

Fast Indexing Engine for Data Identified by a Hashed ID and Stored in an Immutable File

CYRIL DEVER
Edgewhere

November 26, 2018

Abstract

We define here an algorithm for indexing the system based on identifiers that are hashed values which is at the same time very powerful to the writing and the reading. We call it the Treee™.

I. INTRODUCTION

THE challenge was to set up a powerful yet safe search engine to use when the data is some linked list of items that could be themselves connected to each other in subchains, indexed through their identifiers that are only made of hashed values (like SHA-256 string representations), and all stored in an immutable file.

Its best application is for a blockchain file where an *item* is a transaction embedding a smart contract, and each subchain of items the subsequent uses and/or modifications of this smart contract.

This present document describes such indexing engine.

Definition 1 (Item). An *item* is the actual object recorded in the immutable file subject to indexing.

As previously stated, it is initially meant to be a transaction or a block in a blockchain file.

Definition 2 (Leaf). A leaf λ embeds the information helping to retrieve one or more possibly linked *items*.

II. FORMAL DESCRIPTION

1. ACYCLIC GRAPH

Treee™ is an algorithm for indexing *items* recorded in an immutable file based on their identifiers that are hashed values.

It is constructed as an acyclic graph (a tree T), each node containing either a node address (its sons) or a set of *leaf*.

Definition 3 (Hashed Identifier). We call ι a hashed identifier (or hashed value) the hexadecimal string representation of the result of a data d passed through a hashing function $h()$ such as:

$$\iota := h(d) \quad (1)$$

The passed data d could be anything but it must be unique if it were to be used as identifier per se.

In blockchains we operate, this data d is usually the *item* itself, ie. a transaction or a block.

The hashing function could use any cryptographic hashing algorithm as long as it is set beforehand and once and for all in the targeted system¹. The number \mathfrak{N} represents the number of bits of the returned hash, eg. 256 for SHA-256.

The number of sons of a node is deterministic and depends on the depth of the tree.

Let p_k be the number of sons of a node N_k at depth k .

The goal is to create a balanced tree whose width is adaptive to decrease depth and optimize performance.

We are looking to index numbers, in this case the numerical value of the *item*'s unique identifiers $\iota_i (\forall i \in \mathbb{N})$.

¹We currently use the SHA-256 algorithm because of its wide adoption in both end-user and back-end environments

Recall. An identifier is at its core a hashed value, that is its digest is fundamentally a byte array that could represent any positive integer.

2. INDEX

We now explain the course of the index.

Let ι_i^b be the value of the hashed identifier for the *item* i in binary form, eg.

$$\iota_i^b := \text{"a1"} \mapsto 10100001$$

indicating its position in the tree T .

At each step $j := [0..n] \mid n \leq \aleph$, we would pass to child 0 if the j -th bit of ι_i^b equals 0; otherwise, we would pass to child 1.

Let $R_j^{\iota_i}$ be the value of this representation of ι_i^b at step $j < k$.

For a full tree T , we build a representation of this number at each step and traverse the tree the same way.

At the step j of depth k , we pass to child 0 if $R_j^{\iota_i}$ equals 0, we pass to child 1 if $R_j^{\iota_i}$ equals 1, ..., we pass to child p_{k-1} if $R_j^{\iota_i}$ equals p_{k-1} . We stop when the node is an empty leaf λ_j .

Definition 4 (Representative). To construct $R_j^{\iota_i}$, this representative at step j , we will successively take the modulo of prime numbers, each step j using the j -th prime number in the ordered sequence of all prime numbers $\mathcal{P} := [1, 2, 3, 5, 7, 11, 13, \dots]$, starting at 0.

This construction ensures that $R_j^{\iota_i}$ be unique.

Proof. According to the Chinese remainder theorem[1], each number has a unique representative that could be written as the continuation of these modulus.

Indeed, a number n can be written in the following form:

$$n \equiv n \bmod P_i$$

where P_i is the i -th modulo in \mathcal{P} .

Modulos are calculated in $O(1)$ for fixed-sized integers. Since the multiplication is faster than the division (necessary for the calculation of the modulo), one may use multiplications by means of floating:

$$P_i \times \left(n - \left\lfloor n \times \frac{1}{P_i} \right\rfloor \right)$$

This writing in the form of a sequence allows to uniquely define each integer n . \square

Given the random nature of the numbers (or pseudo-random, since the identifiers of the *items* are generated by cryptographic hashing technologies), the tree T is balanced.

To unbalance it in a malicious way, it would be necessary to be able to generate hashes whose modulo follows a particular trajectory.

However, the difficulty of such an operation increases exponentially (in the order of $e^{(k \times \log(k))}$ where k is the depth of T).

As a reminder, the product of the first 16 prime numbers already equals:

$$32, 589, 158, 477, 190, 044, 730 \simeq 3 \times 10^{19}$$

Therefore, as soon as the index contains a reasonably large amount of data, unbalancing the tree in a malicious way would become more and more impossible, if at all possible.

3. LEAF IS INFORMATION ON ITEM

Let s be a suchain of linked *items*. For example, it could be a sequence of transactions between two stakeholders defining the progressive evolution of the terms of their smart contract.

And let s_0 be the first *item* in a subchain of linked *items*.

A leaf $\lambda_{i_s} \in T$ contains the following list of information about an *item* i referred to by its identifier ι_{i_s} in subchain s :

- Identifier (ID) of the current *item* (as a hash string):

$$\lambda_{i_s}^{ID} := \iota_{i_s}$$

- Position, ie. start address of the current *item* in the file, eg.

$$\lambda_{i_s}^{Pos} := 12080$$

- Size (in bytes) of the saved item in the file, eg.

$$\lambda_{i_s}^{Size} := 2074$$

- Origin, ie. the unique identifier of the *item* that is at the start of the *item*'s subchain (if any):

$$\lambda_{i_s}^{Origin} := \iota_{i_{s_0}}$$

- Previous, ie. the optional unique identifier of the previous *item* chained to it:

$$\lambda_{i_s}^{Prev} := l_{i_{s-1}}$$

- Next, ie. the optional unique identifier of the next *item* chained:

$$\lambda_{i_s}^{Next} := l_{i_{s+1}}$$

A leaf whose next *item* field is empty is the last item in the subchain:

$$\lambda_{i_s}^{Next} = \emptyset \iff l_{i_{s+1}} \notin T \quad (2)$$

A leaf whose origin *item* field is equal to the identifier of the current *item* is necessarily the origin of the subchain:

$$\lambda_{i_s}^{Origin} = \lambda_{i_s}^{ID} \iff s = s_0 \quad (3)$$

As such, it has a particular operating since, if there were to be one or more *items* thereafter, the last *item* of the subchain will be identified here as its previous item. Therefore, let s_z be the last index in a subchain of linked *items*, we have:

$$\left\{ \begin{array}{l} \lambda_{i_s}^{Origin} = \lambda_{i_s}^{ID} \\ \lambda_{i_s}^{Prev} \neq \emptyset \end{array} \right\} \iff \lambda_{i_s}^{Prev} := l_{i_{s_z}} \quad (4)$$

The last three fields of the leaf therefore transforms s as a circular linked list.

III. IMPLEMENTATION

1. NODE CREATION

Each step j is paired with the j -th prime number in \mathcal{P} , eg. the prime number used is 11 on one step 5^2 .

At run time, a node is either a parent or a leaf, the latter being an end to a branch in the tree.

Definition 5 (Parent node). A parent node is a node containing other nodes, either leaves or other parent nodes.

A leaf is not a parent node by definition.

²We consider step 0 with prime number 1 being the root of the tree T , hence not being counted.

A new parent node must be created every time a representative number walks through the same path as a previous one up to the existing node, extending the branch by one step.

For example, at step k , if a node contains the leaf for $R_k^{l_x}$ for *item* x and if, for a new *item* y , $R_k^{l_y} = R_k^{l_x}$, then both *items* x and y will see their leaf move to step $k + 1$ (where the two new representatives $R_{k+1}^{l_x}$ and $R_{k+1}^{l_y}$ would give the coordinates of each respective leaf). At step k now lies a new node with two children: the leaf λ_x and the leaf λ_y .

Should another *item* z have $R_k^{l_z} = R_k^{l_y} = R_k^{l_x}$ at step k , either a third leaf (λ_z) would be added in $k + 1$, or the leaf with similar path at $k + 1$ would become a new parent node and the two leaves would move to $k + 2$. So on and so forth...

2. PERFORMANCE AND MEMORY

To reduce the initial depth of the tree, one may use a subset $\mathcal{P}' \subset \mathcal{P}$ with the first prime number in \mathcal{P}' being greater than 2, eg. $\mathcal{P}' := [101, 103, 107, 109, \dots]$. This would avoid some expensive initial walkthroughs.

Besides, we can also use other tricks to improve performance.

For example, as seen before, the modulo operation is of complexity $O(1)$ if the number is of fixed size. However, we can optimize this in two different ways since the operations of multiplying and moving bits are much less expensive in number of operations than the operation of division:

- Method 1:

$$P_i \times \left(n - \left\lfloor n \times \frac{1}{P_i} \right\rfloor \right) \quad (5)$$

- Method 2:

$$P_i \times n \gg 32 \quad (6)$$

For method 2, for a n of size less than 256 (8 bytes), we would need 16 bytes. So, the trade-off is about speed vs memory.

Table 1 gives an example of method 2 for a n of 4 bytes.

Table 1: *Memory management*

Octet	1	2	3	4	5	6	7	8
n	–	–	–	–	n	n	n	n
$n \times P_i$	–	–	–	$n \times P_i \gg 32$	$n \times P_i$	$n \times P_i$	$n \times P_i$	$n \times P_i$

3. USAGE

3.1 Write

To add an element to the tree:

- The new leaf is written in the index;
- We update the λ^{Next} field of the leaf that previously corresponded to the last *item* of the subchain;
- We modify the λ^{Prev} field of the leaf of the origin *item* by writing the identifier of the current *item* to the latter.

3.2 Read

To read/search an item in the index:

- We find in the tree the leaf corresponding to the identifier of the searched item:
 - If the λ^{Next} field of the leaf is empty, this is the last item of the subchain;
 - Otherwise, we go to the next step;
- We find the leaf corresponding to the identifier of the field λ^{Origin} ;
- We use the λ^{Prev} field of this leaf to find the last item of the subchain.

When using the index, we can see that we would perform at most 3 reads or 3 writes plus index runs of $O(\log(n))$ order, where n is the number of *items* in the index.

REFERENCES

- [1] Gauss. *Disquisitiones Arithmeticae*, translated by Arthur A. Clarke, Springer, 1986.