

# Binary search

## doing it less wrong

Paul Khuong

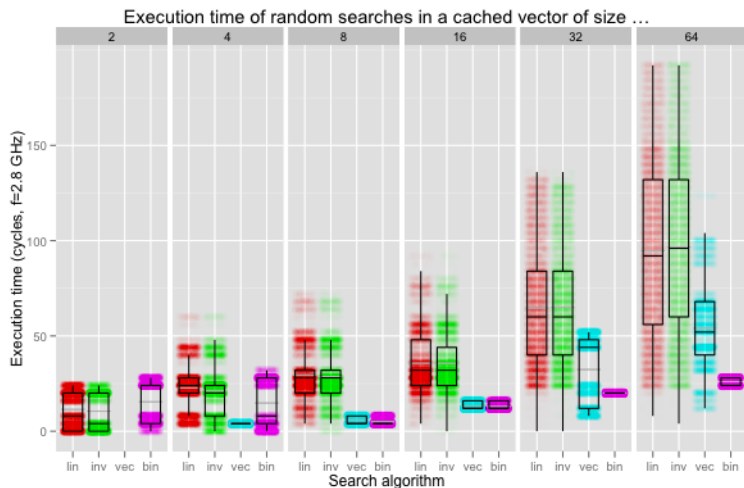
October 30, 2014

## “Binary search is slow”

- ▶ “Linear search is faster for small  $n$ ” (branches)
- ▶ “Fancy layouts scale better” (caches)

# No branch, no misprediction

Data dependent conditional move.



<http://pvk.ca/Blog/2012/07/03/binary-search-star-eliminates-star-branch-mispredictions/>

## Don't... look for early matches

```
if (*mid == needle) {  
    return mid;  
}
```

## Don't... try to adjust bounds tightly

```
if (needle < *mid) {  
    len = half;  
} else {  
    low = mid + 1;  
    len -= half + 1;  
}
```

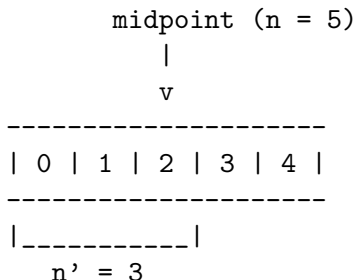
G++ STL

## Don't...do both

```
if (comparison < 0) {  
    high = mid;  
} else if (comparison > 0) {  
    low = mid + 1;  
} else {  
    return mid;  
}
```

glibc, FreeBSD libc

## Simple binary search




```
while ((half = n / 2) > 0) {
    mid = low + half;
    low = (*mid < needle) ? mid : low;
    n -= half;
}
```

So simple, it's AVX2-able!

## Assume a decent compiler

```
loop:  
  lea    (%rdx,%rcx,4), %rdi  
  cmp    (%rdi), %esi  
  cmovge %rdi, %rdx  
  sub    %rcx, %rax  
  mov    %rax, %rcx  
  shr    %rcx  
  jnz    loop
```





# Microbenchmark


Implementations:

- ▶ branch (STL)
- ▶ both (libc)
- ▶ early (only early termination)
- ▶ simple (cmov)

Input: 32 bit ints (random,  $\approx 5\%$  density)

- ▶ 8, 16, 32, ..., 1024
- ▶ 10, 50, 100, 200, ..., 1000

Report average of 128 lookups (median, 1st/99th percentile)



## Caching: $n = 2^k(-1) \Rightarrow$ aliasing issues

Midpoints:

0x200000

0x100000

0x080000

0x040000

0x020000

0x010000


.....

- ▶ Run proper microbenchmarks
- ▶ Offset “mid” point  $((n / 2) + (n / 64))$
- ▶ 3-way (“ternary”) search

In the wild: Bentley & Saxe dynamisation or  $2^k \leq n < 2^{k+1}$ .

<http://pvk.ca/Blog/2012/07/30/binary-search-is-a-pathological-case-for-caches/>

## Practical use case: sparse bitmatrix mult + projection



## a.k.a. (pre)sorted equijoin

Inner loop:

- ▶ branch-free (simple)
- ▶ branches (STL)
- ▶ unrolled (3ary branch-free)

Reuse results

- ▶ roving lower bound
- ▶ galloping search

# Gather phase around peak time



# Sorted array search: a decent finger search

Let  $\Delta = k_i - k_{i-1}$

Galloping search:  $\approx 2 \lg \Delta$  comparisons

3ary search:  $\approx 2 \log_3 \Delta$  D\$ misses


Roving search:  $\approx \lg n$  D\$ misses

```

:          0000000000000003f890 <sparse_tsearch_upto_14>:
1.07:   63f800: 48 c1 e6 06      shl     $0x6,%rsi
0.03:   63f804: 55              push   %rbp
0.43:   63f805: 4d 8d 84 3c 00 00 f8      lea    0xffffffff8000000(%rsi,%rdi,1),%rax
0.00:   63f808: ff              fi
0.28:   63f80d: 4d 89 e5        mov    %rsip,%rbp
0.20:   63f80e: 39 10          cmp    %edx,%rax
0.87:   63f822: 4d 0f 42 f8     movb  %rax,%rdi
1.44:   63f825: 39 97 c8 aa 02 00     cmp    0x2aa0c(%rdi),%edx
3.67:   63f82c: 4d 84 8f 40 35 85 80     lea    0x3580(%rax),%rcx
0.00:   63f833: 4d 84 8f c0 aa 82 00     lea    0x2aa0c(%rdi),%rcx
0.03:   63f83a: 4d 0f 46 c7     movb  %rdi,%rax
1.44:   63f83e: 3b 11          cmp    %rcx,%edx
0.00:   63f840: 4d 0f 46 c6     movb  %rcx,%rcx
0.97:   63f844: 39 91 c3 83 00 00     cmp    0xa30(%rcx),%edx
3.46:   63f84a: 4d 84 81 00 c7 01 00     lea    0xc700(%rcx),%rcx
0.15:   63f851: 4d 8d b1 c0 e3 00 00     lea    0xe30(%rcx),%rsi
0.00:   63f858: 4d 0f 46 f1     movb  %rcx,%rsi
1.73:   63f85c: 3b 10          cmp    %rax,%edx
0.03:   63f85e: 4d 0f 46 c6     movb  %rsi,%rcx
1.49:   63f862: 39 90 00 4c 00 00     cmp    0x4c0(%rcx),%edx
3.89:   63f868: 4d 84 88 c9 97 00 00     lea    0x970(%rcx),%rcx
0.00:   63f86f: 4d 8d b0 00 4c 00 00     lea    0x4c0(%rcx),%rsi
0.00:   63f876: 4d 0f 46 f0     movb  %rcx,%rsi
1.86:   63f87a: 3b 11          cmp    %rcx,%edx
0.10:   63f87c: 4d 0f 46 ce     movb  %rsi,%rcx
1.55:   63f880: 39 91 80 19 00 00     cmp    0x190(%rcx),%edx
4.38:   63f906: 4d 84 81 80 32 00 00     lea    0x3200(%rcx),%rcx
0.10:   63f90d: 4d 8d b1 00 19 00 00     lea    0x190(%rcx),%rsi
0.00:   63f914: 4d 0f 46 f1     movb  %rcx,%rsi
1.36:   63f918: 3b 10          cmp    %rax,%edx
0.95:   63f91a: 4d 0f 46 c6     movb  %rsi,%rcx
1.86:   63f91e: 39 90 80 00 00 00     cmp    0x80(%rcx),%edx
4.89:   63f924: 4d 8d 88 00 11 00 00     lea    0x1100(%rcx),%rcx
0.00:   63f92b: 4d 8d b0 00 08 00 00     lea    0x800(%rcx),%rsi
0.00:   63f932: 4d 0f 46 f0     movb  %rcx,%rsi
2.84:   63f936: 3b 11          cmp    %rcx,%edx
0.33:   63f938: 4d 0f 46 ce     movb  %rsi,%rcx
1.58:   63f93c: 39 91 00 03 00 00     cmp    0x300(%rcx),%edx
7.66:   63f942: 4d 84 81 00 05 00 00     lea    0x500(%rcx),%rcx
0.05:   63f949: 4d 8d b1 00 03 00 00     lea    0x300(%rcx),%rsi
0.00:   63f950: 4d 0f 46 f1     movb  %rcx,%rsi
1.83:   63f954: 3b 10          cmp    %rax,%edx
1.20:   63f956: 4d 0f 46 c6     movb  %rsi,%rcx
1.23:   63f95a: 39 90 00 01 00 00     cmp    0x100(%rcx),%edx
11.38:  63f960: 4d 8d 88 00 02 00 00     lea    0x200(%rcx),%rcx
0.00:   63f967: 4d 8d b0 00 01 00 00     lea    0x100(%rcx),%rsi
0.03:   63f96e: 4d 0f 46 f0     movb  %rcx,%rsi
1.91:   63f972: 3b 11          cmp    %rcx,%edx
3.26:   63f974: 4d 0f 46 ce     movb  %rsi,%rcx
1.12:   63f978: 39 91 80 00 00 00     cmp    0x80(%rcx),%edx
13.82:  63f97e: 4d 8d b1 80 00 00 00     lea    0x80(%rcx),%rsi
0.00:   63f985: 4d 0f 46 f1     movb  %rcx,%rsi
7.21:   63f989: 3b 5c 40        cmp    %rsi,%edx
11.90:  63f98c: 4d 84 46 40     lea    0x40(%rsi),%rcx
0.00:   63f990: c9              leaveq %rcx,%rsi
0.94:   63f991: 4d 0f 46 c6     movb  %rsi,%rcx

```

# And some spooky action at a distance





# Sorted arrays work well on contemporary $\mu$ arch

Don't be (too) clever:

- ▶ Careful with branches
- ▶ Avoid cache aliasing/bad benchmarks
- ▶ Reuse bounds when repeating searches

Cleverness is useful, but getting simple right goes a long way.