marburg. Overall 279 ASes have been selected as ’systemically relevant’ with sectoral attribute attached.

We admit that this step includes manual pre-definition of keywords for sectors. However, mapping ASes to sectors is a fine-grained process, which requires specific information that cannot be derived completely automatically. Our taxonomy, methodology, and tools can be applied to other countries based on an updated keyword list. The creation of the list needs local knowledge.

2.3.3 Constructing Spanning AS Routing Graphs

Following the identification of all ASes relevant for the German Internet infrastructure and a classification of key players, we derive their interconnects. We limit the building of AS graphs to the construction of inter-AS paths without considering individual prefixes. This modeling step is meaningful for our purposes: Even though BGP policies and regional optimizations may lead to varying paths for different IP prefixes announced by the same origin AS, recent studies [152] show that multiple prefixes are reachable via the same AS path for 75 % of origin ASes. Additionally, we focus on a regionally bound network, which is densely meshed by peering points. International redirections within service provider networks are mainly outside the scope of our perspective. From this point of view, restricting the routing on the AS level is a valid approximation.

We identify an AS routing graph for each sector, the bilateral exchange between two sectors, as well as the AS graph of all DE ASes based on the weighted next hop matrix provided by the NECLab topology project [1]. This data is calculated using the continuously updated mea-

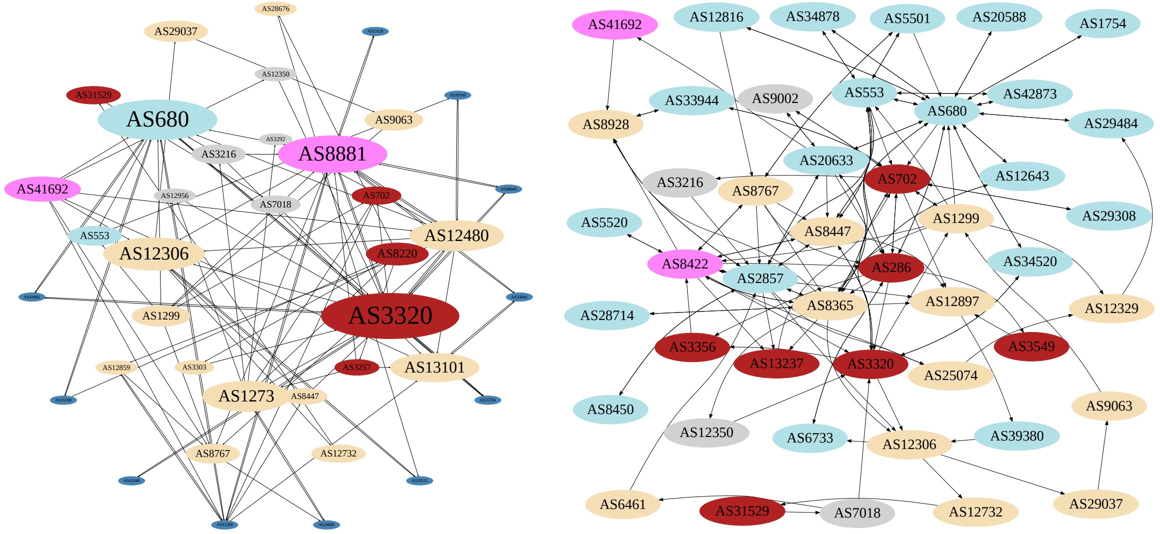

(a) Governmental organizations, node size represents betweenness

(b) Kamada-Kawai drawing for sector research and culture

Figure 2.2: [Best viewed in color] Different visualizations of two sectors

surements by the UCLA [158] and reflects BGP policy decisions [152]. To exclude incomplete paths, we omit row column values of -1 for distinct indices during matrix processing. Note, this occurs very rarely ($\ll 0.3\,\%$). Naturally, the set of ASes in the routing graphs has been extended by intermediate ASes that we have not assigned before to the nationally relevant part of the Internet. These transit nodes are required to link nation-state subsets that would remain isolated otherwise.

2.4 Visualization

Adding AS-level routing relations to our selected and classified AS sets allows us to study and visualize very specific topological set-ups. Although sectoral graphs contain only a very small fraction of ASes as compared to the full Internet, a meaningful visualization of the structure is still a challenge due to the large amount of links.

We created four graph types, a hierarchical, a circular, the Kamada-Kawai, and a flow model. We decided for these four types for two reasons. First, the Kamada-Kawai model is a common approach to visualize complex networks. Second, the other three types have been designed after discussion with potential users. This visualizations allow easily to realize the basic information (e.g., which ASes are important). It is worth noting that additional benefit will be achieved when these graphs are integrated in an interactive software. For example, to zoom into the graph or present additional meta information. All graphs are plotted by GraphViz [6]. The layout processors of GraphViz has been extended to reflect the Internet specific properties.

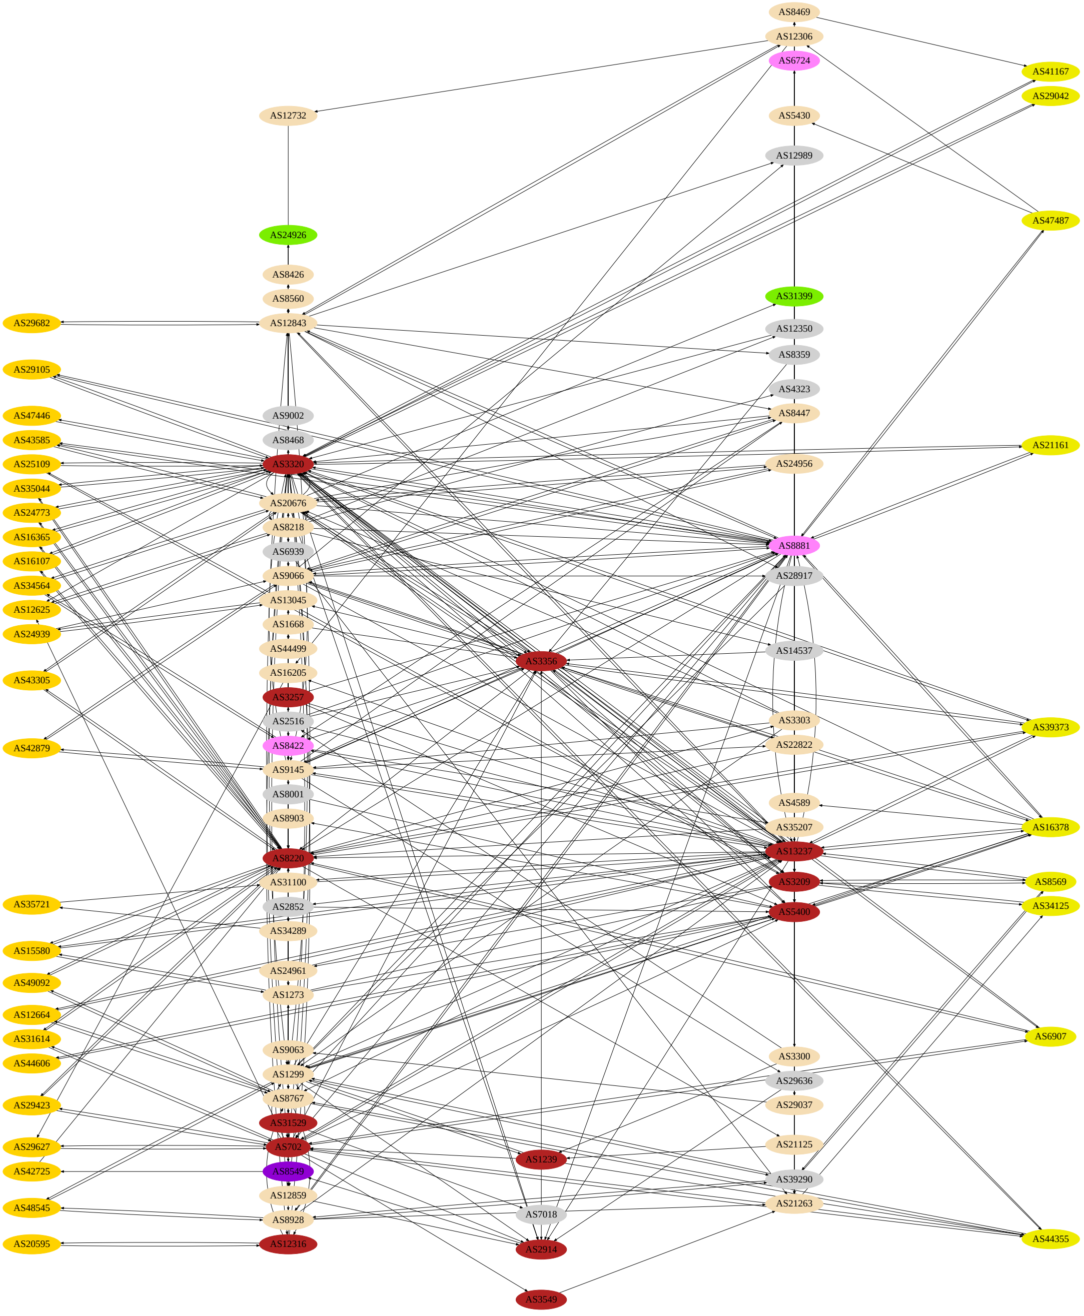

Figure 2.3: [Best viewed in color] Communication flows between traders (right) and financial services (left)

Figure 2.2(a) depicts the AS routing graph of German governmental organizations. The nodes are placed concentrically according to their tier level. Tier 1 providers shape the inner circle. The width of the vertices is scaled with the importance of the AS in delivering data between sectoral members, colors describe the sectoral classification (red: large ISPs, pink: access providers etc.). All grey nodes ASes are only required for transit and do not host a DE IP block.

To reflect meshing, we apply the Kamada-Kawai layout [83]. It models edges as springs and transforms the complete graph into an energetic minimum. Figure 2.2(b) illustrates the spanning graph for research and culture. It reveals that the role of the German National Research and Education Network (AS 680) is less central as commonly assumed. Regional networks (e.g., AS 553, AS 8365 or AS 8422) and commercial ISPs (e.g., AS 702) provide significant connectivity to the community, as well.

The network graph in Figure 2.3 displays the Internet information flows between financial services (leftmost column) and traders (rightmost column). Here it should be noted that the minimal spanning routing system of our AS sets adds transit ASes to the graph. These intermediate hops are visualized as follows. Tier1 providers are arranged in the middle column, while the remaining ISPs are placed at the two intermediate columns. Positions left and right of the center were chosen according to the number of direct links with ASes from each sector. AS 3320 (DTAG), for example, was placed next to financial services, as it provides the majority of upstream links to this sector, while the opposite holds for AS 13237 (Lambdanet).

Despite the relatively small number of ASes for these two sectors, it takes quite many transit providers to interconnect all selected networks. Moreover, taking a closer look at the graph reveals that peering density is clearly asymmetric. Selected ASes (like DTAG for banks or Lambdanet for traders) play a dominant role within individual sectors, while at the same time mutual peerings remains completely absent.

2.5 Analysis of the AS-Structure

In this section, we investigate structural properties of the derived AS routing graphs and measure the relevance of members in sectors and in the overall DE AS graph. It is worth noting that we keep BGP policy modeling by a per path analysis, each path derived from the *weighted* next hop matrix of the NEC topology project. All measurements are relative to allow for comparing sectors of different sizes. Unfortunately, the underlying next hop matrix does not provide edges to connect the five ASes of the medical sector.

2.5.1 Node Centrality

Intermediate nodes between source and receiver attain a relevant role from serving as transits. The number of shortest paths passing through a node m is quantified by the betweenness $B(m)$. If the total number of shortest paths between two nodes i and j is $B(i,j)$, and the number of

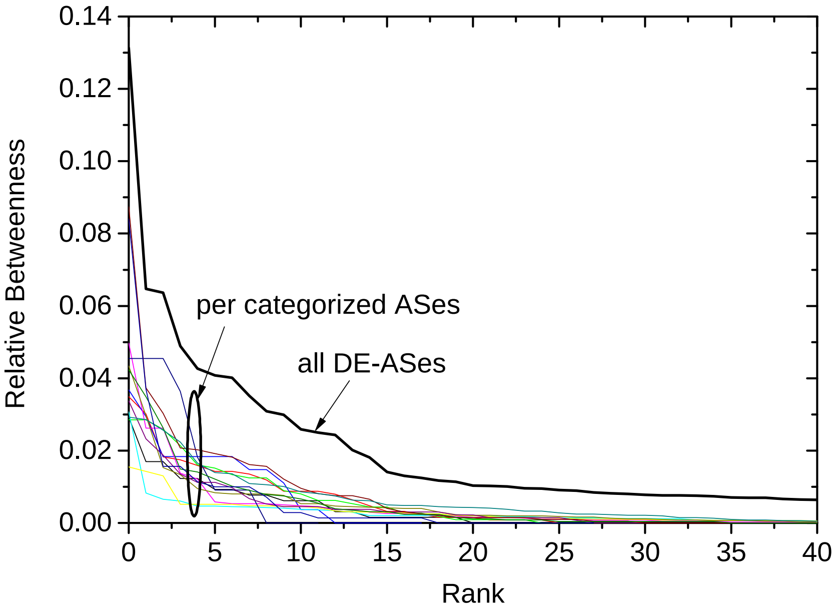

Figure 2.4: Properties of the Internet relevant for Germany and its categorized subgraphs: Relative betweenness

All DE		**Large ISPs**		**Research**	
Ranked AS	Betweenness	Ranked AS	Betweenness	Ranked AS	Betweenness
1. DTAG	0.131	1. DTAG	0.031	1. DFN	0.087
2. Level 3	0.065	2. Lambdanet	0.008	2. Verizon	0.037
3. Lambdanet	0.064	3. Telekom–AT	0.007	3. Manda	0.030
4. Colt	0.049	4. France Telecom	0.006	4. BELWUE	0.021

Table 2.2: Relative betweeness of the top ranked ASes for selected sectors

these paths going through node m is $B(i,m,j)$, then the betweenness of m is defined as the ratio: $B(m) = \sum_{i \neq m \neq j, i \neq j} \frac{B(i,m,j)}{B(i,j)}$. This measurement quantifies also the load at intermediate ASes. The betweenness is normalized by $(|V|-1)(|V|-2)$.

The *term shortest path* refers to the routing path that is actually taken. Our underlying BGP routing model reflects policies [152]. Using the NEC matrix, there exists exactly one effective path between two ASes. However, as discussed in Section 2.3, in our context of locally bound routing this is not a restriction. Independent of the nation-state view, BGP policies may lead to a violation of the triangle inequality. As the routing paths are based on a weighted graph, this property is preserved.

We calculate the betweenness of a node for the routing graphs under discussion. Figure 2.4 shows the relative betweenness, where ASes are ranked in decreasing order. Details for selected

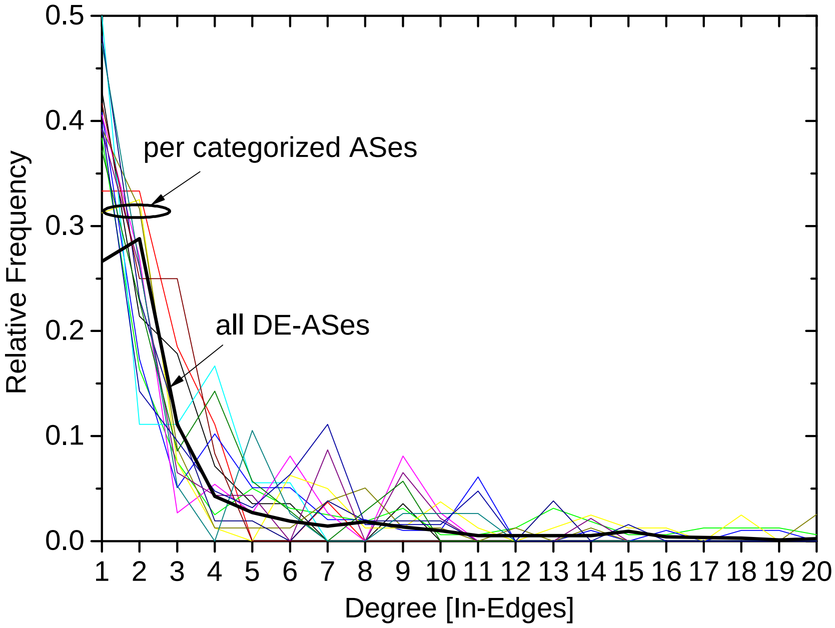

Figure 2.5: Properties of the Internet relevant for Germany and its categorized subgraphs: Indegree

sectors are listed in Table 2.2. In 80% of the cases, this measurement exhibits sharp peaks at the transition from the top most important AS to the second one. This means that in the selected category a dedicated AS is part of a significant number of shortest paths and thus attains a major role in data forwarding. However, the decay from the top most ranked ASes is less steep in the overall German AS graph, showing a more evenly distributed relevance due to increased peering links. Looking at the actual rank orders reveals a relatively stable number of ASes among the top five in each category. For example, AS 3320 (Deutsche Telekom) has in 80% of the cases at least rank 5 and in 48% the highest betweenness.

2.5.2 Degree Distribution

The degree of a node denotes the number of its one-hop neighbors. Figure 2.5 shows the indegree distribution. For visibility, we cut the tail at 20 edges. Overall, the relative frequency decays polynomially for all networks. Thus, there is a higher probability to maintain only a quite limited number of peering relations, but a non-vanishing likelihood for high peering numbers. The distribution of the full DE AS graph decays smoothly, while sectoral groups exhibit systematic peaks for selected node degrees between four and 13. Consequently, specific networks within the sectoral subgraphs are more densely connected than the full graph. These additional weights indicate regional star topologies in sectoral networks.

When comparing the topology within sectors to the complete DE network, we find more pronounced betweenness' and irregular peaks at increased node degrees. Jointly, these two

Sector (# ASes)	$\langle X\rangle$	σ_X	**Sector (# ASes)**	$\langle X\rangle$	σ_X
Transit providers (55)	2.41	0.92	Industry (28)	3.19	0.87
Trading (10)	2.69	0.87	Financial services (32)	3.21	0.89
Science & Culture (22)	2.77	1.12	Shipping & transportation (15)	3.22	0.8
Eyeball ISPs (23)	2.83	1.16	Public administr. & justice (14)	3.22	1.14
Peering points (8)	2.87	0.97	*DE All (1,471)*	3.23	1.04
Public services (4)	3.00	0.6	Energy (11)	3.34	0.79
Media & publishers (19)	3.08	0.94	Other public services (7)	3.40	0.73
Software and systems (31)	3.18	0.89	Medical services (5)	–	–

Table 2.3: Absolute number of ASes per sector and DE graph as well as mean ($\langle X\rangle$) and standard deviation (σ_X) of the distance distributions for corresponding routing graphs

structural metrics indicate that individual ASes provide enhanced connectivity within the specific communities as opposed to direct interconnects. A closer look on the corresponding AS graphs supports this observation. The majority of financial services, for example, tend to peer via Deutsche Telekom (AS 3320) and Colt (AS 8220), while no mutual peering is visible at all. Surprisingly, the governmental federation follows the same pattern. Governmental organizations are mainly interconnected by Deutsche Telekom and Versatel (AS 8881), but a small group uses Plusline (AS 12306) as upstream provider. The latter organizations require the external tier1 ISP AT&T to serve as inter-connect to the remainder of this sector.

2.5.3 Distances

The distance distribution of shortest paths measures the probability that two randomly selected nodes of a network are connected at distance k. This metric describes routing performance and usually follows a Gaussian law. This observation is also reflected in the analysis of the sectors and DE routing (cf., Figure 2.6) with average values between 2.4 and 3.4 (cf., Table 2.3). Routing distances, thus, largely depend on the sector under investigation. Naturally, connections between ISPs are shorter as compared to other branches. Surprisingly, ASes of the trading sector are significantly short, as well. In this group, the majority of members are connected via the same ISPs. Deutsche Telekom (AS 3320) and Vodanet (AS 3209) play a dominant role for transit. Even though there is no bilateral peering within the sector, many traders (e.g., Ebay AS 6907) maintain extensive peering relations. Paths that consist of only one transit hop are easily established. In general, our results show a similar behaviour compared to the global AS topology [100]. Even in the relatively small sectors, interconnects are not significantly denser and path lengths are not generally reduced. Most of the members from the same sector seem eager to stay at distance to each other.

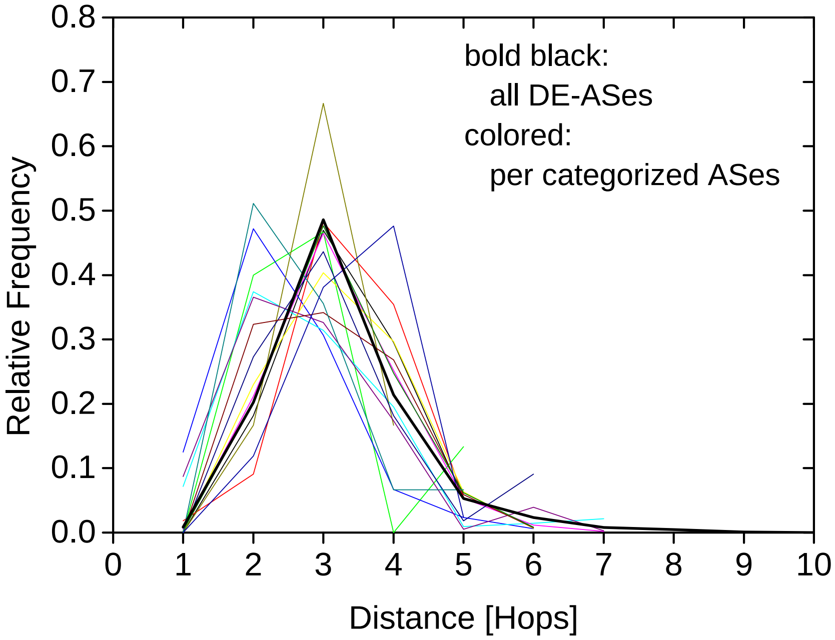

Figure 2.6: Properties of the Internet relevant for Germany and its categorized subgraphs: Distance distribution

2.5.4 Context Dependent Peer Selection

To analyze the peering behaviour in more detail, we answer the following question: How likely does a member of a sectoral group chooses its upstream peer dependent on the sector it communicates with? For each member of a sector, we count the number of different upstream peers relatively to the overall number of paths towards members of distinct sectors. We quantify the relative frequency of corresponding diversity classes over all members of a sector. For ASes that peer with many ISPs, it indicates high probability for high first hop neighbor diversity.

The calculated upstream peer diversity is very heterogeneous and may appear as a characteristic feature of the sectors. Figure 2.7 presents the measurements for selected sectors. Members of the financial services, for example, exhibit constant paths in about 50 % of the cases with enhanced probability (cf., Fig. 2.7(b)). In contrast to this, 80 % of the transit providers select above 80 % of the time neighbors dependent on the target. In general, target specific peering is dominant, as 8 of 10 ASes choose their one-hop neighbor with respect to the destination.

Combining the results with our previous findings indicates that multilateral peering has dominant routing effects on the Internet subpart relevant for Germany. For sectors, however, this higher amount of interconnections does neither result in more densely meshed inter-AS links nor in shorter paths.

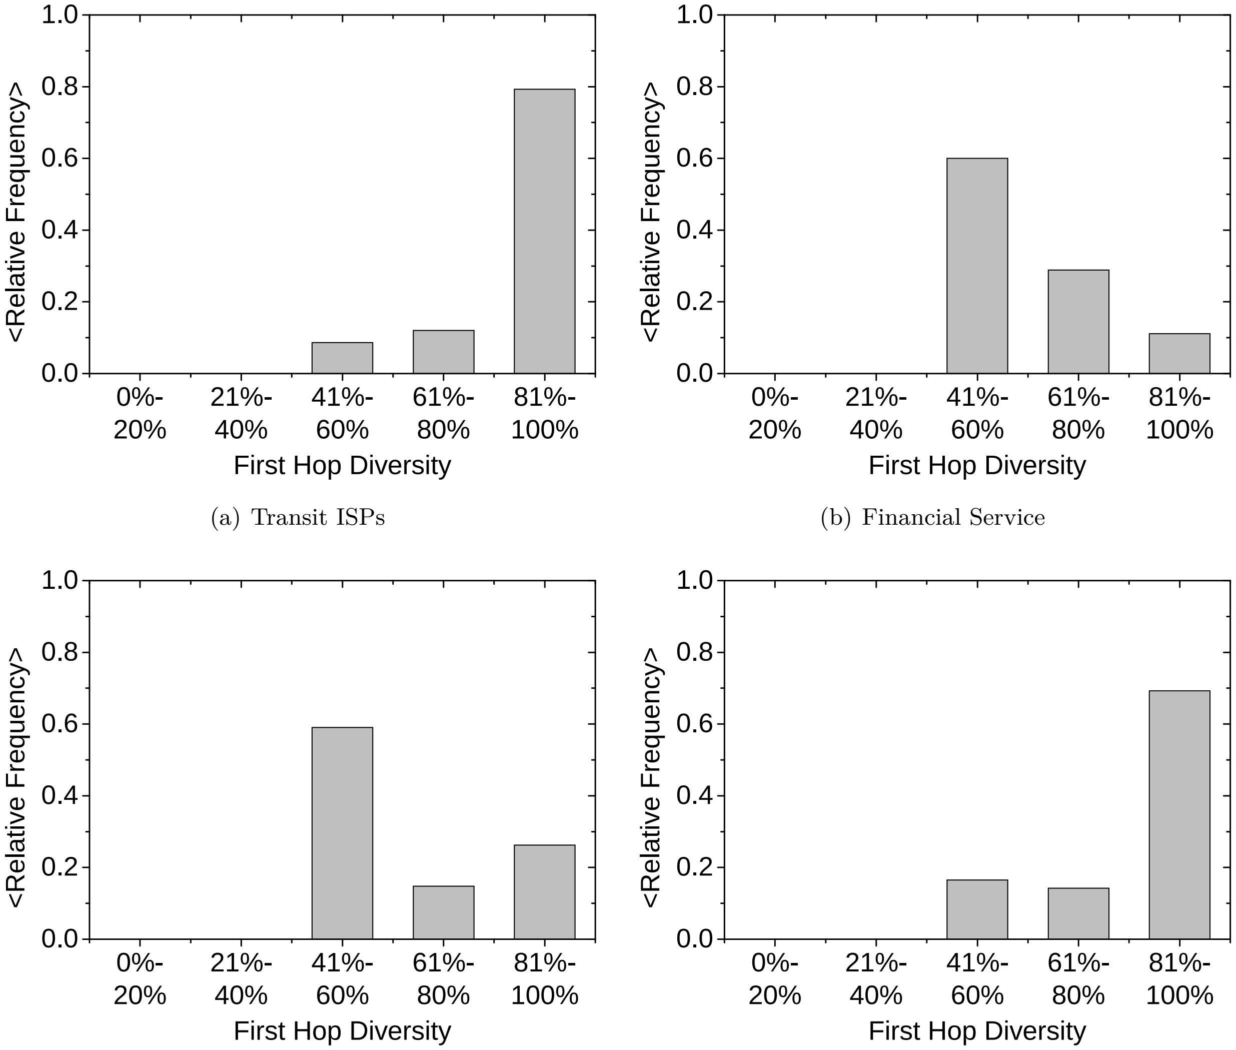

(a) Transit ISPs

(b) Financial Service

(c) Industry

(d) All DE

Figure 2.7: Relative upstream diversity for selected sectors

2.6 Conclusion and Outlook

A clear understanding of the inter-AS structure at the country-level is needed to cope with the interdependencies and intrinsic vulnerability of the Internet. In this chapter, we presented a methodology to identify and classify the relevant ASes of a country. This led to a fine-granular view onto meaningful subsets of Internet stakeholders and a detailed analysis. To the best of our knowledge, this is the first inspection of a country and its key players at the Internet routing sub-structure. The evaluation was exemplified for our home country Germany (DE), and created a list of relevant DE ASes including administrative data and sectoral classification, which will be publicly available.

We associated Autonomous Systems with a country whenever they host IP address space for an organization from there. Our approach outperformed prefix-based techniques by identifying 25 % more ASes. In particular, we were able to spot parts of the public sector hosted by international providers. Our analysis further revealed that members of the same public or business sector tend to not peer with each other, but interconnect via some few national and also international ASes. Deutsche Telekom, Level 3, Lambdanet, Colt, and Versatel were found to be the most important transit ASes for intra-DE communication. Multilateral peering was seen to have dominant routing effects on the German Internet, but the degree of variable upstream selection strongly depends on the sector.

Our future work will extend the current results in both directions, structural analysis and further countries, including their interdependencies. We expect structural properties on a fine-grained basis. In addition, we will extend our analysis towards IPv6. Regarding integration, we will employ current aggregation techniques for ASes belonging to the same organization to derive a condensed national Internet AS ecosystem. Finally, we will also concentrate on the application of our work to existing monitoring systems (e.g., [39]), which may help to reduce complexity due to selected observation points.

Chapter 3

Quantifying New Security Mechanisms in the Current Internet Backbone

3.1 Introduction

The Internet is a critical infrastructure in many countries [34]. Almost all common infrastructures depend on it. Public services, large business sectors etc. are based on a working communication backend. As an illustrative example we refer to the blackout in Italy 2003 where a power outage resulted in a failure of the Internet communication, which finally caused an additional power outage [27]. These cascading effects are complex. In this chapter, we focus on much simpler mechanisms to disturb Internet communication.

The Internet fundamentally rests upon the routing backbone created by the Border Gateway Protocol (BGP). BGP announces IP prefixes to enable inter-domain routing between so-called Autonomous Systems (ASes). One major problem of BGP is its lack of verifiable information exchange. Several prominent incidents highlighted the consequences [136, 26]: An AS incorrectly claims to own an IP prefix and thereby redirects traffic, which not only may lead to traffic interruption, but can be used to intercept and tap data streams. It is worth noting that not only prefixes but also autonomous systems can be hijacked [127], which is not part of this thesis.

Several approaches (e.g., [74]) have been proposed to identify prefix hijacks heuristically. Only recently, countermeasures for hijacks have been deployed in the form of the Resource Public Key Infrastructure (RPKI) and related protocols, which allows for a cryptographical strong binding between prefix and origin AS. A successfully deployed RPKI origin validation would have immediately disclosed the prefix hijacks referenced above—a rigorous route rejection of the invalid updates would have prevented the incidents entirely. However, the second, operationally thrilling step of RPKI-based route selection can face reality only after a high degree of confidence in a reliable deployment has been established.

From the research as well as the operator perspective, a fundamental change such as the introduction of the RPKI needs a careful observation and evaluation from the beginning. RPKI-enabled routers may introduce new security threats by complexity attacks, for example. In addition, they will identify invalid prefix updates independent of its cause. In the future, when

RPKI has become an integral part of the Internet, they should reject them, even if the reason for the incorrect announcement is a misconfiguration of the prefix origin attestation. For the research community this requires the adaption of routing models. For operators, practical insights are needed to overcome operational pitfalls and to identify optimization potentials.

In this chapter, we derive in-depth insights in the local prefix origin validation at routers as well as the currently deployed RPKI infrastructure. We analyze the impact of prefix origin validation on BGP peers and the potential consequences for the current BGP-based route propagation. The contributions of this chapter are in detail:

1. We propose and analyze a threat model for RPKI-enabled routers. We show that there is a variety of options to harm the local system by enabling prefix origin validation. This ranges from rather unlikely attacks to simple threat actions including complexity attacks that are easy to implement.

2. Based on live BGP update streams representing more than 100 peering neighbors, we present a long-term measurement highlighting two months that verifies 420 million IP prefixes against the available ROA data. We observed that the current RPKI system does not suffer from specific attacks but from reduced data quality in the RPKI. Most of the invalid prefix announcements are most likely due to misconfiguration of the prefix attestation objects.

3. We extract the key lessons learned from the data observed during our measurement period and derive advice for ISPs on future operational use of the RPKI.

4. We present a real-time compliant, highly efficient implementation to secure inter-domain routing at BGP peers. This open-source software is a reference implementation of the latest IETF protocol standards written in C. It features a flexible architecture and can be used to extend existing BGP daemons at real routers but also to write new tools in the context of RPKI/BGP research, or to further evaluate RPKI-based prefix origin validation at routers.

5. We systematically explore the overhead of prefix origin validation at commodity router hardware. Enabling RPKI will require $\approx 5\%$ more RAM compared to the storage of the global BGP routing table. The CPU load does depend on the RPKI-deployment state. We also characterize the overhead a BGP router experiences from validating the BGP live streams.

6. We contribute the software and our data collection to the research community for future work.

Up until now, the research on RPKI is quite limited. The Internet Routing Registries provide looking glasses into the RPKI, but a systematic analysis of the RPKI deployment state and its effect on BGP routers is missing. There are no results from monitoring BGP routers that have life RPKI-validation enabled. To the best of our knowledge, no open-source implementation of the RPKI router part exists that can validate in real-time on operational BGP peers. There is no analysis about attacking an RPKI-enabled router.

The remainder of this chapter is structured as follows. In Section 3.2 we present the background on BGP vulnerability and standardized countermeasures. Section 3.3 discusses related work in more detail. We present a threat model for RPKI-based prefix origin validation at routers in Section 3.4. We introduce a library for RPKI router support including performance and experimental threat analysis in Section 3.5. Section 3.6 presents our long-term measurement of prefix origin validation and analyses the data in detail. We conclude and give an outlook in Section 3.7.

3.2 Background

3.2.1 BGP Vulnerability – Prefix Hijacking

The Border Gateway Protocol (BGP) [119] is the inter-domain routing protocol for the Internet. It implements the dynamic exchange of network layer reachability information (i.e., IP prefixes) between Autonomous Systems (ASes). An AS represents a set of prefixes that it originates, but may also operate as pure IP transit. ASes are identified by unique Autonomous System Numbers (ASNs). Larger ISPs service multiple IP prefixes as well as several ASes [32].

Prefix information will be carried in BGP update messages. To make a network globally visible within the Internet backbone, a BGP speaker announces the corresponding IP prefix to its AS neighbors, which will re-distribute the information to adjacent ASes according to local policies. Based on the path attribute within a BGP announcement, a receiving AS is aware of all intermediate ASNs the update message has traversed. For example, if AS 30 receives the update `10.20.0.0/16: AS20-AS10-AS1`, it knows that IP addresses 10.20.0.0 – 10.255.255 are reachable via AS 20 and that AS 1 claims to be the origin AS of this prefix. As a BGP router commonly peers with multiple ASes, it may receive different routes to a single prefix. To ensure uniqueness, it chooses the preferred path for one prefix for inclusion in its forwarding information base (FIB). Data forwarding then follows a longest common prefix match, i.e., the most specific prefix will be selected from the FIB (e.g., 10.20.0.0/16 instead of 10.0.0.0/8 for destination 10.20.1.1) and data will be relayed to the next hop AS.

BGP is based on mutual trust. The currently deployed BGP version 4 does not provide any function to cryptographically verify BGP data carried in update messages. Consequently, BGP is vulnerable from two basic attacks: (a) a BGP router incorrectly originates an IP prefix, (b) a BGP speaker falsely modifies the AS path. Both attack vectors misguide traffic to an incorrect AS. In this paper, we focus on the deployment of upcoming standards to overcome the first problem, as well as the analysis of the corresponding infrastructure.

In BGP, a speaker may announce any IP prefix and receiving peers cannot validate the origin of the prefix. The injection of incorrect BGP data into the Internet backbone by intention or by misconfiguration is easy by design. Some heuristic prediction methods [94] have been developed for discovering illegitimate announcements but tend to fail in cases of multiple origin ASes (MOAs). The Internet Routing Registry (IRR) [79] is a collection of distributed routing databases, where routing data is derived from the RIR databases or commercial DBs. Some ISPs use the IRR to automatically configure announcement filters or to adjust policies. However, the data in the IRR is inaccurate and not complete. This leads to failures. On the other hand, as the correctness of the data cannot be proved cryptographically, many providers do not use it at all.

There are several prominent examples for illegitimate claims of IP prefixes. In February 2008, Pakistan Telecom announced by accident a more specific prefix for YouTube's network (208.65.153.0/24 instead of 208.65.152.0/22), which led to an outage of YouTube for more than 1.5 hours [26]. In April 2010, 15% of the Internet traffic has been redirected to China Telecom, which incorrectly claimed origin of ≈ 37,000 IP prefixes [136, 243 ff.].

To overcome the attack model of prefix hijacking, a cryptographically strong certification system is required that attests the ownership of ASNs and prefixes as well as the legitimate origination of an IP prefix by an AS.

3.2.2 Towards Secure Inter-Domain Routing

Securing BGP requires validation procedures for (i) prefix origins and (ii) AS paths in public announcements and has been introduced by S-BGP [85]. However, the changes required to migrate the global routing system to cryptographically validated updates have never been adopted by the community. Only recently, the IETF SIDR group has developed a standard proposal to authenticate the prefix-to-AS origin relationship without changing the BGP 4 protocol that also offers a straight-forward path to incremental deployment.

RPKI The Resource Public Key Infrastructure (RPKI) [97] is a PKI framework dedicated to securing the Internet routing infrastructure. It includes certificates that prove the ownership of Internet number resources (ASNs and IP prefixes). In general, resource distribution within the Internet follows an assignment hierarchy. To establish a trust chain, each resource certificate includes the delegated resource, the public key of the resource holder and is signed using the private key of the resource assignment instance. Issuing resource certificates follows the way of resource allocation in the Internet. The IANA issues certificates for RIRs (e.g., RIPE), RIRs for LIRs (e.g., ISPs), and LIRs for end users [77].

All certificates are included in a distributed, openly accessible repository. At the moment, each RIR provides a single repository. To create a complete view on the RPKI, these data sources needs to be merged.

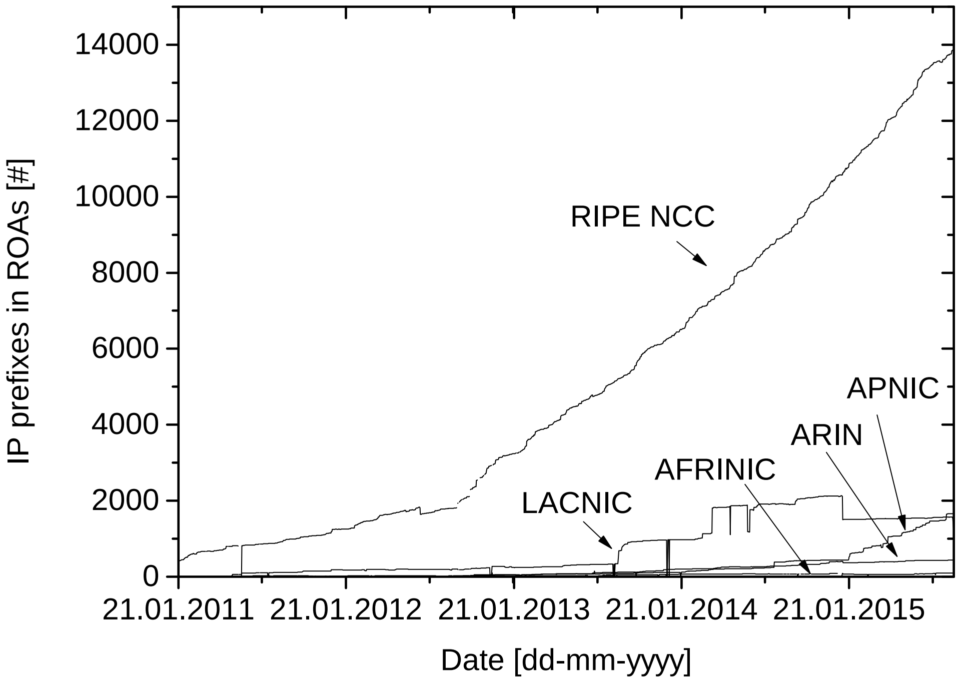

Figure 3.1: IPv4/v6 prefixes registered in the RPKI [121]

Resource certificates attest that the holder is legitimate to use the described resources. The Internet routing decouples the ownership of resources and their deployment. A prefix owner may ask an ISP to route its prefix, for example. Further on, a resource certificate does not bind the mapping of an IP prefix to an origin AS. This is implemented by Route Origin Authorization (ROA) objects.

Route Origin Authorization (ROA) A ROA object authorizes an autonomous system to originate one or multiple prefixes. It essentially contains a set of IP prefixes with an optional maximal prefix length and the authorized AS number. The ROA is signed with the private key of a RPKI certificate that cryptographically confirms the signer as the legitimate holder of these prefixes. A ROA will be created by the prefix holder. The corresponding infrastructure was deployed in January 2011 by all RIRs except ARIN. ARIN started deployment in September this year.

From a BGP perspective, the RPKI is complete when for each valid prefix/AS mapping a corresponding ROA exists. The roll-out of the RPKI marks a significant operational change for the Internet operators. It needs time until a complete RPKI is available. However, there is a continuous increase of RPKI-protected IP prefixes (cf., Figure 3.1), which may indicate a successful acceptance.

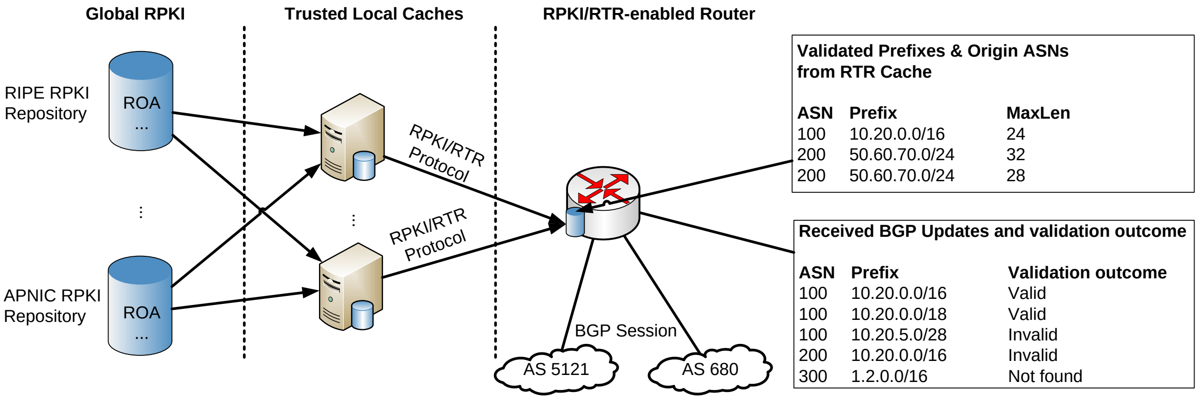

Validated Prefixes & Origin ASNs from RTR Cache

ASN	Prefix	MaxLen
100	10.20.0.0/16	24
200	50.60.70.0/24	32
200	50.60.70.0/24	28

Received BGP Updates and validation outcome

ASN	Prefix	Validation outcome
100	10.20.0.0/16	Valid
100	10.20.0.0/18	Valid
100	10.20.5.0/28	Invalid
200	10.20.0.0/16	Invalid
300	1.2.0.0/16	Not found

Figure 3.2: The RPKI/RTR architecture and the validation outcome of BGP updates

3.2.3 Securing BGP with RPKI

RPKI-enabled routers do not store ROAs themselves but only the *validated* content of these authorities. The validation of ROAs will be performed off-line by trusted cache servers, which are deployed at the network operator site. The RPKI/RTR protocol [29] defines a standard mechanism to maintain the exchange of the prefix/origin AS mappings between the cache server and routers (cf., Figure 3.2). In combination with an origin validation scheme for IP prefixes [104], a router is able to verify BGP updates it receives without suffering from cryptographic complexity. It simply compares the network layer reachability information and the origin AS of the BGP update with the data learned from the trusted cache server.

A BGP update is *valid* if at least one valid ROA exists that covers the announced prefix and matches the BGP origin AS. Whenever valid ROAs exist that cover the BGP prefix but no ROA matches the BGP origin AS, or the netmask of the BGP prefix is longer than the maximum length of the ROA, a BGP update is *invalid*. A BGP update is evaluated as *not found*, if no covering ROA exists. An example for the validation outcome is shown in Figure 3.2.

The ROA scheme implements the concept of positive attestation. This has the following practical implication for operators: As soon as a prefix owner starts to add the prefix to the RPKI, it must attest *any* legitimate origin AS for the announcement of this prefix or sub-prefix. An ISP running multiple ASes that originate the same prefix must create two ROAs. If only one AS has been certified for such a prefix, the other BGP update becomes invalid. Additionally, any BGP update that carries a prefix more specific than allowed by the maximum length in its ROA turns invalid even if the BGP origin AS conforms to the ROA origin AS. In this paper, we study the validation outcome for real BGP updates of a large AS.

It should be emphasised that the RPKI/RTR approach allows for incremental deployment and does not change the BGP specification. Thus peering neighbors need not be RPKI-aware.

3.3 Related Work

Several work revealed that a BGP router is sensitive to external events, which implicitly decrease the performance of the local system. Instabilities of the network may increase CPU and memory at routers significantly and thus may intensify the problem of reduced forwarding service [92]. Internet worms may challenge BGP routers as they introduce an avalanche of BGP updates [151], [95]. Those cases make weak points on the BGP protocol and implementations visible. They challenge the routing behaviour due to insufficient system performance. With the advent of RPKI an additional system component is deployed on BGP peers, which introduces a new vulnerability. This has not been studied so far. In this paper, we present new attacks that may harm an RPKI-enabled router.

Secure inter-domain routing has been discussed since several decades in the research community [31]. The recent deployment of the RPKI allows for real-world analysis of the mechanisms, evolvement, and operational challenges. Up until now, there is no study of the deployed RPKI. The Regional Internet Registries provide looking glasses to monitor the creation of ROAs [121] and to request the validation outcome of prefix origination [122, 93] but an in-depth analysis of the interplay between the RPKI and BGP is missing. It is worth noting that the looking glasses are based on 6 hours dumps of external BGP data and thus does not provide real-time validation but off-line testing. In this work, we conduct a long-term measurement of BGP update verification against the RPKI. This explores the effects of enabling RPKI-based origin verification from the operator perspective. After publishing our first results on the analysis of invalid BGP updates [142], measurements including a larger time period have been presented by other researchers in [78].

Support for the RPKI/RTR cache is available and already deployed [30]. The number of RPKI/RTR clients, though, is limited [28]. Cisco and Juniper already implement RPKI-based prefix origin validation in their (beta) firmware. The BGP–SRx framework [106] provides RPKI security extensions for Quagga that do not comply with the current IETF specifications, though, as the validation of BGP updates is delegated to a specific SRx server. Furthermore, this software is not intended to be used in production systems. To the best of our knowledge, rpki.net/rtr-origin [22] is the only open-source implementation of the RPKI/RTR router part in addition to our library. However, rtr-origin was explicitly developed for debugging in Python and not designed for real-time operations or router integration. In combination with rycnic [22], rtr-origin can also be used as RPKI/RTR cache server.

In this work, we present the first full-fledged C implementation of the RPKI/RTR client protocol along with a performance analysis of BGP prefix origin validation on commodity hardware. Our implementation can operate on life BGP updates. It is suitable for extending BGP daemons as well as for real-time monitoring purposes. Based on this we evaluate our threat model for RPKI-based routers and contrast it with real world measurements.

3.4 Threats on an RPKI/RTR-enabled Router

In this section we discuss threats and attack scenarios on an RPKI-enabled BGP router. From a very essential perspective, an attacker may try to manipulate the AS-to-prefix mapping or to increase basic load, which is introduced by the RTR protocol (i.e., data structure and protocol operations). We start our analysis with more theoretical threats and continue with more likely attacks.

3.4.1 Manipulating ROA Data

An attacker can try to inject malicious AS-prefix pairs, which is on global scale non-trivial. The integrity of the data within the distributed RPKI is ensured by public key cryptography. Trusted cache servers provide this data via an insecure or secure communication channel. Although not recommended, an unprotected TCP connection is allowed by the RPKI/RTR protocol due to the variety of heterogeneous vendor implementations. Those sessions might be captured and used to deliver incorrect data to the router. Threat actions for TCP hijacks are well-known. However, they are most likely *not* the appropriate option to attack the local router. Basically, the corresponding threat models require IP spoofing, including recent work on this topic [116], [117]. We assume that ISPs using unprotected TCP between cache and router will carefully design this network configuration (e.g., activating access control lists, direct interconnects). All other ISPs that are not constrained by plain TCP transport may use SSH, TCP-AO etc., which do not allow to hijack TCP sessions.

In contrast to the injection of incorrect data on the way from the cache server to the router, an attacker may change locally stored information at the router. This requires system access by breaking passwords or taking advantage of system exploits, for example. Again, this attack can be considered as very unlikely for two reasons. First, remote access to backbone routers is typically subject to specific security conditions (e.g., out-of-band management). Second, if an attacker gains control over the router, he is able to change more relevant configuration. Overall, this threat action is not particularly related to the RPKI/RTR protocol but affects the complete router system.

3.4.2 Denial of Service

Any operation of the prefix origin validation mechanism is related to the prefix validation table, which stores trustable AS-to-prefix mappings at the router. Maintenance (add, delete, replace) as well as lookup operations (find) need a traversal of the data structure. An attacker may harm the BGP router by actions that cause an extraordinary load. We distinguish between *simple overload attacks* and *complexity attacks*.

Simple overload attacks

A worst-case action is a reloading of the complete ROA data. An attacker may reset the TCP connection to the cache server [66] or overload the cache itself by a longer DDoS [132]. Both result in connection loss. In case of a backup cache server, the router will establish a new connection and sends a reset query, which triggers the new cache to deliver all data records. Note that this attack is effective also for protected TCP connections such as SSH. The robustness of the system depends on the processing time to load a complete table.

This delay is also important in case of a reboot: A router can start verifying BGP updates not before the data is included in the prefix validation table. Depending on the implementation of the BGP daemon, this means that the router will pause processing BGP updates or start propagating announcements without validation. The first stops the forwarding service for an additional time period. The latter would allow an adversary to exploit the time gap by injecting malicious prefixes. This is a serious problem in particular in an incremental deployment environment, where neighboring peers do not perform prefix origin validation but rely on the upstream.

In addition to an enforcement of a table reload, an attacker may produce a storm of BGP updates. All of the updates need to be verified against the local data. This implicates a high number of lookups. The robustness of the local system is coupled to an efficient lookup strategy.

Complexity attacks

A complexity attack [47], [24] aims at algorithmic operations that lead to worst-case performance. A common example is the creation of collisions in a hash-based data structure. In general for data structures, the idea is either to viciously adapt insert/delete/lookup requests with respect to the stored data set, or to create malicious entries that result in bad performance for the typical set of requests. Obviously, an adversary benefits from particular knowledge about the data in use.

In our case, an attacker will (a) add dedicated ROAs to the RPKI or (b) initiate specific BGP updates that harm the routers, or (c) both. Note that BGP data as well as ROAs are publicly available provided by route monitors or looking glasses, and there are no read restrictions to the RPKI repository. This information can be used to fine-tune the attack. We should also note that a BGP peer can submit an arbitrary prefix update. This might lead to an invalid BGP announcement but the update needs to be processed, though. In contrast to this, putting ROAs successfully into the RPKI is restricted to the valid ownership of the prefix. Nevertheless, the allocation of IP(v6) prefixes is not a challenge. Then, the owner of an IP prefix is allowed to assign any AS number to this prefix. In the worst case, the attacker assigns all AS numbers currently available in the global routing system.

We summarize our observations in Table 3.1.

Objective	Attack	Chance	Our Focus
Infiltrate incorrect AS-prefix entry into prefix validation table	Hijack of cache-router session	−	✘
	System intrusion	−	✘
Complete reload of prefix validation table	Reset of cache-router session	+	✓
	DDoS on cache server	+	✓
Long-term increase of CPU load	BGP update storm	+	✓
	Malicious mixture of ROA data	○	✓

Table 3.1: Overview of threats on an RPKI-enabled router (+ likely, ○ neutral, − unlikely)

3.4.3 Attack Model

In the remainder of this paper, we assume the following attacker. The attacker tries to disturb the forwarding service of the RPKI-enabled router and is able to (1) disconnect router and cache server, (2) send a high amount of BGP updates, (3) add cryptographically valid ROAs to the RPKI repository. The attacker is the legitimate owner of the prefix(es) he adds to the RPKI.

3.5 RTRlib: A Library for RPKI Router Support

To extend routing by an RPKI-based prefix origin verification, the RPKI/RTR protocol needs to be implemented on routers. We argue that a realistic but open implementation is necessary to conduct an in-depth analysis of the vulnerability space. Closed implementations provided by common router vendors do not allow for fine-grained system level access, which is required to explore threat consequences.

In this section, we present the first full-fledged open-source implementation that is suitable for testing purposes as well as production use. We assembled the required functions as an external independent library, which simplifies code reuse. Existing BGP daemons can be extended by simply integrating the library or parts of it. The same code base may also be used to build tools for researchers or ISPs (e.g., to monitor the RPKI).

3.5.1 Design

Our implementation of the RPKI/RTR protocol follows the subsequent design goals:

Broad system integration The library shall run on different system environments and thus minimize dependencies on specific operating system calls and third party tools.

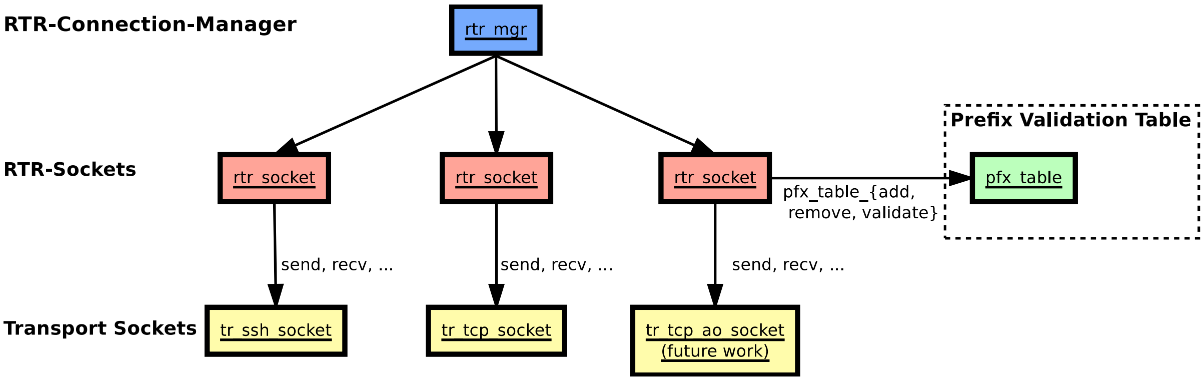

Figure 3.3: Software architecture of the RPKI client lib

We implement the library in C because several BGP daemons are written in C (e.g., BIRD, Quagga) or C++ (e.g., XORP). In contrast to C++, C-functions can be invoked or adopted into other C/C++ programs without modifications. To facilitate the smooth migration to a variety of operating systems, our library is based on POSIX interfaces.

Interoperability The implementation shall be able to exchange data with existing and upcoming RTR cache servers.

The presented RPKI/RTR library implements the latest version of the protocol specification. We performed interop tests with the RPKI/RTR cache servers available [30], which helped to reveal errors on both sites.

Extensibility The library shall be easy to modify for supporting upcoming protocol changes and extensions as well as specific user demands.

To support this requirement, the library consists of separate components with high cohesion and low coupling. Sufficient abstraction allows to exchange the modules with low complexity.

Efficiency BGP routers are confronted with high BGP update rates and must currently store more than 400,000 IP prefixes. With respect to memory and CPU consumptions, the RPKI/RTR implementation shall be prepared to handle these data even though not all prefixes are part of the RPKI at the moment. In addition, the implementation should minimize the internal overhead.

We designed our library to scale very well with current and upcoming BGP data. It does not require specific hardware, but runs on commodity devices and thus supports an easy migration to RPKI origin validation without introducing costs.

3.5.2 Architecture

Our implementation follows a flexible design. The software architecture includes different layers to simplify the extension or an exchange of individual parts as shown in Figure 3.3. The lowest layer of the architecture is built by the *transport sockets*. They allow for the implementation of different transport channels that provide a common interface to exchange PDUs with the cache (i.e., the RPKI/RTR server). The current version of the library supports unprotected TCP and SSH.

On top of the transport layer, the *RTR socket* uses a transport socket for RTR-specific data exchange with the RPKI/RTR server. The RTR socket implements the RPKI/RTR protocol, i.e., fetches validation records and stores them in a prefix table data structure.

The *prefix validation table* stores validated prefix origin data. This abstract data structure provides a common interface to add and delete entries as well as to verify a specific prefix. Its implementation is crucial as the data structure stores all prefixes received from the cache servers (i.e., low memory overhead required) and is responsible to perform prefix lookup for the BGP updates (i.e., find validated IP prefixes very fast). Our library implements a Longest Prefix First Search Tree (LPFST) [154], but can be extended to other data structures. In contrast to common data structures for IP prefix lookup such as Tries or Patricia, the LPFST needs fewer memory access and exhibits lower memory overhead [154].

Internally, the library uses two separate prefix validation tables, one for IPv4 records and one for IPv6 records. This makes tree operations (insert, delete, find) more efficient as the height per tree is lower in contrast to a combined IPv4/v6 tree. The appropriate prefix validation table will be chosen according to the IP version.

On top of the modular architecture, the *RTR connection manager* maintains the connection to multiple RTR servers. This includes failover mechanisms. It represents the main interface for users of the library.

3.5.3 Performance Evaluation

In this section, we analyse the runtime performance and the scaling behaviour of our library at a system level. We deploy our experiments on a dual-core AMD Opteron 280 (2.4 GHz), equipped with 8 GB RAM. The underlying operating system is Linux Ubuntu with 2.6.32-33 kernel. Measurements and analysis of RPKI validation results from real-world data will follow in Section 3.6.

We explicitly note that any comparison with other implementations of the RPKI/RTR router part, which are currently available, would be unfair. The Python implementation rpki.net was not designed for real-time purposes. This is in contrast to RTRlib. On the other hand, professional router implementations that support RPKI such as Cisco or Juniper do not allow for RPKI-specific system measurements; but our microbenchmarks require a very fine-grained analysis.

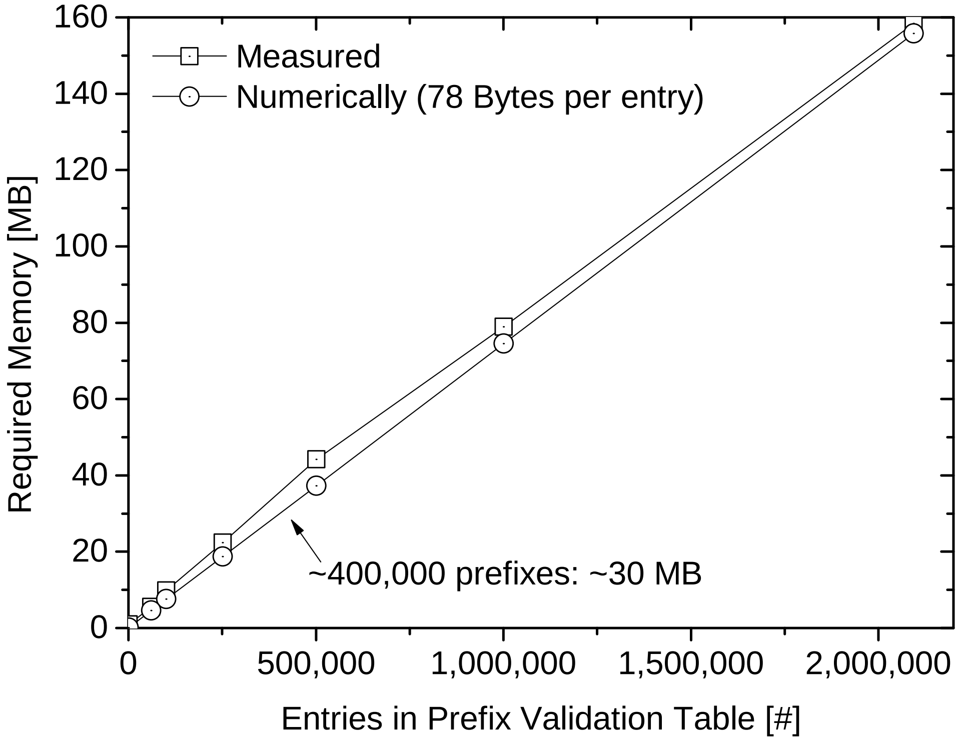

Figure 3.4: Memory consumption of the RPKI/RTR library

Memory Consumption

The memory consumption of the library mainly depends on the number of prefixes inserted into the prefix validation table. Considering a 64 bit architecture with 8 bytes per pointer, a single record within the prefix validation table consumes 78 bytes in our implementation of the longest prefix first search tree. Padding bytes, which maybe be inserted by the compiler, are omitted in this calculation. Note that a common BGP route entry requires between 100 and 200 bytes [52].

To measure the memory required on a real system, we added randomly generated prefixes to the prefix validation table. Figure 3.4 displays the scaling behaviour for different table sizes. The overall memory consumption scales linearly as expected. ROAs for all ≈ 400,000 active IP prefixes included in current BGP routing tables would result in additional ≈ 30 MB of RAM for an RPKI/RTR-enabled router. Cisco suggests to equip their devices with 512 MB of RAM for storing a complete global BGP routing table [41]. Thus a full RPKI validation table would lead to a 5% increase of RAM – a relatively small overhead. In particular, BGP services implemented on commodity hardware, where several gigabytes of RAM are common, should not suffer from RPKI requirements.

CPU Consumption

The processing overhead of RPKI/RTR on the router is dominated by the complexity that results from update and lookup operations on the data structure holding the valid ROA information. Update operations on the prefix validation table are triggered by new, modified, or deleted ROAs, whereas lookups follow BGP updates. In general, it is reasonable to assume that ROA changes occur on a significantly larger time scale than the frequency of BGP updates. Our measurement study confirms this and shows that the CPU load correlates directly with the BGP update rate (cf., Section 3.6.2). To further explore the scaling behaviour of our RPKI/RTR router implementation, we thus focus on (a) the delay to load a new, complete ROA data set and (b) the processing of BGP updates in dependence of available ROA data, i.e., the evolving RPKI deployment state.

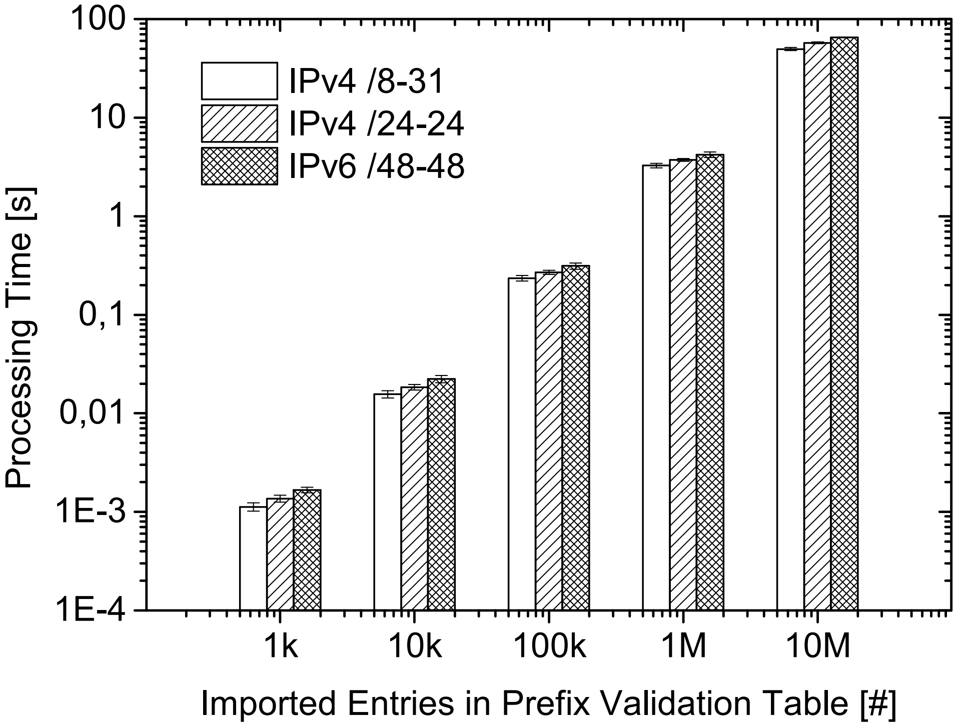

Figure 3.5: Delay to load a bulk of ROA data into the BGP router for different mask lengths. Standard deviation is included but negligible.

Load delay In the first step, we quantify the amount of time required to load a fresh ROA data set to the router (i.e., creation of the prefix validation table). Importing the *all* valid prefix-to-AS mappings is necessary after a reboot or changing cache servers. With the first connection to the RPKI/RTR cache server, the complete set of valid ROA data must be (a) transferred to the router and (b) added to the prefix tree. Performance of the first part depends on the network topology, which is reflected by the distance of router from cache server. The second part is only related to the local system. To quantify this, our measurement node reads a varying number of ROA prefixes from a file and inserts them into an empty prefix table. In addition to the number of entries, we also vary the ROA parameters minimal and maximal prefix length. Figure 3.5 shows the average processing overhead and the corresponding standard deviation. Overall, a polynomial increase is visible with the number of entries, which does not depend on the network mask. For 1 million entries, the operation consumes ≈ 4 seconds. Even a very large set of data of about 10 million entries can be loaded in less than one minute. Note that the time complexity is at least $O(n)$ independent of the data structure in use, where n is the number of entries. Additional complexity arises due to the rebalancing of the tree to increase further lookup speed. The performance is still moderate. Nevertheless, an attacker may acquire a larger IP address space (e.g., a common IPv6 /48 range) and create dedicated host entries (in our example the worst-case is 2^{80} entries). We argue that a maintainer of the RPKI repository should take this into account and implement accordingly 'intrusion' prevention approaches.

Validation complexity The RPKI deployment state is measured by the ratio of valid, or invalid, versus not found IP prefixes. This holds for BGP routing entries but also with respect to the performance overhead. For example, even a prefix that is not part of the RPKI requires a lookup within the prefix validation table. In the worst case, the complete height of the prefix tree must be traversed. In the following, we observe the CPU load for a varying mixture of validation outcome. Note that as the *effects* on the CPU characterize the dependencies on the RPKI deployment state, the analysis does not stress specific values but the overall scaling behaviour. For this reason, we do not consider the CPU load but the amount of executed CPU operations (i.e., ticks or jiffies).

The performance of tree data structures correlates with the tree shape, which is influenced by the inserted data. For our measurements we could derive a future ROA prefix distribution from the current Internet backbone routing table. However, the Internet itself is a continuously evolving structure, which makes predictions quite hard [125]. In addition to our previous argument we want to keep the measurement setup simple and robust. For discussions of the CPU load based on a real ROA data set and live BGP streams we refer to Section 3.6.2.

In proceeding this way, we evaluate the CPU load by randomly generating 100,000 ROA IPv4 prefixes with a fixed minimum and maximum /24-netmask. The AS number of the authorized origin AS is randomly generated, as well. We consider ratios per validation state of 0%, 25%, 50%, 75%, and 100%. For each combination of all possible validation outcomes (e.g., 25% valid, 50% invalid, and 25% not found), we create a total of 2,000 BGP prefixes that match the required states. We emulate a very high BGP update rate of 2,000 verifications per second. Each measurement is sampled with the same prefix data until it is converged. This required approximately 10,000 runs. We average the results over all samples of the same settings.

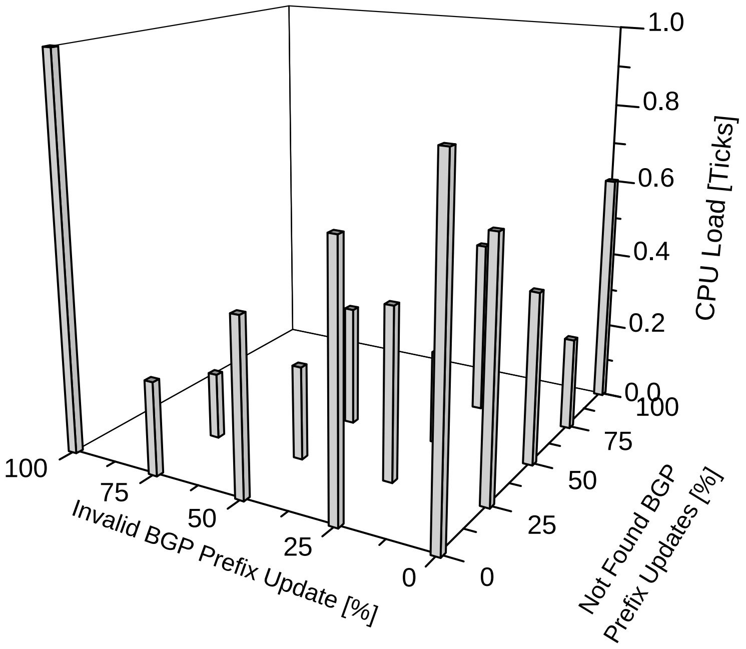

Figure 3.6: CPU overhead for varying validation outcomes

Figure 3.6 visualizes the CPU load for a varying ratio of invalid and not found prefix announcements. The number of valid announcements follows implicitly. The CPU load ranges between 1.0% and 0.17% . It is worth noting that we measure on a very fine-grained scale. A strong dependency on the validation mixture is clearly visible. In contrast to a not found prefix, an invalid prefix update need to be searched completely. In addition, in case of multiple ROAs per prefix all ROA data need to be checked. Overall, although the dependency on the deployment state is visible, the CPU load does not threaten the system. Verifying this observation on different router platforms despite limited system measurement capabilities will be part of our future work. In the subsequent section, we analyze the currently visible RPKI-based prefix origin validation based on a life BGP stream.

3.6 RPKI in the Wild

3.6.1 Measurement Setup

In this section, we analyze the behaviour of an RPKI/RTR-enabled router based on real BGP Updates. We seek insight into what an ISP experiences if operators enable RPKI/RTR at their routers.

The measurement setup consists of three parts (cf., Figure 3.7): Receiving a real time BGP Update stream, obtaining currently deployed and validated ROAs from an RTR cache server,

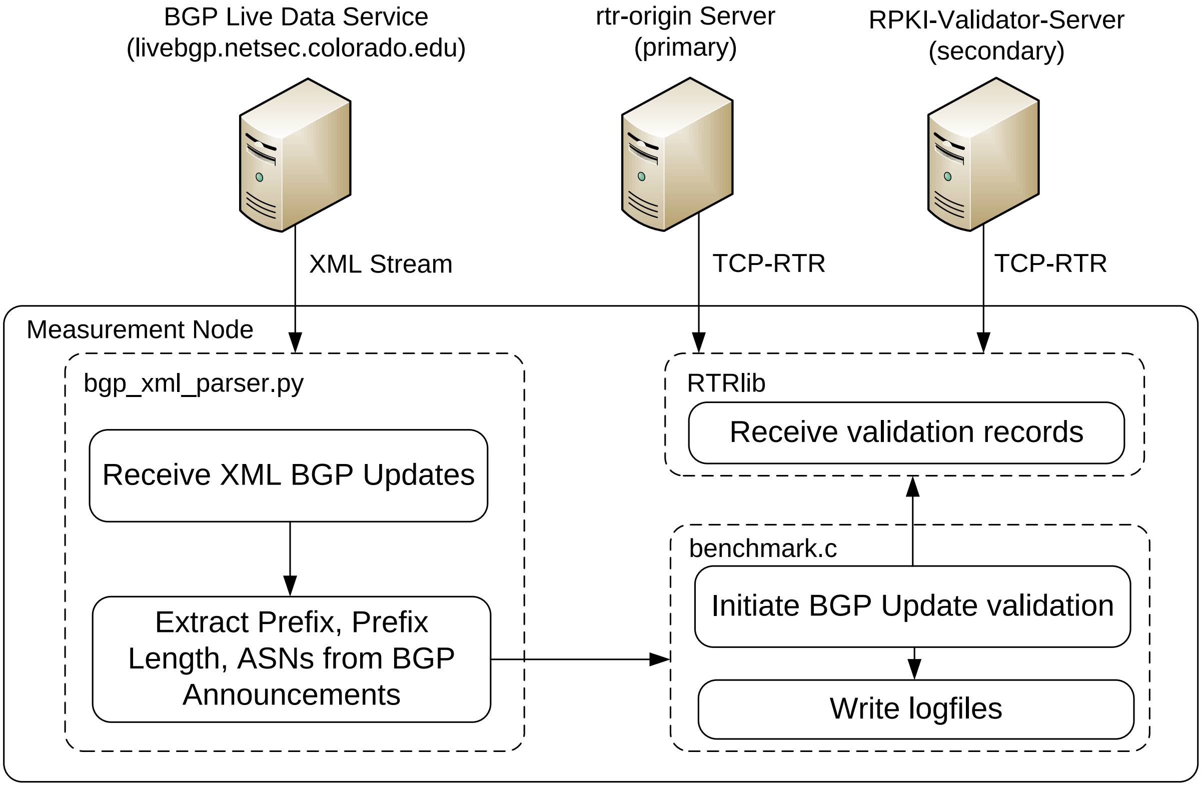

Figure 3.7: Live measurement (January 2012) setup

and verifying the origin ASes for prefixes included in the BGP Updates against the valid ROA data set.

Our measurement node corresponds to a BGP router that runs RPKI/RTR on commodity hardware, similar to our experiments (see Section 3.5.3). To receive live BGP data, we use BGPmon [155, 102], which provides real-time BGP updates in XML format. BGPmon maintains nine direct peerings (e.g., SWISSCOM and Tiscali) as well as indirect peerings via three collectors of the RouteViews project. The indirect peering includes routing updates from more than 100 peers such as AOL, Hurricane Electric, and AT&T. Receiving full BGP Updates from *multiple* peers grants two advantages: First, it helps us to see BGP Updates from multiple vantage points within the Internet and thus to reduce the immanent problem of an incomplete view [109]. Second, the amount of peering relations allows us to experience live BGP Update rates similar to those seen at a larger ISP.

All information necessary for the prefix origin validation will be extracted from the BGP update stream and passed on to a benchmark program. The benchmark program measures the CPU load of the RPKI/RTR operations and logs the validation outcome of each prefix announcement. For each update, we also record the IP prefix, mask length, and AS path, as well as the validated ROA data if available. This will allow for a detailed analysis of the verification decision in our subsequent evaluation. It is worth noting that ROA data do not exist for prefix updates that will be evaluated as “not found”.

We use our library presented in Section 3.5 as implementation of the RPKI/RTR router part. The benchmark program creates an RTR connection manager, which establishes connections to two RPKI/RTR cache servers using plain TCP. We deploy rcynic+rtr-origin [22] as primary cache server and the RIPE NCC RPKI Validator [122] as secondary (backup) cache server. Each cache server considers the trust anchors AfriNIC, APNIC, ARIN, LACNIC, and RIPE NCC. This setup reflects a typical deployment case.

For the following detailed evaluation, we selected two months, January and May. This choice was made for two reasons. At first, we want to follow the evolutionary process and compare characteristic phenomena. At second, these months account well for two extremes in the speed of roll-out. While the deployment remains fairly stable in January, a rapid increase of ROAs is visible in May (see Figure 3.1).

3.6.2 Evaluation

In total we received more than 420 million prefix updates in January and May, which corresponds to $\approx$ 680,000 prefix updates on average per day.

Overhead on Routers

The CPU load corresponds to the number of prefix validations (i.e., the BGP update rate). Figure 3.8(a) visualizes both measurements per minute for January 5, 2012. All other days show the same qualitative behaviour and we restricted the plot only for visibility reasons. During the measurement period, we observed a maximum of 92,549 prefix announcements per minute and a maximum CPU load of 0.41%. The average CPU load per day was 0.02% with a standard deviation of 0.04%. To analyze the correlation between the time series in more detail, we plot both measures depending on each other in Figure 3.8(b). A linear increase is clearly visible, indicating that a low amount of BGP updates induces low CPU overhead, whereas CPU overhead increases moderately with higher update rates.

Overall, the overhead that results from RPKI origin validation is negligible. To perform *basic* BGP operations typical CPU processing load in a tier-1 network may range between 25% and 60% on older but specialized hardware [10]. Such relative load is not considered as a threat [10]. In contrast to basic BGP processes, origin validation is two orders of magnitude lower on recent commodity hardware.

The results also nicely illustrate that our benchmark program works correctly. No peak deviations between CPU load and update rate are visible, which could indicate abnormal behaviour.

Valid vs. Invalid Prefix Updates

Most of the announced IP prefixes are stable in the sense that they remain visible over a longer period of time. They are advertised continuously. To prevent prefix hijacking, each RPKI-enabled BGP router must process any single prefix advertisement (in our case 680,000 prefixes on average per day) and verify the origin AS against the currently valid ROA data. However, not all valid prefixes initiate a change in the Routing Information Base (RIB). In particular, prefix updates that are already part of the RIB need not to be evaluated again if received in a short period of time. Caching might be used to optimize access time, for example. From this perspective we quantify the difference between the information a router *sees* in the updates and the new information a router *learns*.

Figure 3.9(a) and Figure 3.9(b) show the total number of valid and invalid prefixes aggregated per day for January and May, respectively. On average, the number of successful prefix origin validations is larger than the invalid prefixes. There are, however, days where invalids prefix announcements exceed valids. In contrast to January, the amount of valids increased and the number of invalids decreased in May. This is due to an enhanced training of the ISPs by the RIRs, who also notified explicitly owners of (supposedly) misconfigured ROAs.

It is worth noting that this quantity directly reflects the BGP update rate but does not reveal the number of newly learned validation outcomes. A small set of the same invalids may be announced quite frequently and thus outranges valid announcements.

To correct for effects of BGP update rates, we quantify the validation outcome of each prefix only once per day for each distinct state. Practically, this means that multiple valid (or invalid) announcements of each prefix will be counted as a single update independent of the time they occur during the day. This allows us to single out what a router newly learns from the RPKI perspective. Figures 3.9(c),(d) plot the results over the measurement period. With

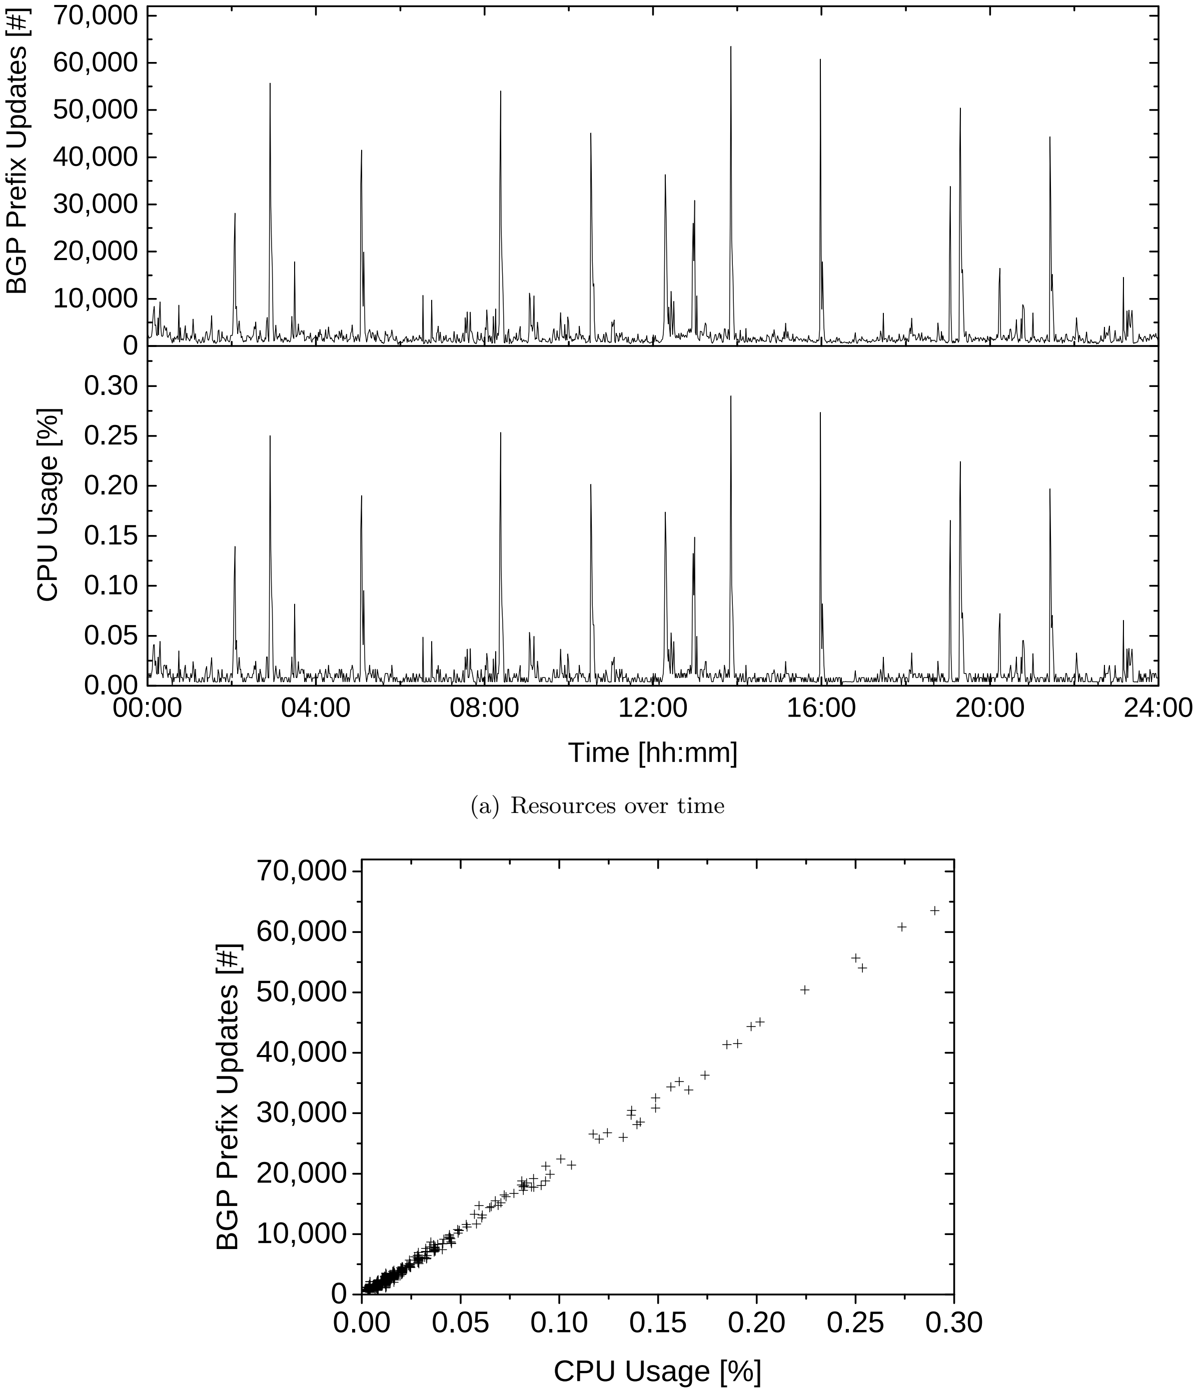

(a) Resources over time

(b) Correlation of CPU load and prefix updates

Figure 3.8: Characteristic CPU load for prefix validation and the number of received prefixes for January 5, 2012

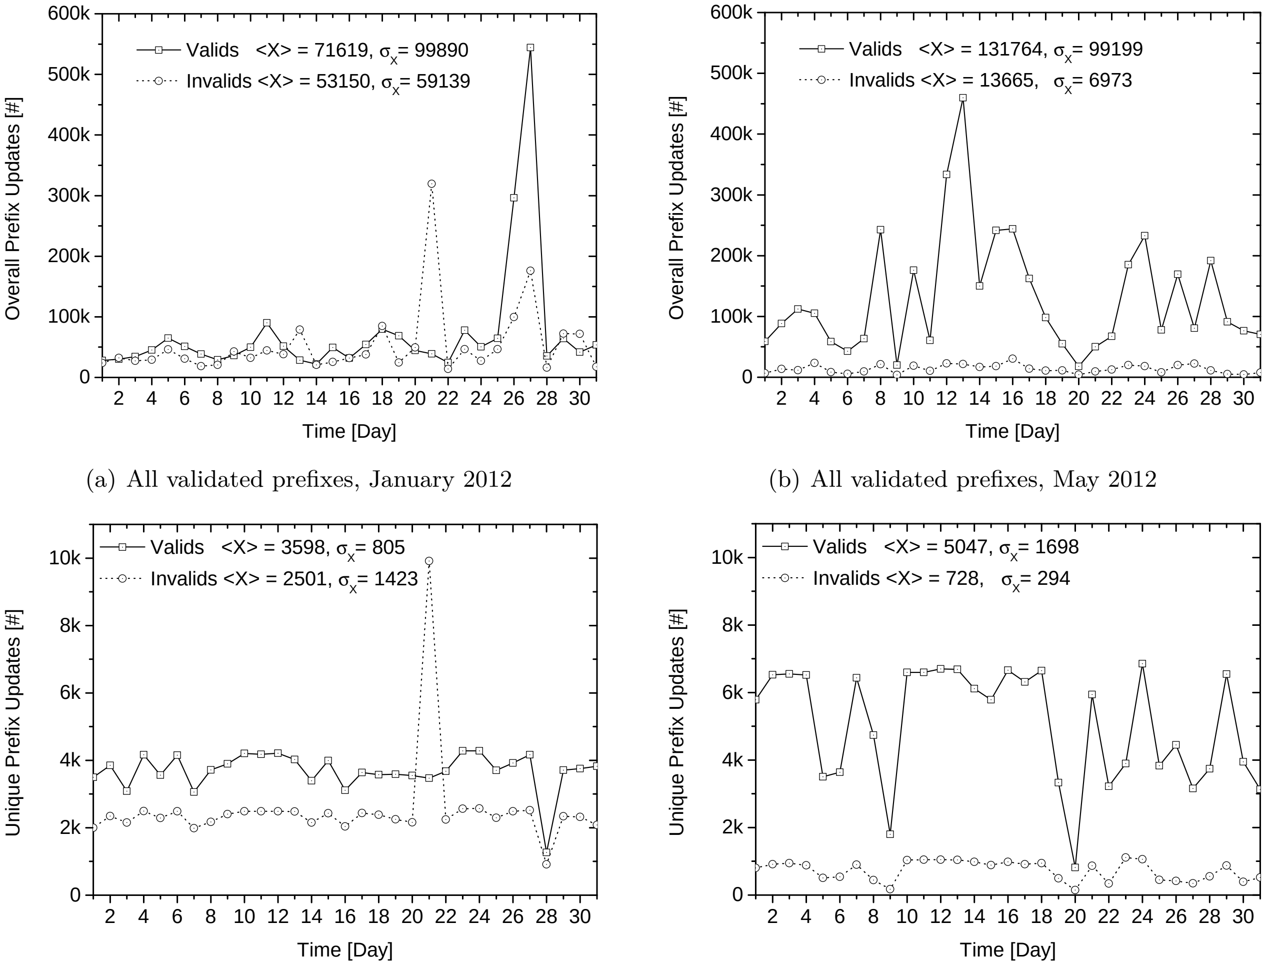

(a) All validated prefixes, January 2012

(b) All validated prefixes, May 2012

(c) Unification of the same prefixes, January 2012

(d) Unification of the same prefixes, May 2012

Figure 3.9: Prefix origin validation outcome (valids and invalids) per day with average ($\langle X \rangle$) and standard deviation (σ_X)

the exception of January 21, the router experiences significantly more unique valid prefixes than invalid announcements. January 21 illustrates a typical incident in which a single ROA invalidates a large set of (customer) prefixes; this we explain in detail later in this chapter analyzing the problem of customer legitimation based on the ATT case. This metric clearly shows when the RPKI data set changes or new prefixes have been injected via BGP. Combining the results of Figures 3.9(a),(b) and Figures 3.9(c),(d), it indicates that most validation storms are due to a high frequency of BGP updates and not based on changes of the prefix/RPKI data.

3.6.3 Invalid Prefix Originations in Detail

We inspect the unsuccessful verifications of prefix origins in more detail to understand the reasons behind incorrect BGP updates. On the one hand, this may be of practical interest as it helps to identify pitfalls that might happen during ROA configuration. On the other hand, it is necessary to model RPKI ROA composition in a more precise way at least from the current deployment state.

General Observations

A BGP prefix update covered by the RPKI may turn invalid due to two reasons: (a) no ROA origin AS matches the BGP origin AS, (b) the BGP prefix update is more specific than attested by the ROA (max length violation). It is worth noting that a max length violation can only be meaningful for prefixes that have a correct origin AS within at least one matching ROA. If the ROA origin does not match, one cannot conclude on an incorrect prefix length as this is then ambiguous.

Figure 3.10(a) plots the ratio of both causes for all unified, invalid prefixes per day in January 2012. Except on January 21 between 80% and 90% of the invalid prefixes are due to an incorrect origin AS. The spike on January 21 results from a de-aggregation failure, which we explain in detail later in this section.

If all ROAs have been configured correctly, the validation outcome indicates a prefix hijack because none of the authorised ASes originates the IP prefix. However, the analyzed state of deployment of the RPKI infrastructure also includes incomplete configurations of the ROAs. During our measurement in January, only the minority of the invalid prefixes result from a more specific announcement.

In case of invalid origin ASes, further analysis reveals that in more than 90% the ROA origin AS is only one hop away from the BGP origin AS.[1] This is surprising and indicates misconfiguration of the ROAs as the legitimate origin AS would most likely not further distribute the incorrect prefix origination instead of its own.

[1] Note that we resolve path prepending, i.e., we collapse subsequent AS hops with the same AS number.

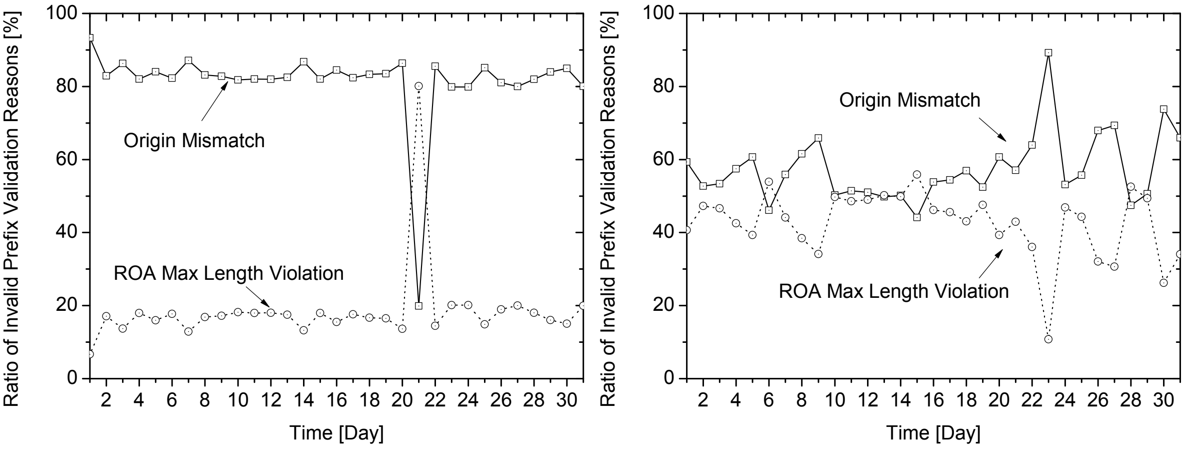

(a) January 2012

(b) May 2012

Figure 3.10: Relative number of reasons for an invalid prefix origin verification per day

The following scenario illustrates a typical misconfiguration that may be the reason for the observations described above: A large ISP owns a super-block (e.g., 8.0.0.0/8) and delegates sub-prefixes to customers (e.g., 8.10.20.0/24 to $customer_A$ and 8.100.200.0/24 to $customer_B$). Each customer announces his sub-prefix using his own AS and chooses the large ISP as upstream provider (i.e., the ISP is only one hop away from the origin AS). If the ISP creates a ROA for the super-block *prior to* the customers create corresponding ROAs, all BGP updates of the customers will be interpreted as invalid. According to discussions with the RIRs and ISPs this type of misconfiguration is the main reason for the previous observation of invalid updates with a ROA origin on path. The incorrect ROA configurations were partly fixed over the time (cf., Figure 3.9), where origin mismatches in general were reduced to 57% on average for May as shown in Figures 3.10(a) and 3.10(b).

Analyzing the origin mismatches in more detail, we also found that the amount of incidents in which a ROA origin is one hop away from the BGP origin decreased on average from 90% to 52%.

Another cause for an origin mismatch based on misconfiguration may occur in a large company that maintains multiple ASes [32] that are administrated by different departments (called sibling ASes). If the departments do not coordinate the authorization of ASes in parallel, the unauthorized ASes will lead to invalid BGP prefix updates. We found those cases as well. Note that an automatic identification of siblings is intricate because only heuristics can be applied. As siblings cannot be detected reliably, we omit further quantifications as this would introduce additional error effects.

In the subsequent sections, we analyze invalid prefix announcements with respect to specific incidents.

(De-)Aggregation

On January 21, 2012 an unusually high amount of prefixes were announced incorrectly. A detailed inspection reveals that 77% of the 9,917 unified invalid prefix announcements belong to AS 12322 (ProXad). The invalidation occurs due to an incorrect max prefix length within the ROA. ProXad owns several large IP blocks. They have corresponding ROAs with a small max length (e.g., 78.192.0.0/10-10 or 82.224.0.0/11-12) and they usually announce each block as aggregated prefix (e.g., 78.192.0.0/10). However, on January 21 for 30 minutes the aggregation dissolved and most of the sub-prefixes were advertised with a network mask between /22 and /24. Consequently, these BGP updates are more specific than allowed by the corresponding ROAs max length.

There may be multiple reasons for such an incident that appeared not to be an origin prefix hijack since the origin AS was still correctly announced as AS 12322 in the invalid updates. An operator should always configure ROAs such that they cover the complete address space. The ROA should reflect the most fine-grained prefix structure an operator is prepared to announce. Setting the maximum length within the ROA properly helps to overcome issues caused by BGP dynamics.

Missing Customer Legitimation

During the testing phase of our measurement setup, which did not span a continuous time interval over several days and thus has been excluded in our previous discussions, we found on December 13, 2011 that approximately 50% of the invalid prefix announcements concerned a single (super-)block. Only one ROA was issued that covered the entire address space of 12.0.0.0/8:

ASN	Prefix	Maximum	Trust
		Length	Anchor
7018	12.0.0.0/8	9	ARIN Test Lab

Only AS 7018 (ATT-Internet4) has been authorised to announce 12.0.0.0/8 and corresponding sub-prefixes. However, more than 2,400 sub-prefixes seem to belong to different ASes. This was not a problem since the ROA was part of a test trust anchor. However, it illustrates two things: First, the RTR Server should select the right trust anchors, otherwise the router will be fed with cryptographically valid but doubtful data. Second, any creator of a ROA must ensure that legitimate sub-allocations have been authorised as well. Any ISP that delegates IP prefixes to customers must ensure that customers have been authorised and started to announce their prefixes before the ISP creates a ROA for the covering prefix. In this case, a ROA might also be created for

ASN	Prefix
27487 (FHLBNY-AS -FEDERAL HOME LOAN BANK OF NY)	12.0.19.0/24

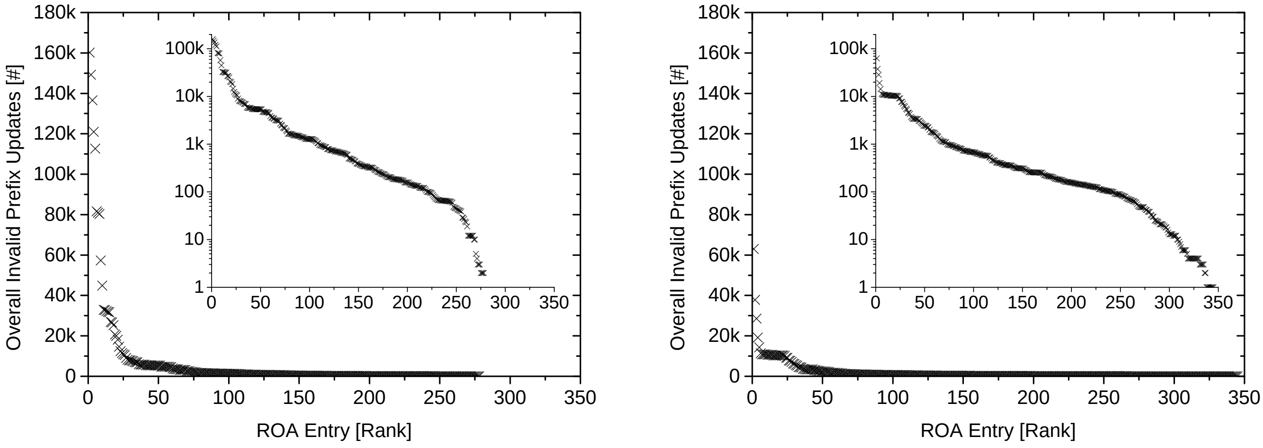

(a) January 2012

(b) May 2012

Figure 3.11: ROA entries ranked according to invalid prefix updates (logarithmic details shown in inserts)

The same should be applied when the sub-prefix will be announced by another AS of the company, e.g.:

```
ASN                                Prefix
2386 (INS-AS - AT&T Data           12.1.216.0/24
     Communications Services)
```

This measurement has been performed at December 13, 2011 while the RPKI data set was still under development towards a complete, correct view. The AT&T case is just an example and the operators continuously refine their data. In fact as reported by RIPE, AT&T has removed the ROA for AS 7018 on December 21, 2011.

3.6.4 Lessons Learned

In our long-term measurement study that focuses on two characteristic month for the current state of deployment, we found two main observations:

1. The additional load at routers by enabling RPKI is negligible. From this point of view RPKI is ready to use.

2. The quality of the current ROA data needs further improvement. Fixing a few hundred ROAs already helps to reduce the number of invalid prefix updates by several orders of magnitude.

The **CPU overhead** introduced by RPKI follows mainly the BGP update rate. Even with a wider deployment (i.e., additional prefixes) the RPKI will exhibit small footprint that is in the order of magnitude observed during our measurements. An increased number of IP prefixes in the RPKI will lead to a higher prefix tree at routers requiring a longer traversing time within the data structure. However, this data structure can efficiently be implemented and makes common operations (add, delete, and lookup) insignificant. It is worth noting that the RPKI data will be more stable in the far future and thus reduces interaction with the RTR cache server. On the other hand, the overhead introduced by the validation of BGP updates is independent of the ROA data set (valid versus invalid versus not found). Thus, the measured CPU load on our commodity hardware will not change significantly in the future deployment state.

The current (low) degree of the **ROA data quality** is surprising and at the moment a severe problem for enabling routing policies based on RPKI. To weight ROAs that are responsible for invalids, we sort each covering ROA based on the number of incorrect prefix updates it triggers (cf., Figure 3.11). In January and May 2012, most of the invalid announcements are distributed over the majority of ROA entries and only ≈ 30% – 40% are related to the top five ROA entries in each case. This not only illustrates that the validation outcome depends on a larger group of ROAs but also that cleaning data is bound to multiple ISPs. However, aligning the top-ten responsible ROAs with the currently announced BGP updates will substantially help to reduce the number of invalid updates.

The reasons for misconfigurations range between little understanding of the ROA concept (e.g., max length problem) and insufficient knowledge of the internal network structure (e.g., missing siblings). We thus highly emphasize two recommendations. First, operators need in-depth training, which is provided by the RIRs and showed already beneficial impact; *and* second, operators should perform an incremental deployment to overview the effects of ROAs, starting with specific subsets (i.e., long netmask). From a practical point of view the latter requires real-time monitoring and evaluation of BGP update validation to identify failures, which we hope to contribute with both our RPKI/RTR implementation and the identification of common pitfalls (see Figure 3.12).

3.7 Conclusion and Outlook

This chapter presented a first practical exploration of the Resource Public Key Infrastructure (RPKI) recently released by the IETF. In an evolutionary approach, RPKI allows for authenticating prefix-to-AS mappings in BGP route advertisements without altering the Internet backbone routing.

We introduced a new threat model for routers that perform prefix origin validation. To analyze the vulnerability space under real conditions, we presented the first full-fledged RPKI/RTR router implementation that is available for public download. Our thorough performance analysis revealed its readiness not only for research and monitoring, but also for production-type services. Nevertheless, we also discussed that an RPKI BGP peer is open for (complexity) attacks, which aim to degrade the system performance. Introducing a specific mixture of ROAs

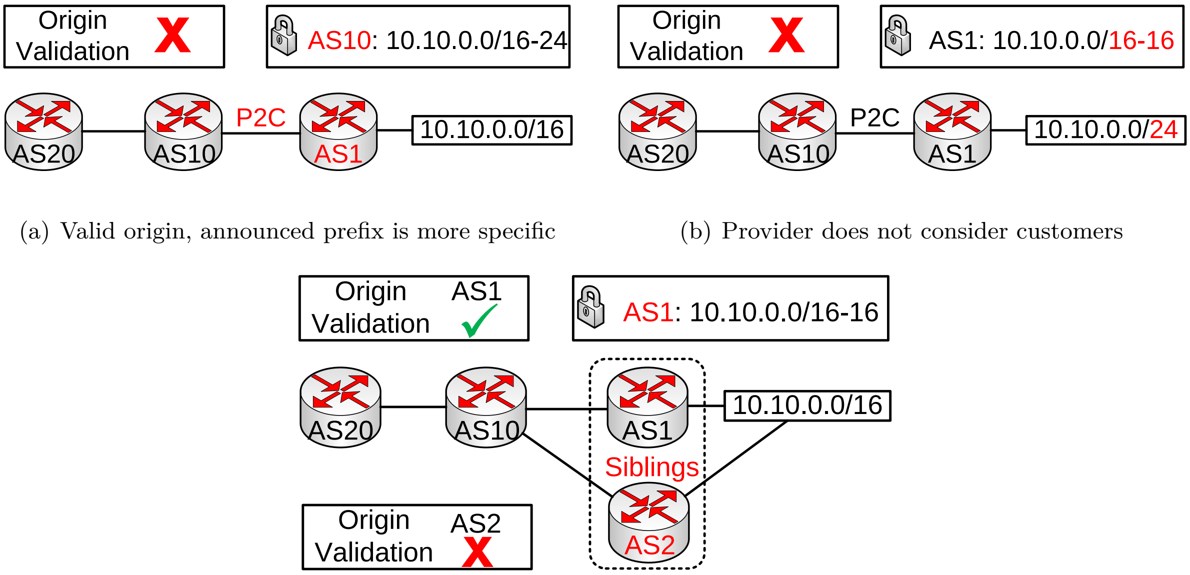

(a) Valid origin, announced prefix is more specific

(b) Provider does not consider customers

(c) Additional ASes of a company not authorized

Figure 3.12: Illustrating potential RPKI/BGP misconfiguration

may lead to visible increase of system load. To our surprise the current RPKI deployment needs to overcome completely different problems.

The second part of our work was dedicated to a long-term measurement and analysis of real-world Route Origin Authorization (ROA) management and the validation of prefixes in real-time BGP streams. This is important to contrast our threat analysis with the current deployment. While monitoring an emerging deployment of operators, we could identify a number of pitfalls and common misconfigurations that bear the potential to hinder a successful rollout of RPKI. So far explicit attacks on the RPKI are not visible. However, we believe that the amount of attacks will most likely change with an increased number of ROAs and RPKI-enabled routers.

Currently we work on establishing an online monitoring service that displays the status of RPKI prefix validation in near real-time. It is our hope that these tools will aid the community in planning, testing, and improving the deployment of RPKI ROAs for the benefit of a future Internet backbone of enhanced security and reliability. Furthermore, based on the insights presented in this chapter it is doable to proceed with a more coarse-grained analysis of closed routers. We will look on this in the future, as well.

Chapter 4

Analysing Effects of Content Delivery Infrastructure on Network Security

4.1 Introduction

Websites often disappear, when prefix hijacking (or BGP misconfiguration) occur [26, 136, 163]. In the worst case, traffic hijacking will be used to implement forged SSL certificates. For example, consider a trusted certificate authority (CA) that cooperates with the attacker, and that this CA creates a certificate to bind the attacker's public key with the victim's domain name, then the attacker can successfully impersonate the victim's web server as soon as traffic has been redirected. At Black Hat conference 2015 another example was introduced. By using prefix hijacking the attacker intercepts traffic between the CA and the victim to gain a valid certificate for the victim's domain [61]. Even though such incidents have not been clearly documented [76] and the SSL vulnerability can be mitigated by DANE [72], it does not solve BGP vulnerability, which reduces service availability of the Web ecosystem—this has been experienced multiple times.

The relevance of websites in social and business critical operations challenges a comprehensive security provisioning. In particular, the network operators of popular sites should be given protective mechanisms at hand. The Resource Public Key Infrastructure (RPKI) is a new framework that enables BGP routers to perform prefix origin validation, and when active prevents incidents such as the YouTube hijack. On the same time, the web is a complex ecosystem with multiple players. The content owner has only limited influence on the delivery infrastructure. This is foremost valid for popular sites as their deployment is usually outsourced to CDNs, which distribute content into many autonomous systems. We know from IPv6 deployment that such companies are slower adaptors of new Internet technologies than smaller ISPs or webhosters. In our case, this leads to less protected popular content but better protected unpopular content.

In this chapter, we conduct a first quantitative analysis of the deployment of securing web servers by RPKI. We map the IP addresses of 1M Alexa domains to IP prefixes and evaluate if these prefixes are part of the RPKI infrastructure. Our findings can be summarized as follows:

1. Less popular websites are commonly better secured than websites with many visitors.
2. CDNs tend to ignore RPKI, whereas ISPs and webhosters started RPKI deployment. For the 199 ASes of popular CDNs we only discovered four entries in the RPKI.
3. Content from CDNs that is not placed in CDN networks but in third party networks benefits from these earlier adopters.
4. CDN deployment policies are the likely cause for a reduced security level at prominent websites.

Furthermore, we present a software extension which performs prefix origin validation in the web browser of end users. The browser extension shows the RPKI validation outcome of the web server infrastructure for the requested web domain. It follows the common plug-in concepts and does not require special modifications of the browser software. It operates on live data and helps end users as well as operators to gain better insight into the Internet security landscape.

The remainder of this chapter is structured as follows. Section 4.3 surveys related work. Section 4.2 details the problem space and our measurement methodology. Section 4.5 presents our findings. Section 4.7 introduces our web browser extension. Section 4.8 concludes with an outlook.

4.2 Background

Why the Web Ecosystem Challenges Network Security Measurements

The web ecosystem basically consists of the following components. The *end user* requests a *web page* from the *content infrastructure*, which is delivered via the *underlying network*. The web page belongs to a specific *domain name* (e.g., `www.google.com`). To successfully reach the content infrastructure two steps are necessary: (a) a valid mapping of the domain name to a host address and (b) the correct routing of the host's prefix within the BGP. DNSSEC [7] ensures valid name to address mapping. This paper focuses on the routing layer.

In the simplest case, a web page is hosted on a single web server situated in a single autonomous system. The prefix owner then needs to create a single RPKI entry for the prefix-AS pair, which hosts the domain. However, highly popular content (i.e., web pages with many visitors) is often distributed among several web servers to increase availability and performance. With the advent of Content Delivery Networks (CDN), these web servers are not only reachable via different IP prefixes but also placed in different ASes. To fully secure the web server infrastructure of the domain, all prefix-AS pairs need to be included into the RPKI. It is worth noting that content provided by CDNs is not necessarily located in the AS of the CDN, in which case the CDN has no control over the authorization of this AS. We will show that end users benefit from diverse deployment, since CDNs—in contrast to larger ISPs—do not protect their IP prefixes so far.

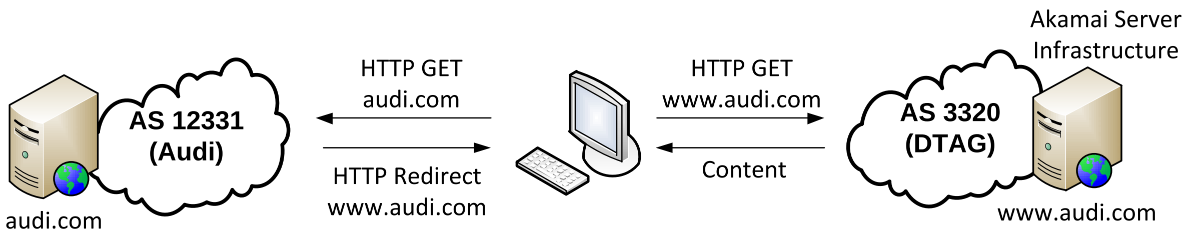

Figure 4.1: Example of accessing CDN-content

In addition to asking where content is located in the network, the question about relevance is of concern. The Alexa list [17] ranks websites according to their popularity in terms of estimated visitors. Unfortunately, Alexa aggregates all sub-domain names to the second level. `maps.google.com` for example is merged into `google.com`. In this example, the two domains obviously offer different content and it may be reasonable to assume a distinct hosting.

On the contrary, it is also common to use different domain names for the same content, most prominently adding the prefix `www` to the second level domain name. Usually, one of the names redirects to the other. The redirection is not necessarily based on DNS canonical names, but also on HTTP redirects [55]. In such cases, the initial HTTP GET request on a w/o www domain may point to a different infrastructure, even though the content is finally provisioned by the same servers that host the www domain (or vice versa).[1] Our analysis focuses on RPKI protection and thus needs to consider the first contact point where misconfiguration may become effective. The distinction between www and w/o www domain may be of lesser importance for measurements that only care about actual content location [12].

We summarize our discussion with a concrete example of the complexity of the current web ecosystem and its implications for RPKI measurements (cf., Figure 4.1). When a web client requests `http://audi.com` from a host at Berlin IXP, the domain is resolved to `143.164.100.143`. The corresponding IP prefix is originated by AS 12331 (Audi). The web server sends an HTTP redirect to `62.156.238.22`, which is operated by AS 3320 (DTAG). In contrast, resolving `www.audi.com` directly leads to `62.156.238.22` via several canonical DNS names. The web server behind `62.156.238.22` is operated by Akamai. From a deployment perspective this illustrates that a CDN has only limited influence in achieving full RPKI protection for a domain name. From a methodology perspective it clarifies that both domain names should be analyzed.

[1] Throughout this paper, we will use the term “w/o www domain” when we refer to a *<domain>* instead of *www.<domain>*.

4.3 Related Work

The deployment of RPKI started in 2011. Several looking glasses and tools exist [93, 139, 107, 87, 141, 120] to inspect the current state of deployment or to do experiments, but up until now only few publications studied the current state of deployment in detail. [142] and [78] analyzes the RPKI validation outcome of entire BGP tables trying to better understand invalid BGP announcements. [45] discusses the risk when RPKI authorities misbehave, and [65] explores the general limitations of current secure inter-domain routing protocols. Nevertheless, large ISPs such as Deutsche Telekom and ATT added their IP prefixes to the RPKI, which motivates the relevance of this new protocol framework. The motivation to adapt new Internet protocols is analyzed in [118], with a special focus on secure inter-domain routing in [37, 64]. Our work complements these insights by clarifying that CDN-content benefits from early RPKI-adoption in large ISP networks.

Several measurement studies discovered the content distribution space (e.g., [133, 75, 137, 12]). We emphasize that the aim of this paper is not to reveal the hosting infrastructure completely but to present a trend analyse of RPKI adoption in the wild and its interplay with the web ecosystem by applying a very basic methodology.

4.4 Methodology

In this chapter, we want to analyze the deployment of RPKI protection for the web server infrastructure of popular web sites. Our measurement study proceeds in four steps: (1) selection of websites, (2) mapping domain names to IP addresses, (3) mapping these IP addresses to IP prefixes routed in the Internet and their origin ASes, (4) validating the BGP information against RPKI. We now describe our methodology, which is meant to be simple and widely reproducible, in detail.

4.4.1 Selecting Domain Names

Domain names in our measurement are taken from the Alexa list, which encompasses 1M entries. Alexa domain names are commonly applied in measurement studies (e.g., [12, 124, 99, 33]), as there is no global directory or other striking alternatives. Using them eases reproducibility. In addition, the Alexa list provides indications of the domain popularity.

4.4.2 Mapping Domains to IP Addresses

The distributed nature of DNS may lead to answers that vary between locations. Requesting the address record of the same domain name from different resolvers may return different IP addresses. This has become even more prevalent in the context of web hosting, where CDNs introduce location-specific DNS replies. We restrict our analysis to data from *public* DNS servers

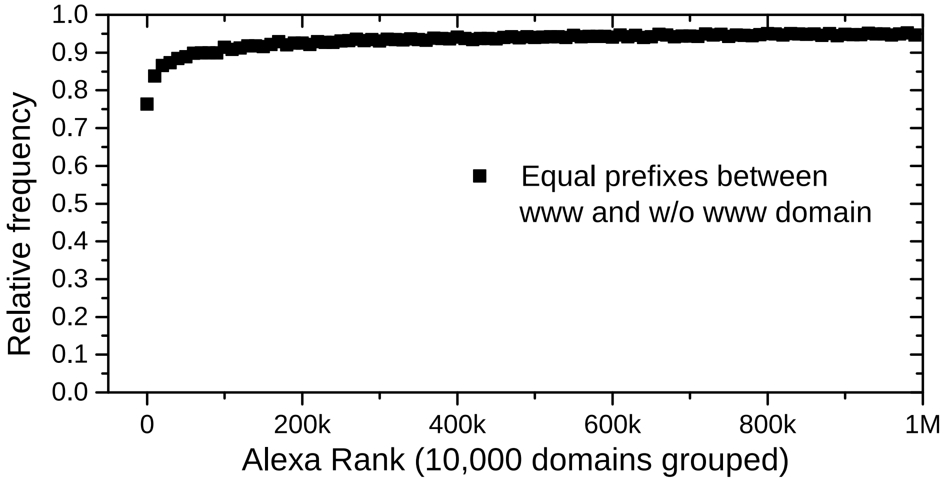

Figure 4.2: Comparison of IP deployment for www and w/o www domain names

like GoogleDNS as they allow for timely and continuous as well as distributed measurements. We decided against launching a data collection campaign or using multiple open recursive DNS servers to request the same name from different vantage points for the following reasons.

First, data collection campaigns are based on volunteers. Resolving two million names is a time-consuming task, which conflicts with time resources participants are usually willing to spend. Establishing a measurement ecosystem such as SETI@home [20], DIMES [131], or BISmark [138] is beyond the scope of this thesis.

Second, most of the Open Recursive DNS servers (ORDNS) are very instable and do not provide reliable answers, because the majority of the entries refer to CPE devices such as home gateways or printers [128]. Our tests showed that the same set of hosts can work well for a short time, but go off-line or return errors a few days later. Based on this unreliable behavior, we exclude ORDNS lists as they do not allow to reasonably reproduce our measurements. In contrast, DNS Looking Glasses and the Google DNS as well as Open DNS are free and stable global DNS resolution services operated for the public. In Section 4.5, we show that our RPKI results remain independent of the DNS server selection. Analyzing the effects of many vantage points will be part of our future work.

We collect all `A`, `AAAA`, and `CNAME` records for the Alexa domain names [17], including the names appended with the prefix `www`. We exclude all invalid DNS answers, i.e., all special-purpose IPv4 and IPv6 addresses reserved by the IANA.

To briefly analyze the overlap between www and w/o www domains, we quantify the amount of equal prefixes per domain. Figure 4.2 shows that for the first 100k domains more than 76% of the IP prefixes equal for both names, for the remaining domains more than 94% of the names refer to the same prefix. To accelerate continuous DNS measurements, it is sufficient to resolve only one type of name.

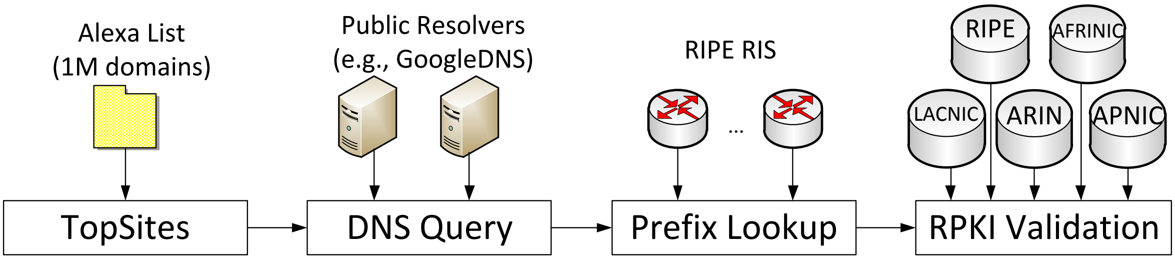

Figure 4.3: Toolchain overview

4.4.3 Mapping IP Addresses to Prefixes and ASNs

We combine dumps of the active tables of the RIPE RIS route servers. For each IP address of a domain name, we extract all covering prefixes and derive the origin AS from the AS path (i.e., the right most ASN in the AS path). Note that entries with an AS_SET are excluded from our study as this leads to an ambiguity of the attribute, why the function is deprecated with the deployment of RPKI [90].

(4) RPKI Validation

For the validation of the BGP data, we follow the necessary steps to perform origin validation at BGP routers. ROA data of all trust anchors (APNIC, AfriNIC, ARIN, LACNIC, and RIPE) are downloaded and validated. Only cryptographically correct ROAs are further used to check the IP prefixes obtained from the BGP table dumps.

To summarize (cf., Figure 4.3), for each domain name in the Alexa list that can be resolved from our DNS vantage point, we create a list of IP prefixes and origin ASes visible within RIPE RIS and assign an RPKI validation state based on the currently deployed set of ROAs. We will make the data publicly available as there are no reasons for privacy or ethical concerns.

Our approach does not measure the protection of embedded web content (e.g., photos that are located under a different domain compared to the landing page). Security incidents in the past [26, 163, 136] showed that routing failures usually affect complete web pages instead of content pieces, which justifies our current focus.

4.5 Results

After resolving 1M Alexa domains from a host in Berlin via GoogleDNS and excluding 0.07% incorrect DNS answers, we gathered 1,167,086 IP addresses for the www domains and 1,154,170 IP addresses for the w/o www domains. These addresses are mapped to 1,369,030 and 1,334,957 different prefix-AS pairs respectively. 0.01% of the IP addresses are not reachable from our BGP vantage points.

We repeated the DNS measurements over several weeks and also resolved the domain names

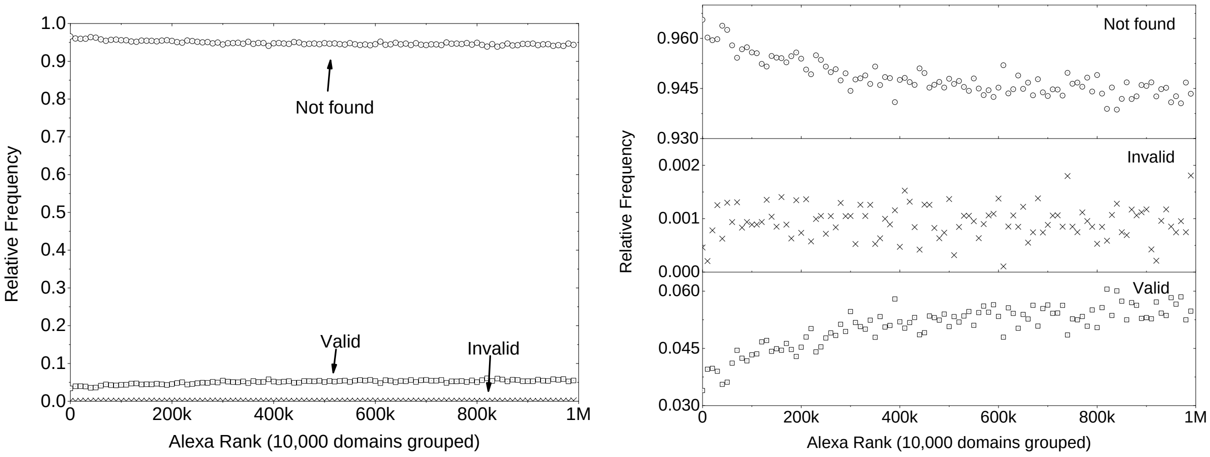

(a) Summary

(b) Zoom

Figure 4.4: RPKI validation outcome for the Alexa domains (*valid* = the origin AS is allowed to announce the prefix, *invalid* = the origin AS is not allowed to announce the prefix, and *not found* = the announced prefix is not covered by the RPKI)

via Open DNS and `us01` of the DNS Looking Glass [57]. For these two different DNS resolvers we found very similar results compared to GoogleDNS.

In the following subsections, we first present the core RPKI validation outcome, and second work out reasons for the observed deployment state. For better visibility, we do not present results per domain but apply a binning of 10k domains in all graphs. As a domain name may refer to multiple IP addresses, which belong to different IP prefixes and ASes, several RPKI states may exist per domain. To represent heterogeneous RPKI deployment, we assign corresponding probabilities to domain names (e.g., 3/5 RPKI coverage of `foo.bar`).

4.5.1 Basic RPKI Insights: Infrastructure of less popular sites is more secured

On average, 6% of the web server prefixes are covered by the RPKI (either correctly or incorrectly announced in the BGP) and 94% are not secured (cf., Figure 4.4(a)). Roughly 0.09% of the prefixes are invalid according to the RPKI prefix origin validation. This observation is in qualitative agreement with the general RPKI deployment. Note that the current invalid BGP announcements do not necessarily indicate hijacking but rather misconfiguration [142]. The amount of invalids is evenly distributed among all web domains.

Looking into RPKI protection in more detail shows that the portion of RPKI coverage correlates with the popularity of websites Figure (4.4(b)). Domains with a low rank (i.e., popular sites) are less likely secured than domains with a high rank. Among the first 100k domains (e.g., `google.com`), for only ≈4.0% of the web server prefixes an origin validation can be performed.

In contrast, for the last 100k domains ≈5.5% are secured. We admit that the absolute numbers are small, but a clear trend is visible and may reflect the deployment strategy of different stakeholders, which we now try to clarify.

4.5.2 CDN-Content benefits from security by third party ISPs

Popular websites are mainly hosted by CDNs such as Akamai, whereas less prominent sites are placed on web servers from a common webhosting service, or on self-maintained servers connected via some third party ISPs such as DTAG or Sprint. Following this, we conjecture that CDNs are more hesitant in deploying RPKI for their IP prefixes. To study this interrelation, we inspect the RPKI repository and search for attestation objects that belong to the ASes of well-known CDNs (i.e., Akamai, Amazon, Cdnetworks, Chinacache, Chinanet, Cloudflare, Cotendo, Edgecast, Highwinds, Instart, Internap, Limelight, Mirrorimage, Netdna, Simplecdn, and Yottaa). If we do not find ROAs including the ASes of the CDNs, the CDNs do not support the RPKI security mechanism. It is worth noting that the results of this approach *do not* depend on DNS measurements and thus *do not* include a bias which might result from our DNS measurement point.

This approach does not consider the case where CDNs own provider independent (PI) address space, which is assigned to ASes of third party ISPs. However, to the best of our knowledge this kind of deployment is not implemented.

For the 199 ASes of the considered CDNs, we only found four entries in the RPKI. These four prefixes are owned by Internap and tied to three origin ASes. Considering the large number of CDN operators, as well as Internap operating at least 41 ASes, this is a very low coverage. On the one hand, we conclude that CDNs tend to not actively participate in the creation of RPKI attestation objects, which is in contrast to webhosters or common ISPs. On the other hand, the results do not show that all CDN-content is only accessible via unprotected prefixes.

CDN content is not exclusively located within CDN network infrastructure, but also placed in third party ISP networks, and thus inherits RPKI deployment there. To quantify this collateral effect, we conduct a very basic classification of CDN domains. Often CDNs use CNAME chains (Canonical Names) to direct DNS requests to their internal end points. We say a domain is served by a CDN, if the IP address of its domain name is indirectly accessed via two or more CNAMEs (e.g., `www.huffingtonpost.com` → `www.huffingtonpost.com.\edgesuite.net` → `a495.g.akamai.net` → `212.201.100.136`). We question this rough heuristic by comparing its results with an independent classification provided by HTTPArchive. HTTPArchive classifies the first 300k Alexa domains based on DNS pattern matching of CNAMEs, which is distinct from our test of DNS indirections. Furthermore, the HTTPArchive monitoring agent is located in Redwood City, CA, USA, and thus a geographically separated vantage point.

Figure 4.5 compares the distributions of CDN-hosted web domains as determined by our classification approach and HTTPArchive. The two almost identically shaped curves clearly

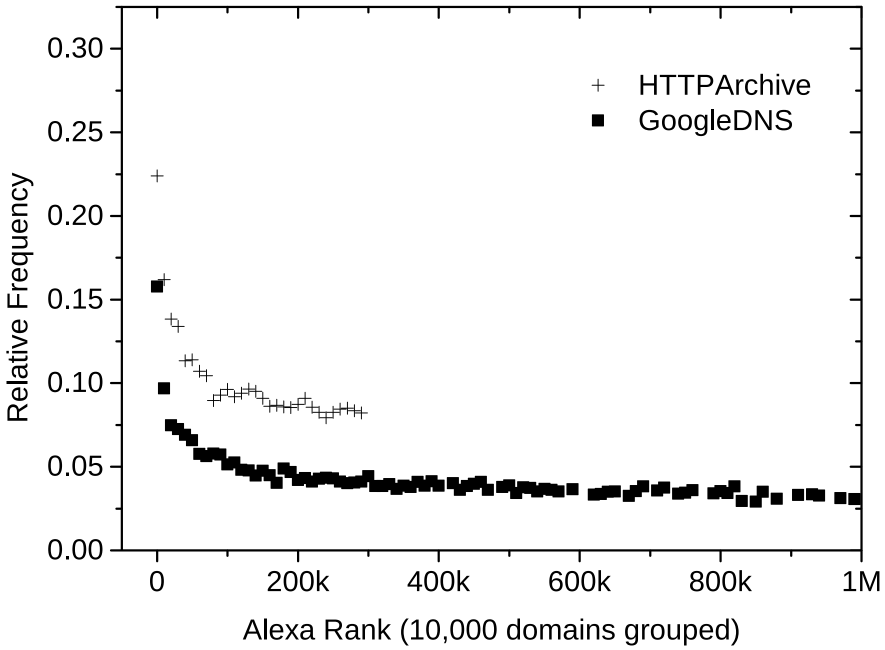

Figure 4.5: Popularity of CDNs—comparison of CDN detection heuristics for 1M Alexa domains

indicate that popular websites are more likely to be served by CDNs. Quantitatively, our approach indicates fewer CDNs than HTTPArchive. This is not surprising, since there is CDN deployment without CNAME chains. However, a conservative (under)-estimate of CDN domains sharpens our view on the RPKI-protection of CDN domains: (Over-)enlarging the set of CDN domains will mix deployment cases and diffuse the overall picture.

4.5.3 CDNs likely to cause reduced security of the popular Web

We now focus our analysis on the relation between RPKI-enabled and CDN-served content. We want to examine our earlier conjecture that hesitant deployment at CDNs is the dominant reason for the observed trend that popular sites are less protected. Figure 4.6 depicts the distribution of RPKI-enabled content under the condition that it is CDN-served. In contrast to the Alexa domains at large, RPKI deployment is fairly independent of the rank for CDNs. Results fluctuate around an average of ≈ 9 ‰. This is almost an order of magnitude lower than the overall RPKI deployment rate, which is plotted for comparison.

Combining the line of arguments, we could show that (a) CDN deployment is strongly enhanced for popular domains, but (b) RPKI deployment on CDN content is low—independent of its popularity. As a result, a high density of CDN sites reduces the RPKI-enabled portion of domains. This holds for those ranks of the Alexa list where CDNs are more common: the low ranks of high popularity. Our argumentation is roughly reflected in numbers. The presence of CDNs is enhanced by about 2 % at the lower Alexa ranks, where RPKI deployment is reduced

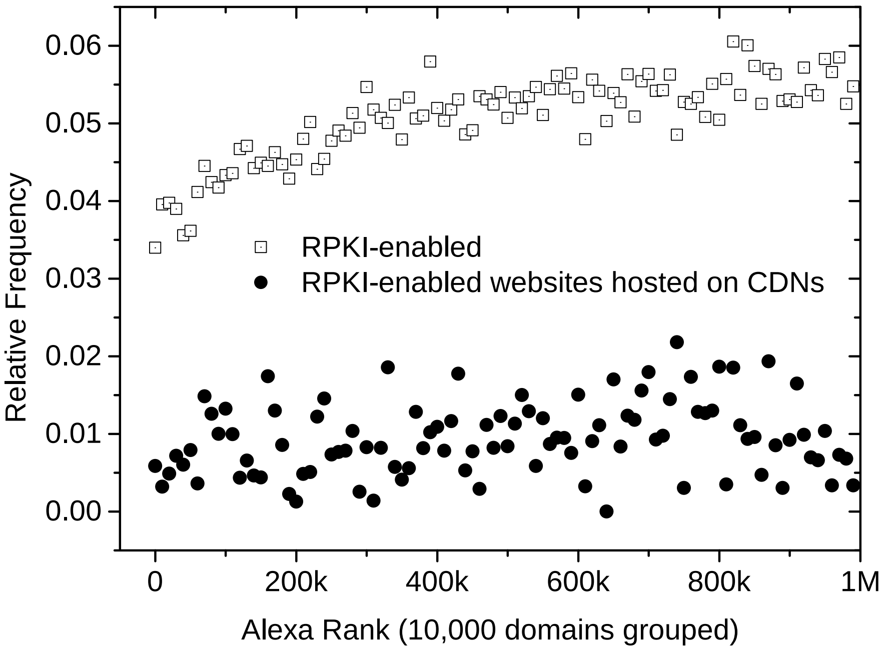

Figure 4.6: RPKI deployment statistics on CDNs and for the unconditioned Web

by about the same 2 %. In summary, we take this as a strong support of our hypothesis that the observable degradation of routing security for popular websites is caused by the resistance of CDN operators to adopt RPKI.

4.6 Reasons for Reduced Deployment

The need for RPKI must be more critically communicated to web providers—particularly large CDNs that serve disproportionate amounts of web traffic. This is largely a process of minimising barriers (e.g., cost), alongside offering incentives. During our study, we have spoken to many network operators to gauge their opinion of RPKI. We have found several reasons why operators have not deployed it. Often this relates to a lack of perceived need, combined with insufficient manpower and expertise in the area. This is common across many technologies. More interestingly, we also have discovered RPKI-specific factors that dissuade adoption.

RPKI is a proactive security solution that implements a positive attestation model. As such, RPKI exposes information that organisations may be wary of revealing. Most worrying to some is the potential of revealing their business relations. In RPKI, prefix owners must proactively create ROAs before any attack occurs. As soon as at least one ROA for an IP prefix exists, *all* valid origin ASes for this IP prefix need to be assigned in the RPKI *before* route updates are processed (otherwise a BGP update including the prefix and missing ASes becomes invalid). Although a prefix owner can assign any AS without asking for approval, it is very likely that the ROA information indicates a business relation between prefix owner and authorized origin AS. This might be a serious concern, from our discussions with several operators.

To illustrate a business policy conflict, imagine that two large CDNs serve secretly as backups for each other. As changes within the DNS are slow, they could quickly redirect traffic using BGP. Similarly, smaller CDNs that rely on third party networks (e.g., using Verisign for external DoS mitigation) may see such information as damaging to their reputation. Despite this, in both cases, RPKI would publicly reveal these setups.

It is worth noting that RPKI data differs from public routing data such as BGP collectors or looking glasses. Those sources also provide insights into peering relations but only after the event has occurred. Furthermore, the data analysis follows an exploratory approach because not all vantage points report the same. In contrast to this, the RPKI represents a catalog which does not only allow for easy browsing but also documents information in advance. In case of a very unlikely or never occurring event (e.g., a backup incident), the RPKI exposes more information. We therefore argue for changes to RPKI that address these business concerns. Whereas privacy-aware protocols are rife in other fields (e.g., application layer), they have always been seen as less important in the routing layer. The above observations undermine this assumption. Arguably, not recognising this issue could be extremely damaging to RPKI deployment.

4.7 RPKI Validation in Web Browsers

4.7.1 Design

The design of our solution is driven by real-time analysis and flexibility. To verify the BGP prefix of the web server a URL resolves to, basically the following steps are necessary: (a) DNS resolution of the web domain, (b) mapping of the IP address to prefix and origin AS visible in BGP, (c) comparison of the prefix/origin AS with ROA data. It is worth noting that the DNS resolution as well as the BGP data depend on the location of the client. However, in contrast to the name to address mapping in the DNS, there is no standard mechanism to request the prefix/AS pair for an IP address with respect to the customer's ISP routing table.

There are two options to implement origin validation in web browsers: (a) the browser extension implements the full router part (i.e., receives valid ROAs from cache server and performs origin validation of BGP data), (b) the extension resolves only the IP address of the web domain and a remote back-end performs the origin validation. We decide for the latter as this allows for easy applicability in most browser platforms, which usually provide add-on concepts based on JavaScript. Back-end and front-end communicate via HTTP as this is native in browsers and does not conflict with most firewall settings.

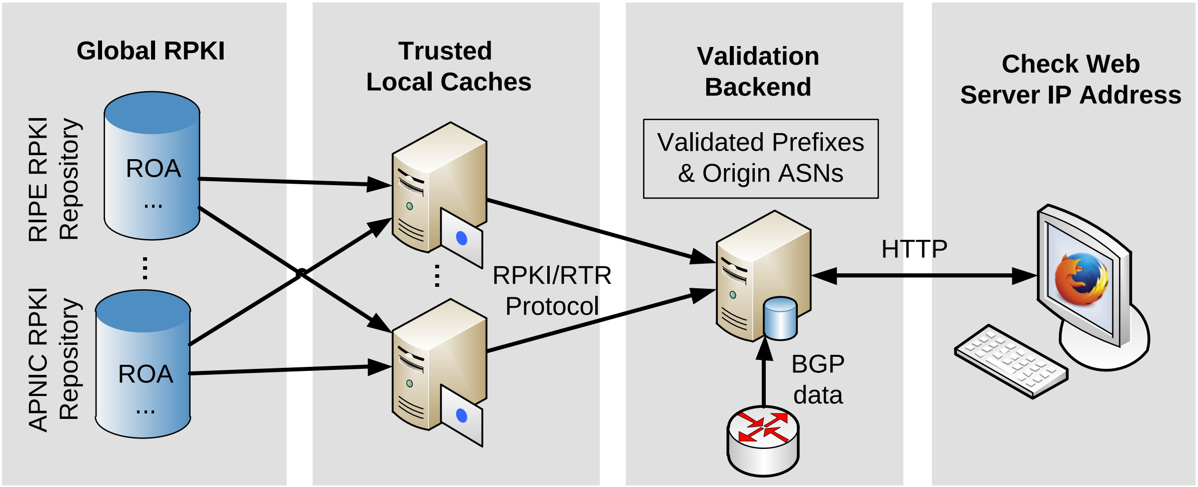

Figure 4.7: System architecture

4.7.2 Implementation

Back-end

Using the Team Cymru community service [5], the back-end resolves the IP prefix of the web server IP address and the corresponding origin AS. We admit that the result does not necessarily comply with the BGP entry of the client's upstream but the vantage points of Team Cymru provide a good coverage.

To fetch ROAs and to validate the BGP information, we deploy the RTRlib [141], an open source implementation of the RPKI/RTR router part. This C library is very efficient with respect to memory and processing resources. Per default, the implementation establishes RTR sessions to two cache servers for fallback reasons. However, end users can configure their own end point for a cache server in the browser extension, and multiple instances of the RTRlib will be started.

Front-end

The browser extension is implemented as dynamic add-on for Mozilla Firefox and Chrome. Other browsers can be easily supported as the browser extension only needs to support the web interface to the back-end. The source code is available on Github[2]. The extension visualizes three states: green (the web server prefix is valid in the BGP), orange (the prefix was not found in the RPKI), and red (the prefix is invalid, the website might be suspicious), see Fig. 4.8. Advanced users can request information about the autonomous system/the IP prefix and configure the host address and port of the RPKI cache server, which will be used by the back-end.

[2] `https://github.com/rtrlib`

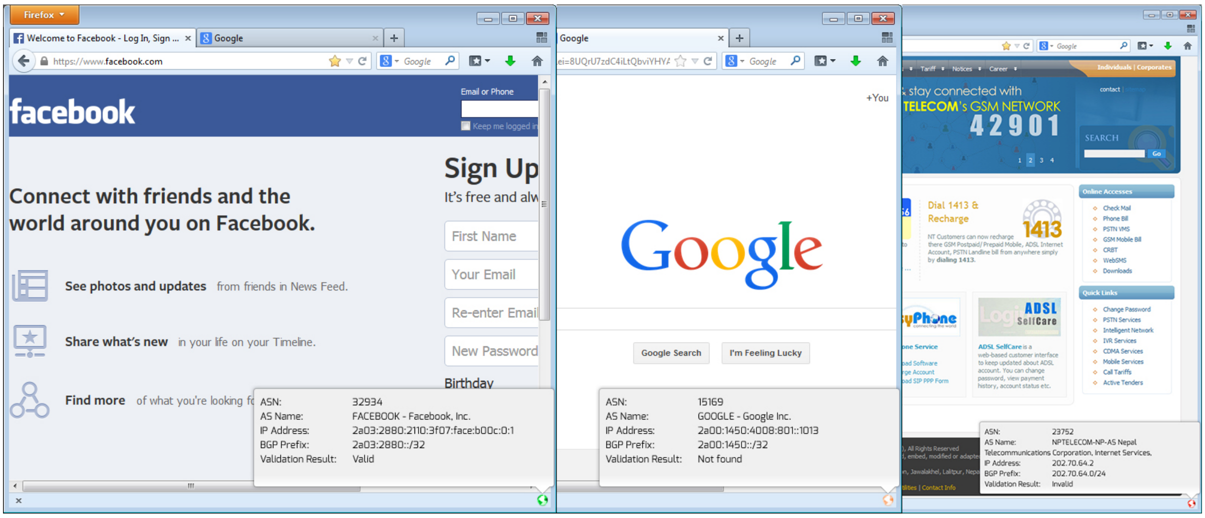

```
ASN:                 32934
AS Name:             FACEBOOK - Facebook, Inc.
IP Address:          2a03:2880:2110:3f07:face:b00c:0:1
BGP Prefix:          2a03:2880::/32
Validation Result:   Valid
```

```
ASN:                 15169
AS Name:             GOOGLE - Google Inc.
IP Address:          2a00:1450:4008:801::1013
BGP Prefix:          2a00:1450::/32
Validation Result:   Not found
```

```
ASN:                 23752
AS Name:             NPTELECOM-NP-AS Nepal
Telecommunications Corporation, Internet Services,
IP Address:          202.70.64.2
BGP Prefix:          202.70.64.0/24
Validation Result:   Invalid
```

Figure 4.8: RPKI validation in Mozilla Firefox for the requested websites with different validation outcome

Future work

Our current solution checks the BGP data for the web server infrastructure of the landing page. Many web pages include embedded content linked via different domains. Analyzing the HTML content fast and combine the different results to a complete picture for the whole web page will be part of our future work.

4.8 Conclusion and Outlook

In this chapter, we analyzed the RPKI-protection of websites. We resolved 1M domain names from the Alexa ranking (including names appended with the `www` prefix), mapped the IP addresses to IP prefixes and origin ASes visible in the global BGP routing table, and validate each prefix-AS pair against the currently deployed RPKI data.

We found that RPKI security deployment is significantly degraded for the more popular websites, which led us to applying a new, initial methodology for discovering its reasons.

Our findings revealed that CDN hosters are the likely cause for this operational bias. Their enhanced provenience at prominent web domains on the one hand, and their obvious reluctance towards RPKI deployment on the other hand strongly indicate that prominent websites would be better protected against routing attacks without CDNs.

In future work, we will aim to explore why CDNs implement this operational behavior. We will compare RPKI deployment with the adoption of other core protocols such as DNSSEC. Furthermore, we will look in more detail if provider independent address space is actually not used by CDNs in the web ecosystem.

Chapter 5

Disclosing Internet Attacks on Mobile Devices

5.1 Introduction

Scanning Internet hosts to initiate a denial of service, to find an exploit, or to discover an unsecured remote access is typically the first step of an attack towards Internet devices. In former times those attacks have been reserved to traditional server systems [18]. Today, not only desktops but also mobile devices (e.g., smartphones) offer intentionally external services.

Mobile phones are particularly threatened by attacks. They are almost always connected with the Internet. Their limited resources do not allow the application of commonly used security mechanisms. In addition, many end users disable security barriers, which have been introduced by vendors, when they root or jailbreak their mobile [112]. From this perspective it is reasonably to assume that attackers specifically target on mobile devices.

Antivirus companies focus on identifying malware. A plethora of research discussed network-based vulnerability of mobiles and proposed solutions (e.g., [69]), but up until now unsolicited remote accesses to mobiles have not been studied in detail. In this chapter, we argue that a measurement infrastructure is required which aims to quantify and to analyse the amount of remote attacks on mobiles.

A common technique to study attack behavior is the deployment of honeypot. A honeypot is a trap for collecting data from unauthorized system access—in this analysis via IP—initiated by remote parties. However, the term "mobile honeypot" is not well-defined and there is only very limited work on the design of a measurement system that allows for both, the analysis of the mobile as well as non-mobile world.

In this chapter we make the following core contributions:

1. We introduce the detailed design and implementation concepts of a mobile honeypot. The principles of the system have been developed independently by two groups, the Deutsche Telekom and us. Unknowingly that both groups worked on the same topic we made the same design choices, and started collaboration later. The honeypot has been running for approximately 1.5 years and is part of the early warning system of one of the largest ISPs

in Europe. The honeypot system abstracts from unnecessary mobile aspects, which allows us to deploy the same software base on standard PCs that are connected to different types of Internet access.

2. We report on preliminary measurement results. This includes a summary of our current observations of attack behaviour on smartphones, as well as a statistical analysis of unsolicited traffic [164]. The traffic measurement presents data from November 2012, which we compare with common wired Internet access and with our results from January 2012 [148]. Surprisingly, we do not find a significant amount of attacks specific to mobiles, which indicates that adversaries operate almost independently of the actually captured host.

The vulnerability of smartphones is based on multiple aspects. This chapter concentrates on remote attacks via the Internet. One might argue that a mobile operator usually do not assign public IP addresses to mobiles and that NAT techniques protect the systems against malicious access. We think the mobile environment needs an early and continuous analysis as well as appropriate tools. There are operators providing public IP addresses. With an increased deployment of IPv6 more public IP addresses will be assigned to end users as several application scenarios requiring direct access without NAT traversal. Note that the attackers then might alter their strategies but our system will still provide the necessary data to explore this transition.

The remainder of this chapter is structured as follows. In Section 5.2 we introduce background and discuss related work in the context of mobile honeypots. We present the design space, implementation, and deployment aspects of the mobile honeypot system in Section 5.3. Our measurement study is discussed in Section 5.4. We conclude with an outlook in Section 5.5.

5.2 Background and Related Work

5.2.1 Trapping Attackers with a Honeypot

In contrast to other security measures that ultimately try to keep the attacker out of the system, honeypots are meant to be compromised. Their value lies in luring the attacker into entering a system and collecting information on how this is done.

A honeypot is typically classified as low interaction or high interaction honeypot and client or server honeypot. A *low interaction honeypot* primarily collects information about the attacker and detects known attacks. The limited level of interaction between attacker and target is achieved by not providing fully functional services but only emulations thereof with known exploits. On the other hand, a *high interaction honeypot* provides a fully functional system. They are used to reveal current and new attacks that do not have to be catered for when setting up the honeypot. Since the high-interaction honeypot is a fully functional system, it has to be closely monitored for successful attacks to prevent the attacker from using the honeypot to target other systems on the network. *Server honeypots* provide vulnerable services to malicious clients. Their focus is in protocol and service specific vulnerabilities. A server honeypot does not offer any legitimate services, any connection by a client can be treated as an attack. *Client honeypots* take the role of a vulnerable client trying to find malicious servers.

5.2.2 Wireless versus Mobile Honeypots

Physical and virtual honeypots [115] have been studied in detail, however, there is only little work in the field of mobile-related honeypots. Mobile honeypots have to be distinguished from *wireless honeypots* [134], [15], which focus on the attacks on the wireless technology. The term *mobile honeypot* is used here referring to honeypots that focus on attacks on mobile devices.[1] They can either be mobile themselves in running on the mobile device in which case they would usually be low interaction honeypots used for deception and detection of known attacks. This also greatly reduces the possibility of the device itself being compromised. On the other hand they can be dedicated devices up to high interaction solutions set up to expose unknown attacks. Mobile honeypots in the sense of honeypots focussing on mobile devices are for example developed by the Chinese Chapter of the Honeynet Project [73]. They are using prototype deployments of honeypots for Bluetooth, WiFi, and MMS. TJ OConnor and Ben Sangster built honeyM [108], a framework for virtualized mobile device client honeypots, which emulates in particular wireless technologies. Mulliner *et al.* [105] propose HoneyDroid, a specific mobile honeypot that exclusively runs on smartphones. We argue that those approaches complicate the measurement across different types of systems. In addition, they are only required if the hardware characteristics are relevant for the study.

Honeypots have been deployed as an important tool to identify attacks, not only in research but also in commercial products. Unfortunately, commercial honeypot installations are usually private without sharing data publicly. We are aware of the UMTS honeypot of the Deutsche Telekom, a large operator, which we collaborate with.

5.3 Mobile Honeypot System

Our primary goal is the design of a measurement system that captures traffic characteristics of malicious behaviour on mobile devices and allows for comparison with non-mobile environments. In addition to these statistical observations, we are interested in the more detailed procedure of potential attackers (e.g., which software do they infiltrate). A common technique is the application of a honeypot. In this section, we discuss appropriate levels of abstraction to cover the mobile environment without losing comparability with non-mobile setups. The mobile

[1] Note that the term “mobile honeypot” is also used to describe other scenarios. Balachander Krishnamurthy [89] uses it to describe prefixes of darknet address space that (1) are advertised to upstream ASes, making the information mobile, and (2) change aperiodically, moving the darknet in the address space.

honeypot has been designed and implemented coincidently by two different groups. Both groups approached completely independently at the same conclusion.

5.3.1 Attacker Model

In this paper, we concentrate on a system that analyzes malicious access via the Internet on smartphones. We argue for a typical attacker model. The attacker tries to compromise the smartphone via unsolicited remote connections [69], or captures the mobile using malware and initiates further denial of service attacks to other mobiles or non-mobile hosts [140], [54]. In any case, such remote attacks are bound to the network layer and moreover do not address specifics of mobile hardware, but solely target at the system level. The adversary actively tries to find vulnerable nodes and may use additional information such as IP topology data or web server logs to differentiate mobile and non-mobile networks.

5.3.2 Design

The term “mobile honeypot” is not well-defined. The general design space is based on the following three questions: (Q1) Is it necessary that the probe runs on a mobile device—if yes which device type (notebook versus smartphones versus ...)? (Q2) Is it necessary that the honeypot runs on a mobile operating system (PC emulation versus mobile device)? (Q3) To which network is the mobile honeypot connected (DSL network versus UMTS network versus ...)?

According to the attacker model, there is no need to operate the mobile honeypot on real smartphones. This reduces complexity in building the honeypot and simplifies long-term operation.

As underlying operating system we decided for Linux. This has two advantages: (1) Most of the currently deployed smartphones use the Android OS. We conducted fingerprinting tests using the well-known tools Nmap and Xprobe, which try to guess the operating system. Both tools cannot distinguish Android from current Linux versions. (2) Using Linux enables us to reuse existing honeypot tools independently of the deployment in mobile or non-mobile scenarios. This allows us the ensure comparison between different systems.

An important change is the adjustment of the virtual file system that is presented to the attacker. It should reflect the directory structure of a typical Android system.

To increase the *attractiveness* for an adversary, we account for “rooted” (or “jailbreaked”) devices. A rooted device grants additional system access to the user. It allows post installation of additional services such as HTTP or file sharing. A honeypot system that intends to capture tools introduced by the attacker need to emulate a rooted smartphone. Note, considering rooted devices provides the attacker with *supplementary* features and thus does not exclude off-the-shelf mobiles.

Regarding the third question we argue that the mobile probe should connect to a real mobile network. Otherwise, an attacker could detect performance differences (e.g., network delay) in advance. In addition, a connection via a real mobile operator ensures the assignment of a topological correct IP address. Note, for an attacker it is easy to identify relevant IP blocks, either by testing or analysing meta data in the Internet registries.

As we are mainly interested in the analysis of statistical effects and not on dedicated attacks, the mobile honeypot is primarily based on low-interaction honeypots.

5.3.3 Implementation

Software

To implement the proposed honeypot system, we use multiple well-known honeypot tools. The mobile honeypot consists of Honeytrap and the following different sub-honyepots: Kippo, Glastopf, and Dionaea.

Honeytrap is used to detect generic attacks. It listens on all other transport ports and is particularly useful to analyse statistical effects. Worms, for example, do not need a protocol compliant negotiation of transmission parameters but send data via an existing TCP connection without waiting for corresponding replies.

Kippo is a dedicated SSH honeypot that emulates remote terminal sessions. Login access is secured by a trivial password, which allows an attacker to gain easily access to the system. The user account is granted administrator privileges. An attacker can execute common programs, as well as download and install additional tools. The honeypot records downloaded files in the background for later analysis. To protect the honeypot against compromising operations, all infiltrated actions are only valid within the current attack session and the execution of newly installed programs is prohibited. Note, this does not conflict with our objectives, as we are interested in the principle behaviour of the attacker.

Glastopf implements a web-based media server providing an upload form. Uploaded data can be stored in a simulated smartphone file system. This honeypot emulates typical vulnerabilities of a web system.

Dionaea is used to emulate TFTP and FTP services.

Network Connectivity

Several mobile operators provide only private IP addresses. Nevertheless, there is a continuous demand for public IP addresses. In particular with an increased deployment of IPv6, we expect a significant change, which will enable mobile nodes to participate in the Internet without NAT traversal.

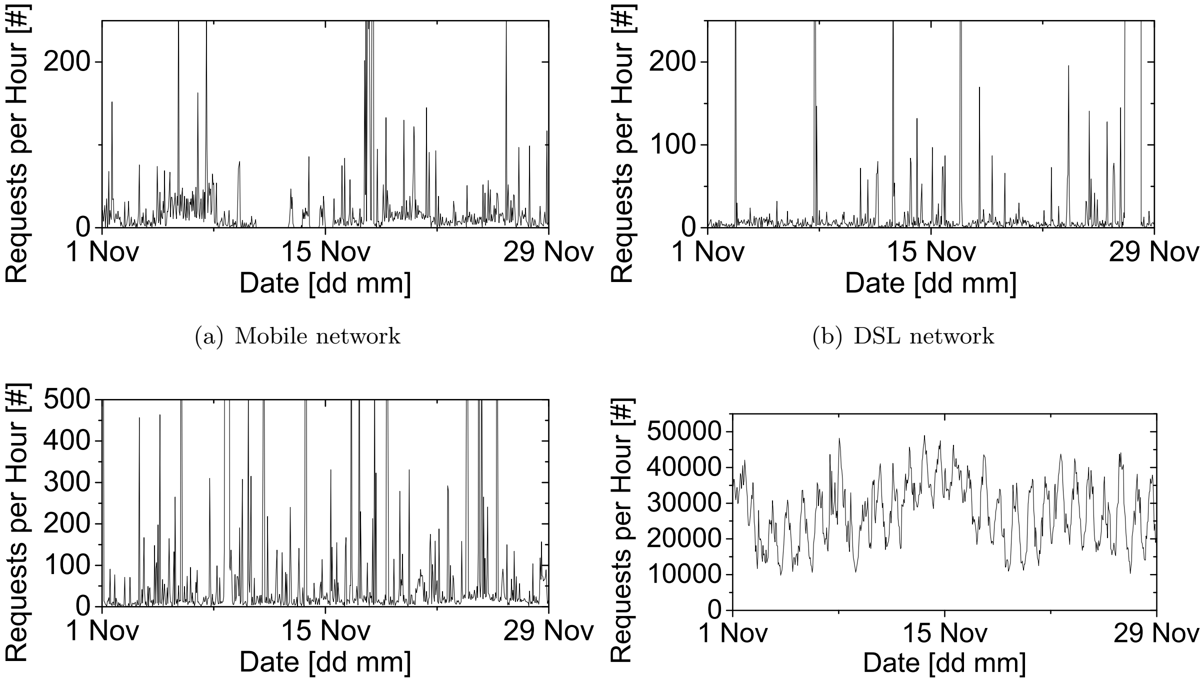

(a) Mobile network

(b) DSL network

(c) Darknet

(d) University network

Figure 5.1: Comparing amount of attacks on different between mobile and non-mobile honeypot probes, Nov. 2012

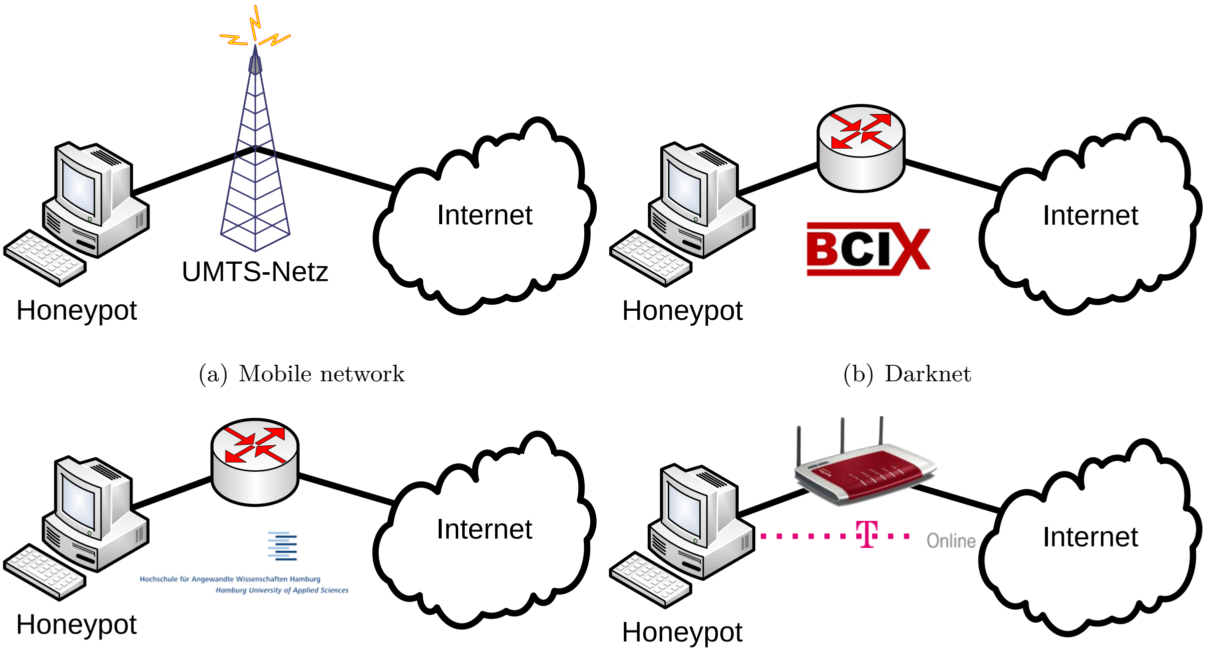
Honeypot
UMTS-Netz
Internet
Honeypot
BCIX
Internet
Honeypot
Hochschule für Angewandte Wissenschaften Hamburg
Hamburg University of Applied Sciences
Internet
Honeypot
Online
Internet

(a) Mobile network

(b) Darknet

(c) University network

(d) DSL network

Figure 5.2: Deploying the (mobile) honeypot system at different network accesses

In addition, many mobile operators, at least in Germany, prevent the communication between end devices per default in NAT domains. For this reason, the deployed mobile honeypot presented in this paper is connected via the Deutsche Telekom, one of the largest telecommunications companies in Europe. They allow to choose an alternative Access Point Name (APN) that provides public IP addresses and thus intra-domain communication.

Note, the proposed honeypot system can be connected to any other type of network access, such as a DSL home network or business Internet access. This allows us to use the same system in different network environments (i.e., wired and wireless infrastructures), and to monitor attack behaviour from different vantage points in the Internet without loosing reproducibility.

5.3.4 Deployment

We started the deployment of the honeypot systems at common PC hardware mid of 2011. Since then they run continuously and surprisingly well. The mobile honeypots of both independent groups include one iOS and two Android probes. They are connected via an USB stick to the UMTS network. All data is exported in a five minute interval to the early warning system of one of European's largest telecommunication companies. To prevent interference and preserve bandwidth of the UMTS link, log data is transmitted using a separate LAN connection. Data from or to the log server is excluded from further analysis.

In addition to the measurement probes that use a mobile Internet access, we deployed the

	# Attacked ports per transport protocol				# Attacks per transport protocol			
	UMTS	Darknet	DSL	University	UMTS	Darknet	DSL	University
TCP	111	133	89	252	14,954	55,378	32,781	22,445,580
UDP	76	71	96	22	637	5,583	8,254	480

Table 5.1: Amount of malicious requests per transport protocol, November 2012

Rank	1	2	3	4	5	6	7	8	9	10
UMTS	22	1433	3306	5900	6666	3389	*1080*	23	*5060*	80
	SSH	MSSQL	MSSQL	VNC		RDP	*SOCKS*	Telnet	*SIP*	HTTP
Darknet	22	139	110	25	3306	*91*	5901	5900	3389	53
	SSH	NetBIOS	POP3	SMTP	MSSQL			VNC	RDP	DNS
DSL	*51099*	22	5900	110	25	3389	143	6666	1433	23
		SSH	VNC	POP3	SMTP	RDP	IMAP		MSSQL	Telnet
Univ.	*445*	139	80	22	110	3389	5900	3306	143	*5902*
	MS AD	NetBIOS	HTTP	SSH	POP3	RDP	VNC	MSSQL	IMAP	

Table 5.2: Top-10 of the most attacked ports (single events emphasized), November 2012

same system at three nodes connected to different non-mobile networks. In detail, the network access is (1) a university network, which reflects a stable and well-known open access; (2) a DSL network, which represents a common home uplink; (3) a darknet, which highlights background noise, because it does not announce any service. These characteristic access types allow for comparison of the mobile measurements with non-mobile environments.

5.4 Measurement Study

In this section, we present measurement results of one month of the mobile honeypot system. We select November 2012 as this represents a complete month without disturbances. This snapshot already reveals interesting insights. A long-term analysis over several years will be part of our future work. In the following discussion, we count any external connect to the honeypot system as an *attack*, its source IP address is called the *attacker*.

5.4.1 General Observations

In general, the number of attacks targeting the mobile probe do not significantly differ from honeypots connected to the Internet via typical wired access. It seems that the attackers scan the Internet without considering specific network types but try to exploit as many devices as possible.

The procedure of the attacker is almost identical to the wired probes. After gaining successfully a shell login and executing some common commands, an adversary usually downloads malicious software and tries to integrate the honeypot into an IRC-botnet. The attacker initiates commands almost independently of the local system properties even if this leads to conflicts (e.g., non-existing directories). We frequently observed that the adversary navigates through the file system directories following the common Linux structure. Specific Android processes have been ignored.

To our surprise, we observed very rarely an intruder that conducted a mobile-specific attack. For example, after establishing an SSH connection to the mobile honeypot, one adversary targeted on the address book as well as the stored photos of the emulated mobile system. Those attacks are usually performed manually and not based on scripts. However, the mobile honeypot did not captured Android- or iOS-specific malware or Exploits.

5.4.2 Comparative Detail Analysis

For our subsequent analysis we focus on network traffic and compare effects on the mobile honeypot with non-mobile systems. We consider the measurement period of November 2012.

Most of the external requests are related to the university network (cf., Table 5.1). The DSL and UMTS honeypots measure on average 21 and 55 attacks per hour, respectively. More surprisingly, the darknet experiences on average about 83 external requests. Around 90% of the attacks use TCP. The prominent ports are 22 (SSH), 1433/3306 (MSSQL), and 80 (HTTP). We summarize details in Table 5.2.

Attacks per AS

To explore the topological location of the attacks, we map the source IP addresses of the adversaries to their origin autonomous system (AS) using the common IP to ASN lookup service provided by Team Cymru. We rank each AS separately per network access.

Overall, most of the attacks are initiated from IP prefixes that belong to the same small set of ASes (cf., Figure 5.3(a)). The top-5 ASes are primarily based in China and Russia and do not cover mobile operators. The distribution of attacks is enhanced for the university network. The darknet and the DSL home network follow a similar shape, in which the mobile network exhibits a more narrowed distribution. For all network types, it is clearly visible that already a small number of ASes have a significant impact on the attack experiences.

Attackers per AS

In our second statistical analysis, we measure the number of different source IP addresses (i.e., attackers) per AS (cf., Figure 5.3(b)). Again, we calculate the rank separately for each network access type. This analysis allows us to estimate the amount of different attack sources and to balance the intensity of each attacker. Consequently, the maximal values are three to two

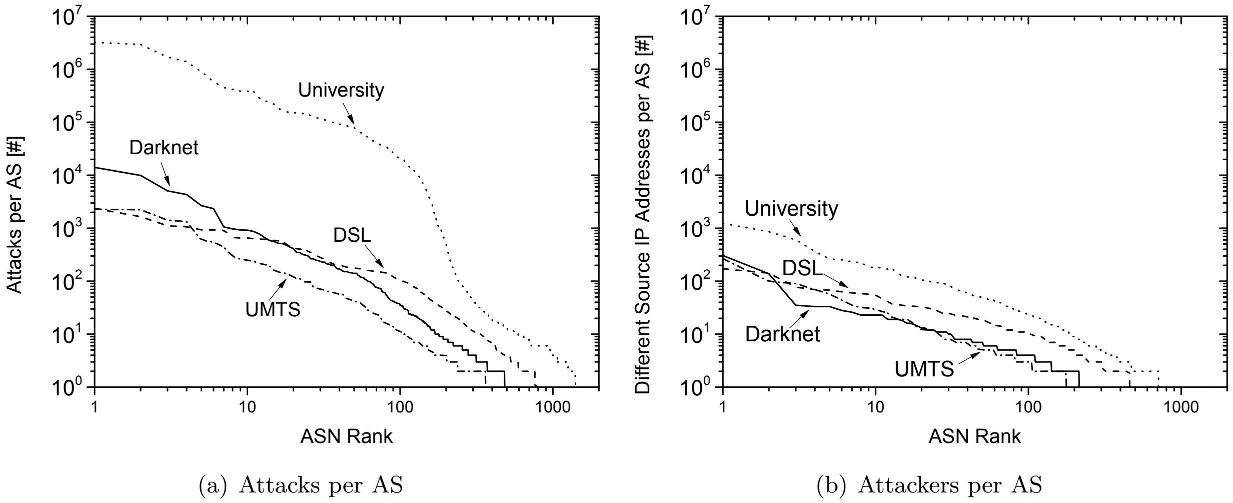

(a) Attacks per AS

(b) Attackers per AS

Figure 5.3: Comparison of requests per autonomous system separately ranked by network access, Nov. 2012

orders of magnitude less compared to the number of attacks. Nevertheless, the characteristic shape of the curves in Figure 5.3(a) still exists.

Comparison with Previous Measurements

In previous measurements from January 2012 we found similar results. The most significant difference is the absolute number of attacks on the mobile system and the darknet. In November 2012 the darknet experienced a surprisingly high amount of attacks, indicating that it is more attractive to an attacker compared to the UMTS and home network. This is in general not true as the IP address space of the darknet is officially not related to any external network service. Looking into this in more detail reveals that a small set of nodes that connect to the darknet initiating a large portion of requests. This observation is also highlighted by our analysis per attacker.

Similar to this, the UMTS network is not spared by attacks in general. In January 2012, the UTMS nodes suffered on average on the same amount of requests compared to the home network. Interestingly regions and originators of attacks were better pronounced and operate at higher intensity in the beginning of this year. We consider this as an indicator for the liveliness of the mobile regime, which needs further analysis in the future.

5.5 Conclusion and Outlook

In this chapter, we presented a mobile honeypot system that allows for a detailed analysis of mobile-specific attacks. A key insight is the abstraction from unnecessary mobility aspects. To study common attacks, the honeypot is neither required to run on a real smartphone, nor on a full-fledged mobile operating system. The mobile honeypot is operated on standard PCs running Linux. This enables the analysis of malicious traffic across different network environments and bears the advantage of simplified long-term maintenance as the same tool basis can be re-used.

We deployed our concept on probes connected to a mobile network, as well as monitoring nodes connected to different types of wired Internet access (i.e., university network, darknet, DSL home network). To our surprise we did not find a relevant ratio of remote attacks that specifically target on the mobile system, neither from non-mobile nor mobile networks. From this observation that attackers tend to ignore the network access in their attack strategy, we conclude that mobile devices need more specific protection against malicious applications (e.g., trojan horse) compared to external, unsolicited requests via the Internet.

In the future we will still maintain our honeypot setup. We will concentrate on more subtle correlation analysis of how specific groups of attackers behave with the aim to identify individual patterns of mobility related aggressions. We will also analyse attacks per port in more detail, and conduct time series analysis. Due to limited statistics, though, these considerations will require a much longer range of observation. Estimating the error of IP spoofing events on our results is also part of our future work.

Chapter 6

Identifying Potentials and Limits of ICN to Protect a Future Internet

6.1 Introduction

One major dedication of today's Internet is the global distribution of content in huge amounts. Content distribution networks (CDNs) facilitate an efficient, wide-area replication of static data for selected content providers, whereas the end-to-end design of TCP/IP does not foresee implicit replication and in-network storage. There is no openly available network standard for the asynchronous, global replication of popular content in the current Internet.

Inspired by the use case of widely deployed Content Delivery Networks (CDNs), current trends of *Information-Centric Networking (ICN)* shift the Internet towards data awareness. In ICN, consumers shall retrieve content by name directly from a network that provides storage, caching, content-based rendezvous, and searching at times. Thereby data sets become first class routable objects and content names require exposure to the control plane.

Several proposals have been presented in recent years [14], among them TRIAD [68], NDN [159, 81], DONA [88], PSIRP [82], and NetInf [13], which differ in several design choices. As we are interested in the stability and security of ICN infrastructures, we will concentrate on the aspects of routing and forwarding.

Essentially two approaches to routing exist in current ICN proposals, an evolutionary path that resolves names to locators and routes on IP (or a related location scheme), and 'clean slate' concepts that route directly on content names. NetInf extends the current Internet by a resolution service that maps content names to topological IDs like IP addresses, but alternatively supports name-based routing. TRIAD, DONA, and NDN perform content retrieval by routing on names. Route responses and the data itself are then forwarded along reverse paths (RPF), either by using IP as a lower layer, or without IP but by dedicated RPF states. PSIRP publishes content objects to a resolution system that incloses full knowledge of the network topology. Requesters trigger the mapping system to generate source routing identifiers in the form of Bloom filters that aggregate IDs of forwarding links.

All solutions operate on the content itself, and force the network infrastructure into a content awareness. A mapping service is not only required to resolve *file* names to source locations, but must answer a request by advising a nearby replica, the existence of which it learned from the data distribution system. Content routers need to rely on (often aggregated) names in its interface tables and—for RPF-based forwarding schemes—a reverse state for every data unit. This control information is highly dynamic and requires regular updates from the data plane. The ICN paradigm thereby opens up the control plane to continuous modifications from the data plane. This is in contrast to the current Internet, where DNS and routing states remain unaltered when a Web page is published, a file is transfered, or data is cached.

In this chapter, we study the impact of traffic conditions on the control plane. We are in particular interested in threats to the stability and security of the ICN infrastructure, whose impacts we evaluate in a theoretical analysis, experimental trials, and simulations based on real-world topologies. Experiments are performed in test networks running PARC's CCNx software. We want to stress, though, that our tests only attribute for the core concepts of content routing and do not evaluate implementation properties of the CCNx prototype. Following the basic insights gained from theoretical and practical analysis, we contribute a sample set of attacks that ground on this correlation of data with control states. We argue that the novelty of these exposures derives from an intrinsic binding to ICN concepts so that attacks—even if reminiscent from today's Internet—cannot be mitigated by simple protocol provisions.

The remainder of this chapter is organized as follows: The specific problems in protecting the ICN infrastructure are stated in Section 6.2. Related work on ICN security is presented in Section 6.3. We theoretically analyze basic threats to stability in Section 6.4 and discuss related implications. Based on practical experiments, threatening scenarios and their effects on the routing system are demonstrated in Section 6.5. Correlation effects in stateful ICN routing are studied in simulations in Section 6.6. These general insights lead to concrete attack scenarios in Section 6.7. The paper concludes with a discussion in Section 6.8.

6.2 Problem Statement

6.2.1 ICN System Model

Information-centric networking involves two functional blocks within the network infrastructure, (1) content publications or announcements, and (2) content subscriptions or (asynchronous) access. Throughout this paper, we assume a generic ICN system model that is composed of these two subsystems, both of which introduce routing or forwarding states at the network layer. Even though not all ICN proposals are constructed equally pronounced in both parts, they all update corresponding table entries in response to data operations of the network infrastructure. In addition, we assume that universal caching is implemented in the content-centric routing system. Universal caching is common to all ICN solutions.

Content requests and delivery do not follow an end-to-end design, but require a dynamic set-up of paths between the requester and a (nearby) copy of the data. Commonly, this is done by

Reverse Path Forwarding (RPF), at which each content request triggers a trail of 'bread crumb' states on routers along the path (NDN, DONA, NetInf). Alternative approaches that route on an underlying routing substrate like IP (NetInf), or constructs source routing identifiers based on complete knowledge of the topology (PSIRP), are not considered further in this work.

Some ICN implementations (e.g., CCN) signal error information when the content is not available. From this perspective they might be considered as request/response scheme. However, in the light of this article, publish/subscribe and request/response are not disjoint categories.

6.2.2 Why Information-Centric Networking is Challenged by Design

Publishing and subscribing in current ICN solutions introduces network control states that generate the following management problems.

1. Addressable content items need advertisement in the route resolution system. Consequently, any end user who can publish requires admission to modify the control plane.
2. Content is conceptually delocalized by universal caching. Data replication thus imposes updates of the routing systems—a change of control state initiated by the data plane.
3. Reverse Path Forwarding requires state initiation and consumption at routers along the path. Corresponding control state updates are not only driven by the data plane, but require processing in wire-speed.

These state operations raise the following threat classes in ways that are unique to ICN.

Resource Exhaustion: Infrastructural entities need to offer accumulating resources like memory and processing power for provisioning, maintaining and exchanging content states. They are therefore threatened by resource exhaustion due to misuse or uncontrolled load. In addition, the asymmetry in size between data requests and delivery leads to traffic amplification when exploited in DoS attacks.

State Decorrelation: The asynchronous nature of publish/subscribe content delivery places the enhanced burden of assuring consistency among distributed data states. Data states that require correlation are situated in distributed mapping systems, which also need to consistently reflect actual content placements, and in forwarding states at routers that define the paths hop-by-hop from a supplier to the requester. Failures in state coherence lead to service disruptions or unwanted traffic flows.

Path & Name Infiltration: The infrastructure relies on the integrity and correctness of content routing and is therefore threatened by poisonous injections of paths and names, in particular. The replicative ICN environment distributes content copies to many, commonly untrusted locations and thereby makes it particularly hard to authenticate valid origins of state insertion requests.

All of these threats bear the potential to seriously degrade the ICN service and lead to insufficient or erroneous data dissemination. A major risk for the ICN infrastructure—and from a general perspective for the ICN concept—results from the power that an end user gains over an ICN distribution backbone.

The Problem of an Open Control Plane

ICN opens the control plane of backbone routers for content consumers and suppliers on a fine-grained base. Granting end users access to the routing and forwarding subsystems is a fundamental step away from the current Internet design and bears significant risks. Current concerns in the context of routing mainly focus on state explosion due to the large amount of content items. One might argue that those resource exhaustions will be solved by more powerful hardware in the future. We will discuss options and limitations of related core aspects in Section 6.4. Still, binding the integrity of the routing infrastructure to the courtesy of *all users* is intrinsic to current ICN approaches—and presumably to the overall ICN concept.

6.3 Related Work

Content Suppliers

Related work on ICN security has primarily focused on validating content correctness and authenticity. Commonly, self-certifying security credentials are included in 'secure names' that facilitate mechanisms for verifying authors, origins, and content integrity [153, 49, 62, 56]. Thus a receiver can be sure to obtain the correct content and an intermediate cache can validate the correctness of the security credentials, which prevents traditional DoS on the ICN system [63]. Nevertheless, having created (or learned) a valid name, any ICN member can re-announce this in the route resolution service, thereby injecting poisonous routes or artificial names into the system.[1] Similar vulnerabilities of DNS and BGP are known from today's Internet infrastructure [31], but remain restricted to (topology) *providers*. ICN opens the liberty of route injection to every content supplier. In an open Internet model, this can be any *end user*. We will discuss threats unique to ICN in Section 6.4.1.

Content Consumers

Before we started our work in [145, 147], little attention has been given to the effects of state management in ICN. Arianfar *et al.* [21] discuss design choices for an ICN router. They concentrate on the content cache and explicitly do not consider per request states. Perino and Varvello [113] have evaluated requirements for content routers that hold content information bases in Bloom filters and reverse paths in pending interest tables (PITs). Under the assumptions of

[1] As a countermeasure, DONA introduces certificates of publishers on the price of per cache-instance varying names. Content routing then works on wildcarding names, which re-introduces the threat of route poisoning.

valid content requests propagated on *homogeneous* network links with a *maximum global* RTT of 80 ms, average PIT sizes are identified in the order of 1 Gbit/s for current line speeds. FIB sizes and lookup complexity were shown to depend nonlinearly on prefix numbers and name lengths. Lauinger [96] explicitly addresses the threat of DoS attacks by filling the available memory of a router with pending interest states.

Such attacks on hardware resources may be mitigated by limiting overall table sizes. However, securing router resources by table limits does degrade network utilization and cannot abandon resource exhaustion problems. In the presence of a table limit, an attacker could initiate massive drops of pending Interests from a router's table and thus disrupt data delivery to regular receivers. The author in [96] proposes to drop Interests at the head of the PIT, which however may easily be misused to DoS-attacking neighbors, or to use Bloom filters instead of PIs. If applied without strict capacity limits, the latter approach is vulnerable to flooding attacks as interface filters degrade their selectivity. In the following section, we will evaluate these effects in detail.

Request state management and related security issues have been raised in [40, 146, 156]. Gasti *et al.* [59] address core issues of route hijacking, state overload, and cache pollution in NDN. They propose countermeasures by extending interface functions, e.g., for limiting rates and survey content delivery. Without considering protective measures in BGP, the authors compare BGP with NDN security and argue that the NDN approach reduces vulnerability to black-holing, as routers can identify unresolved content requests and rank/re-route per prefix and interface. Authors miss that on the one hand RPKI secures BGP against hijacking attacks in a straight-forward manner, while on the other hand proposed countermeasures in ICN cannot prevent attacks of interception and redirection with service degradation.

The idea of limiting the number of incoming Interests, either per prefix (e.g, [60], [150]), per interface (e.g., [8], [42]), or per router (e.g., [48], [42], [8]) to mitigate Interest flooding in NDN has gain larger attention over the last years, after we published our core contribution [145] on this topic, which is part of this chapter. Unfortunately, current approaches cannot solve the problem for general Internet scenarios [16], nicely illustrating that the subsequent analysis is still important.

6.4 Basic Threats to Stability

In this section, we theoretically examine the implications at the control plane for the different data operations and discuss resulting threats that inherently arise at the infrastructure level.

6.4.1 Routing or Mapping Resources

The common view on routing is that of a topological resolution service: Routing guides the paths to hosts. As ICN abandons the host-centric paradigm to address content objects directly, routes to content items attain the role of traditional topological directives.

State and Update Complexity

In ICN, each content item (file) needs retrieval and therefore must be accessible via some resolution service. This may either be implemented by a distributed routing system, or by a mapping service that provides an indirection to topological locators of publishers or content caches. Whenever off-path caching is enabled, the average complexity of the corresponding management operations reads $\langle \#\ of\ content\ items\rangle \cdot \langle \#\ of\ cached\ replica\rangle \cdot \langle update\ frequency\rangle$ ($\langle x\rangle$ denotes average value of x) and must be considered a severe challenge.[2] Solutions that are restricted to on-path caching reduce this complexity to $\langle \#of\ content\ items\rangle \cdot \langle update\ frequency\rangle$. In both cases, the request routing/mapping system is stressed by adding and updating name or – if applicable – cache entries at high frequency, the details of which depend on the implementation of the service.

Cache Announcements

Route maintenance in ICN consists of propagating content publishers (i.e., default paths) as well as cache instances. While the first task is known to generate a high volume of data and frequent updates, caching is expected to largely exceed default announcements in number and update frequency. As a countermeasure, data replication may be limited to caching along default paths, which remarkably reduces the complexity for the routing system. On-path cache replica are met implicitly when requests are routed towards the source. They need not be advertised in the routing or mapping service. On the downside, restricting the caching to default paths will drastically reduce its effectiveness, and a corresponding strategy falls behind today's CDN solutions. Godsi et al. [63] discussed the caching problems in detail. The authors came to the conclusion that on-path caching is merely a warm-up of traditional web proxies.

Route Integrity

ICN, like the current Internet, relies on the integrity of its routing system. A bogus route may block or degrade services, lead to incorrect content delivery, or violate privacy. These core concerns are well-known from BGP [31], where effective countermeasures exist. However, in addition to those vulnerabilities known from BGP routing, threats uniquely arise from data-driven state management in content-centric routing.

The reason for easily implementing malicious routes is inherited from universal caching. An explicit authorization of caches as common in the CDN market is in conflict with open publication and not applicable in general ICN approaches—any node in the ICN network can cache and thus announce any (forged) name. Even if we consider scenarios where only a subset

[2] A global request routing system will need to host at least the amount of the Google index base ($\mathcal{O}(10^{12})$) at a much enhanced update frequency (for timely content access and caching). For comparison, today's DNS subsumes $\mathcal{O}(10^{8})$ names at a very low change rate of $\approx 10^{5}$ alterations per day.

Symbol	Meaning
R_i	The i-th Router
C_i	Capacity of the link between R_i and R_{i+1}
U_i	Utilization of the link between R_i and R_{i+1}
S_i	# of content request states of R_i at its interface towards R_{i+1}
α_i	Content request rate at interface $R_i \rightarrow R_{i+1}$
ω_i	Content arrival rate at interface $R_i \leftarrow R_{i+1}$
T_i	Request timeout at interface $R_i \rightarrow R_{i+1}$
l	Packet length
$\langle \cdot \rangle$	Average value of $\cdot$
$\sigma(\cdot)$	Standard deviation of $\cdot$

Table 6.1: Glossary of Notations

of caches is allowed to distribute content, origin validation measures like RPKI [104] or DNS-based accountability [98] cannot be applied because ICN renounces end point identifiers.

The implication of this problem emerges directly from state maintenance at routers. As the delivery infrastructure is vulnerable to increased delays and delay variations in content supply (see Section 6.4.2), route redirections may be applied to slow down content delivery or to jitter response times to finally harm the complete system. Following the first observation, any intermediate cache can—purposefully or accidentally—threaten its neighborhood.

6.4.2 Forwarding Resources

Traditional routers in the Internet consist of a central processing unit and main memory that are available to the control plane, mainly to learn and determine new routes, as well as FIB memory that is fed by the route selection process. Data forwarding remains bound to FIB lookup and packet processing at line-cards. This design choice purposefully decouples forwarding capacities from control processing and—with equal importance—protects control states from (bogus) data packets.

Current concepts of content-centric data forwarding break with this separation paradigm, and introduce—similar to IP multicast—an additional reverse path forwarding table, also called PIT. Unlike in multicast, this table is updated *packet-wise* on line speed by data-driven events. In the following subsections, we concentrate on the consequences for routing resources in detail. We will consider a chain of routers R_i along a data path and use the notation summarized in Table 6.1.

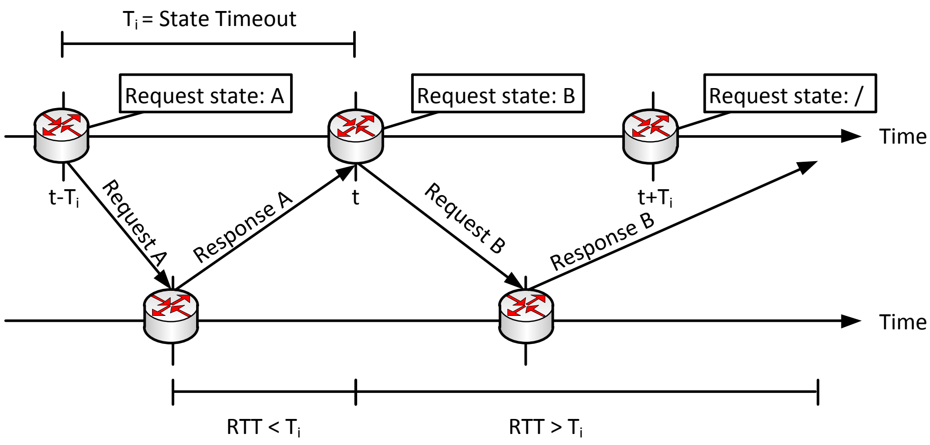

Figure 6.1: Content request states depending on state timeout and round trip time variation

Content Request States versus Content Request Rates versus Network Utilization

Content request states are the essential building block to control flows in a content-centric distribution system that operates hop-by-hop. Each request state will trigger a data packet on return, while the number of open request states corresponds to data arrival at this interface after the transmission time.

We now consider a typical ICN scenario. Router R_i is connected to router R_{i+1} via a point-to-point interface, which is in steady operation. Each router has a predefined request timeout. Note that the timeout is valid for the whole router but activated per interface. For R_i we denote this timeout T_i.

We first want to derive the relation between routing request states at time t and network utilization, where the request timeout is arbitrary. In the absence of request retransmissions, packet loss, and erroneous state removal, the total amount of states increases linearly by newly arriving requests α_i and decreases by content arrivals or timeout ω_i. Hence, the basic rate equation reads

$$S_i(t) = S_i(t - T_i) + \int_{t-T_i}^{t} \alpha_i(\tau) - \omega_i(\tau) d\tau$$

To simplify the term, we denote the time delay of the packet arrival process by $\pi(\cdot)$ and thus derive

$$S_i(t) = S_i(t - T_i) + \int_{t-T_i}^{t} \alpha_i(\tau) - \alpha_i(\pi(\tau)) d\tau$$

Finally, we denote the random variable of packet round trip times by RTT, which is assumed independent of the requests and packet rates, then we derive the mean

$$S_i(t) \quad = \quad \langle \alpha_i \rangle \cdot \min(\langle RTT \rangle, T_i) + \mathcal{O}\left(\sigma(\alpha_i) \cdot \sigma(\min(RTT, T_i))\right) \tag{6.1}$$

From the first part of Equation (6.1), we can immediately deduce that timeout values below the (varying) *RTTs* limit the number of request states, but at the same time will block data forwarding. A second view reveals the strong dependence of routing state on the *RTT* variation, indicated by the error estimation in the $\mathcal{O}$ notation. An example which summarizes the effects is illustrated in Figure 6.1. A similar phenomenon is well-known from TCP [80], but has been overlooked in corresponding previous work on ICN resource considerations [113, 156, 59].

Henceforth we will address the case of data flowing unhindered by the state timeout T_i and assume T_i large enough. Furthermore—for a steady-state scenario—it is assumed that the content request rate fluctuates on a stationary scale. Equation (6.1) then simplifies to

$$\begin{aligned} S_i(t) &\approx \langle \alpha_i \rangle \cdot (\langle RTT \rangle + \kappa\, \sigma(RTT)) && (6.2) \\ &\approx U_i(t)/\langle l \rangle \cdot (\langle RTT \rangle + \kappa\, \sigma(RTT))\,, && (6.3) \end{aligned}$$

with an estimating parameter κ for the mean deviation. The well-known term $(\langle RTT \rangle + \kappa\, \sigma(RTT))$ represents a *retransmission timeout.*[3] For the last step, we roughly assumed that content requests and content arrival are in stationary equilibrium.

Approximation (6.3) yields the desired coupling of the link utilization U_i and the state management resources at a router: On a single point-to-point link without state retransmissions and in flow balance, state requirements are proportional to the network utilization, enhanced by a factor of a *global retransmission timeout.* From a practical point of view this means the following. If an ICN router delivers data but does not introduce interest retransmission on global scale, hardware resources for request states need to be aligned with (a) the interface utilization and (b) with the global delay distribution. At switched interconnects or in bursty communication scenarios, conditions are expected to grow much worse. In more detail, the following observations are noteworthy.

1. Unlike in TCP that estimates a single end-to-end connection, content request states at routers subsume various prefixes and numerous flows. Moreover, content items (prefixes) are explicitly not bound to end points. Thus rapidly varying RTTs are characteristic to interfaces and even to individual flows in content-centric routing. The presence of chunk caching may further increase the *RTT* variation. Hence, no convergent estimator for a round trip time can be reasonably given.

2. In the current Internet, the variation of *RTT* is commonly larger than its average. End-to-end delays are known to approximately follow a heavy-tailed Gamma distribution [25]. PingER [114] reports means and standard deviations of about 250 ms, with maxima up to 5,000 ms. For a constant content request rate of 125k packets/s these RTTs generate the

[3] The corresponding (over-)estimator in TCP is commonly set to 4. However, it is well known that standard TCP algorithms and parameters are inefficient at rapidly changing round trip times, which are characteristic for interface conditions in content-centric routing.

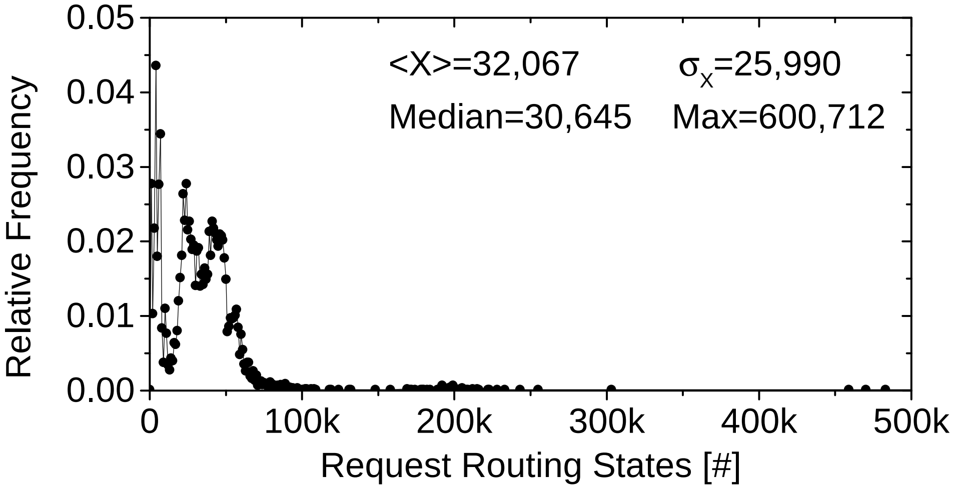

Figure 6.2: Distribution of forwarding states at routers with a 1 Gbit/s link, covering global RTTs that were measured in the PingER project [114] in March 2012

state distribution visualized in Figure 6.2. Consequently, ICN data delivery will suffer remarkably from heterogeneous hardware deployment. To cover worst case scenarios, ICN routers require over-provisioning by more than one order of magnitude compared the average scenarios.

3. Limiting the absolute size of the content request table imposes a strict bound on network utilization. However, the sustained rates are mainly determined by actual RTTs and are hardly predictable. Similar arguments hold for defining timeout values.

4. Applying rate limits to content requests does not change the picture. For an 'on average' optimal limit $C_i \cdot \langle RTT \rangle / \langle l \rangle$, the variation of content replies in time may lead to large over- and under-utilization of network resources that goes along with large fluctuations in request table sizes.

Memory Requirements

A content-centric router that is designed to fully utilize its link capacities, requires sufficient table space for content requests under varying network conditions. Equation (6.3) approximates the corresponding resources when applied to the maximum link capacity C_i. Using the conservative value of $\kappa = 4$ as for TCP, a packet length $l =$ 1,000 bytes, and RTT values from PingER as cited in the previous section, we derive

$$S_i = 1.25\, s/8,000\, bit \cdot C_i \approx 1.6 \cdot 10^{-4} s/bit \cdot C_i \tag{6.4}$$

For a line-speed of 1 to 100 Gbit/s, 160k to 16,000k content request entries then need to be installed per interface at minimum. Compared to previous results [113, 156] our findings are up to one order of magnitude more accurate as we consider *RTT* variation. Still they are merely a rough *lower estimate*, as larger fluctuations of round trip times may significantly increase resource demands.

It is noteworthy that Equation (6.4) holds for any router in a content-centric Internet. Unlike today, where full BGP tables are only required at AS border routers, and interior devices operate on a very small routing table, ICN access routers already demand for a full table memory, the size of which is determined by its interface capacities. In practice, this significantly increases router costs, as any fast interface must co-locate a large block of fast memory.

CPU Load from Table Management

An ICN router maintains states according to user data requests. For any content request, it needs at line speed to (1) insert a state in its request table. On the arrival of any data packet, it needs to (2) search and (3) delete on success in the same table. In addition, a router has to (4) maintain timers of all (soft) states in its request table.

The paradigmatic shift of ICN opens the state table, which is responsible for data delivery, to the end users. Any content requester may create unexpected access patterns. Consequently, to guarantee robustness, an implementation of the request table not only needs to perform dictionary operations very efficiently *on average* but also in worst-case. As there is no final design of an ICN router, we discuss the general design space for hash tables in this context.

The most efficient implementation is the usage of on-chip content-addressable memory (CAM), which is the complete hardware counterpart of a dictionary data structure. However, costs, energy, and limited size prohibit its deployment for expected PIT tables sizes. Typically, state tables for request routing (or caching) are implemented using hash tables. In its elementary form, these data structures can perform all basic dictionary operations (i.e., insert, search, and delete) at constant complexity, but experience hash collisions. Collisions cause a conflict: An implementation that ignores hash collisions will overwrite data. This limits the field of application for ICN, as dealing with collisions increases complexity, or making the system directly threatened by DoS.

In ICN, Perino and Varvello [113], for example, propose to use HC-basic [23], a collision-prone scheme without avoidance mechanisms. Such a design choice makes the ICN system vulnerable to (also purposefully) replacing valid content requests. In realistic deployments of ICN, an implementation of hash tables will either provide mechanisms to prevent or to handle hash collisions.

Essentially four solutions are known to overcome hash collisions: (1) Hash chaining or open addressing, (2) perfect hashing, (3) cryptographic hash functions, and (4) universal hashing.

Hash chaining (i.e., concatenating conflicting keys) or open addressing (i.e., deterministic probing for an alternate location) [46] circumvent collisions on the price of enhanced update costs. The worst case complexity increases to $\mathcal{O}(N)$ for a table of size N. This introduces well-known vulnerabilities, as any pattern that creates collisions will result in such linear complexity instead of amortized $\mathcal{O}(1)$. Crosby and Wallach [47] analyzed this for current software systems (e.g., the Bro IDS). For the more sophisticated and widely deployed hardware hash tables Peacock and Cuckoo, Ben-Porat *et al.* [24] recently studied the structural vulnerability and observed significant performance degradation, as well. Applying these approaches to ICN introduces an obvious threat.

"Perfect hashing" [46] supports constant complexity in the worst-case, but requires a static key set. It is thus not suitable for dynamic content requests that are characteristic in ICN.

Collision resistant cryptographic hash functions can be applied, but lead to a prohibitive increase in memory and CPU consumptions. Enabling cryptographic operations at backbone routers has been discussed continuously in the context of securing BGP [31] and did not succeeded so far. In contrast to requests in ICN, BGP updates occur rarely (even if we consider update storms) and will be sent by known peers. ICN would have to apply cryptographic hashing to all Interest packets. Cryptographic hash functions are—with respect to the packet processing requirements in ICN and current, long-term hardware development—out–of–scope also for a future implementation of the ICN paradigm.

A trade-off between performance and worst-case avoidance can be implemented by universal hashing [35]. It has been introduced to obfuscate the hash function by randomly selecting the hashing algorithm at start-up. Unfortunately, universal hashing is difficult to implement in hardware. Moreover, it still shows a collision probability which is small compared to basic hashing but high compared to cryptographic hash functions. Threats based on fingerprinting make the system additionally vulnerable. It should also be noted that universal hash tables cannot be deployed in a distributed fashion, but are confined to strictly local use cases due to the random selection of hash function. Collaborating request management is thus challenged.

For a detailed overview on state-of-the-art hash tables for high-speed packet processing, we refer to the excellent work by Kirsch *et al.* [86]. Even though wire-speed hash tables are deployed in the current Internet, it is by no means obvious that these approaches provide sufficient robustness in daily ICN operation. Current solutions exhibit serious attack vectors, either by DoS collision overwrite or by a non-constant worst-case performance. In addition, even a constant complexity does not guarantee appropriate robustness as extra memory accesses can prevent wire speed performance.

6.5 Experiments on State-based Forwarding

In this section, we present the results of straight-forward experiments that show the outcome of the core threats as theoretically discussed in Section 6.4. For a detailed presentation of advanced attack scenarios we refer to Section 6.7.

In particular, we concentrate now on system and performance implications of the data-driven state management at infrastructure devices. Even though the measurements mainly relate to the NDN implementation `ccnd`, we should emphasize that we do not evaluate the

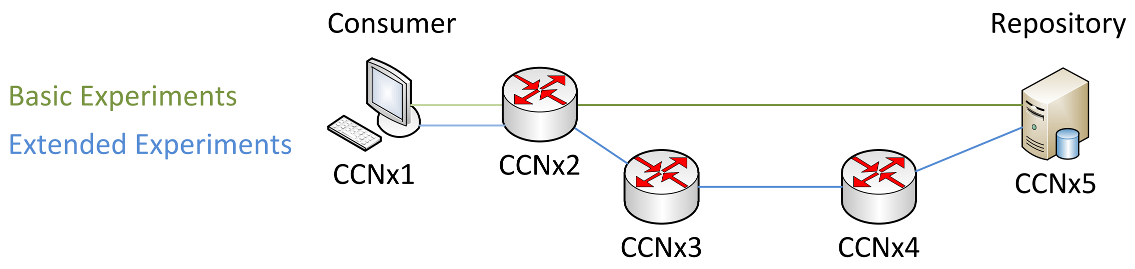

Figure 6.3: Topologies of the experimental settings

implementation itself, but use it as one real-world instance of the information-centric network deployment to illustrate the routing protocol mechanisms. Following this spirit, we do not interpret or discuss absolute performance values, which surely can be improved by optimized software and hardware in the future, but focus on structural and asymptotic analysis.

6.5.1 Core Measurement Setup

In our measurement study, we intentionally deploy *simple* communication scenarios between one content requester and one publisher. The basic network topology is represented by a Daisy chain of directly interlinked CCNx routers with 100 Mbit/s, one end connects the content consumer and the other the content repository (see Fig. 6.3). The *basic topology* consists of two hops and the *extended topology* of five nodes. It is noteworthy that more complex settings, e.g., a Dumbbell topology popular to represent backbone network effects, would enforce the effects, which we already see in our simpler and more transparent examples.

We use the CCNx implementation version 0.5.1 [110], i.e., the client library to announce content Interests, the content repository to store data, and the `ccnd` to forward subscription and data. The following analysis focuses on the effects on the router side. For obtaining a fine-grained view, we concentrate on the local system as well as inter-router dependencies.

We keep default values for all CCNx parameters (i.e., 4770 Byte chunk size and fix pipeline size). In particular, routers do not follow a specific strategy layer, as this would twist robustness towards specific limits as discussed in Section 6.4.2. CCNx routers communicate via TCP (preserving packet order in the basic experiments) or UDP (extended experiments). The convergence times of the different experiments range between three and 18 hours.

6.5.2 Basic Experiments: Resource Consumption

A Fast Path to Resource Exhaustion

An elementary threat intrinsic to data-driven state management arises from the overloading of routers by Interest requests (i.e., Interest flooding). This is most easily provoked by initiating requests for content that does *not* exist. In our scenario, the consumer issues 2,000 Interest messages for *non*-existing content, waits 6 seconds, and repeats these steps until overall 150,000 Interests have been sent.

Figure 6.4 shows the local resource consumptions on the first hop of the content receiver. The number of entries in the Pending Interest Table (PIT), the CPU load, and the required memory increase linearly with subsequent bulks of Interest messages until the system is saturated. In this case, the router reaches its limits of processing and memory resources when storing $\approx$120,000 PIT entries. While sending Interests, the initiating node retransmits previous announcements to keep states fresh at the router. Even though the retransmission timer is below the expiration timer and network delays are very short, the PIT size fluctuates as entries drop due to overloading. After all initial Interest messages have been distributed, the content consumer only retransmits subscriptions.

Our experiment illustrates several problems: A router may easily exhaust PIT space, when content arrives late or not at all. However, even if it was able to store all entries, it would suffer from a 'retransmission only' phase. The retransmissions agglomerate over time and create a continuous stream of signaling that consumes CPU cycles. When the update rate is higher than the processing capabilities permit, retransmissions require buffering, which leads to additional memory overhead (cf., Figure 6.4). A high system load increases the probability of dropping a PIT entry even if its refresh message has been signaled in time. This again causes additional refreshs of the PIT data structure (add/delete calls) and fosters load.

In a recent publication, Yi *et al.* [157] propose to mitigate this threat by signaling content unavailability back to the original requester. Such `NACK` will cure the Interest retransmission effects discussed above for truly unavailable content. However, this workaround has limited effect, as `NACK` suppression introduces a new attack vector at the content supplier side, while a bogus requester can still harm the routing infrastructure (in particular its designated router) by iterating Interest messages over various names of unavailable content.

Chunk-based State Multiplication

To analyze the performance of content consumption, we conduct a bulk file transfer. At this, the content receiver initiates the parallel download of multiple 10 Mbit files over a constant time. We consider three scenarios, the request of 2 files, 10 files, and 100 files per second, which correspond to an underutilized, a fully loaded, and an overloaded link. Figure 6.5 shows the start and completion time of the download per file (top graph), as well as the PIT size, the effective number of Interest retransmissions, and the traffic load including the mean goodput at the first hop. For visibility reasons, we rescaled the y-axis of PI in Figure 6.5(a).

With an increasing number of parallel downloads, not only the download times increase significantly, but also the interval of the request and receive phase grows in the scenarios of (over-)load. While the download time is almost constant for two files per second (cf., Fig. 6.5(a)), the time-to-completion grows non-linearly for the downloads in cases of exces-

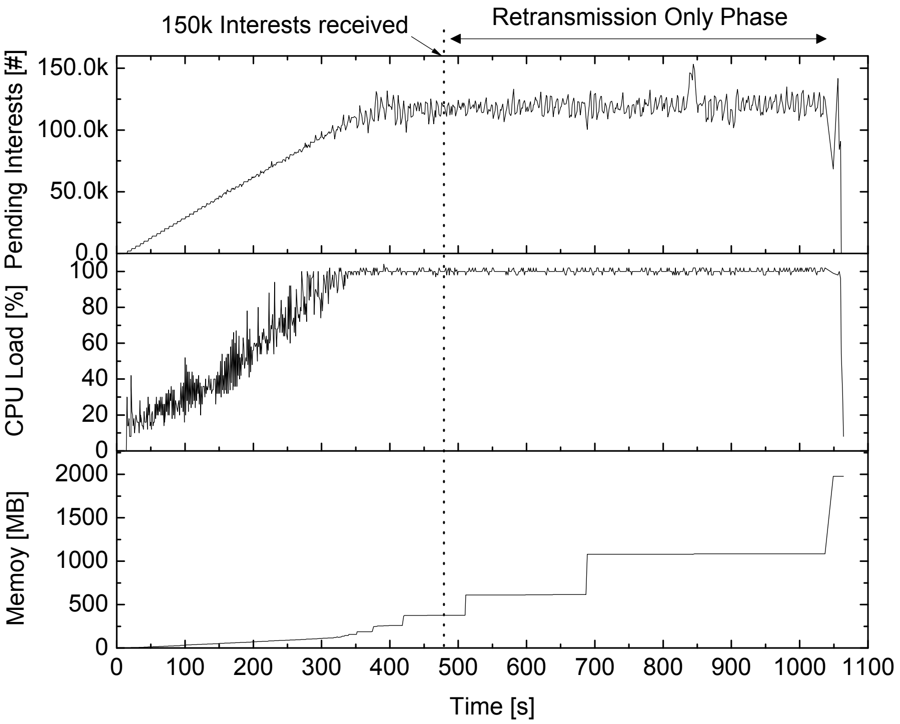

Figure 6.4: Load at the designated router of the receiver while requesting non-existing content

sive parallelism (cf., Fig. 6.5(b),(c)). 150 s are needed to download *each* single file in the worst case (Fig. 6.5(c)), while the link capacity would permit to retrieve *all* files in about 10 s.

The reason for this performance flaw is visualized in the subjacent graphs. A higher download frequency leads to an increasing number of simultaneous PIT entries, which require coordination with the data plane. Each file request will be split into requests of multiple chunks, in which the generation of corresponding Interest messages will be pipelined. In contrast to Section 6.5.2, content exists. As soon as the content traverses, Interest states dissolve and thus release memory. These operations cause a simultaneous burst in CPU load (not shown) and result in growing Interest retransmits after droppings or timeouts (shown in second lowest graphs). This also leads to retransmissions of data chunks. As an overall net effect, the network utilization fluctuates significantly, but does not adapt to actual user demands: Even though data requests could fill the links easily, the average load remains about constant at 30 % of the total network capacity.

In this example we demonstrated that insufficient processing and memory resources will strictly prevent a proper link utilization. This problem cannot be mitigated by rate limiting, as reduced Interest transmission rates will simultaneously reduce network utilization even further (see Section 6.4.2).[4]

The only visible way to assure proper utilization of network resources requires appropriate

[4] We should remind that applying Interest rates in NDN is a mechanism of flow control, and *not* for system

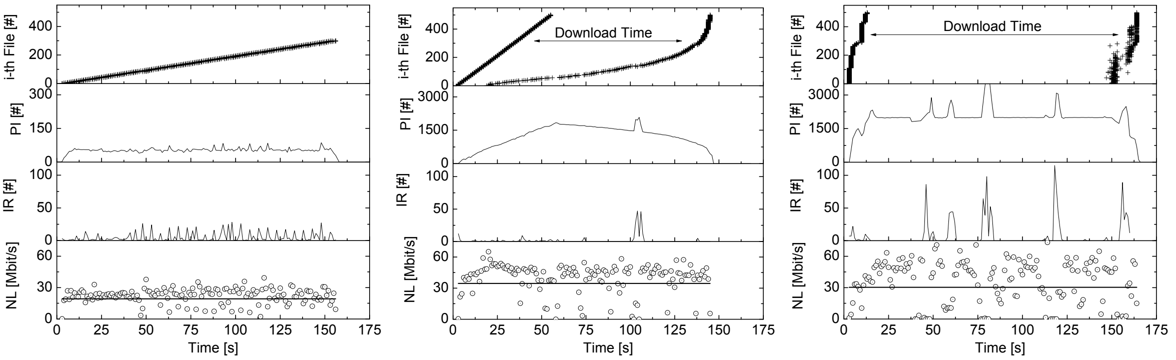

(a) 2 files per second

(b) 10 files per second

(c) 100 files per second

Figure 6.5: Parallel download of 10 Mbit files: Start and stop time of the download per file at the receiver & resource consumption at its designated router [Pending Interests (PI), Interest Retransmits (IR), and Network Load (NL) including the mean goodput (straight line)]

routing resources, i.e., a PI table implementation that is sufficiently large and reliably operates at line speed. As we learned from the analysis in Section 6.4.2, corresponding solutions are not available today. At the current state of the art, an attacker can always reproduce the performance degradations by either blowing up RTT and its variation, or by injecting states that degrade the performance of the PI hash table of the routers.

6.5.3 Extended Experiments: State Propagation and Correlation

In our extended experiments, we take a closer look at hop-by-hop routing performance using the five node routing chain displayed in the lower part of Figure 6.3. Intermediate nodes are numbered from the designated router of the content receiver (first hop) to the router of the content repository (fifth hop). In the following three experiments, we specifically concentrate on correlation effects of the routing resources by controlling the environment using parametrizable virtual machines.

A Homogeneous Network

In this first extended experiment, we simply move our previous picture to the larger topology. All forwarding nodes offer the same resources, two cores@2.4 GHz, 3 GB RAM, and link capacities of 100 Mbit/s. A content requester downloads 500 files of size 10 Mbit at an average rate of 100 files per second. We observe a flattening of Interest propagation towards the source, as states resolve earlier from faster packet delivery (cf., Fig. 6.9(a)).

resource protection. Intermingling these two aspects is likely to produce unwanted performance flaws and leads to new attacks (cf., Section 6.7).

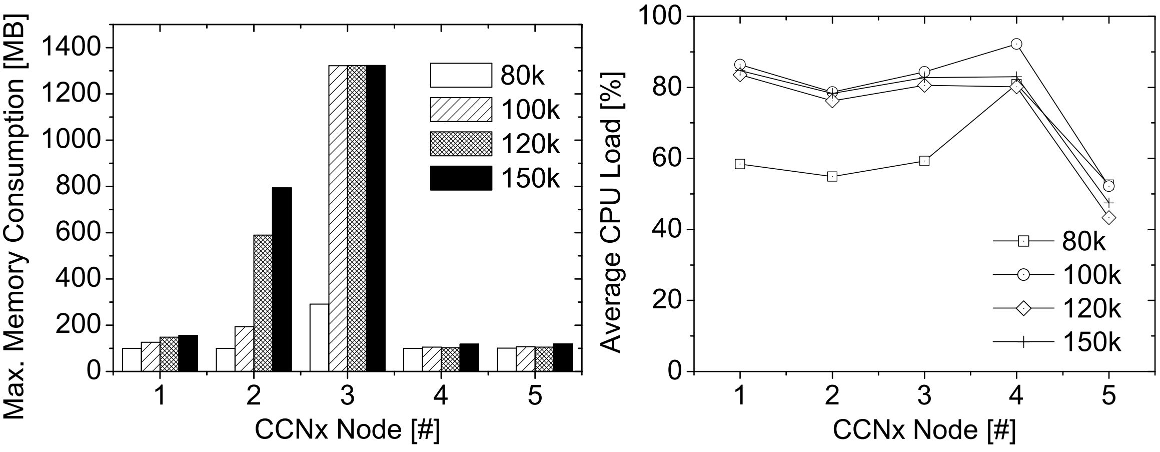

(a) Memory Consumption

(b) Average CPU Load

Figure 6.6: Load per hop for a chain of 5 routers while initiating a 80k, 100k, 120k, and 150k different Interests for non-existing content

A Single Point of Weakness

It is a valid assumption that the content distribution system will consist of heterogeneous devices in terms of all performance metrics. In this second experiment, we introduce device heterogeneity by weakening a single router, the 4th hop (CCNx4), in a controlled way. We want to study the reaction of state management and network performance to this well-defined degradation.

For an initial observation of the dependency on the weakest node, we reduce the CPU capacity of CCNx4 to 25 % (600 MHz) and recap the scenario from Section 6.5.2 for 80k, 100k, 120k, and 150k subscriptions of non-existing content. Independent of the capacity of the network infrastructure, the consumer initiates content subscriptions and continuously refreshes its Interests, which then propagate towards the content repository.

Figure 6.6 shows the maximal memory consumption and the average CPU load per hop during the measurement period. It is clearly visible that the required memory mainly depends on the position of the node within the topology. Memory requirements on the single path fluctuate by two orders of magnitude. The predecessor of the node with the lowest processing capacities (i.e., the 3rd hop) needs 50% – 500% more memory than any other nodes.

We now take a closer look on gradual effects of routing heterogeneity. We observe corrective mechanisms of the network (i.e., Interest retransmissions) depending on router asymmetry. Interest retransmissions serve as the key indicator for timeouts due to router overload. For this task, we configure CCNx4 with four different processing capacities related to the other CCNx routers: 2,400 MHz (homogeneous), 1,200 MHz (50 % capacity), and 600 MHz (25 % capacity).

Surprising results are shown in Figure 6.7. Evidently we see an instability in the forwarding

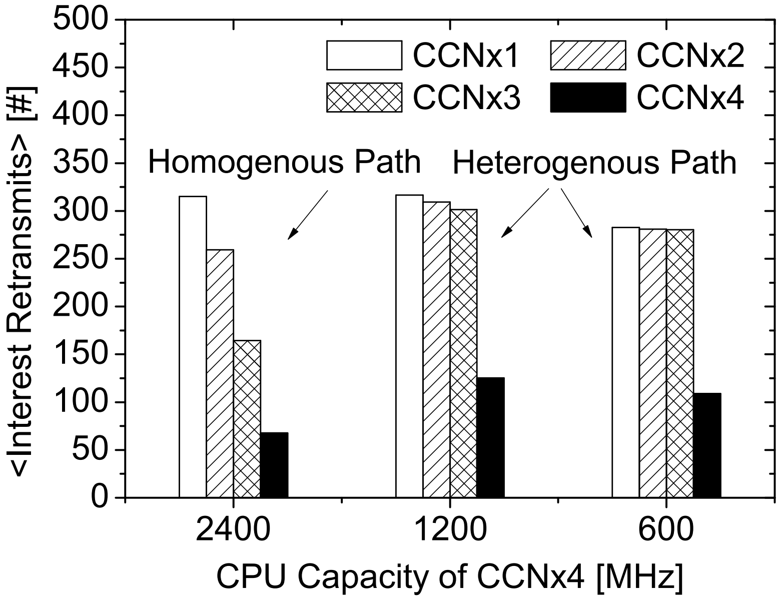

Figure 6.7: Effect of routing heterogeneity on Interest trading

behaviour of the network. The characteristic picture of a balanced network is a steady decay of Interest retransmits towards the source, as data delivery gets faster and more reliable in proximity to the publisher. However, at the first occasion of a ‘bottleneck’—independent of its strength—the picture flips. Interest retransmission drastically increases and all routers except for the bottleneck equally see about the maximal rate of retransmissions in this scenario. State retransmissions at the weak forwarder (CCNx4) instantaneously doubles to the maximal level of managed states this router can cope with.

This experiment clearly shows how sensitive content-centric routing reacts to varying network resources. A light disturbance of the state propagation process reveals the instability of a steady-state flow by immediately turning content transport into a significantly different condition of maximal error management. To compensate this, network operators would need to adjust hardware resources. Doing this carefully is very challenging in a heterogeneous Internet. The rather unpredictable nature of Internet communication in terms of delay would require a high overprovisioning of routers.

Complex Inhomogeneities

In our final experiment concerned with content routing, we explore situations of largely decorrelated network conditions. In a real-world meshed ICN backbone, this may well occur from different side traffics. The objective of this study is to analyse the vulnerability of hop-by-hop state maintenance in ICN routing. Therefore we configure all routers to admit fast changing resources occurring in anti-cycles. In detail, each router (CCNx1, . . . , CCNx5) is forced into a

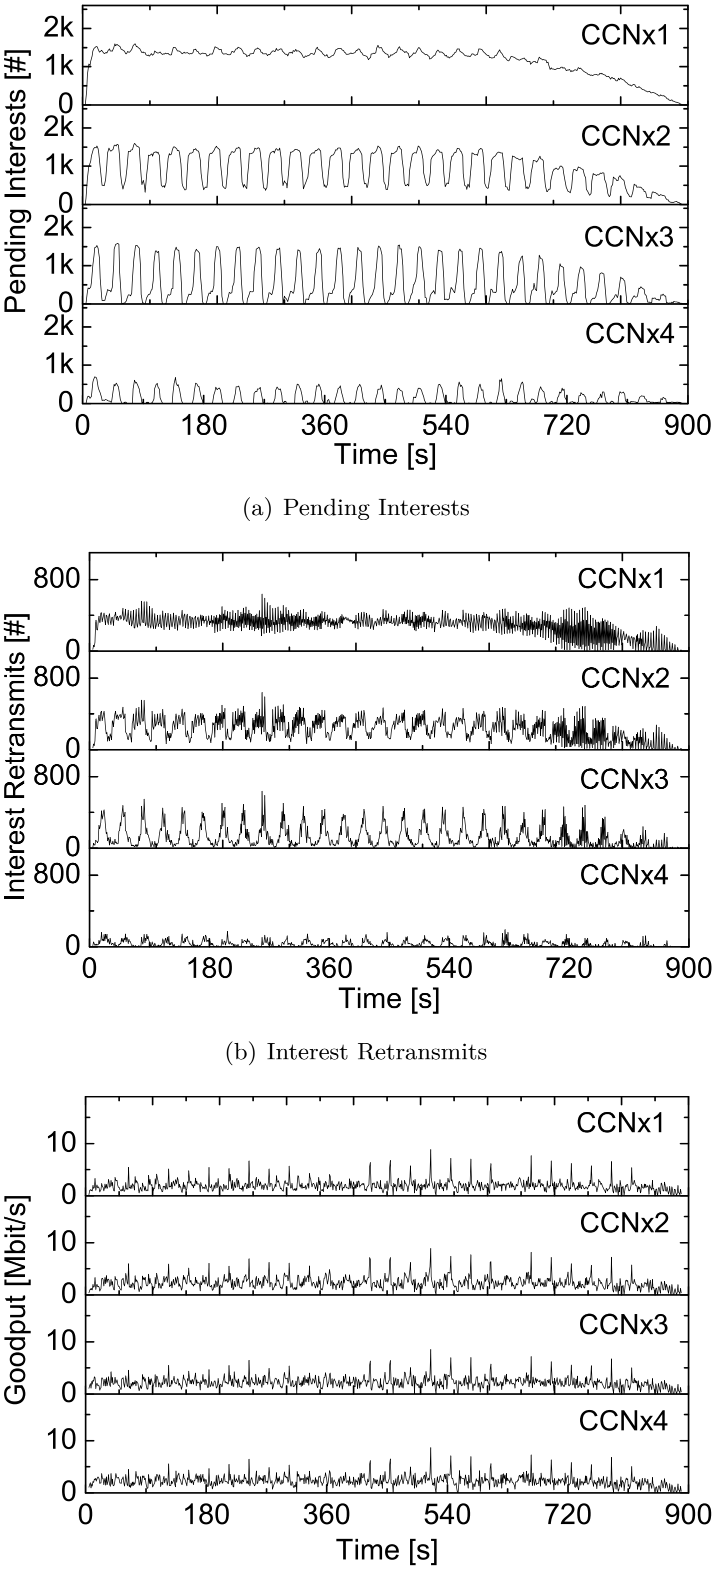

(a) Pending Interests

(b) Interest Retransmits

(c) Network Utilization

Figure 6.8: Vulnerability against side traffic effects: Routing and forwarding performance in a five-hop network with *alternating* CPU reductions among all nodes

10 s periodic CPU reduction by 90 %. Resource reduction periods were shifted between routers at a rate of 10 s so that at least one of the three routers in the forwarding chain was kept in challenged conditions. The objective of this repelling setup, which similarly may well occur from different side traffics in a meshed backbone, is to analyse the vulnerability of hop-by-hop state maintenance in ICN routing.

Results, i.e., Pending Interests, retransmits, and network throughput of this alternating resource scenario are displayed in Figure 6.8. The course of pending Interests as well as Interest retransmissions open a distinguished view on the fine-grained sensitivity of content routing to neighboring router conditions. State provisioning fluctuates on the resource resolution scale of 30 s throughout the network. More importantly, data transmission rates drop down to about 2.4 Mbit/s, while the overall load of Interest states remains compatible to the homogeneous network. Uncoordinated network resource availability thus leads to a low overall performance in conjunction with high network resource consumptions. Time-to-completion for each file download correspondingly explodes to 900 s for the same 10 Mbit files as in our initial experiment. It should be recalled that network capacities do allow for a simultaneous download of all 500 files within 10 s.

A comparative result of the different scenarios in our experimentally-driven analysis is presented in Figure 6.9. We contrast the load imposed onto the infrastructure by Interest states with the average network performance in the three experimental scenarios, homogeneous network, single point of weakness, and alternating resources at routers. The striking picture in all three settings is that the efficiency of network utilization is low on the overall, but drastically drops whenever inhomogeneities occur. The hop-by-hop forwarding performance thus appears rather fragile. In contrast, network state propagation attains various patterns, but always remains at compatible level at the router of maximal load.

These observations suggest the following rule of thumb for CCN routing performance: State maintenance always follows the maximal requirements, while forwarding performance will adapt to the weakest resource in place. This overall picture is clearly inefficient and future work on ICN solutions would largely benefit from improving this behaviour.

6.6 Simulation of Complex Networks

In our previous experimental evaluation, we have concentrated on simple topologies and on an in-depth analysis of individual router behaviour under data-driven state management. We will now focus on the overall performance of complex networks built from real-world topologies that we import into discrete event simulations. It is noteworthy that discrete event simulations do not experience load when managing states, but solely account for the interplay of request-routing and forwarding in the overall networking system.

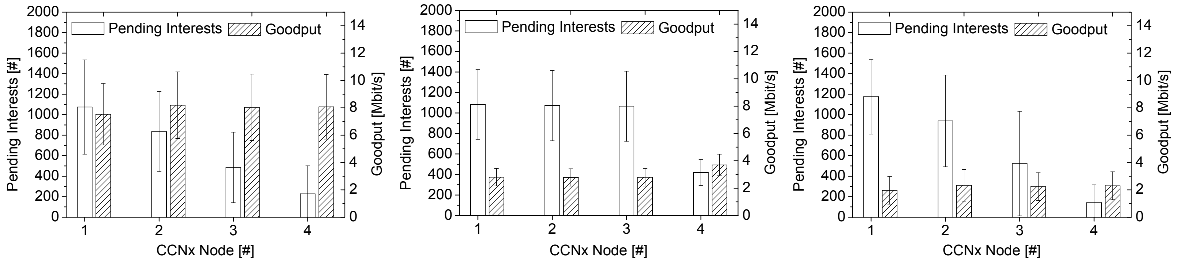

(a) Homogeneous Network (b) Single Point of Weakness (c) Alternating Resources

Figure 6.9: Comparison of state management and forwarding performance in different network scenarios (mean and standard variation)

6.6.1 Simulation Setup

Network simulations are based on ndnSIM [9], release 6 Nov. '12, an NDN implementation for NS-3. Our reference topology is built from the Rocketfuel [101] data set, which is commonly used in the ICN context [38, 162]. In detail, we started from the Sprintlink topology (# 1239) as core network of 315 nodes. These core routers are interconnected by point-to-point links with latencies obtained from the data set and homogeneous bandwidths of 10 Mbit/s. This backbone topology is extended by adding three additional edge nodes at each core router. The connections between core routers and associated edge nodes are via direct links of 1 Mbit/s with a latency of 10 ms. We should emphasize that bandwidths have been assigned with the aim of balancing the network. Absolute values carry no meaning, as we study effects of relative network performance.

Every simulation node is provided with a protocol stack consisting of the link-layer Face *(ndn::NetDeviceFace)*, and the NDN protocol *(ndn::L3Protocol)* implementation. *ndn::BestRoute* implementation is used as Forwarding Strategy, whereas the Content Store module is not in use, and left uninstantiated in our configuration. For analyzing the coherence of states in the complex network, we limit the PIT sizes at all routers to 100. This value corresponds to a balanced utilization of the network core at 10 Mbit/s for the given average delay of 80 ms in our topology. We apply and compare the two PIT replacement strategies *persistent*, which keeps table entries and drops newly arriving requests on overload, and *random*, which randomly replaces entries of a full table.

Communication in the network is between *ndn::Producer* and *ndn::ConsumerCbr* applications. The consumer application issues Interests at a configurable frequency, and thus initiates data transfers. The producer applications are configured to reply with a data packet of 1024 Bytes in response to each arriving Interest that addresses a matching name. In each simulation run, we create 20 producers that are randomly placed either on edge or on core nodes. Regardless of its position, there is at most one producer per node. Consumers are always placed at edge nodes in a random fashion.

The experiments run until they converge (i.e., more than eight hours). It is worth noting that the time average and the ensemble average are the same for the random process under investigation.

During simulation runs, we monitor PIT states, Interest retransmits, and data forwarding at each node. We extract the following metrics: The *maximal numbers of Interest drops and retransmits* per router, which serve as indicators of state decorrelation and network stress, the *overall goodput* jointly attained at all receivers, and the *maximal transfer time* taken over all data chunks that are delivered at a time of measurement. We decided for the analysis of maximal values instead of average values because to illustrate the change of worst-case performance.

6.6.2 Results

In our simulations, we study communication scenarios that are balanced at network edges, while backbone links—like in today's Internet—may be overbooked. Figure 6.10 displays the typical network performance for such a case: 20 producers are placed at core nodes, and 20 consumers per content source request 128 chunks per second for a simulation time of 500 seconds. These data requests saturate the access link capacities of 1 Mbit/s at each receiver and lead to immediate process of dropping and retransmitting Pending Interest at core nodes. It is clearly visible that dropping and retransmission frequencies are lower for the persistent PIT management (Figure 6.10(a)), since request states that have entered routers on a complete path will remain present until data is forwarded. In contrast, the random replacement strategy may erase Interests on an established path and thus has a higher likelihood of decorrelating router states.

The overall forwarding capacity of the simulated scenarios is about 250 Mbit/s. Our simulation experiments both start out with a short peak in network utilization (≈ 50 Mbit/s), but forwarding goodput quickly decays below 5 % of the capacity as soon as PIT overloads occur. Correspondingly, the slowest chunks arrive in times that grow linearly with the simulation time as a result of state decorrelation. Certain Interest states do not succeed in guiding a packet transfer until the backbone stress ceases, while timeouts and retransmits superimpose data forwarding in a mutually obstructive way. We observed the latter behaviour in all of our simulation runs with varying number of producers or request frequency. It is noteworthy that the persistent state strategy at routers avoids a dropping of data packets, since any path already established from receiver to source remains intact for forwarding. As a consequence, the overall data goodput in the network corresponds to successful chunk deliveries, whereas a significant number of data packets (≈ 0.5 per dropped Interest) is lost on path while the random replacement strategy is in operation. The enhanced goodput results in Figure 6.10(b) are caused by undelivered packets, which have only partially passed from the source to the receivers. At first glance, this might be misleading. However, it is worth noting that the ICN architecture is more complex compared to the current Internet. Even if the data did not reach the original receiver,

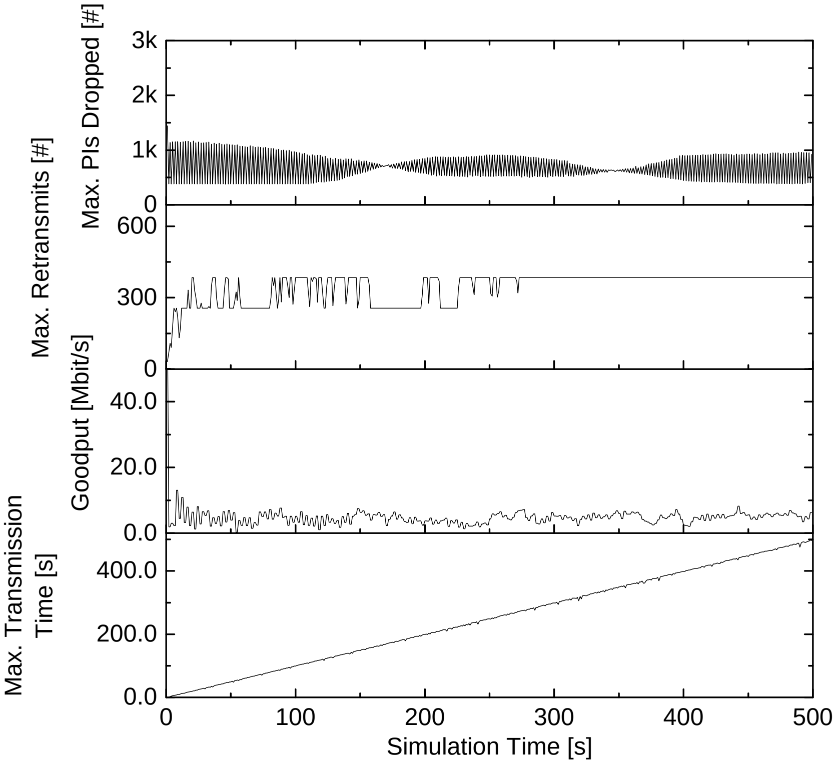

(a) Persistent PIT Management

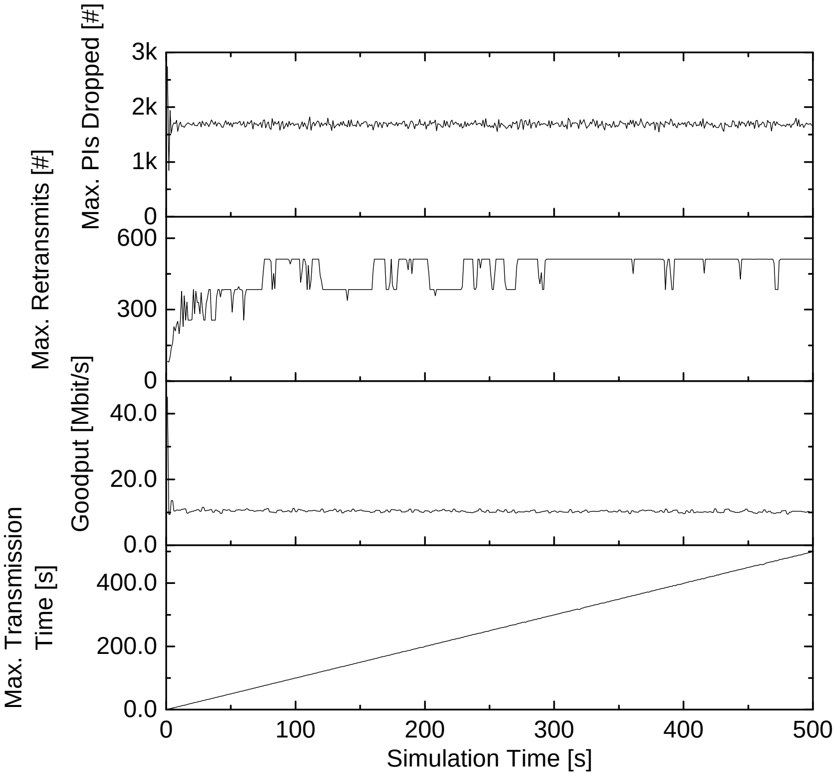

(b) Random PIT Management

Figure 6.10: Parallel chunk download: Stress and performance in a real-world topology

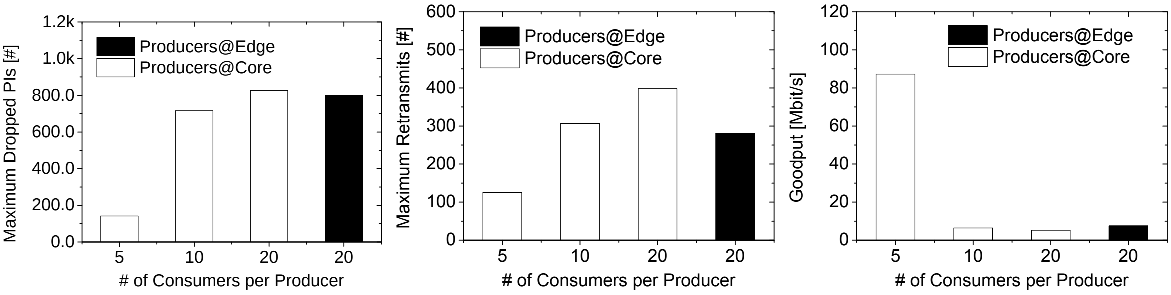

(a) Interest Drops (b) Interest Retransmits (c) Network Goodput

Figure 6.11: Variation of producer position and number of consumers

caching may help subsequent consumers that are located on the path from the source to the original destination.

Next we study the scalability and robustness of network performance with respect to varying capabilities and loads of producers for the case of persistent PIT management. We compare producers attached to the edge with producers in the core at increasing numbers of consumers. The results displayed in Figure 6.11 show a clear dependency of network integrity on request intensity. The more consumers request content services, the stronger the network decorrelates (see Figures 6.11(a,b)). At the same time, the overall network performance drastically decreases as visualized in Figure 6.11(c). This comes at no surprise, as unsatisfied Pending Interests simply accumulate in the network with increasing probability of dropping on its path to the source. Placing producers at the edge hardly changes measurement results, even though the topological network balance is inverted. This confirms the observation that the network operations degrade due to state decorrelation along the path prior to arriving at the source.

Surprising in some parts, these simulation studies reveal a significant dependency of network performance and stress on the use patterns, independent of the specifics of the topology. Even though the discrete event simulations are robust with respect to system resources of router and—unlike in our experiments—performance remains unaffected by state maintenance, content-centric routing tends to not efficiently exploit transmission resources in realistic networks of intermediate size. Instead, our findings indicate that a state-wise uncoordinated hop-by-hop forwarding behaviour has unstable performance and can be easily degraded by adverse use patterns of consumers. This must be seen as a severe threat to the infrastructure, as a flooding of interests by end users can result in a denial of service attack at the remote core of the network.

Related phenomena of resource decorrelation are well known from Bittorrent-like systems [70], where a randomised resource trading often leads to a service degradation dominated by the weakest constituent. This is in contrast to the current Internet, whose paradigm of “Best Effort” actually defines a stable dynamic of maximized resource availability within the networking system. We suggest that future ICN research shall optimize the dynamics of network conditions under varying resource conditions.

6.7 Examples of Attack Scenarios

In this section, we briefly introduce attack scenarios for each threat enumerated in Section 6.2.2. Some attacks are unique to ICN, others—even though known from the Internet—gain a new level of severity by exploiting ICN intrinsics. We concentrate on attackers who exploit vulnerabilities of the infrastructure by generating various subscription resp. publication states as discussed in Sections 6.4 and 6.5.

6.7.1 Attacks Related to Resource Exhaustion

As shown in the previous section, routing and forwarding capacities of the infrastructure can be easily compromised by overloading its content request or interest tables. As non-aggregable name requests for locally unavailable content propagate through the network, resource exhaustion attacks can be transparently initiated from the remote. For this purpose it remains completely indifferent whether hardware resources are drained at unrestricted PIT sizes or table space exhausts according to various limiting configurations. Correspondingly, FIB overflows at routers occur in response to excessive publications or updates of names. Details of the attacker's effects depend on the state-dropping strategy—for simplicity we assume dropping tails as used by CCNx. In addition, virtual resources may also be depleted. The injection of bogus Interests disturbs ICN flow control mechanisms at routers, for example, because it reduces request limits.

Remotely Initiated Overload An attacker that controls one or more machines (a botnet) may initiate massive requests for unavailable content based on Interest flooding (see Section 6.5.2). Corresponding interests propagate towards the publisher and eventually accumulate at some content router causing overload conditions. Depending on its intensity, this attack will lead to a service impairment or DoS for the (remote) content distribution tree(s) branching at the degraded router, unless the networking system is able to re-route the requests. It is worth noting that timeouts at regular users on the subtree will initiate retransmission 'storms' and thereby amplify the attack.

Piling Requests due to a Slow Source Performance of a content source may be degraded by artificially high numbers of direct requests causing slowed down responsiveness. Alternatively, a captured source or its overloaded access router may drastically increase response times of content delivery. In slowing down a (popular) source, an attacker lowers the data return rate and thereby the extinction of pending interests at all routers on paths to receivers. Thus attacking a single point may result in a widely increased load at the routing infrastructure.

Mobile Blockade A mobile node may issue a large number of invalid (or slow) Interests that block the state table of the access router for the period of state timeout. In a shared link-layer environment that cannot easily detect its departure, the mobile adversary can traverse neighboring networks on circular routes and continue to offload its interest bundle with the effect of a blockade of the regionally available networks. Initial countermeasures are difficult to apply, as the retransmission of Interests is part of the regular mobility pattern in ICN.

Fooling Rate Limiting Current ICN approaches [156] propose rate limiting to restrict the number of Interest states. Its main purpose is flow control to avoid congestion. In contrast to common believes (see [59] for discussions) we argue that this is not an appropriate countermeasure to protect the ICN distribution system against attacks. An attacker can easily create an interest storm that exceeds the anticipated interest limit. The dedicated router will throttle the number of accepted interests per interface or interface+prefix, and finally ignore subsequent interests. Consequently, a single end user blocks a prefix or harms all members of its domain. We started analysing these effects in more detail in our on-going work [16].

Note, applying rate limiting per end host (or user) is non-trivial in ICN. ICN explicitly discontinues the concept of host identifiers (e.g., due to security reasons). Thus, a router cannot track particular sources that perform Interest flooding. Assuming end point identifiers exist and routers are enabled with a tracking function, an attacker can spoof addresses. The same holds for push-back mechanisms [59], which signal an overload towards the source and thus try to isolate an attacker. In addition, an attacker that would receive such a control message can ignore it.

6.7.2 Attacks Related to State Decorrelation

ICN requires consistent states during the request routing phase *and* the asynchronous content delivery. While bogus announcements or flapping of routes may introduce loops or increase the likelihood thereof, incoherent forwarding paths may result in partial content transmission that uses network resources without success in data delivery.

Heterogeneity Attack An attacker that controls several machines (e.g., a botnet) may direct requests to accumulate at a specific router in the network and generate a point of performance degradation in the core. Heterogeneity will cause a significant service depletion for all crossing flows (see Section 6.5.3), if the network does not reroute. In the presence of rerouting, the adversary may use the same attack to trigger route flipping with corresponding jitter enhancements, which—in contrast to the Internet—will degrade access router performance for consumers.

Infringing Content States An attacker that controls end systems or content routers could announce updates of content or cache appearances at a high frequency that exceeds the routing convergence time. As a consequence, the overloaded route resolution service will be unable to correctly process the updates of proper content sources or caches with the effect of incomplete content representation and erroneous data replication states. Content requesters will be thus led into false retrievals or access failures. As content announcement is commonly built on soft-state approaches, failures will timeout after a period of undisclosed inconsistency, which the adversary could initiate in a momentary attack.

Jamming Attack A node on a shared link may issue a large number of content requests without maintaining the Interests at its own (loosing interest). Content will then arrive at the local link without a receiver. This is particularly harmful in mobile environments of limited bandwidth. A mobile attacker can jam a region by traversing shared radio links while requesting bulk data.

6.7.3 Attacks Related to Path and Name Infiltration

ICN raises content names and cache locations to first class objects and must therefore remain open to naming and placing data. The request routing system carries routes to names in its FIB or a mapping service, both of which are vulnerable to resource exhaustion and route poisoning. While an explosion in the pure number of names may be mitigated in parts by aggregation according to some authoritative naming conventions like in today's domain names, bogus route infiltration must be considered the more delicate issue.

Route-to-Death An adversary that controls a cache system may redirect routes to it and slow down content delivery or jitter response times. As the routing infrastructure is vulnerable to increased delays and delay variations, resource exhaustion threats apply to the requesting infrastructure (see Section 6.4.2). In the presence of universal caching, reasonable counter measures to using a valid, but alienating cache are difficult to define.

Route Set Inflation An adversary may announce bogus routes to cached copies of any content object. Content requests from its vicinity are then directed towards an erroneous location and—if unanswered or retarded— lead to long-lasting forwarding states and a possible DoS. This threat can be mitigated by resource-intensive attempts to route towards multiple locations that become increasingly painful when an attacker controls a botnet and injects invalid routes at large scale.

6.8 Conclusion and Outlook

In this chapter, we have analyzed network instabilities and threats in information-centric networks that are caused by (a) backbone control states initiated by end users and (b) data-driven state management.

Some threats are easy to anticipate (e.g., resource exhaustion), others are more intricate due to the complex interplay of distributed management (e.g., state decorrelation). For the latter previous practical insights in the design of (conceptually related) multicast protocols already revealed good and bad design options. One of the major design goals of Bidirectional PIM [71], for example, was "eliminating the requirement for data-driven protocol events"—after the operating experiences with data-driven DVMRP or PIM-SM. With these results, we want to stimulate the discussion about basic security in content-centric backbone routing.

As an insightful view onto the stability of the ICN infrastructure, we contrast our findings with the vulnerability of the current Internet and recall basic countermeasures from a high level perspective.

In the current Internet, an attack may target at the Internet backbone, i.e., BGP, as well as at end hosts. BGP is vulnerable in many aspects [43]. Bringing BGP routers to their knees is complicated and rare [44], [160], [36], [129]—a collateral damage from the data to the control plane based on sharing links. ICN eases such attacks by allowing immediate control changes from the data plane. Malicious path or name infiltration is similar to IP prefix hijacking, the core attack on BGP. Countermeasures introduce a cryptographically strong binding (e.g., [85], [104], [98]) between the (legitimate) advertiser and the announced resource. ICN democratizes content caching by design, consequently any ICN node is allowed to announce a cache copy and thus to implement a routing state. Proposals for the current Internet will not help to overcome route interception, route set inflation, or route-to-death in ICN.

Today, (D)DoS attacks are usually directed towards end hosts. In this chapter, we have shown that ICN extends these threats to the backbone by design, and that existing countermeasures against both, DDoS and incorrect distribution states fail in the ICN field.

Defending from DDoS is already complicated in the Internet and becomes more intricate in ICN. From the conceptual perspective, the core challenge is not in deploying accountability (e.g., [135], [19]) but identifying an attack. Attack detection approaches [111] usually make application specific assumptions about traffic patterns, which cannot be applied to a generic Internet service for content delivery. We showed that the very fluctuating Internet delay space challenges resource provision in ICN (cf., Section 6.4.2). As content states will accumulate in the network (cf., Section 6.5), and inter-provider deployment almost surely will lead to a heterogeneous, unbalanced design, rate limiting may milden, but cannot effectively prevent the resource exhaustion problems discussed in this work.

Current CDN deployments remain agnostic of these infringements by running under proprietary regimes. Present ICN proposals do not seem to have taken up the battle of standing in the wild. Very recently countermeasures have been proposed to mitigate Interest flooding attacks [48], [42], [150], [8]. In future work, we will systematically compare countermeasures against Interest flooding. We will continue our first steps [16] not only to cover more complex application scenarios but also to complement simulations by long-range real-world experiments.

Furthermore, we will investigate hybrid approaches, which combine different detection points (per interface and per prefix) to increase accuracy and performance while limiting PIT entries.

Finally, we should note that in an open Internet, threats are built on the worst scenarios, not on average cases. If we want an information-centric Internet to remain open and reliable, a major redesign of its core architecture appears inevitable.

Chapter 7

Summary

In this thesis, we started from the observation that the Internet is a critical infrastructure, which needs severe protection. We took a practical view on Internet security, considering the whole ecosystem including the network, end devices, and services, protected and threatened by current and future Internet protocols. We contributed tools, methodologies, and measurement results to improve the current state of art as well as operational practice.

Chapter 2 presented a methodology to derive a nation-centric view on the Internet. This methodology has been exemplified for Germany. Using publicly available registry data and careful manual verification, we identified important ASes for Germany. Applying passive BGP measurements we analyzed then the interconnections between these ASes. Surprisingly there were only few direct interconnections between ASes of the same business sector, visible from public monitor points. Local peering increased over the last years but should be more widely applied.

Although the Internet is completely decentralized and a multi-stakeholder network, which on the first glance does not allow for country and business-specific classification of their autonomous systems, there is an increasing demand to continue studying this. Our analysis showed that it is crucial to start from the less aggregated set of Internet resources. Network providers operate autonomously. They split an original assigned address space among multiple customers, which may be located in different countries. Considering IP blocks instead of IP prefixes allowed us to identify 25% more ASes.

Chapter 3 analyzed the deployment of the RPKI from the router perspective. We introduced the first full-fledged open-source implementation of the RPKI router protocol that is suitable in real-world deployment. This library enables now RPKI in two of the most popular open-source BGP daemons, Quagga and BIRD.

Using this tool, we conducted a first detailed analysis of the local system performance when a BGP router validates prefix. The additional load introduced by RPKI is negligible on commodity hardware. The choice of the local data structure to store verification data is important yet. Implementation choices usually should consider deployment scenarios. We found that the

CPU load may increase up to one order of magnitude if the updates only lead to invalids. Still, commodity hardware is able to deal with this.

After testing that the overhead of RPKI origin validation does not prevent from deploying, we analyzed the validation of real-world BGP updates. We found a surprisingly high amount of invalids, in particular during the beginning of RPKI deployment. Further analysis revealed that for 90% of the invalid updates the authorized AS is the direct upstream of the incorrect origin. This clearly is not an attack on the BGP control plane as the legitimate AS could easily filter the invalid announcement. Based on an advanced analysis we introduced heuristics to identify common pitfalls when deploying the RPKI.

Chapter 4 extended the scope from a pure backbone perspective towards the relationship between network security and content distribution services deployed in the current Web. We conducted the first quantitative analysis of the deployment of RPKI by web providers. Analysing 1 million web domains we found that 6% of the web server prefixes are secured by the RPKI and the remaining 94% are unprotected. Furthermore, there is a correlation that popular sites tend to be not protected at all as they more likely served by CDNs. Only 4 entries exist in the RPKI for the 199 ASes of the studied CDNs.

To improve the security experience for web users, we presented a web browser extension which verifies the prefix origin AS of the requested web server infrastructure, and shows the state intuitively to the user. The verification is currently based on the data of predefined BGP vantage points but can be extended to consider control plane data from the actual upstream of the users. It is worth noting that we lack standard mechanisms for such a location-dependent view. The currently developed BGP Monitoring Protocol [130] as well as appropriate approximation schemes may solve this problem in the future. However, our design and implementation is first step that helps to verify that content is reliably and securely delivered.

Chapter 5 explored the attack on end devices via the Internet even further by separating remote attacks based on the network access. Due to the increasing relevance and evolution of mobiles, we focused on this regime. We designed a honeypot that efficiently reflects characteristics of mobile systems. We then measured unsolicited network traffic by deploying four honeypots (i.e., a mobile honeypot, a darknet honeypot, a DSL honeypot, and a university honeypot) to analyze attack behavior. Surprisingly, we did not monitored a significant number of attacks dedicated to mobiles. This indicates that most attackers operate almost independently of the actually captured host, or at least assumed a heterogeneous set of systems.

Chapter 6 took a radical different view. Instead of protecting the current Internet infrastructure, we analyzed the potentials of a future Internet architecture. Information-centric networks inherently prevent DoS attacks towards end devices as they only address content instead of hosts. However, we found that this design choice will harm the infrastructure instead. By enabling end users to directly change states of the control plane, they can easily perform a DoS attack on the core. Independently of an explicit attacker, we showed that the robustness of an ICN network is in conflict with heterogeneous, (unpredictable) environments such as the Internet.

Future Work We discussed detailed aspects of future work in the separate sections. From a high-level perspective, we will focus on both improving existing security solutions to increase deployment and the design of network architectures with inherent protection mechanisms, without conflicting with the openness of the Internet.

List of Figures

List of Tables

Bibliography

[1] "Internet AS-level topology construction & analysis." [Online]. Available: http://topology.neclab.eu/

[2] "Marktanteile der führenden Anbieter von Breitbandinternet in Deutschland vom 1. Quartal 2011 bis zum 4. Quartal 2014 ." [Online]. Available: http://de.statista.com/statistik/daten/studie/196770/umfrage/marktanteile-der-fuehrenden-breitband-anbieter-in-deutschland/

[3] "MaxMind – GeoLite Country." [Online]. Available: http://www.maxmind.com/app/geoip_country

[4] "RIPE Routing Information Service (RIS)." [Online]. Available: http://www.ripe.net/projects/ris/rawdata.html

[5] "Team Cymru." [Online]. Available: http://www.cymru.com/

[6] "Graphviz - Graph Visualization Software," http://www.graphviz.org/, 2010.

[7] D. E. 3rd, "Domain Name System Security Extensions," IETF, RFC 2535, March 1999.

[8] A. Afanasyev, P. Mahadevan, I. Moiseenko, E. Uzun, and L. Zhang, "Interest Flooding Attack and Countermeasures in Named Data Networking," in *Proc. of IFIP Networking*. Piscataway, NJ, USA: IEEE Press, 2013.

[9] A. Afanasyev, I. Moiseenko, and L. Zhang, "ndnSIM: NDN simulator for NS-3," NDN, Technical Report NDN-0005, October 2012. [Online]. Available: http://www.named-data.net/techreport/TR005-ndnsim.pdf

[10] S. Agarwal, C.-N. Chuah, S. Bhattacharyya, and C. Diot, "Impact of BGP Dynamics on Router CPU Utilization," in *Proc. of Passive and Active Network Measurement*, ser. LNCS, C. Barakat and I. Pratt, Eds., vol. 3015. Berlin Heidelberg: Springer, 2004, pp. 278–288.

[11] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger, "Anatomy of a Large European IXP," in *Proc. of the ACM SIGCOMM*. New York, NY, USA: ACM, 2012, pp. 163–174.

[12] B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig, "Web Content Cartography," in *Proc. of ACM SIGCOMM IMC*. New York, NY, USA: ACM, 2011, pp. 585–600.

[13] B. Ahlgren *et al.*, "Second NetInf Architecture Description," 4Ward EU FP7 Project, Tech.report D-6.2 v2.0, 2010.

[14] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, "A Survey of Information-Centric Networking," *IEEE Communications Magazine*, vol. 50, no. 7, pp. 26–36, July 2012.

[15] N. Al-Gharabally, N. El-Sayed, S. Al-Mulla, and I. Ahmad, "Wireless Honeypots: Survey and Assessment," in *Proceedings of the 2009 conference on Information Science, Technology and Applications (ISTA '09).* New York, NY, USA: ACM, 2009, pp. 45–52.

[16] S. Al-Sheikh, M. Wählisch, and T. C. Schmidt, "Revisiting Countermeasures Against NDN Interest Flooding," in *2nd ACM Conference on Information-Centric Networking (ICN 2015), Poster Session.* New York: ACM, Oct. 2015, pp. 195–196.

[17] Alexa Internet Inc., "Top 1M Sites." [Online]. Available: http://s3.amazonaws.com/alexa-static/top-1m.csv.zip

[18] M. Allman, V. Paxson, and J. Terrell, "A Brief History of Scanning," in *Proc. of the ACM IMC.* New York, NY, USA: ACM, 2007, pp. 77–82.

[19] D. G. Andersen, H. Balakrishnan, N. Feamster, T. Koponen, D. Moon, and S. Shenker, "Accountable Internet Protocol (AIP)," in *Proc. of the ACM SIGCOMM.* New York, NY, USA: ACM, 2008, pp. 339–350.

[20] D. P. Anderson, J. Cobb, E. Korpela, M. Lebofsky, and D. Werthimer, "SETI@Home: An Experiment in Public-resource Computing," *Communications of the ACM*, vol. 45, no. 11, pp. 56–61, Nov. 2002.

[21] S. Arianfar, P. Nikander, and J. Ott, "On Content-Centric Router Design and Implications," in *Proc. of ReARCH workshop.* New York, NY, USA: ACM, 2010.

[22] R. Austein, "rpki.net," http://subvert-rpki.hactrn.net/trunk/, 2012.

[23] A. Badam, K. Park, V. S. Pai, and L. L. Peterson, "HashCache: Cache Storage for the Next Billion," in *Proc. of USENIX NSDI.* Berkeley, CA, USA: USENIX Assoc., 2009, pp. 123–136.

[24] U. Ben-Porat, A. Bremler-Barr, H. Levy, and B. Plattner, "On the Vulnerability of Hardware Hash Tables to Sophisticated Attacks," in *Proc. of IFIP Networking*, ser. LNCS, vol. 7289. Berlin, Heidelberg: Springer–Verlag, 2012, pp. 135–148.

[25] C. J. Bovy, H. T. Mertodimedjo, G. Hooghiemstra, H. Uijterwaal, and P. V. Mieghem, "Analysis of End to end Delay Measurements in Internet," in *Proc. of the Passive and Active Measurement Workshop-PAM*, March 2002.

[26] M. A. Brown, "Pakistan hijacks YouTube – Renesys Blog," February 2008. [Online]. Available: http://www.renesys.com/blog/2008/02/pakistan-hijacks-youtube-1.shtml

[27] S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, and S. Havlin, "Catastrophic cascade of failures in interdependent networks," *Nature*, vol. 464, no. 7291, pp. 1025–1028, 2010.

[28] R. Bush, R. Austein, K. Patel, H. Gredler, and M. Waehlisch, "Resource Public Key Infrastructure (RPKI) Router Implementation Report," IETF, RFC 7128, February 2014.

[29] R. Bush and R. Austein, "The RPKI/Router Protocol," IETF, Internet-Draft – work in progress 26, February 2012.

[30] R. Bush, R. Austein, K. Patel, H. Gredler, and M. Waehlisch, "RPKI Router Implementation Report," IETF, Internet-Draft – work in progress 01, January 2012.

[31] K. Butler, T. Farley, P. McDaniel, and J. Rexford, "A Survey of BGP Security Issues and Solutions," *Proc. of the IEEE*, vol. 98, no. 1, pp. 100–122, January 2010.

[32] X. Cai, J. Heidemann, B. Krishnamurthy, and W. Willinger, "Towards an AS-to-Organization Map," in *Proc. of the 10th ACM IMC*. New York, NY, USA: ACM, 2010, pp. 199–205.

[33] S. Calzavara, G. Tolomei, M. Bugliesi, and S. Orlando, "Quite a Mess in My Cookie Jar!: Leveraging Machine Learning to Protect Web Authentication," in *Proc. of the 23rd WWW*. New York, NY, USA: ACM, 2014, pp. 189–200.

[34] G. Carle, J. Schiller, S. Uhlig, W. Willinger, and M. Wählisch, Eds., *The Critical Internet Infrastructure (Dagstuhl Seminar 13322)*, vol. 3, no. 8. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2013.

[35] J. L. Carter and M. N. Wegman, "Universal classes of hash functions," *Journal of Computer and System Sciences*, vol. 18, no. 2, pp. 143–154, 1979.

[36] L. Cavedon, C. Kruegel, and G. Vigna, "Are BGP Routers Open to Attack? An Experiment," in *Proc. Int. Conf. on Open Research Problems in Network Security*. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 88–103.

[37] H. Chan, D. Dash, A. Perrig, and H. Zhang, "Modeling Adoptability of Secure BGP Protocol," in *Proc. ACM SIGCOMM*. New York, NY, USA: ACM, 2006, pp. 279–290.

[38] J. Chen, M. Arumaithurai, X. Fu, and K. K. Ramakrishnan, "G-COPSS: A Content Centric Communication Infrastructure for Gaming Applications," in *Proc. of IEEE ICDCS*. Los Alamitos, CA, USA: IEEE Computer Society, 2012, pp. 355–365.

[39] Y.-J. Chi, R. Oliveira, and L. Zhang, "Cyclops: The AS-level Connectivity Observatory," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 5, pp. 5–16, 2008.

[40] Y. Chung, "Distributed Denial of Service is a Scalability Problem," *ACM SIGCOMM CCR*, vol. 42, no. 1, pp. 69–71, 2012.

[41] Cisco, "BGP: Frequently Asked Questions," http://www.cisco.com/image/gif/paws/5816/bgpfaq_5816.pdf, 2012.

[42] A. Compagno, M. Conti, P. Gasti, and G. Tsudik, "Poseidon: Mitigating Interest Flooding DDoS Attacks in Named Data Networking," ArXiv e-prints, Tech. Rep. 1303.4823, March 2013.

[43] S. Convery, "An Attack Tree for the Border Gateway Protocol," IETF, Internet-Draft – work in progress 01, September 2003.

[44] S. Convery and M. Franz, "BGP Vulnerability Testing: Separating Fact from FUD v1.1," in *NANOG28 Meeting 2003 and Black Hat USA 2003*, June/July 2003.

[45] D. Cooper, E. Heilman, K. Brogle, L. Reyzin, and S. Goldberg, "On the Risk of Misbehaving RPKI Authorities," in *Proc. of HotNets–XII*. New York, NY, USA: ACM, 2013.

[46] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms - 2nd Edition*. MIT Press, 2001.

[47] S. A. Crosby and D. S. Wallach, "Denial of Service via Algorithmic Complexity Attacks," in *Proc. of USENIX Security Symposium*. Berkeley, CA, USA: USENIX Assoc., 2003, pp. 29–44.

[48] H. Dai, Y. Wang, J. Fan, and B. Liu, "Mitigate DDoS Attacks in NDN by Interest Traceback," in *Proc. of IEEE INFOCOM NOMEN Workshop*. Piscataway, NJ, USA: IEEE Press, 2013.

[49] C. Dannewitz, J. Golićlic, B. Ohlman, and B. Ahlgren, "Secure Naming for a Network of Information," in *Proc. of the IEEE Global Internet Symposium*. Piscataway, NJ, USA: IEEE, 2010.

[50] X. Dimitropoulos, D. Krioukov, G. Riley, and K. Claffy, "Revealing the Autonomous System Taxonomy: The Machine Learning Approach," in *Proc. of the PAM Conf. 2006*, M. Allman and M. Roughan, Eds. Web, 2006, pp. 91–100. [Online]. Available: http://www.pamconf.net/2006/papers/pam06-proceedings.pdf

[51] J. Dischler, "Building for the next moment," http://adwords.blogspot.de/2015/05/building-for-next-moment.html, Google, Google's official blog for news, tips and information on AdWords, May 2015.

[52] K. Dooley and I. Brown, *Cisco IOS Cookbook*, 2nd ed. Sebastopol, USA: O'Reilly, 2006.

[53] M. Faloutsos, P. Faloutsos, and C. Faloutsos, “On Power-Law Relationships of the Internet Topology,” in *Proc. of ACM SIGCOMM’99.* New York, NY, USA: ACM Press, 1999, pp. 251–262.

[54] A. P. Felt, M. Finifter, E. Chin, S. Hanna, and D. Wagner, “A Survey of Mobile Malware in the Wild,” in *Proc. of ACM CCS Workshop SPSM.* New York, NY, USA: ACM, 2011, pp. 3–14.

[55] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach *et al.*, “Hypertext Transfer Protocol – HTTP/1.1,” IETF, RFC 2616, June 1999.

[56] N. Fotiou, G. F. Marias, and G. C. Polyzos, “Publish–Subscribe Internetworking Security Aspects,” in *Trustworthy Internet*, N. Blefari-Melazii, G. Bianchi, and L. Salgarelli, Eds. Heidelberg: Springer, 2011, pp. 3–15.

[57] Frederic Cambus, “Multilocation DNS Looking Glass.” [Online]. Available: http://www.dns-lg.com/

[58] L. Gao, “On Inferring Autonomous System Relationships in the Internet,” *IEEE/ACM Trans. Netw.*, vol. 9, no. 6, pp. 733–745, 2001.

[59] P. Gasti, G. Tsudik, E. Uzun, and L. Zhang, “DoS and DDoS in Named-Data Networking,” ArXiv e-prints, Tech. Rep. 1208.0952, August 2012.

[60] P. Gasti, G. Tsudik, E. Uzun, and L. Zhang, “DoS and DDoS in Named Data Networking,” in *Proc. of ICCCN.* IEEE, 2013, pp. 1–7.

[61] A. Gavrichenkov, “Breaking HTTPS with BGP Hijacking,” in *Black Hat. Briefings*, 2015. [Online]. Available: https://www.blackhat.com/us-15/briefings.html

[62] A. Ghodsi, T. Koponen, J. Rajahalme, P. Sarolahti, and S. Shenker, “Naming in Content-oriented Architectures,” in *Proceedings of the ACM SIGCOMM workshop on Information-centric networking*, ser. ICN ’11. New York, NY, USA: ACM, 2011, pp. 1–6.

[63] A. Ghodsi, S. Shenker, T. Koponen, A. Singla, B. Raghavan, and J. Wilcox, “Information-Centric networking: Seeing the Forest for the Trees,” in *Proc. of the 10th ACM HotNets Workshop*, ser. HotNets-X. New York, NY, USA: ACM, 2011.

[64] P. Gill, M. Schapira, and S. Goldberg, “Let the Market Drive Deployment: A Strategy for Transitioning to BGP Security,” in *Proc. of ACM SIGCOMM.* New York, NY, USA: ACM, 2011, pp. 14–25.

[65] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford, “How secure are secure interdomain routing protocols,” in *Proc. of ACM SIGCOMM.* New York, NY, USA: ACM, 2010, pp. 87–98.

[66] F. Gont and S. Bellovin, "Defending against Sequence Number Attacks," IETF, RFC 6528, February 2012.

[67] R. Govindan and A. Reddy, "An Analysis of Internet Inter-Domain Topology and Route Stability," in *Proc. of the IEEE INFOCOM'97*. Washington, DC, USA: IEEE Computer Society, 1997, pp. 850–857.

[68] M. Gritter and D. R. Cheriton, "An Architecture for Content Routing Support in the Internet," in *Proc. USITS'01*. Berkeley, CA, USA: USENIX Association, 2001, pp. 4–4.

[69] C. Guo, H. J. Wang, and W. Zhu, "Smart-Phone Attacks and Defenses," in *Proc. of HotNets–III*. New York, NY, USA: ACM, 2004.

[70] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang, "Measurements, Analysis, and Modeling of BitTorrent-like Systems," in *Pro. of 5th ACM SIGCOMM conference on Internet Measurement (IMC)*. Berkeley, CA, USA: USENIX Association, 2005, pp. 4–4.

[71] M. Handley, I. Kouvelas, T. Speakman, and L. Vicisano, "Bidirectional Protocol Independent Multicast (BIDIR-PIM)," IETF, RFC 5015, October 2007.

[72] P. Hoffman and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA," IETF, RFC 6698, August 2012.

[73] honeynet Project, "The Honeynet Project Chinese Chapter Status Report (Period Apr 2007 to Dec 2008)," 2009. [Online]. Available: http://www.honeynet.org/node/336

[74] X. Hu and Z. M. Mao, "Accurate Real-time Identification of IP Prefix Hijacking," in *Proceedings of the 2007 IEEE Symposium on Security and Privacy (SP '07)*. Washington, DC, USA: IEEE Computer Society, 2007, pp. 3–17.

[75] C. Huang, A. Wang, J. Li, and K. W. Ross, "Measuring and Evaluating Large-scale CDNs," in *Proc. of the ACM IMC*. New York, NY, USA: ACM, 2008, pp. 15–29.

[76] L.-S. Huang, A. Rice, E. Ellingsen, and C. Jackson, "Analyzing Forged SSL Certificates in the Wild," in *Proc. of the IEEE Symposium on Security and Privacy*. Los Alamitos, CA, USA: IEEE Computer Society, 2014.

[77] G. Huston, "Resource Certification," *The Internet Protocol Journal*, vol. 12, no. 1, pp. 13–26, March 2009.

[78] D. Iamartino, C. Pelsser, and R. Bush, "Measuring BGP route origin registration validation," in *Proc. of PAM*, ser. LNCS. Berlin: Springer, 2015, pp. 28–40.

[79] "Internet Routing Registry," http://www.irr.net, 2010.

[80] V. Jacobson, “Congestion Avoidance and Control,” *SIGCOMM Comput. Commun. Rev.*, vol. 18, no. 4, pp. 314–329, August 1988.

[81] V. Jacobson, D. K. Smetters, J. D. Thornton, and M. F. Plass, “Networking Named Content,” in *Proc. of the 5th Int. Conf. on emerging Networking EXperiments and Technologies (ACM CoNEXT’09).* New York, NY, USA: ACM, Dec. 2009, pp. 1–12.

[82] P. Jokela, A. Zahemszky, C. E. Rothenberg, S. Arianfar, and P. Nikander, “LIPSIN: Line Speed Publish/Subscribe Inter-networking,” in *Proc. of the ACM SIGCOMM 2009.* New York, NY, USA: ACM, 2009, pp. 195–206.

[83] T. Kamada and S. Kawai, “An Algorithm for Drawing General Undirected Graphs,” *Inf. Process. Lett.*, vol. 31, no. 1, pp. 7–15, 1989.

[84] J. Karlin, S. Forrest, and J. Rexford, “Nation-State Routing: Censorship, Wiretapping, and BGP,” arXiv.org/CoRR, Tech. Rep. abs/0903.3218, March 2009.

[85] S. Kent, C. Lynn, and K. Seo, “Secure Border Gateway Protocol (S-BGP),” *Selected Areas in Communications, IEEE Journal on*, vol. 18, no. 4, pp. 582–592, 2000.

[86] A. Kirsch, M. Mitzenmacher, and G. Varghese, “Hash-Based Techniques for High-Speed Packet Processing,” in *Algorithms for Next Generation Networks*, G. Cormode and M. Thottan, Eds. London: Springer-Verlag, 2010, pp. 181–218.

[87] S. Knight, H. X. Nguyen, O. Maennel, I. Phillips, N. J. G. Falkner, R. Bush *et al.*, “An Automated System for Emulated Network Experimentation,” in *Proc. of ACM CoNEXT.* New York, NY, USA: ACM, 2013, pp. 235–246.

[88] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker *et al.*, “A Data-Oriented (and beyond) Network Architecture,” *SIGCOMM Computer Communications Review*, vol. 37, no. 4, pp. 181–192, 2007.

[89] B. Krishnamurthy, “Mohonk: MObile Honeypots to Trace Unwanted Traffic Early,” in *Proc. of the ACM SIGCOMM Workshop on Network Troubleshooting (NetT).* New York, NY, USA: ACM, 2004, pp. 277–282.

[90] W. Kumari and K. Sriram, “Recommendation for Not Using AS_SET and AS_CONFED_SET in BGP,” IETF, RFC 6472, December 2011.

[91] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, “Internet Inter-domain Traffic,” in *Proc. of the ACM SIGCOMM ’10.* New York, NY, USA: ACM, 2010, pp. 75–86.

[92] C. Labovitz, G. R. Malan, and F. Jahanian, “Internet Routing Instability,” *IEEE/ACM Trans. Netw.*, vol. 6, no. 5, pp. 515–528, Oct. 1998.

[93] L. Labs, “Origin Validation Looking Glass,” 2014. [Online]. Available: http://www.labs.lacnic.net/rpkitools/looking_glass/

[94] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang, “PHAS: A Prefix Hijack Alert System,” in *Proc. of the 15th Conference on USENIX Security Symposium.* Berkeley, CA, USA: USENIX Assocication, 2006.

[95] M. Lad, X. Zhao, B. Zhang, D. Massey, and L. Zhang, “Analysis of BGP Update Surge During Slammer Worm Attack,” in *Proc. of IWDC*, ser. LNCS, vol. 2918. Berlin, Heidelberg: Springer-Verlag, 2003, pp. 66–79.

[96] T. Lauinger, “Security & Scalability of Content-Centric Networking,” Master’s thesis, TU Darmstadt, Darmstadt, Germany, 2010.

[97] M. Lepinski and S. Kent, “An Infrastructure to Support Secure Internet Routing,” IETF, RFC 6480, February 2012.

[98] A. Li, X. Liu, and X. Yang, “Bootstrapping Accountability in the Internet We Have,” in *Proc. of the 8th NSDI.* Berkeley, CA, USA: USENIX Association, 2011.

[99] R. Lychev, S. Goldberg, and M. Schapira, “Bgp security in partial deployment: Is the juice worth the squeeze?” in *Proc. of ACM SIGCOMM.* New York, NY, USA: ACM, 2013, pp. 171–182.

[100] P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos, K. C. Claffy *et al.*, “The Internet AS-Level Topology: Three Data Sources and One Definitive Metric,” *ACM SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 1, pp. 17–26, January 2006.

[101] R. Mahajan, N. Spring, D. Wetherall, and T. Anderson, “Inferring link weights using end-to-end measurements,” in *Proc. of the 2nd ACM SIGCOMM Workshop on Internet measurment (IMW’02)*, ser. IMW ’02. ACM, 2002, pp. 231–236.

[102] D. Massey, H. Yan, M. Strizhov *et al.*, “BGPmon. Next generation BGP Monitor,” http://bgpmon.netsec.colostate.edu/, 2012.

[103] S. Meiling, T. C. Schmidt, and T. Steinbach, “On Performance and Robustness of Internet-Based Smart Grid Communication: A Case Study for Germany,” in *6th IEEE Int. Conf. on Smart Grid Communications (SmartGridComm’15)*, Nov. 2015.

[104] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein, “BGP Prefix Origin Validation,” IETF, Internet-Draft – work in progress 10, October 2012.

[105] C. Mulliner, S. Liebergeld, and M. Lange, “Poster: HoneyDroid - Creating a Smartphone Honeypot,” 2011, poster at IEEE Security & Privacy.

[106] NIST, "BGP Secure Routing Extension (BGP-SRx)," http://www-x.antd.nist.gov/bgpsrx/, 2012.

[107] NIST, "Global Prefix/Origin Validation using RPKI," 2014. [Online]. Available: http://www-x.antd.nist.gov/rpki-monitor/

[108] T. OConnor and B. Sangster, "honeyM: A Framework for Implementing Virtual Honeyclients for Mobile Devices," in *Proc. of the third ACM WiSec.* New York, NY, USA: ACM, 2010, pp. 129–138.

[109] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "The (in)Completeness of the Observed Internet AS-level Structure," *IEEE/ACM Trans. Netw.*, vol. 18, no. 1, pp. 109–122, 2010.

[110] PARC, "The CCNx Homepage," http://www.ccnx.org/, 2012.

[111] T. Peng, C. Leckie, and K. Ramamohanarao, "Survey of Network-Based Defense Mechanisms Countering the DoS and DDoS Problems," *ACM Comput. Surv.*, vol. 39, no. 1, 2007.

[112] S. Perez, "Behind The Scenes Of The iPhone 5 Jailbreak," *Techcrunch*, 2013. [Online]. Available: http://techcrunch.com/2013/01/21/behind-the-scenes-of-the-iphone-5-jailbreak/

[113] D. Perino and M. Varvello, "A Reality Check for Content Centric Networking," in *Proc. of the ACM SIGCOMM WS on Information-centric Networking (ICN '11).* New York, NY, USA: ACM, 2011, pp. 44–49.

[114] "PingER. Ping end–to–end reporting," http://www-iepm.slac.stanford.edu/pinger/, 2012.

[115] N. Provos and T. Holz, *Virtual Honeypots. From Botnet Tracking to Intrusion Detection*, 2nd ed. Upper Saddle River, NJ: Addison–Wesley, 2008.

[116] Z. Qian and Z. M. Mao, "Off-Path TCP Sequence Number Inference Attack – How Firewall Middleboxes Reduce Security," in *Proc. of the IEEE Symposium on Security and Privacy.* Los Alamitos, CA, USA: IEEE Computer Society, 2012, pp. 347–361.

[117] Z. Qian, Z. M. Mao, and Y. Xie, "Collaborative TCP Sequence Number Inference Attack — How to Crack Sequence Number Under a Second," in *Proc. of ACM CCS.* New York, NY, USA: ACM, 2012, pp. 593–604.

[118] S. Ratnasamy, S. Shenker, and S. McCanne, "Towards an Evolvable Internet Architecture," in *Proc. of ACM SIGCOMM.* New York, NY, USA: ACM, 2005, pp. 313–324.

[119] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," IETF, RFC 4271, January 2006.

[120] A. Reuter, M. Wählisch, and T. C. Schmidt, "RPKI MIRO: Monitoring and Inspection of RPKI Objects," in *Proc. of ACM SIGCOMM, Demo Session.* New York: ACM, August 2015, pp. 107–108. [Online]. Available: http://dx.doi.org/10.1145/2785956.2790026

[121] "RIPE NCC. RIR Trust Anchor Statistics," https://www.ripe.net/lir-services/resource-management/certification/rir-trust-anchor-statistics, 2012.

[122] RIPE NCC, "Making Better Routing Decisions Through RPKI Validation ," http://www.ripe.net/lir-services/resource-management/certification/making-better-routing-decisions-through-rpki-validation, 2012.

[123] H. Roberts and D. Larochelle, "Mapping Local Internet Control," Berkman Center, Harvard University, Tech. Rep., 2010. [Online]. Available: http://cyber.law.harvard.edu/netmaps/mlic.pdf

[124] F. Roesner, T. Kohno, and D. Wetherall, "Detecting and Defending Against Third-party Tracking on the Web," in *Proc. of the 9th USENIX NSDI.* Berkeley, CA, USA: USENIX Association, 2012.

[125] M. Roughan, W. Willinger, O. Maennel, D. Perouli, and R. Bush, "10 Lessons from 10 Years of Measuring and Modeling the Internet's Autonomous Systems," *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 9, pp. 1810–1821, 2011.

[126] J. H. Saltzer, D. P. Reed, and D. D. Clark, "End-to-End Arguments in System Design," *ACM Trans. Comput. Syst.*, vol. 2, no. 4, pp. 277–288, Nov 1984.

[127] J. Schlamp, G. Carle, and E. W. Biersack, "A Forensic Case Study on AS Hijacking: the Attacker's Perspective," *SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 2, pp. 5–12, April 2013.

[128] K. Schomp, T. Callahan, M. Rabinovich, and M. Allman, "On Measuring the Client-Side DNS Infrastructure," in *Proc. of ACM IMC.* New York, NY, USA: ACM, 2013.

[129] M. Schuchard, E. Vasserman, A. Mohaisen, D. Foo Kune, N. Hopper, and Y. Kim, "Losing Control of the Internet: Using the Data Plane to Attack the Control Plane," in *Proc. of Network & Distributed System Security Symposium (NDSS).* Internet Society, Feb. 2011.

[130] J. Scudder, R. Fernando, and S. Stuart, "BGP Monitoring Protocol," IETF, Internet-Draft – work in progress 16, November 2015.

[131] Y. Shavitt and E. Shir, "DIMES: Let the Internet Measure Itself," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 5, pp. 71–74, 2005.

[132] R. Shirey, “Internet Security Glossary,” IETF, RFC 2828, May 2000.

[133] C. A. Shue, A. J. Kalafut, and M. Gupta, “The Web is Smaller Than It Seems,” in *Proc. of ACM IMC*. New York, NY, USA: ACM, 2007, pp. 123–128.

[134] R. Siles, “HoneySpot: The Wireless Honeypot. Monitoring the Attacker’s Activities in Wireless Networks. A design and architectural overview,” The Spanish Honeynet Project, Research Project, December 2007. [Online]. Available: http://honeynet.org.es/papers/honeyspot/HoneySpot_20071217.pdf

[135] D. R. Simon, S. Agarwal, and D. A. Maltz, “AS-Based Accountability as a Cost-effective DDoS Defense,” in *Proc. of Workshop on Hot Topics in Understanding Botnets*. Berkeley, CA, USA: USENIX Association, 2007.

[136] D. M. Slane, C. Bartholomew *et al.*, “2010 Report to Congress,” U.S.–China Economic and Security Review Commission, Annual Report, November 2010. [Online]. Available: http://www.uscc.gov/annual_report/2010/annual_report_full_10.pdf

[137] A.-J. Su, D. R. Choffnes, A. Kuzmanovic, and F. E. Bustamante, “Drafting behind akamai: Inferring network conditions based on cdn redirections,” *IEEE/ACM Trans. Netw.*, vol. 17, no. 6, pp. 1752–1765, 2009.

[138] S. Sundaresan, S. Burnett, N. Feamster, and W. de Donato, “BISmark: A Testbed for Deploying Measurements and Applications in Broadband Access Networks,” in *Proc. of USENIX Annual Technical Conference*. Berkeley, CA, USA: USENIX Assoc., 2014, pp. 383–394.

[139] Surfnet, “RPKI Dashboard,” 2014. [Online]. Available: http://rpki.surfnet.nl/

[140] P. Traynor, M. Lin, M. Ongtang, V. Rao, T. Jaeger, P. McDaniel *et al.*, “On Cellular Botnets: Measuring the Impact of Malicious Devices on a Cellular Network Core,” in *Proc. of ACM CCS*. New York, NY, USA: ACM, 2009, pp. 223–234.

[141] M. Wählisch, F. Holler, T. C. Schmidt, and J. H. Schiller, “RTR-lib: An Open-Source Library in C for RPKI-based Prefix Origin Validation,” in *Proc. of USENIX Security Workshop CSET’13*. Berkeley, CA, USA: USENIX Assoc., 2013. [Online]. Available: https://www.usenix.org/conference/cset13/rtrlib-open-source-library-c-rpki-based-prefix-origin-validation

[142] M. Wählisch, O. Maennel, and T. C. Schmidt, “Towards Detecting BGP Route Hijacking using the RPKI,” in *Proc. of ACM SIGCOMM, Poster Session*. New York: ACM, August 2012, pp. 103–104. [Online]. Available: http://conferences.sigcomm.org/sigcomm/2012/paper/sigcomm/p103.pdf

[143] M. Wählisch, R. Schmidt, T. C. Schmidt, O. Maennel, S. Uhlig, and G. Tyson, “RiPKI: The Tragic Story of RPKI Deployment in the Web Ecosystem,” in *Proc. of Fourteenth ACM Workshop on Hot Topics in Networks (HotNets)*. New York: ACM, Nov. 2015, pp. 11:1–11:7. [Online]. Available: http://dx.doi.org/10.1145/2834050.2834102

[144] 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 the AS Level,” in *13th Passive and Active Measurement Conference (PAM)*, ser. LNCS, vol. 7192. Berlin Heidelberg: Springer-Verlag, 2012, pp. 200–210.

[145] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, “Backscatter from the Data Plane — Threats to Stability and Security in Information-Centric Networking,” Open Archive: arXiv.org, Technical Report arXiv:1205.4778, 2012. [Online]. Available: http://arxiv.org/abs/1205.4778

[146] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, “Bulk of Interest: Performance Measurement of Content-Centric Routing,” in *Proc. of ACM SIGCOMM, Poster Session*. New York: ACM, August 2012, pp. 99–100. [Online]. Available: http://conferences.sigcomm.org/sigcomm/2012/paper/sigcomm/p99.pdf

[147] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, “Backscatter from the Data Plane – Threats to Stability and Security in Information-Centric Network Infrastructure,” *Computer Networks*, vol. 57, no. 16, pp. 3192–3206, Nov. 2013. [Online]. Available: http://dx.doi.org/10.1016/j.comnet.2013.07.009

[148] M. Wählisch, S. Trapp, C. Keil, J. Schönfelder, T. C. Schmidt, and J. Schiller, “First Insights from a Mobile Honeypot,” in *Proc. of ACM SIGCOMM, Poster Session*. New York: ACM, August 2012, pp. 305–306. [Online]. Available: http://conferences.sigcomm.org/sigcomm/2012/paper/sigcomm/p305.pdf

[149] M. Wählisch, A. Vorbach, C. Keil, J. Schönfelder, T. C. Schmidt, and J. H. Schiller, “Design, implementation, and operation of a mobile honeypot,” Open Archive: arXiv.org, Technical Report arXiv:1205.4778, 2013. [Online]. Available: http://arxiv.org/abs/1301.7257

[150] K. Wang, H. Zhou, H. Luo, J. Guan, Y. Qin, and H. Zhang, “Detecting and mitigating interest flooding attacks in content-centric network,” *Security and Communication Networks*, vol. 7, no. 4, pp. 685–699, April 2013.

[151] L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin *et al.*, “Observation and Analysis of BGP Behavior Under Stress,” in *Proc. of the 2nd ACM SIGCOMM Workshop on Internet measurement*, ser. IMW ’02. New York, NY, USA: ACM, 2002, pp. 183–195.

[152] R. Winter, "Modeling the Internet Routing Topology – In Less than 24h," in *Proc. of the 2009 ACM/IEEE/SCS 23rd Workshop on Principles of Advanced and Distributed Simulation (PADS '09).* Washington, DC, USA: IEEE Computer Society, 2009, pp. 72–79.

[153] W. Wong and P. Nikander, "Secure Naming in Information-centric Networks," in *Proc. of Re-Architecting the Internet Workshop (ReARCH '10).* New York, NY, USA: ACM, 2010, pp. 12:1–12:6.

[154] L.-C. Wuu, T.-J. Liu, and K.-M. Chen, "A longest prefix first search tree for IP lookup," *Computer Networks*, vol. 51, no. 12, pp. 3354–3367, August 2007.

[155] H. Yan, R. Oliveira, K. Burnett, D. Matthews, L. Zhang, and D. Massey, "BGPmon: A real-time, scalable, extensible monitoring system," in *Proc. of the 2009 Cybersecurity Applications & Technology Conference for Homeland Security (CATCH'09).* Washington, DC, USA: IEEE Computer Society, 2009, pp. 212–223.

[156] C. Yi, A. Afanasyev, I. Moiseenko, L. Wang, B. Zhang, and L. Zhang, "A Case for Stateful Forwarding Plane," PARC, Tech. Rep. NDN-0002, July 2012.

[157] C. Yi, A. Afanasyev, L. Wang, B. Zhang, and L. Zhang, "Adaptive Forwarding in Named Data Networking," *SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 3, pp. 62–67, 2012.

[158] B. Zhang, R. Liu, D. Massey, and L. Zhang, "Collecting the Internet AS-level Topology," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 1, pp. 53–61, 2005.

[159] L. Zhang, D. Estrin, J. Burke, V. Jacobson, and J. D. Thornton, "Named Data Networking (NDN) Project," NDN, Tech.report ndn-0001, 2010.

[160] Y. Zhang, Z. M. Mao, and J. Wang, "Low-Rate TCP-Targeted DoS Attack Disrupts Internet Routing," in *Proc. of Network & Distributed System Security Symposium (NDSS).* Internet Society, Mar. 2007.

[161] S. Zhou, G.-Q. Zhang, and G.-Q. Zhang, "Chinese Internet AS-Level Topology," *IET Communications*, vol. 1, no. 2, pp. 209–214, 2007.

[162] Z. Zhu, C. Bian, A. Afanasyev, V. Jacobson, and L. Zhang, "Chronos: Serverless Multi-User Chat Over NDN," NDN, Technical Report NDN-0008, Oct. 2012.

[163] E. Zmijewski, "Indonesia Hijacks the World – Renesys Blog," May 2014. [Online]. Available: http://www.renesys.com/2014/04/indonesia-hijacks-world/

[164] T. Zseby and kc claffy, "Workshop report: darkspace and unsolicited traffic analysis (DUST 2012)," *SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 5, pp. 49–53, Sep. 2012.

Appendix A

Publications of the Author (Last 5 Years)

The following list includes publications authored or co-authored, edited or co-edited by the author of this thesis. The work has been published in peer-reviewed journals, conference or workshop proceedings, as technical report, or as IETF/IRTF RFC. For a complete list of publications, we refer to `http://www.inf.fu-berlin.de/~waehl/publications/index.html`.

[1] M. Wählisch, R. Schmidt, T. C. Schmidt, O. Maennel, S. Uhlig, and G. Tyson, "RiPKI: The Tragic Story of RPKI Deployment in the Web Ecosystem," in *Proc. of Fourteenth ACM Workshop on Hot Topics in Networks (HotNets).* New York: ACM, 2015.

[2] S. Al-Sheikh, M. Wählisch, and T. C. Schmidt, "Revisiting Countermeasures Against NDN Interest Flooding," in *Proc. of 2nd ACM Conference on Information-Centric Networking (ICN). Poster Session.* New York: ACM, 2015, pp. 195–196.

[3] T. C. Schmidt, S. Wölke, N. Berg, and M. Wählisch, "Partial Adaptive Name Information in ICN: PANINI Routing Limits FIB Table Sizes," in *Proc. of 2nd ACM Conference on Information-Centric Networking (ICN). Poster Session.* New York: ACM, 2015, pp. 193–194.

[4] G. Pellegrino, C. Rossow, F. J. Ryba, T. C. Schmidt, and M. Wählisch, "Cashing out the Great Cannon? On Browser-Based DDoS Attacks and Economics," in *Proc. of 9th USENIX Security Workshop on Offensive Technologies (WOOT).* Berkeley, CA, USA: USENIX Assoc., 2015.

[5] M. Wählisch and T. C. Schmidt, "See How ISPs Care: An RPKI Validation Extension for Web Browsers," in *Proc. of ACM SIGCOMM, Demo Session.* New York: ACM, August 2015, pp. 115–116.

[6] A. Reuter, M. Wählisch, and T. C. Schmidt, "RPKI MIRO: Monitoring and Inspection of RPKI Objects," in *Proc. of ACM SIGCOMM, Demo Session.* New York: ACM, August 2015, pp. 107–108.

[7] T. Markmann, T. C. Schmidt, and M. Wählisch, "Federated End-to-End Authentication for the Constrained Internet of Things using IBC and ECC," in *Proc. of ACM SIGCOMM, Poster Session.* New York: ACM, August 2015, pp. 603–604.

[8] J. Schlamp, J. Gustafsson, M. Wählisch, T. C. Schmidt, and G. Carle, "The Abandoned Side of the Internet: Hijacking Internet Resources When Domain Names Expire," in *Proc. of 7th International Workshop on Traffic Monitoring and Analysis (TMA)*, ser. LNCS, M. Steiner, P. Barlet-Ros, and O. Bonaventure:, Eds., vol. 9053. Heidelberg: Springer-Verlag, 2015, pp. 188–201.

[9] H. Petersen, M. Lenders, M. Wählisch, O. Hahm, and E. Baccelli, "Old Wine in New Skins? Revisiting the Software Architecture for IP Network Stacks on Constrained IoT Devices," in *Proc. of ACM MobiSys. IoT-Sys WS.* New York: ACM, 2015.

[10] P. Rosenkranz, M. Wählisch, E. Baccelli, and L. Ortmann, "A Distributed Test System Architecture for Open-source IoT Software," in *Proc. of ACM MobiSys. IoT-Sys WS.* New York: ACM, 2015.

[11] R. Hiesgen, D. Charousset, T. C. Schmidt, and M. Wählisch, "Programming Actors for the Internet of Things," *Ercim News*, vol. 101, pp. 25–26, April 2015.

[12] E. Baccelli, C. Mehlis, O. Hahm, T. C. Schmidt, and M. Wählisch, "Information Centric Networking in the IoT: Experiments with NDN in the Wild," in *Proc. of 1st ACM Conference on Information-Centric Networking (ICN).* New York: ACM, 2014, pp. 77–86.

[13] F. Jäger, T. C. Schmidt, and M. Wählisch, "How Dia-Shows Turn Into Video Flows: Adapting Scalable Video Communication to Heterogeneous Network Conditions in Real-Time ," in *Proc. of the 39th Annual IEEE Conference on Local Computer Networks (LCN'14).* Piscataway, NJ, USA: IEEE Press, 2014, pp. 218–226.

[14] M. Vallentin, D. Charousset, T. C. Schmidt, V. Paxson, and M. Wählisch, "Native Actors: How to Scale Network Forensics," in *Proc. of ACM SIGCOMM. Demo session.* New York: ACM, 2014, pp. 141–142.

[15] T. C. Schmidt, S. Wölke, and M. Wählisch, "Peer my Proxy – A Performance Study of Peering Extensions for Multicast in Proxy Mobile IP Domains," in *Proc. of 7th IFIP Wireless and Mobile Networking Conference (WMNC 2014).* Piscataway, NJ, USA: IEEE Press, May 2014.

[16] O. Hahm, E. Baccelli, H. Petersen, M. Wählisch, and T. C. Schmidt, "Demonstration Abstract: Simply RIOT – Teaching and Experimental Research in the Internet of Things," in *Proc. of the 13th ACM/IEEE International Symposium on Information Processing in Sensor Networks (IPSN).* Piscataway, NJ, USA: IEEE Press, 2014, pp. 329–330.

[17] G. Bartl, L. Gerhold, and M. Wählisch, "Towards a theoretical framework of acceptance for surveillance systems at airports," in *Proc. of 11th International Conference on Information Systems for Crisis Response and Management (ISCRAM)*, S. R. Hiltz, M. S. Pfaff, L. Plotnick, and P. C. Shih, Eds. The Pennsylvania State University, USA, 2014, pp. 299–303. [Online]. Available: http://iscram2014.ist.psu.edu/sites/default/files/misc/proceedings/p180.pdf

[18] E. Baccelli, O. Hahm, and M. Wählisch, "Spontaneous Wireless Networking to Counter Pervasive Monitoring," in *Proc. of W3C/IAB workshop on Strengthening the Internet Against Pervasive Monitoring (STRINT)*, 2014. [Online]. Available: https://www.w3.org/2014/strint/papers/26.pdf

[19] A. Förster, C. Sommer, T. Steinbach, and M. Wählisch, Eds., *Proceedings of the 1st OMNeT++ Community Summit, Hamburg, Germany, September 2, 2014*, no. arXiv:1409.0093. Open Archive: arXiv.org, 2014. [Online]. Available: http://arxiv.org/html/1409.0093

[20] E. Baccelli, F. Juraschek, O. Hahm, T. C. Schmidt, H. Will, and M. Wählisch, Eds., *Proceedings of the 3rd MANIAC Challenge, Berlin, Germany, July 27 - 28, 2013*, no. arXiv:1401.1163. Open Archive: arXiv.org, 2014. [Online]. Available: http://arxiv.org/html/1401.1163

[21] T. C. Schmidt, M. Waehlisch, R. Koodli, G. Fairhurst, and D. Liu, "Multicast Listener Extensions for Mobile IPv6 (MIPv6) and Proxy Mobile IPv6 (PMIPv6) Fast Handovers," RFC Editor, IETF, RFC 7411, November 2014. [Online]. Available: https://www.rfc-editor.org/rfc/rfc7411.txt

[22] T. C. Schmidt, S. Gao, H.-K. Zhang, and M. Waehlisch, "Mobile Multicast Sender Support in Proxy Mobile IPv6 (PMIPv6) Domains," RFC Editor, IETF, RFC 7287, June 2014. [Online]. Available: http://www.rfc-editor.org/rfc/rfc7287.txt

[23] R. Bush, R. Austein, K. Patel, H. Gredler, and M. Waehlisch, "Resource Public Key Infrastructure (RPKI) Router Implementation Report," RFC Editor, IETF, RFC 7128, February 2014. [Online]. Available: http://www.rfc-editor.org/rfc/rfc7128.txt

[24] M. Wählisch, R. Schmidt, T. C. Schmidt, O. Maennel, and S. Uhlig, "When BGP Security Meets Content Deployment: Measuring and Analysing RPKI-Protection of Websites," Open Archive: arXiv.org, Technical Report arXiv:1408.0391, August 2014. [Online]. Available: http://arxiv.org/abs/1408.0391

[25] H. Petersen, E. Baccelli, M. Wählisch, T. C. Schmidt, and J. Schiller, "The Role of the Internet of Things in Network Resilience," Open Archive: arXiv.org, Technical Report arXiv:1406.6614, June 2014. [Online]. Available: http://arxiv.org/abs/1406.6614

[26] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, "Backscatter from the Data Plane – Threats to Stability and Security in Information-Centric Network Infrastructure," *Computer Networks*, vol. 57, no. 16, pp. 3192–3206, Nov. 2013.

[27] T. C. Schmidt, M. Wählisch, D. Charousset, and S. Meiling, "On Name-based Group Communication: Challenges, Concepts, and Transparent Deployment," *Computer Communications*, vol. 36, no. 15–16, pp. 1657–1664, Sep.-Oct. 2013.

[28] D. Charousset, T. C. Schmidt, R. Hiesgen, and M. Wählisch, "Native Actors – A Scalable Software Platform for Distributed, Heterogeneous Environments," in *Proc. of the 4th ACM SIGPLAN Conference on Systems, Programming, and Applications (SPLASH '13), Workshop AGERE!* New York, NY, USA: ACM, Oct. 2013, pp. 87–96.

[29] M. Wählisch, F. Holler, T. C. Schmidt, and J. H. Schiller, "RTRlib: An Open-Source Library in C for RPKI-based Prefix Origin Validation," in *Proc. of USENIX Security Workshop CSET'13.* Berkeley, CA, USA: USENIX Assoc., 2013.

[30] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, "Lessons from the Past: Why Data-driven States Harm Future Information-Centric Networking," in *Proc. of IFIP Networking.* Piscataway, NJ, USA: IEEE Press, 2013.

[31] S. Meiling, T. Steinbach, T. C. Schmidt, and M. Wählisch, "A Scalable Communication Infrastructure for Smart Grid Applications using Multicast over Public Networks," in *Proc. of ACM Symposium on Applied Computing (SAC'13).* New York: ACM, March 2013, pp. 690–692.

[32] O. Hahm, E. Baccelli, M. Günes, M. Wählisch, and T. C. Schmidt, "RIOT OS: Towards an OS for the Internet of Things," in *Proc. of the 32nd IEEE INFOCOM. Poster Session.* Piscataway, NJ, USA: IEEE Press, 2013, pp. 2453–2454.

[33] M. Landsmann, H. Perrey, O. Ugus, M. Wählisch, and T. C. Schmidt, "Topology Authentication in RPL," in *Proc. of the 32nd IEEE INFOCOM. Poster Session.* Piscataway, NJ, USA: IEEE Press, 2013, pp. 2447–2448.

[34] G. Carle, J. Schiller, S. Uhlig, W. Willinger, and M. Wählisch, Eds., *The Critical Internet Infrastructure (Dagstuhl Seminar 13322)*, vol. 3, no. 8. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2013.

[35] A. Ghodsi, B. Ohlmann, J. Ott, I. Solis, and M. Wählisch, Eds., *Information-centric networking – Ready for the real world? (Dagstuhl Seminar 12361)*, vol. 2, no. 9. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2013.

[36] M. Waehlisch, T. C. Schmidt, and S. Venaas, "A Common API for Transparent Hybrid Multicast," RFC Editor, IRTF, RFC 7046, December 2013. [Online]. Available: http://www.rfc-editor.org/rfc/rfc7046.txt

[37] H. Perrey, M. Landsmann, O. Ugus, M. Wählisch, and T. C. Schmidt, "Topology Authentication in RPL," Open Archive: arXiv.org, Technical Report arXiv:1312.0984, December 2013. [Online]. Available: http://arxiv.org/abs/1312.0984

[38] M. Wählisch, "Conference Reports. Sec '13: 22nd USENIX Security Symposium. Large Scale Systems Security III," *;login:*, vol. 38, no. 6, pp. 29–31, Dec. 2013, electronic supplement.

[39] E. Baccelli, F. Juraschek, O. Hahm, T. C. Schmidt, H. Will, and M. Wählisch, "The MANIAC Challenge at IETF 87," *the IETF Journal*, vol. 9, no. 2, pp. 27–29, Nov. 2013.

[40] M. Wählisch, A. Vorbach, C. Keil, J. Schönfelder, T. C. Schmidt, and J. H. Schiller, "Design, Implementation, and Operation of a Mobile Honeypot," Open Archive: arXiv.org, Technical Report arXiv:1301.7257, Jan 2013. [Online]. Available: http://arxiv.org/abs/1301.7257

[41] M. Wählisch, E. Baccelli, J. Schiller, A. Voisard, T. C. Schmidt, S. Pfennigschmidt, M. Palkow, U. Weigmann, and U. Hanewald, "Technische Dimensionen der Flughafensicherheit," *Crisis Prevention*, no. 1, pp. 15–16, Jan. 2013.

[42] S. Meyer, M. Wählisch, and T. C. Schmidt, "Exploring Reachability via Settlement-Free Peering," in *Proc. of the ACM SIGCOMM CoNEXT. Student Workshop.* New York: ACM, Dec. 2012, pp. 49–50.

[43] D. Charousset, T. C. Schmidt, and M. Wählisch, "Actors and Publish/Subscribe: An Efficient Approach to Scalable Distribution in Data Centers," in *Proc. of the ACM SIGCOMM CoNEXT. Student Workshop.* New York: ACM, Dec. 2012, pp. 53–54.

[44] S. Meiling, T. C. Schmidt, and M. Wählisch, "Large-Scale Measurement and Analysis of One-Way Delay in Hybrid Multicast Networks," in *Proc. of the 37th Annual IEEE Conference on Local Computer Networks (LCN'12).* Piscataway, NJ, USA: IEEE Press, 2012.

[45] F. Jäger, T. C. Schmidt, and M. Wählisch, "Predictive Video Scaling – Adapting Source Coding to Early Network Congestion Indicators," in *2nd IEEE International Conference on Consumer Electronics - Berlin (ICCE-Berlin 2012).* Piscataway, NJ, USA: IEEE Press, Sep. 2012.

[46] S. Zagaria, T. C. Schmidt, S. Meiling, and M. Wählisch, "A Monitoring Framework for Hybrid Multicast Networks," in *2nd IEEE International Conference on Consumer Electronics - Berlin (ICCE-Berlin 2012).* Piscataway, NJ, USA: IEEE Press, Sep. 2012.

[47] M. Wählisch, S. Trapp, J. Schiller, B. Jochheim, T. Nolte, T. C. Schmidt, O. Ugus, D. Westhoff, M. Kutscher, M. Küster, C. Keil, and J. Schönfelder, "Vitamin C for your

Smartphone: The SKIMS Approach for Cooperative and Lightweight Security at Mobiles," in *Proc. of ACM SIGCOMM. Demo.* New York: ACM, August 2012, pp. 271–272.

[48] M. Wählisch, S. Trapp, C. Keil, J. Schönfelder, T. C. Schmidt, and J. Schiller, "First Insights from a Mobile Honeypot," in *Proc. of ACM SIGCOMM. Poster.* New York: ACM, August 2012, pp. 305–306.

[49] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, "Bulk of Interest: Performance Measurement of Content-Centric Routing," in *Proc. of ACM SIGCOMM. Poster.* New York: ACM, August 2012, pp. 99–100.

[50] M. Wählisch, O. Maennel, and T. C. Schmidt, "Towards Detecting BGP Route Hijacking using the RPKI," in *Proc. of ACM SIGCOMM. Poster.* New York: ACM, August 2012, pp. 103–104.

[51] Y. Yang, M. Wählisch, Y. Zhao, and M. Kyas, "RAID the WSN: Packet-based Reliable Cooperative Diversity," in *Proc. of the IEEE International Conference on Communications (ICC).* Piscataway, NJ, USA: IEEE Press, 2012, pp. 371–375.

[52] 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 the AS Level," in *Proc. of the 13th Passive and Active Measurement Conference (PAM)*, ser. Lecture Notes in Computer Science, N. Taft and F. Ricciato, Eds., vol. 7192. Heidelberg: Springer, 2012, pp. 200–210.

[53] E. Baccelli, L. Gerhold, C. Guettier, U. Meissen, J. Schiller, T. C. Schmidt, G. Sella, A. Voisard, M. Wählisch, and G. Wittenburg, "SAFEST: A Framework for Early Security Triggers in Public Spaces," in *Proc. of WISG 2012 – Workshop Interdisciplinaire sur la Securite Globale*, January 2012. [Online]. Available: http://hal.inria.fr/docs/00/66/66/98/PDF/WISG2012-SAFEST-JOINT-PAPER-3.pdf

[54] S. Meiling, D. Charousset, T. C. Schmidt, and M. Wählisch, "HAMcast – Evaluierung einer systemzentrierten Middleware-Komponente für einen universellen Multicast-Dienst im Future Internet," *Praxis der Informationsverarbeitung und Kommunikation (PIK)*, vol. 35, no. 2, pp. 83–89, Mai 2012.

[55] E. Baccelli, T. C. Schmidt, and M. W"ahlisch, Eds., *Proceedings of The 1st ACM International Workshop on Sensor-Enhanced Safety and Security in Public Spaces, SESP'12 (co-located with MobiHoc'12).* New York, NY, USA: ACM, 2012.

[56] E. Baccelli, O. Hahm, M. Wählisch, M. Günes, and T. C. Schmidt, "RIOT: One OS to Rule Them All in the IoT," INRIA, Research Report RR–8176, Dec. 2012. [Online]. Available: http://hal.inria.fr/hal-00768685

[57] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, "Backscatter from the Data Plane — Threats to Stability and Security in Information-Centric Networking," Open Archive: arXiv.org, Technical Report arXiv:1205.4778v1, May 2012. [Online]. Available: http://arxiv.org/abs/1205.4778v1

[58] S. Trapp, M. Wählisch, and J. Schiller, "Bridge the Gap: Measuring and Analyzing Technical Data for Social Trust between Smartphones," Open Archive: arXiv.org, Technical Report arXiv:1205.3068, May 2012. [Online]. Available: http://arxiv.org/abs/1205.3068v1

[59] T. C. Schmidt and M. Wählisch, "Why We Shouldn't Forget Multicast in Name-oriented Publish/Subscribe," Open Archive: arXiv.org, Technical Report arXiv:1201.0349v1, January 2012. [Online]. Available: http://arxiv.org/abs/1201.0349v1

[60] M. Wählisch, T. C. Schmidt, and G. Wittenburg, "On Predictable Large-Scale Data Delivery in Prefix-based Virtualized Content Networks," *Computer Networks*, vol. 55, no. 18, pp. 4086–4100, Dec. 2011.

[61] T. C. Schmidt, M. Wählisch, B. Jochheim, and M. Gröning, "WiSec 2011 Poster: Context-adaptive Entropy Analysis as a Lightweight Detector of Zero-day Shellcode Intrusion for Mobiles," *ACM SIGMOBILE Mobile Computing and Communications Review (MC2R)*, vol. 15, no. 3, pp. 47–48, July 2011.

[62] H. L. Cycon, T. C. Schmidt, M. Wählisch, D. Marpe, and M. Winken, "A Temporally Scalable Video Codec and its Applications to a Video Conferencing System with Dynamic Network Adaption for Mobiles," *IEEE Transactions on Consumer Electronics*, vol. 57, no. 3, pp. 1408–1415, August 2011.

[63] T. C. Schmidt, M. Wählisch, M. de Brühn, and T. Häberlen, "Ein Routing-Atlas für die strukturelle und visuelle Exposition des deutschen Internets," *Praxis der Informationsverarbeitung und Kommunikation (PIK)*, vol. 34, no. 2, pp. 60–72, June 2011.

[64] T. C. Schmidt, G. Hege, M. Wählisch, H. L. Cycon, M. Palkow, and D. Marpe, "Distributed SIP Conference Management with Autonomously Authenticated Sources and its Application to an H.264 Videoconferencing Software for Mobiles," *Multimedia Tools and Applications*, vol. 53, no. 2, pp. 349–370, June 2011.

[65] S. Trapp, M. Wählisch, and J. Schiller, "Short Paper: Can Your Phone Trust Your Friend Selection?" in *Proc. of the 1st ACM CCS Workshop on Security and Privacy in Mobile Devices (SPSM)*. New York: ACM, 2011, pp. 69–74.

[66] D. Charousset, S. Meiling, T. C. Schmidt, and M. Wählisch, "A Middleware for Transparent Group Communication of Globally Distributed Actors," in *Proc. of the Workshop on*

Posters and Demos Track. ACM/IFIP/UNSENIX Middleware. New York, USA: ACM, Dec. 2011.

[67] A. Knauf, G. Hege, T. C. Schmidt, L. Grimm, T. Kluge, P. Pogrzeba, and M. Wählisch, "Eine mobile VoIP Anwendung auf Basis eines RELOAD P2P Overlays," in *Wireless Communication and Information, Digital Divide and Mobile Applications*, J. Sieck, Ed. Boizenburg, Germany: Verlag Werner Hülsbusch, Oct. 2011, pp. 131–136.

[68] S. Meiling, D. Charousset, T. C. Schmidt, and M. Wählisch, "Implementierung und Performance-Evaluierung einer systemzentrierten Middleware-Komponente für einen universellen Multicast-Dienst," in *Report 298, 6. GI/ITG Workshop Leistungs-, Zuverlässigkeits- und Verlässlichkeitsbewertung von Kommunikationsnetzen und verteilten Systemen (MMBnet11)*. Hamburg, Germany: Universität Hamburg, Dept. Informatik, Sep 2011, pp. 80–88.

[69] S. Meiling, D. Charousset, T. C. Schmidt, and M. Wählisch, "HAMcast: Entwicklung und Evaluierung einer Architektur zur universellen Gruppenkommunikation im Internet," in *4. DFN-Forum Kommunikationstechnologien, Beiträge der Fachtagung*, ser. Lecture Notes in Informatics, P. Müller, B. Neumair, and G. D. Rodosek, Eds., vol. 187. German Informatics Society, June 2011, p. 149.

[70] H. Schwetlick, T. C. Schmidt, H. L. Cycon, and M. Wählisch, Eds., *Proceedings of the 1st IEEE International Conference on Consumer Electronics (ICCE-Berlin 2011)*. Piscataway, NJ, USA: IEEE Press, 2011.

[71] M. Wählisch, "One Day in the Life of RPKI," RIPE Labs, Community note, Dec. 2011. [Online]. Available: https://labs.ripe.net/Members/waehlisch/one-day-in-the-life-of-rpki

[72] M. Wählisch, "Beta Version of the RPKI RTR Client C Library Released," RIPE Labs, Community note, Sep. 2011. [Online]. Available: https://labs.ripe.net/Members/waehlisch/beta-version-of-the-rpki-rtr-client-c-library-released

[73] T. C. Schmidt, M. Waehlisch, and S. Krishnan, "Base Deployment for Multicast Listener Support in Proxy Mobile IPv6 (PMIPv6) Domains," RFC Editor, IETF, RFC 6224, April 2011. [Online]. Available: http://www.rfc-editor.org/rfc/rfc6224.txt

Appendix B

Invited Talks

In addition to talks at conferences, workshops, and IETF/IRTF meetings, the author of this thesis presented his work at the following events:

Dec. 2014 *How to Protect from Prefix Hijacking Using RPKI*, meeting of the expert group Infrastructure Security, eco/DE-CIX, Frankfurt/Main, Germany

Sep. 2013 *Prefix Origin Validation on Routers: Tools and Measurements*, National Provider Workshop organized by Federal Office for Information Security (BSI), Bonn, Germany

Mar. 2013 *RPKI in the Wild. Prefix Origin Validation on BGP Routers*, CAIDA, UC San Diego, USA

Mar. 2013 *The Internet – A Critical Infrastructure. About a Nation-Centric View on the Internet and the Importance of ASes*, CAIDA, UC San Diego, USA

Feb. 2013 *Updates from the Internet Backbone: An RPKI/RTR Router Implementation, Measurements, and Analysis*, short talk at the Network and Distributed System Security Symposium (NDSS), San Diego, USA

Nov. 2012 *Secure Inter-Domain Routing*, BCIX Technical Workshop “Securing the Internet’s Routing Infrastructure”, Berlin, Germany

Jan. 2012 *(Mobile) Internet-Kommunikation – Aktuelle Themen zur Sicherheit*, IT Security Seminar, Technische Hochschule Wildau, Germany

Oct. 2011 *Wie sicher ist eigentlich der Cyberspace*, 4. Nacht des Wissens, Hamburg, Germany

Oct. 2011 *HAMcast – A system-centric architecture to enable a universal multicast service in the Future Internet*, Institute of Telematics, KIT, Karlsruhe, Germany

Appendix C

Supervised Bachelor's and Master's Theses

The author of this thesis identified several topics for Bachelor's and Master's theses. He was the principle supervisor of the following theses:

1. Andreas Reuter: *Monitoring and Inspection of RPKI Repositories.* Bachelor's Thesis, Institute of Computer Science, Freie Universität Berlin, May 2015.

2. Jan-Christopher Pien: *Entwicklung und Evaluierung eines opportunistischen Verschlüsselungsverfahren auf Basis von Social Trust.* Bachelor's Thesis, Institute of Computer Science, Freie Universität Berlin, December 2014.

3. Fabrice Jean Ryba: *Implementing and Analysing sFlow measurements at an Internet Exchange Point.* Bachelor's Thesis, Institute of Computer Science, Freie Universität Berlin, July 2014.

4. Robert Schmidt: *Schutz wichtiger Webseiten durch RPKI. Messung und Analyse.* Bachelor's Thesis, Institute of Computer Science, Freie Universität Berlin, January 2014.

5. Samir Al-Sheikh: *Vergleichende Analyse von Abwehrmethoden gegen Interest Flooding Attacks in Named Data Networking.* Bachelor's Thesis, Institute of Computer Science, Freie Universität Berlin, January 2014.

6. Paul Wolpers: *Entwurf und Entwicklung eines Modells für die Analyse der Datenbanken der Regional Internet Registries.* Bachelor's Thesis, Institute of Computer Science, Freie Universität Berlin, January 2014.

7. Michael Mester: *Untersuchung und Optimierung der Leistungsfähigkeit der Prefix-Origin-Validation in einer realen BGP-Umgebung.* Master's Thesis, Institute of Computer Science, Freie Universität Berlin, December 2013.

8. Raphael Wutzke: *Analyse, Entwicklung und Implementierung eines Schutzes vor Portangriffen auf Smartphones unter Nutzung von Multipath TCP.* Bachelor's Thesis, Institute of Computer Science, Freie Universität Berlin, July 2013.

9. Marcel Kölbel: *Untersuchung der Qualität von Antworten im Amazon Mechanical Turk.* Bachelor's Thesis, Institute of Computer Science, Freie Universität Berlin, February 2013.

10. Marcin Nawrocki: *Entwurf und Implementierung eines Frameworks für die Analyse von Ad-hoc-Hotspot-Kommunikation.* Bachelor's Thesis, Institute of Computer Science, Freie Universität Berlin, October 2012.

11. Dennis Lampert: *Vergleichende Analyse von Private Set Intersection Protokollen.* Bachelor's Thesis, Institute of Computer Science, Freie Universität Berlin, October 2012.

12. Maximilian Schmidt: *Autonome Vertrauensimplementierung zwischen Home Gateways und Smartphones.* Bachelor's Thesis, Institute of Computer Science, Freie Universität Berlin, October 2012.

13. Robert Schlenz: *Entwurf und Implementierung einer Applikation zur Kategorisierung von Kontakten basierend auf Kommunikationsdaten.* Bachelor's Thesis, Institute of Computer Science, Freie Universität Berlin, August 2012.

14. Michael Zettelmann: *Ein Dienst zur Präfixgenerierung für P2P-Overlay-IDs basierend auf BGP-Daten an Internet Exchange Points.* Diploma Thesis, Institute of Computer Science, Freie Universität Berlin, July 2012.

15. Dominik Weidemann: *Design and implementation of a protocol for establishing ad-hoc trust between smartphones.* Bachelor's Thesis, Institute of Computer Science, Freie Universität Berlin, June 2012.

16. Christopher Flach: *Die strukturelle, zeitliche Analyse der relevanten deutschen IPv6 Internet-Infrastruktur.* Master's Thesis, Institute of Computer Science, Freie Universität Berlin, February 2012.

17. Fabian Holler: *Konzeption und Entwicklung einer Client-seitigen RPKI-RTR Library zur Validierung der Präfix-Zugehörigkeit von autonomen Systemen in BGP-Routen.* Bachelor's Thesis, Department Informatik, Hamburg University of Applied Sciences, November 2011.