

Hashing

- hash functions
- collision resolution
- applications

References:
Algorithms in Java, Chapter 14
<http://www.cs.princeton.edu/introalgsds/42hash>

1

Summary of symbol-table implementations

implementation	guarantee			average case			ordered iteration?
	search	insert	delete	search	insert	delete	
unordered array	N	N	N	N/2	N/2	N/2	no
ordered array	lg N	N	N	lg N	N/2	N/2	yes
unordered list	N	N	N	N/2	N	N/2	no
ordered list	N	N	N	N/2	N/2	N/2	yes
BST	N	N	N	1.39 lg N	1.39 lg N	?	yes
randomized BST	7 lg N	7 lg N	7 lg N	1.39 lg N	1.39 lg N	1.39 lg N	yes
red-black tree	3 lg N	3 lg N	3 lg N	lg N	lg N	lg N	yes

Can we do better?

2

Optimize Judiciously

More computing sins are committed in the name of efficiency (without necessarily achieving it) than for any other single reason - including blind stupidity. - William A. Wulf

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

We follow two rules in the matter of optimization:

- Rule 1: Don't do it.
- Rule 2 (for experts only). Don't do it yet - that is, not until you have a perfectly clear and unoptimized solution.

- M. A. Jackson


Reference: Effective Java by Joshua Bloch.

3

Hashing: basic plan

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

Hash function. Method for computing table index from key.



Issues.

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

Classic space-time tradeoff.

- No space limitation: trivial hash function with key as address.
- No time limitation: trivial collision resolution with sequential search.
- Limitations on both time and space: **hashing (the real world).**

4

► hash functions

► collision resolution

► applications

5

Computing the hash function

Idealistic goal: scramble the keys uniformly.

- Efficiently computable.
- Each table position **equally likely** for each key.

← thoroughly researched problem,
still problematic in practical applications

Practical challenge: need different approach for each type of key

Ex: Social Security numbers.

- Bad: first three digits.
- Better: last three digits.

573 = California, 574 = Alaska
assigned in chronological order within a
given geographic region

Ex: date of birth.

- Bad: birth year.
- Better: birthday.

Ex: phone numbers.

- Bad: first three digits.
- Better: last three digits.

6

Hash Codes and Hash Functions

Java convention: all classes implement `hashCode()`

`hashCode()` returns a 32-bit `int` (between -2147483648 and 2147483647)

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

First try:

```
String s = "call";  
int code = s.hashCode();  
int hash = code % M;
```

7121 → code
8191 → M
3045982 → code % M

Bug. Don't use $(code \% M)$ as array index

1-in-a billion bug. Don't use $(Math.abs(code) \% M)$ as array index.

OK. Safe to use $((code \& 0x7fffffff) \% M)$ as array index.

← hex literal 31-bit mask

7

Java's `hashCode()` convention

Theoretical advantages


- Ensures hashing can be used for every type of object
- Allows expert implementations suited to each type

Requirements:

- If `x.equals(y)` then `x` and `y` must have the same hash code.
- Repeated calls to `x.hashCode()` must return the same value.

Practical realities

- True randomness is hard to achieve
- Cost is an important consideration



Available implementations

- default (inherited from `Object`): Memory address of `x` (!!!)
- customized Java implementations: `String`, `URL`, `Integer`, `Date`.
- User-defined types: **users are on their own**

← that's you!

8

A typical type

Assumption when using hashing in Java:

Key type has reasonable implementation of `hashCode()` and `equals()`

Ex. Phone numbers: (609) 867-5309.

exchange extension

```
public final class PhoneNumber
{
    private final int area, exch, ext;
    public PhoneNumber(int area, int exch, int ext)
    {
        this.area = area;
        this.exch = exch;
        this.ext = ext;
    }
    public boolean equals(Object y) { // as before }
    public int hashCode()
    { return 10007 * (area + 1009 * exch) + ext; }
}
```

sufficiently random?

Fundamental problem:

Need a **theorem** for each data type to ensure reliability.

9

A decent hash code design

Java 1.5 string library [see also Program 14.2 in Algs in Java].

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

*i*th character of s

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

- Equivalent to $h = 31^{L-1} \cdot s_0 + \dots + 31^2 \cdot s_{L-3} + 31 \cdot s_{L-2} + s_{L-1}$.
- Horner's method to hash string of length L: L multiplies/adds

Ex.

```
String s = "call";
int code = s.hashCode();
```

$$3045982 = 99 \cdot 31^3 + 97 \cdot 31^2 + 108 \cdot 31 + 108 \cdot 31^0 \\ = 108 + 31 \cdot (108 + 31 \cdot (99 + 31 \cdot (97)))$$

Provably random? Well, no.

10

A poor hash code design

Java 1.1 string library.

- For long strings: only examines 8-9 evenly spaced characters.
- Saves time in performing arithmetic...

```
public int hashCode()
{
    int hash = 0;
    int skip = Math.max(1, length() / 8);
    for (int i = 0; i < length(); i += skip)
        hash = (37 * hash) + s[i];
    return hash;
}
```

but great potential for bad collision patterns.

```
http://www.cs.princeton.edu/introcs/13loop/Hello.java
http://www.cs.princeton.edu/introcs/13loop/Hello.class
http://www.cs.princeton.edu/introcs/13loop/Hello.html
http://www.cs.princeton.edu/introcs/13loop/index.html
http://www.cs.princeton.edu/introcs/12type/index.html
```

Basic rule: need to use the whole key.

11

Digression: using a hash function for data mining

Use **content** to characterize documents.

Applications

- Search documents on the web for documents similar to a given one.
- Determine whether a new document belongs in one set or another


Approach

- Fix **order** k and **dimension** d
- Compute `hashCode()` % d for all k-grams in the document
- Result: d-dimensional vector **profile** of each document
- To compare documents:

Consider angle θ separating vectors

$\cos \theta$ close to 0: not similar

$\cos \theta$ close to 1: similar



$$\cos \theta = a \cdot b / |a| |b|$$

Effective for literature, genomes, Java code, art, music, data, video

12

Digression: using a hash function for data mining

```
% more tale.txt
it was the best of times
it was the worst of times
it was the age of wisdom
it was the age of
foolishness
...
```

```
% more genome.txt
CTTTCGGTTTGAAC
GAAGCCGCGCTCT
TGCTGCTGCAGC
ATCGTTC
...
```

$\cos \theta$ small: **not** similar

$k = 10$
 $d = 65536$

tale.txt			genome.txt		
i	10-grams with hashCode() i	freq	10-grams with hashCode() i	freq	
0		0		0	
1		0		0	
2		0		0	
...					
435	best of ti foolishnes	2	TTTCGGTTTG TGCTGCTGC	2	
8999	it was the	8		0	
...					
12122		0	CTTTCGGTTT	3	
...					
34543	t was the b	5	ATCGGTCGA	4	
...					
65535					
65536					

profiles

13

Digression: using a hash function to profile a document for data mining

```
public class Document
{
    private String name;
    private double[] profile;
    public Document(String name, int k, int d)
    {
        this.name = name;
        String doc = (new In(name)).readAll();
        int N = doc.length();
        profile = new double[d];
        for (int i = 0; i < N-k; i++)
        {
            int h = doc.substring(i, i+k).hashCode();
            profile[Math.abs(h % d)] += 1;
        }
    }
    public double simTo(Document other)
    {
        // compute dot product and divide by magnitudes
    }
}
```

14

Digression: using a hash function to compare documents

```
public class CompareAll
{
    public static void main(String args[])
    {
        int k = Integer.parseInt(args[0]);
        int d = Integer.parseInt(args[1]);
        int N = StdIn.readInt();
        Document[] a = new Document[N];
        for (int i = 0; i < N; i++)
            a[i] = new Document(StdIn.readString(), k, d);
        System.out.print(" ");
        for (int j = 0; j < N; j++)
            System.out.printf(" %.4s", a[j].name());
        System.out.println();
        for (int i = 0; i < N; i++)
        {
            System.out.printf("%.4s ", a[i].name());
            for (int j = 0; j < N; j++)
                System.out.printf("%8.2f", a[i].simTo(a[j]));
            System.out.println();
        }
    }
}
```

15

Digression: using a hash function to compare documents

Cons	US Constitution
TomS	"Tom Sawyer"
Huck	"Huckleberry Finn"
Prej	"Pride and Prejudice"
Pict	a photograph
DJIA	financial data
Amaz	Amazon.com website .html source
ACTG	genome


```
% java CompareAll 5 1000 < docs.txt
```

	Cons	TomS	Huck	Prej	Pict	DJIA	Amaz	ACTG
Cons	1.00	0.89	0.87	0.88	0.35	0.70	0.63	0.58
TomS	0.89	1.00	0.98	0.96	0.34	0.75	0.66	0.62
Huck	0.87	0.98	1.00	0.94	0.32	0.74	0.65	0.61
Prej	0.88	0.96	0.94	1.00	0.34	0.76	0.67	0.63
Pict	0.35	0.34	0.32	0.34	1.00	0.29	0.48	0.24
DJIA	0.70	0.75	0.74	0.76	0.29	1.00	0.62	0.58
Amaz	0.63	0.66	0.65	0.67	0.48	0.62	1.00	0.45
ACTG	0.58	0.62	0.61	0.63	0.24	0.58	0.45	1.00

16

- ▶ hash functions
- ▶ **collision resolution**
- ▶ applications

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



The diagram shows 16 vertical red lines representing bins, indexed 0 to 15 on a horizontal axis. The distribution of balls is as follows:

Bin Index	Number of Balls
0	1
2	1
4	2
5	1
7	1
10	1
11	1
14	3

Birthday problem.

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

Coupon collector.

Expect every bin has ≥ 1 ball after $\Theta(M \ln M)$ tosses.

Load balancing.

After M tosses, expect most loaded bin has $\Theta(\log M / \log \log M)$ balls.

18



Coupon collector.

Load balancing.

Collisions

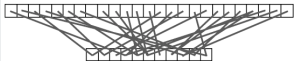
Collision. Two distinct keys hashing to same index.

Conclusion. Birthday problem \Rightarrow can't avoid collisions unless you have a ridiculous amount of memory.

Challenge. Deal with collisions efficiently.

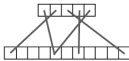
Approach 1:
accept multiple collisions

25 items, 11 table positions
~2 items per table position



Approach 2:
minimize collisions

5 items, 11 table positions
~.5 items per table position



19

Challenge. Deal with collisions efficiently.

5 items, 11 table positions
~ .5 items per table position

The diagram shows a 2D hexagonal lattice with 11 positions. A top row of 5 positions contains 5 items, represented by small squares with diagonal lines. Below this, a second row of 6 positions is shown, with lines connecting the top row to the second row to form a triangular pattern. The bottom row consists of 11 positions, with lines connecting the second row to the bottom row to form a larger triangular pattern.

1. Separate chaining. [H. P. Luhn, IBM 1953]
Put keys that collide in a list associated with index.

2. Open addressing. [Amdahl-Boehme-Rochester-Samuel, IBM 1953]
When a new key collides, find next empty slot, and put it there.

The diagram illustrates two hash table implementations. On the left, 'separate chaining' shows an array of pointers (st[0] to st[8190]) where each pointer points to a linked list of keys. For example, st[0] points to 'jocularly', which points to 'seriously'. st[1] points to 'listen'. st[2] is null. st[3] points to 'suburban', which points to 'untravelling', which points to 'considering'. st[8190] points to 'browsing'. On the right, 'linear probing' shows an array of slots (st[0] to st[30001]) where keys are stored in the next available slot. For example, st[0] contains 'jocularly', st[1] is null, st[2] contains 'listen', st[3] contains 'suburban', st[30001] contains 'browsing', and all other slots are null.

```
graph LR
    subgraph Separate_Chaining [separate chaining]
        st0[st[0]] --> jocularly[jocularly]
        jocularly --> seriously[seriously]
        st1[st[1]] --> listen[listen]
        st2[st[2]] --> null2[null]
        st3[st[3]] --> suburban[suburban]
        suburban --> untravelling[untravelling]
        untravelling --> considering[considering]
        st8190[st[8190]] --> browsing[browsing]
    end

    subgraph Linear_Probing [linear probing]
        st0[st[0]] --> jocularly2[jocularly]
        st1[st[1]] --> null1[null]
        st2[st[2]] --> listen2[listen]
        st3[st[3]] --> suburban2[suburban]
        st30001[st[30001]] --> browsing2[browsing]
        st4[st[4]] --> null4[null]
        st5[st[5]] --> null5[null]
        st6[st[6]] --> null6[null]
        st7[st[7]] --> null7[null]
        st8[st[8]] --> null8[null]
        st9[st[9]] --> null9[null]
        st10[st[10]] --> null10[null]
        st11[st[11]] --> null11[null]
        st12[st[12]] --> null12[null]
        st13[st[13]] --> null13[null]
        st14[st[14]] --> null14[null]
        st15[st[15]] --> null15[null]
        st16[st[16]] --> null16[null]
        st17[st[17]] --> null17[null]
        st18[st[18]] --> null18[null]
        st19[st[19]] --> null19[null]
        st20[st[20]] --> null20[null]
        st21[st[21]] --> null21[null]
        st22[st[22]] --> null22[null]
        st23[st[23]] --> null23[null]
        st24[st[24]] --> null24[null]
        st25[st[25]] --> null25[null]
        st26[st[26]] --> null26[null]
        st27[st[27]] --> null27[null]
        st28[st[28]] --> null28[null]
        st29[st[29]] --> null29[null]
        st30[st[30]] --> null30[null]
        st31[st[31]] --> null31[null]
        st32[st[32]] --> null32[null]
        st33[st[33]] --> null33[null]
        st34[st[34]] --> null34[null]
        st35[st[35]] --> null35[null]
        st36[st[36]] --> null36[null]
        st37[st[37]] --> null37[null]
        st38[st[38]] --> null38[null]
        st39[st[39]] --> null39[null]
        st40[st[40]] --> null40[null]
        st41[st[41]] --> null41[null]
        st42[st[42]] --> null42[null]
        st43[st[43]] --> null43[null]
        st44[st[44]] --> null44[null]
        st45[st[45]] --> null45[null]
        st46[st[46]] --> null46[null]
        st47[st[47]] --> null47[null]
        st48[st[48]] --> null48[null]
        st49[st[49]] --> null49[null]
        st50[st[50]] --> null50[null]
        st51[st[51]] --> null51[null]
        st52[st[52]] --> null52[null]
        st53[st[53]] --> null53[null]
        st54[st[54]] --> null54[null]
        st55[st[55]] --> null55[null]
        st56[st[56]] --> null56[null]
        st57[st[57]] --> null57[null]
        st58[st[58]] --> null58[null]
        st59[st[59]] --> null59[null]
        st60[st[60]] --> null60[null]
        st61[st[61]] --> null61[null]
        st62[st[62]] --> null62[null]
        st63[st[63]] --> null63[null]
        st64[st[64]] --> null64[null]
        st65[st[65]] --> null65[null]
        st66[st[66]] --> null66[null]
        st67[st[67]] --> null67[null]
        st68[st[68]] --> null68[null]
        st69[st[69]] --> null69[null]
        st70[st[70]] --> null70[null]
        st71[st[71]] --> null71[null]
        st72[st[72]] --> null72[null]
        st73[st[73]] --> null73[null]
        st74[st[74]] --> null74[null]
        st75[st[75]] --> null75[null]
        st76[st[76]] --> null76[null]
        st77[st[77]] --> null77[null]
        st78[st[78]] --> null78[null]
        st79[st[79]] --> null79[null]
        st80[st[80]] --> null80[null]
        st81[st[81]] --> null81[null]
        st82[st[82]] --> null82[null]
        st83[st[83]] --> null83[null]
        st84[st[84]] --> null84[null]
        st85[st[85]] --> null85[null]
        st86[st[86]] --> null86[null]
        st87[st[87]] --> null87[null]
        st88[st[88]] --> null88[null]
        st89[st[89]] --> null89[null]
        st90[st[90]] --> null90[null]
        st91[st[91]] --> null91[null]
        st92[st[92]] --> null92[null]
        st93[st[93]] --> null93[null]
        st94[st[94]] --> null94[null]
        st95[st[95]] --> null95[null]
        st96[st[96]] --> null96[null]
        st97[st[97]] --> null97[null]
        st98[st[98]] --> null98[null]
        st99[st[99]] --> null99[null]
        st100[st[100]] --> null100[null]
        st101[st[101]] --> null101[null]
        st102[st[102]] --> null102[null]
        st103[st[103]] --> null103[null]
        st104[st[104]] --> null104[null]
        st105[st[105]] --> null105[null]
        st106[st[106]] --> null106[null]
        st107[st[107]] --> null107[null]
        st108[st[108]] --> null108[null]
        st109[st[109]] --> null109[null]
        st110[st[110]] --> null110[null]
        st111[st[111]] --> null111[null]
        st112[st[112]] --> null112[null]
        st113[st[113]] --> null113[null]
        st114[st[114]] --> null114[null]
        st115[st[115]] --> null115[null]
        st116[st[116]] --> null116[null]
        st117[st[117]] --> null117[null]
        st118[st[118]] --> null118[null]
        st119[st[119]] --> null119[null]
        st120[st[120]] --> null120[null]
        st121[st[121]] --> null121[null]
        st122[st[122]] --> null122[null]
        st123[st[123]] --> null123[null]
        st124[st[124]] --> null124[null]
        st125[st[125]] --> null125[null]
        st126[st[126]] --> null126[null]
        st127[st[127]] --> null127[null]
        st128[st[128]] --> null128[null]
        st129[st[129]] --> null129[null]
        st130[st[130]] --> null130[null]
        st131[st[131]] --> null131[null]
        st132[st[132]] --> null132[null]
        st133[st[133]] --> null133[null]
        st134[st[134]] --> null134[null]
        st135[st[135]] --> null135[null]
        st136[st[136]] --> null136[null]
        st137[st[137]] --> null137[null]
        st138[st[138]] --> null138[null]
        st139[st[139]] --> null139[null]
        st140[st[140]] --> null140[null]
        st141[st[141]] --> null141[null]
        st142[st[142]] --> null142[null]
        st143[st[143]] --> null143[null]
        st144[st[144]] --> null144[null]
        st145[st[145]] --> null145[null]
        st146[st[146]] --> null146[null]
        st147[st[147]] --> null147[null]
        st148[st[148]] --> null148[null]
        st149[st[149]] --> null149[null]
        st150[st[150]] --> null150[null]
        st151[st[151]] --> null151[null]
        st152[st[152]] --> null152[null]
        st153[st[153]] --> null153[null]
        st154[st[154]] --> null154[null]
        st155[st[155]] --> null155[null]
        st156[st[156]] --> null156[null]
        st157[st[157]] --> null157[null]
        st158[st[158]] --> null158[null]
        st159[st[159]] --> null159[null]
        st160[st[160]] --> null160[null]
        st161[st[
```

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

```

st[0]      jocularly  →  seriously
st[1]      listen
st[2]      null
st[3]      suburban  →  untravalled  →  considering
⋮
st[8190]   browsing

```

easy extension of linked list ST implementation

linear probing ($M = 30001, N = 15000$)

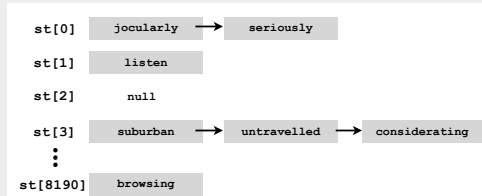
20

Collision resolution approach 1: separate chaining

Use an array of $M < N$ linked lists.

← good choice: $M = N/10$

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



key	hash
call	7121
me	3480
ishmael	5017
seriously	0
untravelling	3
suburban	3
...	...

21

Separate chaining ST implementation (skeleton)

could use doubling →

no generics in arrays in Java →

```

public class ListHashST<Key, Value>
{
    private int M = 8191;
    private Node[] st = new Node[M];

    private class Node
    {
        Object key;
        Object val;
        Node next;
        Node(Key key, Value val, Node next)
        {
            this.key = key;
            this.val = val;
            this.next = next;
        }
    }

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

    public void put(Key key, Value val)
    // see next slide

    public Val get(Key key)
    // see next slide
}
    
```

compare with linked lists

22

Separate chaining ST implementation (put and get)

```

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, value, first);
}

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 (Value) x.val;
    return null;
}
    
```

Identical to linked-list code, except hash to pick a list.

23

Analysis of separate chaining

Separate chaining performance.

- Cost is proportional to length of list.
- Average length = N / M .
- **Worst case:** all keys hash to same list.

Theorem. Let $\alpha = N / M > 1$ be average length of list. For any $t > 1$, probability that list length $> t \alpha$ is exponentially small in t .

↑
depends on hash map being random map

Parameters.

- M too large \Rightarrow too many empty chains.
- M too small \Rightarrow chains too long.
- Typical choice: $\alpha = N / M \approx 10 \Rightarrow$ **constant-time** ops.

24

Collision resolution approach 2: open addressing

Use an array of size $M \gg N$.

← good choice: $M = 2N$

- Hash: map key to integer i between 0 and $M-1$.

Linear probing:

- Insert: put in slot i if free; if not try $i+1, i+2$, etc.
- Search: search slot i ; if occupied but no match, try $i+1, i+2$, etc.

-	-	-	S	H	-	-	A	C	E	R	-	N
0	1	2	3	4	5	6	7	8	9	10	11	12

-	-	-	S	H	-	-	A	C	E	R	I	-
0	1	2	3	4	5	6	7	8	9	10	11	12

insert I
hash(I) = 11

-	-	-	S	H	-	-	A	C	E	R	I	N
0	1	2	3	4	5	6	7	8	9	10	11	12

insert N
hash(N) = 8

25

Linear probing ST implementation

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

    private int hash(Key key) // as before

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

    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;
    }
}
```

← standard ugly casts

← compare with elementary unordered array implementation

← standard array doubling code omitted (double when half full)

26

Clustering

Cluster. A contiguous block of items.

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

-	-	-	S	H	A	C	E	-	-	-	X	M	I	-	-	-	P	-	-	R	L	-	-
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

cluster

Knuth's parking problem. Cars arrive at one-way street with M parking spaces. Each desires a random space i : if space i is taken, try $i+1, i+2$, ... What is mean displacement of a car?



Empty. With $M/2$ cars, mean displacement is about $3/2$.

Full. Mean displacement for the last car is about $\sqrt{\pi M / 2}$

27

Analysis of linear probing

Linear probing performance.

- Insert and search cost depend on length of cluster.
- Average length of cluster = $\alpha = N / M$.
- Worst case: all keys hash to same cluster.

← but keys more likely to hash to big clusters

Theorem. [Knuth 1962] Let $\alpha = N / M < 1$ be the load factor.

Average probes for insert/search miss

$$\frac{1}{2} \left(1 + \frac{1}{(1-\alpha)^2} \right) = (1 + \alpha + 2\alpha^2 + 3\alpha^3 + 4\alpha^4 + \dots) / 2$$

Average probes for search hit

$$\frac{1}{2} \left(1 + \frac{1}{(1-\alpha)} \right) = 1 + (\alpha + \alpha^2 + \alpha^3 + \alpha^4 + \dots) / 2$$

Parameters.

- Load factor too small \Rightarrow too many empty array entries.
- Load factor too large \Rightarrow clusters coalesce.
- Typical choice: $M \approx 2N \Rightarrow$ constant-time ops.

28

Hashing: variations on the theme

Many improved versions have been studied:

Ex: Two-probe hashing

- hash to two positions, put key in shorter of the two lists
- reduces average length of the longest list to $\log \log N$

Ex: Double hashing

- use linear probing, but skip a variable amount, not just 1 each time
- effectively eliminates clustering
- can allow table to become nearly full

29

Double hashing


Idea Avoid clustering by using second hash to compute skip for search.

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

Second hash. Map key to nonzero skip value k .

Ex: $k = 1 + (v \bmod 97)$.

↑
hashCode ()



Effect. Skip values give different search paths for keys that collide.

Best practices. Make k and M relatively prime.

30

Double Hashing Performance

Theorem. [Guibas-Szemerédi] Let $\alpha = N / M < 1$ be average length of list.

Average probes for insert/search miss

$$\frac{1}{(1 - \alpha)} = 1 + \alpha + \alpha^2 + \alpha^3 + \alpha^4 + \dots$$

Average probes for search hit

$$\frac{1}{\alpha} \ln \frac{1}{(1 - \alpha)} = 1 + \alpha/2 + \alpha^2/3 + \alpha^3/4 + \alpha^4/5 + \dots$$

Parameters. Typical choice: $\alpha \approx 1.2 \Rightarrow$ constant-time ops.

Disadvantage. Delete cumbersome to implement.

31

Hashing Tradeoffs

Separate chaining vs. linear probing/double hashing.

- Space for links vs. empty table slots.
- Small table + linked allocation vs. big coherent array.

Linear probing vs. double hashing.

		load factor α			
		50%	66%	75%	90%
linear probing	get	1.5	2.0	3.0	5.5
	put	2.5	5.0	8.5	55.5
double hashing	get	1.4	1.6	1.8	2.6
	put	1.5	2.0	3.0	5.5

number of probes

32

Summary of symbol-table implementations

implementation	guarantee			average case			ordered iteration?	operations on keys
	search	insert	delete	search	insert	delete		
unordered array	N	N	N	N/2	N/2	N/2	no	<code>equals()</code>
ordered array	$\lg N$	N	N	$\lg N$	N/2	N/2	yes	<code>compareTo()</code>
unordered list	N	N	N	N/2	N	N/2	no	<code>equals()</code>
ordered list	N	N	N	N/2	N/2	N/2	yes	<code>compareTo()</code>
BST	N	N	N	$1.38 \lg N$	$1.38 \lg N$?	yes	<code>compareTo()</code>
randomized BST	$7 \lg N$	$7 \lg N$	$7 \lg N$	$1.38 \lg N$	$1.38 \lg N$	$1.38 \lg N$	yes	<code>compareTo()</code>
red-black tree	$2 \lg N$	$2 \lg N$	$2 \lg N$	$\lg N$	$\lg N$	$\lg N$	yes	<code>compareTo()</code>
hashing	1*	1*	1*	1*	1*	1*	no	<code>equals()</code> <code>hashCode()</code>

* assumes random hash code

33

Hashing versus balanced trees

Hashing

- simpler to code
- no effective alternative for **unordered** keys
- faster for simple keys (a few arithmetic ops versus $\lg N$ compares)
- (Java) better system support for strings [cached hashCode]
- does your hash function produce **random** values for your key type??

Balanced trees

- stronger performance guarantee
- **can support many more operations for ordered keys**
- easier to implement `compareTo()` correctly than `equals()` and `hashCode()`

Java system includes both

- red-black trees: `java.util.TreeMap`, `java.util.TreeSet`
- hashing: `java.util.HashMap`, `java.util.IdentityHashMap`

34

Typical "full" ST API

```
public class *ST<Key extends Comparable<Key>, Value>
{
    *ST()                create a symbol table
    void put(Key key, Value val) put key-value pair into the table
    Value get(Key key)    return value paired with key
                        (null if key is not in table)
    boolean contains(Key key) is there a value paired with key?
    Key min()             smallest key
    Key max()             largest key
    Key next(Key key)     next largest key (null if key is max)
    Key prev(Key key)     next smallest key (null if key is min)
    void remove(Key key)  remove key-value pair from table
    Iterator<Key> iterator() iterator through keys in table
}
```

Hashing is **not** suitable for implementing such an API (no order)

BSTs are **easy** to extend to support such an API (basic tree ops)

Ex: Can use LLRB trees implement priority queues for distinct keys

35

- hash functions
- collision resolution
- applications

36

Set ADT

Set. Collection of distinct keys.

<code>public class *SET<Key extends Comparable<Key>, Value></code>	
<code>SET()</code>	create a set
<code>void add(Key key)</code>	put key into the set
<code>boolean contains(Key key)</code>	is there a value paired with key?
<code>void remove(Key key)</code>	remove key from the set
<code>Iterator<Key> iterator()</code>	iterator through all keys in the set

Normal mathematical assumption: collection is **unordered**
 Typical (eventual) client expectation: **ordered** iteration

Q. How to implement?

A0. Hashing (our ST code [value removed] or `java.util.HashSet`)

A1. Red-black BST (our ST code [value removed] or `java.util.TreeSet`)

unordered iterator
 $O(1)$ search
 ordered iterator
 $O(\log N)$ search

37

SET client example 1: dedup filter

Remove duplicates from strings in standard input

- Read a key.
- If key is not in set, insert and print it.

```
public class DeDup
{
    public static void main(String[] args)
    {
        SET<String> set = new SET<String>();
        while (!StdIn.isEmpty())
        {
            String key = StdIn.readString();
            if (!set.contains(key))
            {
                set.add(key);
                StdOut.println(key);
            }
        }
    }
}
```

No iterator needed.
 Output is in same order
 as input with
 dups removed.

```
% more tale.txt
it was the best of times
it was the worst of times
it was the age of wisdom
it was the age of
foolishness
...

% java DeDup < tale.txt
it
was
the
best
of
times
worst
age
wisdom
foolishness
...
```

Simplified version of FrequencyCount (no iterator needed)

38

SET client example 2A: lookup filter

Print words from standard input that are found in a list

- Read in a list of words from one file.
- Print out all words from standard input that are in the list.

```
public class LookupFilter
{
    public static void main(String[] args)
    {
        SET<String> set = new SET<String>();
        In in = new In(args[0]);
        while (!in.isEmpty())
            set.add(in.readString());

        while (!StdIn.isEmpty())
        {
            String word = StdIn.readString();
            if (set.contains(word))
                StdOut.println(word);
        }
    }
}
```

create SET

process list

print words that
are not in list

39

SET client example 2B: exception filter

Print words from standard input that are **not** found in a list

- Read in a list of words from one file.
- Print out all words from standard input that are **not** in the list.

```
public class LookupFilter
{
    public static void main(String[] args)
    {
        SET<String> set = new SET<String>();
        In in = new In(args[0]);
        while (!in.isEmpty())
            set.add(in.readString());

        while (!StdIn.isEmpty())
        {
            String word = StdIn.readString();
            if (!set.contains(word))
                StdOut.println(word);
        }
    }
}
```

create SET

process list

print words that
are not in list

40

SET filter applications

application	purpose	key	type	in list	not in list
dedup	eliminate duplicates		dedup	duplicates	unique keys
spell checker	find misspelled words	word	exception	dictionary	misspelled words
browser	mark visited pages	URL	lookup	visited pages	
chess	detect draw	board	lookup	positions	
spam filter	eliminate spam	IP addr	exception	spam	good mail
trusty filter	allow trusted mail	URL	lookup	good mail	
credit cards	check for stolen cards	number	exception	stolen cards	good cards

41

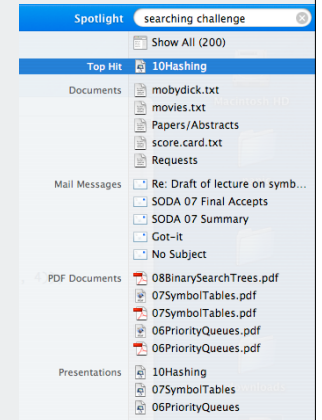
Searching challenge:

Problem: Index for a PC or the web

Assumptions: 1 billion++ words to index

Which searching method to use?

- 1) hashing implementation of SET
- 2) hashing implementation of ST
- 3) red-black-tree implementation of ST
- 4) red-black-tree implementation of SET
- 5) doesn't matter much



42

Index for search in a PC

```
ST<String, SET<File>> st = new ST<String, SET<File>>();
for (File f: filesystem)
{
    In in = new In(f);
    String[] words = in.readAll().split("\\s+");
    for (int i = 0; i < words.length; i++)
    {
        String s = words[i];
        if (!st.contains(s))
            st.put(s, new SET<File>());
        SET<File> files = st.get(s);
        files.add(f);
    }
}
```

← build index

```
SET<File> files = st.get(s);  
for (File f: files) ...
```

← process lookup

← process lookup request

43


Searching challenge:

Problem: Index for a book

Assumptions: book has 100,000+ words

Which searching method to use?

- 1) hashing implementation of SET
- 2) hashing implementation of ST
- 3) red-black-tree implementation of ST
- 4) red-black-tree implementation of SET
- 5) doesn't matter much



44

Index for a book

```
public class Index
{
    public static void main(String[] args)
    {
        String[] words = StdIn.readAll().split("\\s+");
        ST<String, SET<Integer>> st;
        st = new ST<String, SET<Integer>>();

        for (int i = 0; i < words.length; i++)
        {
            String s = words[i];
            if (!st.contains(s))
                st.put(s, new SET<Integer>());
            SET<Integer> pages = st.get(s);
            pages.add(page(i));
        }

        for (String s : st)
            StdOut.println(s + ": " + st.get(s));
    }
}
```

read book and create ST

process all words

print index!

Requires ordered iterators (not hashing)

45

Hashing in the wild: Java implementations

Java has built-in libraries for hash tables.

- java.util.HashMap = separate chaining implementation.
- java.util.IdentityHashMap = linear probing implementation.

```
import java.util.HashMap;
public class HashMapDemo
{
    public static void main(String[] args)
    {
        HashMap<String, String> st = new HashMap<String, String>();
        st.put("www.cs.princeton.edu", "128.112.136.11");
        st.put("www.princeton.edu", "128.112.128.15");
        StdOut.println(st.get("www.cs.princeton.edu"));
    }
}
```

Null value policy.

- Java HashMap allows null values.
- Our implementation forbids null values.

46

Using HashMap

Implementation of our API with java.util.HashMap.

```
import java.util.HashMap;
import java.util.Iterator;

public class ST<Key, Value> implements Iterable<Key>
{
    private HashMap<Key, Value> st = new HashMap<Key, Value>();

    public void put(Key key, Value val)
    {
        if (val == null) st.remove(key);
        else st.put(key, val);
    }

    public Value get(Key key) { return st.get(key); }
    public Value remove(Key key) { return st.remove(key); }
    public boolean contains(Key key) { return st.contains(key); }
    public int size() { return st.size(); }
    public Iterator<Key> iterator() { return st.keySet().iterator(); }
}
```


47

Hashing in the wild: algorithmic complexity attacks

Is the random hash map assumption important in practice?

- Obvious situations: aircraft control, nuclear reactor, pacemaker.
- Surprising situations: denial-of-service attacks.

malicious adversary learns your ad hoc hash function
(e.g., by reading Java API) and causes a big pile-up in
single address that grinds performance to a halt



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.

Reference: <http://www.cs.rice.edu/~scrosby/hash>

48

Algorithmic complexity attack on the Java Library

Goal. Find strings with the same hash code.

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

Key	hashCode()
Aa	2112
BB	2112

Key	hashCode()
AaAaAaAa	-540425984
AaAaAaBB	-540425984
AaAaBBBb	-540425984
AaAaBBBB	-540425984
AaBBbAaA	-540425984
AaBBbAaB	-540425984
AaBBBBbA	-540425984
AaBBBBBB	-540425984
BBbAaAaA	-540425984
BBbAaAaB	-540425984
BBbAaBBb	-540425984
BBbAaBBB	-540425984
BBBBbAaA	-540425984
BBBBbAaB	-540425984
BBBBBBbA	-540425984
BBBBBBBB	-540425984

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

Does your hash function
produce **random** values
for your key type??

49

One-Way Hash Functions

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

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

insecure

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

// prints bytes as hex string
```

Applications. Digital fingerprint, message digest, storing passwords.

Too expensive for use in ST implementations (use balanced trees)

50