

Fathom: A Protocol for Decentralized Assessment

Abstract

Fathom is a decentralized protocol to create meaningful credentials through the consensus of knowledge communities. Currently, most credentials rely on being backed by large institutions in order to be meaningful, which has led towards credentials that can be conveniently mass-customized and issued. This left most people in a situation where they can only reliably communicate a small part of their skill-set as dictated by the institutions. With fathom, anyone can create a credential for any kind of skill. In order to earn that credential, the fathom protocol defines an assessment-game, where the truthful evaluation of an applicant's skill is the schelling point. The economic incentivization of the fathom protocol makes users congregate around credentials that are well defined, in that they allow individuals to play the assessment game with a positive outcome. As such any community of any field can create their own credentials, use them in order to self-organize more effectively and its members can independently communicate their skills to the outside world.

Contents

1	Introduction	3
1.1	Problem Statement	3
1.2	Introducing Fathom	3
1.3	Outline	4
2	The Fathom Protocol	5
2.1	Architecture Overview	5
2.1.1	Concept	5
2.1.2	The Concept-Tree	6
2.1.3	The Assessment Game	6
2.2	Assessment Process	7
2.2.1	Setup	7
2.2.2	Assessment of the candidate	8
2.2.3	Committing a Score	9
2.2.4	Steal and Reveal	9
2.2.5	Determining the Outcome	9
2.2.6	Payout of Assessors	10

3	Security	12
3.1	Incentive Structure	12
3.1.1	Incentives for members to confirm or decline an assessment	12
3.1.2	Incentives for assessors to grade truthfully	12
3.2	Attack Vectors	13
3.2.1	Sybil Attacks	13
3.2.2	Simple Trolling	14
3.2.3	P + epsilon attack	14
4	Status & Roadmap	16
4.1	Implementation	16
4.1.1	Smart Contracts	16
4.1.2	Front-end	16
4.2	How to Contribute	16
4.3	Instantiating the network	17
5	Vision	17
6	Glossary	18
6.0.1	Concept	18
6.0.2	Parents	18
6.0.3	Assessor	18
6.0.4	Assessee	18
6.0.5	Tokens	18
6.0.6	Assessment	18
6.0.7	Weight	19
6.0.8	Table of Contents	19

1 Introduction

Society functions on knowing what people can do. Everybody needs to be able to communicate their skills to others in order to coordinate. This used to be a social process within a local community, but society has grown and largely outsourced that function to institutions. Today, people need to communicate a greater diversity of skills and experiences than ever before, over long timescales and across geographic, cultural and linguistic barriers. It is our belief that institutions cannot provide that service and will be unable to provide it in the future.

We seek to provide an alternative to institutional credentials by specifying a social protocol with economic incentives to allow knowledge communities to define their own standards and individuals to be assessed resulting in meaningful, verifiable and durable credentials

1.1 Problem Statement

The coupling of learning and assessment in current institutional models is unscalable and creates a set of perverse incentives for both educators and students. The bureaucracy of centralized institutions makes them resource intensive and slow to adapt to changes. As a consequence they are only able to offer a small set of experiences, defaulting to those that can be mass-produced.

Because communicating one's experiences is so essential in today's society, it is in an individual's best interest to actively mold their experiences towards what they can communicate, instead of towards what they can aspire to. Therefore, relying on institutions to be the arbiter of people's abilities has a chilling effect on societal progress.

1.2 Introducing Fathom

Fathom is a protocol to create and assess meaningful credentials through the consensus of knowledge communities.

It allows anyone to create a credential and anyone to be assessed in it. The core process involves a jury of randomly assembled assessors with relevant experience, as previously proven by the protocol, playing an 'assessment-game' in which they are economically incentivized such that an accurate assessment is the schelling point.

The protocol makes no assumptions about what is being assessed, instead it allows communities to form their own definitions and rules, to be then carried out collectively.

Implemented on a public blockchain, it will be possible to distribute the work necessary for assessments to scale far beyond what institutions are capable of.

Furthermore, blockchains can enable a truly inclusive, accessible and extensible credential-ecosystem, which is censor-resistant, durable and leaves individuals in full control of their identities: Able to learn skills and accumulate experiences towards their unique aims as well as to shape and strengthen the network by participating as an assessor.

1.3 Outline

The purpose of this document is two-fold: to provide a formal specification of the fathom-protocol, and to describe the infrastructure necessary to deploy it on a large scale. The second part will also lay out the launching process that we think will be favorable for a widespread adoption and for the overall integrity of the system. While the core protocol is unlikely to change much, the infrastructure around it is still being developed and as such this document is subject to change. We believe in transparency and working out in the open and therefore invite the reader to engage with us: By contributing ideas or comments, or by actively taking part in the development. Together, we seek to enable a wide variety of real-world applications based on this new layer of trust for digital social relationships.

< >

2 The Fathom Protocol

This page has four parts: First, an overview is given of the different elements that make up the fathom network. Second, the process of creating certificates and being assessed in them is presented step-by-step. The third part lays out the incentive structure that makes it economically disadvantageous to deviate from the protocol. The last part presents potential attack vectors and how they are mitigated.

2.1 Architecture Overview

This section will introduce the three main components of a fathom-assessment: First and second, a *concept*, representing an assessable quality and the *concept-tree*, which relates all concepts to each other. Third, *assessments* which draw individuals from the concepts to be assessors and, upon a positive assessment, add new individuals to them.

2.1.1 Concept

Concept is an umbrella term to capture any kind of skill, quality, piece of knowledge or fact that can be established about an individual. Therefore, each concept C has the following properties:

- **Parent Concepts:** This is the set of concepts P that C is a subset of. For example, the concept ‘Math’ could be a parent of ‘Linear Algebra’. If there is no suitable parent, a concept can be located beneath the ‘mew’-concept - a specially designated concept with no skills associated to it and with no parents.
- **Connection-strength(s):** For each parent $p \in P$, a concept C denotes a connection strength c_p from 0 to 1, specifying the degree of similarity or difference. Upon a positive assessment, c_p is used to determine how much someone who masters the child-concept should be considered competent in the parent-concept (see the incentive section for a more detailed accounting of how the c -value is determined).
- **Expiration time:** Concepts can specify expiration times e_c to reflect that some skills become outdated, need to be maintained, or change over time. An example would be concepts related to taxation-laws, which are changed on a relatively frequent basis and where false or outdated information can lead to significant losses. Members who have been assessed in a concept longer ago than specified by the expiration time do not lose their certificates, but can no longer take part in the process of assessing others.

- **Members:** a set of individuals who have passed an assessment in the concept in question or in one of its children
- **Component Weights:** For each member the concept stores a set of component weights, a positive integer and date, corresponding to assessments in the concept or its children. The sum of valid component weights, those which do not have a date longer ago than the concept's expiration time e_c , is used to probabilistically call a member to act as assessors in assessments¹.

2.1.2 The Concept-Tree

As all concepts have at least one parent, the entirety of all concepts forms a tree. The 'new'-concept is at the root and as one moves down concepts become more specific. The weighted parent-child connection between concepts is also reflected in the weights of their members: Individuals who have been assessed in a child concept are also added as members of the parent concept(s), with their weight(s) reduced by the connection strength. The set of individuals in a child concept can therefore be seen as a specific subcommunity of the parent-community. Individuals are therefore members of many more concepts than those they have been assessed in, although those will be the ones where their weight will be highest.

This nesting of knowledge-communities is valuable when sampling members to create a pool of potential assessors. If a concept has not enough members for the pool they can be drawn from the parent concepts. The members of a parent concept will have related experience (as it has been populated by the concept and siblings), and hence have a higher probability of being able to assess the concept in question.

2.1.3 The Assessment Game

An assessment is the process by which a jury of qualified individuals (assessors) decides whether or not some candidate (assessee) fulfills the necessary conditions to become a member of a concept. When initiating an assessment in a concept, the assessee decides how many assessors they want and how much they are willing to pay to each one of them. That offer is forwarded to potential assessors (see setup for drawing specifics) who must stake the offered amount in order to accept. Thus, a market forms around assessments, allowing the system to scale from easy to assess, and hence cheap, concepts, to more involved, complicated, and hence expensive ones.

Upon completion of the assessment, assessors are paid the price offered by the assessee and a proportion of their stake if they come to consensus around the

¹Although it's possible to repeat an assessment, only the result of the most recent assessment will be taken into account for the weight.

applicants skill. The proportion of the stake being paid back is proportional to the assessor's proximity to the average score of the biggest cluster of scores and will also be added a portion of the stake of dissenting assessors, should there be any (see payout for details).

Also, the mechanism by which assessors log in their scores is designed such that colluding assessors can double cross each other, thereby creating a coordination problem in an adversarial environment, where the only point of coordination (schelling point) left is a truthful assessment. In case the majority vote of the assessors is positive, the assessed candidate will get i) a score in the assessed concept, similar to a grade in university or school, ii) a weight in the concept and iii) a weight in all parent-concepts, proportionally reduced by their respective connections strenghts.

2.2 Assessment Process

A fathom assessment goes through five phases: A setup phase, where the assessors are called from the concept tree, the assessment, where the assessors determine the assessee's skill, commit- and reveal-phases, where the assessors log in their score and, at last, the calculation of the result.

For each phase this section is gonna depict the choices of the involved participants, their interactions and what happens if they deviate from the protocol.

2.2.1 Setup

2.2.1.1 Creating the assessment:

Wanting to be certified in a concept C , the assessee needs to specify the following parameters:

1. A timeperiod during which they would like the assessment to start and end (latest start- and end-time).
2. The number of assessors N_a to be assessed by. While there is a minimum number of five assessors to guarantee a fair voting, the assessee might want to be assessed by a bigger number in order to receive a higher weight and higher chances to become assessors themselves (given that they end the assessment with a passing score).
3. The price $cost_a$ that each assessor will be paid.

2.2.1.2 Calling assessors from the concept-tree:

A pool of potential assessors is created by probabilistically drawing members from the concept and its parents. The selection of potential assessors happens according to a tournament-selection of size 2, starting at the assessed concept:

- Two members are picked at random and their weights are compared.
- Only if the weight of the first member is bigger than the second one, the first is added to the pool, given that they are not already in there.

Thus it is assured that each member has a chance of being called as assessor (the chance of being drawn first), whilst giving a higher chance to members with higher weights).

To make it hard to predict who will be in the pool, no more than half of the members of each concept can be called as assessors. Therefore, the maximum number of tournaments per concept c is limited by the number of members in the concept. Specifically Y is:

$$Y = \min(N_{req}, \frac{m_c}{2})$$

,with N_{req} being the number of required assessors and m_c the number of members in the concept c .

After Y attempts, (on average resulting in $Y/2$ members being added) this selection process is recursively repeated for each of the n_{p_c} 's parent concepts of c , with N_{req} being set to N_{req}/n_{p_c} . The recursive calling stops if a pool size of 200 times the desired number of assessors or the mew-concept is reached.

2.2.1.3 Assessors confirm by staking:

Each assessor that is being called, can decide to participate in it by staking the offered price. Once the desired number of assessors has confirmed, the assessment moves to the next stage. Assessors from the pool self-select whether they think would be competent judges on concept in question. If so, they signal their intent to participate by staking the offered price. More considerations why assessors would or wouldn't want to confirm are elaborated in the incentive section. If not enough assessors can be found before the desired start-time of the assessment, the Assessment is cancelled and everybody who deposited collateral is refunded.

2.2.2 Assessment of the candidate

In a fathom-assessment there is no notion or form what constitutes a test and the form or procedure of how candidates are evaluated is left to each individual assessor. Ultimately, assessors express their verdict of assessee's skill as a number on a scale (e.g. between 0 and 100) - with everything above half being considered a passing score.

Yet, what exactly defines a failing, passing or barely passing assessment can be different for each concept as well and should be agreed upon by the community. Moreover, the assessment could also be the place to put up some sybil protection

mechanism in the form of extra requirements that make it hard to repeat an assessment (see sybil-attack for more details on how this could work).

2.2.3 Committing a Score

Sending in a score follows the commit-reveal procedure common in blockchain applications. Assessors signal that they have decided on a score by concatenating it with a secret element, also referred to as ‘salt’ and submitting its hashed value ($\text{hash}=\text{sha3}(\text{score}+\text{salt})$).

If any assessor fails to commit a score before the assessment ends their stake is being burned. If, as a consequence, less assessors than would be required for the minimum size of a viable assessment have committed, the assessment is cancelled and everyone is being refunded. Otherwise, the assessment progresses to the next stage.

2.2.4 Steal and Reveal

To end the assessment, the assessors reveal their verdict by submitting their score and salt separately. Any assessor (or external person) who knows about another assessor’s score and salt, can do so as well, thereby stealing half of the assessor’s stake, burning the rest and eliminating him/her from the assessment game. This prevents the assessors from credibly guaranteeing each other their commitment to logging in a specific score, thus making it harder to collude.

While stealing is possible at all times after an assessor has committed (even if others have not yet), revealing will only be possible after all assessors have committed and a buffer period of 12 hours has elapsed. The buffer ensures that there is time to challenge someone’s commit, even if they waited until moments before the end of the assessment period to send it in. Should any assessor fail to reveal, their stake is burned and they are eliminated from the assessment game². If the number of assessors decreases below the necessary minimum, the rest of the participants is being refunded and the assessment ends without a score.

2.2.5 Determining the Outcome

In order for an assessment to result in a final score, one score must be in consensus with enough other scores to form a 51% majority. Two scores are considered to be in consensus if their difference is less than the *consensus-distance* ϕ . If such a score s_{origin} exists, the final score s_{final} is computed as the average of all scores that are in consensus with s_{origin} .

²This should be unlikely, as at that point assessors have nothing to lose but rewards to gain.

Should there be two scores with majorities of equal sizes, the one that will result in a lower final score wins. If there is no point of consensus, the assessment is considered invalid and all stakes are burned. Otherwise the assessors' payments will be computed as described in the following section.

Also, in case the result s_{final} is a passing score, the assessee is registered as new member of the concept with a weight $w_i = s_{final} * N_{in}$. They get a weight in all parents and grandparent-concepts, reduced by the respective connection weights until their weight is below a certain threshold.

2.2.6 Payout of Assessors

Payments to assessors consist of two parts, their returned stake and the assessee's reward. Both are attributed differently, depending on whether or not the assessor is inside out outside the majority cluster of winning assessors.

Therefore, an assessor i 's distance $dist_i$ from the final score s_{final} is measured against the consensus range ϕ :

$$dist_i = \frac{|s_i - s_{final}|}{\phi}$$

Outside assessors ($1 < dist_i < 2$) only get back a part of their stake, reduced linearly in relation to their distance from the final cluster. Thus, an outside assessor j 's payout is computed as:

$$payout_{out_j} = stake_j * \frac{(2 - dist_j)}{2}$$

Inside assessors ($dist_i \leq 1$) get back their entire stake, any stake that is not returned to outside assessors (distributed equally among them) and a share of the assessee's reward, proportional to their proximity to the final score:

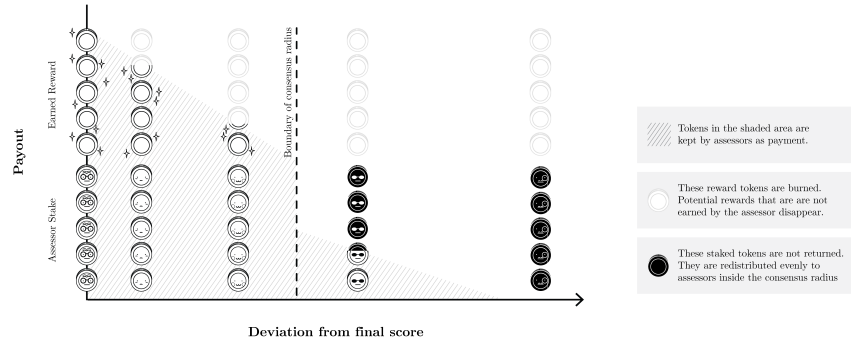
$$payout_{in_i} = stake_i + \frac{1}{N_{in}} * \sum_j^{N_{out}} (stake_j - payout_{out_j}) + r_a * (1 - dist_i)$$

,with N_{out} and N_{in} denoting the number of assessors outside and inside the winning cluster and r_a being the reward of the assessee.

Thus, the best case scenario for an assessor is to be inside the winning cluster, close to the final score, with a large minority outside of it.

Any payout that is not returned to the assessors will be burned (if it were redistributed, assessors could collude to cover a range of scores and redistribute amongst them without any loss).

This figure summarizes the payout mechanism in a single graph:



< >

3 Security

The fathom protocol derives security from both its technical implementation and incentive structure. There are specific attack vectors that it mitigates against as well as some areas for further research.

3.1 Incentive Structure

While traditional credentials are meaningful because they are backed by reputable institutions, a fathom credential is meaningful because it is the result of many individuals having undergone a financial risk in the assessment game in order to create it. This section will lay out the decisions fathom users face when participating in that game, as well as the economic risks associated to them. Specifically, we show i) when assessors are likely to participate in an assessment in the first place and ii) why they can not collude with each other in order to shortcut the work associated with a truthful assessment.

3.1.1 Incentives for members to confirm or decline an assessment

This case is especially relevant for creating new concepts - as those will be initially empty and rely on members of its parent to participate in as assessors. A member of a concept who has received an offer to be an assessors will consider whether...

- They feel competent enough in their abilities to come to the same conclusion as a group of other, randomly selected assessors that are confident in *their* abilities.
- The concept in question is well enough defined so that assessors with similar impressions of the assessee's skill will be able to translate these into similar scores.
- Whether its connection strength is appropriately chosen. A passing score in the concept in question will result in the assessee being added to the assessors concept, and therefore affect the assessor's chances in future assessment games, as well as the quality of the community. If the connection strength would be high but the concept comparatively easy to earn, the assessor's community could be flooded and their chances of being called as an assessor reduced.

3.1.2 Incentives for assessors to grade truthfully

As truthfully assessing someone requires effort and the assessors payout is pegged to their alignment with each other, there is a motivation for them to

collude, e.g. by agreeing ahead of time which score to commit. The creation of an adversarial environment between assessors is thus vital for the protocol to function as intended. Therefore, several mechanisms are put in place: First, assessors are paid out more if there are dissenting assessors (see Figure 1) Consequently, any assessor taking part in a collusion of X assessors, must be afraid that they will be double crossed by a subgroup of more than $X/2$ assessors. These are motivated to do so because they would be rewarded with part of the crossed assessors stake. Moreover, it is not possible for assessors to credibly prove to another assessor that they have actually committed to a collusion and logged in a previously agreed-upon score. In order to do so, the proving party would have to reveal their score and salt to the other assessors. Yet with this information, the other assessors could simply steal the assessors stake, which would eliminate the former assessor from the assessment and directly transfer the half of the revealed assessor's stake to the revealer.

3.2 Attack Vectors

This section will outline some of the general classes of attacks against the protocol and a subjective view of their complexity, severity and to what degree they are considered to be mitigated.

3.2.1 Sybil Attacks

In a sybil attack, the attacker creates many false identities and then uses them to subvert the system, e.g. by controlling most of the identities in a concept, giving him control over who will be accepted and the ability to create assessments for himself in order to steal the stakes of other assessors.

To set up such an attack the attacker would, instead of being assessed by many assessors in one assessment, create multiple assessments with fewer assessors. This would be the same amount of work but result in four identities in the concept. Repeating the procedure, the attacker could count on some of his identities being called as assessors in which case the subsequent repetitions would become cheaper and less time-consuming until they have the majority in the concept or are called multiple times as assessor such that they can set up a 51% attack on individual assessments. In such a scenario, the attacker could control the outcome of the assessment and steal the stake of the other assessors.

Severity of attack: While a sybil attack does cost a fair amount of money to set up, the potential benefits are big enough to incentivize a try. As a compromised concept can potentially 'poison' its parent concepts as well and thus potentially effect the entire tree, we consider it to very severe.

Complexity: While a sybil attack is fairly complex, it can be effectuated by a single attacker, which is why it would be careless to assume that the degree of complexity will be a deterrent factor.

Degree of Protection: One possible mitigation that is not yet part of the protocol, will be to split the certificate and the right to be an assessor in two separate assessments. While this does not address the fundamental issue, it makes it easier for the sybil-protection measures to be integrated into the assessment process. For example, the assessment to become an assessor could ask the to-be-assessors for some piece of their own work or something that is new and can not be readily found on the internet as would be the case with the mere knowledge or skill required in the concept.

3.2.2 Simple Trolling

A troll, for arbitrary reasons, might try to poison the fathom network by creating a bunch of bogus assessments or concepts or by behaving irrationally while being an assessor. In all cases, such behavior is expensive and ineffective, as his stakes are burned (when not following through with an assessment) or redistributed to others (when logging in bogus scores).

Bogus concepts will simply incur costs on the troll and be filtered out by assessors (see incentives). Creating bogus assessments as assessee will be even more costly (transaction costs and the fees for assessors). The worst effect a troll can have is to become an assessors and to prematurely end the assessment, if as a consequence of their behavior, its size is reduced below the minimum of five. In that case all other participants will be refunded, though.

Complexity: Behaving irrationally is simple and so is attacking the system this way.

Severity: With no financial costs to other participants these kinds of attacks are not considered severe. An exception might be the creation of concepts, which if done by a well-resourced attacker, amounts to spamming the system.

Degree of protection: We consider simple trolling to be sufficiently discouraged because of the associated costs. If such behavior would be escalated into a spam-attack of greater proportions, the degree of protection will depend on the users or the fathom frontends ability to filter concepts and assessments by meaningful criteria.

3.2.3 P + epsilon attack

In a P + epsilon attack, the attacker circumvents the incentivization by creating a mechanism that others can trust in because it gives them a credible guarantee about the attacker's behavior. While this would have been difficult in a pre-blockchain era, smart contracts are nearly ideally suited to implement such mechanisms.

The attack works like this: In a schelling-point game, the assessors are being paid out the same amount P . regardless of the result (option A, B, C or any

other...). The attacker, let's say wanting to push for a certain option A, will credibly guarantee anyone voting for A that he will be paid $P+\epsilon$, if they vote A and the majority doesn't. Assuming a system that is not dominated by altruistic actors, voting A is now the game-theoretically best option (guaranteed maximal payout). Therefore, the majority will vote A and the attacker will have taken over the mechanism - at zero cost.

Although there exist some protection mechanisms that can increase the attackers risk (size of the needed bribe) and some counter-coordination mechanisms that come close to defeating such an attack, there is currently no guaranteed countermeasure.

Complexity: As the crucial element of this attack is the mechanism by which the attacker commits to his intention to paying out in case the bribed voter is not in the majority, the complexity is proportional to the difficulty of construing such a mechanism. In the case of fathom, the difficulty to reconstruct the relevant information (did an assessor really vote for the desired option A?). Currently, this is rather simple, so setting up this attack would not be very complex.

Severity: As this attack can disrupt the system at potentially zero-cost, we consider it to be very severe.

Degree of Protection: As of right now, the protocol is not protected against such measurs. Future versions of it could implement some more complicated schemes in order to keep the scores of individual assessors secret and make it harder to retrieve the individual assessors' scores.

< >

4 Status & Roadmap

This section will describe the current state of the different components the fathom-team is developing, how one can contribute and what we plan to build in the future. It also describes our plans to instantiate and initially support the network.

4.1 Implementation

4.1.1 Smart Contracts

The assessment process of the fathom protocol as described in this document is fully implemented in solidity. The code, a testsuite and some supplementary simulations are in our gitlab repository.

4.1.2 Front-end

We are also working on various user-facing applications to interact with the fathom protocol. Our first project was an arbitrary chat room application allowing the creation of social environments based on on-chain data, which was a prototype for a platform for concept communities in fathom, as well as for a space to conduct assessments. The code for this project is in this repository. While we still subscribe to the idea of permissioned chatrooms, work is currently underway to refactor the work into different, more abstract parts that can be reused in a modular fashion. The best way to stay up-to-date is to follow in our blog or in the org-repo in which we keep an updated roadmap.

4.2 How to Contribute

It's our philosophy to work out in the open as much as possible, so this section will lay out the documents we are using to structure our work and the process by which we invite volunteers to contribute.

All parts of our project are in our GitLab repository. The documents in the org-repo are a good place to start.

Also, we publish an approximately bi-weekly progress report on our blog to keep you updated about the development and about ways to engage with us and to help us prepare fathom for a successful launch.

We appreciate any interest in fathom and seek input from individuals with all kinds of educational- or cryptoeconomic-related backgrounds. Feel free to reach out to us at *contact@fathom.network* or to create issues in our repo.

4.3 Instantiating the network

In order to bootstrap the network, we will run a decentralized learning bootcamp, aka the *fathom playground*. In brief, it will bring together a group of motivated learners that will receive the initial batch of network-tokens and use them to create credentials in order to coordinate their own learning. A detailed explanation of the playground will soon be published on our website and our blog.

5 Vision

We believe that a participatory protocol tied directly to the communities practicing skills and defining ideas will diminish the gap between credentials that can be credibly assessed and issued and the wide variety of skills and abilities that people are capable of.

Tying economic incentives to this social process and ontology, such that they are both visible to everybody and aligned amongst all those participating, allows for fathom-credentials to be trustworthy and transparent.

By distributing the work required to communities, allows the system to scale and be accessible to anyone, no matter their previous records, achievements or socio-economic circumstances.

We believe that through these traits fathom enables a world where people are free to shape their own experiences, communicate them to others, and organize to achieve shared ambitions.

<

6 Glossary

6.0.1 Concept

Represents the shared knowledge, skillset, or other attribute that the users who have attained the concept have in common. Examples could range from Calculus proficiency to English fluency to marathon completion.

6.0.2 Parents

A parent of a concept is one more general than it. For example Calculus is a child of Math and a parent of Differentiation. When an individual is successfully earned a weight in a concept they also own a lower weight in its parents.

6.0.3 Assessor

A member of an Assessment who has previously acquired the Concept the Assessee(s) are trying to attain.

6.0.4 Assessee

The subject and initiator of an assessment, who wants to become a member of a concept.

6.0.5 Tokens

An ERC20 token that is used to pay for assessments, and is earned by assessors. It is created upon successful assessments where more is paid out than what is paid by the assessee.

6.0.6 Assessment

The process by which an Assessee/Assesseees attain Concepts and Assessors attain tokens. The Assessee/Assesseees initiate an Assessment by paying a number of Tokens and inputting a Concept. A number of Assessors, who have previously been favourably assessed on the Concept, are put into communication with the Assessee(s). Each Assessor then scores the Assessee(s) from -100 to 100 and the scores are averaged. Assessors who score near the average are rewarded with a greater Token payout. If the score is positive the Assessee/Assesseees attain the Concept.

6.0.7 Weight

Upon a successful assessment an individual earns a weight in a concept. This is calculated as a function of their score, the number of assessors in the largest cluster, and the time the assessment was taken.

6.0.8 Table of Contents

1. Introduction
 - 1.1 Problem Statement
 - 1.2 Introducing Fathom
 - 1.3 Outline 2. The Fathom Protocol
 - 2.1 Architecture Overview
 - 2.1.1 Concept
 - 2.1.2 The Concept Tree
 - 2.1.3 The Assessment Game
 - 2.2 Assessment Process
 - 2.2.1 Setup
 - 2.2.2 Assessment of the Candidate
 - 2.2.3 Committing a Score
 - 2.2.4 Steal and Reveal
 - 2.2.5 Determining the Outcome
 - 2.2.6 Payout of Assessors
 3. Security
 - 3.1 Incentive Structure
 - 3.1.1 Incentives to Confirm or Decline an Assessment
 - 3.1.2 Incentives for Assessors to Grade Truthfully
 - 3.1 Attack Vectors
 - 3.2.1 Sybil Attacks
 - 3.2.2 Simple Trolling
 - 3.2.2 P+epsilon Attack
 4. Vision
- >