


# Linear Virtual Hashing

by. Tim Henderson (tadh@Case.edu).

hackthology.com [github.com/timbadh](https://github.com/timbadh)



## References

Database Systems - The Complete Book

by H. Garcia-Molina, J.D. Ullman, and J. Widom (2002)

Linear Hashing: a new tool for file and table addressing  
by Litwin, witol

in Very Large Databases (VLDB)

Volume 6, Pages 212-223

1780

# Classical Hashing - A review.

I expect all of you to be familiar with the classic hashtable. However, some of the finer points may be fuzzy. Let's review!

ADT:

Size(): int - how many entries?

has (key: Hashable): boolean - is the key in the table?

get (key: Hashable): Object throws KeyNotFound

put (key: Hashable, value: Object)


rm (key: Hashable) throws KeyNotFound.

Hashable

hash(): int

insert

key  $\rightarrow$  hash()  $\rightarrow$   $\frac{\text{hash} \% \text{table-size}}$



Search

key  $\rightarrow$  hash()  $\rightarrow$   $\frac{\text{hash} \% \text{table-size}}$

remove

Key  $\rightarrow$  Hash()  $\rightarrow$  hash % table-size

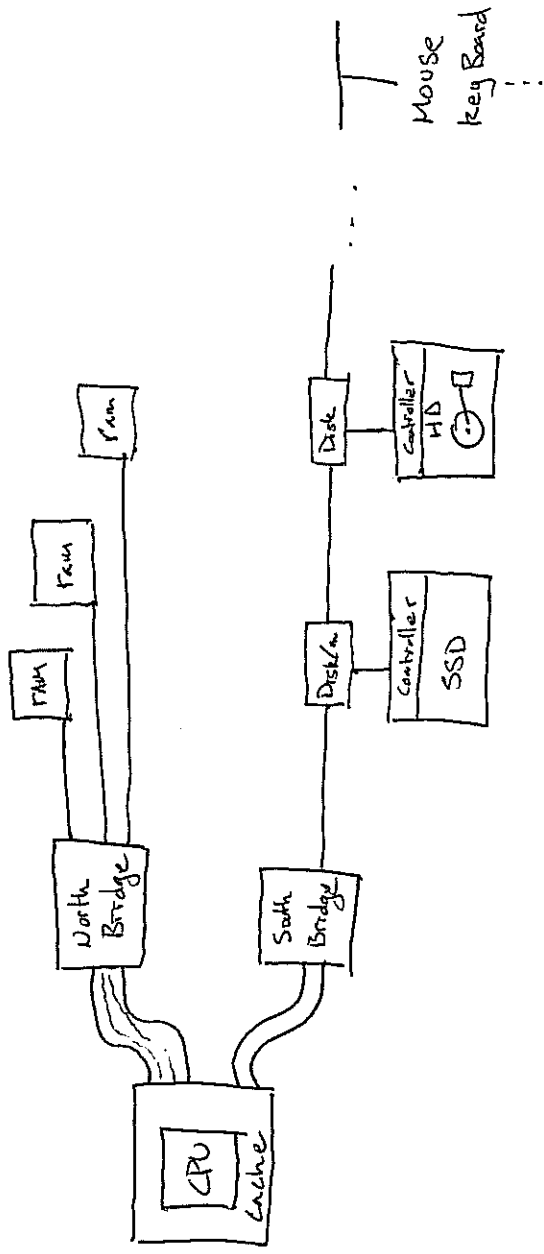
resize

1. Double Table size by allocating a new array.
2. Copy all entries from old array to the new one
  - a. Scan old table to get each entry
  - b. entries must be rehashed

Properties

1. Amortized (on average) lookup, insert, remove costs of  $O(1)$ .
2. Lots of wasted space. Experience suggests expansion at .6 utilization.
3. All entries must be copied on expansion
4. Great for in memory lookup tables
5. Terrible for secondary memory tables. Why?

# Secondary Storage, The Rules Change



1. Secondary Storage is slower as a medium
2. The Bus is slower
3. Many Peripherals hang of the South Bridge
4. Depending on interface, they may be daisy chained  
↳ bus contention.

To deal with all these factors

1. We read and write pages ~~that~~ which are Blocks of size 4096 bytes.
2. In a good world we would read contiguous runs and write back individual blocks.
3. writes are batched.
4. you don't "read a byte", you read several blocks and get just the byte you want.
5. you employ caching heavily.
6. Data-structures are now measured in terms of # of Disk Accesses. (Block reads/writes) - 2


Let's go Back To Hashing -

First Adjustment: Let's hash into Blocks

Key = "Wizard"

↓  
Hash()

hash = 0x1617b1ae →  $\frac{\text{hash}}{\text{\# Blocks}}$



Blocks are sorted Arrays of keys/Values. So it takes  $\frac{1}{2} (\text{Block Size} / \text{record Size})$  to get an item.

But... how do we expand?

1. The classic "Doubling" Scheme no longer looks so good.
  - a. Allocates a huge amount of space
  - b. probably writes to all of it, assuming a good hash function.
2. We can't just "add" blocks
3. We need to be able to still find old records.

Solution - Linear Hashing.

note. There are other systems but LH is the best.

## Linear Hashing - Properties.

1. Mean Accesses per Lookup	on Successful <del>Insert</del> lookup	insert
Utilization .6 ~ 1.03	.75 ~ 1.27	.75 ~ 2.62
.9 ~ 1.35	.9 ~ 2.37	.9 ~ 3.73

2. Grows at a linear rate

at most 2 new blocks created on insert.

3. Requires little dynamic re-arrangement.

only on block split and only 1 block's worth of records.

4. Dynamically shrinks and grows

5. Does not need address translation (necessarily)

when translation is used it becomes "virtual"

6. In comparison a b+ tree of reasonable size may need at least 4 disk accesses per random lookup. (of course a b+ tree will perform much better for a range scan).

7. Simple Algorithm. Esp. compared to B+ trees.



# Linear Hashing.

1. Key Insight: Incrementally use more bits of the hash function  $H(\cdot)$ . This allows us to grow the table by almost "just adding a block".  
ex.



$$H(\text{wizard}) \rightarrow 1100_2$$

We will use  $\lg(n)$  bits of  $H(\cdot)$  where  $n$  is the # of Blocks in the file.

File  $n=2$   $i=1$   $r=2$




2. When a bucket is full chain another Block



3. Split when  $\text{Record} > \text{UTILIZATION} * \# \text{ Buckets} * \# \text{ Records per Block}$   
 $5 > .8 \cdot 2 \cdot 3$  So we need to split  
 $5 > 4.8$  but how do we do that?

# Splitting

First We Add a Bucket



A bit<sup>more</sup> of the hash function is now used.

00  
01  
10

The bucket we added is

$$1a_i = 10 = a_1, a_2$$

The bucket we need to split is


$$0a_i$$

Why?

In general if we add  $1a_2 \dots a_i$  we split  $0a_2 \dots a_i$

The split bucket is thus 00. After the split we get...

File  $n=3$   $i=2$   $r=5$





# Linear Hashing Algorithm Summary.

## Insert(key)

```
hash = Hash(key)
bkt_idx = let
    hash = sxxxxa1a2...ai
    m = a1...ai-1
    in if m < n then
        else
            m - 2i-1 = 0 a2...ai
            bkt = get_bucket(bkt_idx)
            bkt.put(key, value)
            r++
    if r > UTILIZATION * n (records per block) then
        split()
```

We use this algorithm everytime we need to compute the bucket index let's call it bucket\_idx().

## Get(key)

```
hash = Hash(key)
bkt_idx = bucket_idx(hash)
bkt = get_bucket(bkt_idx)
return bkt.get(key).
```

## Remove(key)

```
hash = Hash(key)
bkt_idx = bucket_idx(hash)
bkt = get_bucket(bkt_idx)
return bkt.remove(key)
r--
```

## Split()

```
bkt_idx = n % (1 << (i-1))
bkt-a = get_bucket(bkt_idx)
bkt-b = alloc()
n++
if n > (1 << i) then
    i++
for each entry in bkt-a
    put entry in bkt-b if ai-1 = 1 in key.
```