Complicating Complexity

Algorithm Performance in the New Machine Age

JavaOne, October 2017

Maurice Naftalin

@mauricenaftalin

Maurice Naftalin

Java 5

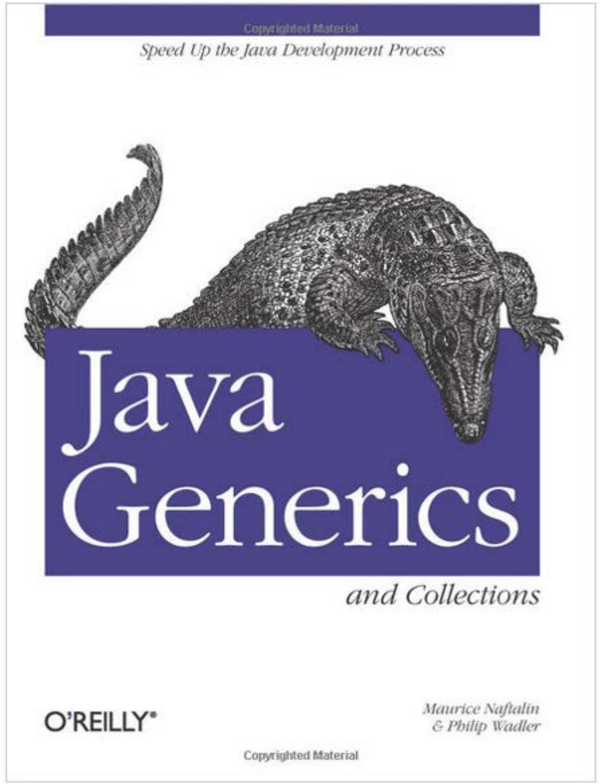

Java 8

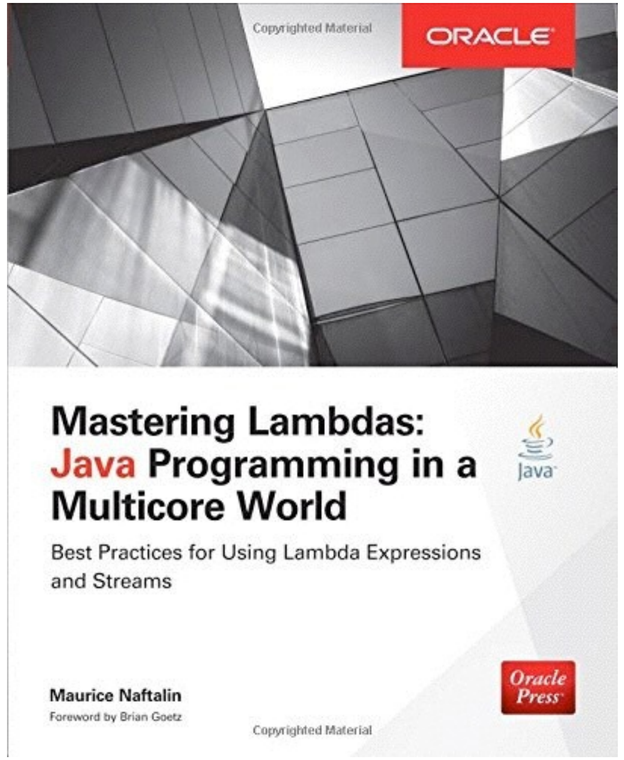

ORACLE
Developer
Champion

2013 2014 2015

Not Your Father's Complexity

- Original title for this talk
- Mental model of program performance
  - Number of instructions executed
- Why we need to change this, and how

Who *Is* Your Father?

…professionally speaking!

Don Knuth (1938 –), Creator of TeX,
Turing Award laureate,…

Author of *The Art of Computer Programming*

What He Really Thinks

"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"

What is Computational Complexity?

- an algorithm with time complexity $O(g(n))$ will complete in less than $c * g(n)$ steps, for some c and sufficiently large n

- e.g. Knuth, on multiple list insertion sort: execution time is

$$3.5N^2 + 24.5N + 4M + 2$$

  - but we only care about the N^2 term!

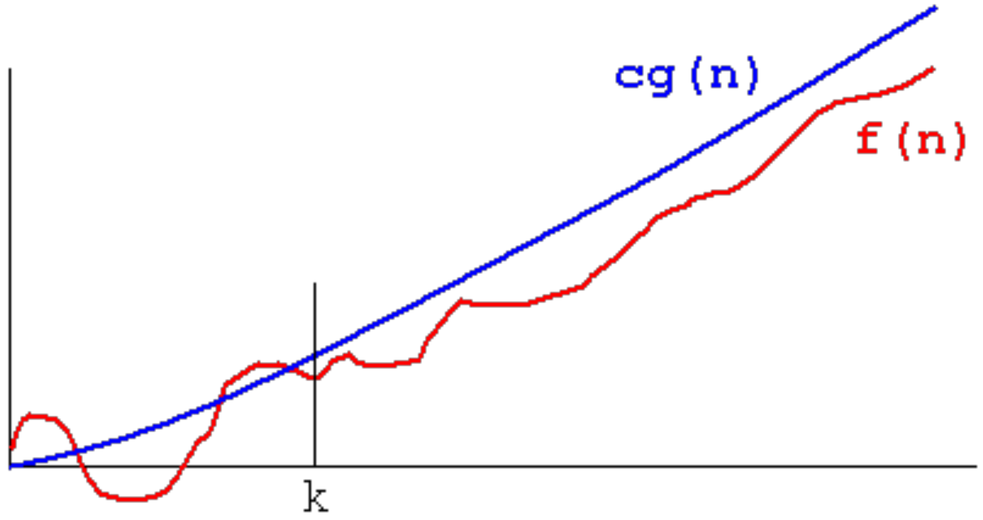

Why Bother?

It's the traditional way of evaluating algorithms!

Complexity	Effect of doubling N
$O(1)$	Unchanged
$O(\log N)$	Increased by a constant
$O(N)$	Doubled
$O(N \log N)$	Doubled + an amount proportional to N
$O(N^2)$	Increased fourfold
…	…
…	…
…	…

Who *Is* Your Father?

Another Candidate:

Joshua Bloch (1961 –), Author of *Effective Java*

Author of the Java Collections Framework

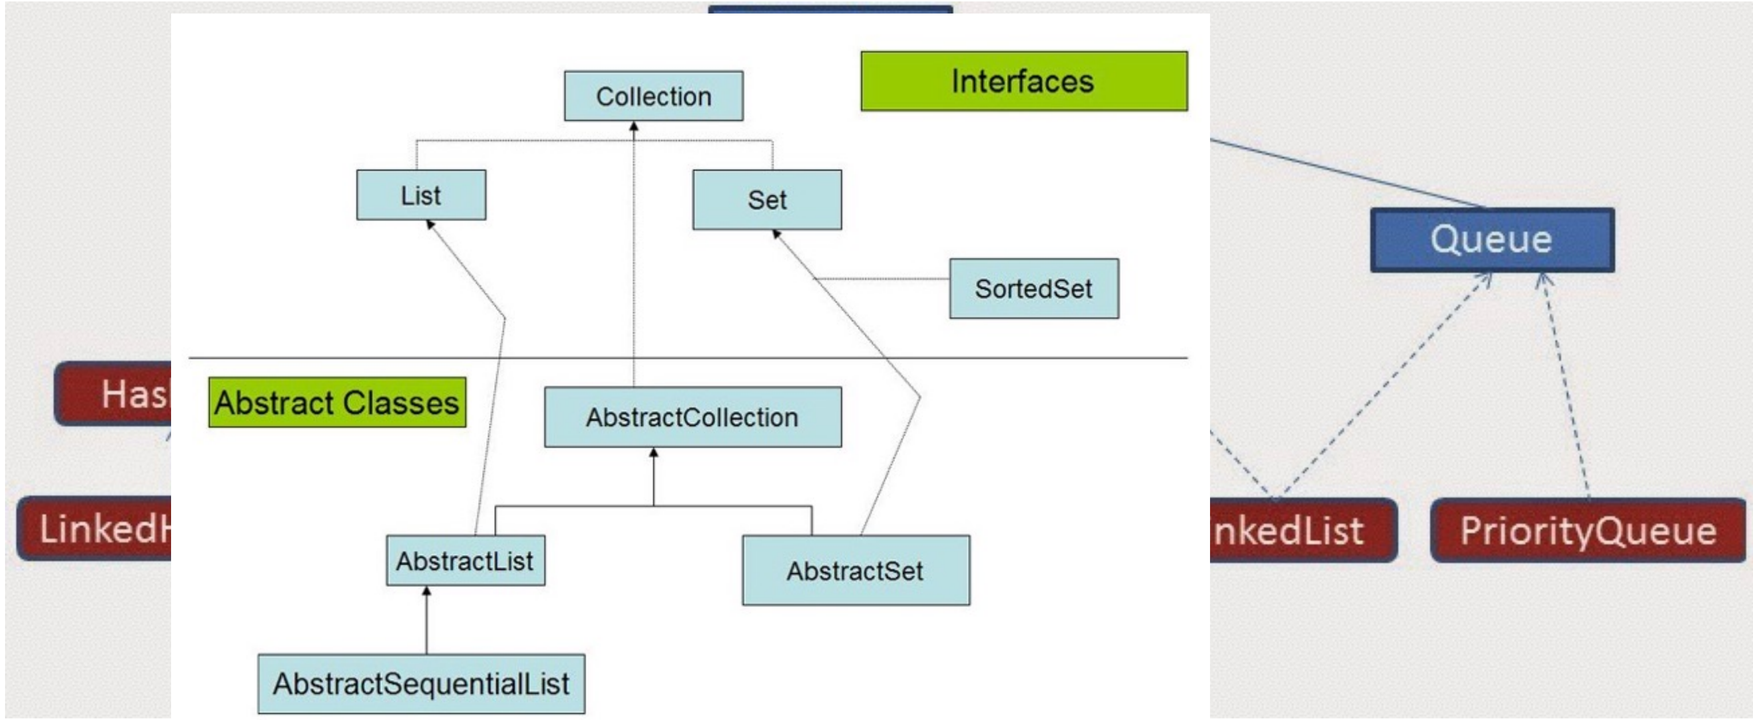

What Josh Told Me…

- the Java Collections Framework is *interface-based*
- choose the interface (`Set`, `List`, `Queue`) that meets your requirement
- then choose the implementation with best performance for your usage scenario
- for example: which `List` implementation?
  - `ArrayList`, `LinkedList`, `CopyOnWriteArrayList`?

Which List Implementation?

	get()	add()	remove(0)
ArrayList	*O*(1)	*O*(1)	*O*(N)
LinkedList	*O*(N)	*O*(1)	*O*(1)
COWArrayList	*O*(1)	*O*(N)	*O*(N)

Did this ever make sense?

— Yes, on these assumptions:

- can ignore constant factors
- all instructions have the same duration
- memory doesn't matter
- instruction execution dominates performance

Was Complexity Study Ever Worth It?

Of course it was!

But instruction execution is only one bottleneck. Many others:

- Disk/Network
- Garbage Collection
- Resource Contention

and more…

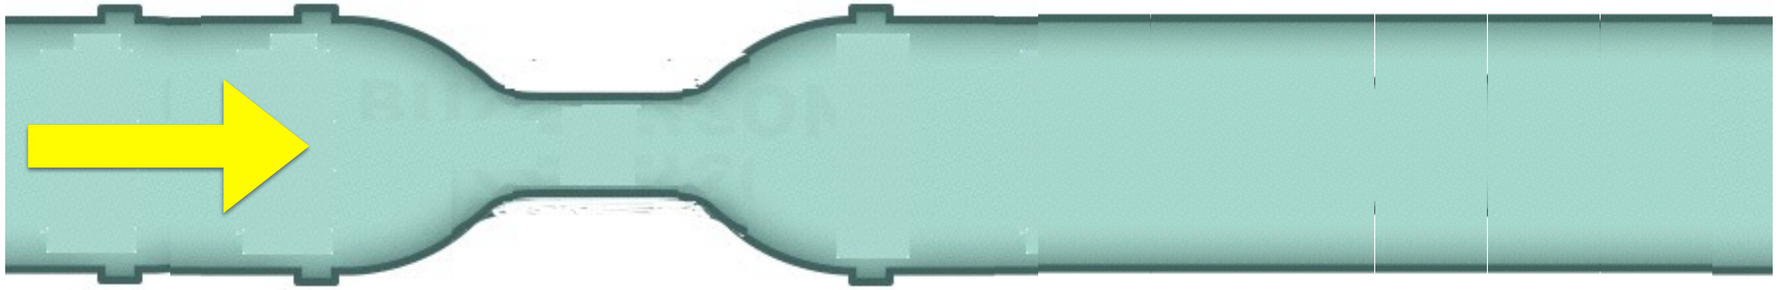

Was Complexity Study Ever Worth It?

Of course it was!

But instruction execution is only one bottleneck. Many others:

- Disk/Network
- Garbage Collection
- Resource Contention

and more…

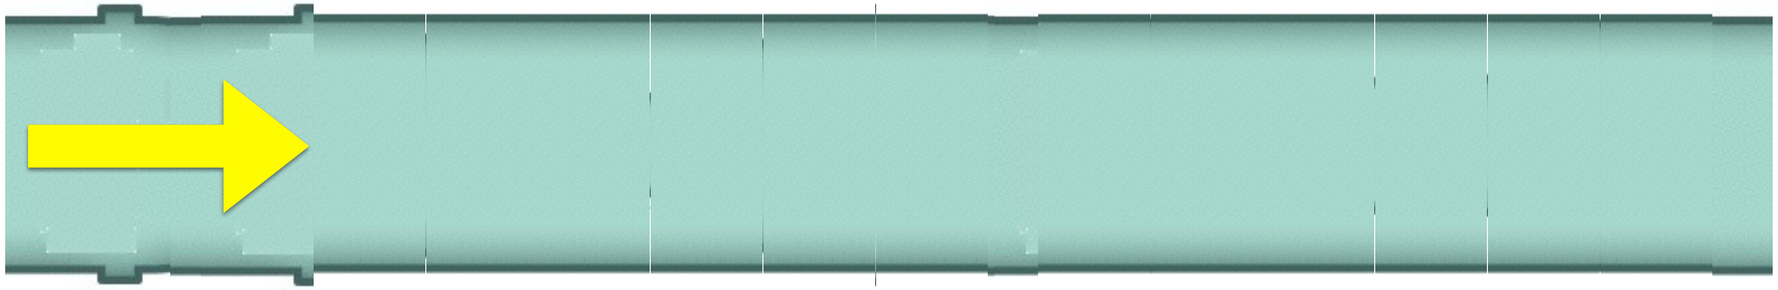

Was Complexity Study Ever Worth It?

Of course it was!

But instruction execution is only one bottleneck. Many others:

- Disk/Network
- Garbage Collection
- Resource Contention

and more…

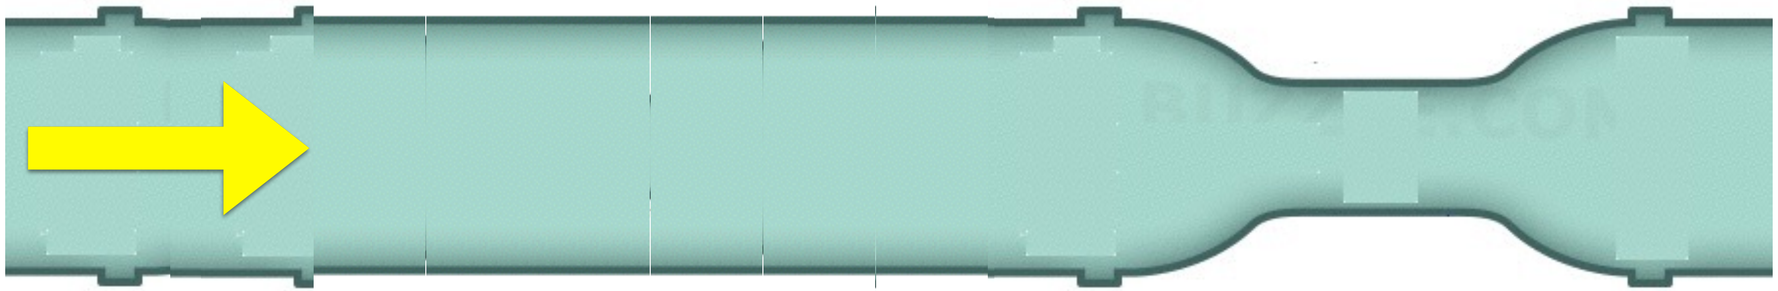

pre-1980

pre-1980

The Golden Age of Complexity

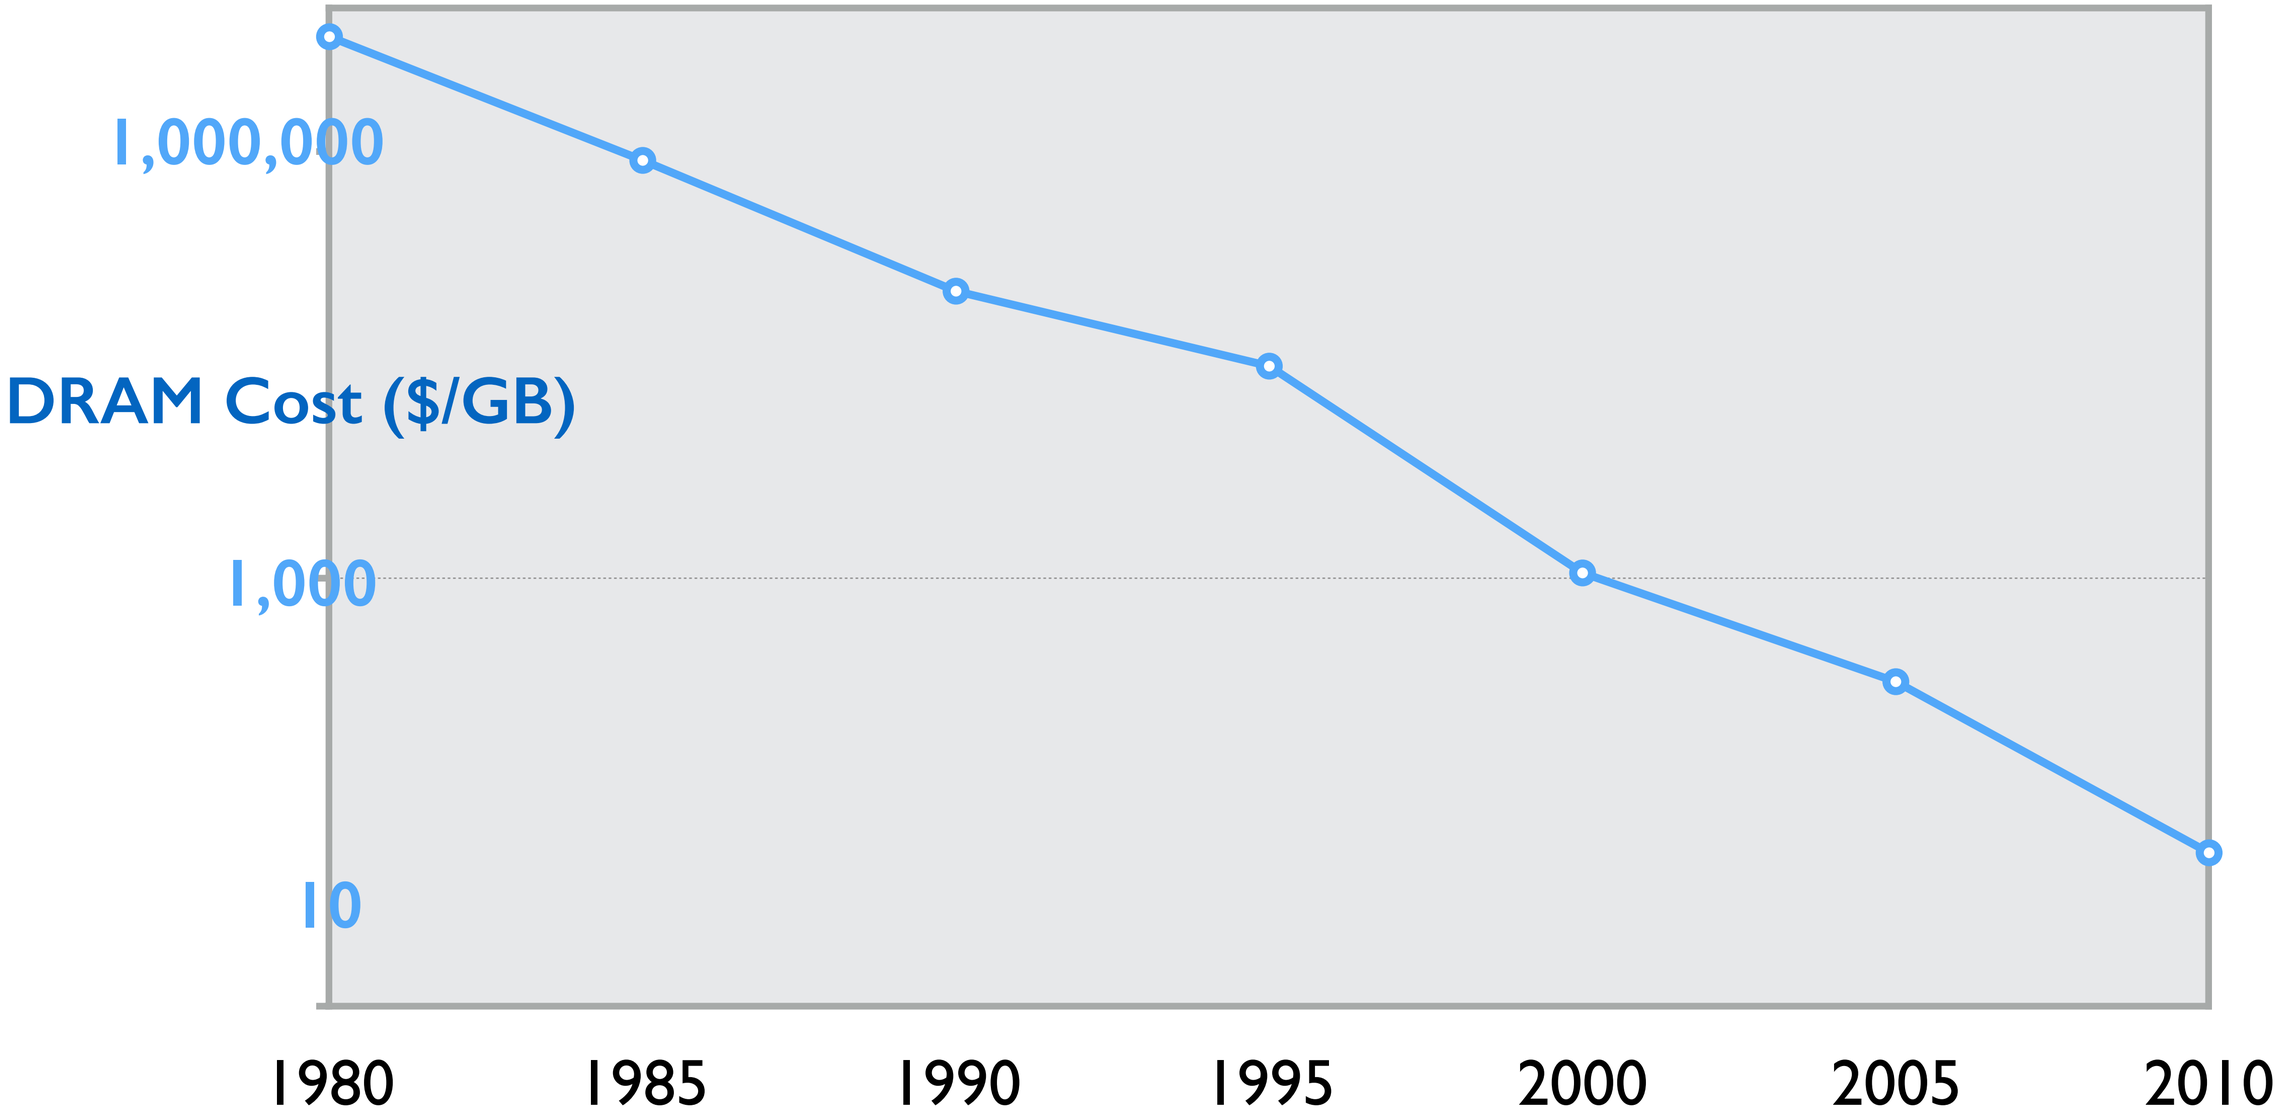

The Golden Age of Complexity

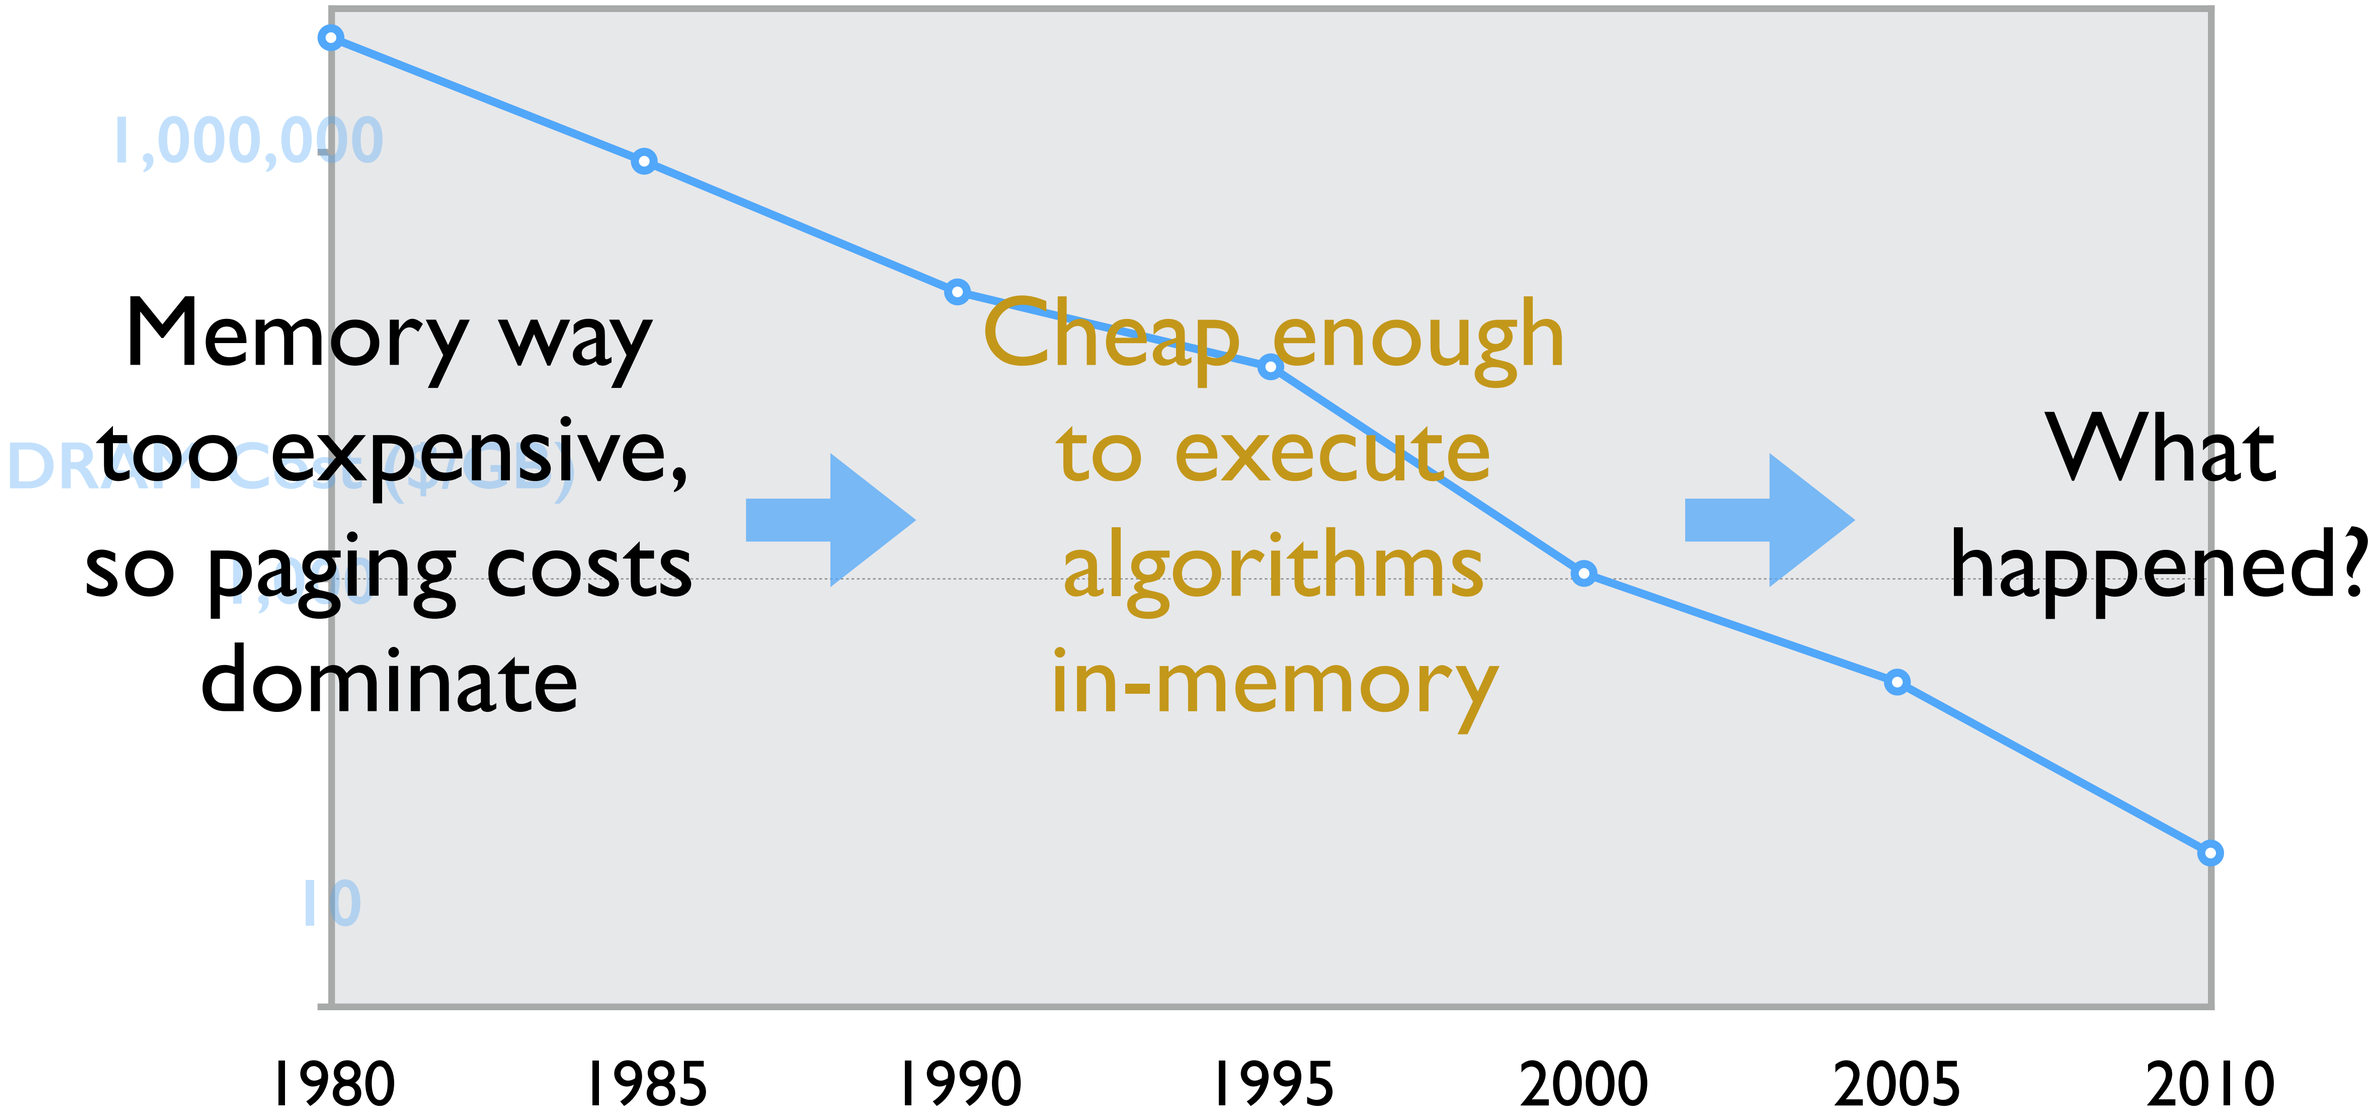

The Golden Age of Complexity

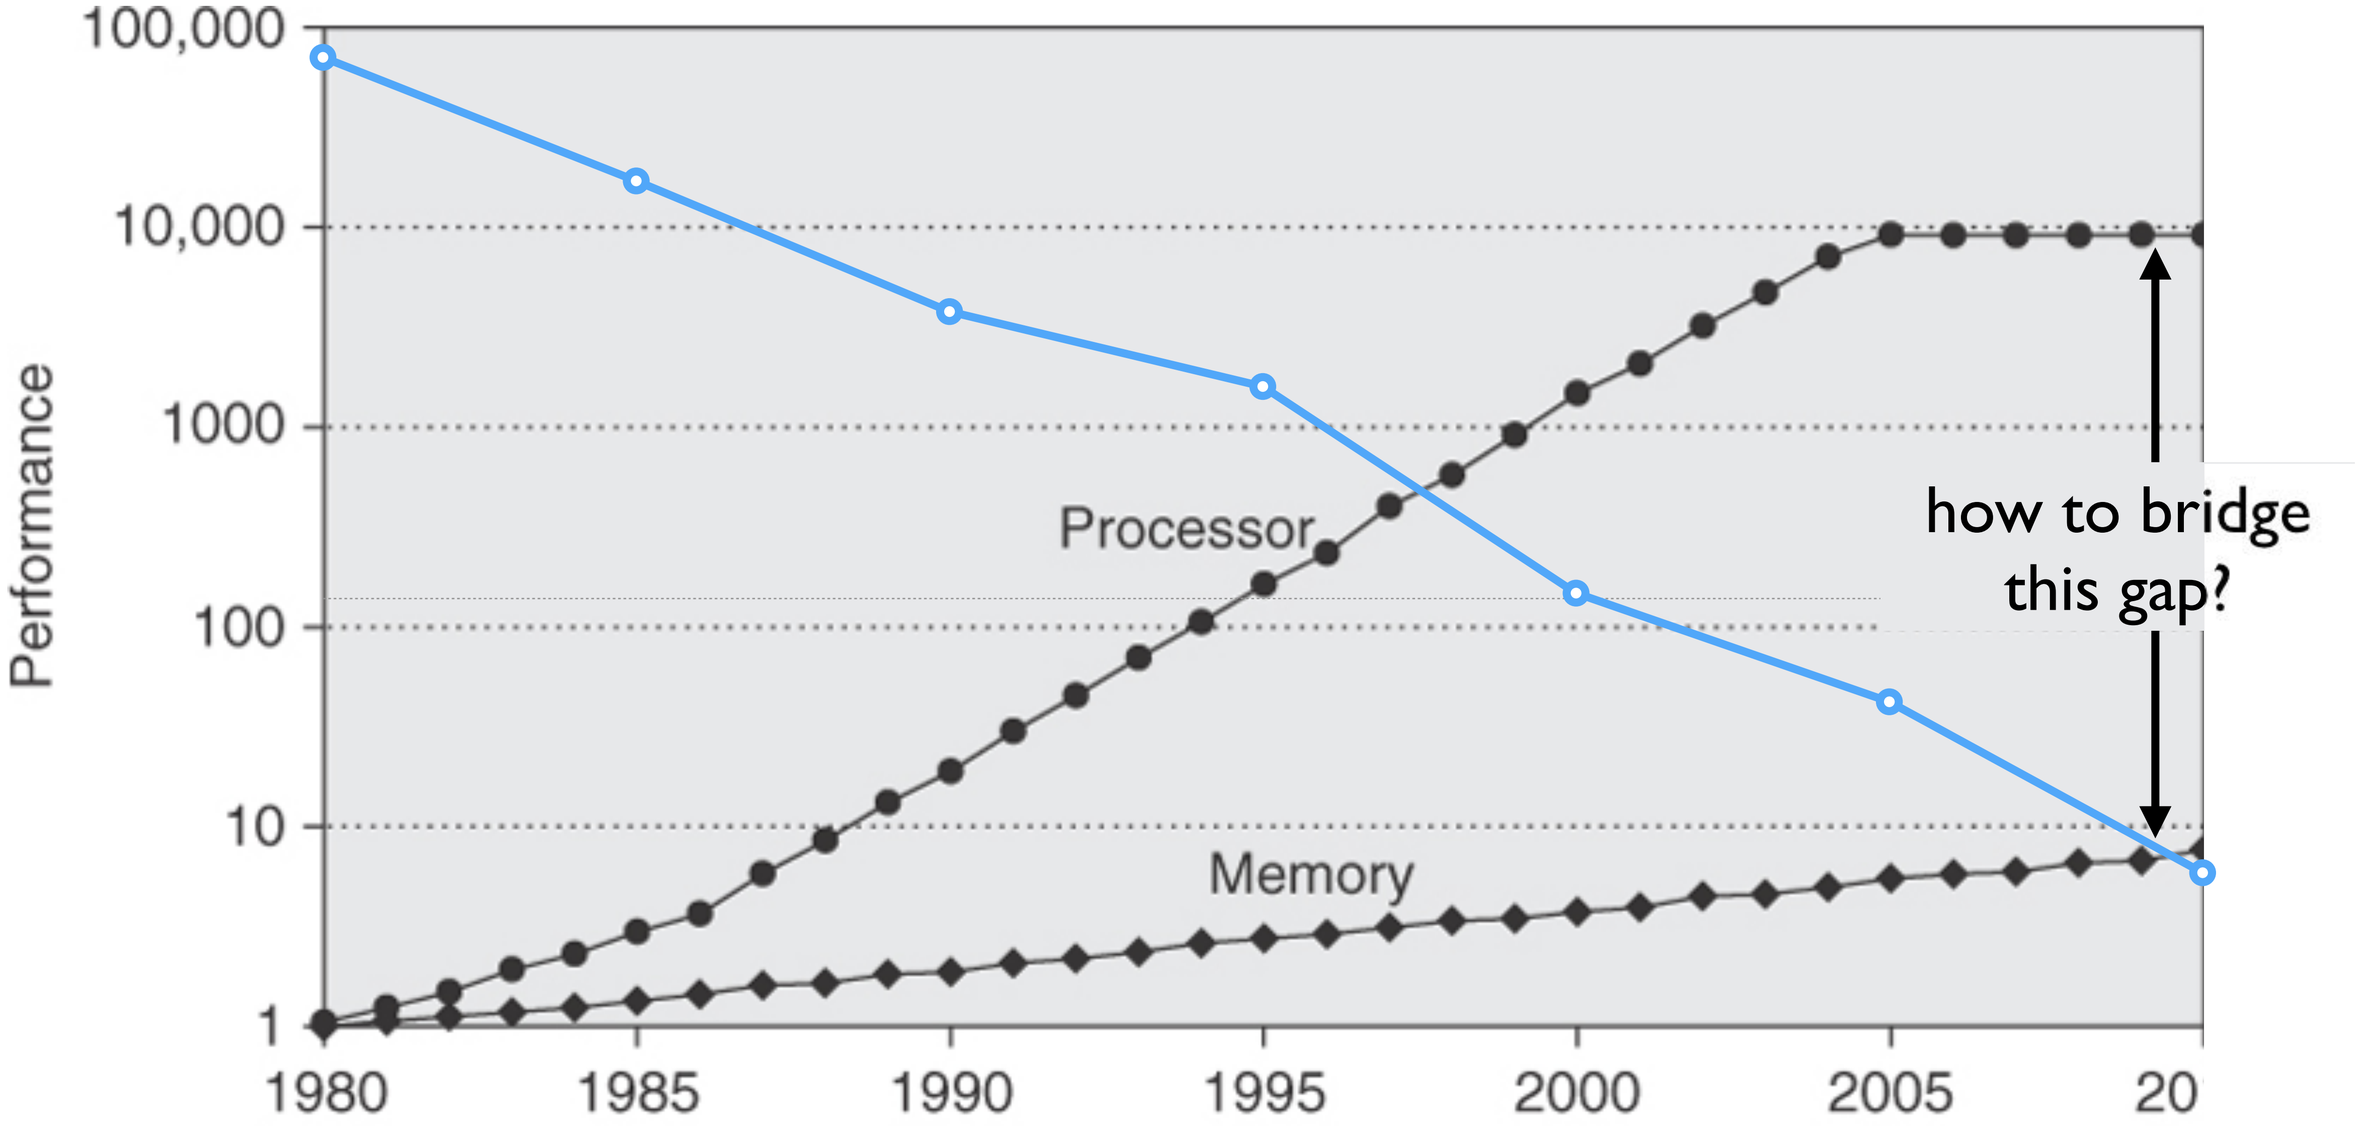

Keeping the Cores Running Today

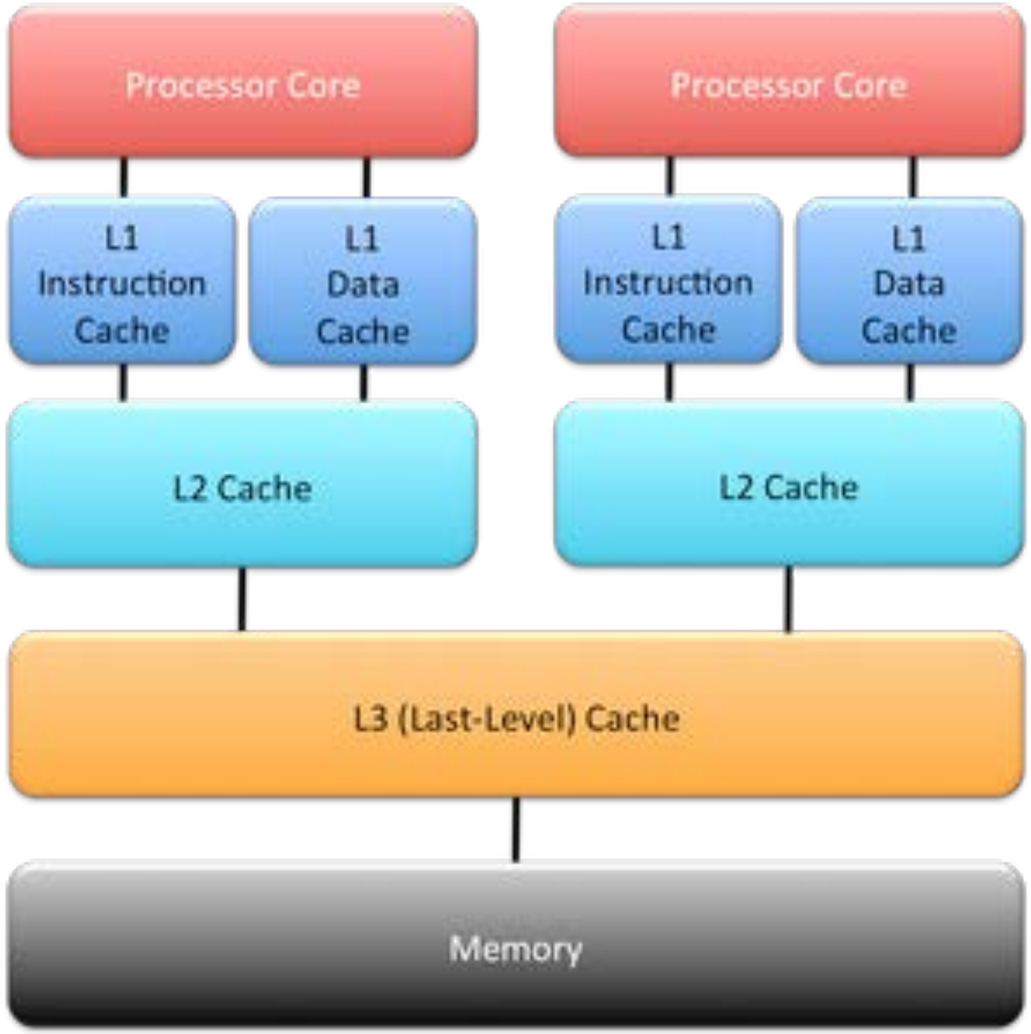

The Memory Hierarchy

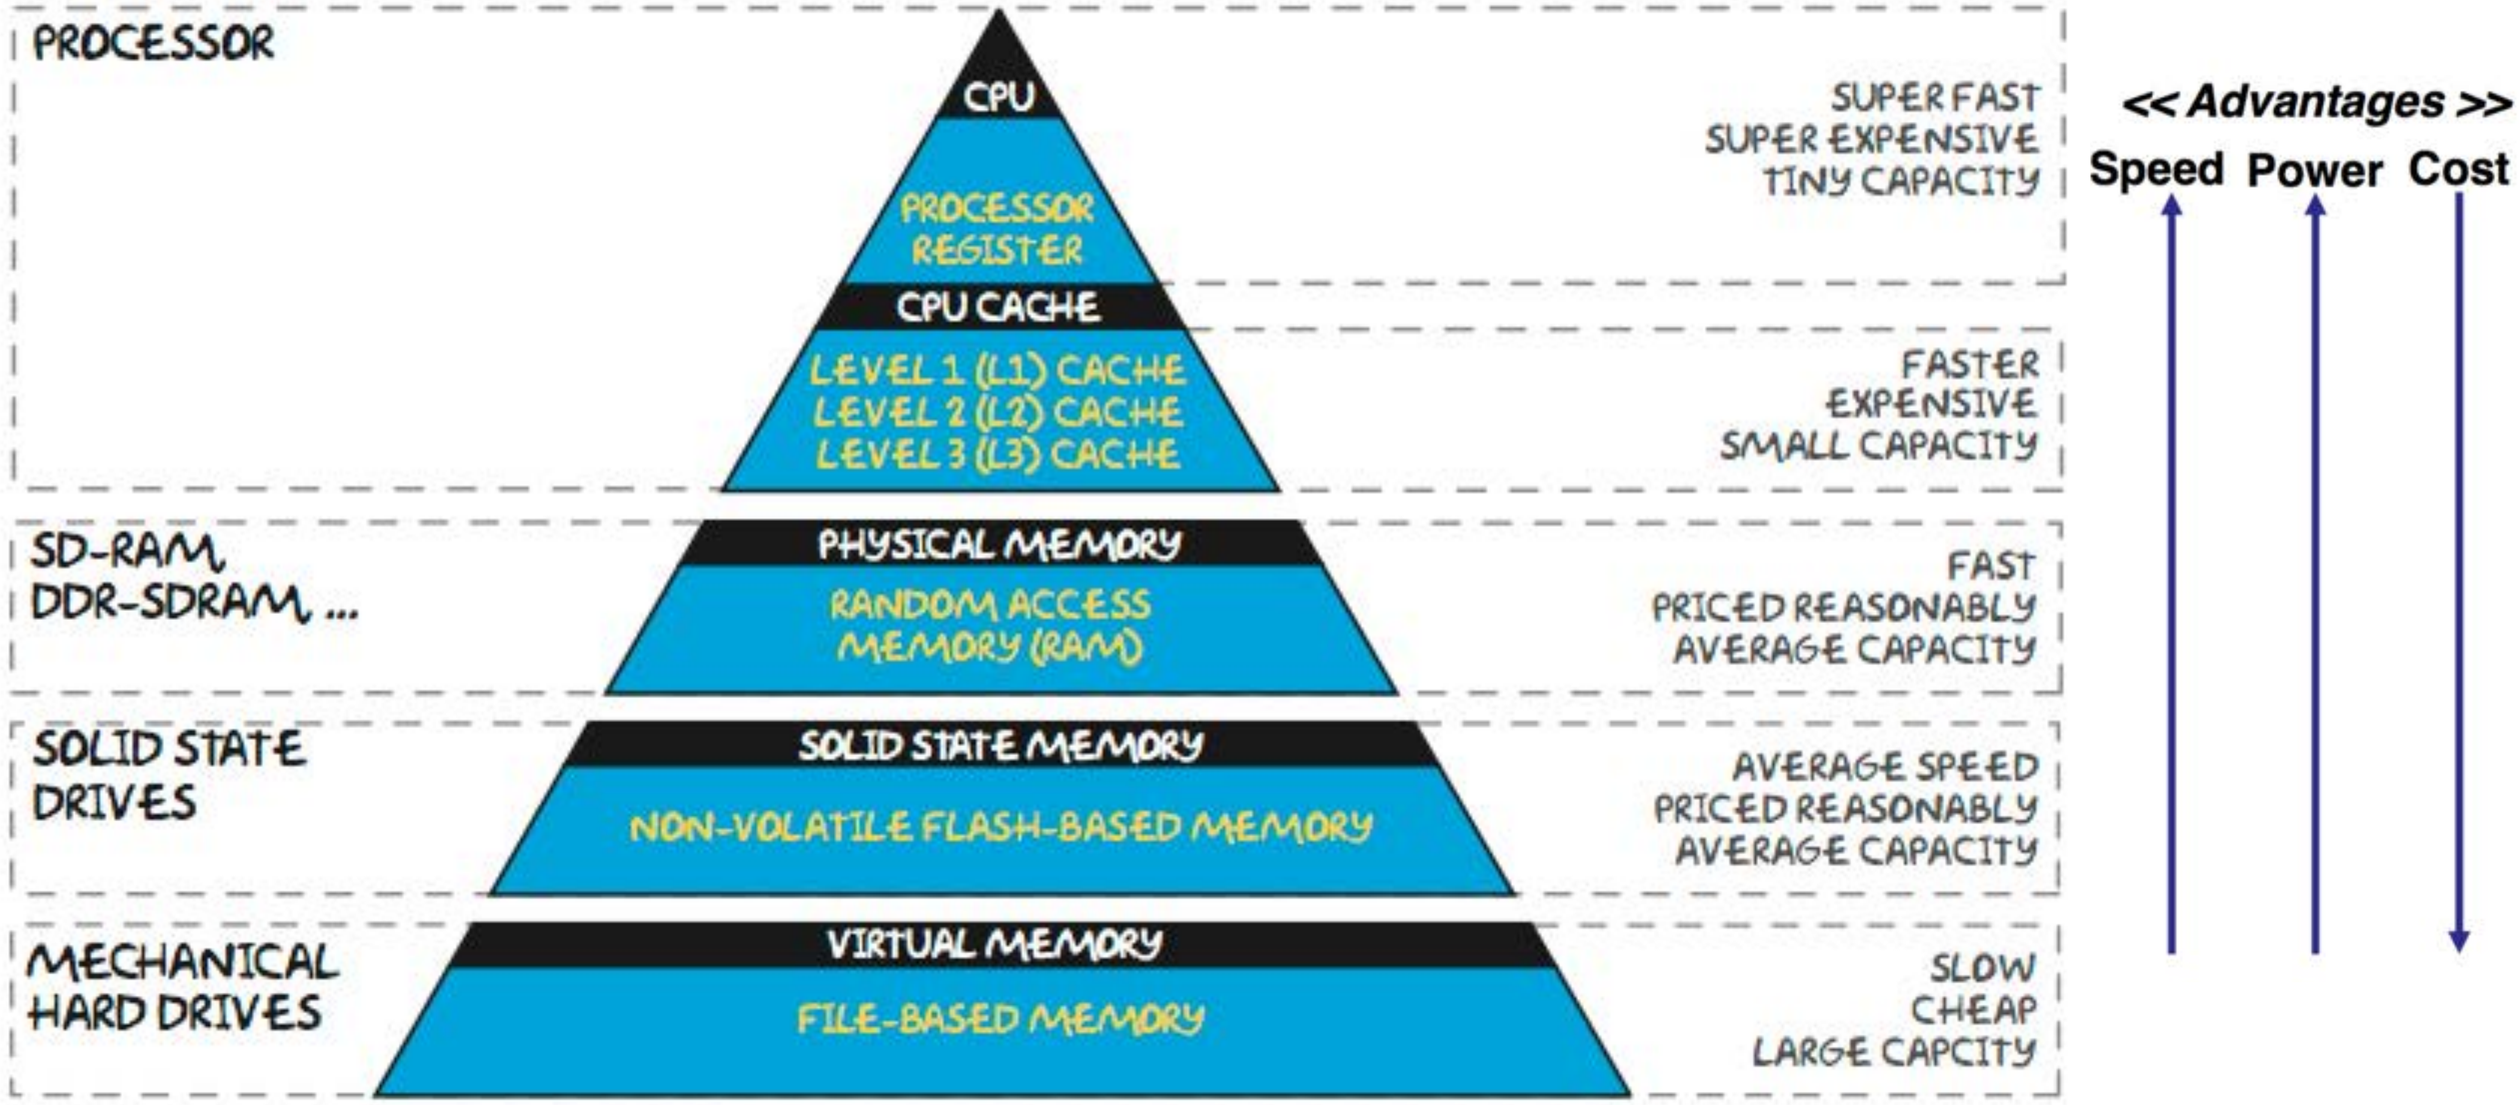

The Memory Hierarchy

Cache Effects Often Dominate

- Main memory retrieval costs ~100x L1 access
    - 2-300x register access
- Typical programs have 95% hit rate
    - it's the other 5% that hurts
- Why cache misses? Two possible reasons (many others):
    - insufficient capacity
    - failure of prefetching
        - unpredictable data access patterns!

Stride Prefetching

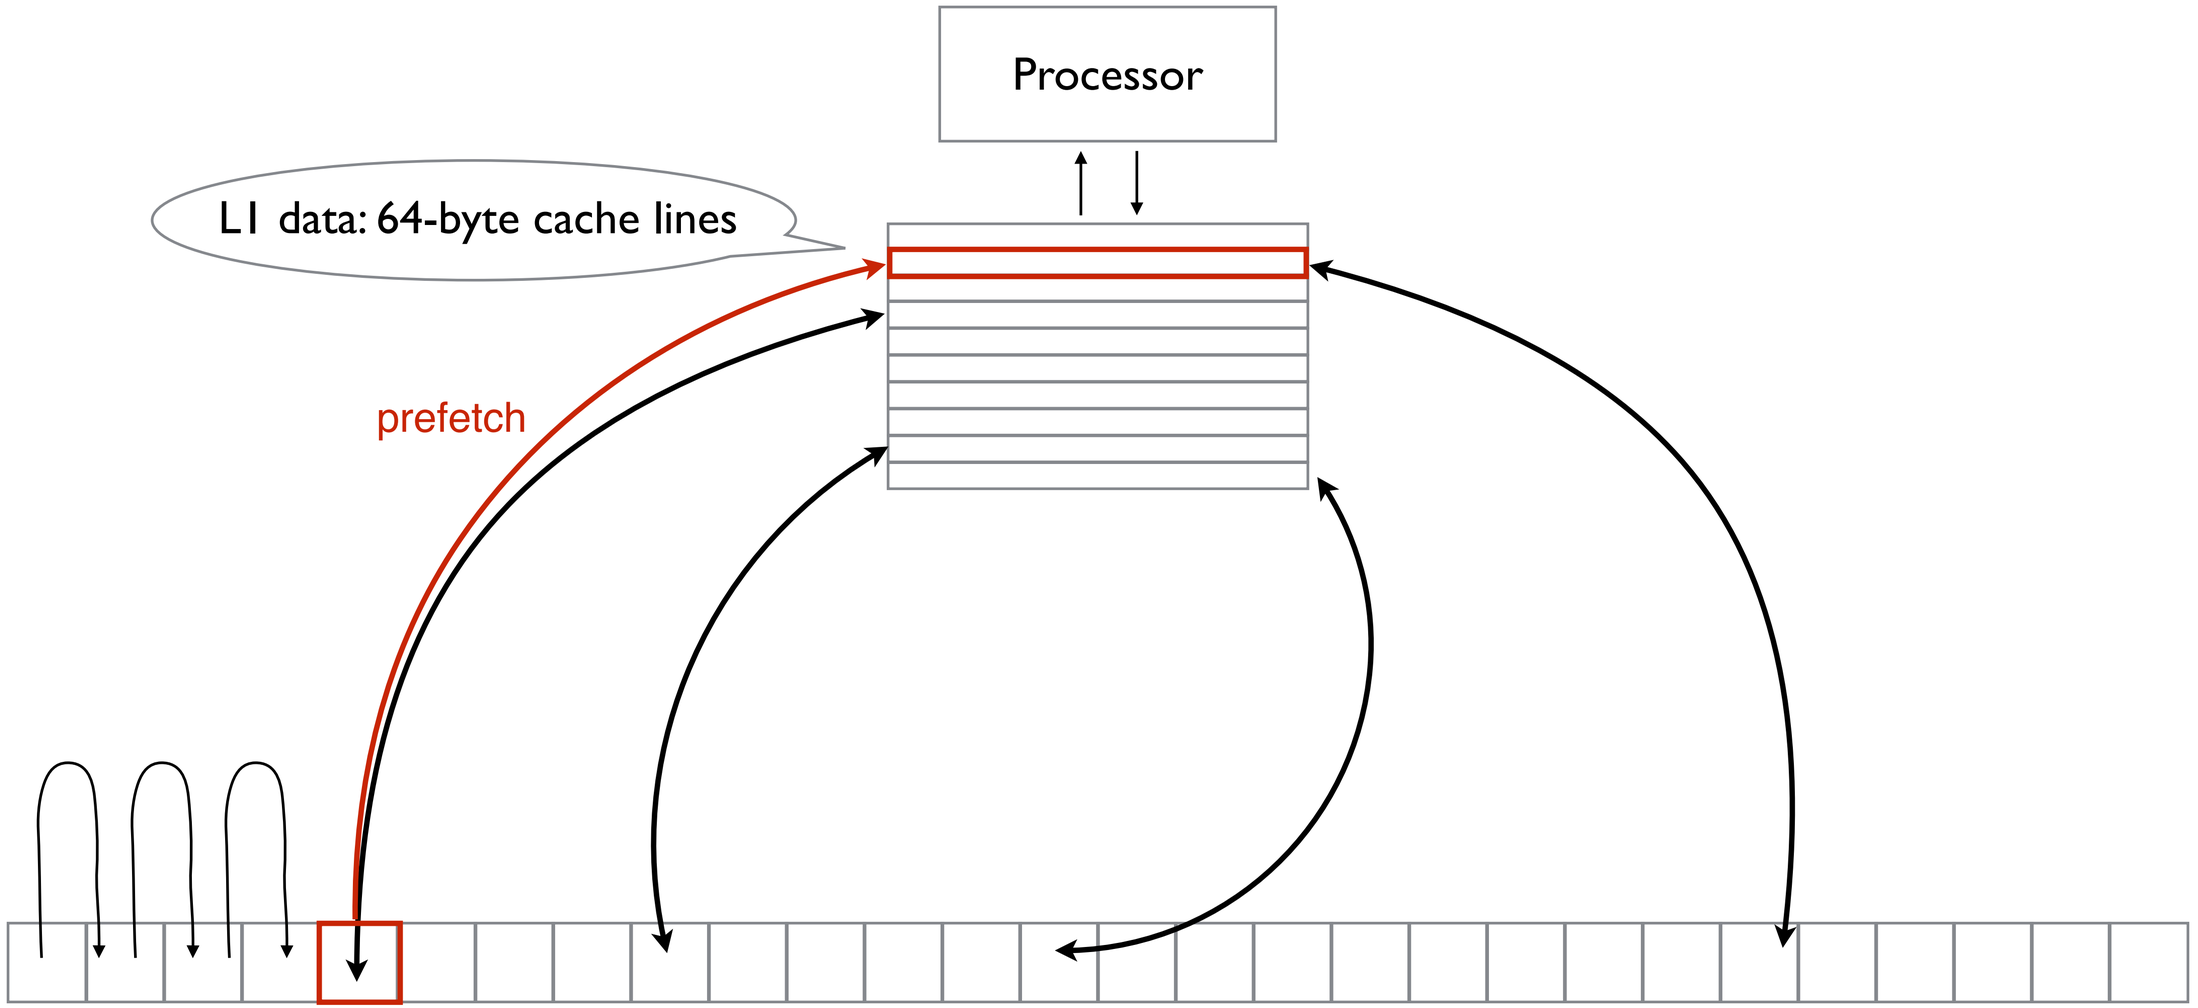

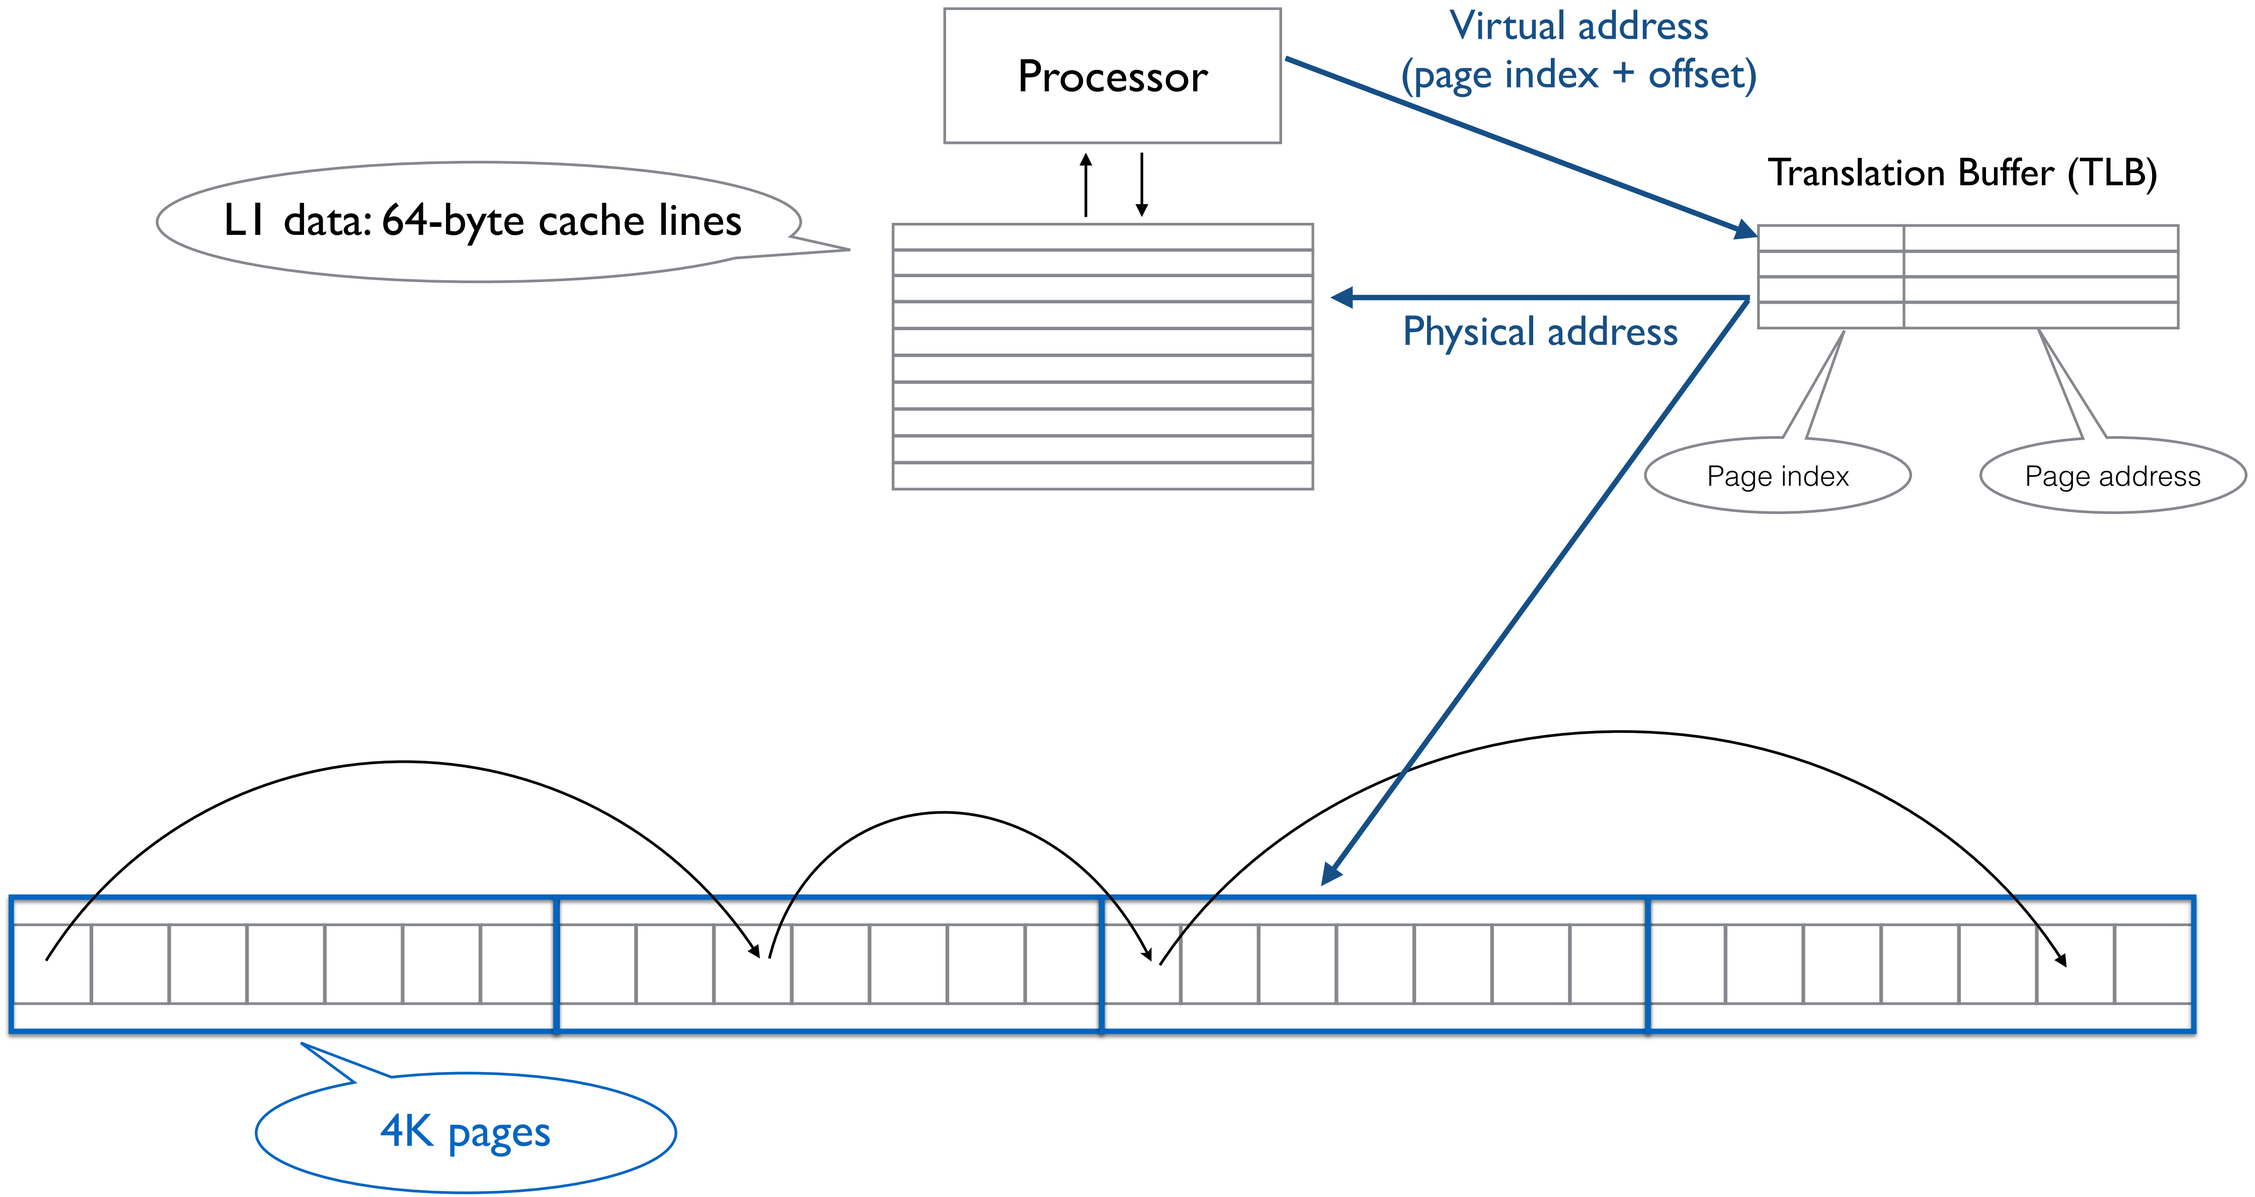
Processor
Virtual address
(page index + offset)
Translation Buffer (TLB)
L1 data: 64-byte cache lines
Physical address
Page index
Page address
4K pages

How Does Caching Play With Complexity?

Sample case: traversing a list

- *O*(n), obviously

First issue: data size. Let's compare:

- `LinkedList`
- primitive array

LinkedList

list length (K)	1	7	63	511
performance (ns/op)	7.25	9.03	20.87	29.07
CPI (clockticks/instrn)	0.32	0.41	0.93	1.33
	events/operation			
cycles	17.97	22.77	51.32	72.66
instructions	56.08	55.88	54.96	54.49
L1-dcache-load-misses	1.18	1.83	1.87	2.65
L1-dcache-loads	18.94	19.39	18.88	18.22
L1-dcache-stores	12.00	12.18	11.99	11.15
LLC-load-misses	0	0	0.41	1.31
LLC-loads	0	0.72	1.33	1.56
dTLB-load-misses	0	0	0	0.90
dTLB-loads	19.05	19.00	19.15	18.09
dTLB-stores	12.04	12.09	12.16	11.02

What's Going On?

`LinkedList`: node size is 24 bytes

Running on Intel Core i5:

L1 data	32K
L2	256K
L3	3M

Each new list item is 40 bytes (24 + 16)

– L1 cache will be full at <1K items

`ArrayList` is better, but not much: each new item is 20 bytes

Primitive Array

list length (K)	1	7	63	511
performance (ns/op)	3.62	3.65	3.65	3.66
CPI	0.30	0.30	0.30	0.31
	events/operation			
cycles	9.09	9.16	9.10	9.13
instructions	30.24	30.13	29.94	29.85
L1-dcache-load-misses	0.00	0.01	0.06	0.06
L1-dcache-loads	12.00	12.00	11.97	12.14
L1-dcache-stores	6.00	6.02	6.02	6.04
LLC-load-misses	0.00	0.00	0.00	0.00
LLC-loads	0.00	0.00	0.00	0.00
dTLB-load-misses	0.00	0.00	0.00	0.00
dTLB-loads	12.17	12.00	11.90	12.06
dTLB-stores	6.03	5.99	5.98	6.05

How Does Caching Play With Complexity?

Second issue: data locality

Two different problems:

- Data density: how much of your 64-byte cache lines is data that you actually need?
- Prefetching: are you giving the processor a chance to help?

Populating `LinkedList` “Naturally”

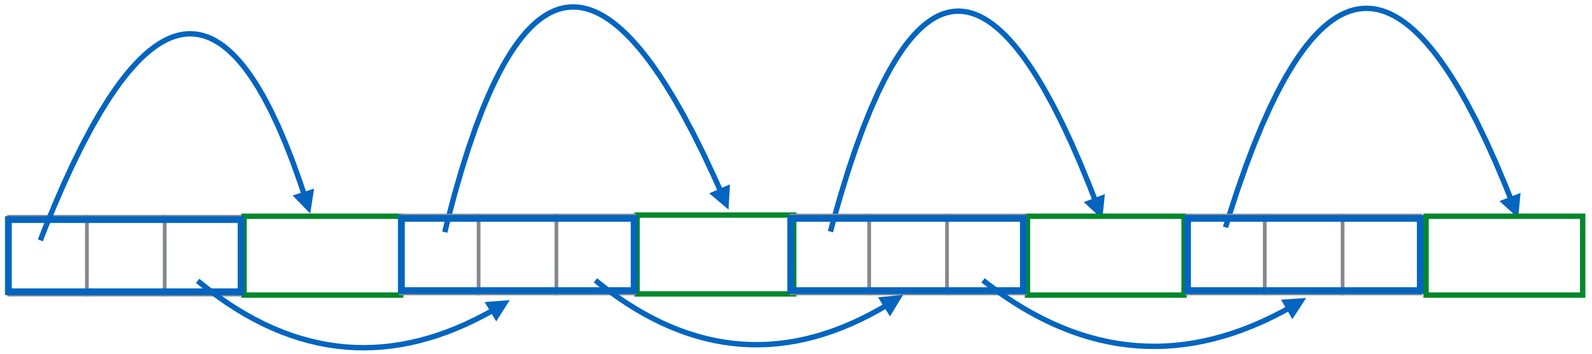

```
for (int i = 0; i < LIST_LENGTH; i++) {
     linkedList.add(random.nextInt());
}
```

Populating `LinkedList` Randomly

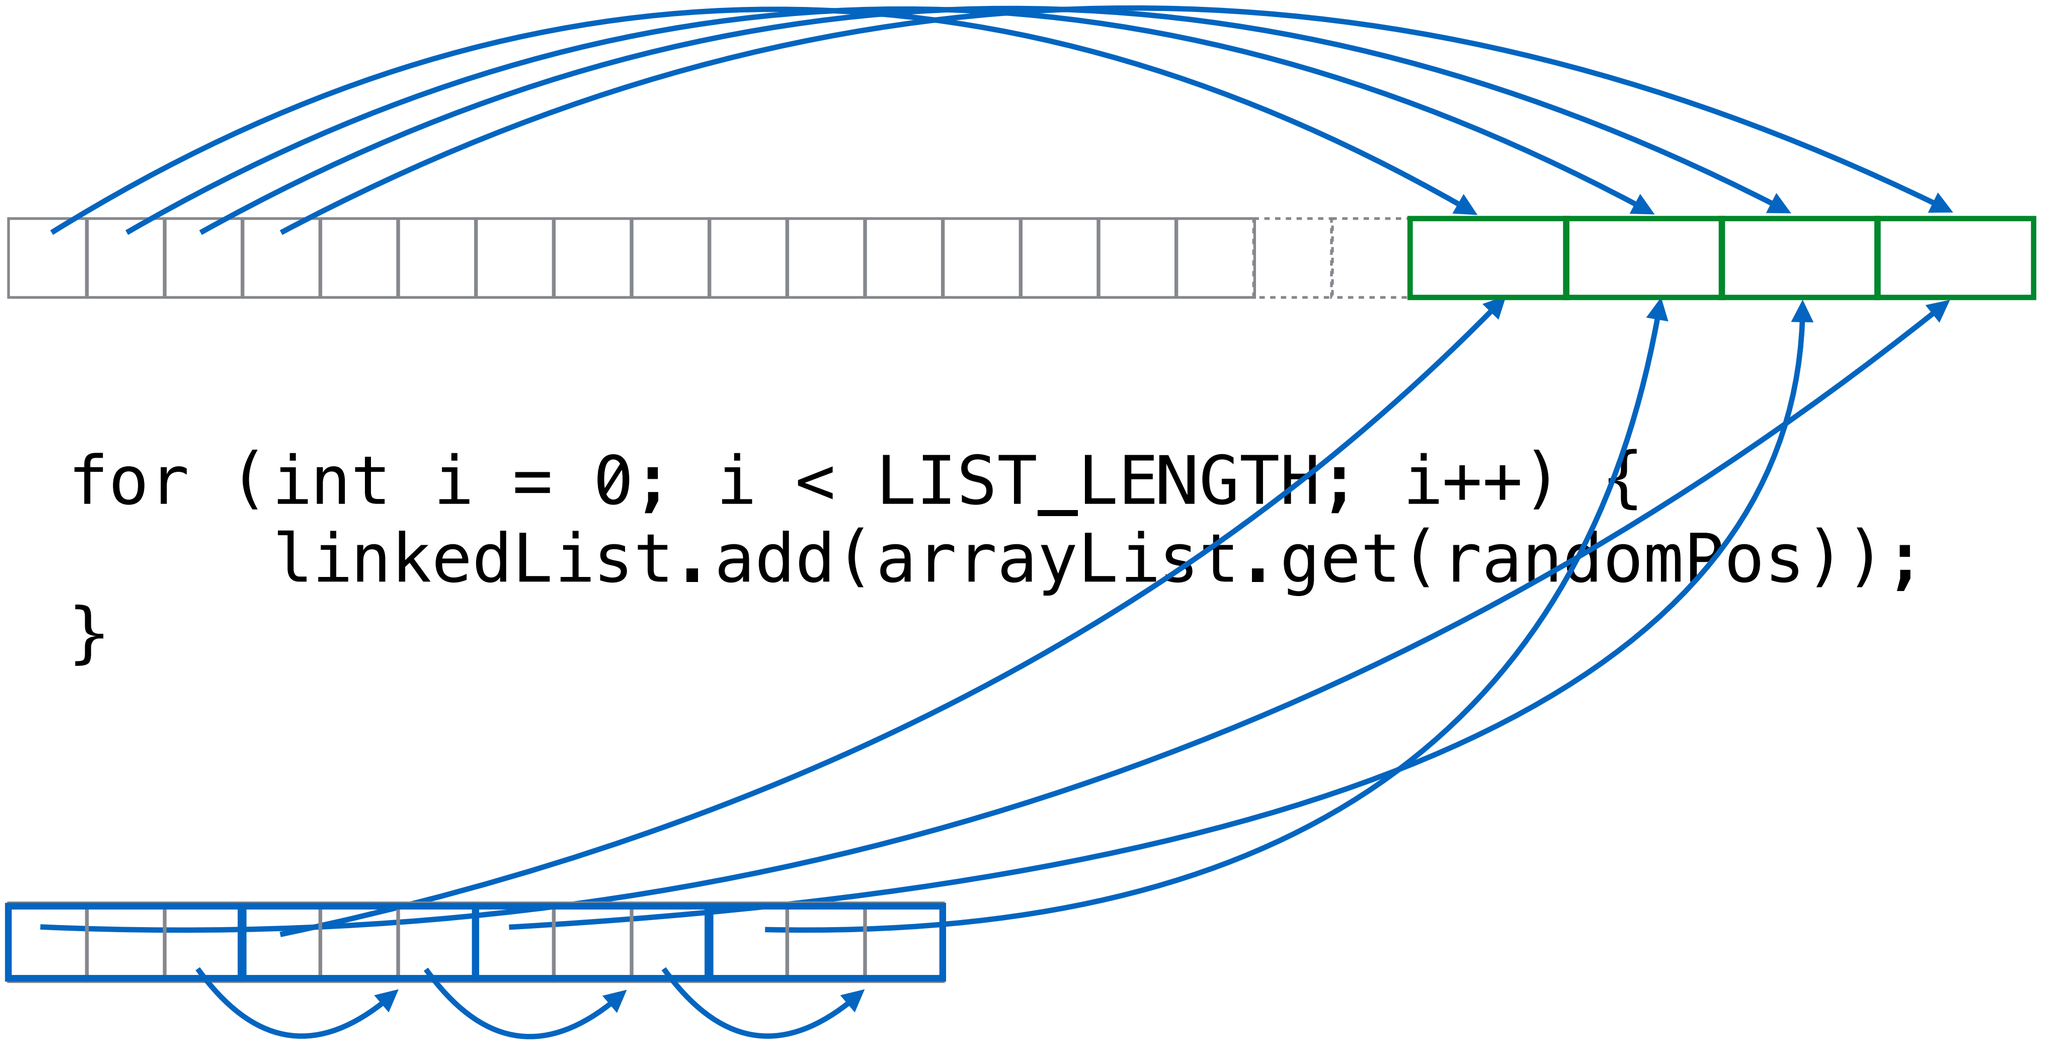

```
for (int i = 0; i < LIST_LENGTH; i++) {
      linkedList.add(arrayList.get(randomPos));
}
```

“Demo”

Poor Unloved `LinkedList`…

HashMap

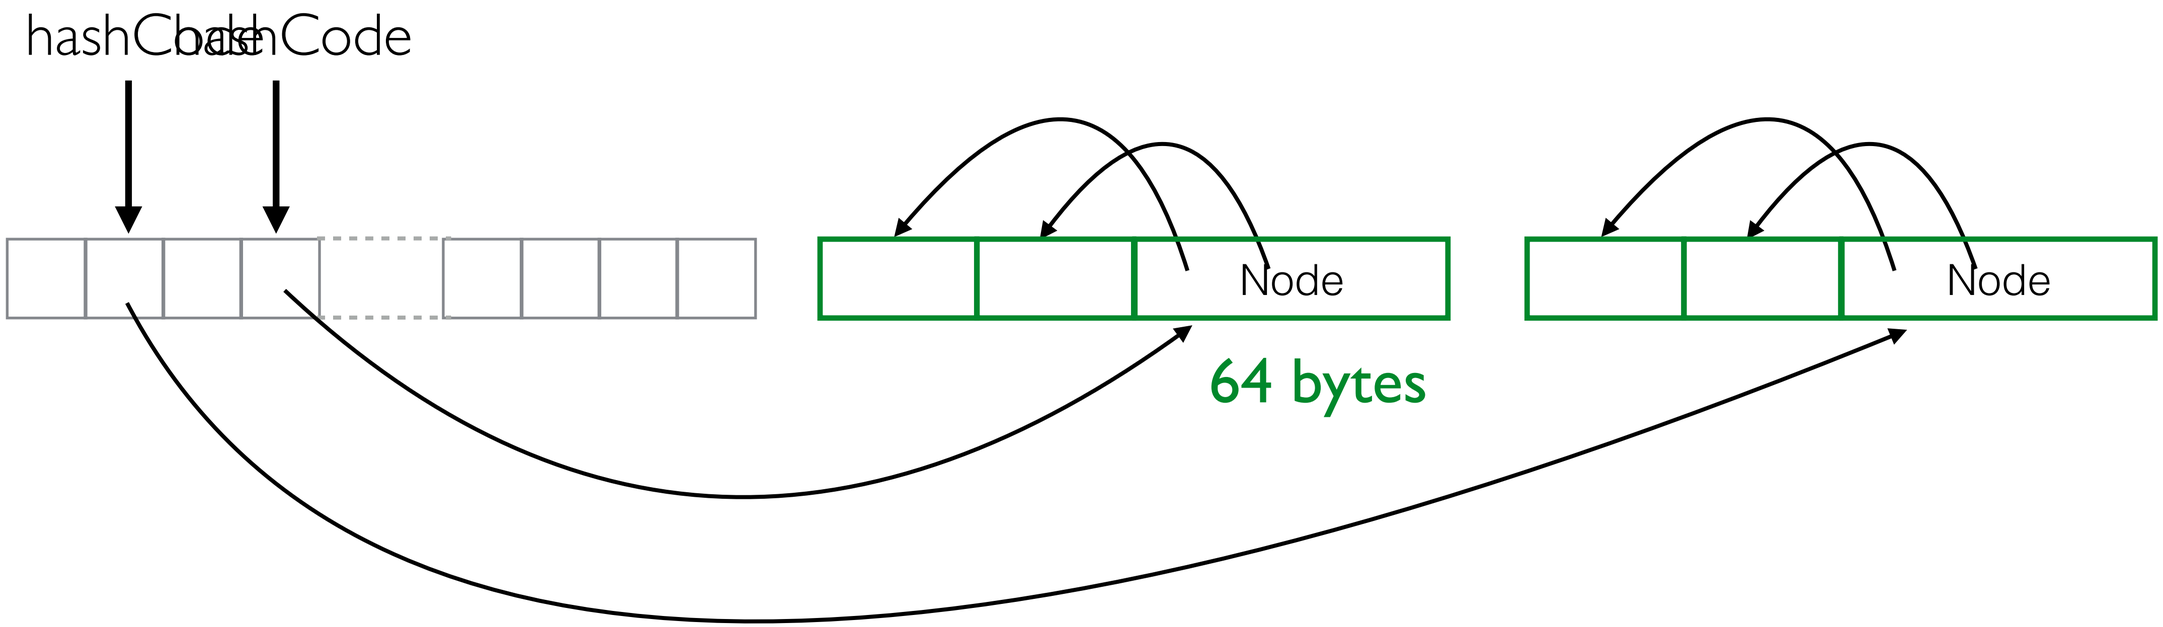

HashMap collision

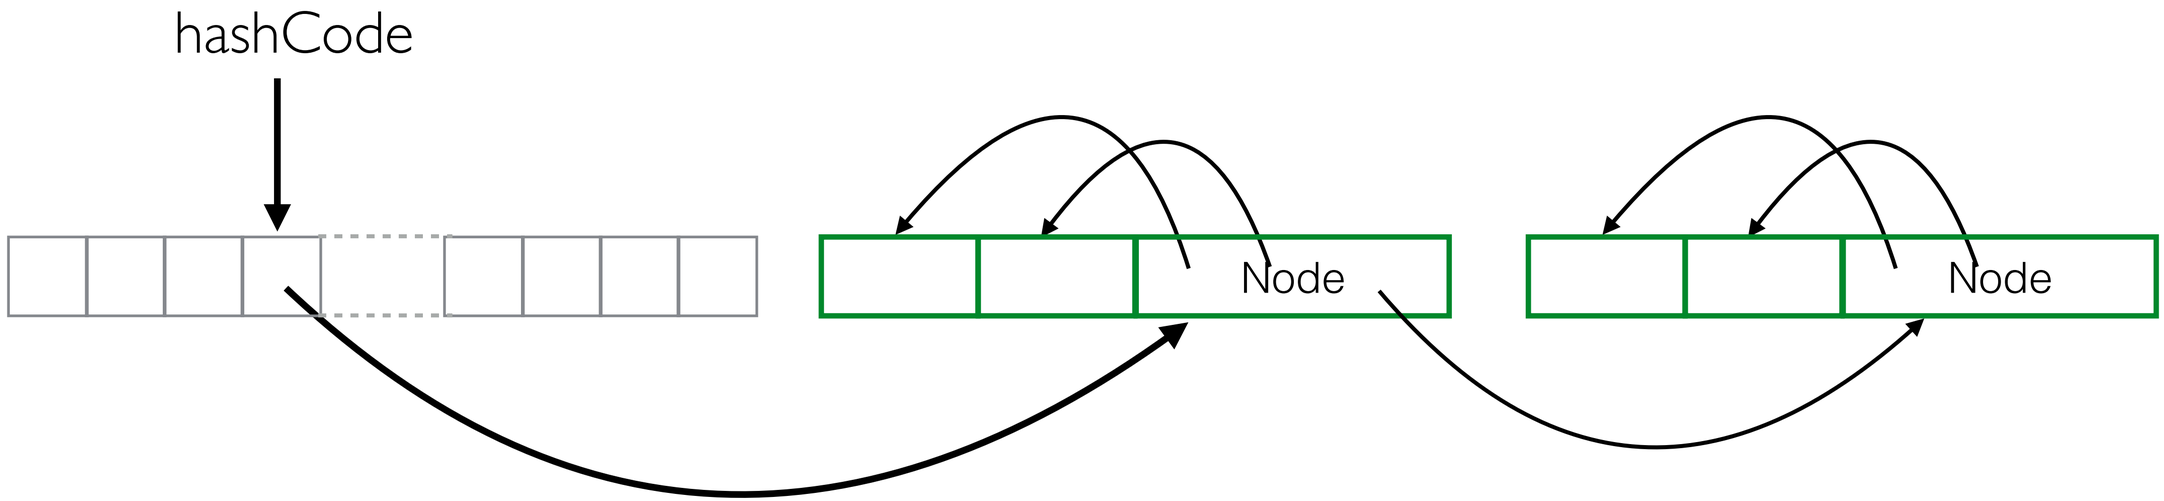

ImmutableCollections.MapN

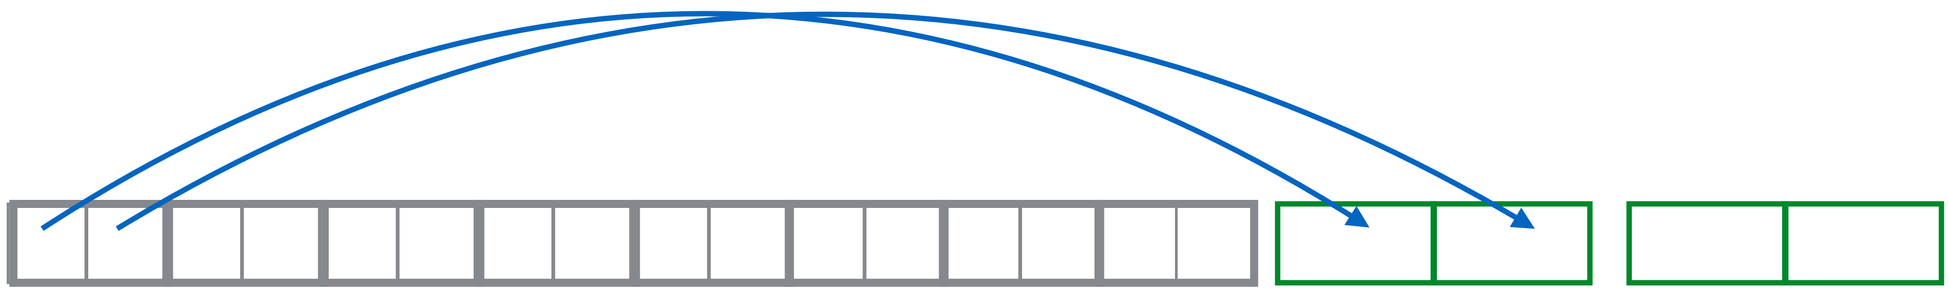

```
Map<Integer,Integer> immutableMap = Map.ofEntries(
   Map.entry(512,1024),
   Map.entry(513,1026),
   Map.entry(514,1028));
```

Reducing Memory Footprint

Third-party collections frameworks usually have a focus on low memory footprint:

- Eclipse Collections
- fastutil
- Vavr (formerly Javaslang)
- Apache Commons Collections
- Guava
- Trove
- Argon

Improving Data Locality

- 3rd-party frameworks (previous slide) often support primitive collections

- ObjectLayout
  - `StructuredArray` – like a C-style "array of struct"
  - also arrays as part of objects, and cohered aggregates

- Roaring Bitmaps
  - compressed bitmaps, very fast

- Project Valhalla
  - language-level solution for value objects and primitive collections

Conclusion, of sorts…

Performance **mostly doesn't matter**

… but when it does matter, it really matters!

Every performance improvement represents a tradeoff

Algorithm complexity is still important

… but so is

- network/database performance,
- GC,
- resource contention,
- caching

So actually it **is** your father's complexity – just a lot more complex than before!