

The Universal Trust Machine

A survey on the Web3 path towards enabling long term digital cooperation through decentralised trust

Rohan Madhwal

5568412

TU Delft

Delft, Netherlands

R.Madhwal@student.tudelft.nl

Abstract—Since the dawn of human civilization, trust has been the core challenge of social organization. Trust functions to reduce the effort spent in constantly monitoring others’ actions in order to verify their assertions, thus facilitating cooperation by allowing groups to function with reduced complexity. To date, in modern societies, large scale trust is almost exclusively provided by large centralized institutions. Specifically in the case of the Internet, Big Tech companies maintain the largest Internet platforms where users can interact, transact and share information. Thus, they control who can interact and conduct transactions through their monopoly of online trust. However, as recent events have shown, allowing for-profit corporations to harness so much power and act as gatekeepers to the online world comes with a litany of problems. While so far ecosystems of trust on the Internet could only be feasibly created by large institutions, Web3 proponents have a vision of the Internet where trust is generated without centralised actors. They attempt to do so by creating an ecosystem of trust constructed using decentralised technology. This survey explores this elusive goal of Web3 to create a “Universal Trust Machine”, which in a true decentralised paradigm would be owned by both nobody and everybody. In order to do so, we first motivate the decades-old problem of generating trust without an intermediary by discussing Robert Axelrod’s seminal research on the evolution of cooperation in the iterated prisoner’s dilemma. Next, we present the infrastructural and social challenges that a hypothetical Universal Trust Machine would have to overcome in order to enable long term cooperation in a decentralised setting. We proceed to present various Reputation Systems, all of which present promising technique for encouraging trustworthy behaviour in a decentralised network through indirect reciprocity. After this, we discuss the family of emerging Distributed Ledger technologies whose secure transaction facilitating and privacy preserving techniques promise to be a good complement to the current limitations of vanilla Reputation Systems. Then, we present other promising alternative techniques. Finally, we conclude by discussing how a combination of the presented techniques could be used to potentially create the desired Universal Trust Machine.

I. INTRODUCTION

Humans in a society rely on trust in every stage of their life, in every action they perform. Children trust that their parents will nurture and guide them, adults trust that their family and loved ones won’t deceive them. When crossing the street on a zebra crossing, we trust that motorists will obey the traffic

laws, when buying items at the market, we trust in the quality of the goods being provided to us. Regardless of whether one believes that society is a function of divine order or of a social contract, trust between its members is the very fabric of its organising foundation. [1]

More generally, consider an agent, such as a human or a robot, who is required to use limited agency to navigate and take actions in a world with limited direct information available to it at any given moment. In such a world, trust is an important social heuristic that allows the agent to make wagers on the predictive benevolence of other agents. [2]

Ecologist Garrett Hardin defines trust as “encapsulated interest”, since it facilitates peaceful and stable social relations that form the basis of collective behavior and productive cooperation. Thomas Hobbes, considered by many to be one of the founders of modern political philosophy, argues that the natural state of humans is nasty and brutish, however, trust helps to convert that into something peaceful and efficient. In his book “A treatise of human nature”, enlightenment philosopher David Hume discusses the importance of trust to the functioning of a society. According to sociologist and philosopher Niklas Luhmann, trust effectively reduces complexity and risks, allowing for coordination with increased performance. [3] This is easy to understand intuitively since trusting individuals and groups reduces the effort one would spend in constantly monitoring the actions of others in order to verify their assertions. It is easy to conclude that a society without a notion of trust would find it hard to function effectively, or to exist at all. [4]

The growth of human civilization from small-scale hunter-gatherer societies to thriving economies of nation states is testament to the benefits provided by the growth of trust and cooperation inside societies. However, history reminds us that the requirement of trust for facilitating cooperation also leads to the growth of large centralized institutions since these institutions historically provided the best defense in economic transactions against the untrustworthy. [5]

While trust might be fundamental to cooperation in a society, underlying every social transaction is the desire to further one’s personal gain by abusing the trust of an unsuspecting op-

ponent and defecting against the expected trustworthy action. [1] For example, in a transaction where a merchant pre-pays a farmer for their produce at the end of the year, the farmer may be tempted to keep the payment and not provide the promised crops, or provide crops of a lower quality than was agreed upon.

According to Margart Levi, “good defenses make good neighbors”. Hence, the need for such defenses in economic transactions necessitated institutional bases of reaching agreement and resolving disputes that might result from them. Institutions that were able to provide third party enforcement in a transaction were hence able to ensure personal security and the security of the transaction. Thus, they were able to encourage cooperation and grow immensely as a result of their importance in doing so. [5]

However, allowing profit driven institutions to amass so much power comes with its own set of problems. The financial crisis of 2008 which was primarily attributed to failure of trusted institutions such as banks and other financial institutions has led to a growing distrust in such institutions. [6] This was most notably witnessed by the recent growth of blockchain technology and adoption of cryptocurrencies such as Bitcoin and Ethereum as a decentralised alternatives to large financial institutions.

The Internet is the most remarkable addition to how social capital can be built in the world. Collaborative work performed on the internet is continuously changing how humans think about social interaction. It is worth considering how trust on the Internet differs from and is similar to trust in general.

Since users on the Internet often know virtually nothing about each other, all they can rely on is the immediate record of the other party’s behaviour in past interactions with them to decide whether they can be trusted. However, this inability to judge different providers of services on the Internet is not very different from the general inability to judge the quality of other services that are required in the real world, such as doctors or lawyers. Similar to the real world, providers of service on the Internet care not only about the current interaction, but also the result of the interaction on their future reputation. Hence, building and maintaining one’s reputation by acting in a trustworthy manner is a requirement both in the real world and on the Internet. [7]

On the other hand, the most notable difference between the two is caused by the Internet’s unique capacity to allow collaboration and interactions at a global level. Take for example the case of buying an item from a local store, while doing so, trust is generally not an issue and often, all that matters is the perceptible competence of the store in the goods and services it provides. However, buying an item from a seller on an online marketplace like eBay requires a markedly different level of trust to allow the transaction to occur. Since, in addition to simple competence of the seller, the buyer would also require that reliable behaviour is guaranteed. Given two online sellers that sell the exact same item at the same price, a buyer would prefer the seller that has a large number of reviews/testimonials. Therefore, even in simple transactions,

due to the global scale of the internet, the risk of fraud is substantial and hence, methods of generating trust to allow collaboration are required. [7]

Even though the Internet was built on distributed protocols, large scale cooperation was consolidated around a few centralised services where social trust was created and enforced by large profit driven institutions. [8] Specifically, in two key functions of the web, web-publishing and discovery of content, technological institutions such as Google, Meta and Twitter slowly became curators and gatekeepers for the information being published on the Internet and people who were allowed to interact with it. As a result of this, the platforms accrued the power to control and own a large share of the information published and consumed on the Internet.

Recently however, abuses of information and communication technology by such institutions for surveillance, spreading of disinformation and coercion of the public have come to light. Notable examples include Google’s deepening involvement with Egypt’s repressive government and Twitter enabling the Chinese government to promote disinformation on the repression of Uighurs. [9]

Such propensity of Big Tech organisations to abuse their ecosystems of trust for their own profit through privacy violations and misinformation is leading to a shift in the general attitude towards large centralised information platforms. The requirement of a large centralised authority or platform owner to maintain and enforce trust in sociotechnical systems is increasingly being viewed more as a hindrance rather than a help. [9]

A growing alternative to the existing model of the platform driven Internet is the idea of Web3 which is motivated by the idea of using decentralised technologies such as blockchain. It is hard to exactly define Web3 since there is a lack of consensus even among researchers on what the idea of Web3 means. In section IV we attempt to clearly define what Web3 refers to in the context of the paper. On a high level, Web3 can be thought of as an ecosystem of applications which aims to generate trust purely through decentralised technology and mathematical primitives. Thus, Web3 aims to be a “Universal Trust Machine”, eliminating the need for profit driven organisations and allowing for the creation of a “commons” [10] where everybody is free to publish, read, react, and interact with content.

However, as shown in section V, fostering cooperation in a community with the presence of bad actors is not a trivial problem. In a centralised system, it is possible to govern in an ad-hoc manner, altering rules of the system as new problems and trust issues arise. This is obviously not possible in decentralised systems since no one single party can instruct everyone how to act. Therefore, all rules of interactions among the independently acting, self-interested parties must be explicitly and clearly defined before any interactions occur. Further, these rules should reasonably incentivise cooperation and disincentivise cheating to foster long-term cooperation.

This problem of cooperation has been studied in the field of game theory and analysing studies in this field could help

motivate how to develop systems where the best course of actions for neighbours is to cooperate for mutual good.

Plethora of research also exists on models and mathematical primitives for generating trust in decentralised systems, most notably, reputation systems have gained prominence as a way to create safe and trustable communities in decentralised networks. [11]

This survey attempts to explore such mechanisms for generating trust in Web3. In section II we discuss some principles in the work of Evolution of Cooperation which help motivate how long term cooperation could come about naturally. Next, in section III we attempt to define what decentralised networks are and how the decentralised movement came about. In section IV we explain the motivation behind Web3 and the technologies associated with it. After this, we discuss problems one faces when designing a decentralised system which fosters long term cooperation in section V. In section VI we discuss Reputation Systems for decentralised systems and present some promising systems in literature. We proceed to discuss the limitations of Reputation Systems and present Distributed Ledger Technologies in section VII which help cover a lot of these limitations. Finally, we conclude with a future roadmap for the construction of a Universal Trust Machine.

II. EVOLUTION OF COOPERATION

The history of humanity is one filled with conflict, destruction and war. The pursuit of peaceful cooperation is more than just a hippie dream, it has attracted a great deal of research across multiple fields. We posit that the goal of the desired “Universal Trust Machine” is to build a digital utopia where such peaceful cooperation can occur and persist over a long-term time period.

One of the foundational works investigating how cooperation can emerge and persist without a third party is “The Evolution of Cooperation”, a 1984 book written by political scientist Robert Axelrod which expanded upon the highly influential paper he co-authored with evolutionary biologist W.D. Hamilton [12]. The book’s central question is “Under what conditions will cooperation emerge in a world of egoists without central authority?”.

Axelrod held two computer simulation tournaments where multiple strategies for playing an iterated two-player Prisoner’s Dilemma game were solicited from professionals across multiple disciplines. The Prisoner’s Dilemma is a popular game analyzed in game theory where two rational agents are faced with a dilemma, they are arrested by the police and have to individually decide to either cooperate with the police or stay silent. The dilemma was originally framed by Merrill Flood and Melvin Dresher in 1950. A key requirement of the game is that: $t > r > p > s$ and $2 \times r > t$ where t , r , p and s represent payoffs for the different outcomes of the game. If both players choose to stay silent i.e. they cooperate with each other, they are each awarded r , on the other hand if both players defect, they are each awarded s . If one player stays silent while the other defects, the player who defects is rewarded t while the player who chose to stay silent is paid s . Fig. 1 demonstrates

this payoff matrix visually. Hence, although the decision to collectively stay silent is overall the most optimal, individually, the best decision is to defect.

Further, in an iterated Prisoner’s Dilemma game there is a probability w that two players will interact in the next round. [13] Contestants who submitted algorithms to play the tournament accrued points in each round according to the shown payoff matrix by playing against other strategies. The tournament consisted of five iterated prisoner’s dilemma games in total with each game consisting of 200 rounds each.

The Darwinian theory of evolution would suggest that the most selfish strategy would perform the best and while indeed, in a single iteration defecting is always the best strategy, in the iterated Prisoner’s Dilemma the strategy that ended up performing the best in both rounds was a simple “Tit For Tat” strategy. As the name suggests, this strategy was based on the concept of direct reciprocity, the next move of an agent following the strategy is determined by the last move of the opposing agent, if the opposing agent cooperated, the agent following Tit For Tat would cooperate too and vice versa.

Based on the results of the tournament, Axelrod identified four characteristics that he believed led Tit For Tat to perform the best of all strategies:

1) **Niceness**

By being nice, Tit For Tat can benefit from long term mutual cooperation with other strategies that are also nice. However, it is important to note that niceness alone would lead to exploitation from other strategies who are not nice

2) **Forgiveness**

Strategies that are not forgiving are doomed to be locked into mutual destruction after a single defection from an opponent. Tit For Tat allows an opponent to start cooperating again after defecting initially which makes it forgiving

3) **Retaliation**

As pointed out earlier, niceness alone leads to exploitation by uncooperative strategies. By retaliating when the other strategy doesn’t cooperate as expected, Tit For Tat avoids being exploited by such strategies

4) **Certainty**

By being easy to understand, Tit For Tat makes it easy for other strategies to understand what it’s doing thus allowing them to come to a mutually beneficial strategy much faster

Axelrod’s analysis thus provides an interesting set of prescriptions for designing strategies for nodes on a decentralised network. Keeping in mind that not all interactions need to be zero-sum and it may be possible for all cooperating parties to benefit on the long term by cooperating and not being the first to defect seem to work as good principles which suggest that cooperation could indeed organically grow in a pool of egoistic nodes. However, being too nice also has its downsides and any effective strategy should be quick to retaliate to prevent exploitation. Finally, keeping it simple seems to be effective advice otherwise the strategy might risk confusing potentially cooperative neighbours.

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	$r = 3$	$t = 5$
	Defect	$s = 0$	$p = 1$

$t =$ temptation to defect
 $r =$ reward for mutual cooperation
 $p =$ punishment for mutual defection
 $s =$ sucker's payoff
 $t > r > p > s$

Fig. 1. A typical payoff matrix of a 2 player prisoner's dilemma [14]

Further, there are lessons for designers of Web3 applications, the most important being having a large “shadow of the future”, i.e. a sufficiently large w which guarantees that nodes interact with each other more durably and frequently so they have time to develop a mutually cooperative strategy and since they are more likely to defect if the probability of meeting a node again is low. This can be done in many ways including using spatiotemporal structures e.g. clustering of small groups in space [14]

However, there are limitations to Axelrod's results:

1) **Assumptions are too simplified**

Not all real world interactions are as simple as an Iterated Prisoner's Dilemma game. Often participants can communicate with each other and hence collaboration through other means may be a better strategy. Further, it may not be possible for real world participants to necessarily perceive credible threat, or respond to it rapidly and accurately

2) **Not necessarily universal**

In his 2000 paper “Twenty Years on: The Evolution of Cooperation Revisited”, Hoffman [15] showed that Axelrod's tournament was sensitive to the initial population composition and the potential for strategies to make mistakes. Under different initial compositions and assumptions, other strategies were shown to perform better than Tit For Tat

3) **Does not consider indirect reciprocity**

While direct reciprocity is a powerful mechanism, it relies on repeated encounters between individuals. However this is too simplifying an assumption to model human interactions where exchanges are often asymmetric and fleeting. Indirect Reciprocity is more representative of real human exchanges where we help people even if they've never directly helped us before based on some indirect exchange. For example, we may donate to a charity which helps other people. [16]

III. (DE)CENTRALISED

Before considering more contemporary solutions to the problem of enabling long-term cooperation, it is important

to clarify what it means to create “decentralised” trust, what the aims of the Web3 movement are and to identify the main problems that a Universal Trust Machine should solve in order to be considered successful. In this section we discuss what decentralisation is, then in the following two sections we proceed with a similar discussion of Web3 and the inherent problems in creating decentralised trust.

Decentralisation is not a novel concept and has been prevalent in research even outside the sciences. In the social sciences, it boasts a 200 year history and has been a popular concept across multiple disciplines. Examples include concepts such as subsidiarity, democracy, liberty and equality in political science, systems theory and self determination in management and decision science, fiscal decentralisation in economics. [17]

In technology, the concepts of technological decentralisation have been evolving for over half a century. [17] A popular example of a decentralised IT movement is the open source software movement which represents a radical retake on copyright law and involves developing and sharing software in a decentralized and collaborative way, relying on peer review and community production.

The importance and the success of this movement is demonstrated by the domination of multiple areas of software by open source projects. Popular examples are the the open source Apache projects which dominates the market of server software over commercial alternatives from Microsoft, Sun etc and the Linux operating system which has seen popular use being embedded in a range of devices from mobile phones, recording devices to large scale servers in data centers. [18]

The concept of a “decentralised network” was first coined by Paul Baran, one of the inventors of packet switching. In general, networks can be classified as two components, “star” or centralized and “grid”/“mesh” or distributed. In a star/centralized network, all nodes are connected to a single node, hence, each participant needs to go through a central component to interact with each other. While in a distributed network on the other hand, there is no such central node and each node can communicate with each other without going through a centralised point. In practice, a combination of

these components is used to form a network, Baran called such a mixed network "decentralised" because there was no single, central point of failure. [19] Fig. 2 demonstrates these networks visually.

In contemporary modern literature, the term decentralised network is used to refer to networks where the technology, content and infrastructure on the network is controlled by participants and contributors rather than large central platforms. This control is manifested in various ways, such as participants controlling parts of the infrastructure like servers and routers, collaborators owning data in their own private data silos which is queried by the network during discovery, participants possessing the autonomy to decide the operational details of the network, what content needs to be publicised and what needs to be deleted etc. [8] In this context, a popular example of a centralised network would be Twitter which owns all the content that users publish on it, while an example of a decentralised network is Tribler, a peer to peer file sharing system which improves upon the BitTorrent protocol which enables users to share content with keyword search and boasts a reputation-management system to encourage collaboration. [20]

Over the past decade, decentralised networks have received a reinvigorated interest due to the emergence in popularity of cryptocurrencies such as Bitcoin and Ethereum. In his whitepaper proposing Bitcoin, Satoshi Nakamoto proposed a novel decentralised peer-to-peer network protocol which facilitates an electronic payment system. [21] The popularity of these cryptocurrencies has also resulted in explosion of blockchain and decentralised technologies and projects. Proponents and developers of these technologies wish to see a shift in the publishing and discovery of content and information over the Internet away from a few profit-driven Big Tech corporations and into the hands of the users who generate them, guaranteeing the privacy of their data and also ensuring that everyone has a fair and equal voice.

For example, "78 days", a collaborative project between the Starling Lab and Reuters uses decentralised ledgers to preserve historical data important to humanity. The goal of the project is to curb misinformation. It achieves so by ensuring the integrity and authenticity of the information as it captured and stored using a system called Content Authenticity Initiative. It also uses a storage system built on blockchain called Filecoin that requires data providers to prove that they are holding the authentic data and not a tampered version. Most importantly it ensures that the contributors of the information have a way to maintain their creation of the content through the records stored with the data. [22]

IV. WEB3

The term "Web2.0" was first coined by Tim O'Reilly in 2007 to describe an Internet where platforms enabled users to publish, consume and interact with content, and with each other. [23] It was supposed to expand upon the first iteration of the Internet or "Web1.0" which largely consisted of static pages meant only to display information. So while "Web1.0"

was "the read web", "Web2.0" aimed to be the "the read-write web" (coined by Richard McManus in 2003).

Critics of Web2.0, such as the inventor of the World Wide Web, Tim Berners-Lee feel that Web2.0 failed to achieve the vision of the Internet as a secure, decentralised exchange of public and private data, with users' data being increasingly stored in corporate data silos. Instead, to guarantee security of their data, they want users to own their own data. [24]

The term "Web3.0" was coined by Polkadot and Ethereum co-founder Gavin Wood in 2014, he used it to describe an Internet that is decentralised, open and transparent. [25]

The current Web3 movement aims to transform the platform oriented Web2.0 Internet into a decentralised web ecosystem which: 1) avoids monopoly of content discovery and propagation by large centralised actors 2) prevents the spread of misinformation and fake news 3) provides its users the ability to create, exchange and react to information in a secure, private and free manner 4) supports immersive web development [17]

Liu et al [26] define Web3 as a movement which agnostic of any specific overarching applications or underlying infrastructures will usher in "an era of computing where the critical computing of applications is verifiable", that is, an application that conforms to the idea of Web3 is one where all stakeholders are able to verify the execution of the application based on predetermined terms without the presence of an intermediary.

Packy McCormick defines Web3 as "the internet owned by the builders and users, orchestrated with tokens" [27] Defining Web3 with its key property being user ownership is a common approach taken by a majority of research papers on the topic. Hence, Web3 is positioned as the "read, write, own" web. While Web2.0 was a frontend revolution that allowed users to create and interact with created content online, Web3 is instead a backend revolution which aims to change how the created content is stored. Instead of keeping data on centralised data silos, Web3 aims to provide data storage to users in a distributed manner in a way that users can own and monetise the content they created. Thus, it aims for the disintermediation of existing parties such as large big tech companies in data governance. [28]

V. THREATS TO LONG TERM DIGITAL COOPERATION

In order to enable the dream of Web3, it is fundamental to be able to create a commons with communities of users interacting with each other through decentralised networks, free to read, publish and interact with content. However, two broad classes of threats make creating long term digital cooperation in decentralised networks a non-trivial task: Social and Infrastructural threats. In the following sections we briefly cover these threats and establish why they pose a problem to cooperation.

A. Infrastructural Threats

In section I, we motivated why trust is fundamental to achieving cooperation inside communities. Since in a Web3 application based on a decentralised network there are no third parties for enforcing trust, before using a service to




Fig. 2. a) Centralised network b) Decentralised network c) Distributed network [19]

cooperate with other nodes in the network, users look for assurance that the other party is trustable. This is especially true for applications that depend on blockchain technology due to the immutable nature of transactions making it incredibly hard to punish bad actors. [29] Therefore, in addition to social problems, there are also several infrastructural problems stemming from the presence of bad actors who wish to abuse the trust of their neighbours for their own benefits makes the problem of achieving long term cooperation in a decentralised network a non-trivial task. A system that will be able to achieve the stated dreams of Web3 should be effectively able to tackle these problems, below is a brief description of a few of these problems:

1) *Sybil Attack*: In a distributed network, if an entity can control a large number of nodes and hence obtain a large number of node identifiers, they can use this dominance of identities to control the network and undermine the mechanisms of the network which results in a network with less robustness and freedom. Such an attack is often referred to in literature as a *Sybil Attack*, where a *Sybil* is the fake identity of an entity. [30]

The Sybil Attack was first mentioned by Doucer in [31]. In this paper, Doucer argues that only a central authority can prevent a Sybil Attack under realistic assumptions of resource distribution and coordination.

While the intuitive solution to making a network robust against a Sybil Attack seems to be to make it expensive to create new identities in the network, doing so increases the social cost of the network by making it hard for new users to join it.

In the context of Reputation Systems in decentralised networks, a colluding group of malicious nodes could also increase the reputation of its nodes by itself and hence threaten the integrity of the network.

2) *Free riding*: In order to encourage successful long term cooperation, it is important that enough peers are providing sufficient resources for the system to become large and truly useful. In the absence of a third party monitoring each user, it is possible that some users stop contributing and only consume resources being generated by other users. *Free riders* are peers that eagerly consume resources without reciprocating any in return. It is easy to see how free riders diminish the quality of service for other peers, but more importantly, by making contributing peers feel exploited they disincentivise cooperation in the system and thus threaten the existence of the whole system, especially systems that are predicated on the foundation of sharing.

However, in the context of a decentralised network the most important problem created by free riding is that if only a few users are providing resources, they end up acting as centralised servers, this threatens the security of the network and defeats the very goal of the Web3 application.

Gnutella is a popular peer-to-peer file sharing platform which allows users private access to information. In their paper “Free Riding on Gnutella”, Eytan Adar and Bernardo A. Huberman [32] showed that 70% of Gnutella users were not sharing any files and nearly 50% of responses for file discovery were being returned by the top 1% of sharing hosts.

Similarly, Locher et al [33] were able to create “BitThief”, a free riding BitTorrent agent that was able to achieve high download rates even without seeding any data in return. They were also able to demonstrate that sharing communities which originally intended to promote cooperation among peers ultimately provide many incentives to cheat.

3) *Pollution Attack*: In 2005, Liang et al [34] showed that it was possible for an attacker in a decentralised network to corrupt certain targeted content, rendering it unusable and then making it available to the network in a large quantity. Since users on the network are unable to distinguish between the

polluted and the original content through content discovery alone, users download the polluted content and further share it with other peers, resulting in the polluted content spreading through the network.

In their analysis of the FastTrack peer to peer sharing system, it was found that as many as 50%-80% of copies of popular content were polluted.

4) *Index Poisoning*: Often resource sharing in decentralised networks is conducted through indices, which allow users to conveniently discover the location of their desired content. Depending on the architecture of the system, the index could be distributed over a fraction of the file sharing nodes (as in FastTrack) or over all the nodes.

In an Index Poisoning attack an attacker inserts bogus records into the index, for example, by inserting random identifiers that do not correspond to any address into the index. This way, when a user attempts to download a file they are unable to locate its content, leading to them finally abandoning the search. [35]

While the Pollution attack described earlier requires the attacker to obtain high-bandwidth to make sufficient versions of the corrupted copies available in the network, the Index Poisoning attack is easier in that it requires less resources to pull off.

5) *Slandering*: Under Sybil Attack, we discussed that it may be possible for a colluding group of malicious nodes to do *self promotion* to increment their own reputation in a reputation based decentralised system. On the other hand, it may also be possible for a group to coordinate to reduce the reputation of a victim, such an attack is called *slandering* [36]

6) *White Washing*: Nodes that have accrued a bad reputation by acting in an undesired manner can "clean" a bad reputation through *white washing* to avoid the negative effects of the disincentive system [36]

7) *Denial of service*: Cooperating nodes can work to block the functioning of a decentralised system, preventing other peers from utilizing its services

B. Social Threats

As seen by recent events, the rise of populist movements stands to be the biggest threat to the state of democracy worldwide. Many observers, especially journalists have suggested that the rise and spread of these movements has been massively aided through social media. [37] While social media can be a powerful tool for spreading information, when left unregulated, it can also lead to multiple social issues which greatly threaten long term cooperation. Some of these issues are:

1) *Echo Chambers and Polarisation*: "Echo Chambers" are used to describe the mechanism by which people on sociotechnical platforms are exposed to large or exclusively pro-attitudinal communication. Such grouping of like minded people on social networks ('homophily') is believed to arise from preferential connection to like minded individuals when creating/breaking bonds and also from peer influence which results in connected individuals growing more similar. [38]

The presence of an Echo Chamber could support populist messages that support rejection of expertise and reasoned debate among different views and lead to the emphasis of popularity of people or ideas over substance of their views. Therefore, Echo Chambers can lead to an insulation of users from the truth and even more perniciously, to be exposed to fake news.

In their study on Echo Chambers in the context of COVID-19 discussions on Twitter, Jiang et al [39] found strong evidence of political echo chambers on the topic on both ends of the political spectrum, but particularly so in the right-winged community. They found that tweets by right leaning users were almost exclusively retweeted by users who were also right leaning. Further, from random walk simulations, it was found that information in right leaning bubbles rarely travelled out of that bubble, forming a "small, yet intense political bubble". In another study on Climate Change discussions on Twitter, Williams et al "found a high degree of polarisation in attitudes, consistent with self selection bias" [38]

Studies have suggested that echo chambers could lead to polarisation of users and thus to users retreating into like-minded networks [40], which creates segmentation in networks and thus poses a large challenge to long term cooperation.

2) *Inequality and Social Divide*: While the idea of a digital democracy is appealing, it is hampered by findings of socio-economic inequality which prevent usage of the platforms by certain stratas of society. Beyond inability to access platforms, it is possible that members of society lack the skills to express their views or consume information that is being shared by other members. [41]

A lack of participation by different members of society could lead to the propagation of biased views or misinformation against the underrepresented members. Thus, it constitutes a credible threat to long term cooperation.

However, diffusion theories predict inequality at the outset of any innovation which is narrowed as time progresses and adoption rate spreads.

VI. REPUTATION SYSTEMS

Instead of only relying on direct reciprocation in decentralised systems, we can allow users that help each other out to establish a good reputation which can be used to reward them in some other way. After all, this is more representative of real social interactions, while we are interested in how people interact with us, we are also interested in the actions of others which we learn about from social channels such as gossip. In taking actions, we don't only take into account our direct experiences but also experiences we've learnt about from indirect sources. Similarly, when choosing to assist someone we also consider how it affects our reputation in society.

Although animals possess simple mechanisms for indirect reciprocity, only humans engage in complex reputation systems. [16] This seems to be because such systems require a substantive cognitive load, not only does it require a memory of all transactions but also requires the ability to monitor the dynamically changing social network of the group. Hence, the

strategies required to succeed in indirect reciprocity are also understandably a lot more complex than the simple Tit For Tat strategy that succeeds in direct reciprocity.

In their paper on reputation systems, Resnick et al [11] define a reputation system as one that “collects, distributed and feedback about participants’ past behavior ... these systems help people decide whom to trust, encourage trustworthy behavior, and deter participation by those who are unskilled or dishonest.”

As mentioned before, users on decentralised networks look for some form of assurance that their transactions on the network will be successful. The reputation of a user in reputation systems serves as a “shadow of the future” to each transaction, creating an expectation for what a user can expect when dealing with another user.

Consider the example of eBay’s reputation system, the “Feedback Forum”: after a transaction is completed, a buyer or seller can rate each other (1, 0 or -1) and leave comments. A participant in eBay accumulates such points over time which are displayed next to their screen name. A buyer can view a seller’s points and comments left by other users to create a “shadow of the future” into the transaction they can expect to have if they buy an item from the seller. Many other online forums and marketplaces such as Amazon and Stack Overflow rely on similar Reputation Systems.

According to Resnick, a reputation system must meet three challenges: [42]

- 1) Provide information that should allow users to distinguish between trustworthy and non trustworthy users,
- 2) Encourage users to be trustworthy, and
- 3) Discourage participation from users who aren’t

In addition to the above, a successful reputation system should also be able to avoid avoid issues mentioned in V

The following are a few notable Reputation Systems which attempt to accomplish the objectives stated above:

A. PageRank

One of the most widely known Reputation Systems in the World is Google’s PageRank. PageRank determines a rough estimate of the relative importance of a website by computing a ranking for every web page. The underlying assumption of PageRank is that a website that is more important is more likely to receive links from other websites than a website that is less important. PageRank is an interesting example of a Reputation Mechanism since while it may not be the exclusive algorithm used by Google, it has inspired many other reputation algorithms.

The calculation of PageRank of a website can be simplified to the below equation:

$$\sum \frac{\text{PageRank of Inbound Link}}{\text{Number of Outgoing Links on that Page}} \quad (1)$$

Hence, if a website a with a high PageRank has a link to another website b , website b will receive a large boost to its PageRank. However, the contribution of a ’s PageRank to b ’s PageRank will be reduced if a has a lot of outgoing links, this

is ensured by dividing the contribution of each inbound link by the number of outgoing links on that page.

Through this simple idea, Google was able to very successfully rank websites in terms of relevance. The idea was so revolutionary that PageRank is still used in Google today (along with 200 other more complex algorithms). [43] However, PageRank relies on a Trusted Oracle model which requires a centralised service, dependency on such an oracle to provide reputation introduces points of failure and does not scale well. Further, the original version of PageRank is susceptible to Sybil attacks. [44]

B. WikiTrust

WikiTrust [45] is the reputation system used for one of the largest collaborative applications known to mankind: the writing of articles on Wikipedia. It is a content-driven reputation system, that is, it relies on automated analysis of the content generated by the user and the collaboration process to derive the reputation of the user, rather than explicit feedback provided by users on other users. It is possible to use such a reputation system since the applications it’s catered for is entirely content driven.

The goals of *WikiTrust* are to incentivise lasting, meaningful contributions from users, help increase the quality of content being produced, spot vandals and to offer users an indicator of the quality of the content they are consuming. To achieve these goals, WikiTrust maintains different reputations for users and the content they create.

If a user makes a contribution that is meaningful and its content is preserved in future edits, they gain reputation, on the other hand, if their contributions are wholly or partially undone by future edits, then they lose reputation. Content starts with no reputation, if they are revised by users with high-reputation, it gains reputation. On the other hand, if the text is disturbed by too many edits, indicating that the content may not be trustworthy, it loses reputation.

In order to estimate how much each contribution is preserved or removed as required for the above, WikiTrust relies on an edit distance function $d(r, r')$ which is computed based on how many words have deleted, inserted, replaced and displaced from the edit that led from r to r' . Relying on such a distance functions allows the reputation system to be language independent. Finally, the value of an edit is calculated using the function:

$$q(b|a, c) = \frac{d(a, c) - d(b, c)}{d(a, b)} \quad (2)$$

Where b is the edit being evaluated, a is the revision before the edit and c is the revision after it. $q(b|a, c)$ outputs a value between -1 and +1; it is equal to -1 if $a = c$ and hence implying that b was entirely reverted, on the other hand, it is equal to +1 if the change from a to b was entirely preserved. However, a limitation of this approach is that since it requires subsequent revisions, it is unable to judge newly created revisions.

WikiTrust only considers not negative reputation values, new users are assigned a reputation very close to 0, this ensures that vandals cannot white wash themselves since their new identities would have a similar reputation to their vandal identity. Also, due to the content driven nature of the system, creating Sybils is harder than in a system where identities can simply be used to promote each other.

C. EigenTrust

While the reputation systems listed so far possess many interesting properties, both of them require a centralised “oracle” which acts as an intermediary for all nodes in the network, aggregating and providing trust values when a node requests them. Such an oracle is antithetical to the design of a decentralised application.

EigenTrust [46] could be described as a distributed version of PageRank. The algorithm allows the calculation of a unique *global trust value* \vec{t} for each peer i in the network which reflects the experience of all peers in the network with peer i . \vec{t} provides i a trust value for each peer in the network which it can refer to in order to establish how much it can trust another peer, ensuring it only conduct transactions with trustworthy peers. Further, it also has mechanisms to ensure that a malicious group of cooperating peers cannot lie for their own benefit.

Similar to eBay’s reputation system, the system requires each peer to rate another peer after it conducts a transaction with them. This results in the creation of a local trust value s for each peer where s_{ij} reflects how much i trusts peer j based on its transactions with them. It is suggested that one way of calculating s_{ij} is using:

$$s_{ij} = sat(i, j) - unsat(i, j) \quad (3)$$

Where $sat(i, j)$ and $unsat(i, j)$ represent the number of satisfactory and unsatisfactory transactions that i had with j respectively. These localised trust values s_{ij} are further normalised to produce c_{ij} to ensure that the trust values are between 0 and 1, to ensure that malicious peers can’t simply assign arbitrarily high local trust values to other malicious peers and low values to other peers, abusing the system.

The local trust values are then aggregated in each peer to produce t as below:

$$t_{ik} = \sum_j c_{ij} c_{jk} \quad (4)$$

This is equivalent to i asking its acquaintances how much they trust their acquaintances. However, so far t only reflects the experience of i and its acquaintances. This process needs to be repeated again in order to reflect i ’s acquaintances’ acquaintances experience and so on. The authors of the paper are able to prove that the final trust vector \vec{t}_i will converge to the same vector \vec{t} for every peer i in the network, namely the left principal eigenvector of the matrix $[c_{ij}]$

\vec{t} is calculated in a distributed manner. The authors prove that in a network with where the number of active peers are small, this can be done relatively efficiently since each peer

has limited transactions. In a network with a large number of peers, the algorithm can be performed efficiently by limited the local number of local trust values c_{ij} (t) at each peer can report. Further, a decay factor a can be used to reduced the influence of malicious cooperating peers ensuring that the same group of nodes vouching for each other is not as significant in the resulting trust values as a diverse group.

The main challenge facing the design of distributed reputation systems is the aggregation of local trust values into global trust values. EigenTrust’s translated the notion of *transitive trust* where if a peer trusts another peer it also trust its trusted network into a distributed algorithm which can run efficiently without congesting the network.

While EigenTrust is a powerful method, it requires the presence of a prior notion of trust i.e. a group of peers that are known to be trustworthy. The authors suggest that this could be the first few peers that join the network since the designers and early users of a P2P network are less likely to want to cheat in a network that they helped create. However, this assumption is a significant disadvantage of this Reputation Mechanism.

D. HonestPeer

EigenTrust’s assumption of a static group of trusted peers marginalises other peers, resulting in them being ranked much lower despite them potentially being honest. It also leads to potential poisoning vulnerabilities since if a trusted peer downloads a poisoned file from a malicious peer, it could result in the network also downloading the file. Finally, relying on a selected group of nodes comes with a lot of the same problems as relying on a single central entity.

HonestPeer [47] is an enhanced version of EigenTrust which tackles this problem by giving peers with a high reputation value a role in calculating the global reputation of other peers. Hence, instead of solely relying on a static group of peers, the algorithm selects a group of highly trusted peers dynamically, making it more robust and less centralised. While several improvements have been suggested to the original EigenTrust algorithm, HonestPeer is notable because it is able to reduce the algorithm’s dependency on pre-trusted peers without sacrificing the simplicity of the algorithm, an important requirement for effective cooperation strategies as seen in section II

The implementation follows the same approach of calculating trust values for each peer as EigenTrust. The trust values calculated in each run are used to find a honest peer h for each node, where:

$$t_h^k = \max_i (t_i^k) \quad (5)$$

This honest peer, h is then used in the calculating the proliferation parameter a , where:

$$a = \begin{cases} t_h^k & \text{if } t_h^k > 0.5 \\ 1 - t_h^k & \text{if } t_h^k \leq 0.5 \end{cases} \quad (6)$$

Based on this, the current reputation of peer i is then calculated as:

$$t_i^{k+1} = \begin{cases} a \times p_i + (1 - a) \times \sum_{x=1}^{x=n} c_{xi} t_x^k & \text{if } h \in P \\ (1 - a) \times p_i + a \times \sum_{x=1}^{x=n} c_{xi} t_x^k & \text{if } h \notin P \end{cases} \quad (7)$$

Where P is the group of pre-trusted peers. As a result of this modification, the influence of the pre-trusted peers is high if h is a part of them, otherwise their effect on the reputation is marginalised. Through simulation in a p2p file sharing network, the paper's authors were able to demonstrate that HonestPeer reduced the percentage of invalid files and increased the success rate of good files downloaded. Further, by allowing h to be dynamically replaceable, the algorithm ends up more scalable too, which was also demonstrated in simulation.

E. PeerTrust

Like all the Reputation Systems mentioned above, PeerTrust [48] evaluates a node's trustworthiness by taking in consideration the feedback a peer has obtained from other peers. However, in addition to simply considering the feedback, PeerTrust also takes into account certain other factors:

1) Credibility of Feedback

It may be possible for a peer to lie in their feedback due to malicious motives. Therefore, the credibility of a node is taken into account when deciding how much to value their feedback of another node.

2) Transaction Context Factor

Not all transactions are equal, for example, in an application that involves transactions between users, transactions of greater economic value should influence a user's trust more than transactions of small value since otherwise a user could behave in a trustworthy manner in a lot of transactions of small value but cheat in one transaction of large value and still end up with a positive reputation. Beyond this, just because a node can provide good services in a certain context, this does not necessarily imply that they can provide comparable services in a completely different context, for example, a node that provides good information about tourism should not also automatically be trusted to provide equally good medical advice. Therefore, the context of the transaction is made a factor when calculating trust.

3) Community Context Factor

In order to deal with community specific issues and vulnerabilities such as lack of incentive to provide feedback and collaboration of malicious peers to manipulate their trust scores, community contexts are added as a factor when calculating trust.

Hence, the trust value $T(u)$ of a node u is defined as:

$$T(u) = \alpha \times \sum_{i=1}^{I(u)} S(u, i) \times Cr(p(u, i)) \times TF(u, i) + \beta \times CF(u) \quad (8)$$

Where,

- $I(u)$ denotes the total number of transactions of u with all other nodes in a recent time window
- $p(u, i)$ denotes the other node that participated in peer u 's i th transaction
- $S(u, i)$ is the normalised amount of satisfaction that u received from node $p(u, i)$ in the i th transaction
- $Cr(v)$ is the credibility of the feedback submitted by v
- $TF(u, i)$ is the adaptive transaction context factor for node u 's i th transaction
- $CF(u)$ is the adaptive community context factor for node u

F. MeritRank

MeritRank [49] uses a merit based tokenomics model which aims to bound the benefits of Sybil attacks instead of preventing them altogether. The system is based on the assumption that peers observe and evaluate each others' contribution, similar to the reputation system used in eBay. Each peer's evaluation is stored in a personal ledger and modelled in a feedback graph where the feedback to each user is modelled as a special token value which accumulates over time. It is also assumed that each peer is able to discover the feedback graph, for example, through a gossip protocol. MeritRank manages to achieve this Sybil tolerance by imposing the following constraints on how reputation can be gained inside the feedback graph:

1) Relative Feedback

This constraint places a bound on how much feedback a single entity can provide to another entity by the degree of the entity i.e. the size of the set of its neighbours. This constraint assists in limiting a single entity from creating multiple parallel Sybils

2) Transitivity α decay

This constraint limits the ability of an entity to create a serial Sybil attack by terminating random walks in the feedback graph with a probability α

3) Connectivity β decay

Sybil attack edges in a feedback graph are often bridges i.e. their cut creates two separate components. This constraint introduces a punishment for a node for being in a separate component

A trust graph modelled using these MeritRank's constraints will satisfy:

$$\lim_{|S| \rightarrow \infty} \frac{w^+(\sigma_s)}{w^-(\sigma_s)} \leq c \quad (9)$$

where, $w^+(\sigma_s)$ is the profit gained by the Sybil Attack σ_s , $w^-(\sigma_s)$ is the cost of the Sybil attack, S is the set of Sybils and c is some constant value such that $c > 0$. Thus MeritRank is able to provide a Reputation Mechanism with feedback which is Sybil tolerant.

Table I summarises the Reputation Systems discussed so far. The survey of Reputation Systems above is not an exhaustive list, for a more comprehensive treatment of the subject the

reader is referred to the following surveys: [50]–[53]. However, the Reputation Systems stated above provide a decent summary of how indirect reciprocity can be used to create trust in a decentralised network while also tackling a lot of trust issues inherent to the decentralised setting.

While indirect reciprocity through Reputation Systems is a powerful tool for tackling the infrastructural threats in V such as Sybil attacks, whitewashing etc., in of themselves, Reputation Systems cannot serve as an “Universal Trust Machine” due to the following limitations:

1) **Not Privacy Preserving**

While a lot of Reputation Systems presented so far may be confidentiality preserving, they are not privacy preserving, i.e. they do not prevent the discovery of users who contributed to a reputation rating. For example, if a node goes offline between two reputation queries, the difference in the aggregated reputation score across the two queries can reveal the user’s contribution. [54]

2) **Do not provide a mechanism to carry out transactions**

While Reputation Systems are great for applications such as P2P file exchange, enabling cooperation in other domains such as e-commerce and IoT requires mechanisms beyond simple provision of trust such as the ability to carry out transactions.

3) **Lack of Flexibility**

Guaranteeing trust in the real world is often not as simple as simply identifying trustable users, centralised institutions provide bespoke functionality such as financial contracts and escrows in order to enable cooperation between users who want to perform transactions with each other. In section VII, we show how Smart contracts can allow for such functionality in the distributed world.

4) **Requires all entities to remain online**

Most of the decentralised solutions suggested above suffer from the problem of requiring trusted entities to always be online and connectable in order to enable network discovery and furthermore, have a valid address where they can be contacted. This is too large an assumption in a lot of domains such as IoT where nodes are constantly going offline.

5) **Do not solve social threats**

It is important to note that Reputation Systems only tackle infrastructural threats. Due to their limited scope, they are unable to tackle the social threats listed in section V.

In the next section, we present Distributed Ledger Technology (DLT), a technology that offers a solution to a lot of the limitations listed above.

VII. DISTRIBUTED LEDGER TECHNOLOGY

Any technology that leverages a ledger in order to store data distributed across multiple nodes is referred to as a DLT. The recent rise in popularity of Bitcoin has led to prominence of Blockchain technology and multiple other technologies that leverage Blockchain such as Ethereum, Hyperledger, Cardano etc. However, Blockchain is not the only DLT, some other

examples of DLTs include: Tangle, Hashgroup and Sidechain. [55] presents an extensive comparison of these technologies.

At their core, DLTs are data structures that allow recording of transactions and functions for their manipulation. Generally, all DLTs are based on three well-known foundational technologies: [55]

1) **Public Key Cryptography**

Since DLTs operate in insecure, distributed environments, Public Key Cryptography allows for the establishment of secure digital identities and communications between nodes. The digital identities also allow for the enforcement of ownership of resources in the network and thus helps facilitate transactions.

2) **Distributed Peer to Peer Network**

Operating in a distributed network allows for a highly scalable network without a single point of failure (as in a centralised network).

3) **Consensus Mechanism**

The presence of such a mechanism allows for all nodes in the network to converge on a single version of global truth without a trusted intermediary.

DLTs are similar to Reputation Systems in that often, their main goal is to allow interactions between users that do not trust each other without a trusted third party. [56] By design, DLTs offer a high degree of transparency, traceability and security which allows them to offer security, privacy and trustworthiness inside a diverse set of applications. Below is a survey of four DLT technologies which helps to demonstrate the benefits provided by them:

A. Bitcoin

Bitcoin [21] is a cryptocurrency that leverages Blockchain technology to offer a distributed and immutable ledger that stores transaction history. The core technology consists of a linked list of blocks that are connected together, with each block referencing the previous block in the chain. Transactions are continuously being appended in blocks to the chain and are visible to all participants in the network.

Bitcoin uses Proof of Work (PoW), a form of cryptographic proof as a consensus mechanism which allows a party to prove to other nodes in the network that a specific amount of computational power was expended. The nodes in the network are able to verify this expenditure with minimal effort. The purpose of using PoW as consensus mechanism is to deter manipulation of data in the ledger by imposing an infeasible energy and hardware requirement in order to do so, thus guaranteeing security of the ledger and allowing nodes in the distributed network to conduct transactions with each other.

However, multiple papers have criticised the inefficiency of PoW and Bitcoin [57], [58], notably the requirement of wasting enormous amount of electricity through expensive mining equipment. Further, transactions on the network require large transaction fees which have to be paid to miners and a large confirmation time for transactions, making Bitcoin a poor choice for applications that require a large amount of transactions, quickly.

Year	Reputation Mechanism	Trust Function	Decentralised	Sybil Resistant
1999	PageRank	Trustworthiness of a node is determined by counting the number and quality of links to it.	×	×
2003	EigenTrust	Distributed PageRank, allows for calculation of a global trust value which reflects the experience of all peers in the network.	✓	×
2004	PeerTrust	To determine trustworthiness, in addition to the feedback from other peers, also takes into consideration other factors such as credibility of feedback, the transaction and community context.	✓	×
2012	WikiTrust	Trustworthiness of user is calculated by measuring meaningfulness of their contribution This meaningfulness is calculating using an edit distance function.	×	×
2015	HonestPeer	Enhanced EigenTrust, reduces dependency of original algorithm on a set of pre-trusted peers by selecting a group of highly selected peers dynamically	✓	×
2022	MeritRank	Provides a set of decay based strategies which help to provide Sybil resistance to otherwise Sybil prone strategies	✓	✓

TABLE I
OVERVIEW OF MENTIONED REPUTATION SYSTEMS

B. Ethereum

Similar to Blockchain, Ethereum [59] is based on Blockchain technology, however, as of 2022, Ethereum differs from Bitcoin in that it uses Proof of Stake (PoS) as a consensus mechanism instead of PoW. In PoS, the node with the highest stake and not the highest computing power obtains the right to book-keeping, where the stake is a reflection of a node's ownership of a specific amount of currency. [60] Therefore, it solves the problem of wasted computing power present in PoW to a certain extent and can reduce the time required to reach a consensus. However, a drawback of a stake based consensus mechanism is that it can result in centralisation in extreme cases.

Ethereum is also different from Bitcoin in that it is a programmable blockchain platform; using Smart Contracts, users on the Ethereum Platform can not only perform simple transactions, but can also create complex transactions.

The term *Smart Contract* was first termed by N. Szabo who defined it as a "computerized transaction protocol that executes the terms of a contract" [61]. In Ethereum, a smart contract represents a deterministic, Turing-complete program which consists of a collection of code (functions) and data (state) which are deployed to the Ethereum network and run as programmed. Smart Contracts allow users on the network to define complex rules for facilitating interactions. In Ethereum, Smart Contracts are programmed using the Solidity programming language.

However, it is worth noting that Smart Contracts in Ethereum suffer from certain limitations:

1) Vulnerability Prone

Smart Contract code is prone to vulnerabilities which can be costly when exploited. In a famous example, \$50 million was stolen from a DAO in an attack that exploited a concurrency-based vulnerability. A trivial error is bound to be exploited by hackers, though this limitation is true for all open source code in general.

2) Expensive

Code in a Smart Contract running on the Ethereum network is significantly costlier per instruction than the same code running on a typical cloud server.

Further, even though Ethereum uses PoS, it still suffers from high transaction fees similar to Bitcoin.

C. IOTA

IOTA is a cryptocurrency meant for the IoT industry. IOTA is different from the technologies listed so far in that it does not use Blockchain as its underlying ledger and instead uses the *Tangle* [62], a directed acyclic graph (DAG) for storing transactions. Tangle is both a decentralised data storage architecture and a consensus protocol where each node in its DAG represent a transaction and the connections between nodes represent validations of the transactions.

Unlike Bitcoin and Ethereum, IOTA removes the dichotomy between transaction miner and validator, instead requiring users who are adding transactions to the network to validate other transactions. Thus, in practice, in order to add a transaction to the network, a user needs to choose and validate two other transactions, using Hashcash as the PoW algorithm with a lowered difficulty. The rationale behind using PoW is to prevent spam transactions. Through this mechanism, IOTA is able to reduce or get rid of transaction costs and allow a much higher transaction per second count than Bitcoin and Ethereum. Further, since it relies on network users to validate transactions, as the number of users in the network increases, the validation time of submitted transactions also decreases, rendering it much more scalable than both Bitcoin and Ethereum.

However, as of the time of writing this paper, IOTA relies on a centralised coordinator to assist in adding transactions to the network and further, it does not have support for smart contracts. (Though this is planned to change in a future release). Further, since the technology is still relatively new, it is still experimental and may contain security threats. [63]

D. TrustChain

A large drawback of the solutions listed so far is that none of them tackle the infrastructural threats presented in section V, most notably, they are not Sybil resistant, hence, while they are useful in isolation for facilitating transactions, they cannot be used for the use case of creating trust. TrustChain [64] is unique in this sense, since it is a Sybil-resistant, scalable blockchain. TrustChain relies on *NetFlow* a Reputation Mechanism, to calculate reputation using an interaction graph and the max-flow algorithm, thus allowing it to be Sybil resistant.

In TrustChain, each node is responsible for storing their own Blockchain and hence, transaction history. In addition to containing a single transaction's data, each block in the Blockchain also contains two references, one to the last block in its own Blockchain and another to the last block in the transacting party's Blockchain. While this mechanism doesn't allow the prevention of double spending, it allows detection of it since all blocks must have two outgoing and two incoming links. Further, the additional pointer to the counterparty's chain makes it hard to remove or alter blocks in one's chain, resulting in the Blockchains being tamper proof. TrustChain blocks are also exchanged between nodes using gossiping mechanisms and hence, replicated network wide, allowing the network to be resilient against nodes going offline.

TrustChain's local storage also allows it to be highly scalable since there is no global consensus mechanism, also removing the need for transaction fees and leading to a large amount of transaction per second.

Though the lack of a global Blockchain and hence, lack of a currency earned by validators in the network could also be seen as a drawback since a large reason for growth behind popular cryptocurrencies is the investment of the users in it since they possess a vested interest in it. Further, unlike Ethereum, TrustChain does not possess a mechanism for creating complex contracts and hence, lacks flexibility in trust generation.

Table II provides an overview of the DLTs covered in this section. [54] provides an extensive survey of how DLTs, like Blockchain are being used in Distributed Trust and Reputation Management systems.

VIII. OTHER MECHANISMS

Besides direct reciprocity and indirect reciprocity, there are also other mechanisms that should be considered when understanding how cooperation could evolve in a decentralised network.

A. Network Reciprocity

While the analysis so far relies on a well-mixed population, in reality the spatial structures of social connections are not well mixed, instead certain groups interact with each other more often than others. In such a setting, it may be able to form network cluster of cooperators who help each other out resulting in a "Network Reciprocity" which is a generalisation of "Spatial Reciprocity". [16]

In their paper "The WebEngine - A Fully Integrated, Decentralised Web Search Engine", Mario M. Kubek and Herwik Unger [65] suggest an idea of constructing "content overlay networks". This involves creating social graphs with nearby and distant neighbours, where nearby neighbours are neighbours that share similar content.

B. Machine Learning based Reputation Systems

[66] suggests a Reputation System that utilizes SVMs in order to establish trustworthiness of nodes in a decentralised network and demonstrates its effectiveness

IX. FUTURE ROADMAP

Reputation Systems are a powerful mechanism for preventing infrastructural threats to long term cooperation such as Sybil Attacks, however, in isolation they possess multiple limitations that prevent them from providing a decentralised trusted ecosystem. Distributed Ledger Technologies are a strong complement to Reputation Mechanism as seen in TrustChain.

However, none of the solutions in the existing literature for generating trust attempt to tackle Social Issues inherent in decentralised systems. While in a centralised system it is possible to solve these issues in an ad-hoc manner, decentralised systems require all rules of the system to be explicitly stated upfront and hence a "Universal Trust Machine" would have to tackle these issues in order to create an ecosystem that can rival a centralised system.

X. CONCLUSION

This survey presents the progress of Web3 on its road towards creating a "Universal Trust Machine", a hypothetical ecosystem where decentralised trust is generated without the aid of any trusted third parties. In order to do so, we first motivate the problem of creating cooperation in a sea of adversaries by discussing Robert Axelrod's research on the Evolution of Cooperation. Next, we clarify terminology by presenting context on the concepts of *Decentralisation* and *Web3*. We then present issues that such a machine would have to tackle in order to foster long term cooperation. Finally, we present contemporary technologies such as Reputation Systems and Distributed Ledger Technologies which in conjunction could be used to construct the desired machine.

REFERENCES

- [1] L. Mui, "Computational models of trust and reputation: Agents, evolutionary games, and social networks," 05 2014.
- [2] L. Christov-Moore, D. Bolis, J. Kaplan, L. Schilbach, and M. Iacoboni, "Trust in social interaction: From dyads to civilizations," Jun 2022. [Online]. Available: psyarxiv.com/urav8
- [3] M. T. F. Zanini and C. P. Migueles, "Trust as an element of informal coordination and its relationship with organizational performance," *Economia*, vol. 14, no. 2, pp. 77–87, 2013. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1517758013000064>
- [4] K. Newton, "Trust, social capital, civil society, and democracy," *International Political Science Review*, vol. 22, no. 2, pp. 201–214, 2001. [Online]. Available: <https://doi.org/10.1177/0192512101222004>
- [5] K. S. Cook, M. Levi, and R. Hardin, *Whom can we trust?: How groups, networks, and institutions make trust possible*. Russell Sage Foundation, 2009.
- [6] T. Earle, "Trust, confidence, and the 2008 global financial crisis," *Risk analysis : an official publication of the Society for Risk Analysis*, vol. 29, pp. 785–92, 05 2009.
- [7] R. Hardin, *Trust*, ser. Key Concepts. Wiley, 2006. [Online]. Available: <https://books.google.nl/books?id=jB2SBK5izRcC>
- [8] G. Korpil and D. Scott, "Decentralization and web3 technologies," 5 2022. [Online]. Available: https://www.techrxiv.org/articles/preprint/Decentralization_and_web3_technologies/19727734
- [9] P. De Filippi, M. Mannan, and W. Reijers, "Blockchain as a confidence machine: The problem of trust and challenges of governance," *Technology in Society*, vol. 62, p. 101284, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0160791X20303067>
- [10] J. Hofmokr, "The internet commons: towards an eclectic theoretical framework," *International Journal of the Commons*, vol. 4, no. 1, pp. 226–250, 2010.

Year	DLT	Ledger	Low Fees	High Scalability	Smart Contracts	Sybil Resistance
2008	Bitcoin	Global Blockchain	×	×	×	×
2014	Ethereum	Global Blockchain	×	×	✓	×
2018	IOTA	Tangle	✓	✓	×	×
2018	TrustChain	Locally Stored Linked Blockchains	✓	✓	×	✓

TABLE II
OVERVIEW OF MENTIONED DLTs

- [11] P. Resnick, K. Kuwabara, R. Zeckhauser, and E. Friedman, "Reputation systems," *Commun. ACM*, vol. 43, no. 12, p. 45–48, dec 2000. [Online]. Available: <https://doi-org.tudelft.idm.oclc.org/10.1145/355112.355122>
- [12] R. Axelrod and W. D. Hamilton, "The evolution of cooperation," *science*, vol. 211, no. 4489, pp. 1390–1396, 1981.
- [13] R. Axelrod, *The Evolution of Cooperation*, 1984.
- [14] J. Barker, "Robert axelrod's (1984) the evolution of cooperation," 01 2017.
- [15] R. Hoffmann, "Twenty years on: The evolution of cooperation revisited." *J. Artificial Societies and Social Simulation*, vol. 3, 03 2000.
- [16] M. A. Nowak, "Five rules for the evolution of cooperation." *Science*, vol. 314, no. 5805, pp. 1560–1563, 2006. [Online]. Available: <https://www.science.org/doi/abs/10.1126/science.1133755>
- [17] L. Cao, "Decentralized ai: Edge intelligence and smart blockchain, metaverse, web3, and descI," *IEEE Intelligent Systems*, vol. 37, no. 3, pp. 6–19, 2022.
- [18] J. Lerner and J. Tirole, "The economics of technology sharing: Open source and beyond," *Journal of Economic Perspectives*, vol. 19, no. 2, pp. 99–120, June 2005. [Online]. Available: <https://www.aeaweb.org/articles?id=10.1257/0895330054048678>
- [19] P. Baran, "On distributed communications networks," *IEEE Transactions on Communications Systems*, vol. 12, no. 1, pp. 1–9, 1964.
- [20] J. Pouwelse, P. Garbacki, J. Wang, A. Bakker, J. Yang, A. Iosup, D. Epema, M. Reinders, M. van Steen, and H. Sips, "Tribler: a social-based peer-to-peer system," *Concurrency and Computation: Practice and Experience*, vol. 20, pp. 127 – 138, 02 2008.
- [21] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," *Cryptography Mailing list at https://metzdowd.com*, 03 2009.
- [22] [Online]. Available: <https://www.starlinglab.org/78days/>
- [23] T. O'Reilly, "What is web 2.0: Design patterns and business models for the next generation of software," *University Library of Munich, Germany, MPRA Paper*, vol. 65, 01 2007.
- [24] "Home · solid." [Online]. Available: <https://solidproject.org/>
- [25] G. Wood. [Online]. Available: <http://gavwood.com/dappsweb3.html>
- [26] Z. Liu, Y. Xiang, J. Shi, P. Gao, H. Wang, X. Xiao, B. Wen, Q. Li, and Y.-C. Hu, "Make web3.0 connected," *IEEE Transactions on Dependable and Secure Computing*, vol. 19, no. 5, pp. 2965–2981, 2022.
- [27] J. M. Garon, "Legal implications of a ubiquitous metaverse and a web3 future," *SSRN Electron. J.*, 2022.
- [28] A. Park, M. Wilson, K. Robson, D. Demetis, and J. Kietzmann, "Interoperability: Our exciting and terrifying web3 future," *Business Horizons*, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0007681322001318>
- [29] E. Bellini, Y. Iraqi, and E. Damiani, "Blockchain-based distributed trust and reputation management systems: A survey," *IEEE Access*, vol. 8, pp. 21 127–21 151, 2020.
- [30] J. Dinger and H. Hartenstein, "Defending the sybil attack in p2p networks: taxonomy, challenges, and a proposal for self-registration," in *First International Conference on Availability, Reliability and Security (ARES'06)*, 2006, pp. 8 pp.–763.
- [31] J. R. Douceur, "The sybil attack," in *Revised Papers from the First International Workshop on Peer-to-Peer Systems*, ser. IPTPS '01. Berlin, Heidelberg: Springer-Verlag, 2002, p. 251–260.
- [32] E. Adar and B. A. Huberman, "Free riding on gnutella," *First Monday*, vol. 5, no. 10, Oct. 2000. [Online]. Available: <https://firstmonday.org/ojs/index.php/fm/article/view/792>
- [33] T. Locher, P. Moore, S. Schmid, and R. Wattenhofer, "Free riding in bittorrent is cheap," in *5th Workshop on Hot Topics in Networks (HotNets)*, 2006. [Online]. Available: <http://eprints.cs.univie.ac.at/5696/>
- [34] J. Liang, R. Kumar, Y. Xi, and K. Ross, "Pollution in p2p file sharing systems," in *Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies.*, vol. 2, 2005, pp. 1174–1185 vol. 2.
- [35] J. Liang, N. Naoumov, and K. W. Ross, "The index poisoning attack in p2p file sharing systems," in *Proceedings IEEE INFOCOM 2006. 25TH IEEE International Conference on Computer Communications*, 2006, pp. 1–12.
- [36] V. Agate, A. De Paola, G. Lo Re, and M. Morana, "A simulation framework for evaluating distributed reputation management systems," in *Distributed Computing and Artificial Intelligence, 13th International Conference*, S. Omatu, A. Semalat, G. Bocewicz, P. Sitek, I. E. Nielsen, J. A. García García, and J. Bajo, Eds. Cham: Springer International Publishing, 2016, pp. 247–254.
- [37] J. Bartlett, J. Birdwell, and M. Littler, *The new face of digital populism*. Demos, 2011.
- [38] H. T. Williams, J. R. McMurray, T. Kurz, and F. Hugo Lambert, "Network analysis reveals open forums and echo chambers in social media discussions of climate change," *Global Environmental Change*, vol. 32, pp. 126–138, 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0959378015000369>
- [39] J. Jiang, X. Ren, and E. Ferrara, "Social media polarization and echo chambers in the context of covid-19: Case study," *JMIRx Med*, vol. 2, no. 3, p. e29570, Aug 2021. [Online]. Available: <https://med.jmirx.org/2021/3/e29570>
- [40] P. Norris and R. Inglehart, *Cultural backlash: Trump, Brexit, and authoritarian populism*. Cambridge University Press, 2019.
- [41] J. Schradie, "The trend of class, race, and ethnicity in social media inequality," *Information, Communication & Society*, vol. 15, no. 4, pp. 555–571, 2012. [Online]. Available: <https://doi.org/10.1080/1369118X.2012.665939>
- [42] P. Resnick, "Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system," *Advances in Applied Microeconomics*, vol. 11, 10 2002.
- [43] L. Page, S. Brin, R. Motwani, and T. Winograd, "The PageRank Citation Ranking: Bringing Order to the Web," Stanford Digital Library Technologies Project, Tech. Rep., 1998. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.1.31.1768>
- [44] Cheng, Alice, Friedman, and Eric, "Manipulability of pagerank under sybil strategies," 01 2006.
- [45] L. De Alfaro, A. Kulshreshtha, I. Pye, and B. T. Adler, "Reputation systems for open collaboration," *Commun. ACM*, vol. 54, no. 8, p. 81–87, aug 2011. [Online]. Available: <https://doi-org.tudelft.idm.oclc.org/10.1145/1978542.1978560>
- [46] S. Kamvar, M. Schlosser, and H. Garcia-molina, "The eigentrust algorithm for reputation management in p2p networks," *The EigenTrust Algorithm for Reputation Management in P2P Networks*, 04 2003.
- [47] H. A. Kurdi, "Honestpeer: An enhanced eigentrust algorithm for reputation management in p2p systems," *Journal of King Saud University - Computer and Information Sciences*, vol. 27, no. 3, pp. 315–322, 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1319157815000440>
- [48] L. Xiong and L. Liu, "Peertrust: supporting reputation-based trust for peer-to-peer electronic communities," *IEEE Transactions on Knowledge and Data Engineering*, vol. 16, no. 7, pp. 843–857, 2004.
- [49] B. Nasrulin, G. Ishmaev, and J. Pouwelse, "Meritrank: Sybil tolerant reputation for merit-based tokenomics," 2022. [Online]. Available: <https://arxiv.org/abs/2207.09950>
- [50] A. Jøsang, R. Ismail, and C. Boyd, "A survey of trust and reputation systems for online service provision," *Decision Support Systems*, vol. 43, no. 2, pp. 618–644, 2007, emerging Issues in Collaborative Commerce. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167923605000849>
- [51] S. Ruohomaa, L. Kutvonen, and E. Koutrouli, "Reputation management survey," *Ruohomaa , S , Kutvonen , L Koutrouli , E 2007 , ' Reputation management survey ' , pp. 103-111 ., 01 2007.*

- [52] H. Yu, Z. Shen, C. Leung, C. Miao, and V. R. Lesser, "A survey of multi-agent trust management systems," *IEEE Access*, vol. 1, pp. 35–50, 2013.
- [53] L. Mui, M. Mohtashemi, and A. Halberstadt, "Notions of reputation in multi-agents systems: A review," 01 2002, pp. 280–287.
- [54] E. Bellini, Y. Iraqi, and E. Damiani, "Blockchain-based distributed trust and reputation management systems: A survey," *IEEE Access*, vol. 8, pp. 21 127–21 151, 2020.
- [55] N. El Ioini and C. Pahl, *A Review of Distributed Ledger Technologies: Confederated International Conferences: CoopIS, CTC, and ODBASE 2018, Valletta, Malta, October 22-26, 2018, Proceedings, Part II*, 10 2018, pp. 277–288.
- [56] M. Walport, "Distributed ledger technology: Beyond blockchain. uk government office for science," *Tech. Rep. Tech. Rep.*, 2016.
- [57] A. Urquhart, "The inefficiency of bitcoin," *Economics Letters*, vol. 148, pp. 80–82, 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0165176516303640>
- [58] T. Xue, Y. Yuan, Z. Ahmed, K. Moniz, G. Cao, and C. Wang, "Proof of contribution: A modification of proof of work to increase mining efficiency," in *2018 IEEE 42nd Annual Computer Software and Applications Conference (COMPSAC)*, vol. 01, 2018, pp. 636–644.
- [59] W. Ethereum, "Ethereum whitepaper," *Ethereum*. URL: <https://ethereum.org> [accessed 2020-07-07], 2014.
- [60] W. Li and M. He, "Comparative analysis of bitcoin, ethereum, and libra," in *2020 IEEE 11th International Conference on Software Engineering and Service Science (ICSESS)*, 2020, pp. 545–550.
- [61] N. Szabo, "Formalizing and securing relationships on public networks," *First monday*, 1997.
- [62] S. Popov, "The tangle," *White paper*, vol. 1, no. 3, p. 30, 2018.
- [63] M. Conti, G. Kumar, P. Nerurkar, R. Saha, and L. Vigneri, "A survey on security challenges and solutions in the iota," *Journal of Network and Computer Applications*, vol. 203, p. 103383, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1084804522000467>
- [64] P. Otte, M. de Vos, and J. Pouwelse, "Trustchain: A sybil-resistant scalable blockchain," *Future Generation Computer Systems*, vol. 107, pp. 770–780, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167739X17318988>
- [65] M. M. Kubek and H. Unger, "The webengine: A fully integrated, decentralised web search engine," in *Proceedings of the 2nd International Conference on Natural Language Processing and Information Retrieval*, ser. NLPPIR 2018. New York, NY, USA: Association for Computing Machinery, 2018, p. 26–31. [Online]. Available: <https://doi-org.tudelft.idm.oclc.org/10.1145/3278293.3278294>
- [66] R. Akbani, T. Korkmaz, and G. Raju, "Emltrust: An enhanced machine learning based reputation system for manets," *Ad Hoc Networks*, vol. 10, no. 3, pp. 435–457, 2012. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1570870511001867>