# F1 Fee Distribution

**Dev Ojha**
Tendermint, Berkeley, CA, USA
UC Berkeley, CA, USA
dojha@berkeley.edu

**Christopher Goes**
Tendermint, Berlin, Germany
`https://pluranimity.org/about`
cwgoes@tendermint.com

## Abstract

In a proof of stake blockchain, validators need to split the rewards gained from transaction fees each block. Furthermore, these fees must be fairly distributed to each of a validator's constituent delegators. Delegators accrue this reward throughout the entire time which they are delegated, and they have a special operation to withdraw accrued rewards.

The F1 fee distribution scheme works for any algorithm to split fees and inflation between validators each block, with minimal iteration, and the only approximations being due to finite decimal precision. Per block there is a single iteration over the validator set, to enable reward algorithms that differ by validator. No iteration is required to delegate or to withdraw. The state usage is one state update per validator per block and one state entry per active delegation. F1 can optionally handle arbitrary inflation schemes, auto-bonding of rewards, and varying validator commission rates.

**2012 ACM Subject Classification** Theory of computation → Algorithmic mechanism design

**Keywords and phrases** Proof of Stake, Fee Distribution, Cosmos

**Digital Object Identifier** 10.4230/OASIcs.Tokenomics.2019.10

**Supplement Material** Implementation: `https://github.com/cosmos/cosmos-sdk/tree/develop/x/distribution`

## 1 Introduction

In a proof of stake blockchain, each validator has an associated stake. Transaction fees are rewarded to validators based on the incentive scheme of the underlying proof of stake model. However, only rewarding the proposers as in many proof-of-work incentive models causes incentive problems. See these prior works discussing this problem. [1] [3] This fee distribution problem occurs in delegated proof-of-stake blockchains, as there is a need to distribute a validator's fee rewards and inflation proportionally to its delegators according to amount of stake each has delegated. The trivial solution of just paying the rewards to each delegator every block is too expensive to perform on-chain, as it would require reading and writing all delegator accounts. Instead fee distribution algorithms must require that delegators perform a withdraw action, which when performed yields the same total amount of fees as if they had received them at every block.

This paper details F1, an approximation-free, slash-tolerant fee distribution algorithm which allows validator commission-rates, inflation rates, and fee proportions, which can all efficiently change per validator, every block. The algorithm requires iterating over the bonded validators every block, which is cheap due to staking logic already requiring iteration over all validators, which causes the expensive state-reads to be cached. Withdraws require no iteration.


© Dev Ojha and Christopher Goes;
licensed under Creative Commons License CC-BY
International Conference on Blockchain Economics, Security and Protocols (Tokenomics 2019).
Editors: Vincent Danos, Maurice Herlihy, Maria Potop-Butucaru, Julien Prat, and Sara Tucci-Piergiovanni;
Article No. 10; pp. 10:1–10:6

OpenAccess Series in Informatics
OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The key point of how F1 works is that it tracks how much rewards a delegator with 1 stake delegated to a given validator would be entitled to if it had bonded at block 0 until the latest block. When a delegator bonds at block $b$, the amount of rewards a delegator with 1 stake would have if bonded at block 0 until block $b$ is also persisted to state. When the delegator withdraws, they receive the difference of these two values. Since rewards are distributed according to stake-weighting, this amount of rewards can be scaled by the amount of stake a delegator had delegated. The following paragraph describes this in more detail, with a demonstration of equivalence to the inefficient iterative algorithm by reduction. Section 2 details how to adapt this algorithm to handle commission rates, slashing, inflation, and auto-bonding of fees.

### Base algorithm

In this section, we show that the F1 base algorithm gives each delegator rewards identical to that which they'd receive in the naive and correct fee distribution algorithm that iterated over all delegators every block.

Even distribution of a validators rewards amongst its validators weighted by stake means the following: Suppose a delegator delegates $x$ stake to a validator $v$ at block $h$. Let the amount of stake the validator has at block $i$ be $s_i$ and the amount of fees they receive at this height be $f_i$. Then if a delegator contributing $x$ stake decides to withdraw at block $n$, the rewards they receive are

$$\sum_{i=h}^{n} \frac{x}{s_i} f_i = x \sum_{i=h}^{n} \frac{f_i}{s_i}$$

Note that $s_i$ does not change every block, it only changes if the validator gets slashed, or if any delegator alters the amount they have delegated. We'll relegate handling of slashes to Subsection 2.2, and only consider the case with no slashing here. We can change the iteration from being over every block, to instead being over the set of blocks between two changes in validator $v$'s total stake. Let each of these set of blocks be called a period. A new period begins every time that validator's total stake changes. Let the total amount of stake for the validator in period $p$ be $n_p$. Let $T_p$ be the total fees that validator $v$ accrued in period $p$. Let $h$ be the start of period $p_{init}$, and height $n$ be the end of $p_{final}$. It follows that

$$x \sum_{i=h}^{n} \frac{f_i}{s_i} = x \sum_{p=p_{init}}^{p_{final}} \frac{T_p}{n_p}$$

Let $p_0$ represent the period which begins when the validator first bonds. The central idea to the F1 model is that at the end of the $k$th period, the following is stored at a state location indexable by $k$: $\sum_{i=0}^{k} \frac{T_i}{n_i}$. Let the index of the current period be $f$. When a delegator wants to delegate or withdraw their reward, they first create a new entry in state to end the current period. Then this entry is created using the previous entry as follows:

$$Entry_f = \sum_{i=0}^{f} \frac{T_i}{n_i} = \sum_{i=0}^{f-1} \frac{T_i}{n_i} + \frac{T_f}{n_f} = Entry_{f-1} + \frac{T_f}{n_f}$$

Where $T_f$ is the fees the validator has accrued in period $f$, and $n_f$ is the validators total amount of stake in period $f$.

The withdrawer's delegation object has the index $k$ for the period which they ended by bonding. (They start receiving rewards for period $k+1$) The reward they should receive when withdrawing is:

$$x \sum_{i=k+1}^{f} \frac{T_i}{n_i} = x \left( \left( \sum_{i=0}^{f} \frac{T_i}{n_i} \right) - \left( \sum_{i=0}^{k} \frac{T_i}{n_i} \right) \right) = x \left( Entry_f - Entry_k \right)$$

It is clear from the equations that this payout mechanism maintains correctness, and requires no iterations. It just needed the two state reads for these entries.

$T_f$ is a separate variable in state for the amount of fees this validator has accrued since the last update to its power. This variable is incremented at every block by however much fees this validator received that block. On the update to the validators power, this variable is used to create the entry in state at $f$, and is then reset to 0.

This fee distribution proposal is agnostic to how all of the blocks fees are divied up between validators. This creates many nice properties, for example it is possible to only rewarding validators who signed that block.

# 2 Additional add-ons

## 2.1 Commission Rates

Commission rates are the idea that a validator can take a fixed $x\%$ cut of all of their received fees, before redistributing evenly to the constituent delegators. This can easily be done as follows:

In block $h$ a validator receives $f_h$ fees. Instead of incrementing that validators "total accrued fees this period variable" by $f_h$, it is instead incremented by $(1 - commission\_rate) * f_p$. Then $commission\_rate * f_p$ is deposited directly to the validator's account. This allows for efficient updates to a validator's commission rate every block if desired. More generally, each validator could have a function which takes their fees as input, and outputs a set of outputs to pay these fees too. (i.e. x% going to themselves, y% to delegators, z% burnt)

## 2.2 Slashing

Slashing is distinct from withdrawals, since it lowers the stake of all of the delegator's by a fixed percentage. Since no one is charged gas for slashes, a slash cannot iterate over all delegators. Thus we can no longer just multiply by $x$ over the difference in stake. This section describes a simple solution that should suffice for most chains needs. An asymptotically optimal solution is provided in section 2.4.

The solution here is to instead store each period created by a slash in the validators state. Then when withdrawing, you must iterate over all slashes between when you started and ended. Suppose you delegated at period 0, a y% slash occured at period 2, and your withdrawal creates period 4. Then you receive funds from periods 0 to 2 as normal. The equations for funds you receive for periods 2 to 4 now uses $(1-y)x$ for your stake instead of just $x$ stake. When there are multiple slashes, you just account for the accumulated slash factor.

There is a griefing attack[2] a validator can perform on its delegators in this model. The validator can make itself be slashed "n" times, with a linear increase in the cost to withdraw for its constituent delegators. It is anticipated that the slashing penalty is sufficiently high that this won't be a practical concern.

## 2.3 Inflation

Inflation is the idea that we want every staked coin to create more staking tokens as time progresses. The purpose being to drive down the relative worth of unstaked tokens. Each block, every staked token should produce $x$ staking tokens as inflation, where $x$ is calculated from a function $inflation$ which takes state and the block information as input. Let $x_i$ represent the evaluation of $inflation$ in the $i$th block. The goal of this section is to auto-bond inflation in the fee distribution model without iteration. This is done by preserving the invariant that every state entry contains the rewards one would have if they had bonded one stake at genesis until that corresponding block.

In state a variable should be kept for the number of tokens one would have now due to inflation, given that they bonded one token at genesis. This is $\prod_0^{now}(1+x_i)$. Each period now stores this total inflation product along with what it already stores per-period.

Let $R_i$ be the fee rewards in block $i$, and $n_i$ be the total amount bonded to that validator in that block. The correct amount of rewards which 1 token at genesis should have now is:

$$Reward(now) = \sum_{i=0}^{now} \left( \prod_{j=0}^{i} 1 + x_j \right) * \frac{R_i}{n_i}$$

The term in the sum is the amount of stake one stake becomes due to inflation, multiplied by the amount of fees per stake.

Now we cast this into the period frame of view. Recall that we build the rewards by creating a state entry for the rewards of the previous period, and keeping track of the rewards within this period. Thus we first define the correct amount of rewards for each successive period, proving correctness of this via induction. We then show that the state entry that gets efficiently built up block by block is equal to this value for the latest period.

Let $start, end$ denote the start/end of a period.

Suppose that $\forall f > 0$, $Reward(end(f))$ is correctly constructed as

$$Reward(end(f)) = Reward(end(f-1)) + \sum_{i=start(f)}^{end(f)} \left( \prod_{j=0}^{i} 1 + x_j \right) \frac{R_i}{n_i}$$

and that for $f = 0$, $Reward(end(0)) = 0$. (With period 1 being defined as the period that has the first bond into it) It must be shown that assuming the supposition $\forall f \leq f_0$,

$$Reward(end(f_0+1)) = Reward(end(f_0)) + \sum_{i=start(f_0+1)}^{end(f_0+1)} \left( \prod_{j=0}^{i} 1 + x_j \right) \frac{R_i}{n_i}$$

Using the definition of $Reward$, it follows that:

$$\sum_{i=0}^{end(f_0+1)} \left( \prod_{j=0}^{i} 1 + x_j \right) * \frac{R_i}{n_i} = \sum_{i=0}^{end(f_0)} \left( \prod_{j=0}^{i} 1 + x_j \right) * \frac{R_i}{n_i} + \sum_{i=start(f_0+1)}^{end(f_0+1)} \left( \prod_{j=0}^{i} 1 + x_j \right) \frac{R_i}{n_i}$$

Since the first summation on the right hand side is $Reward(end(f_0))$, the supposition is proven true. Consequently, the reward for just period $f$ adjusted for the amount of inflation 1 token at genesis would produce, is:

$$\sum_{i=start(f)}^{end(f)} \left( \prod_{j=0}^{i} 1 + x_j \right) \frac{R_i}{n_i}$$

Note that

$$\sum_{i=start(f)}^{end(f)} \left( \prod_{j=0}^{i} 1 + x_j \right) \frac{R_i}{n_i} = \left( \prod_{j=0}^{end(f-1)} 1 + x_j \right) \sum_{i=start(f)}^{end(f)} \left( \prod_{j=start(f)}^{i} 1 + x_j \right) \frac{R_i}{n_i}$$

By definition of period, and inflation being applied every block,
$n_i = n_{start(f)} \left( \prod_{j=start(f)}^{i} 1 + x_j \right)$. This cancels out the product in the summation, therefore

$$\sum_{i=start(f)}^{end(f)} \left( \prod_{j=0}^{i} 1 + x_j \right) \frac{R_i}{n_i} = \left( \prod_{j=0}^{end(f-1)} 1 + x_j \right) \frac{\sum_{i=start(f)}^{end(f)} R_i}{n_{start(f)}}$$

Thus every block, each validator just has to add the total amount of fees (The $R_i$ term) that goes to delegates to some per-period term. When creating a new period, $n_{start(f)}$ can be cached in state, and the product is already stored in the previous periods state entry. You then get the next period's $n_{start(f)}$ from the consensus' power entry for this validator. This is thus extremely efficient per block.

When withdrawing, you take the difference as before, calculating the difference between the reward entry at the withdrawing height and the bonding height. This yields the amount of rewards you would have obtained with $(\prod_0^{begin\ bonding\ period} 1 + x)$ stake from the block you began bonding at until now. $(\prod_0^{begin\ bonding\ period} 1 + x)$ is known, since its included in the state entry for when you bonded. You then divide the entitled fees by $(\prod_0^{begin\ bonding\ period} 1 + x)$ to normalize it to being the amount of rewards you're entitled to from 1 stake at that block to now. Then as before, you multiply by the amount of stake you had initially bonded.

In addition to the above, the withdrawer also needs rewards due to inflation itself. This can be done by taking the accumulated inflation factor, and dividing it by the inflation factor until the beginning of the bonding period. This factor is $(\prod_{begin\ bonding\ period}^{now} 1 + x)$, and then that gets scaled by how much they initially bonded.

The inflation function could vary per block, and per validator if ever a need arose. If the inflation rate is the same for all validators then there can be a single state entry for the entries corresponding to the product of inflations. Inflation creation can trivially be epoched as long as inflation isn't required within the epoch, through changes to the $inflation$ function.

## 2.4 Withdrawing with no iteration over slashes

Notice that a slash is the same as a negative inflation rate for a validator in one block. For example a 20% slash is equivalent to a $-20\%$ inflation for a validator in a block. Given correctness of auto-bonding inflation with different inflation rates per-validator, it follows that handling slashes can be correctly done by simply setting the validators inflation factor in that block to be the negative of the slash factor. This significantly simplifies the withdrawal procedure.

## 2.5 Auto bonding fees

Auto bonding of fees also follows from the correctness of the inflation model. Split up the rewards into one component with only the staking token, and one component with the remaining tokens. Add to the inflation rate for that block for that validator, $\frac{amount_of_staking_token}{n_i}$, $n_i$ being the validators stake in that block. Set the rewards to then just be the remaining tokens.

## 2.6 Delegation updates

Updating your delegation amount is equivalent to withdrawing earned rewards and a fully independent new delegation occurring in the same block. The same applies for redelegation. From the view of fee distribution, partial redelegation is the same as a delegation update and a new delegation.

# 3 State Requirements

State entries can be pruned quite effectively. Suppose for the sake of exposition that there is at most one delegation / withdrawal to a particular validator in any given block. Then each delegation is responsible for one addition to state. Only the next period, and this delegator's withdrawal could depend on this entry. Thus once this delegator withdraws, this state entry can be pruned. For the entry created by the delegator's withdrawal, that is only required by the creation of the next period. Thus once the next period is created, that withdrawal's period can be deleted.

This can be easily adapted to the case where there are multiple delegations / withdrawals per block, by maintaining a reference count in each period starting state entry.

The slash entries for a validator can only be pruned when all of that validator's constituent delegators have their bonding period starting after the slash. This seems ineffective to keep track of, thus it is not worth it. Each slash should instead remain in state until the validator unbonds and all delegators have their fees withdrawn.

Thus, with reference counting, it will always be the case that the total reference count for a particular validator is equal to the number of active delegations (each keeping a reference to the period ended by their delegation) plus the number of slashes (each keeping a reference to the period ended by the slash) plus one (for the most recent period).

# 4 Implementers Considerations

We have heretofore described F1, a pragmatic fee distribution algorithm with many benefits. The overhead per block is a simple iteration over the bonded validator set, which will often occur anyways due to underlying proof of stake logic (such as to check whether any validators have entered or left). Consequently it can be implemented with minimal additional, as the state entry reads and writes can be cached. All calculations are exact, modulo minor errors resulting from fixed precision decimals. F1 supports arbitrary inflation and fee models, which can vary per validator per block (which enables desirable incentive mechanisms, for example paying only validators which signed the block), as can the commission rates of the individual validators. The simplicity of the scheme lends itself well to implementation. F1 has been implemented in the Cosmos SDK and will be used in the Cosmos Hub blockchain.

## References


1 Sunny Aggarwal. Cosmos Proof of Stake. Crypto Economics Security Conference 2018, 2018. URL: https://www.youtube.com/watch?v=XxZ04w2x4nk.

2 Vitalik Buterin. Discouragement Attacks. ETH research, 2018. URL: https://github.com/ethereum/research/blob/367507e0785f488cc269e0a3a61f49ce3c000327/papers/discouragement/discouragement.pdf.

3 Giulia Fanti, Leonid Kogan, Sewoong Oh, Kathleen Ruan, Pramod Viswanath, and Gerui Wang. Compounding of Wealth in Proof-of-Stake Cryptocurrencies. Arxiv, October 2018. arXiv:1809.07468.