Gridcoin Blue Paper Section 1: Expected Time to Stake and Net Weight

James C. Owens, ILikeChocolate, Jacob T. Held

October 22, 2018

Abstract **– This document provides an overview of Gridcoin staking as implemented in the version 8 and above protocol (V8+) and a statistical analysis of the expected time to stake. The security aspects of staking and specific algorithms involved will be covered in a future document. While this document specifically covers staking from the perspective of Gridcoin, the analysis is also adaptable to other proof-of-stake coins that use unspent transaction output value to stake.**

1 Overview

Gridcoin (GRC)[1] is a decentralized proof-of-stake cryptocurrency utilizing proof of stake version 2 (PoSv2) that incentivizes participation in distributed scientific computing through the Berkeley Open Infrastructure for Network Computing (BOINC).[2] As of the V8+ protocol, Gridcoin follows PoSv2 with a target block spacing of 90 seconds.[3]

From a high-level perspective, the staking process for Gridcoin is a series of Bernoulli trials ("dice rolls" with only two possible outcomes: "success" or "failure") of staking probability p_{stake} for each eligible unspent transaction output (UTXO) at a controlled rate. Each of these staking trials is an independent event. Therefore, the probability distribution of the number of staking trials required for a successful stake is a compound geometric distribution with a p_{stake} for each UTXO.[4] Hereinafter, the staking probability of an individual UTXO i will be referred to as p_i, which is proportional to the UTXO's value, v_i.

A wallet's UTXOs participate in the staking process according to the following criteria:

1. The wallet is online.
2. The wallet is unlocked for staking (if encrypted).
3. The UTXO age (i.e. time since last transaction) is greater than or equal to 16 hours.
4. The UTXO value is greater than or equal to 1/80 GRC.
5. If a balance reserve for spending is set, then only the UTXOs whose values are less than or equal to the balance minus reserve can be staked.

During the staking process, the stake miner loop in miner.cpp generates a staking hash for each UTXO, followed by a sleep of 8 seconds in a continuous loop as long as there are qualified coins to stake. Additionally, the staking transaction is subject to a 16 second granularity time mask, which renders stakes within the same 16 second masking period degenerate, resulting in distinguishable staking trials every 16 seconds in the wallet.

For the purposes of this document, the cryptographic details of stake hashes and block validation will be put aside for future discussion. Ignoring these details, the stake hash can be treated as a random 256 bit number with values from 0 to $2^{256} - 1$. The stake miner performs the staking trial for each UTXO by comparing two numbers: the staking hash and the weighted target for that UTXO. If the staking hash is less than or equal to the weighted target for that UTXO, then the stake will be successful for that UTXO if it passes validation by the staking node and other nodes in the network. This inequality is the "coin toss" and forms the heart of the staking process.

[1] Gridcoin Github repository: https://github.com/gridcoin-community/Gridcoin-Research
[2] Gridcoin Whitepaper: https://gridcoin.us/assets/img/whitepaper.pdf
[3] Adapted from Blackcoin: https://blackcoin.org/blackcoin-pos-protocol-v2-whitepaper.pdf
[4] See Appendix A for more information on UTXOs.

2 Basic Staking Probability

2.1 Difficulty and Staking

The Unit Difficulty Target T_U defines the base staking target and is approximately 2.70×10^{67}.[5] The unweighted target of the current block, T, is adjusted every block by the retargeting algorithm (Section 2.3). The staking difficulty D depends on T_U and T:

$$D = \frac{T_U}{T} \tag{1}$$

The maximum possible target T_L for Gridcoin is $2^{(256-20)} \approx 1.10 \times 10^{71}$, called the proof of stake limit; therefore, the minimum D in Gridcoin (corresponding to $T = T_L$) is approximately 2.44×10^{-4}.

In PoSv2, as implemented with Gridcoin V8+, the Stake Target T_S is the maximum possible hash value that will result in a successful stake (i.e. hash values $\in [0, T_S]$ will stake). T_S is defined as $T \cdot w_i$, where $w_i = floor(80 \cdot v_i)$ is the coin weight[6], and v_i is the value in GRC of the i^{th} UTXO.[7] Thus, the probability p_i of UTXO i staking in a staking trial is

$$p_i = \frac{T_S}{H_{max}} = \frac{T \cdot w_i}{H_{max}} = \frac{T}{H_{max}} \cdot \text{floor}(80 \cdot v_i) \tag{2}$$

where H_{max} is the maximum possible hash value: $2^{256} = 256^{32} \approx 1.16 \times 10^{77}$.

Substituting Eqn. 1 into Eqn. 2 to include only values for constants readily available in the wallet UI, p_i becomes

$$p_i = \frac{T_U}{D \cdot H_{max}} \cdot \text{floor}(80 \cdot v_i) \tag{3}$$

Thus, for a wallet of n UTXOs, the total staking probability p_t of the wallet ignoring cooldown is

$$\begin{aligned} p_t = \sum_{i=1}^{n} p_i &= \frac{T_U}{D \cdot H_{max}} \cdot \sum_{i=1}^{n} \text{floor}(80 \cdot v_i) \\ &= \frac{T_U}{D \cdot H_{max}} \cdot w_t \end{aligned} \tag{4}$$

where w_t is the total weight of all eligible UTXOs in the wallet.

Specifically, w_t is:

$$w_t = \sum_{i=1}^{n} floor(80 \cdot v_i) \tag{5}$$

In the case where a reserved balance is specified in the wallet to be spendable at any time, the reserved coins do not participate in staking, limiting v_i to:

$$v_i \leq B - R \tag{6}$$

where B is the wallet balance and R is the reserved balance.

In practice, because UTXOs can be reorganized to ensure that the floor function has a negligible impact, and the sum of UTXOs with less than $\frac{1}{80}$ GRC have a negligible contribution to the total staking balance, v_t, this sum can be approximated with little error. Thus

$$v_t = \sum_{i=1}^{n} v_i = B \approx \frac{w_t}{80} \tag{7}$$

and

$$p_t = \frac{T_U \cdot 80}{H_{max}} \cdot \frac{v_t}{D} \tag{8}$$

2.2 Estimated Time to Stake (ETTS) by the Expected Value Method

The expected value of the random variable X representing the number of Bernoulli trials required for success is the reciprocal of the trial probability because the number of trials required for success is geometrically distributed:

$$E(X) = \frac{1}{p_t} \tag{9}$$

Distinct staking trials occur every 16 seconds, so the ETTS $\hat{t}_s$, measured in days, is

$$\hat{t}_s = \frac{16}{3600 \cdot 24} \cdot E(X) = \frac{16}{3600 \cdot 24} \cdot \frac{1}{p_t} \tag{10}$$

[5] The Unit Difficulty Target T_U and Current Unweighted Target T are expressed as a 32 bit compact representation (known as "nBits") of the 256 bit number with value 0x1D00FFFF for T_U. The first two hex digits of this number are the exponent and the next six are the coefficient, so $1D00FFFF = 0x00FFFF \cdot 256^{0x1D-3} = (256^2 - 1) \cdot 256^{26}$. The choice of this value for T_U is largely historical, with its roots in Bitcoin, upon which much of the Gridcoin code is based, where this value is the maximum target (minimum difficulty).

[6] In the Gridcoin code it is implemented as int64 division of v_i, which is internally represented in "Halfords" (1/100000000 GRC), by 1250000, which achieves the equivalent of the floor function as presented here.

[7] As a result of this rounding, any UTXO i smaller than $\frac{1}{80}$ GRC will have $w_i = 0$ and cannot stake.

Substituting Eqn. 8 for p_t yields:

$$\hat{t}_s = \frac{16}{3600 \cdot 24} \cdot \frac{H_{max}}{T_U \cdot 80} \cdot \frac{D}{v_t} \tag{11}$$

Let

$$G = \frac{16}{3600 \cdot 24} \cdot \frac{H_{max}}{T_U \cdot 80} \simeq 9942.2056 \approx 10000 \text{ days} \tag{12}$$

Then

$$\hat{t}_s = G \cdot \frac{D}{v_t} \approx 10000 \cdot \frac{D}{v_t} \text{ days} \tag{13}$$

This provides a simple and useful method of estimating the time to stake for a given wallet balance.

2.3 Retargeting Algorithm

In response to variations in block spacing caused by staking wallets coming online or going offline and the stochastic nature of staking, the wallet adjusts the unweighted target of the current block j, T_j, to maintain an average block spacing of $t_B = 90$ seconds. This retargeting adjustment is performed for each block and is a function of the target of the previous block, T_{j-1}, and the time elapsed between the two previous blocks, t_{j-1}. To avoid large, sudden fluctuations in T_j, the wallet defines a target timespan of $t_T = 960$ seconds as a smoothing constant. In the wallet, t_T is converted into an interval I in blocks, so $I = \frac{t_T}{t_B} = \frac{960}{90} = 10\frac{2}{3}$ blocks. Then

$$\begin{aligned} T_j &= T_{j-1} \cdot \frac{(I-1) \cdot t_B + 2 \cdot t_{j-1}}{(I+1) \cdot t_B} \\ &= T_{j-1} \cdot \frac{t_T - t_B + 2 \cdot t_{j-1}}{t_T + t_B} \end{aligned} \tag{14}$$

Rearranging the first line of 14,

$$T_j = T_{j-1} \cdot \left(\frac{I-1}{I+1} + \frac{2}{I+1} \cdot \frac{t_{j-1}}{t_B} \right) \tag{15}$$

Thus, the wallet adjusts the target, and consequently the difficulty, as follows:

If	Then	Result
$t_{j-1} < t_B$	$T_j < T_{j-1}$	Difficulty Increases
$t_{j-1} = t_B$	$T_j = T_{j-1}$	No Change
$t_{j-1} > t_B$	$T_j > T_{j-1}$	Difficulty Decreases

This is true for all blocks, and so staking time between blocks will act as a damped harmonic oscillator, fluctuating and eventually converging to an equilibrium of t_B.

Due to the smoothing built into this algorithm, estimations of the time to stake and the network volume lag slightly behind changes in active nodes in the network. Additionally, the v_{net} that is displayed in the wallet is averaged over the previous 40 blocks (1 hour) to further smooth out short term fluctuations.[8]

2.4 Net Weight vs. Difficulty

All UTXOs online and not cooling down undergo independent staking trials every 16 seconds. Thus the entire network's ETTS follows Eqn. 13, with v_t in this case being equal to the net weight v_{net}, the sum of all actively staking UTXOs across the network. As described in Section 2.3, to maintain an average block spacing of $t_B = 90$ seconds (= 40 blocks/hour = 960 blocks/day), the network adjusts D by changing T_C every block. Consequently, the relationship between v_{net} and D is obtained by setting $\hat{t}_s = 90$ seconds and $v_t = v_{net}$ in Eqn. 11:

$$\begin{aligned} \frac{90}{3600 \cdot 24} &= \frac{16}{3600 \cdot 24} \cdot \frac{H_{max}}{T_U \cdot 80} \cdot \frac{D}{v_{net}} \\ \implies v_{net} &= \frac{H_{max}}{450 \cdot T_U} \cdot D \\ &\simeq 9544517.4 \cdot D \\ &\approx 10^7 \cdot D \end{aligned} \tag{16}$$

Hence each unit of difficulty represents a net weight of approximately 10 million actively staking coins. Furthermore, substituting the second line of Eqn. 16 into Eqn. 11 gives

$$\hat{t}_s = \frac{1}{960} \cdot \frac{v_{net}}{v_t} \tag{17}$$

which shows the inverse relationship between staking time and the value of the wallet, and also verifies ~ 960 blocks per day upon setting $v_t = v_{net}$.

[8] See the netweight sim tab of the spreadsheet linked in Appendix B, Supplementary Materials, for a demonstration of how the estimations follow changes in the network.

3 Cooldown

3.1 Cooldown Correction for ETTS

After UTXO i stakes, it undergoes a 16 hour cooldown period τ during which it is ineligible to stake again;[9] this cooldown must be included in the calculation of the ETTS, $\hat{t}_s$. Since the staking probability of UTXO i is 0 during its cooldown period and p_i otherwise, the cooldown-corrected ETTS for some UTXO i, $\hat{t}_i^C$, is the sum of τ and $\hat{t}_s$ for UTXO i:[10]

$$\hat{t}_i^C = \tau + \frac{G \cdot D}{v_i} \tag{18}$$

where

$$\tau = 16 \text{ hours} = \frac{2}{3} \text{ days} \tag{19}$$

The frequency of UTXO i staking including cooldown is $f_i^C = (\hat{t}_i^C)^{-1}$. Since each UTXO stakes independently, the frequency of staking for a wallet including cooldown f_s^C is the sum of the staking frequencies of each UTXO:

$$f_s^C = \sum_{i=1}^{n} f_i^C = \sum_{i=1}^{n} \frac{1}{\hat{t}_i^C} = \sum_{i=1}^{n} \frac{1}{\tau + \frac{G \cdot D}{v_i}} \tag{20}$$

Considering a wallet split into n UTXOs of equal value $\bar{v}$ such that:

$$\bar{v} = \frac{v_t}{n} = v_i \ \forall i \tag{21}$$

simplifies Eqn. 20 to:

$$f_s^C = \frac{n}{\tau + \frac{G \cdot D}{\bar{v}}} \tag{22}$$

Furthermore, since $f_s^C = (\hat{t}_s^C)^{-1}$, the ETTS of the whole wallet corrected for cooldown is

$$\hat{t}_s^C = \frac{\tau + G \cdot \frac{D}{\bar{v}}}{n} = \frac{\tau}{n} + \hat{t}_s \tag{23}$$

Therefore, the number of UTXOs "dilutes" the cooldown, where $\frac{\tau}{n}$ is the cooldown correction to $\hat{t}_s$.

3.2 Staking Efficiency

The staking efficiency η of a wallet is the ratio of the corrected expected staking frequency to the ideal expected staking frequency:

$$f_s = \frac{1}{\hat{t}_s} \tag{24}$$

$$\eta = \frac{f_s^C}{f_s} = \frac{\hat{t}_s}{\hat{t}_s^C} = \frac{\hat{t}_s}{\frac{\tau}{n} + \hat{t}_s} \tag{25}$$

As expected, η increases as n increases, or, equivalently, as $\bar{v}$ decreases. A useful version version of this equation may be obtained by substituting Eqns. 13 and 23 into Eqn. 25 and simplifying:

$$\eta = \frac{D \cdot G}{D \cdot G + \tau \cdot \bar{v}} \tag{26}$$

Likewise, solving for $\bar{v}$ yields:

$$\bar{v} = \frac{G \cdot D}{\tau} \cdot \left(\frac{1}{\eta} - 1 \right) \tag{27}$$

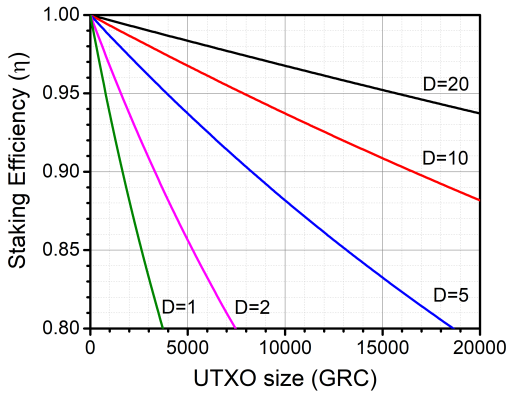

Figure 1: Efficiency as a function of average UTXO size ($\bar{v}$) from Eqn. 26 for a variety of difficulties (D). Note that under normal network conditions, D is expected to be between 2 and 10.

Provided healthy (i.e. $D > 1$) network conditions, splitting a wallet into UTXOs of ≤ 1700 GRC is enough to achieve a staking efficiency of $> 90\%$, as shown in Figure 1. However, it should be noted that this efficiency is compounded over time; higher efficiency of staking will result in a higher wallet balance, which has a higher probability of staking. Using an ideal approximation from Eqn. 23, the effects of UTXO efficient splitting can be seen over a period of months/years, as shown in Figure 2. In reality, more scatter would occur due to changes in difficulty and statistical deviations from the

[9] Additionally, coins in UTXO i cannot be spent for 100 blocks (100 confirmations) after staking.

[10] Wallet staking is dynamic, so every staking cycle includes a cooldown period. Furthermore, it is assumed for analytical purposes that D remains constant as UTXOs begin/conclude cooldown periods, which is a good approximation in a network in quasi-equilibrium with a large set of stakeable UTXOs.

expected time to stake, but this illustrates the simplest, ideal case.

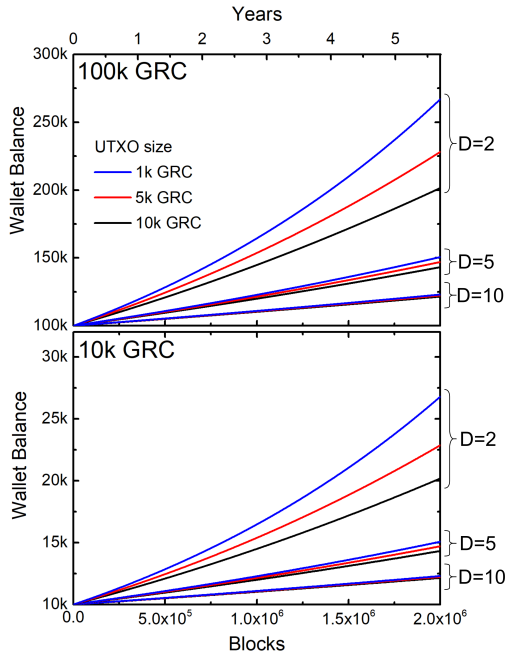

Figure 2: Ideal projected balances following Eqn. 23 for wallets with starting balances of 100k and 10k GRC over a period of 2 million blocks (approximately 5.7 years) for D = 2, 5, and 10. Here, with $D = 2$, a UTXO size of 10k GRC represents an efficiency of 75%, whereas a UTXO size of 1k GRC has an efficiency of 97%.

4 ETTS in the Gridcoin Wallet Version 3.7.12.0+

4.1 Hybrid Approach for ETTS

The method implemented by J. C. Owens in the Gridcoin wallet version 3.7.12.0+[11] uses a more sophisticated and realistic approach to predict $\hat{t}_s$ than either the equilibrium equations discussed above or the previous algorithm in the wallet. While the expected value as calculated by Eqn. 23 is arrived at *over a large number of stakes*, additional information about the time left in cooldown for each recently staked UTXO can be used to refine this approximation.

While UTXO staking is probabilistic, the remaining time in cooldown after staking is deterministic, that is, 16 hours minus the time since staking. This leads to the algorithm that is implemented in wallet version 3.7.12.0+, where a hybrid calculation is employed that considers both the deterministic nature of cooldown periods in progress as well as the memoryless, probabilistic nature of the Bernoulli trials for UTXOs off of cooldown.

Consider the wallet represented in Figure 3, where the total balance is split into four equally sized UTXOs and each bar represents a 16 hour cooldown period during which the probability to stake is zero. In this hypothetical situation, UTXO A has not staked recently and has a normal probability of staking as calculated by Eqn. 13. Likewise, B staked over 16 hours ago, has thus come off cooldown, and has the same probability to stake as A because it has the same value. C, on the other hand, staked less than 16 hours ago and has a known cooldown period equal to 16 hours minus the time since staking. Similarly, D just staked and has a full 16 hours of cooldown left before it is eligible to stake again.

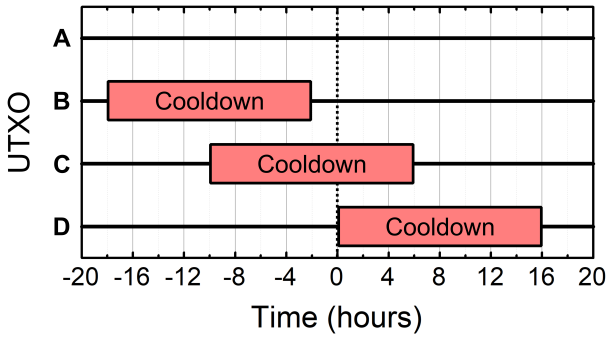

Figure 3: Illustrative chart of a wallet divided into UTXOs A, B, C, and D, where each has staked at a different point in time and is subject to a 16 hour cooldown period, denoted with red bars, during which the probability of staking is zero.

In this scenario, during the period up to the expiration of the cooldown of UTXO C, only A and B can stake. From that point, only A, B, and C can stake until D's cooldown has expired, at which point the entire wallet will resume staking according to Eqn. 13. This creates a stepwise change in probability of staking that must be considered to offer an aggregate ETTS.

To take this into account, the code walks through these UTXOs using a nested loop with two levels. In the outer loop, which can be thought of as the vertical line in Figure 3 for the time being considered, the UTXOs are collected and ordered by their cooldown expiration. An ordered set of events is constructed which consists of the current time followed by the time that each UTXO comes off of cooldown. A timestamp mask of 16x the normal mask (which means 256 second granularity instead of 16 seconds) is used to reduce the work in the outer portion of the nested loop, so that a 16 hour interval will have a maximum of 225 events. The inner loop goes through each UTXO. If UTXO i is in cooldown

[11] https://github.com/gridcoin-community/Gridcoin-Research/pull/1044

[12] See the ETTS tab of the spreadsheet linked in Appendix B, Supplementary Materials, for a complete example calculation using the method detailed below, including a complete example scenario similar to Figure 3.

during the interval between events, its p_i is 0. If it is off of cooldown, its p_i is determined by Eqn. 13.[12]

4.2 Hybrid Algorithm for ETTS

The probability of a wallet staking will remain constant until either: 1) a UTXO stakes (then goes on cooldown); 2) a UTXO comes off of cooldown; or 3) a UTXO is spent. Since UTXOs coming off of cooldown are predictable events, a series of intervals between UTXOs coming off of cooldown may be constructed during which the probability of staking is constant — i.e. an interval of constant v_t.

Let p_i be the probability that UTXO i stakes during a single staking loop, with $i \in [1, n_k]$, where n_k is the number of eligible UTXOs in interval k. Let t_k be the time duration of interval k, quantized to multiples of 256 seconds. The probability that the first i' UTXOs *do not stake* during a single staking loop is:

$$\begin{aligned} \bar{P}_{i'} &= (1-p_1)\cdot(1-p_2)\cdot\ldots\cdot(1-p_{i'}) \\ &= \Pi_{i=1}^{i'}(1-p_i) \end{aligned} \tag{28}$$

Thus, the probability that *none* of the UTXOs stake during the loop is:

$$\bar{P}_{n_k} = \Pi_{i=1}^{n_k}(1-p_{n_k}) \tag{29}$$

and the probability that one of the UTXOs *will* stake is:

$$\begin{aligned} P_{i'} &= 1-\bar{P}_{i'} \\ &= 1-\Pi_{i=1}^{i'}(1-p_i) \end{aligned} \tag{30}$$

However, the stake loop occurs every 16 seconds; many loops can occur in an interval of constant v_t. Consider stake loop j_k within interval k. The cumulative probability $\bar{L}_{j_k}$ that no UTXO stakes within j_k loops is:

$$\bar{L}_{j_k} = (\bar{P}_{n_k})^{j_k} \tag{31}$$

Let m_k be the total number of stake loops in interval k:

$$m_k = \frac{t_k}{16 \text{ seconds}} \tag{32}$$

such that $j_k \in [1, m_k]$. Then the probability $\bar{I}_k$ that no UTXO stakes within interval k is:

$$\bar{I}_k = \bar{L}_{m_k} = (\bar{P}_{n_k})^{m_k} \tag{33}$$

and the probability that a UTXO does stake within this interval is:

$$I_k = L_{m_k} = 1-(\bar{P}_{n_k})^{m_k} \tag{34}$$

This covers the probability to stake within a single interval. To iterate over multiple intervals, consider F_k, the probability that a UTXO stakes by interval k. The probability that no UTXOs stake by this interval is:

$$\begin{aligned} \bar{F}_k &= \bar{I}_1\cdot\bar{I}_2\cdot\ldots\cdot\bar{I}_k \\ &= \Pi_{k'=1}^{k}\bar{I}_{k'} \end{aligned} \tag{35}$$

or recursively:

$$\bar{F}_k = \bar{F}_{k-1}\cdot\bar{I}_k \tag{36}$$

Thus, the probability that a UTXO will stake by the k^{th} interval is

$$\begin{aligned} F_k &= 1-\bar{F}_k = 1-\Pi_{k'=1}^{k}\bar{I}_{k'} \\ &= 1-\Pi_{k'=1}^{k}(\bar{P}_{n_{k'}})^{m_{k'}} \end{aligned} \tag{37}$$

Recursively,

$$F_k = 1-(1-F_{k-1})\cdot(1-I_k) \tag{38}$$

If the desired cumulative probability of staking, F_D,[13] has not yet been reached by stake loop $m_{k_{max}}$, the final interval $k_{max}+1$, in which all UTXOs are online and staking, must be considered. Note here that $m_{k_{max}+1}$ will be the number of loops required to achieve F_D. The wallet is still able to stake after this point (in fact, as $F_D \to 1$, the number of required trials $\to \infty$), but the interval $k_{max}+1$ ends once F_D is reached. To determine $m_{k_{max}+1}$, we use Eqns. 34 and 38:

$$\begin{aligned} I_{k_{max}+1} &= 1-(\bar{P}_{n_{k_{max}+1}})^{m_{k_{max}+1}} \\ F_{k_{max}+1} = F_D &= 1-(1-F_{k_{max}})\cdot(1-I_{k_{max}+1}) \end{aligned} \tag{39}$$

Combining and solving for m_{k+1} yields:

$$m_{k_{max}+1} = \frac{\ln(1-F_D)-\ln(1-F_{k_{max}})}{\ln(1-P_{n_{k_{max}+1}})} \tag{40}$$

[13] $F_D = 80\%$ as of version 3.7.12.0. See Section 4.3 for a comparison of the expected value and hybrid approaches, and the reasoning for choosing $F_D = 80\%$.

If F_D has not been reached by k_{max}, then $F_{k_{max}}$ is used for F_k and m_{k+1} represents the number of additional trials necessary to achieve F_D. If F_D is reached or exceeded at an interval k that is before k_{max}, then an "overshoot" has occurred. In that situation, Eqn. 40 is used at interval k rather than k_{max} in modified form:

$$m_k = \frac{\ln(1 - F_D) - \ln(1 - F_k)}{\ln(1 - P_{n_k})} \tag{41}$$

F_k will be greater than F_D, and the result will be negative, which represents the number of trials to "backtrack" for the overshoot. This is far more efficient than iterating over staking loops individually to avoid the overshoot.

To arrive at the ETTS, the number of trials needs to be summed for all intervals. Because the wallet undergoes a staking loop every 16 seconds, the corresponding ETTS in days for F_D is:

$$\hat{t}_s = \frac{16}{24 \cdot 3600} \cdot \sum_{k=1}^{k'} j_{max}^k \tag{42}$$

where k' is the interval in which F_D is reached or exceeded, $k' \in [0, k_{max} + 1]$.

4.3 Expected Value vs. Hybrid Approach

To facilitate the comparison of the expected value approach (Section 2.2) and the hybrid approach (Section 4.1), we can consider an ideal case where the entire wallet balance always contributes to the staking probability p.[14] This reduces the hybrid approach to a simpler cumulative distribution function (CDF)[15]:

$$P(p, n_C) = 1 - (1 - p)^{n_C} \tag{43}$$

where n_C is the number of trials required to achieve some total probability P from some individual trials of probability p. Solving for n_C gives

$$n_C = \frac{\ln(1 - P)}{\ln(1 - p)} \tag{44}$$

However, the expected number of trials n_E is known from the expected value approach:

$$n_E = \frac{1}{p} \tag{45}$$

Let R denote the ratio between then number of trials from the hybrid approach and the expected value approach:

$$\begin{aligned} R = \frac{n_C}{n_E} = \frac{\frac{\ln(1-P)}{\ln(1-p)}}{\frac{1}{p}} &= \frac{\ln(1 - P)}{\frac{1}{p} \cdot \ln(1 - p)} \\ &= \frac{\ln(1 - P)}{\ln((1 - p)^{\frac{1}{p}})} \end{aligned} \tag{46}$$

Since each p_i is small, R may be approximated:

$$\begin{aligned} R \approx \lim_{p \to 0} R = \lim_{p \to 0} \frac{\ln(1 - P)}{\ln((1 - p)^{\frac{1}{p}})} &= \frac{\ln(1 - P)}{\ln(e^{-1})} \\ &= -\ln(1 - P) \end{aligned} \tag{47}$$

where the identity $\lim_{p \to 0}(1 - p)^{\frac{1}{p}} = e^{-1}$ is used. Substituting $P = F_D = 0.80$ into Eqn. 47, we see that the ETTS provided by the hybrid method at the ideal limit (ignoring cooldown) is $\sim 61.1\%$ greater than the one provided by the expected value method. This means that the ETTS provided by the wallet in 3.7.12.0+, ignoring cooldown, is about 61% greater than the estimate in Eqn. 13.

Additionally, letting $n_C = n_E = \frac{1}{p}$ and $p \to 0$ in Eqn. 43,

$$\begin{aligned} \lim_{p \to 0} P\left(p, \frac{1}{p}\right) &= \lim_{p \to 0}(1 - (1 - p)^{\frac{1}{p}}) \\ &= 1 - \frac{1}{e} \simeq 63.2\% \end{aligned} \tag{48}$$

Note also that $R \approx 1$ when $P \simeq 0.632$ (Eqn. 47). Therefore the expected value method is essentially equivalent to the CDF method at $F_D = 63.2\%$.

$F_D = 80\%$ was chosen to reduce the frequency of the hybrid approach underestimating the actual time required to stake, such that the actual time to stake will exceed the ETTS 20% of the time over a large number of stakes. In comparison, using Eqn. 48, the expected value approach provides a result that will be exceeded $1 - 63.2\% = 36.8\%$ of the time over a large number of stakes. While the choice of F_D is arbitrary, 80% provides a good middle ground ETTS for smaller wallets that stake infrequently, reducing the probability that the ETTS is exceeded, surprising the wallet owner.[16]

[14] While practically impossible, this condition may be approximated using a reasonably large number of UTXOs.

[15] See https://steemit.com/gridcoin/@hotbit/confident-time-to-stake, which bears some similarity in its treatment of ETTS

[16] This is important when considering solo mining Gridcoin through BOINC computation, where a block must be staked to claim research rewards and the rewards expire after 180 days. This subject will be covered in depth in a future Gridcoin Bluepaper section.

4.4 Variable-P $\hat{t}_s, \hat{t}_s^C$

Since R gives the ratio of the hybrid ETTS to the expected value ETTS, we can multiply $\hat{t}_s$ by R to modify it to account for varying P.

To begin, combine Eqn. 13 and Eqn. 47:

$$\begin{aligned} \hat{t}_s(P) &= R \cdot \hat{t}_s \\ &= \frac{n_C}{n_E} \cdot \hat{t}_s \end{aligned} \tag{49}$$

Taking the limit as $p \to 0$ gives

$$\hat{t}_s(P) = -\ln(1-P)\frac{G \cdot D}{v_t} \tag{50}$$

Note that $\hat{t}_s(P) = \hat{t}_s$ when $n_C = n_E$ and when $P = 0.632$. Restating Eqn. 13 and Eqn. 23 at $P = F_D = 0.80$,

$$\hat{t}_s|_{P=0.80} \simeq \frac{16000}{v_t} \cdot D \tag{51}$$

$$\hat{t}_s^C|_{P=0.80} = \frac{(\frac{2}{3})}{n} + \frac{16000}{v_t} \cdot D \tag{52}$$

In this way, Eqn. 50 provides a variable-P version of Eqns. 13 and 23, which can be used as a quick rule of thumb for estimating the time to stake.

5 Future Improvements

Two shortcomings of the ETTS calculation will be addressed in future versions of the wallet. First, the current calculation of ETTS is null when all UTXOs are on cooldown, even though some UTXOs could be coming off of cooldown shortly. This will be addressed by changing the use of the staking flag in the ETTS calculations. Second, the D used in the ETTS algorithm is averaged over 40 blocks, which may be too short for infrequent stakers. This potentially can be addressed by adjusting the averaging interval recursively based on the wallet balance and difficulty.

6 Acknowledgments

The authors would like to thank the Gridcoin community members who aided in the preparation and editing of this document, particularly h202, hotbit, and Nutney. J.C.O. would like to acknowledge the developer funding provided by the Gridcoin Foundation which facilitates the development and documentation of Gridcoin.

Appendices

A Supplemental Notes

UTXOs: Gridcoin, like other coins similar to Bitcoin, utilizes transactions between addresses recorded in the blockchain. Just like cash transactions, there are inputs and outputs to each transaction. Since everything in the blockchain must be traceable, one or more outputs from a prior transaction which are "unspent" (not sent to an address in another wallet by the transaction), expressed with a unique ID and a value, are used as the inputs for the next transaction by the wallet. These unspent "outputs" are called Unspent Transaction Outputs (UTXOs). From the perspective of an individual wallet, the total value of the UTXOs in the wallet is the wallet balance. From a network-wide perspective, the total value across all UTXOs in the network is the total supply minus the number of Gridcoins "burned" (sent to an irretrievable address).

B Supplementary Materials

The following link is to a spreadsheet which presents the block spacing t^i, and D as a function of v_{net}.[17] It also provides a calculation of staking efficiency and ETTS.

`https://docs.google.com/spreadsheets/d/1FGJewsTIWMhxHf--GjfOOa57npoY8uC8Mpux7kURZas/edit?usp=sharing`

C List of Relevant Functions in the Source Code

Function	Source Code Location(s)
bool CreateCoinStake(CBlock &blocknew, CKey &key, vector ⟨const CWalletTx*⟩ &StakeInputs, uint64_t &CoinAge, CWallet &wallet, CBlockIndex* pindexPrev)	miner.h, miner.cpp
static unsigned int GetNextTargetRequiredV2(const CBlockIndex* pindexLast, bool fProofOfStake)	main.h, main.cpp
double GetEstimatedNetworkWeight(unsigned int nPoSInterval)	main.h, main.cpp
double GetAverageDifficulty(unsigned int nPoSInterval)	main.h, main.cpp
double GetEstimatedTimetoStake(double dDiff, double dConfidence)	main.h, main.cpp

[17] This is not a full Monte Carlo simulation as it substitutes theoretical input block frequencies (spacing) based on a "God parameter" of the actual amount of Gridcoin online, but nevertheless, it shows very effectively the response of the network, and the measured parameters, to that "God parameter." In a full Monte Carlo simulation, the scatter would be higher, but the behavior will follow the same trends.