

CZECH TECHNICAL UNIVERSITY IN PRAGUE
FACULTY OF ELECTRICAL ENGINEERING
DEPARTMENT OF COMPUTER SCIENCE



Diploma Thesis

Trust Models on Adversarial Distributed Security Agents

Bc. Dita Hollmannová

Supervisor: Ing. Sebastián García, Ph.D.

Study programme: Open Informatics

Specialisation: Cyber Security

August 2020

I. Personal and study details

Student's name: **Hollmannová Dita** Personal ID number: **456871**
Faculty / Institute: **Faculty of Electrical Engineering**
Department / Institute: **Department of Computer Science**
Study program: **Open Informatics**
Specialisation: **Cyber Security**

II. Master's thesis details

Master's thesis title in English:

Trust models on adversarial distributed security agents

Master's thesis title in Czech:

Modely důvěry pro distribuované sdílení dat v nedůvěryhodné síti

Guidelines:

The goal of this work is to propose a protocol for sharing data in a decentralized network of peers, where each node gains reputation for their actions. Information from nodes with low reputation may be discarded, while nodes with high reputation will be heard. This serves as a protection, because malicious nodes would first have to gain trust of the network before they could affect it.

There are multiple approaches to compute reputation [2, 3], but they rely mostly on adherence to the protocol, uptime and other simple features. The trust model used by the Sality botnet [4] simply measures how many „good“ interactions a node had with it's neighbor. There are numerous attacks that an adversary can use to gain trust of the network [1, 5]. In this thesis, the trust model will not only use data from the protocol itself, but also network monitoring and statistics provided by SLIPS [6]. The student will analyze different trust models and options to attack them. A new trust model that uses data from SLIPS will be proposed, and its performance will be evaluated. Finally, the model will be implemented as a module inside SLIPS, and will enable sharing said network data with other nodes running SLIPS.

Bibliography / sources:

- [1] Eleni Koutrouli, Aphrodite Tsalgatidou: Taxonomy of attacks and defense mechanisms in P2P reputation systems - Lessons for reputation system designers, 2010
- [2] Jingpei Wang, Jie Liu: The Comparison of Distributed P2P Trust Models Based on Quantitative Parameters in the File Downloading Scenarios, 2016
- [3] Tigist Abera, Ferdinand Brasser, Lachlan J. Gunn, David Koisser, Ahmad-Reza Sadegh: SADAN: Scalable Adversary Detection in Autonomous Networks, 2017
- [4] Nicolas Falliere: Sality: Story of a Peer-to-Peer Viral Network, 2011
- [5] Emmanouil Vasilomanolakis, Jan Helge Wolf, Leon Böck, Shankar Karuppayah, Max Mühlhäuser: I Trust my Zombies: A Trust-enabled Botnet, 2017
- [6] SLIPS: Stratosphere labs intrusion prevention system, <https://github.com/stratosphereips/StratosphereLinuxIPS/>

Name and workplace of master's thesis supervisor:

Ing. Sebastián García, Ph.D., Artificial Intelligence Center, FEE

Name and workplace of second master's thesis supervisor or consultant:

Date of master's thesis assignment: **05.02.2020** Deadline for master's thesis submission: **14.08.2020**

Assignment valid until: **30.09.2021**

Ing. Sebastián García, Ph.D.
Supervisor's signature

Head of department's signature

prof. Mgr. Petr Páta, Ph.D.
Dean's signature

III. Assignment receipt

The student acknowledges that the master's thesis is an individual work. The student must produce her thesis without the assistance of others, with the exception of provided consultations. Within the master's thesis, the author must state the names of consultants and include a list of references.

Date of assignment receipt

Student's signature

Declaration

I declare that I elaborated this thesis on my own and that I mentioned all the information sources and literature that have been used in accordance with the Guideline for adhering to ethical principles in the course of elaborating an academic final thesis.

In Prague on August 14, 2020

.....

Acknowledgements

Writing this thesis has been a long journey, and I am very grateful to the people who have supported me along the way. A huge *thank you* goes to my supervisor, Ing. Sebastián García, Ph.D., for always taking the time to discuss any issues and new ideas with me, and for his valuable feedback. I also wouldn't be able to finish the thesis without the encouragement, patience, and loving support of Tomáš Palider and other members of my family.

I am also thankful to all my friends at the Stratosphere Laboratory, for proofreading this thesis, helping me when I hit a tough patch, sharing their knowledge, and in general, for always being there for me. I am forever grateful for the opportunity to be part of the Stratosphere team.

Abstract

Intrusion Prevention Systems (IPSs) are used to detect and block attacks in the network, and it is not uncommon to have multiple devices with an IPS present in a network. However, as far as we know in the current state-of-the-art IPSs, none of them help each other by sharing the attacks they have encountered.

In this thesis, we propose, design and implement a complete P2P protocol, called Dovecot, to share detection information between IPSs. Individual IPSs become peers in the Dovecot P2P network, and they share their detections and use data from others to improve their blocking decisions. Although sharing data is important, the main problem is how much that data can be trusted.

To protect the peers from adversaries that are trying to manipulate their decisions, we design the Ω -Trust trust model for Dovecot, that computes the trust of a peer based on its network behavior. Identifying malicious peers is important, because otherwise adversaries could corrupt the decisions of the IPS. Contrary to other solutions, our trust computation can not be externally manipulated by the remote peers. The trust on each peer is used to weight the data they send and compute the local Ω -score and Ω -confidence of each IP address. These values represent what the network believes about the IP address and how trustworthy the information is, and they are later used by the local IPS as a network opinion to either block or not block the IP address.

The proposed methods were implemented as a Dovecot module for the free software Stratosphere Linux IPS (Slips), along with a networking library named Pigeon. This enables Slips to share detection data and take advantage of shared knowledge, while making the implementation available for others to build on. The module is written in Python, Pigeon is written in GoLang, and the two processes interact with each other and Slips through Redis channels.

We evaluated the performance of Dovecot in a simulated environment, testing multiple attack strategies and malicious peer strategies. We assume the attacker has multiple devices, some of them running the attack and others sharing false reports to confuse the IPS. We simulate 5 scenarios, each with a different strategy of the malicious reporters. The number of malicious reporter peers varies from none to nine and they use three different types of reputation attacks.

The detection accuracy of the Dovecot framework was compared to the performance of an IPS without any P2P mechanism in place, which was 75%. In the worst-case scenario where *all* the peers in the p2p network are malicious and they lie about *every* data they send, Dovecot can still give the same performance as an IPS alone. In a normal network, where it is likely that no malicious peers will be encountered, Dovecot can have a performance of 92.5% if a small amount of false positives is allowed. The same results are achieved even

with malicious peers in the network, as long as the majority of peers is honest. If the user assumes that the probability of a malicious peer joining the Dovecot network is low, then Dovecot can greatly improve the detections of IPS systems.

Key words

Intrusion Prevention System, Peer-to-Peer, Trust Model

Abstrakt

Systémy pro odhalení průniku (IPS) slouží k hledání a blokování síťových útoků. Ačkoliv může v jedné síti být více zařízení vybavených systémem IPS, v současném stavu techniky nevíme o žádném systému IPS, který by spolupracoval s ostatními a sdílel své nálezy. V této diplomové práci jsme navrhli a vytvořili P2P protokol Dovecot pro sdílení dat o síťových útocích. V našem protokolu se jednotlivé systémy IPS stávají členy Dovecot P2P sítě, posílají si navzájem své nálezy a používají je ke zpřesnění a zlepšení vlastních rozhodnutí. Ačkoliv předmětem protokolu je posílání dat, daleko důležitější je správné vyhodnocení toho, nakolik jsou získaná data důvěryhodná. Kvůli ochraně sítě před škodlivými členy jsme navrhli model důvěry Ω -Trust, který určuje důvěryhodnost každého člena na základě jeho síťového provozu. Pokud by škodliví členové zůstali v síti bez povšimnutí, mohli by rozepisovat nepravdivé nálezy a tím znehodnotit rozhodovací schopnosti IPS. Výstupem protokolu pro danou IP adresu jsou hodnoty Ω -score a Ω -confidence, získané sloučením nálezů od členů sítě, s přihlédnutím k jejich důvěryhodnosti. Tyto hodnoty vyjadřují názor členů sítě na danou IP adresu, a nakolik si jsou členové sítě touto hodnotou jisti. Systém IPS obě hodnoty využije při rozhodování o blokování konkrétní IP adresy. Metody popsané v této práci jsme implementovali a zveřejnili. První částí implementace je modul Dovecot pro svobodný IPS software Slips v jazyce Python, druhou částí je síťová knihovna Pigeon, psaná v jazyce Go. Obě komponenty spolu komunikují skrz databázový systém Redis. Implementací obou částí byla do systému Slips přidána možnost připojit se do P2P sítě a využívat nálezy ostatních členů. Účinnost protokolu Dovecot jsme ověřili v uměle vytvořeném prostředí, kdy jsme zkoušeli různé taktiky útočníka i škodlivých členů sítě. Pro porovnání přesnosti jsme použili samostatný systém IPS bez P2P sdílení, který byl schopný správně rozlišit útočníka od běžného zařízení s úspěšností 75 procent. V sítích, kde byli všichni členové sítě škodliví, lze vybrat takové parametry, aby přesnost pod 75 procent nikdy neklesla. Pokud předpokládáme, že většina členů sítě nemá škodlivé úmysly, lze parametry nastavit tak, že systém dokáže rozlišit útočníka od běžného zařízení s přesností až 92.5 procenta.

Klíčová slova

Systém pro odhalení průniku (IPS), Peer-to-Peer, Model důvěry

Contents

1	Introduction	1
2	Previous Work and Background	3
2.1	P2P Architecture	3
2.2	P2P File Sharing Networks	4
2.3	Byzantine Generals Problem	5
2.4	Salinity Botnet	6
2.5	Game Theory	6
3	Our Ω-Trust Model	9
3.1	Model Overview	9
3.2	Computing the Trust in a Peer	12
3.3	Ω -score and Ω -confidence	12
3.3.1	Ω -confidence	13
3.3.2	Ω -score	13
3.4	Using the Ω -score and Ω -confidence in the IPS	14
3.5	Linking PeerIDs to IP Addresses	15
3.6	Common Attacks on P2P Trust Models	16
3.6.1	Unfair Recommendations	17
3.6.2	Inconsistent Behavior	18
3.6.3	Identity Management	18
3.7	Avoiding Relevant Attacks in our Setting	19
4	Implementation of the Trust Model	21
4.1	Parts of the Dovecot Protocol	21
4.2	IPS and Communications with it	23
4.3	Pigeon and Communication with it	23
4.3.1	Sending Data to Peers	24
4.3.2	Receiving Data from Peers	24
4.3.3	Peer Updates and Reliability	25
4.4	Communication between Peers	26
4.4.1	Data Request	26
4.4.2	Report	27

5	Design of the Evaluation	29
5.1	Simulating Attacks and Adversaries	29
5.2	Overview of the Experiments	29
5.2.1	Networking	30
5.2.2	Traffic	30
5.2.3	Simulating the IPS	30
5.2.4	Peer Interaction	32
5.2.5	Results Extraction	32
5.3	Types of Devices and Peers	32
5.4	Evaluating the Results	33
6	Description of each Scenario in the Experiments	35
6.1	Malicious Device Attack Plans	36
6.2	Scenario 1 - IPS only, no P2P Network	38
6.3	Scenario 2 - P2P is Working, no Malicious Peers	38
6.4	Scenario 3 - Malicious Peers Praise the Malicious Device	39
6.5	Scenario 4 - Malicious Peers Lie about Everything	40
6.6	Scenario 5 - Malicious Peers Lie about Everything and also Attack	40
7	Experiment Results	41
7.1	Scenario 1 - IPS only, no P2P Network	41
7.2	Scenario 2 - P2P is Working, no Malicious Peers	44
7.2.1	Setup 2A - All Peers are Attacked, no Malicious Peers	44
7.2.2	Setup 2B - Peers are Attacked in Sequence, no Malicious Peers	45
7.2.3	Setup 2C - A Subset of Peers is Attacked, no Malicious Peers	47
7.3	Scenario 3 - Malicious Peers Praise the Malicious Device	48
7.3.1	Setup 3A - All Peers are Attacked	49
7.3.2	Setup 3B - Peers are Attacked in Sequence	50
7.4	Scenario 4 - Malicious Peers Lie about Everything	50
7.5	Scenario 5 - Malicious Peers Lie about Everything and also Attack	53
7.6	Results Discussion	55
8	Conclusion	61
8.1	Future Work	62
8.1.1	Better Prediction Computation	62
8.1.2	Manually Added Friends	63
8.1.3	Initial Reliability for New Peers, Cool Boot Start	64
8.1.4	Correlate Trust with Reports	64
8.1.5	Further Improvements	65
8.1.6	Experiments	66
A	Detailed accuracy tables	71

Glossary

benign describes a device or a peer that is participating in the network without the intention to run attacks or exploit other devices. We also assume the benign device is not infected. The opposite of benign is malicious. xi, xii

confidence is a value from 0 to 1 generated by the IPS. It is one of the values reported in the detection, along with the score. High values mean the IPS places a high degree of certainty on the score value, low values mean the score is uncertain. xi, xii, 23

detection is the result of the work of the IPS and is always made in relation to an IP address. The IPS claims in the detection whether the observed IP address is benign or malicious. In this work, we assume a detection from an IPS consists of two values: a score and a confidence. xi, xii, 1

device is used here to describe any machine (phone, computer, IoT device) that is capable of connecting to the internet. A device that has the ability to connect to the P2P network is also a peer. xi, 1, 9

IPS (Intrusion prevention system) is a software that monitors traffic and attempts to detect attacks on the device or the network. The outcome of the work of the IPS is a detection that consists of a score and a confidence. xi, xii, 1

malicious describes a device or peer that is running attacks, or plans on running attacks in the future. The opposite of malicious is benign. xi, xii

peer is a device that can communicate using the Dovecot P2P protocol. xi, xii, 9

prediction is a value about an IP address, that is computed from both the IPS detection (score and confidence) and the network data (the Ω -score and Ω -confidence). It is used to make the blocking decision. 10

reliability is a value from 0 to 1 that describes the quality of the connection with a remote peer. Higher values mean better connection quality. Reliability is computed in the Pigeon networking library from several simple statistics such as ping timeouts and it doesn't describe the intentions of the peer in the network. 21, 25

scenario is an experiment setting with a given strategy of malicious peers. 35

score is a value from -1 to 1 generated by the IPS. It is part of the detection, along with confidence. Positive values mean the observed IP address is likely benign, negative values mean the observed IP address is likely malicious. xi, 23

setup is a precise configuration of an experiment. 35

trust describes the faith a given peer places into another peer. Low trust means that the peers haven't interacted yet and don't *know* each other, or the interactions they had were not satisfactory. High trust means the peers have a long history of satisfactory interactions. The precise definition of trust changes in different papers researching trust models. 2

Chapter 1

Introduction

Computers connected to a network are constantly threatened by attackers. To protect the computers, administrators use various types of detection and protection software. However, these protection programs do not communicate with each other to share information in real time, and therefore do not take advantage of the knowledge of other detection systems in the network.

One of the defensive mechanisms used to mitigate attacks are the network intrusion prevention systems (IPSs). An IPS is a device or program that monitors network traffic and decides whether to block or allow traffic from different devices based on the data captured. An IPS can use different techniques to detect attacks: whitelists, blacklists, behavior analysis, anomaly detection, machine learning, or a combination of multiple techniques. In case of an attack, an IPS will detect the device running the attack, and then block all traffic from it. The malicious device can then pick a new target and run more attacks, until that target's IPS detects it and blocks it as well.

The attacker may have enough time to attack every device in the network before being detected, because current state-of-the-art IPSs don't share the detections with each other. Each IPS must do the same work to observe, identify, and block all of the attackers even though other IPSs have already done the same work. This process is probably not-optimal because, for the most part, IPSs don't cooperate. As far as we know, the non-cooperation of IPSs is still an open problem in our community.

A common solution for cooperation between IPSs is the use of collaborative databases; for example, blacklists of malicious IP addresses [36], or the MISP project [34] for threat sharing. Despite being simple, this is the most advanced sharing method the community currently has. With collaboration in place, it is harder for the attacker to succeed, since once they are detected for the first time and published in the list, they will be blocked immediately by all participating IPSs. While the shared threat databases improve the performance of IPSs, they are centralised, and updating threat data can be slow. The database must be maintained, and new entries are either added manually or by a limited group of selected contributors. Moreover, it is not clear how the data is curated or forgotten.

We propose the creation of a P2P network protocol, called Dovecot, where the IPSs running on personal computers send each other information about network devices. This would enable them to share detections immediately after they are made, without relying on a database or other centralised system.

With such a sharing approach, a new problem arises since P2P networks are in principle open for anyone to join, therefore making it necessary to differentiate between legitimate peers and malicious peers. A malicious peer is a member of the P2P network that lies about detecting attackers. Mechanisms called *trust models* have been designed to cope with this situation, and they can be deployed in different types of networks.

P2P trust models are mostly used in file sharing networks [17, 32, 35] to find the best peer for downloading a given file. Those models are not well adaptable to our setting, because they are designed to build a hierarchy of file sources, not evaluate a given peer. Another use of trust models is in autonomous cars [1], where the model is built on using signed data for verification. Models for network intrusion detection were also proposed [9, 11, 12, 20], but not implemented into an existing IPS.

In this work, we design and implement a trust model called Ω -Trust for the Dovecot P2P network of IPSs. The model must be balanced: a weak model would allow attackers to skew the decision of an IPS, but a model that is too strict will ignore reports from honest peers, defeating the purpose of the trust model. The structure of our trust model is heavily inspired by the Sality botnet [10], where each peer computes the trust values locally from first-hand experience. This makes the network harder to manipulate since there is no mechanism to recommend the trust value of another peer.

In the Dovecot network, each participating peer computes a local trust for every remote peer, using data obtained from the local IPS. Reports coming from remote peers are weighted with the reporters' trust values and the reports from all peers are aggregated, to get one *network opinion* value called P2P_detection. Information from peers with high trust is valued more than information from untrusted peers.

The protocol proposed in this thesis is implemented as the Dovecot module in the free software Stratosphere Linux IPS (Slips) [13]. The module is written in Python [27], and connects to a P2P networking library named Pigeon, which is written in Go language [15].

An important part of the Dovecot protocol is the message format based on JSON [4], which specifies not only the messages between peers, but also the communication between both parts of the software. Any IPS following the format can connect to the Dovecot network and share detections with other peers. Also, because trust values are always computed locally, different implementations can choose to use different trust models. More importantly, our Dovecot protocol shares detection scores, not alerts, so it doesn't have to understand the inner workings of the IPS.

Dovecot was tested in a simulated environment, and the detection rate of the participating IPSs was compared against the detections made by an isolated IPSs. Experiments have shown that the Dovecot doesn't jeopardize the security of the IPS, even in the least favorable network setting. If we assume that honest peers have a majority in the network, Dovecot can improve the decision accuracy from 75% to up to 92.5%.

In this thesis, we have implemented Dovecot, a modular software that enables P2P communication of intrusion prevention systems, and that can be easily added to existing IPSs. We have designed the Ω -Trust trust model that uses IPS detections, as well as connectivity data, to compute the trust of a peer. We have proposed a simple method that decides whether an IP address should be blocked, based on IPS detections and reports from the network. In a set of experiments, we have simulated several attacker strategies and measured the accuracy of our blocking decisions.

Chapter 2

Previous Work and Background

The field of adversarial networks and establishing trust relationships has been studied heavily, and the results have been applied in different scenarios. The solution or approach is always tailored to the given use case, and cannot be used in our environment directly. Several problems that are similar to ours have been tackled before. This chapter introduces them briefly, and discusses their similarities and differences when compared to our setting.

2.1 P2P Architecture

To better understand the following sections, it is first necessary to introduce the Peer-to-Peer architecture. The principles of P2P networking are explained in detail in [22], and this section only provides a short summary.

P2P architecture is an alternative to Client-Server architecture, which is typically used by HTTP, FTP, SMTP and other protocols. In Client-Server models, one server is running a service, and all the clients connect to it with their requests. This creates a single point of failure, and consumes resources on the server. In contrast to that, in a Peer-to-Peer model, as the name suggests, communication happens between individual peers, which have the role of both the client and the server. For example, a peer that is downloading a large file using BitTorrent may be at the same time sending parts of the file to other peers that want to download it - which saves resources of the original file source [22].

P2P doesn't necessarily mean that no server is needed. In the file sharing networks Napster and BitTorrent, a central server is used to hold metadata about files and where to download them, even though the server is not needed for the actual file downloads. P2P also doesn't mean that all peers must be equal in their functionality and purpose. In Kazaa, Gnutella2 and JXTA, there are two types of peers [22]. Exact terms vary in each of the networks, and so does their exact purpose, but simply put, there are peers and super-peers. Peers are *leaves* in the network, and keep only a few connections to super-peers. The super-peers are the *trunk*, they act as a proxy and relay communication to other peers.

2.2 P2P File Sharing Networks

In P2P file sharing, one peer requests a file and other peers reply if they can share it. Then, one remote peer is selected as the sender, and a *transaction* takes place, where the parties connect and exchange the file. The exact peer communication is described in detail for example in [18]. The requesting peer can only verify authenticity of the file if it knows the hash, and only once the full file has been downloaded. This opens opportunities for attackers to share spam or malicious files.

Several approaches to the identification of malicious peers in file sharing networks have been proposed. Mostly, they are based on assigning a trust value for each peer, and updating it - a good transaction increases the trust, and a suspicious or bad one decreases it [17, 32, 35]. This is similar to our setting, but there are differences that cannot be overlooked:

Finding the Best Match In file sharing, peers often need to make transactions with peers they have never encountered before. The trust models in such networks often revolve around a thorough procedure for requesting peer trust from the network, and then picking the best peer to interact with.

In contrast, IPS peers cannot choose the devices they will share the network with. They are part of the network, and an IPS can only decide whether it is secure to communicate with them or not.

Usage Peer interaction in file sharing consists of individual file transfers - unlike in our case, where interaction is continuous and lasts the whole time the peer is connected to the network.

Stability of Peers While the interactions in file sharing environments may be isolated, the peers often stay in the network for extended periods of time, for example the *seeders* in BitTorrent[6].

Some approaches like R-Chain [21] take advantage of this and choose peers - *witnesses* - to save the trust externally and guarantee authenticity. Sadly, in our setting, the peers may be laptops that jump networks frequently, and therefore there are no stable peers to rely on. This also applies to networks that are not connected to the Internet.

Pre-trusted Peers The Eigentrust [17] model relies on a set of pre-trusted peers, a group of peers that are controlled by the authors and guaranteed to be honest¹. Because our peer network should be independent and maintain itself, it is not acceptable for it to rely on a given set of peers. Also, pre-trusted peers cannot be reached from networks without Internet connection.

¹The authors do not discuss the possibility of the pre-trusted peers getting infected and becoming malicious.

Information Trust models often work with scarce information, for example, the user's binary evaluation of a file sharing transaction [17], and adherence to the protocol. Our setting has a great advantage, because apart from the trust protocol itself, our peers are IPSs and can detect attacks in the network. It is likely that a peer running on an IP address that behaves maliciously is malicious as well, even if it follows the protocol perfectly. Our peers may thus make more informed decisions thanks to this.

2.3 Byzantine Generals Problem

The problem of the Byzantine generals is a theoretical one, discussing communication between several generals that have been separated and must agree on an issue by sending messengers. However, some generals have been captured, subdued and are sending inconsistent messages. The problem has applications in distributed computing, where individual communicating components can be either faulty or malicious.

The PBFT (Practical Byzantine Fault Tolerance) [5] algorithm proposes a protocol for providing a secure and reliable distributed system, where each peer has a copy of the service the system provides. The protocol assumes that less than a quarter of the peers are malicious - precisely, $3f + 1$ reliable peers are needed to balance f faulty or malicious ones.

In this protocol, the client (or one of the peers) sends a request to a designated peer, which distributes it through the network. The distribution of requests is complicated, because the protocol assumes all peers have the same internal states and data - the exact steps are beyond the scope of this chapter. Once each peer has the request, they run it on the service, and reply directly to the client. When the client receives more than f identical responses, they can be sure that the answer is correct, because the protocol assumes that at most f peers are faulty.

There are assumptions in the PBFT approach that cannot be made in our scenario:

Same State, Same Data All PBFT peers have a copy of the same service in the same state with the same underlying data, and thus all of them should always come to the same results. Network monitors may have different information because they may be listening at different places in the network, thus their conclusions might also be different based on the detection model they are using. We cannot assume that honest and functioning peers will come to the same result.

Client-Server Infrastructure In PBFT, the peers are providing a service as a group, and can be perceived as a server. The client is an external entity that makes requests to the *server*. This can be overcome by simply allowing the peers to act as clients and initiate requests. However, the strict separation of roles between clients and the serving peers can make other operations IPSs need to do difficult.

No Server-initiated Sharing It is important for our model that peers that detect malicious behavior can share it immediately with their peers. PBFT has no mechanism for a peer in the *server* part to share unsolicited information - the communication always starts with a query, and all peers reply to it.

Message Complexity During the synchronization part of PBFT, where all peers agree on the order of the requests and internal service state, each pair of peers exchanges messages, leading to a $\Theta(n^2)$ complexity. This is not good enough for a network that may grow in size.

2.4 Sality Botnet

Sality [10] is a Windows malware used to distribute and download other malicious code. Because of its P2P nature, it is relevant to this work, and some design decisions were inspired by it. Machines infected with Sality become peers in the Sality botnet. When the bot-master wants to distribute new malware, they build an URL pack (payload) and sign it. The pack is then pushed to the network; peers verify its signature, distribute it, and download the malware to their host machines.

Sality peers use a value called **goodcount** as a simple trust mechanism to maintain a list of reliable neighbors. The **goodcount** starts at zero and is increased after every interaction that follows the protocol format, and lowered after each bad interaction. Peers with high **goodcount** are recommended to others, and peers with low **goodcount** are erased from the peer list. The **goodcount** is a strictly private value that is never sent over the network - this is not because it is confidential, but because computing it locally guarantees its integrity. An adversary cannot influence the **goodcount** a peer has for them (or anyone else) by manipulating the protocol. The only way to boost the **goodcount** is to interact honestly for an extended period of time.

Different versions of Sality have appeared in the Internet since 2003; the P2P component was introduced in 2008 [10], and we are not familiar with any successful attempts to infiltrate its network. Our protocol is inspired by this, and therefore also computes trust values strictly locally.

It is important to note the difference between trust in our protocol and the **goodcount** in Sality. Sality uses the **goodcount** solely to maintain a peer list and promote well-behaving peers. The payload itself is always verified against the provided signature, and the **goodcount** of the sending peer is irrelevant. On the contrary, the trust of a peer in our protocol can be interpreted as the credibility of the data received from that peer. There is no other method of verifying the data since there is no authority that could send it and sign it.

2.5 Game Theory

Game theory is the study of games, where a set of players (or agents) interact. Each player knows the rules of the game, can make moves and has a goal to reach. Game theory is a wide field, there are many types of games and many fields where they can be applied. In our case, each IPS (peer) is an agent that can collaborate with other agents, but will never know if they are malicious or not. Shillo et al. [30] propose a game theoretical solution to identifying deceitful agents in a network by collecting local experience as well as requesting information from other agents, however, this assumes the deceit can be identified.

Game theory is used to motivate honest peers by providing incentives. In a network where peers can share data, some peers may decide not to do so. These are *free-riders* on the network - they collect community knowledge while not putting any effort in.

Zhu et al. [37] propose a method to motivate IPS peers to share their detection data with others in an environment where attackers are not part of their network. While giving incentive is out of scope of this thesis, this research may be relevant in future work.

Game theory can also be used may also be to improve overall detection quality, when all players are honest and not adversarial [2]. A group of IPSs collaborate in order to protect the network against an external attacker. The IPSs collaborate to pick the best strategy, as each IPS only has the capacity to deploy a limited number of detection algorithms, and cooperation allows each sensor to specialize on one detection type, improving the overall performance.

Chapter 3

Our Ω -Trust Model

Dovecot, the trust protocol we propose, is built to share information between IPSs and help them make better blocking decisions. We chose the P2P approach, because it is fast; attacker's IP addresses are usually used shortly, and waiting for an IP address to be added to a blacklist takes too long. Furthermore, P2P can be used to report attackers inside the local network, which couldn't be reported on a public blacklist. It also removes the need for maintaining a central repository.

Because the very purpose of an IPS is to provide security, the protocol itself must be robust and secure. In P2P networks, anyone can become a peer, thus the environment must be considered untrusted, or adversarial. A trust protocol should be resilient against malicious actors by default, and during the design process, possible attack surfaces should be considered. The Dovecot protocol will be used by peers to share useful information about network security. If there is a flaw in such protocol, it will enable attackers to spread false information, which could lead the IPSs to either block an innocent device (effectively running a Denial-of-Service against it), or ignore warnings and keep communicating with a malicious device, thus defeating the purpose of the IPS itself.

Before discussing the trust model of Dovecot in detail, it is first necessary to define the two main terms used to identify participants in the network: device and peer. A device is anything that can communicate over the network, acting normally or attacking others (like a computer, phone, router, etc). A device is identified by its IP address. We use the term device to describe some generic entity that exists in the network.

A peer is a device that is capable of connecting to the P2P network (but not necessarily connected at the moment). Peers are identified by their peer IDs, which are further explained in Section 4.3. A peer group is a set of Peers currently connected to the P2P network.

3.1 Model Overview

Dovecot is used to share data between IPSs. The data is shared in the form of reports; a peer may send a report after making a significant detection in the IPS, or when asked for it by other peers. The trust model in Dovecot, the Ω -Trust, aggregates reports from multiple peers into Ω -score and Ω -confidence values, which represent what the network believes about an IP address at the moment. The values will later be combined with a detection from the

local IPS to compute a prediction, and based on the prediction, the local IPS will decide to block or not block the IP address.

While processing the reports, it is important to differentiate which peer the report is coming from. The concept of trust is introduced, where benign peers should have high trust, and malicious peers should have lower trust. The trust is used when aggregating the reports, and the benign peers should get more *voting power*.

Trust for a peer in Ω -Trust is based on data from two sources: the network behavior (IPS detection of the evaluated peer), and how much the peer adheres to the protocol. On top of that, the trust is only computed locally, directly from data sent by the evaluated peer. This is in contrast with trust recommendations or distributed trust computation like in [17]. This design decision was inspired by the Sality botnet mentioned in Section 2.4, where the `goodcount` is also strictly local and cannot be manipulated by other peers. With such a design, it is not possible to manipulate the trust, since it is computed locally based on first hand interactions with the evaluated peer. Note that the trust is computed only for peers, not for all devices.

Figure 3.1 shows the data flow in our model. Let's say that there is a device Z in the network, and three peers: a local peer and peers A and B . The local peer needs to compute a prediction for the device Z , to see if it should be blocked or not.

The device Z is sending traffic (packets) to all peers, and the IPS in each of the three peers makes a detection based on this. A detection consists of two values: a score that states whether the IPS detected the IP address as malicious, and a confidence that shows how much confidence the IPS has in the score value. The peers A and B are also sending traffic over the network, and the local IPS makes a detection for them, too.

The peers A and B share the detections they made about device Z with the local peer by sending a report. The local peer receives the reports and starts computing the prediction for device Z . The prediction is always computed with respect to a set of peers that sent a report - the reporters P . In this case, there are two reporters, $P = \{A, B\}$.

First, trust of each reporter is computed. As we mentioned before, there are two inputs in the trust computation: one is the reliability of the peer (adherence to the protocol) and the other is the `IPS_detection` made for the peer. The trust values of the reporters are combined with the confidence they reported to get an Ω_c value. The Ω_c is the *confidence* the IPS should consider when working with the reports - it is high when the peers are trusted and they are confident in the scores they report, and it is low when the peers are untrusted or not confident in the reports.

The trust values are further processed; they are normalized and transformed, to get a weighted trust (*voting power*) of each peer, that is proportional to the relative trust the peer has in P . The weighted trust is used to combine the score values from the reports and get a *score from the network*: Ω_s , which is computed from scores reported by peers A and B .

Now the local IPS has four values about device Z : score and confidence captured by the local IPS, and the Ω_s and Ω_c computed from the network reports. The IPS uses these values to compute a prediction, and can decide to block (or not to block) device Z based on it. When the local peer shares data about Z (or A and B) with other peers, it only shares the score and confidence from the local IPS. The trust, reliability, Ω_s and Ω_c are internal values, and they should not be propagated.

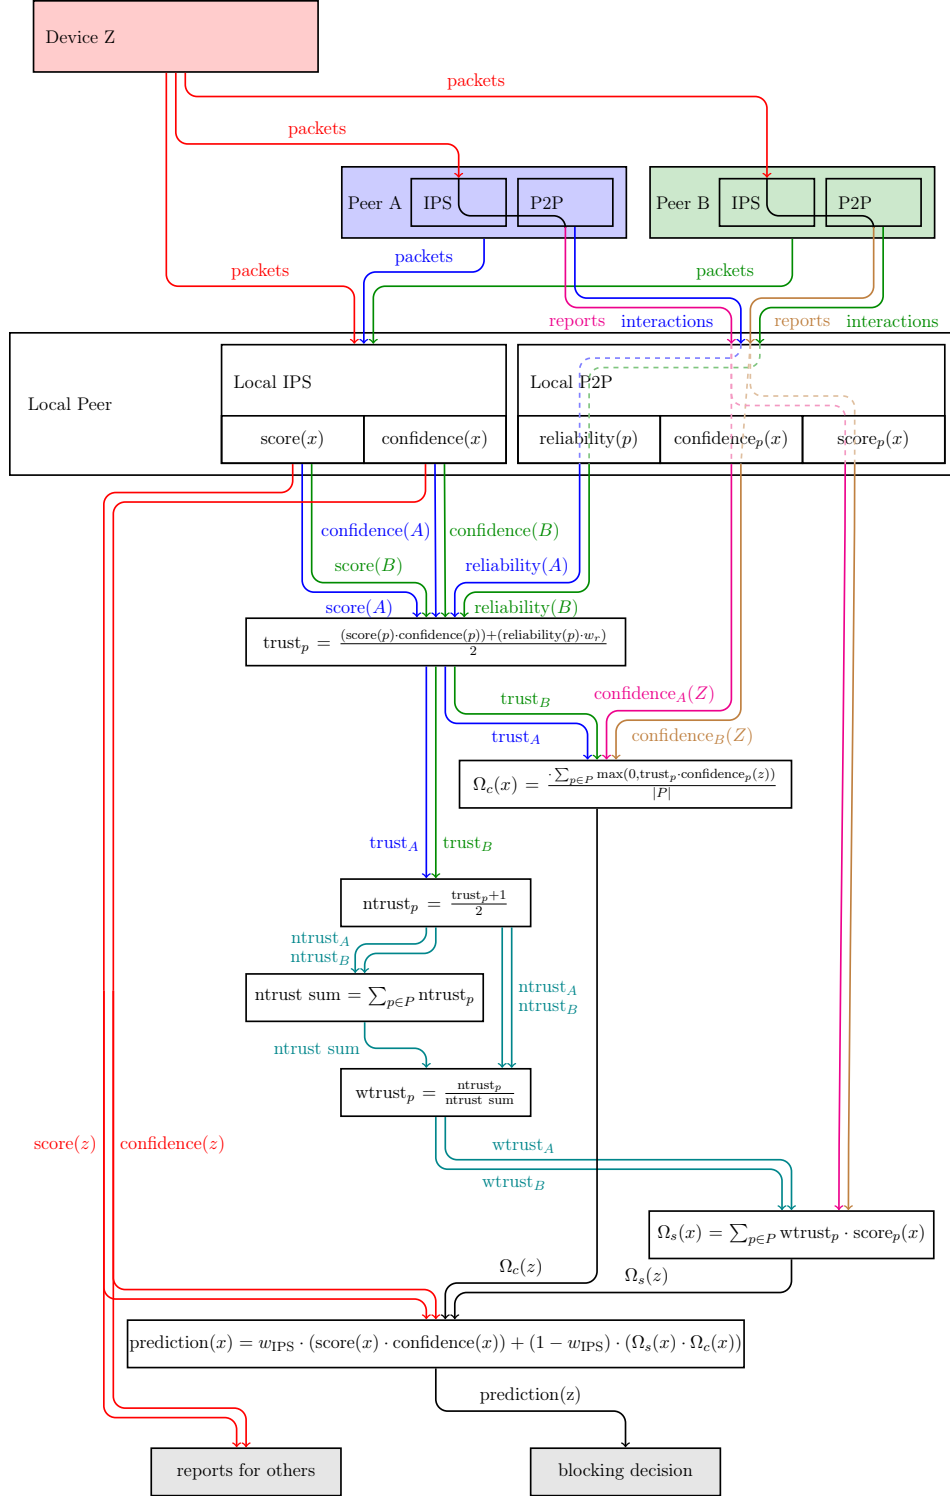


Figure 3.1: Schema of how predictions are computed in Dovecot. The schema follows how data is sent through the network, where it is captured and processed, and all computations needed to get the prediction value.

3.2 Computing the Trust in a Peer

A peer's trust is computed from the detections made by the local IPS about that peer, or, to be precise, from the most recent detection for any IP address claimed by that peer (see Section 3.5 to learn about IP address ownership). The reliability, a value that describes how well the peer behaves in the P2P network, is also used.

$$\text{trust}_p = \frac{(\text{score}(p) \cdot \text{confidence}(p)) + (\text{reliability}(p) \cdot w_r)}{2} \quad (3.1)$$

The trust computation for a peer p is shown in Equation 3.1. The score is a value from the IPS, computed with a function $\text{score}(p)$. It is a float from the $[-1, 1]$ interval, where -1 means that p is malicious, and 1 means it is benign. The confidence, computed with a function $\text{confidence}(p)$, is a value from the $[0, 1]$ interval, and it describes how confident the IPS is in assigning the score value, where 0 means not confident, and 1 means very confident.

Another input used to compute the trust of peer p is the reliability (computed with a function $\text{reliability}(p)$), a value from the $[0, 1]$ interval that describes the quality of the communication with p . Peers that follow the protocol will have a high reliability, and peers that do not follow the protocol will have a reliability close to zero. Reliability is never transferred over the network, as it is only used inside a peer. Section 4.3.3 discusses reliability in detail.

The value w_r , inside the $[0, 1)$ interval, is a weight applied to the reliability. This is a measure put in place to lower the influence of the reliability, which is computed in a more basic way compared to the score and confidence. The w_r also makes sure that the trust value will not reach the exact values 1 and -1, as this would lead to a division by zero in the following computations. This means that the trust values are inside the $[-0.5, 1)$ interval, specifically $[-0.5, \frac{1+w_r}{2}]$, depending on the weight parameter.

3.3 Ω -score and Ω -confidence

The Ω -score and the Ω -confidence on IP address z are the outputs of the trust model. The Ω -score describes what the network is reporting about an IP address, and it is similar to the score in the detection from the IPS. The Ω -confidence is slightly different from the confidence in the IPS, because it holds not only information about what confidence the peers reported, but also how trusted the reporters are. This means that the Ω -confidence can be used to decide how trustworthy the Ω -score is. In short, we refer to the Ω -score of IP address z as $\Omega_s(z)$, and to the Ω -confidence of that score as $\Omega_c(z)$.

When computing the $\Omega_c(z)$ and $\Omega_s(z)$, all reports about z ever received from remote peers are considered, and the most recent report from each reporting peer is used in the computation. A report about an IP address z sent by a peer p contains two values, a score ($\text{score}_p(z)$) and a confidence associated with that score ($\text{confidence}_p(z)$).

To compute the $\Omega_c(z)$ and $\Omega_s(z)$, first a set P of all the peers that sent a report about IP address z is assembled. The $\Omega_c(z)$ and $\Omega_s(z)$ are always computed with respect to a given P . If more peers send the reports, the values have to be recomputed for a new P .

3.3.1 Ω -confidence

The computation of the Ω -confidence is shown on Equation 3.2. The trust of every reporting peer p is taken and combined with the confidence the peer is reporting. This means that an untrusted peer that reported a high confidence and a trusted peer that is not confident with its report will both contribute to lowering the $\Omega_c(z)$.

$$\Omega_c(z) = \frac{1}{|P|} \cdot \sum_{p \in P} \max(0, \text{trust}_p \cdot \text{confidence}_p(z)) \quad (3.2)$$

In case the peer is not trusted at all, the trust may be below zero. The formula $\max(0, \text{trust}_p \cdot \text{confidence}_p(z))$ will be zero for such peers, and therefore will lower the Ω_c . This makes sense, because if malicious peers were contributing to the Ω_s , the confidence for that value should be low. It also ensures that the confidence will not be negative. The upper bound for Ω_c is $\frac{1+w_r}{2}$, even when all the peers are reporting high confidences, as it is limited by the trust (see Section 3.2).

The reported confidences are processed separately from the scores, to provide a countermeasure against malicious peers that report high confidences. Less trusted peers cannot manipulate the result by reporting high confidences, as the high confidences will instead boost scores reported by the more influential, highly trusted peers. Malicious peers reporting a fake high confidence may end up supporting a score that they don't want to endorse.

3.3.2 Ω -score

Recall that Ω -confidence is computed with respect to a set P of all peers that sent reports about the IP address. When computing the Ω -score using reports from a set of peers P , we do not only want to know how high their trust is (which is already represented in Ω_c), but rather what the relationship between their trust values is. If all peers are equally trusted, they should have the same *voting power*. If some peers are more trusted than others, they should get a higher proportion of the vote. We achieve this by normalizing the values, then finding the combined *amount of trust* of the reporters in P , and computing the proportion of trust a peer has inside P . We call this value the weighted trust of peer p , or $w\text{trust}_p$. The weighted trust values are then used to compute $\Omega_s(z)$.

The trust values of each peer are normalized by transforming the interval the trust values are in, $[-1, 1]$, to the $[0, 1]$ interval. This is shown on Equation 3.3. Note that because the trust doesn't occupy the full $[-1, 1]$ interval but only $[-0.5, \frac{1+w_r}{2}]$, the $n\text{trust}_p$ will also not occupy the whole $[0, 1]$ interval.

$$n\text{trust}_p = \frac{\text{trust}_p + 1}{2} \quad (3.3)$$

The normalized trust values are summed to get the amount of trust in a given set of reporters P . This is seen on Equation 3.4. The proportion of trust contributed by a peer p is the weighted trust of p , and it is computed as shown on Equation 3.5.

$$n\text{trust sum} = \sum_{p \in P} n\text{trust}_p \quad (3.4)$$

$$\text{wtrust}_p = \frac{\text{ntrust}_p}{\text{ntrust sum}} \quad (3.5)$$

The Ω -score is computed using the wtrust values computed earlier as the weights. The computation is shown on Equation 3.6.

$$\Omega_s(z) = \sum_{p \in P} \text{wtrust}_p \cdot \text{score}_p(z) \quad (3.6)$$

As weighted trust values are in the $[0, 1]$ interval, and scores are in the $[-1, 1]$ interval, the Ω_s will be in the $[-1, 1]$ interval.

3.4 Using the Ω -score and Ω -confidence in the IPS

The IPS needs to make the *block / don't block* decision about an IP address, using the score and confidence values from both the IPS and the network data represented by the Ω -score and Ω -confidence. There are various methods that can be used to ensemble the values, and finding the perfect method is out of scope of this thesis.

For experiment purposes, we worked with a very simple model where the P2P and IPS detections are multiplied to get one prediction value. A decision to block the IP address is made if the prediction value falls below a given threshold.

The detection value from the IPS is computed by simply multiplying the score and confidence values, as shown in Equation 3.7. Note that score is a value from -1 to 1, and confidence ranges from 0 to 1, so the result will also be in the -1 to 1 range. An $\text{IPS_detection}(z)$ value close to zero indicates that the confidence for the detection is low, or the score itself is low. An $\text{IPS_detection}(z)$ value of 1 means that the IP address z is not malicious, and the confidence about this is high. On the other hand, $\text{IPS_detection}(z) = -1$ means that the IPS is sure that the IP address z is malicious. If the value is close to zero, it means that the IPS is not sure about the score it is reporting.

$$\text{IPS_detection}(z) = \text{score}(z) \cdot \text{confidence}(z) \quad (3.7)$$

Computing the P2P_detection value is done in a similar way, only instead using the Ω -score and Ω -confidence, and it is shown in Equation 3.8. The meaning is also similar to the $\text{IPS_detection}(z)$, but receiving a $\text{P2P_detection}(z)$ close to zero probably means that the Ω -confidence was low, which means that the peers in the network (and therefore the detection from the network as well) should not be trusted. This is more serious than getting the $\text{IPS_detection}(z)$ close to zero, because in P2P, this can mean that there are malicious peers manipulating with the protocol.

$$\text{P2P_detection}(z) = \Omega_s(z) \cdot \Omega_c(z) \quad (3.8)$$

The prediction is a one final value computed by the IPS, that evaluates an IP address. It ranges from -1 to 1, where low values mean that the IP address is malicious and should be blocked, and high values mean that the IP address is benign. The prediction is computed

by giving a weight w_{IPS} to the $IPS_detection$ and a weight $w_{P2P} = 1 - w_{IPS}$ to the $P2P_detection$. The computation is shown in Equation 3.9.

$$\text{prediction}(z) = w_{IPS} \cdot \text{IPS_detection}(z) + (1 - w_{IPS}) \cdot \text{P2P_detection}(z) \quad (3.9)$$

The w_{IPS} can be used to configure the *significance* of the IPS data. Setting the weight to 1 means that the $P2P_detection$ is ignored, and only data from the IPS is used. On the contrary, setting the weight to 0 means that the $IPS_detection$ is not used, and the prediction is based only on P2P data. Simply put, a $w_{IPS} = 0.4$ (below 0.5) means that the IPS will trust its own detection slightly less than data from the network. This is later simulated in the experiments, to see how much can the IPS trust the network before the security is jeopardized.

For the blocking decision, there is a threshold put in place, which is used to make the blocking decision. This is shown in the 3.10 below.

$$\text{block}(z) = \begin{cases} \text{true}, & \text{if } \text{prediction}(z) \leq \text{threshold} \\ \text{false}, & \text{otherwise} \end{cases} \quad (3.10)$$

As these parameters directly affects the final result reached by the system, we simulate different values for both the IPS weight and the threshold in the experiments, and propose optimal values. This essentially means we are trying to find how sensitive the system should be when blocking (threshold), and how much it should follow recommendations of the network (IPS weight).

3.5 Linking PeerIDs to IP Addresses

Ω -Trust computes a trust value for each peer, where each peer is identified by its PeerID. PeerIDs are further explained in Chapter 4, but for now, it is sufficient to know that a peer generates their PeerID when joining the network, and then all P2P communication is associated with that PeerID. The PeerID is preserved when the peer reboots or changes IP address, but malicious peers may choose to generate new PeerIDs whenever they choose to.

The P2P communication and trust computation works on the PeerID level, however, an IPS only links detections to IP addresses. Because peers may change the IP address they are running on, it is necessary to create a mapping that links one PeerID to all of the peer's actions across various IP addresses.

In our database, the Ω -Trust protocol has the timestamps of when each PeerID appeared on an IP address (claimed it). We say that the IP address *belongs* to a peer from the time that peer first connected from that IP address. The IP address is released when another peer claims the IP address, or when the same peer connects from a different IP address. All actions done by the IP address while it *belongs* to a peer are considered to be actions of the peer, and can be used to compute its trust.

Note that if a peer leaves the network and a malicious device joins on the same IP address, the malicious actions performed by the device will still be associated with the original peer. This may seem unfair, but not linking the actions on the IP address to the last peer that

used it would make it easy for attackers to keep a high trust value despite attacking. The attackers could simply shut down the P2P layer, then run some attacks, then run the P2P again, all from one device without the need to join and leave the network.

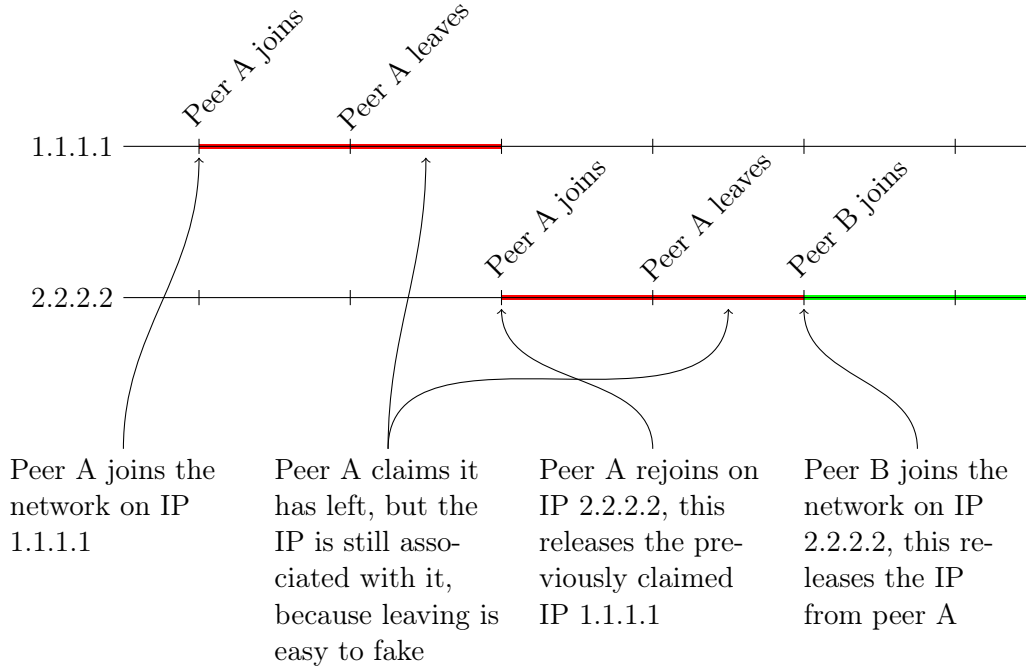


Figure 3.2: Ownership schema of IP addresses and PeerIDs in the Ω -Trust protocol. The red line marks the time when the Ω -Trust protocol believes an IP address still belongs to peer A. Green is similar for peer B.

The IP address ownership schema can be seen in Figure 3.2. The sample network has two IP addresses, 1.1.1.1 and 2.2.2.2, and two peers A and B. The timeline goes from left to right. When computing trust for peer A, detections made in the red area can be associated with peer it, and detections in the green area belong to peer B.

3.6 Common Attacks on P2P Trust Models

When designing the Dovecot protocol, we consulted a paper by Koutrouli and Tsalgatidou [19] on P2P reputation systems. The paper provides an overview of different trust models, highlighting their different purpose and design decisions. It describes common attacks on trust models, splitting them in three categories (unfair recommendations, inconsistent behavior and identity management related attacks) and suggests what trust model designers can use to mitigate these attacks. The following sections show some of the attacks in each category, and discuss their relevance in the Dovecot P2P IPS network.

3.6.1 Unfair Recommendations

In the trust models reviewed by the paper, the peers may send inaccurate or just plain wrong recommendations, which may change how peers treat each other as well as who they interact with. For example, say the attackers want to promote the malicious peer E as a file source. They will send a falsely positive recommendation about E to promote it. After receiving multiple recommendations about E, a peer may choose to download files from E, because of its positive reputation.

Recall that the trust of a peer in Dovecot describes how much we believe the reports from that peer about an IP address. We never use reports from a peer to compute trust of other peers. Trust is always computed directly from data provided by the peer itself. Therefore, it is not possible for malicious peers to affect the trust of another peer by sending unfair reports. This means that a malicious peer cannot make someone less (or more) trusted. It would seem that because of this, unfair recommendations are not an issue. However, the reports from peers are used to compute the Ω -score and Ω -confidence, which are then used in the decision to block or not block an IP address. This means that unfair recommendation attacks are still relevant and cannot be ignored, even though in this case, the unfair recommendations are about a device and not a peer.

The following list describes unfair recommendation attacks, applied to our proposal where possible.

Badmouthing Falsely claiming that a benign IP address is actually malicious. This may damage the targeted IP address by forcing peers to stop communicating with it. It is a serious problem and it is considered in our Ω -Trust model and simulated in the experiments.

Unfair Praises Falsely claiming that a malicious IP address is actually benign. This is very similar to Badmouthing, and it is also part of our experiments.

Random Opinions Sending random reports to save computing power and meet report quotas. This is an issue in models where reports are rewarded or demanded. The problem here is when the incentive is too strong and rewards the quantity, not the quality of reports. As there is no incentive in Dovecot, randoms opinions are not an issue.

Inaccurate Recommendations Sending reports without having enough information. The paper suggests this can be handled by providing a confidence value with the report, which Dovecot does.

Unfair Praises under Pressure Giving high rating in hopes of receiving a high rating as well. This has been observed on Ebay [7], where sellers and buyers quickly give high ratings to guilt their counterpart into also giving them a high rating. This is not relevant for Dovecot, because there is no social aspect, and also a peer can't access the reports others are sharing about it.

3.6.2 Inconsistent Behavior

Traitors Malicious peers stay in the network for a long time behaving correctly, gain high trust, and then exploit the trust for malicious intentions. This is an issue in Dovecot, but there is no way to know a peer's intentions beforehand. The countermeasures can include making the trust hard to gain and easy to lose, or assigning higher weight to recent events. In the experiments, we explore both traitor devices, and traitor peers.

Discrimination Peers behave nicely towards majority of peers, but focus their attacks on a selected target. This way, they can run the attack, but because they are only being reported by one victim, the overall trust will not be affected. In our setting, this can manifest in two ways, both of which are relevant and simulated in our experiments.

- The malicious peer will attack one of the peers, and this peer's IPS will detect the attack. The victim will likely block the attacker, and report the observed behavior to other peers. However, because no other peer is detecting or reporting the attack, the attacker will likely not be blocked by anyone else.
- The malicious peer is not attacking directly, but is sharing unfair information with a selected victim peer. This should not affect the model, as the victim will still get honest reports from other peers from the network.

Free-riding Peers are taking advantage of data from others, without contributing themselves. This is an unaddressed issue in Dovecot, and peers can choose to not share anything and will not be punished for it. Because there is no mechanism in place, there is no point in verifying it experimentally. It also does not damage the performance of the protocol.

3.6.3 Identity Management

Sybil Attack The attacker creates multiple identities (peers), thus owning majority of peers in the network. The malicious collective can then coordinate and execute previously mentioned attacks. This is a problem in protocols where there is no registration authority, and our protocol is the case. This attack is simulated in the experiments.

Whitewashing The attackers lose reputation by running an attack. They then discard their identity and get a new one. In some models, this could give them a better trust, because the initial trust for new peers is higher than the attackers' current trust. Switching identities clears their record, and the attackers can continue their attacks, until the reputation gets low and a new identity is generated.

Impersonation In models where no authentication is present, an attacker may impersonate another peer. The Dovecot network is built on the libp2p library, where the identification (the `PeerID`) is secured by a private key [26], so straight forward impersonation is not possible on the P2P layer. However, the score of the underlying IP address can be damaged by running attacks while faking the IP address, which is addressed in Section 3.7.

Man in the Middle P2P traffic tunneled through other peers may be subject to MitM attacks, where the middleman could alter reports sent between peers. Dovecot uses libp2p, and relies on the relay mechanisms provided by the library [24].

3.7 Avoiding Relevant Attacks in our Setting

Badmouthing and Unfair Praises Malicious reports cannot be avoided, and the trust model must be robust enough to withstand them. We address this by giving higher significance to reports from more trusted peers (see Section 3.3).

Traitors Detecting traitors is nearly impossible, because they are indistinguishable from honest peers for long periods of time. One countermeasure is to make trust easy to lose and hard to gain, which will mean that the traitor's trust will drop quickly upon discovery and will be hard to earn back. However the trade-off here is that a honest peer may misbehave by accident and its trust will suffer greatly because of it.

Discrimination As explained before, there are two types of discrimination: selectively attacking the victim and sharing false information with the victim.

Attacking the victim is tightly connected to Badmouthing. If peers are reluctant to believe a report about an attack, badmouthing will not affect them. However, they will also not consider reports from a targeted victim to be true. On the other hand, if peers are easily convinced, they will uncover actual attacks faster, but be fooled by badmouthing attacks more often. When traitor peers are included, the whole situation is very complicated to balance.

Sharing false information can trick the victim into making a wrong blocking decision. Since trust is local in Dovecot, nothing unusual is happening in the eyes of other peers, and the victim handles the fake reports just like in badmouthing or unfair praises, not knowing that it is the only target of the attack. As trust is never shared, other peers will not learn about this.

Free-riding Free-riding can be prevented by providing a reward when peers share information. Peers may choose to free-ride when the cost of computing the information is too high, they are malicious, or when they benefit from it - all service may go to the highest trusted peer, and it may not want to lose this position by *giving* trust to others. There is no incentive model in place in Dovecot, as the stated reasons are not very significant:

- There is no cost in computing the reports, they are part of the IPS anyway and no additional computation is needed.
- Malicious peers have much more impactful strategies at hand to damage the network.
- A highly trusted peer does not benefit by getting all the service, because there is no concept of *getting service*. Also, it can't affect the trust of other peers by not sharing.

Sibyl Attack Generating multiple identities can be prevented by having a central registration authority, or at least slowed down by demanding a proof-of-work, as explained in [16], for each joining PeerID. There is no countermeasure currently put in place to prevent Sybil attacks, however, the resilience of the protocol for different amounts of malicious peers is tested in the experiments.

Whitewashing Whitewashing is an issue in models where trust can go below the initial value. It is important to propose an initial value that is not too big of a win for an attacker with damaged reputation.

Impersonation As mentioned earlier, PeerIDs are not easy to spoof, as one would have to steal the encryption keys stored inside the peer. However, a malicious entity could pretend to use a peer's IP address, send malicious traffic, and therefore get the peer's trust lowered, or the IP address blocked. Attackers can spoof their IP address in different ways, and this needs to be detected in order for our protocol to be useful. There are already methods that detect this [28, 33], and we assume there is a reliable mechanism in the IPS that doesn't let spoofed packets pass through.

Chapter 4

Implementation of the Trust Model

A significant part of this thesis work was devoted to the implementation. The Dovecot protocol will be available to the community in the Stratosphere Linux IPS (Slips) [13] as a trust module. The IPS implementation itself was not part of the thesis.

This chapter describes how Dovecot implementation is structured and what are the roles of its individual parts. It also describes the format each individual part uses to communicate and discusses the design decisions.

4.1 Parts of the Dovecot Protocol

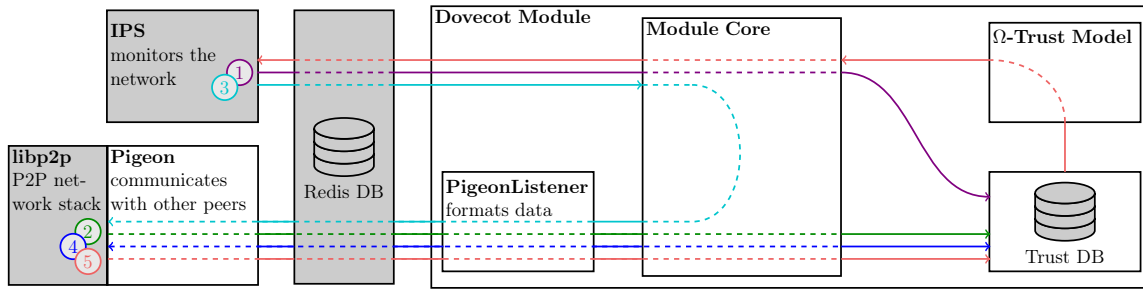
The complete implementation of the Dovecot protocol has three layers: the Stratosphere Linux IPS, which was already written in Python, the Dovecot module for Slips, written in Python, and our P2P networking library called Pigeon, which is implemented in GoLang. The implementation, as of the date this thesis is submitted, uses GoLang version 1.14.4, Python version 3.8.3, and the latest code from the `develop` branch of Slips.

The layers communicate by sending data through Redis channels. Redis was chosen because it is an easy way to connect multiple programming languages, and because it is already widely used in Slips. The components of the code and data flows between them are shown in Figure 4.1.

The IPS saves all updates to the Redis database. It notifies the Dovecot module through a Redis channel whenever information for any IP address changes. It can also request the module (via the Redis channel) to find more information about an IP address. The outputs from the Dovecot module are saved into Redis, where the IPS can access them.

Pigeon serves as a networking layer. It is used to send messages between peers, but it also computes the reliability of peers using up-time and ping, which ranges from 0 to 1. Pigeon sends all the messages it receives from other peers to a Redis channel where the module can process them. It also keeps track of the currently connected peers, their IP addresses and reliability. If any information changes, Pigeon sends an update to the Dovecot module to notify it.

The Dovecot module is the core part of the protocol. It listens for updates from the IPS and Pigeon, and saves all detections in a SQL database called `TrustDB`. This means



- ① When the **IPS makes a detection**, it saves it in the Redis database. It also notifies the Dovecot module about this (through the Redis), and the module copies the data to an internal database.
- ② **A new peer joins the network**. Pigeon sends information about this peer to the Dovecot module, and the module saves it. The same data flow happens when peer information is updated.
- ③ The **IPS requests information** about a given IP address from the Dovecot module. The module forwards this request to Pigeon, which then sends it to all known peers.
- ④ In a different peer, a **request is received by Pigeon**. Pigeon forwards it to the module, which processes the request, fetches data from the database, and formats it to send it back to the peer that asked.
- ⑤ After a **response from a remote peer** is received, it is added to the internal database. After a timeout, it is assumed that all responses have arrived, and the Dovecot module core asks the trust model to give a result. At this point, the trust model processes data about all known peers, computes their trust, aggregates their reports, and computes the Ω -score and Ω -confidence. These values are then saved to Redis by the Dovecot module, and the module notifies the IPS that the data it requested is ready.

Figure 4.1: Structure of the implementation including data flows. A sample data flow is shown in colored lines and described. The areas marked in gray were not developed as part of this thesis.

that the same information is stored redundantly (in the IPS's Redis, and in the TrustDB), but this allows for faster trust computation. The most important part of the module is the Ω -Trust model, which works with data from the TrustDB to compute the trust of individual peers as well as the Ω -score and Ω -confidence of each IP address. The Dovecot module is structured so that the trust model is independent, and can be improved or replaced easily without disrupting the module.

An important part of the implementation is the message format used by the individual parts. The design is highly modular - any IPS can be easily modified to use Redis and participate in the Dovecot network. Advanced developers may deploy custom networking layers to replace Pigeon. The crucial part of the module, the trust model, can be also replaced, with only slight changes to the module code. This is done to allow easy development

and improvement of the protocol.

4.2 IPS and Communications with it

The current implementation of Dovecot works with the Stratosphere Linux Intrusion Prevention System (Slips) [13]. Slips is an intrusion prevention system that is based on behavioral detections and machine learning algorithms. Its core is to separate the traffic into profiles for each IP address, and then separate the traffic further into time windows to analyze traffic features in different ways [14].

Slips is built around a Redis database, and uses the hashset `IPsInfo` to hold information about all IP addresses the IPS has encountered. Data for each IP address in this hashset is a dictionary with the relevant keys for our module being `score` and `confidence`, which describe a detection made by the IPS. Score is a value that describes if the IPS believes a peer is malicious or not, ranging from -1 (malicious) to 1 (benign). Confidence of the detection describes the certainty the IPS has in the score value, ranging from 0 (not sure at all) to 1 (certain).

Slips uses two Redis channels to communicate with the module: `ip_info_change` and `p2p_data_request`. The first channel is used to notify the module when score or confidence of an IP address change, and the message contains the affected IP addresses. The second channel is used by Slips when it wants to get an opinion on an IP address from the Dovecot network. The message also contains the IP address in question.

Return values from the module (the Ω -score and Ω -confidence) are saved into the `IPsInfo` hashset for the corresponding IP address under the key `p2p4slips`, as seen in Listing 4.1. The values are computed by aggregating reports from other peers as explained in Section 3.3. Additionally, a unix timestamp is saved with the values to mark the time when the values were computed (which is not necessarily same as the time when the reports were received from peers).

```
{
  "p2p4slips":
  {
    "network_score": 0.9,
    "network_confidence": 0.6,
    "timestamp": 154900000
  }
}
```

Listing 4.1: Example of a saved module output (the Ω -score and Ω -confidence) in Redis, where the IPS can access it.

4.3 Pigeon and Communication with it

Pigeon is the P2P layer of the protocol and it is built around the `libp2p` [23] library. It is implemented in GoLang, unlike the rest of the project, because the Python implementation

of the P2P library is still under heavy development, and doesn't yet support full functionality of the libp2p standard. Libp2p is part of the Inter-Planetary File System [3] suite.

The purpose of Pigeon is to manage connections to other peers and send messages from the module to peers (and from peers back to the module). Pigeon identifies different peers by a PeerID, which is a value that each peer creates for themselves when joining the libp2p network [25]. The PeerID is represented by a base58 encoded string, for example `QmYyQSo1c1Ym7orWxLYvCrM2EmxFTANf8wXmmE7DWjhx5N`. PeerIDs are hashes of public encryption keys, and because of that, they can be used to securely identify a peer. Pigeon can always read the sender's PeerID from an incoming message and start communication with other peers by knowing their PeerID.

Pigeon is only a wrapper around libp2p. It does not need to directly control any P2P processes (message forwarding, checking neighbors in hash tables etc), as libp2p conveniently takes care of all this in the background. All messages received by Pigeon are meant for the local peer.

4.3.1 Sending Data to Peers

Pigeon is listening on the `p2p_pygo` Redis channel for messages to be sent to other peers. The messages sent over the channel are in JSON, an example of a message is shown in Listing 4.2. The field `recipient` is the PeerID of the peer that should receive the message, and the message is the data to send. The PeerID can also be replaced by the wildcard represented by the character `*` in order to have the message sent to all known peers. Notice that the message is a `base64` encoded string (shortened for convenience in the example). Messages are always `base64` encoded to have a clearly defined charset, and Pigeon doesn't decode them. The actual contents of the message are discussed in Section 4.4.

```
{
  "message": "ewogICAgImtleV90eXB1IjogImlw.....jYKfQ==",
  "recipient": "QmYyQSo1c1Ym7orWxLYvCrM2EmxFTANf8wXmmE7DWjhx5N"
}
```

Listing 4.2: Example of the data sent from the module to Pigeon (the message field is shortened)

4.3.2 Receiving Data from Peers

Pigeon is also listening for messages being sent by other peers in the network. Each incoming message is processed by Pigeon only briefly. Sender's PeerID and the Unix timestamp of the message are collected, but the data inside the message is not unpacked.

Raw message, as well as the sender's PeerID and the timestamp are formatted as JSON and sent to the `p2p_gopy` Redis channel, where the Dovecot module is listening. The JSON string also contains a `message_contents` field with the value `go_data`. This is to differentiate a network message from an update that initiated inside the Pigeon, which is discussed in the following section. The message sent to the module looks like Listing 4.3.

```

{
  "message_type": "go_data",
  "message_contents":
    {
      "reporter": "QmVZgyY8Usx5DGidoxpDRaSRA1u9JBXy4nP9ymWGMBBjp3",
      "report_time": 154900000,
      "message": "ewogICAgImtleV90eXB1IjogImlw.....jYKfQ=="
    }
}

```

Listing 4.3: An example report from another peer, after being processed by Pigeon (the message field is shortened)

4.3.3 Peer Updates and Reliability

To have an overview of how the peers are behaving, Pigeon computes the reliability of each peer. This is done by rating each interaction with a number from 0 (bad) to 1 (good), and then averaging over all ratings.

The peers follow a simple protocol using a basic `hello` message containing the version and a `ping` message. The peer will get only positive interaction reviews if it follows the protocol, but will lose reliability if the hello message is missing parameters, there are no ping replies, or the pings are repeated too often. Currently, the reliability is averaged over all interactions, which gives new peers the advantage of having reliability 1 (this is no longer possible after any transgression, for example one missed ping).

Apart from reliability, Pigeon also tracks the IP addresses of peers. Whenever a peer makes a new connection, its IP address is saved. This linking of the peer to an IP address is important for the module because it uses PeerIDs to compute trust, but IP addresses to make detections. The rules for linking IP addresses to PeerIDs were described in Section 3.5.

The reliability and IP address changes are both important, and the Dovecot module should be notified. Messages about such changes in the network are sent to the module in a JSON format through the `p2p_gopy` channel. The message must contain the `peerid` field, but the `ip` and `reliability` fields can be omitted, as it is possible that only one of the values is new. The message type `peer_update` is added to the message, to differentiate from the `go_data` message. The message may look like Listing 4.4.

```

{
  "message_type": "peer_update",
  "message_contents":
    {
      "peerid": "QmYyQSo1c1Ym7orWxLYvCrM2EmxFTANf8wXmmE7DWjhx5N",
      "ip": "1.2.3.4",
      "reliability": "0.3"
    }
}

```

Listing 4.4: Example message with updated peer data sent from Pigeon to the module

4.4 Communication between Peers

Peer communication is initiated inside the Dovecot module, transmitted through Pigeon into the P2P network, where another Pigeon receives it and passes it to its module for processing. In the following text, we will refer to this as the peers communicating, but keep in mind that Pigeons represent the networking layer, and the logic is happening inside the modules.

There are two types of messages that the peers exchange: reports and requests. Reports are shared by peers that have made a new detection and want to share it with the rest of the network. On the other hand, requests are sent by peers when they lack detection data, and they ask the rest of the nodes in the network for more information.

Requests and reports are always sent in relation to a **key**. The key is a device identifier, usually an IP address. The protocol requires the **key_type** to always be specified. This allows for the model to be easily adapted to a new identifier in the future. Currently, the only supported **key_type** is **ip**.

In Chapter 3, we mentioned that reports from peers contain the score and the confidence their IPS detected. We have chosen to share only the two values, as problems arise with other solutions presented in the literature (raw capture samples, processed detection data, to final detection values etc.).

- Sharing captures could potentially leak sensitive information. Reliably anonymizing data is complicated in real-time and even anonymized data should not be shared with untrusted parties [8].
- Larger amounts of data would drain network resources and storage on other peers.
- All IPSs that decide to join the network would have to support one complex data format. Cooperation between different types of IPSs would be difficult, as it would be time consuming to modify each IPS to implement the correct format.

The format we are using is short and convenient to transfer, and can be extracted from complex IPS outputs. Apart from that, it can be easily sent through the network, even to adversarial peers, because it doesn't convey any sort of confidential data. The only leaked information about the peer that is sending the data is whether it has interacted with the given IP address or not.

In the future, other data formats may prove to be better suited for detection sharing. The Dovecot protocol is prepared for such changes, and similarly as with **key_type**, a field to specify data type is provided. This is the **evaluation_type**, and the only currently supported type is **score_confidence**.

4.4.1 Data Request

The request message is used when a peer wants to get information about a given device. It consists of four fields:

- **message_type** Type of the message, in this case this will always be **request**

- `key_type` Type of the key that identifies the remote device: `ip`
- `key` The identifier of the device the peer is asking about
- `evaluation_type` Type of data (detection) the peer is expecting: `score_confidence`

An example request message can be seen in Listing 4.5. Note that this is not the actual message that is sent to Pigeon. The message is first encoded in base64, the recipient PeerID is added to it (or * to send it to all known peers) and lastly, it is formatted as shown in Listing 4.2.

```
{
  "message_type": "request",
  "key_type": "ip",
  "key": "1.2.3.4",
  "evaluation_type": "score_confidence"
}
```

Listing 4.5: Request the module sends through Pigeon to other peers, to ask for data about the IP address 1.2.3.4. The request is formatted using base64 before being sent to Pigeon.

Peers will always reply with report messages when they receive a request. This is not optional, and if the peers have no data to share, they will send an empty `evaluation`.

4.4.2 Report

Reports are sent in response to requests, or independently when new detections are processed in the IPS. The fields in a response message are same as in the request message, except for:

- `message_type` Type of the message, in this case this will always be `report`
- `evaluation` Data given by the reporter's IPS

The inner format of the evaluation can differ. The implemented `score_confidence` detection expects to have `None` (when the peer has no data to share) or a dictionary with two more values:

- `score` Score given by the reporter's IPS
- `confidence` Confidence given by the reporter's IPS

The report may look as shown in Listing 4.6, and before sending, it will also be encoded with base64.

```
{
  "message_type": "report",
  "key_type": "ip",
  "key": "1.2.3.4",
  "evaluation_type": "score_confidence",
  "evaluation":
  {
    "score": 0.9,
    "confidence": 0.6
  }
}
```

Listing 4.6: Example report with data about 1.2.3.4. It will be encoded using base64 and sent to other peers (or the one peer that sent a request) through Pigeon.

Chapter 5

Design of the Evaluation

To evaluate Dovecot, we run several experiments. In them, we model that different devices are present in the network and sending some traffic, and there are also peers sharing reports about their detections. We manipulate how many of the network entities are malicious, and what their strategies are. This is done to measure how the trust model contributes to the overall IPS detection accuracy, and to find if there are any constraints needed for a good performance - for example on the maximum number of malicious peers. Knowing the limitations is crucial, because our proposal will be used by security software, and must be therefore resilient to attacks. In this chapter, we describe the environment in which the simulations are run.

5.1 Simulating Attacks and Adversaries

To evaluate an IPS, the models could in theory be tested with a real IPS on captured traffic, which would provide very accurate results. There are several datasets [31] that could be used to run such experiments. However, this could only be used to check model performance against an attacker who has no knowledge of the trust model, and doesn't participate in the P2P network. An informed attacker would likely schedule the attacks and participate as a peer in the P2P network in order to affect the IPS more effectively. However, there are no datasets crafted for this purpose, since it is not only about traffic, but about IPS detections and strategies of attackers as well.

Because of this, we decided to simulate the network communication and the IPS detection process, without actually running the network layer. The devices in our network simply choose a behavior profile from a defined set, and the IPS makes a detection based on that behavior. This is discussed in detail in Section 5.2.3 further in this chapter.

Before the individual experiments can be discussed, it is first necessary to understand how the experiment environment is designed and what it can and cannot simulate.

5.2 Overview of the Experiments

The experiments are run in rounds in order to simplify the parameterization. Each round represents one time window of network traffic. Twenty rounds are repeated in each experi-

ment - this is to enable behavior changes in time, and to analyze different attack strategies. Each round is structured as follows, and in some steps of the round, the participants can make decisions and change their behavior:

1. **Networking:** Devices join or leave the network and change IP addresses. They also become members of the peer group or leave the peer group.
2. **Traffic:** Devices pick their behavior profile in the network (malicious or benign).
3. **IPS detections:** The IPS of each device processes the network behavior it "captured" and computes the detection score and confidence for other devices.
4. **Peer interaction:** Peers use Dovecot to communicate and update internal values.
5. **Results extraction:** Experiment controller polls some peers to get the network status.

5.2.1 Networking

At the beginning of each round, the devices may choose their network status (joining the network, leaving the network, getting a new IP address), which is in turn tracked by the experiment controller. It is useful to be able to specify that different devices join at different moments in time or their IP addresses change, but the current experiments work with all devices joining at the same time, and this feature is not used.

5.2.2 Traffic

After picking a network status, the devices choose their action (the network traffic they send). Based on this action, the simulated IPS makes a detection. There are currently two actions available.

- **Benign** The device sends some communication that should be identified as *not dangerous* by an IPS. However, this may result in a falsely positive detection from some simulated IPSs.
- **Attack** The device runs an attack which should be identified as *malicious* by an IPS, but it can also be missed by some simulated IPSs.

5.2.3 Simulating the IPS

After all devices choose a network action, the IPS detection process is simulated. We need to generate two values as the detection result from the IPS for each IP address: a score and a confidence. The basic idea is that the simulated IPS takes the previous score and confidence for an IP address into account, and modifies it based on recently observed traffic. In the IPS we simulate, the score slightly improves after a **Benign** behavior and heavily drops after **Attack** behavior. The confidence changes based on the score. If the score is positive and slightly increasing, we make the confidence increase. If the score drops, it means that an attack was detected and the confidence is set to the maximal value of 1.0. On the other

hand, if the score is negative and it starts increasing, it means that attacks were detected in the past, but they stopped, and the monitored device is behaving normally again. This is suspicious, so the confidence drops to zero. Confidence stays zero, until the score reaches positive values, after which it starts increasing again.




Figure 5.1: Score, confidence and IPS_detection from a simulated IPS. The monitored device executed three attacks on rounds 8, 9, and 10, and sent benign traffic in all other rounds. At the beginning of the experiment, the device is building a good reputation by not attacking. At round 8, an attack it detected, which sends the score below zero, and the confidence to one. This continues until round 10, when the last attack is run. At round 11, the device sends benign traffic again, so the score starts increasing, and because of the suspicious change, the confidence drops to zero. Confidence stays at zero until round 15, where the score gets positive, so the confidence starts to rise again.

Figure 5.1 shows how the score and confidence values change over time, and how this affects the IPS_detection coming from the IPS. The monitored device behaved normally (chose the **Benign** action) for the first 8 rounds of the experiment, and then attacked three times on rounds 8, 9, and 10, before going back to normal behavior. In the IPS, the attack on round 8 is reflected by a significant drop in the score, along with a confidence value 1.

On round 10, the confidence still is 1, while the score has dropped all the way to 0. After that, the device starts behaving normally again, and the score slowly grows, and when it reaches positive values, the confidence starts slowly growing as well. The IPS_detection, shown in yellow, is computed by multiplying the score and confidence (see Equation 3.7).

On round 15, the score is still increasing, but it is no longer negative. This means that the monitored device is no longer considered suspicious by the IPS, and the confidence is increased. Because the confidence in the current as well as the previous round is no longer zero, the product of score and confidence grows above zero as well.

5.2.4 Peer Interaction

During the *peer interaction* step of the round, members of the Dovecot network can exchange messages and discuss which IP addresses are malicious. This is the crucial part of the round, since the trust model is applied here.

Because networking is simulated only as a set of actions in the experiments, there is not enough data to compute the reliability. The peers in the simulated network therefore all have a fixed reliability value of 1. The reliability weight w_r is set to 0.7.

5.2.5 Results Extraction

During this step, the experiment controller polls chosen benign peers and asks them about IP addresses of selected devices (both benign and malicious). The query can trigger more P2P communication if the polled peer doesn't have enough reports. The results show the Ω -score and Ω -confidence inside the polled peer at the round, as well as a score and confidence from the IPS. The experiments are time-consuming, so the raw values are saved, and different decision thresholds and other parameters are tested separately at the end of the experiment.

5.3 Types of Devices and Peers

Devices in the experiment can pursue various strategies, depending on their interests. The strategy defines how the device behaves in the network, and how it behaves in the peer group (in case it is a peer). It can also specify when the device connects, disconnects or changes IP addresses.

The strategy is a long term plan that is assigned to the device at the beginning of the experiment. The behavior of the strategy can depend on the experiment round - for example, malicious devices may choose to build good reputation for some time and attack at a later stage in the experiment. Strategy doesn't change throughout the experiment.

There are four basic profiles a device strategy may fit:

1. **Benign device outside the peer group** This is a regular device in the network, that doesn't participate in Dovecot. It's network behavior is either inactive or benign.
2. **Benign peer** This is the member of the Dovecot network. In the network behavior, it is either inactive or sending benign data. In the P2P scheme, it responds honestly to any requests from other peers.
3. **Malicious device outside the peer group** This is a device that belongs to the attacker. It can choose different network behaviors according to the strategy.
4. **Malicious peer** This device also belongs to the attacker, and can choose different network behaviors. On top of that, it is a member of the Dovecot network, and it can send fake information to other peers to manipulate their decisions.

Profiles 1 and 2 are straightforward and their strategies are very basic. The strategies of the malicious devices can be elaborate and specify different behaviors for different rounds. Note that the malicious peer may choose to join or leave the P2P network at will. The precise strategies of all participants in the experiment are defined and described later in Chapter 6.

5.4 Evaluating the Results

During the experiments, selected peers are polled about specified IP addresses, and the Ω -score, Ω -confidence as well as the score and confidence from the peer's IPS are saved. The final prediction is computed for each round based on that data, and a blocking decision is made, as explained in Section 3.4. Because the experiment environment is fully controlled, we can compare the values to the ground truth to measure accuracy. The decision is wrong in case of false positives (a benign IP address was blocked), or false negatives (a malicious IP address was not identified).

$$accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (5.1)$$

The accuracy is computed considering the overall number of detections and by extracting the values True Positive (TP), False Positive (FP), True Negative (TN) and False Negative (FN), as shown on Equation 5.1. In our experiments, we can block a device and then unblock it later because its IPS _{detection} improved. Note that in reality, this would not happen, as the value for a blocked device cannot change. For simplicity, we assume that blocking is just an output that doesn't affect future interactions.

Chapter 6

Description of each Scenario in the Experiments

In this chapter, we describe the types of experiments we executed and the goal of these experiments. There are multiple experiment types, broken down by the strategy of the malicious peers. We call these experiment types *scenarios*. The scenarios are run multiple times, each time with a different amount of malicious peers. The attack strategy of the malicious device may also slightly vary across scenarios. A precise configuration of the scenario contains the intentions, strategy and attack plans of each device and peer, and it is called a *setup*.

The scenarios reflect common attacks the protocol may face, and their design was inspired by the attack types presented in [19], which we further discussed in Sections 3.6 and 3.7. There is one *control* peer that is always benign, which is used to get results from the Dovecot network. We call it the *Observer* peer. No other peers are directly polled for results, although the Observer may communicate with them to get data from their IPSs.

Throughout the experiments, we have 10 peers and 2 devices, as shown in Table 6.1. The Observer always has the IP address 1.1.1.12. There are always two devices, a malicious device on 1.1.1.10 and a benign device on 1.1.1.11. The detections for those two devices are polled from the Observer to get the results. The peers on IP addresses 1.1.1.1 to 1.1.1.9 may be benign or malicious, and this is specified in each experiment setup. When we say that *all peers are malicious*, this means that the Observer is still benign, however all peers it communicates with are malicious.

We assume that all malicious devices belong to the same attacker, know about each other, and coordinate the attacks. This might not always be the case in reality, but it is the worst case scenario for the honest peers, and if the protocol withstands it, there is no need to simulate unrelated attackers, since they can never have a bigger impact.

We also assume that the attacker is rational, and knows how Dovecot works (running attacks lowers the trust). Malicious peers will (with the exception of scenario 5) never run the attack themselves and reveal their IP addresses. All attacks will be executed from the malicious device, and the malicious peers will use their high trust values to manipulate the Ω -score and Ω -confidence computed by the Observer.

IP address	Type	Intentions
1.1.1.12	Peer	Benign
1.1.1.1	Peer	Benign
1.1.1.2	Peer	
1.1.1.3	Peer	
1.1.1.4	Peer	
1.1.1.5	Peer	
1.1.1.6	Peer	
1.1.1.7	Peer	
1.1.1.8	Peer	
1.1.1.9	Peer	
1.1.1.10	Device	Malicious
1.1.1.11	Device	Benign

Table 6.1: IP addresses, types and intentions of experiment participants. Peers on IP addresses 1.1.1.1 to 1.1.1.9 can be either benign or malicious, depending on the setup.

The benign device is present in the network to control how the model performs on benign devices. Measuring it on one of the peers would be inconvenient since the intentions of the peers may change, and in some setups, there are no benign peers at all (except for the Observer). It is also necessary to measure the performance on a benign device in order to see how many false positives occur.

6.1 Malicious Device Attack Plans

The malicious device can choose different patterns to attack others. However, its behavior towards the Observer is the same across all attack plans. The malicious device is always benign towards the Observer in the first half of the experiment, and always attacks the Observer in the second half of the experiment. This means that the malicious device is a traitor, because there were no attacks and then the behavior changed. Having the exact same behavior across all experiments allows us to compare the results. If the P2P model is not used at all, the Observer will always give the same results for all attack plans.

Note that in several setups, some of the peers are malicious. This doesn't change the attack plan of the malicious device. In the malicious strategies that we propose, the malicious peers always report that the malicious device is benign, independent of the traffic they receive from it. Simply put, attack plans do not consider the presence of malicious peers, as that would be confusing, so a static behavior is set for every experiment round and every IP address. If a peer on that IP address is malicious, the value will be ignored.

In our experiments, the malicious device can choose from three attack plans. In all those plans, the Observer is attacked in the second half of the experiment:

- A) The malicious device targets all the peers for the entire duration of the experiment, except for the Observer. The Observer is attacked later, in the second half of the experiment. The peers should warn the Observer about the malicious device.

- B) The malicious device targets the peers one by one in sequence, eventually attacking the Observer. Again, peers should warn the Observer.
- C) The malicious device targets a subset of peers. The Observer will be added to the list of targets, and therefore get attacked, in the second half of the experiment.

Attack plan A is straightforward. Peers with IP addresses from 1.1.1.1 to 1.1.1.9 are attacked in every round of the experiment. The malicious device chooses benign action towards the Observer in the first half of the experiment, and then attacks it in the second half. This should favor Dovecot, because the Observer will receive accurate information from the network at all times. Detailed actions of the malicious device with plan A are shown in Table 6.2

round \ IP	0	1	2	3	4	5	6	7	8	9	10 - 19
Observer	B	B	B	B	B	B	B	B	B	B	A
1.1.1.1	A	A	A	A	A	A	A	A	A	A	A
1.1.1.2	A	A	A	A	A	A	A	A	A	A	A
1.1.1.3	A	A	A	A	A	A	A	A	A	A	A
1.1.1.4	A	A	A	A	A	A	A	A	A	A	A
1.1.1.5	A	A	A	A	A	A	A	A	A	A	A
1.1.1.6	A	A	A	A	A	A	A	A	A	A	A
1.1.1.7	A	A	A	A	A	A	A	A	A	A	A
1.1.1.8	A	A	A	A	A	A	A	A	A	A	A
1.1.1.9	A	A	A	A	A	A	A	A	A	A	A

Table 6.2: Attack plan A of the malicious device. *A* stands for *Being Attacked*, *B* stands for *Being Benign*.

Plan B is more complex, as there is a sliding window of attacks. Progress of the attack in time can be seen in Table 6.3. The attack on each of the peers lasts 3 rounds, which is the amount of time needed for the IPS_detection in our simulated IPS to reach the score of -1. In contrast to plan A, no peers are attacked in the second half of the experiment, which means their opinion of the malicious device will improve over time. This attack plan will be less favorable for Dovecot, because the peers will be reporting conflicting information.

Attack plan C is a variation of plan A. In C, only a subset of n peers are attacked, and the attacks last throughout the whole experiment, just like in A. The Observer is attacked only in the second half of the experiment. For convenience, we refer to the plan as C_n , based on the number n of attacked peers. For example, attack plan C6, where 6 of the peers are attacked, is shown in Table 6.4. A special case of C is C9 (all 9 peers are attacked), in which case it becomes the same as plan A. Because the subset of attacked peers doesn't change throughout the experiment, some peers will be sure of the bad intentions of the malicious device, while other peers will report it as benign. This may also result in lower performance, the goal being to see how many peers need to be reporting in order to correctly block the malicious device.

round \ IP	0	1	2	3	4	5	6	7	8	9	10	11 - 19
Observer	B	B	B	B	B	B	B	B	B	B	A	A
1.1.1.1	A	A	A	B	B	B	B	B	B	B	B	B
1.1.1.2	B	A	A	A	B	B	B	B	B	B	B	B
1.1.1.3	B	B	A	A	A	B	B	B	B	B	B	B
1.1.1.4	B	B	B	A	A	A	B	B	B	B	B	B
1.1.1.5	B	B	B	B	A	A	A	B	B	B	B	B
1.1.1.6	B	B	B	B	B	A	A	A	B	B	B	B
1.1.1.7	B	B	B	B	B	B	A	A	A	B	B	B
1.1.1.8	B	B	B	B	B	B	B	A	A	A	B	B
1.1.1.9	B	B	B	B	B	B	B	B	A	A	A	B

Table 6.3: Attack plan B of the malicious device. *A* stands for *Being Attacked*, *B* stands for *Being Benign*. In rounds 11 to 19, the malicious device attacks only the Observer on IP address 1.1.1.12.

6.2 Scenario 1 - IPS only, no P2P Network

This scenario measures IPS performance without any P2P scheme, so it is possible to establish a baseline. Because there are no parameters to adjust, and all attack plans appear the same to the Observer, this scenario only has one setup. There is an Observer on IP address 1.1.1.12, the malicious device on IP address 1.1.1.10 and a benign device on IP address 1.1.1.11. There are 20 rounds, where the malicious device waits the first 10 rounds, and then attacks the Observer for the rest of the experiment. The benign device has benign interaction with the Observer throughout the experiment. There are no other devices, as they are not important. There are also no peers, because this is a baseline experiment where no P2P mechanism is deployed. Although the benign device may seem to be unimportant, it must be part of the network so that the IPS's reaction to both types of devices can be monitored, as mentioned earlier.

There is no need to repeat the experiment with different attack plans, since in all plans, behavior of both devices towards the Observer is the same. When there are no peers that would report different attacks they have experienced, the results would be the same.

6.3 Scenario 2 - P2P is Working, no Malicious Peers

In this scenario, the Dovecot is used, and there are no malicious actors inside the peer group. There are 10 communicating peers (including the Observer), and there are the same devices as before running on 1.1.1.10 and 1.1.1.11, with the same behavior towards the Observer as they had in scenario 1. However, in this case, the malicious device also uses one of the three possible attack plans to attack other peers, so this scenario has 10 setups, labelled A, B, and C1 to C8, like the attack plans.

The goal of this scenario is to see how the other peers report to the Observer, whether the Observer will be able to make better blocking decisions because of Dovecot, and how

round \ IP	0	1	2	3	4	5	6	7	8	9	10 - 19
Observer	B	B	B	B	B	B	B	B	B	B	A
1.1.1.1	A	A	A	A	A	A	A	A	A	A	A
1.1.1.2	A	A	A	A	A	A	A	A	A	A	A
1.1.1.3	A	A	A	A	A	A	A	A	A	A	A
1.1.1.4	A	A	A	A	A	A	A	A	A	A	A
1.1.1.5	A	A	A	A	A	A	A	A	A	A	A
1.1.1.6	A	A	A	A	A	A	A	A	A	A	A
1.1.1.7	B	B	B	B	B	B	B	B	B	B	B
1.1.1.8	B	B	B	B	B	B	B	B	B	B	B
1.1.1.9	B	B	B	B	B	B	B	B	B	B	B

Table 6.4: Attack plan C6 of the malicious device in. *A* stands for *Being Attacked*, *B* stands for *Being Benign*. C6 means that 6 of the peers are attacked, and these are the peers 1.1.1.1 to 1.1.1.6.

different attack plans affect this. Plan C1 is an important setup in this case, because it shows the impact of a single peer reporting the attack.

6.4 Scenario 3 - Malicious Peers Praise the Malicious Device

Malicious peers are introduced to the network, and their strategy is to send unfair praises of the malicious device to the Observer. The experiment is run for different numbers of malicious peers (1 to 9) and attack plans A and B. There is no need to test this in attack plan C, because the results would be the same as results for attack plan A.

If we were to run this scenario with attack plan C, there would be two parameters: the number of peers that are benign, and the number of peers that are attacked. These numbers are not complementary, because malicious peers may still be attacked, and benign peers may not be attacked. However, neither of these situations makes sense.

If any malicious peers are attacked, they do not report the attacks, as they are claiming everyone is benign anyway, so there is no point in having more attacked peers than benign peers. Also, if some benign peers are present that are not attacked, they would report the malicious device as benign, which is the same thing the malicious peers do (the scores and confidences would be slightly different during the first rounds, but that is insignificant here). The only setups that make sense are when the attacked peers are same as the benign peers. This essentially means that there is a number of peers that are correctly reporting the malicious device, and the rest of the peers are reporting that it is benign, no matter what traffic they receive. And this is the exact same thing that happens when attack plan A is used with varying number of malicious peers.

This scenario tests the resilience of Dovecot against the Unfair Praises attack mentioned in Section 3.7, both individual (when only 1 peer is malicious) and collaborative (when 2 or more peers are malicious).

6.5 Scenario 4 - Malicious Peers Lie about Everything

In this scenario, the malicious peers want to block the good device by sending unfair bad reports about it to the Observer. This is the Badmouthing attack, as presented in Section 3.7. The peers are also unfairly praising the malicious device, meaning that they share the exact opposite of what the ground truth is. This will likely be the most challenging scenario, because in the extreme case when all peers are malicious, the protocol will only receive data that is entirely wrong. This experiment is only run with attack plan A.

Recall scenario 2 with attack plan C, where only a subset of peers are benign and reporting that the malicious device is attacking them. This is very similar to 4A from the point of view of the Observer. The Observer is not detecting any attacks from a device, but it is getting reports that the device is actually malicious. In case of 2C, the Observer should trust the network and should block the reported device - and in the second half of the experiment, the device reveals that it is malicious. In 4A, the peers are lying and the device stays benign throughout the experiment.

This poses a challenge, because clearly the Observer can get two sets of exact same reports, but the correct decision is different and impossible to find. Only one of these cases can be solved: if the Observer trusts the network, it will block a benign device in 4A. If the Observer does not trust the network, it will not block the malicious device in 2A.

6.6 Scenario 5 - Malicious Peers Lie about Everything and also Attack

In the previous scenarios, we have assumed that the attacker is clever, wants to keep the trust of its peers as high as possible, and the attacks are only run from the malicious device. In this scenario, the attacker is not concerned with building high trust for the peers. The malicious peers directly attack the Observer, which lowers their trust. They also use the most aggressive reporting strategy, and report that the benign device is malicious, and the malicious device is benign (same as in Scenario 4, Section 6.5).

The malicious device uses attack plan A, and the malicious peers are benign in the first 5 rounds of the experiment, and then attack the Observer until the end of the experiment. This is similar to the Traitor attack, as the peers seem benign and then start attacking and their trust lowers. The goal of this scenario is to see the performance of Dovecot when there are easy-to-discover malicious peers in the network.

Chapter 7

Experiment Results

In this chapter, we discuss the results of experiments from Chapter 6, and propose parameters that work the best across all setups.

There are two parameters we work with, which change the way the detections are combined and processed in the IPS:

- **Threshold** to block an IP address: The threshold is used to make the blocking decision. An IP address that has a prediction lower (not equal) than the threshold will be blocked. The lower the threshold, the more *evidence* the IPS needs to make the blocking decision. If the threshold is too low, the IPS will wait too long before blocking an attacker. Otherwise, if the threshold is too high, the IPS becomes strict and it might block benign devices as well.
- **IPS weight**: The IPS weight, or w_{IPS} , is used to weight the IPS_detection when computing the prediction. Note that the P2P_detection is weighted with $1 - w_{IPS}$. The weight configures how much the IPS should trust the locally captured data compared to data from the network. If the weight is set to 1, the IPS will not use the P2P_detection at all. If the weight is set to 0, the IPS will only work with P2P_detection and the IPS_detection will be ignored.

The threshold is taken from the $[-1, 1]$ interval, and the weight is taken from the $[0, 1]$ interval. We compute the accuracy and test the protocol with a 0.1 step in both variables, as this appeared to be granular enough to see trends in the data.

7.1 Scenario 1 - IPS only, no P2P Network

To understand all experiments, it is important to realize that the accuracy of the IPS on the benign and malicious devices in our experiments can never exceed 0.75. This is because the accuracy is measured using data from all 20 rounds in each experiments, and the malicious device doesn't run any attacks until halfway through the experiment. Even with the perfect IPS we simulate, it is impossible to correctly detect the attacker, because it is indistinguishable from the benign device.

It may seem unfair that the IPS is penalized for not identifying an attacker that is not attacking it. This is by design, and there are multiple reasons for this. The IPS we simulate is essentially perfect, as it has 100 percent accuracy (if the rounds where the attacker doesn't attack are not considered). With a 100 percent accuracy, there is no point in listening to a P2P network as there is nothing to be gained, but in reality no IPS is perfect, and the P2P approach may be useful. It would be complicated to simulate an imperfect IPS, so we compensate this by having a perfect IPS that doesn't see the attacker in the first 10 rounds. This also simulates the situation where other peers already know about the attacks and are sending warnings about it. It is better to think about this as identifying the *concept* of an attacker, not the attacks themselves. An IPS can only detect the attacks an attacker is sending, but a cooperating network may help a peer detect the attacker before a direct attack is even sent.

Another interesting data property is that it takes precisely half of the experiment for the values in the IPS to reach maximal possible IPS_detection value - the values for both devices grow gradually until reaching the maximal value (1) on round 9, at which point the malicious device strikes and its IPS_detection value drops. This can be seen in Figure 7.1, where IPS_detection for the benign device is shown in green, and for the malicious device in red. In the first ten rounds, both devices send benign traffic, and the IPS_detection for them grows. At round 10, the malicious device starts attacking and its IPS_detection drops to -1. The benign device continues sending benign traffic, and its IPS_detection stays at 1. Note that these values will be the same in all experiment scenarios that we run, because the behavior of the devices towards the Observer is always the same, so the results can be easily compared across experiments.

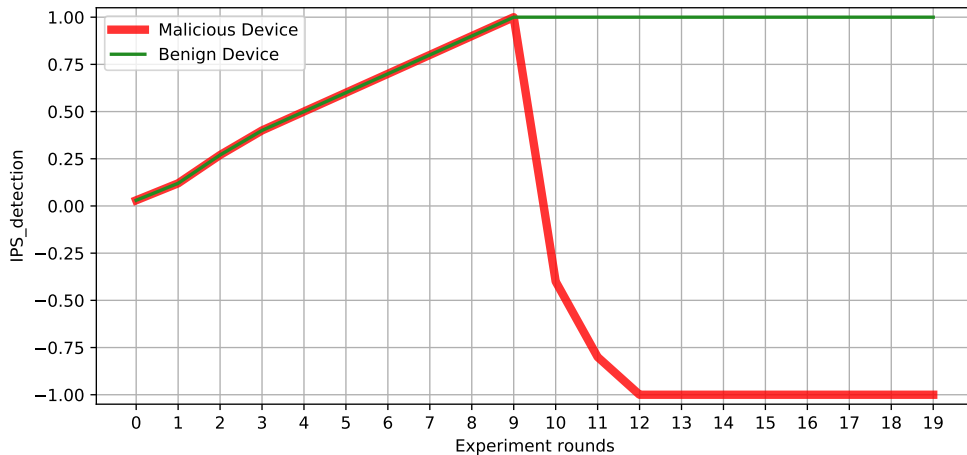


Figure 7.1: IPS_detection values made by the Observer about the benign device (green) and malicious device (red) throughout all experiments. The x-axis shows the timeline of the experiment, y-axis are the IPS_detection values for each round.

The accuracy of blocking decisions on both IP addresses (individually and combined) is shown in Table 7.1. In each cell of the table, the accuracy over all 20 rounds of the experiment is computed. The accuracy is computed either for an individual device, or for

both of them combined, and it changes based on the threshold that is used.

IP \ T	-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0
1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5
All	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table 7.1: Accuracy of the Observer’s IPS with various thresholds. Accuracy is computed separately for each of the devices, as well as for both of them together: the rows show accuracy values for the malicious device (IP address 1.1.1.10), the benign device (IP address 1.1.1.11) and combined accuracy for both devices. The highlighted columns show optimal threshold selection.

The blocking thresholds can be interpreted as the IPS sensitivity level. Note that we talk about prediction values, but in scenario 1 where no P2P is used, it holds that $\text{IPS_detection} = \text{prediction}$:

- A threshold of -1 means that the IPS is very tolerant, and needs a prediction value lower than -1 to block an IP address. The values never go below -1, so the IPS will never block any connection, because the blocking happens when the value is strictly lower than the threshold. In this case, the malicious device will be blocked with an accuracy of 0 (because it is never blocked), and the benign device will be blocked with an accuracy of 1, because not blocking it is the correct choice. Combined accuracy is therefore 0.5.
- A threshold of 1 means that the IPS is suspicious and will only allow communication with IP addresses that have detection values of 1 or higher. This means the attacking device will always be blocked, because its prediction value does not reach 1. The benign device will be blocked in the first half of the experiment, until its prediction value is 1.

For example, consider the threshold 0. From the IPS_detections on Figure 7.1, it can be seen that the IPS_detection values of the benign device are always positive. Because they are above the threshold, the benign device will never be blocked, which is the correct decision, and the accuracy on the benign device is therefore 1. The IPS_detection values of the malicious device are positive in the first 10 rounds and negative in the remaining 10 rounds. This means that the malicious device will not be blocked for 10 rounds (because the values are above the threshold), but it will be blocked in the remaining 10 rounds, because the values are below the threshold. This leads to an accuracy 0.5 on the malicious device. To get the overall accuracy, we see that there were 10 errors and 30 correct decisions, which means the accuracy when using threshold 0 is 0.75.

When analysing the combined accuracy in Table 7.1, it may seem that the IPS can set an arbitrary threshold higher than -0.3 and the accuracy will be the same. This is not the case. The accuracy computation as introduced in Section 5.4, Equation 5.1, is very primitive, and doesn’t consider the cost of individual actions. In our case, we want to avoid blocking benign devices, even at the cost of not identifying attackers in the early stages of their attack.

Setting the threshold to 0.1 or higher would lead to successful blocking of the malicious device, but also unwanted blocking of the benign device. This is why the IPS threshold

should be set somewhere between -0.3 and 0, where these errors are minimal. The significant result of this experiment is the fact that for the given malicious device behavior, the IPS reaches at most 0.75 accuracy.

7.2 Scenario 2 - P2P is Working, no Malicious Peers

The second scenario is designed to study how Dovecot will perform when no adversaries are interfering with it. There is a malicious device using various attack plans, but there are no malicious peers sending false reports. The scenario was evaluated using different decision thresholds and different weights for combining P2P_detection with the IPS_detection when computing the prediction. This is the ideal situation for the IPS and Dovecot.

7.2.1 Setup 2A - All Peers are Attacked, no Malicious Peers

In this setup, IPS_detections made by the Observer are the same as explained in scenario 1 (Figure 7.1); the IPS_detection values for both devices grow, and at round 10, the IPS_detection of the malicious device starts dropping to -1, while IPS_detection of the benign device stays at 1. The other peers, however, are targeted from the beginning until the end of the experiment. The IPS_detection values computed by their IPSs are compared to the Observer's IPS_detection values in Figure 7.2. It can be seen that other IPSs detect the malicious device sooner. IPS_detection values for the benign device (IP address 1.1.1.11) are the same in all peers, so they are not shown here.




Figure 7.2: IPS_detection values for the malicious device (1.1.1.10) made from the Observer (orange) that is not targeted until round 10, compared to IPS_detection values for the same device made by other peers that are targeted since the first round of the experiment.

The prediction values are computed from the Observer's IPS_detection and from the P2P_detection based on reports by other peers, and they are visualized in Figure 7.3. These are predictions from the experiment 2A, for a threshold value of 0, and IPS weight 0.4. An IPS

weight of 0.4 means that 40 percent of the prediction is computed from the `IPS_detection`, and the remaining 60 percent from the `P2P_detection`. This means that whatever the peers are reporting will be more significant than what the local IPS detected.

In 2A (see Figure 7.3), the peers have detected and reported the malicious device, and thanks to the low IPS weight, the prediction for the malicious device is below zero. A horizontal line on 0 (the threshold) can be drawn to separate the predictions for both of the devices perfectly, which leads to accuracy 1 if these parameters are chosen in this setup.




Figure 7.3: Prediction values for both devices throughout the experiment. The threshold was set to 0 and the IPS weight was set to 0.4. It is clear that the peers were reporting the malicious device and kept the prediction below zero. In round 10, the attacker also targets the Observer, so it becomes more sure about the blocking. Note that the prediction values never reach either 1 or -1. This is because the `P2P_detection` (which the prediction is computed from) can't reach those values.

As seen in Table 7.2, the Observer achieved much better results when cooperating with other peers. This table format is very similar to Table 7.1, but the IPS weight parameter is included as well. Only the combined accuracy for both devices is shown in the table, to keep this chapter readable. To see the accuracy of individual devices, a full table is provided in Appendix A, Table A.1. Notice the last row of the table, where the weight of the IPS is 1. This means that P2P is not used, and the IPS is relying on local data only. The accuracies in the last row are therefore same as accuracies presented in Table 7.1.

It is clear from this experiment that Dovecot was successful in informing the Observer about the devices in the network, and helped improve its security by raising the accuracy by 0.25. The improvement was the most significant if a threshold was set to 1 and the IPS weight was less than 0.5.

7.2.2 Setup 2B - Peers are Attacked in Sequence, no Malicious Peers

The attack plan in this experiment differs significantly from the plan used in Section 7.2.1, as the peers are not attacked all at once, but sequentially. This causes the `IPS_detection`

$w \backslash T$	-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.5	0.5	0.775	0.85	0.9	0.925	0.95	0.975	0.975	0.975	1.0	0.95	0.925	0.9	0.85	0.825	0.825	0.8	0.775	0.5	0.5
0.1	0.5	0.5	0.75	0.75	0.8	0.875	0.95	0.95	0.975	0.975	1.0	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5
0.2	0.5	0.5	0.725	0.75	0.75	0.75	0.85	0.95	0.975	0.975	1.0	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5
0.3	0.5	0.5	0.725	0.75	0.75	0.75	0.75	0.75	0.95	0.975	1.0	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5
0.4	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.975	1.0	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.775	0.5
0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.8	0.925	0.95	0.95	0.925	0.9	0.85	0.825	0.825	0.8	0.775	0.5
0.6	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.825	0.825	0.9	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5
0.7	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.8	0.775	0.825	0.85	0.875	0.875	0.85	0.825	0.8	0.775	0.5
0.8	0.5	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.8	0.8	0.825	0.825	0.825	0.8	0.775	0.5
0.9	0.5	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.75	0.775	0.775	0.775	0.775	0.775	0.775	0.5
1.0	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5

Table 7.2: Accuracy of decisions in setup 2A - all peers except for the Observer are attacked, no malicious peers are present. Cells where the prediction matches or outperforms the baseline are highlighted. It is clear that multiple parameter pairs exist that improve the accuracy over the IPS alone. In this *best-case scenario* of all honest peers being attacked, it may be possible to reach 100% accuracy, but is an idealistic situation. For detailed accuracy, see Table A.1.

values made by each of the peers to be different. Figure 7.4 shows how IPSs inside the peers detected the malicious device. IPS_detections for the benign device are not shown, because they are the same as in the Observer in setups 1 and 2A.

Notice that the IPS_detection values rise right after the attack stops. This is because of the way the confidence value is used to compute the IPS_detection value. If two consecutive data points differ significantly, the confidence drops to zero. And multiplying by zero causes the result (final IPS_detection) to also be zero. The Observer receives the values before processing, and can take the low confidence into account.

$w \backslash T$	-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.625	0.75	0.925	0.9	0.85	0.825	0.825	0.8	0.775	0.5	0.5
0.1	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.75	0.9	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5
0.2	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.75	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5
0.3	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.75	0.75	0.875	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5
0.4	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.75	0.85	0.875	0.85	0.825	0.8	0.775	0.775	0.5
0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.775	0.825	0.85	0.85	0.825	0.825	0.8	0.775	0.5
0.6	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.775	0.8	0.825	0.85	0.825	0.8	0.775	0.5
0.7	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.8	0.8	0.825	0.8	0.775	0.5
0.8	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.8	0.775	0.5
0.9	0.5	0.5	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5
1.0	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5

Table 7.3: Accuracy of decisions in setup 2B - each peer is attacked for three rounds, no malicious peers are present. Cells where the prediction matches or outperforms the baseline are highlighted. In this *good-case scenario*, it can be seen that Dovecot can have an 92.5% accuracy if some False Positives are allowed. The False Positives came from the fact that at the beginning of the experiments the benign device has a neutral trust that is above the prediction threshold. Only after the time passes the benign device is not detected anymore. For detailed accuracy, see Table A.2.

The results are significantly worse than in setup 2A, as is clear from Table 7.3. The best possible accuracy is 0.925, but this comes at the cost of falsely blocking the benign device. It is understandable that the benign device is blocked if the threshold is above zero. The IPS_detection values for the benign device start close to zero and grow as the experiment progresses. If the threshold is above zero, it the IPS will block the benign device in the initial rounds of the experiment, before the IPS_detection grows. If this is to be avoided,




Figure 7.4: IPS_detection values for the malicious device made by other peers. The peers with low IP addresses were attacked first, so their detections move immediately towards -1, which means malicious. The peers with higher IP addresses have the detection move to -1 later in the experiment. See that the IPS_detection immediately grows after the attacks stop, and in some peers, it grows all the way up to 1, meaning the device is considered benign again.

a threshold of 0 or lower needs to be set, as seen in Table A.2. This brings the accuracy to the same value of 0.75 that the IPS has on its own.

7.2.3 Setups 2C - A Subset of Peers is Attacked, no Malicious Peers

In this case, the malicious device targets only a selected subset of n peers throughout the experiment. For low n values, the message the Observer gets isn't very strong, and the malicious IP address will probably not be blocked.

This problem arises for $n = 1$. With $n = 1$, only one peer is targeted and correctly reports about the malicious device. Other peers are reporting that the malicious device is benign. Accuracies in this setup are shown in Table 7.4.

If we want to listen to the report and improve the accuracy by using it, we may set the threshold to 0.7 and the weight to 0 or 0.1. This results in an accuracy of 0.8, but it includes unintentionally blocking the benign device, same as in Section 7.2.2 (see Table A.3 for detailed accuracy). Parameters without this flaw reach a lower accuracy of 0.75, which is same as the IPS alone. In short, if only one peer is reporting an attack and we set the parameters to use it, we are also blocking the benign devices.

Another thing to keep in mind is that if one honest peer can manipulate the IPS, one malicious peer will be able to do it, too. This is shown to be true later in Section 7.4. By not listening to the reports, we may be losing valuable data, but there may not be a better option.

To achieve perfect accuracy (keep the accuracy of 1 on the benign device while improving accuracy of the malicious device to 1), at least 5 peers (a majority in a group of 9) need to

$w \backslash T$	-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	
0.0	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.525	0.525	0.525	0.525	0.55	0.8	0.775	0.5	0.5	
0.1	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.525	0.525	0.725	0.775	0.8	0.775	0.5	0.5	
0.2	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.775	0.775	0.775	0.775	0.5	0.5
0.3	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.775	0.775	0.775	0.775	0.5	0.5
0.4	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.5
0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.775	0.775	0.775	0.75	0.75	0.775	0.775	0.775	0.775	0.5
0.6	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.5
0.7	0.5	0.5	0.5	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.5
0.8	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5
0.9	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5
1.0	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5

Table 7.4: Accuracy of decisions in setup 2C1 - the Observer is attacked in the second half of the experiment, only 1 other peer is attacked throughout the experiment. No malicious peers are present. Cells where the prediction matches or outperforms the baseline are highlighted. For detailed accuracy, see Table A.3.

$w \backslash T$	-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	
0.0	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	1.0	0.95	0.925	0.9	0.85	0.825	0.825	0.8	0.775	0.5	0.5	
0.1	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.75	0.85	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5	
0.2	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.8	0.9	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5	
0.3	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.725	0.75	0.75	0.8	0.875	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5	
0.4	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.775	0.8	0.85	0.875	0.85	0.825	0.8	0.775	0.775	0.5	
0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.8	0.825	0.85	0.85	0.825	0.825	0.8	0.775	0.5	
0.6	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.8	0.825	0.825	0.85	0.825	0.8	0.775	0.5	
0.7	0.5	0.5	0.5	0.7	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.8	0.8	0.825	0.8	0.775	0.5	
0.8	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.75	0.775	0.775	0.775	0.775	0.8	0.775	0.5	
0.9	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.5	
1.0	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5

Table 7.5: Accuracy of decisions in setup 2C5 - the Observer is attacked in the second half of the experiment, 5 peers are attacked throughout the experiment. No malicious peers are present. Cells where the prediction matches or outperforms the baseline are highlighted. For detailed accuracy, see Table A.7.

be sending the reports. This situation is shown in Table 7.5. In this particular case, when the threshold is set to 0 and the weight of the IPS_detection is zero (which means that only the P2P_detection is used to compute the prediction), the model also correctly blocks the malicious device, and the accuracy is 1. Increasing the IPS weight lowers the prediction accuracy, but keeps the benign device from being blocked.

There are 9 possible n values to run the experiment with, and as more peers are attacked, the accuracy of 1 can be reached with a wider range of parameters. When all peers are attacked, the attack plan C9 matches attack plan A, and those results are shown in Table 7.2. Listing all setup details at this place would be confusing and not very helpful. Tables from 2C are provided in Appendix A, Tables A.3 to A.11.

7.3 Scenario 3 - Malicious Peers Praise the Malicious Device

The malicious peers are claiming that all devices in the network are benign. This means they are unfairly praising the malicious device, and accurately praising the benign device. The reports sent by them will have slightly higher impact than the reports from honest peers at the beginning of the experiment, because the honest peers send slightly increasing scores and increasing confidences, while the malicious peers outright send that the score is 1 and confidence is 1 as well.

This scenario is designed to see how the malicious peer can manipulate the predictions about the malicious device by praising it. The overall accuracy is not the main focus in this case, although it is still important to keep in mind.

7.3.1 Setup 3A - All Peers are Attacked

In this setup, peers were praising the malicious device, while all honest peers were reporting the exact opposite at all times. In Table 7.6, we see the accuracy if no honest peers are present in the network, and all 9 peers are malicious. It is possible to block the malicious device with accuracy 1, but that leads to lowering the accuracy on the benign device to 0.5 (see Table A.12), which is not acceptable. Parameters can be adjusted to keep the overall accuracy of 0.75, which lowers the accuracy on the malicious device to 0.5, but is precise about the benign device.

An accuracy of 0.5 for the malicious device essentially means that it will not be blocked by the IPS until it attacks the IPS directly. This makes sense, because there are no honest peers to report the attack sooner.

$w \backslash T$	-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0		
0.0	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	
0.1	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.75	0.5	0.5
0.2	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.5	0.5
0.3	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.5	0.5
0.4	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5
0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5
0.6	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5
0.7	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5
0.8	0.5	0.5	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5
0.9	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5
1.0	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table 7.6: Accuracy of decisions in setup 3A1 - all peers are malicious and report that the attacking device is benign. Cells where the prediction matches or outperforms the baseline are highlighted. For detailed accuracy, see Table A.12.

When two malicious peers are replaced with benign ones, accuracy increases for lower IPS weights. Table 7.7 shows this data. However, when consulting the detailed data in Table A.14, it is clear that while the attacker is identified correctly (with accuracy 1), the benign device is blocked (0.8 accuracy) and the combined accuracy is 0.9. Parameters that would keep the benign device unblocked lead to the combined accuracy of 0.75, like in Table 7.6 or in Section 7.1. It needs to be stressed that 0.75 is the best accuracy any system working alone can achieve, since the attacker is not attacking all the time. In this case, there are only two peers reporting the attack, and 7 malicious peers that report that there is no attack.

The benign device is perfectly classified when the Observer communicates with 5 malicious and 4 benign peers. At this point, as shown in Table 7.8, accuracy can reach 0.975 with IPS weight 0 and threshold 0.1. While this result looks promising, the same parameters were not successful in blocking the attacker when all 9 peers were malicious (Table 7.6), and resulted in an accuracy of 0.5. We see that at least 50 percent of the peers had to be honest to get a perfect classification of the benign device.

$w \backslash T$	-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	
0.0	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.55	0.625	0.9	0.85	0.825	0.775	0.5	0.5	
0.1	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.525	0.825	0.875	0.85	0.825	0.775	0.5	0.5	
0.2	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.775	0.8	0.85	0.85	0.825	0.8	0.5	0.5	
0.3	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.725	0.775	0.775	0.8	0.825	0.825	0.825	0.8	0.5	0.5
0.4	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.775	0.775	0.775	0.8	0.8	0.825	0.8	0.775	0.5	
0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.775	0.775	0.775	0.775	0.8	0.775	0.8	0.8	0.775	0.5	
0.6	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.7	0.725	0.725	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.775	0.8	0.775	0.5
0.7	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.75	0.775	0.775	0.775	0.5	
0.8	0.5	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.75	0.75	0.75	0.775	0.5	
0.9	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.75	0.75	0.75	0.75	0.5	
1.0	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	

Table 7.7: Accuracy of decisions in setup 3A3 - all peers except for the Observer are attacked, there are three benign peers including the Observer. The malicious peers report that the attacking device is benign. Cells where the prediction matches or outperforms the baseline are highlighted. For detailed accuracy, see Table A.14.

$w \backslash T$	-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.975	0.975	0.95	0.925	0.875	0.85	0.8	0.775	0.5	0.5
0.1	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.825	0.975	0.95	0.925	0.875	0.85	0.8	0.8	0.775	0.5	0.5
0.2	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.75	0.825	0.85	0.95	0.925	0.875	0.85	0.825	0.775	0.5	0.5	
0.3	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.8	0.825	0.9	0.925	0.875	0.85	0.825	0.8	0.5	0.5	
0.4	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.8	0.8	0.825	0.875	0.875	0.85	0.825	0.8	0.775	0.5
0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.775	0.8	0.85	0.85	0.85	0.825	0.8	0.775	0.5
0.6	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.775	0.775	0.775	0.825	0.8	0.825	0.825	0.8	0.775	0.5
0.7	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.8	0.775	0.8	0.8	0.8	0.775	0.5
0.8	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.775	0.5
0.9	0.5	0.5	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.75	0.75	0.75	0.75	0.75	0.5
1.0	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table 7.8: Accuracy of decisions in setup 3A5 - all peers except for the Observer are attacked, there are five benign peers including the Observer. The malicious peers report that the attacking device is benign. Cells where the prediction matches or outperforms the baseline are highlighted. For detailed accuracy, see Table A.16.

7.3.2 Setup 3B - Peers are Attacked in Sequence

When attacks are less reported by the network and at the same time the attacker is supported by unfair praising, the attacks are very hard to detect and block. It takes at least 5 good peers to 4 malicious peers to get an accuracy of 0.9 (see Table 7.9). This scenario is hard for the P2P network, because same as in Section 7.2.2, other peers are only targeted briefly and are reporting the malicious device with low confidence for a short time, eventually even reporting it is benign. This becomes even harder when some of the peers are malicious.

7.4 Scenario 4 - Malicious Peers Lie about Everything

All of the malicious peers are sharing scores that are opposite to the ground truth, with high confidence values from the beginning. This is challenging, because the model will perform badly when part of the decision is *outsourced* to the network. The difficulty of this setup is clear from Table 7.10, a setup with no honest peers, where the combined accuracy never exceeds the baseline value 0.75. While all peers are malicious, it is still possible to have the IPS weight set to 0.8 without impacting the accuracy.

Recall the results of scenario 2C1, shown in Table 7.4, where only one peer is targeted by the attacker and is reporting it to the Observer. For the reports to be considered by the IPS,

7.4. SCENARIO 4 - MALICIOUS PEERS LIE ABOUT EVERYTHING

$w \backslash T$	-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	
0.0	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.575	0.65	0.65	0.675	0.675	0.825	0.8	0.775	0.5	0.5	
0.1	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.525	0.65	0.65	0.9	0.875	0.825	0.8	0.775	0.5	0.5	
0.2	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.6	0.875	0.875	0.875	0.85	0.8	0.775	0.5	0.5	
0.3	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.825	0.85	0.85	0.85	0.85	0.825	0.775	0.5	0.5	
0.4	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.75	0.825	0.85	0.85	0.85	0.825	0.8	0.775	0.5	
0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.8	0.825	0.825	0.825	0.825	0.8	0.775	0.5
0.6	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.775	0.8	0.8	0.8	0.8	0.8	0.775	0.5	
0.7	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.8	0.8	0.8	0.775	0.5	
0.8	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.5	
0.9	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5	
1.0	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	

Table 7.9: Accuracy of decisions in setup 3B6 - each peer is attacked for three rounds, then the Observer is targeted. There are six benign peers including the Observer. Malicious peers report that the attacking device is benign. Cells where the prediction matches or outperforms the baseline are highlighted. For detailed accuracy, see Table A.26.

$w \backslash T$	-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.5	0.5	0.225	0.15	0.1	0.075	0.025	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.025	0.075	0.1	0.15	0.225	0.5	0.5
0.1	0.5	0.5	0.5	0.5	0.2	0.125	0.05	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.075	0.125	0.375	0.45	0.5	0.5
0.2	0.5	0.5	0.5	0.5	0.5	0.5	0.15	0.025	0.0	0.0	0.0	0.0	0.0	0.025	0.05	0.275	0.35	0.4	0.45	0.5	0.5
0.3	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.0	0.0	0.0	0.0	0.0	0.225	0.275	0.325	0.375	0.4	0.45	0.5	0.5
0.4	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.475	0.0	0.0	0.0	0.225	0.275	0.325	0.35	0.375	0.4	0.45	0.475	0.5
0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.425	0.525	0.225	0.25	0.3	0.325	0.35	0.375	0.4	0.45	0.475	0.5
0.6	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.7	0.65	0.6	0.575	0.3	0.325	0.35	0.375	0.4	0.45	0.475	0.5
0.7	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.7	0.675	0.65	0.625	0.6	0.35	0.375	0.4	0.45	0.475	0.5
0.8	0.5	0.5	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.725	0.725	0.725	0.675	0.675	0.65	0.65	0.4	0.45	0.475	0.5
0.9	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.725	0.725	0.725	0.75	0.7	0.7	0.7	0.7	0.725	0.475	0.5
1.0	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table 7.10: Accuracy of decisions in setup 4A1 - all peers except for the Observer are attacked. Only the Observer is benign. Malicious peers are reporting that the malicious device is benign, and that the benign device is malicious. This is the worst case scenario. Cells where the prediction matches or outperforms the baseline are highlighted. The best parameter pair to be used here is a weight of 0.8, and the threshold -0.1. For detailed accuracy, see Table A.30.

it is necessary to set the weight to 0.1 and the threshold to 0.7. From Table 7.10, it is clear that those parameters are not acceptable - when all peers in the network are malicious, they can exploit them and the prediction accuracy will drop to 0.325. The parameters must be chosen carefully, even at the cost of losing reports from individuals, as the primary purpose of the IPS is security.

As the proportion of honest peers in the network increases, more possible parameter combinations can be used, but the accuracy doesn't exceed 0.75 until there are 5 benign to 4 malicious peers in the network. The data for this setup is seen in Table 7.11, where there are two big changes: the highest possible accuracy jumps from 0.75 to 0.925, and there is a large area of low scores for threshold 0.

Sudden spikes when the parameters are both zero can be explained after we understand how the data looks throughout the experiment. If more peers lie than tell the truth, the prediction value for the benign device will be below zero, and the prediction value for the malicious device will be above zero (exactly the same amount, but inverted). These two prediction values create an interval. Setting the threshold outside the interval will classify one of the devices correctly, and always misclassify the other device. Setting the threshold inside the interval will misclassify both devices (because each is at the wrong side of the

$T \backslash w$	-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.925	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.1	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.75	0.825	0.875	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.2	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.725	0.875	0.825	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.3	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.725	0.75	0.725	0.775	0.825	0.8	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.4	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.725	0.775	0.8	0.8	0.8	0.5	0.5	0.5	0.5	0.5	0.5
0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.725	0.75	0.775	0.775	0.775	0.775	0.5	0.5	0.5	0.5	0.5
0.6	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.725	0.75	0.75	0.775	0.775	0.775	0.775	0.5	0.5	0.5	0.5
0.7	0.5	0.5	0.5	0.7	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.75	0.775	0.5	0.5	0.5
0.8	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.75	0.75	0.75	0.775	0.5	0.5
0.9	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.5
1.0	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table 7.11: Accuracy of decisions in setup 4A6 - all peers except for the Observer are attacked. Six devices including the Observer are benign, 4 are malicious. The malicious peers are reporting that the malicious device is benign, and that the benign device is malicious. Cells where the prediction matches or outperforms the baseline are highlighted. For detailed accuracy, see Table A.35.

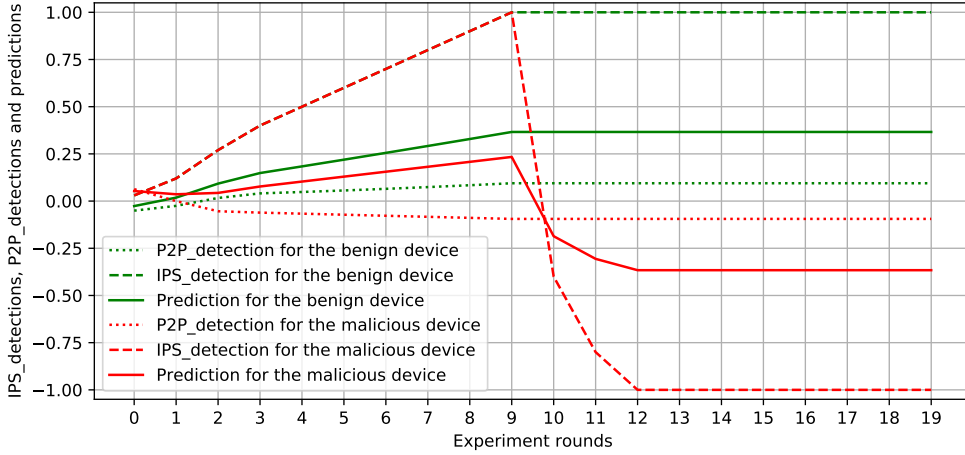


Figure 7.5: P2P_detection, IPS_detection and final prediction for both benign and malicious device. The prediction was computed with IPS weight set to 0.3.

threshold). This is clear for example in Table A.32. Note that the edge of the interval is not always sharp. This is because the Ω -score and Ω -confidence values (explained in Section 3.3) used in the computation are low at the beginning of the experiment, and needs a couple of rounds to stabilize.

Building on this, as honest peers become a majority in the network and 0 is set as the threshold, the prediction values for the two devices are at the *correct side*, and will both be blocked/not blocked correctly. On the other hand, when a threshold at a side of the interval is selected, one of the devices is still classified correctly, so the accuracy is 0.5. This is why on Table 7.11, the unexpectedly high value for threshold 0 and weight 0 appears.

It is also worth it to examine the *hole* in Table 7.11 for threshold 0 and IPS weights from 0.2 to 0.6. Figure 7.5 shows how the P2P_detection, IPS_detection and final prediction for both benign and malicious device change during the experiment rounds. Notice the P2P_detection for both devices. There are 5 benign peers that report honestly and 4 peers

that report the opposite of truth with high confidence. At the beginning of the experiment, the honest peers report slowly increasing scores with low confidences. Because they only have a weak majority, the score reported with high confidence by the malicious peers sways the P2P₋detection.

When IPS weight is set to 0.3, as the green solid lines in Figure 7.5 show, the benign device is misclassified in the first round (which is a mistake the IPS alone doesn't make), and the malicious device isn't classified as malicious until round 10 (which is same as when using the IPS alone). Because of this, the accuracy is slightly worse than accuracy of the IPS alone. If the IPS weight is set to 0, the accuracy improves, because the model is able to block the attacker correctly after the first round. Unfortunately, this also blocks the benign device in the first two rounds.

With increasing number of honest peers, the impact of lying peers fades, and if only one malicious peer is left in the network, the combined accuracy of 1 can be reached when using threshold 0 and weights 0 to 0.4. This is seen in Table 7.12.

$w \backslash T$	-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	
0.0	0.5	0.5	0.5	0.5	0.5	0.5	0.825	0.925	0.95	0.975	1.0	0.95	0.9	0.825	0.8	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.1	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.8	0.95	0.975	1.0	0.95	0.9	0.85	0.8	0.775	0.5	0.5	0.5	0.5	0.5	0.5
0.2	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.75	0.75	0.975	1.0	0.95	0.925	0.875	0.825	0.8	0.5	0.5	0.5	0.5	0.5	0.5
0.3	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.775	1.0	0.95	0.925	0.875	0.85	0.8	0.775	0.5	0.5	0.5	0.5	0.5
0.4	0.5	0.5	0.5	0.5	0.725	0.725	0.75	0.75	0.75	0.75	0.875	0.925	0.925	0.9	0.85	0.825	0.8	0.5	0.5	0.5	0.5	0.5
0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.825	0.875	0.9	0.875	0.825	0.8	0.775	0.5	0.5	0.5	0.5
0.6	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.825	0.85	0.85	0.85	0.825	0.825	0.8	0.5	0.5	0.5	0.5
0.7	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.8	0.8	0.825	0.825	0.825	0.8	0.775	0.5	0.5	0.5
0.8	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.8	0.8	0.8	0.8	0.5	0.5	0.5
0.9	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.5
1.0	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table 7.12: Accuracy of decisions in setup 4A8 - all peers except for the Observer are attacked. Eight peers including the Observer are benign, one is malicious. The malicious peers reporting that the malicious device is benign, and that the benign device is malicious. Cells where the prediction matches or outperforms the baseline are highlighted. For detailed accuracy, see Table A.37.

7.5 Scenario 5 - Malicious Peers Lie about Everything and also Attack

In these experiments, the malicious peers are also attacking. This means that the prediction accuracy should not be affected very much by their presence. Note that this scenario is not very realistic, as the experiments do not simulate blocking yet. In reality, attacks from malicious peers would be detected, the IPS would block the malicious peers, and they would not be able to send any more malicious reports (and continue attacking).

From Table 7.13, it is clear that the IPS weight should be set close to 1 to get a good accuracy. In this setup, all the peers are malicious and reporting lies during the first five rounds. After that, they start attacking and the network confidence drops to zero. This may damage the prediction value, but the impact is far less significant than if the same reports are sent by highly trusted peers (see setup 4A1, Table 7.10). The progress of the P2P₋detections is clear from Figure 7.6.

$w \backslash T$	-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.5	0.5	0.5	0.5	0.475	0.45	0.4	0.375	0.375	0.35	0.35	0.35	0.375	0.375	0.4	0.45	0.475	0.5	0.5	0.5	0.5
0.1	0.5	0.5	0.5	0.5	0.5	0.5	0.425	0.375	0.375	0.375	0.6	0.6	0.375	0.375	0.425	0.45	0.5	0.5	0.5	0.5	0.5
0.2	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.4	0.375	0.6	0.625	0.6	0.6	0.4	0.425	0.45	0.5	0.5	0.5	0.5	0.5
0.3	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.6	0.625	0.625	0.6	0.6	0.65	0.425	0.45	0.5	0.5	0.5	0.5	0.5
0.4	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.7	0.625	0.625	0.625	0.6	0.65	0.7	0.475	0.5	0.5	0.5	0.5	0.5
0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.675	0.625	0.625	0.65	0.65	0.7	0.725	0.5	0.5	0.5	0.5	0.5
0.6	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.725	0.675	0.625	0.65	0.675	0.675	0.725	0.75	0.5	0.5	0.5	0.5
0.7	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.7	0.675	0.65	0.675	0.675	0.725	0.75	0.75	0.5	0.5	0.5
0.8	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.725	0.725	0.725	0.675	0.7	0.7	0.75	0.75	0.75	0.5	0.5
0.9	0.5	0.5	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.725	0.725	0.725	0.75	0.7	0.725	0.75	0.75	0.75	0.75	0.5
1.0	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table 7.13: Accuracy of decisions in setup 5A1 - all peers except for the Observer are attacked. The Observer is the only benign peer. The peers are lying and attacking. The malicious peers are reporting the opposite to the truth about the devices, and they attack the Observer in rounds 5 to 19. Cells where the prediction matches or outperforms the baseline are highlighted. For detailed accuracy, see Table A.39.

If there are 4 benign peers and 5 malicious peers reporting to the Observer, the accuracy reaches 0.775 for IPS weight 0 and threshold -0.1. This can be seen in Table 7.14. As more malicious peers are removed and replaced with benign peers, there are more well-performing parameter pairs, with high accuracies when the IPS weight is low. Setting the threshold above zero results in higher accuracies, but they come at the cost of blocking the benign device.

$w \backslash T$	-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.775	0.75	0.75	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.1	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.75	0.65	0.8	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.2	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.7	0.775	0.825	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.3	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.725	0.75	0.7	0.7	0.8	0.8	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.4	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.725	0.725	0.725	0.8	0.8	0.5	0.5	0.5	0.5	0.5	0.5
0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.725	0.725	0.725	0.775	0.775	0.775	0.5	0.5	0.5	0.5	0.5
0.6	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.725	0.725	0.75	0.725	0.775	0.775	0.775	0.5	0.5	0.5	0.5
0.7	0.5	0.5	0.5	0.7	0.7	0.725	0.725	0.75	0.75	0.75	0.725	0.725	0.725	0.75	0.75	0.775	0.775	0.75	0.775	0.5	0.5
0.8	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.725	0.75	0.75	0.725	0.75	0.75	0.75	0.775	0.5	0.5
0.9	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.5
1.0	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table 7.14: Accuracy of decisions in setup 5A5 - all peers except for the Observer are attacked. There are 5 benign peers, including the Observer. The malicious peers are reporting the opposite to the truth about the devices, and they attack the Observer in rounds 5 to 19. Cells where the prediction matches or outperforms the baseline are highlighted. For detailed accuracy, see Table A.43.

Results in this section are not good enough, considering that the peers are openly malicious, and still at least 4 benign peers are needed to achieve a better result than the IPS alone. Here are some solutions that can be applied to further mitigate impact of malicious peers in the network. If the Ω_c is zero, there is no point in using the P2P_detection in prediction computation at all, because it is zero and it only lowers the impact of the IPS_detection. This essentially damages the performance of the IPS, even if none of the reports from malicious peers are used.

For lower Ω -confidence values, it is possible that the P2P_detection will damage the prediction. The way the prediction is computed should be adjusted, to give the IPS_detection a higher voting power if the Ω_c is low. A low Ω_s can mean the peers are not known well enough yet, or some of them are attacking. If we know that some of the peers in the network




Figure 7.6: IPS_detection and P2P_detection for the benign and malicious devices, in a network where all peers are malicious, report direct lies, and start attacking the Observer at round 5.

are malicious, it is more secure to disable the P2P trust system entirely, as it likely means someone is trying to attack the network, and it would be safer to just trust the IPS.

7.6 Results Discussion

In this section, we discuss results from experiment scenarios 2, 3, and 4, where the malicious peers (if present) do not directly attack the Observer. From the experiments, it is clear that Dovecot significantly improves the decisions in some cases. However, for Dovecot to be resilient, it is necessary to choose parameters that work in the worst case scenario as well.

To do this, we have compared accuracies across all setups and found their minimal values. These values are shown in Table 7.15, and they are the worst accuracy the parameter pair can achieve when faced with all setups in our experiment scenarios 2, 3 and 4. Right away, it is obvious that the parameters around $[0, 0]$, that looked so promising from setups with a lot of honest peers (Table 7.2, Table 7.5, Table 7.11), can not be used, because the accuracy drops as far as to 0 in less hospitable networks (Table 7.4, Table 7.6, Table 7.10). Parameter pairs that can be used so the accuracy doesn't drop below 75 percent in any of the setups are highlighted in Table 7.15.

From Table 7.15, a set of possible parameter pairs arise, all having worst-case accuracy 0.75. To avoid false positives, we choose those with accuracy 1 on the benign device. There are six of those: threshold -0.1 for IPS weight 0.8, thresholds -0.2 and -0.3 for IPS weight 0.9, and thresholds -0.3 to -0.1 for IPS weight 1. These are the parameter candidates - using them, and assuming the network simulations are accurate, it is guaranteed that the accuracy will not be worse than the accuracy of the IPS alone.

From the set of candidate parameter pairs, we want to choose a pair that not only does

		T																				
		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.15	0.2	0.3	0.45	1.0	1.0
	1.1.1.11	1.0	1.0	0.45	0.3	0.2	0.15	0.05	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.225	0.15	0.1	0.075	0.025	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.025	0.075	0.1	0.15	0.225	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.15	0.25	0.75	0.9	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	0.4	0.25	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.2	0.125	0.05	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.075	0.125	0.375	0.45	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.1	0.55	0.7	0.8	0.9	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	0.3	0.05	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.15	0.025	0.0	0.0	0.0	0.0	0.0	0.025	0.05	0.275	0.35	0.4	0.45	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.55	0.65	0.75	0.8	0.9	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.0	0.0	0.0	0.0	0.0	0.225	0.275	0.325	0.375	0.4	0.45	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.55	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.475	0.0	0.0	0.0	0.225	0.275	0.325	0.35	0.375	0.4	0.45	0.475	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.85	0.65	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.425	0.525	0.225	0.25	0.3	0.325	0.35	0.375	0.4	0.45	0.475	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.55	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.75	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.65	0.6	0.575	0.3	0.325	0.35	0.375	0.4	0.45	0.475	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.65	0.55	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.7	0.675	0.65	0.625	0.6	0.35	0.375	0.4	0.45	0.475	0.5
0.8	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.6	0.55	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.725	0.725	0.725	0.675	0.675	0.65	0.65	0.4	0.45	0.475	0.5
0.9	1.1.1.10	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0
	All	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.725	0.725	0.725	0.75	0.7	0.7	0.7	0.7	0.725	0.475	0.5
1.0	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table 7.15: Minimal accuracy across all our experiment setups. Clearly, parameters exist where the accuracy doesn't go below the baseline. The best combination seems to be $w_{IPS} = 0.8$, $T = -0.1$, meaning we trust the IPS 80 percent and the network 20 percent. Setting both parameters to zero leads to an accuracy of 0.

		T																					
		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	
0.0	1.1.1.1.0	0.0	0.0	0.55	0.7	0.8	0.85	0.9	0.95	0.95	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.8	0.7	0.55	0.0	0.0	
	All	0.5	0.5	0.775	0.85	0.9	0.925	0.95	0.975	0.975	0.975	1.0	1.0	0.975	0.95	0.925	0.9	0.85	0.825	0.775	0.5	0.5	
0.1	1.1.1.1.0	0.0	0.0	0.5	0.5	0.6	0.75	0.9	0.9	0.95	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.7	0.6	0.0	0.0	
	All	0.5	0.5	0.75	0.75	0.8	0.875	0.95	0.95	0.975	0.975	1.0	1.0	0.975	0.95	0.925	0.9	0.85	0.825	0.8	0.5	0.5	
0.2	1.1.1.1.0	0.0	0.0	0.45	0.5	0.5	0.5	0.7	0.9	0.95	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.6	0.0	0.0	
	All	0.5	0.5	0.725	0.75	0.75	0.75	0.85	0.95	0.975	0.975	1.0	0.975	0.95	0.95	0.925	0.9	0.85	0.825	0.8	0.5	0.5	
0.3	1.1.1.1.0	0.0	0.0	0.45	0.5	0.5	0.5	0.5	0.5	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.6	0.0	0.0	
	All	0.5	0.5	0.725	0.75	0.75	0.75	0.75	0.75	0.95	0.975	1.0	0.975	0.95	0.925	0.925	0.9	0.85	0.825	0.8	0.5	0.5	
0.4	1.1.1.1.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0	
	All	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.975	1.0	0.95	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	
0.5	1.1.1.1.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.6	0.85	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0	
	All	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.8	0.925	0.95	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	
0.6	1.1.1.1.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.65	0.75	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0	
	All	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.825	0.825	0.9	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	
0.7	1.1.1.1.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.6	0.65	0.75	0.85	0.95	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0	
	All	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.8	0.8	0.825	0.85	0.875	0.875	0.85	0.825	0.8	0.775	0.5	
0.8	1.1.1.1.0	0.0	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.65	0.75	0.8	0.9	0.95	1.0	1.0	1.0	1.0	1.0	
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0
	All	0.5	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.8	0.8	0.825	0.825	0.825	0.8	0.775	0.5	
0.9	1.1.1.1.0	0.0	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.6	0.65	0.65	0.75	0.8	0.85	0.9	0.95	1.0	1.0	
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.0	
	All	0.5	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.5
1.0	1.1.1.1.0	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table 7.16: Maximal accuracy across all our experiment setups. The best parameter pair seems to be $w_{IPS} = 0.4$, $T = 0$, meaning we trust the IPS 40 percent and the network 60 percent, where we get an accuracy of 1. This is however the best-case scenario, and the minimal values in Table 7.15 need to be considered when picking parameter pairs.

not make the accuracy worse, but also improves it. To get an overview of parameters with the highest potential, we have created Table 7.16 with the best possible accuracies for the given parameters, across all experiment setups.

Unfortunately, all six parameter pairs selected from Table 7.15 do not reach higher accuracies even when the best possible results are chosen. They have an accuracy of 0.75 in both the best-case and worst-case scenario. This means that the accuracy when using them is 0.75 at all times. This makes our model resilient against badmouthing, unfair praises, and other types of attacks, however, the model performs as well as the IPS alone, and therefore doesn't bring any improvement.

In the experiments, we have observed that if at least 50 percent of peers are reporting correctly, or at least 50 percent of the peers are honest, the combined IPS prediction can significantly outperform the individual IPS. If we make an assumption that the honest peers have a majority, and only consider setups with 6 or more honest peers, the worst-case scenario accuracies improve, as seen in Table 7.17. It is clear right away that there are far more acceptable parameter pairs than in Table 7.15, and better accuracies can be reached.

We want to select parameters that would help the accuracies the most. This is hard to quantify, because each parameter pair may be performing well in a different setup, and we would have to know the distribution of the setups. To work around this, we simply choose the parameters which have the highest accuracy across all setups (Table 7.16).

To avoid false positives, a good choice is to pick a threshold of 0 (as seen in Table 7.16), and set the IPS weight as low as possible, which is 0.7 in this case. Parameters $T = 0$, $w_{IPS} =$

CHAPTER 7. EXPERIMENT RESULTS

		T																				
		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.2	0.25	0.35	0.4	0.45	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.15	0.25	0.3	0.75	0.9	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.75	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.15	0.2	0.75	0.85	0.9	0.95	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.65	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.55	0.65	0.75	0.85	0.9	0.95	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.75	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.55	0.6	0.7	0.75	0.8	0.85	0.9	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.8	0.7	0.6	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.5	0.5	0.5	0.5	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.55	0.6	0.7	0.75	0.8	0.85	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.65	0.55	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.775	0.775	0.75	0.5	0.5	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.55	0.6	0.65	0.7	0.8	0.85	0.9	0.95	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.7	0.65	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.5	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5	0.5	0.5
0.8	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.8	0.7	0.65	0.6	0.55	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5	0.5
0.9	1.1.1.10	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5
1.0	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5	0.5
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table 7.17: Minimal accuracies across a subset of experiments, where it is guaranteed that there are more honest peers than malicious peers. Clearly, there are much more parameter pairs to choose from, and some of them never perform worse than the IPS alone.

0.7 will never have accuracy worse than 75 percent (in setups where the majority of peers is honest), and the best accuracy the parameters can reach is 80 percent. The only setup in which the accuracy 80 percent is reached is 2A (everyone except the Observer is attacked, and everyone is reporting it honestly, Table 7.2). This scenario is probably not very realistic, because the attack is visible to all the peers except the Observer.

When false positives are acceptable, the parameters can have higher thresholds, and an accuracy of up to 0.925 can be achieved. This happens when using threshold 0.3 and IPS weights 0.4 to 0.6. In this case, far more setups are improved, and a complete list of those is in Table 7.18. From the table, we propose to use the threshold 0.3 and the weight 0.5, because those parameters reach the best score, 0.925, in six of the setups. This means that IP addresses with predictions below 0.3 will be blocked, and the predictions are computed

setup		No malicious peers										Malicious peers praise the attacking device				Malicious peers always lie							
		peers targeted by the malicious device (1 to 8)										number of honest peers (6 to 9)				number of honest peers (6 to 9)							
w _{IPS}		2A	2B	2C1	2C2	2C3	2C4	2C5	2C6	2C7	2C8	3A6	3A7	3A8	3A9	3B6	3B7	3B8	3B9	4A6	4A7	4A8	4A9
	0.4		0.9	0.85	0.75	0.75	0.775	0.8	0.85	0.9	0.9	0.9	0.875	0.925	0.925	0.925	0.825	0.85	0.85	0.875	0.8	0.875	0.9
0.5		0.925	0.825	0.775	0.775	0.775	0.8	0.825	0.875	0.925	0.925	0.825	0.875	0.925	0.925	0.8	0.825	0.825	0.825	0.775	0.85	0.9	0.9
0.6		0.925	0.775	0.75	0.75	0.775	0.775	0.8	0.825	0.85	0.9	0.8	0.825	0.85	0.9	0.775	0.775	0.775	0.775	0.775	0.8	0.85	0.9

Table 7.18: Prediction accuracy in various setups (experiments) for three of the most promising parameter pairs. All parameter pairs share the threshold of 0.3, and the IPS weights are 0.4 (40 percent of the prediction is provided by the IPS, 60 percent by the network), 0.5 and 0.6 (60 percent of the prediction is provided by the IPS, 40 percent by the network) respectively. In all setups in this table, the honest peers have a majority in the network.

by combining equal parts of the IPS_detection and P2P_detection.

Chapter 8

Conclusion

We have presented, implemented and evaluated Dovecot, a P2P protocol that enables collaboration of IPSs in the network. Dovecot creates a Peer-to-Peer network, where each machine with an IPS is a peer. The peers can then share with each other the detections they make, and warnings about attackers in the network. Reports from peers are collected and each IPS uses them to block or not block an IP address.

Since the P2P network is open, there is no guarantee that the peers will send honest and reliable information. An adversary can join Dovecot network and send fake reports. For this reason, data received from the network cannot be fully trusted.

To address this issue, we created the Ω -Trust trust model that aggregates data from all the peers, taking into account the trust level of each peer. Peers with higher trust will have a greater say in the final prediction. The trust of each peer is computed based on the peer's behavior in the Dovecot protocol, and detection data about the peer's IP address.

The protocol was implemented to work with the Stratosphere Linux IPS (Slips) software, and will be published with Slips as a free software module. It has two parts: the Dovecot module for Slips with the built-in trust model, and Pigeon, our networking layer that makes sure the module can communicate with other peers to share IPS data. An important part Dovecot is the message format, which specifies communications between individual peers as well as between Pigeon and the module. This design allows other IPS developers to integrate Dovecot easily, without having to write the networking layer themselves.

The Dovecot protocol, including our Ω -Trust model, was evaluated with simulated IPS and attackers in scenarios that varied in attack strategy of the malicious peers, attack strategy of the malicious device, number of targeted peers, and proportion of malicious peers. In setups where the majority of peers are honest and have enough information about the attacker, Dovecot improves prediction accuracy, while in setups where all peers are malicious, the configuration of Dovecot can be changed to perform as good as an IPS alone. To make the final prediction, it is necessary to pick two parameters, a weight for the IPS and a threshold, that are used at all times. The parameters can be chosen based on the amount of malicious peers we expect to encounter in the network.

The first approach assumes that all of the peers are honest. This is realistic since creating and operating an adversarial peer is very costly for the attacker, and it is likely that the network will never encounter a malicious peer during its lifetime. We may choose a threshold

of 0 and the IPS weight of 0.7, and if there are no malicious peers in the network, this improves the accuracy of the IPS from 75% to 80% if we don't want any False Positive errors, or up to 92.5% if we allow some False Positives. In Future Work (subsection 8.1.1), we discuss how the prediction computation can be improved to avoid false negatives, which would lead to a 100% accuracy in some setups.

The second approach is to assume that there are malicious peers in the Dovecot network, but their numbers are limited. We have observed that if less than 50% of the peers are malicious, Dovecot can still improve prediction accuracy. This is done by using the same parameters as before, and leads to the same accuracy improvement of 92.5% as with the previous assumption.

The last approach is to expect that a dedicated attacker will join the Dovecot network eventually, and will try to exploit the protocol as a malicious peer. To prevent this, we choose the threshold of 0 and IPS weight of 0.8. These parameters were shown to have a minimal accuracy 75% across all evaluated setups. This means that with those parameters, the accuracy will never be worse than 75%, even when all the peers in the network are malicious and directly lying (as shown in experiment setup 4A1). However, these parameters are too conservative, and essentially lead to the data reported by the network to be ignored. Because of this, the accuracy with these parameters stays at 75% even when the network is full of honest peers that are reporting accurately. If we want to be sure that an attacker controlling all peers in the Dovecot network will not tamper with our IPS, then Dovecot can be configured with a threshold of 0 and an IPS weight of 0.8 to have the same performance of an IPS alone. We believe that the use of Dovecot can greatly improve detection in the local network if configured correctly.

8.1 Future Work

Performance of Dovecot as described in this thesis depends heavily on the parameters that are used, and the protocol can be configured as needed to account for the risk expected in the network. The results hint that having some input from external entities can be helpful. In this section, we discuss ideas that could be implemented to improve the performance and limit the impact of malicious peers. The experiments should also be expanded, to further test Dovecot in other scenarios.

8.1.1 Better Prediction Computation

The computation of the prediction value, explained in Section 3.4, Equation 3.9, is not directly part of the trust model. It is rather a part of the IPS that describes how the Ω_c and Ω_s values will be used. The straightforward formula with one parameter was chosen to allow for easy and quick experiments. It has been revealed in Section 7.5 that it is far from perfect, and it should be improved.

The biggest flaw in the current computation is that the Ω_c (confidence in the network) is simply multiplied by the Ω_s , without giving enough attention to the value itself. However, the value of Ω_c is crucial, as numbers close to zero mean that some peers in the network are not known yet, or are already known to be malicious. And having malicious peers in the

network should result in removing their reports from the Ω_s , or shutting Dovecot down, as there are probably malicious peers present that are trying to send fake reports.

The other flaw is that the computation doesn't compare the score values enough. If the local IPS is not detecting an attacker, but the peers are reporting it, the prediction should consider how serious the report is, how confident and trusted the reporters are, and maybe decide to block the device based on the reports. The current prediction computation does this correctly. However, the opposite situation is very different. If the local IPS is detecting a device as malicious and the peers are reporting that it is benign, it probably means that the peers don't have enough information. If the IPS is sure about the attack, no amount of positive reports from the network should be able to unblock the device.

A simple mechanism would be to always set the prediction value lower or equal to the `IPS_detection`. This way, if a device is detected as malicious, the `IPS_detection` will be below zero, and the prediction will stay below zero independently of positive reports from other peers. As long as we are working with a perfect IPS that does not have any false positives, this approach will be effective. For an imperfect IPS, a better prediction computation should be proposed and tested.

In our experiments, this improvement will be most apparent in setup 2C and all the setups using attack plan B. In these experiments, many peers, even if they are honest, are not affected by the attacker, and are therefore reporting that it is benign. This may lead to the IPS unblocking the attacker, which is undesirable, and could be prevented.

8.1.2 Manually Added Friends

In some networks (offices, home networks, schools etc), often the people who bring their laptops are friends who know each other and trust each other. It is only logical that in the P2P model, a peer (laptop) belonging to a friend should have higher trust than a peer randomly encountered in the network. Introducing a list of friends that are trusted more than other peers by the trust model could help, because it makes it harder for the attackers to gain a majority in the network.

A possible method for combining data from general public and friends is to have a proportion (let's say 50 percent) of the vote always go to friends, and the rest always go to the network. This way, if the peer had one friend in the network sending a report, and 10 malicious peers sending the opposite report, the attacker would still not be able to manipulate the network, as half of the vote would still go to the friend. The trust would still be used to weight reports inside the group of friends, and to weight reports from foreign peers.

Splitting the vote in half between friends and general public makes Sybil attacks very hard, if not impossible. If the peer has any friends, the attacker could spawn as many malicious peers as needed, and they would still not be able to affect more than 50 percent of the vote, and as the experiments show, having more than half of the vote is critical.

However, it must be stressed that this doesn't prevent the attacks completely. If the attacker is able to compromise some friends, majority of the network would become malicious. The attacker can also manipulate what the friends will be reporting, by carefully sending either benign or malicious communication to them, and thus making them report what the attacker wants.

8.1.3 Initial Reliability for New Peers, Cool Boot Start

Choosing the default reliability (called trust in other models) must be done carefully, because if the default is too high, malicious peers would rather generate new identities than work on improving their trust. To avoid this, the default value is usually set to the lowest possible value.

Current Dovecot implementation doesn't address this issue. New peers start with a reliability value of 1, which is the highest possible¹. This is not good, as a malicious peer may choose to create a new identity to get a better reliability. In the future, the implementation will be changed to give initial reliability 0 to new peers.

With initial reliability zero, the trust in new peers is low. As reliability grows over time, so does the peer's trust, but this takes a long time. Setting the initial reliability low damages the local IPS, because it may take a long time before reports from a peer are taken seriously. The problem here is that we want new peers to get better starting point, but at the same time make it useless for the malicious peers to reset their identity.

We think this could be done by changing the default reliability depending on the network state. This is inspired by human behavior in different environments. When a person is in a new city and has met only friendly people so far, a newly introduced person will also seem friendly. Contrary to that, if the person had their luggage stolen, was scammed at the airport and overcharged in the taxi, they will be suspicious towards any other people in that city.

Back in the network setup, this could be simulated by setting the default reliability *a little lower than the worst trust or reliability recently encountered in the network*. This would mean that if the network is trusted and all peers are working correctly, new peers would also be given a relatively high default reliability. On the other hand, if there are untrusted (malicious) peers in the network, any new peers would also not be trusted.

With this mechanism put in place, the attacker could get a new identity in hopes of earning a better trust, but in reality, the model would remember that there was a peer with low trust in the network, and set default reliability for the new identity even lower. This wouldn't damage the benign peers, as the adaptive reliability could never go lower than 0.

There are some questions to answer before putting this mechanism in place:

- how long should the network remember a peer with low trust?
- we want the "older" peers to be more trusted than new peers with the same behavior. How would the trust model react to this? Wouldn't it make treason type attacks easier?

8.1.4 Correlate Trust with Reports

Two levels of a trust model are differentiated in [19]: service and recommendations. In environments where a trust model is needed, each peer is usually providing a service (sharing a file, or communicating in a network in our case), and also providing recommendations

¹Recall that reliability describes how the peers are behaving in the Dovecot protocol, if they are spamming, and so on. It is one of the inputs in the trust computation, the other input is the score and confidence of that peer from the IPS.

(sharing good file sources, or sending reports in our case). The reputation of a peer can be computed separately with regards to each layer, and a peer may have a *service reputation* and a *recommendation reputation*. The advantage of separating the two reputations is clear for example in file sharing, as a peer with low bandwidth may have bad (slow) service, but still provide valuable and accurate recommendations.

Trust in Dovecot would be considered a *service reputation*, because service in our case means network behavior, and the trust computation in Dovecot has two inputs: the IPS detection of the remote peer (score and confidence) and the reliability, both of which are the *service*.

The obvious drawback of Dovecot is that we are using this *service reputation* (trust) to evaluate the quality of recommendations (reports). On top of that, malicious peers can easily provide good service (gain high trust) and use it to send false recommendations (reports).

Ideally, the quality of reports should also be used in the trust computation, but this process is complicated: Reports can be evaluated by comparing data from all peers and lowering trust of those whose reports don't match the majority. However, this will not work when only one peer is witnessing the attack, as instead of considering the warning, the reporter would be punished. It would also become even more dangerous for malicious peers to gain majority in the network, as they would be able to share something fake and thus lower trust the remaining honest peers have in each other. It also goes against the principle of not using anything from other peers to compute trust for a given peer.

The reports could also be compared with the local detections, but this is exploitable as well. Attackers could run an attack from a malicious device, then report it from a malicious peer, thus increasing the trust of that malicious peer. They could also attack a legitimate peer and make it report the attack, but then never attack anyone else. This would make it seem like the legitimate peer was lying about the attack.

8.1.5 Further Improvements

Promoting recent data Dovecot uses only the most recent data for a peer to compute its trust. It would also be possible to process all data. To aggregate data over a time period, a fading mechanism is proposed in [32]. The idea is to make recent information more valuable than older information. This is obviously useful, and also has the effect of preventing traitor-type attacks. A malicious peer that is building a good trust over a long time would only do one transgression and the trust would immediately drop, because the recent attack would be more significant than the long history of good behavior.

Global Networks Current Dovecot protocol and networking model only support local networks, which is not perfect, because many attacks may come from the internet, and sharing them with peers in other networks could be useful as well. The libp2p library already supports global connections, but the Dovecot would have to be adjusted. Currently, there is no limit on the number of neighbors a peer can have, and reports are sent to all known peers, so scalability would be an issue. Going global could also open new interaction possibilities, but opportunities for attacks as well.

Incentive There is currently no incentive in Dovecot to motivate peers to share honestly. This is because there is barely any cost to sharing data that had to be computed anyway, and the reference implementation shares by default. Modifying it to remove the sharing would require a person to do programming without any useful outcome. Still, if more IPSs start supporting the protocol, it might be necessary to motivate the developers to add the sharing feature.

Challenge Insertion To verify intentions of a remote peer, it is possible to insert network traffic and then ask the peer for detections about that traffic. This was proposed in [29] as a measure to monitor performance of intrusion detection systems. However, malicious peers may still be able to bypass the challenge test - a malicious peer may be configured to badmouth one specific IP address, and it will report honestly about all other IP addresses.

8.1.6 Experiments

Current experiments, although they investigate various scenarios, need to be expanded to further test Dovecot and conclude how helpful it is and what the limitations are.

One principle in many trust models is that peers that are in the network longer should have higher trust than the peers that just joined. This is not modelled in our experiments, as all peers join at the same time. Experiments should be run that investigate how late-joining affects the malicious peers, and how long the peer would have to stay before its trust is the same as trust of others.

All experiments were run using the same simulated IPS in all peers. More experiments should be run where the IPS is worse than now, better than now, or varies across the peers. This could be interesting, because it is possible that Dovecot would be more useful when the local IPS performance is low.

Current experiments use a hypothetical very primitive model of an IPS. However, the IPS is the crucial part of the experiments. The protocol should be tested on real captures, and evaluated in practice.

Bibliography

- [1] Tigist Abera, Ferdinand Brasser, Lachlan J Gunn, David Koisser, and Ahmad-Reza Sadeghi. Sadan: Scalable adversary detection in autonomous networks. *arXiv preprint arXiv:1910.05190*, 2019.
- [2] Karel Bartos, Martin Rehak, and Michal Svoboda. Self-organized collaboration of distributed ids sensors. In *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*, pages 214–231. Springer, 2012.
- [3] Juan Benet. Ipfs-content addressed, versioned, p2p file system. *arXiv preprint arXiv:1407.3561*, 2014.
- [4] Tim Bray. The JavaScript Object Notation (JSON) Data Interchange Format. RFC 8259, December 2017.
- [5] Miguel Castro, Barbara Liskov, et al. Practical byzantine fault tolerance. In *OSDI*, volume 99, pages 173–186, 1999.
- [6] Bram Cohen. The bittorrent protocol specification. https://www.bittorrent.org/beps/bep_0003.html, 2008.
- [7] Chrysanthos Dellarocas, Ming Fan, and Charles A Wood. Self-interest, reciprocity, and participation in online reputation systems. *MIT Sloan Working Papers*, 2004.
- [8] Niels Van Dijkhuizen and Jeroen Van Der Ham. A survey of network traffic anonymisation techniques and implementations. *ACM Comput. Surv.*, 51(3), May 2018.
- [9] Claudiu Duma, Martin Karresand, Nahid Shahmehri, and Germano Caronni. A trust-aware, p2p-based overlay for intrusion detection. In *17th International Workshop on Database and Expert Systems Applications (DEXA '06)*, pages 692–697. IEEE, 2006.
- [10] Nicolas Falliere. Sality: story of a peer-to-peer viral network. *White paper, Symantec Security Response*, 2011.
- [11] Carol J Fung, Olga Baysal, Jie Zhang, Issam Aib, and Raouf Boutaba. Trust management for host-based collaborative intrusion detection. In *International Workshop on Distributed Systems: Operations and Management*, pages 109–122. Springer, 2008.
- [12] Carol J Fung, Jie Zhang, Issam Aib, and Raouf Boutaba. Dirichlet-based trust management for effective collaborative intrusion detection networks. *IEEE Transactions on Network and Service Management*, 8(2):79–91, 2011.

- [13] Sebastián García and Kamila Babayeva. <https://www.stratosphereips.org/stratosphere-ips-suite>.
- [14] Sebastián García and Kamila Babayeva. <https://github.com/stratosphereips/StratosphereLinuxIPS/blob/master/README.md>.
- [15] The go programming language. Retrieved June 2, 2020, from <https://golang.org>.
- [16] Markus Jakobsson and Ari Juels. Proofs of work and bread pudding protocols. In *Secure information networks*, pages 258–272. Springer, 1999.
- [17] Sepandar D Kamvar, Mario T Schlosser, and Hector Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *Proceedings of the 12th international conference on World Wide Web*, pages 640–651, 2003.
- [18] Tor Klingberg and Raphael Manfredi. The gnutella rfc, version 0.6, 2002.
- [19] Eleni Koutrouli and Aphrodite Tsalgatidou. Taxonomy of attacks and defense mechanisms in p2p reputation systems—lessons for reputation system designers. *Computer Science Review*, 6(2-3):47–70, 2012.
- [20] Wenjuan Li, Weizhi Meng, et al. Design of intrusion sensitivity-based trust management model for collaborative intrusion detection networks. In *IFIP International Conference on Trust Management*, pages 61–76. Springer, 2014.
- [21] Lintao Liu, Shu Zhang, Kyung Dong Ryu, and Partha Dasgupta. R-chain: A self-maintained reputation management system in p2p networks. In *ISCA PDCS*, pages 131–136, 2004.
- [22] Lu Liu and Nick Antonopoulos. From client-server to p2p networking. In *Handbook of Peer-to-Peer Networking*, pages 71–89. Springer, 2010.
- [23] Protocol Labs. Libp2p - a modular network stack. <https://libp2p.io>.
- [24] Protocol Labs. Libp2p - circuit relay. <https://docs.libp2p.io/concepts/circuit-relay/>.
- [25] Protocol Labs. Libp2p - peer identity. <https://docs.libp2p.io/concepts/peer-id/>.
- [26] Protocol Labs. Libp2p - security considerations. <https://docs.libp2p.io/concepts/security-considerations/>.
- [27] Python Software Foundation. The python programming language. Retrieved May 14, 2020, from <https://python.org>.
- [28] Vivek Ramachandran and Sukumar Nandi. Detecting arp spoofing: An active technique. In *International conference on information systems security*, pages 239–250. Springer, 2005.
- [29] Martin Reháč, Eugen Staab, Volker Fusenig, Michal Pěchouček, Martin Grill, Jan Stiborek, Karel Bartoš, and Thomas Engel. Runtime monitoring and dynamic reconfiguration for intrusion detection systems. In *International Workshop on Recent Advances in Intrusion Detection*, pages 61–80. Springer, 2009.

- [30] Michael Schillo, Petra Funk, and Michael Rovatsos. Using trust for detecting deceitful agents in artificial societies. *Applied Artificial Intelligence*, 14(8):825–848, 2000.
- [31] Stratosphere. Stratosphere laboratory datasets, 2015. Retrieved June 29, 2020, from <https://www.stratosphereips.org/datasets-overview>.
- [32] Zhenhua Tan, Xingwei Wang, and Xueyi Wang. A novel iterative and dynamic trust computing model for large scaled p2p networks. *Mobile Information Systems*, 2016, 2016.
- [33] Kai Tao, Jing Li, and Srinivas Sampalli. Detection of spoofed mac addresses in 802.11 wireless networks. In *International Conference on E-Business and Telecommunications*, pages 201–213. Springer, 2007.
- [34] Cynthia Wagner, Alexandre Dulaunoy, Gérard Wagener, and Andras Iklody. Misp: The design and implementation of a collaborative threat intelligence sharing platform. In *Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security*, pages 49–56. ACM, 2016.
- [35] Li Xiong and Ling Liu. Peertrust: Supporting reputation-based trust for peer-to-peer electronic communities. *IEEE transactions on Knowledge and Data Engineering*, 16(7):843–857, 2004.
- [36] Lenny Zeltser. Free blocklists of suspected malicious ips and urls. <https://zeltser.com/malicious-ip-blocklists/>, 2019.
- [37] Quanyan Zhu, Carol Fung, Raouf Boutaba, and Tamer Basar. A game-theoretical approach to incentive design in collaborative intrusion detection networks. In *2009 International Conference on Game Theory for Networks*, pages 384–392. IEEE, 2009.

BIBLIOGRAPHY

Appendix A

Detailed accuracy tables

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	
		1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10
0.0	1.1.1.10	0.0	0.0	0.55	0.7	0.8	0.85	0.9	0.95	0.95	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.0	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.7	0.65	0.65	0.6	0.55	0.0	0.0
0.0	All	0.5	0.5	0.775	0.85	0.9	0.925	0.95	0.975	0.975	0.975	1.0	0.95	0.925	0.9	0.85	0.825	0.825	0.8	0.775	0.5	0.5	
0.1	1.1.1.10	0.0	0.0	0.5	0.5	0.6	0.75	0.9	0.9	0.95	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.1	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
0.1	All	0.5	0.5	0.75	0.75	0.8	0.875	0.95	0.95	0.975	0.975	1.0	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5	
0.2	1.1.1.10	0.0	0.0	0.45	0.5	0.5	0.5	0.7	0.9	0.95	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.2	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
0.2	All	0.5	0.5	0.725	0.75	0.75	0.75	0.85	0.95	0.975	0.975	1.0	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5	
0.3	1.1.1.10	0.0	0.0	0.45	0.5	0.5	0.5	0.5	0.5	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.3	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
0.3	All	0.5	0.5	0.725	0.75	0.75	0.75	0.75	0.75	0.95	0.975	1.0	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5	
0.4	1.1.1.10	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.4	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.55	0.0	
0.4	All	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.975	1.0	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.775	0.5	
0.5	1.1.1.10	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.6	0.85	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.5	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.8	0.7	0.65	0.65	0.6	0.55	0.0	
0.5	All	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.8	0.925	0.95	0.95	0.925	0.9	0.85	0.825	0.825	0.8	0.775	0.5	
0.6	1.1.1.10	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.65	0.75	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.6	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	
0.6	All	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.825	0.825	0.9	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	
0.7	1.1.1.10	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.65	0.75	0.85	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.7	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	
0.7	All	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.8	0.775	0.825	0.85	0.875	0.875	0.85	0.825	0.8	0.775	0.5	
0.8	1.1.1.10	0.0	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.6	0.65	0.75	0.8	0.9	0.95	1.0	1.0	1.0	1.0	1.0
0.8	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	
0.8	All	0.5	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.8	0.8	0.825	0.825	0.825	0.8	0.775	0.5	
0.9	1.1.1.10	0.0	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.6	0.65	0.65	0.75	0.8	0.85	0.9	0.95	1.0	1.0	1.0
0.9	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	
0.9	All	0.5	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.75	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.5
1.0	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0
1.0	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5	
1.0	All	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table A.1: Accuracy of decisions in setup 2A - all peers except for the Observer are attacked, no malicious peers are present.

APPENDIX A. DETAILED ACCURACY TABLES

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	
		0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.7	0.65	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.55	0.95	0.925	0.9	0.85	0.825	0.825	0.8	0.775	0.5	0.5	
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.5	0.65	0.75	1.0	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.825	0.925	0.9	0.875	0.85	0.825	0.825	0.8	0.775	0.5	0.5	
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.5	0.65	0.75	1.0	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.75	0.775	0.85	0.85	0.9	0.875	0.85	0.825	0.825	0.8	0.775	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.65	0.75	0.9	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.775	0.8	0.85	0.875	0.85	0.825	0.825	0.8	0.775	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.5	0.6	0.7	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.775	0.8	0.825	0.85	0.825	0.8	0.775	0.775	0.5	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.6	0.65	0.75	0.85	0.95	1.0	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.7	0.65	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.8	0.825	0.825	0.825	0.8	0.775	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.65	0.7	0.8	0.85	0.95	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.775	0.8	0.825	0.825	0.825	0.8	0.775	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.8	0.8	0.8	0.775	0.5	0.5
0.8	1.1.1.10	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.6	0.65	0.7	0.8	0.85	0.9	0.95	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	
	All	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.5	0.5
0.9	1.1.1.10	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	
	All	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5	0.5
1.0	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5	
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	

Table A.6: Accuracy of decisions in setup 2C4 - the Observer is attacked in the second half of the experiment, 4 peers are attacked throughout the experiment. No malicious peers are present.

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
		0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.7	0.65	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	1.0	0.95	0.925	0.9	0.85	0.825	0.825	0.8	0.775	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.5	0.7	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.75	0.85	0.95	0.925	0.9	0.875	0.85	0.825	0.825	0.8	0.775	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.6	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.8	0.9	0.925	0.9	0.875	0.85	0.825	0.825	0.8	0.775	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.45	0.5	0.5	0.7	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.6	0.7	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.75	0.775	0.8	0.85	0.875	0.85	0.825	0.825	0.8	0.775	0.775
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.6	0.7	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.65	0.6	0.55	0.0
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.8	0.825	0.85	0.85	0.825	0.825	0.8	0.775	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.6	0.65	0.75	0.85	0.9	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5
0.7	1.1.1.10	0.0	0.0	0.4	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.65	0.7	0.75	0.85	0.9	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.8	0.8	0.825	0.8	0.775	0.5
0.8	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.65	0.75	0.85	0.9	1.0	1.0	1.0	1.0		

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0		
		1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.15	0.2	0.3	0.45	1.0	1.0		
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.8	0.7	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.15	0.25	0.75	0.9	1.0	1.0		
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.7	0.6	0.0	0.0		
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.5	0.5		
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.1	0.55	0.7	0.8	0.9	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.6	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.5	0.5	
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.55	0.65	0.7	0.75	0.8	0.9	0.95	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.6	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.75	0.75	0.5	0.5	
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.55	0.65	0.7	0.75	0.8	0.9	0.95	1.0		
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.5	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0		
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.55	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0			
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0		
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95	1.0				
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0		
	All	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5	0.5
0.8	1.1.1.10	0.0	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0			
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0		
	All	0.5	0.5	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5	0.5
0.9	1.1.1.10	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0			
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.0		
	All	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5	0.5
1.0	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0				
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5			
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table A.12: Accuracy of decisions in setup 3A1 - all peers except for the Observer are malicious and report that the attacking device is benign.

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0		
		1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.15	0.25	0.4	1.0	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.65	0.55	0.0	0.0		
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.525	0.525	0.55	0.575	0.825	0.775	0.5	0.5		
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.15	0.65	0.85	1.0	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.65	0.6	0.0	0.0		
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.525	0.75	0.8	0.825	0.8	0.5	0.5		
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.6	0.75	0.85	0.95	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.8	0.75	0.65	0.6	0.0	0.0		
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.525	0.75	0.775	0.8	0.8	0.8	0.5	0.5		
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.55	0.65	0.75	0.85	0.9	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.65	0.6	0.0	0.0		
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.775	0.8	0.775	0.8	0.5	0.5		
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.6	0.65	0.75	0.8	0.9	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0		
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.775	0.775	0.775	0.8	0.775	0.5	0.5	
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.55	0.6	0.65	0.75	0.8	0.9	0.95	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0		
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.5	0.5	
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.45	0.5	0.55	0.6	0.65	0.75	0.8	0.85	0.95	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0		
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.75	0.775	0.775	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.6	0.65	0.75	0.8	0.85	0.9	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0																					

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0		
		1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.6	0.55	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.975	0.975	0.95	0.925	0.875	0.85	0.8	0.775	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.5	0.65	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.6	0.55	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.825	0.975	0.95	0.925	0.875	0.85	0.8	0.775	0.5	0.5	0.5	0.5	
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.5	0.65	0.8	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.75	0.7	0.65	0.55	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.75	0.825	0.85	0.95	0.925	0.875	0.85	0.825	0.775	0.5	0.5	0.5	0.5	
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.6	0.75	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.75	0.7	0.65	0.6	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.8	0.825	0.9	0.925	0.875	0.85	0.825	0.8	0.5	0.5	0.5	0.5	
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.5	0.6	0.7	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.85	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.8	0.8	0.825	0.875	0.875	0.85	0.825	0.8	0.775	0.5	0.5	0.5	
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.55	0.65	0.75	0.85	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.8	0.85	0.85	0.85	0.85	0.825	0.8	0.775	0.5	0.5	
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.55	0.65	0.7	0.8	0.85	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.825	0.825	0.8	0.825	0.825	0.8	0.775	0.5	0.5	
0.7	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.7	0.75	0.8	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.8	0.775	0.8	0.8	0.8	0.8	0.775	0.5	0.5
0.8	1.1.1.10	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.7	0.8	0.85	0.9	0.95	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.5	0.5	0.5
0.9	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0	
	All	0.5	0.5	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.5	0.5	0.5
1.0	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5	0.5	0.5	
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table A.16: Accuracy of decisions in setup 3A5 - all peers except for the Observer are attacked, there are five benign peers including the Observer. The malicious peers report that the attacking device is benign.

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0		
		1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.975	1.0	0.95	0.925	0.9	0.875	0.825	0.8	0.775	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.5	0.75	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.75	0.875	1.0	0.95	0.925	0.9	0.875	0.825	0.8	0.775	0.5	0.5	0.5	0.5	
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.5	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.75	0.925	0.95	0.925	0.9	0.875	0.85	0.8	0.775	0.5	0.5	0.5	0.5	
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.45	0.45	0.5	0.5	0.5	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.55	0.0	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.725	0.725	0.75	0.75	0.75	0.825	0.9	0.925	0.9	0.875	0.85	0.825	0.775	0.5	0.5	0.5	0.5	
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.5	0.5	0.5	0.65	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.8	0.85	0.875	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5	0.5	
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.6	0.7	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.775	0.8	0.825	0.85	0.875	0.85	0.825	0.8	0.775	0.5	0.5	0.5	
0.6	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.6	0.65	0.75	0.85	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.5	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.4	0.4	0																		

APPENDIX A. DETAILED ACCURACY TABLES

$w \backslash T$		T																				
		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.75	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.875	0.95	0.975	0.975	0.95	0.925	0.9	0.85	0.825	0.8	0.775	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.75	0.95	0.975	0.975	0.95	0.925	0.9	0.85	0.825	0.8	0.775	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.4	0.5	0.5	0.5	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.75	0.75	0.75	0.975	0.975	0.95	0.925	0.9	0.875	0.825	0.8	0.775	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.5	0.5	0.65	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.75	0.75	0.825	0.95	0.95	0.925	0.9	0.875	0.85	0.8	0.775	0.5	0.5	
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.5	0.5	0.55	0.75	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	
	All	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.775	0.85	0.9	0.925	0.9	0.875	0.85	0.825	0.8	0.775	
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.65	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.8	0.85	0.875	0.9	0.875	0.85	0.825	0.8	0.775	
0.6	1.1.1.10	0.0	0.0	0.0	0.4	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.65	0.7	0.8	0.9	1.0	1.0	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	
	All	0.5	0.5	0.5	0.7	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.8	0.825	0.85	0.875	0.85	0.825	0.8	0.775	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.65	0.75	0.8	0.9	0.95	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	
	All	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.775	0.8	0.8	0.825	0.825	0.825	0.8	0.775	
0.8	1.1.1.10	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	
	All	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.8	0.8	0.8	0.8	0.775	
0.9	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	
	All	0.5	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	
1.0	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	

Table A.18: Accuracy of decisions in setup 3A7 - all peers except for the Observer are attacked, there are seven benign peers including the Observer. The malicious peers report that the attacking device is benign.

$w \backslash T$		T																			
		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.65	0.85	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.825	0.925	0.95	0.975	1.0	0.975	0.95	0.9	0.8	0.875	0.85	0.825	0.8	0.775	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.45	0.5	0.6	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.8	0.95	0.975	1.0	0.975	0.95	0.925	0.875	0.85	0.825	0.8	0.775	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.45	0.5	0.5	0.5	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.75	0.75	0.975	1.0	0.975	0.95	0.925	0.875	0.85	0.825	0.8	0.775	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.5	0.5	0.5	0.5	0.55	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.75	0.975	0.975	0.95	0.925	0.9	0.85	0.825	0.8	0.775	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0
	All	0.5	0.5	0.5	0.5	0.725	0.725	0.75	0.75	0.75	0.75	0.875	0.95	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775
0.5	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.75	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55
	All	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.85	0.9	0.925	0.9	0.875	0.85	0.825	0.8	0.775
0.6	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.65	0.75	0.85	0.95	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0
	All	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.8	0.825	0.85	0.875	0.875	0.85	0.825	0.8	0.775	0.5
0.7	1.1.1.10	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.7	0.75	0.85	0.95	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0
	All	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.8	0.8	0.825	0.85	0.85	0.825	0.8	0.775	0.5
0.8	1.1.1.10	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.65	0.7	0.75	0.85	0.9	0.95	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0
	All	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.8	0.8	0.8	0.8	0.8	0.775
0.9	1.1.1.10	0.0	0.4	0.4	0.45	0															

$w \backslash T$		T																				
		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	1.1.1.1.0	0.0	0.0	0.0	0.0	0.6	0.75	0.85	0.9	0.95	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55
	All	0.5	0.5	0.5	0.5	0.8	0.875	0.925	0.95	0.975	0.975	1.0	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5
0.1	1.1.1.1.0	0.0	0.0	0.0	0.0	0.5	0.5	0.7	0.9	0.95	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55
	All	0.5	0.5	0.5	0.5	0.75	0.75	0.85	0.95	0.975	0.975	1.0	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5
0.2	1.1.1.1.0	0.0	0.0	0.0	0.4	0.5	0.5	0.5	0.65	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.7	0.75	0.75	0.75	0.825	0.95	0.975	1.0	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5
0.3	1.1.1.1.0	0.0	0.0	0.0	0.45	0.45	0.5	0.5	0.5	0.5	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.725	0.725	0.75	0.75	0.75	0.75	0.975	1.0	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5
0.4	1.1.1.1.0	0.0	0.0	0.0	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.8	0.7	0.65	0.6	0.55	0.55	0.0
	All	0.5	0.5	0.5	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.975	0.95	0.95	0.925	0.9	0.85	0.825	0.8	0.775	0.775	0.5
0.5	1.1.1.1.0	0.0	0.0	0.0	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.7	0.85	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0
	All	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.85	0.875	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5
0.6	1.1.1.1.0	0.0	0.0	0.0	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.6	0.7	0.8	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0
	All	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.8	0.825	0.85	0.9	0.9	0.875	0.85	0.825	0.8	0.775	0.5
0.7	1.1.1.1.0	0.0	0.0	0.0	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.65	0.7	0.8	0.9	0.95	1.0	1.0	1.0	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0
	All	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.8	0.8	0.825	0.85	0.85	0.85	0.825	0.8	0.775	0.5
0.8	1.1.1.1.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.6	0.65	0.7	0.8	0.85	0.95	1.0	1.0	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.8	0.8	0.825	0.825	0.8	0.775	0.5
0.9	1.1.1.1.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.6	0.6	0.65	0.7	0.8	0.85	0.9	0.95	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.5
1.0	1.1.1.1.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5

Table A.20: Accuracy of decisions in setup 3A9 - all peers except for the Observer are attacked. There is only one malicious peer and it reports that the attacking device is benign.

$w \backslash T$		T																				
		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	1.1.1.1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.15	0.2	0.3	0.45	1.0	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.8	0.7	0.55	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.1	1.1.1.1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.15	0.25	0.75	0.9	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.7	0.6	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.5	0.5
0.2	1.1.1.1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.1	0.55	0.7	0.8	0.9	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.6	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.5	0.5
0.3	1.1.1.1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.55	0.65	0.75	0.8	0.9	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.6	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.5	0.5	0.5
0.4	1.1.1.1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.55	0.65	0.7	0.75	0.8	0.9	0.95	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.5
0.5	1.1.1.1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5
0.6	1.1.1.1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.55	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.6	0.55	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5
0.7	1.1.1.1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.6	0.55
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.5
0.8	1.1.1.1.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0
	1.1.1.1.1	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.6	0.55		

APPENDIX A. DETAILED ACCURACY TABLES

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	
		0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.2	0.25	0.35	0.4	0.55	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.65	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.525	0.55	0.55	0.55	0.525	0.55	0.5	0.5	
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.2	0.25	0.35	0.9	0.95	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.65	0.6	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.55	0.55	0.55	0.775	0.775	0.5	0.5	
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.2	0.7	0.85	0.9	0.95	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.8	0.75	0.65	0.6	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.525	0.55	0.75	0.8	0.775	0.775	0.5	0.5	
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.5	0.65	0.75	0.8	0.9	0.95	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.65	0.6	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.75	0.775	0.775	0.775	0.775	0.5	0.5	
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.5	0.55	0.7	0.75	0.8	0.85	0.95	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.775	0.775	0.775	0.775	0.75	0.775	0.775	0.5	
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.55	0.6	0.7	0.75	0.8	0.85	0.9	1.0	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.775	0.775	0.775	0.75	0.75	0.775	0.5	0.5	
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.55	0.6	0.7	0.75	0.8	0.85	0.9	0.95	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.75	0.75	0.5	
0.7	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.6	0.7	0.75	0.8	0.85	0.9	0.95	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.75	0.75	0.75	0.5	
0.8	1.1.1.10	0.0	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.75	0.75	0.75	0.5	
0.9	1.1.1.10	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95	1.0	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.65	0.6	0.55	0.0	
	All	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.75	0.75	0.75	0.5	
1.0	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5	0.5	
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	

Table A.22: Accuracy of decisions in setup 3B2 - each peer is attacked for three rounds, then the Observer is targeted. There are two benign peers including the Observer. Malicious peers report that the attacking device is benign.

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
		0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.2	0.3	0.35	0.45	0.5	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.8	0.7	0.65	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.575	0.6	0.575	0.575	0.575	0.775	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.2	0.3	0.35	0.85	0.95	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.7	0.65	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.55	0.575	0.575	0.575	0.775	0.8	0.775	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.15	0.25	0.8	0.9	0.95	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.7	0.65	0.6	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.525	0.525	0.55	0.8	0.8	0.8	0.8	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.6	0.75	0.8	0.9	0.95	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.7	0.65	0.6	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.75	0.8	0.8	0.8	0.8	0.8	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.5	0.55	0.6	0.65	0.75	0.8	0.85	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.7	0.65	0.6	0.55	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.775	0.775	0.8	0.8	0.775	0.775	0.775	0.775	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.55	0.6	0.65	0.7	0.8	0.85	0.9	0.95	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.7	0.65	0.6	0.55	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.775	0.775	0.775	0.8	0.775	0.775	0.775	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.6	0.65	0.7	0.75	0.85	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.8	0.7	0.65	0.6	0.55	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.5
0.7	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.75	0.75	0.75	0.5
0.8	1.1.1.10	0.0	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0																	

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0		
		1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.15	0.45	0.55	0.6	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.575	0.7	0.725	0.7	0.875	0.85	0.825	0.8	0.775	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.3	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.625	0.95	0.925	0.875	0.85	0.825	0.8	0.775	0.5	0.5	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.65	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.8	0.9	0.925	0.875	0.85	0.825	0.8	0.775	0.5	0.5	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.5	0.6	0.85	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.75	0.775	0.875	0.9	0.9	0.85	0.825	0.8	0.775	0.5	0.5	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.55	0.65	0.85	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.775	0.85	0.875	0.875	0.85	0.825	0.8	0.775	0.5	0.5	
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.5	0.55	0.65	0.8	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.775	0.825	0.85	0.85	0.85	0.825	0.8	0.775	0.5	0.5	
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.55	0.6	0.7	0.85	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.825	0.825	0.825	0.825	0.8	0.775	0.5	0.5	
0.7	1.1.1.10	0.0	0.0	0.0	0.4	0.4	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.75	0.85	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.8	0.8	0.8	0.8	0.8	0.775	0.5	0.5	
0.8	1.1.1.10	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.7	0.8	0.85	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.775
0.9	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75
1.0	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5	0.5	
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table A.28: Accuracy of decisions in setup 3B8 - each peer is attacked for three rounds, then the Observer is targeted. There are eight benign peers including the Observer. Malicious peers report that the attacking device is benign.

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0		
		1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.2	0.5	0.6	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.6	0.7	0.725	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.85	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.875	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.5	0.6	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.75	0.75	0.9	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.6	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.75	0.75	0.875	0.875	0.875	0.85	0.825	0.8	0.775	0.5	0.5	
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.5	0.6	0.65	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.75	0.75	0.775	0.875	0.875	0.85	0.825	0.8	0.775	0.5	0.5	0.5	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.45	0.5	0.5	0.6	0.65	0.8	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.825	0.85	0.85	0.825	0.8	0.775	0.5	0.5	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.6	0.7	0.8	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.75	0.775	0.8	0.825	0.825	0.825	0.8	0.775	0.5	0.5	0.5

APPENDIX A. DETAILED ACCURACY TABLES

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.15	0.2	0.3	0.45	1.0	1.0
	1.1.1.11	1.0	1.0	0.45	0.3	0.2	0.15	0.05	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.225	0.15	0.1	0.075	0.025	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.025	0.075	0.1	0.15	0.225	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.15	0.25	0.75	0.9	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	0.4	0.25	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.2	0.125	0.05	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.075	0.125	0.375	0.45	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.1	0.55	0.7	0.8	0.9	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	0.3	0.05	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.15	0.025	0.0	0.0	0.0	0.0	0.0	0.0	0.025	0.05	0.275	0.35	0.4	0.45	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.55	0.65	0.75	0.8	0.9	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.0	0.0	0.0	0.0	0.0	0.0	0.225	0.275	0.325	0.375	0.4	0.45	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.55	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.475	0.0	0.0	0.0	0.225	0.275	0.325	0.35	0.375	0.4	0.45	0.475	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.85	0.65	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.425	0.525	0.225	0.25	0.3	0.325	0.35	0.375	0.4	0.45	0.475	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.55	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.75	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.7	0.65	0.6	0.575	0.3	0.325	0.35	0.375	0.4	0.45	0.475	0.5
0.7	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.65	0.55	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.7	0.675	0.65	0.625	0.6	0.35	0.375	0.4	0.45	0.475	0.5
0.8	1.1.1.10	0.0	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.725	0.725	0.725	0.675	0.675	0.65	0.65	0.65	0.4	0.45	0.475	0.5
0.9	1.1.1.10	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0
	All	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.725	0.725	0.725	0.75	0.7	0.7	0.7	0.725	0.725	0.475	0.5
1.0	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table A.30: Accuracy of decisions in setup 4A1 - all peers except for the Observer are attacked. Only the Observer is benign. Malicious peers are reporting that the malicious device is benign, and that the benign device is malicious.

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.15	0.25	0.4	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	0.4	0.3	0.2	0.05	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.2	0.15	0.1	0.025	0.0	0.0	0.0	0.0	0.0	0.025	0.075	0.125	0.2	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.05	0.15	0.65	0.85	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	0.3	0.15	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.15	0.075	0.0	0.0	0.0	0.0	0.0	0.025	0.075	0.325	0.425	0.5	0.5	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.6	0.75	0.85	0.95	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.4	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.2	0.0	0.0	0.0	0.0	0.0	0.05	0.3	0.375	0.425	0.475	0.5	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.55	0.65	0.75	0.85	0.9	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.0	0.0	0.2	0.275	0.325	0.375	0.425	0.45	0.5	0.5	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.6	0.65	0.75	0.8	0.9	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.85	0.55	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.425	0.475	0.225	0.25	0.3	0.325	0.375	0.4	0.45	0.5	0.5	0.5	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.55	0.6	0.65	0.75	0.8	0.9	0.95	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.65	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.675	0.65	0.55	0.275	0.3	0.325	0.375	0.4	0.45	0.475	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.75	0.8	0.85	0.95	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.8	0.7	0.55	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.675	0.65	0.625	0.575	0.325	0.375	0.4	0.425	0.475	0.5	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.6	0.65	0.75	0.8	0.85	0.9	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.7	0.6	0.55	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.725	0.725	0.725	0.675	0.625	0.65	0.4	0.425	0.45	0.5	0.5	0.5
0.8	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.75	0.8	0.85	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1																		

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	
		0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.15	0.35	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	0.4	0.25	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.2	0.125	0.0	0.0	0.0	0.0	0.0	0.0	0.075	0.175	0.5	0.5	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.15	0.8	0.95	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.4	0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.2	0.1	0.0	0.0	0.0	0.0	0.0	0.075	0.4	0.475	0.5	0.5	0.5	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.65	0.75	0.9	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.0	0.0	0.0	0.0	0.225	0.325	0.375	0.45	0.5	0.5	0.5	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.65	0.75	0.85	0.95	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.45	0.0	0.225	0.25	0.325	0.375	0.425	0.475	0.5	0.5	0.5	0.5	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.6	0.65	0.7	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.55	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.675	0.65	0.5	0.3	0.325	0.35	0.4	0.45	0.5	0.5	0.5	0.5	0.5	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.5	0.6	0.65	0.7	0.8	0.9	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.8	0.65	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.675	0.65	0.625	0.325	0.35	0.4	0.425	0.475	0.5	0.5	0.5	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.4	0.45	0.45	0.5	0.5	0.6	0.65	0.7	0.75	0.85	0.9	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.65	0.55	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.7	0.725	0.675	0.675	0.65	0.625	0.375	0.425	0.425	0.45	0.5	0.5	0.5	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.65	0.55	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.725	0.725	0.7	0.675	0.65	0.4	0.45	0.475	0.5	0.5	0.5	0.5
0.8	1.1.1.10	0.0	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.7	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.7	0.7	0.7	0.45	0.5	0.5	0.5
0.9	1.1.1.10	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.7	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.725	0.725	0.725	0.725	0.475	0.5
1.0	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5	0.5	0.5
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table A.32: Accuracy of decisions in setup 4A3 - all peers except for the Observer are attacked. Three devices including the Observer are benign. Malicious peers are reporting that the malicious device is benign, and that the benign device is malicious.

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	
		0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.2	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.15	0.0	0.0	0.0	0.1	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.6	0.85	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.05	0.0	0.0	0.3	0.425	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.65	0.8	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.475	0.0	0.225	0.325	0.4	0.475	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.65	0.75	0.85	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.85	0.55	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.65	0.525	0.325	0.375	0.425	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.6	0.7	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.75	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.675	0.625	0.6	0.35	0.4	0.45	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.55	0.6	0.7	0.75	0.85	0.95	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.7	0.7	0.675	0.65	0.375	0.425	0.475	0.5	0.5	0.5	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.55	0.6	0.65	0.75	0.8	0.9	0.95	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.8	0.7	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.725	0.725	0.7	0.675	0.675	0.4	0.45	0.475	0.5	0.5	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.55	0.6	0.65	0.7	0.8	0.85	0.9	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.6	0.55	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.7	0.7	0.725	0.725	0.75	0.725	0.725	0.725	0.7	0.7	0.7	0.7	0.7	0.45	0.5	0.5	0.5	0.5
0.8	1.1.1.10	0																					

		T																				
		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.75	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.65	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.875	0.95	0.95	0.925	0.825	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.75	0.6	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.95	0.95	0.925	0.875	0.8	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.4	0.5	0.5	0.5	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.8	0.65	0.55	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.75	0.75	0.75	0.95	0.95	0.9	0.825	0.775	0.5	0.5	0.5	0.5	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.5	0.5	0.65	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.8	0.7	0.6	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.75	0.75	0.8	0.925	0.9	0.85	0.8	0.5	0.5	0.5	0.5	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.5	0.5	0.55	0.75	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.65	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.775	0.825	0.875	0.875	0.825	0.8	0.5	0.5	0.5	0.5	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.65	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.7	0.65	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.825	0.85	0.85	0.825	0.775	0.5	0.5	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.4	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.65	0.7	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.8	0.825	0.825	0.8	0.775	0.5	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.65	0.75	0.8	0.9	0.95	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.75	0.7	0.65	0.6	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.8	0.775	0.8	0.8	0.8	0.5	0.5	0.5
0.8	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	0.0	0.0
	All	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.775	0.5	0.5	0.5
0.9	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.6	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.5
1.0	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5	0.5
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table A.36: Accuracy of decisions in setup 4A7 - all peers except for the Observer are attacked. Seven devices including the Observer are benign. Malicious peers are reporting that the malicious device is benign, and that the benign device is malicious.

		T																				
		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.65	0.85	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.8	0.65	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.825	0.925	0.95	0.975	1.0	0.95	0.9	0.825	0.8	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.45	0.5	0.6	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.8	0.7	0.6	0.55	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.8	0.95	0.975	1.0	0.95	0.9	0.85	0.8	0.775	0.5	0.5	0.5	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.5	0.5	0.5	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.65	0.6	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.75	0.75	0.975	1.0	0.95	0.925	0.875	0.825	0.8	0.5	0.5	0.5	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.4	0.45	0.5	0.5	0.5	0.5	0.55	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.7	0.6	0.55	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.775	1.0	0.95	0.925	0.875	0.85	0.8	0.775	0.5	0.5	0.5	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.45	0.45	0.5	0.5	0.5	0.5	0.75	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.7	0.65	0.6	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.725	0.725	0.75	0.75	0.75	0.75	0.875	0.925	0.925	0.9	0.85	0.825	0.8	0.5	0.5	0.5	0.5
0.5	1.1.1.10	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.75	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.825	0.875	0.9	0.875	0.825	0.8	0.775	0.5	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.65	0.75	0.85	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.75	0.7	0.65	0.6	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.825	0.85	0.85	0.85	0.825	0.8	0.5	0.5	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.7	0.75	0.85	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.8	0.7	0.65	0.6	0.55	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.8	0.8	0.825	0.825	0.825	0.8	0.775	0.5	0.5	0.5
0.8	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.65	0.7	0.75	0.85	0.9	0.95	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.0	0.0	0.0
	All	0.5																				

APPENDIX A. DETAILED ACCURACY TABLES

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
		0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.6	0.75	0.85	0.9	0.95	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	All	0.5	0.5	0.5	0.5	0.8	0.875	0.925	0.95	0.975	0.975	1.0	0.95	0.9	0.875	0.825	0.8	0.775	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.5	0.5	0.7	0.9	0.95	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.7	0.65	0.6	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.75	0.75	0.85	0.95	0.975	0.975	1.0	0.95	0.925	0.875	0.85	0.825	0.8	0.5	0.5	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.4	0.5	0.5	0.5	0.65	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.7	0.65	0.6	0.55	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.7	0.75	0.75	0.75	0.825	0.95	0.975	1.0	0.95	0.925	0.9	0.85	0.825	0.8	0.775	0.5	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.45	0.45	0.5	0.5	0.5	0.5	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.725	0.725	0.75	0.75	0.75	0.75	0.975	1.0	0.95	0.925	0.9	0.875	0.825	0.8	0.775	0.5	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.725	0.725	0.75	0.75	0.75	0.75	0.775	0.975	0.95	0.925	0.9	0.875	0.85	0.825	0.8	0.5	0.5	0.5
0.5	1.1.1.10	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.7	0.85	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.85	0.875	0.925	0.9	0.875	0.85	0.825	0.8	0.775	0.5	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.7	0.8	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.8	0.75	0.65	0.6	0.55	0.0	0.0	0.0
	All	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.8	0.8	0.85	0.9	0.9	0.85	0.825	0.8	0.775	0.5	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.65	0.7	0.8	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.0	0.0	0.0
	All	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.8	0.825	0.85	0.85	0.85	0.85	0.8	0.5	0.5	0.5
0.8	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.65	0.7	0.8	0.85	0.95	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.8	0.8	0.825	0.825	0.8	0.775	0.5
0.9	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.6	0.65	0.7	0.8	0.85	0.9	0.95	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.775	0.5
1.0	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5	0.5
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table A.38: Accuracy of decisions in setup 4A9 - all peers except for the Observer are attacked. Only one of the peers is malicious. It is reporting that the malicious device is benign, and that the benign device is malicious.

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
		0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.7	0.75	0.75	0.8	0.9	0.95	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	0.95	0.9	0.8	0.75	0.75	0.7	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.475	0.45	0.4	0.375	0.375	0.35	0.35	0.35	0.375	0.375	0.4	0.45	0.475	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.5	0.7	0.75	0.75	0.85	0.9	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	0.85	0.75	0.75	0.75	0.7	0.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.425	0.375	0.375	0.375	0.6	0.6	0.375	0.375	0.425	0.45	0.5	0.5	0.5	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.5	0.7	0.8	0.85	0.9	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.8	0.75	0.75	0.75	0.7	0.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.4	0.375	0.6	0.625	0.6	0.6	0.4	0.425	0.45	0.5	0.5	0.5	0.5	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.5	0.5	0.5	0.8	0.85	0.9	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.75	0.75	0.75	0.7	0.7	0.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.6	0.625	0.625	0.6	0.6	0.65	0.625	0.45	0.5	0.5	0.5	0.5	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.45	0.5	0.5	0.5	0.5	0.65	0.9	0.95	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.75	0.75	0.75	0.75	0.7	0.65	0.5	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.7	0.625	0.625	0.625	0.6	0.65	0.7	0.475	0.5	0.5	0.5	0.5	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.8	0.95	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.85	0.75	0.75	0.75	0.7	0.6	0.5	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.675	0.625	0.625	0.65	0.65	0.7	0.725	0.5	0.5	0.5	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.85	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.75	0.75	0.75	0.75	0.7	0.6	0.5	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.725	0.675	0.625	0.65	0.675	0.675	0.725	0.75	0.5	0.5	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.8	0.9	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.75	0.7	0.65	0.6	0.5	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.7	0.675	0.65	0.675	0.675	0.725	0.75	0.75	0.5	0.5	0.5
0.8	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.6	0.6	0.65	0.7	0.85	0.9	1.0	1.0	1.0
	1.1.1.11	1.0	1.0																			

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0		
		1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.7	0.75	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.0	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.8	0.75	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	All	0.5	0.5	0.5	0.5	0.5	0.5	0.475	0.4	0.375	0.35	0.0	0.35	0.375	0.4	0.45	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.75	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.1	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.75	0.75	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.1	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.45	0.375	0.375	0.575	0.25	0.375	0.4	0.45	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.5	0.6	0.85	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.2	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.75	0.75	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.2	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.375	0.6	0.625	0.6	0.3	0.425	0.45	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.45	0.5	0.5	0.5	0.75	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.3	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.75	0.75	0.7	0.65	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.3	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.6	0.625	0.6	0.575	0.375	0.45	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.5	0.5	0.55	0.6	0.85	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.4	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.85	0.75	0.75	0.7	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.4	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.675	0.625	0.625	0.625	0.6	0.425	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.55	0.6	0.75	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.5	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.75	0.75	0.7	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.5	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.675	0.625	0.65	0.65	0.675	0.475	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0
0.6	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.8	0.75	0.75	0.7	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.6	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.7	0.65	0.65	0.675	0.675	0.75	0.475	0.5	0.5	0.5	0.5	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.6	0.6	0.65	0.8	0.9	0.95	1.0	1.0	1.0	1.0	1.0
0.7	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.0	0.0	0.0	0.0	0.0	0.0
0.7	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.7	0.675	0.7	0.675	0.675	0.725	0.75	0.475	0.5	0.5	0.5	0.5	0.5
0.8	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.75	0.85	0.9	0.95	1.0	1.0	1.0	1.0
0.8	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.0	0.0	0.0	0.0	0.0
0.8	All	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.725	0.725	0.725	0.725	0.7	0.725	0.75	0.75	0.475	0.5	0.5	0.5	0.5
0.9	1.1.1.10	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.75	0.8	0.85	0.9	0.95	1.0	1.0	1.0
0.9	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.0	0.0	0.0	0.0
0.9	All	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.725	0.725	0.75	0.75	0.725	0.75	0.75	0.75	0.75	0.475	0.5	0.5	0.5
1.0	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0	1.0
1.0	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5	0.5	0.5
1.0	All	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table A.40: Accuracy of decisions in setup 5A2 - all peers except for the Observer are attacked. There are 2 benign peers, including the Observer. The malicious peers are reporting the opposite to the truth about the devices, and they attack the Observer in rounds 5 to 19.

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0		
		1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.7	0.75	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.0	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.75	0.75	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.0	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.375	0.375	0.0	0.35	0.375	0.45	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.5	0.5	0.75	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.1	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.75	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.1	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.475	0.375	0.6	0.25	0.375	0.45	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.5	0.65	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.2	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.75	0.75	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.2	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.6	0.625	0.6	0.325	0.45	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.45	0.5	0.5	0.5	0.55	0.8	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.3	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.75	0.75	0.65	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.3	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.675	0.625	0.625	0.6	0.4	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.5	0.5	0.5	0.6	0.65	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.4	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.75	0.7	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.4	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.675	0.625	0.65	0.625	0.45	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.6	0.65	0.8	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0
0.5	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.8	0.75	0.7	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0.5	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.7	0.65	0.675	0.675	0.7	0.475	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.6	0.65	0.7	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0
0.6	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.75	0.7	0.6							

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0		
		1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.6	0.75	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.7	0.55	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.8	0.875	0.925	0.85	0.775	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.5	0.5	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.75	0.65	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.75	0.75	0.9	0.875	0.825	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.5	0.5	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.7	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.75	0.725	0.925	0.85	0.8	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.4	0.45	0.5	0.5	0.5	0.5	0.75	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.75	0.65	0.55	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.725	0.8	0.875	0.825	0.775	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.65	0.85	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.8	0.7	0.65	0.55	0.0	0.0	0.0	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.725	0.775	0.825	0.85	0.825	0.775	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.6	0.7	0.85	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.75	0.7	0.6	0.55	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.75	0.725	0.75	0.775	0.8	0.85	0.8	0.775	0.5	0.5	0.5	0.5	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.65	0.75	0.85	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.7	0.65	0.6	0.0	0.0	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.75	0.725	0.75	0.775	0.775	0.775	0.8	0.8	0.5	0.5	0.5	0.5	0.5	0.5
0.7	1.1.1.10	0.0	0.0	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.5	0.6	0.65	0.7	0.8	0.85	0.95	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.7	0.65	0.55	0.0	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.8	0.775	0.5	0.5	0.5	0.5	0.5
0.8	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.6	0.6	0.65	0.75	0.8	0.85	0.95	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.9	0.85	0.8	0.7	0.65	0.6	0.55	0.0	0.0	0.0	0.0	
	All	0.5	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.75	0.75	0.775	0.775	0.5	0.5
0.9	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.95	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.8	0.75	0.7	0.65	0.6	0.55	0.0	0.0	0.0	
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.775	0.775	0.5
1.0	1.1.1.10	0.0	0.4	0.4	0.45	0.45	0.45	0.5	0.5	0.5	0.5	0.55	0.6	0.65	0.65	0.7	0.75	0.8	0.9	0.95	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.85	0.85	0.8	0.75	0.7	0.6	0.55	0.5	0.5	0.5	
	All	0.5	0.7	0.7	0.725	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Table A.44: Accuracy of decisions in setup 5A6 - all peers except for the Observer are attacked. There are 6 benign peers, including the Observer. The malicious peers are reporting the opposite to the truth about the devices, and they attack the Observer in rounds 5 to 19.

$w \backslash T$		-1.0	-0.9	-0.8	-0.7	-0.6	-0.5	-0.4	-0.3	-0.2	-0.1	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0		
		1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11	All	1.1.1.10	1.1.1.11
0.0	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.65	0.75	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.65	0.55	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.825	0.875	0.95	0.95	0.925	0.825	0.775	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.1	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.4	0.5	0.6	0.9	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.85	0.7	0.6	0.55	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.7	0.75	0.8	0.95	0.95	0.925	0.85	0.8	0.775	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.2	1.1.1.10	0.0	0.0	0.0	0.0	0.0	0.45	0.5	0.5	0.5	0.5	0.5	0.65	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.8	0.65	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.5	0.5	0.725	0.75	0.75	0.75	0.95	0.95	0.9	0.825	0.8	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.3	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.65	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.95	0.9	0.8	0.7	0.65	0.55	0.0	0.0	0.0	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.5	0.5	0.5	0.5	0.5
0.4	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.55	0.8	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.75	0.65	0.55	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.775	0.85	0.925	0.875	0.825	0.8	0.775	0.5	0.5	0.5	0.5	0.5	0.5
0.5	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.65	0.8	0.95	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.7	0.65	0.55	0.0	0.0	0.0	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.775	0.825	0.875	0.85	0.825	0.775	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.6	1.1.1.10	0.0	0.0	0.0	0.4	0.45	0.45	0.5	0.5	0.5	0.5	0.65	0.7	0.8	0.9	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
	1.1.1.11	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.9	0.85	0.8	0.75	0.65	0.55	0.0	0.0	0.0	0.0	0.0	0.0	
	All	0.5	0.5	0.5	0.7	0.725	0.725	0.75	0.75	0.75	0.75	0.75	0.775	0.775	0.8	0.825	0.825	0.8	0.775	0.5	0.5	0.5	0.5	0.5

