Rep on the block : A next generation reputation system based on the blockchain

Richard Dennis
School of Computing
University of Portsmouth
Portsmouth, United Kingdom
Richard.dennis@port.ac.uk

Gareth Owen
School of Computing
University of Portsmouth
Portsmouth, United Kingdom
Gareth.owen@port.ac.uk

***Abstract*— This paper presents the first generalized reputation system that can be applied to multiple networks that is based on the blockchain. We first discuss current reputation systems, conducting a critical analysis of their current security vulnerabilities, before looking at how new blockchain based technologies are used. We propose an innovative new reputation system that is based on blockchain technologies which aims to solve many unanswered questions in today's current generation reputation systems. We then consider the limitations of such a system, before using simulations and analyses to demonstrate methods of overcoming these limitations. We conclude by suggesting areas for future studies, and summarizing our findings.**

Blockchain, reputation systems, cryptographic protocols, distrubted networks, peer-to-peer, Bittorent

I. INTRODUCTION

Reputation measures how much the community trusts you, and is calculated on your previous transactions and interactions with the community. The greater your reputation, the more trustworthy you are seen to be on the network and, with a user's reputation on the line, users choose to behave more honestly on the network [1].

At present, eBay has the most widely used reputation system and processes over a billion transactions per day [2]. Each transaction could result in two reputation scores being left (one from the buyer, the other from the seller); it is therefore essential that reputation systems can handle a large number of transactions, and have adequate sources to handle this level of data.

E-commerce reputation systems often implement multi-dimensional reputations, which allow the user to rate the seller on a range of factors such as postage cost and quality of communications. All major E-commerce websites use the traditional client-server model, where the reputation data is centrally stored, calculated and distributed on a centralized server and all clients can request to see this data from the central server.

In eBay's system, the positive feedback percentage is calculated based on the total number of positive and negative feedback ratings for transactions in the last 12 months, excluding repeat feedback from the same member for purchases made within the same calendar week [3].

The reputation score is calculated centrally by the E-commerce website, which has the negative effect of the company being able to change the reputation calculation algorithm and force the deployment of this to all users without their knowledge. For example, eBay recently prevented sellers from leaving negative feedback about buyers.

Although successful reputation systems have been implemented on multiple web services, they are all based on the centralised server model which makes them unsuitable for deployment in a Peer-to-Peer (P2P) networks, where the principle purpose is decentralisation of control away from a single entitle. Thus far, the effective communication and sharing of unmodified information relating to trust and reputation remains an unsolved issue [4].

There are several reputation systems implemented in peer-to-peer networks, which aim to provide users of the network with an incentive to behave honestly and to deter "freeloaders" for example, on the Gnutella network - the most popular P2P network - an estimated 70% of all peers behave in such a way [5]. Freeloaders are peers who download content from the network, but who do not distribute any content.

There are various implementations of reputation systems on peer-to-peer networks; some require the implementation of a trusted central server, much like the E-commerce model, which records and calculates all users' ratings, whilst other systems try to distribute the reputation system with a distributed database that all peers on the network have an updated copy of. The final implementation of reputation systems on a P2P network only records reputation of peers it has interacted with.

Unlike E-commerce reputation systems where participation is mandatory, enrolment in a P2P reputation system is optional and many nodes are concerned about the loss of privacy or the additional resources that are required.

P2P reputation systems are single-dimensional systems, with each peer only leaving one bit of data about the transaction that has taken place; this enhances efficiency and also reduces load on the network.

The calculation of reputation differs from implementation to implementation, however the general calculation method for each peer is that their reputation is the sum of all reputation feedback received.

All reputation systems, no matter how they are deployed or what type of network they are deployed over, face the same fundamental issues. The ability to link an identity to a single user, and to prevent that user from obtaining more than one identity, is key to preventing a user exploiting the system by creating multiple identities and transacting between them.

Another limitation that is central to all reputation systems and which remains an open question is how to quantify reputation? Furthermore, how can we ensure the reputation left by a user is accurate and is based on a real transaction?

This paper is organised as follows: section two describes related work in this area, focusing on reputation systems implemented in peer-to-peer systems. Section three discusses our proposed reputation system along with some of the technologies used in it, whilst section four summarises our approach with simulation and comparison of our network compared to currently implemented reputation systems, before concluding with suggestions for future work and summarising the contribution of this paper.

II. Related Work

A. Existing "decentralized" reputation systems

Peer-to-peer based reputation systems have been around almost as long as peer-to-peer networks themselves, with the first system mentioned in literature in 2003 [6]. Reputation systems on peer-to-peer networks all have different goals; from choosing reliable resources, ensuring peers behave honestly, and rating the quality of content of a shared file.

Reputation systems in peer-to-peer networks have to contend with the known issues of reputation systems in general, with the additional complexity a peer-to-peer network adds. Additional issues, such as how to keep data up to date, accurate and distributed to a large set of peers which changes dynamically are faced when deploying a reputation system over a P2P network [7].

Wang proposes a reputation system based on the Bayesian model, which aims to rate file shares based on the quality of the file they are sharing, as well as the trust in each peer [8]. This is a novel system, separating trust and reputation, whereas previous systems tend to combine these into a single rating [9]. The separation of these ratings allows for all users to gain a trust rating by acting honestly on the network. However, one key issue with Wang's model is that they assume all users are honest in their ratings of each peer, which is an unlikely occurrence in the real world. The system does not try to provide each user with a global view of the network, instead reputations are collected by each peer based on previous transactions. The reputation and trust is based on a binary system, and for each successful transaction a reputation score of 1 is given. The trust and reputation scores are the sum of all scores. Several issues with this type of reputation exist, the assumption that all scores are genuine for a transaction that actually took place and that there are no malicious actors trying to profit from the system is perhaps the most major one. The reputation scores and calculation done by each peer is however a good workaround to the known issue of distributing data to all nodes in the network; Makan argues that repeated communication between peers at separate times is unlikely to occur, rendering this type of reputation system useless [10].

Gupta et al. [5] take a radically different approach to Wang. Instead of localized reputations, they implement a centralized server model on a P2P network. This implementation does not require all nodes in the network to use the reputation service. The centralization of the stored reputation is an effective way of ensuring all users can gain access to up to date reputation data, solving the client synchronization issue of distributing this data across nodes. However, this model once again assumes there to be no malicious actors in the network. Like Wang this model uses a binary scoring system and the reputation score is just the sum of reputations received for each peer. It is assumed that a negative reputation would not be transmitted. One innovative system in this implementation to prove a peer sent a piece of data is the creation of a receipt. A piece of data containing the file name, identities, etc. of both the parties involved is generated and signed by both parties' private keys, this is then sent to the central authority who can award reputation to both users. This would allow a multidimensional reputation system where a user can be rated on their actions on individual files.

Both of the systems proposed by Wang and Gupta fail to address both the issue of identity management – to ensure users can only obtain a single identity - and the possibility that peers may collude together in order to profit from the system to increase their own reputations.

B. Attacks on decentralised reputation systems

Attacking a reputation system can lead to significant benefits for an attacker. On an Ecommerce website, a user with high reputation can expect to receive an 11.2% premium on all goods they sell, which provides a motive for attack.

Perhaps the most challenging attack to prove and prevent is the unfair ratings attack. In this attack, an attacker provides ratings that do not reflect their genuine opinion of the rater in an attempt to lower the peer's reputation. Jøsang and Golbeck describe a possible defense against such an attack by comparing ratings of users to ratings left by higher trusted users on the network [11]. However, they fail to consider an attack first described by Lou whereby the peer is selectively malicious [12]; their proposed defense being that this type of attack would go undetected, and potentially penalize the honest node instead. This is an attack that has been conducted by GCHQ in an attempt to discredit selected targets [13].

Collusion is another popular attack that is common in reputation systems. This attack is based on a group of nodes who collude between each other with the aim of lowering a target node's reputation. One solution against an ongoing colluding attack is to calculate the reputation score based on the average of all reputations received from a peer.

The collusion attack is often deployed in conjunction with the Sybil attack. The Sybil attack is where a single user gains access to multiple legal identities. While Jøsang and Golbeck do not describe any countermeasures [11], Douceur describes how the success of a Sybil attack depends on the cost of obtaining an identity [14], and clearly shows how the effectiveness of a Sybil attack is reduced when the cost of

generating a new identity increases. The most effective countermeasure is to link the identity to a real world identity, as described by Yu et al [15]. However, the disadvantage of such a countermeasure is while this all but prevents a Sybil attack, it makes entrance to the network expensive for the network, due to the resources required to verify every user, a solution that would not scale well.

The re-entry attack also exploits the cost of entry to a network. With this attack, an attacker can choose to behave maliciously; once they have a low reputation that impacts their attack, they stop using that account and generate a new account and use this, this method is constantly repeated. Prêtre rightly appraises this attack as efficient not only because of the low cost of entry to the network, but the network sees a user with zero reputation scores as higher than a user with negative scores, providing the user with an incentive to dispose of the account [16].

While the majority of reputation systems currently deployed are vulnerable to these - and more - attacks, Jøsang and Golbeck question whether it is necessary for the reputation system to be perfectly secure [11]. They argue that, in the majority of situations, there is little incentive to attack the network, and the value of a reputation system lies elsewhere.

III. Our approach

We propose a general blockchain based reputation system that aims to solve several major challenges that the previous generations of reputation systems have failed to resolve, as well as preventing attacks that are possible on current generation reputation systems. We will focus on the application of this system on a peer-to-peer network, although it is also just as easily deployed on a classic E-commerce website.

Blockchain technology is a novel peer-to-peer approach to linking a sequence of transactions or events together in a way that makes them immutable. This was originally described by Nakamoto and implemented for the virtual currency Bitcoin [17]. In Bitcoin, users exchange money using transactions much like in real life. When a user creates a transaction he broadcasts this to all peers in the network. A special group of peers, called miners, collect broadcast transactions and attempt to incorporate them into a block that satisfies a cryptographic hash function. The process of producing a block is computationally intensive and probabilistic. Given a proposed block, each miner has a fixed and independent probability of successfully producing a block which satisfies the hash function for each unit of computation time. Whilst producing a block is hard, verification of a correct block is not.

Blocks are also linked together by chaining the hash of the previous block with each subsequent one. Thus, an attacker must control a significantly proportion of the computation power (typically 51%) to produce one false block and faking transactions back into the past is exponentially hard.

The collection of blocks (and their transactions) is called the ledger in Bitcoin and is publically inspectable by any peer. Thus a peer can see and verify any transaction from any point in time.

The blockchain was first described by Nakamoto in his paper describing the Bitcoin protocol [17]. The blockchain is a public leger of all transactions that have ever been completed since the first “genesis” block. Each transaction from the Bitcoin protocol is broadcast to all nodes in the network which are maintaining the blockchain, known as miners.

These miners check the transactions were valid (e.g., sender has enough coins to send) and then package all the valid transactions into a block. All nodes have a complete copy of the blockchain and keep this up to date. The block must contain a cryptographic hash of the previous block, this is the method used to cryptographically link every block in the blockchain to its previous block, all the way back to the first, genesis block. Once the block has been assembled, all miners on the network undertake a challenge of finding a nonce, so that the hash of the current block contains a set amount of zeros at the start. This process is commonly referred to as mining. Mining is a competition between all miners on the network, and the first miner to find the nonce and publish this confirmed block to the network receives a set amount of Bitcoins.

The use of previous hashes in each block prevents any attack where the contents of a block is changed, as if this were to happen that block and all subsequent blocks hashes would not match up. The only method a user would be able to use to change data in a previous block is to control 51% of all computational power on the network. Known as the 51% attack, this attack requires a majority of the computational power to be used to “re-mine” each block from the block that was altered. This would require a substantial amount of computing power, as the Bitcoin network currently has 510,000,000 GH/S [18] of computational power solely dedicated to mining, which is 256 times more powerful than the combination of the top 500 supercomputers in the world [19].

It is this property that makes the blockchain into a very secure ledger, which will remain secure to all adversaries who control less than 51% of the computational power of the network, as the cost of resources required to control 51% would outweigh the potential rewards.

IV. Description of our approach

We propose a new reputation system based on the blockchain technology. To reduce load on the current Bitcoin blockchain and to reduce inflation of the blockchain, we will create an entirely new blockchain, the sole purpose of which is to store reputation from completed transactions.

The proposed network has two goals – to withstand previously documented attacks on reputation systems and to provide a generalised reputation system that can be implemented into any network.

In a peer-to-peer network environment, we propose to solve the issue of quantifying reputation by removing the human opinion from the transaction. Our system will only store single dimensional reputation, with each user leaving either a 1 for a positive transaction, or a 0 for a non-satisfactory transaction. A positive transaction is classified as a transaction in which the user received the file they requested.

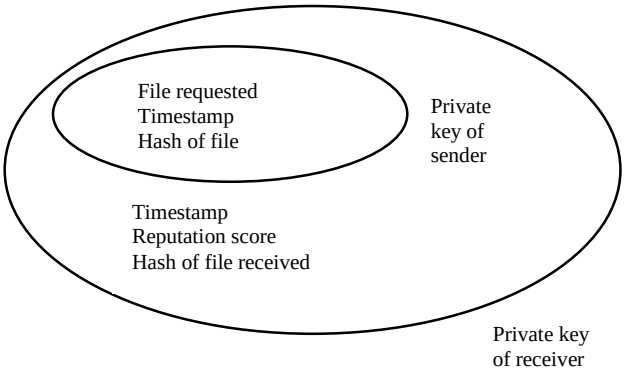

Figure 1*: Receipt of transaction sent to the miners*

We classify a transaction as the sending a piece of data, such as a file, signed by the sender's private key to a user who requested it.

Upon receiving the correct file, the user sends a transaction consisting of the reputation score, a timestamp, and a hash of the received file. This data is then encrypted with the receiver's private key and is sent to the miners. This ensures the reputation left by a user is based on a real transaction, a major issue in current generation reputation systems. The unfair ratings attack is now no longer possible since there is now cryptographic proof the user sent a requested file, and the user received it.

Fig 1 is a diagram of the format of a transaction which would be sent to the miners

The miners check the validity of the transaction by contacting each user involved in the transaction, and requests a signed proof, containing the file hash and a random nonce sent by the miner to be included. This is to prove each user sent/received the file, however this does have the drawback of requiring the users to still be online, for the miners to verify the transaction. The miners then assemble these verified transactions into a block of other transactions before confirming them in a method identical to current Bitcoin implementation. Fig 2 shows some pseudo code detailing how a miner would verify a transaction.

A method to ensure users cannot generate multiple identities cheaply is to link the indemnity creation to the IP address of a user. IPV4 addresses are becoming more expensive to purchase, as there is a lack of them available. While this method does not prevent an attacker from creating multiple identities, it makes the cost of doing so much more expensive, thus deterring all but the most well-funded attacker.

Identity based encryption systems with the ability to generate a public key based on an email address were also evaluated and tested; this was a desirable feature, however the requirement of a centralized server to generate all public/private keys made this option unsuitable for our system.

The ability to prevent multiple identities from a single machine, is key in preventing a Sybil attack, this combined with the expensive cost of entrance [20] to our network, makes it unviable for all but the most powerful adversary to conduct a Sybil attack on the network. To adapt this system for an E-commerce network, the data sent to the miners would be the Bitcoin transaction hash, the public key of the sender of the item and the public key of the receiver.

```
Miner Verification

For each transaction
        Connect (tx, ip)
        Send (random nonce)
Response ← receive
Verify (response, tx, nonce)
 ➔ Add to block mining queue ()
```

Figure 2: *Pseudo code for miner verification of a transaction*

To reduce malicious transactions on the network, we also propose a proof-of-stake system, where a user with a low, or no, reputation stakes a small amount of currency (Bitcoins) into a triple signed wallet. A triple signed wallet is a wallet created with three sets of keys, one from the sender, one from the receiver and one from an impartial third party. When a low-reputation user wants to share a file, they demonstrate they are honest by sending a small amount of currency to the wallet set up especially for this transaction; this would mean if the user were to behave dishonestly and send a malicious file, the amount stored in the wallet would be sent to a pool which the network uses to act as a reward for miners finding blocks. This is chosen to discourage any user from trying to profit from this feature. If the transaction were to be conducted honestly, the file sender would receive the amount they staked back.

To ensure this network cannot be affected by a 51% attack in the early days of deployment we utilize the power of the Bitcoin network by using merge mining. Merge mining allows all miners on the Bitcoin network to use their hashing power on our reputation system. This does not reduce the hashing power of the Bitcoin network, but does increase the total hashing power of the reputation and thus the security of the reputation system, as now to conduct the 51% an attacker would need to control the majority of computing power of both the Bitcoin and reputation network.

As well as the distributed blockchain, which ensures every peer has a full copy of the blockchain, eliminating client synchronisation issues as faced on previous distributed reputation systems, we also use the "friend peer reputation" model. As well as publishing all reputation about transactions onto the blockchain, the client also stores reputation from peers it has had previous interactions with. This can be multi-dimensional reputation, such as speed of the transaction, quality of file, etc. This information is not published to the blockchain as it would increase the cost of storage required per transaction and more importantly it is subjective from a user's perspective.

The final component of our reputation system is how to calculate reputation score of each peer. Reputation scores are not published on the blockchain. Unlike most previous generation reputation systems where the reputation client is community controlled, our proposed reputation system is client controlled. The client can calculate the reputation score based on parameters set by them. For example, a user could only view reputations from users on a specific network. To prevent against the collusion attack, where multiple users trade between

themselves multiple times in order to unfairly gain reputation, each user will only be given a reputation score based on the average of all their reputation score. This ensures if two nodes are transacting together, they will get the same reputation scores whether they send one transaction or a thousand transactions to each other.

For the network to have the property of temporal adaptability the client could only rate users from reputation over a short period of time. Josang et al. [21] demonstrate a user's behaviour in the last few days is a more accurate indicator of the user's future behaviour than analysing all previous behaviour on the network.

To select a user, they wish to download a file from, for example, a user finds all the peers which are hosting the file, the client then calculates the reputation for each peer using data from the blockchain and also using the friend peer reputation data to calculate a list a of the most reputable peers. Only requiring the client to calculate reputation of a small subset of peers reduces the computational resources required by the client. Once the user has calculated the most reputable client they can initialize the download. This method of peer selection can be used for E-commerce and other type of networks.

V. LIMITATIONS, ANALYSIS AND SOLUTIONS

As with any network there are some limitations in the deployment and use of this network. The majority of the limitations we faced were due to fundamental flaws in the architecture of the blockchain protocol

Unlike the majority of peer-to-peer networks, where network growth is uncapped, and will continue to grow as long as new nodes join and stay in the network, a blockchain based network has a hard limit on the number of transactions that can be processed per second. EBay currently process on average 23,148 reputation transactions a second, however due to requirement of a block being mined every ten minutes, and a maximum block size, our network would only be able to process 10 transactions a second. This is a significant reduction in the transactions our proposed network is able to process a second compared to a more traditional, previous-generation reputation system.

If the network were to receive more than 10 transactions a second, the miners would be forced to queue the reputation scores which would be included in a later block. This is not just an inconvenience to users who are relying on the network, it could also open the door for a denial of service where malicious colluding nodes would spam the miners with transactions, forcing miners to conduct computationally expensive verification of these transactions and forcing genuine users' transactions to be queued and delayed.

The "hard limit" on the number of transactions that can be processed a second also limits growth of the network and could render this application useless for some scenarios.

We will look at solutions to this problem later on in this section.

Another limitation on how effective and successful the reputation system is to be is the global deployment and adoption.

Currently, in addition to the issues mentioned above, other issues stopping this network from being deployed and implemented on a large scale is that the required resources on each node make this expensive to implement, with the proposed 1MB block size, (the same as the Bitcoin network) the blockchain could increase at a rate of 144MB a day (53GB a year).

These properties make it unlikely that a network with a high amount of low resourced users, such as mobile users, would implement this reputation system. This is a critical part of the success of the reputation system

It would take several months from deployment for the reputation system to become effective, gaining the necessary data and feedback from users that would allow other users on the network to make informed decisions on the trustworthiness of a peer. It would therefore take several months from deployment before the full potential of this reputation system would be noticed.

While we have proposed a system that solves a number of known issues with current generation reputation systems, and which secures them using cryptographic functions, the risk of unknown technical flaws in the cryptography used could undermine security on the network.

The final limitation of the proposed network is un-defendable attacks, such as an intelligent colluding attack. While we have proposed countermeasures for such an attack, it might still be possible for an attacker to profit from the system. The impact of such an attack should be low, and with all the aforementioned countermeasures implemented such an attack would be very expensive to conduct, but we will never be able to defend against all possible attacks with 100% success rate.

VI. ANALYSIS OF LIMTAITONS

In the section we will conduct analysis of various methods to reduce the limitations of the network through simulations and calculations. We will also compare our proposed solutions to the current implementation and compare the results to other networks. One simple solution to increase the number of transactions per second would be to remove the maximum size of a block. This would increase the number of transactions per second the network would be able to compute, for example an increase to a 5MB block size would allow for 50 transactions per second. However, for this system to match EBay's 23,148 reputation transactions per second the block size would need to be 2.351GB causing the blockchain to increase in size by

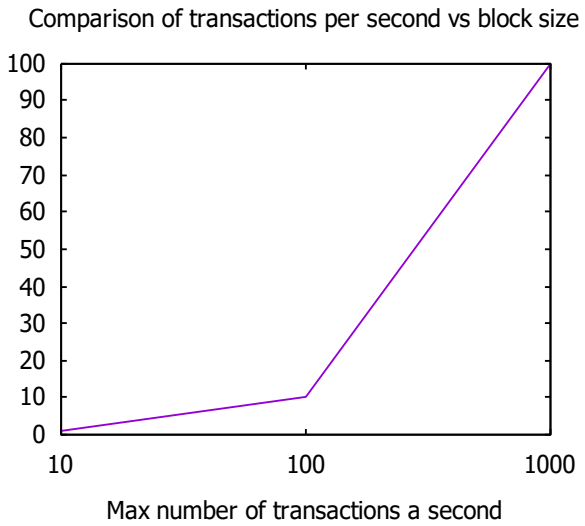

Figure 3: *Graph showing relationship between block size and transactions a second*

339GB a day; this is unstainable and shows that increasing the block size is not the solution.

Fig 3 shows how the increased block size increases the number of transactions per second.

Another method to increase the transactions per second, is to reduce the time required for each block to be mined. Currently the difficult of the proof of work is calculated such that a block is confirmed every ten minutes. This could be reduced to 5 minutes, or even a single minute, to increase the transactions per second the network is able to process.

Both methods of increasing the block size would increase the resources required by the user, such as more storage space to save the blockchain, as well as greater bandwidth to receive blocks at an increased rate. This would also further limit the participation of low-resourced nodes such as mobile devices. We therefore propose that each node is no longer required to download the entire blockchain, instead only the miners would be required to download and keep up to date the entire blockchain. This would change how reputations for users are calculated by the client; they would now be required to contact a pool of miners requesting the data for a specific user. A pool of miners will be used to prevent a malicious miner sending incorrect data to the requester, as in a pool, a majority of the miners would need to be malicious for this to occur.

We calculated the probability of randomly selecting a malicious pool (where 50%+ of the pool is malicious) for varying amounts of network compromisation, in comparison to randomly selecting a single miner using the equation below. We then simulated this model in python before plotting the results on a graph as seen in Fig 4.

$$\rho(m) = \binom{k}{m} p^m (1-p)^{k-m}$$

As shown in Fig 4, this method is very effective up to 40% of malicious nodes in the network, and effectively solves two

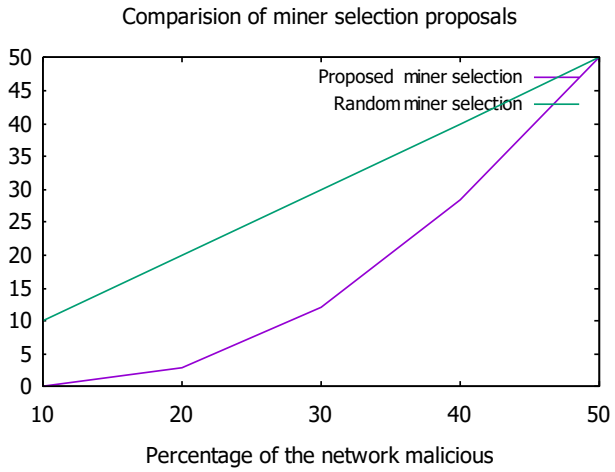

Figure 4: *Comparison of miner selection algorithms*

limitations by allowing low-resourced users to join, as well as increasing the number of transactions per second. This demonstrates our proposed network is able to handle double the amount of malicious nodes supplying malicious reputation data as the reputation system proposed by Zhou [22]

A. *Analysis of results*

We have looked at the reputation system proposed and described some limitations faced during implementation. To ensure these limitations were mitigated, we developed a series of countermeasures to ensure the proposed network is as deployable as possible, in order for it to be successful.

The solutions to the limitation issues have now improved the scalability of the network. The countermeasures proposed and simulated in this paper could be implemented into any blockchain based application which is having scalability issues.

Changing the block confirmation time from ten minutes to five not only aids with scalability of the network, doubling the number of transactions that can be processed per second, it also increases security, as now a malicious peer would be able to be detected 50% faster than before. This increase in detection time was an unexpected benefit.

There could however be negative impacts caused by our recommended changes to solve the limitation issues. The increased resources (storage space for the blockchain) on the miners could result in fewer miners on the network; this would in turn lower the security of the network, however the blockchain of the reputation system would still be significantly smaller than Bitcoin's blockchain for at least the first two years of deployment, so we do not see this actually happening. Another perceived negative impact is the calculation time for a peer to calculate a user's reputation will be higher, this is due to the peer now needing to request this data from a pool of miners. The network latency and processing of this request would add a small delay, but this would not be significant enough for the user to notice.

VII. Conclusion

In this paper we have discussed a next generation reputation system based on the blockchain, we have shown how a generalized reputation system that could be implemented into various networks is possible. We discuss in detail how the reputation system would be implemented and demonstrate how our proposed system solves many of the issues faced by current reputation systems. We conducted analysis on the limitations faced by our system before describing how these could be overcome.

Overall, this paper aimed to propose a reputation system which solves the majority of issues faced in current reputation systems. However, this is just the foundation of the idea and there is a lot more research to be conducted in the future in various areas to ensure this reputation system is capable of replacing all reputation systems in the real world.

VIII. Future work

This paper has shown how a reputation system could be easily implemented on a blockchain, and how our proposed reputation system theoretically solves the majority of issues faced by current generation systems. However, this is just the beginning of development of this network, and there are still many avenues of research left to pursue in this area.

The most important piece of work to conduct in the future is to make this proposed network live. This will then let us examine in greater detail if the assumptions in this paper hold true on a real world.

We cannot yet answer questions such as whether a user who acts honestly on one network can be assumed to act honestly on all networks they interact with, or when does past reputation for a user become irrelevant, but with more research we hope to be able to resolve these questions and more besides.

The deployment onto a live network would also enable more accurate analysis of how users interact with the reputation system to allow a more accurate algorithm for calculating reputation scores to be refined.

The deployment onto a real world network would also allow us to see if our solutions to known issues and limitations hold true, or if new issues surface.

This paper has so far assumed a user does not worry about privacy, however there is a growing consensus that privacy is a critical factor in using any web application, so it would be a very interesting research area to consider if privacy can be implemented on a reputation system without succumbing to attacks which exploit the weak links between identity and users.

We have focused on two applications for this system; an Ecommerce eBay type application where users can rate if they received the item, and also a peer-to-peer network, where users can rate each other peer if they have provided the correct file, in an attempt to detect any malicious nodes spreading malicious files through the network. It would be beneficial to the future success of this network if other implementations in these applications where possible. For example, instead of just rating a peer on whether it sent the correct file in a peer-to-peer network, could this system be adapted to bittorent and used to provide each client with the optimum download and upload speed, allowing each users to rate a series of other criteria to provide a better service to the client.

The final area for future research is how to optimize the blockchain. Could pruning the blockchain be a possibility in this situation, this would allow the network to scale higher due to the lower resources needed.

These are just some of the interesting research areas that we have yet to fully analyze, and with more research this project could be the next generation of reputation systems.

References

[1] G. Prisco. (2015, May 14).*The World Table Launches a Quantified Reputation System* [online]. Available: https://bitcoinmagazine.com/articles/world-table-launches-quantified-reputation-system-1431633676 (Access date: 03 Decemeber 2015)

[2] WSO2. (2011). *EBay uses 100% Open Source WSO2 Enterprise Service Bus to Process more than 1 Billion Transactions per Day* [online]. Available: http://wso2.com/download/wso2-ebay-case-study.pdf (Access date: 03 Decemeber 2015)

[3] Ebay. (2015). *Changes to Feedback – FAQ* [online]. Available: http://pages.ebay.co.uk/help/sell/feedback-faq.html (Access date: 03 Decemeber 2015)

[4] P. Dewan and P. Dasgupta, "Securing reputation data in peer-to-peer networks", in Proc. of Parallel and Distributed Computing and Systems (PDCS 2004), Cambridge, MA, 2004.

[5] M. Gupta, P. Judge and M. Ammar. (n.d.). *A Reputation System for Peer-to-Peer Networks* [online]. Available: https://www.cs.indiana.edu/~minaxi/pubs/reputation.pdf (Access date: 03 Decemeber 2015)

[6] E. Damiani et al. (n.d.). "A Reputation-Based Approach for Choosing Reliable Resources in Peer-to-Peer Networks," in Proc. of the 9th ACM conference on Computer and Communications Security, [2002]© ACM. doi: 10.1145/586110.586138

[7] K. Walsh and E. Gün Sirer. (2006). *Experience with an Object Reputation System for Peer-to-Peer Filesharing* [online]. Available: http://static.usenix.org/event/nsdi06/tech/walsh/walsh_html/ (Access date: 03 Decemeber 2015)

[8] Y. Wang and J. Vassileva. (n.d.). *Trust and Reputation Model in Peer-to-Peer Networks* [online]. Available: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.89.2915&rep=rep1&type=pdf (Access date: 03 Decemeber 2015)

[9] L. Xiong and L. Liu, "PeerTrust: Supporting Reputation-Based Trust for Peer-to-Peer Electronic Communities", *IEEE Trans. Knowl. Data Eng,* vol. 16, no, 7, pp. 843-857, Jul. 2004

[10] J. Makan and M. Kutar. (n.d.). *Trustbusters: Enforcing Account Creation* [online]. Available: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.385.9031&rep=rep1&type=pdf (Access date: 03 Decemeber 2015)

[11] A. Jøsang and J. Golbeck, "Challenge for Robust Trust and Reputation Systems," in Proc. of 5th International Workshop on Security and Trust Management, [2009] © Elsevier Science B.V, Sept. 2009

[12] X. Lou and K. Hwang, “Collusive Piracy Prevention in P2P Content Delivery Networks,” *IEEE Transactions Comput.*, vol. 58, no. 7, pp. 970-983, Jul. 2009

[13] G. Greenwald. (2014, February 24). *How covert agents infiltrate the internet to manipulate, deceive, and destroy reputations* [online]. Available: https://theintercept.com/2014/02/24/jtrig-manipulation/ (Access date: 03 Decemeber 2015)

[14] J. Douceur. “The Sybil Attack,” In *IPTPS '01 Revised Papers from the First International Workshop on Peer-to-Peer Systems*, Cambridge, MA, 2002, pp. 251-260

[15] H. Yu et al, “SybilLimit: A Near-Optimal Social Network Defense against Sybil Attacks,” in *2008 IEEE Symposium on Security and Privacy*, [2008], © IEEE. doi: 10.1109/SP.2008.13

[16] B. Prêtre, “Attack on Peer-to-Peer Networks,” Semester Thesis, Dept. of Computer Science, Swiss Federal Institute of Technology (ETH) Zurich, 2005

[17] S. Nakamoto. (n.d.). *Bitcoin: A Peer-to-Peer Electronic Cash System* [online]. Available: https://bitcoin.org/bitcoin.pdf (Access date: 03 Decemeber 2015)

[18] Blockchain. (2015, October). *Hash Rate* [online]. Available: https://blockchain.info/charts/hash-rate (Access date: 03 Decemeber 2015)

[19] R. Cohen. (2013, November 28). *Global Bitcoin Computing Power Now 256 Times Faster Than Top 500 Supercomputers, Combined!* [online]. Available: http://www.forbes.com/sites/reuvencohen/2013/11/28/global-bitcoin-computing-power-now-256-times-faster-than-top-500-supercomputers-combined/ (Access date: 03 Decemeber 2015)

[20] A. Mohaisen and J. Kim, “The Sybil Attacks and Defenses: A Survey,” *Smart Computing Review*, vol. 3, no. 6, pp. 480-489, Dec. 2013

[21] A. Jøsang et al. (n.d.). *Simulating the Effect of Reputation Systems on e-Markets* [online]. Available: http://folk.uio.no/josang/papers/JHF2003-iTrust.pdf (Access date: 03 Decemeber 2015)

[22] R. Zhou and K. Hwang, "PowerTrust: A Robust and Scalable Reputation System for Trusted Peer-to-Peer Computing," *IEEE Trans. Parallel Distrib. Syst.*, vol. 18, no. 4, pp. 460-473, Apr. 2007