A Probabilistic Model Revealing Shortcomings in Lua's Hybrid Tables

Pablo Rotondo

LIGM, Université Gustave Eiffel

Joint work with Conrado Martínez (UPC), Cyril Nicaud (LIGM)

Séminaire LIGM, Champs-Sur-Marne, 4 April, 2023.

Introduction

- Aim: study and model actual implementations
 - Engineers sometimes choose innovative implementations e.g., TimSort in Python.
 - Study choices in depth, make recommendations.

Introduction

- Aim: study and model actual implementations
 - Engineers sometimes choose innovative implementations e.g., TimSort in Python.
 - Study choices in depth, make recommendations.
- The Lua programming language
 - Scripting language widely used in the gaming industry,
 - Efficient, lightweight (few Kb of C code!), embeddable.

Introduction

- Aim: study and model actual implementations
 - Engineers sometimes choose innovative implementations e.g., TimSort in Python.
 - Study choices in depth, make recommendations.
- The Lua programming language
 - Scripting language widely used in the gaming industry,
 - Efficient, lightweight (few Kb of C code!), embeddable.
- \Rightarrow Lua 5.0 introduced several innovations, among them a new Table structure.

The Lua programming language¹



What is Lua?

Lua is a powerful, efficient, lightweight, embeddable scripting language. It supports procedural programming, object-oriented programming, functional programming, data-driven programming, and data description.

Lua combines simple procedural syntax with powerful data description constructs based on associative arrays and extensible semantics. Lua is dynamically typed, runs by interpreting bytecode with a register-based virtual machine, and has automatic memory management with incremental garbage collection, making it ideal for configuration, scripting, and rapid prototyping.

Why choose Lua?

Lua is a proven, robust language

Lua has been used in many industrial applications (e.g., Adobe's Photoshop Lightroom), with an emphasis on embedded systems (e.g., the Ginga middleware for digital TV in Brazil) and games (e.g., World of Warcraft and Angry Birds). Lua is currently the leading scripting language in games. Lua has a solid reference manual and there are several books about it. Several versions of Lua have been released and used in real applications since its creation in 1932. Lua featured in HOPL III, the Third ACM SIGPLAN History of Programming Languages Conference, in 2007. Lua won the Front Line Award 2011 from the Game Developers Magazine.

¹Copyright © 1994–2022 Lua.org, PUC-Rio.

The Lua programming language¹



What is Lua?

Lua is a powerful, efficient, lightweight, embeddable scripting language. It supports procedural programming, object-oriented programming, functional programming, data-driven programming, and data description.

Lua combines simple procedural syntax with powerful data description constructs based on associative arrays and extensible semantics. Lua is dynamically typed, runs by interpreting bytecode with a register-based virtual machine, and has automatic memory management with incremental garbage collection, making it ideal for configuration, scripting, and rapid prototyping.

Why choose Lua?

Lua is a proven, robust language

Lua has been used in many industrial applications (e.g., Adobe's Photoshop Lightroom), with an emphasis on embedded systems (e.g., the Ginga middleware for digital TV in Brazil) and games (e.g., World of Warcraft and Angry Birds). Lua is currently the leading scripting language in games. Lua has a solid reference manual and there are several books about it. Several versions of Lua have been released and used in real applications since its creation in 1932. Lua featured in HOPL III, the Third ACM SIGPLAN History of Programming Languages Conference, in 2007. Lua won the Front Line Award 2011 from the Game Developers Magazine.

Lua is fast

Lua has a deserved reputation for performance. To claim to be "as fast as Lua" is an aspiration of other scripting languages. Several benchmarks show Lua as the fastest language in the realm of interpreted scripting languages. Lua is fast not only in fine-tuned benchmark programs, but in real life too. Substantial fractions of large applications have been written in Lua.

If you need even more speed, try LuaJIT, an independent implementation of Lua using a just-in-time compiler.

¹Copyright © 1994–2022 Lua.org, PUC-Rio.

- Only data-structuring mechanism in Lua
 - assignment H[x]=y, any types of x and y.

- Only data-structuring mechanism in Lua
 - assignment H[x]=y, any types of x and y.
- Implementation
 - originally a simple hash-table up to Lua 4.

- Only data-structuring mechanism in Lua
 - assignment H[x]=y, any types of x and y.
- Implementation
 - originally a simple hash-table up to Lua 4.
 - Lua 5 introduced a hybrid hash-array implementation,

- Only data-structuring mechanism in Lua
 - assignment H[x]=y, any types of x and y.
- Implementation
 - originally a simple hash-table up to Lua 4.
 - Lua 5 introduced a hybrid hash-array implementation,

The implementation of tables in Lua involves some clever algorithms. Every table in Lua has two parts: the *array* part and the *hash* part. The array part stores entries with integer keys in the range 1 to n, for some particular n. (We will discuss how this n is computed in a moment.) All other entries (including integer keys outside that range) go to the hash part.

Figure: Extract from the book Lua Programming Gems.

- Only data-structuring mechanism in Lua
 - assignment H[x]=y, any types of x and y.
- Implementation
 - originally a simple hash-table up to Lua 4.
 - Lua 5 introduced a hybrid hash-array implementation,

The implementation of tables in Lua involves some clever algorithms. Every table in Lua has two parts: the *array* part and the *hash* part. The array part stores entries with integer keys in the range 1 to n, for some particular n. (We will discuss how this n is computed in a moment.) All other entries (including integer keys outside that range) go to the hash part.

Figure: Extract from the book Lua Programming Gems.

In our work we study this mechanism.

Running time:

Running time:

insertions and lookups work in amortized O(1)
even if table is full.

Running time:

insertions and lookups work in amortized O(1)

even if table is full.

• we show there is a degradation if deletions are allowed.

Running time:

insertions and lookups work in amortized O(1)

even if table is full.

• we show there is a degradation if deletions are allowed.

Consider sequences of T insertions/deletions from empty table

Running time:

insertions and lookups work in amortized O(1)

even if table is full.

• we show there is a degradation if deletions are allowed.

Consider sequences of \boldsymbol{T} insertions/deletions from empty table

Proposition: worst case

There is sequence of operations giving time $\Theta(T^2)$.

Running time:

insertions and lookups work in amortized O(1)

even if table is full.

• we show there is a degradation if deletions are allowed.

Consider sequences of \boldsymbol{T} insertions/deletions from empty table

Proposition: worst case

There is sequence of operations giving time $\Theta(T^2)$.



• Example requires an unlikely cycle of delete-insert.

Running time:

insertions and lookups work in amortized O(1)

even if table is full.

• we show there is a degradation if deletions are allowed.

Consider sequences of \boldsymbol{T} insertions/deletions from empty table

Proposition: worst case

There is sequence of operations giving time $\Theta(T^2)$.



- Example requires an unlikely cycle of delete-insert.
- A problem for more realistic scenarios ?

Our Simple Probabilistic model

Consider $p > \frac{1}{2}$. A random sequence of T insertion/deletions:

- with probability p insert a new element,
- with probability 1 p delete an element.

Our Simple Probabilistic model

Consider $p > \frac{1}{2}$. A random sequence of T insertion/deletions:

- with probability p insert a new element,
- with probability 1 p delete an element.



Main result: Lua hash-table With high probability, complexity is $\Omega(T \log T)$.

Our Simple Probabilistic model

Consider $p > \frac{1}{2}$. A random sequence of T insertion/deletions:

- with probability p insert a new element,
- with probability 1 p delete an element.



Main result: Lua hash-table With high probability, complexity is $\Omega(T \log T)$.

٢

Our Simple Probabilistic model

Consider $p > \frac{1}{2}$. A random sequence of T insertion/deletions:

- with probability p insert a new element,
- with probability 1 p delete an element.



Main result: Lua hash-tableWith high probability, complexity is $\Omega(T \log T)$. \odot

We propose potential fixes...

Our Simple Probabilistic model

Consider $p > \frac{1}{2}$. A random sequence of T insertion/deletions:

- with probability p insert a new element,
- with probability 1 p delete an element.



Main result: Lua hash-tableWith high probability, complexity is $\Omega(T \log T)$.

We propose potential fixes... which we will see and implement during the talk

Plan of the talk

- 1. The Lua hashmap
- 2. The probabilistic model
- 3. Hybrid Tables and insertions
- 4. Conclusions and further work

Problem: represent set from large universe \mathcal{K} (e.g., IP addresses)

Minimal interface: init, insert, lookup, delete

Problem: represent set from large universe \mathcal{K} (e.g., IP addresses)

- Minimal interface: init, insert, lookup, delete
- Implementations: LinkedList (☞), Array (☞), balanced tree (☺),

Problem: represent set from large universe \mathcal{K} (e.g., IP addresses)

- Minimal interface: init, insert, lookup, delete
- ► Implementations: LinkedList (☞), Array (☞), balanced tree (☺), hash-tables (★)

Problem: represent set from large universe \mathcal{K} (e.g., IP addresses)

- Minimal interface: init, insert, lookup, delete
- Implementations: LinkedList (☞), Array (☞), balanced tree (☺), hash-tables (★)

Hash-tables solve the problem of the array "solution"

- array H of size M much smaller than $|\mathcal{K}|$,
- hash function $h: \mathcal{K} \to \mathbb{Z}_{\geq 0}$,
- store $x \in \mathcal{K}$ in slot $H[h(x) \mod M]$.

Problem: represent set from large universe \mathcal{K} (e.g., IP addresses)

- Minimal interface: init, insert, lookup, delete
- Implementations: LinkedList (आ), Array (आ), balanced tree (☺), hash-tables (★)

Hash-tables solve the problem of the array "solution"

- array H of size M much smaller than $|\mathcal{K}|$,
- hash function $h: \mathcal{K} \to \mathbb{Z}_{\geq 0}$,
- store $x \in \mathcal{K}$ in slot $H[h(x) \mod M]$.

... but there are surely many y with $h(x) \equiv h(y) \mod M$?

Problem: represent set from large universe \mathcal{K} (e.g., IP addresses)

- Minimal interface: init, insert, lookup, delete
- Implementations: LinkedList (☞), Array (☞), balanced tree (☺), hash-tables (★)

Hash-tables solve the problem of the array "solution"

- array H of size M much smaller than $|\mathcal{K}|$,
- hash function $h: \mathcal{K} \to \mathbb{Z}_{\geq 0}$,
- store $x \in \mathcal{K}$ in slot $H[h(x) \mod M]$.

... but there are surely many y with $h(x) \equiv h(y) \mod M$? \Rightarrow collision

Problem: represent set from large universe \mathcal{K} (e.g., IP addresses)

- Minimal interface: init, insert, lookup, delete
- Implementations: LinkedList (☞), Array (☞), balanced tree (☺), hash-tables (★)

Hash-tables solve the problem of the array "solution"

- array H of size M much smaller than $|\mathcal{K}|$,
- hash function $h: \mathcal{K} \to \mathbb{Z}_{\geq 0}$,
- store $x \in \mathcal{K}$ in slot $H[h(x) \mod M]$.

... but there are surely many y with $h(x) \equiv h(y) \mod M$? \Rightarrow collision

Hash-functions must

- avoid collisions as much as possible,
- be fast to compute.

We still have to decide what to do in case of a collision:

- ▶ Separate chaining: use a linked-list at each slot *H*[*i*],
- ▶ Internal chaining: put key somewhere else (where?) in *H*.

We still have to decide what to do in case of a collision:

- ▶ Separate chaining: use a linked-list at each slot *H*[*i*],
- ▶ Internal chaining: put key somewhere else (where?) in *H*.

but what if

- Separate chaining: linked-lists are very long ?
- ▶ Internal chaining: array *H* is full ?

We still have to decide what to do in case of a collision:

- ▶ Separate chaining: use a linked-list at each slot *H*[*i*],
- ▶ Internal chaining: put key somewhere else (where?) in *H*.

but what if

- Separate chaining: linked-lists are very long ?
- Internal chaining: array H is full ?
- \dots rehash into larger array H'

Hashmap mechanism

Lua's hashmap consists of

• an array H of size $M = 2^m$,

Hashmap mechanism

Lua's hashmap consists of

- an array H of size $M = 2^m$,
- entries: key, value, index of next entry in chain
Hashmap mechanism

Lua's hashmap consists of

- an array H of size $M = 2^m$,
- entries: key, value, index of next entry in chain



 \Rightarrow internal chaining

Hashmap mechanism

Lua's hashmap consists of

- an array H of size $M = 2^m$,
- entries: key, value, index of next entry in chain



⇒ internal chaining

• for x integer, hash function $h(x) = x \mod M$,

Hashmap mechanism

Lua's hashmap consists of

- an array H of size $M = 2^m$,
- entries: key, value, index of next entry in chain



⇒ internal chaining

• for x integer, hash function $h(x) = x \mod M$,

516 SEARCHING

6.4

when K is odd, and this will lead to a substantial bias in many files. It would be even worse to let M be a power of the radix of the computer, since $K \mod M$ would then be simply the least significant digits of K (independent of the other digits). Similarly we can argue that M probably shouldn't be a multiple of 3;

Figure: Extract from The Art of Computer Programming (vol. 3)

Hashmap mechanism: hash function

Not immediate that h may be a bad choice

- fast for integers: $x \mod M = x \& (M-1)$,
- more involved for strings:

```
unsigned int luaS_hash (const char *str, size_t l,
    unsigned int seed) {
    unsigned int h = seed ^ cast_uint(l);
    for (; l > 0; l--)
        h ^= ((h<<5) + (h>>2) + cast_byte(str[l - 1]))
        ;
    return h;
}
```

Hashmap mechanism: hash function

Not immediate that h may be a bad choice

- fast for integers: $x \mod M = x \& (M-1)$,
- more involved for strings:

```
unsigned int luaS_hash (const char *str, size_t l,
    unsigned int seed) {
    unsigned int h = seed ^ cast_uint(l);
    for (; l > 0; l--)
        h ^= ((h<<5) + (h>>2) + cast_byte(str[l - 1]))
        ;
    return h;
}
```

In this talk we do not discuss the choice of the hash-function h,

Base assumption

The function $x \mapsto h(x) \mod M$ is roughly uniform

Insertions work as follows: we want to insert key x

• if position h(x) free \Rightarrow insert



h(a) = 4, h(b) = 4,

Insertions work as follows: we want to insert key x

• if position h(x) free \Rightarrow insert

... else position h(x) is occupied by key y,

• if $h(y) = h(x) \Rightarrow put x$ into a free position, update chain

$$pos(y) \rightarrow pos(z) \rightarrow \dots$$
 to $pos(y) \rightarrow pos(z) \rightarrow pos(z) \rightarrow \dots$



$$h(a) = 4, h(b) = 4, h(d) = 4,$$

Insertions work as follows: we want to insert key \boldsymbol{x}

• if position h(x) free \Rightarrow insert

... else position h(x) is occupied by key y,

• if $h(y) = h(x) \Rightarrow put x$ into a free position, update chain

$$pos(y) \rightarrow pos(z) \rightarrow \dots$$
 to $pos(y) \rightarrow pos(z) \rightarrow pos(z) \rightarrow \dots$



$$h(a) = 4, h(b) = 4, h(d) = 4, h(e) = 7,$$

Insertions work as follows: we want to insert key x

• if position h(x) free \Rightarrow insert

... else position h(x) is occupied by key y,

• if $h(y) = h(x) \Rightarrow$ put x into a free position, update chain

 $pos(y) \rightarrow pos(z) \rightarrow \dots$ to $pos(y) \rightarrow pos(z) \rightarrow pos(z) \rightarrow \dots$

if h(y) ≠ h(x) ⇒ we migrate y into a free position, updating its chain and put x at position h(x).



$$h(a) = 4, h(b) = 4, h(d) = 4, h(e) = 7, h(f) = 7$$

Insertions work as follows: we want to insert key x

• if position h(x) free \Rightarrow insert

... else position h(x) is occupied by key y,

• if $h(y) = h(x) \Rightarrow$ put x into a free position, update chain

 $pos(y) \rightarrow pos(z) \rightarrow \dots$ to $pos(y) \rightarrow pos(x) \rightarrow pos(z) \rightarrow \dots$

if h(y) ≠ h(x) ⇒ we migrate y into a free position, updating its chain and put x at position h(x).



$$h(a) = 4, h(b) = 4, h(d) = 4, h(e) = 7, h(f) = 7$$

Insertions work as follows: we want to insert key x

• if position h(x) free \Rightarrow insert

... else position h(x) is occupied by key y,

• if $h(y) = h(x) \Rightarrow$ put x into a free position, update chain

 $pos(y) \rightarrow pos(z) \rightarrow \dots$ to $pos(y) \rightarrow pos(z) \rightarrow pos(z) \rightarrow \dots$

if h(y) ≠ h(x) ⇒ we migrate y into a free position, updating its chain and put x at position h(x).



h(a) = 4, h(b) = 4, h(d) = 4, h(e) = 7, h(f) = 7

Insertions work as follows: we want to insert key x

• if position h(x) free \Rightarrow insert

... else position h(x) is occupied by key y,

• if $h(y) = h(x) \Rightarrow$ put x into a free position, update chain

 $pos(y) \rightarrow pos(z) \rightarrow \dots$ to $pos(y) \rightarrow pos(z) \rightarrow pos(z) \rightarrow \dots$

if h(y) ≠ h(x) ⇒ we migrate y into a free position, updating its chain and put x at position h(x).



h(a) = 4, h(b) = 4, h(d) = 4, h(e) = 7, h(f) = 7

Finding free position: pointer from right to left.

Insertions work as follows: we want to insert key x

• if position h(x) free \Rightarrow insert

... else position h(x) is occupied by key y,

• if $h(y) = h(x) \Rightarrow$ put x into a free position, update chain

 $pos(y) \rightarrow pos(z) \rightarrow \dots$ to $pos(y) \rightarrow pos(z) \rightarrow pos(z) \rightarrow \dots$

if h(y) ≠ h(x) ⇒ we migrate y into a free position, updating its chain and put x at position h(x).



h(a) = 4, h(b) = 4, h(d) = 4, h(e) = 7, h(f) = 7

Finding free position: pointer from right to left. If pointer exits, rehash.

Deletions are simple

value is marked as nil : equivalent to H[x] = nil,

Deletions are simple

- value is marked as nil : equivalent to H[x] = nil,
- coherent with semantics: if $x \notin H \Rightarrow H[x]$ evaluates nil,

Deletions are simple

- value is marked as nil : equivalent to H[x] = nil,
- coherent with semantics: if $x \notin H \Rightarrow H[x]$ evaluates nil,
- chaining (next cell) kept intact.

Deletions are simple

- value is marked as nil : equivalent to H[x] = nil,
- coherent with semantics: if $x \notin H \Rightarrow H[x]$ evaluates nil,
- chaining (next cell) kept intact.

Deleted position k

• can be reused to insert x when h(x) = k,

Deletions are simple

- value is marked as nil : equivalent to H[x] = nil,
- coherent with semantics: if $x \notin H \Rightarrow H[x]$ evaluates nil,
- chaining (next cell) kept intact.

Deleted position k

- can be reused to insert x when h(x) = k,
- not taken into account by free position pointer
 mecessary to keep previous chaining

Deletions are simple

- value is marked as nil : equivalent to H[x] = nil,
- coherent with semantics: if $x \notin H \Rightarrow H[x]$ evaluates nil,
- chaining (next cell) kept intact.

Deleted position k

- can be reused to insert x when h(x) = k,
- not taken into account by free position pointer
 mecessary to keep previous chaining

... deleted spots are cleaned up during rehashing

Finding a free position: use pointer starting at end and moving to the left. If pointer exits, we rehash.

• Hashtable is then full, maybe with deleted cells.

Finding a free position: use pointer starting at end and moving to the left. If pointer exits, we rehash.

- Hashtable is then full, maybe with deleted cells.
- Count actually used cells n,

Finding a free position: use pointer starting at end and moving to the left. If pointer exits, we rehash.

- Hashtable is then full, maybe with deleted cells.
- Count actually used cells n,
- ▶ Set size $M = 2^m$, with smallest m such that $n + 1 \le 2^m$, $\implies +1$ for inserted element.

Finding a free position: use pointer starting at end and moving to the left. If pointer exits, we rehash.

- Hashtable is then full, maybe with deleted cells.
- Count actually used cells n,
- ▶ Set size $M = 2^m$, with smallest m such that $n + 1 \le 2^m$, $\implies +1$ for inserted element.

Worst-case scenario construction

- insert until filling hashtable of size $M = 2^m$,
- ▶ do *M* iterations: (deletion+insertion),
- insertions induce rehash unless deleted cell is picked (p = 1/M)

Expected complexity $\Theta(M^2)$ for 3M operations.

Finding a free position: use pointer starting at end and moving to the left. If pointer exits, we rehash.

- Hashtable is then full, maybe with deleted cells.
- Count actually used cells n,
- ▶ Set size $M = 2^m$, with smallest m such that $n + 1 \le 2^m$, $\implies +1$ for inserted element.

Worst-case scenario construction

- insert until filling hashtable of size $M = 2^m$,
- ▶ do *M* iterations: (deletion+insertion),
- insertions induce rehash unless deleted cell is picked (p = 1/M)

Expected complexity $\Theta(M^2)$ for 3M operations.

... but it is not very realistic, users do not behave this way (?)

We set a more interesting yet simple model

Probabilistic model

Fix $p > \frac{1}{2}$ and apply T insertion/deletions from an empty table:

- \blacktriangleright with probability p insert a new element,
- with probability 1-p delete an element among present ones.

We set a more interesting yet simple model

Probabilistic model

Fix $p > \frac{1}{2}$ and apply T insertion/deletions from an empty table:

- with probability p insert a new element,
- with probability 1 p delete an element among present ones.

Hashtable tends to grow: # keys $\approx pT - (1-p)T = (2p-1)T$

We set a more interesting yet simple model

Probabilistic model

Fix $p > \frac{1}{2}$ and apply T insertion/deletions from an empty table:

- with probability p insert a new element,
- with probability 1 p delete an element among present ones.

Hashtable tends to grow: # keys $\approx pT - (1-p)T = (2p-1)T$

Theorem (Martínez, Nicaud, R 2022)

With high probability, Lua uses $\Omega(T \log T)$ time for this process.

We set a more interesting yet simple model

Probabilistic model

Fix $p > \frac{1}{2}$ and apply T insertion/deletions from an empty table:

- with probability p insert a new element,
- with probability 1 p delete an element among present ones.

Hashtable tends to grow: # keys $\approx pT - (1-p)T = (2p-1)T$

Theorem (Martínez, Nicaud, R 2022)

With high probability, Lua uses $\Omega(T \log T)$ time for this process.

Intuition: Large number of useless rehashes

We set a more interesting yet simple model

Probabilistic model

Fix $p > \frac{1}{2}$ and apply T insertion/deletions from an empty table:

- with probability p insert a new element,
- with probability 1 p delete an element among present ones.

Hashtable tends to grow: # keys $\approx pT - (1-p)T = (2p-1)T$

Theorem (Martínez, Nicaud, R 2022)

With high probability, Lua uses $\Omega(T \log T)$ time for this process.

- Intuition: Large number of useless rehashes
- Each rehash costs linear time $\Theta(M)$.

٢

When rehashing table of size ${\cal M}$

free position pointer does not take into account deleted cells,

When rehashing table of size ${\cal M}$

- free position pointer does not take into account deleted cells,
- if $0 < \delta < M/2$ deleted cells remain,

When rehashing table of size ${\cal M}$

- free position pointer does not take into account deleted cells,
- if $0 < \delta < M/2$ deleted cells remain,

... new hashtable has same size,

When rehashing table of size ${\cal M}$

- free position pointer does not take into account deleted cells,
- if $0 < \delta < M/2$ deleted cells remain,

 \ldots new hashtable has <u>same size</u>, even if δ small

When rehashing table of size ${\cal M}$

- free position pointer does not take into account deleted cells,
- if $0 < \delta < M/2$ deleted cells remain,

 \ldots new hashtable has <u>same size</u>, even if δ small

 \ldots and we have f = $\delta-1$ free cells after rehash

When rehashing table of size ${\cal M}$

- free position pointer does not take into account deleted cells,
- if $0 < \delta < M/2$ deleted cells remain,

 \ldots new hashtable has <u>same size</u>, even if δ small

... and we have $f = \delta - 1$ free cells after rehash

Example: Before and after rehash (insertion -18), $\delta = 1$



When rehashing table of size ${\cal M}$

- free position pointer does not take into account deleted cells,
- if $0 < \delta < M/2$ deleted cells remain,

 \ldots new hashtable has same size, even if δ small

 \ldots and we have f = $\delta-1$ free cells after rehash

Example: Before and after rehash (insertion -18), $\delta = 1$


Intuition: useless rehash

When rehashing table of size ${\cal M}$

- free position pointer does not take into account deleted cells,
- if $0 < \delta < M/2$ deleted cells remain,

 \ldots new hashtable has same size, even if δ small

 \ldots and we have f = $\delta-1$ free cells after rehash

Example: Before and after rehash (insertion -18), $\delta = 1$



such rehash is not of much use: only cleans (few) deleted cells

Intuition: useless rehash

When rehashing table of size ${\cal M}$

- free position pointer does not take into account deleted cells,
- if $0 < \delta < M/2$ deleted cells remain,

 \ldots new hashtable has same size, even if δ small

 \ldots and we have f = $\delta-1$ free cells after rehash

Example: Before and after rehash (insertion -18), $\delta = 1$



such rehash is not of much use: only cleans (few) deleted cells

Intuition: useless rehash

When rehashing table of size ${\cal M}$

- free position pointer does not take into account deleted cells,
- if $0 < \delta < M/2$ deleted cells remain,

 \ldots new hashtable has same size, even if δ small

 \ldots and we have f = $\delta-1$ free cells after rehash

Example: Before and after rehash (insertion -18), $\delta = 1$



such rehash is not of much use: only cleans (few) deleted cells

 \ldots useless rehashes go on until we hit a rehash with δ = 0

Theorem (Martínez, Nicaud, R 2022)

With high probability, Lua uses $\Omega(T\log T)$ time for this process.

Number of keys in hashmap after t operations ≈ (2p − 1)t, with high probability size M only increases

Theorem (Martínez, Nicaud, R 2022)

With high probability, Lua uses $\Omega(T \log T)$ time for this process.

- Number of keys in hashmap after t operations $\approx (2p-1)t$, with high probability size M only increases
- At some point rehash into size $M = 2^m$ of order $\Theta(T)$: $\implies \# \text{keys} = 2^{m-1} + 1 \text{ and } 2^{m-1} - 1 \text{ free spots.}$

Theorem (Martínez, Nicaud, R 2022)

With high probability, Lua uses $\Omega(T\log T)$ time for this process.

- ▶ Number of keys in hashmap after t operations $\approx (2p-1)t$, with high probability size M only increases
- At some point rehash into size $M = 2^m$ of order $\Theta(T)$: $\implies \# \text{keys} = 2^{m-1} + 1 \text{ and } 2^{m-1} - 1 \text{ free spots.}$

• But then random deletions slow down growth of M:

 \implies from f free spots, we obtain γf after next rehash.

Theorem (Martínez, Nicaud, R 2022)

With high probability, Lua uses $\Omega(T \log T)$ time for this process.

- ▶ Number of keys in hashmap after t operations $\approx (2p-1)t$, with high probability size M only increases
- At some point rehash into size $M = 2^m$ of order $\Theta(T)$: $\implies \# \text{keys} = 2^{m-1} + 1 \text{ and } 2^{m-1} - 1 \text{ free spots.}$
- But then random deletions slow down growth of M: \implies from f free spots, we obtain γf after next rehash.

Lemma

If hashmap has size M and just after a rehash it contains $f \gg \sqrt{M}$ free spots, at the next rehash it still has size M and contains at least γf free spots (whp).

Theorem (Martínez, Nicaud, R 2022)

With high probability, Lua uses $\Omega(T \log T)$ time for this process.

- ▶ Number of keys in hashmap after t operations $\approx (2p-1)t$, with high probability size M only increases
- At some point rehash into size $M = 2^m$ of order $\Theta(T)$: $\implies \# \text{keys} = 2^{m-1} + 1 \text{ and } 2^{m-1} - 1 \text{ free spots.}$
- But then random deletions slow down growth of M: \implies from f free spots, we obtain γf after next rehash.

Lemma

If hashmap has size M and just after a rehash it contains $f \gg \sqrt{M}$ free spots, at the next rehash it still has size M and contains at least γf free spots (whp).

 \ldots at least $\log M$ rehashes to increase M

With (very) high probability:

- the hashtable is never empty after t = 0,
- we rehash at some point.

With (very) high probability:

- the hashtable is never empty after t = 0,
- we rehash at some point.

Under these *simplifying assumptions*, between two rehashes, the number of deleted cells satisfies the recurrence (starting from $\delta_{t_0} = 0$)

$$\delta_{t+1} = \begin{cases} \delta_t - 1 & \text{with probability } \frac{p\delta_t}{M} & [\text{insertion at deleted key}], \\ \delta_t & \text{with probability } p\left(1 - \frac{\delta_t}{M}\right) & [\text{insertion at free cell}], \\ \delta_t + 1 & \text{with probability } 1 - p & [\text{deletion}]. \end{cases}$$

$$\delta_{t+1} = \begin{cases} \delta_t - 1 & \text{with probability } \frac{p\delta_t}{M} & [\text{insertion at deleted key}], \\ \delta_t & \text{with probability } p\left(1 - \frac{\delta_t}{M}\right) & [\text{insertion at free cell}], \\ \delta_t + 1 & \text{with probability } 1 - p & [\text{deletion}]. \end{cases}$$

► Equilibrium point at
$$\delta_t \approx \frac{1-p}{p}M$$
,
 \circledast when $\delta_t < \frac{1-p}{p}M$ tendency to increase,
 \circledast when $\delta_t > \frac{1-p}{p}M$ tendency to decrease,

$$\delta_{t+1} = \begin{cases} \delta_t - 1 & \text{with probability } \frac{p\delta_t}{M} & [\text{insertion at deleted key}], \\ \delta_t & \text{with probability } p\left(1 - \frac{\delta_t}{M}\right) & [\text{insertion at free cell}], \\ \delta_t + 1 & \text{with probability } 1 - p & [\text{deletion}]. \end{cases}$$

► Equilibrium point at
$$\delta_t \approx \frac{1-p}{p}M$$
,
 \circledast when $\delta_t < \frac{1-p}{p}M$ tendency to increase,
 \circledast when $\delta_t > \frac{1-p}{p}M$ tendency to decrease,

- Rehash occurs before reaching equilibrium
 - \ldots at the beginning δ_t increases linearly
 - ... as we approach equilibrium, increase weakens

Evolution of number of deleted cells δ_t : linear increase



Evolution of number of deleted cells δ_t : linear increase



Evolution of number of deleted cells δ_t : linear increase



 $\delta_{t+1} = \begin{cases} \delta_t - 1 & \text{with probability } \frac{p\delta_t}{M} & [\text{insertion at deleted key}], \\ \delta_t & \text{with probability } p\left(1 - \frac{\delta_t}{M}\right) & [\text{insertion at free cell}], \\ \delta_t + 1 & \text{with probability } 1 - p & [\text{deletion}]. \end{cases}$

Let free₀ free (unused) cells after last rehash (set t = 0)

Remark: there is a linear regime for a proportion of time After $t = \lfloor \frac{1-p}{p} \texttt{free}_0 \rfloor < \texttt{free}_0$ steps $\frac{p\delta_t}{M} \leq (1-p) \frac{\texttt{free}_0}{M} \leq \frac{1-p}{2},$ $\Delta \delta_t = 1$ still twice as likely as $\Delta \delta_t = -1$.

Evolution of number of deleted cells δ_t : stopping time

- By time $t = \lfloor \frac{1-p}{p} \texttt{free}_0 \rfloor$, number of deleted cells δ_t increased linearly.
- It remains to show that it will not decrease

Evolution of number of deleted cells δ_t : stopping time

- By time $t = \lfloor \frac{1-p}{p} \texttt{free}_0 \rfloor$, number of deleted cells δ_t increased linearly.
- It remains to show that it will not decrease

Lemma

If $c > \frac{1}{2p-1},$ number of operations before next rehash τ satisfies

 $\tau \leq c \cdot \texttt{free}_0$

with high probability.

... and so only a linear number of steps remain

Evolution of number of deleted cells δ_t : stopping time

- By time $t = \lfloor \frac{1-p}{p} \texttt{free}_0 \rfloor$, number of deleted cells δ_t increased linearly.
- It remains to show that it will not decrease

Lemma

If $c > \frac{1}{2p-1},$ number of operations before next rehash τ satisfies

 $\tau \leq c \cdot \texttt{free}_0$

with high probability.

... and so only a linear number of steps remain

Lemma

With high probability equilibrium has never been reached by time τ .

... and at worse δ_t looks like a $\frac{1}{2} - \frac{1}{2}$ random walk (close to τ)

Potential solutions:

implement real deletions

Potential solutions:

implement real deletions

... requires several changes (update chains, list of available cells)

Potential solutions:

- implement real deletions
 - ... requires several changes (update chains, list of available cells)
- keep a minimum proportion of free cells after rehashing

Potential solutions:

implement real deletions

... requires several changes (update chains, list of available cells)

keep a minimum proportion of free cells after rehashing

... we will implement it now in 2 minutes for $\geq 20\%$

Potential solutions:

implement real deletions

... requires several changes (update chains, list of available cells)

keep a minimum proportion of free cells after rehashing

... we will implement it now in 2 minutes for $\geq 20\%$

Potential solutions:

implement real deletions

... requires several changes (update chains, list of available cells)

keep a minimum proportion of free cells after rehashing

... we will implement it now in 2 minutes for $\geq 20\%$



Potential solutions:

implement real deletions

... requires several changes (update chains, list of available cells)

keep a minimum proportion of free cells after rehashing

... we will implement it now in 2 minutes for $\geq 20\%$



Ensuring a proportion $\beta \in (0,1)$ of empty cells

• we require at least βM operations before rehashing again

Potential solutions:

implement real deletions

... requires several changes (update chains, list of available cells)

keep a minimum proportion of free cells after rehashing

... we will implement it now in 2 minutes for $\geq 20\%$



Ensuring a proportion $\beta \in (0,1)$ of empty cells

- we require at least βM operations before rehashing again
- amortized #insertions per operation $\leq (M + \beta M)/(\beta M) = 1 + \beta^{-1}$

 \circledast Rehashes into same size hashtables pile up to $\Omega(T\log T)$,

- \circledast Rehashes into same size hashtables pile up to $\Omega(T\log T)$,
- \otimes More precisely it is $\Theta(T \log T)$.

- \circledast Rehashes into same size hashtables pile up to $\Omega(T\log T)$,
- \otimes More precisely it is $\Theta(T \log T)$.

And without deletions?

& Problem arises when considering effects of deletions.

- \circledast Rehashes into same size hashtables pile up to $\Omega(T\log T)$,
- \otimes More precisely it is $\Theta(T \log T)$.

And without deletions?

- Problem arises when considering effects of deletions.
- The hybrid data-structure presents a similar issue: using the array-part "simulates" deletions on the hash-part

- \circledast Rehashes into same size hashtables pile up to $\Omega(T\log T)$,
- \otimes More precisely it is $\Theta(T \log T)$.

And without deletions?

- Problem arises when considering effects of deletions.
- The hybrid data-structure presents a similar issue: using the array-part "simulates" deletions on the hash-part

Proposition [only insertions]

Inserting n elements into Lua's table takes $\Theta(n \log n)$ in the worst case.

- \circledast Rehashes into same size hashtables pile up to $\Omega(T\log T)$,
- \otimes More precisely it is $\Theta(T \log T)$.

And without deletions?

- Problem arises when considering effects of deletions.
- The hybrid data-structure presents a similar issue: using the array-part "simulates" deletions on the hash-part

Proposition [only insertions]

Inserting n elements into Lua's table takes $\Theta(n \log n)$ in the worst case.

Example: inserting
$$-(2^{k}-1), -(2^{k}-2), \dots, -1, 0, 1, \dots, 2^{k}$$

• array-part corresponds to interval [1, n], $n = 2^{j}$,

- array-part corresponds to interval [1, n], $n = 2^{j}$,
- the interval must be more than *half-full*, > n/2 elements,

- array-part corresponds to interval [1, n], $n = 2^{j}$,
- the interval must be more than *half-full*, > n/2 elements,
- maximum $n = 2^j$ chosen during rehash.

- array-part corresponds to interval [1, n], $n = 2^j$,
- the interval must be more than *half-full*, > n/2 elements,
- maximum $n = 2^j$ chosen during rehash.


Hybrid Tabes mechanism: principle

- array-part corresponds to interval [1, n], $n = 2^j$,
- the interval must be more than *half-full*, > n/2 elements,
- maximum $n = 2^j$ chosen during rehash.



... rehash just emulated a deletion in the hash-part !

Example: inserting $-(2^k - 1), -(2^k - 2), \dots, -1, 0, 1, \dots, 2^k$

▶ keys $-(2^k - 1), -(2^k - 2), \dots, 0$ go into hash-part

Example: inserting $-(2^k - 1), -(2^k - 2), \dots, -1, 0, 1, \dots, 2^k$

▶ keys $-(2^k - 1), -(2^k - 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full

- ▶ keys $-(2^k 1), -(2^k 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full
- key 1 induces rehash

- ▶ keys $-(2^k 1), -(2^k 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full
- key 1 induces rehash \dots but goes into the array-part

- ▶ keys $-(2^k 1), -(2^k 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full
- ▶ key 1 induces rehash ... but goes into the array-part ⇒ array-part $A = 2^0$, hash-part $M = 2^k$.

- ▶ keys $-(2^k 1), -(2^k 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full
- ▶ key 1 induces rehash ... but goes into the array-part ⇒ array-part $A = 2^0$, hash-part $M = 2^k$.
- key 2 induces rehash

- ▶ keys $-(2^k 1), -(2^k 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full
- ▶ key 1 induces rehash ... but goes into the array-part ⇒ array-part $A = 2^0$, hash-part $M = 2^k$.
- \blacktriangleright key 2 induces rehash \ldots and goes into the array-part

- ▶ keys $-(2^k 1), -(2^k 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full
- ▶ key 1 induces rehash ... but goes into the array-part ⇒ array-part $A = 2^0$, hash-part $M = 2^k$.
- ▶ key 2 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^1$, hash-part $M = 2^k$.

- ▶ keys $-(2^k 1), -(2^k 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full
- ▶ key 1 induces rehash ... but goes into the array-part ⇒ array-part $A = 2^0$, hash-part $M = 2^k$.
- ▶ key 2 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^1$, hash-part $M = 2^k$.
- key 3 induces rehash

- ▶ keys $-(2^k 1), -(2^k 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full
- ▶ key 1 induces rehash ... but goes into the array-part ⇒ array-part $A = 2^0$, hash-part $M = 2^k$.
- ▶ key 2 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^1$, hash-part $M = 2^k$.
- ▶ key 3 induces rehash ... and goes into the array-part

- ▶ keys $-(2^k 1), -(2^k 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full
- ▶ key 1 induces rehash ... but goes into the array-part ⇒ array-part $A = 2^0$, hash-part $M = 2^k$.
- ▶ key 2 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^1$, hash-part $M = 2^k$.
- ▶ key 3 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^2$, hash-part $M = 2^k$.

- ▶ keys $-(2^k 1), -(2^k 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full
- ▶ key 1 induces rehash ... but goes into the array-part ⇒ array-part $A = 2^0$, hash-part $M = 2^k$.
- ▶ key 2 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^1$, hash-part $M = 2^k$.
- ▶ key 3 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^2$, hash-part $M = 2^k$.
- key 4 is inserted directly into array-part $[1, 2^2]$.

- ▶ keys $-(2^k 1), -(2^k 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full
- ▶ key 1 induces rehash ... but goes into the array-part ⇒ array-part $A = 2^0$, hash-part $M = 2^k$.
- ▶ key 2 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^1$, hash-part $M = 2^k$.
- ▶ key 3 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^2$, hash-part $M = 2^k$.
- key 4 is inserted directly into array-part $[1, 2^2]$.
- key 5 induces rehash

- ▶ keys $-(2^k 1), -(2^k 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full
- ▶ key 1 induces rehash ... but goes into the array-part ⇒ array-part $A = 2^0$, hash-part $M = 2^k$.
- ▶ key 2 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^1$, hash-part $M = 2^k$.
- ▶ key 3 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^2$, hash-part $M = 2^k$.
- key 4 is inserted directly into array-part $[1, 2^2]$.
- \blacktriangleright key 5 induces rehash ... and goes into the array-part

- ▶ keys $-(2^k 1), -(2^k 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full
- ▶ key 1 induces rehash ... but goes into the array-part ⇒ array-part $A = 2^0$, hash-part $M = 2^k$.
- ▶ key 2 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^1$, hash-part $M = 2^k$.
- ▶ key 3 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^2$, hash-part $M = 2^k$.
- key 4 is inserted directly into array-part $[1, 2^2]$.
- ▶ key 5 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^3$, hash-part $M = 2^k$.

. . .

- ▶ keys $-(2^k 1), -(2^k 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full
- ▶ key 1 induces rehash ... but goes into the array-part ⇒ array-part $A = 2^0$, hash-part $M = 2^k$.
- ▶ key 2 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^1$, hash-part $M = 2^k$.
- ▶ key 3 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^2$, hash-part $M = 2^k$.
- key 4 is inserted directly into array-part $[1, 2^2]$.
- ▶ key 5 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^3$, hash-part $M = 2^k$.
- ▶ keys $2^j + 1, j \ge 0$, induce rehash, rest inserted in array-part directly

. . .

- ▶ keys $-(2^k 1), -(2^k 2), \dots, 0$ go into hash-part $\Rightarrow M = 2^k$ full
- ▶ key 1 induces rehash ... but goes into the array-part ⇒ array-part $A = 2^0$, hash-part $M = 2^k$.
- ▶ key 2 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^1$, hash-part $M = 2^k$.
- ▶ key 3 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^2$, hash-part $M = 2^k$.
- key 4 is inserted directly into array-part $[1, 2^2]$.
- ▶ key 5 induces rehash ... and goes into the array-part ⇒ array-part $A = 2^3$, hash-part $M = 2^k$.
- ▶ keys $2^j + 1, j \ge 0$, induce rehash, rest inserted in array-part directly ⇒ time $\Omega(k \cdot 2^k)$

Permutations of $[1, \ldots, n]$ are a natural choice:

- they come up in practice in many contexts,
- they give the array-part a good shot at being used, yet ...

Permutations of $[1, \ldots, n]$ are a natural choice:

- they come up in practice in many contexts,
- they give the array-part a good shot at being used, yet ...

- mechanism is exactly as in previous example,
- ▶ keys $2 \cdot 2^k + 1, 2 \cdot 2^k + 2, \dots, 3 \cdot 2^k$ go into hash-part

Permutations of $[1, \ldots, n]$ are a natural choice:

- they come up in practice in many contexts,
- they give the array-part a good shot at being used, yet ...

- mechanism is exactly as in previous example,
- ▶ keys $2 \cdot 2^k + 1, 2 \cdot 2^k + 2, \dots, 3 \cdot 2^k$ go into hash-part $\Rightarrow M = 2^k$ full,

Permutations of $[1, \ldots, n]$ are a natural choice:

- they come up in practice in many contexts,
- they give the array-part a good shot at being used, yet ...

- mechanism is exactly as in previous example,
- ▶ keys $2 \cdot 2^k + 1, 2 \cdot 2^k + 2, \dots, 3 \cdot 2^k$ go into hash-part $\Rightarrow M = 2^k$ full,
- keys $1, 2, 3, 5, 9, \ldots$ induce rehash, while rest inserted into array-part

Permutations of $[1, \ldots, n]$ are a natural choice:

- they come up in practice in many contexts,
- they give the array-part a good shot at being used, yet ...

- mechanism is exactly as in previous example,
- ▶ keys $2 \cdot 2^k + 1, 2 \cdot 2^k + 2, \dots, 3 \cdot 2^k$ go into hash-part $\Rightarrow M = 2^k$ full,
- keys $1, 2, 3, 5, 9, \ldots$ induce rehash, while rest inserted into array-part
- ... but this is rather unlikely for a permutation

Random model starting from an empty table:

• Consider random permutation π of [n]

Random model starting from an empty table:

- Consider random permutation π of [n]
- Insert elements in order $\pi(1), \pi(2) \dots, \pi(n)$.

Random model starting from an empty table:

- Consider random permutation π of [n]
- Insert elements in order $\pi(1), \pi(2) \dots, \pi(n)$.

Random model starting from an empty table:

- Consider random permutation π of [n]
- Insert elements in order $\pi(1), \pi(2) \dots, \pi(n)$.

Fix $c < \frac{1}{2}$ and consider an arbitrary $g(n) \rightarrow \infty$, e.g., $g(n) = \log^*(n)$.

Lemma

For every time $t \leq cn$, none of $S_j = [1, 2^j]$ for $2^j \geq g(n)$ is half-full with probability tending to one.

Random model starting from an empty table:

- Consider random permutation π of [n]
- Insert elements in order $\pi(1), \pi(2) \dots, \pi(n)$.

Fix $c < \frac{1}{2}$ and consider an arbitrary $g(n) \rightarrow \infty$, e.g., $g(n) = \log^*(n)$.

Lemma

For every time $t \leq cn$, none of $S_j = [1, 2^j]$ for $2^j \geq g(n)$ is half-full with probability tending to one.

 \implies array part not really used

Random model starting from an empty table:

- Consider random permutation π of [n]
- Insert elements in order $\pi(1), \pi(2) \dots, \pi(n)$.

Fix $c < \frac{1}{2}$ and consider an arbitrary $g(n) \rightarrow \infty$, e.g., $g(n) = \log^*(n)$.

Lemma

For every time $t \leq cn$, none of $S_j = [1, 2^j]$ for $2^j \geq g(n)$ is half-full with probability tending to one.

⇒ array part not really used and complexity essentially linear

Theorem (Martínez, Nicaud, R 2022)

Inserting random permutation takes time $\mathcal{O}(n g(n))$ provided that n does not approximate powers of two from above^a.

^aFix $b \in (1, 2)$, $n \in \mathbb{N}_b := \bigcup_{j \ge 0} \{k : 2^j b < k \le 2^{j+1}\}.$

Recap and conclusions

- & Lua's hybrid data-structure is an interesting idea.
- We have presented a simple and natural probabilistic model revealing shortcomings in Lua's hashtables.

Recap and conclusions

- & Lua's hybrid data-structure is an interesting idea.
- We have presented a simple and natural probabilistic model revealing shortcomings in Lua's hashtables.
- Issue can be fixed by ensuring more room when rehashing.
- $\ensuremath{\mathfrak{B}}$ This would also fix the hybrid part.

Recap and conclusions

- & Lua's hybrid data-structure is an interesting idea.
- We have presented a simple and natural probabilistic model revealing shortcomings in Lua's hashtables.
- Issue can be fixed by ensuring more room when rehashing.
- & This would also fix the hybrid part.

Conclusions

- ⊗ Will Lua conceptors take this into account?
- \circledast Important to model and study algorithms implemented in practice.

Thank you!