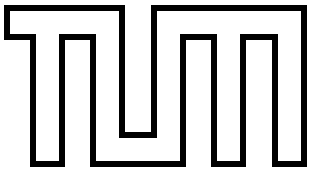

DEPARTMENT OF INFORMATICS

TECHNICAL UNIVERSITY OF MUNICH

Bachelor's Thesis in Informatics

Technical Analysis of the Tangle in the IOTA-Environment

Bennet Breier

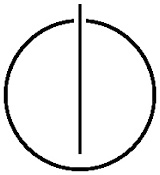

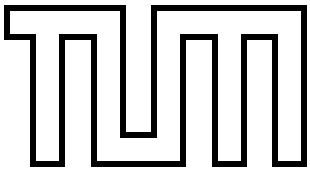

DEPARTMENT OF INFORMATICS

TECHNICAL UNIVERSITY OF MUNICH

Bachelor's Thesis in Informatics

Technical Analysis of the Tangle in the IOTA-Environment

Technische Analyse des Tangle in der IOTA-Umgebung

Author:	Bennet Breier
Supervisor:	Prof. Dr. Florian Matthes
Advisor:	Patrick Holl
Date:	November 6, 2017

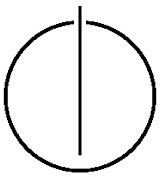

I confirm that this bachelor’s thesis is my own work and I have documented all sources and material used.

Munich, 6 November 2017 Bennet Breier

Acknowledgments

First and foremost, I am very grateful for the discussions with my advisor Patrick Holl. He helped me set the right frame for the thesis and motivated me to dig deeper into the technology.

Many thanks to Prof. Dr. Florian Matthes who proposed this interesting research topic and enabled me to complete this thesis at his chair.

A very special gratitude goes out to Paul Handy who has not only freed up time to give a long interview, but also kept answering my questions via Slack. He was always happy to shed light on the implementation and functioning of the tangle whenever I reached out to him.

Equally open about the ideas behind IOTA did Alexander Renz help me by having an inspiring interview with him. Thank you very much, Alexander!

Special thanks also to Serguei Popov who replied to my mathematical questions instantly even though he was traveling around the globe.

Abstract

Distributed Ledger Technology has become popular with the creation of the Bitcoin Blockchain in 2007. However, when considering its applicability for the Internet of Things, issues like scalability, transaction fees, offline accessibility, and quantum security have not been resolved. The IOTA Foundation has developed and published an alternative to Blockchain which claims to resolve these issues: the Tangle.

As its major part, this thesis first examines the ingredients of the Tangle and the theoretical aspects of the Peer-to-Peer protocol. Then it compares Tangle and Blockchain along the following characteristics: data structure, scalability, immutability, fee structure, offline capability, privacy, and energy consumption. Furthermore, it gives an overview over the IOTA Foundation and examines some specificities of its Tangle-implementation.

The analysis is based on literature and online research as well as two expert interviews with members of the IOTA Foundation. It allows for a better technical understanding and assessment of the Tangle-technology.

Contents

Outline of the Thesis

Chapter 1: Introduction

This chapter motivates the analysis of Distributed Ledgers and specifically the Tangle by illustrating an exemplary use-case. The example builds on two main attributes of Distributed Ledgers: trust & immutability. Subsequently, it points out the specific advantages of the Tangle, focusing on the Internet of Things. After stating the three research questions, it ends with briefly explaining the research approach and structure of this thesis.

Chapter 2: Analysis of the Tangle

At first, this chapter introduces the basic principles of the protocol in the P2P-network. Then it digs into the ingredients necessary to build and use a Tangle and explains how to interact with the Tangle. Afterwards, it answers why the ledger is stable and how it can be protected against a selection of four kinds of double-spending attacks: Simple Large Weight Attack, Parasite Chain Attack, Splitting Attack, and Sybil Attack. The next sections deal with certain resulting attributes which the Tangle possesses, namely offline capability, scalability, lack of privacy, and quantum resistance. The chapter ends with summarizing two add-on features building on the Tangle, namely Smart Contracts and Masked Authenticated Messaging.

Chapter 3: Comparison of Tangle and Blockchain

After a concise dive into the basic principles behind the Blockchain-technology, this chapter compares the Tangle to the Blockchain alongside seven attributes. Not always are aspects completely identical or different in either system, but sometimes can only parallels be established between the two technologies which are partly similar, partly different.

Chapter 4: The Tangle in the IOTA-environment

The IOTA Foundation initiated and developed the only existing Tangle at the time. Since they are so closely intertwined, there is a short section about the Foundation itself. Subsequently, this chapter addresses concepts which are implemented in the currently existing version of the Tangle, but do not belong to an abstracted view on the technology.

Chapter 5: Outlook

This chapter gives a short overview over further developments and improvements by IOTA as well as open research questions. They could be addressed in further works of research. The most important points involve smart contracts, cryptographic security, concrete use-cases, and energy consumption.

Chapter 6: Conclusions

In theory, the Tangle is a mature Distributed Ledger, which reaches consensus in a scalable fashion without the need for fees. It is based on profound mathematical research. Focusing on IoT-use-cases makes perfect sense, because the Tangle can solve exactly the problems IoT-systems would have if they used the Blockchain. In practice however, energy consumption, proven security, and the closed-source Coordinator pose considerable challenges. Most of all, the Tangle will need expanding adoption so that the system can build on the benefits of High Load.

1. Introduction

Distributed Ledger Technology (DLT) has become popular since the release of Bitcoin in 2007. What it enables is a shared and cryptographically secured database which stores the transaction (tx) of tokens from one address to another. No party owns the database, instead every node in the network keeps a copy of it. On top of that, no party can change the history of txs retrospectively. This makes txs between parties possible even though they might not trust each other and even without an intermediary, because it solves the Byzantine Generals Problem described by Lamport et al. [32] in 1982.

One key concept that a distributed ledger can achieve is digital identification. In a distributed infrastructure that the internet is, authentication must be ensured in order to establish trust between two or more parties. Such a ledger could store a tamper-proof digital identity of Internet-of-Things (IoT)-devices and even human individuals. The history of the digital identity history would be immutable. This means that an identity would need to be approved by some authority only once initially, but could yield trust throughout the rest of its life. New mobility concepts emerge where modes of transportation are connected to enhance mobility itself. Since different companies offer various mobility services, the ecosystem would benefit strongly from authentication across sub-systems. Currently, each different car- or bike-sharing service, for example, requires the user to register, perhaps handing in the same documents, like driver's license, several times. This takes time and must happen in advance instead of allowing for ad-hoc access. Equally beneficial would a distributed ledger be for mere Machine-to-Machine communication, especially for sensors, e.g. vibration sensors in cars, selling data to multiple companies [12]. In an attempt to build such a comprehensive database for the Internet of Things, the IOTA Foundation was established.

The most common DLT that could enable comprehensive digital identification to date is called Blockchain. The two assets which have gained the highest market capitalization are Bitcoin and Ether which are transacted over the Bitcoin- and the Ethereum-Blockchain, respectively. However, despite its popularity the Blockchain-technology has two inherent short-comings when it comes to high-load scenarios, such as in Internet-of-Things:

- limited scalability
- fees per tx

Moreover, the two named assets apply cryptographic functions which are not quantum-secure (yet). Facing these three issues the IOTA Foundation has developed the Tangle-technology, in order to equip IoT with a suitable distributed ledger. A scalable ledger is necessary, because millions of devices will send an enormous number of txs. This

includes micro-txs which only send minimal amounts of tokens/money, as well as zero-value txs, like messages or sensor-data. This means that if a device had to pay a fee for every piece of data it sends to peers, numerous IoT-scenarios would be uneconomic. The Tangle has very similar features as the Blockchain, but allegedly solves the before-mentioned problems.

This thesis aspires to answer the following three research questions:

1. What is the theoretical foundation of the Tangle?
2. What are the key similarities and differences between Tangle and Blockchain?
3. How does IOTA use and advance the Tangle in its environment?

With regard to the research approach, for historically proven concepts the thesis collects information from academic papers. However, beyond Google Scholar, much of the gathered information stems from questions in IOTA's online communities, i.e. IOTA Forum, Slack-Team, StackExchange, reddit, various blog-posts and interviews published online. Nevertheless, the thesis is profoundly based on two expert interviews with Paul D. Handy, core-developer since December 2016, and Alexander Renz, business advisor of IOTA since July 2017, conducted specifically for this thesis, as well as direct code review.

This thesis is structured along the above-mentioned research questions. The first chapter after the introduction covers data structure and the protocol, which make up the Tangle. It answers questions like: how to interact with the Tangle, how does the Tangle grow, which features result from it, how are attacks deflected, and what is advanced functionality building on the basics. In the second chapter, the reader gets to know how the Blockchain works differently or similarly like the Tangle. For this it is important to have understood most of the previous sections. Finally, before giving a conclusion and further outlook, chapter 4 shortly examines the IOTA Foundation as well as concepts which do not inhere the Tangle-technology, but are IOTA-specific.

2. Analysis of the Tangle

This chapter digs into the technical concepts making up the technology. After examining the general protocol and data structure, the section on tx processing elaborates on how to interact with the Tangle. Then there is a brief mathematical analysis on why the system remains stable and an explanation on how the Tangle reaches consensus. When the deflection of certain attacks has been examined, the following sections cover why the Tangle is offline-capable, scalable, and quantum resistant, and why it allows for hardly any privacy. The chapter concludes by describing two advanced concepts building on the basics.

2.1. Overview & definitions

The Tangle is an asynchronous protocol in a peer-to-peer (P2P) network to facilitate trustless tx-processing on an immutable, distributed, decentralized ledger secured by cryptographic measures. It scales with the number of txs and is potentially able to run without an inherent crypto-coin [17]. The following paragraph describes the general functioning of the Tangle: The Tangle is comprised of sites/txs and edges which form a Directed Acyclic Graph. The P2P-network consists of nodes. An edge indicates that one site directly approves another. A path symbolizes indirect approval.

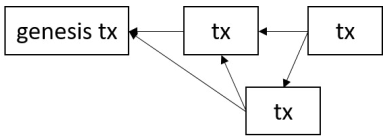

Figure 2.1.: Genesis of the Tangle

The entire supply of tokens gets generated in the genesis tx. This leads to a fixed supply of IOTA-tokens of exactly $(3^{33}-1)/2 = 2,779,530,283,277,761$ [17]. It was chosen for hardware efficiency and in order to be a sufficiently high number in Internet of Things scenarios. A startup called Jinn Labs was originally developing hardware that was going to use 32 trits. But, the amount is planned to be increased in some future snapshot, because it is now going to use 81 trits, which could accommodate $(3^{82}-1)/2$ tokens [13]. The genesis tx is the only tx allowed to send tokens from an address with 0 tokens to another address and approve no other tx. No tokens are created afterwards. As a consequence, the currency is deflationary, because tokens can get lost, e.g. if a user loses their access-key, called "seed". The current implementation by IOTA uses 81-character-seeds composed only of the characters A-Z and the number 9, so 27 possible characters. A tx is the transfer of x tokens from address A to address B. If x = 0, then the tx is also called message. An address is the public key of an asymmetrical encryption scheme (e.g. RSA, ECC), with the private key as the access to the address's funds. Like in standard Public-key

cryptography, all tx data is inputted into a cryptographic function using the private key which produces the signature. The signature is appended to the actual data to proof that this tx is indeed issued by A. B can check this by decrypting the signature using A's public key and comparing it to the sent data. Theoretically, the whole tx could be encrypted with B's public key to ensure privacy if A has access to it, e.g. if there is some PKI-infrastructure. However, other nodes could not read the tx and verify whether the amount sent from A to B is valid [38].

Whenever a tx should be added to the Tangle, a node must first create a bundle including this tx. Then it selects two txs it wants this tx to validate and reference in the Tangle. In fact, it has to do this for every tx inside the bundle. It is valid to choose only one tx, like shown in Figure 2.1. However, in almost all cases, nodes select two txs because this maximizes the likability for the issued tx to be selected itself in the MCMC algorithm (section 2.4.2). The process of selecting two txs is called Tip Selection (section 2.4.2) and is at its best if every node uses roughly the same tip selection strategy. In order to verify a tx, each tx directly and indirectly referenced by this tx must be checked whether the transacted funds are sufficient - a process called validation (section 2.4.3). If a tx is added without verifying the referenced txs, no other tx will reference this tx if it is invalid and so it would get orphaned. For referencing, the issuing node must perform some Proof of Work (PoW)(section 2.4.5). Then the tx can be broadcasted to the neighbors.

In short, a node must perform the following five steps to issue a tx (section 2.4):

1. bundling & signing
2. tip selection
3. validation
4. PoW
5. publishing

Apparently, an issuing node does not pay any fee for performing this tx, except for the processing costs. When comparing it to the Blockchain, one usually focuses on the cost for the PoW. A node must participate in the propagation of txs, because otherwise its neighbors will quickly realize that it is "lazy" and therefore remove it from the list of neighbors. Consequently, it would not be able to issue txs itself anymore. [44]

The following definitions will be used throughout this thesis:
All txs in the Tangle = $\mathbb{T}$
For txs $x, y \in \mathbb{T}$:
y approves x directly = y references x = $x \hookleftarrow y$
y approves x indirectly = $x \not\hookleftarrow y \wedge \exists z \in \mathbb{T} : z \hookleftarrow y \wedge x \rightsquigarrow z$
y approves x = y approves x directly or indirectly = $x \rightsquigarrow y$
tip = $x \in \mathbb{T}$, where $\nexists y \in \mathbb{T} : x \rightsquigarrow y$

own weight of a tx x = measure for the amount of work done for x (in theoretical analyses set to 1) = $w(x) = 1$
For efficiency in trinary systems, own weights can only assume values 3^n, with $n \in \mathbb{N}$ [38]. Currently, in the implementation by IOTA the own weight (= minWeightMagnitude) is set to a constant $81 = 3^4$ for hardware efficiency.
Cumulative weight $H(t)$ of x at time t:

$$H(t) = w(x) + \sum_{y \in Approving} w(y) = 1 + |Approving|, \quad Approving = \{z \in \mathbb{T} \,|\, x \rightsquigarrow z\}$$

[38]

An example for a Tangle is depicted in Figure 2.5. Currently, IOTA uses RocksDB as a persistence provider which is a simple Key-Value-Store with Bloom-Filters. This is because IOTA wanted an embedded database which "was performant and allowed for concurrent access" [13].

2.2. Seeds, Addresses & Transactions

It is important to understand what exactly seeds, addresses, and txs are. A seed should be generated as randomly as possible, because this acts like the password to a user's account. There is no specific algorithm behind this generation. In the end, the seed only has to be a combination of 81 chars of letters A-Z and the number 9. The wallet software uses the IOTA API[1] which generates a new address from the seed and a unique address index along a standard algorithm. The address index starts at 0 and is incremented for each new address. "Address Index can be any positive integer" [18], including zero. The wallet prohibits creating a new address before the old one has been attached. This is due to the way it searches for all addresses associated with this wallet / seed. It does not store the balances of all addresses by itself but receives this information from the Tangle. The way the search works is that it starts with the address index of 0 and keeps incrementing until it does not find a corresponding tx anymore. This implies that if an address is not attached, it cannot be found. Therefore, the wallet must prohibit gaps of unattached addresses, otherwise it would not find all associated addresses. This explains why after a snapshot (section 4.5) the account balance appears as 0. Since the wallet software does not store associated addresses itself, it must recalculate associated addresses starting at address index 0 and look for associated balances in the snapshot-file first [18].

Attaching an address means publishing this new address on the Tangle. The wallet achieves this by performing a tx of 0 tokens to this address. To be more precise, the bundle contains only one output-tx associated with this address and does not need to be signed by an input address. For this step, it must reference two txs on the Tangle and publish the bundle, which means attaching this tx to the Tangle. Put differently, attaching a tx to the Tangle means validating and doing the PoW for two txs on the Tangle. Besides, users can help the security of the network by spamming the Tangle

[1] https://github.com/iotaledger

with 0-value txs, because this requires doing benevolent PoW, i.e. PoW that works for the main-Tangle, and would therefore require an attacker to acquire more resources.

An address carries a number of tokens and can be used both as input to a bundle (i.e. the tokens are spent) or as output (i.e. the address receives tokens). Because of the W-OTS (section 2.4.4) the address may be used as input only once, since the address must be accompanied by the correct signature. Accordingly, an address is basically the aggregation of one or numerous UTXOs whose combined value is the address balance. To calculate the account balance of a user, the balances of all addresses belonging to the user must be summed together. This will be discussed further in the next section.

Figure 2.2 summarizes the anatomy of a tx stored on the Tangle. The hash is calculated from all fields below, including the nonce. Since *trunkTransaction* and *branchTransaction* are the hashes of the referenced txs, all information from these two txs are included into the hash of the referencing tx. What is interesting is that there is only one address-field. This is because a tx is either used as the input or the output of a bundle and does not act like a transfer of tokens all by itself. This means that the address belongs to the sender when used as an input, or to the recipient when used as an output. Furthermore, the *signatureMessageFragment* is the longest part of the tx using 2187 trytes. This is because it stores the signature of the sender, but can also store message content. The three fields *currentIndex*, *lastIndex*, and *bundle* are for locating the tx within a bundle [44].

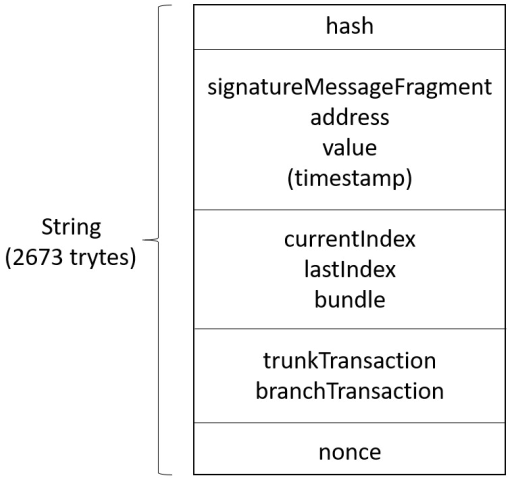

Figure 2.2.: One tx on the Tangle

2.3. Account-balances vs. UTXO-scheme

On the Tangle, the two concepts account-balances and UTXO-scheme are difficult to distinguish.

First of all, one user does not only have one address, but can generate several addresses from his seed (section 2.2). This makes the scheme look like a UTXO scheme,

but it is in fact using account balances. Therefore, one address corresponds to one account, while all addresses derived from one seed combined make up one user account. IOTA does not use an actual UTXO-scheme, because all incoming txs of one address are summed up to one single account balance. As a consequence, you do not input single txs into a bundle and sign them, but you input the entire balance of your address and sign it. In other words, there are three layers of complexity: txs, addresses, and users/seeds. Moreover, the format of snapshots has the balance of one address as an attribute (section 4.5). Due to the clear facts for account-balances, one of the founders of IOTA only stated: "IOTA uses a UTXO-like scheme." [42]

The UTXO-scheme is a concept usually discussed for Blockchains. In the case of UTXOs, you would select a sufficient number of UTXOs you own and input them into an output UTXO owned by the recipient. This kind of bundling happens when working with the Tangle, too, because you can aggregate multiple addresses/account balances into one bundle. Part of this entirety is sent to the designated recipient's address, the remainder is outputted as another UTXO to the current user's address. Since a user cannot simply send an arbitrary amount from his balance, this scheme is not called account-balance. Put another way, there are only two layers of complexity: UTXOs and users/addresses. In this sense, addresses on the Tangle are very similar to UTXOs, while txs are not.

Like in a UTXO-scheme, the user can benefit from parallelization. He can generate multiple addresses for different threads. Also, it allows for some privacy advantages. Unlike UTXOs, account balances are not stateless and since the wallet-software tracks related txs on the Tangle automatically, it is not more complicated [15].

In conclusion, transfers must deplete the entire balance of utilized accounts. Because addresses are used as inputs instead of single txs, this implies that IOTA uses an account-balances-scheme.

2.4. Transaction Processing

When a node wants to send a message or tokens from one account to another, it must perform five steps: bundling & signing, tip selection, validation, PoW, and publishing. These steps are explained in the following.

2.4.1. Bundling & Signing

For a node to make a tx it must have the private keys to addresses storing sufficient funds. Since one address might not suffice, a transfer of tokens is published as a so-called bundle, combining multiple txs as inputs into one atomic tx. In fact, a node can issue txs only within a bundle. So, a node needs to choose which of its addresses to use for the transfer, eventually constructing the bundle. These addresses are added as input addresses, which implies that their values must be negative. The recipient's address serves as the output of the bundle, so the value of this tx is positive. If not the entire input-value is exhausted, then the rest is stored in a fresh address controlled by

the sender, called remainder address. The corresponding tx has a positive value. This explains the crucial condition which any bundle must fulfill: The sum over the values of all included txs must be 0. As a side-note, this makes bundles possible which include only one output-tx with value 0, without any input txs. This complex scheme is applied because addresses may be used as inputs only once, since the signature is generated with the Winternitz-One-Time-Signature-Scheme (section 2.4.4). The necessity of a remainder address is the main reason why bundling was devised. If we recall the structure of a tx, we realize that each tx has only one address field. Knowing that a tx is either used as input or output, this makes a lot of sense. However, in terms of terminology it might make more sense to rename txs into sites and bundles into txs.

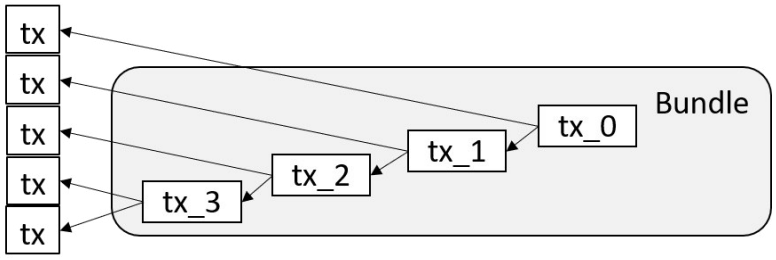

Figure 2.3.: Example for a bundle of four txs stored on the Tangle

For each tx included in the bundle, the issuing node must do PoW. In figure 2.3, this would amount to doing PoW four times. By design, the next tx in a bundle is always stored as the trunk-tx of the previous tx. In this way, the whole bundle can be extracted from the Tangle only by finding the first tx of the bundle and traversing down the trunk-tx fields [44].

The bundled txs might by chance not be approved by other txs on the Tangle. In this case the node has two options: Either it reattaches the tx or it promotes it. Reattaching/Replaying means issuing the same tx again with a different hash, because different trunk- and branch-txs are selected, meaning to redo the PoW for newly selected tips. This might be necessary if the issued tx does not get approved by enough txs to get confirmed, so the user attaches it to the Tangle in another spot. The previously attached tx is left behind and gets orphaned.

Promoting refers to attaching a tx on top of another tx in order to increase the probability of approval. One leg approves the previous tx, the other leg approves another tip. This roughly doubles the chance for the previous tx to be approved (directly or indirectly). Furthermore, if a node has doubts that its tx was propagated through the network, e.g. due to connectivity issues, it can rebroadcast the tx. This means "sending the exact same transaction [...] again" [22]. However, the same problem could arise due to the dynamic throttling mechanism described in section 2.4.5.

2.4.2. Tip Selection

When the addresses for the new bundle are chosen, two tips per tx in the bundle must be selected for being referenced. The algorithm by which a node performs this selection is called Tip Selection Strategy. Theoretically, we could speak of tx selection, because in general it is allowed to approve any tx. However, if nodes selected normal txs for approval, the number of tips would grow indefinitely and the system would not be stable. So nodes should be incentivized to only approve tips.

Using the "Random Tip Selection Strategy" means that all tips have equal probability to be selected, like in a Laplace Experiment. The random strategy is interesting only for mathematical analyses because in practice it does not protect against double-spend txs dealt with in section 2.7.

Therefore, a more sophisticated approach is proposed in the whitepaper and currently implemented in the majority of nodes. The Markov Chain Monte Carlo (MCMC) algorithm should make sure that tips are selected non-deterministically along the path of the largest cumulative weight in a reasonable amount of time. A node has to run it every time before it can attach a tx.

First of all, the node places a certain number of N, e.g. 10, walkers sufficiently deep in the Tangle. As a heuristic it can use, for example, the set of txs having cumulative weights between W and 2W. It can also instead use the set of txs which it has received in the time period $[t_0, 2t_0]$. Each walker then moves along the references toward the set of tips, i.e. the walker can walk from tx x to y iff $x \hookleftarrow y$.

The transition probabilities can be calculated as follows:

$$P_{xy} = \frac{e^{-\alpha(H_x - H_y)}}{\sum_{z:x \hookleftarrow z} e^{-\alpha(H_x - H_y)}} \quad x, y \in \mathbb{T}, x \hookleftarrow y$$

The idea is that the smaller the difference between H_x and H_y, the less cumulative weight the walker loses if he chooses y (instead of z or another tx). The proposed density function is a valid Riemann density, because the following two properties hold true [26]:

1. $f \geq 0$
2. $\sum_{y:x \hookleftarrow y} P_{xy} = 1$

Figure 2.4.: Transition probabilities in MCMC

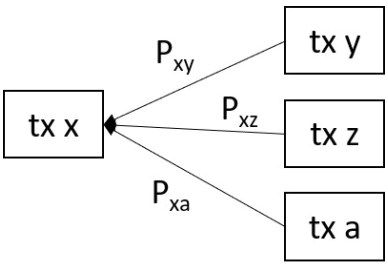

The first property holds true, because $e^x > 0, \forall x \in \mathbb{R}$. The second is fulfilled, because $\sum_{y:x \hookleftarrow y} P_{xy}$ adds up to the same sum as the denominator, which can be canceled to 1.

$H_x > H_y$ holds true for all $x, y \in \mathbb{T}$ so that the exponent is always negative. The parameter α scales the exponential function horizontally. The lower alpha, the more stretched the exponential function and the less sharp the differences in probabilities. If a node would like to rather frequently follow the path of largest cumulative

weight, then it should set α rather high. It should not be too high, however, to retain non-determinism. This formula is specified in the whitepaper, but is not yet proven to be optimal. For example, one could use x^{-3} instead of the exponential function. In general, it should be a rapidly decreasing function, because the steeper this function, the more large cumulative weight is prioritized over small cumulative weight, eventually making the algorithm a bit more deterministic. So the Tangle growth becomes more directed.

When the walkers have reached a tip, the node can decide among them. To make selection faster, the node could simply select the walker which found a tip first. But then old tips might be favored, which are naturally closer to where the walkers started their walks. Consequently, it might be best to either wait, for example, for the third walker or to ignore any walker who finished "too fast" (this criterion is not further specified in the whitepaper) [38, p.20].

Nodes are not required to follow one Tip Selection Strategy, but they benefit from aligning their strategy to some non-deterministic "reference" rule for the following reasons:

- Nodes want to maximize the velocity by which, in turn, their own issued txs are approved. If the probability distribution of another strategy deviates very much from the default one, then these txs are in general less likely to be selected by subsequent txs. In other words, this keeps the probability of selecting a "bad" tip small.
- If all nodes were able to follow a deterministic strategy, all of them would end up choosing the same tip and so there would be much competition for subsequent approvals.
- A superior strategy might include finding out the "best" tips, but this is hardly possible, because plenty of walkers would have to be calculated which is time consuming. Then, once a result is found, the Tangle has already changed.

2.4.3. Validation

Validation is done after the tips are selected. Verifying a site on the tangle is a recursive process, because each site builds on a verified sub-tangle. The result of validation again is a valid sub-tangle. In order to verify a tx, a node must ensure that this tx only references verified txs and that it does not set any account balance to negative. This means that an address A can be used as an input only if it stores sufficient funds. Furthermore, a valid tx must include a hash which fulfills the Hashcash requirements, meaning that the PoW was done on this tx (section 2.4.5). But how can validation of a Tangle be done correctly if there is no time-order of txs? Suppose address A has 10 tokens and addresses B and C have 0 tokens. Then two txs happen:
A $\rightarrow$ B (10)
B $\rightarrow$ C (10)
In this scenario, a node would only approve the second tx if it sees the first tx on the

Tangle, as well. If B wants to send 10 tokens to C, the ledger must first keep a record of a tx sending at least 10 tokens to B. Moreover, this implies that verification does not require that the first tx temporally happened before the second. A node can issue a tx with tokens it does not own yet, but which it knows it will receive at a later point in time. However, his tx is not referenced until the necessary tx has popped up. This then makes the likability of his tx to be chosen in the MCMC-algorithm. Therefore, such behavior would not allow premature spending of tokens. Still it can save time, because PoW is already done when the required tx pops up.

Validation is a time consuming process and requires computing resources. What is the incentive for nodes to validate the two txs they reference? "If you don't follow the protocol then your transactions won't be confirmed nor even broadcast to the others." [5]. Put differently, if a node issues a new tx that approves conflicting tx, then others will not approve this new tx, and so it will fall into oblivion.

2.4.4. Hashing

Even though it is not a step by itself, this chapter elicits background knowledge on hashing in general. Hashing is used primarily in the following situations:

- Pow / referencing txs
- Seed & address generation
- Signatures
- Identification of txs and bundles

A hash function is used to produce one unique output for desirably every unique input. In the case of PoW, a node takes specific attributes of the tx, aggregates them to one package, and inputs this into the hash function. The output of the function is a number which can only be calculated if one uses this exact input, meaning it is impossible to find a different input with the same output. This would be called a collision. Furthermore, nobody can retrieve the input just from the output. The aforementioned explanations do not imply that a hash function does not create collisions, as the following formalized, slightly more exact version of the explained three properties specifies:

- Preimage resistance (or one-wayness) means that given an output y it is "computationally infeasible" to find an input x', s.t. $h(x') = y$.
- 2nd-preimage resistance (weak collision resistance) means that given an input x it is "computationally infeasible" to find another input $x' \neq x$, s.t. $h(x) = h(x')$.
- collision resistance (or strong collision resistance) means that it is "computationally infeasible" to find two inputs $x \neq x'$, s.t. $h(x) = h(x')$. This property is stronger, because an attacker would have one more degree of freedom.

Computational infeasibility depends on the context. [35, p.323f.]
On top of these three conditions, it is usually desirable to calculate the output of the hash function quickly once given the input. In the context of PoW, this makes verification of the work done faster. Moreover, the output of the hash function has the same size for all inputs.

When a node wants to send tokens from one seed to another, it signs the bundle with the private key of each address inputted using the Winternitz-One-Time-Signature-Scheme (W-OTS) [39]. This is why an address may be used as input only one, whereas it can receive tokens from an unlimited number of bundles. That is also the reason why surplus tokens in a bundle must not be sent back to one of the inputs. For each reuse of the address, the security-level of the signature is halved [8].

The W-OTS is a generalization of the Merkle OTS, which uses the Lamport-Diffie-OTS [39]. The W-OTS trades more computation time for less required space. More specifically, it needs about twice as long, but requires only slightly more than half the output size for the same level of security as a Lamport-Diffie-OTS.

W-OTS is quantum-resistant, whereas signature schemes based on factorization or the discrete logarithm problem, such as DSA and ECDSA, can be broken using Shor's algorithm for quantum computers. The W-OTS is a hash-based signature scheme to sign the data and is combined with Merkle's tree authentication scheme [39], "which reduces the authenticity of many one-time verification keys to the authenticity of a single public key." [14, p.364] "The Winternitz OTS (W-OTS) is most suitable for combining it with Merkle's tree authentication scheme because of the small verification key size and the flexible trade-off between signature size and signature generation time." [14, p.364]

Furthermore, IOTA has "added the ability to choose between 3 levels of signature security": 81-trit (128-bit), 162-trit (256-bit), and 243-trit (384-bit), depending on the computing resources of the device. In general, IOTA uses both the Curl and the KECCAK hash functions from the sponge family of hash functions. Here, the cost-function is the KECCAK hash function (known as SHA-3).

2.4.5. Proof of Work

Before issuing a tx, a node has to perform a Proof of Work (PoW) as step number four It is similar to Bitcoin. This means that a hash function must be calculated over and over while each time incrementing a nonce until a certain criterion is fulfilled, just the way Hashcash [10] introduced the concept. The resource-intensive task yields a token which proves that the task has been performed, so-called PoW. Checking the integrity of the PoW should be a quick task, which is the reason why a hash function is used, in the case of IOTA Curl-P-81. Along the structure of a transaction depicted in Figure 2.2, the hash is calculated as follows:
hash = Curl(signatureMessageFragment, address, ..., trunkTransaction, branchTransaction, nonce)

"Hashcash was originally proposed as a mechanism to throttle systematic abuse of un-metered internet resources such as email, and anonymous remailers in May 1997."

[10, p.1] The original benefit of using Hashcash is Denial of Service (DOS) protection. In the Tangle, if there were no work necessary to issue txs, an attacker could partition the network strongly by spamming it. This is because a node would have to relay all these txs, since the attacker is sending valid txs. Moreover, since this process is parallelizable, this would resemble a DDOS-attack against the whole network. But if the node is able to realize that the attacker has not done some necessary work, then it refrains from relaying the txs. This means for the attacker that he must do the work. But if resources are necessary to calculate the cost function, then computing it numerous times becomes very costly. Therefore, spamming would become a high investment, while normal users are hardly affected by the increase of required resources. The trade-off between posing a high hurdle for spammers due to high difficulty and leaving normal users mainly unaffected makes Hashcash very delicate.

On top of DOS-protection, Hashcash elicits the mechanism of PoW which makes immutability possible (section 3.4) and defends against double-spending attacks (section 2.7). As defense against attacks, spamming by honest nodes is desired because of two reasons:

- The more txs are issued, the faster txs are approved and therefore processed.
- There is a race between "good resources" and attacker resources, described in section 2.7.

[10] describes Hashcash for both interactive and non-interactive cost-functions, which refers to whether a server issues a challenge or not. The PoW in the Tangle is non-interactive and uses a "publicly auditable", "bounded probabilistic", and "trapdoor-free" cost-function. As a cost-function, a hash-function can be used where the objective is to find partial hash collisions. Publicly auditable means that the cost-function can be easily verified by anybody. Bounded probabilistic means that the effort to compute the function is non-deterministic but takes finitely long time. The PoW in the Tangle is bounded because the set of possible outputs of the hash function is finite. Furthermore, trapdoors are pieces of information which allow breaking preimage resistance, but the KECCAK hash function used by IOTA is by default trapdoor-free. Also, there is no need for a hash function which allows for trap-doors into the other direction, because this would mean that nodes would be able to skip PoW.

The higher the difficulty, the more resources or the more time it takes to perform the task. However, this does not affect validation-time of the PoW done. There exists a dynamic throttling mechanism on the Tangle which changes the difficulty and therefore the resources needed to perform PoW. It works differently than on the Blockchain. Its purpose on the Blockchain is to keep the number of blocks per second constant. This is achieved by lowering the number of acceptable hashes. On the Tangle however, it is supposed to make sure that the network is not flooded and therefore partitioned by an enormous number of txs. The difficulty "is supposed to self-adjust according to the network topology" [43]. The way dynamic throttling works in the Tangle is that nodes queue incoming txs for propagation and drop txs if a certain queue size has been reached. The dropping is not arbitrary but favors the txs with highest own weight,

meaning the ones which have received the most work are more important. If a node's txs are not echoed back by its neighbors, it realizes that its txs are not relayed through the network. In this case it must increase the weight (minWeightMagnitude) granted to its txs by doing more work. Devices with small computing power would still be able to send txs, albeit less quickly, because they would focus their power on a single tx in order to reach sufficient weight.

The following two functions perform the prioritization for propagation of heavier txs. They are part of the class com.iota.iri.network.Node.java.

```
private static ConcurrentSkipListSet<TransactionViewModel> weightQueue() {
   return new ConcurrentSkipListSet<>((transaction1, transaction2) -> {
      if (transaction1.weightMagnitude == transaction2.weightMagnitude) {
         for (int i = Hash.SIZE_IN_BYTES; i-- > 0;) {
            if (transaction1.getHash().bytes()[i] !=
                transaction2.getHash().bytes()[i]) {
               return transaction2.getHash().bytes()[i] -
                   transaction1.getHash().bytes()[i];
            }
         }
         return 0;
      }
      return transaction2.weightMagnitude - transaction1.weightMagnitude;
   });
}

public void broadcast(final TransactionViewModel transactionViewModel) {
   broadcastQueue.add(transactionViewModel);
   if (broadcastQueue.size() > QUEUE_SIZE) {
      broadcastQueue.pollLast();
   }
}
```

The broadcastQueue of type ConcurrentSkipListSet stores all txs that wait for propagation. The first function sorts the txs by weightMagnitude, which is the own weight, descending. The second function queues another tx and drops the last tx after sorting.

In conclusion, PoW serves three purposes:

- DDOS-protection
- immutability
- protection against double-spending

Dynamic throttling works by prioritizing txs for propagation which received the most work.

2.5. Stability

This paragraph explains mathematically why the number of tips $L(t)$ is stable (or "positive recurrent" [40, p.215]) for $t \to \infty$ and does not go to infinity:

$$\lim_{t\to\infty} \mathbb{P}(L(t) = k) = c \quad k \in \mathbb{N}, c \in (0; 1]$$

Furthermore, it finds a term for its expected value L_0. For the analysis we assume Random Tip Selection. This means that every tip has equal probability to be chosen, like in a Laplace-Experiment. Honest nodes should not apply the Random Tip Selection Strategy, however, because it does not defend against lazy and malicious nodes. Sites which are not tips should have 0 probability. However, because of race-conditions due to computation and propagation delay, they can still have a small probability. This is because it can happen that a node approves a tip after another node somewhere in the network has already approved it. This happens in High Load situations more frequently than in Low Load. Defined informally, Low Load means that the "typical number of tips is small" [38, p.9], while High Load refers to a situation with a typically high number of tips. This distinction between Low and High Load is important because the coordinator of IOTA is only active until the network has reached High Load (section 4.2). However, there is no exact threshold to distinguish between the two situations.

$h(L, N) = h$ is the time it takes for a node to issue a tx, including computing time etc., in a situation where there are L tips and N sites. We define a random variable $X(t)$ which counts the number of incoming txs during time interval $[0, t]$. $X(t)$ is a Poisson Counting Process with parameter λ, because it fulfills the following seven requirements [40, p.312f.]:

1. $X(t) \geq 0$, because the number of tips cannot be negative.

2. $X(t)$ is integer valued, because tips cannot be divided.

3. If $s < t$, then $X(s) \leq X(t)$, because incoming tips are counted irrespective whether they are approved later or not.

4. For $s < t$, $X(t) - X(s)$ equals the number of txs that occur in the time interval $(s, t]$, holds true because $X(t)$ simply increments for every incoming tip at exactly the point in time where it comes in. Consequently, $X(t)$ is the number of incoming txs $\in [0, t]$ and $X(s) \in [0, s]$ and therefore $X(t) - X(s) \in (s, t]$.

5. $X(0) = 0$, because the counting starts at $t = 0$.

6. The counting process has independent increments, because one incoming tx does not influence another. For this we must assume that there is a sufficiently large number of nodes issuing txs independently.

7. The number of incoming txs in the time interval of length t is Poisson distributed with mean λt.

Furthermore, we define another random variable $Y(t)$ which counts the number of approvals which a given tip receives during time interval h, with $N = |\mathbb{T}|$. $Y(t)$ is Poisson distributed, because it can be derived directly from $X(t)$. The rate of this Poisson process is $\alpha = \frac{2\lambda}{L} = \mathbb{E}(Y)$, because each incoming tx, arriving at rate λ, can approve the given tx twice and there are L tips which can possibly be approved. Consequently, it holds true that

$$\mathbb{P}(Y = 0) = \frac{(\alpha h)^0}{0!} e^{-\frac{2\lambda h}{L}} = e^{-\frac{2\lambda h}{L}}$$

This allows us to find a term for $\mathbb{E}(\Delta L)$ when we observe a node issuing one tx and referencing two txs. As a reminder, the process of issuing takes an amount of time given by h:

$$\mathbb{E}(\Delta L) = 1 - 2\mathbb{P}(Y = 0)$$

$$= \textit{ expected number of issued tips } - \textit{ expected number of erased tips}$$

$$= 1 - 2e^{-\frac{2\lambda h}{L}}$$

The factor 2 comes from the fact that each node references (up to) two txs, potentially erasing two tips.

With the above formula in mind we understand that $L(t)$ is a continuous-time random-walk over $\mathbb{N}$. This means that $L(t)$ either increments, decrements, or stays constant for each time step $h(L, N)$. Moreover, we can observe in the above formula that $\mathbb{E}(\Delta L)$ is negative for large L, because $e^x \to 1$ for $x \to 0$. On the other hand, $\mathbb{E}(\Delta L)$ is positive for small L, because $e^x \to 0$ for $x \to -\infty$. This means that $L(t)$ increments if it is small and decrements if it is large. This keeps $L(t)$ fluctuating around an expected value L_0.

L_0 can be determined as follows. For a stable L, we know that $\mathbb{E}(\Delta L) \to 0$ for $t \to \infty$. If this did not hold true, $L(t)$ would eventually grow infinitely large. A diminishing $\mathbb{E}(\Delta L)$ implies that $0 \approx 1 - 2e^{-\frac{2\lambda h}{L_0}}$ which yields

$$\mathbb{E}(L) = L_0 = \frac{2\lambda h}{\ln 2} \approx 2.885\lambda h$$

On top of that, we can conclude that the expected duration necessary for a given tip to be approved for the first time equals $\alpha^{-1} = \frac{L_0}{2\lambda} \approx 1.443h$.

2.6. Consensus

Even though txs can approve arbitrary txs in the Tangle after validating them, consensus must not be arbitrary. After a while, all nodes must agree on one state of the ledger, otherwise users could not have trust that their txs are universally accepted. Like in a Blockchain, consensus can take a while. However, the required time increases in Blockchains with increasing number of participants because network propagation time indirect proportions to the number of orphaned blocks. With increasing number of participants in the Tangle, the time to reach consensus decreases because the Tangle grows faster. The whitepaper calls this time adaptation period [38, p.14]. Consensus on the Tangle is reached on the set of txs which are directly or indirectly referenced by all relevant tips. There is no clear definition for relevant tips given, but one could define a threshold for the minimum probability of a tip to be selected by the MCMC-algorithm. Figure 2.5 depicts an example-Tangle. Consensus is reached on the set of txs in dark gray, while tips are marked light gray. Apparently, the tip at the bottom is irrelevant, because it is extremely unlikely to be selected by the MCMC-algorithm [29].

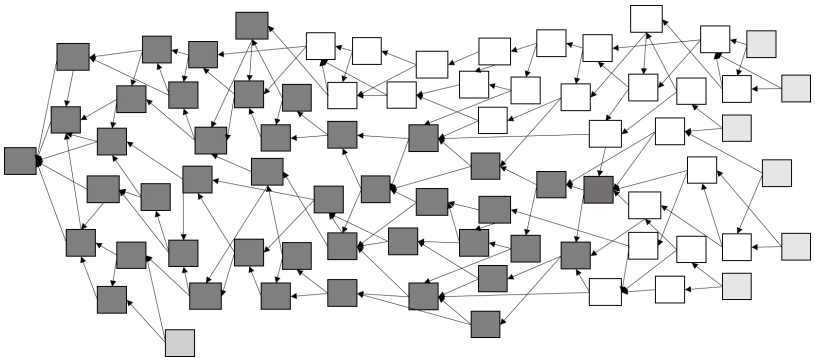

Figure 2.5.: Consensus is reached for the txs in dark gray, tips are marked light gray

The following examines what exactly merchants wait for until they accept a payment. In high load regime: The whitepaper proposes merchants to wait for a "sufficiently large cumulative weight" [38, p.15]. In the actual implementation, they wait for the following quotient to reach a certain percentage:
(rate of growth of cumulative weight of the issued tx) / ($\lambda * w$)
$\lambda * w$ = Rate of growth of cumulative weight of the genesis tx
This is because w is the mean weight of a generic tx and λ is the rate of txs coming in. In the terms of the whitepaper, the merchant waits until a certain percentage of the adaptation period of this tx is over [38, p.14] In other words, a merchant waits for a certain percentage of all tips to (indirectly) approve the respective tx.
"[Y]ou set confirmation level for yourself. You may decide that if 95 tips of 100 returned by Monte Carlo reference your tx then it's legit. Some merchants may wait for 99 of 100." [29]

In low load regime: As long as the coordinator (section 4.2) is active, the merchant simply checks whether the coordinator indirectly approves the tx.

2.7. Attack vectors for double-spending

For the designers of a cryptographic system it is crucial to understand all possible attack vectors and the ways to deflect them. The attacks described in the following give only a short insight into possible scenarios. All of them are directed at performing a double-spend. "So what is a double-spend? It is most simply described as any transaction that brings an address balance to a negative value. For example, spending the total balance two different times." [30] The following does not attempt to be a comprehensive list. There are other possible attacks, e.g. the eclipse attack.

2.7.1. Simple Large Weight Attack / 34%-Attack

The attacker tries to issue a double-spend tx while forking off the originally accepted tx. Two goals have to be attained: First, the merchant must accept the original tx and deliver the purchased goods, and second, the network must be convinced to build upon the attacker's sub-Tangle. The way to achieve these goals is to wait long enough for the merchant to accept the payment. During this time, the attacker can build his sub-Tangle offline. Furthermore, the attacker must outpace the growth of cumulative weight in the main-Tangle by doing excessive PoW on the sub-Tangle. Outpacing the main-Tangle would be easier if the attacker could issue just one tx of large weight. [38, p.16] That is why the own weight of a tx is limited. Consequently, the attacker must issue numerous txs on the sub-Tangle which reference only txs in the sub-Tangle, but not in the main-Tangle.

Eventually, for the attacker it boils down to having more than a third, i.e. 34%, of the total hashing power to force the enough walkers from the MCMC algorithm into walking along his Tangle. Described with the Byzantine Generals Problem [113] this means that more than a third of the generals are traitors, which results in a faulty system.

2.7.2. Parasite Chain Attack

The parasite chain attack equals an extended large weight attack. The attacker tries to build a side-Tangle quickly to issue a double-spend tx. At first, by issuing a legitimate tx on the main Tangle, he pays a merchant and receives the corresponding asset. Afterwards, he is then able to make the previous tx obsolete by constructing a side-Tangle which must eventually become the main Tangle. The difference here is that the side-Tangle builds upon the main Tangle by occasionally referencing it This allows it to reach even higher score and height. Therefore, it is called a parasite.

This attack can be prevented when nodes use some type of MCMC tip selection strategy if we assume that the main Tangle has more hashing power than the attacker. Consequently, the main-Tangle receives more cumulative weight than the side-Tangle, which makes walkers of the MCMC-algorithm more likely to walk through the main-Tangle. After the legitimate tx on the main-Tangle, no txs approve tips from both Tangles at once, because the legitimate tx conflicts with the double-spend tx on the

side-Tangle. From this point on, it is the same as a simple large weight attack, except that the side-Tangle has a slightly higher probability in the MCMC-algorithm. [38]

2.7.3. Splitting Attack

The Splitting Attack is very similar to a large weight attack, except that the attacker does not need to run a race against the computing power of the entire network all by himself. Instead, he tries to balance the weights of the two emerging sub-Tangles, so that half of the network grows one sub-Tangle, and the other half the other. The attacker issues two or more non-conformant txs, thus splitting the network. Without these conflicting txs some node would eventually approve one tx of either sub-Tangle and therefore merge them together. The more balanced the two sub-Tangles grow by themselves, the less work the attacker has to do, because the network works on the two sub-Tangles equally. This lets them grow long enough that a merchant would accept the attacker's payment. His benefit is that he can spend his funds on both sub-Tangles, resulting in a double-spend.

To prevent this attack from being successful, the honest nodes have to make it too difficult for an attacker to balance the two sub-Tangles. To achieve this, the MCMC-algorithm is designed in a certain way. First of all, the entry point of the random walk must start deep enough in the Tangle to make sure it starts before the split. Otherwise, the walkers would not be able to make a decision between the two sub-Tangles. And secondly, the algorithm must decide sharply, i.e. rather deterministically, for the sub-Tangle with greater total weight. This second point can be achieved by making the transition probability more dependent on H_x, which would make the algorithm more deterministic for older txs. [38, p.24]

Another interesting defense would be if one potent party issued plenty of txs at once, thus unbalancing the sub-Tangles. However, this seems less inherently secure, because it requires some policeman.

A factor making it even harder for the attacker to accurately balance the sub-Tangles is network latency. Txs take time to propagate through the mesh network all the way to servers owned by the attacker and vice-versa. This forbids any node to be omni-present so that an attacker could not react on changes of the two sub-Tangles quickly enough.

In conclusion, the splitting attack can easily be averted by a correct MCMC-algorithm. Physical conditions make such an attack even more unlikely.

2.7.4. Sybil Attack

In a Sybil Attack, the attacker uses a plethora of identities in a P2P-network in order to take over the network. According to [20], peers can issue an arbitrary number of Sybils according to their processing power, if there is no central "certification authority" [20, 5]. For example, if a node had only malicious Sybils, it would obtain bad data, because it is effectively separated from the honest network. A clear distinction of single identities is not needed and not even desired on the Tangle, though. Using asymmetric cryptography, only the private key allows for access to funds and the public key is the

unforgeable address of the receiver. Yet, there are two aspects which disallow Sybil Attacks: Since the Tangle requires a Hashcash-like PoW for each tx, the power of each entity in the network is solely dependent on their computing power. On top of that, full nodes need to have neighbors (since it is a P2P-system). The way they are found might potentially induce a vector for a Sybil attack. But in the case of IOTA, neighbors must be found explicitly by a human using social networks. Having thousands of humans add Sybils as neighbors would be a great feat of Social Engineering.

2.8. Offline Capability

According to the CAP-theorem, a database can never fully realize all of the following three attributes simultaneously: consistency, availability, and partition tolerance (CAP). The first conventional databases, e.g. MySQL databases, used to focus on C and A. This means all data is constantly synchronized among all servers so that there exists only exactly one truth in the network. At the same time, the database always delivers results when queried. Later it became popular to develop databases combining A and P, e.g. MongoDB for streaming services. The big difference is that there may exist several truths of one single database, meaning one truth per network partition. Usually such databases achieve only eventual consistency. The Tangle is such a database. The Tangle is always available, because a node can attach txs to the local copy of the Tangle even when it is offline. This is because building on a Tangle in older versions does not forbit merging it with the newest version later on. Consequently, the Tangle tolerates when there are different versions passing around in the network, because eventually these truths can be merged into one consistent truth.

In offline-mode, a node builds a sub-Tangle which references at least one tx on the main-Tangle. Both the sub-Tangle and the main-Tangle keep growing, but when the sub-Tangle is then added to the main-Tangle, new txs are able to reference txs of either Tangle. In doing so, the two Tangles melt together. For this to happen, the sub-Tangle must neither conflict within itself nor with txs from the main-Tangle.

However, the sub-Tangle has one disadvantage: its total weight is significantly lower than the weight of the main-Tangle. As a result, hardly any node using the MCMC-algorithm would walk through the sub-Tangle and approve a tip in the sub-Tangle. By intention, this is actually exactly the mechanism by which choosing old tips is discouraged. To make the tips of the sub-Tangle more likely to be approved, one can issue a tx approving one tip each from the main- and sub-Tangle. In this case, the probability of a successful merge would be half the probability of other tips to be selected. This probability can be doubled by issuing another tx which approves the previous tx as well as a different tx from the main-Tangle, as depicted in Figure 2.6. Compared to a tip approving two txs on the main-Tangle, the sub-Tangle now has the same probability to be approved as any other tip from the main-Tangle.

In conclusion, an offline sub-Tangle can easily be re-integrated into the main-Tangle even after a long offline-period by performing some additional PoW.

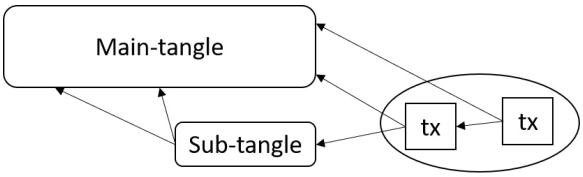

Figure 2.6.: Two additional txs necessary to promote sub-Tangle

2.9. Scalability

Scalability is the ability of a system to handle a linearly increasing load with linearly increasing resources. For DLTs we care about the number of txs per second which can be handled by the network. One can make the distinction between horizontal and vertical scalability. Horizontal means, scaling by allowing for better concurrency and providing more processing units. Vertical means, scaling by adding more computing power and storage space to existing processing units.

The only real limitations to scalability of the Tangle are bandwidth and storage size of the Tangle, but they are no inherent limitations from its data structure. Instead, the more txs are issued, the faster payments get confirmed. This is because the cumulative weight of a tx grows faster. The Tangle achieves this by allowing for horizontal scaling, because vertical scaling is undesirable in IoT-systems relying on light-weight embedded systems [21]. The question of how an enormous number of txs can be handled with limited resources is an ongoing debate.

2.10. Privacy

Without any further techniques, all txs, including the number of tokens, addresses, and messages, are visible to the entire network. This is necessary for every node to verify that one tx does not conflict with another and the PoW has been done. Users concerned with their data and especially their resulting meta-data (such as frequency of txs between certain addresses etc.) would have to hope that IOTA implements Private Messaging (section 6). Nevertheless, there is one method called Masked Authenticated Messaging (section 2.12.2) to establish an encrypted connection between parties via the Tangle, albeit this covers only messages, i.e. zero-value-txs.

2.11. Quantum Resistance

IOTA aspired to build a quantum secure system because it is not clear when there will exist quantum capable computers or whether there already are any. They consider the possibility that in secret some state-actor has already invented something capable of breaking common encryption techniques. Core-developer Paul Handy compared it to the incidence around SHA-1 or MD5 which showed that publicly available crypto-functions were corrupted by state-actors years before this information got publicly

available [13].

Furthermore, IOTA argues that trinary software running on trinary hardware will be more efficient. Therefore all code uses the trinary numeral system and provides functions for converting bytes and trytes. The Radix Economy (efficiencies of a number system) to the base of 3 is more efficient than 2 [2].

Asked on whether the trinary system might be less secure since it is an usual way of programming, founder Sergey Ivancheglo replied: "Trinary software is at least as secure as binary, it's just different numeral systems. Here is an example where binary fails while trinary doesn't: What is the result of "-X" if X is a 8-bit number of value -128?" This example insinuates that 128 does not fit into a signed byte while -128 does.

To the dismay of adopters of IOTA, the aspiration to be quantum secure has resulted in intense controversies because the Tangle has used their own hash functions instead of peer-reviewed ones, so-called Curl-P. This resulted in harsh criticism by a team from the Digital Currency Initiative at the MIT Media Lab in July 2017. "We found that IOTA's custom hash function Curl is vulnerable to a well-known technique for breaking hash functions called differential cryptanalysis" [37]. This means that the hash-function lost its weak and strong collision resistance. They found a vulnerability in the Curl-function, documented it and informed the IOTA Foundation which immediately implemented a patch. Curl was used for hashing messages as part of the signature algorithm. Therefore, the team allegedly could have forged signatures of txs, meaning they could have spent tokens from addresses they do not own [27]. Since the coordinator is not open source, it is hard to tell whether the coordinator would recognize such attacks.

Instead of Curl-P, IOTA now uses KECCAK (termed Kerl by IOTA), so a hard-fork was necessary, which was done during the snapshot on 8 August 2017 [37]. IOTA replied in a blog post and proved the impracticality of an attack proposed by the team [37]. On top of that, one of the founders and creators of Curl-P stated that the vulnerability has been implemented as a countermeasure against copy-cats [28].

In conclusion, IOTA strives for quantum resistance with good intentions but has opened up fierce controversy about the security of their system.

2.12. Advanced Functionality

2.12.1. Smart Contracts

A smart contract is a piece of code which is stored into the ledger as the content of a tx. Programmers can use the normal control structures like conditions and loops to simulate a state machine and change the state of the ledger. Nodes can interact with the contract by invoking functions of the contract and sending or receiving tokens to or from it.

Smart Contracts are inherently limited in Directed Acyclic Graphs (DAGs), because there is no absolute temporal order among the txs. Therefore, only such smart contracts can possibly be executed on the Tangle which do not care about whether one tx happened before the other. This would mean that one cannot model the following: I will

send you 1 token as soon as you have sent me 2 tokens. Nevertheless, IOTA is working on kinds of smart contracts [6]. They are working on enforcing correct timestamps in the Tangle. The basic problem is that everybody does their own PoW, so they can forge their timestamps, or their clocks are simply inaccurate. To solve this problem, the IOTA Foundation is working on implementing oracles which can provide timestamps of transactions. "One of the main ways to extend the utility and applications of IOTA is through oracles. Through this one can feed outside data, such as timestamps, into the IOTA network. You can expect some exciting announcements here." [45]. IOTA does not make public yet, whether it will provide a language similar to Solidity by Ethereum [13].

Dr. Serguei Popov, the author of IOTA's whitepaper, published a brief analysis of two ways of enabling trustworthy timestamps. The problem with timestamps is that a malicious node may issue a tx x with a tampered timestamp. So honest nodes need a way make sure that a timestamp is more or less correct. Neither proposed way can give a really exact estimate for the timestamp of x. However, they can prevent total outliers.

For the first way, which is deterministic, a node computes a confidence interval for the correct time T_x when x was issued. It is derived from the given timestamps t_i of all txs which are independent from x. Independent means that they neither reference x, nor are they referenced by x (directly or indirectly), i.e. $(x \not\leftarrow y) \wedge (y \not\leftarrow x)$. As illustrated in Figure 2.7, dependent txs are already in a partial order, so there is nothing to estimate about their relationship to one another. But to relate T_x to the timestamps of independent txs, one must first collect and sort ascending the set of the timestamps of all independent txs from x. Then one chooses a $\beta \in (0; 0.5)$ and calculates the β- and $(1-\beta)$-quantiles. Finally, one knows that $T_x \in [q_\beta; q_{1-\beta}]$ holds true to a confidence level of $1 - 2\beta$. Figure 2.7 shows a small example. Here, we have ten independent txs with timestamps $t_1, ..., t_{10}$, with $t_i \leq t_{i+1}$. If we choose, for example, $\beta = 0.21$, then $T_x \in [t_3; t_8]$. If t_1 happened to be a timestamp which was given an illegally small value, it is simply cut off. This estimate is quite rough, since the position of T_x in the sorted time series might not only deviate due to inaccurate timestamps, but also due to a difference in the number of independent txs which happened before and after T_x. The author recommends $\beta \in [0.2; 0.3]$, but cannot yet say whether there is a proof of optimality.

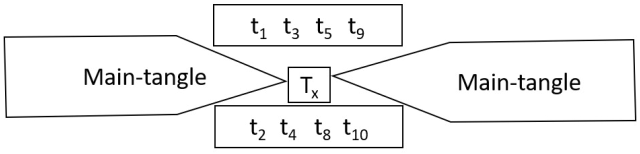

Figure 2.7.: Example of ten timestamps yielding a confidence interval of $[t_3; t_8] \ni T_x$ if $\beta = 0.21$

The second way is unfortunately elaborated only sporadically, but the basic idea is to create a non-deterministic/random confidence interval. If a malicious node wanted to

set a timestamp "from the future", honest nodes would not reference this tx. Therefore, the tx would not gain cumulative weight and so the MCMC-algorithm would be very unlikely to walk to it.

2.12.2. Masked Authenticated Messaging

Masked Authenticated Messaging (MAM) equals an encrypted RSS-feed on the Tangle [13]. A node can thereby broadcast messages in a stream which are signed asymmetrically for authentication and encrypted symmetrically. Since the payload is encrypted, no tokens can be sent via MAM. The sending node opens the stream by broadcasting and signing one tx to one of its addresses. This tx is the entry point to the stream. Every tx in the stream points to the ID of the subsequent tx. This allows nodes holding the public key of this address to trace forward all txs belonging to this channel and thereby listen in on the stream. They can find the txs even though they are scattered throughout the Tangle. On top of the public key, listening nodes must possess the symmetric key. The payload gets encrypted using a shared symmetric key, because symmetric encryption does not bloat payload size while being highly secure. This implies that only authorized parties gain access to the channel.

It is possible to serve several MAM-channels at once, to fork-off from one channel and even to allow new nodes to enter without them being able to view older messages. Since these are very recent inventions, they are not covered in this paper anymore.

This feature will, for example, enable sensors to securely and continuously send lots of data to authorized devices without sending the data multiple times, but instead storing it securely in the Tangle only once. Devices can then be made to pay for gaining access to the data, resulting in a marketplace for data [13].

3. Comparison of Tangle and Blockchain

After understanding the principles behind the Tangle, the reader can follow this chapter to get a clear notion of what makes it different from the Blockchain and where it exhibits parallels. This chapter assumes a Blockchain with a consensus protocol based on Proof-of-Work, not Proof-of-Stake or something else. In more concrete examples, the Ethereum-Blockchain is consulted.

3.1. The Fundamentals of Blockchain

To conduct a reasoned comparison, one must understand the principles behind Blockchain. The Blockchain is a decentralized, distributed, immutable, trustless ledger of txs in a heterogeneous P2P-network. It is comprised of blocks where every block references exactly one previous block, thus creating one single chain. Every block contains a variable number of txs, while its total size is bound to a certain limit, e.g. 5MB. A new tx is added to a network-spanning pool of pending txs and competes with the other txs for being included into the next block. A reference between two blocks is established by computing the hash of the previous block including some nonce and storing this hash in the current block. Only the genesis/first block has a null-reference. The hash must fulfill a constraint like in Hashcash, which results in PoW (section 2.4.5). Computing the PoW is called mining. The most popular crypto-currencies using Blockchains, Bitcoin and Ethereum, currently apply PoW, but there are other consensus protocols, e.g. Proof-of-Stake, Proof-of-Attention, Proof-of-Burn, Proof-of-Capacity, etc. Nodes participating in the P2P-network must propagate blocks so that every node contains the same, single truth.

Theoretically, every one of these nodes can perform the PoW and eventually find a block by finding a valid hash for the previous block (if the Blockchain is permissionless). However, in larger Blockchains, the PoW has become so difficult that only pools of computers or ASICs (Application-Specific Integrated Circuits) or server farms are capable of reaching a sufficiently large number of hashes per second to make finding a block probable. This is because the Hashcash-protocol involves a dynamic throttling mechanism by which difficulty is increased when the hashing power throughout the network grows. Dynamic throttling has the purpose of keeping the number of blocks found per second constant.

It happens on a regular basis that the Blockchain splits into two (or even more) chains due to network latency. Nodes follow the rule to build on the longer chain or rather, even more detailed, the chain with more PoW done. This means that for one block there are two (or more) successors competing for miners to build on them. The

winner becomes the main chain, the other gets orphaned, as depicted in Figure 3.1. Txs of orphaned blocks are stored back into the pool of pending txs. This splitting of the chain into two can potentially result in a double-spend if a merchant accepts a payment before the competition between different chains has been settled. As a result, the confirmation time is composed of two phases: The first phase is completed when the issued tx is included into a block, the second phase when a specified number of blocks have built on this one.

Since nodes are not obliged to perform any PoW themselves to issue txs, the network is split into two types of actors: users and miners. Consequently, the system relies on a compensation mechanism for the work of the miners. Compensation can take place by two ways: block rewards or tx fees. All tokens are created in the genesis block, but block rewards cause an increase of tokens, because their creation "out of thin air" is simply accepted in the protocol. Since tokens can get lost, this keeps the amount of tokens in circulation roughly constant. [36][1]

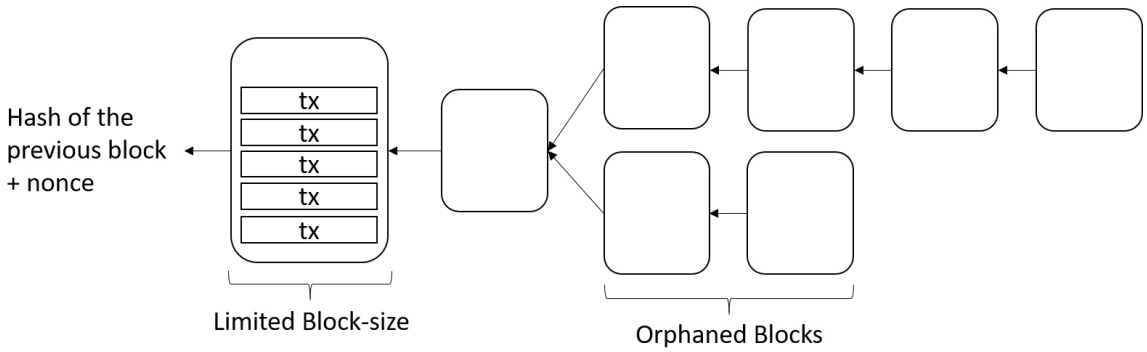

Figure 3.1.: Schematic Blockchain

3.2. Data Structure

The data structures of Blockchains and Tangles are not too far apart. These ledgers store transfers of tokens between addresses while one entity hashes a previous entity. On the Tangle, one entity means one tx which hashes one or two previous ones. On the Blockchain, it means a bunch of txs aggregated into one block which hashes only one previous block. This means that on the Tangle a node has a choice which entities it wants to reference, which destroys the inherent absolute chronological order among txs.

Another minor difference lies in the structure of bundles. A tx on the Tangle is not an actual payment, but only an input or output of a specific bundle. Eventually, a payment is made trough bundles, while on the Blockchain it is made in a simple tx inside a block. This does not imply any similarity of bundles and blocks, though.

3.3. Scalability

The most striking and far-reaching difference between the Tangle and the Blockchain stems from the fixed supply of blocks per second. This inherent restraint limits the number of txs per second to a fixed value and therefore acts as a bottleneck on the Blockchain. This prohibits both horizontal and vertical unlimited scalability. The Ethereum Blockchain, for example, keeps the number relatively constant at approx. 6,000 blocks per day, which is depicted in Figure 3.2. As a side node, the decline since April 2017 stems from a so-called difficulty bomb designed by the Ethereum Foundation. It was introduced to prepare the switch from PoW to Proof-of-Stake.

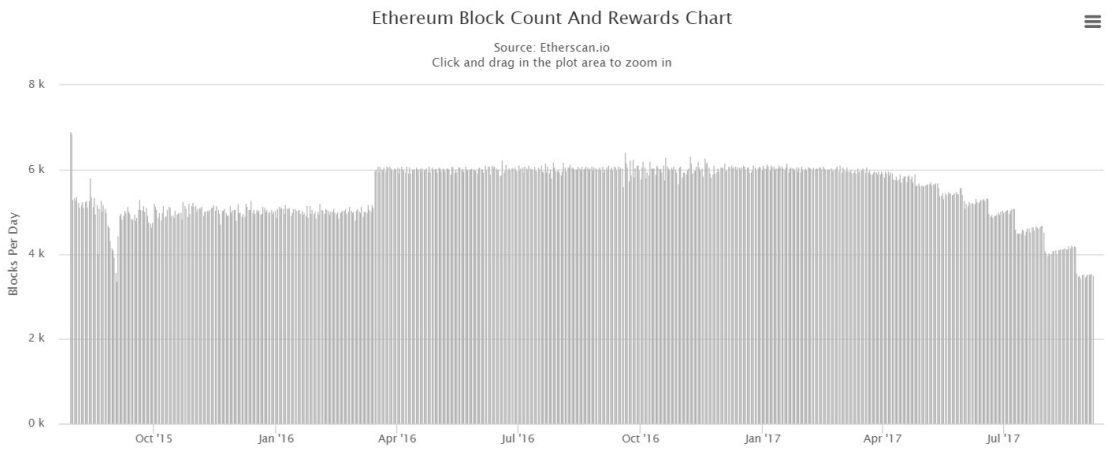

Figure 3.2.: Constant number of blocks per day in the Ethereum-Blockchain

[4]

As explained above, the Blockchain is inherently limited to a specified number of txs per day:
number of txs per day = number of blocks per day $\times$ number of txs per block (block size). By increasing the difficulty of the PoW to find a new block, the speed of finding new blocks is kept at a fixed rate. The reason for this procedure lies in the prevention of inflation. Since every new block yields a certain number of coins, an increasing rate of new blocks would increase inflation and therefore make the coin-price plummet. Inflation would not be all too extreme, because the mining rewards decrease with increasing number of mined blocks and some coins might get lost over time. Also, roughly only 20% of all coins have not been mined yet. Still, inflation would take place until the mining reward approaches zero. Possibly, the actual crux of the matter lies in the fact that tx-facilitation and coin-creation are inherently intertwined. Therefore, the only infinitely scalable solution for a Blockchain might be to relinquish block rewards and lessen the increase of difficulty. The incentive for miners would then be narrowed down to the tx fees while the number of mined blocks per day would increase and with it the total amount of tx fees. However, the Blockchain might still be less scalable

due to its partition intolerance (section 3.6). The other solution that has even been attempted already is increasing the block size. However, this solution merely adjusts the inherent limit of txs per day and leaves the Blockchain still unscalable.

By classifying the network into users and miners, the Blockchain incentivizes centralization of mining resources, which slows down the network further. Differently, the Tangle has made mining/validation "an intrinsic property of utilizing the network" which makes it scalable as described in section 2.9[21].

3.4. Immutability

Immutable in general means that a tx cannot be altered whatsoever after it has been published on the network. As soon as a tx has been written into a block and the confirmation time has passed, this tx can practically not be altered anymore. Two mechanisms are responsible for this. First, every (full) node holds a copy of the Blockchain and can compare incoming blocks from their neighbors whether they agree with its own copy. Second, the tx is indirectly included into the hashes of all subsequent blocks, because each hash includes the hash stored in the previous block. So if this tx is altered, all subsequent hashes would be altered, as well. Otherwise, a node would instantly notice that the hash is incorrect. Therefore, a malicious node would not only need to outpace the hashing power of the rest of the network, but on top of that make up for the deficit in the length of his version of the Blockchain.

Just like the Blockchain, the Tangle is immutable for exactly the same reasons. If a node altered one tx, the hashes stored in the directly approving txs would not match with the actual hash of the tx. All indirectly approving txs would equally store a differing hash. Nodes would therefore simply not accept the alteration of the tx. Instead, for the alteration to become valid, the attacker would need to build a side-Tangle on top of the altered tx, just like on the Blockchain.

3.5. Fee Structure & Time to Confirmation

A merchant waits for confirmation before he accepts a payment and delivers the goods. There are two similar concepts in both Tangle & Blockchain which determine the confirmation time: On the Tangle it is (1) the priority of propagation through the network and (2) the growth of cumulative weight on top of the issued tx. Similarly, on the Blockchain it is (1) being accepted into a block and (2) the number of blocks building on top of this one.

Aspect (1) indeed depends directly on the fee paid by the buyer: On the Tangle this means that the more PoW the buyer does on this tx, the higher its own weight and therefore the higher priority it receives to be relayed through the network. On the Blockchain there is a pool of open txs which compete to be accepted by a miner. The higher fee the buyer attaches to this tx, the more willing miners will be to include it into their next block. As a consequence, if a buyer wants to get faster confirmation, he would need to increase the fee he is willing to pay for this tx.

When regarding aspect (2), the buyer has no influence on confirmation time. Rather does aspect (2) depend on the level of security a merchant requires for his tx. This is because the more PoW is done on top of a tx, the less likely it becomes that another sub-Tangle/sub-chain will overtake this Tangle/chain. On the Tangle this means that, put simply, he specifies a percentage of how much consensus must be reached until he accepts the payment. (section **??**) On the Blockchain this means that he accepts the payment only after a specified number of blocks have built on top of the block containing the tx.

How buyer and merchant can tweak confirmation time is summarized in the following table:

	Buyer	Merchant
Tangle	The higher the own weight, the higher the probability of being propagated.	Specific percentage of consensus.
Blockchain	The higher the fee paid, the more likely to be accepted by a miner.	Specific number of blocks building on top of this block.

In conclusion, the time to confirmation follows similar principles in both Tangle and Blockchain.

However, currently the buyer must do the PoW himself when using the Tangle, because facilitators taking over the PoW for clients have yet to emerge. On the other hand, he has no practical way of doing it himself on the Blockchain, which requires him to pay a miner. These fees can become overly expensive due to increasing demand versus constant supply. This is because the number of blocks per second is fixed, while more and more txs are issued. A quote by Dominik Schiener underlines this jeopardy: "Ein ironischer Aspekt von Blockchain ist folgender: je beliebter die einzelne Blockchain wird, desto schwieriger wird es, sie wirklich zu verwenden - da die Transaktionsgebühren immer teurer werden." ("An ironical aspect of the Blockchain is the following: the more popular a single Blockchain becomes, the more difficult does it become to really use it - because the tx fees become more and more expensive") [25]

Another difference achieved by the Tangle is that the more txs are issued, the faster confirmation occurs, because consensus is reached faster. On the Blockchain, more txs means more competition for blocks and therefore ceteris paribus the time to enter a block and with it confirmation time increases. However, this is a result of the difference in scalability discussed previously.

3.6. Offline Capability

A DLT is only truly offline capable if a node can disconnect from the network, continue issuing txs into the ledger, reconnect, and then merge its ledger with the public ledger seamlessly. On the Blockchain, such behavior would not even be possible if the user is both user and miner. In offline-mode, a node could issue any txs, but eventually these

txs would become part of the ledger only after reconnecting to the network. Building blocks by itself would not help because it would merely build a side-chain which could not be merged with the public Blockchain later on. This means that a node can indeed prepare txs, but not add them to the ledger in offline-mode[1].

Assessed with the CAP-theorem (section 2.8), the Blockchain focuses on consistency and does not allow partitions of Blockchains. Only one true Blockchain exists in the entire network at all times and a node can only participate if it builds upon the latest blocks. Also, miners only perform work on the last block added, to prevent working on an obsolete chain. While the network is consistent on this truth, the information of the other truths, i.e. partitions, does not get integrated into the eventual truth. Consequently, the Blockchain is not offline capability. Contrary to Blockchains, the Tangle is offline capable, as explained in section 2.8.

3.7. Privacy

Both in Tangle and Blockchain, txs are by default stored in plain-text. Nevertheless, there are concepts to make txs more private and veil meta-data on the Blockchain, e.g. Zero-Knowledge-Proofs or the Hawk-System for Smart Contracts described in [31]. On the Tangle, users can already make use of MAM (section 2.12.2) while Private Messaging is in development.

3.8. Energy Consumption

The security of the network requires honest nodes to support the network with a large amount of computing resources. Otherwise, an attacker could easily launch a 34%-attack by purchasing enough hardware and energy. Since PoW consumes energy, the total energy consumption naturally rises severely as the rate of txs grows. Unfortunately, to date these computations serve no other purpose than securing the network, e.g. calculating some scientific problems. Neither technology has yet solved this environmental issue exhaustively.

As an example, Bitcoin miners use approx. 19 TWh per year regarding Sept. 2017. This equals half the total electricity consumption of the Republic of Ireland in 2014, with a population of over 4.7 Million people. [19]

The energy consumption of the Tangle would become comparably large if it uses standard CPU-bound PoW. This is because the main-Tangle must defend against resourceful attackers. For IoT-use-cases one can argue that a stupendous number of tiny devices combined will be able to reach a high number of hashes per second without requiring each single device to consume much energy individually. So each single device would lose hardly any battery life due to PoW. Furthermore, IOTA argues that it will develop and spread special trinary hardware which will produce enormous hash-rates at low energy consumption[21]. However, it should not be forgotten that a potential attacker would benefit from this hardware as well. Consequently, in total the energy consumption would be just as high. Nevertheless, just like on the Blockchain, the utility of one

successful double-spend tx compared to the tremendous investment necessary makes such attacks questionable.

Network-bound PoW might function as a sustainable alternative to CPU-bound PoW. Paul Handy mentioned that IOTA might introduce such a scheme in the future. It works using guided tour puzzles[9]. Since its adoption is not clear yet, it is not part of this thesis.

In conclusion, both Blockchain and Tangle waste enormous amounts of electricity in current implementations for keeping their ledgers secure.

4. The Tangle in the IOTA-environment

This chapter first collects information on the IOTA Foundation itself, its vision, mission, and team. Afterwards, a few concepts specific to the IOTA Tangle are examined. The chapter ends with a summary of how the Tangle will be further developed.

4.1. The IOTA-Foundation

The IOTA Foundation is a nonprofit organization (German: "gemeinnützige Stiftung") based in Berlin, Germany, founded by David Sønstebø, Sergey Ivancheglo, Serguei Popov, and Dominik Schiener. Developers have been working on an implementation of the Tangle since 2015. Like any other crypto-currency, the IOTA token were issued in an Initial Coin Offering (ICO). It took place on 13 June 2017 very successfully on the crypto-exchange-platform Bitfinex [41]. These quotes from members phrase the vision and mission of the IOTA Foundation:

- "IOTA was initiated with a very clear and focused vision of enabling the paradigm shift of the Internet of Things, Industry 4.0 and a trustless 'On Demand Economy' through establishing a de facto standardized 'Ledger of Everything' [45].
- "The goal of the IOTA Foundation is it to build a flourishing Machine Economy, where machines seamlessly interact and transact with each other." [24]
- "Most of general public will be using IOTA without even suspecting that. We need to reach only manufacturers." [5]
- "The backbone of IoT is here" [7]

Summarized, this implies that IOTA does not attempt to supersede Bitcoin as a currency, or Ethereum as a general DLT-facilitator. Rather does it focus specifically on IoT use-cases, most of all in the automotive sector, where the Tangle can substantiate its specific advantages over the Blockchain.

IoT scenarios usually encompass the following:

- nodes are mostly "specialized chips with pre-installed firmware" [38, p.3]
- a huge number of nodes participates in the network
- micro-payments are easy and often necessary for machine-2-machine interaction

When talking about IoT, Fog-Computing becomes increasingly important. The term refers to the introduction of another layer between devices and the cloud. Its benefits

are to "conserve network bandwidth", "minimize latency", minimize the time between data collection, analysis, and reaction, "collect and secure data across a wide geographic area with different environmental conditions", and "better security" [16]. Put simply, especially extremely small devices, e.g. mere sensors, would rather rely on close-by fog-servers than distant and overloaded cloud servers.

IOTA has been set up as a nonprofit foundation in order to unbiasedly negotiate with companies and create standards around DLTs in the IoT-sector. This is the reason why IOTA tends to prioritize machine-friendliness over human-friendliness. Originally, it emerged from a stealth-startup called "Jinn Labs" which develops energy-efficient trinary hardware specifically for IoT-devices. The plan for devices relying on the Tangle will be to equip them with this unprecedented ASIC. Existing devices could still participate in the network, but significantly less efficient and slower. In other words, the Foundation aspires to develop its own hardware standard. According to founder David Sønstebø, the idea of introducing a new hardware-component does not follow outlandish expectations, but has been thought through by hardware manufacturers themselves. Allegedly some of these manufacturers have already actively requested IOTA for more information. [21]

When examining the team behind IOTA, one can hardly find out who has something close to an employment contract. But on their blog blog.iota.org, the IOTA Foundation welcomes new people to IOTA, implying some kind of binding relationship between them. Furthermore, the IOTA community has listed affiliated people on iotasupport.com. Combining these sources, the total number of team members sums up to 32 at present, including the four founders, at least five developers, designers, mathematicians, and business-people. The founders regularly call for involved private developers to participate openly and officially. Furthermore, IOTA has a fast growing Slack-Team of more than 26,000 members.

IOTA's funds of currently approx. 10 Mio. USD come from the community (private supporters), corporate supporters and alliances, as well as German and Swiss government grants, according to founder Dominik Schiener [3]. Still, it is in an advanced startup-phase and looking for more cooperations to gain resources. Therefore, the community raised approx. 3% of all tokens for an initiative called the "Big Deal". IOTA can apply these funds freely to compensate corporate contributors for resources, such as developers. Moreover, such collaborations can lead to a gain in reputation, possibilities for Venture Capital, and opening a door to the Asian crypto market and media. For the adoption of the global Tangle it is important to attract both established corporations as well as newly emerging startups in the IoT-field. IOTA makes few of its partnerships public, because for companies it usually means attempts to find new strategic positions. But allegedly, IOTA is highly active in the automotive industry and partners with big automotive companies in Europe and the USA.

About competitors, founder Dominik Schiener comments: "I think Visa and MasterCard are the biggest competitors because they're also trying to get into that IoT space but they're completely centralized, and they have transaction fees, ongoing costs, etc. There is no real M2M [Macine-to-Machine] transaction today that can be taken seri-

ously. Most of the time it is still done centrally, over the cloud, where the machine simply triggers a request to a server and the server then does the payment, not the machine directly." [11] Since it is still early days of the technology, they also see the risk of copy-cats. However, Alexander Renz sees the networking effect that comes with expanding adoption as a copy protection. [12]

In conclusion, the foundation aspires to reach broad adoption of the Tangle by forming partnerships and building a supportive online-community. This will possibly allow them to establish standards in the IoT-sector, especially in the field of mobility.

4.2. The Coordinator

The Tangle relies on a high number of txs per second to ensure security. So it needs some kind of bootstrapping concept. In the first years, a coordinator node is responsible for making any kind of attacks, especially 34%-attacks, impossible, by frequently issuing milestones. Honest nodes somehow rely on them to detect faulty nodes. Speaking in the terms of their whitepaper, the coordinator is necessary in a Low Load Regime. The foundation does not specify a deadline until when the coordinator will be removed from the system. This causes some resentment in the community, because firstly the system has a centralized component, namely the coordinator, and secondly its code is not open-source. Therefore, it is neither totally clear what tasks exactly the coordinator takes over nor how it achieves that. Nevertheless, every tx must be approved by the coordinator in some way and once this is done the tx can be viewed as confirmed. Consequently, not all described mechanisms of the Tangle are in use at the time, e.g. how consensus is reached (section 2.6).

4.3. Peer Discovery

The Tangle runs via a mesh network which requires every full node to have a couple of neighbors which it exchanges updates on the ledger with. Currently, the users themselves must find the recommended number of seven neighbors via community platforms, like Slack. The IOTA Foundation tested automatic peer discovery in 2016, but realized that it slows down the network noticeably since it wastes an exponentially increasing share of bandwidth. Moreover, txs which are part of earlier snapshots can be rebroadcasted. Also, the network-wide difficulty of PoW cannot be adjusted automatically anymore. [43] Furthermore, manual tethering makes Sybil-attacks and therefore 34%-attacks significantly more difficult, because the attacker would need to convince hundreds of human users to bilaterally tether with them. This argument holds true only with certain caveats addressed in [33].

4.4. Types of Nodes

To be considerate of different capabilities and requirements of nodes, IOTA provides software for three types of nodes: full-node, light-node, perma-node.

- A full-node is the standard node on the P2P-network running the IRI. Therefore, it is required to connect to neighbors and propagate txs, otherwise its neighbors remove it from their list. It stores the validity only of the last snapshot plus all subsequent transactions as the current version of the Tangle, because for most use-cases this suffices. It can issue txs and do the necessary validation all by itself.

- A light-node, light-client, or light-wallet is not an actual node on the network but relies on a full-node to act as a server. Since it does not store the Tangle, it must request the full-node to provide txs it can do the PoW on. Communication takes place only with this Light-Wallet-Server. Usually it secures the network if honest nodes spam the network with txs, doing PoW. However, it does not help the network to spam with a light node, because the resources of the connected server are spammed. [22]

- A perma-node is essentially a full-node, but stores not only the latest snapshot but the entire history of the Tangle. There are certain use-cases which require knowing the exact course of action in the past Tangle.

	full-node	light-node	perma-node
Stores the whole Tangle	✗	✗	✓
Stores the Tangle since the latest snapshot	✓	✗	✓
Finds neighbors & communicates with them	✓	✗	✓
Bundling & signing	✓	✓	✓
Tip selection	✓	✗	✓
Validation	✓	✗	✓
PoW	✓	✓	✓
Publishing	✓	✗	✓

[21]

Examples for Light-Wallet-Servers: http://iota.bitfinex.com:80, http://eugene.iota.community:14265, http://service.iotasupport.com:14265

One practical question on nodes remains: If there is no mining and people can connect via a light-node, what is the incentive for anyone to run a node (which costs energy, setup, maintenance, etc.)? The following reasons are taken from [23]:

- "You are aware of the fact that running the full-node is beneficial for the Tangle topology and you want to help. [...]

- You have lots of transactions to make and don’t want to rely on a light node-server, as there is no guarantee that they are online when you need them.

- You have a web app running and need the stable connection.
- You want to have maximum speed, so you choose the full-node.
- You want to have a copy of the Tangle database, that is generated when using a full-node. [...]
- In the future, maybe you provide a service and earn money for a full node. [...]
- You invested and want to support the Tangle as much as possible."

4.5. Snapshotting

To keep required storage capacity for full-nodes low, the coordinator performs a signed snapshot once in a while. Currently, snapshots are triggered manually, but it is planned to make them fully automatic. [43] A new snapshot deletes all txs from the Tangle and only stores the balance of each address. Therefore, a full-node must merely store the snapshot and all subsequent txs in order to be able to validate txs. In the past, IOTA used snapshots to coincidentally introduce profound updates to the software, e.g. the IRI. For example, the snapshot on 8 August 2017 required all users to send their tokens to a freshly generated seed, because a different hash-function was introduced after the publications of Neha Narula et al.. Snapshots are similar to "pruning" on the Blockchain, but allow for compression of multiple txs into one record if the same recipient address is used. [18] A snapshot is stored and transmitted as a simple array of JSON-objects, having the attributes "address" and "balance". An excerpt from the snapshot on 8 August 2017 looks like this:

```
[
   {
      "address":
         "ZHITHKLKRZEY9HJCWH9DBLIHZLWB9OUMSKZHNAEQVMPSMQYWYJHUJRZHIJI
          GJBUSHLXLWETVWNFWLPZAL",
      "balance": "50000000"
   },
   {
      "address":
         "DWVNBUBSTUTSEB9KZPATIKHDJPZGEEFATMIGZFKQZTAYVZ99GNQAMQVNRBS
          UATBNVDOPOLPUYBQUXWHUO",
      "balance": "200000000000"
   },
   ...
]
```

A snapshot has to be confirmed by a minimum number of nodes to be viable as the basis for the Tangle. They do that by checking all txs and balances. So, it is not possible to manipulate it easily.

5. Conclusion

This thesis allows for a comprehensive understanding of the Tangle-technology and gives background information on the environment it is embedded in. The analysis has shown that it is indeed a highly scalable, immutable, (quantum) secure, feeless, and offline capable Distributed Ledger Technology. Even though something called "tangle" sounds non-transparent, it is actually a very transparent and comprehensible system. The theory behind the technology seems well-conceived. Its application is a great example for how mathematical research invents the foundation for completely new systems. On top of that, the thesis shows it is a promising idea to focus on IoT-use-cases, because they require a light-weight and resource efficient way of issuing micro-transactions without high monetary fees. Compared to the Blockchain, these requirements are the main advantages of the Tangle.

Nevertheless, in the comparison of Tangle and Blockchain it becomes apparent that their concepts are very similar. This leads me to think of the Tangle as a partition-tolerant version of the Blockchain, when consulting the CAP-theorem. The concepts of immutability, fee structure, and time to confirmation, exhibit profound parallels because they build on Hashcash-like Proof-of-Work. On the other hand, the data structures as well as the roles of nodes in the P2P-network differ significantly.

Even though the technology is well conceived theoretically, there are certain caveats to the promises. First of all, limited privacy and extensive energy consumption are problems both Tangle and Blockchain are yet to resolve. Moreover, the cryptographic functions applied in the implementation of IOTA as well as the trinary approach have raised doubts on the security of the IOTA-implementation of the Tangle. Especially the coordinator as a bootstrapping-concept demands a certain level of trust in the foundation, because it is closed-source. Since the IOTA-Tangle is the only existing implementation of this technology, the system in High Load has still to prove its proper functionality in practice.

6. Outlook

By the time this thesis was written, new features and improvements of current functionality are being implemented. For example, reattachment will work "under the hood" so that users can be sure the tx will eventually be confirmed by itself [34]. Furthermore, txs carry lots of meta-data, e.g. the frequency of txs between certain addresses. With Private Networking the IOTA Foundation is exploring means of covering even meta-information, similar to Zero-Knowledge-Proofs [45]. Moreover, startups might emerge which provide fog-computing services doing the PoW for the IoT-devices of other companies. It might be called PaaS (PoW-as-a-Service). Event-publishing for P2P-nodes is currently in beta-mode. But, once working as intended, it would for example enable light-clients to receive feedback on whether their tx is propagating properly. So, they would be able to adjust the required weight. Lastly, if IOTA really reaches a number of txs per second in the magnitude of millions, then nodes might have problems relaying all of this data. However, according to founder Sergei Ivanchov: "Nodes can split the burden and solve this problem in swarm-like manner. IOTA is designed with this approach in mind." [29]

After getting an understanding for the principles behind the Tangle, further research questions arise. First of all, as the ongoing work done by Dr. Serguei Popov suggests, research on enabling smart contracts in DLTs based on DAGs continues. Beside enforcing timestamps, there might emerge other possible ways to achieve contracts. Furthermore, when considering the debate around cryptography and trinary systems, a comprehensive analysis of the cryptographic functions might be necessary. It could be of interest whether trinary programming leads to more mistakes in the process of programming, and if so, whether it indeed enhances efficiency while decreasing energy consumption. Moreover, researchers and corporates might be interested in concrete use-cases and a classification of scenarios in which the Tangle surpasses the Blockchain. Finally, substantial effort should be invested into realizing reductions of the energy consumption of DLTs. One proposal which the IOTA Foundation is working on is network-bound PoW instead of CPU-based PoW. Its puzzle would work like non-parallelizable Hashcash [13]. Whether this solves the issue remains an open question.

Appendix

A. Interview Transcripts

Interview Transcript

In the following interview, PH stands for the interviewee Paul D. Handy and BB for the interviewer Bennet Breier.

BB:

What do you like about the IOTA community when you're interacting with them and maybe one thing that you don't like?

PH:

Ah, what do I like about the community, I mean I came out of the community, so, it's kind of hard for me to say. I like myself I guess, but, … there's a lot of really active people, like, not everyone on the community obviously is active, but a lot of the ones who have been around long enough to kind of just get how the community woks, right? A lot of them voluntarily contribute a lot of help. There's a lot of people who go into the chats, you know, They're not getting paid or anything. They would help the other community members solve their own problems. The best examples of programmers is that,

If I were to hire, that I have seen had been people from the community who have been from themselves gone out, demonstrated some ingenuity, some unique thing.

BB:

Is that the way that you started with IOTA that you tried things out, downloaded some of their stuff?

PH:

Yeah so, when I was a little bit greener in the community I was certain to think that I want to make an automated, kind of like a shape shift service for making Bitcoin big going between Bitcoin and IOTAs, that was like a year ago or so. I was, I had free time, because I was on fraternal leave, and looking into that I have come to notice that there are a lot of places where the libraries needed help where there is a lot of programming that still needs to be done. And noticed that programmatically moving my funds around was a little more difficult than, you know I had initially expected. After I got bored with my old job and quit that I found that Dom was asking people for help and was offering to pay for stuff so I just asked them: Hey what do you need help with? I am an electrical engineer and I can program at least good enough to get myself into trouble and he threw me a bone and … It was doing the proof of work for the native stuff that was the first project I had and …

BB:

Right, you were working on the Curl-lib, right? Making it more efficient.

PH:

Ccurl was the first one, then Curl, JS, I dad that a few months later, but it's just, they like the way I work, my work ethic, and the kind of just kept asking me to do stuff. And here I am, a year, well not quite a year later, but many months later and playing fireman

BB:

Since we are already talking about this topic, I already read in your vita on Xing, kind of, you wrote that you initiated RocksDB for IOTA? For what reasons, what design considerations led you to use RocksDB?

PH:

We wanted an embedded database. One that was performant and allowed for concurrent access. So that ordinarily does pretty well, not a whole lot. There is a lot of scary stories about LevelDB correcting stuff and Rocks is a fork of LevelDB. And it has solved most of the corruption issues and there is another choice that I would have gone with, but it is a proprietary embedded database owned by Intel that I am not even sure that Intel knows they have. It is not open-source. I know someone who, well, my father was involved in it and it is a database that originally came out of the national lab. It was originally done in Ada. I am not sure if you are familiar with Ada?

BB:
No, unfortunately not. But, so, Rocks DB allows for concurrency and embedded …

PH:

It allows for concurrent access, it is embedded, it does Bloom Filters. It is essentially just enough of what we needed.

BB:

And it is also just a simple Key-Value-Store right? So you are not using it for analyses or something like that

PH:
Well, yes and no. The reason why we moved away from the previous database, that was just memory mapped files, is that it did not allow for concurrent access essentially. It was not able to keep up.

[BB giving a delayed introduction of his persona and the purpose of the interview]

PH:

One thing about the community which I don't like, and it's generally not the community but it's people who only occasionally come into the community is there's some people who have an entitled attitude that, because they bought some tokens on some exchange, somebody else owes them free

labor. I don't like that. That's something that happens sometimes. It is fine to ask for help if you're polite, but sometimes there's people who come in demanding all sorts of things. It is not stock in a company. The IOTA foundation is not a company issuing stock like a company. This is people who are just working on a technology and trying to help the technology gain adoption.

BB:
I also do own some tokens, but nevertheless I am just as happy that you could make it for a while for this interview.

Do you have the pdf open?

PH:
I'll open it. […]

BB:

The first difficult question to consensus is: Why does nobody have an incentive to do 34%-attack? This question arose when I listened to an interview by David Sonstebo who said that the Tangle is inherently resistant to such attacks. It rather strengthens the network. I don't understand why it strengthens the network and if it really strengthens the network, why would we need Proof of work in that case?

PH:
So the Proof of Work is the 34%. I hope you understand that the Proof of Work does not necessarily equate computational Proof of Work. It can also be network bound Proof of Work.

BB:

The way I understand Proof of Work is that a bit like in Bitcoin where you like put a lot of energy into servers and they calculate hashes and in the end you find a hash

PH:

Yes, That's how CPU-bound Proof of Work works. But for IOTA the end goal is not necessarily CPU bound Proof of Work. For Internet of Things environments, it is more likely move toward network bound Proof of Work. So network bound Proof of Work is where… if you look up the Wikipedia Proof of Work variants, you'll see there is network stuff there. It is essentially just, you use the fact that the physical latency that is required for you to go from you to someone else back to you. Because the point of Proof of Work Is that it takes time and increasing difficulty to accomplish a task.

BB:
So this is kind of related to the second question; about energy consumption. To me it's a bit like, it's a race of computational improvement on the malicious side vs. the good/honest side. The honest side is the community and all people who spam the network beneficially. And there might be an attacker who wants to have a server farm, but it's really costly to have such a server farm.

PH:

It is. And so, I guess if we come back to the first question, well we can come back to the question about energy consumption later, right? What is the attacker's incentive to perform a 34% attack? If we look at the 34% attack in Bitcoin, it is a misnomer to call it a 51% attack. It is a 34% attack. IOTA just doesn't cherry-pick 51%, because it is really 51% of the rest of the network.

BB:

So if the attacker has 34% of the hashing rate on the network, he can easily do a double-spend, right?

PH:

Yes, it is easier. But that requires omnipresence. That requires the attacker to be able to know at all points in the network to know what transactions are coming up. By nature of a couple of things, first the Monte-Carlo-Algorithm that everyone is running and also the physical network topology. If an attacker wants to keep a double-spend fork. It keeps growing to forks growing simultaneously, he has to constantly keep those balanced exactly. A small change in the weight in one of those is an exponential increase to the heavier one that the network hash rate is gonna go to that one.

BB:

OK, that sounds a bit like what the white paper describes.

PH:

Right, that's what's explained in the whitepaper. So the whitepaper essentially assumes omnipotence and does not go into the fact of the physical network. With a physical network, you cannot connect to all of the nodes. You cannot have an instance

BB:

It's also related to the CAP-Theorem. So is it possible to say that there is some kind of omnipresence in Bitcoin because it uses not Partition-Tolerance but C and A?

PH:

Well, not really, this is how the Chinese Miners were able to do their ... attack partially.

BB:

Maybe it gets clearer if you explain how the tangle could absorb such an attack?

PH:

Yeah. [...] First off, an attacker, if he is doing a 34% attack what he wants to be able to do is overload the tangle, right? But in order to do that he needs to be able to transmit all his transactions to the tangle. But the way that the Proof of Work works is that, more proof of work does not increase your own weight when the Monte Carlo Algorithm is run. More Proof of Work prioritizes your transaction to be transmitted across the network. So the transmission is ordered by Proof of Work and once you

start filling your neighbors' broadcast keys up quickly, then you have to go to the next higher MinWeightMagnitude, so that your transactions that you're building on top of that propagate. And then if you are going so fast and drop off all of your old ones from your neighbors' broadcast keys, you essentially flood yourself out. Like trying to pierce through a pinhole really hard.

BB:

I'll have to think about that later. But it's not like imagining that these two forks are combined again later? Because the attacker wants to fork off, right?

PH:

Yeah the attacker wants to fork off and what he wants to do is he wants to build them up evenly until he is convinced that someone essentially that he has spent with them and then switch off to the other subtangle.

BB:

Maybe a related question on page 2: What specifically wait for when waiting for confirmation of payment and why? I guess it's one of these two options: Cumulative weight or percentage of approving tips.

PH:

There is high load regime and low load regime in the whitepaper. The coordinator is used in a low load regime. In that mode it is easy to say when the coordinator indirectly approved the transaction.

BB:

Ok, then the merchant is fine with it. That makes sense.

PH:

In the high load regime, it's the lambda w in the whitepaper. You have that adaption period and then you have the linear growth of weight by lambda w. So essentially when you say you are confirmed when the growth rate of your transaction's weight is equal to the growth rate of the tangle or close to that. It can be something like, it's like a percentage confidence thing right?

BB:

So it's not plain the cumulative weight, but more sophisticated.

PH:

It's essentially listening on the network and seeing how many of these indirectly approving my transaction. You are listening all the new transactions that are coming in. Are these approving me or approving some other part of the tangle. If you just count that, you'll see in the beginning maybe 25% are approving you. You don't wanna cut that at that point. You wait a little bit longer, maybe you see 60% are approving you. Yeah, maybe your merchant will accept that. Wait a little bit longer, seeing 75% and eventually you will have something close to 100% of all incoming tips, all incoming transactions are approving your transaction.

BB:

And at the point where 100% of all incoming transactions approve of one specific transaction then there is consensus reached on that transaction.

PH:

You could say it like that. If you have the tangle split into four, quite difficult do an eclipse attack on the tangle because you don't know who the participants on the tangle are. A very sophisticated attack might be able to do it, but if it's in a physical mesh network, it gets more and more difficult to do that.

BB:

Maybe just a quick question on energy consumption. Is it gonna rise like in Bitcoin on the size of a country's consumption?

PH:

If network bound Proof of Work is used which is Sergei's essentially stated that is what is probably gonna be done in Internet of Things, then it's not going to be energy consumption but time consumption. I'll just be based on latency. There are small computations but the overall thing will be based only [...]. And there might be places where it makes more sense to have it energy based. Other places where it makes more sense to have it latency based. But I think it will most likely move towards network bound in general case.

BB:

What are the advantages and disadvantages of doing it latency based? One disadvantage could be that there are really intense partitions on the network which take a lot of time, … These partitions are really far apart from each other.

PH:

A naïve form of network bound Proof of Work could be: You want to send transactions to the tangle. You ask me for a token, I give you a token, you use that token to generate some message authentication code [...] that can be shown that I gave that token to you for that transaction and then you have proved that you have reached out to me, got a token for that, and I gave you a token for that. That's a naïve way of how that works. Sergei knows it better than I do.

BB:

Maybe you have some ideas on advantages and disadvantages?

PH:

The advantage of the CPU bound one is that it is more straightforward to implement it's just pure cryptographic riddles. Obviously the disadvantage is that it requires more energy, exponentially. With IOTA the growth of the Proof of Work required is not related to people competing for a block reward. It is related to the local state of the network. If there is a lot of transactions passing through your local network, your local part of the IOTA mesh topology, then you have to do more Proof of Work. But that's a transient thing. It might go up it might go down.

BB:

Is there something I can read about it? Is it still hashcash?

PH:

Right now, it is still hashcash.

BB:

And if we do the latency based one, is it also hashcash? Maybe it is based on some scientific paper. That would be perfect!

PH:

Could you remind me about that tomorrow?

BB:

I haven't read about it anywhere, not even in the roadmap. There is a lot of things in the roadmap, but…

PH:

For network-bound Proof of Work stuff you need to be following the tanglemath channel. That's generally a troll-fest over there.

BB:
So it's not really official?

PH:

It's not like an announced. This is what we're going to do. This is on the brainstorm-map, so to say. I think it's been planned longer than I have known it.

BB:

Do you think this is likely to be introduced? Because otherwise the energy consumption would increase similar to bitcoin I suppose.

PH:

Similar to Bitcoin, but not quite the same as bitcoin. It is more like Hashcash than Bitcoin's hashcash, because Bitcoin is a competition for a block. IOTA's hashcash is more like hashcash for email. Hashcash for email does not increase exponentially. It'll go up and down, but that depends essentially on the self-regulating denial of service protection.

BB:

I think, that is already enough on that topic. Let's move on to the next question. Page 2 would be about smart contracts. General Smart Contracts are not possible, because only a partial order can be established on the tangle. How will Smart Contracts be made possible and how important do you estimate them for the adoption of IOTA's tangle compared to Ethereum?

PH:

The assertion that is made in the first sentence, I am not sure if I can really agree with. I can't say a lot about it right now. It's under wraps for the moment. How will smart contracts be made possible? They will not be executed directly. Think of the tangle kind of like the IP-stack. You don't run programs over IPv4. You sent messages over IPv4, right? Smart contracts run on IOTA are not run as part of that internal messaging part. But they take advantage of that eventually. Probably can't say much more than that.

BB:

IOTA would assume that they are similarly important to IOTA's growth. Because Ethereum is mainly based on smart contracts.

PH:

It's bitcoin plus for-loops.

BB:

Ethereum would not grow as much if it did not have smart contracts. But IOTA would still be an important technology without smart contracts. So how important do you think it is?

PH:

I think the role of them will grow more. I think it will be a significant thing in 5 years. Or well, I think it will be a significant thing in 1 year. I think that smart contracts will be a big thing in IOTA and will probably be a thing that will be talked about as one of the main points more than Masked Authentication Messaging or Flash-Channels.

BB:

Maybe one follow-up question on that: I think Ethereum has Solidity as its own language. Would IOTA create some standard for that also?

PH:

Can't comment on that.

BB:

Lets move to the next question. How do I make sure that when I reconnect to the internet, so I was disconnected, issuing some transactions on the tangle and then later I want to push them to the tangle online. How do I make sure that my chain/tangle of transactions is approved by the main tangle? I find it kind of a really difficult question because if everybody uses the MCMC algorithm, then it's really unlikely that my transactions are approved because they are kind of old, compared to the main tangle.

PH:

If you pick good tips to merge your subtangle in to the main tangle, then so long as your subtangle is consistent, it has a fair share of getting merged in. If your subtangle is inconsistent with the maintangle, it is going to be rejected.

BB:

Absolutely. The way I could imagine it is that I have one chain that is not connected to the main tangle except for the points where it was originally still online and then later at the end of my own chain I could put another transaction which approves my chain and the main tangle. But then still the likability is still only half of what other transactions would have.

[…]

PH:

Not exactly. Lets say you pick a really good tip from the main tangle. Lets say you pick the one, the very best tip on the main tangle. The one that is most likely traversed by everyone else. Lets just pretend that there is 2 tips. One has the probability of 60, the other of 40. I mean .6, .4. You reference .6 when you merge in your subtangle. People are more likely to walk to the .6 and then if you are the only transaction there, they are definitely going to walk to your transaction.

BB:

But isn't everybody else doing the same? So is it like a race for the best tips?

PH:

No, it kind of washes up. What you would probably do is… Really the risk when you're trying to merge subtangles together is having a lazy tip. A lazy tip… I doesn't depend on the length of the previous stuff so much as how low you are by the time everyone else gets your transaction. It doesn't depend on the size of your subtangle that you're merging.

BB:
Yeah right, it rather depends on how old the transaction is, right?

PH:

It depends on how old the transaction in the main tangle you are referencing is. Not how old the one in your subtangle is.

BB:
In a way also, because it has less cumulative weight then.

PH:
No, your tips when you create it has a cumulative weight of one. All the tips have a cumulative weight of one.

BB:

But you look at the path.

PH:

The path is going through the main tangle and someone is going to walk to your transaction. If they walk to your transaction and it is consistent with everything else that they have, they will select your transaction. The probability of selecting your transaction has nothing to do with the transactions that your transaction approves, as far as their cumulative weight. Because cumulative weight is everything on top of you, not everything below you.

BB:

Isn't it exactly this mechanism that discourages approving old transactions? Because in that case my transaction wouldn't be chosen by anybody else, since going the path from the main tangle to my tip, this one step would only have half…

PH:

Yeah I guess you could say that you have a 50% reduction in your probability because only one of your legs is …

[...]

So you could have that 50% reduction, because they are not going to walk through your subtangle. They are going to walk through the main tangle to your transaction. Something you could do is, if it does not work with one transaction, you could promote that transaction with more transactions on top of it. So you make one transaction where one leg points to yours, the other leg points to the main and then you make another transaction that points to your merging one and the main. So it's pointing to 2 main-tangle-things.

BB:
So that is not reattaching but putting on top?

PH:

I think the word in the whitepaper is promoting.

BB:

Does reattaching mean taking it away again?

PH:

No. It's just left behind essentially.

BB:

I think I have understood what you mean on that question. Because of the flash network, I haven't found much online.

PH:

I haven't written much of it.

BB:

So maybe I'll just leave that out.

PH:

I can give you a short overview.

[...]

The skinny on flash networks is, you use multi signatures. You do not want to use multi signatures many times, but a couple of times might be ok.

BB:
And how does it make the network faster?

PH:
These transactions don't happen on the network. So here is what happens. You and I, we make a multisignature between the two of us, two of two. So you control one of the keys, I control one of the keys. It's probably gonna be a 4 of 4, but effectively a two of two. Security-wise it would be a 4 of 4. In order to spend from that multisig you must sign and I must sign a bundle. You would propose to me a bundle and I would either say "fuck off" or I might say "great, I'll sign that". To prevent some certain situations…

BB:

And the bundle I send is the transaction I want to get on the tangle?

PH:

Maybe not. We will get there in just a sec. I am just making sure we are clear on our foundation. We first start by funding our multisignature channel. You and I we both send deposits to that. The amount of deposits that we send is relative to our level of trust between each other. If you trust me fully and you are only paying to me, then we will say you fund 100% and I fund nothing into the channel. I am just going to probably blindly sign anything you said to me, because it's only going to make me richer or not more rich, but it is not going to make me poorer. For the reasons for that, we will get there in just a moment. But if you don't trust me, we will sign a 50-50. So that we both have the same amount to lose, if one of us acts against the economic interest of the party, both of us. After we fund that and the funding is confirmed. We have published to the tangle, we have sent from our personal addresses to the pool address, the funding address for the flash channel.

BB:

What do we send there?

PH:

Money. We send tokens.

So that we fund the channel on the tangle. Individually. You fund it so much, I fund it so much. And after that all of our changes to our state are going to happen off-tangle. And we'll only publish the latest to the tangle. Because the first rule, don't use your signature many times, we pick high-security keys and we are only going to use each of our multi-sigs two or three times max., depending on certain cryptographic factors.

BB:

And which transactions would be recorded on the tangle finally?

PH:

We'll get to that in just a minute. We start out by making. Once we have our channel funded, we make another few multi-signatures. Just like the first one. But the tangle is probably never gonna hear about them. We're going to construct a tree and we're going to send from our funding one, our first one, we're gonna call that the root. We send from our root down the full amount that the root

has, down to the last one that we created. And the last one is going to send some to your outputs, some to my outputs and the rule is that whatever is not sent to yours or mine gets sent back to another multisig that we both control. Like a refund address, no, a change address. Down that bottom line is where we actually make the change to the state. When you want to send me five IOTAs, there's a couple of things we have to do. We release from your deposit five IOTAs, because those are going to me. And we keep that stake-incentive equal, we're also going to release 5 IOTAs from my stake, if we funded 50-50. If 50 IOTAs, 50 IOTAs, then you end up with 45 left, I end up with 5 left in my stake and 10 to my output, which means that I gained 5 IOTAs. If you were to publish this on the tangle, you would still have control over 45 and I would still have control over 45, but I would have a net 5 IOTAs. As you are moving forward in time and you are just building up a tree with these bundles where you're pass the entire amount down to the leaves. The leaves spend and then you construct a new part of the tree, for a new leaf. Then at the end you just publish the branch following to the latest leaf. So the amount you are sending to the refund address only goes down, the amount in your output only goes up.

BB:

And all the other transactions are not actually transactions on the tangle, but just between the two of you. And are you only reaching a leaf as soon as the deposit is depleted?

PH:

You make a bunch of bundles and you make and sign bundles that pass from the root to parent to parent to leaf, or whatever, from child to child to leaf. All of those intermediate bundles aren't going to get published to the tangle unless that leaf is published to the tangle. That leaf bundle. That's pretty much it, that's probably not super-clear but,… The idea in the end is, depending on the rule that you're following, two uses of the key or three uses of the key, the maximum number of bundles you can have to publish is going to be a log base-2 or a log base-3 is the number of transactions that you decided beforehand that you are going to make in your channel. What is log base-2 of a million, something probably 15 or so

BB:

And that would be the number of transactions we could do via this channel?

PH:

Ah it is 20. So you decide before you start transacting in your channel roughly for how long you want your channel to live. If you decide that you want your channel to live for 2 million transactions, then you are going to, the depth of your tree is going to be 20 nodes.

BB:

And this would enable a million transactions via this channel?

PH:

Yes. A million state changes. In your channel. And every time you make a state change you can assume that it is as good as confirmed instantly.

BB:

A state change?

PH:

A state change is when you publish all the things down to a leaf and your leaf bundle makes some change to your local state – who gets what output. And in the end, the maximum number of bundles that you'll have to publish is 20 bundles for what would have been a million bundles. Instead of publishing a million, you publish 20.

BB:

But just for understanding, it is not some kind of doing Proof of Work in advance for publishing it afterwards?

PH:

Right, you only do Proof of Work at the very end. When you are going to publish it to the tangle. So if you wanted to do a billion transactions, that would be 30, under the 2-use-rule.

BB:

The next question is already done, I have found that in the whitepaper.

PH:

Ah yeah, you can reference the same thing twice.

BB:

Exactly, yes!

That should be an easy question: Is it problematic to send out a new transaction before the previous transaction confirms?

PH:

It is not necessarily problematic. You know it is possible to merge them but I think the best thing for that is to indirectly reference the previous transaction when you publish your new transaction that depends on it.

BB:

We're almost done. Because of deflationary… There is a total supply in the beginning, then some coins might get lost. So basically the price must increase over time, because the amount is just decreasing. That's basically the first question in there, why is it actually this supply, why is it that number? Because if I took the MaxSafeInteger in JavaScript I could get a higher number than the one that was chosen. So for example, why is it not 3 to the 33 minus 1?

PH:

Because that would require 34 trits.

BB:

Ok?

PH:

To represent 3 to the 33 minus 1 requires 34 trits and they are originally designing some hardware that was gonna use 32 trits, it is going to use 81 trits, 81-trit-words, to my knowledge, which is the reason you have probably seen on a forum, for example reddit, that there is a plan to increase this supply. And I don't know if it's announced but that would be to 81-trit supply. Or should be 3 to the 81 over 2.

BB:

But that would be much than the MaxSafeInteger in JavScript, right?

PH:

Yes, we would be using BigInt in JavaScript or something like that.

BB:
I am not quite sure about the reasons there, is it not enough money?

PH:

The reason is hardware efficiency. The Jinn hardware, that IOTA was born out of, there's a hardware overflow which you can trigger if you have 81 trits and add the overflows and that's more efficient than having a checked overflow. It's for reasons of hardware efficiency.

BB:

It'll probably happen while the coordinator is still on and while we can do snapshots.

PH:

Yeah, I think so. I don't know about the exact date. If you look at the transaction structure, the space is already there. But it is only 32 bits that is used.

BB:

Interesting, I didn't see that yet. The last difficult question: maybe you wanna comment on quantum computers, because it's probably one of the more interesting things about IOTA, except for that it's a much more efficient technology compared to the Blockchain, but it still has this revolutionary thing that it works with trinary instead of binary. Do you think quantum computers or are you just doing that for …

PH:

I do think that it is a bit naïve to, if you read just even current popular literature, quantum computers are already a thing. Trillion qubit quantum computers, you don't see those, but I think it's naïve to assume that some state actor would not already have some form of

quantum computer. And I say that because the stuff that leaded out from the NSA many decades ago was showing that the NSA was decades ahead of commodity hardware and if you take into account Moore's law, I think from that standpoint and from the standpoint of the increasing capability of states to use coalition to keep this kind of secrets, I think it is quite naïve to think that there is not already at least some adequate capability of quantum computing.

BB:

So you would not only say it is gonna be a reality soon, but it might even be a reality already for certain parties, for example some National Security Agencies or stuff like that?

PH:

I think it is likely to come soon to the public I think it's with high probability held with some state actor as well yeah.

BB:

Amazing, when people talk about quantum computing it always sounds like, like fusion. Kind of – would be cool, but nobody get it done.

PH:

I think back in 2013, I remember some Australian researchers, who came out with some method of doing some pretty efficient "on-die"(?) gate, but my personal heuristic is if it is out in the public literature, it has probably been in the secret literature for a lot longer.

BB:

Thank you for answering all those questions!

[...]

BB:

Maybe about the Curl Function. Is it going to stay in use for certain things, because actually it works properly except for this collision resistance.

PH:

As Sergei said, he designed it to allow for practical collisions. He has also never stated that it is a cryptographic hash function, but it is a hashfunction, with 81 rounds, as were actually heading to the snapshot, it is probably a cryptographic hashfunction at that point. 81 rounds should move it outside the brute-force window, move it out to where brute-forcing is efficient as using any other analytical technique to create collisions.

BB:

Is it gonna stay in the code?

PH:

Curl-P will probably be obviated for Kerl, which is still in the middle of development. There is a team of cryptographers working on reviewing, …, I don't personally know what exactly they are reviewing, I obviously know who some of them are, but I can't disclose that. They are world-renowned cryptographers revealing the new Kerl which has been planned for a long time. The new Kerl is the not-yet fully defined, there is still something like 81 trillion variants of it. It's being designed by genetic algorithm.

BB:

So is it like rolling your own crypto? Can you say it like that? Because you are inventing your own functions instead of using the normal ones?

PH:

I'd refer you to Sergei's post. He talks about the need for a paradigm shift. He posted this to reddit. The time for a paradigm shift has come. Pretty good layman explanation for the reasons for what he did. Ah yes, it is pinned on there. I probably couldn't say much better than what he had already said except.. What is a standard hashing function that you know of? And what's the […]

BB:

Well, I think there is this new standard as you said before Keccak, the sponge based hash functions. They were reviewed for 9 years or even longer.

PH:
How long was SHA-1 reviewed and MD5.

BB:

MD5 shouldn't be in use anymore. I don't know how long…

PH:
But it was in use, but it was broken. I think that it is a fallacy that something that has been studied for many years is thereby secure. Inertia doesn't make security. Does that make sense?

BB:

I can't say anything against or in favor of it.

PH:

SHA-1 was a big standard. Before it was known to be broken. It was broken by NSA years ago. Many years before it was publicly broken. Or maybe with the CIA, anyway. So I mean a few things, I think to assume that they haven't been involved in the creation of other hashing algorithms like SHA-256 and may or not have compromised that. I think it is naïve to

not take that possibility into account. For one. For two, Kerl is designed near to be something efficient for internet of things and even Curl-P is very efficient. I don't have the exact numbers for it, nevertheless it is quite efficient. I've seen it perform on FPGA very efficient, very fast. In Sergei's paradigm shift post he pretty well lays out that the reasons for the design decisions of Curl, he designed it to be simple, to be easy to analyze. Such that the difficulty for a human to analyze it is going to be roughly the same as the difficulty for an AI to analyze it. So an AI is not likely to come out with a more sophisticated analysis of a cryptographic hash function than the human is. But other hashing functions, they get more complicated, they do more complicated S-Boxes, they do longer S-Boxes, and those things become more and more infeasible to analyze. The S-Box of Curl is A + B + AAB – 1. Done with trinary logic. So if you overflow you go to negative 1. If you underflow you go to one. It's a circle or whatever you called.

Let's go over a couple of incorrect things here. IOTA is not UTXO.

BB:

It's not?

PH:

It's not. You do not spend from transaction outputs, you spent from account balances. Account balances being a one time signature balance. So you could send many times to an address, and when you send from that address, that full amount of all those previous ones, your bundle is the exact same size as it would be to send from only being sent two ones. But there are many addresses and their one time use, because the one-time use signature require one-wayness in the hashing function. Merkle addresses would require collision resistances. It's more efficient to do it with a one-time signature. It allows you to do very intuitive multi-signature schemes.

BB:

The reason why it is not UTXO is because it's One-Time-Signatures and…

PH:

You could just say it is account balances, but the accounts only live for one spend.

BB:

So it is a bit like a hybrid.

PH:

It's a bit like a hybrid. It's pretty much account based, but it looks like UTXO, because you move from one account to another.

BB:

Exactly, because you actually need UTXOs for a bundle as an input and the output is a UTXO again, isn't it?

PH:

A UTXO refers to a transaction that happened on the Blockchain. A UTXO would say, from this Blockhash, or from this transaction in this Block. With IOTA it is only, from this address, does this address have it's balance. It never references the transaction. So it is not at all UTXO. It just looks that way because of the way that the signatures are used.

BB:
You could say that it has all the advantages of UTXO, though, e.g. parallelization? One account is one seed and one seed has lots of private keys, so one thread could for example use half of the private keys and the other half is used by another thread.

PH:

Yeah you could do that. You can also do one to many, many to one, all these types of things, but your bundle size doesn't bloat because of that. Well, it bloats when you make larger bundles. But when you go to spend from a many to one, you don't have a bloated bundle size for that resulting one, like you do in Blockchain. Because in blockchain,… I was a miner, for a very little time, it was a net loss, but got some number of coins. When I wanted to go and gather all those up, it cost me around 50 dollars, it was one address but it was 20 different transactions and I had to gather all those into a new address, because the Bitcoin cash stuff. That cost me like 50 dollars to pull all those in. Because it's from different transactions, from different blocks. But with IOTA, since for that address that received all of those funds, all that matters is the end value of the address. That's why it is more account based. It works really well in Internet of Things.

BB:

I'll have to read a bit more about UTXOs I guess.

PH:

I am probably not good at commenting on the coordinator. I don't know a whole lot about it. I know a little bit, but not more than you could find out by reading the client code. And MAM is not Zero-Knowledge-Proofs.

BB:

It's kind of similar, isn't it?

PH:

No, it's like a message stream, it's like RSS. It's Masked Authenticated Messaging. You take a message, you sign that message, and then you encrypt that payload. You publish that to the tangle. You share the address, or actually the merkle root, or let's just say the public key,

that you used to sign that with. You share the public key with someone else on the tangle. They can then find your message and they will be able to trace your message forward each time that you publish more messages in the future.

BB:

But if you encrypt it, then nobody know what the value was in the transaction, right?

PH:

It's not a transaction, it's a message. There's no value.

BB:

So it's just like RSA scheme?

PH:

The Masked Authentication part, it's about this channel thing. I can give you a 100 Byte key, half of it is symmetric encryption key, the other half of it is the public key, so you can find it on the tangle. That's like an entry point into my channel, to start reading messages that I publish. That second part of the entry point, that root, changes throughout time. That part is transient. I am gonna make new merkle trees as the future goes. At some point I could go to someone else, I could give them a different key, which brings them in into the latest state. But they can't go backwards in time and look at the older stuff.

BB:

It's only for messaging. If I wanted to send tokens with MAM, that wouldn't work?

PH:

No that doesn't happen.

BB:
It's like PGP messaging?

PH:

It's like PGP messaging, but PGP is only the encryption part. This is encryption and that stream of messaging going forward, so it's not only about how do I read it, but also how do I find the next one. And how do I find all of the ones in the future. You could start with a small key and find stuff all the way through the future. It's not sending it to you or anything. It's more like a way to listen to a stream on the tangle. Almost as if it's a radio stream. An encrypted radio stream.

Proof of Stake on the tangle, I don't think that works. Tanglemath I think Sergei made some comments on it. Essentially, just the way that tangles work and the 2-D nature of them, Proof of Stake is very difficult.

BB:

The other question, though, would make sense, that startups would emerge, which do the Proof of Work for companies. It's a bit like fog-computing, right?

PH:

Yeah, I think that may happen. The network-bound stuff might even end up including something like that. I don't know for sure, I don't know enough about it.

BB:

Thank you so much, Paul!

Interview Transcript

In the following interview, AR stands for the interviewee Alexander Renz and BB for the interviewer Bennet Breier.

BB:

You describe yourself as New Mobility Enthusiast. Are you also a car or motorbike enthusiast?

AR:

So let me just respond in English since you ask the question in English. Maybe we can switch around, I guess you don't care anyway. When we talk about New Mobility, I used to be really interested in cars, but frankly I am not that interested in cars, I mean I like a nice car that doesn't give me a headache. I like beautiful cars, but I am not enthusiastic by a long shot. But when I say New Mobility, it is really the future of mobility which we describe alongside 5 core themes. One is connected mobility, autonomous vehicles, electric mobility, shared mobility and urban mobility. New mobility, when you look at the united States, it's typically handicapped people and wheelchairs. That was the original connotation of New Mobility, but we have now established New Mobility as a new term, so the industry now refers to this New mobility and we actually also have trade marks around the mobility world and so on. When we talk about mobility we really talk about how to move people in good surround and thinking about the digital transformation of mobility and transportation. And many ways the car is in some ways maybe a problem. We realize that the German industry and trhe German well-being is very much dependent on the car. So we are motivated to transition it into a sustainable future.

BB:

I guess you yourself have a car and somehow, would you expect more from your car or what made you think we have to have something new. Is it from your personal perspective. The benefits you see for yourself or what made you go into, or, trying to get the future of mobility right?

AR:

I used to work for different technology companies. My first job out of college was Bosch but I got bored pretty quickly there. I worked in Thailand, that was interesting, because that was a new emerging market in the 90s. Working in the headquarters was always very boring and painful to me. So in 1998 I joined SAP which was a very big startup at the time. I was employee number 28 thousand something. But I started with IoT at SAP in 2000. I initiated the sponsorship of the MIT OOID center. The same meeting we learned about the center and RFID and the IoT, that guy Steve Davy, the CIO of Procter and Gamble also introduced us to a company called BIOS group that was a spin off of the Sanifer Institute. At the Sanifer Institute there was a guy called Stuart Kaufman who

was the founder and let's say the godfather of complexity science. It's like chaos theory on steroids. He had this consulting and software development company that was a spin-off that dealt with swarm intelligence, autonomous agent based syystems. That was for me the more interesting thing about, and RFID would just be one sensor amongst many other sensors that would follow to create what I would like to actually term as real-world aware applications.

BB:

From the way you got to different companies it somehow got you to futuristic technologies.

AR:

I know it's long winded but it's relevant what I am saying. I worked on autonomous adaptive systems. We built prototypes for Compac and Procter and Gamble. But the missing thing was, we didn't have an infrastructure to deploy such distributed systems, because they were, by their very definition, they were distributed across sites, across entities in terms of legal entities. There was no infrastructure. So we looked at open-grid infrastructure and other technology which very advanced in the early 2000s. Then I went to Microsoft and so forth. But this whole issue of how do you really manage such systems at scale and how do you create things that become their own economic agents.

BB:

What you mean with systems it's not a system that one company provides but it's systems by different companies that work together.

AR:

One thing we built for example, would do replenishment. The idea was that Walmart would have all these smart shelves and it would have its distribution centers enabled with RFID and adaptive agents. And then Procter and Gamble would produce and replenish against actual demand. So we created an agent that would calculate the utility function at any one stocking point, like an intelligent shelve or a store at the aggregated level, a distribution center and so on. For incremental replenishments. We looked at probabilities since you have supply risk, demand risk. We looked at what would be the utility for individual nodes in the supply network so that they would then negotiate then for replenishment and priorities. There was no longer a static priority. For example that Walmart is more important than target like in existing systems. But it would be based on the current context. The rate of demand, the inventory levels and so on, that we would replenish one versus the other store. It would be a negotiation across enterprise boundaries. I really ran this.

When I worked with BC, that was the thing I had been looking for for like 15 years. Identities now, asset transfer, monetary exchange. The perfect environment. The reason why we got into mobility is, I grew up in a car family. My father trained the former "Vorstandsvorsitzende" the head of RnD of Daimler, Thomas Weber, to become an industrial tool maker. And then he went on to study. We always thought I go to school with the CEO of Mercedes Benz AG back in those days. So we thought we would end up in Daimler, but luckily we ended up in the Software industry. When I joined Microsoft, my general manager was Satya Nadella, the current CEO of Microsoft. He was my sponsor for IoT and all that kind of stuff there. Then he went to the online division. When Steve Balmer

realized that Google is not a house of cards, as he liked to call it. This online division to fight Google around with Bing and online advertising and so on. Microsoft back in those days was still the most powerful company. Almost unlimited financial resources. Technical talent and so on. Satya was the guy to burn 2.5 billion dollars a year, for several years. And he did not really make much progress at all against Google.

BB:

What were the 2.5 billion burnt on?

AR:

Just building data centers, buying companies, building the bing search engine, building advertising exchanges and so on. But of course the thing is once you have a platform like Google. For example search is a scale game. The more search you have, the more queries you have to, the better your relevance will be and people will use your search engine. The more people use your search engine, the more people will advertise on it.

BB:

It's a reinforcing system.

AR:

Exactly. The more people advertise on it, since it is a bidding system, the higher the returns will be because you have a liquid market place. The more value about the customer the more valuable the impression will be. This is important because the problem now of the future of mobility becoming digital now. And there will be more and more data-driven business models in the future. My question is how will Daimler, BMW, Bosch, VW and so on, how are they going to compete with a Google or an Amazon or Facebook. They all have no interest in building a car but in really getting to consume the time that people will spend ultimately in an autonomous vehicle. And they will have nothing better to do than searches on Google and buy stuff. There is a risk of commoditizing the auto industry, like we have seen with media for example. The media industry has been in this game for many years now where facebook, google only aggregate content that other people write. 96% of growth in online-advertising goes either to google or facebook. Life sucks for a media company or journalist.

BB:

Am I getting it right that that's the reason why you see high potential in the automotive industry. To make it digital, to collect data from sensors. Is that the reason why you are consulting the IOTA Foundation.

AR:

The whole notion here is that, when we look at the media industry or the automotive industry, the problem is today that there is not a level playing field. All of the technology capability and financial resources aside, the legal framework is very different for facebook compared to Bild-Zeitung. Likewise, when Google does not build a car they say, "we are not going to compete with the auto industry, we wanna partner with the auto-industry". And why is that, because they do not have any

interest in building a car because it is one of the most regulated industries out there. It is very complex to build a car and there is all the safety-regulations that you need to standardize. Therefore, it is a lot harder to build a physical thing like a car compared to building an algorithm that you deploy into the cloud or a mobile-device that you updated over the air. That aside, just the legal environment is very different. Bild-Zeitung, as Karl Dieckmann said in another panel, facebook offered to Bild-zeitung to look at their audience analytics. Facebook has of course very fine-grained abilities to look at data of individual readers to target them with specific ads or messages. But Bild-Zeitung cannot even by German law look at this analytics. The thing is now we have no level playing-field. That is on the one hand regulatory. But the other question is really who is holding data in our future and who controls what happens with that data. When you think about a vehicle as a platform like it will be in the future. Like your phone is a platform. That cannot only sense the in-vehicle data but the environment, the surroundings, with all the lidar camera systems, but also now the inside of the vehicle, because there are cameras that are passenger facing. We will have sensors that measure the heart rate, a camera that will analyze the mood of the passenger. Tell us if somebody is happy, depressed. We can classify somebody's sex, age.

BB:

Do you have some example for a specific use case which would create some benefit for customers.

AR:

Just the important thing, just to close that thought is that, when you think about all this data. There is we don't want a future where all the data is owned by facebook and google and the fact that you are depressed all the time goes to your insurance company. If you ask them for life insurance they say, you will be disguised suicidal. They would rather not insure you. Why I am so passionate about BCs is that it is a way for us to give data ownership and control to people and things and then be able to monetize and give a person a chance to do so. To monetize this data because it is very valuable. But very important is of course is that people have a say and the decision making. Like day-to-day that they eat healthy food. In future we need people to understand the value of their data and they have to make smart conscious as to who do they share this data with and what happens with the data and with tokenization of data you can of course discrete pieces and you can share it for a specific purpose and you can define a clear monetary value and attach it to the data. That's the thing. How can we create friction for the internet services. Because now you say, the consumer owns the data and it is not a free for all, where all the big platforms have access to the data. And then of course in the future they can generate, using the data, all this artificial intelligence. The internet has become very centralized. It was of course not the intention to create the internet we have today. But it was at some point all about democracy and creating equal opportunities, but now the world is very different. We have five big companies and some in China that control the internet. We have nobody in Europe that has a significant role. For me IoT and let's say DLTs should be very strategic to Europe to create the next generation internet and ultimately create these equal opportunities for people, around data and industry 4.0, around AR, VR, AI and so on. But really based on a different model where we don't create global champions that then control all the data and all the key services but where the BMWs, the Boschs and so on can benefit and create value and we as consumers can live in a society where, it's a free society and democratic society, where it's not like you can blackmail

anyone you want to blackmail, because there's big honey pots of data that intelligence agencies can tap into, and know everything about every citizen.

BB:

So do you see the IOTA Foundation as the hub for data security?

AR:

Maybe IOTA Foundation's role is to create the ecosystem for the open-source IOTA protocol. Similar to the linux foundation, the idea behind the IOTA foundation is to create an ecosystem and foster the future where people build applications on this technology.

BB:

About the applications building on this technology. Theoretically, there could be competitors of IOTA. For example, of course Ethereum with the Blockchain, copycats of the tangle, also Amazon or Google could also build a tangle. Why or how does IOTA make sure that these projects that you also mentioned definitely build on the IOTA tangle.

AR:

Of course it is still early days and as of today the tangle is a unique approach and pretty much the only technology currently available that people can get the benefits of a blockchain without the limitations of the blockchain. It is all about fostering the community. I think IOTA has a large, biggest Slack community. There is a lot of people working on it. It is not only getting people to work on proof of concepts and getting familiar with the technology which means really building this ecosystem. Putting out the developer resources and all that. Getting people to build that stuff. And of course at some point you have this network effect. Ethereum at this stage has a larger following simply by the fact that it has been around for longer and it's a great platform to build things. But of course it has limitations when you actually want to move something into production. Ultimately everybody will be happy if the IOTA tangle is winning or is seeing broad adoption. It is also about … Nobody at IOTA would say that IOTA is the only ledger, the only viable technology. It is still very early days. It is still a really small community that thinks to understand this technology. One of the key issues is how do you also make it more tangible for decision makers? The typical board member in an automotive company today has just about understand halfway what the cloud is.

BB:

What do you tell them? I could imagine that it is the best if you tell them a specific use-case, like what is this technology actually for, especially in the automotive sector. Because from my perspective it is basically a ledger which is trustless and connects companies which share one ledger, one database. What does the user get from it, so to say?

AR:

I know you want a use-case and you will get it. Let's say, first of all, what's important is, why this technology is so interesting and why at the same time it is so challenging to get people to understand it and take the lead, trying to embrace it. If you look at today's automotive value chain, it is a very

siloed value chain. It's like tier-2 suppliers, tier-1 suppliers that supply the OEM with system components. Sample the car, the car works, there are no quality effects. Now they ship it to the distributor, the exchange with the dealer and on to the customer and that's when the customer drives. Nobody really cares if you are using all the infotainment features that the dealer has sold you. If you are like me you thank god that you can actually drive the car because if you were to sit and had to work your way through the infotainment system, you might become suicidal pretty quickly. There is one thing that is changing already, when the car really gets connected, the OEM needs to maintain this ongoing relationship with the customer. And would try to monetize the entire lifecycle of the vehicle. And sell new services, new features to the consumer. There is a bunch of stuff there that we talk about how the blockchain can enable that. What's more important is, this horizontal value chain, siloed value chain, is something that is today a key asset for the auto industry. They have developed these relationships and have this engineering competence. But in the future the user must be in control and his digital life has mobility needs. There will now be a new ecosystem, business network that is forming now in the future of mobility to address the individual mobility needs of you, me, whoever. What that means is, we will have different modes transportation.

BB:

It's not just connecting card, but connecting all means of mobility.

AR:

For example when you think about one-demand shared mobility. You might do ride-sharing today, car-sharing tomorrow, ride-haling. You may use your own car. The next day you use public transportation. And do the last mile with a bike sharing service. Imagine now in this world you had one shared identity. This is why I mentioned Microsoft earlier on. What you need is an identity now so you can authenticate yourself across all of these different services. Because the auto industry may still think today that they can establish a Volkswagen-id or something like that and become a mobility service provider like Ford is trying. It's gonna fail. Microsoft tried to do Microsoft Pass, but today if we have any identity out there in the internet, it is facebook id, google id. Whenever you use such ids you give away data and you are fucked. My view is, we need a digital wallet with a self-sovereign identity. That is implemented in the blockchain. Now you have attestation. For example like the government can attest that you have a driver's license and your insurance can attest that you have an insurance. But now when you look for car-sharing, maybe one day you use drivenow, the other day you will find there's a car2go right there where you are, but today you don't have an account with car2go. There's not a friction today, that you need to get your driver's license validated and all this stuff. If you have a self-sovereign identity you could now decide, give to this car2go my credentials, but only the data that you need to share to get this car rented. Not your life-history, not who your girlfriend is, and god knows what. You are in control of data. Actually if you look at cyber-security risk. All these companies have no freaking clue about cyber security. They think…

BB:

One quick thing to the car2go thing. Actually I have car2go and also drivenow and I had to validate the driver's license with both of them. So what would be enabled would be that when I have it confirmed once, that's saved on the ledger, in my wallet or id.

AR:

Yes, so you get that. The other thing, when it comes to data-control and data-ownership, that you heard in Germany about the Equifax hack. It's like a big financial credit rating company. Hundreds of millions of records got hacked. Just think about it. The people at all these OEMs and tier-ones, they would in private conversations admit to you there is no fucking clue about cyber security. They think they should be the ones to be all these data into their cloud, because the kind of understand that the data is valuable even though they don't really know how to monetize it. But they know the data is valuable and we should be the ones to put it in our cloud and own it and control it. But there will be a hacker that will breach these systems. That's gonna be a very bad day to whoever it's happening. Daimler, Bmw, Volkswagen, whoever it is. It's much better for all of us if we can separate that and put the consumer in control of their data and we store it in different vaults depending on the security levels. Your financial data would be in a high, more secure vault, than maybe your mobility data. But you control the vaults and you decide who gets access who gets a public key of that data. You also wanna separate data from applications. Only when we wanna rent this car2go, they need to know, yes he got a driver's license, here is the credit card data. Ideally I don't have to store this data in my own systems, because it is a liability. And companies that are smart understand that. The other thing is, we as Europeans have to think that people […]. Look what is Uber? It is modern day slavery. I still use Uber because for me it is like market research, but an Uber is a centralized platform that does nothing but match demand and supply. They set the terms. The drivers make less and less money per mile. At the point they would now tell me. I lost my black limo service and I had my own company. I lost that to Uber. Now I am driving for Uber with a Toyota Prius, but the rates go down and down. I need to buy the car, maintain the car. And the money I get now per mile. I cannot really make a living. Since mobility is so important for a society in terms of access to jobs, access to education and so on. To make sure that it is sustainable but also affordable for everyone. You don't want to have a centralized platform to really have everybody by the balls when it comes to how do I get from A to B. Much better would be now if we had a peer to peer true sharing economy. Because when you think about what does AirBnB or what does Uber, what does Lift do really. They essentially establish trust. Because it is a trusted brand so I think I can get, … When I use this system there is billing and other services, but I know that the guy I am riding with has a driver's license and insurance and has a background check that was done. Imagine if I wanna go to the airport next week, I am sure that in my community in west Seattle there is somebody who is going to the airport at the same time. If I could now establish trust through the network using a distributed ledger where you have computation and you have trust established by previous txs, then we could essentially establish a true Peer to Peer sharing economy where I find the people in my community that go the same route and we set our own terms, e.g. how much should I get because I give the other person a ride and vice versa. Likewise if I knew that I could trust the people, if I could rent my car which sits 99.5% of the time out there. I could generate revenue with my car.

BB:

Wouldn't you still some kind of platform where you could look for offers? There must be some platform that brings together demand and supply.

AR:

In a distributed fashion. You could decide, e.g., once you think about you own and control your data, then you could run your own digital assistant. Not like some Google assistant that basically does nothing but to learn what to sell you next. But it could be your personal assistant who learns your preferences. And the system could predict when you need mobility. And just suggest. You go to a beer with your friends, it knows the parking situation in Munich, it knows the price for ride haling. If you want to meet at 8 oclock according to your calendar, you should leave now and take the S-Bahn. Because you know you will be drink, actually your wearable will be connected to it. And it knows that you get hammered, based on your heart-rate. All these sensors tell you when you should stop drinking, so that you don't get a hang-over.

BB:

That's actually a really important point, because IOTA is about IoT, right? It's about Machine to machine. And the thing with this digital identity is also very important but it's not the IoT aspect, right?

AR:

Maybe this identity thing is even more important when you think about IoT because this whole IoT is a bunch of bullshit unless you can establish a trusted connectivity between things. People put stuff into their homes and all these echo and google loudspeakers that basically listen to your conversations all the time. That is one thing. But people put highly questionable security models, like cameras that surveil your house. That can be easily hacked. And since all of these things are connected, we talked earlier about this emerging business network. Just think about it. It's very strategic when you think about this mobility ecosystem where all these different modes of transportation create one system that provide this seamless access and […] any mode of transportation. The communication infrastructure is one element there that enables the connectivity in the communication and data flows but also there is an energy infrastructure, because the future will be electric mobility. You generate energy using your rooftop solar. And now you maybe need to feed that energy back into the grid. You maybe wanna do energy trading. You wanna address the electric charging infrastructure. Why wouldn't I open up my power socket to somebody who wants to charge their car on my house? If I could peer to peer trade that energy with them. But also you have the vehicle to grid, so the vehicle could be a storage device.

BB:

This would make the payment, … I remember Dominik Schiener talked about this, that there is the charging station and the car and they can communicate which would be a really simple, nice use-case for the tangle.

AR:

When you think about the devices out there: The electric charging station, the wifi access point, the smart lighting. The camera. All of these are entry points for hackers. If you could hack one of these devices, you could hack the entire system. The future warfare will be based on information wars and cyber wars. Imagine if you are, as Germany for example, you have this convergence of the infrastructure. Somebody is hacking this thing. Now you have the most critical infrastructure, energy,

transportation, communication, could be infected. You need to make sure that you have trusted activity between these different devices, so that you actually know, … If you think about the platooning scenario in automotive. You know what that is?

BB:

No I don't

AR:

Platooning is if you have different trucks that form like a train. Only the guy in the front ultimately may have to drive, the other guys can stay in the bunk bed. But of course the idea is that you improve the fuel efficiency. You need a vehicle to vehicle connectivity so that the lead truck tells: Yes you follow me. How do you make sure that this lead truck, or the thing that claims that it is a truck, is actually a truck, and not a fake id?

BB:

How do you make that sure? I don't know.

AR:

You can make that sure with trusted identities on the Blockchain.

BB:

And that has to be confirmed by the producer or the company that coordinates the trucks.

AR:

You can issue this guid and all this kind of stuff, but then you have these point to point connections, where Bosch or whatever decide on the guid that we establish. I think what we need is fluid point to point multi thing connectivity that is ad-hoc. And where you not only rely on a once issued identity that could be faked or copied or whatever, but one where you really leverage the ongoing tx history and validation on the network to establish that trust. When you think about this platooning now you solve with the blockchain or a distributed ledger like iota you solve one problem, but, and this is why it is not so easy to, somebody could say: I could do that with some other means. Yeah, sure you could do that. But how do you make sure you have economic incentives between these different trucks? The first guy has the bad deal. He has a driver that needs to drive. He has the least gain in fuel efficiency. He gains a little because there is less drag, but then the second truck is really in the best shape, the third truck is also in a good position. What you really want to enable such collaborative business models is that you really need to do monetary exchange with those trucks. With IOTA you could use IOTA to share the benefits of this platoon equally. When this driver then reaches their maximum driving time, then they need to go to a resting area.

BB:

Are the trucks by different companies? Then you would need something like that.

AR:

And even if you were to say it's eternal, you maybe wanna have a record of when a driver was sleeping or were able to sleep because they were in the third truck on the platoon. The same truck, he needs to go to a resting area. Bosch is building these fenced Raststätten. Now the truck comes, it should be identified automatically using some identity of sorts. They enter, they booked it in advance. Now the identity establishes. Yes, this is how they are, open the gate. Depending on how long I stay in this resting facility that I can all track of course with IoT and GPS. Automatically bill the truck for the time that they spent there. And if it spent any off-the market services, like changed tires or fixed stuff, it goes into the digital twin of that truck, so it's updated, then we do billing and so on. As we talk about these use-cases, once you think about, … We were talking about why is use-cases and why is IoT, no we talked about identities, exactly. This trusted identity is huge. It's not just people, there is no limitation here, we could take people or machines. But especially long-term when you think about autonomous cars then. What you talked about with Dominik, you will have the ability for inductive charging. That's a very good example. That when you stop at a traffic light and now you need identities established between the vehicle and the infrastructure. Now there might be depending on how long you stop at this traffic light it might be very small exchange that happens here. If you come with Prof. Matthes, come with SAP model to manage this tx, the tx is something from an energy company. They basically realize the cannot make a tx on SAP that costs less than 2 euros. It's so little energy and it's microtxs and you really wanna send ten or fifty cents and now comes your transactional system it now costs two dollars, it's not a lot of fun.

BB:

Alright. I think I've understood a lot of the ideas that you have for the automotive industry also with …

AR:

We are not done with that. We just talked about this ecosystem and really think about this blockchain or distributed ledger as the new fabric that connects all of these different, energy, communication, mobility, everyday things. This identity and monetary exchange. But now let's talk about maybe some more concrete example about automotive.

BB:

That would be great! Even though it was pretty concrete already.

AR:

We did something that the car guy could understand. For example, when we are talking now not about the future of autonomous yet, but let's talk about more down to earth examples. For example, we have vehicle to vehicle communication. Vehicle to infrastructure communication. There you need this trusted connectivity, you need to be able to transfer data in a secure way and one application that is really concrete that is over-the-air-updates. You basically wanna have a record, because there are certain policies involved defining that certain ECUs should be updated. Before you even start that process, you need to configure out between a content provider tier-1, OEM and so on, what updates do we have and are they approved for an update, are they released and all that kind of stuff. It's a

multi party, multi entity issue. The blockchain, or the DL could be a multi-entity-record to keep track of all that stuff.

BB:

Which parties are involved in this? It's not only the distributer of the software that should be updated?

AR:

No, for example, take for example BMW. Would be the one who initiates this updated to your BMW. But it could be that for example Bosch that makes the diesel injection system, that they say: Hey we have this diesel problem kind of thingy, so we might be the ones who have developed the software update. That now has to be sent to BMW to get reviewed, approved in terms of a software lifecycle. From there, it is distributed down to the vehicle. The car owner also has to agree to it. And one thing that is gonna be important, especially in the future when you have autonomous driving, is whenever you transfer data you have data integrity. You have a machine-learning model that you developed in the cloud and you now want to send that to the vehicles so they can use this product as in-vehicle AI. It would not be very brilliant if somebody intercepted that communication and introduced some malware or manipulated that model. And now these cars hit trees and stuff.

BB:

It's not a use-case for the tangle specifically, but for DLT in general, like Blockchain, right?

AR:

Yeah. Many use cases can be done with a blockchain or a DL or a tangle. The practicalities are then more the performance requirements, the scalability, the practicality with tx fees or not, and so on. Another big are that I think is very concrete is when you think about sharing data from a car today, there is this debate who owns the data. The OEMs think, which is starting to change thanks to me, who should own the data. I of course think it is the consumer who owns that data. But the notion is today, the OEM says: Look, I will bring the data from the vehicle into my cloud and then I aggregate it and give access to a tier-1 supplier or an insurance company or whatever. But the OEM is very much eager to control the access to the in-vehicle data and keep the data in their cloud and hold it close to their chest. Imagine a system where the blockchain would be the system that helps you define as an OEM which ECUs, which sensor data could be shared with third parties? And the consumer may then ultimately decide, what data could be shared from their vehicle. The point I am getting to is, imagine for example, a breaking system. A breaking system may generate data that tells you about the friction on the read surface and may tell you stuff about is it slippery or not, and stuff like that. It might basically tell you about near accidents. I am sure there is a ton of stuff that somebody like Bosch might be very interested in, of low-level data of this ECU that Daimler for example may not really find all that interesting. They might filter out that data, they might not even bring it into their cloud and if they do they might aggregate it up. Just to reduce the data footprint, because it is also cost. What you can do for example with IOTA MAM and of course now with a distributed data marketplace is that you could open up a sensor and share the data peer to peer via a secure channel. And essentially stream that data directly to a tier-1 partner. What we are also talking about a lot to

different departments of transportation in the united states, is this value in connected vehicle data to help them maintain the road infrastructure. It might find that a sign needs to be replaced because the reflector does not work anymore or in the winter where should we deploy our plowing machines based on road conditions. When you think about real time maps and high definition maps for autonomous driving like for Here maps. The question is how can the connected vehicle update these maps in real time based on sensor data that you, with the camera and lidar and so on, that senses the environment. One OEM has the brilliant idea to say: Let's take the mobile camera and update Here maps based on some general terms and conditions that any of our customers will have to once accept for the car to actually become useful to them. Whenever we see certain information we send this to Here, so they can update the map. The will say: We don't store any personally identifiable data. But as you know, as a computer scientist, you can pretty easily use patterns to recreate who is who, even if there is no personally identifiable data attached to it. All I am saying is, the real way, the best way to do this is to use something like in IOTA MAM. With a different system, it's like in real AI to detect a relevant event. Now you could essentially update and share this data with the department of transportation, with Here, whoever it is. In a way that it is anonymized. So they don't know your identity. Your privacy when you see your second girlfriend or your third, is protected. It would not share location data or so. It is protecting your privacy but at the same time, the department of transportation or Here would be sure that it comes from a trusted source of data. This is very important when we talk about the platforms and the ownership of data that we have today. I am sure you know these people also. Data Scientists or people who like to work on AI to go to work for companies that have a lot of data. Because you need data to do that kind of stuff. But imagine if we had all of these things in the blockchain on the iota tangle, the iota tangle could become the biggest source of high-quality trusted data that could drive machine-learning and AI with. That would not give up our control and ownership. We as humans could decide what happens, what is the future of AI, how does it help us.

BB:

That is actually already a really good outlook. I don't wanna take too much of your time. I realize we already have one hour.

AR:

For me this is, you are going to be a multiplier. You are a young guy, we need people, like, I could go on forever on this.

BB:

You have already mentioned a lot of use cases and explained to me. I find it always really mind-boggling. I will have to listen to the recording again, because you mentioned so many things that are, I have to think about them again to …

AR:

It's always simple stuff that you can just use to get people to understand the concept. For example in Germany there is a lot of odometer fraud. Kilometerstandbetrug. Billions of value are destroyed, or created depending on how you look at it. It is kind of mind-blowing to think why wouldn't you as an

industry also, why wouldn't you really want to put this data onto a DL and in the same way you really want a digital twin for every car. This digital twin, we are talking about an identity. It would have its complete configuration to start with. If you are, as an OEM, you will ultimately become liable for a car once it's in autonomous mode. That's pretty certain that is gonna happen. As an OEM you are highly motivated to make sure that you got this under control. Today there is a lot of fake after market parts. People make all this shit in China and so on. Bosch and so forth they know they lose a lot of revenue, because of it. But they also don't want to confuse the consumer too much and make them aware that there are fake parts out there.

BB:

But eventually they would be responsible if there was a fake part and if this part was responsible for an accident, then they would need to know.

AR:

Yeah exactly, but more important is of course how you prevent it. When you think about a lidar. The lidar system that is one key component in an autonomous car, or a camera, like a mobile eye camera. You wanna make sure that when it comes to the chip for example, when the chip as a core component of the system is created, that you issue that identity. That you have all the test cycles recorded for that chip. So you have really that quality assurance and all that kind of stuff. Then it moves through the supply chain. If for example your vehicle detects using in-vehicle AI one of these components is malfunctioning. Now you wanna replace that lidar. That replacement part you wanna basically track from the source all the way through to this shop. Then in the shop you wanna make sure that all the test equipment that Bosch and all these companies provide, to test these automotive systems. That they are all calibrated correctly. That they run the latest software. You wanna make sure that the technician that is performing this work has the training and skills to perform it. As you replace that lidar this becomes a new part in the digital wallet, identity, twin of that car. Now it is a part of a new system. Your whole after market history is recorded in that. If there is an accident, you will have a black-box. How do you make it so that it cannot be manipulated?

BB:

Aren't there black-boxes in airplanes and so?

AR:

Yes. The thing is always could somebody, it's one airplane, but once you have millions of cars out there, how could you make sure that nobody could change that blackbox after the fact. And of course you know the answer.

BB:

Well, yeah you store it encrypted with an encryption key somehow…

AR:

You put it on a DL.

BB:

It depends. The blackbox just collects data that it gets from the different sensors and you could just store it locally.

AR:

But it would be an added benefit if you hashed that blackbox every now and then on a DL.

BB:

It depends on for what purpose you have a blackbox. If you only have it for the case of an accident, then you would need the data only for analyzing the accident.

AR:

Once you put it on a stupid ledger, it is an immutable ledger. If somebody tampered with the blackbox, because they tried to whatever, prove that they are not to be blamed for the accident. Somebody might have an interest in changing it afterwards. This could just be one of the security mechanisms to avoid tampering with a blackbox. Similarly, what you can also think of is, in cyber-security, you have different ECUs in a car that are all very insecure. Or in general you have a system that run a certain firmware and they have a certain software status and so on. So if you hashed that data and put it on a DL you could detect when somebody has tampered with that firmware.

BB:

It is not much of an effort even. And not much data. Sounds good.

AR:

When you think about this you see that it is all about such an ecosystem that has overlaps. These data-silos.

BB:

The use-cases are all in the same area, but they give different benefits. And if you have the overall system all these use-cases would be included.

AR:

Yeah, but also you could do for example, imagine electric mobility. We talked about charging. Another thing is now, the value of a battery will be detmined by its charge cycles and how these charge cycles happened. When you wanna sell a car, one interesting thing will be what is the value of the battery still. How much life does it still have in it? If you were to store all of these charging cycles in a DL, you could basically have an immutable ledger of how many times has this battery been charged and how and so on. That could become the new reading in the future.

BB:

For electric vehicles that is super-important. Because that is perhaps that heart or, the liver, of the electric car.

AR:

An important thing to me is, and this is what I am trying to tell the executives. It's like, they all feel proud that they understand how Uber works and how you can monetize data and create new personalized experiences maybe. But I just have to tell them, you are going to be fucked. Because if Microsoft with their resources and their money, would you make two and a half billion in losses for a couple of years? To compete with google and facebook for digital services in the car? The important thing is we need to think how they can create this level playing field or change the game in favor of those companies. The best way to do that is really this data control and ownership by the consumer, because it creates friction for the platforms and protects our privacy, makes us aware of the value and for example, when you use your BMW to go to the city and go shopping, BMW will be able to find you a parking spot now. Finally. It took a long time and this is the sad part that it took so long and other people to show them why parking is important for the car ownership experience. And why this navigation crap isn't really cutting it. In the future BMW based on the data that they have access to doesn't really know what you wanna do in the city. Do you wanna buy what, where, what do you wanna see. Google knows all of this. If you think about a DL in the future where, somebody else, I don't know what your top retailer is in Munich, but somebody at the Kaufinger Strasse. If they know that you wanna buy a Gucci bag for your girlfriend now that you have graduated and you are now rich. They might wanna say, sure you wanna buy a Gucci bag, they will park your car for you. And we will pay for it if you buy a 600 euro bag. If you have a world where the consumer decides, do I give BMW access to that data? He might get something back. Chances are much better for BMW to ever see that data compared to a world where facebook, google and amazon basically have access to that data and you don't, you are only the hardware. The other thing is then, what's also gonna happen in the industry is if it moves more and more to that shared ownership model and less car ownership. Even in the world of car ownership. It is not going to work that you equip a car with all of these features that nobody ends up using. But what you have is the ability now with all these passenger facing sensors to be able to actually predict when somebody may need a certain service. We have demonstrated with a company called OSR at Frankfurt how you can use all of these sensors to classify people by age, sex, mood. Analyzing gestures to understand what are they doing. Tracking eye gaze.

BB:

So classify them into moods or into certain personalities?

AR:

Yes, are you happy, said blablabla. What are you doing, are you on the phone now. There is a certain gesture to tell me you are on the phone. Are you buckled up, e.g. in a taxi. I could use a sensor, I could use a camera to analyze have you fastened your seatbelt. Is the passenger getting in still the old grandpa, is the person still entering the car, so I shouldn't drive off now, right? All of this stuff you will need. How will a robotic taxi find somebody in a big crowd on timesquare? I could be done with facial recognition. But I wanna make sure that my biometric data is controlled and owned by me and not by Uber. Or sell it to the highest bidder. If you have all of this sensor data in the car, you can argue so that this is kind of creepy and so on. The other day I drove from my house here to the store and I realized that I am really driving aggressively. And I didn't realize, I didn't have any problems

today, I didn't have any bad thing going on, why am I driving like an asshole now? And I realized my girlfriend usually tells me that I am hungry before I am hungry, because I act differently, I am hangry. I realized I am actually hungry now that's why I am driving that aggressively. If my car could detect for example that I am hangry and I am not driving very safely, it could tell me, hey why don't I take over. What if I road rage, cursing around and whatever. The autonomous car could say why don't I take over and by the way you wanna eat some Thai food today and go to a Thai place. If you spend your night with your friends in the beer place and you have a wearable, the wearable can, I have actually seen it, one of my friends in Switzerland has built wearable, before he ordered another glass of wine, he looked at the analysis of his heart rate and blablabla, he says no I shouldn't have another one because otherwise I will have a hangover. I am fine now but if I drink another glass of wine then I will be feeling it tomorrow.

BB:

That's an analysis he did for himself.

AR:

With a highly accurate biometric device that's measuring his vitals and so on. How his heart rate changed over the course of the evening and god knows what other measures he takes there. Let's assume your wearable has sensed there was some drinking going on. We know you were at the Hofbräuhaus. Then your sleep is being analyzed. I don't sleep very well when I drink. Most people don't. Now you get up in the morning and your wearable connects to your car. Usually your car would say, when he goes to work he wants to be productive, so the lighting, the applications that we present you, your email or your text messages and so on. You are in working mode and be productive as the car drives you to work. But now given the context that your other IoT devices have created, the system now knows that hey today it would be really good to get another twenty minutes of rest. When you get into your car, it will be dark windows, sleeping position, you nap on the way, maybe drive a little bit different, so you can actually sleep. When your girlfriend goes into the car to pick up the kids, it knows that kids are now in the car, because it senses that with a camera and the location, we can predict that already. The Terminator that you watched earlier is no longer running. Now it is Biene Maya and the Boehse Onkel song is changed to Michael Jackson Heal the World. Google knows all those things, but as long as you own and control the data, you can think of a lot of personalization and the future, and this is where I am going with it, when you come back from your hike from the Allgäuer Alpen, then your car, let's say you buy a massage function in your car, but then when you are actually driving, you don't wanna switch it on and it's all too complicated. And the future BMW may build all of these features into the car because it is reducing the complexity supply chain and those little heating thingies don't really cost that much money. Now they are by default in the car. You have no more buttons, hardware, like we have today, but you have a big touch display. Now what I can do I can sense all of these environmental and personal factors, I can see hey that guy really could use a massage now and now I deliver that massage function on demand and you pay by the minute and not like today 980 euros when you order the car. Now you deliver this over the air, we talked about trusted connectivity, secure data integrity, but also there will be a small tx that has to happen.

BB:

That sounds so futuristic.

AR:

An autonomous car is not far away, it's gonna be robotic taxis and that's the first thing that will happen. Robotic taxis.

BB:

There is a lot of intelligence in there, though. It is both. It is the connection of systems as well as analyzing, or having some overall analysis run over what the sensors from different systems generated.

AR:

When you have a robotic taxi, you will need a passenger facing camera, you will use lidars and other camera systems to enable the entire onboarding of the passenger. Then the camera would need to report when is the grandpa sitting down when is he buckled up and so on. But then of course you could imagine people sitting in the car and have nothing to do, now there will be monetization models. In car advertising, because I know where you start, where you end your trip. So I know this context. If I were a digital service, like an uber, I may have access to browsing history and text messages, so I know pretty much who you are gonna see, for what purpose. I can have in car contextual advertising and if I know that you are going to Hofbräuhaus or some bar or a club on a Saturday. I know that you are gonna make a decision for a beer in a few minutes. I could run an add that says, Weissbier. As an advertiser I could not only run this ad in this context which is a more interesting impression than hitting you up at 8 oclock in the morning unless you are an alcoholic, at 8 am in the morning Erdinger Weissbier. But I could actually imagine, it is a bit futuristic, but actually six people are working on I could track now, are you actually watching or are you looking outside the window while my ad is running. I could imagine using a cryptocurrency and a token as an incentive, I could pay you for your attention. That's my future that I want. You can say yeah let's watch this ad and get a token for it. Maybe get a discount for when you order it in the store.

BB:

That's really difficult to forecast whether it is gonna work. Whether consumers would actually, use it. Because different people have different, money is relative. For some people it might be interesting,

AR:

It is also a matter of what your financial situation and your level of stupidity is. In the United States for example, whenever people do research on shared mobility versus car ownership. When I graduated from university I go myself a C-class Mercedes. A new one. Not 20 years old. A new one. If you look at the typical student in America, when they graduate from university, they don't have to worry about a C-class Mercedes, they worry about first of all finding a job that kind of pays decent money so they can deleverage themselves and pay off their student loans and their debt that they start their life into. When you have no money, then Uber sounds like a great idea, when you can afford your own car. Doesn't really mean that they wouldn't like the idea of owning a car.

BB:

We should get to an end some time. What was that?

AR:

And you have to listen to that again.

BB:

I'll listen to it and write all of it down. But that's good because then I can actually understand all of what you have said. Because if I am honest it is quite mind-boggling. It is really cool and I am happy I can write this down in the bachelor's thesis also. Maybe just one closing question, you don't have to comment on it for too long, but I am interested in smart contracts which are a problem on the IOTA tangle. Do you believe that IOTA will become the backbone of IoT even if smart contracts where not possible?

AR:

First of all, I mean there is something you need to be, I would not make it a big deal here, and I am saying this now because I don't know if that is public information. So can't write this in your thesis unless you annotate it with Dominik Schiener or something like that. Smart Contracts is a roadmap item. They wanna build a smart contract and they are working on one. I think it is recognized that something like a smart contract is needed that is, … And you could even imagine a world where IOTA deals with the IoT world and lower level lean devices and partitioning the blockchain and enabling offline txs. That ultimately feed into a Ethereum based blockchain where you could have smart contracts, where you could aggregate this data, but that could also be, there is no scenario where you could not have the different systems interacting, co-exist. Smart contracts is in the works.

BB:

Paul also told me about it Then I asked whether they wanna create a language like Solidity that is what Ethereum uses for smart contracts. And he said, can't comment on that. So I guess there is something going on on that. I also know of possibilities or ways of enabling smart contracts in a certain way. There is no time order on the tangle like in the blockchain. So you need some kind of work-around. For example also Serguei Popov, the guy who released the whitepaper, he also did more research on timestamps which would be necessary for general smart contracts. So I guess there will be something about that. But we don't know it yet so I can't write it in the thesis.

Thank you very much, that you took your time, Alexander!

Bibliography

[1] Block chain - Bitcoin Wiki. https://en.bitcoin.it/wiki/Block_chain. Accessed: 2017-10-03.

[2] Building A Base 3 Computer. https://www.youtube.com/watch?v=EbJMtJq20NY. Accessed: 2017-10-29.

[3] Dominik Schiener's IAA presentation today. http://ethereum-infochain.blogspot.de/2017/09/dominik-schieners-iaa-presentation-today.html. Accessed: 2017-10-30.

[4] Ethereum Block Count And Rewards Chart. https://etherscan.io/chart/blocks. Accessed: 2017-09-09.

[5] IOTA AMA. https://www.reddit.com/r/Iota/comments/6yvpfo/iota_ama_september_8th/. Accessed: 2017-10-04.

[6] IOTA Developer Hub. https://dev.iota.org/. Accessed: 2017-10-30.

[7] IOTA Homepage. http://iota.org/. Accessed: 2017-10-08.

[8] Lamport signature: How many signatures are needed to forge a signature? https://crypto.stackexchange.com/questions/2640/lamport-signature-how-many-signatures-are-needed-to-forge-a-signature. Accessed: 2017-10-06.

[9] M. Abliz and T. Znati. A Guided Tour Puzzle for Denial of Service Prevention. In *2009 Annual Computer Security Applications Conference*, 2009.

[10] Adam Back. Hashcash - a denial of service counter-measure. Technical report, 2002.

[11] Jules Besnainou. Blockchain and IoT: A Conversation with Dominik Schiener of the IOTA Foundation. https://www.cleantech.com/blockchain-and-iot-a-conversation-with-dominik-schiener-of-the-iota-foundation/. Accessed: 2017-10-04.

[12] Bennet Breier. Telephone Interview with Alexander Renz. Conducted: 2017-10-12.

[13] Bennet Breier. Telephone Interview with Paul D. Handy. Conducted: 2017-09-22.

[14] Johannes Buchmann, Erik Dahmen, Sarah Ereth, Andreas Hülsing, and Markus Rückert. *On the Security of the Winternitz One-Time Signature Scheme*, pages 363–378. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.

[15] Vitalik Buterin. Thoughts on UTXOs. https://medium.com/@ConsenSys/thoughts-on-utxo-by-vitalik-buterin-2bb782c67e53. 2017-09-29.

[16] Cisco. Fog Computing and the Internet of Things - Extend the Cloud to Where the Things Are. Technical report, 2015.

[17] IOTA Community. An introduction to IOTA. https://iotasupport.com/whatisiota.shtml. Accessed: 2017-10-07.

[18] IOTA Community. How addresses are used in IOTA. https://iotasupport.com/how-addresses-are-used-in-IOTA.shtml. Accessed: 2017-09-08.

[19] Digiconomist. Bitcoin Energy Consumption Index. https://digiconomist.net/bitcoin-energy-consumption. Accessed: 2017-10-29.

[20] John R. Douceur. The Sybil Attack. Technical report, Microsoft Research, 2002.

[21] Arthur Falls and Sønstebø. Ether Review #69 - IOTA & the Post-Blockchain Era. https://soundcloud.com/arthurfalls/ether-review-69-iota-the-post-blockchain-era. Accessed: 2017-08-26.

[22] Steffen for The Tangler. AMA - Answers for newcomers. http://www.tangleblog.com/2017/06/05/ama-answers-for-newcomers/. Accessed: 2017-09-08.

[23] Steffen for The Tangler. The incentive to run a full node for iota. http://www.tangleblog.com/2017/06/27/incentive-run-fullnode-iota/. Accessed: 2017-10-16.

[24] Lewis Freiberg. Instant & Feeless - Flash Channels. https://blog.iota.org/instant-feeless-flash-channels-88572d9a4385. Accessed: 2017-09-29.

[25] Philipp Giese and Dominik Schiener. Blockchains ohne Blöcke und Chain - ein Interview mit Dominik Schiener. https://www.btc-echo.de/blockchains-ohne-bloecke-und-chain-ein-interview-mit-dominik-schiener/. Accessed: 2017-10-03.

[26] Stephan Haug. Statistik für Betriebswirtschaftslehre. Technische Universität München, July 2016.

[27] Ethan Heilman, Neha Narula, Thaddeus Dryja, and Madars Virza. IOTA Vulnerability Report: Cryptanalysis of the Curl Hash Function Enabling Practical Signature Forgery Attacks on the IOTA Cryptocurrency. https://github.com/mit-dci/tangled-curl/blob/master/vuln-iota.md. Accessed: 2017-09-08.

[28] Sergey Ivancheglo. IOTA Cofounder Sergey Ivancheglo aka Come-from-Beyond's Responses to the ongoing FUD about so called 'vulnerabilities' in IOTA Code which never really existed. https://medium.com/@mistywind/iota-cofounder-sergey-ivancheglo-aka-come-from-beyonds-responses-to-the-ongoing-fud-about-so-ea3afd51a79b. Accessed: 2017-10-14.

[29] Scott J. IOTA Consensus Masterclass. https://forum.iota.org/t/iota-consensus-masterclass/1193. Accessed: 2017-08-26.

[30] Scott J. IOTA Double-Spending Masterclass. https://forum.iota.org/t/iota-double-spending-masterclass/1311. Accessed: 2017-08-26.

[31] A. Kosba, A. Miller, E. Shi, Z. Wen, and C. Papamanthou. Hawk: The Blockchain Model of Cryptography and Privacy-Preserving Smart Contracts. In *2016 IEEE Symposium on Security and Privacy (SP)*, 2016.

[32] Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine Generals Problem. In *ACM Transactions on Programming Languages and Systems*, volume 4, 1982.

[33] Frank Li, Prateek Mittal, Matthew Caesar, and Nikita Borisov. SybilControl: Practical Sybil Defense with Computational Puzzles. In *STC'12*, 2012.

[34] matthewwinstonjohnson. Reattach vs Rebroadcast. https://matthewwinstonjohnson.gitbooks.io/iota-guide-and-faq/content/getting-started/reattach-vs.html. Accessed: 2017-10-10.

[35] Alfred J. Menezes, Scott A. Vanstone, and Paul C. Van Oorschot. *Handbook of Applied Cryptography*. CRC Press, Inc., Boca Raton, FL, USA, 1st edition, 1996.

[36] Satoshi Nakamoto. A P2P electronic cash system. Technical report, 2008.

[37] Neha Narula. Cryptographic vulnerabilities in IOTA. https://medium.com/@neha/cryptographic-vulnerabilities-in-iota-9a6a9ddc4367. Accessed: 2017-09-07.

[38] Serguei Popov. The tangle. IOTA Whitepaper. Technical report, Jinn Labs, June 2017.

[39] Merkle R.C. A Certified Digital Signature. In *Brassard G. (eds) Advances in Cryptology - CRYPTO'89 Proceedings. CRYPTO 1989. Lecture Notes in Computer Science*, volume 435. Springer, New York, NY, 1990.

[40] Sheldon M. Ross. *Introduction to Probability Models. Tenth Edition*. Elsevier, 10 edition, 2010.

[41] Dominik Schiener. A Primer on IOTA. https://blog.iota.org/a-primer-on-iota-with-presentation-e0a6eb2cc621. Accessed: 2017-09-07.

[42] Dominik Schiener. Bundles. https://domschiener.gitbooks.io/iota-guide/content/chapter1/bundles.html. Accessed: 2017-10-19.

[43] Dominik Schiener. Discussion Removing Peer Discovery. https://forum.iota.org/t/discussion-removing-peer-discovery/939. Accessed: 2017-09-07.

[44] Dominik Schiener. The Anatomy of a Transaction. https://domschiener.gitbooks.io/iota-guide/content/chapter1/transactions-and-bundles.html. Accessed: 2017-08-16.

[45] David Sønstebø. IOTA Development Roadmap. https://blog.iota.org/iota-development-roadmap-74741f37ed01. Accessed: 2017-08-04.