

## 3.4 HASH TABLES

- ▶ hash functions
- ▶ separate chaining
- ▶ linear probing
- ▶ context

<http://algs4.cs.princeton.edu>


Last updated on Mar 4, 2015, 9:56 AM

## Premature optimization

*“ Programmers waste enormous amounts of time thinking about, or worrying about, the speed of noncritical parts of their programs, and these attempts at efficiency actually have a strong negative impact when debugging and maintenance are considered.*

*We should forget about small efficiencies, say about 97% of the time: premature optimization is the root of all evil.*

*Yet we should not pass up our opportunities in that critical 3%. ”*



## Symbol table implementations: summary

implementation	guarantee			average case			ordered ops?	key interface
	search	insert	delete	search hit	insert	delete		
sequential search (unordered list)	$N$	$N$	$N$	$N$	$N$	$N$		<code>equals()</code>
binary search (ordered array)	$\log N$	$N$	$N$	$\log N$	$N$	$N$	✓	<code>compareTo()</code>
BST	$N$	$N$	$N$	$\log N$	$\log N$	$\sqrt{N}$	✓	<code>compareTo()</code>
red-black BST	$\log N$	$\log N$	$\log N$	$\log N$	$\log N$	$\log N$	✓	<code>compareTo()</code>


Q. Can we do better?

A. Yes, but with different access to the data.

## Hashing: basic plan

Save items in a **key-indexed table** (index is a function of the key).

**Hash function.** Method for computing array index from key.




**Issues.**

- Computing the hash function.
- Equality test: Method for checking whether two keys are equal.
- Collision resolution: Algorithm and data structure to handle two keys that hash to the same array index.

**Classic space-time tradeoff.**

- No space limitation: trivial hash function with key as index.
- No time limitation: trivial collision resolution with sequential search.
- Space and time limitations: hashing (the real world).



### 3.4 HASH TABLES


- ▶ hash functions
- ▶ separate chaining
- ▶ linear probing
- ▶ context

### Computing the hash function

**Idealistic goal.** Scramble the keys uniformly to produce a table index.

- Efficiently computable.
- Each table index equally likely for each key.

thoroughly researched problem, still problematic in practical applications



**Ex. Social Security numbers.**

- Bad: first three digits. ← 573 = California, 574 = Alaska (assigned in chronological order within geographic region)
- Better: last three digits.

**Practical challenge.** Need different approach for each key type.

### Hash tables: quiz 1

Which of the following would be a good hash function for U.S. phone numbers to integers between 0 and 999?

- A. First three digits.
- B. Second three digits.
- C. Last three digits.
- D. Either B or C.
- E. I don't know.




### Java's hash code conventions

All Java classes inherit a method `hashCode()`, which returns a 32-bit int.

**Requirement.** If `x.equals(y)`, then `(x.hashCode() == y.hashCode())`.

**Highly desirable.** If `!x.equals(y)`, then `(x.hashCode() != y.hashCode())`.



**Default implementation.** Memory address of `x`.

**Legal (but poor) implementation.** Always return 17.

**Customized implementations.** Integer, Double, String, File, URL, Date, ...

**User-defined types.** Users are on their own.

## Implementing hash code: integers, booleans, and doubles

### Java library implementations

```
public final class Integer
{
    private final int value;
    ...

    public int hashCode()
    { return value; }
}
```

```
public final class Double
{
    private final double value;
    ...

    public int hashCode()
    {
        long bits = doubleToLongBits(value);
        return (int) (bits ^ (bits >>> 32));
    }
}
```

```
public final class Boolean
{
    private final boolean value;
    ...

    public int hashCode()
    {
        if (value) return 1231;
        else      return 1237;
    }
}
```

convert to IEEE 64-bit representation;  
xor most significant 32-bits  
with least significant 32-bits

Warning: -0.0 and +0.0 have different hash codes

9

## Implementing hash code: strings

Treat string of length  $L$  as  $L$ -digit, base-31 number:

$$h = s[0] \cdot 31^{L-1} + \dots + s[L-3] \cdot 31^2 + s[L-2] \cdot 31^1 + s[L-1] \cdot 31^0$$

```
public final class String
{
    private final char[] s;
    ...

    public int hashCode()
    {
        int hash = 0;
        for (int i = 0; i < length(); i++)
            hash = s[i] + (31 * hash);
        return hash;
    }
}
Java library implementation
```

char	Unicode
...	...
'a'	97
'b'	98
'c'	99
...	...

**Horner's method:** only  $L$  multiplies/adds to hash string of length  $L$ .

```
String s = "call";
s.hashCode(); ← 3045982 = 99·313 + 97·312 + 108·311 + 108·310
                = 108 + 31·(108 + 31·(97 + 31·(99)))
```

10

## Implementing hash code: strings

### Performance optimization.

- Cache the hash value in an instance variable.
- Return cached value.

```
public final class String
{
    private int hash = 0; ← cache of hash code
    private final char[] s;
    ...

    public int hashCode()
    {
        int h = hash; ← return cached value
        if (h != 0) return h;
        for (int i = 0; i < length(); i++)
            h = s[i] + (31 * h);
        hash = h; ← store cache of hash code
        return h;
    }
}
```

Q. What if hashCode() of string is 0? ← hashCode() of "pollinating sandboxes" is 0

11

## Implementing hash code: user-defined types

```
public final class Transaction implements Comparable<Transaction>
{
    private final String who;
    private final Date when;
    private final double amount;

    public Transaction(String who, Date when, double amount)
    { /* as before */ }

    ...

    public boolean equals(Object y)
    { /* as before */ }

    public int hashCode()
    {
        int hash = 17; ← nonzero constant
        hash = 31*hash + who.hashCode(); ← for reference types, use hashCode()
        hash = 31*hash + when.hashCode(); ← for primitive types, use hashCode() of wrapper type
        hash = 31*hash + ((Double) amount).hashCode();
        return hash; ← typically a small prime
    }
}
```

12

## Hash code design

### "Standard" recipe for user-defined types.

- Combine each significant field using the  $31x + y$  rule.
- If field is a primitive type, use wrapper type `hashCode()`.
- If field is `null`, use 0.
- If field is a reference type, use `hashCode()`. ← applies rule recursively
- If field is an array, apply to each entry. ← or use `Arrays.deepHashCode()`

**In practice.** Recipe above works reasonably well; used in Java libraries.

**In theory.** Keys are bitstring; "universal" family of hash functions exist.

awkward in Java since only one (deterministic) `hashCode()`

**Basic rule.** Need to use the whole key to compute hash code; consult an expert for state-of-the-art hash codes.

13


## Hash tables: quiz 1

Which of the following is an effective way to map a hashable key to an integer between 0 and  $M-1$  ?

- A. 

```
private int hash(Key key)
{ return key.hashCode() % M; }
```
- B. 

```
private int hash(Key key)
{ return Math.abs(key.hashCode()) % M; }
```
- C. Both A and B.
- D. Neither A nor B.
- E. *I don't know.*



14

## Modular hashing

**Hash code.** An int between  $-2^{31}$  and  $2^{31} - 1$ .

**Hash function.** An int between 0 and  $M - 1$  (for use as array index).

typically a prime or power of 2

```
private int hash(Key key)
{ return key.hashCode() % M; }
```

bug

```
private int hash(Key key)
{ return Math.abs(key.hashCode()) % M; }
```

1-in-a-billion bug

`hashCode()` of "polygenelubricants" is  $-2^{31}$

```
private int hash(Key key)
{ return (key.hashCode() & 0x7fffffff) % M; }
```

correct



15

## Uniform hashing assumption

**Uniform hashing assumption.** Each key is equally likely to hash to an integer between 0 and  $M - 1$ .

**Bins and balls.** Throw balls uniformly at random into  $M$  bins.



**Birthday problem.** Expect two balls in the same bin after  $\sim \sqrt{\pi M / 2}$  tosses.

**Coupon collector.** Expect every bin has  $\geq 1$  ball after  $\sim M \ln M$  tosses.


**Load balancing.** After  $M$  tosses, expect most loaded bin has  $\sim \ln M / \ln \ln M$  balls.

16

## Uniform hashing assumption

**Uniform hashing assumption.** Each key is equally likely to hash to an integer between 0 and  $M - 1$ .

**Bins and balls.** Throw balls uniformly at random into  $M$  bins.




Java's String data uniformly distribute the keys of Tale of Two Cities

## Collisions

**Collision.** Two distinct keys hashing to same index.

- Birthday problem  $\Rightarrow$  can't avoid collisions. ← unless you have a ridiculous (quadratic) amount of memory
- Coupon collector  $\Rightarrow$  not too much wasted space.
- Load balancing  $\Rightarrow$  no index gets too many collisions.



**Challenge.** Deal with collisions efficiently.

## 3.4 HASH TABLES



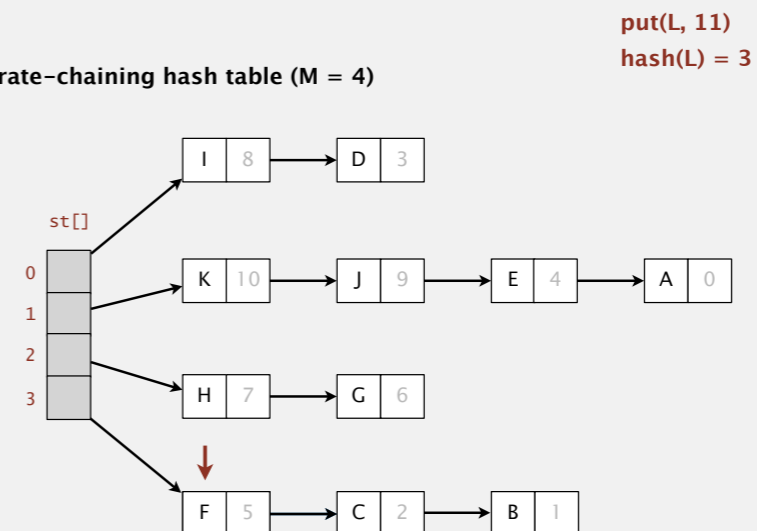
- ▶ hash functions
- ▶ separate chaining
- ▶ linear probing
- ▶ context

## Separate-chaining symbol table

Use an array of  $M < N$  linked lists. [H. P. Luhn, IBM 1953]

- Hash: map key to integer  $i$  between 0 and  $M - 1$ .
- Insert: put at front of  $i^{\text{th}}$  chain (if not already in chain).
- Search: sequential search in  $i^{\text{th}}$  chain.

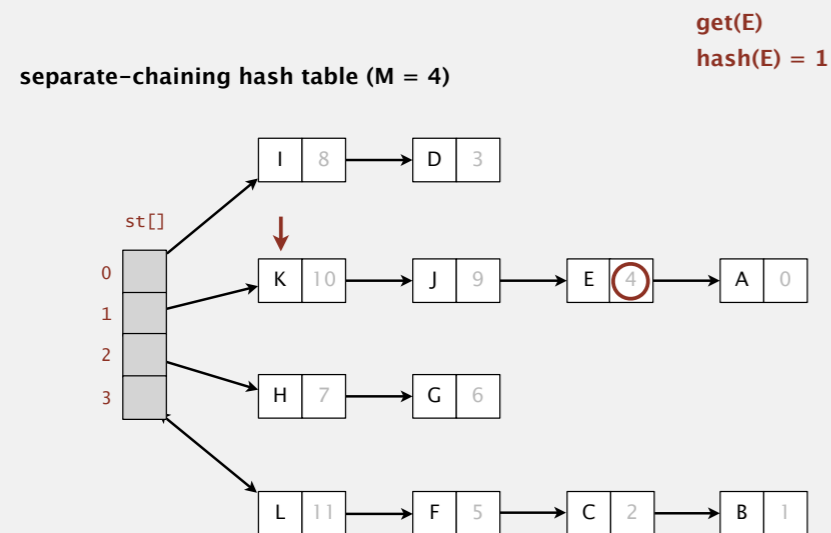
separate-chaining hash table ( $M = 4$ )



## Separate-chaining symbol table

Use an array of  $M < N$  linked lists. [H. P. Luhn, IBM 1953]

- Hash: map key to integer  $i$  between 0 and  $M - 1$ .
- Insert: put at front of  $i^{\text{th}}$  chain (if not already in chain).
- Search: sequential search in  $i^{\text{th}}$  chain.



21

## Separate-chaining symbol table: Java implementation

```
public class SeparateChainingHashST<Key, Value>
{
    private int M = 97;           // number of chains
    private Node[] st = new Node[M]; // array of chains

    private static class Node
    {
        private Object key;
        private Object val;
        private Node next;
        ...
    }

    private int hash(Key key)
    { return (key.hashCode() & 0x7fffffff) % M; }

    public Value get(Key key) {
        int i = hash(key);
        for (Node x = st[i]; x != null; x = x.next)
            if (key.equals(x.key)) return x.val;
        return null;
    }
}
```

array doubling and halving code omitted

no generic array creation  
(declare key and value of type Object)

22

## Separate-chaining symbol table: Java implementation

```
public class SeparateChainingHashST<Key, Value>
{
    private int M = 97;           // number of chains
    private Node[] st = new Node[M]; // array of chains

    private static class Node
    {
        private Object key;
        private Object val;
        private Node next;
        ...
    }

    private int hash(Key key)
    { return (key.hashCode() & 0x7fffffff) % M; }


    public void put(Key key, Value val) {
        int i = hash(key);
        for (Node x = st[i]; x != null; x = x.next)
            if (key.equals(x.key)) { x.val = val; return; }
        st[i] = new Node(key, val, st[i]);
    }
}
```

23

## Analysis of separate chaining

**Proposition.** Under uniform hashing assumption, prob. that the number of keys in a list is within a constant factor of  $N/M$  is extremely close to 1.

**Pf sketch.** Distribution of list size obeys a binomial distribution.



equals() and hashCode()

**Consequence.** Number of probes for search/insert is proportional to  $N/M$ .

- $M$  too large  $\Rightarrow$  too many empty chains.
- $M$  too small  $\Rightarrow$  chains too long.
- Typical choice:  $M \sim \frac{1}{4} N \Rightarrow$  constant-time ops.

M times faster than sequential search


24

## Resizing in a separate-chaining hash table


**Goal.** Average length of list  $N/M = \text{constant}$ .

- Double size of array  $M$  when  $N/M \geq 8$ ;  
halve size of array  $M$  when  $N/M \leq 2$ .
- Note: need to rehash all keys when resizing.  $\leftarrow x.\text{hashCode}()$  does not change;  
but  $\text{hash}(x)$  can change

before resizing ( $N/M = 8$ )



after resizing ( $N/M = 4$ )




25


## Deletion in a separate-chaining hash table

- Q.** How to delete a key (and its associated value)?
- A.** Easy: need to consider only chain containing key.

before deleting C



after deleting C



26

## Symbol table implementations: summary

implementation	guarantee			average case			ordered ops?	key interface
	search	insert	delete	search hit	insert	delete		
sequential search (unordered list)	$N$	$N$	$N$	$N$	$N$	$N$		<code>equals()</code>
binary search (ordered array)	$\log N$	$N$	$N$	$\log N$	$N$	$N$	✓	<code>compareTo()</code>
BST	$N$	$N$	$N$	$\log N$	$\log N$	$\sqrt{N}$	✓	<code>compareTo()</code>
red-black BST	$\log N$	$\log N$	$\log N$	$\log N$	$\log N$	$\log N$	✓	<code>compareTo()</code>
separate chaining	$N$	$N$	$N$	$1^*$	$1^*$	$1^*$		<code>equals()</code> <code>hashCode()</code>

\* under uniform hashing assumption

27

ROBERT SEDGWICK | KEVIN WAYNE  
<http://algs4.cs.princeton.edu>

## 3.4 HASH TABLES

- ▶ hash functions
- ▶ separate chaining
- ▶ linear probing
- ▶ context

## Collision resolution: open addressing

**Open addressing.** [Amdahl-Boehme-Rochester-Samuel, IBM 1953]

- Maintain keys and values in two parallel arrays.
- When a new key collides, find next empty slot, and put it there.

linear-probing hash table (M = 16, N = 10)

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
keys[]	P	M			A	C		H	L		E				R	X
vals[]	11	10			9	5		6	12		13				4	8

put(K, 14)  
hash(K) = 7

K  
14

29

## Linear-probing hash table summary

**Hash.** Map key to integer  $i$  between 0 and  $M - 1$ .

**Insert.** Put at table index  $i$  if free; if not try  $i + 1$ ,  $i + 2$ , etc.

**Search.** Search table index  $i$ ; if occupied but no match, try  $i + 1$ ,  $i + 2$ , etc.

**Note.** Array size  $M$  **must be** greater than number of key-value pairs  $N$ .

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
keys[]	P	M			A	C	S	H	L		E				R	X

M = 16



30

## Linear-probing symbol table: Java implementation

```
public class LinearProbingHashST<Key, Value>
{
    private int M = 30001;
    private Value[] vals = (Value[]) new Object[M];
    private Key[] keys = (Key[]) new Object[M];

    private int hash(Key key) { /* as before */ }

    private void put(Key key, Value val) { /* next slide */ }

    public Value get(Key key)
    {
        for (int i = hash(key); keys[i] != null; i = (i+1) % M)
            if (key.equals(keys[i]))
                return vals[i];
        return null;
    }
}
```

array doubling and halving code omitted

sequential search in chain i

31

## Linear-probing symbol table: Java implementation

```
public class LinearProbingHashST<Key, Value>
{
    private int M = 30001;
    private Value[] vals = (Value[]) new Object[M];
    private Key[] keys = (Key[]) new Object[M];

    private int hash(Key key) { /* as before */ }

    private Value get(Key key) { /* prev slide */ }

    public void put(Key key, Value val)
    {
        int i;
        for (i = hash(key); keys[i] != null; i = (i+1) % M)
            if (keys[i].equals(key))
                break;
        keys[i] = key;
        vals[i] = val;
    }
}
```

sequential search in chain i

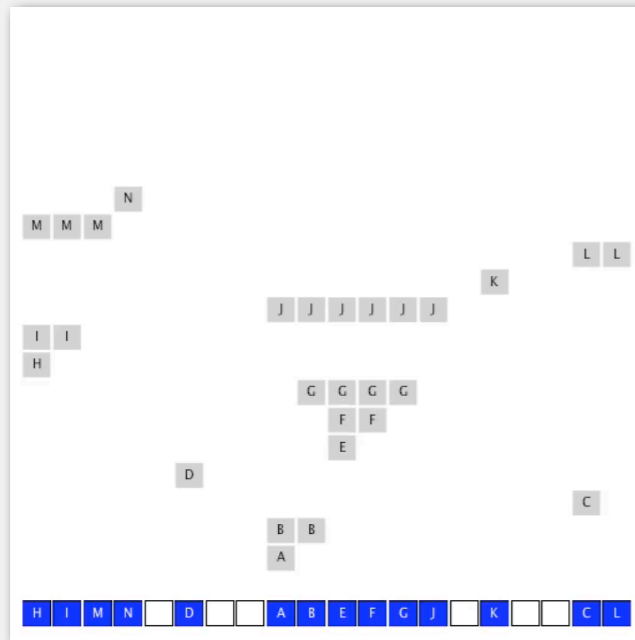
32



## Clustering

**Cluster.** A contiguous block of items.

**Observation.** New keys likely to hash into middle of big clusters.



33

## Knuth's parking problem

**Model.** Cars arrive at one-way street with  $M$  parking spaces.

Each desires a random space  $i$ : if space  $i$  is taken, try  $i + 1, i + 2$ , etc.

**Q.** What is mean displacement of a car?



**Half-full.** With  $M/2$  cars, mean displacement is  $\sim 5/2$ .

**Full.** With  $M$  cars, mean displacement is  $\sim \sqrt{\pi M/8}$ .

**Key insight.** Cannot afford to let linear-probing hash table get too full.

34


## Analysis of linear probing

**Proposition.** Under uniform hashing assumption, the average # of probes in a linear probing hash table of size  $M$  that contains  $N = \alpha M$  keys is:

$$\sim \frac{1}{2} \left( 1 + \frac{1}{1-\alpha} \right) \quad \sim \frac{1}{2} \left( 1 + \frac{1}{(1-\alpha)^2} \right)$$

search hit                      search miss / insert

**Pf.**



**Parameters.**

- $M$  too large  $\Rightarrow$  too many empty array entries.
- $M$  too small  $\Rightarrow$  search time blows up.
- Typical choice:  $\alpha = N/M \sim 1/2$ .  $\leftarrow$  # probes for search hit is about  $3/2$   
# probes for search miss is about  $5/2$

35

## Resizing in a linear-probing hash table

**Goal.** Average length of list  $N/M \leq 1/2$ .

- Double size of array  $M$  when  $N/M \geq 1/2$ .
- Halve size of array  $M$  when  $N/M \leq 1/8$ .
- Need to rehash all keys when resizing.

before resizing

	0	1	2	3	4	5	6	7
keys[]		E	S			R	A	
vals[]		1	0			3	2	

after resizing

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
keys[]					A	S				E					R	
vals[]					2	0				1					3	

36

## Deletion in a linear-probing hash table

Q. How to delete a key (and its associated value)?

A. Requires some care: can't just delete array entries.

before deleting S

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
keys[]	P	M			A	C	S	H	L		E				R	X
vals[]	10	9			8	4	0	5	11		12				3	7

after deleting S ?

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
keys[]	P	M			A	C		H	L		E				R	X
vals[]	10	9			8	4		5	11		12				3	7

doesn't work, e.g., if  $\text{hash}(H) = 4$

37

## ST implementations: summary

implementation	guarantee			average case			ordered ops?	key interface
	search	insert	delete	search hit	insert	delete		
sequential search (unordered list)	$N$	$N$	$N$	$N$	$N$	$N$		<code>equals()</code>
binary search (ordered array)	$\log N$	$N$	$N$	$\log N$	$N$	$N$	✓	<code>compareTo()</code>
BST	$N$	$N$	$N$	$\log N$	$\log N$	$\sqrt{N}$	✓	<code>compareTo()</code>
red-black BST	$\log N$	$\log N$	$\log N$	$\log N$	$\log N$	$\log N$	✓	<code>compareTo()</code>
separate chaining	$N$	$N$	$N$	$1^*$	$1^*$	$1^*$		<code>equals()</code> <code>hashCode()</code>
linear probing	$N$	$N$	$N$	$1^*$	$1^*$	$1^*$		<code>equals()</code> <code>hashCode()</code>

\* under uniform hashing assumption

38

## 3-SUM (REVISITED)

**3-SUM.** Given  $N$  distinct integers, find three such that  $a + b + c = 0$ .

**Goal.**  $N^2$  expected time case,  $N$  extra space.

39



### 3.4 HASH TABLES

- ▶ hash functions
- ▶ separate chaining
- ▶ linear probing
- ▶ context

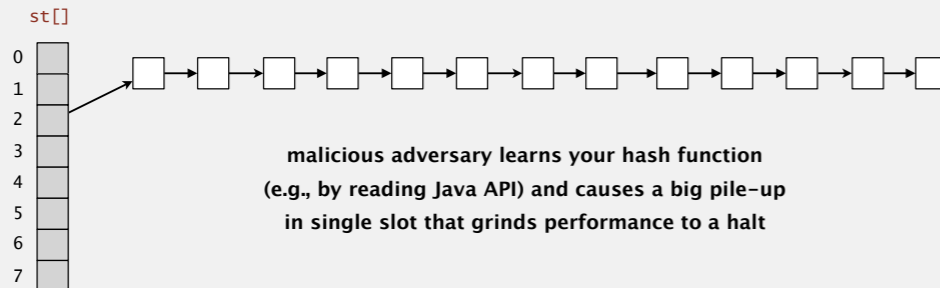
ROBERT SEDGWICK | KEVIN WAYNE  
<http://algs4.cs.princeton.edu>

## War story: algorithmic complexity attacks

Q. Is the uniform hashing assumption important in practice?

A. Obvious situations: aircraft control, nuclear reactor, pacemaker, HFT, ...

A. Surprising situations: **denial-of-service** attacks.



Real-world exploits. [Crosby-Wallach 2003]

- Bro server: send carefully chosen packets to DOS the server, using less bandwidth than a dial-up modem.
- Perl 5.8.0: insert carefully chosen strings into associative array.
- Linux 2.4.20 kernel: save files with carefully chosen names.

41

## War story: algorithmic complexity attacks

A Java bug report.

Jan Lieskovsky 2011-11-01 10:13:47 EDT

Description

Julian Wälde and Alexander Klink reported that the String.hashCode() hash function is not sufficiently collision resistant. hashCode() value is used in the implementations of HashMap and Hashtable classes:

<http://docs.oracle.com/javase/6/docs/api/java/util/HashMap.html>  
<http://docs.oracle.com/javase/6/docs/api/java/util/Hashtable.html>

A specially-crafted set of keys could trigger hash function collisions, which can degrade performance of HashMap or Hashtable by changing hash table operations complexity from an expected/average  $O(1)$  to the worst case  $O(n)$ . Reporters were able to find colliding strings efficiently using equivalent substrings and meet in the middle techniques.

This problem can be used to start a **denial of service attack** against Java applications that use untrusted inputs as HashMap or Hashtable keys. An example of such application is web application server (such as tomcat, see [bug #750521](#)) that may fill hash tables with data from HTTP request (such as GET or POST parameters). A remote attack could use that to make JVM use excessive amount of CPU time by sending a POST request with large amount of parameters which hash to the same value.

This problem is similar to the issue that was previously reported for and fixed in e.g. perl:

[http://www.cs.rice.edu/~scrosby/hash/CrosbyWallach\\_UsenixSec2003.pdf](http://www.cs.rice.edu/~scrosby/hash/CrosbyWallach_UsenixSec2003.pdf)

42

## Algorithmic complexity attack on Java

Goal. Find family of strings with the same hashCode().

Solution. The base-31 hash code is part of Java's String API.

key	hashCode()	key	hashCode()	key	hashCode()
"Aa"	2112	"AaAaAaAa"	-540425984	"BBaAaAaAa"	-540425984
"BB"	2112	"AaAaAaBB"	-540425984	"BBaAaAaBB"	-540425984
		"AaAaBBaAa"	-540425984	"BBaAaBBaAa"	-540425984
		"AaAaBBBB"	-540425984	"BBaAaBBBB"	-540425984
		"AaBBaAaAa"	-540425984	"BBBBaAaAa"	-540425984
		"AaBBaAaBB"	-540425984	"BBBBaAaBB"	-540425984
		"AaBBBBaAa"	-540425984	"BBBBBBaAa"	-540425984
		"AaBBBBBB"	-540425984	"BBBBBBaBB"	-540425984

$2^N$  strings of length  $2N$  that hash to same value!

43

## Diversion: one-way hash functions

One-way hash function. "Hard" to find a key that will hash to a desired value (or two keys that hash to same value).

Ex. MD4, MD5, SHA-0, SHA-1, SHA-2, WHIRLPOOL, RIPEMD-160, ....

known to be insecure

```
String password = args[0];
MessageDigest sha1 = MessageDigest.getInstance("SHA1");
byte[] bytes = sha1.digest(password);

/* prints bytes as hex string */
```

Applications. Crypto, message digests, passwords, Bitcoin, ....

Caveat. Too expensive for use in ST implementations.

44


## Separate chaining vs. linear probing

### Separate chaining.

- Performance degrades gracefully.
- Clustering less sensitive to poorly-designed hash function.

### Linear probing.

- Less wasted space.
- Better cache performance.



	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
keys[]	P	M			A	C	S	H	L		E				R	X
vals[]	10	9			8	4	0	5	11		12				3	7

45

## Hashing: variations on the theme

Many improved versions have been studied.

### Two-probe hashing. [ separate-chaining variant ]

- Hash to two positions, insert key in shorter of the two chains.
- Reduces expected length of the longest chain to  $\sim \lg \ln N$ .

### Double hashing. [ linear-probing variant ]

- Use linear probing, but skip a variable amount, not just 1 each time.
- Effectively eliminates clustering.
- Can allow table to become nearly full.
- More difficult to implement delete.

### Cuckoo hashing. [ linear-probing variant ]

- Hash key to two positions; insert key into either position; if occupied, reinsert displaced key into its alternative position (and recur).
- Constant worst-case time for search.



46

## Hash tables vs. balanced search trees

### Hash tables.

- Simpler to code.
- No effective alternative for unordered keys.
- Faster for simple keys (a few arithmetic ops versus  $\log N$  compares).
- Better system support in Java for `String` (e.g., cached hash code).

### Balanced search trees.

- Stronger performance guarantee.
- Support for ordered ST operations.
- Easier to implement `compareTo()` correctly than `equals()` and `hashCode()`.

### Java system includes both.

- Red-black BSTs: `java.util.TreeMap`, `java.util.TreeSet`.
- Hash tables: `java.util.HashMap`, `java.util.IdentityHashMap`.

↑  
linear probing

↑  
separate chaining

47