

UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL
INSTITUTO DE INFORMÁTICA
PROGRAMA DE PÓS-GRADUAÇÃO EM COMPUTAÇÃO

MÁRCIO BARBOSA DE CARVALHO

**Judging Traffic Differentiation as Network
Neutrality Violation According to Internet
Regulation**

Thesis presented in partial fulfillment of the
requirements for the degree of Doctor of
Computer Science

Advisor: Prof. Dr. Lisandro Zambenedetti
Granville

Porto Alegre
December 2022

CIP — CATALOGING-IN-PUBLICATION

de Carvalho, Márcio Barbosa

Judging Traffic Differentiation as Network Neutrality Violation According to Internet Regulation / Márcio Barbosa de Carvalho. – Porto Alegre: PPGC da UFRGS, 2022.

132 f.: il.

Thesis (Ph.D.) – Universidade Federal do Rio Grande do Sul. Programa de Pós-Graduação em Computação, Porto Alegre, BR–RS, 2022. Advisor: Lisandro Zambenedetti Granville.

1. Network Neutrality. 2. Traffic Differentiation. 3. Internet Regulation. I. Granville, Lisandro Zambenedetti. II. Título.

UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL

Reitor: Prof. Carlos André Bulhões

Vice-Reitora: Prof^a. Patricia Pranke

Pró-Reitor de Pós-Graduação: Prof. Celso Giannetti Loureiro Chaves

Diretora do Instituto de Informática: Prof^a. Carla Maria Dal Sasso Freitas

Coordenador do PPGC: Prof. Claudio Rosito Jung

Bibliotecária-chefe do Instituto de Informática: Beatriz Regina Bastos Haro

AGRADECIMENTOS

Agradeço à minha família pelo apoio e incentivo em todos estes anos de estudo. Obrigado pai (*in memoriam*) e mãe por tudo! Nem consigo descrever tudo que fizeram por mim. Acho que o sacrifício para que eu fizesse os primeiros cursos de Informática valeram a pena! Obrigado vó Eva (*in memoriam*) pela dedicação ao cuidar tão bem de mim desde os 5 anos de idade, pelos ensinamentos e pelos valores que me passou. Obrigado mana Raquel pela parceria em todas as horas. Obrigado esposa Camila pelo companheirismo, pela tua dedicação ao cuidar de nós e pelas tantas vezes que tu te manteve ao meu lado quando precisei, inclusive, muitas vezes deixando a tua vida de lado. **Te amo!** Obrigado minha filha Aurora por iluminar e completar as nossas vidas. Esses 5 meses juntos já nos trouxeram muitas alegrias. **Te amo, filha!**

Agradeço ao Prof. Lisandro pela dedicação na minha orientação desde a graduação, passando pelo mestrado e agora durante o doutorado. Registro um agradecimento especial pela oportunidade de realizar o doutorado sob tua orientação. Sei que ter um orientando com dedicação parcial não é a situação ideal. Obrigado por todas as oportunidades oferecidas e especialmente pela oportunidade de fazer o estágio-sanduíche no IT-Aveiro em Portugal. Foi uma experiência que ultrapassou o aspecto acadêmico. Minha esposa e eu temos outra percepção de vida depois da estada por lá. Tenho muito orgulho de ser teu orientando! Ainda sobre o sanduíche, obrigado Raniery pelo apoio durante o projeto e pela parceria no AINA 2018 na Cracóvia. Obrigado também ao TRE-RS, principalmente ao Daniel Wobeto, por viabilizar a licença para o sanduíche.

Obrigado à malta do IT-Aveiro pela receptividade e ao Prof. Rui Aguiar pela oportunidade de realizar o estágio-sanduíche no IT-Aveiro. Obrigado Prof. Daniel Corujo pela orientação durante o projeto, principalmente pelos socorros quando estávamos em impasse sobre como conduzir o trabalho. Obrigado aos colegas de laboratório Eduardo Sousa e Anishur Rahman pelas conversas variadas nos intervalos. Obrigado Eduardo da Silva (IFC) pelas reflexões profundas sobre a vida e pelo apoio (incluindo o logístico). Obrigado Vitor Cunha pela grande parceria, pelo grande aprendizado técnico que tive durante os projetos, pelas conversas durante as caminhadas pela Universidade de Aveiro (incluindo as noturnas em tempos de deadline), pelo ensinamento sobre a cultura portuguesa e pela paciência em repetir e ensinar as palavras que eu não entendia, hehehehe. Sobre os supostos workshops entre políticos brasileiros e portugueses, espero que o nosso "workshop" tenha sido tão valioso para si como foi para mim.

ABSTRACT

Network Neutrality (NN) is a principle that establishes that traffic generated by Internet applications should be treated equally and it should not be affected by arbitrary interference, degradation, or interruption. Despite this common sense, NN has multiple definitions spread across the academic literature, which differ primarily on what constitutes the proper equality level to consider the network as neutral. NN definitions may also be included in regulations that control activities on the Internet. However, the regulations are set by regulators whose acts are valid within a geographical area, named jurisdiction. Thus, both the academia and regulations provide multiple and heterogeneous NN definitions. In this thesis, the regulations are used as guidelines to detect NN violations, which are, by this approach, the adoption of traffic management practices prohibited by regulators. Thereafter, the solutions can provide helpful information for users to support claims against illegal traffic management practices. However, state-of-the-art solutions adopt strict academic definitions (*e.g.*, all traffic must be treated equally) or adopt the regulatory definitions from one jurisdiction, which is not realistic or does not consider that multiple jurisdictions may be traversed in an end-to-end network path, respectively. An impact analysis showed that, under certain circumstances, from 39% to 48% of the detected Traffic Differentiations (TDs) are not NN violations when the regulations are considered, exposing that the regulatory aspect must not be ignored. In this thesis, a Regulation Assessment step is proposed to be performed after the TD detection. This step shall consider all NN definitions that may be found in an end-to-end network path and point out NN violation when they are violated. A service is proposed to perform this step for TD detection solutions, given the unfeasibility of every solution implementing the required functionalities. A Proof-of-Concept (PoC) prototype was developed based on the requirements identified along with the impact analysis, which was evaluated using information about TDs detected by a state-of-the-art solution. The verdicts were inconclusive (the TD is an NN violation or not) for a quarter of the scenarios due to lack of information about the traversed network paths and the occurrence zones (where in the network path, the TD is suspected of being deployed). However, the literature already has proposals of approaches to obtain such information. These results should encourage TD detection solution proponents to collect this data and submit them for the Regulation Assessment.

Keywords: Network Neutrality. Traffic Differentiation. Internet Regulation.

Julgando Diferenciação de Tráfego como Violação da Neutralidade da Rede de acordo com a Regulação

RESUMO

Neutralidade da rede (NR) é um princípio que estabelece que o tráfego de aplicações e serviços seja tratado igualmente e não deve ser afetado por interferência, degradação, ou interrupção arbitrária. Apesar deste senso comum, NR tem múltiplas definições na literatura acadêmica, que diferem principalmente no que constitui o nível de igualdade adequado para considerar a rede como neutra. As definições de NR também podem ser incluídas nas regulações que controlam as atividades na Internet. No entanto, tais regulações são definidas por reguladores cujos atos são válidos apenas dentro de uma área geográfica denominada jurisdição. Assim, tanto a academia quanto a regulação fornecem definições múltiplas e heterogêneas de NR. Nesta tese, a regulação é utilizada como guia para detecção de violação da NR, que nesta abordagem, é a adoção de práticas de gerenciamento de tráfego proibidas pelos reguladores. No entanto, as soluções adotam definições estritas da academia (por exemplo, todo o tráfego deve ser tratado igualmente) ou adotam as definições regulatórias de uma jurisdição, o que pode não ser realista ou pode não considerar que várias jurisdições podem ser atravessadas em um caminho de rede, respectivamente. Nesta tese, é proposta uma etapa de Avaliação da Regulação após a detecção da Diferenciação de Tráfego (DT), que deve considerar todas as definições de NR que podem ser encontradas em um caminho de rede e sinalizar violações da NR quando elas forem violadas. Uma análise de impacto mostrou que, em determinadas circunstâncias, de 39% a 48% das DTs detectadas não são violações quando a regulação é considerada. É proposto um serviço para realizar a etapa de Avaliação de Regulação, visto que seria inviável que todas as soluções tivessem que implementar tal etapa. Um protótipo foi desenvolvido e avaliado usando informações sobre DTs detectadas por uma solução do estado-da-arte. Os veredictos foram inconclusivos (a DT é uma violação ou não) para 75% dos cenários devido à falta de informações sobre os caminhos de rede percorridos e sobre onde a DT é suspeita de ser implantada. No entanto, existem propostas para realizar a coleta dessas informações e espera-se que os proponentes de soluções de detecção de DT passem a coletá-las e submetê-las para o serviço de Avaliação de Regulação.

Palavras-chave: Neutralidade da Rede, Diferenciação de Tráfego, Regulamentação da Internet.

LIST OF ABBREVIATIONS AND ACRONYMS

ANR	Average Normalized Ranks
AP	Application Provider
API	Application Programming Interface
AR	Argentina
AS	Autonomous System
ASN	Autonomous System Number
AU	Australia
BEREC	Body of European Regulators for Electronic Communications
BGP	Border Gateway Protocol
BR	Brazil
CA	Canada
CDF	Cumulative Distribution Function
CDN	Content Delivery Network
CIM	Common Information Model
CO	Colombia
CP	Content Provider
CPU	Central Process Unit
CRUD	Create Read Update Delete
DE	Germany
DE-HE	Hessen, Germany (DE)
DMTF	Distributed Management Task Force
DNS	Domain Name System
DPI	Deep Packet Inspection
DSL	Digital Subscriber Line
DT	Drop Tail
EDF	Early Deadline First
ES	Spain
EU	European Union
FCC	Federal Communications Commission
FCFS	First Come First Served
FIN	Finlayse
FTC	Federal Trade Commission

FQDN	Fully Qualified Domain Name
GB	United Kingdom
GH	Ghana
HTTP	Hyper Text Transfer Protocol
HU-JN	Jász-Nagykun-Szolnok, Hungary
IATA	International Air Transport Association
ICANN	Internet Corporation for Assigned Names and Numbers
ICMP	Internet Control Message Protocol
IE	Ireland
IETF	Internet Engineering Task Force
IL	Israel
IN	India
IoT	Internet of Things
IP	Internet Protocol
IPTV	Internet Protocol (IP) Television
IPv4	IP version 4
IPv6	IP version 6
ISO	International Organization for Standardization
ISP	Internet Service Provider
IT	Information Technology
JP	Japan
KL	Kullback-Leibler
K-S	Kolmogorov-Smirnov
KDE	Kernel Density Estimation
L-PLR	Lower Limit Packet Loss Rate
MVC	Model View Controller
NANO	Network Access Neutrality Observatory
NCSL	National Conference of State Legislatures
NETCONF	Network Configuration Protocol
NFV	Network Function Virtualization
NML	Network Mark-Up Language
NN	Network Neutrality
NRA	National Regulatory Authority
NS	Network Slice

NTP	Network Time Protocol
OGF	Open Grid Forum
OONI	Open Observatory of Network Interference
OS	Operating System
OSI	Open System Interconnection
OTT	Over-the-Top
P2P	Peer-to-Peer
PCAP	Packet Capture
PECF	Power and Exponential Composite Function
PoC	Proof-of-Concept
PL	Poland
PLR	Packet Loss Rate
QoE	Quality of Experience
QoS	Quality of Service
RIB	Routing Information Base
REST	REpresentational State Transfer
RIB	Routing Information Base
SSL	Secure Sockets Layer
SLA	Service Level Agreement
SMIv2	Structure of Management Information Version 2
SMTP	Simple Mail Transfer Protocol
SNI	Server Name Indication
SP	Strict Priority
SP	Super Peer
SQL	Structured Query Language
RED	Random Early Detection
RIPE	Réseaux IP Européens
RO-TR	Teleorman, Romania
RPC	Remote Procedure Call
RQ	Research Question
RS	Serbia
RSP	Rogue Super Peer
RST	Reset
RTSP	Real Time Streaming Protocol

RTT	Round Trip Time
SE	Sweden
SG	Singapore
SIP	Session Initiation Protocol
TCP	Transmission Control Protocol
TD	Traffic Differentiation
TH	Thailand
TLS	Transport Layer Security
TTL	Time to Live
UDP	User Datagram Protocol
UML	Unified Modeling Language
URL	Uniform Resource Locator
USA	United States of America
US-CA	California, USA
US-CO	Columbia, USA
US-VT	Vermont, USA
VoD	Video on demand
VoIP	Voice over IP
VPN	Virtual Private Network
VXDL	Virtual private eXecution infrastructure Description Language
YANG	Yet Another Next Generation
W3C	World Wide Web Consortium
WFQ	Weighted Fair Queue
WRED	Weighted Random Early Detection
ZA	South Africa

LIST OF FIGURES

Figure 1.1	End-to-end network path traversing multiple jurisdictions	14
Figure 1.2	NN violation detection flow.....	15
Figure 3.1	Percentage of verdicts of Glasnost tests	47
Figure 3.2	Distribution of countries in Glasnost tests.....	48
Figure 3.3	Distribution of countries whose Glasnost tests detected TDs	49
Figure 3.4	Distribution of countries with relative % of TD over tests	50
Figure 3.5	Top 10 jurisdictions with detected TDs	51
Figure 3.6	Top 10 jurisdictions with relative % of TD over tests	52
Figure 3.7	Agreements and disagreements of regulatory instructions.....	55
Figure 3.8	Influence of regulatory instructions in verdicts	56
Figure 3.9	Influence of NN regulation rollback in the United States of America (USA)57	
Figure 4.1	Relations between network topology, communication graph, and occurrence zones	62
Figure 4.2	Information Model - Topology classes	64
Figure 4.3	Information Model - Differentiation classes.....	65
Figure 4.4	Information Model - Regulation classes.....	66
Figure 4.5	Excerpt from the proposed nn-regulation Yet Another Next Generation (YANG) model	70
Figure 4.6	Excerpt from the proposed tdo YANG model	71
Figure 4.7	Excerpt from the proposed tdo YANG model - Remote Procedure Call (RPC) statement definition	72
Figure 5.1	JurisNN deployment scenario.....	74
Figure 5.2	JurisNN architecture	76
Figure 6.1	JurisNN prototype architecture.....	87
Figure 6.2	Place of the TD deployment: results distribution	95
Figure 6.3	Place of the TD deployment: verdicts distribution.....	96
Figure 6.4	German test: results distribution.....	98
Figure C.1	Response time	127
Figure C.2	CPU usage.....	128
Figure C.3	Memory usage.....	129
Figure C.4	jetconf/data.py code snippet.....	130

LIST OF TABLES

Table 2.1	The state-of-the-art	42
Table 2.2	The criteria adopted by each solution	43
Table 3.1	Verdicts numbers	47
Table 3.2	NN definitions that were considered	53
Table 4.1	Suitable data models analysis.....	67
Table 4.2	Existing YANG data models	68
Table 4.3	Proposed new YANG data models and augmentations to the existing ones ...	69
Table 5.1	Auxiliary functions used by Judgment Algorithms.....	83
Table 6.1	RESTCONF Application Programming Interface (API)	88
Table 6.2	Applications affected by the TDs detected by Wehe between January, 1st 2021 and February, 28th 2021	89
Table 6.3	TD information, collected network paths, and occurrences	90
Table 6.4	AS numbers and the respective AS names	90
Table 6.5	Jurisdictions found in source, destination, and network path.....	91
Table 6.6	Regulations about throttling in the found jurisdictions valid from 2021- 01-01 onwards.....	93
Table 6.7	Targeting test: results distribution	97

CONTENTS

1 INTRODUCTION	13
1.1 Problem statement	15
1.2 Main hypothesis and research questions	16
1.3 Main contributions.....	17
1.4 Thesis roadmap	17
2 BACKGROUND AND STATE-OF-THE-ART	19
2.1 Internet regulation and the jurisdictions in the cyberspace	20
2.2 Traffic Differentiation.....	23
2.3 Network Neutrality	26
2.4 State-of-the-art	30
2.5 Summary.....	44
3 IMPACT ANALYSIS	45
3.1 Glasnost dataset	46
3.2 NN definitions from regulations.....	50
3.3 Results and discussion	54
3.3.1 Introducing regulatory instructions into Glasnost’s results	54
3.3.2 Discussion	58
3.4 Summary.....	59
4 MODELING	61
4.1 Concepts, assumptions, and solution objectives.....	61
4.2 Information model	63
4.3 Investigation of suitable data models	66
4.4 Data models based on YANG	68
4.5 Summary.....	73
5 PROPOSAL	74
5.1 Deployment scenario.....	74
5.2 JurisNN architecture	75
5.3 Judgment algorithms	77
5.3.1 Jurisdiction establishment criteria and auxiliary functions.....	77
5.3.2 Examples of Judgment Algorithms.....	82
5.4 Summary.....	84
6 EVALUATION	86
6.1 JurisNN prototype.....	86
6.2 Evaluation data	88
6.3 JurisNN judgment results	94
6.3.1 Analysis caveats	99
6.4 Summary.....	101
7 CONCLUSION	102
REFERENCES	106
APPENDIX A — RESUMO EXPANDIDO	115
APPENDIX B — PROPOSED YANG MODELS	118
APPENDIX C — PERFORMANCE CHARACTERIZATION	126
APPENDIX D — SCIENTIFIC PRODUCTION	131
D.1 Published Papers	131
D.2 Collaborations	131

1 INTRODUCTION

Network Neutrality (NN) is a principle that establishes that traffic generated by Internet applications should be treated equally and it should not be affected by arbitrary interference, degradation, or interruption. Despite this common sense about NN, it has multiple definitions spread across the academic literature (GARRETT et al., 2018b). These definitions differ primarily on what constitutes the proper equality level to consider a network as neutral. For instance, they range from those that state that all traffic packets should be treated equally (GARRETT et al., 2018b) to those that allow justifiable Traffic Differentiation (TD) practices (JORDAN, 2009). NN definitions may also be included in the regulations that control the activities on the Internet. These regulations are set by legislators and regulatory agencies whose acts are valid just within a geographical area named *jurisdiction*, which usually encompasses a state, a country, or a region. Therefore, each jurisdiction may have its NN definitions. Thus, both academic literature and regulations provide multiple and heterogeneous NN definitions.


Several solutions have been proposed and deployed to detect violation of the NN principle. In general, they measure the network traffic searching for abnormal behaviors that may indicate that the Internet Service Provider (ISP) is deploying TD practices (*e.g.*, throttling, blocking, or prioritization). For most of the solutions, the detected TDs are implicitly pointed out as NN violations (KANUPARTHY; DOVROLIS, 2010) (ZHANG; MAO; ZHANG, 2009) (DISCHINGER et al., 2010). Therefore, these solutions apply a strict definition of NN that prohibits any TD. Besides, a few solutions of the state-of-the-art use definitions from the regulation to point out the detected TD as NN violation (Body of European Regulators for Electronic Communications, 2016) (BUSTOS-JIMÉNEZ; FUENZALIDA, 2014). However, these solutions follow only the regulation stated in the jurisdictions of their proponents.

This thesis advocates that the regulations should be used as guidelines to build solutions designed to detect NN violation. In this case, an NN violation would be the adoption of traffic management practices that were prohibited by the legislators and regulatory agencies (instead of a violation of academic definitions). Being based on regulations, the solutions may provide legally actionable evidence to support customer complaints against ISPs, for deliberation in the competent judicial authority (BUSTOS-JIMÉNEZ; FUENZALIDA, 2014). Also, the regulatory instructions tend to be more detailed because they are used to support the activities of the regulatory agencies. Indeed, they detail which

TD techniques are prohibited (or allowed) to be applied over traffic from which applications, protocols, or services, and in which situations. For instance, is the throttling of Peer-to-Peer (P2P) protocols allowed without restriction, confined just to peak hours, or is it forbidden? Is zero-rating a mobile application allowed when the competitors are still charged? Can Internet contents be blocked, and if so, is a court order required, or can anything be blocked arbitrarily?

The fact that regulations are valid just within the jurisdiction of the legislators, or the respective regulatory agencies, leads to the scenario depicted in Fig. 1.1. Each jurisdiction (Jurisdiction 1, Jurisdiction 2, ..., Jurisdiction n) may have its NN definitions (NN 1, NN 2, ..., NN n , respectively). Thus, the end-to-end network path between the end-user and the application server may traverse different jurisdictions with different NN definitions. Also, even when the traffic is confined within one jurisdiction, it should be evaluated according to the regulation stated in that jurisdiction. However, as mentioned earlier in this thesis, most of the state-of-the-art solutions point out the detected TDs as NN violations, regardless of the regulation stated in the jurisdiction where they happened. In turn, those solutions that use definitions from regulations are concerned only to the jurisdiction where they are deployed, thus, ignoring that multiple NN definitions may be found along with the end-to-end network path.

Figure 1.1 – End-to-end network path traversing multiple jurisdictions



Source: (CARVALHO et al., 2020)

NN violations can be deployed anywhere in the path between two communication points on the Internet. It can be deployed, for example, by the ISPs that connect end-user devices to the network, or even at the core of the Internet backbone (ZHANG; MAO; ZHANG, 2009). Therefore, the detection of NN violation must explicitly consider where, in an end-to-end network path, the violation was deployed. However, since NN definitions vary from jurisdiction to jurisdiction when regulations are considered, some behaviors in one part of the Internet may be considered an NN violation, while in other parts, the same behaviors may not represent an NN violation at all. Thus, to be effective, the NN violation detection shall consider the diverse definitions of NN found in the jurisdictions traversed in the end-to-end network path.

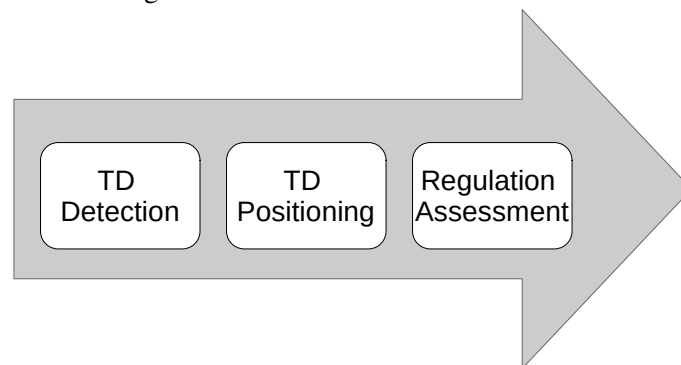
1.1 Problem statement

Most of the solutions from the state-of-the-art point out NN violation adopting strict definitions from academic literature instead of following the regulations stated by legislators and regulatory agencies. Those that apply the regulations ignore that multiple definitions for NN may be found along with an end-to-end network path. By considering the above facts, the problem addressed in this thesis can be stated as follows.

Problem definition: *The problem addressed in this thesis is the lack of solutions for detection of NN violation that follow the definitions stated on regulations and also consider the multiple definitions that may be found along with an end-to-end network path due to the jurisdictions of legislators and regulatory agencies.*

This thesis proposes the addition of one extra step (Regulation Assessment) to the existing steps (TD detection and positioning) already performed by state-of-the-art solutions for detection of NN violation, as depicted in Figure 1.2. The Regulation Assessment step encompasses the management of information related to the NN definitions stated on regulations around the world, the processing of TD information to compare it to these definitions, and the pointing out of the NN violation when the TD violates these definitions. However, having every solution committed to NN violation detection implementing such requirements would be impractical. Thus, such requirements could be provided by an additional solution committed only to the Regulation Assessment step. This thesis addresses the problem by proposing a solution designed as a service that implements the Regulation Assessment step to be consulted by existing and future solutions or by users interested on evaluating detected TDs against the established regulation.

Figure 1.2 – NN violation detection flow



Source: adapted from (CARVALHO; SCHAURICH; GRANVILLE, 2018)

1.2 Main hypothesis and research questions

Given the lack of solutions that point out NN violation considering the multiple regulations that may be found along with an end-to-end network path and due to the difficulty to implement the requirements to perform such assessment in every solution committed to NN violation detection, this thesis presents the following hypothesis.

Hypothesis: a service performing the Regulation Assessment step can provide the requirements to point out NN violation considering multiple regulations found along with an end-to-end network path.

In order to guide the investigations conducted in this thesis, the following Research Questions (RQs) associated with this hypothesis are defined and presented.

RQ 1: *How much the consideration of regulations impacts the results that have been achieved by state-of-the-art solutions for NN violation detection?*

RQ 2: *How to enable the pointing out of NN violation according to multiple definitions stated on the regulations set by legislators and regulatory agencies?*

RQ 3: *Is the TD information collected by TD detection and positioning solutions enough to point out NN violation considering multiple regulations found along with an end-to-end network path?*

The RQ 1 is related to the quantification of the impact of the proposed Regulation Assessment step on the results achieved by state-of-the-art solutions. The need for such a step is verified and confirmed by quantifying such impact. The analysis conducted considers the TDs detected by Glasnost (DISCHINGER et al., 2010) evaluating whether these TDs still are pointed out as NN violations when regulations are considered.

The RQ 2 is related to the identification of the requirements to perform the Regulation Assessment step. After that, it is also related to the modeling, proposal, and evaluation of a solution to perform the Regulation Assessment step. The modeling and solution proposals are conducted, whose architecture is implemented on a Proof-of-Concept (PoC) prototype that is evaluated using artificial crafted TD information.

The RQ 3 is related to the integration of solutions designed to TD detection and positioning with the Regulation Assessment solution. Such integration is evaluated from the perspective of the required information to perform the assessment by submitting the information about TDs collected by Wehe (LI et al., 2019) analyzing the conclusiveness of the achieved results (the TD is considered an NN violation or not).

1.3 Main contributions

While solving the research problem and answering the related research questions, many advances on the state-of-the-art have been provided. This thesis presents the following contributions:

1. An analysis to quantify the influence of considering the regulations when pointing out NN violation based on evidence detected by state-of-the-art solution. This analysis quantifies how wrong are the results provided by these solutions by ignoring the regulatory aspect of the Internet.
2. The identification of the requirements to perform the Regulation Assessment step accompanied by proposals of information and data models using the findings of the conducted investigation. These models may help to unify the effort from multiple solutions devoted to the detection of NN violation considering the multiple NN definitions found along with an end-to-end network path.
3. An architecture for a possible solution addressing the identified requirements and following the proposed models accompanied by a PoC prototype that may serve as an initial solution towards a definitive solution to the research problem.
4. An analysis based on evidence provided by the state-of-the-art solution to evaluate the conclusiveness of the results achieved by the proposed solution, discussing whether the information collected by TD detection and positioning solutions is enough to properly point out NN violation considering the regulations.

1.4 Thesis roadmap

The remainder of this thesis is organized as follows.

In Chapter 2, this thesis describes relevant aspects related to NN and to the new methodology proposed to point out NN violation. In this regard, concepts related to Internet regulation, Traffic Differentiation (TD), and NN definitions are presented. The state-of-the-art of solutions committed to NN violation and TD detection is also provided.

In Chapter 3, the investigation addresses the RQ 1. The research starts evaluating the results publicized by Glasnost (DISCHINGER et al., 2010) analyzing the detected TDs to verify whether they are, in fact, NN violations considering the corresponding regulations. The evaluation uses the public dataset of 2016 (Measurement Lab, 2016),

which is the most recent complete year available in the dataset. Using this data, three hypothetical scenarios are designed to help to answer the RQ 1.

In Chapter 4, the findings of addressing the RQ 1 help to identify the requirements to perform the Regulation Assessment step. These requirements guided the design of information models and the investigation of existing data models to represent them. The choice for models based on the Yet Another Next Generation (YANG) language is explained along with the presentation of both the existing YANG models used and the proposal of new ones to complement the information from the existing ones.

In Chapter 5, the deployment scenario of the proposed solution is presented, detailing the actors and their contributions. A conceptual architecture is presented for the proposed solution, detailing its modules and the respective responsibilities. One of the modules is responsible for the judgment itself, helped by *Judgment Algorithms*. These algorithms are responsible for using the information from regulations to judge TD as NN violation. Three Judgment Algorithms are detailed and presented.

In Chapter 6, a PoC prototype based on the conceptual architecture is presented. The data used to evaluate the judgment performed by the prototype is presented, which includes the information about TDs detected by Wehe in 2021 (LI et al., 2019) (MEASUREMENT LAB, 2021), the network paths collected by Réseaux IP Européens (RIPE) Atlas platform (to complement the Wehe dataset), and the regulatory instructions configured into the prototype. The conclusiveness of the results of the three algorithms is discussed, and the analysis caveats are presented.

In Chapter 7, the final remarks and the future work in the context of this thesis are presented. The research questions are revisited to discuss the provided answers and point out the remaining gaps subject to future investigations.

2 BACKGROUND AND STATE-OF-THE-ART

In this chapter, aspects of the NN principle related to this thesis are presented. This thesis proposes the use of NN definitions stated on regulations to point out violation of the NN principle. In this sense, Internet regulation aspects are briefly presented (Section 2.1), such as how the Internet is regulated nowadays and the jurisdiction concept. This concept drives the choice about which regulation must be applied in each case. The jurisdiction is established by tests that courts adopt to decide whether they have competence to judge a case. In turn, the regulations define which traffic management practices are prohibited to be adopted by Internet operators. These practices are commonly referred as Traffic Differentiation (TD), whose details are briefly presented (Section 2.2), such as their triggers, classification, differentiation mechanism, and perceived discrimination.

There is not a commonly accepted definition for the NN principle (GARRETT et al., 2018b), which lead to the statement of several definitions biased to the NN aspect being evaluated by its proponent (*e.g.*, inter-network criteria, economical aspects) (Section 2.3). Although this lack of a commonly accepted definition, the essence of the NN principle is the freedom of choice by users on the Internet use for legal purposes. In this sense, in a neutral network, legal uses of the network should not face obstacles of any nature, such as performance interference, disruption, or economical differentiation. Indeed, legal uses of the network has been defined within Internet regulation frameworks to enforce NN (GARRETT et al., 2022). In turn, Internet access is a service offered by operators that need to plan the network capacity by estimating the service usage. In this sense, the operators may offer Internet access with distinct capacities (transmission rate and data cap) without violating the essence of the NN principle, since users have freedom of choice on the use of the contracted Internet capacity.

The state-of-the-art solutions for detection of TD and violation of the NN principle are presented (Section 2.4). The focus is on the detection objective (TD or NN violation), the type of TD detected (*e.g.*, throttling, blocking), the technique adopted in the analysis (single test or aggregated data), and the criteria adopted to detect TD (traffic characteristic, statistical inference, or regulation). Finally, a summary is presented highlighting important concepts presented along with this chapter (Section 2.5).

2.1 Internet regulation and the jurisdictions in the cyberspace

In this section, concepts related to Internet regulation are presented, including aspects about the establishment of jurisdiction in cyberspace.

Until the rise of the Internet to the main public in the beginning of the 1990s, the Internet was mainly an academic network (KRÄMER; WIEWIORRA; WEINHARDT, 2013). In this phase, the Internet was maintained and operated by the collaborative work from its stakeholders (*e.g.*, researchers, university staff). For instance, the design and operation of numbering and naming structures conducted by Jonathan B. Postel (COHEN-ALMAGOR, 2013). As long as the network started to grow and gain relevance in other fields, such as entertainment and business, the collaborative work was not enough given the Internet shift of scale. In this sense, a few organizations started to establish rules for technical aspects of the Internet: the Internet Corporation for Assigned Names and Numbers (ICANN) for Internet Protocol (IP) address space and Domain Name System (DNS) management, the Internet Engineering Task Force (IETF) for Internet standards and protocols, and the World Wide Web Consortium (W3C) for Web standards. Such rules must be followed by the Internet operators in order to keep the network running.

In the 1990s, the Internet continued to grow, gaining importance in many human activities, augmenting the chance of conflicts over resources and services offered through the network. Such conflicts paved the way for the discussion of whether and how the Internet should be regulated (EKO, 2001) (MAYER-SCHÖNBERGER, 2002). One approach discussed is self-regulation, in which the Internet actors elaborate the rules that must be followed, which could increase the legitimacy of the rules. However, such an approach is impacted by the difficulty of imposing punishments due to the private nature of the Internet actors (providers, companies, and organizations) who lack the power to enforce the rules. Another approach discussed is the regulation by national laws established by regulators (*e.g.*, legislators, regulatory agencies), which could use the state power to enforce the rules. However, such an approach is impacted by the Internet characteristics (*e.g.*, decentralized structure, lack of territorial limits) that hinder the regulator legitimacy in the whole network. Another approach discussed is the regulation by international treaties, in which countries could agree about the establishment of rules. This approach could increase regulation legitimacy and enforceability. However, this approach is impacted by the required consensus to define the rules, given that many countries should agree on the treaty to be effective. However, countries may not agree about some issues (*e.g.*, freedom

of speech, censorship) hindering the treaty establishment.

Nowadays, Internet regulation is a mix of self-regulation, national-scope regulation, and international treaties. Self-regulation is applied in technical aspects of the Internet by organizations like ICANN, IETF, and W3C. In this sense, self-regulation helps in the legitimacy of the technical recommendations stated by these bodies because they have significant representation from the operational actors. National laws and policies are being established by many countries, regulating the economic and social aspects of the Internet, such as network neutrality, auditing, and civil rights (GARRETT et al., 2022). In this sense, national laws help achieve the required enforceability of the rules since the state has the power to impose punishments. International treaties are rarer due to the difficulty of their establishment. Even so, the European Union (EU) has established a single Internet regulation for the region. However, regulation enforcement is delegated to the state members. This mix of regulations increases the legitimacy and enforceability of Internet regulation (MAYER-SCHÖNBERGER, 2002).

The approaches mentioned above are examples of regulation based on rules defined *a priori* (KELLER, 2019), which is named an *ex-ante* approach. However, in some countries, the supposed harmful conducts performed by ISPs are evaluated *after* their occurrence, which is named an *ex-post* approach. In this approach, these conducts are claimed in the regular judicial system or a specialized bureau. The claims are judged according to jurisprudence, *i.e.*, the previous understanding of certain conduct and the specific harm caused. For instance, this regulation approach has been adopted in the United States of America (USA) since mid-2018 for the NN regulation (Federal Communications Commission, 2018). This lack of federal regulation in the USA opened the opportunity for USA states to establish their NN regulations.

The development of these multiple models for Internet regulation is intrinsically related to the discussion of who has the power to establish the Internet regulation and where such regulation can be valid, which is named the regulation *jurisdiction* (MAY; CHEN; WEN, 2004). The jurisdiction concept is related to how the power given by the national sovereignty is partitioned between the state actors (*e.g.*, lawmakers, organizations, bodies, agencies, states of a federation, municipalities). The partition defines which actors have the competence over specific matters (*e.g.*, transit, health, Internet) at which level (*e.g.*, national, regional, local). This partition may be *horizontal* in which the power is exclusively exercised by one actor (*e.g.*, a national body or agency) or may be *vertical* in which a hierarchy of actors is defined (*e.g.*, national-, regional-, and local-scope actors).

There are three types of jurisdiction: *jurisdiction to prescribe* (when the actor has the competence to establish laws on the specific matter), *jurisdiction to adjudicate* (when the court can decide cases of the specific matter), and *jurisdiction to enforce* - also referred as the scope of jurisdiction (SVANTESSON, 2020) (when the decisions of a court have an effect on that jurisdiction) (GLADSTONE, 2003).

The jurisdiction concept then is closely dependent on information related to geographical localization because this information is required to assess who has competence over a matter in an area. However, due to this localization dependence, the jurisdictions are hard to establish on the Internet because of its characteristics (*e.g.*, decentralized structure, lack of territorial limits). Beyond, there are many components and actors to track the localization (*e.g.*, servers, routers, endpoints, users, ISPs, Content Providers (CPs)) that make this task even harder. For instance, Chertoff *et al.* (2015) propose four ways to establish the jurisdiction of a case about data based on: data creator citizenship, data subject citizenship, data holder citizenship, and location of the harm, exposing the complexity of the matter. In turn, Jiménez and Lodder (2015) use a metaphor based on the high seas to simplify the discussion about jurisdiction on the Internet (JIMÉNEZ; LODDER, 2015). The authors represent the endpoints of communication as *harbors* and the whole communication infrastructure between them as the *high seas*. Harbors A and B represent the local jurisdictions of the origin and destination of the communication. The high seas represent the usage of international laws. Using this metaphor, they discuss cases in the USA, Germany (DE), and the Netherlands, pointing which jurisdiction (Harbor A, Harbor B, or high seas) each actor (*e.g.*, plaintiffs, defendants, courts) adopted.

In order to establish the jurisdictions, the courts usually adopt tests that assess characteristics of the case to decide whether they have competence over it, which is referred to as admissibility control. The straightforward way to establish jurisdiction is by assessing where the harmful action took place. Therefore, if the harmful action happened within the court's jurisdiction, the court can judge the case. Courts around the world commonly use this test. In turn, USA courts adopt, among others, the *Targeting test* (WANG, 2008) that consists in assessing the place where harmful action had *effects*. For instance, a company established in jurisdiction A committing harmful actions against its clients in jurisdiction B can be claimed in jurisdiction B (where the harmful actions had effects). In turn, German courts adopt both tests: where the harmful action happened and where the injury has incurred (JIMÉNEZ; LODDER, 2015).

The jurisdiction impacts the decision of which regulation should be applied in a

case on the Internet. Therefore, such an aspect must be considered by a solution devoted to assessing NN violation based on the regulations. In this sense, the jurisdiction tests presented above must be considered when establishing the jurisdiction of a TD in order to apply the right regulation. In the next section, concepts related to Traffic Differentiation that may be deployed by ISPs are presented.

2.2 Traffic Differentiation

In this section, the TD definitions that are related to aspects and definitions of NN are briefly presented. In this sense, Traffic Differentiation (TD) is characterized by its *Triggers*, *Traffic Classification*, *Differentiation Mechanism*, and *Perceived Discrimination* (GARRETT et al., 2018b). Each of these aspects is detailed in this section.

Triggers are properties of the flow (*e.g.*, related application, source, or destination path) or network conditions (*e.g.*, congestion) that leads an ISP to deploy the TD. A decade ago, the flows from P2P applications (*e.g.*, BitTorrent) were the main targets of TD deployed by ISPs motivated by multiple factors (*e.g.*, high network utilization, an avenue to illegal content sharing). In the past years, the high network utilization from Over-the-Top (OTT) services (*e.g.*, video streaming) shifted the deployment of TD against their applications (LI et al., 2019). Other triggers related to the business of the providers and the operation of their networks are also pointed as motivations to deploy TD. Some providers included high-level services (*e.g.*, Voice over IP (VoIP), Video on demand (VoD) platforms) to complement their Internet access services, which is referred to as vertical integration (SCHULZRINNE, 2018). Then, providers start to deploy TD against their competitors in the vertical services to favor the quality of their services. The provider may also introduce TD to avoid or reduce the utilization of operationally expensive links (*e.g.*, bad peer or transit agreements) (DISCHINGER et al., 2008).

Traffic Classification uses properties of the flow (*e.g.*, packet header or payload, behavior) to determine the flow priority. The headers of packets provide helpful information to classify flows, such as the source and destination IP addresses and the transport ports, which are used to determine its application. However, the transport ports used by an application are easy to modify, turning easy to circumvent the classification. In turn, traffic classification based on the content of packet payload, known as Deep Packet Inspection (DPI), is more effective in recognizing the signature of many applications. However, they require much more computational resources, which usually are not available in

regular network devices. In order to overcome this problem, DPIs are usually deployed through hardware specialized in traffic processing (appliances) and, recently, through virtual appliances based on Network Function Virtualization (NFV) (MIJUMBI et al., 2016). However, nowadays, most of the Internet traffic is encrypted (Google, 2022) hiding the packet payload from DPIs and, consequently, hindering its effectiveness. As a countermeasure, DPIs track the Transport Layer Security (TLS) handshake to collect the Server Name Indication (SNI) field of the certificates used to start the encrypted session, which stores the Fully Qualified Domain Name (FQDN) of the application server. This information may allow identifying the application if such FQDN is used only for one application. However, protocols like TLS 1.3 and QUIC also encrypt SNI fields, hindering, even more, the ability of DPI to identify applications. In the future, due to the expected increase in encrypted traffic, content-based classification still may be used for flow identification but using only DNS lookup information (LI et al., 2019). However, traffic classification may require the adoption of behavioral-based classification techniques (GU; ZHANG; XUE, 2011), which already are available but having much less effectiveness than content-based classification. Machine learning classification may be used to improve the accuracy of behavioral-based classification (JONATHAN; MISRA; OSAMOR, 2021).

The Differentiation Mechanism is how the ISP interferes with the flow (*e.g.*, block, delay, drop, modify) to implement the TD. Blocking mechanisms interrupt the differentiated communication simply by not forwarding its packets or interfering with it inducing its termination (*e.g.*, sending crafted packets setting the FIN or RST flags of the Transmission Control Protocol (TCP)). The delaying mechanism interferes with the communication affecting its performance, which may be influenced by the deployment of traffic queues associated with schedulers (*e.g.*, Strict Priority (SP), Weighted Fair Queue (WFQ), Weighted Random Early Detection (WRED)). These mechanisms may differentiate the traffic delaying or prioritizing its communication. Dropping mechanisms interrupt the communication discarding its packets, which is also known as blackholing (NAWROCKI et al., 2019). Modifying mechanisms may both act over packet headers or payload to affect its performance (*e.g.*, interfering with transmission and congestion control of TCP) or its content (*e.g.*, proxies that reduce the quality or modify images (ZIPROXY, 2022)). Another kind of TD is related to the economic aspects of the Internet service (*e.g.*, zero-rating). Zero-rating is a common practice adopted by mobile operators that offer a bundle of applications whose traffic does not consume the monthly data cap. In this sense, the users have an incentive to use the applications within the bundle. Therefore, this practice

influences the user choice interfering in the market because users adopt applications based on the economic incentive instead of adopting them based on their functionalities.

Perceived Discrimination is the way users and detection solutions perceive the TD, which may be large delays, increased jitter, throttling, blocked traffic, or integrity violation. The detection of these effects is the typical approach used by solutions to determine the deployment of TD by ISPs. The techniques adopted to detect these effects are detailed in Section 2.4, where the state-of-the-art is presented.

The last mile ISP is commonly seen as one guilty for the deployment of TD. However, the TD can be deployed anywhere along with an end-to-end network path by different Internet operators, which may influence the proper technique for detection. Of course, the last mile ISP may be responsible for the deployment of TD. However, the TD may also be deployed by ISPs that operate in the core of the Internet, motivating proposals for the detection of TD by these operators (ZHANG; MAO; ZHANG, 2009).

TD also has been reported in distinct scenarios, such as Internet of Things (IoT) and 5G. IoT architectures may be highly affected by end-to-end delay, packet loss, and low throughput (GARRETT et al., 2018a). For instance, an IoT device or application that has its communication throttled may present worse Quality of Experience (QoE) to users. Thereafter, they would face difficulty to compete against favored or not impaired devices and applications (GARRETT et al., 2018a). 5G requires several architectural network changes to achieve its established challenges, such data rates 10 to 100 times higher and end-to-end latency lower than 5ms. Traffic management is among the techniques to accomplish such challenges with reasonable cost and better QoE (GUPTA; JHA, 2015). Another architectural network change related to 5G is the adoption of Network Slices (NSs) that are computational resources deployed within the network infrastructure to process traffic near to users. For instance, a NS could be deployed within a cellular antenna to process surveillance camera traffic without the need to contact the related cloud service. Such kind of network architectural change may also open opportunity to newer TD deployment scenarios. For instance, there are concerns of deployment of TD in NS environments by the provisioning of constrained resources to slices that process the traffic that need to be degraded (SMIRNOVA et al., 2019).

The TD may be motivated and triggered by many aspects affecting the communications in many ways. In the next section, these aspects and effects are presented within the NN debate context.

2.3 Network Neutrality

Network Neutrality (NN) is a principle that does not have a unique and widely accepted definition, which also exposes the controversy around the NN principle (MAILLÉ; REICHL; TUFFIN, 2012). Among the most used definitions is the one from the academia proposed by Tim Wu in 2003 (one of the earlier works to explore NN) and the one stated by the Federal Communications Commission (FCC), the USA regulatory agency, in 2015. Along with this section, definitions from academia and regulatory agencies are presented.

The most used definition for the NN principle was proposed by Tim Wu in 2003. At that time, the community was discussing a regulatory remedy, named Open Access regulation, to prohibit the vertical integration of cable operators with ISPs to bundle cable and Internet services. Such bundling could foreclose competition on the ISP market because cable operators would have the incentive to prioritize their Internet services or degrade competitors' services (KRÄMER; WIEWIORRA; WEINHARDT, 2013). Based on the results of a survey, Tim Wu finds that network operators (cable or Digital Subscriber Line (DSL)) would continue to have an interest in controlling Internet usage within their networks independently of the inability to offer Internet services by vertical integration (as prohibited by the Open Access regulation). The interest in controlling Internet usage is due to other aspects also identified in the survey, such as the prevention of the use of Home Networking and Virtual Private Network (VPN) services. Such findings contrast with the Open Access regulation objectives and the idea of an NN principle. Therefore, NN should require a broader remedy than Open Access regulation. Towards the statement of this broader remedy and an NN definition, Tim Wu explores the fact that ISPs are members of two networks: the local broadband network (their own network) and the inter-network (interconnection with other operators that, in the end, construct the Internet). He argues that operators should have the right to police the local broadband network because restrictions in this network are usually required for good network management. In contrast, restrictions on the inter-network, *i.e.*, in the traffic exchange with other operators, affect the whole network. Based on these arguments, Tim Wu's definition for the NN principle is as follows: "absent evidence of harm to the local network or the interests of other users, broadband carriers should not discriminate in how they treat traffic on their broadband network on the basis of inter-network criteria." (WU, 2003). This definition explores the NN principle concern of ISPs introducing TD based on unreasonable criteria, such as business arrangements or unfair competition. Indeed, it allows reasonable

network management activities to keep the local network healthy (*e.g.*, prevention of network misuse, avoidance of malware dissemination, and the establishment of data caps). However, it does not allow broadband discrimination based on inter-network aspects that could hinder the choice and access of applications by users (*e.g.*, blocking of VPN and e-mail services). He argues in a subsequent work that this NN principle could favor competition and innovation in the network edge (content and applications) without hindering the competition and innovation in the core network (WU, 2004).

The economic aspects of the Internet also leverage concerns in the NN debate. For instance, there is a concern about side payments that ISPs could charge from CPs and Application Providers (APs) to prioritize or deliver their traffic to users. These side payments could hinder innovation because new companies (CPs and APs) would have a higher cost to enter the market with similar Quality of Service (QoS) of their established competitors. Another concern is about extra payments that ISPs could charge from users to access specific contents, which are called specialized services or specialized networks. These extra payments could lead to the Internet fragmentation because users only would have access to portions of the network permitted by the ISPs. In the end, the Internet could become similar to a television package bundling that give access to specific content and applications according to the contracted package. In this sense, Hahn and Wallsten (2006) adopt a definition that tackles these concerns: “Net neutrality usually means that broadband service providers charge consumers only once for Internet access, do not favor one content provider over another, and do not charge content providers for sending information over broadband lines to end users” (HAHN; WALLSTEN, 2006). However, Hahn and Wallsten criticize a regulation that follow such an NN definition (seen by them as a price regulation). They discuss the distortions introduced by a price regulation from the point of view of the economic efficiency.

The NN principle is broader than economical concerns and the unreasonable traffic management practices. There are also concerns to preserve the open Internet nature. For instance, the prevention of architectural arrangements that hinder devices for joining the network, which is inspired in the Carterfone case (LESK, 2010). A similar architectural concern is the preservation of the end-to-end nature of the Internet that prevents that network core elements interfere with the communication of the endpoints. Another architectural concern is the recognition that the best effort Internet characteristic is not enough for new services that require better QoS parameters, such as IP Television (IPTV) and VoIP. In this sense, Crowcroft (2007) establishes a meta definition listing the ar-

chitectural aspects that a definition for the NN principle should cover. Crowcroft's meta definition is divided into four parts tackling the components that an NN principle definition should address: "*I - Connectivity Neutrality* must be defined concerning end-to-end service at every layer. *II - Performance Neutrality* must define rules for Service Level Agreements (SLAs) (existing ones and new ones with appropriate delay bounding services for IPTV) in a measurable, comprehensible and transparent fashion. *III - Service Neutrality* must define rules for availability of new net services such as multi-home, multi-cast, mobility, etc. in a way that allows cross-provider/cross-platform differences to exist until these services have sufficiently matured. *IV - Cross Layer Neutrality* must define how combinations of services are built and how the consumer gets to choose between them." (CROWCROFT, 2007). This meta definition covers architectural elements of the Internet, such as the end-to-end principle of every layer, the transparency of established SLAs commitments (that could require TD to be satisfied), the open standard of protocols and services (to guarantee interoperability), and the interactions between ISPs (to guarantee that users in one ISP can access services available on any part of the Internet).

The regulation also provides NN principle definitions established by regulators (legislators and regulatory agencies) within their jurisdictions. These definitions reflect the regulators' view about NN, harmonizing the interests of the Internet stakeholders. Therefore, these definitions tend to be more heterogeneous, reflecting aspects as national and corporation interests, social welfare, and maturity of the regulatory framework. Next, a few NN definitions from regulations are presented, exposing such heterogeneity.

Brazil has NN stated as an Internet principle since 2014 (Presidência da República, 2014). The Brazilian NN regulation states that ISPs must treat all packets equally. Few exceptions are allowed for traffic discrimination and degradation, which include emergency services and indispensable technical requirements. The discrimination and degradation cannot harm users and must be proportional, isonomic, and transparent. The ISPs are also prohibited from monitoring, blocking, filtering, or analyzing the packet contents. In 2016, the referred indispensable technical requirements that allow discrimination and degradation were detailed in a complementary law (Presidência da República, 2016), which are the technical requirements to handle security and traffic congestion issues.

Chile was one of the first countries to establish an NN regulation, which was established in 2010 (BUSTOS-JIMÉNEZ; FUENZALIDA, 2014). The Chilean NN regulation states that ISPs may not arbitrarily block, interfere with, discriminate, hinder or restrict the right of any Internet user to use, send, receive, or offer any content, application, or

service over the legal Internet and any other legal activity or use performed through the network (BUSTOS-JIMÉNEZ; FUENZALIDA, 2014).

The EU has an NN regulation since 2016 when the Body of European Regulators for Electronic Communications (BEREC) stated the regulatory instructions for all EU state members (Body of European Regulators for Electronic Communications, 2016). The European NN regulation states that ISPs must treat all packets equally. The regulation allows day-to-day traffic management, if it can be technically justified. Few exceptions are allowed for traffic blocking, discrimination, and degradation, including legal obligations, the integrity of the network, congestion management, and exceptional temporary situations. The regulation prohibits ISPs from adopting traffic prioritization.

The USA established its NN regulation in 2015 (Federal Communications Commission, 2015) which remained in force until 2018 (Federal Communications Commission, 2018). This regulation states three principles: no blocking, no throttling, and no paid prioritization. The no blocking principle guarantees that users have access to any lawful content. The no throttling principle guarantees that lawful content, applications, services, and non-harmful devices do not have their traffic impaired or degraded. The no paid prioritization principle prohibits ISPs from prioritizing traffic to earn side payments or other benefits from third parties. These three principles are frequently used by academic works as a definition to present the NN principle.

The aforementioned NN definitions from regulatory agencies are examples of regulation established *a priori*, also referred to as *ex-ante* regulation. However, the regulatory agencies and courts may also judge claims of supposed unfair traffic management practices by analyzing the nuances of each case against the jurisprudence (past decisions on similar cases) or the specific harms caused by the TD deployment. For instance, they can use anti-trust concepts from commercial regulation to judge whether ISPs deployed illegal traffic management practices because they incur unfair competition. The application of NN principles *post factum* is referred to as *ex-post* regulation.

The USA rolled back the FCC's 2015 NN regulation in 2018 (Federal Communications Commission, 2018), adopting an *ex-post* regulation without explicit NN definitions. The lack of explicit NN regulation drives the allegations of NN violations to the Federal Trade Commission (FTC), which is the bureau responsible for judging claims regarding anti-trust allegations. Such claims usually are related to unfair competition conducted by companies. Therefore, by being judged by the FTC, the traffic management adopted by ISPs should not hurt anti-trust commercial rules. However, non-neutral traffic

management actions (*e.g.*, blocking, throttling, paid prioritization) can be allowed if they can justify their motivations within an anti-trust claim proceeding.

The existence of divergent NN definitions both from the academia and regulatory agencies becomes evident. The definitions from the academia differ in their objectives ranging from inter-network criteria, economic effects, and Internet architectural concerns. The definitions from regulatory agencies vary in a broad spectrum ranging from explicit rules to implicit rules or no rule at all. Among the explicit ones, there are different ranges of which traffic management motivations are allowed for degradation and discrimination. Therefore, the NN principle has distinct definitions according to the objectives of who defines it (academia or regulator) or the place where it is being enforced. However, these multiple definitions share the NN principle essence that is the freedom of choice by users on the Internet use for legal purposes. In this sense, in a neutral network, legal uses of the network should not face obstacles of any nature, such as performance interference, disruption, or economical differentiation. In the next section, the state-of-the-art of solutions designed to detect Traffic Differentiation and NN violation is presented.

2.4 State-of-the-art

In this section, the solutions designed to detect TD and NN violation are presented. This thesis proposes that the detection of NN violation should be supported by regulations. Adopting this criterion, the detection of NN violations without the regulatory support is classified as TD detection. However, along with this section, the solution's goal (detect TD or NN violation) is presented as stated by their authors as well as the type of TD detected (*e.g.*, throttling, blocking). The techniques adopted by the solutions (single or aggregated tests) are presented as well as the criteria (traffic characteristic, statistical inference, or regulation) used to detect TD and NN violations.

It is important to note that solutions designed for measurement of performance (speed tests), QoS parameters, or SLA fulfillment that do not detect TD or NN violation, or that are not related to NN regulation are out-of-scope. Next, the solutions of the state-of-the-art are presented according to their criteria to detect TD or NN violations (traffic characteristic detection, statistical inference, and regulation).

One approach to detect TD is to look for traffic characteristics that indicate a differentiation is in place. For instance, a solution may look for the presence of TCP Reset (RST) packets artificially introduced in the communication, which is a common tech-

nique adopted for traffic blocking; a solution may test the connectivity of the application transport port; or a solution may collect metrics that directly points out to differentiation, such as content-related counters for content modification detection. Next, a few solutions that detect TD by inspecting traffic characteristics are presented.

BTTest (DISCHINGER et al., 2008) is a system to detect TD (blocking) of BitTorrent flows by detecting the injection of forged TCP RST messages that quits the communication between the endpoints. The system is composed of a Java applet (client that emulates the BitTorrent protocol) and a measurement server. On the client-side, due to the lack of administrative privileges, BTTest monitors the appearance of “IOException” messages accompanied by other error messages, such as “Connection reset by peer” or “An existing connection was forcibly closed by the remote host,” that indicates that a TCP RST message was received. On the server-side, BTTest captures network traces similar to tcpdump (TCPDUMP, 2022) to detect RST messages. The methodology involves varying 3 factors: TCP port (BitTorrent 6881 or unattributed 4711 port), direction (upload or download), and payload (BitTorrent protocol or random content). The combination of port and payload factors allow to detect how the ISP identifies the BitTorrent flow (port- or content-based classification). Three conditions are checked to point out that the ISP blocks BitTorrent traffic by injecting forged RST messages: an “IOException” accompanied by the messages above in the client-side, at least one incoming RST packet in the server-side, and no outgoing RST packet from the server-side before a Finalyse (FIN) or RST packet (to ensure the communication was not quit by the server by any reason).

Netalyzr (KREIBICH et al., 2010) is a system designed to perform tests to detect several network issues, which includes a few tests to detect TD (blocking and content modification). The system is composed of a Java applet client, one front-end server, and several back-end servers. The front-end server receives the browser’s user request and randomly chooses one back-end server (server from now on) to be responsible for the tests. The browser gets from the server the applet instrumented to perform tests against the same server. The applet tests connectivity to 25 well-known ports to detect blocking, including applications fully implemented by the system (Hyper Text Transfer Protocol (HTTP) (TCP/80) and DNS (User Datagram Protocol (UDP)/53)) and others that only a request/response application was implemented instead of the protocol assigned to the port. The applet also detects blocking based on content type by trying to download multiple file types: executable, MP3, torrent, and EICAR test file (EICAR - European Expert Group for IT-Security, 2022). The applet detects content modification by downloading

two images from the server and comparing them to the expected sizes and dimensions.

Another approach to detect TD is the application of statistical tests that infers that a differentiation is in place. For instance, a solution may compare the measured distribution of distinct flows to discover they have distinct performance; a solution may perform a confound factor analysis to discover that users from a specific ISP experience degraded performance when accessing an application; a solution may compare the results of a statistical measure from a sample to detect the differentiation; or a solution may build a system of equations with measurement results that only is solvable when the network is neutral. Next, a few solutions that detect TD by using statistical tests are presented.

Chkdif (RAVAIOLI; BARAKAT; URVOY-KELLER, 2012) is a tool for detection of NN violation introduced by Access ISPs. It replays traffic collected from the ordinary network usage of the user to perform measurements in both directions (upstream and downstream). Such network usage contains flows from multiple applications adopted by the user. Therefore, the tool does not need to be instrumented *a priori* to be able to measure a specific application. For the upstream measurement, the tool crafts the Time to Live (TTL) field of the replayed packets to target a specific router within the ISP network to evaluate only the access network. When the replayed packet arrives on the targeted router, it responds using the TTL exceeded Internet Control Message Protocol (ICMP) message. Then, Chkdif measures the Round Trip Time (RTT) of the upstream communication. For the downstream measurement, the tool sends the collected traffic to a server that replays it forging its IP identification field to allow the detection of the replayed traffic (among the usual traffic). The traffic is replayed at a constant rate that allow Chkdif to measure one-way delay of the downstream communication. The collection of measurements of all application flows for each direction is used to build Dirichlet random distributions (EBRAHIMI; SOOFI; ZHAO, 2011), which capture the behavior of other distributions. For each application flow, the measured RTT and the one-way delay also are used to build histograms for each direction. The probability that the histogram of one application flow belongs to the Dirichlet random distribution of all application flows determines whether the flow is facing TD (prioritization or throttling) in that direction. The rationale is that all application flows should have a similar distribution of RTT or one-way delay within the access ISP network.

NetPolice (ZHANG; MAO; ZHANG, 2009) (formerly NVLens (ZHANG; MAO; ZHANG, 2008)) is a system designed to detect TD (prioritizing and throttling) introduced by ISPs that operate at the core of the Internet. NetPolice crafts the TTL field of probe

packets to target the ISP's ingress and egress routers to detect TD introduced inside the ISP network. The solution measures the Packet Loss Rate (PLR) introduced within the ISP network by subtracting the egress router PLR by the ingress router PLR. NetPolice detects TD introduced based on routing information by selecting probes from distinct Autonomous Systems (ASs) as source or destination. It also detects TD introduced based on content information by varying the payload using 5 application patterns (HTTP, BitTorrent, Simple Mail Transfer Protocol (SMTP), PPLive, and VoIP). The samples are grouped in distributions according to the TD detection: by routing information (source or destination) or by content information (each application). The distributions are compared using the Kolmogorov-Smirnov (K-S) statistical test using $\alpha = 95\%$ to determine whether the distributions are distinct, which denotes that the ISP is performing TD.

Differentiation Detector (KAKHKI et al., 2015b) (also referred as Mobile Replay (KAKHKI et al., 2015a)) is a system designed to detect TD (throttling) in mobile networks. The system is composed of an Android App, a VPN server, and a Replay server. The Android App uses preset replay scripts of applications (*e.g.*, Netflix, Skype) to generate traffic within a VPN tunnel (control) and outside the tunnel (regular traffic). Other applications can also be measured by routing their regular traffic through a VPN server controlled by the solution to generate a replay script. These replay scripts recreate the application communication by mimicking its behavior: preserving the order of application messages, byte streams, and the inter-message times. The Replay server also records the communication to extract measurement information (*e.g.*, throughput, RTT, jitter) that are used to detect differentiation. The authors propose a statistical test named *Area test* that measures the area between the two Cumulative Distribution Functions (CDFs) resulted from the measured metrics for each traffic (regular and control). The measured area is normalized by the minimum peak of the two CDFs. In order to detect differentiation, the solution uses the results from both the K-S test (also used by NetPolice) and the Area test. If the K-S test detects differentiation and the normalized area is higher than 0.1 or 0.2, then the application is being differentiated by the ISP. The paper points out that these values achieve a good accuracy, but the chosen value for the experiments is not presented.

Wehe (LI et al., 2019) is a system designed to detect NN violation (throttling). Its architecture is very similar to Differentiation Detector (KAKHKI et al., 2015b) and Mobile Replay (KAKHKI et al., 2015a) consisting of a smartphone app and a replay server. However, the methodology is slightly different. Wehe fills the payload of control traffic with bit inverted application payload (contrasting with the use of VPN tunnels for trans-

mitting control traffic on earlier works). Wehe also replays the traffic in both directions (contrasting with replay only from the server to clients on earlier works). The K-S test is also adopted, but Wehe compares its results to results achieved using the Jackknife re-sampling method. If the average throughput from regular and bit inverted traffic differs at least in 10% and results from the whole sample and re-sampling are similar, then Wehe detects the throttling. Another difference is that Wehe also performs analysis using aggregated data, which is pointed out by authors to solve a few confound factors that may affect tests (*e.g.*, bandwidth volatility). In the aggregated analysis, the results of several user tests are grouped by ISP and tested application. Then, Wehe measure statistics about each ISP that help to analyze the adopted traffic management practices (*e.g.*, prevalence, changes over the time). This analysis also allows Wehe to determine the rate limit adopted by the ISPs, performing a Kernel Density Estimation (KDE) analysis that identifies the throughput that is most present in ISP-application pair results.

Li *et al.*'s work (LI et al., 2016) proposes a method to detect NN violation based on the PLR. The method estimates the PLR using a transformation of the Mathis model (which is designed to model throughput). For the PLR estimation, due to the Mathis model transformation, the method requires the throughput and RTT. The authors use the data collected by Neubot (BASSO; SERVETTI; DE MARTIN, 2011) hosted in M-Lab (MEASUREMENT LAB, 2022) to analyze their method by comparing the results for "speed test" (HTTP) and BitTorrent. They propose a Power and Exponential Composite Function (PECF) model to fit the speed test and BitTorrent PLR curves. Then, they estimate the point where both curves differ most. This difference is compared against the Lower Limit Packet Loss Rate (L-PLR) established by the authors by analyzing the Neubot data. They estimated that when L-PLR is greater than 0.10, there is a high probability of NN violation. By detecting the variations between the PLRs of two flows, the method can detect instances of throttling and prioritization.

Network Access Neutrality Observatory (NANO) (TARIQ et al., 2009) is a system designed to detect TD (throttling) using passive measurements. The system is composed of agents and servers. The agents are responsible for collecting information from user regular communications to estimate the performance of the adopted applications. These estimations are grouped with complementary information that may influence the application performance: about the user's machine (*e.g.*, Operating System (OS), Central Process Unit (CPU) and memory usages) and provided by the user (*e.g.*, Geographic location, ISP contract, SLA). In the server, the collected information from multiple users is used to

perform a confound factor analysis, establishing a causal inference between the observed performance and the ISPs. NANO builds strata for the results collected for each factor that reflect similar network and user machine conditions that allow the comparison of the observed performance. Within a stratum, the performance for one ISP is compared against the average performance (baseline) for the other ISPs in the same stratum. If the difference on the performance is statistically significant, then the ISP is deploying TD. When NANO detects differentiation, it can identify the discrimination criteria by using a decision tree-based classification. For this identification, it uses the confound factors to compose the feature set and the verdicts (the ISP performs discrimination or not) as the target variable. The criteria for the discrimination are identified by following the generated classification rules of the decision tree.

POPI (LU et al., 2010) is a system designed to detect TD (throttling and prioritization), measuring the PLR for 23 applications (transport ports). The system measures the PLR by saturating the available bandwidth with bursts containing traffic of all applications, revealing the differences among their forwarding priorities by observing different PLR. The observed PLR for each application is ordered to derive ranks from the higher to lower rates. When all applications are treated equally, the ranks should vary randomly. If an application is degraded (prioritized), its ranks tend to be lower (higher). The authors propose a method based on the Average Normalized Ranks (ANR) to group applications that experience similar PLR and may have similar forwarding priority. If more than one group of applications is identified, then the ISP is performing TD. The system also operates in two modes: end-to-end and host-by-host. The host-by-host mode can locate the devices that are introducing the TD, which are referred to as the *spot of difference*.

Packsen (WEINSBERG; SOULE; MASSOULIE, 2011) is a system designed to detect TD (throttling and prioritization) by comparing the performance of a baseline (*e.g.*, HTTP) and a measured flow (*e.g.*, BitTorrent). These flows are sent interleaved to face similar network conditions. The system uses three methods employed in a layered approach: one method to detect TD and two methods to infer the traffic shaping characteristics. To detect TD, the first method uses the Mann-Whitney U statistical test to compare the distributions of inter-packet arrival times of each flow. The test estimates whether the distributions are the same based on the ranks of the measures (a similar approach adopted by POPI). If the test infers that the samples have different distributions (using $p\text{-value} < 0.05$, *i.e.*, the distributions differ outside the 95% confidence interval), then the TD is detected. To detect the traffic shaper characteristics when a TD is detected, the other

two methods are performed depending on the cross-traffic. When there is no cross-traffic, Packsen adopts a method that consists of establishing the bandwidth ratio perceived by each flow to infer if the ISP deployed an SP scheduler (one flow dominates) or WFQ (there is a ratio between the flows). However, cross-traffic may influence the ratios observed, making this method inaccurate. To overcome the presence of cross-traffic, the authors propose another method that consists of assuming that a WFQ scheduler is configured to achieve a specific bandwidth ratio between the flows and that the cross-traffic arrives on the shaper following a Poisson distribution. Using the Poisson model, the measured expectancy, the measured variance, and other parameters yielded by the assumed WFQ scheduler, the method infers the capacity attributed to each flow and the cross-traffic rate.

ShaperProbe (KANUPARTHY; DOVROLIS, 2011) is a system designed to detect TD (throttling) by identifying the presence of traffic shapers along with the network path. Shapers usually adopt a queue management method (*e.g.*, leaky bucket) to conform the traffic to desired rates. These methods allow traffic to be transmitted in higher rates until a certain amount of data be transmitted (maximum burst rate). After this point, the traffic is transmitted on the desired lower rate (token generation rate). ShaperProbe detects the presence of shapers by observing the shift on the transfer rate from the higher to the lower one. The system is composed by clients (executable or Vuze BitTorrent plugin) and by servers hosted in the M-Lab platform. The client generates traffic at the rate of the narrower link (which is estimated by the system). On the server-side, the transfer rates (throughput) are measured within time ranges of 300 ms. The measures are compared online to detect transfer rate level shifts of at least 10% from higher to lower rate, indicating the presence of the shaper. When the system detects the shaper, it also identifies its parameters: maximum burst size and the token generation rate (conformed rate). If the system detects the shaper, then the ISP is performing TD.

Diffprobe (KANUPARTHY; DOVROLIS, 2010) is a tool that aims to detect TD using network tomography mechanisms (CASTRO et al., 2004), which consists of discovering internal aspects of the network using external observations (end-to-end measurements). By measuring delay and PLR, Diffprobe detects which kind of packet scheduling forwarding scheme (First Come First Served (FCFS), SP, WFQ) and packet dropping scheme (Drop Tail (DT), WRED) are deployed in the ISP routers. For the measurements, it uses an application flow (A) and a probing flow (P) (slightly different from flow A just enough to trick the Traffic Classification modules). To detect the packet scheduling forwarding scheme, DiffProbe applies a Kullback-Leibler (KL) divergence test (LEXA,

2004) to assess how distinct are the delay distributions (for A and P). Distinct delay distributions indicate that a discriminatory packet scheduling forwarding scheme is in place. Then, it identifies which scheme is deployed (SP or WFQ) by analyzing the delay variability. The SP scheduler transmits high priority flows first, therefore they have lower delay distributions. In contrast, the WFQ scheduler transmits low and high priority flows according to the specified weights, therefore high priority flows would have higher delay distributions (compared to a SP scheduler) because the scheduler transmits lower priority flows eventually. When packet dropping is detected, the tool identifies whether a discriminatory dropping scheme (WRED) is in place. The non-discriminatory schemes (DT, Random Early Detection (RED)) tend to drop packets at the same rate independently of the flow. If the flows have different drop rates, then the ISP deployed a WRED packet dropping scheme that tend to drop more packets from a lower priority flow. For DiffProbe, if a discriminatory packet scheduling forward or dropping scheme is detected, then the ISP is performing TD.

Zhang *et al.*'s work (ZHANG; MARA; ARGYRAKI, 2014) proposes a methodology to detect NN violation (throttling) also inspired in network tomography, which tries to infer internal network performance characteristics of links by external observations (end-to-end measurements), *i.e.*, without explicitly measuring them. The method builds a system of equations ($\vec{y} = A \cdot \vec{x}$) using external observations as the outputs (\vec{y}), the relationship between links and paths (A), and the properties that are being inferred (\vec{x}). The properties of internal links and paths are inferred by solving the system. In order to detect NN violation, Zhang *et al.* assume that if the network is neutral, then the system has a solution; otherwise, the network is not neutral. The authors also developed a methodology to identify the non-neutral links using a system of equations restricted to the links and paths of interest. Therefore, the criteria to detect NN violation is the presence of non-neutral links (which turn the system of equations unsolvable).

The approaches presented above (based on the detection of traffic characteristics or on statistical tests) can be mixed. Next, a few solutions that mix them are presented.

Glasnost (DISCHINGER *et al.*, 2010) is a system designed to detect TD (throttling and blocking) comparing the throughput and fail rate of application and reference flows. The reference flow mimics the application flow (*e.g.*, BitTorrent, HTTP) differing its port or content to trick the Traffic Classification. Glasnost varies the port between the usual application port (*e.g.*, TCP/6881 for BitTorrent) and a neutral port that is not assigned to any application (*e.g.*, TCP/10009). Glasnost also varies the content (payload)

between the standard application messages and a similar pattern (same order and size) but filled with random bytes. Each combination of port and content is tested in upstream and downstream directions and also is repeated twice. For throttling detection, Glasnost considers that the performance of flows should differ at least 50% (a high threshold to decrease the false-positive rate). The comparison of the throughput of flows varying the port and content allows Glasnost to identify the Traffic Classification mechanism (content- or port-based classification). For blocking detection, it performs similar comparisons using the transmitted bytes of flows or the information that the flow failed (*e.g.*, injection of forged TCP RST messages). The verdict presented to user (blocking, throttling, or no TD detected) is based on results collected in a single test, *i.e.*, the system does not aggregate results from distinct users to achieve the verdict.

Bonafide (BASHKO et al., 2013) is a system designed to detect TD (throttling and blocking) in mobile networks. The measurement methodology is inspired by Glasnost (DISCHINGER et al., 2010), but adapting the client-side to run on Android devices and tuning the measurement not to consume too much data to preserve the users' data cap. Inspired in Glasnost, the methodology consists of comparing the performance of application and reference flows in multiple measurement cycles. These flows are artificially generated for a set of applications (HTTP, FlashVideo (YouTube), Session Initiation Protocol (SIP), Real Time Streaming Protocol (RTSP), BitTorrent, and VoIP H.323). The reference flow is identical to the application flow in terms of message size and order, but changing the protocol messages and payload by random data. The solution decides if the application is suffering TD using two steps. The first step analyzes the failure rate (ratio of measurement cycles that fail by timeout or connection reset, a kind of blocking mechanism). If the failure rate of the random flow is $< 20\%$ and the failure rate of the application flow is $> 70\%$, then the application has suffered TD. The second step combines the Mann-Whitney U significance test and Student's t distribution. If the calculated U value is minor than an $U_{critical}$ (p -value < 0.05), then the application may have been differentiated. Then, the Student's t distribution is used for further examination of the measured values against the average, but the authors do not detail the used parameters.

Gnutella Rogue Super Peer (RSP) (BEVERLY; BAUER; BERGER, 2007) is a system designed to detect TD (blocking) inducing clients of a P2P overlay to perform connectivity tests of transport ports. The system deploys a Rogue Super Peer that is contacted by clients trying to join the overlay. The RSP responds to all requests with a "busy message" but referring to an IP and port to contact another Super Peer (SP).

However, the referred IP and port are from the Measurement Server that listens to all ports and redirects requests to actual SPs. The system associates clients to their ISP using their IP address and Border Gateway Protocol (BGP) Routing Information Base (RIB). When one client from one ISP (BGP prefix) connects to one transport port of the Measurement Server, then this ISP does not block this port. However, if the client does not contact the Measurement Server in the referred port, it does not mean that its ISP blocks such port. For instance, the client may choose not to follow the referral or may give up to join the overlay. Therefore, the authors use a probabilistic inference assuming that 90% of the distributed referrals are not followed by clients. They found that 50 referrals must fail to have 99.5% of probability that the ISP is indeed blocking a port. Thus, if the RSP distributes 50 referrals indicating a single port to clients of the same ISP and none of these clients contacted the Measurement Server in such port, then the ISP is blocking the port.

NeutMon (GREGORI; LUCONI; VECCHIO, 2018b)(GREGORI; LUCONI; VECCHIO, 2018a) is a system designed to detect NN violation (blocking, throttling, and prioritization) in mobile networks. The system detects differentiations that block traffic, limit the application bandwidth or route the application's packets through slower or more congested links. For the measurements, the system uses a synthetic BitTorrent traffic and a control traffic that mimics the BitTorrent protocol behavior (same payload size and inter-packet time) but filled with random bytes. For blocking detection, if the BitTorrent traffic fails on the standard TCP/6881 port, the system performs further tests: BitTorrent traffic on a random port and control traffic on the standard BitTorrent and the same random port. The combination of the results allow to identify port- and content-based blocking. For the detection of bandwidth-based differentiation, the system compares the throughput (speed test) achieved by the BitTorrent and control traffic. For the detection of routing-based differentiation, the system identifies the hops along with the path between the client and the server using an approach similar to traceroute (TTL manipulation). The system takes advantage of the connection used in the speed test to perform the path identification by crafting the TTL field of packets, in contrast to existing traceroute approaches that do not establish the connection at the transport level. In this way, the mechanism identifies the hops that the application traffic is traversing because the identification uses the same connection. The system builds sets of addresses traversed by the application and control traffic in the i^{th} hop. Then, the system computes the cardinality of the sets of addresses traversed by the application traffic, excluding the addresses traversed by the control traffic and vice-versa. High set cardinalities indicate the ISP adopts routing-based differentiation

using transport or application level information.

The above solutions detect NN violation or TD using criteria that are intrinsically related to the methodology employed in the detection mechanism (presence of traffic characteristic and/or statistical inference). In turn, the regulation establish the traffic management practices that are considered NN violations and exceptional situations that such practices are allowed. Therefore, the regulation establish higher level criteria to point out NN violation that complement the lower level criteria adopted by detection solutions. Next, a few solutions that consider the higher level criteria established on the regulation to point out NN violation are presented.

Adkintun (BUSTOS-JIMENEZ et al., 2013) (BUSTOS-JIMÉNEZ; FUENZALIDA, 2014) is a platform composed of servers and probes designed to monitor the Chilean NN regulation. The probes may be deployed as a client application or as an instrumented wireless router. The probes measure a set of metrics (latency, jitter, bandwidth, IP changes, availability, and packet loss) to characterize the broadband Internet quality according to the regulation. End-users may access a portal to get aggregated measures for their ISP, service level, and the quality perceived by other users. However, there is no detail about the methodology adopted by Adkintun to detect the NN violation. In addition, by considering only the Chilean NN regulation, Adkintun cannot be used to point out violation to NN definitions stated elsewhere. Therefore, it cannot be considered as a solution to point out NN violation in an end-to-end network path that may traverse multiple jurisdictions having distinct NN definitions.

ISPANN (SCHAURICH; CARVALHO; GRANVILLE, 2018) is a tool designed to detect NN violation (blocking, throttling, and prioritization) by auditing the configuration of ISPs network devices to find instructions that may violate the NN principle. ISPANN collects such configurations directly from the network devices through management interfaces (*e.g.*, OpenFlow). For the audit assessment, the ISP network administrator or an auditor indicate the country where the ISP operates, which selects the rules established on the NN regulation for that country. For each rule, the tool has an algorithm to detect NN violation by inspecting the network configuration. ISPANN detects blocking by finding a drop packet rule and prioritization by finding an Openflow flow priority rule for certain application. It also measures latency and port load for the user flow paths. Then it compares the user flow path latency and port load against a threshold latency and port load (average + standard deviation of other flow paths). If it finds another alternative flow path with lower latency or port load, then it detects that ISP is degrading traffic by routing it

through worst paths. It is important to note that this is the first work that recognizes the importance of accounting for the multiple definitions for the NN principle that may be found in different countries. However, ISPANN is a network configuration auditor that look for instructions that may violate the NN principle. Therefore, it is not a solution for the detection of end-to-end NN violation.

There are also proposals that are not directly designed for the detection of TD or NN violations but that can detect related information. Biczók *et al.* (BICZÓK; YOUNG; KUZMANOVIC, 2008) proposes a system to identify if the ISP deploys traffic filters and shapers (middleboxes). Kuzmanovic and Knightly's work (KUZMANOVIC; KNIGHTLY, 2001) proposes a statistical inference method to detect and characterize (discovering the classes and its weights) the deployment of rate limiters (leaky bucket) and scheduler mechanisms (SP, WFQ, Early Deadline First (EDF)) deployed in multi-class systems. Open Observatory of Network Interference (OONI) (FILASTÒ; APPELBAUM, 2012) is a framework devoted to hosting network tests that can be designed for detection of TD, such as blocking, censorship, and traffic interference.

In Table 2.1, the state-of-the-art is summarized, focusing on the aspects relevant to the discussion performed in this thesis. Next, these aspects are presented.

The first aspect is the goal of the detection. There are solutions focused on detecting TD or NN violation. This thesis advocates that an NN violation need to be pointed out according to definitions from regulations. It is important to note that the table reflects the goals stated by the authors in their respective work. Therefore, the state-of-the-art is composed of solutions (Chkdiff, Wehe, Li *et al.*'s work, Zhang *et al.*'s work, and Neutmon) that detect NN violation but without the support of regulations.

The second aspect is the type of TD the solution is able to detect. It is important to note that throttling and prioritization reflect similar traffic differentiation objectives. For instance, to prioritize one traffic, others may need to be throttled; by throttling one traffic, others are being prioritized. However, the classification in the Table 2.1 reflect the objectives declared by solution authors.

The third aspect is the detection technique that can be based on a single test or on aggregated data analysis. This aspect influences in what step the regulations should be assessed. Solutions based on a single test present their results to users just after the end of the test. Therefore, the regulations should be assessed just after the test and before presenting the results to users. Solutions based on aggregated data perform their analysis in batch. Therefore, the regulations should be assessed along with the batch analysis.

Table 2.1 – The state-of-the-art

Solution	Detection		TD				Technique		Criteria		
	TD	NNV	T	P	B	CM	S	A	C	SI	R
BTTest	✓				✓		✓		✓		
Netalyzr	✓				✓	✓	✓		✓		
Chkdiff		✓	✓	✓			✓			✓	
NetPolice	✓		✓	✓				✓		✓	
Differentiation Detector (Mobile Replay)	✓			✓			✓			✓	
Wehe		✓	✓				✓	✓		✓	
Li <i>et al.</i> 's work		✓	✓	✓				✓		✓	
NANO	✓		✓					✓		✓	
POPI	✓		✓	✓			✓			✓	
Packsen	✓		✓	✓			✓			✓	
ShaperProbe	✓		✓				✓			✓	
DiffProbe	✓		✓	✓			✓			✓	
Zhang <i>et al.</i> 's work		✓	✓					✓		✓	
Glasnost	✓		✓		✓		✓		✓	✓	
Bonafide	✓		✓		✓		✓		✓	✓	
Gnutella RSP	✓				✓			✓		✓	
Neutmon		✓	✓	✓	✓		✓		✓	✓	
Adkintun		✓	✓	✓	✓			✓			✓
ISPANN		✓	✓	✓	✓		✓		✓	✓	✓

Legend: TD - Traffic Differentiation, NNV - NN Violation, T - Throttling, P - Prioritization, B - Blocking, CM - Content Modification, S - Single, A - Aggregated, C - Characteristic, SI - Statistical Inference, R - Regulation

Source: the author

The fourth aspect is the criteria used to detect TD or NN violation. The Table 2.2 provides further details about the criteria adopted and the metrics observed by solutions. The criteria can be the presence of a particular traffic characteristic. For instance, the presence of TCP RST packets. The criteria may be statistical inference depending on the methodology deployed, usually defined by a threshold in the adopted metric. The criteria may be the definitions established on regulations. Although both Adkintun and ISPANN consider the definitions stated on the regulation, they are not general approaches to tackle the problem of pointing out NN based on regulations. ISPANN is a network auditor designed to help administrators to assess their network configurations according to the regulation. The jurisdiction issue is solved, in this context, by the administrator choosing the correct regulation to assess the network configuration according to the country where the ISP operates. In turn, for Adkintun, the solution only considers the Chilean regulation. Therefore, the pointing out of NN violation based on regulations stated in other jurisdictions is not possible.

Table 2.2 – The criteria adopted by each solution

Solution	Metrics	Criteria
BTTest	–	“IOException” and related messages in the client and one TCP RST in the server
Netalyzr	Connectivity and image attributes	Connectivity to 23 applications and a few content types and comparison to expected size and dimensions of two images
Chkdifff	RTT (upstream) and one-way delay (downstream)	Application flow does not belong to the Dirichlet random distribution of all flows
NetPolice	PLR	K-S test ($\alpha = 0.95$)
Differentiation Detector (Mobile Replay)	Throughput	K-S and Area tests (normalized area higher than 0.1 or 0.2)
Wehe	Throughput	K-S test (p-value < 0.05) and JackKnife achieves the same result in 95% of time and difference in throughput is higher than 10%
Li <i>et al.</i> 's work	PLR	Difference between the PLR of two flows is higher than 0.10
NANO	Throughput	Statistically significant performance difference between strata
POPI	PLR	Authors' method detects more than one group of application with similar PLR
Packsen	Inter-packet arrival time	Mann-Whitney U test to detect difference among flow distributions (p-value < 0.05)
ShaperProbe	Throughput	Rate at least 10% lower than the higher rate
DiffProbe	Delay and PLR	KL test for delay distributions and analysis of the delay variability; differences among the PLRs
Zhang <i>et al.</i> 's work	Additive performance metrics	System of equations built using external observations and relationships between links does not have a solution
Glasnost	Throughput and fail rate	Throughput of application and reference flows differ at least 50% for throttling; fail rate for blocking
Bonafide	Failure rate (timeout and presence of RST packets) and goodput	Failure rate of reference flow $< 20\%$ and application flow $> 70\%$; goodput distributions has Mann-Whitney U significance test (p-value < 0.05) higher than an $U_{critical}$ and Student t test to evaluate the goodput averages (no detail about the parameters)
Gnutella RSP	Failure rate	50 distributed referrals to users of the same ISP to a port that does not connect to the server
Neutmon	Fail rate and throughput	Fail rate comparison for blocking; throughput CDF comparison for throttling; and addresses set cardinalities for routing-based differentiation
Adkintun	Latency, jitter, throughput, IP changes, availability, PLR	No detail about the detection methodology
ISPANN	Latency, port load, priority	Drop packet configuration, latency/port load higher than average + standard deviation latency/port load of all paths and better alternative path, OpenFlow flow priority field

Source: the author

2.5 Summary

In this chapter, concepts related to this thesis were presented. First, aspects related to Internet regulation were presented because this thesis proposes that it should be the source of definitions to point out NN violations. The jurisdiction concept is also presented because it determines what regulation must be considered. Second, concepts related to the traffic differentiation that may violate the NN principle are presented. For instance, this thesis presents the triggers that lead an ISP to deploy TD, the traffic classification mechanisms that define the traffic priority, the differentiation mechanism that is how the ISP deploy the TD, and the perceived discrimination that is the effect perceived by users and applications when the TD is in place. Third, definitions for the NN principle stated by the academia and regulatory agencies are presented. Both sources of definitions impose heterogeneity because there is no common and widely accepted definition for the NN principle by academia and definitions from regulators are valid only on their jurisdictions. Last, the state-of-the-art solutions designed to detect TD and NN violation are presented. The focus is on the goal of the proposed solution (detect TD or NN violation), the kind of TD detected, the technique adopted to perform the detection, and the criteria used to detect TD and NN violation.

This thesis advocates that NN violation should be pointed out according to NN definitions established on regulations. In the next chapter, an impact analysis of this approach is performed by analyzing the results achieved by a state-of-the-art solution using the NN definitions from the regulation.

3 IMPACT ANALYSIS

This thesis proposes a methodology that is based on the NN definitions established by the regulation to point out NN violation. However, this methodology may impact the results that has been achieved by state-of-the-art solutions. In this chapter, an analysis is conducted to quantify such impact by evaluating the results achieved by state-of-the-art solutions but interpreting them using definitions from the regulation. The analysis aims to answer the following questions:

- *(i)* How much different regulations agree whether a particular TD is considered as an NN violation?
- *(ii)* How many of the NN violations detected by state-of-the-art solutions remain as violations considering the regulation?
- *(iii)* How much influence may the interpretation of regulatory instructions have over the results?
- *(iv)* How much influence the changes in the regulatory instructions had along with the time?

In order to answer these questions, the results publicized by state-of-the-art solutions for the detection of NN violation are analyzed checking whether the detected TDs are, in fact, NN violations considering the corresponding regulation. In Section 3.1, the public dataset of Glasnost (Measurement Lab, 2016) is presented. At the time that this analysis was performed, it was the only one that provides the network captures and meta-data that caused the TD verdicts. The data from 2016 was evaluated because it is the most recent complete year available in the dataset. Three hypothetical scenarios were designed to help answer the previous questions. These scenarios explore changes in two factors: the regulatory interpretations and the date that the TD was detected, which affects the regulatory instructions that need to be considered. In Section 3.2, the regulatory instructions from the jurisdictions traversed by the tests (countries and regions) that were considered in the analysis are presented. In Section 3.3, the quantitative results of this impact analysis are presented, discussing also its relevant aspects and the findings that are used to define the requirements for the models proposed in Chapter 4.

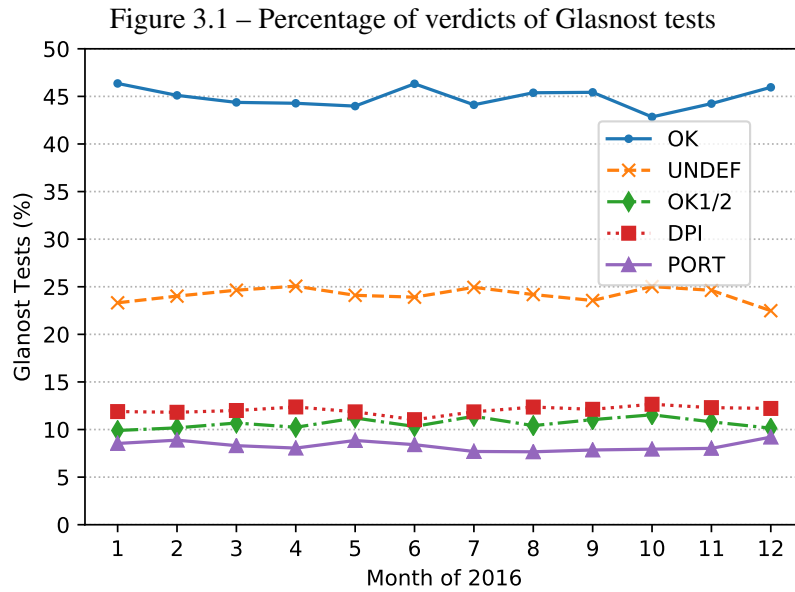
3.1 Glasnost dataset

In this section, the Glasnost dataset is presented, focusing on the aspects that are important for the analysis conducted in Section 3.3, in which the influence of regulatory instructions over the detection of NN violation is quantified. By focusing on these aspects, most of the results presented here have not been explored in the literature before.

Glasnost detects throttling and blocking. For the deployment of these TDs, traffic classification may be applied based on packet content (DPI) or application port usage (PORT). Throttling is detected based on the comparison of the performance of an application flow against reference flows. The reference flows are modified in their content (to bypass DPIs) or their ports (to bypass port-based classification). Blocking is detected when one or more of these flows fail to connect. Application and reference flows are tested in both directions (upload and download).

Glasnost was hosted on M-Lab from 2009 to 2017 (Measurement Lab, 2016). The collected data is available along with the parser to build a dataset of detected TDs. At the time that this analysis was conducted, Glasnost was the only solution that provided network captures of the detected TDs alongside all metadata information, which allows building a suitable dataset to research the detected TDs. Later, Wehe (LI et al., 2019) also publicized its dataset of detected TDs that allows a similar analysis. Indeed, the Wehe dataset is used in the analysis conducted in Chapter 6. In turn, the Glasnost dataset was analyzed in the impact analysis conducted in this chapter using data ranging from January to December 2016, the most recent 12-month window with a stable amount of tests, which contains more than 2TiB of data. The dataset consists of test logs and dumps (PCAP format). The test log has performance information of the flows during the tests. The dump is kept as evidence when the test detects a TD.

The parser (Max Planck Institute for Software Systems, 2022) needs just the test logs to generate the dataset of detected TDs, which comprises 362GiB of logs for the 12 months of 2016 chosen for the impact analysis conducted in this chapter. The parser has three main tasks. *Import logs* reads the data from logs and imports them into the database. *Update geo&asn* annotates the database with the country and Autonomous System Number (ASN) of clients based on their IP addresses using GeoLite2 (Maxmind, 2022) and PyASN (Economics of Cybersecurity Research Group, Delft University of Technology, 2022), respectively. Finally, *parse&analyze* annotates the database with the verdicts of the tests.



Source: (CARVALHO et al., 2020)

Table 3.1 – Verdicts numbers


Conclusion	Verdicts	# of Tests	% Verdicts	% Conclusion
No violation	OK	27954	45.2	45.2
Inconclusive	UNDEF	14823	24.0	34.6
	OK1/2	6520	10.6	
TD detected	DPI	7372	11.9	20.2
	PORT	5108	8.3	

Source: (CARVALHO et al., 2020)

After the parser’s processing, the dataset has the verdicts of the 61773 Glasnost tests conducted during 2016. The parser judges the tests as “OK,” “UNDEF,” “OK1/2,” “DPI,” or “PORT.” “OK” means that no evidence of a TD was detected. “UNDEF” can mean that the test was noisy or had the port changed during the test. “OK1/2” means that the test was “OK” in one direction (upload or download) but faced a problem in the other direction. “DPI” means that a TD (throttling or blocking) was detected based on the packet’s content. “PORT” means that a TD was detected based on application ports.

In Figure 3.1, the percentage of verdicts of Glasnost tests along the year of 2016 is presented. These percentages remained steady throughout the year. The values are very close to the overall percentages for the year presented in Table 3.1. The table also presents the conclusions made over these verdicts: TD detected (DPI and PORT-based) at 20.2%, inconclusive tests (OK1/2 and UNDEF) at 34.6%, and no TD detected (OK) at 45.2%. Although not presented in Figure 3.1 and Table 3.1, within the 12480 TDs detected by Glasnost in 2016, the proportion of blocking (2781, $\approx 22\%$) and throttling events (9699, $\approx 78\%$) is also relevant and influences the impact analysis results.

Figure 3.2 – Distribution of countries in Glasnost tests



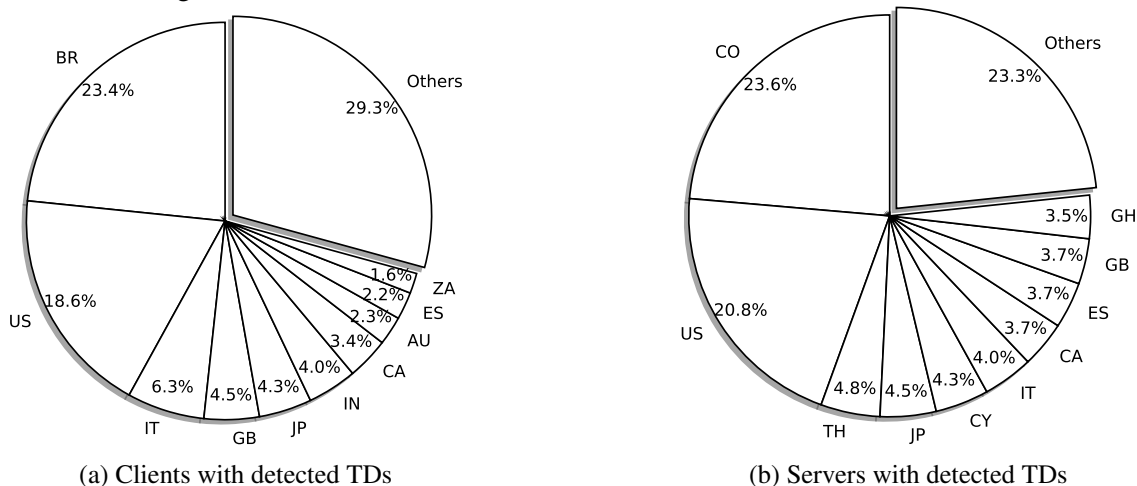
Source: (CARVALHO et al., 2020)

The dataset comprises TDs that affected different applications represented by different proportions: BitTorrent (63.0%), FlashVideo (17.3%), HTTP (11.3%), NNTP (3.5%), eMule (1.8%), SSH (1.3%), POP3 (0.8%), IMAP (0.5%), and Gnutella (0.3%). These differences do not mean that one application is most differentiated than others. Instead, it is related to the application selected by users in Glasnost’s interface (BitTorrent is the default test). As the Glasnost is a tool similar to a speed test, the selected application is the only one tested, *i.e.*, there is no agent on the user’s machine performing collection of other application’s traffic. However, this skew in the proportion is important to take into account when different interpretations of the regulatory instructions related to P2P traffic are applied, as presented in Section 3.3 since they represent 65.1% of the tests.

The Glasnost parser performs the Geo Location of clients based on their IP addresses and ASNs but such information is not available for servers. However, the dataset specifies the M-Lab server involved in the test that is identified by the International Air Transport Association (IATA) code of the closest airport. Based on this characteristic, a table was built with the respective country of each M-Lab server using PyAirports (Data61, 2022), allowing to determine the country of the client and server of each Glasnost test.

In Figures 3.2a and 3.2b, the top countries in the number of tests acting as clients or servers, respectively, is presented. The dataset comprises tests from 196 countries as clients and 29 as servers. The results show that the distribution is slightly off between clients and servers. Brazil is the second country as a client, but it does not appear in the ranking as a server because there was no M-Lab server in Brazil in 2016. The tests from Brazil were likely routed to Colombia, which ranks second as a server but does not

Figure 3.3 – Distribution of countries whose Glasnost tests detected TDs



Source: (CARVALHO et al., 2020)

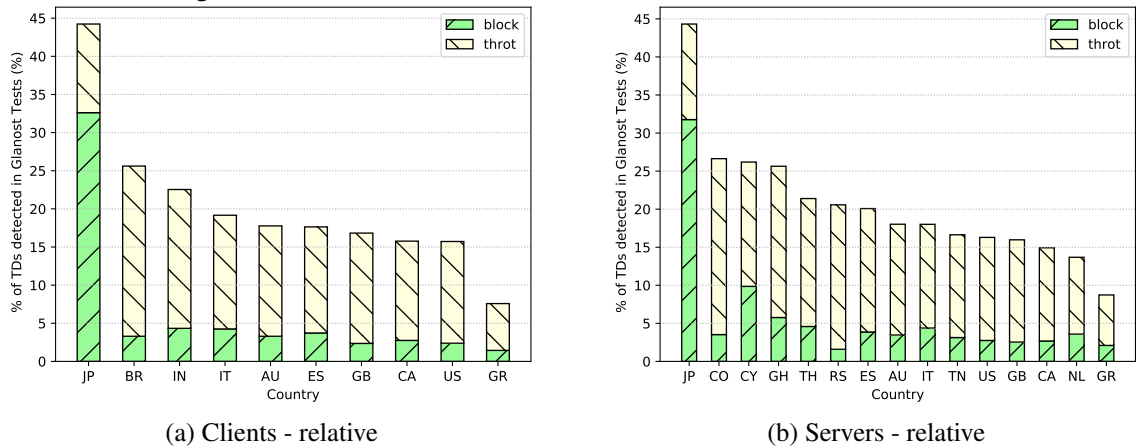
rank at the top of clients. The same applies to India that had its tests routed to Thailand. This discrepancy between the number of clients and servers exposes that a representative number of tests traverse country borders, thus subjected to multiple NN definitions.

In Figures 3.3a and 3.3b, the top countries in the number of detected TDs as clients and servers, respectively, is presented. Glasnost detected TDs in tests of clients of 165 countries and servers of 29 countries. Comparing Figures 3.2a and 3.2b against Figures 3.3a and 3.3b, the results show that their distributions are very different. Brazil and Colombia became the first countries as clients and servers in detected TDs, respectively. Japan also climbs several positions in the top countries with detected TDs. These facts expose that the proportion of tests that detected TDs is different in each country.

In Figures 3.4a and 3.4b, the top countries with the highest percentages of detected TDs as client and server, respectively, is presented. Countries that had less than 1000 tests were filtered out to avoid noise because there are countries with few tests and a very high proportion of TDs. Japan is the first one in the proportion of TDs as clients and as servers. The majority of blocking is also interesting, contrasting with the overall proportion of blocking and throttling presented before (22% vs. 78%). Just Japan showed this behavior, exposing a massive appetite for blocking by Japanese ISPs. Few countries have proportions above the overall proportion of TDs ($\approx 20\%$): Japan, Brazil, and India as clients; and Japan, Colombia, Cyprus, Ghana, Thailand, Serbia, and Spain as servers.

This section shows that there is a representative number of tests that traverse country borders. Therefore, these tests need to be analyzed, considering the definitions stated in multiple regulations. The results also showed that the proportion of detected TDs and

Figure 3.4 – Distribution of countries with relative % of TD over tests



Source: (CARVALHO et al., 2020)

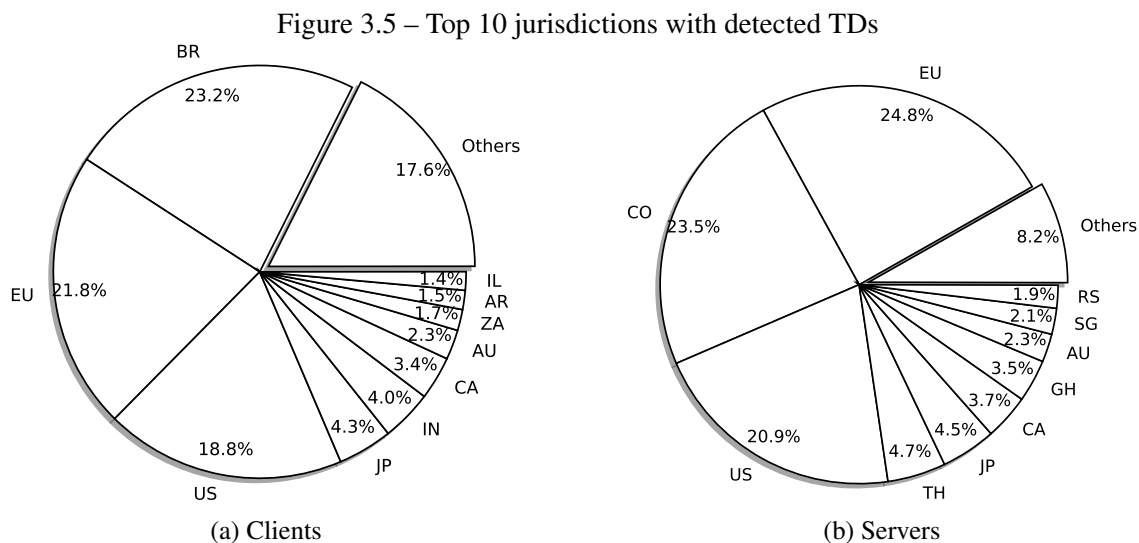
even the types of TDs deployed are different in each country. In the next section, the regulatory instructions from the top jurisdictions with detected TDs are presented, showing the heterogeneity of these definitions. Indeed, an analysis ignoring the existence of multiple and heterogeneous definitions may be inaccurate or even incorrect.

3.2 NN definitions from regulations

In this section, the regulatory instructions established in the jurisdictions considered in the analysis conducted in Section 3.3 are presented. As Glasnost detects throttling and blocking, only the regulatory instructions that are specific to these TDs are presented.

Figures 3.5a and 3.5b present the jurisdictional view of the data related to the tests that detected TD, previously presented in Figures 3.3a and 3.3b without minding for jurisdiction. For instance, the EU countries and their territories are grouped into the EU jurisdiction in these charts. The top 10 jurisdictions for clients and servers are presented in Figures 3.5a and 3.5b, respectively. The results show differences from the data presented in Figures 3.3a and 3.3b. In Figures 3.5a and 3.5b, EU is the second jurisdiction for clients, and the first one for servers, because now it groups all tests related to EU countries. Despite accounting for the TDs of the USA territories, the percentage of TDs remains unchanged due to their small contribution (just 23 TDs). As the countries and territories were grouped into their jurisdictions, other countries emerge in the Top 10 jurisdictions.

In Figures 3.6a and 3.6b, the top jurisdictions in the proportion of detected TDs is presented, which was previously presented in Figures 3.4a and 3.4b without minding for jurisdiction. EU appears as the fifth jurisdiction for clients, and seventh for servers,



Source: (CARVALHO et al., 2020)


because now it groups the results of all EU members, even those that do not appear in Figures 3.4a and 3.4b due to a small number of tests.

In Section 3.3, the analysis considers the regulatory instructions from the Top 10 jurisdictions for servers and clients (Figures 3.5a and 3.5b). As some jurisdictions appear in both Top 10s (clients and servers), the analysis considers regulatory instructions from 15 jurisdictions as detailed in Table 3.2 covering $\approx 81\%$ of the detected TDs.

The academic literature was researched to find information about the regulation of each jurisdiction. When this research did not yield satisfactory results, information from the regulatory agencies responsible for the jurisdiction was researched. The regulations that are written in English, Portuguese, and Spanish were inspected in their original form. For other languages, they were translated using Google services. Due to this methodology, it is essential to note that inconsistencies may be found, which may impose a negligible impact on the results. Despite of this, the argument that heterogeneous definitions of NN exist across different regulations still applies. The point is that different definitions need to be considered, without mattering which particular instances. The work (GARRETT et al., 2022) referenced in a few cases is an abundant source of references and served as an index to the different regulatory frameworks available worldwide. The regulatory instructions of each jurisdiction are presented next.

Some jurisdictions do not specify exceptions in their regulations or have not established regulations. Brazil (BR) is the only jurisdiction that prohibits throttling and blocking (Presidência da República, 2014) without exceptional situations either. India (IN) established its former regulation focusing on Internet fees (Ministry of Communications,

Figure 3.6 – Top 10 jurisdictions with relative % of TD over tests



Source: (CARVALHO et al., 2020)

2016), thus allowing throttling and blocking without exceptions. For some countries, evidence of the lack of regulation was found, such as Australia (AU) (GARRETT et al., 2022), Serbia (RS) (BDK Advokati, 2017), and South Africa (ZA) (SETENARESKI, 2017)¹. For others, evidence either of the lack of regulation or its existence was not found, such as Ghana (GH) and Thailand (TH), where the lack of regulation is supposed. Thus, TD is considered allowed in these countries in the analysis.

Some jurisdictions prohibit throttling and blocking while still allowing exceptions for both. Argentina (AR) includes exceptions for blocking, such as judicial claims or user requests (*e.g.*, parental control) (Senado y Cámara de Diputados de la Nación Argentina, 2014). Colombia (CO) includes exceptions for throttling (*e.g.*, congestion avoidance, security) and blocking (*e.g.*, prohibited or restrict use content, parental control) (Colômbia. Congreso Nacional, 2011). Israel (IL) includes exceptions that may be defined by the prime minister (KNESSET, 2014). IN includes exceptions for throttling and blocking (*e.g.*, emergencies, restrictions on unlawful content, security, and integrity of the network) in their regulation stated in 2018 (Telecom Regulatory Authority of India, 2017).

Some other jurisdictions prohibit throttling or blocking but allow exceptions only for one of them. EU establishes that throttling may be allowed under certain situations that the National Regulatory Authority (NRA) shall evaluate and decide (Body of European Regulators for Electronic Communications, 2016). The USA allowed blocking of content deemed illegal in its former NN regulation.

Some jurisdictions allow throttling and blocking, as long as they can be considered justifiable, such as Canada (CA) (CRTC - Canadian Radio-Television and Telecommuni-

¹After this analysis, Garrett et al. 2022 identified the South African NN regulation

Table 3.2 – NN definitions that were considered

Jurisdiction	NN regulatory instructions			
	begin validity	end validity	throttling	blocking
AR	2014-12-18	still valid	X*	X*
AU	–	–	✓	✓
BR	2014-04-23	still valid	X	X
EU	2015-06-30	2016-08-29	X*	X
	2016-08-30	still valid	X*	X
CA	2009-10-21	still valid	✓*	✓*
CO	2011-06-16	still valid	X*	X*
GH	–	–	✓	✓
IL	2014	still valid	X*	X*
IN	2016-02-08	2018-07-10	✓	✓
	2018-07-11	still valid	X*	X*
JP	2006	still valid	X*	✓
RS	–	–	✓	✓
SG	2011-06-16	still valid	✓*	X*
TH	–	–	✓	✓
US	2015-02-26	2018-06-10	X	X*
	2018-06-11	still valid	✓*	✓*
ZA	–	–	✓	✓

allowed (✓), prohibited (X), has exceptions (*)

EU jurisdiction = AT, AW, AX, BE, BG, BM, CW, CY, CZ, DE, DK, EE, ES, FI, FO, FR, GB, GF, GI, GL, GP, GR, HR, HU, IE, IT, KY, LT, LU, LV, MQ, MT, NC, NL, PF, PL, PT, RE, RO, SE, SI, SK, SX, TC, VG, YT
 US jurisdiction = GU, PR, US, VI

Source: (CARVALHO et al., 2020)

cations Commission, 2009) and USA after 2018-06-11. Thus, providers can perform the traffic management deemed necessary without explicit prohibitions established *a priori*.

Some jurisdictions prohibit one practice but allow the other. Japan (JP) prohibits throttling but allows exceptions such as for P2P or high demand users (CARTER et al., 2010). Evidence of the prohibition of blocking in Japan was not found, which may explain the already mentioned appetite for blocking by Japanese ISPs. Singapore (SG) allows throttling as long as justifiable (reasonable traffic management) but prohibits the blocking of legal content (Info-communications Development Authority of Singapore, 2011).

General comments about the regulatory instructions presented need to be pointed. Despite the more restrictive instructions in the EU regulation established in 2016 (compared to 2015), the conclusion is that there is no significant difference about throttling and blocking between both regulations. Japan has NN regulation established since 2006. Evidence that this regulation remains the same since 2006 was found because there are guidelines, discussed in 2018, to establish new rules (JITSUZUMI, 2018). The USA faced a hot debate around NN. For the sake of simplicity, the analysis considers just the two major eras in the USA NN regulation, the FCC regulation era and the current lack of regulatory instructions era, ignoring the legal battles that suspended or re-established each regulation in specific periods (GARRETT et al., 2022).

After inspecting these regulations, a wide range of regulatory instructions established around the world is shown, differing in their allowances, prohibitions, and exceptional situations. This finding confirms the concerns that motivated this thesis that the NN violation detection needs to be based on regulatory definitions due to such differences.

3.3 Results and discussion

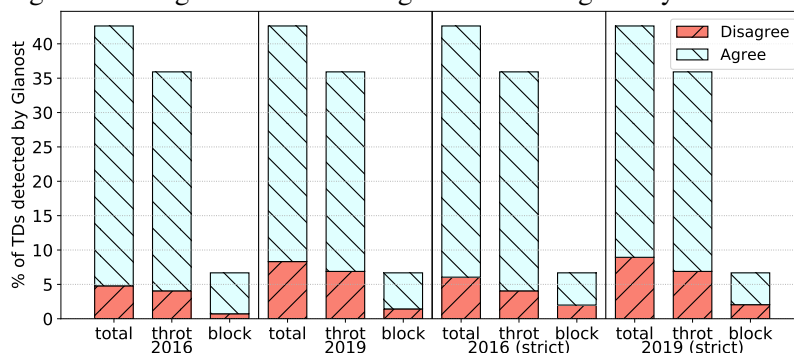
In this section, the analysis of the Glasnost dataset introducing NN definitions from the regulations is presented. How much these regulations influence the verdicts is quantified, given that TDs were pointed out as NN violations without the regulatory interpretation. All results were gathered through Structured Query Language (SQL) queries in a PostgreSQL database that holds the parsed Glasnost dataset.

3.3.1 Introducing regulatory instructions into Glasnost's results

First, the number of detected TDs whose clients and servers were in the same jurisdiction is analyzed against those that were in distinct jurisdictions. As for this analysis the regulations are not required, just the jurisdictional area, all the TDs in the dataset are considered, instead of reducing the analysis to those whose the regulatory instructions were researched (Table 3.2). The analysis found that in 49% of detected TDs the clients and servers are in the same jurisdiction, while conversely, in 51% of detected TDs they are in distinct jurisdictions. These values demonstrate a clear split between the detected TDs happening in single or multiple jurisdictions. Despite the M-Lab infrastructure having broad global coverage and service that routes tests against the server closer to the client, half of the detected TDs has the client in one jurisdiction and the server in another. On the Internet, where the scale of Content Delivery Networks (CDNs) is much bigger, this proportion shall decrease. However, a representative quantity of traffic traversing multiple jurisdictions still may be found, thus facing multiple NN definitions.

Then, the TDs are analyzed introducing the regulatory instructions presented in Section 3.2. The analysis is restricted to the TDs whose regulations of their jurisdictions were examined (Table 3.2), thus reducing the dataset from 12480 to 10048 TDs ($\approx 81\%$). Four scenarios were evaluated, as follows. As the Glasnost dataset comprises tests of 2016, the regulatory instructions enforced at the time of the tests were considered, which

Figure 3.7 – Agreements and disagreements of regulatory instructions



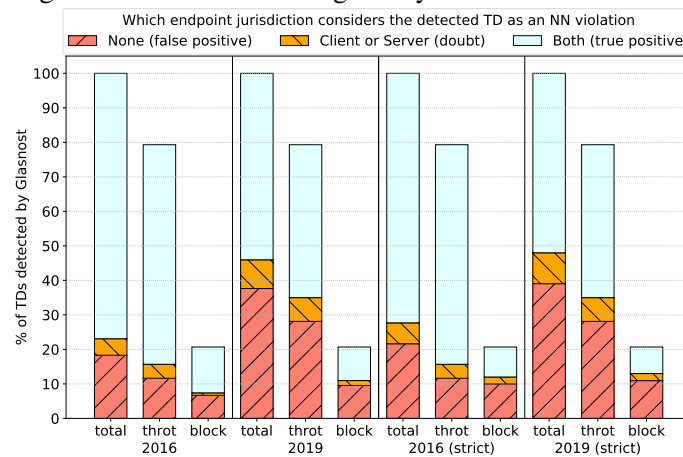
Source: (CARVALHO et al., 2020)

is the most realistic scenario. After this consideration, three hypothetical scenarios are introduced. One scenario is a time shift as if the tests were performed in 2019 (after the USA NN regulation rollback and stricter NN rules in India). The other two hypothetical scenarios apply a stricter interpretation of the regulation (in 2016 and 2019) for those definitions that allow throttling and blocking unlawful content. In this sense, all tests related to P2P (BitTorrent, eMule, and Gnutella) are considered as being related to unlawful content. This interpretation was introduced only to analyze its impact on the results. Thus, it does not reflect the opinion of the author. Finally, these scenarios are named as 2016, 2019, 2016 (strict), and 2019 (strict), respectively.

As different NN definitions may be stated for client and server jurisdictions, how much these definitions agree about a TD is investigated. In other words, whether they consider the TD as a violation or not is irrelevant for this analysis, but whether both reach the same verdict. Only the TDs where divergences may appear were considered, *i.e.*, when clients and servers are in different jurisdictions, which corresponds to $\approx 42\%$ of TDs. It is important to remember that the TDs were restricted to those whose regulations are available in Table 3.2. Thus, this percentage is slightly different from the presented before (49%), considering all detected TDs. The results are presented in Figure 3.7.

In Figure 3.7, a significant number of agreements of different regulations can be seen. The top three jurisdictions drive the vast source of agreements for clients (BR, EU, and the USA) and servers (EU, CO, and the USA) that agree about throttling and blocking practices in 2016. In 2019, a rise in disagreements can be seen due to the NN regulation rollback in the USA. Adding the interpretation of P2P as unlawful content also increases the number of disagreements due to the lack of this exception in Brazil (throttling and blocking) and Colombia (throttling). As P2P is involved in a massive part of detected TDs (65.1%), this interpretation impacts the results, especially in blocking

Figure 3.8 – Influence of regulatory instructions in verdicts



Source: (CARVALHO et al., 2020)

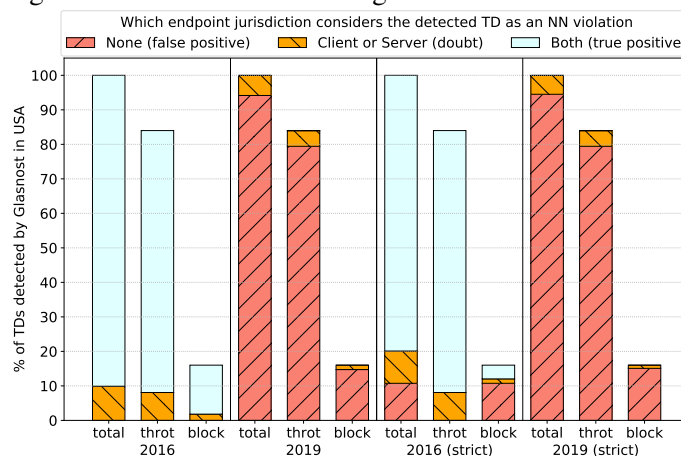
because regulations mostly allow blocking it. As the content is deemed unlawful, it does not make sense to allow throttling (slow down) access to the content.

The amount of disagreements presented in Figure 3.7 is small but representative, exposing regulatory discrepancies and their impact on the results. These results show that one cannot judge whether the TD is a violation just with the information available. More information is needed to decide in which jurisdiction the TD occurred so that one can judge the TD under the proper regulatory instructions.

The influence of the regulations when TDs are judged also was analyzed. Thus, the analysis is concerned with how much of the TDs detected by Glasnost cannot be considered NN violations under the regulatory perspective. The results are presented in Figure 3.8, that shows in which jurisdictions (client or server) the detected TD is considered an NN violation. The legend "None" means that the TD is not considered a violation in client or server jurisdictions; thus, it is a "false positive." They are indeed TDs, but they were interpreted implicitly as NN violations, which is wrong according to the regulation enforced on the jurisdictions of the endpoints. The remaining legends ("Client or Server" and "Both") are self-explanatory.

In 2016, the results show that $\approx 18\%$ of detected TDs were false positives. Additionally, $\approx 5\%$ of TDs are not considered violations on the client or on the server jurisdiction of the TD, thus, may or may not be false positives (doubts). Therefore, introducing the regulatory perspective, the analysis found that from 18% to 23% (adding the 5% that might be) of TDs detected by Glasnost could not be considered NN violations. The proportions are slightly different for throttling and blocking, having more false positives for blocking ($\approx 32\%$).

Figure 3.9 – Influence of NN regulation rollback in the USA



Source: (CARVALHO et al., 2020)

In 2019, a steep rise is seen in the number of false positives ($\approx 37\%$). The TDs that are not considered violations on server or client jurisdictions also increased up to $\approx 8\%$. Thus, after the USA NN regulation rollback, from 37% to 45% of TDs detected by Glasnost could not be considered NN violations. As seen in 2016, the TDs related to blocking also presents more false positives ($\approx 46\%$).

Adding the interpretation of P2P being the avenue for illegal content, an increase is seen in the false positives for 2016 (strict) and 2019 (strict). This fact happens because more situations are introduced where a TD is allowed. In 2016 (strict), the number of false positives is $\approx 21\%$ (compared to $\approx 18\%$ observed in 2016). The situation that the TD is not considered a violation in client or server jurisdiction also increases to $\approx 6\%$ (compared to $\approx 5\%$ observed in 2016). The overall result is that from 21% to 27% of detected TDs could not be considered NN violations when this interpretation is added.

When adding the interpretation above to the 2019 results, an increase in false positives and doubts can also be seen. In 2019 (strict), the number of false positives is $\approx 39\%$ (compared to $\approx 37\%$ observed in 2019). The situation that the TD is not considered a violation in client or server jurisdiction also increases to $\approx 9\%$ (compared to $\approx 8\%$ observed in 2019). The overall result is that from 39% to 48% of detected TDs could not be considered NN violations. It is important to note that the proportion is even higher for blocking, achieving $\approx 52\%$ for false positives and $\approx 10\%$ for doubts, thus, achieving a surprising range of 52% to 62% of false positives related to blocking.

In Figure 3.9, a similar analysis is presented but just considering the USA jurisdiction to quantify how the NN regulation rollback from 2018 influences the results. Its legends follow the same pattern of Figure 3.8. The Figure shows a clear influence by com-

paring the scenarios of 2016 (2016 and 2016 (strict)) and 2019 (2019 and 2019 (strict)). Indeed, the confirmed NN violations (Both) exchanges with false positives (None). The doubts (client or server) decrease because, in a reasonable amount of tests, the clients from other countries were targeting servers in the USA. As the TDs somehow related to USA (client or server) represents $\approx 25\%$ of the TDs presented in Figure 3.8, the representative influence of the USA NN regulation rollback can be noted in the overall results.

3.3.2 Discussion

The research was impacted by the difficulty to find a reliable source of information about the NN definitions stated in each jurisdiction. The original regulations were inspected in some cases, but they are hard to find and be sure that they are the most up-to-date version; additionally, they are usually written in a foreign language. Access to this information should be easier for foreign researchers. It would also help to build systems based on reliable and easy to achieve regulations.

There is a broad spectrum of possible interpretations of detected TDs as NN violations. On one side is the absence of NN regulation in which all TDs are allowed and, therefore, there is no NN violation. On the other side is the presence of a strict NN regulation in which all TDs are prohibited and, therefore, all of them are NN violations. The exceptions introduced in the regulatory instructions provide interpretations of NN between these two sides. In turn, the state-of-the-art solutions for NN violation detection assume a strict regulation that all TDs are classified as NN violations. This analysis explores different NN interpretations, moving across time (after the USA regulation rollback) and adding more exceptional situations (P2P as unlawful content). As the regulation is alive along the time (see regulatory changes in Table 3.2), the solutions designed for NN violation detection also need to consider these changes over time and the exceptions that each regulation may introduce.

It is important to note that the presented results arise from an optimistic analysis of the Glasnost dataset under the regulatory perspective. As the dataset provides information only about the client and server involved in the test, only the jurisdiction of the endpoints can be evaluated. Thus, the assessment of the effects of any hidden jurisdictions along the network path of a test is not possible. While the client and server may be hosted in the same jurisdiction, the network path that connects them during the test may still traverse other jurisdictions with different legal allowances than initially predicted. For instance,

when considering a test with both the client and the server in the USA, the packets of this test may be transiting at some point through a router in Canada or even Mexico. The same may happen for clients and servers that are already in different jurisdictions. For instance, a client in Brazil performing a test against a server in Colombia may be routed through routers in Peru or Bolivia. The introduction of these hidden jurisdictions in the analysis could increase disagreements and false positives, thus impacting the results found.

The current USA NN regulation is also another source of possible contradictory definitions, which may end in more disagreements and false positives. In this analysis, the regulation instructions that are considered were stated by countries or regions that define the jurisdictions. However, after the USA NN regulation rollback, some USA states established their NN regulation due to a lack of Federal regulation. Future work in this context should account for this situation in the USA.

To help to overcome the two points above, the modeling of solutions based on regulations should consider more complex traffic definition scenarios where the TDs take place. The models should represent the whole path between the client and application server that could expose these hidden jurisdictions (even smaller jurisdictions like states). However, the traffic between the client and the application server can follow multiple paths, thus, possibly transversing even more jurisdictions. Even worst, paths are dynamic and may be valid only for a few hours. Therefore, such models should allow the representation of the whole topology that was traversed by the application packets involved in a TD. Beyond this, solutions designed for detection of NN violation should have better accuracy when pointing where the TD is happening (at least to a country or state level) to help to establish the proper jurisdiction.

3.4 Summary

In this chapter, the influence of regulations over the results of NN violation detection systems was quantified. Using NN definitions from the regulation, the results available in the Glasnost dataset about TDs (throttling and blocking) detected in 2016 were analyzed to answer the questions. As these NN definitions are valid just within their jurisdictions, the Geo Localization information of clients and M-Lab servers were used to apply the correct definitions over the TDs detected by Glasnost. The answers to the research questions are as follows.

- (i) How much different regulations agree whether a particular TD is considered

as an NN violation? The analysis found that the regulations of the endpoints agree (is a violation or not) about the verdicts of 91% to 95% of the detected TDs. This convergence is driven by jurisdictions that represent a vast amount of the detected TDs agreeing about blocking and throttling practices. However, the number of disagreements (5% to 9%) is not negligible and indicates situations that TDs cannot be judged as violations with the available information. The pointing out of where the detected TD occurred is needed, at least at a country or state level.

(ii) How many of the NN violations detected by state-of-the-art solutions remain as violations considering the regulation? The analysis found that under certain circumstances, from 39% to 48% of the detected TDs are not NN violations according to regulations. This result confirms the assumption that solutions must consider the regulation. Otherwise, their results have a good chance of being wrong.

(iii) How much influence may the interpretation of regulatory instructions have over the results? In order to answer this question, TDs related to P2P applications were interpreted as being illegal content because a few regulations allow the blocking of such content. The analysis found that from 16% to 22% of the detected blocking practices were no longer considered as NN violations. This finding exposes that solutions for NN that consider the regulation must model interpretations to be aligned to regulations.

(iv) How much influence the changes in the regulatory instructions had along with the time? The analysis focused on the regulation change that happened in the USA in 2018. The former regulation prohibited throttling and blocking but allowed the blocking of illegal content. The current regulation does not establish prohibitions *a priori*. Thus, all TD practices are permitted since they do not violate the antitrust principle. The analysis found that the TD practices considered violations drop from 90% to 0%. We expected this considerable change due to the paradigm shift between the former and current regulations. However, it exposes the need to model changes (especially the smaller ones that may be more probable) in the regulation that may happen and its validity over time.

The findings of this analysis expose that solutions for detection of NN violation must consider regulations, or they need the support of another solution that is aware of them. These findings seem to be obvious, but state-of-the-art solutions do not tackle them. In the next chapter, information and data models are presented to represent the information identified by this analysis. Such models may be used by solutions aware of the regulatory perspective that may be used by state-of-the-art solutions to assess the detected TDs under this perspective.

4 MODELING

In this chapter, the modeling aspects required by the solution to address the problem of pointing out NN violations according to regulations are presented. In Section 4.1, the concepts identified during the research and impact analysis are presented along with the modeling assumptions and the established solution objectives. In Section 4.2, the information model that was designed to support the concepts and solution objectives is presented. In Section 4.3, the investigation of candidate data models to represent the information model is presented along with the rationale that guided the decision for YANG data models. In Section 4.4, the data models based on YANG are presented, which includes the YANG models from IETF and the new ones that were designed in this thesis.

4.1 Concepts, assumptions, and solution objectives

In this section, the concepts, assumptions, and solution objectives that guided the modeling decisions are presented.

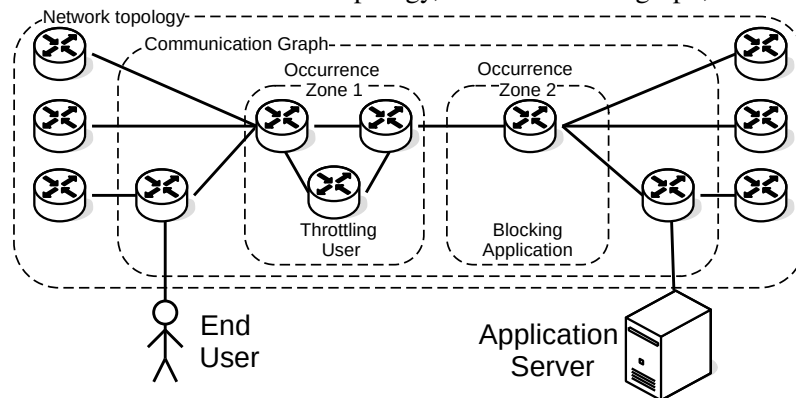
The proposed solution is designed to complement state-of-the-art solutions providing the capability of judging detected TD as NN violation considering the regulation. In such a solution, the state-of-the-art solution provides the information to be represented in the model. The model should be technologically agnostic and easily extendable to accommodate new information because each solution may identify specific information.

Jurisdictions are the geographical areas where legislators or regulatory agencies have the competence to establish regulations. The side effect is that each regulation is valid only on the jurisdiction of its legislator. This fact leads to the situation depicted in Figure 1.1 that exposes that the communication can be under distinct NN regulations along with an end-to-end network path. Thus, modeling must consider geographical information, such as the place where the regulations are established as well as the place where the TDs are deployed, to be able to establish jurisdictions properly.

Usually, links are seen as non-neutral (ZHANG; MARA; ARGYRAKI, 2014). However, the non-neutral behaviors attributed to links are introduced by nodes (*e.g.*, routers, switches), where traffic management mechanisms dictate how the links shall perform. Therefore, the node is where the NN violation may happen and it can be seen as the criminal responsible for violating the NN regulation¹. There are a few possibilities to

¹Of course, these nodes were configured by a network administrator that is the final responsible.

Figure 4.1 – Relations between network topology, communication graph, and occurrence zones



Source: (CARVALHO; GRANVILLE, 2022)

determine the jurisdiction of a case, as presented in Section 2.1. Indeed, *the place where the node is deployed* is one of the factors considered to establish the jurisdictions.

There is a need to represent the regulations within the solution. As seen in Table 3.2, the regulatory instructions generally have two types of regulatory commands: prohibition or allowance. Each command may establish exceptional situations, which also should be represented. As discussed in Section 3.3.2, states are smaller jurisdictions with their regulations, but these regulations are under a broader regulation (*e.g.*, the national one). Therefore, there is a hierarchy between them: state regulations are under national regulations, national regulations are under regional regulations. There are also global regulation efforts (*e.g.*, Digital Constitutionalism) or first-order constitutional principles (*e.g.*, “everything which is not forbidden is allowed,” “everything which is not allowed is forbidden” principles) (HADLEY; ANDENÆS; FAIRGRIEVE, 2000) that are foundations of legal systems. Therefore, this regulation hierarchy should be represented.

As discussed along with Section 3.3, there is the need to represent regulatory *interpretations* within the solution because regulatory commands usually refer to classes of traffic. For instance, a possible regulatory instruction could be “traffic blocking is prohibited, except for unlawful content.” However, there is no information about which traffic should be considered “unlawful content.” Besides, what is considered “unlawful content” in one jurisdiction may not be considered the same in another. Therefore, the solution should represent the interpretations of classes of traffic for each regulation.

There is also the need to represent the topology involved in the communication to map the jurisdictions traversed by the traffic. Therefore, the solution should represent nodes (*e.g.*, routers, switches) and links. Additionally, the solution should be agnostic about the topology structure or the level of information that TD detection solutions may

collect. In this sense, multiple types of connectivity links should be represented: data link level (*e.g.*, Ethernet), network-level (*e.g.*, IP version 4 (IPv4), IP version 6 (IPv6)), or transport/application level connectivity (*e.g.*, VPN tunnels).

A few solutions try to point out where the TD is deployed within the network path (GARRETT; BONA; DUARTE, 2021). These solutions lack accuracy, usually referring to a set of links and nodes where the TD may have been deployed. In order to represent this, the concept of *Occurrence Zone* is defined, as depicted in Figure 4.1. The Occurrence Zone is the set of links and nodes (a connected subset of the communication graph, which is the part of the topology involved in the communication), the differentiation (*e.g.*, throttling, blocking), and the entity that was differentiated (*e.g.*, user, application). As multiple differentiations in distinct parts of the communication graph may be found, the model must represent several Occurrence Zones within one communication graph.

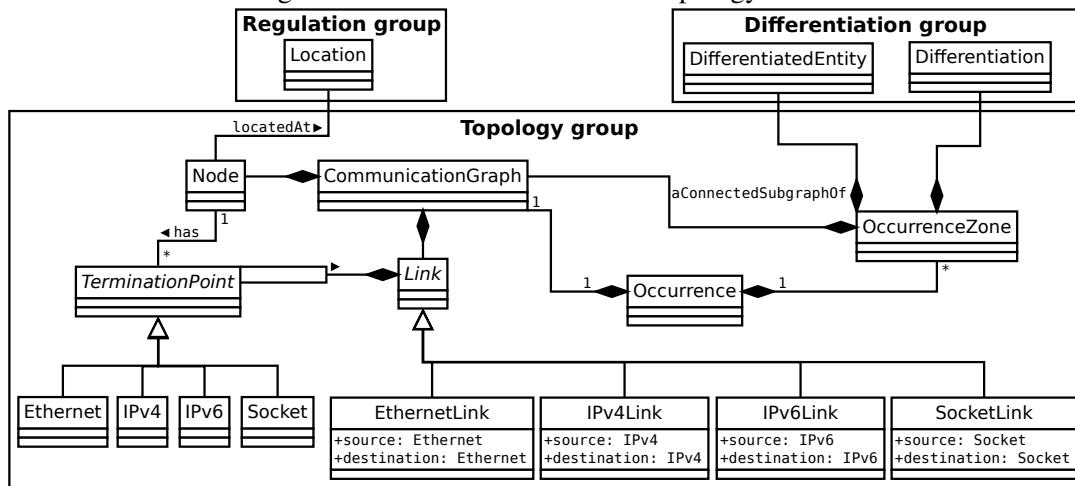
In the next section, the information model derived from the above concepts, assumptions, and solution objectives is presented.

4.2 Information model

In this section, the information model designed to represent the concepts complying with the requirements and objectives presented in the previous section is presented. In order to help the explanation and presentation of the information model, it was divided into three groups: topology, differentiation, and regulation. The information model was represented by Class diagrams from the Unified Modeling Language (UML) (BAUMANN; GRÄSSLE; BAUMANN, 2005) and most of the attributes were omitted here for simplicity. Indeed, a few attributes are presented in the Data Model section (Section 4.4) and all attributes are available in the Appendix B.

In Figure 4.2, the Topology group of the information model is presented. The *Occurrence* class is composed of one *CommunicationGraph* class and many *OccurrenceZone* classes. The *CommunicationGraph* class is responsible for representing the topology involved in the communication. It is composed of *Nodes* and *Links*. A Node may have several *TerminationPoints*, which may be multiple TerminationPoints from the same Open System Interconnection (OSI) layer (*e.g.*, IPv4), such as routers with multiple IP addresses, or may be TerminationPoints of different layers (*e.g.*, Ethernet and IPv4) for devices that connect with different connectivity technologies, such as Cable Modems or routers. Nodes are located at one Location class, which is presented in the Regulation

Figure 4.2 – Information Model - Topology classes



Source: (CARVALHO; GRANVILLE, 2022)

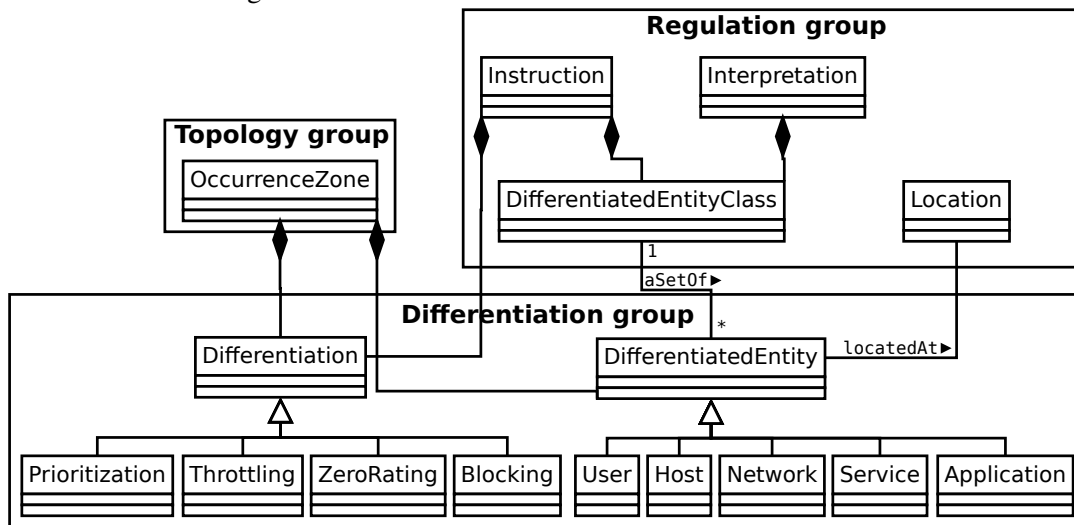
group. A Link is a self-relation of two distinct TerminationPoints of the same type (*e.g.*, IPv4). As the proposed solution is designed to complement the functionality of state-of-the-art solutions, the type of technology that such solutions may identify may vary. Thus, the information model allows being expanded through generalizations of TerminationPoints and their associated Link types.

The *OccurrenceZone* class is composed of a connected sub-graph of the *CommunicationGraph*. This sub-graph represents the granularity and accuracy of the placement of the TD performed by the state-of-the-art solutions that may identify several Nodes and Links or just a single Node as the Occurrence Zone. The *OccurrenceZone* class is also composed of the *DifferentiatedEntity* and *Differentiation* classes presented in the *Differentiation* group, representing who was differentiated and what kind of differentiation was faced. Therefore, the *OccurrenceZone* is a portion of the topology identified by the state-of-the-art solution where a *DifferentiatedEntity* faces *Differentiation*.

In Figure 4.3, the *Differentiation* group of the information model is presented. The *Differentiation* class represents all sorts of TD that may affect communications. Through generalizations, the model allows to include new types of TDs. The *Differentiation* is part (composition) of the *OccurrenceZone* class from the *Topology* group. It is also part of the *Instruction* class, which is presented in the *Regulation* group. The *DifferentiatedEntity* class represents all sorts of entities that may have the communication affected. Through generalizations, the model also allows to include new types of *DifferentiatedEntities*. It is part (composition) of the *OccurrenceZone* class from the *Topology* group.

In Figure 4.4, the *Regulation* group of the information model is presented. The *Location* class represents the geographical areas where jurisdictions may be established.

Figure 4.3 – Information Model - Differentiation classes



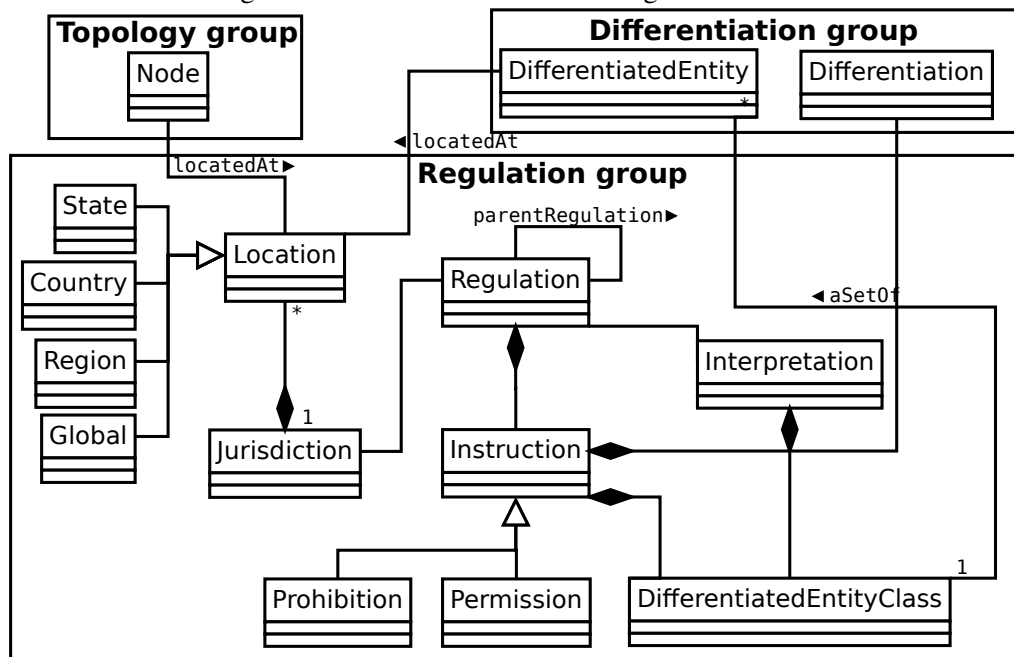
Source: (CARVALHO; GRANVILLE, 2022)

There is an implicit hierarchy between them: Global, Region, Country, and State. Although there are not global scope regulations established nowadays, the Global class can be used to represent approaches such as Digital Constitutionalism (PADOVANI; SANTANIELLO, 2018). However, even these Global regulation efforts may require agreements from participants, which are hard to achieve wide coverage in practice. In turn, the *Global* class is used also to model first-order constitutional principles (e.g., “everything which is not forbidden is allowed,” “everything which is not allowed is forbidden” that are foundations of legal systems. The *Jurisdiction* class is a composition of Location classes.

The *Regulation* class is related to one Jurisdiction class. The Regulation class has a self-relation (parentRegulation) to express the regulation hierarchy: state regulations are related to country regulations, country regulations may be related to regional or global regulations, and global regulations are at the top of the hierarchy (they relate to themselves). The Regulation class is composed of *Instruction* classes that express the individual regulatory instructions that compose the regulation. These Instructions can be of two types: *Prohibition* or *Permission*. They are related to DifferentiatedEntityClass and a Differentiation. For instance, the regulatory instruction may express: “blocking” of “lawful content” is “prohibited,” which are the Differentiation, DifferentiatedEntityClass, and Instruction, respectively. The *Interpretation* class represents which DifferentiatedEntities are interpreted/considered from one DifferentiatedEntityClass for one regulation. For instance, what DifferentiatedEntities are considered “lawful content” in the example.

In the next section, the investigation of existing data models to represent the information model is presented along with the reasons for the choice for YANG.

Figure 4.4 – Information Model - Regulation classes



Source: (CARVALHO; GRANVILLE, 2022)

4.3 Investigation of suitable data models

A comprehensive search was performed to find suitable data models to represent the information model that was designed for the solution to judge TD as NN violation according to NN regulation definitions. Beyond the feasibility to represent the information model, other criteria were also assessed, such as easiness to design the data model, extensibility of the standard, and expected longevity of the standard. Among the surveyed candidates are those based on Network Search (UDDIN; STADLER; CLEMM, 2013) (UDDIN et al., 2014) (UDDIN; STADLER, 2016), Structure of Management Information Version 2 (SMIv2) (MCCLOGHRIE; PERKINS; SCHOENWAELDER, 1999), and Virtual private eXecution infrastructure Description Language (VXDL) (KOSLOVSKI; PRIMET; CHARÃO, 2009). After a deeper investigation, three data models were chosen for further investigation: Common Information Model (CIM), Network Mark-Up Language (NML), and YANG. In Table 4.1, the modeling goals are presented along with information whether these data models can accomplish them. In the next paragraphs, these data models are presented, highlighting their weaknesses to accomplish the goals.

CIM (Distributed Management Task Force, 2022a), maintained by the Distributed Management Task Force (DMTF) (Distributed Management Task Force, 2022b), is a standard based on object orientation to represent all sort of elements of an Information Technology (IT) environment, such as systems, applications, and networks. In terms of

Table 4.1 – Suitable data models analysis

Goal	CIM	NML	YANG
Topology representation	+/-	✓	✓
Several links per node	✓	✓	✓
Links on different OSI layers	✓	+/-	+/-
Unidirectional links	✗	✓	✓
Occurrence Zone representation	+/-	✓	✓
Localities representation	✓	✓	✗
Extensible	✓	+/-	✓

Source: the author

topology representation, CIM has an experimental class named `TopologyGraph` that is a collection of `ConnectivityCollections`, whose members of one collection share connectivity. After that, the topology can be built connecting the nodes that share the same `ProtocolEndpoints` within and across the `ConnectivityCollections`. Therefore, the model stores the data required to build the topology, but it needs to be built interactively by inspecting the `ProtocolEndpoints` of each node. The represented connectivity is bi-directional within the `ConnectivityCollections`, hindering the representation of unidirectional links. The issues of topology representation hinder the Occurrence Zone representation as well.

NML (HAM et al., 2013), maintained by the Open Grid Forum (OGF) (Open Grid Forum, 2020), is a standard based on object orientation to represent network topology. Regarding the representation of links on different OSI layers, the NML model is technology agnostic and lacks attributes to represent the technologies on each layer. In order to overcome this limitation, one possible solution could be the use of name conventions for the Ports and Links to identify the technology involved. However, the solution using the model would have the duty to ensure restrictions related to technologies. For instance, the Ports that are related to one Link must use the same technology. In terms of NML extensibility, which is required to introduce the concepts related to regulations, the participation in the OGF is required to extend the model. However, the last NML standard document is from May 2013, indicating that it is not receiving extensions for a long time.

Finally, YANG (BJORKLUND, 2016), maintained by the IETF (Internet Engineering Task Force, 2022), is a modeling language that is the standard to represent information for the Network Configuration Protocol (NETCONF) (ENNS et al., 2011). Several data models using YANG, which are also kept as IETF standards, are proposed to represent information for specific purposes, such as network topology representation (CLEMM et al., 2018b). The analysis in Table 4.1 considered the available data models that are standards and drafts from the IETF. For the representation of links on different OSI layers,

Table 4.2 – Existing YANG data models

Model	Reference	Goal
ietf-network-topology	RFC 8345	To represent general purpose network topologies. It augments the ietf-network module proposed on a previous draft.
ietf-l3-unicast-topology	RFC 8346	To augment the ietf-network and ietf-network-topology modules to represent layer 3 related information (<i>e.g.</i> , IP addresses).

Source: the author

there is only a data model that represents Layer-3 connectivity (CLEMM et al., 2018a). In order to represent other kinds of connectivity, a new data model augmenting the existing ones should be designed. For the locality representation, there is no data model available that represents localities. Therefore, a new data model or an extension to existing ones also should be designed. The YANG data models have the advantage of being easily extended through augmentations (over existing data models) or through the proposal of new data models, which could also be part of a contribution to IETF. Given that such extensions can solve the pointed limitations, YANG was chosen as the data model to represent information of the proposed solution. In the next section, new data models based on YANG are presented along with the augmentations to existing YANG models that were used to represent information within the proposed solution.

4.4 Data models based on YANG

In this section, the models proposed in this thesis and augmentations to the existing YANG models are presented. The proposed models represent the information required to point out NN violation according to regulations, which includes information about NN regulations and TD occurrences. In this section, only excerpts of the proposed models are detailed. However, the complete models are available in Appendix B.

In Table 4.2, the existing YANG models used as the base for the proposed models are listed. The ietf-network-topology (CLEMM et al., 2018b) is an augmentation to the ietf-network module that only represents networks, network-refs, nodes, and node-refs. The ietf-network-topology model augmentation adds elements such as termination points and links that make possible to represent technological agnostic network topologies. Therefore, the elements of these two modules represent the topology structure without technical details. The technical information shall be introduced by complementary

Table 4.3 – Proposed new YANG data models and augmentations to the existing ones

Model	Goal
nn-regulation	NN regulation representation
tdo	TD occurrence zones representation

Source: the author

models that augment these two base models. The *ietf-l3-unicast-topology* (CLEMM et al., 2018a) is such a complementary model that augments the base models to introduce technical information related to layer 3 connectivity (*e.g.*, IP addresses).

In Table 4.3, the proposed YANG models and the proposed augmentations are listed. The *nn-regulation* model is designed to represent the regulations and related information, such as the jurisdictions, the instructions, the interpretations, the traffic differentiations, and the entities who suffer the differentiations. This model is proposed in this thesis, and it does not rely on information from existing models. The *tdo* model is designed to represent the traffic differentiation occurrence zones and related information, such as the network topology, the communication graph, and the elements (nodes and links) where the differentiation is taking place. This model is proposed in this thesis and it relies on information from *ietf-network*, *ietf-network-topology*, and *ietf-l3-unicast-topology* models. The *nn-regulation* and the *tdo* models are detailed next.

In Figure 4.5, an excerpt of the proposed *nn-regulation* model is presented. The excerpt presents the list of regulations (lines 37-39) and the information of each regulation (lines 1-34). Each regulation points to its parent through the *parent-regulation-ref* attribute (lines 11-18). The goal of this attribute is twofold: (i) to represent the intrinsic hierarchy of regulations regarding to jurisdictions, representing the relationship between state- and national-scope regulations, and national- and regional-scope regulations; and (ii) to represent the relationship between regulations and first-order constitutional principles (HADLEY; ANDENÆS; FAIRGRIEVE, 2000) used by countries to base their legal system. The regulation validity is represented in the *begin-validity* and *end-validity* attributes (lines 19-29). The *jurisdiction-list* (line 30) represents the International Organization for Standardization (ISO) codes of the states and countries where the regulation is valid. The *instruction-list* (line 31) represents the instructions of the regulation, which is the command (prohibition or permission) of one differentiation (*e.g.*, blocking, prioritization, degradation) over a differentiated entity class (*e.g.*, user, network, application).

The *interpretation-list* (line 32) represents information used to relate differentiated entities to their respective differentiated entity classes. This interpretation may vary ac-

Figure 4.5 – Excerpt from the proposed nn-regulation YANG model

```

1      grouping regulation-list {
2          description "The list of regulations.";
3          list regulation {
4              key "regulation-id";
5              leaf regulation-id {
6                  type regulation-id;
7              }
8              leaf name {
9                  type string;
10             }
11             leaf parent-regulation-ref {
12                 type leafref {
13                     path "/nnr:regulations/nnr:regulation/nnr:regulation-id";
14                 }
15                 description
16                 "Reference to the parent regulation or itself when root regulation.";
17                 mandatory true;
18             }
19             leaf begin-validity {
20                 type yang:date-and-time;
21                 description
22                 "The begin of validity of the regulation.";
23                 mandatory true;
24             }
25             leaf end-validity {
26                 type yang:date-and-time;
27                 description
28                 "The end of validity of the regulation.";
29             }
30             uses jurisdiction-list;
31             uses instruction-list;
32             uses interpretation-list;
33         }
34     }
35
36     // data definition statements
37     container regulations {
38         uses regulation-list;
39     }

```

Source: the author

ording to the regulation. For instance, one regulation may interpret P2P applications as unlawful content while others may not interpret them in this way. The list is composed of abstract differentiated-entity objects that have attributes used to classify the differentiated entities. One mandatory attribute is the “regulation-class” that indicates the differentiated-entity-class, while other specified attributes are used in the classification process. For instance, in China the blocking and censorship of adult pages is allowed (DARER; FAR-NAN; WRIGHT, 2018). In such a case, the interpretation list may have several Application differentiated-entity objects with the attribute regulation-class as “adult-page” and each object with an attribute “url” with the web page Uniform Resource Locator (URL). Then, the Chinese regulation can have one instruction with “permission” of “blocking/censorship” the “Application of regulation-class adult-page”, which are the instruction-type, differentiation, and differentiated-entity of the regulation-class, respectively.

The topology information (depicted in Figure 4.1) is represented in the tdo model. The Network topology is represented by the ietf-network, ietf-network-topology, and ietf-l3-unicast-topology models. The communication graph is the part of the Network topology traversed by the communication (*e.g.*, application traffic). Therefore, the communication graph is built referencing the nodes and links of the Network topology. The

Figure 4.6 – Excerpt from the proposed tdo YANG model

```

1 grouping occurrence-zone-list {
2   description
3     "The zones within the topology where Traffic Differentiations are happening (occurrence)";
4
5   list occurrence-zone {
6     key "occurrence-zone-id";
7     leaf occurrence-zone-id {
8       type "occurrence-zone-id";
9       description
10        "The id of occurrence zone.";
11    }
12    leaf differentiation {
13      type "nrr:differentiation";
14      description
15        "The Traffic Differentiation that is happening within the occurrence zone.";
16    }
17
18    uses "nrr:differentiated-entity";
19
20    list nodes {
21      key "node-ref";
22      leaf node-ref {
23        type leafref {
24          path "/tdo:occurrences/tdo:occurrence[tdo:occurrence-id=current()/../../../../tdo:occurrence-id]/tdo:communication-graph/tdo:nodes/tdo:node-ref";
25        }
26        description
27          "References the nodes within the communication graph that are in the occurrence zone
28          .";
29      }
30    }
31    list links {
32      key "link-ref";
33      leaf link-ref {
34        type leafref {
35          path "/tdo:occurrences/tdo:occurrence[tdo:occurrence-id=current()/../../../../tdo:occurrence-id]/tdo:communication-graph/tdo:links/tdo:link-ref";
36        }
37        description
38          "References the links within the communication graph that are in the occurrence zone
39          .";
40      }
41    }
42    description
43      "The occurrence zone definition.";
44  }

```

Source: the author

occurrence zone where the TD is taking place is defined by the nodes and links that are part of the communication graph. Therefore, these elements are modeled as references to the nodes and links of the communication graph.

One occurrence is composed of the communication graph and a list of occurrence zones. In Figure 4.6, one excerpt of the tdo model is presented. The excerpt presents the list of occurrence zones (lines 1-43) and the information of each occurrence zone (lines 5-42). Each occurrence zone is about one traffic differentiation (lines 12-16) over one differentiated entity (line 18) that is happening on one part of the communication graph that is defined as lists of node references (lines 20-29) and link references (lines 30-39) where the TD is taking place. These node and link references point to the communication graph as defined by the path attribute (lines 24 and 34).

The tdo model also augments the ietf-network model adding the attribute "located-at" (ISO code of the country or state) to nodes. Further details about network topology, communication graphs, and the augmentation to the node element are not present in the

Figure 4.7 – Excerpt from the proposed tdo YANG model - Remote Procedure Call (RPC) statement definition

```

1 // rpc statements
2 rpc analyze {
3   description
4     "Analyze whether the traffic differentiation occurrence is an NN violation according to the regulation.";
5   input {
6     leaf occurrence-id {
7       type "occurrence-id";
8       mandatory true;
9       description
10        "The id of the occurrence.";
11    }
12    leaf occurrence-zone-id {
13      type "occurrence-zone-id";
14      mandatory true;
15      description
16        "The id of occurrence zone.";
17    }
18  }
19  output {
20    container results {
21      list result {
22        key "id";
23        leaf id {
24          type uint32;
25          mandatory true;
26        }
27        leaf name {
28          type string;
29        }
30        container admissibility {
31          leaf rationale {
32            type string;
33          }
34          leaf admissible {
35            type boolean;
36          }
37        }
38        container judgment {
39          leaf metric-description {
40            type string;
41          }
42          leaf metric-value {
43            type decimal64 {
44              fraction-digits 2;
45            }
46          }
47          leaf rationale {
48            type string;
49          }
50          leaf is-violation{
51            type boolean;
52          }
53        }
54      }
55    }
56  }
57 }

```

Source: the author

excerpt of Figure 4.6 but are available in Appendix B.

In Figure 4.7, an excerpt of the tdo model is presented that has the definition of the RPC method named analyze. The method analyze receives the occurrence-id and the occurrence-zone-id as inputs (lines 6-11 and lines 12-17, respectively) and executes all Judgment Algorithms (detailed in the next chapter) configured in the system. The output is a container named results (line 20-55) composed of a list with information about each Judgment Algorithm. The algorithm name is represented in the name attribute (lines 27-29). The container admissibility (lines 30-37) represents attributes about the admissibility step. The rationale of the admissibility step result is represented in the rationale attribute

(lines 31-33), whose result whether the occurrence zone is admitted for being judged by the algorithm is represented in the admissible attribute (lines 34-36). If the occurrence zone was admitted to be judged by the algorithm, the judgment step result is represented by in the judgment container (lines 38-53). The judgment step result is expressed by a metric, whose description is represented in the metric-description attribute (lines 39-41) and its value is represented in the metric-value attribute (lines 42-46). The judgment rationale is represented in the rationale attribute (lines 47-49) and provides information about how the judgment result was computed. The is-violation attribute (lines 50-52) presents the verdict whether the TD is considered an NN violation by the Judgment Algorithm.

In this section, a few excerpts from the proposed YANG models were presented. However, the complete models are available in Appendix B.

4.5 Summary

In this chapter, the requirements that guided the design of the proposed information and data models were presented, such as the need to represent network topologies, communication graphs, occurrence zones, and regulatory information. The designed information models were presented as Class Diagrams along with a brief explanation of the classes and the relationship between them. The information model established the requirements for the investigation of the suitable data models to represent it. A comprehensive search for candidate data models was conducted. This search selected three models for a deeper analysis of their capabilities: CIM, NML, and YANG languages. The YANG modeling language and the available IETF YANG-based models (as standards or drafts) were chosen to base the data models for the proposed solution. A new YANG data model named tdo was designed to represent the network topology, the communication graph, and the occurrence zones. This model uses existing IETF YANG models to represent topology information (ietf-network, ietf-network-topology, and the ietf-l3-unicast-topology). A new YANG data model named nn-regulation was designed to represent the regulatory information, such as instructions, interpretations, differentiations, and differentiated entities. The nn-regulation model augments the ietf-network model to add the located-at attribute to nodes that is used to establish their jurisdiction.

In the next chapter, the proposed solution for pointing out NN violation according to the regulation is presented. This solution relies on the information represented in the data models presented in this chapter.

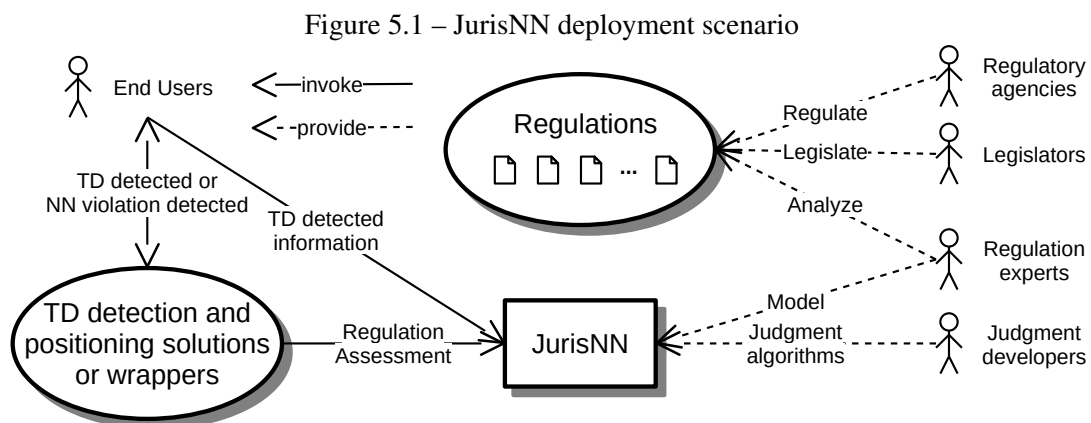
5 PROPOSAL

In this chapter, the proposed solution named JurisNN for the detection of NN violation according to regulations is presented. In Section 5.1, JurisNN is presented focusing on aspects of how it is intended to be deployed, including information about the actors and systems that interact with the solution. In Section 5.2, the JurisNN conceptual architecture is presented, which is modeled as a framework due to its possible interactions with other systems. In Section 5.3, a few Judgment Algorithms are presented, which are one of the components of the JurisNN architecture. These Judgment Algorithms use distinct criteria to judge whether a TD occurrence is an NN violation.

5.1 Deployment scenario

In this section, the JurisNN deployment scenario is presented. This scenario is depicted in Figure 5.1, presenting the actors and systems that interact with JurisNN. These interactions influence the JurisNN architecture design, which is presented in Section 5.2.

JurisNN aims to provide to End Users the verdict whether the TD they are facing is considered an NN violation according to the established regulation. Therefore, JurisNN may be used by anyone interested in assess her/his network regarding the NN principle. By the way, the *End Users* may interact with JurisNN indirectly or directly. The indirect way is through the *TD detection and positioning solution* used to assess their network. These solutions may request the assessment of regulation to JurisNN on behalf of End Users, providing the required information about the TD. However, such a feature needs to be implemented into these solutions. In order not to impose changes in such solutions,



Source: the author

wrappers could be developed. These wrappers can execute the TD detection and positioning solution, submit the TD information to JurisNN, and then return the result of the regulation assessment to the End User. The direct way is through the interaction of the End User with the JurisNN by submitting the information collected by the TD detection and positioning solutions to perform the regulation assessment. Considering these possibilities, the wrapper approach seems to be the most suitable. However, it requires that TD detection and positioning solutions expose the information collected about the detected TD in a way that can be processed by the wrappers.

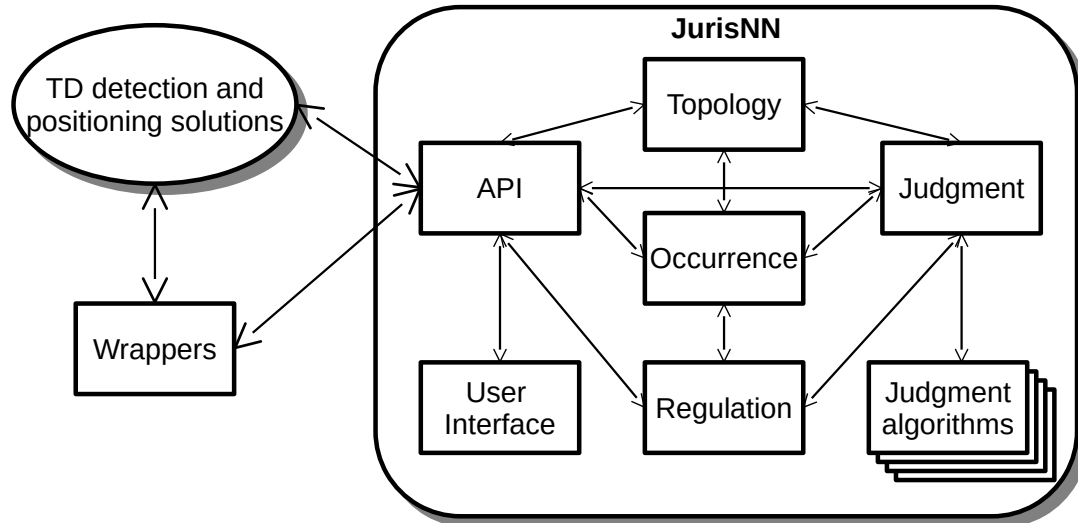
The *Regulatory agencies* and *Legislators* provide the regulations of their jurisdictions that specify the circumstances in which a detected TD can be classified as an NN violation. Although in Figure 5.1 the *Regulations* are presented in an organized set, it is not the truth in the wild. These regulations are spread across each jurisdiction's law or regulatory system, as pointed in Subsection 3.3.2. In this deployment scenario, the *Regulation experts* have the responsibility of identifying the right regulation artifacts, interpret them, and model their regulatory instructions into the JurisNN. They also are responsible for introducing into JurisNN the interpretations of each regulatory instruction within the jurisdiction. They need to provide the information required to classify the applications into their regulatory classes following these interpretations. For instance, when the regulatory instruction refers to a regulation class named "unlawful content," they need to enumerate which specific traffic is considered "unlawful content."

The information about the TDs submitted to JurisNN needs to be processed to judge whether it is an NN violation, which is made through *Judgment algorithms*, as presented in the next section and detailed in Section 5.3. These algorithms are provided by *Judgment developers* that need to know how the TD and regulatory information are represented into JurisNN. A few examples of such algorithms are presented in Section 5.3. In the next section, the JurisNN architecture is presented.

5.2 JurisNN architecture

In this section, the JurisNN architecture is presented. The architecture is depicted in Figure 5.2. The architecture design reflects the observations made in the Discussion subsection (Subsection 3.3.2) of the Impact analysis chapter and the objectives presented in the Modeling goals section (Section 4.1) of the Modeling chapter. Therefore, the module descriptions presented next are brief because they already have been mentioned.

Figure 5.2 – JurisNN architecture



Source: (CARVALHO; GRANVILLE, 2022)

The Application Programming Interface (API) module is responsible for all external interactions with JurisNN, including the interaction with TD detection and positioning solutions and the wrappers. The same API is used by the User Interface module that external actors may use to interact with the system, as depicted in Figure 5.1. Therefore, the API module needs to interact with all system models to expose their functionalities.

The Topology module is responsible for keeping the information about nodes, links, and their relationships. The Occurrence module is responsible for the information about the TD, which includes the kind of detected TD, the entity that is suffering the TD, the communication graph referencing nodes and links from the Topology, and occurrence zones referencing nodes and links from the communication graph. The Regulation module is responsible for the information about regulations registered in the system, including their jurisdictions, instructions, and interpretations.

The Judgment module is responsible for the regulation assessment processing using information from the Topology, Occurrence, and Regulation modules. The regulation assessment depends on definitions about how jurisdictions are established when judging a case, as presented in Section 2.1. When the TD positioning is inaccurate, the occurrence zone may span multiple jurisdictions without pointing to the exact one where the TD is deployed. In such a case, the system may not point accurately whether the TD may be considered an NN violation, as discussed in Subsection 3.3.2. Therefore, there is a level of uncertainty in the regulation assessment processing. In order to accommodate these several possibilities, the regulation assessment is delegated to Judgment Algorithms that are detailed in the next section.

5.3 Judgment algorithms

In this section, a few examples of Judgment Algorithms are detailed. These algorithms are responsible for accommodating the subjectivity and uncertainty related to the judgment of TD as NN violation according to the NN regulations. One source of subjectivity and uncertainty is the jurisdiction establishment, as presented in Subsection 2.1. The jurisdiction may be established by observing several factors, such as where the injury of a TD has incurred or where the harmful action took place. Thus, the suitability of each algorithm concerning jurisdictional issues needs to be assessed before applying its judgment. Therefore, the first step of one Judgment Algorithm is assessing the TD occurrence to decide whether it is suitable to be judged by the algorithm, which is named *admissibility control*. Another source of subjectivity and uncertainty is the judgment itself, which may be affected by several factors, such as the inaccuracy of the TD positioning or interpretations of the regulatory instructions. The assessment of such factors is the central part of the Judgment Algorithm, whose step is named the *judgment*.

It is important to note that jurisdictional issues affect both admissibility control and judgment steps. Therefore, the concerns about jurisdictions span the whole Judgment Algorithm. It is also important to point out that distinct Judgment Algorithms may adopt the same jurisdiction establishment criterion, requiring similar processing. Next, the functions that perform the processing of the criteria identified in Section 2.1 (where the injury of a TD has incurred or where the harmful action took place) are presented, which are used by Judgment Algorithms that adopt these criteria. Indeed, in Subsection 5.3.2, the developed Judgment Algorithms (Place of the TD deployment, Targeting test, and German test) are presented, indicating the functions used in their steps.

5.3.1 Jurisdiction establishment criteria and auxiliary functions

The Judgment Algorithms are divided into two steps: admissibility control and judgment. These steps are very dependent on the criterion used to establish jurisdictions, such as where the injury of a TD has incurred or where the harmful action took place. However, different Judgment Algorithms may adopt the same criterion, requiring similar processing to judge TD as NN violation. Next, the functions that implement these criteria are presented. Judgment Algorithms use these functions to perform the admissibility control and judgment steps following the above criteria.

Procedure 1 EffectsFeltWithinLocalityCountry

Input: Occurrence, OccurrenceZone, Locality

```

1: countries ← ∅
2: differentiated_entity ← getDifferentiatedEntity(Occurrence, OccurrenceZone)
3: if differentiated_entity.type ∈ {User, Host} then
4:   countries ← countries ∪ {getCountry(differentiated_entity.located_at)}
5: end if
6: if differentiated_entity.type ∈ {Network, Service, Application} then
7:   communication_graph ← getCommunicationGraph(Occurrence)
8:   source ← getSource(communication_graph)
9:   destination ← getDestination(communication_graph)
10:  countries ← countries ∪ {getCountry(source.located_at)}
11:  countries ← countries ∪ {getCountry(destination.located_at)}
12: end if
13: if getCountry(Locality) ∈ countries then
14:   return true
15: else
16:   return false
17: end if

```

The admissibility control step needs to evaluate whether the jurisdiction of a TD can be attributed to a specific locality according to the jurisdiction establishment criteria (the place of the injury or the place of the harmful action). By the place of the injury criterion, the place where the effects of the TD are felt is considered to establish the jurisdiction, as presented in the EffectsFeltWithinLocalityCountry function (Procedure 1). The place where the effects of a TD are felt depends on the differentiated entity. When users and hosts suffer differentiation, the effects of the TD are felt by themselves. Therefore, only their locations are considered (lines 3-5). However, when networks, services, or applications suffer differentiation, the effects of the TD are felt by the endpoints of the communication. For instance, when an application is differentiated, its provider and its users are affected. Therefore, the location of the communication endpoints (source and destination) is considered (lines 6-12). As the jurisdictional rules are established at the country level (*e.g.*, the Targeting test is adopted in the USA, and the German test is adopted in DE), the locations are considered at the country level, even when the locality is a state. If the country of the Locality is in the set of countries (line 13), then the effects of the TD are felt within that Location (line 14). In turn, by the place of the harmful action criterion, the place of the nodes (where they are deployed) that are present in the Occurrence Zone is considered. Such processing is done by the HarmfulAction-WithinLocalityCountry function, which was omitted for the sake of simplicity. Indeed, its processing is very similar to Procedure 1 but building a set of countries where the nodes within the occurrence zone are deployed and testing whether the Locality country is within this set of countries.

Procedure 2 IsViolationByRegulation

Input: TD, regulation

- 1: $is_prohibited \leftarrow false$; $is_permitted \leftarrow false$; $is_defined \leftarrow false$;
- 2: $instructions \leftarrow getInstructions(regulation)$
- 3: $differentiated_entity_classes \leftarrow getEntityClasses(TD, regulation)$
- 4: **for all** instruction \in instructions **do**
- 5: **if** instruction.regulation_class \in differentiated_entity_classes **then**
- 6: **if** instruction.type = Prohibition **then**
- 7: $is_prohibited \leftarrow true$; $is_defined \leftarrow true$;
- 8: **end if**
- 9: **if** instruction.type = Permission **then**
- 10: $is_permitted \leftarrow true$; $is_defined \leftarrow true$;
- 11: **end if**
- 12: **end if**
- 13: **end for**
- 14: $is_violation \leftarrow is_prohibited$ **and not** $is_permitted$
- 15: **return** $is_violation, is_defined$

The judgment step needs to evaluate whether the TD is considered an NN violation according to the regulation, which is helped by two auxiliary functions presented next. The first auxiliary function is IsViolationByRegulation (Procedure 2), which assesses whether the TD is considered an NN violation considering what is established in one specific regulation. This assessment considers the interpretations stated for the regulation. For instance, one regulation may point that “blocking” is prohibited, but the “blocking” of “unlawful content” is allowed, and therefore it is not considered a violation. However, the regulation does not enumerate what is considered “unlawful content,” requiring interpretations to clarify when one regulatory instruction applies to a TD by inspecting the DifferentiatedEntity. The $getEntityClasses(TD, regulation)$ (line 3) returns the regulation classes that the differentiated entity pertains, considering the regulation interpretations. This classification uses attributes of the differentiated entities (*e.g.*, users, applications) that indicate that they are a member of a regulation class. For instance, if the regulation considers P2P applications as “unlawful content,” one application could have an attribute “type = P2P” to indicate this. For all regulatory instructions (lines 4-13), the instruction is checked whether it is related to the classes that the differentiated entity pertains to (line 5). If the instruction is about one of the regulation classes of the differentiated entity, the boolean variables are adjusted to specify if the instruction establishes prohibition (lines 6-8) or permission (lines 9-11). If the TD is prohibited and there is no permission established, it is considered an NN violation by the regulation (line 14).

The second auxiliary function is named IsViolationInLocality (Procedure 3) that assesses whether the TD is considered an NN violation considering all regulations that one locality may be subjected to. The evaluation looks for the regulations that may be

Procedure 3 IsViolationInLocality

Input: TD, Locality

```

1: regulation_levels ← [state, country, region, global]
2: regulation_hierarchy ← getRegulationHierarchy(Locality)
3: is_violation ← false
4: for all regulation_level ∈ regulation_levels do
5:   if regulation_hierarchy[regulation_level] is defined then
6:     regulation ← getRegulation(regulation_hierarchy[regulation_level])
7:     is_violation, is_defined ← IsViolationByRegulation(TD, regulation)
8:     if is_defined then
9:       return is_violation
10:    else
11:      continue
12:    end if
13:  end if
14: end for
15: return is_violation

```

applied for that locality following the bottom-up approach from state, national, regional, and global level regulations (line 1 order). The `getRegulationHierarchy(Locality)` (line 2) returns an array with the established regulation for each level. For each level, the `IsViolationByRegulation` function (Procedure 2) is invoked (line 7) to check whether the TD is considered a violation for that regulation. The processing stops when one regulation specifies whether the TD is an NN violation (`is_defined = true`, line 8). Therefore, the regulations from the lower levels (smaller scope) are preferred, which is usual in legal systems. If no regulation gives a verdict whether the TD is an NN violation, it is not considered an NN violation (line 15). These auxiliary functions are used by the functions that implement the judgment step following each jurisdiction establishment criterion.

By the place of the injury criterion, the function for the judgment step is named `IsViolationByPlaceOfInjury` (Procedure 4). Its structure is similar to the function of the admissibility control for the same criterion (Procedure 1) because it also varies according to the differentiated entity. When users or hosts suffer differentiation, their locations are used in the assessment (lines 3-5). When networks, services, or applications suffer differentiation, the location of the communication endpoints is used instead. However, their full locations are used in the judgment process, *i.e.*, the locality is not restricted to the country level, as performed in the admissibility control step. For instance, if the location is a state, this state may have its NN regulation that must be considered in the assessment. The source and destination country are rechecked (lines 11 and 14) because the TD may have been admitted because only one endpoint is within the country that adopts the criterion, but the criterion only applies to the endpoint within that country. It is important to note the use of the auxiliary function `IsViolationInLocality` (Procedure 3)

Procedure 4 IsViolationByPlaceOfInjury

```

Input: Occurrence, OccurrenceZone, Locality
1: differentiated_entity ← getDifferentiatedEntity(Occurrence, OccurrenceZone)
2: TD ← getTD(Occurrence, OccurrenceZone)
3: if differentiated_entity.type ∈ {User, Host} then
4:   is_violation ← IsViolationInLocality(TD, differentiated_entity.located_at)
5: end if
6: if differentiated_entity.type ∈ {Network, Service, Application} then
7:   communication_graph ← getCommunicationGraph(Occurrence)
8:   source ← getSource(communication_graph)
9:   destination ← getDestination(communication_graph)
10:  is_violation_source ← false; is_violation_destination ← false
11:  if getCountry(source.located_at) == getCountry(Locality) then
12:    is_violation_source ← IsViolationInLocality(TD, source.located_at)
13:  end if
14:  if getCountry(destination.located_at) == getCountry(Locality) then
15:    is_violation_destination ← IsViolationInLocality(TD, destination.located_at)
16:  end if
17:  is_violation ← is_violation_source or is_violation_destination
18: end if
19: return is_violation

```

in lines 4, 12, and 15. It is important to note also that this function returns a boolean that indicates the TD is considered an NN violation or not. Indeed, the admissibility control and judgment steps for the place of injury criterion (Procedures 1 and 4, respectively) use the locations of the endpoints or the differentiated entities, which are known and do not impose uncertainty on the judgment process. Therefore, the result whether the TD is an NN violation by this criterion is conclusive by design.

By the place of the harmful action criterion, the judgment process needs to evaluate the location of nodes within the occurrence zone. In turn, the process needs to accommodate situations where the occurrence zone spans multiple localities with distinct regulatory instructions for the same TD. For instance, the TD is considered an NN violation in one locality and it may not be considered a violation in another. The function named *IsViolationByPlaceOfHarmfulAction* (Procedure 5) computes the result of the judgment step for this criterion. For each locality within the occurrence zone, it counts how many nodes are located there (*getLocationCounts()*, line 2). For each location, it invokes the auxiliary function *IsViolationInLocality* (Procedure 3) to assess whether the TD is considered a violation (line 5). As one locality may consider the TD an NN violation while others do not, the judgment process needs to point out the level of similarity in the verdicts within the occurrence zone. It uses an adaptation of the Jaccard similarity index (IVCHENKO; HONOV, 1998) to compare the number of nodes that have the same verdict against all nodes within the occurrence zone (line 11). The MAX is used because the best similarity

index between the two possible verdicts (the TD is an NN violation or not) represents the similarity of the occurrence zone. When the occurrence zone similarity index is 1.0, the verdict is conclusive because it is the same for all nodes within the occurrence zone. Otherwise, for a few nodes within the occurrence zone, the regulations may disagree about the verdict and, therefore, the TD cannot be judged for sure as an NN violation or not.

Procedure 5 IsViolationByPlaceOfHarmfulAction

Input: Occurrence, OccurrenceZone

- 1: is_violation_count \leftarrow 0; is_not_violation_count \leftarrow 0
- 2: location_counts \leftarrow getLocationCounts(Occurrence, OccurrenceZone)
- 3: TD \leftarrow getTD(Occurrence, OccurrenceZone)
- 4: **for all** locality \in location_counts **do**
- 5: **if** isViolationInLocality(TD, locality) **then**
- 6: is_violation_count += location_counts{locality}
- 7: **else**
- 8: is_not_violation_count += location_counts{locality}
- 9: **end if**
- 10: **end for**
- 11: OZ_similarity \leftarrow $\frac{MAX(is_violation_count, is_not_violation_count)}{is_violation_count + is_not_violation_count}$
- 12: **if** is_violation_count \geq is_not_violation_count **then**
- 13: is_violation \leftarrow **true**
- 14: **else**
- 15: is_violation \leftarrow **false**
- 16: **end if**
- 17: **return** OZ_similarity, is_violation

These functions implement the two criteria presented in Section 2.1 for the two steps performed by Judgment Algorithms. Next, the functions used by each Judgment Algorithm to perform each step are listed.

5.3.2 Examples of Judgment Algorithms

The functions presented previously are used to implement the admissibility control and judgment steps of Judgment Algorithms. Each algorithm adopts distinct criteria to establish jurisdiction and to decide whether the TD is an NN violation. In Table 5.1, the functions used by each Judgment Algorithm (Place of TD deployment, Targeting test, and German test) in each step (admissibility control and judgment) are listed.

The Place of TD deployment: the straightforward way of judging TD as NN violation considering the regulation is using the place where the TD is deployed to establish the jurisdiction, *i.e.*, the place where the harmful action took place. As almost all courts around the world accept this criterion, the admissibility control step does not need to check any condition about the TD occurrence. If one jurisdiction does not accept this cri-

Table 5.1 – Auxiliary functions used by Judgment Algorithms

Auxiliary function	Judgment algorithms		
	Place of TD deployment	Targeting test	German test
Admissibility control step:			
EffectsFeltWithinLocalityCountry		✓	✓
HarmfulActionWithinLocalityCountry	✓		✓
Judgment step:			
IsViolationByPlaceOfInjury		✓	✓
IsViolationByPlaceOfHarmfulAction	✓		✓

terion, the admissibility control could be amended to exclude TD occurrences that span such jurisdiction using the HarmfulActionWithinLocalityCountry function. The judgment process uses the IsViolationByPlaceOfHarmfulAction function (Procedure 5). As there may be a level of uncertainty on the results of IsViolationByPlaceOfHarmfulAction function due to the occurrence zone characteristics, the Judgment Algorithm needs to evaluate the similarity index of the occurrence zone before pointing out that the TD is an NN violation or not. If the similarity index is 1.0, the algorithm can consider the verdict as conclusive. Otherwise, the verdict is inconclusive.

The Targeting test: the USA courts adopt tests to decide whether they have the competence to judge a claim based on where the effects of action occur. For instance, if one organization located at a jurisdiction A targets its business to a jurisdiction B, the organization acts that affected clients/users in jurisdiction B can be claimed that jurisdiction. The Judgment Algorithm adopting the Targeting test is detailed, which is one of the tests based on the effects of an action. Since this test applies to the USA, the admissibility control uses the EffectsFeltWithinLocalityCountry function to evaluate whether the effects of the TD are felt within the USA. The judgment step uses the IsViolationByPlaceOfInjury function to achieve the verdict about the TD. As this criterion is based on the locations of the endpoints or differentiated entities, which are known and do not impose any uncertainty, the results are conclusive by design.

The German test: in Germany, the jurisdiction may be established by two criteria: where the harmful action happened and where the injury has incurred. Therefore, Germany adopts, at the same time, the criteria adopted by the above two Judgment Algorithms. The admissibility control is the junction of EffectsFeltWithinLocalityCountry function (Procedure 1) to evaluate the place of injury of the TD criterion and HarmfulActionWithinLocalityCountry function to evaluate the place of the harmful action criterion. In the end, if Germany is found as the country of the affected entities (where the injury

has incurred) or as the country of at least one node within the occurrence zone (the place of the harmful action), then the admissibility control accepts to judge the TD. The same happens with the judgment process that is also the junction of the two functions to evaluate whether the TD is an NN violation by each criterion: `IsViolationByPlaceOfInjury` and `IsViolationByPlaceOfHarmfulAction` (Procedures 4 and 5, respectively). This algorithm also must check the level of uncertainty on the results of the `IsViolationByPlaceOfHarmfulAction` function, evaluating the similarity index of the occurrence zone before pointing out whether the TD is an NN violation when such criterion is adopted.

In this section, representative examples of Judgment Algorithms were presented. They adopt different approaches to establish jurisdictions and to judge TD as NN violation considering the regulation. These Judgment Algorithms output their decision that may be accompanied by one metric to indicate its level of certainty (*e.g.*, similarity index). Other algorithms could be designed using the criteria used in the presented examples but outputting a different metric that uses further information or processes the TD information in another way. Other algorithms could also be designed to accomplish other criteria to establish jurisdiction. The point is that given the factors that affect the judgment, a unique Judgment Algorithm is unfeasible to accommodate all these possibilities.

5.4 Summary

In this chapter, the proposed solution to judge TD as NN violation according to definitions from the NN regulations, named `JurisNN`, was presented. The `JurisNN` deployment scenario was presented, highlighting the interactions of the system with TD detection and positioning solutions, wrappers, and end-users. The deployment scenario also highlighted the actors that provide information for the system, including regulation experts that model regulatory information into the system model and judgment developers that design the Judgment Algorithms. The `JurisNN` architecture was briefly presented, given that its modules are significantly related to the concepts presented in the Modeling chapter. In turn, the Judgment Algorithms were detailed because they are responsible for accommodating the nuances in jurisdiction establishment criteria. However, distinct Judgment Algorithms may adopt similar jurisdiction establishment criteria, requiring similar processing. Thus, functions that implement the identified criteria to establish jurisdictions (by the place of the injury or the place of harmful action) for each Judgment Algorithm step (admissibility control and judgment) were presented. For each criterion and step,

the functions were presented highlighting the aspects that influence the jurisdiction establishment and how the information within the models are used to perform the steps. A few Judgment Algorithms were presented (Place of TD deployment, Targeting test, and German test), highlighting the jurisdiction establishment criteria used and, therefore, the functions used to implement them. The verdict conclusiveness of the functions is highlighted, which is influenced by TD information considered by the jurisdiction establishment criteria.

In the next chapter, JurisNN is evaluated from the functional perspective. A PoC prototype was developed following the conceptual architecture presented in this chapter. The functional evaluation is based on TD information collected by a state-of-the-art solution to assess the conclusiveness (the TD is an NN violation or not) of the results.

6 EVALUATION

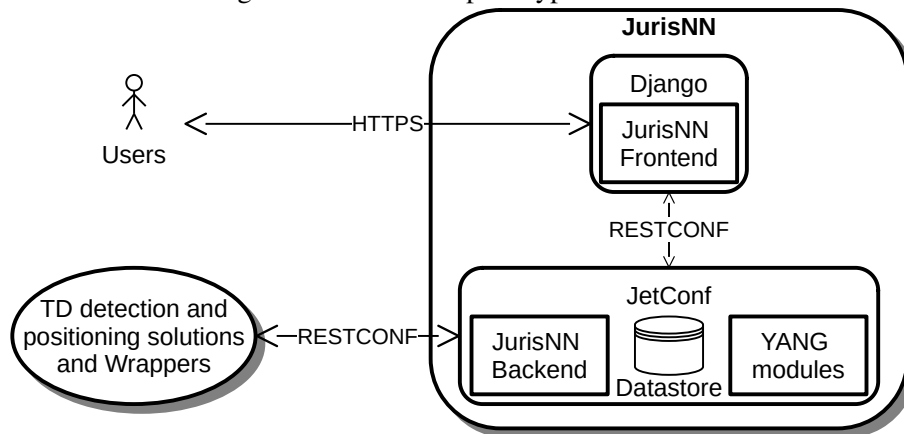
In this chapter, the proposed solution JurisNN is evaluated using information about TDs detected by Wehe (LI et al., 2019). The evaluation aims to answer whether the TD detection and positioning solutions provide enough information to properly judge TD as NN violation using the NN definitions from regulations. In Section 6.1, the architecture of the JurisNN prototype is presented along with its RESTCONF API. This API is used by TD detection and positioning solutions to submit TD information to perform the Regulation Assessment step. In Section 6.2, the dataset of TDs detected by Wehe is presented along with the required steps to complement the dataset with network paths collected using the RIPE Atlas platform. In Section 6.3, the results achieved by the JurisNN prototype are presented, discussing judgment results of each Judgment Algorithm presented in Section 5.3 and a discussion about the caveats of this analysis. In turn, in Appendix C, a performance characterization of the prototype is presented, evaluating its response time, and usage of CPU and memory.

6.1 JurisNN prototype

In this section, the JurisNN prototype that implements the conceptual architecture depicted in Figure 5.2 is presented. The prototype architecture is depicted in Figure 6.1, whose main components are the backend and the frontend, which are detailed next.

The backend was implemented using the JetConf framework (CZ.NIC, 2022). This framework is an implementation of the RESTCONF protocol (CLAISE; CLARKE; LINDBLAD, 2019) based on Python. This framework provides the basic functionality to perform Create Read Update Delete (CRUD) operations in the Datastore based on YANG models. The *JurisNN Backend* is the implementation of controllers, within the JetConf framework, that define the behavior of RESTCONF requests related to data or operations. The requests related to CRUD operations within the Datastore are handled by the data controllers. For instance, all the CRUD operations related to the *ietf-network-topology*, *nn-regulations*, and *tdo* models presented in Tables 4.2 and 4.3 are handled by the respective data controllers. In turn, the RPC operations defined in YANG models are handled by operation controllers. For instance, the RPC *analyze* defined on the *tdo* model, responsible for the judgment of Occurrence Zones, is implemented by an operation controller. It is important to note that RESTCONF is not required for the JurisNN solution. However,

Figure 6.1 – JurisNN prototype architecture



Source: the author

as the JetConf framework implements several facilities to build solutions based on YANG models and the interaction with this component may be based on REpresentational State Transfer (REST), it was used to build the prototype. However, the conceptual architecture depicted in Figure 5.2 can be implemented using another framework.

In Table 6.1, a few paths exposed by the RESTCONF API are exemplified. The containers and items of the YANG models are accessible through the API following the path patterns exemplified on the table. The HTTP method (GET, POST, PUT, DELETE) defines which kind of CRUD operation is performed, according to the REST definitions. The RPC operations are invoked by the POST method. This API is used both by the frontend and by TD detection and positioning solutions to interact with the backend.

The Judgment Algorithms are implemented in Python (JAWORSKI; ZIADÉ, 2021). One class *JudgmentAlgorithm* was developed to provide basic operations related to the admissibility and judgment steps (getters and setters) and defining the interfaces of the admissibility and judge methods. The Judgment Algorithms can inherit such basic operations and implement the interfaces for the admissibility and judge methods.

The *JurisNN Frontend* was developed using the Django framework (SHAW; BADHWAR; BIRD, 2021). The framework provides the Model View Controller (MVC) design pattern (BUSCHMANN; HENNEY; SCHMIDT, 2007) to develop web applications based on Python. As the backend performs the storage of information, only the view and controller capabilities of the Django framework were used to implement the User Interface. One class was developed to perform the communication of the controllers with the backend, which performs the operations through the RESTCONF API. In the next section, the data used to evaluate JurisNN is presented.

Table 6.1 – RESTCONF API

Model	Path
ietf-network	/restconf/data/ietf-network:networks
ietf-network	/restconf/data/ietf-network:networks/network=<id>
nn-regulation	/restconf/data/nn-regulation:regulations
nn-regulation	/restconf/data/nn-regulation:regulations/regulation=<id>
tdo	/restconf/data/tdo:occurrences
tdo	/restconf/data/tdo:occurrences/occurrence=<id>
tdo	/restconf/data/tdo:occurrences/occurrence=<id>/communication-graph
tdo	/restconf/data/tdo:occurrences/occurrence=<id>/occurrence-zone=<id>
tdo	/restconf/operations/tdo:analyze
jetconf	/restconf/operations/jetconf:conf-commit

Source: the author

6.2 Evaluation data

In this section, the data used in the JurisNN evaluation concerning the judgment results is presented. The detected TD information was collected from the Wehe dataset (MEASUREMENT LAB, 2021). As this dataset does not provide the network path between the clients and the servers of tests, the RIPE Atlas platform was used to collect traceroutes between the client and server ASs. As the Wehe tests traversed multiple jurisdictions, the NN definitions found in these jurisdictions are also presented. The details of how these data were used in the evaluation are presented along with this section.

The Wehe solution provides a dataset of conducted tests within the M-Lab platform (MEASUREMENT LAB, 2022) since October 2020 and also within its website (CHOFFNES, 2022) since November 2018. In the dataset provided through the website, the amount of available data fluctuates along with time. In the dataset provided through the M-Lab, the amount of data is stable after December 2020, according to information provided by one of the authors. This analysis considered the data that was collected from January 1, 2021, until February 28, 2021, and publicized in the M-Lab platform (MEASUREMENT LAB, 2021) consisting of 170 GiB of data about 77270 tests.

The Wehe authors provide analysis scripts (CHOFFNES, 2022) to process the dataset achieving the verdicts about the tests. The scripts require the data about the tests organized in one folder with one sub-folder by day with their respective tests (*e.g.*, *wehe/2021-01-01*, *wehe/2021-01-02*). However, the M-Lab platform provides the data organized by one sub-folder for each year, with one sub-folder for each month, with one sub-folder for each day (*e.g.*, *wehe/2021/01/01*, *wehe/2021/01/02*). One script was developed to convert the M-Lab folder structure to the structure expected by the scripts.

Table 6.2 – Applications affected by the TDs detected by Wehe between January, 1st 2021 and February, 28th 2021

Application	TDs
Amazon	78
DisneyPlus	43
Hulu	49
NBCSports	37
Netflix	50
Twitch	28
Vimeo	23
Youtube	161
Total	469

Source: the author

The dataset is analyzed using two scripts provided by Wehe authors. The first script processes the dataset and creates a folder structure organizing the tests by the client ISP, based on the AS responsible for its IP. The second script performs the TD detection characterizing the tests as True Positives, True Negatives, False Positives, and False Negatives for each ISP. One script was developed to collect the information about the tests that detected TDs (True Positives and False Negatives), although only True Positives were found. Table 6.2 lists the applications affected and the number of TDs detected by Wehe.

In Table 6.3, the number of detected TDs for each pair of source ASN (derived from the client subnet) and destination ASN (derived from the server IP) is presented, receiving an id that is used to identify the source/destination ASN pairs in the results. The ASNs were retrieved using PyASN (Economics of Cybersecurity Research Group, Delft University of Technology, 2022) using information collected from the RIB information from 2021-02-01 (within the dataset time interval). The Table 6.4 summarizes the ASNs and their respective AS names.

For the analysis performed by the JurisNN prototype, both the communication graph (nodes and links used to perform the communication) and the occurrence zone information (the nodes and links where the TD was detected) are required. However, Wehe does not collect the network path between the client and the server of the tests nor perform the TD positioning to establish the occurrence zone. To overcome this limitation, the available information about each TD detected by Wehe was used to collect network paths to complement the dataset. The analysis scripts provide the /24 subnet of the client IP (the first 24 bits of the IP address with zeros in the last 8 bits). In turn, the analysis scripts do not provide information about the server that processed the client test. However, the Packet Capture (PCAP) files of each test are also available within the dataset. One script was developed to extract the server IP from the respective PCAP file. Thereafter, the client and server IP addresses are available to collect network paths between them.

Table 6.3 – TD information, collected network paths, and occurrences

id	Src. ASN	Dst. ASN	TDs	Traces	Occurrences (TDs x Traces)	id	Src. ASN	Dst. ASN	TDs	Traces	Occurrences (TDs x Traces)
1	6167	174	21	36	756	17	21928	3257	36	5	180
2	6167	1280	8	12	96	18	21928	3356	34	5	170
3	6167	1299	17	9	153	19	21928	6453	63	4	252
4	6167	3257	11	5	55	20	21928	6461	31	6	186
5	6167	3356	20	6	120	21	21928	6939	3	4	12
6	6167	6453	15	6	90	22	22394	174	13	0	0
7	6167	6461	24	6	144	23	22394	1280	2	0	0
8	12912	3257	1	3	3	24	22394	1299	7	0	0
9	20057	174	7	3	21	25	22394	3257	8	0	0
10	20057	3257	7	0	0	26	22394	3356	7	0	0
11	20057	3356	8	7	56	27	22394	6453	17	0	0
12	20057	6453	9	3	27	28	22394	6461	10	0	0
13	20057	6461	4	5	20	29	57269	1299	4	3	12
14	21928	174	45	20	900	30	57269	3257	3	2	6
15	21928	1280	1	6	6	31	57269	3356	1	2	2
16	21928	1299	24	10	240	32	57269	6453	8	5	40

Source: the author

Table 6.4 – AS numbers and the respective AS names

AS Number	AS Name
174	COGENT-174, US
1280	ISC-AS-1280, US
1299	TELIANET Telia Carrier, SE
3257	GTT-BACKBONE GTT, US
3356	LEVEL3, US
6167	CELLCO-PART, US
6453	AS6453, US
6461	ZAYO-6461, US
6939	HURRICANE, US
12912	TM, PL
20057	ATT-MOBILITY-LLC-AS20057, US
21928	T-MOBILE-AS21928, US
22394	CELLCO, US
57269	DIGISPAINTELECOM, ES

Source: the author

In order to complement the Wehe dataset with network paths between client and servers, traceroutes were collected using the RIPE Atlas platform (Réseaux IP Européens, 2022a). It is a platform composed of probes (hardware and software) to perform network measurements such as ping, traceroute, DNS resolution, and Secure Sockets Layer (SSL), HTTP, and Network Time Protocol (NTP) requests. The probes used to perform a measurement can be selected by criteria such as their area in the globe, country, prefix, or ASN. The platform provides a Python API named Cousteau (Réseaux IP Européens, 2022b) to request measurements.

The Python Cousteau API was used to request traceroute tests from each source ASN from Table 6.3 to one server IP in the corresponding destination ASN. Preliminary tests using the client /24 prefix to select probes did not yield satisfactory results. To

Table 6.5 – Jurisdictions found in source, destination, and network path

id	Jurisdictions
1	S={US, US-CT, US-HI}, D={US}, P={US, US-MA}
2	S={US, US-HI}, D={US}, P={US, US-MA}
3	S={US, US-TN, US-UT, US-VA}, D={SE}, P={SE, US, US-MA}
4	S={US, US-FL, US-PA, US-TN, US-UT}, D={US}, P={US, US-MA}
5	S={US, US-CT, US-MA, US-PA}, D={US, US-VT}, P={US, US-MA}
6	S={US, US-CT, US-FL, US-TN, US-VA}, D={US}, P={US, US-MA}
7	S={US, US-FL, US-HI, US-MA, US-UT}, D={US}, P={US, US-MA}
8	S={PL}, D={US}, P={DE, DE-HE, PL, US}
9	S={US, US-FL, US-OH}, D={US}, P={US, US-CA, US-NC}
10	S={US-CA, US-FL}, D={US}, P={}
11	S={US, US-CA, US-FL, US-TX}, D={US, US-FL}, P={US, US-CO, US-FL, US-NC}
12	S={US, US-FL, US-GA, US-TX}, D={US, US-FL, US-GA, US-TX}, P={US, US-CA, US-FL, US-NC}
13	S={US-FL, US-TX}, D={US, US-IN}, P={US, US-AZ, US-CA, US-NC}
14	S={US-AZ, US-CA, US-FL, US-IL, US-KY, US-MA, US-MI, US-NV, US-NY, US-PA, US-TX, US-WA}, D={US}, P={US, US-RI}
15	S={US-CA}, D={US}, P={SE, US, US-RI}
16	S={US, US-AL, US-CO, US-FL, US-NC, US-NY, US-TX, US-UT, US-WA}, D={SE}, P={SE, US, US-RI}
17	S={US-CO, US-FL, US-IL, US-IN, US-MA, US-NC, US-NY, US-PA, US-WA}, D={US}, P={US, US-RI}
18	S={US-CA, US-FL, US-IL, US-IN, US-KY, US-MA, US-NY, US-PA, US-TX, US-VA, US-WA}, D={US}, P={US, US-RI}
19	S={US, US-AL, US-CA, US-FL, US-GA, US-IL, US-IN, US-MA, US-MN, US-NC, US-NY, US-TX, US-WA}, D={US}, P={US, US-RI}
20	S={US-AZ, US-CA, US-CO, US-FL, US-IN, US-MA, US-MI, US-PA, US-TX}, D={US}, P={CA, ES, GB, IE, US, US-RI}
21	S={US-MI}, D={US}, P={US, US-RI}
22	S={US-CA, US-IL, US-IN, US-NY, US-OH}, D={US}, P={}
23	S={US-CA}, D={US}, P={}
24	S={US-GA}, D={SE}, P={}
25	S={US-IL, US-NY}, D={US}, P={}
26	S={US-CA, US-IN, US-NY, US-TX}, D={US}, P={}
27	S={US-CA, US-GA, US-IL, US-IN, US-NY, US-OH, US-VA}, D={US}, P={}
28	S={US-CA, US-IN, US-NY, US-OH}, D={US}, P={}
29	S={ES}, D={SE}, P={ES, RO-TR, SE}
30	S={ES}, D={US}, P={ES, GB, HU-JN, US}
31	S={ES}, D={GB}, P={ES, GB, RO-TR, US}
32	S={ES}, D={IE}, P={CA, ES, GB, IE, RO-TR, SE, US}

Legend: S - Source, D - Destination, P - Path

Source: (CARVALHO; GRANVILLE, 2022)

overcome this limitation, the probe selection was performed by the source ASN, which also did not found network paths due to the lack of probes for a few source/destination ASN pairs. The traceroute measurements were requested for one week (from 2021-04-12 until 2021-04-19) six times a day (0:00, 4:00, 8:00, 12:00, 16:00, and 20:00 GMT) resulting in 42 requests for each source/destination ASN pair in the measurement period. The traceroute results were compared to only consider different traceroutes between the source/destination ASN pair in the analysis. For the traceroute comparison, the internal network addresses (10.0.0.0/8, 172.16.0.0/16, and 192.168.0.0/16) were discarded from the paths. Traces with the same hop indexes associated with the same addresses are considered equal. Otherwise, they are considered different. In Table 6.3, the number of different traces found by the RIPE Atlas platform for each source/destination ASN pair is presented. The results show that within the possible 42 collected traces, most of the source/destination ASN pairs have less than ten different traces. Only the pair with id 1 has a higher number of different network paths (36).

Each source/destination ASN pair has a number of associated TDs and a number of network paths. For the analysis conducted using the JurisNN prototype, one TD Oc-

currence was submitted for each TD and each possible network path. For instance, the source/destination ASN pair id 1 (source ASN 6167 and destination ASN 174) has 21 associated TDs and 36 different traces connecting these ASs generating the corresponding $21 * 36 = 756$ TD Occurrences. For the source/destination ASN pairs that RIPE Atlas did not return possible network paths, TD Occurrences were not submitted to JurisNN.

In Table 6.5, the jurisdictions found as source (based on the client /24 subnet considering the last octet as 1, *i.e.*, the first address that could be attributed to a host in the subnet), destination (based on the server IP), and the ones found along with the network path (based on the traceroutes) are presented. The Geo Localization was performed using the MaxMind GeoLite2 database (Maxmind, 2022). The database returns the “most specific” location for some IP addresses, which was used to determine state-level jurisdictions (*e.g.*, US-CT, US-HI). When there is no most specific location, the jurisdiction was considered at the country-level (*e.g.*, USA, SE).

In Table 6.5, most of the jurisdictions found (source, destination, and path) are within the USA or USA states. Many USA states were identified as source jurisdictions, given that the identified ASs has a presence in multiple USA states. Some of the jurisdictions found are in the EU. In addition, Canada (CA) and the United Kingdom (GB) also were found. For a few ASN pair ids, the results point out that Wehe tests were originated in one jurisdiction and were processed by a server in a far jurisdiction. For instance, for ASN pair id 3, the Wehe tests were originated in the USA, and the server is located in Sweden (SE). For ASN pair id 8, the tests were originated in Poland (PL) and the server is located in the USA, but the packets also traversed Germany (DE). Similar situations also can be noted in ASN pair ids 15, 16, 20, 29, 30, 31, and 32. The fact of tests traversing multiple jurisdictions (that may be far away in other countries or may be close in a neighbor state) exposes the possibility of these traffics crossing distinct regulatory frameworks with different NN definitions. In Table 6.6, the NN definitions about throttling (the TD type detected by Wehe) are presented for the jurisdictions traversed by Wehe tests, whose definitions are detailed next.

A few jurisdictions allow throttling. The allowance may be explicit as in CA that allows throttling since they can be considered justifiable (CRTC - Canadian Radio-Television and Telecommunications Commission, 2009). Alternatively, the allowance may be implicit given to the lack of explicit prohibitions established *a priori* as in the USA (Federal Communications Commission, 2018). In both cases, ISPs are entitled to perform the traffic management practices deemed necessary.

Table 6.6 – Regulations about throttling in the found jurisdictions valid from 2021-01-01 onwards

Jurisdiction	Throttling	Jurisdiction	Throttling	Jurisdiction	Throttling	Jurisdiction	Throttling
CA	✓*	US	-	US-IN	-	US-RI	-
DE	✗*↑	US-AL	-	US-KY	-	US-TN	-
DE-HE	✗*↑	US-AZ	-	US-MA	-	US-TX	-
ES	✗*↑	US-CA	✗*	US-MI	-	US-UT	-
GB	✗*	US-CO	✗*	US-MN	-	US-VA	-
HU-JN	✗*↑	US-CT	-	US-NC	-	US-VT	✗*
IE	✗*↑	US-FL	-	US-NV	-	US-WA	-
PL	✗*↑	US-GA	-	US-NY	-		
RO-TR	✗*↑	US-HI	-	US-OH	-		
SE	✗*↑	US-IL	-	US-PA	-		

Legend: allowed (✓), prohibited (✗), upper-level regulation (↑), none regulation (-), has exceptions (*)

Source: the author

The lack of country-level regulation in the USA opened the opportunity for USA states to establish their NN regulations. The National Conference of State Legislatures (NCSL) maintains a website (National Conference of State Legislatures, 2021) organizing the NN regulation efforts conducted by USA states, which was used in this research. Several states started discussions around the subject on their legislative organizations. However, only three states (California, USA (US-CA), Columbia, USA (US-CO), and Vermont, USA (US-VT)) finished the legislative proceedings of such regulations, reestablishing the FCC Open Internet rules from 2015 (Federal Communications Commission, 2015), which prohibits the throttling of lawful content, applications, services, and devices.

A few jurisdictions found as the source, destination, and network paths do not have local NN regulations. However, they are within the jurisdiction of a broader upper-level NN regulation. All the jurisdictions in this situation are within the EU jurisdiction (DE, Hessen, DE (DE-HE), Spain (ES), Jász-Nagykun-Szolnok, Hungary (HU-JN), Ireland (IE), PL, Teleorman, Romania (RO-TR), and SE) under the BEREC regulation. This regulation establishes that throttling is prohibited, but it may be allowed under certain situations that the NRA shall evaluate and decide (Body of European Regulators for Electronic Communications, 2016).

The Brexit process has finished on December 31, 2020. Therefore, the GB left the EU on January 1, 2021 (the beginning of the interval considered in the Wehe dataset). However, the Ofcom (the British NRA) still is applying the EU regulation (Ofcom, 2020) that prohibits throttling.

In this section, the data used in the JurisNN evaluation was presented, which consists of the detected TD information (provided by the Wehe dataset), the network path between the client and server ASs (provided by the RIPE Atlas platform), and the NN

definitions found on the jurisdictions traversed by Wehe tests. In the next section, the details of how this data was used to evaluate JurisNN regarding the judgments and the achieved results are presented.


6.3 JurisNN judgment results

This section details how the data presented in the previous section was used to evaluate JurisNN regarding the judgment results. The NN definitions summarized in Table 6.6 were introduced into the regulation datastore respecting the regulation hierarchy (top-, region-, country-, and state-level regulations). No information about the jurisprudence was used to configure the interpretations for such regulations (*e.g.*, what is considered justifiable unlawful content for the Canadian regulation?).

For each TD detected by Wehe, the JurisNN was used to judge whether such TD is considered an NN violation according to the regulation established in the jurisdictions traversed by the Wehe tests. As for most TDs, multiple network paths were found between the client and server ASs using the RIPE Atlas platform, each TD detected by Wehe was associated with each network path found by RIPE Atlas to compose a TD Occurrence for JurisNN. In order to evaluate a TD Occurrence by JurisNN, information about the communication topology and the TD itself need to be submitted, which are detailed next.

Each network path associated with each TD was submitted to JurisNN as one `ietf-network:networks/network`. Each hop in the network path was submitted as one node in the `ietf-network` using the hop index as `node-id` and identified Geo Localization in the `located-at` attribute. As no IP or link information is used by the Judgment Algorithms presented in Subsection 5.3.2, such information was not submitted to JurisNN.

Each TD detected by Wehe was associated to one network path identified by RIPE Atlas. The TD and the associated network path were submitted to JurisNN as one `tdo-occurrences/occurrence`. All the nodes in the associated network path were referenced in the communication-graph associated to the occurrence. Wehe detects the throttling of applications. Thus, one occurrence zone about the Differentiation “Throttling” and the DifferentiatedEntity “Application” was submitted to JurisNN for each detected TD and network path. The name of the application, summarized in Table 6.2, was submitted as the attribute name of the DifferentiatedEntity. Also, as Wehe does not point where in the network path the TD was deployed, all the nodes in the communication-graph of the occurrence were referenced in the occurrence zone. After the submission of the above



Source: (CARVALHO; GRANVILLE, 2022)


information, JurisNN judges the TD by invoking the `tdo:analyze` operation, informing the occurrence-id and the occurrence-zone-id.

The judgment of the TD is performed by the Judgment Algorithms, presented in Subsection 5.3.2. These algorithms have two steps: admissibility control and the judgment itself. Next, the results achieved by the three Judgment Algorithms are presented.

The first algorithm is the Place of the TD deployment, which evaluates the NN definitions established in the place where the nodes within the Occurrence Zone are deployed. As this algorithm applies to any jurisdiction, the admissibility control accepts to judge all TDs. In Figure 6.2, the achieved results using box plots that represent the minimum, first quartile, median, third quartile, and maximum values of the distribution of results are presented. The ASN pair ids presented in the x-axis are those that were admitted by the algorithm, which are all the pairs submitted to JurisNN. The missing pair ids are those for which RIPE Atlas has not identified network paths. Therefore, they were not submitted to JurisNN (ids 10, 22, 23, 24, 25, 26, 27, and 28).

The judgment by the Place of the TD deployment algorithm returns occurrence zone similarity index and the most prevalent verdict about the TD being an NN violation by the algorithm. The similarity index reflects the certainty in the verdict, in which similarity 1.0 indicates that for all nodes within the Occurrence Zone the verdict is the same. For the ASN pair ids 1, 2, 4, 6, 7, and 21, the similarity index for all Occurrence Zones analyzed is 1.0. Looking for the jurisdictions traversed by the Wehe tests associated with these ASN pair ids in Table 6.5, all traversed jurisdictions are in the USA and none of these jurisdictions are those that reestablished the Open Internet rules (US-CA, US-CO, and US-VT). Therefore, the detected TDs associated with these ASN pair ids are not NN

Figure 6.3 – Place of the TD deployment: verdicts distribution



Source: (CARVALHO; GRANVILLE, 2022)

violations when the regulations are considered. For all other ASN pair ids, the similarity index ranges below 1.0, indicating that the achieved verdicts have an uncertainty level, hindering the judgment whether the TD is an NN violation. The high similarity indexes for the ASN pair ids 5, 9, 11, 12, 13, 14, 17, 18, and 19 are explained because their tests traversed jurisdictions within the USA, but a few nodes are in the jurisdictions that reestablished the Open Internet rules. Therefore, the TD is considered an NN violation in a small part of the Occurrence Zone and it is not considered a violation in the another part. The high similarity index for ASN pair 29 is explained because its tests traversed countries in the EU, where the BEREC regulation is in place, prohibiting throttling. However, the Geo Localization process did not return the location for a few nodes, hindering the verdict achievement for these nodes. The wide ranges in the similarity index for the ASN pair ids 15, 16, 20, 31, and 32 are explained because the tests traversed jurisdictions in USA and EU, whose regulations diverge about throttling. The same situation happens for ASN pair ids 3, 8, and 30, but with narrower similarity index ranges, indicating that network paths associated to the ASN pair ids have a homogeneous mix of jurisdictions in USA and EU that agree on the verdict about the TD.

The results expose that, for most scenarios, the pointing out of whether the TD is considered an NN violation is not possible (except those that the similarity index is 1.0 without ranges in the box plots). The overall result is that for only 6 (from 24) ASN pair id scenarios the achieved results are conclusive, indicating the need for better TD positioning. Another analysis without grouping the TDs in the client/server ASN pairs also was performed. In Figure 6.3, the distribution of verdicts without such grouping is presented. This analysis points out to a high amount of conclusive tests also using the whole network path as Occurrence Zone as performed in the previous analysis. The conclusive verdicts (No Violation and NN violation) sum 81.9% of the occurrences submitted to JurisNN. In turn, the inconclusive results sum 18.1% of the occurrences. However, most TDs detected by Wehe involved ASs within the USA, where most jurisdictions allow the

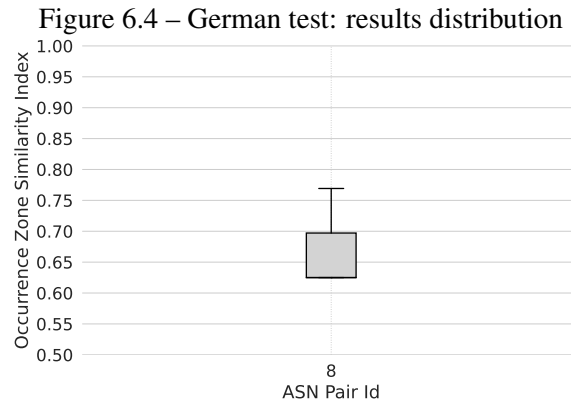
Table 6.7 – Targeting test: results distribution

ASN Pair id	Judgment		ASN Pair id	Judgment	
	No violation	NN violation		No violation	NN violation
1	756	0	13	20	0
2	96	0	14	880	20
3	153	0	15	0	6
4	55	0	16	220	20
5	108	12	17	160	20
6	90	0	18	160	10
7	144	0	19	232	20
8	3	0	20	156	30
9	21	0	21	12	0
11	42	14	30	6	0
12	27	0			

Source: (CARVALHO; GRANVILLE, 2022)

throttling of applications due to the lack of NN regulation. Therefore, given this situation, the Wehe tests performed within the USA are likely to traverse jurisdictions that have the same verdict about the detected TD. Therefore, the dataset may be biased towards situations in which regulations (actually, the lack of) agree about the TD be an NN violation or not. Then, the Occurrence Zones submitted to JurisNN are likely to have high similarity indexes. Thus, the results grouping ASN pair ids as scenarios seem to be more representative. Therefore, considering the better representativeness of scenarios, the data analyzed as described in Section 6.2, and all caveats presented along with this section and in Subsection 6.3.1, this analysis still points to the need for better TD positioning to judge TD practices considering the regulations using the Place of the TD deployment algorithm.

The second algorithm is the Targeting test, which considers the location of the endpoints of the communication affected by the TD to establish the jurisdiction. Wehe detects the throttling of applications. Therefore, the Targeting test only considered TDs whose DifferentiatedEntity is “Application.” In such cases, the algorithm evaluates both endpoints (client and server) because the TD affects both the application provider and users. USA courts adopt the Targeting test. Therefore, the admissibility control checks if the client or server are within the USA to admit judging the TD. In Table 6.7, the results listing the ASN pair ids admitted to being judged by the algorithm are presented. It is important to note that this algorithm does not depend on the TD positioning accuracy because only the jurisdictions in the endpoints are considered. Therefore, the algorithm returns conclusive results by design. For 12 ASN pair ids (1, 2, 3, 4, 6, 7, 8, 9, 12, 13, 21, and 30), the algorithm just pointed out TDs that are not considered as NN violations. For 8 ASN pair ids (5, 11, 14, 16, 17, 18, 19, and 20), the algorithm pointed a few TDs



Source: (CARVALHO; GRANVILLE, 2022)

that are considered NN violations, but most of their associated TDs are not considered NN violations. Only for the ASN pair id 15, all the TDs are considered NN violations because all the TDs have US-CA (that reestablished the Open Internet rules) as source jurisdiction. As this algorithm does not depend on the TD positioning, the information provided by the TD detection solutions was enough to judge the TDs.

The third algorithm is the German test, which applies both criteria adopted by the previous algorithms. The algorithm evaluates the jurisdictions at the endpoints (the place where the injury has incurred) and along with the network path (the place where the harmful action took place). Courts in DE apply this test. Therefore, only the TDs that have the endpoints in DE or that traversed the DE are admitted to be judged by the algorithm. In Figure 6.4, the achieved results using box plots that represent the minimum, first quartile, median, third quartile, and maximum values of the distribution of the results are presented. Only the ASN pair id 8 was admitted to be judged because their associated TDs have the source in PL and destination in the USA but traversed DE along with the network path. Therefore, the results were achieved by the place of the harmful action criterion. As the network path is part in the EU and part in the USA, the similarity indexes range from 0.62 to 0.77, also indicating the divergence about the TD within the Occurrence Zones. Therefore, considering the data analyzed as described in Section 6.2 and all caveats presented along with this section and in Subsection 6.3.1, all these results are inconclusive, indicating also that TD positioning accuracy (the whole network path as occurrence zone) was not enough to judge the TDs by the German test algorithm.

The findings of the analysis are summarized now. The Targeting test algorithm does not depend on the jurisdictions traversed along with the network path because its judgment is based on the jurisdictions of the endpoints established by the Geo Localization of the client and server. Thus, the algorithm is not affected by the positioning of the

TD. Therefore, the available information provided by the state-of-the-art solution was enough for the judgment by this algorithm. Both the Place of the TD deployment and the German test algorithms depend on the network path information because they evaluate the place where the TD was deployed. Therefore, considering the data analyzed as described in Section 6.2, and all caveats presented along with this section and in Subsection 6.3.1, the analysis shows that the judgment performed by such algorithms using the whole identified network path as the Occurrence Zone achieved inconclusive results (similarity index below 1.0). Thus, better TD positioning is required (the whole network path is not enough) to point smaller Occurrence Zones, narrow to the actual TD deployment, to correctly judge the TD as NN violation using the definitions established in the NN regulation. Indeed, Garrett *et al.* (GARRETT; BONA; DUARTE, 2021) propose a method to determine where the TD was deployed at AS-level, which may be a proper level for the Regulation Assessment. Therefore, there is the expectation that the results of this thesis incentive that NN violation detection solutions have the functionalities both to identify network paths and to position the TD in the network path. Indeed, the results presented along with this thesis shown that this is a requirement to properly judge TD as NN violation with support of the regulation, which can help users providing actionable evidence to support claims against unfair traffic management practices deployed by ISPs.

6.3.1 Analysis caveats

The caveats of the performed analysis need to be presented. Most of them are related to the lack of network path information provided by TD detection and positioning solutions and the Geo Localization process.

The Geo Localization was performed using the MaxMind GeoLite2 IP database (Maxmind, 2022), which is one of the most used databases for Geo Localization and it is adopted in several scenarios. For instance, firewalls use its information to block IP addresses from foreign countries (F5, 2022). The database has good coverage. Considering all addresses submitted to JurisNN (39595), for only 2.475 hops (6.3%) the database did not have Geo Location information. However, the database may have inaccurate information. For instance, the ASN pair id 32 is related to a client in ES connecting to a server in IE, both in Europe. However, the identified network path traversed CA or the USA in North America at some point, which may be possible but is unlikely to be true. Therefore, the results may suffer the influence of the inaccuracy in the Geo Localization process.

The network path between the client and the server was collected using the RIPE Atlas platform (Réseaux IP Européens, 2022a) conducting traceroutes from the client AS targeting the server IP a few months after the TD detection. This approach has some caveats. As the network path was collected months after the TD detection, the network has likely changed along with this time. Even the network conditions are different from those when the TD was detected, which may impact the routing. Another issue is that the traceroute uses the ICMP protocol and it may not identify all hops in the network path because the ISP routers can be configured not to respond to ICMP messages or rate-limit such messages. Another issue is that ICMP packets may be routed differently from application packets. Therefore, the identified network path may not reflect the path traversed by the application traffic. However, this may indicate that the ISP is performing TD (by routing application traffic differently), which is prohibited by a few NN regulations. In order to overcome such issues, the TD detection and positioning solution could identify the network path using similar packets used to perform the TD detection (as performed by NeutMon (GREGORI; LUCONI; VECCHIO, 2018a) that crafts the TTL field of probe packets to identify routers along with the network path). Therefore, the network path could be collected simultaneously as the TD detection is performed. Thus, facing the same network conditions and being routed as the application packets were routed.

The RIPE Atlas platform offers methods to select the probe to perform the tests based on the network prefix, ASN, country, or globe area. In the initial phase of this analysis, the same network prefix of the client of the TD was used to request probes to perform tests. However, very few probes were selected using this criterion. Therefore, the criterion was shifted to use the client ASN to request probes, which selected probes for 398 TDs (84.9%) from the 469 TDs. However, the use of the client ASN to select the probes may introduce issues because most of the ASNs listed in Table 6.4 have a presence in multiple jurisdictions (*e.g.*, multiple states within the USA). Therefore, the tests may be originated from a locality different from that of the client of the TD detected by Wehe. This issue could be overcome if the TD detection and positioning solution would have performed the network path identification.

It is important to note that the dataset content impacts the analysis results directly. Therefore, one dataset from another solution or the Wehe dataset analyzed in another time window may produce other results and discussion. However, the main conclusions about the lack of network path information and improper TD positioning accuracy are likely to hold even using another dataset.

6.4 Summary

In this chapter, the JurisNN prototype was evaluated regarding the conclusiveness of the judgment results achieved by the three Judgment Algorithms: Place of the TD deployment, Targeting test, and German test. The Place of the TD deployment results showed that for 6 from 24 of the scenarios (ASN client and server pairs), the judgment provides conclusive verdicts (similarity index of the Occurrence Zones is 1.0). Therefore, considering the data analyzed as described in Section 6.2, and all caveats presented along with Section 6.3 and in Subsection 6.3.1, the results indicate the requirement of better TD positioning information than the whole network path, which was the approach adopted. The Targeting test is not affected by the network path and TD positioning information because it evaluates only the endpoint jurisdictions. Its results showed that most of the TDs are not considered NN violations in accordance to the regulations because the algorithm admits only TDs whose at least one endpoint is within USA, where for most jurisdictions (state and country-level) there is no NN regulation in place. The German test only admitted one scenario (TD client and server pair) whose TD detection test traversed DE. However, the network path traversed jurisdictions in EU and the USA, thus, mixing jurisdictions where the TD is considered an NN violation with jurisdictions that do not. The judgment similarity indexes range from 0.62 to 0.77, which are inconclusive verdicts. Therefore, considering the data analyzed as described in Section 6.2, and all caveats presented along with Section 6.3 and in Subsection 6.3.1, the results also indicate the requirement of better TD positioning information than the whole network path to properly judge TDs by this algorithm.

In the next chapter, the research questions are revisited to discuss the answers provided (using the results of this evaluation) and the remaining gaps subject to future investigations are pointed out.

7 CONCLUSION

This thesis has presented the state-of-the-art solutions for the detection of TD and NN violation, which usually are based on strict NN definitions from academia or on the regulation stated on the jurisdiction of their proponents. The former does not consider the regulatory aspect of the Internet, providing incomplete information to the user to support claims against ISPs. The latter does not consider that users' traffic may traverse multiple jurisdictions, thus having multiple NN definitions along with an end-to-end network path. A novel approach to point out NN violation considering the multiple NN definitions found in an end-to-end network path is introduced, adding the Regulation Assessment step after the TD detection and positioning steps performed by state-of-the-art solutions. In addition, given the difficulty that such assessment be performed by every solution devoted to detecting TD and NN violation, a service is proposed that is responsible for the management of information related to the NN definitions stated on regulations around the world, the processing of TD information to compare it to these definitions, and pointing out the NN violation when the TD violates these definitions. In this context, research was conducted to verify the following hypothesis.

Hypothesis: a service performing the Regulation Assessment step can provide the requirements to point out NN violation considering multiple regulations found along with an end-to-end network path.

This hypothesis motivated the establishment of three Research Questions (RQs) that guided the research. Based on the work conducted in this thesis, the answers to the RQs are detailed next.

RQ 1: *How much the consideration of regulations impacts the results that have been achieved by state-of-the-art solutions for NN violation detection?*

Answer – Such consideration has a high impact on the results.

In this thesis, an impact analysis was conducted by evaluating whether TDs detected by Glasnost still could be pointed out as NN violations when the regulations are considered. The results showed a small but not negligible disagreement (the TD is a violation or not) between the regulations of the endpoint jurisdictions (5% to 9%). Considering the evaluated data, this result indicates that only the information about the jurisdictions of the endpoints is not enough to properly point out NN violation, requiring a better positioning of where the TD happened to define the proper regulations to be applied. In addition,

considering the evaluated data, the results showed also that under certain situations from 39% to 48% of the detected TDs were not considered NN violations by the regulations. This result showed the high impact of such consideration on state-of-the-art results and confirmed the assumption that solutions must consider the regulation. Otherwise, their results have a good chance of being incorrect.

RQ 2: *How to enable the pointing out of NN violation according to multiple definitions stated on regulations set by legislators and regulatory agencies?*

Answer – Judgment Algorithms can be developed to perform the regulation assessment using information from data models.

In this thesis, information models were designed to represent the information required to perform the regulation assessment. The modeling requirements were based on the findings of the conducted impact analysis. These models represent information about the regulations, the Traffic Differentiations, and the network topology that support the affected communication. The latter is essential to represent wherein an end-to-end network path the TD is suspect of being deployed, which is named as *Occurrence Zone*. YANG models (existing and proposed complementary ones) can be used to represent the information within the solution. The investigation identified that the jurisdiction of a case can be established by different methodologies in distinct legal systems, which impacts what regulation should be considered. Judgment Algorithms can be developed to accommodate such nuances. In the end, these algorithms are responsible for inspecting the TD information submitted by TD detection and positioning solutions and for analyzing the appropriated regulatory instructions to point out the NN violation when the submitted TD violates the regulation.

RQ 3: *Is the TD information collected by TD detection and positioning solutions enough to point out NN violation considering multiple regulations found along with an end-to-end network path?*

Answer – No. Considering the evaluated data, there is a lack of information about the traversed network paths and about where in this network path the TD is suspected of being deployed.

In this thesis, a PoC prototype named JurisNN was developed to perform the regulation assessment step. Information about TDs collected by Wehe was used to analyze the conclusiveness of the results achieved by JurisNN with the available information. The prototype performs three Judgment Algorithms that represent the jurisdiction establishment methodologies identified during the investigation: Place of TD deployment, Target-

ing test, and German test. For the Targeting test adopted by USA courts, the provided information was enough to point out NN violation following the regulation. This test establishes the jurisdiction based on who was targeted by the TD, which may be users or application/service providers. Thus, the jurisdiction is established in the endpoints that are easily identified and positioned by the communication source/destination addresses.

The Place of TD deployment and German test were affected by the lack of information. Both Judgment Algorithms depend on the place in the end-to-end network path where the TD was introduced to establish the jurisdiction (for consequence, the proper regulation to be considered). The Wehe dataset does not provide such information. The conducted analysis complemented the dataset information collecting network paths between the source AS of the client and the destination IP of the server involved in each TD detected by Wehe. In this analysis, the whole network path was considered as the Occurrence Zone given the lack of information of where the TD was deployed. For the Place of TD deployment algorithm, the analysis found that for 18 of 24 identified scenarios (source and destination AS pairs), considering the evaluated data and the analysis caveats, the results were inconclusive (similarity index of the verdicts within the Occurrence Zone is below 1.0). For the German test, only one client/destination AS pair was admitted to being judged by it (only one traversed DE). The similarity indexes of the verdicts within the Occurrence Zones range from 0.62 to 0.77, which are also inconclusive. Considering the evaluated data and the analysis caveats, the results of both Judgment Algorithms point to the need of better information about where in the end-to-end network path the TD is being deployed, which already had been identified by the impact analysis but based on less information. It is noteworthy that the network path information had to be collected to complement the information provided by Wehe in order to allow the conducted analysis, thus also indicating the lack of enough information.

The investigations conducted and the results presented in this thesis support the proposed hypothesis. A prototype of a service performing the Regulation Assessment step was developed and evaluated. However, the investigation found that the state-of-the-art solutions do not collect and provide the information that is required to properly point out NN violation based on regulations. However, there are available proposals that could be incorporated by solutions to collect such information. The end-to-end network paths could be collected using the approach of NeutMon (GREGORI; LUCONI; VECCHIO, 2018a). The positioning of the TD could be achieved using Garrett *et al.*'s approach (GARRETT; BONA; DUARTE, 2021). This thesis also showed the importance of

the Regulation Assessment step to provide reliable information to support user's claims. Therefore, given these findings, the authors of TD detection and positioning solutions are expected to be encouraged to incorporate the collection of the required information in their solutions. The adaptation of solutions is also expected in order to submit the collected TD information for the Regulation Assessment service.

The prototype developed as a PoC allowed to collect the results discussed in this thesis. However, improvements are required to have a functional solution. The RESTCONF is designed for configuration management having a transactional semantic to submit and commit new configuration versions into the datastore. When a new configuration is committed, the RESTCONF evaluates whether the whole new datastore version respect the YANG model definitions before committing the transaction. However, the JurisNN operates as a Web Service in which the submitted information does not need to be kept into the system for long time. As long several requests are submitted to JurisNN, the time to commit the transaction becomes higher due to all information residing in the datastore. A new framework could be developed to leverage the YANG models without incurring the burden of the RESTCONF transactional model. JurisNN also needs an effort to collect and model the NN regulation and its interpretations stated worldwide into the system. An observatory that collect and organize such information could help to keep the system up-to-date with the NN regulation established worldwide.

Based on the conducted research, future work opportunities in the context of this thesis are identified. For instance, this thesis focused on the relationship of TD detection and positioning solutions and the Regulation Assessment service. However, there are regulatory instructions that allow TD to be deployed under situational network conditions (*e.g.*, congestion) but requiring that the ISPs publicize such information for transparency. Therefore, the service could be amended to incorporate transparency information within the solution that Judgment Algorithms could use. The development of other Judgment Algorithms is also an open issue. The developed algorithms just considered the nodes in the network topology because none solution provided complex network topology information where the link information could be considered. For instance, an algorithm could consider the weights of each network path on the similarity index calculation when the traffic traverses multiple paths.

REFERENCES

- BASHKO, V. et al. BonaFide: A traffic shaping detection tool for mobile networks. In: **Proceedings of the 2013 IFIP/IEEE International Symposium on Integrated Network Management, IM 2013**. [S.l.: s.n.], 2013. ISBN 9783901882517.
- BASSO, S.; SERVETTI, A.; DE MARTIN, J. C. The network neutrality bot architecture: A preliminary approach for self-monitoring of internet access QoS. In: **Proceedings - IEEE Symposium on Computers and Communications**. [S.l.: s.n.], 2011. p. 1131–1136.
- BAUMANN, H.; GRÄSSLE, P.; BAUMANN, P. **UML 2.0 in Action: A Project-based Tutorial**. [S.l.]: PACKT, 2005. (From technologies to solutions). ISBN 9781904811558.
- BDK Advokati. **Zero rating vs net neutrality – a (still) uncertain future in the EU and Serbia**. [S.l.], 2017. Accessed on August 2022. Available from Internet: <<https://www.lexology.com/library/detail.aspx?g=8df9a669-06cc-4296-b9fc-04a5eb06af46>>.
- BEVERLY, R.; BAUER, S.; BERGER, A. The internet is not a big truck: Toward quantifying network neutrality. In: **Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)**. [S.l.: s.n.], 2007. ISBN 9783540716167. ISSN 16113349.
- BICZÓK, G.; YOUNG, W.; KUZMANOVIC, A. Monitoring network bias. In: **ACM SIGCOMM**. [S.l.: s.n.], 2008.
- BJORKLUND, M. **The YANG 1.1 Data Modeling Language**. Fremont, CA, USA, 2016. (Internet Request for Comments, 7950). Accessed on August 2022. Available from Internet: <<https://www.rfc-editor.org/rfc/rfc7950.txt>>.
- Body of European Regulators for Electronic Communications. **BoR (16) 127 - BEREC Guidelines on the Implementation by National Regulators of European Net Neutrality Rules**. 2016.
- BUSCHMANN, F.; HENNEY, K.; SCHMIDT, D. **Pattern-oriented Software Architecture: On Patterns and Pattern Language**. [S.l.]: John Wiley & Sons, 2007.
- BUSTOS-JIMÉNEZ, J.; FUENZALIDA, C. All packets are equal, but some are more equal than others. In: **Proceedings of the 8th Latin American Networking Conference, LANC 2014**. New York, NY, USA: ACM, 2014. (LANC '14), p. 5:1—5:8. ISBN 978-1-4503-3280-4.
- BUSTOS-JIMENEZ, J. et al. Adkintun: SLA monitoring of isp broadband offerings. In: **Proceedings - 27th International Conference on Advanced Information Networking and Applications Workshops, WAINA 2013**. [S.l.: s.n.], 2013. ISBN 9780769549521.
- CARTER, K. R. et al. A Comparison of Network Neutrality Approaches In: The U.S., Japan, and the European Union. **Ssrn**, p. 1–30, 2010.
- CARVALHO, M. B. et al. Quantifying the influence of regulatory instructions over the detection of network neutrality violations. In: **2020 IFIP Networking Conference (Networking)**. [S.l.: s.n.], 2020. p. 334–342.

CARVALHO, M. B.; SCHAURICH, V. G.; GRANVILLE, L. Z. Considering Jurisdiction When Assessing End-to-End Network Neutrality. **IEEE Internet Computing**, 2018.

CARVALHO, M. B. d.; GRANVILLE, L. Z. Jurisnn: Judging traffic differentiations as network neutrality violations according to the regulation. **Computer Networks**, v. 205, p. 108738, 2022. ISSN 1389-1286.

CASTRO, R. et al. Network tomography: Recent developments. **Statistical Science**, 2004. ISSN 08834237.

CHERTOFF, M.; ROSENZWEIG, P. A Primer on Globally Harmonizing Internet Jurisdiction and Regulations. **Global Commission on Internet Governance, Paper Series**, No. 10, n. 10, p. 1–16, 2015.

CHOFFNES, D. **Wehe website**. 2022. Accessed on August 2022. Available from Internet: <<https://wehe.meddle.mobi/>>.

CLAISE, B.; CLARKE, J.; LINDBLAD, J. **Network Programmability with YANG**. [S.l.]: Addison Wesley, 2019. ISBN 9780135180396.

CLEMM, A. et al. **A YANG Data Model for Layer 3 Topologies**. Fremont, CA, USA, 2018. (Internet Request for Comments, 8346). Accessed on August 2022. Available from Internet: <<https://www.rfc-editor.org/rfc/rfc8346.txt>>.

CLEMM, A. et al. **A YANG Data Model for Network Topologies**. Fremont, CA, USA, 2018. (Internet Request for Comments, 8345). Accessed on August 2022. Available from Internet: <<https://www.rfc-editor.org/rfc/rfc8345.txt>>.

COHEN-ALMAGOR, R. **Internet history**. [S.l.: s.n.], 2013. 19–39 p. ISBN 9781466629325.

Colômbia. Congreso Nacional. **Ley 1.450 de 2011: por lacual se expideelPlan Nacional de Desarrollo, 2010-2014**. 2011. Accessed on August 2022. Available from Internet: <<https://www.alcaldiabogota.gov.co/sisjur/normas/Norma1.jsp?i=43101>>.

CROWCROFT, J. Net Neutrality: The Technical Side of the Debate ~ A White Paper. **International Journal of Communication**, v. 1, p. 567–579, 2007. ISSN 0146-4833.

CRTC - Canadian Radio-Television and Telecommunications Commission. **Telecom Regulatory Policy CRTC 2009-657: review of the Internet traffic management practices of Internet service providers**. 2009. Accessed on August 2022. Available from Internet: <<https://crtc.gc.ca/eng/archive/2009/2009-657.htm>>.

CZ.NIC. **Jetconf**. 2022. Accessed on August 2022. Available from Internet: <<https://jetconf.readthedocs.io/en/latest/>>.

DARER, A.; FARNAN, O.; WRIGHT, J. Automated discovery of internet censorship by web crawling. **WebSci 2018 - Proceedings of the 10th ACM Conference on Web Science**, p. 195–204, 2018.

Data61. **PyAirports**. 2022. Accessed on August 2022. Available from Internet: <<https://github.com/NICTA/pyairports>>.

DISCHINGER, M. et al. Glasnost : Enabling End Users to Detect Traffic Differentiation. In: **Proceedings of the 7th USENIX conference on Networked systems design and implementation (NSDI'10)**. [S.l.: s.n.], 2010.

DISCHINGER, M. et al. Detecting BitTorrent blocking. In: **Proceedings of the ACM SIGCOMM Internet Measurement Conference, IMC**. [S.l.: s.n.], 2008. ISBN 9781605583341.

Distributed Management Task Force. **Common Information Model**. 2022. Accessed on August 2022. Available from Internet: <<https://www.dmtf.org/standards/cim/>>.

Distributed Management Task Force. **Distributed Management Task Force**. 2022. Accessed on August 2022. Available from Internet: <<https://www.dmtf.org/>>.

EBRAHIMI, N.; SOOFI, E. S.; ZHAO, S. Information measures of Dirichlet distribution with applications. **Applied Stochastic Models in Business and Industry**, v. 27, n. 2, p. 131–150, 2011. ISSN 15241904.

Economics of Cybersecurity Research Group, Delft University of Technology. **PyASN**. 2022. Accessed on August 2022. Available from Internet: <<https://pypi.org/project/pyasn/>>.

EICAR - European Expert Group for IT-Security. **EICAR Anti Malware Testfile**. 2022. Accessed on August 2022. Available from Internet: <<https://www.eicar.org/>>.

EKO, L. Many Spiders, One Worldwide Web: Towards a Typology of Internet Regulation. **Communication Law and Policy**, v. 6, n. 3, p. 445–484, 2001. ISSN 1081-1680.

ENNS, R. et al. **Network Configuration Protocol (NETCONF)**. Fremont, CA, USA, 2011. (Internet Request for Comments, 6241). Accessed on August 2022. Available from Internet: <<https://www.rfc-editor.org/rfc/rfc6241.txt>>.

F5. **K22162026: Geolocation network firewall rule with Europe (EU) or Asia Pacific (AP) Address/Region not matching as expected**. 2022. Accessed on August 2022. Available from Internet: <<https://support.f5.com/csp/article/K22162026>>.

Federal Communications Commission. **FCC-15-24: Protecting and Promoting the Open Internet**. 2015. Accessed on August 2022. Available from Internet: <<https://docs.fcc.gov/public/attachments/FCC-15-24A1.pdf>>.

Federal Communications Commission. **FCC-17-166: Restoring Internet Freedom**. 2018. Accessed on August 2022. Available from Internet: <<https://docs.fcc.gov/public/attachments/FCC-17-166A1.pdf>>.

FILASTÒ, A.; APPELBAUM, J. OONI: Open observatory of network interference. **2nd USENIX Workshop on Free and Open Communications on the Internet, FOCI 2012, co-located with USENIX Security 2012**, 2012.

GARRETT, T.; BONA, L.; DUARTE, E. A Holistic Approach for Locating Traffic Differentiation in the Internet. **Computer Networks**, v. 200, p. 108489, 2021. ISSN 1389-1286.

GARRETT, T. et al. Traffic differentiation on internet of things. In: **2018 IEEE Symposium on Service-Oriented System Engineering (SOSE)**. [S.l.: s.n.], 2018, p. 142–151.

GARRETT, T. et al. Monitoring Network Neutrality: A Survey on Traffic Differentiation Detection. **IEEE Communications Surveys and Tutorials**, 2018.

GARRETT, T. et al. A survey of network neutrality regulations worldwide. **Computer Law & Security Review**, v. 44, p. 105654, 2022. ISSN 0267-3649.

GLADSTONE, J. A. Determining Jurisdiction In Cyberspace: The "Zippo" Test or the 'Effects' Test? In: **Proceedings of the 2003 InSITE Conference**. [S.l.: s.n.], 2003.

Google. **Transparency Report: HTTPS encryption on the web**. 2022. Accessed on August 2022. Available from Internet: <<https://transparencyreport.google.com/https/overview>>.

GREGORI, E.; LUCONI, V.; VECCHIO, A. NeutMon: Studying neutrality in European mobile networks. In: **INFOCOM 2018 - IEEE Conference on Computer Communications Workshops**. [S.l.: s.n.], 2018. ISBN 9781538659793.

GREGORI, E.; LUCONI, V.; VECCHIO, A. Studying forwarding differences in european mobile broadband with a net neutrality perspective. In: **24th European Wireless 2018 "Wireless Futures in the Era of Network Programmability"**, EW 2018. [S.l.: s.n.], 2018. ISBN 9783800745609.

GU, C.; ZHANG, S.; XUE, X. Encrypted internet traffic classification method based on host behavior. **International Journal of Digital Content Technology and its Applications**, v. 5, n. 3, p. 167–174, 2011.

GUPTA, A.; JHA, R. K. A Survey of 5G Network: Architecture and Emerging Technologies. **IEEE Access**, IEEE, v. 3, p. 1206–1232, 2015. ISSN 21693536.

HADLEY, G.; ANDENÆS, M.; FAIRGRIEVE, D. **Judicial Review in International Perspective**. [S.l.]: Springer Netherlands, 2000. (Judicial Review in International Perspective, v. 2). ISBN 9789041113788.

HAHN, R.; WALLSTEN, S. The Economics of Net Neutrality. **Economics of Networks**, 2006.

HAM, J. van der et al. **Network Markup Language Base Schema version 1**. [S.l.], 2013. Accessed on August 2022. Available from Internet: <<https://www.ogf.org/documents/GFD.206.pdf>>.

Info-communications Development Authority of Singapore. **Decision Issued by the Info-communications Development Authority of Singapore - Net Neutrality**. 2011. Accessed on August 2022. Available from Internet: <https://www.imda.gov.sg/~media/imda/files/inner/pcdg/consultations/20101111_netneutrality/netneutralityexplanatorymemo.pdf>.

Internet Engineering Task Force. **Internet Engineering Task Force**. 2022. Accessed on August 2022. Available from Internet: <<https://ietf.org/>>.

IVCHENKO, G.; HONOV, S. On the jaccard similarity test. **Journal of Mathematical Sciences**, Springer, v. 88, n. 6, p. 789–794, 1998.

JAWORSKI, M.; ZIADÉ, T. **Expert Python Programming - Fourth Edition: Master Python by Learning the Best Coding Practices and Advanced Programming Concepts**. [S.l.]: Packt Publishing, 2021. ISBN 9781801071109.

JIMÉNEZ, W. G.; LODDER, A. R. Analyzing approaches to internet jurisdiction based on a model of harbors and the high seas. **International Review of Law, Computers and Technology**, v. 29, n. 2-3, p. 266–282, 2015.

JITSUZUMI, T. Zero-Rating and Net Neutrality in the Mobile Market: The Case of Japan. **SSRN Electronic Journal**, n. 2015, p. 1–24, 2018.

JONATHAN, O.; MISRA, S.; OSAMOR, V. Comparative Analysis of Machine Learning techniques for Network Traffic Classification. **IOP Conference Series: Earth and Environmental Science**, v. 655, n. 1, 2021. ISSN 17551315.

JORDAN, S. Four questions that determine whether traffic management is reasonable. In: **2009 IFIP/IEEE International Symposium on Integrated Network Management, IM 2009**. [S.l.: s.n.], 2009. p. 137–140.

KAKHKI, A. M. et al. Identifying traffic differentiation on cellular data networks. In: **Computer Communication Review**. [S.l.: s.n.], 2015. ISBN 9781450328364. ISSN 19435819.

KAKHKI, A. M. et al. Identifying traffic differentiation in mobile networks. In: **Proceedings of the ACM SIGCOMM Internet Measurement Conference, IMC**. [S.l.: s.n.], 2015. ISBN 9781450338486.

KANUPARTHY, P.; DOVROLIS, C. DiffProbe: Detecting ISP service discrimination. In: **Proceedings - IEEE INFOCOM**. [S.l.: s.n.], 2010. ISBN 9781424458363. ISSN 0743166X.

KANUPARTHY, P.; DOVROLIS, C. ShaperProbe: End-to-end detection of ISP traffic shaping using active methods. In: **Proceedings of the ACM SIGCOMM Internet Measurement Conference, IMC**. [S.l.: s.n.], 2011. ISBN 9781450310130.

KELLER, C. I. Between exception and harmonization: The theoretical debate on internet regulation. **Publicum**, v. 5, n. 1, p. 137–166, 2019. ISSN 24477982.

KNESSET. **Communications Law (Telecommunications and Broadcasts), 5742-1982 - Art. 51c**. 2014. Accessed on August 2022. Available from Internet: <https://www.nevo.co.il/law_html/Law01/032_002.htm#med15>.

KOSLOVSKI, G. P.; PRIMET, P. V. B.; CHARÃO, A. S. VXDL: Virtual resources and interconnection networks description language. In: **Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering**. [S.l.: s.n.], 2009. ISBN 3642020798. ISSN 18678211.

KRÄMER, J.; WIEWIORRA, L.; WEINHARDT, C. Net neutrality: A progress report. **Telecommunications Policy**, ELSEVIER SCI LTD, v. 37, n. 9, p. 794–813, 2013. ISSN 0308-5961.

KREIBICH, C. et al. Netalyzr: illuminating the edge network. **Proceedings of the 10th annual conference on Internet measurement - IMC '10 (2010) 246**, p. 246–259, 2010.

- KUZMANOVIC, A.; KNIGHTLY, E. W. Measuring service in multi-class networks. In: **Proceedings - IEEE INFOCOM**. [S.l.: s.n.], 2001. ISSN 0743166X.
- LESK, M. Son of Carterfone: Network Neutrality or Regulation? **Security & Privacy, IEEE**, IEEE, USA, v. 8, n. 3, p. 77–82, 2010. ISSN 1540-7993.
- LEXA, M. **Useful Facts about the Kullback-Leibler Discrimination Distance**. [S.l.], 2004.
- LI, D. et al. A Novel Framework for Analysis of Global Network Neutrality Based on Packet Loss Rate. In: **Proceedings - 2015 International Conference on Cloud Computing and Big Data, CCBD 2015**. [S.l.: s.n.], 2016. ISBN 9781467383509.
- LI, F. et al. A large-scale analysis of deployed traffic differentiation practices. In: **SIGCOMM 2019 - Proceedings of the 2019 Conference of the ACM Special Interest Group on Data Communication**. [S.l.: s.n.], 2019. ISBN 9781450359566.
- LU, G. et al. POPI: A User-Level Tool for Inferring Router Packet Forwarding Priority. **IEEE/ACM Transactions on Networking**, 2010. ISSN 10636692.
- MAILLÉ, P.; REICHL, P.; TUFFIN, B. Internet governance and economics of network neutrality. In: **Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)**. [S.l.: s.n.], 2012. ISBN 9783642303814. ISSN 03029743.
- Max Planck Institute for Software Systems. **Glasnost**. 2022. Accessed on August 2022. Available from Internet: <<https://github.com/marcelcode/glasnost/>>.
- Maxmind. **GeoLite2 Databases**. 2022. Accessed on August 2022. Available from Internet: <<https://dev.maxmind.com/geoip/geoip2/geolite2/>>.
- MAY, B. E.; CHEN, J. C. V.; WEN, K. W. The differences of regulatory models and internet regulation in the european union and the united states. **Information and Communications Technology Law**, v. 13, n. 3, p. 259–272, 2004. ISSN 14698404.
- MAYER-SCHÖNBERGER, V. The shape of governance: Anaylzing the world of Internet regulation. **Virginia Journal of International Law**, v. 43, n. January, p. 605–674, 2002.
- MCCLOGHRIE, K.; PERKINS, D.; SCHOENWAELDER, J. **Structure of Management Information Version 2 (SMIv2)**. [S.l.], 1999. Accessed on August 2022. Available from Internet: <<http://www.rfc-editor.org/rfc/rfc2578.txt>>.
- Measurement Lab. **The M-Lab Glasnost Data Set, 2016-01-01 up to 2016-12-31**. 2016. Accessed on August 2022. Available from Internet: <<https://measurementlab.net/tests/glasnost>>.
- MEASUREMENT LAB. **The M-Lab Wehe Data Set (2021-01-01 – 2021-02-28)**. 2021. Accessed on August 2022. Available from Internet: <<https://measurementlab.net/tests/wehe>>.
- MEASUREMENT LAB. **MEASUREMENT LAB**. 2022. Accessed on August 2022. Available from Internet: <<https://www.measurementlab.net/>>.

MIJUMBI, R. et al. Network function virtualization: State-of-the-art and research challenges. **IEEE Communications Surveys and Tutorials**, 2016. ISSN 1553877X.

Ministry of Communications. **Telecom Regulatory Authority of India Issues ‘Prohibition of Discriminatory Tariffs for Data Services Regulations, 2016’**. 2016. Accessed on August 2022. Available from Internet: <<http://pib.nic.in/newsite/PrintRelease.aspx?relid=136211>>.

National Conference of State Legislatures. **Net Neutrality 2021 Legislation**. 2021. Accessed on August 2022. Available from Internet: <<https://www.ncsl.org/research/telecommunications-and-information-technology/net-neutrality-2021-legislation.aspx>>.

NAWROCKI, M. et al. Down the black hole: Dismantling operational practices of BGP blackholing at IXPS. **Proceedings of the ACM SIGCOMM Internet Measurement Conference, IMC**, p. 435–448, 2019.

Ofcom. **Monitoring compliance with the net neutrality rules**. 2020. Accessed on August 2022. Available from Internet: <<https://www.ofcom.org.uk/research-and-data/internet-and-on-demand-research/net-neutrality>>.

Open Grid Forum. **Open Grid Forum**. 2020. Accessed on August 2022. Available from Internet: <<https://ogf.org/>>.

PADOVANI, C.; SANTANIELLO, M. Digital constitutionalism: Fundamental rights and power limitation in the Internet eco-system. **International Communication Gazette**, v. 80, n. 4, p. 295–301, 2018.

Presidência da República. **Lei N° 12.965, de 23 de abril de 2014**. 2014. Accessed on August 2022. Available from Internet: <http://www.planalto.gov.br/ccivil_03/_ato2011-2014/2014/lei/l12965.htm>.

Presidência da República. **Decreto N° 8.771, de 11 de maio de 2016**. 2016. Accessed on August 2022. Available from Internet: <http://www.planalto.gov.br/ccivil_03/_ato2015-2018/2016/decreto/d8771.htm>.

PYTHON. **tracemalloc — Trace memory allocations**. 2022. Accessed on August 2022. Available from Internet: <<https://docs.python.org/3/library/tracemalloc.html>>.

RAVAIOLI, R.; BARAKAT, C.; URVOY-KELLER, G. Chkdif: Checking traffic differentiation at internet access. In: **CoNEXT Student 2012 - Proceedings of the ACM Conference on the 2012 CoNEXT Student Workshop**. [S.l.: s.n.], 2012. ISBN 9781450317757.

Réseaux IP Européens. **RIPE Atlas**. 2022. Accessed on August 2022. Available from Internet: <<https://atlas.ripe.net/>>.

Réseaux IP Européens. **RIPE Atlas Cousteau**. 2022. Accessed on August 2022. Available from Internet: <<https://ripe-atlas-cousteau.readthedocs.io/en/latest/>>.

SCHAURICH, V. G.; CARVALHO, M.; GRANVILLE, L. Z. ISPAN: A policy-based ISP Auditor for Network Neutrality violation detection. In: **The 32-nd IEEE International Conference on Advanced Information Networking and Applications (AINA-2018)**. Pedagogical University of Kracow, Poland: [s.n.], 2018.

SCHULZRINNE, H. Network neutrality is about money, not packets. **IEEE Internet Computing**, 2018. ISSN 19410131.

Senado y Cámara de Diputados de la Nación Argentina. **Ley 27.078**. 2014. Accessed on August 2022. Available from Internet: <<http://servicios.infoleg.gob.ar/infolegInternet/anexos/235000-239999/239771/norma.htm>>.

SETENARESKI, L. E. **Fiscalização da neutralidade da rede e seu impacto na evolução da internet**. Thesis (PhD) — Federal University of Paraná, Brazil, 2017.

SHAW, B.; BADHWAR, S.; BIRD, A. **Web Development with Django**. [S.l.]: Packt Publishing, 2021. ISBN 9781839212505.

SMIRNOVA, I. et al. Network Slicing in the Scope of Net Neutrality Rules. In: **Progress in Electromagnetics Research Symposium**. [S.l.: s.n.], 2019. v. 2019-June, p. 1516–1521. ISBN 9781728134031. ISSN 19317360.

SVANTESSON, P. Scope of jurisdiction online and the importance of messaging – lessons from Australia and the EU. **Computer Law and Security Review**, v. 38, 2020.

TARIQ, M. B. et al. Detecting network neutrality violations with causal inference. In: **CoNEXT'09 - Proceedings of the 2009 ACM Conference on Emerging Networking Experiments and Technologies**. [S.l.: s.n.], 2009. ISBN 9781605586366.

TCPDUMP. **Man page of TCPDUMP**. 2022. Accessed on August 2022. Available from Internet: <<https://www.tcphelp.org/manpages/tcpdump.1.html>>.

Telecom Regulatory Authority of India. **Recommendations On Net Neutrality**. 2017. Accessed on August 2022. Available from Internet: <https://www.trai.gov.in/sites/default/files/Recommendations_NN_2017_11_28.pdf>.

UDDIN, M.; STADLER, R. A bottom-up approach to real-time search in large networks and clouds. In: **Proceedings of the NOMS 2016 - 2016 IEEE/IFIP Network Operations and Management Symposium**. [S.l.: s.n.], 2016. p. 985–990.

UDDIN, M.; STADLER, R.; CLEMM, A. A query language for network search. **Integrated Network Management (IM 2013), 2013 IFIP/IEEE International Symposium on**, p. 109–117, 2013.

UDDIN, M. et al. Graph search for cloud network management. In: **IEEE/IFIP NOMS 2014 - IEEE/IFIP Network Operations and Management Symposium: Management in a Software Defined World**. [S.l.: s.n.], 2014.

WANG, F. F. Obstacles and solutions to internet jurisdiction a comparative analysis of the EU and US laws. **Journal of International Commercial Law and Technology**, v. 3, n. 4, p. 233–241, 2008.

WEINSBERG, U.; SOULE, A.; MASSOULIE, L. Inferring traffic shaping and policy parameters using end host measurements. In: **Proceedings - IEEE INFOCOM**. [S.l.: s.n.], 2011. ISBN 9781424499212. ISSN 0743166X.

WU, T. Network Neutrality, Broadband Discrimination. **Cyberspace Law eJournal**, 2003.

WU, T. The Broadband Debate, A User's Guide. **Journal on Telecommunications and High Technology Law**, v. 3, n. 1, p. 69–96, 2004.

ZHANG, Y.; MAO, Z. M.; ZHANG, M. Ascertaining the Reality of Network Neutrality Violation in Backbone ISPs. In: **HotNets**. [S.l.]: Association for Computing Machinery, Inc., 2008.

ZHANG, Y.; MAO, Z. M.; ZHANG, M. Detecting traffic differentiation in backbone ISPs with NetPolice. In: **Proceedings of the ACM SIGCOMM Internet Measurement Conference, IMC**. [S.l.: s.n.], 2009. ISBN 9781605587707.

ZHANG, Z.; MARA, O.; ARGYRAKI, K. Network neutrality inference. In: **SIGCOMM 2014 - Proceedings of the 2014 ACM Conference on Special Interest Group on Data Communication**. [S.l.: s.n.], 2014. p. 63–74.

ZIPROXY. **Ziproxy**. 2022. Accessed on August 2022. Available from Internet: <<http://ziproxy.sourceforge.net/>>.

APPENDIX A — RESUMO EXPANDIDO

Julgando Diferenciação de Tráfego como Violação da Neutralidade da Rede de acordo com a Regulação da Internet

Neutralidade da Rede (NR) é um princípio que estabelece que o tráfego de aplicações e serviços deve ser tratado igualmente e não deve ser afetado por interferências, degradações ou interrupções arbitrárias. Apesar deste senso comum sobre o princípio da NR, ele tem várias definições espalhadas pela literatura acadêmica (GARRETT et al., 2018b), que diferem principalmente no que constitui o nível de igualdade adequado para considerar uma rede neutra. As definições de NR também podem ser incluídas na regulamentação que controla as atividades na Internet. No entanto, essas instruções são definidas por reguladores cujos atos são válidos apenas dentro de uma área geográfica chamada de jurisdição. Desse modo, tanto a academia quanto a regulamentação fornecem múltiplas e heterogêneas definições de NR.

Esta tese defende que a regulamentação deve ser usada como guia para construir soluções de detecção de violações da NR. Neste caso, uma violação da NR seria a adoção pelos ISPs de práticas de gerenciamento de tráfego que foram proibidas pelos reguladores, em vez de uma violação de uma definição acadêmica. Desta forma, as soluções podem fornecer informações úteis para os usuários que podem ser utilizadas para apoiar reclamações contra práticas de gestão de tráfego ilegais.

No entanto, as soluções do estado-da-arte ou adotam uma definição acadêmica rigorosa (por exemplo, todo o tráfego deve ser tratado igualmente) ou adotam as definições regulatórias de uma única jurisdição, o que pode não ser realista ou não considera que várias jurisdições podem ser percorridas em um caminho de rede fim-a-fim, respectivamente. Como uma Diferenciação de Tráfego (DT) pode ser implantada em qualquer lugar em um caminho de rede fim-a-fim, as soluções deveriam considerar as múltiplas definições de NR que podem ser encontradas em diferentes partes da Internet, pois uma mesma prática de gestão de tráfego pode ser considerada uma violação da NR em determinada jurisdição e ser permitida em outra.

Nesta tese, é proposta uma etapa de Avaliação da Regulamentação a ser executada após a detecção da DT realizada por soluções. Esta etapa deve considerar todas as definições de NR que podem ser encontradas em um caminho de rede fim-a-fim e sinalizar violações da NR quando estas definições são violadas. É proposto um serviço para realizar a etapa de Avaliação da Regulamentação para soluções de detecção de DT,

visto que seria difícil que todas as soluções tivessem que implementar esta etapa.

Como esta tese propõe uma nova maneira de sinalizar violações da NR, foi realizada uma análise de impacto, em que as DTs detectadas pelo Glasnost (DISCHINGER et al., 2010) foram analisadas conforme as definições presentes na regulamentação. Entre outros resultados, esta análise mostrou que, em certas circunstâncias, de 39% a 48% das DTs detectadas não são violações da NR quando a regulamentação é considerada, expondo que o aspecto da regulamentação da Internet não deve ser ignorado.

Com base em requisitos identificados durante a análise de impacto, foram propostos modelos de informação para representar as informações necessárias para sinalizar violações da NR utilizando definições da regulamentação. Estes modelos são organizados em 3 áreas: topologia de rede, diferenciações de tráfego e regulamentação da Internet. Foi conduzida uma investigação para avaliar modelos de dados que poderiam representar estes modelos de informação. A investigação apontou que os modelos baseados em YANG seriam os mais adequados pela existência de modelos para representar topologias de rede e pela facilidade de estender os modelos existentes para representar outras informações.

Uma arquitetura conceitual para uma solução, chamada de JurisNN, capaz de sinalizar violações da NR considerando a regulamentação foi proposta. Esta arquitetura possui 3 módulos diretamente relacionados aos modelos de dados propostos, topologia, ocorrências e regulamentação; 2 módulos para interação: API e interface do usuário; e um módulo de julgamento. O módulo de julgamento é auxiliado por *algoritmos de julgamento*, que implementam maneiras distintas de estabelecer a jurisdição de casos na Internet. Durante a investigação conduzida nesta tese, foram identificadas 3 maneiras distintas: lugar de implantação da DT, teste de alvo e teste alemão.

Um protótipo de prova de conceito (PoC) para o JurisNN foi desenvolvido seguindo a arquitetura conceitual. Este protótipo foi avaliado utilizando informações sobre DTs detectadas pelo Wehe (LI et al., 2019). Como o Wehe não coleta informação sobre os caminhos de rede percorridos pelos testes, o seu conjunto de dados foi complementado através da coleta de caminhos de rede entre o Sistema Autônomo do cliente e o endereço IP do servidor do teste que detectou DT. Estes caminhos de rede foram coletados utilizando a plataforma RIPE Atlas (Réseaux IP Européens, 2022a). As informações sobre as DTs detectadas pelo Wehe e os caminhos de rede coletados foram submetidos para o JurisNN para julgamento das DTs como violações da NR conforme a regulamentação. Como não há informação sobre onde neste caminho de rede a DT é suspeita de ser implantada (zona de ocorrência), o caminho completo foi considerado como a zona de ocorrência.

Os algoritmos de julgamento alcançaram resultados distintos. Para o teste de alvo, as informações disponíveis foram suficientes para julgar as DTs como violações da NR conforme a regulamentação. Este teste baseia-se nas informações de jurisdições dos dispositivos finais (usuário final ou servidor de aplicação), cujas informações podem ser facilmente obtidas ao analisar os endereços de origem e destino do teste. Por outro lado, os veredictos alcançados foram inconclusivos para os algoritmos de julgamento lugar de implantação da DT (18 dos 24 cenários) e teste alemão (único cenário avaliado). Estes algoritmos dependem de informações do caminho de rede percorrido. Ao considerar todo o caminho como zona de ocorrência, não foi possível concluir se a DT era uma violação da NR ou não. Portanto, informações mais precisas sobre onde a DT é suspeita de ser implantada são necessárias. Desta forma, pode-se restringir quais jurisdições estão sendo afetadas e por consequência pode-se aplicar a regulamentação correta.

Os resultados obtidos pelos algoritmos de julgamento foram afetados pela falta de informações sobre os caminhos de rede percorridos e sobre onde está a zona de ocorrência no caminho da rede. No entanto, existem propostas para realizar a coleta de tais informações. Espera-se que os resultados obtidos na análise de impacto e na avaliação dos algoritmos de julgamento do JurisNN encorajem os desenvolvedores de soluções de detecção de DT a coletar essas informações e submetê-las ao serviço de Avaliação da Regulamentação. Desta forma, as soluções podem complementar os resultados fornecidos aos usuários com os veredictos sobre suas DTs detectadas de acordo com as regulamentações.

APPENDIX B — PROPOSED YANG MODELS

YANG data model for NN regulation

```
1 module nn-regulation {
2
3   yang-version 1.1;
4   namespace "ufrgs:inf:computernetworks:nn-regulation";
5   prefix "nnr";
6
7   // import statements here: e.g.,
8   // import ietf-yang-types { prefix yang; }
9   // import ietf-inet-types { prefix inet; }
10
11  import ietf-yang-types {
12    prefix "yang";
13    revision "2013-07-15";
14    reference "RFC 6991: Common YANG Data Types";
15  }
16
17  description
18    "";
19
20  revision 2021-04-18 {
21    description
22      "Initial revision.";
23  }
24
25  // extension statements
26
27  // feature statements
28
29  // identity statements
30
31  // typedef statements
32
33  typedef location-id {
34    type string {
35      pattern "[A-Z]{2}(-[A-Z]{2})?";
36    }
37    description
38      "The ISO Country Code of a location that may be a country or a state.";
39  }
40
41  typedef regulation-id {
42    type yang:counter32;
43    description
44      "The identification of a normative.";
45  }
46
47  typedef instruction-id {
48    type yang:counter32;
49    description
50      "The identification of a instruction.";
51  }
52
53  typedef interpretation-id {
54    type yang:counter32;
55    description
56      "The identification of a interpretation.";
57  }
58
59  typedef instruction-type {
60    type enumeration {
61      enum "prohibition" {
62        description
63          "The instructions that prohibit differentiations.";
64      }
65      enum "permission" {
```

```

66     description
67     "The instructions that allow differentiations.";
68 }
69 }
70 }
71
72 typedef differentiation {
73     type enumeration {
74         enum "blocking" {
75             description "Blocking.";
76         }
77         enum "prioritization" {
78             description "Prioritization.";
79         }
80         enum "degradation" {
81             description "Degradation.";
82         }
83         enum "throttling" {
84             description "Throttling.";
85         }
86     }
87     description
88     "Traffic Differentiation Types.";
89 }
90
91 // grouping statements
92
93 grouping differentiated-entity {
94     container differentiated-entity {
95         choice differentiated-entity-type {
96             case user {
97                 uses user-attributes;
98             }
99             case host {
100                 uses host-attributes;
101             }
102             case network {
103                 uses network-attributes;
104             }
105             case service {
106                 uses service-attributes;
107             }
108             case application {
109                 uses application-attributes;
110             }
111         }
112         description
113         "Entities whose Traffic is being Differentiated.";
114     }
115 }
116
117 grouping user-attributes {
118     container user {
119         leaf "name" {
120             type string;
121         }
122         leaf "regulation-class" {
123             type string;
124         }
125         leaf "located-at" {
126             type "nrr:location-id";
127             description
128             "The identifier of the location.";
129         }
130     }
131 }
132
133 grouping host-attributes {
134     container host {
135         leaf "name" {

```

```
136     type string;
137   }
138   leaf "regulation-class" {
139     type string;
140   }
141   leaf "located-at" {
142     type "nnr:location-id";
143     description
144       "The identifier of the location.";
145   }
146 }
147 }
148
149 grouping network-attributes {
150   container network {
151     leaf "name" {
152       type string;
153     }
154     leaf "regulation-class" {
155       type string;
156     }
157     leaf "located-at" {
158       type "nnr:location-id";
159       description
160         "The identifier of the location.";
161     }
162   }
163 }
164
165 grouping service-attributes {
166   container service {
167     leaf "name" {
168       type string;
169     }
170     leaf "regulation-class" {
171       type string;
172     }
173     leaf "located-at" {
174       type "nnr:location-id";
175       description
176         "The identifier of the location.";
177     }
178   }
179 }
180
181 grouping application-attributes {
182   container application {
183     leaf "name" {
184       type string;
185     }
186     leaf "regulation-class" {
187       type string;
188     }
189     leaf "located-at" {
190       type "nnr:location-id";
191       description
192         "The identifier of the location.";
193     }
194   }
195 }
196
197 grouping jurisdiction-list {
198   description "The information about a jurisdiction.";
199   leaf-list jurisdiction {
200     type location-id;
201   }
202 }
203
204 grouping instruction-list {
205   description "The information of a regulation instruction.";
```



```

206 list instructions {
207     key "instruction-id";
208     leaf instruction-id {
209         type instruction-id;
210     }
211     leaf instruction-type {
212         type "instruction-type";
213     }
214     leaf differentiation {
215         type differentiation;
216     }
217     uses differentiated-entity;
218 }
219 }
220
221 grouping interpretation-list {
222     description "The list of Interpretations.";
223     list interpretations {
224         key "interpretation-id";
225         leaf interpretation-id {
226             type interpretation-id;
227         }
228         uses differentiated-entity;
229     }
230 }
231
232 grouping regulation-list {
233     description "The list of regulations.";
234     list regulation {
235         key "regulation-id";
236         leaf regulation-id {
237             type regulation-id;
238         }
239         leaf name {
240             type string;
241         }
242         leaf parent-regulation-ref {
243             type leafref {
244                 path "/nnr:regulations/nnr:regulation/nnr:regulation-id";
245             }
246             description
247                 "Reference to the parent regulation or itself when root regulation.";
248             mandatory true;
249         }
250         leaf begin-validity {
251             type yang:date-and-time;
252             description
253                 "The begin of validity of the regulation.";
254             mandatory true;
255         }
256         leaf end-validity {
257             type yang:date-and-time;
258             description
259                 "The end of validity of the regulation.";
260         }
261         uses jurisdiction-list;
262         uses instruction-list;
263         uses interpretation-list;
264     }
265 }
266
267 // data definition statements
268 container regulations {
269     uses regulation-list;
270 }
271
272 // augment statements
273
274 // rpc statements
275

```

```

276 // notification statements
277
278 // DO NOT put deviation statements in a published module
279
280 }

```

YANG data model for TD Occurrences

```

1 module tdo {
2
3   yang-version 1.1;
4   namespace "ufrgs:inf:computernetworks:traffic-differentiation-occurrence";
5   prefix "tdo";
6
7   import ietf-network-topology {
8     prefix "nt";
9     reference
10      "RFC 8345";
11   }
12
13   import ietf-network {
14     prefix "nw";
15     reference
16      "draft-ietf-i2rs-yang-network-topo-20
17      NOTE TO RFC EDITOR:
18      (1) Please replace above reference to
19      draft-ietf-i2rs-yang-network-topo-20 with RFC
20      number when published (i.e. RFC xxxx).
21      (2) Please replace the date in the revision statement with the
22      date of publication when published.";
23   }
24
25   import ietf-yang-types {
26     prefix "yang";
27     reference "RFC 6991: Common YANG Data Types";
28   }
29
30   import nn-regulation {
31     prefix "nnr";
32   }
33
34   description
35     "Communication Graph input for NN Jurisdiction Assessment.";
36
37   revision "2022-09-04" {
38     description
39       "initial revision";
40   }
41
42   // extension statements
43
44   // feature statements
45
46   // identity statements
47
48   // typedef statements
49
50   typedef occurrence-id {
51     type yang:counter32;
52     description
53       "The identification of a TD occurrence.";
54   }
55
56   typedef occurrence-zone-id {
57     type yang:counter32;
58     description
59       "The identification of a occurrence zone.";
60   }
61
62   // grouping statements

```

```

63
64 grouping communication-graph {
65     description
66         "All the nodes and links that are being used in the communication.";
67
68     container communication-graph {
69         leaf network-ref {
70             type leafref {
71                 path "/nw:networks/nw:network/nw:network-id";
72             }
73             description
74                 "Reference to the network-id where the communication graph is inserted.";
75             mandatory yes;
76         }
77         list nodes {
78             key "node-ref";
79             leaf node-ref {
80                 type leafref {
81                     path "/nw:networks/nw:network[nw:network-id=current()../../tdo:network-ref]/nw:node/nw:node-id";
82                 }
83                 description
84                     "References the nodes within the network that participates in the communication.";
85             }
86             leaf is-source {
87                 type boolean;
88                 default false;
89             }
90             leaf is-destination {
91                 type boolean;
92                 default false;
93             }
94         }
95         list links {
96             key "link-ref";
97             leaf link-ref {
98                 type leafref {
99                     path "/nw:networks/nw:network[nw:network-id=current()../../tdo:network-ref]/nt:link/nt:link-id";
100                }
101                description
102                    "References the links within the network that participates in the communication.";
103            }
104        }
105    }
106 }
107
108 grouping occurrence-zone-list {
109     description
110         "The zones within the topology where Traffic Differentiations
111         are happening (occurrence)";
112
113     list occurrence-zone {
114         key "occurrence-zone-id";
115         leaf occurrence-zone-id {
116             type "occurrence-zone-id";
117             description
118                 "The id of occurrence zone.";
119         }
120         leaf differentiation {
121             type "nnr:differentiation";
122             description
123                 "The Traffic Differentiation that is happening within the
124                 occurrence zone.";
125         }
126     }
127     uses "nnr:differentiated-entity";
128
129     list nodes {
130         key "node-ref";
131         leaf node-ref {
132             type leafref {

```

```

133         path "/tdo:occurrences/tdo:occurrence[tdo:occurrence-id=current()../../../../tdo:occurrence-id]/tdo:
           communication-graph/tdo:nodes/tdo:node-ref";
134     }
135     description
136         "References the nodes within the communication graph that are in the occurrence zone.";
137 }
138 }
139 list links {
140     key "link-ref";
141     leaf link-ref {
142         type leafref {
143             path "/tdo:occurrences/tdo:occurrence[tdo:occurrence-id=current()../../../../tdo:occurrence-id]/tdo:
               communication-graph/tdo:links/tdo:link-ref";
144         }
145         description
146             "References the links within the communication graph that are in the occurrence zone.";
147     }
148 }
149 description
150     "The occurrence zone definition.";
151 }
152 }
153
154 grouping occurrences {
155
156     container occurrences {
157         list occurrence {
158             key "occurrence-id";
159             leaf occurrence-id {
160                 type "occurrence-id";
161             }
162             leaf "occurrence-date" {
163                 type "yang:date-and-time";
164                 description
165                     "The date of the occurrence.";
166                 mandatory true;
167             }
168             uses communication-graph;
169             uses occurrence-zone-list;
170         }
171     }
172 }
173
174 // data definition statements
175 uses occurrences;
176
177 // augment statements
178 augment
179     "/nw:networks/nw:network/nw:node" {
180     leaf located-at {
181         type "nr:location-id";
182         description
183             "The identifier of the location.";
184     }
185 }
186
187 // rpc statements
188 rpc analyze {
189     description
190         "Analyze whether the traffic differentiation occurrence is an NN violation according to the regulation.";
191     input {
192         leaf occurrence-id {
193             type "occurrence-id";
194             mandatory true;
195             description
196                 "The id of the occurrence.";
197         }
198         leaf occurrence-zone-id {
199             type "occurrence-zone-id";
200             mandatory true;

```

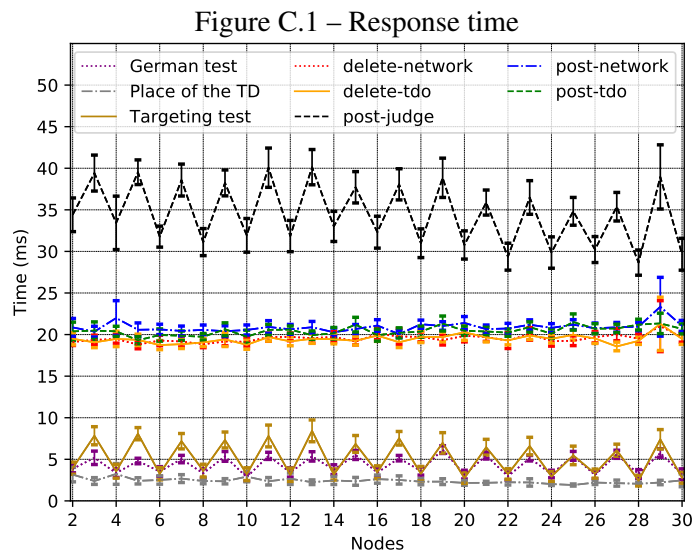
```
201     description
202     "The id of occurrence zone.";
203 }
204 }
205 output {
206     container results {
207         list result {
208             key "id";
209             leaf id {
210                 type uint32;
211                 mandatory true;
212             }
213             leaf name {
214                 type string;
215             }
216             container admissibility {
217                 leaf rationale {
218                     type string;
219                 }
220                 leaf admissible {
221                     type boolean;
222                 }
223             }
224             container judgment {
225                 leaf metric-description {
226                     type string;
227                 }
228                 leaf metric-value {
229                     type decimal64 {
230                         fraction-digits 2;
231                     }
232                 }
233                 leaf rationale {
234                     type string;
235                 }
236                 leaf is-violation{
237                     type boolean;
238                 }
239             }
240         }
241     }
242 }
243 }
244
245 // notification statements
246
247 // DO NOT put deviation statements in a published module
248
249 }
```

APPENDIX C — PERFORMANCE CHARACTERIZATION

In this section, the quantitative aspects of the JurisNN prototype are evaluated, which includes the evaluation of response time of each type of request and of each judgment algorithm, and the usage of CPU and memory. The evaluation scenario is composed of one VirtualBox 6.1 virtual machine with one virtual CPU and 1 GiB of memory running Ubuntu 20.04.2 LTS hosted on a Windows 10 machine based on an Intel(R) Core(TM) i5-10210U CPU@1.60GHz and with 8 GiB of memory. The virtual machine runs Docker version 20.10.6 that runs one container for each JurisNN main component (frontend and backend). The JurisNN is deployed using the docker-compose environment that performs the configuration of the components, including their inter-communication.

A script was developed to perform requests to JurisNN. The TD submitted is throttling of the application named "Netflix." The network topology, communication graph, and occurrence zone are composed of several nodes, varying from 2 to 30, which is the usual limit in the traceroute command. Indeed, the network path collection presented in Section 6.2 collected paths with ≈ 8.9 nodes on average. These nodes have the located-at field adjusted alternating USA and DE: USA in the odd nodes and DE in the even nodes. In this way, the TD is admitted to be judged by the Targeting test algorithm (because at least one endpoint is within the USA) and by the German test algorithm (because or one endpoint is located in DE when the number of nodes is even, or at least one node in the network path is within DE, when the number of nodes is odd). For instance, for two nodes, they are located in the USA and DE, respectively. For three nodes, they are located in the USA, DE, and USA, respectively. The Place of the TD deployment algorithm does not have admissibility conditions. Thus, it admits to judging all submitted TDs independently of the node location. Each TD was submitted to JurisNN 30 times, and the graphs present confidence intervals with 95% of confidence level.

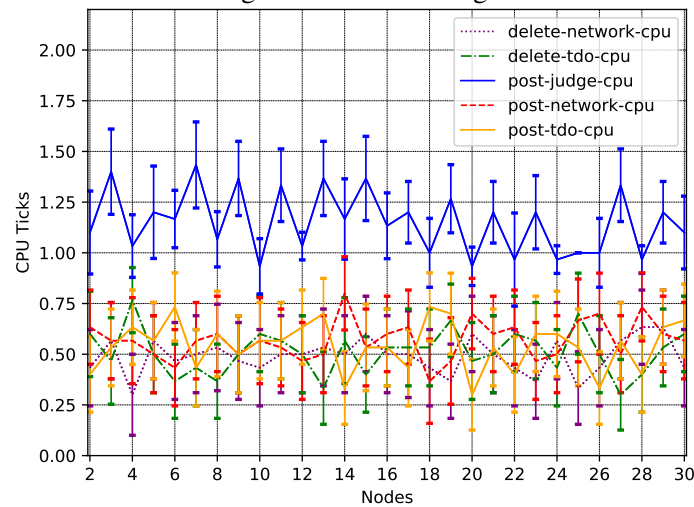
The first quantitative aspect evaluated is the response time of each type of request and judgment algorithm. The response time was calculated by the script that controls the evaluation annotating the time before the request and immediately after the request (for requests) and in a similar way but within the JurisNN Backend code for judgment algorithms. Therefore, for requests, the response time is related to the whole communication, including the internal network stack (the evaluation was conducted within the same virtual machine) and only the internal time for judgment algorithms. The results are presented in Figure C.1. The response time of requests related to submission and deletion of data in



JurisNN, post-network and delete-network (related to ietf-network model), and post-tdo, and delete-tdo (related to tdo model) have a similar value ($\approx 20\text{ms}$) with the confidence intervals overlapping indicating that all requests have the same average response time.

The response time of requests related to the judgment (post-judge) follows a pattern in which when the number of nodes is odd, the value is slightly higher ($\approx 38\text{ms}$) than when the number of nodes is even ($\approx 33\text{ms}$). This pattern is due to how the occurrence zones were artificially crafted alternating nodes within USA and DE and by the regulatory instructions for these jurisdictions, which affects the German test and Targeting test algorithms. For these two algorithms, the response times follow the same pattern. When the number of nodes is odd, the network path starts and ends in the USA (*e.g.*, USA, DE, USA). In this situation, the Targeting algorithm admits the TD in the first condition, checking the source jurisdiction that is within the USA. The judgment step evaluates whether the TD is considered a violation in the source location. However, as in the USA, there is no regulation in place; the algorithm needs to evaluate the instructions of the top-level regulation (Top level permissive regulation). As the TD is not considered a violation in the source location, the algorithm evaluates the regulation in the destination location, which is the same process as the source, thus requiring four regulation assessments (2 for USA and 2 for the Top-level). The German test in such conditions evaluates the source and destination locations, which ends up not to admit the TD because its endpoints are in the USA (conditions 1 and 2). Then, it evaluates the network path (condition 3) admitting the TD because there are nodes within DE. The judgment step checks whether the source and destination are within DE, which is not the case. Then it counts the nodes within each

Figure C.2 – CPU usage

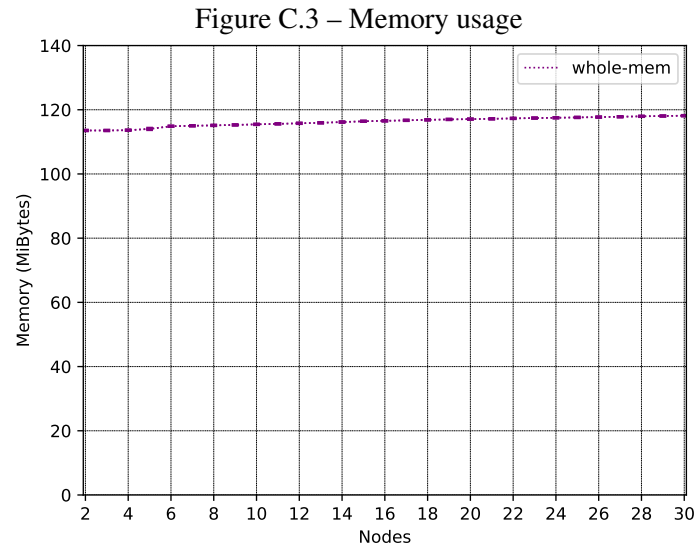


Source: the author

jurisdiction and evaluates only the regulatory instructions of DE (the BEREC regulation that prohibits throttling of applications), thus, requiring only one regulation assessment.

When the number of nodes is even, the network path starts in the USA and ends in DE. In this situation, the difference for the Targeting test is that the judgment step evaluates only the source location, which is within USA, requiring two regulation assessments. For the German test, the difference is that the judgment step only needs to evaluate the destination location, which is within DE, requiring one regulation assessment. Therefore, the number of regulation assessments impacts the algorithms' response time. The Place of TD deployment is not affected by request characteristics because the admissibility control admits all TDs, and the judgment counts the nodes in each jurisdiction and performs the judgment for each jurisdiction, thus, having the same average response time independently of the number of nodes. The Place of TD deployment is faster because its admissibility control almost does not require processing time. The judgment step performs three regulation assessments: USA and Top-level permissive regulation, for USA location, and BEREC regulation for DE location. Despite such observations about the request characteristics impacting the response time, the judgment algorithms (post-judge) requires less than 45ms to perform the Regulation Assessment, which is acceptable. In addition, the PoC was developed using components based on Python, which could be substituted if the response time must be improved.

The second quantitative aspect evaluated is the CPU usage of each type of request. The CPU usage was calculated by the script that controls the evaluation annotating the number of CPU ticks (sum of the kernel and user times, fields 14 and 15 from the `/proc/[pid]/stat`) consumed by the JetConf process (that executes the JurisNN Backend)



Source: the author

before and after each request. The results are presented in Figure C.2. The results are consistent with the response time results: the requests related to submission of network and TD occurrence data having a similar CPU usage and the requests related to the judgment having a higher CPU usage when the number of nodes is odd and lower CPU usage when the number of nodes is even.

The third quantitative aspect evaluated is the memory usage of the JurisNN Backend. The memory usage was calculated by the script that controls the evaluation, annotating the memory (virtual memory, field 23 from `/proc/[pid]/stat`) allocated by the JetConf process (that executes the JurisNN Backend) before and after each whole round of requests (network and TD occurrence data submission, judgment algorithms execution, and network and TD occurrence data deletion). The results are presented in Figure C.3. The results show that the solution achieves a steady memory allocation near to 120 MiB. However, initial evaluations showed a behavior consistent with a memory leak achieving 145 MiBytes at the end of the experiment (30 nodes).

Further evaluation performing 30 requests every 10 seconds for one hour showed that the system allocated 700 MiBytes with a monotonic increase of memory allocation. Inspecting the memory allocation using the `tracemalloc` library (PYTHON, 2022) in Python indicated the JetConf modules responsible for the memory allocation. The JetConf code was inspected finding that it keeps historical data, as presented in Figure C.4, which may be interesting for RESTCONF because it is deemed to handle configuration information. However, JurisNN operates as a web service, in which this feature stores all submitted requests. Such feature was disabled commenting the line 3, and the extended

Figure C.4 – jetconf/data.py code snippet

```
1 # Set a new Instance node as data root, store old root to archive
2 def set_data_root(self, new_root: InstanceNode):
3     #self._data_history.append(self._data)
4     self._data = new_root
```

Source: the author

memory evaluation (one hour experiment) was repeated, showing that JetConf achieved memory allocation stability near to 121 MiBytes, which is consistent with the results in Figure C.3. This memory utilization is considered acceptable for JurisNN.

The JurisNN prototype performance characteristics demonstrated that the solution is feasible, having low response time, CPU and memory usage. The results are considered acceptable for the JurisNN requirements.

APPENDIX D — SCIENTIFIC PRODUCTION

D.1 Published Papers

1. MARCIO BARBOSA DE CARVALHO, Vinícius Garcez Schaurich, Lisandro Zambenedetti Granville. **Considering Jurisdiction When Assessing End-to-End Network Neutrality**. IEEE Internet Computing, vol. 22, no. 6, pp. 27-34, 1 Nov.-Dec. 2018.
 - **Status:** published.
 - **Qualis:** A1.
2. MARCIO BARBOSA DE CARVALHO, Vitor Cunha, Eduardo da Silva, Daniel Corujo; Joao Paulo Barraca, Rui Luís Aguiar, Lisandro Zambenedetti Granville. **Quantifying the Influence of Regulatory Instructions over the Detection of Network Neutrality Violations**. 2020 IFIP Networking Conference (Networking), 2020, pp. 334-342.
 - **Status:** published.
 - **Qualis:** A2.
3. MARCIO BARBOSA DE CARVALHO, Lisandro Zambenedetti Granville. **JurisNN: Judging Traffic Differentiations as Network Neutrality Violation in Accordance to the Regulation**. Elsevier Computer Networks, vol. 205, 2022.
 - **Status:** published.
 - **Qualis:** A1.

D.2 Collaborations

1. Vinícius Garcez Schaurich, MARCIO BARBOSA DE CARVALHO, Lisandro Zambenedetti Granville. **ISPANN: A Policy-Based ISP Auditor for Network Neutrality Violation Detection**. 2018 IEEE 32nd International Conference on Advanced Information Networking and Applications (AINA), 2018, pp. 1081-1088.
 - **Status:** published.
 - **Qualis:** A2.

2. Vitor Cunha, MARCIO BARBOSA DE CARVALHO, Daniel Corujo, Joao Paulo Barraca, Diogo Gomes, Alberto Egon Schaeffer-Filho, Carlos Raniery Paula dos Santos, Lisandro Zambenedetti Granville, Rui Luís Aguiar. **An SFC-enabled approach for processing SSL/TLS encrypted traffic in Future Enterprise Networks.** 2018 IEEE Symposium on Computers and Communications (ISCC), 2018, pp. 01013-01019.
 - **Status:** published.
 - **Qualis:** A2.
3. Eduardo Sousa, Vitor Cunha, MARCIO BARBOSA DE CARVALHO, Daniel Corujo, Joao Paulo Barraca, Diogo Gomes, Alberto Egon Schaeffer-Filho, Carlos Raniery Paula dos Santos, Lisandro Zambenedetti Granville, Rui Luís Aguiar. **Orchestrating an SFC-enabled SSL/TLS traffic processing architecture using MANO.** 2018 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN), 2018, pp. 1-7.
 - **Status:** published.
 - **Qualis:** B2.
4. Vitor Cunha, Eduardo da Silva, MARCIO BARBOSA DE CARVALHO, Daniel Corujo, Joao Paulo Barraca, Diogo Gomes, Alberto Egon Schaeffer-Filho, Carlos Raniery Paula dos Santos, Lisandro Zambenedetti Granville, Rui Luís Aguiar. **A Network Service for Preventing Data Leakage from IoT Cloud-assisted Equipment.** 2019 IEEE Symposium on Computers and Communications (ISCC), 2019, pp. 1-7.
 - **Status:** published.
 - **Qualis:** A2.
5. Vitor Cunha, Eduardo da Silva, MARCIO BARBOSA DE CARVALHO, Daniel Corujo, Joao Paulo Barraca, Diogo Gomes, Lisandro Zambenedetti Granville, Rui Luís Aguiar. **Network slicing security: Challenges and directions.** Wiley Internet Technology Letters, vol. 2, no. 5, pp. e125, 2019.
 - **Status:** published.
 - **Qualis:** not classified.