

Zlto: Increase and Track Positive Behavior using Secure Blockchain Technology

Michael Gu^{1,3}, Allan Van Der Meulen¹, Kurt Appolis¹, Christine Taphel¹, Marlon Parker¹, and Gary Wills²

¹RLabs

²University of Southampton, Cyber Physical Systems Research Group

³University of North Carolina at Chapel Hill, Department of Computer Science

Abstract—Zlto(zlah-toh) is a project to provide secure validation of credentials for young people. Many African youth find it difficult to obtain jobs because of a lack of formal qualification, but they have a breath of informal skills. In this paper we will describe how we implement the Zlto system to provide proof of skills completed through a peer-to-peer evaluation mechanism and stored using blockchain. Since its first deployment in 2016, Zlto has processed over 1 million transactions. We present simulation and analysis of the reputation protocol used to verify work, and further evaluate the peer review process using a think a loud focus group method with a group of peer reviewers in order to improve the system.

I. INTRODUCTION

Many young people in South Africa find it difficult to access jobs and economic opportunities due to a lack of skills and appropriate networks to pursue such opportunities [1]. The Zlto project aims to address the youth unemployment challenge that found in many parts of the world, including South Africa. In particular Zlto is aimed at young people who lack the appropriate skills or economic support to enter the job market, and as a result are disconnected, excluded and discouraged in their job search.

The Zlto project provides youth with access to work opportunities that enables them to develop work experience and also allows them to initiate education related voluntary community work (tutoring children, cleaning schools, voluntary work with partner organizations, etc) in their local communities.

While accessing these training and development opportunities, youth are rewarded Zlto digital currency, allowing them to access a higher level of opportunities such as formal jobs, education, and access to small business finance or financial services.

The digital currency is redeemable for commodities and services such as transport (addressing the challenge of high costs of mobility for job seekers), toiletries and sanitary pads, basic food supplies, mobile phones, daycare services or training toolkits which will enhance their learning experience, particularly for girls.

Because the Zlto platform gives youth autonomy to earn and spend their incentives, they are provided an opportunity to increase financial literacy to improve financial inclusion in South Africa. In using the Zlto platform, youth can learn about digital payments, budgeting, leading to financial literacy and will allow them to build a credible history of

work that they can use as references while seeking for job opportunities. These opportunities includes training, coaching, mentor-ship, business development services, internships, access to finance and financial services. These range of opportunities are offered by the Zlto project and affiliate partner organizations.

The youth on the Zlto platform will also be able to build credible work and learning experiences by having their evaluated work asset being stored on the blockchain [2]. This work asset will be fully under the control of the young person.

The proof of work asset is validated through a peer-to-peer evaluation. The proof of work asset, held on the Zlto platform acts as a virtual referral and proves the youths lifelong learning and work experience history.

The paper first presents an overview of the Zlto project, specially the peer-review process, then we discuss the main algorithms for calculating the reputation, and the reputation flow. Then we explore the advantages of a hybrid blockchain model and how we secure the transactions. We then present the evaluation of the peer review process and discuss the advantages of this approach for e-society. We concluded with a look at the future work of this project.

II. OVERVIEW OF ZLTO

Zlto is a rewards system used to facilitate community impact. Impact is verified by peers in the Zlto network and verifications are then securely stored using blockchain technology. The impact can subsequently be purchased by investors, who directly fund the inventory of rewards available to users. The Zlto backend infrastructure draws from several leading blockchain technologies. The primary aim is to create a highly reliable and scalable system for evaluating work and storing verified work assets. The systems essential principles are immutability, low latency, and a high secure transaction rate.

The Zlto system is composed of 4 key elements:

- 1) Peer-to-peer evaluation protocol deployed as Ethereum smart contracts
- 2) Private blockchain for processing smart contracts, periodically synced with the public main net
- 3) Private Byzantine Fault-tolerant storage nodes
- 4) Zlto client facing applications

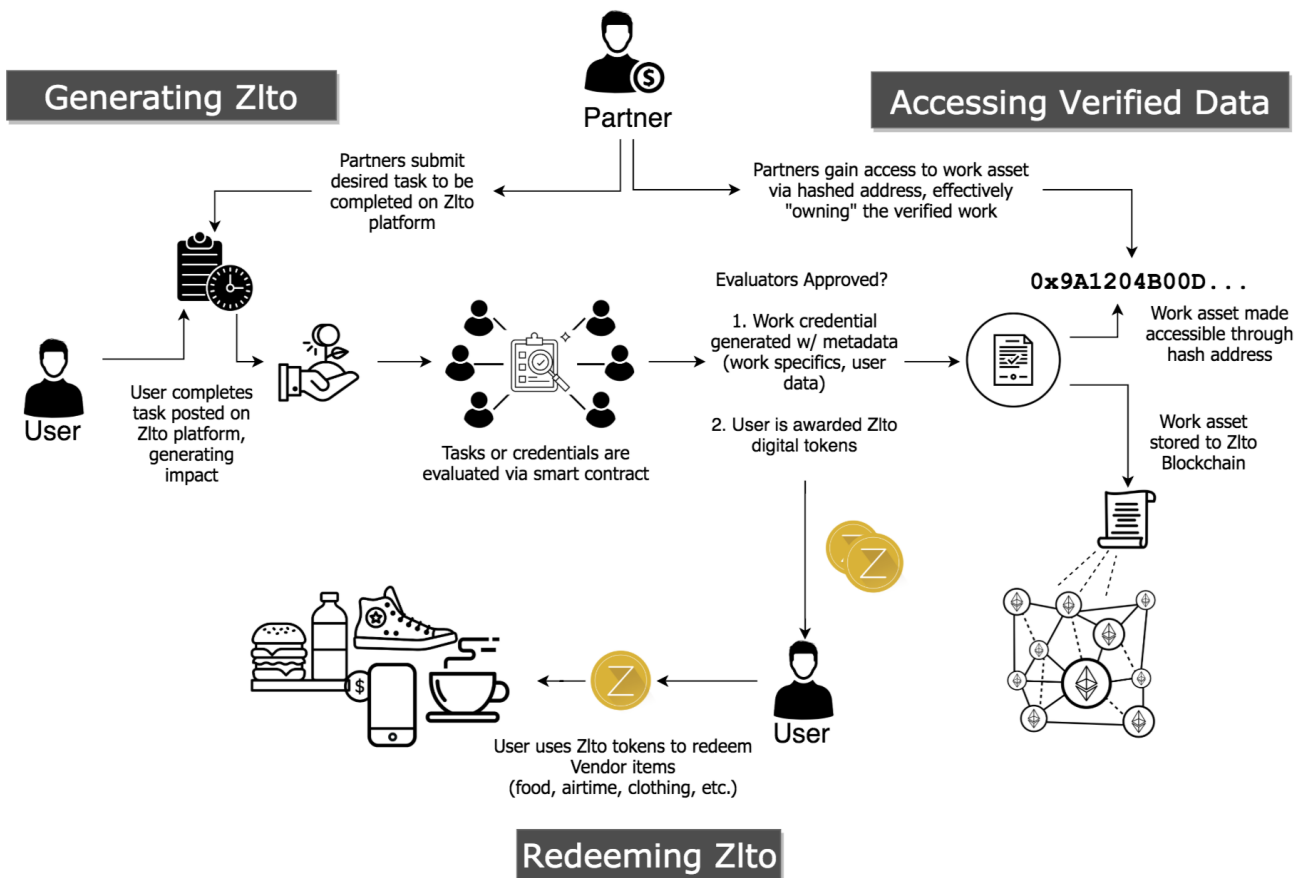


Fig. 1. The interaction of the different community stakeholders with the Zlto platform

III. P2P EVALUATION PROTOCOL

The Zlto system requires verification of the validity of informal work submitted by users in order to properly incentivize positive behavior.

The first Zlto pilot in 2016 attempted to verify work through an internal employee approval process, which quickly became a significant bottleneck to the platform's scalability. To achieve this core functionality without this bottleneck, we present a method for peer-to-peer evaluation of work submitted by Zlto users that are eventually approved or denied.

We present a protocol that attempts to formalize the interaction of peers on a network who collaborate to evaluate the integrity of *work* in a peer-to-peer context. Work is data signalling the completion of a service; it can be fulfilling a work duty, performing community service, or providing a resource.

Each interaction begins with a *requester*, a peer who proposes work to be evaluated. Several *judges* subsequently perform the evaluations and submit a final judgment on the work's integrity. In submitting evaluations, the judges invest a stake of their reputation, which is eventually converted to reputation gain if their evaluation is agreeably accurate as

viewed by other peers.

The evaluation cycle ends once the reputation management protocol believes consensus among judges has been reached, and reputation is rewarded or reallocated.

The capacity of reputation to be sought after depends on the implementation of the reputation management protocol. Specifically, the reputation management protocol can be divided into two compatible processes: reputation stake and reputation flow calculations.

In this way, the Zlto reputation management system serves the purpose of reaching consensus on the integrity of proposed work that a requester submits. The algorithms selected for the protocol accomplish the following goals:

- 1) High expertise judges have a higher probability of reputation gain than those with lower expertise (who are possibly malicious in nature)
- 2) Malicious claims to work are denied with high frequency
- 3) Genuine claims to work are accepted with high frequency

We further attempt to approximate the operational frequency of achieving each goal through simulations found in part IV.

A. Calculating Reputation Stake

The amount of reputation that a judge must stake in the process of submitting an evaluation varies between reputation management protocols.

The Zlto reputation flow protocol is influenced by the Backfeed¹ protocol, where each judge incurs a cost that decreases in proportion to how much reputation has already been committed:

$$Cost_{judge_i} = s * r_i * (1 - (V_k/R)^\alpha) \quad (1)$$

s = stake fraction (in [0,1])

r_i = reputation of judge i before submitting an evaluation

V_k = sum of reputation of judges who submitted an evaluation

R = sum of reputation of all judges in the system

α = reputation payment skewness

It is important to note that the Backfeed protocol operates in an environment called a Decentralized Collaboration(DC), where a fixed number of peers operate in a closed group for duration of the evaluation cycle. In this closed network the sum of reputation, R , can be arbitrarily calculated.

Instead of operating in a closed network, Zlto operates in an environment where N peers may enter and leave the network during the evaluation cycle (N is unknown and potentially unbounded). Moreover the localized nature of the Zlto user base increases the threat of malicious collaboration between requesters and evaluators in approving false claims.

Given these differences in Zlto's operating environment, we modify the Backfeed reputation cost calculation by defining a reputation *bounty* for each request submitted. In practice this value is a function of the number of hours worked and the number of people affected by the user.

This reputation *bounty* is then used when calculating the proportion of reputation committed after each evaluation submission, effectively replacing the R variable defined in Backfeed.

B. Reputation Flow

Reputation flow or re-allocation is a mechanism by which reputation changes for each judge during an evaluation cycle [5]. As subsequent judges submit evaluations, reputation is updated for each prior judge in a manner that depends on how the two align with respect to their evaluation.

When judge n ($n > 0$) submits an evaluation, the reputation is recalculated for all previous judges $\{judge_1, \dots, judge_{n-1}\}$. We update each judge's reputation using the following equation:

$$Update_{judge_{i < n}} = r_i + *(\delta_{i,n} * d * r_n * r_i) / W_k \quad (2)$$

r_i = reputation of judge i at the time judge n submits an evaluation

d = stake distribution fraction

r_n = reputation of judge n after incurring the staking cost

$\delta_{i,n} = 1$ if judge i 's evaluation agrees with judge n 's evaluation, 0 otherwise

W_k = sum of all reputation of those judges who agree with judge n 's evaluation

After each evaluation submission, our reputation flow algorithm aggregates staked reputation for all judges to calculate a new reputation for each judge:

$$NewReputation_{1\dots n} = ReputationFlow(\{stakedReputation_1, \dots, stakedReputation_n\})$$

When the evaluation cycle ends, each judge is left with a final reputation value. Whether a judge's final reputation is higher or lower than before the evaluation cycle should be a reflection of their expertise and the degree that their judgment aligns with subsequent evaluations submitted.

$$FinalReputation_n = NewReputation_{nx} \quad (3)$$

where x represents the number of evaluations submitted during the evaluation cycle when the final judgment for the work in question has been determined.

```

1  {
2      "id": 1,
3      "requesterId": 1,
4      "metadata": {
5          "type": "education",
6          "numHours": 10,
7          "numPeopleImpacted": 50,
8          "repToBeGained": 10,
9          "createdAt": 1533841331
10     }
11 }

```

Listing 1: An example request for evaluation transaction

```

1  {
2      "id": 1,
3      "evaluator": {
4          "name": "Toby",
5          "id": 1,
6          "reputationBefore": 50
7      },
8      "createdAt": 1533841872,
9      "judgment": true
10 }

```

Listing 2: An example peer evaluation transaction

¹<http://backfeed.cc/assets/docs/TechnicalSummary.pdf>

IV. SIMULATION OF REPUTATION FLOW

We devise a method for simulating a real-world application of the protocol to determine the long-term state of our reputation flow algorithm. The simulation results enable us to determine the operational parameter space for our protocol's environment when deployed.

A. Claim Model

For each epoch of the simulation we generate a sequence of claims, where each claim holds a *bounty* and a *trueValue*:

$$claim = \begin{cases} trueValue \in [0, 1] \\ bounty \in [0.1, 10] \end{cases}$$

A *trueValue* of 0 represents a disingenuous claim posing a threat to the system, while a *trueValue* of 1 represents a claim submitted by an honest requester without malicious intent.

In the real-world Zlto system, a *bounty* is calculated using the number of hours worked and number of people that the work has impacted. For the purposes of the simulation we use a range compatible with our definition of threshold defined below.

We further define a *good claim* as one with *trueValue* > 0.5, and a *bad claim* as one with *trueValue* ≤ 0.5.

B. Judge Model

The simulation will proceed by generating judges and subsequently asking each judge whether they wish to risk their reputation and submit a judgement on a claim for potential reputation gain.

We generate 20 judges per claim, with each judge holding a different level of expertise represented by *accuracy* and *precision* variables ∈ [0, 1].

Each judge's center represents the judge's perception of a claim's true value. This perception eventually translates into a decision to join the evaluation.

We define a judge's *center* with relation to a claim's *trueValue* is a function of accuracy:

$$center_i = (1 - accuracy) * rand + accuracy * trueValue \quad (4)$$

Each judge's *estimate* is thus a normal random number with arguments being a judge's center and precision:

$$Estimate_i = normrnd(mu = center_i, sigma = precision_i) \quad (5)$$

To represent expertise differences within the judge pool, we define a *good* judge as having an accuracy = 0.85 and precision/std. dev = 0. In contrast all other judges are generated using a random accuracy and precision.

C. Judge Decision

In addition to generating an estimate of the claim, each judge must generate a *decision* that represents their willingness to risk their reputation to "join" the evaluation cycle and submit a formal evaluation of the claim.

Our chosen protocol only permits judges to submit a boolean decision representing "yes" or "no". However, in the real-world Zlto system, a judge can have an internal decision calculus that ranges between 0 and 1.

Thus we define a *Threshold* value for each judge that is a function of their confidence in their estimate and their risk tolerance.

Assuming that a judge's estimate is unclear if it is closer to 0.5, we define *Confidence* as the absolute distance between a judge's *Estimate* and the center of our scale that represents an unclear estimate (0.5):

$$Confidence_i = \frac{|0.5 - estimate_i|}{0.5} \quad (6)$$

Risk Aversion is defined as the inverse of the claim's bounty:

$$RiskAversion_i = \frac{1}{(bounty)^\alpha} \quad (7)$$

where α = risk constant, experimentally set to 4

Combining these variables, we define each judge's *Threshold* for making a decision as:

$$Threshold_i = Confidence^{RiskAversion} \quad (8)$$

Finally for each judge we generate a random number ∈ [0, 1].

If the random number falls under the *Threshold* calculated, the judge joins the evaluation cycle and stakes her reputation.

$$Decision_i = \begin{cases} 1 \dots rand \leq Threshold_i \\ 0 \dots rand > Threshold_i \end{cases}$$

TABLE I
DISTRIBUTION OF THRESHOLD GIVEN CONFIDENCE AND RISK
AVERSION VALUES

		Risk Aversion					
		0	0.1	0.5	1	2	∞
Confidence	0	1	0	0	0	0	0
	0.1	1	.794	.316	.1	.01	0
	0.25	1	.871	.5	.25	.0625	0
	0.5	1	.933	.707		.25	0
	0.75	1	.971	.866	.75	.5625	0
	0.9	1	.990	.949	.9	.81	0
	1	1	1	1	1	1	0

Table 1 shows the distribution of threshold values given various confidence and risk aversion values. Our definition of *Threshold* enables an incredibly risky judge (*Risk Aversion* = 0) to always risk their reputation and participate in the evaluation. A very risk-averse judge (*Risk Aversion* = infinity) will never join regardless of the level of confidence in their estimate.

D. Simulation Results

Each simulation consists of 100 trials, each with 100 claims and 20 judges. For each trial there is also a random proportion of "good" judges generated.

The simulation is run under several task models to understand the performance of our protocol when faced with different pools of claims.

1) *Task Model 1: Generate each claim with a random trueValue distribution:*

2) *Task Model 2: Generate claims with a bi-modal trueValue distribution:*

E. Simulation Discussion

Figures 2 and 3 show that our reputation flow protocol achieves the three goals of the protocol.

No more than 10% good judges is required in the network such that over 90% of good claims will be accepted and bad claims denied.

Under the bi-modal task model (where the true value of claims are distributed closer to 0 or 1) more than 95% of good claims are accepted and more than 98% of bad claims denied given 10% judges in the network.

We note that in production setting we select judges based on their high rate of activity on the platform. We further monitor their ability to reach a high threshold of our user base and flag suspicious activity. Thus the percent of good judges in production is likely to be significantly higher than 10%, but is indeterminable. Thus our simulation enables us to understand the performance of protocol for any number of good judges in the network.

V. A HYBRID BLOCKCHAIN MODEL

The Zlto system employs a consortium of public and private blockchain services [7]. The smart contracts powering the peer-to-peer evaluation protocol run on a privately deployed Ethereum blockchain using the Go Ethereum implementation (geth). Finalized work assets are also stored on private storage nodes with Tendermint secure replication middleware. Finally, we interface with the public by periodically applying encryption, hashing, and syncing our data with our ERC20 storage contract deployed on the Ethereum mainnet.

Our choice in using primarily private blockchains for the majority of our transactions lies in their scalability. Private blockchains are significantly better than public blockchains [8] and the network parameters, such as network congestion and transaction fees are known in advance. Private blockchains enable protection rights across a distributed, immutable ledger with predefined rules on who can perform changes and validate new blocks. Because validators are known in advance, risks of 51 percent attacks are significantly reduced. Because transactions are verified by only a few nodes with trusted processing power and availability, transactions significantly cheaper. Using the Proof-of-authority consensus algorithm, our private nodes are intentionally designed to be well-connected and ensure finality in less than a second.

Finally, our data is encrypted and hashed before syncing with the Ethereum main network through an ERC20 contract intended for hash storage. This ensures a small footprint on the mainnet, as storage costs can become prohibitive (3 cents/transaction). Storing only the encrypted hashes of data that is replicated privately maintains security of our users data while providing a mechanism to publicly prove our data contents should our system or a specific data history be called into question. Timestamps at each step of the work asset lifecycle, from evaluation cycle to storage and main net sync, ensures the integrity of our data.


VI. SECURE TRANSACTION

Our storage layer lies between the completion of the evaluation protocol and storage onto the Ethereum main network. The storage layer consists of a network of nodes dedicated to storing the complete work asset in a private, immutable manner. These nodes run a BigChainDB server, a blockchain-based database supporting immutability and fast querying, backed by a MongoDB instance sitting on each node for added availability and replication. To create a resilient network among our storage nodes that is resilient to data corruption (e.g. hard drive failure, hacking), we use the Tendermint consensus middleware. Tendermint allows a network of decentralized nodes to be Byzantine Fault Tolerant, ensuring the networks availability even when 1/3 of the nodes are down [9, 10]. Performance benchmarking shows that 64 nodes across 7 data centers on 5 continents can process thousands of transactions per second with commit latency of a few seconds [11]. Even under adversarial conditions such as malicious votes or validator crashes, the performance remains consistent.

Because Tendermint does not allow for forking of history, committed blocks cannot be reverted or invalidated in the future. Each transaction is stored upon cryptographic signature. Changes to the contents of a piece of data changes the signature, and thus changes are detected. Even if a user changes their public key, public keys are all known and thus there exists no proxy to hide data mutations. Every node in the storage network maintains a full copy of the data in its own MongoDB database. If one node is corrupted, it cannot affect the MongoDB sitting at the center of the other nodes; all other nodes will retain their copy of all the data to date, as there is no global MongoDB instance.

Storing the work assets of the youth on the blockchain also provides additional encryption and security [12, 13]. Leveraging the Ethereum blockchain will also ensure that there is no downtime, fraud or third-party interference. Youth will have full control of their work asset that they can leverage by giving access to future employees. Our client-facing application provides a testbed of users for our multi-layer evaluation, and storage protocol. Our users lay at the center of the Zlto system and act as the fuel that separates us from competing blockchain solutions targeting this area.

% Good Judges vs. % Good Tasks Rejected



% Good Judges vs. % Bad Tasks Accepted


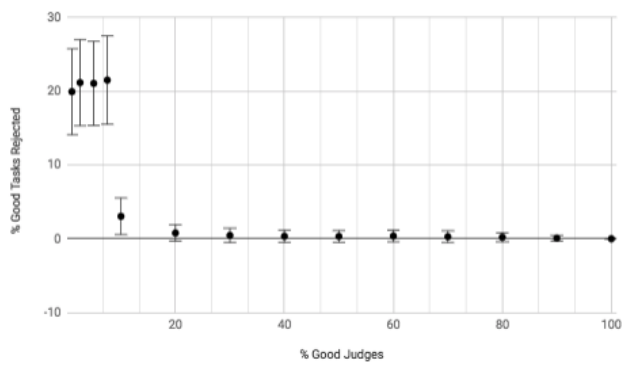


Fig. 2. Task Model 1: Random True Value Distribution

% Good Judges vs. % Good Tasks Rejected



% Good Judges vs. % Bad Tasks Accepted

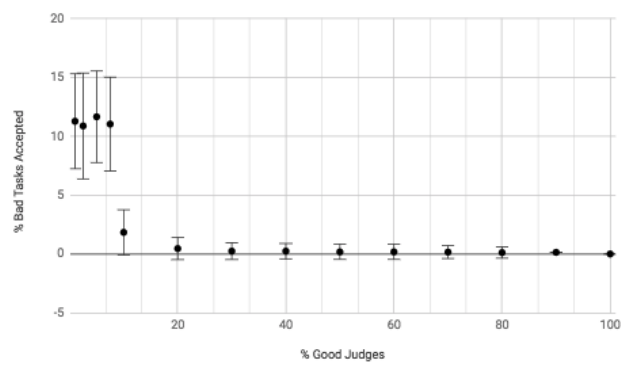


Fig. 3. Task Model 2: Bi-modal True Value Distribution

```
1 function notarizeHashes(  
2     bytes32[] idArray,  
3     bytes32[] hashArray  
4 ) onlyOwner public returns(bool) {  
5     require(msg.sender == owner);  
6     require(idArray.length == hashArray.length);  
7  
8     for (uint i=0; i<idArray.length; i++) {  
9         hashesById[idArray[i]] = hashArray[i];  
10  
11         emit DocumentAdded(idArray[i], hashArray[i]);  
12     }  
13     return true;  
14 }
```

Listing 3: Smart contract function for storing finalized work asset hashes

VII. EVALUATION

The system was evaluated through the use of focus groups of peer reviewers who experienced a deployed version of the protocol described. We begin with the selection of the first group of peer reviewers, the evaluation experience, and the feedback from the focus groups.

A. Selection of candidate for peer reviewers

The Zlto project team examined the records of the top 100 Zlto earners in order to find possible peer reviewers [14]. Peer reviewers are required to understand the various ways in which young people may earn Zlto. The team targeted the top 100 earners who had earned Zltos from a range of activities. For instance we identified reviewers who understood they could earn Zlto for teaching courses, but did not realize that youth could also earn Zlto for attending the course. Thus it was important for reviewers to understand the breadth of activities that youth can undertake to earn Zltos.

In addition the Zlto project team also selected youth who were active in providing feedback to the system. A peer reviewer is expected to contribute to the reviewing process and therefore it was important in the initial phase to obtain peer reviewers who were comfortable with providing feedback. The review process is not arduous, as the process takes around 30 seconds for a new peer reviewer (they generally do not get quicker at undertaking a review) and they are expected to undertake around 20 peer reviews a week (equating to 10 minutes for a new peer reviewer).

The final criteria for selection was a user's willingness to sacrifice 10% of the Zlto balance currently in their Zlto wallet. These sacrificed Zltos are used to generate each peer reviewers initial reputation score. We ensured that they understood their balance would increase as they successfully submitted reviews which are agreeable with those of their peers. Their reputation score may decrease if the peer reviewer does not participate in any reviews for one week or the peer reviewer approves a work asset review that is declined. The name of the peer reviewers are not known to the youth submitting the work assets; this protects the peer reviewer from undue influence from youth submitting the work asset. In addition to an increase in reputational score, a peer reviewer gets a 10% bonus on all the Zlto they earn from each approved work asset (activity) they undertake.

B. Process

From the top 100 earners of Zltos the Zlto project team invited 17 people to participate in the peer-review trial, 12 of whom consented to take part in the trial. The group of 12 were invited for training as peer reviewers. The Zlto team explained the user flow through the system, the various methods youth may earn Zltos, and how they as peer-reviewers earn reputation in the process.

The peer reviewers were trained on how to conduct peer reviews, how to write a good review, and what specifically to look for before approving or declining submitted work. For example someone claiming to have swept the roads but

the photograph submitted depicts a scene of an office, or an incompatible time of day, should be rejected.

The peer reviewers were then asked to rate a selection of work assets selected by the Zlto team. Some of these work assets were intentionally fraudulent claims. This ensures reviewers were experienced in identifying claims which resembled realistic scenarios but were lacking sufficient evidence.

C. Findings from the peer reviewers.

Once the group finished the training they were left for two to three weeks to conduct peer reviews on submitted work assets before being invited for a focus group. The focus groups was conducted in pairs of peer reviewers. The peer reviews were asked to go through the system, thinking aloud the problems they encountered.

Areas of concern for all the peer reviewers include the dichotomous nature of approving or declining work assets. They asked for an "unsure" option in cases where they did not have experience and their reputation score would thus not be affected. Moreover they asked for an "inappropriate" option that flags the submission for review by the Zlto team. An example would include a photograph submitted of work being done at a crche, where the user did not obtain permission to upload the photograph from parents whose children may be pictured.

One feature which all the peer reviewers requested was a channel to provide feedback for each submission in the case that more evidence was necessary. Currently each submission's description is contained in a free text box, but the Zlto team is looking to make this feature a fixed response in order to reduce the need for moderation and the possibility of inappropriate comments.

All the reviewers have remained in the program, but two of the 12 could only complete reviews in an environment containing free WiFi, as they did not have the funds for data charges incurred when downloading high resolution images contained in some submissions. This has now been resolved by compressing the size of images and building incentives for peer reviewers to buy data from their mobile network providers using their Zlto earnings.

VIII. DISCUSSION

Our system provides structured experimentation of blockchain solutions with strategic evaluation of feasibility against our user base and value at stake.

Blockchain technology allows young people to store their credentials from both formal and informal training skills providers as well as work experience that is validated by a network of peers on the Zlto platform. This enables youth who cannot traditionally access formal learning opportunities to validate their informal learning and acquisition of skills, ultimately gaining credibility for prospective employers.

The blockchain allows Zlto to scale by harnessing the power of decentralization. Via smart contracts, Zlto ensures that the validation of the social impact can be managed by

appointed validators on the platform. It will also be key in scaling Zlto into other regions across Africa.

Youth using the Zlto platforms will also be able to create savings through their digital wallet to access products and services that they would not be able to obtain due to their low income or no income backgrounds.

The usage of digital services globally where access to financial service is increasing but usage is low [15]. Thus the data generated from the Zlto digital wallet enables Youth to develop credit scores. By actively tracking both earning and spending habits, our team is developing a credit modeling system based on user data generated. We further collaborate with vendors and merchants to provide credit-worthy access to financial services or a line of credit.

The Zlto project will improve usage by monetizing (attributing digital rewards) a user's lifelong learning journey, enabling them to earn as they learn and work while subsequently spending their earnings on products and services with partner merchants. Storing their work assets using a blockchain-based mechanism allows them to build credibility of work in a secure environment. Each Zlto work asset will be fully under the control of the young person associated with it.

The Zlto platform also provides an opportunity to equip and develop youth as a pipeline for organizations to place youth in job and work opportunities. Zlto provides a strong value proposition to private sector institutions who generally lack incentives to work with youth without work experience.

The Zlto platform partners with social impact investors providing them direct access to validated impact through data tokens. This creates income generating opportunities to sustain the Zlto platform. The proof of work asset, held on the Zlto platform, acts as a virtual referral and proves each youths lifelong learning and work history. The Zlto platform also stimulates spending, as users redeem their earnings and rewards at merchants and informal traders in their local communities - enabling more economic value to these traders in a difficult economic climate.

IX. CONCLUSIONS AND FUTURE WORK

We present a system with the ability to generate an economy of informal work verified using a peer-to-peer reputation protocol.

The Zlto system derives value from informal work by rewarding positive behavior in a secure fashion and by providing partners with insights on emerging markets.

Zlto exists as a client-facing mobile application that has been deployed for 2 years, with 24,000 active users processing over 1 million transactions while logging 138,000 hours in their respective communities. We have partnered with the leading retailer in South Africa, Shoprite, to deliver 150,000 dollars of rewards to users.

The Zlto team is actively building tools in the blockchain space with confidence that they will be used by real people who depend on its availability and stakeholders who depend on its security. The conversion of value in the system proposed has helped us partner and earn funding

from governments, retail vendors, technology and financial institutions.

Equipped with the necessary social capital to drive the adoption of Zlto, we are aiming to use the system in new contexts for various applications. Examples include leveraging Zlto data in building a model of credit to serve as alternative data to financial institutions and also assisting unbanked refugee groups in the UK sustain themselves through informal community work.

The reputation flow mechanism can be enhanced by capturing metadata to evaluate the differences or similarities of the judges. If metadata composed of several variables with specified weights are used to build judge profiles, then one can build a general-purpose reputation flow based on metadata differences and similarities to extract more valuable data when maximizing a final reputation curve. Future investigation is necessary that may utilize machine learning to automate metadata distance calculations.

X. ACKNOWLEDGMENT

We acknowledge the wider RLabs team for partnering with us during the early stages of Zlto and supporting us in our efforts to develop the project, as well as all our partners and our local communities that used the platform.

We also thank the MasterCard Foundation, Naspers, Shoprite and Google for supporting this project through their grants and initiatives.

REFERENCES

- [1] Parker, Marlon, Wills, Julia and Wills, Gary (2013) RLabs: A South African Perspective on a Community-driven Approach to Community Information. [in special issue: Special Issue: CI and the Co-Creation of Innovation Journal of Community Informatics, 9 (3).
- [2] Atlam, Hany F., Alenezi, Ahmed, Alassafi, Madini O. and Wills, Gary (2018) Blockchain with Internet of Things: Benefits, challenges, and future directions. International Journal of Intelligent Systems and Applications, 10 (6), 40-48.
- [3] Alex Pazaitis, Primavera De Filippi Vasilis Kostakis (2017) Blockchain and value systems in the sharing economy: The illustrative case of Backfeed. Technological Forecasting and Social Change. Vol 125, December 2017, Pages 105-115
- [4] Backfeed: Technological Forecasting and Social Change Available at <http://backfeed.cc/assets/docs/TechnicalSummary.pdf> Access Date 6 February 2019
- [5] Paul Resnick, Ko Kuwabara, Richard Zeckhauser, Eric Friedman (2003) Reputation Systems, Communications of the ACM, December 2000, Vol. 43 No. 12, Pages 45-48
- [6] Audun Jsang, Roslan Ismail, Colin Boyd (2007) A survey of trust and reputation systems for online service provision. Decision Support Systems Volume 43, Issue 2, March 2007, Pages 618-644
- [7] A. Kosba, A. Miller, E. Shi, Z. Wen and C. Papamanthou, "Hawk: The Blockchain Model of Cryptography and Privacy-Preserving Smart Contracts," 2016 IEEE Symposium on Security and Privacy (SP), San Jose, CA, USA, 2016, pp. 839-858.
- [8] Vitalik Buterin (2015) On Public and Private Blockchains, Posted by on August 6, 2015 <https://blog.ethereum.org/2015/08/07/on-public-and-private-blockchains/> Access Date 6 February 2019
- [9] M. K. Reiter, J. J. Wylie, G. R. Goodson and G. R. Ganger, "Efficient Byzantine-Tolerant Erasure-Coded Storage," International Conference on Dependable Systems and Networks, 2004(DSN), Florence, Italy, 2004, pp. 135.
- [10] Jian Yin, Jean-Philippe Martin, Arun Venkataramani, Lorenzo Alvisi, Mike Dahlin (2003) Separating agreement from execution for byzantine fault tolerant services SOSP '03 Proceedings of the nineteenth ACM symposium on Operating systems principles Bolton Landing, NY, USA October 19 - 22, 2003, Pages 253-267

- [11] Tendermint Core (2018) <https://tendermint.com/docs/> Access Date 6 February 2019
- [12] Xiaoqi Li, Peng Jiang, Ting Chen, Xiapu Luo, Qiaoyan Wen (2017) A survey on the security of blockchain systems, *Future Generation Computer Systems*, Elsevier 23 August 2017 Harry Halpin, Marta Piekarska (2017) Introduction to Security and Privacy on the Blockchain. IEEE European Symposium on Security and Privacy Workshops (EuroSPW) April 26-28, 2017 in Paris, France
- [13] Creswell, J. (2013). *Research Design: Qualitative, Quantitative, and Mixed Methods Approaches*. Research design Qualitative quantitative and mixed methods approaches (4th ed.). SAGE Publications
- [14] Peters G.W., Panayi E. (2016) Understanding Modern Banking Ledgers Through Blockchain Technologies: Future of Transaction Processing and Smart Contracts on the Internet of Money. In: Tasca P., Aste T., Pelizzon L., Perony N. (eds) *Banking Beyond Banks and Money*. New Economic Windows. Springer, Cham
- [15] Manav Gupta (2017) *Blockchain for Dummies: IBM Limited Edition*, Wiley