

Data structures for a Verkle-powered EVM

Hadrien Croubois - EthCC[7] July 2024

OpenZeppelin's thesis

- There will be a trillion dollar open economy built on blockchains and powered by smart contracts
- This new, open economy will be **built by teams of creative people** developing new applications used by billions of people
- These teams **will need a set of tools, products and services** to make sure that what they are building is **safe and reliable**
- OpenZeppelin will be a **leading provider of these solutions**, allowing teams to **build faster with lower risk**



OpenZeppelin's products





3 - Data structures for a Verkle powered EVM - https://shared.hadriencroubois.com/slides/EthCC2024.pdf





@openzeppelin/contracts@5.0.2

@openzeppelin/contracts-upgradeable@5.0.2





How smart contract state is stored

- Each account comes with a storage trie
- Used for persistent data

ERC-20 balances, ERC-721 ownership, everything

• Storage is divided in slots

Each account has 2²⁵⁶ slots available

- Slots must be "warmed up" before use You only pay the warm-up price once per tx
- All slots are identical



That all changes with the verge

- All storage in the same (verkle) trie
- Position in that shared trie is a combination of account and offset ("slot number")
- Slots are gathered in buckets

For a given account, consecutive slots are in the same bucket (up to 256 slots)

• Buckets are "warmed up", not slots So reusing slots in a given bucket is efficient



Z OpenZeppelin

The grammar of storage layouts (Solidity & Vyper)





Where the grammar fails: Mappings

- Mappings are extensively used in solidity Balances, ownership, allowance, ...
- **Mappings work by pointing to a pseudo-random location** Virtually impossible to get a collision + all slots are equally good
- Data is scattered everywhere

Very poor data locality for verkle, each key point to a different bucket



Where the grammar fails: Arrays

- Arrays have way better data locality then mappings
- Each array access check that length is within bounds Length information is located at the root (different bucket)
- The start of the array is unlikely to coincide with the start of a bucket

Sub-optimal for small arrays, tapers off when the array length increases

	00000000x0
Length	0x00000001
	0x00000002
Item[0]	0xb10e…0cf6
Item[1]	0xb10e…0cf7
Item[2]	0xb10e…0cf8
Item[3]	0xb10e…0cf9
<pre>Item[4]</pre>	0xb10e…0cfa
Item[5]	0xb10e…0cfb
Item[6]	0xb10e…0cfc
Item[7]	0xb10e…0cfd
Item[8]	0xb10e…0cfe
Item[9]	0xb10e0cff
Item[10]	0xb10e0d00

What you can do as a developer: consolidate mappings

- Merge mappings that use the same key Map to a structure that contains all the values
- Easy to do, but will occasionally not work If the key points to the end of a bucket, the structure may span over two buckets
- Limited to some cases

In ERC-20, balances and approvals have different key patterns, so they cannot be consolidated



10 – Data structures for a Verkle powered EVM – https://shared.hadriencroubois.com/slides/EthCC2024.pdf





11 – Data structures for a Verkle powered EVM – https://shared.hadriencroubois.com/slides/EthCC2024.pdf

~ ‡	34 🔳		contracts/token/ERC721/ERC721.sol 🖸
. <u>†</u>		00	-25,13 +25,19 側 abstract contract ERC721 is Context, ERC165, IERC721, IERC721Metadata, IERC721Er
25	25		// Token symbol
26	26		<pre>string private _symbol;</pre>
27	27		
28		-	<pre>mapping(uint256 tokenId => address) private _owners;</pre>
	28	+	struct TokenDetails {
	29	+	address owner;
	30	+	address approval;
	31	+	}
29	32		
30		-	<pre>mapping(address owner => uint256) private _balances;</pre>
	33	+	struct AccountDetails {
	34	+	uint256 balance;
	35	+	<pre>mapping(address => bool) operators;</pre>
04	36	+	}
31	37		manning/uint/255 takanId at address) ariusta FakanAnnravala:
32	20		mapping(uint256 tokenid -> adures) private _tokenaprivats,
22	20	T	mapping(unit200 cokeniu -> tokenbecarts) private _cokens,
3/	33		manning/address numer -> manning/address negrator -> honll) private negratorApprovals.
54	40	+	mapping(address owner => AccountPatalis) private accounts:
35	41		muchania (response and the response of the range in the response of the respon
36	42		/**
37	43		* @dev Initializes the contract by setting a `name` and a `symbol` to the token collection.
÷‡·		00	-58,7 +64,7 @@ abstract contract ERC721 is Context, ERC165, IERC721, IERC721Metadata, IERC721Er
58	64		if (owner == address(0)) {
59	65		<pre>revert ERC721InvalidOwner(address(0));</pre>
60	66		}
61		-	return _balances[owner];
	67	+	return _accounts[owner].balance;
62	68		}
63	69		
64	70		(/**)
+		00	-128,7 +134,7 🝿 abstract contract ERC721 is Context, ERC165, IERC721, IERC721Metadata, IERC721Er
128	134		* @dev See {IERC721-isApprovedForAll}.
129	135		*/
130	136		<pre>function isApprovedForAll(address owner, address operator) public view virtual returns (bool) {</pre>
131		-	return _operatorApprovals[owner][operator];
	137	+	<pre>return _accounts[owner].operators[operator];</pre>
132	138		}
400	400		

TOTAL GAS	EXECUTION 25580	GAS 46%	NON-CODE WITNESS GAS	41%	CODE CHUNK GAS	13%	TOTAL GAS	
Evention							Evention	
Execution							Execution	
EXECUTED INSTRUCTIONS	EXECUTED BYTE	сн 1	IARGED BYTES 5438	EXECU 0.02			EXECUTED INSTRUCTION	NS
Witness charges							Witness charges	
EVENT	GAS	PARAMS					EVENT	GAS
SLOAD	2100	[0x697bcd5513f62773030 0x343ff8127bd64f680be4	0135ca227848c4f499f7e4, e996254dc3528603c6ecd5436	64b4cf956eb	dd28f0028]		SLOAD	2100
SLOAD	2100	[0x697bcd5513f62773030 0xc84cb94819e3894bc65b	0135ca227848c4f499f7e4, b2304132cee4583975460a89f	36ee188487	ef96a9616e4]		SLOAD	200
SSTORE	3500	[0x697bcd5513f62773030 0xc84cb94819e3894bc65b)135ca227848c4f499f7e4, b2304132cee4583975460a89f	36ee188487	ef96a9616e4]		SSTORE	3500
SLOAD	2100	[0x697bcd5513f62773030 0x86a5d9b5c8b0fd7f94fe4)135ca227848c4f499f7e4, 4d996f9c461757fb23594c5cb3	33f827a98c2	59e3bda0]		SLOAD	2100
SSTORE	3500	[0x697bcd5513f62773030 0x86a5d9b5c8b0fd7f94fe4)135ca227848c4f499f7e4, 4d996f9c461757fb23594c5cb3	33f827a98c2	59e3bda0]		SSTORE	3500
SLOAD	2100	[0x697bcd5513f62773030 0x118c1ea466562cb796e)135ca227848c4f499f7e4, 30ef705e4db752f5c39d773d2	2c5efd8d46f	67194e78a]		SLOAD	2100
SSTORE	3500	[0x697bcd5513f62773030 0x118c1ea466562cb796e3)135ca227848c4f499f7e4, 30ef705e4db752f5c39d773d2	2c5efd8d46f	67194e78a]		SSTORE	3500
SSTORE	3500	[0x697bcd5513f62773030 0x343ff8127bd64f680be4	0135ca227848c4f499f7e4, e996254dc3528603c6ecd5436	54b4cf956eb	dd28f0028]		SSTORE	500

[0x63fa0723a30405f10b0aa8c21442d3c3a71595dd,

[0x63fa0723a30405f10b0aa8c21442d3c3a71595dd.

[0x63fa0723a30405f10b0aa8c21442d3c3a71595dd.

[0x63fa0723a30405f10b0aa8c21442d3c3a71595dd,

0x343ff8127bd64f680be4e996254dc3528603c6ecd54364b4cf956ebdd28f0028]

0x343ff8127bd64f680be4e996254dc3528603c6ecd54364b4cf956ebdd28f0029] [0x63fa0723a30405f10b0aa8c21442d3c3a71595dd, 0x343ff8127bd64f680be4e996254dc3528603c6ecd54364b4cf956ebdd28f0029]

0x86a5d9b5c8b0fd7f94fe4d996f9c461757fb23594c5cb33f827a98c259e3bda01

0x86a5d9b5c8b0fd7f94fe4d996f9c461757fb23594c5cb33f827a98c259e3bda0]

0x118c1ea466562cb796e30ef705e4db752f5c39d773d22c5efd8d46f67194e78a] [0x63fa0723a30405f10b0aa8c21442d3c3a71595dd, 0x118c1ea466562cb796e30ef705e4db752f5c39d773d22c5efd8d46f67194e78a]

0x343ff8127bd64f680be4e996254dc3528603c6ecd54364b4cf956ebdd28f0028]

PARAMS

transferFrom on a "Normal ERC-721"

transferFrom on a "Packed ERC-721"



What you can do as a developer: use (custom) arrays

Mapping

```
library HeapMapping {
    using SafeCast for *;
```

```
struct Uint256Heap {
    mapping(uint32 => uint32) tree;
    mapping(uint32 => Node) items;
    uint32 size;
    uint32 nextItemIdx;
}
```

```
struct Node {
    uint256 value;
    uint32 heapIndex; // value -> position
}
```

Array

```
library HeapArray {
    using SafeCast for *;
```

struct Uint256Heap {
 Uint256HeapNode[] data;
}

```
struct Uint256HeapNode {
    uint256 value;
    uint32 index; // position -> value
    uint32 lookup; // value -> position
}
```

```
function _unsafeNodeAccess(
    Uint256Heap storage self,
    uint32 pos
) private pure returns (Uint256HeapNode storage result) {
    assembly ("memory-safe") {
        mstore(0x00, self.slot)
        result.slot := add(keccak256(0x00, 0x20), mul(pos, 2))
    }
}
```

Custom array implementation

```
library HeapArray2 {
    using SafeCast for *;
```

```
struct Uint256Heap {
    bytes32 _placeholder_do_not_use;
}
```

```
struct Uint256HeapNode {
    uint256 value;
    uint32 index; // position -> value
    uint32 lookup; // value -> position
}
```

```
struct Uint256HeapLength {
    uint256 value;
}
function _unsafeNodeAccess(
    Uint256Heap storage self,
    uint32 pos
) private pure returns (Uint256HeapNode storage result) {
    assembly ("memory-safe") {
        mstore(0x00, self.slot)
        result.slot := add(keccak256(0x00, 0x20), add(mul(pos, 2), 1))
    }
}
```

Z OpenZeppelin

```
function _unsafeLengthAccess(
    Uint256Heap storage self
) private pure returns (Uint256HeapLength storage result) {
    assembly ("memory-safe") {
        mstore(0x00, self.slot)
        result.slot := keccak256(0x00, 0x20)
    }
}
```





Z OpenZeppelin

14 – Data structures for a Verkle powered EVM – https://shared.hadriencroubois.com/slides/EthCC2024.pdf

What we ultimately need: compiler support!

Compiler must evolve to optimize for Verkle storage costs ...

- **Replace "keccak256(...)" with "keccak256(...) & ~0xFF"** for Array and mapping derivation
- Move the length of the array to the same "space" as the elements
- Provide in-language mechanism to "extend" a mapping from a parent contract

... but this is breaking

bstra	ct contract ERC721 is Context, ERC165, IERC721, IERC721Metadata, IERC721Errors {
us	ing Strings for uint256;
11	Token name
st	ring private _name;
11	Token symbol
st	ring private _symbol;
ma	<pre>pping(uint256 tokenId => address) private _owners;</pre>
maj	<pre>pping(address owner => uint256) private _balances;</pre>
ma	<pre>pping(uint256 tokenId => address) private _tokenApprovals;</pre>
maj	<pre>pping(address owner => mapping(address operator => bool)) private _operatorApprovals;</pre>
ma	<pre>ct contract EKC/21Enumerable is EKC/21, IEKC/21Enumerable { ppipg(address owner => mapping(wint256 index => wint256)) private ownedTokens:</pre>
maj	<pre>pping(uint256 tokenId => uint256) private _ownedTokensIndex;</pre>
ui	nt256[] private allTokens:
mai	<pre>pping(uint256 tokenId => uint256) private _allTokensIndex;</pre>
_	
hetra	ct contract ERC721URTStorage is TERC4006 ERC721 /
us:	ing Strings for uint256;
	Interface ID as defined in ERC-4906. This does not correspond to a traditional interfa
11	
11	defines events and does not include any external function.
// // by	<pre>defines events and does not include any external function. tes4 private constant ERC4906_INTERFACE_ID = bytes4(0x49064906);</pre>
// // by	<pre>defines events and does not include any external function. tes4 private constant ERC4996_INTERFACE_ID = bytes4(0x49964996); Optional mapping for token URIs</pre>

Z OpenZeppelin

@openzeppelin/contracts
 docs.openzeppelin.com
 forum.openzeppelin.com

Thank you for your attention





Source code & results



POAP