Reputation scoring

Storj DS&A

1 Model summary

We present a model for assigning reputation scores to storage nodes, which can be used to determine node reliability as it pertains to uptime and file audit successes. The score depends on actual node behavior, and is tuneable so that a variety of potential use-cases are addressed. We base our model on the Beta reputation system, as outlined in The Beta Reputation System[1] by Jøsang and Ismail. While we focus our attention on assigning a score to measure both uptime and file audit successes, we remark that the general model has many tuneable parameters and may be tuned to accommodate many of the needs our network may encounter. Focusing on audits and uptime, we present a more concrete implementation for which many of these parameters are fixed. We leave only two parameters free, making the problem of reputation score assignment tractable while still allowing for a degree of flexibility in the score assignment model.

The two aspects of the model that we currently find most attractive are the single value feedback value v enabled by a transaction normalization weight w, and the concept of a forgetting factor λ. The single value feedback normalization weight allows for reputation to be updated by feedback given as a single value v. For our purposes, we fix $w = 1$, and choose transaction rating $v \in [-1, 1] \setminus \{0\}$ as a function of time spent on the network (how long the node has been a part of the network) and a global tuneable scaling parameter c_v which controls the rate at which a satellite builds confidence in new nodes. We remark that values of $v \in [-1, 0)$ are used for negative outcomes (failed uptime or audit checks) and values of $v \in (0, 1]$ are used for positive outcomes, with values of v starting close to 0 and tending toward either extreme (± 1) the longer a node is on the network. The forgetting factor λ is a scalar in $[0, 1]$ that determines how much past information is kept in determining current reputation score. Mathematically, this is achieved by discounting prior reputation scores with a geometrically decaying factor with base λ. Practically, choosing $\lambda = 0$ means that only the most recent transaction rating matters, while choosing $\lambda = 1$ ensures that all prior transaction ratings are used to determine the current reputation score. For the initial scoring model, we fix $\lambda = 0.99$ while remarking that this parameter should be tuned depending on the desired degree of past scores we want the current reputation score to contain.

[1]To briefly summarize this reputation scoring model: the reputation score is taken to be the expectation of the Beta distribution determined by parameters resulting from positive and negative transactions. The expectation will fall between 0 and 1 and may be taken as a measure summarizing the degree of confidence we have in a node's reliability based on prior node behavior.

2 Selected simulations

We plot some figures to illustrate the reputation scores and reputation score histories arising from different choices of the parameters that are tuneable in the model. To read the figures, we recommend estimating that a node conducts 1,000 scorable transactions per month; with this estimate, we produce figures that illustrate the reputation score history developing over the course of 12 months for a given node. In this simulation, we assume that a node responds successfully 90% of the time on average.

2.1 Varying c_v

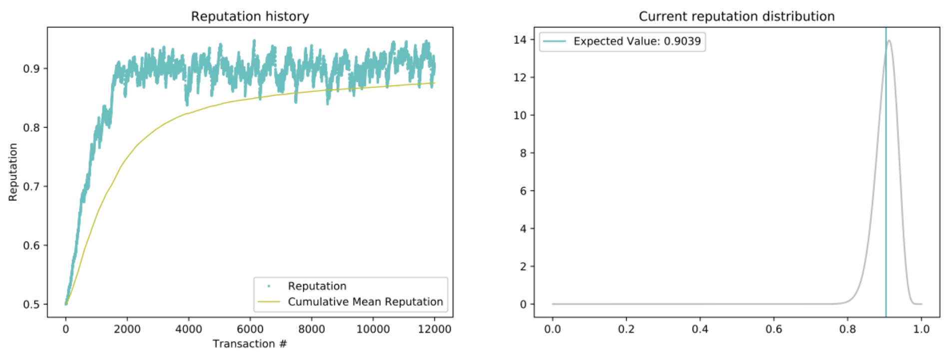

Figure 1: Score history with $\lambda = 0.99$ and $c_v = 1000$.

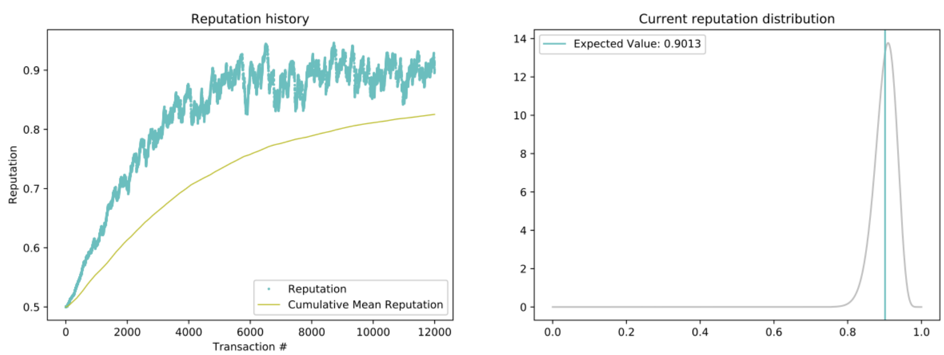

Figure 2: Score history with $\lambda = 0.99$ and $c_v = 3000$.

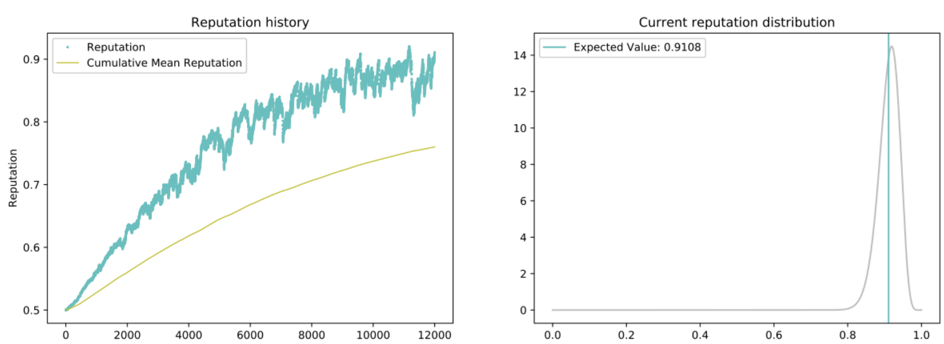

Figure 3: Score history with $\lambda = 0.99$ and $c_v = 6000$.

2.2 Varying λ

Here we plot figures illustrating how different values of the forgetting factor λ affect the reputation score when a node changes its behavior after a given amount of time. In this simulation, we suppose a node changes its behavior from returning a success 90% of the time (on average) to returning a success 80% of the time (on average), with the change occuring after 12 months on the network. We observe that the reputation score is subject to more extreme fluctuations the farther λ is from 1, which is also reflected in the variance of the reputation distribution.

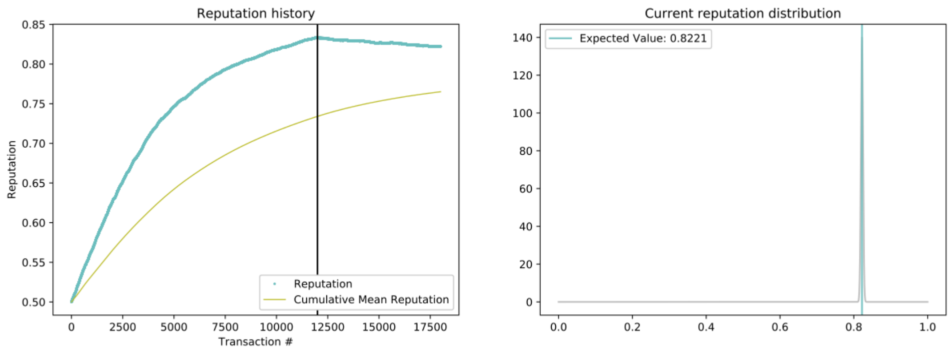

Figure 4: Score history with $\lambda = 1$ and $c_v = 3000$.

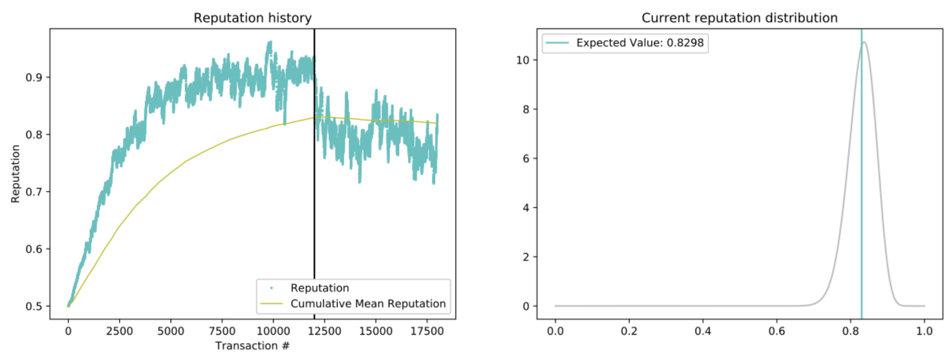

Figure 5: Score history with $\lambda = 0.99$ and $c_v = 3000$.

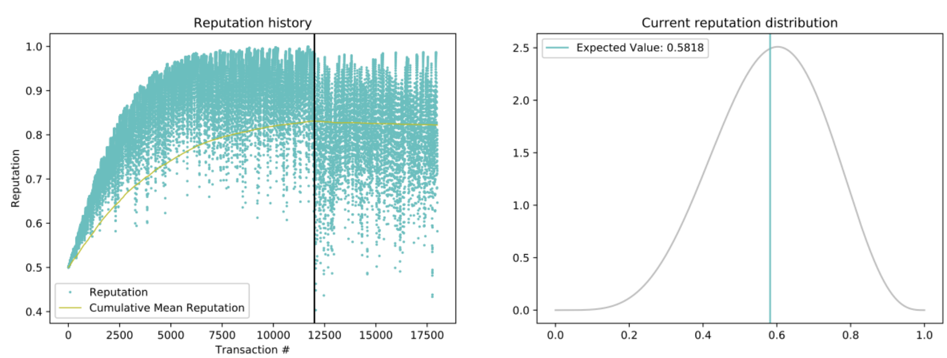

Figure 6: Score history with $\lambda = 0.9$ and $c_v = 3000$.