

# Are Your Incoming Aliases Really Necessary? Counting the Cost of Object Ownership

Alex Potanin, Monique Damitio, James Noble  
Victoria University of Wellington, New Zealand

# Are Your Incoming Aliases Really Necessary? Remembering the Cost of Object Ownership

Alex Potanin, Monique Damitio, James Noble  
Victoria University of Wellington, New Zealand

# Aliasing

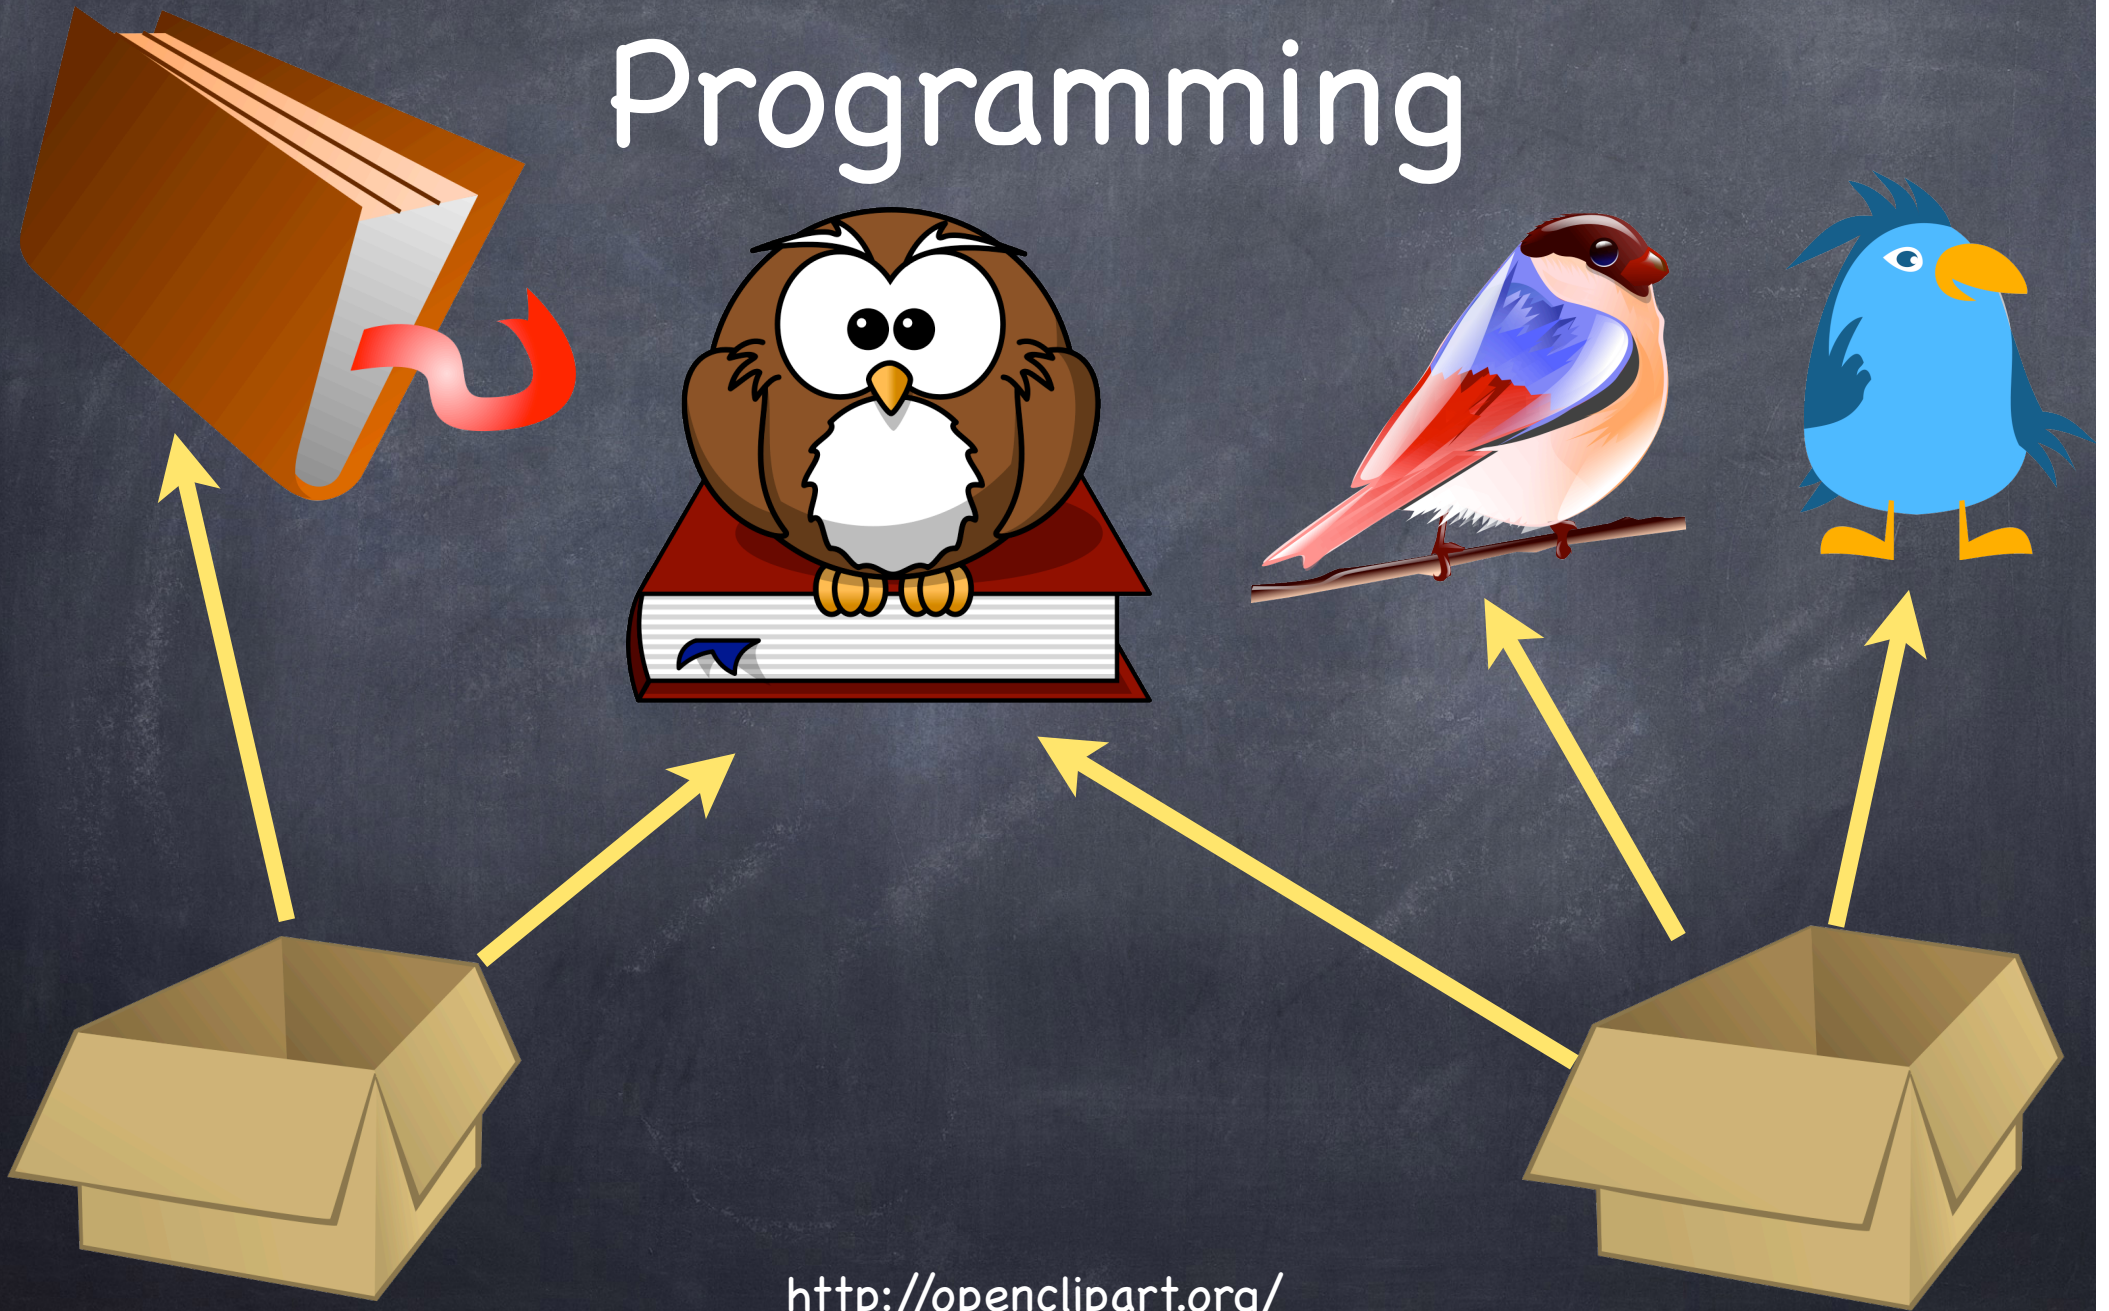
“The big lie of object-oriented programming is that objects provide encapsulation.”

John Hogg

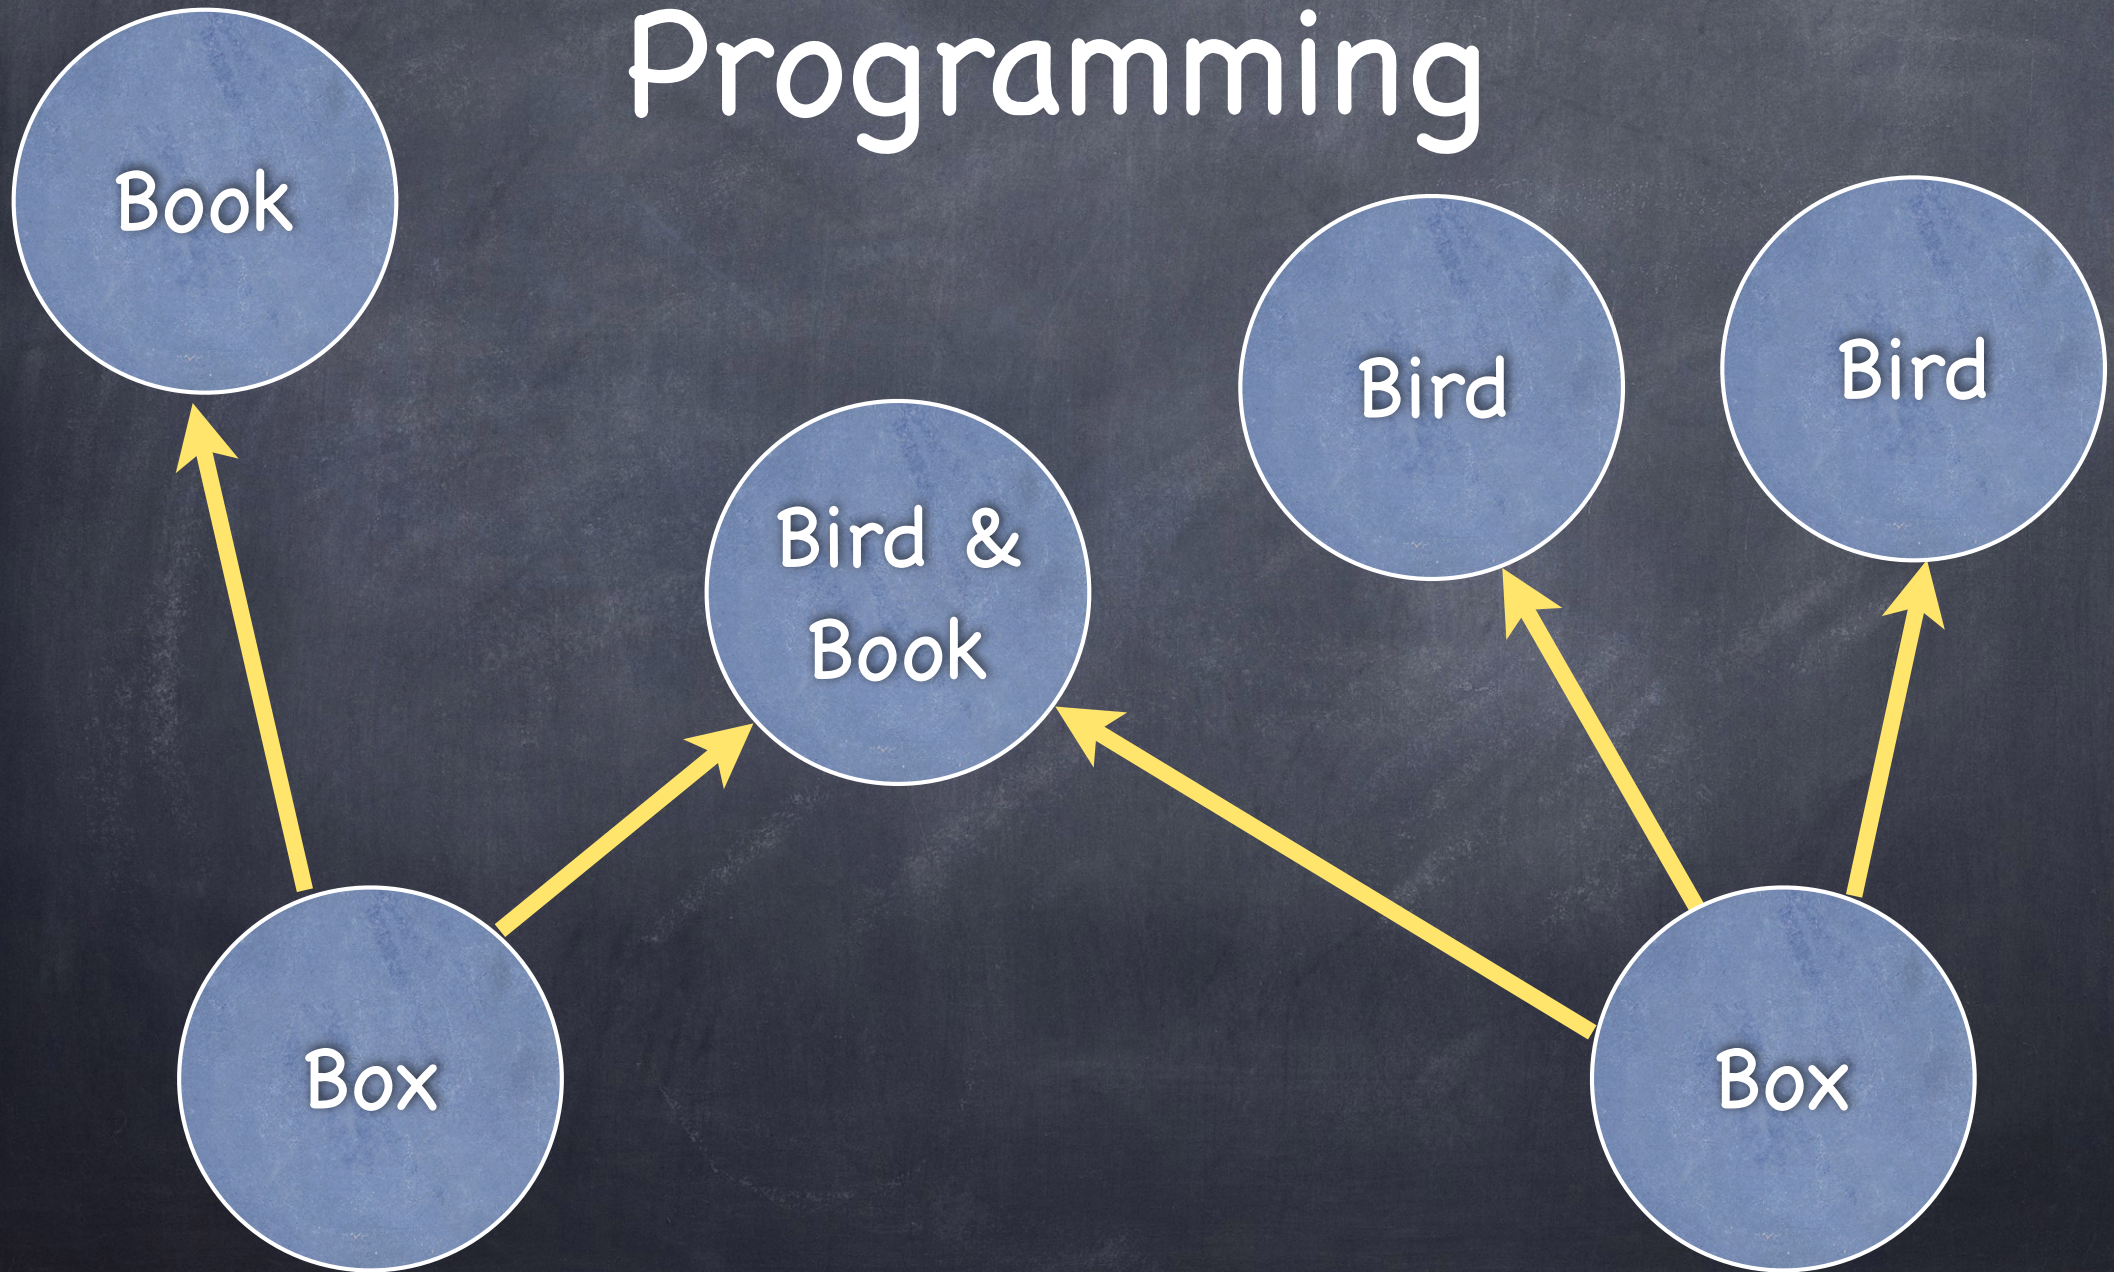
Islands: Aliasing Protection in Object-Oriented Languages

OOPSLA 1991

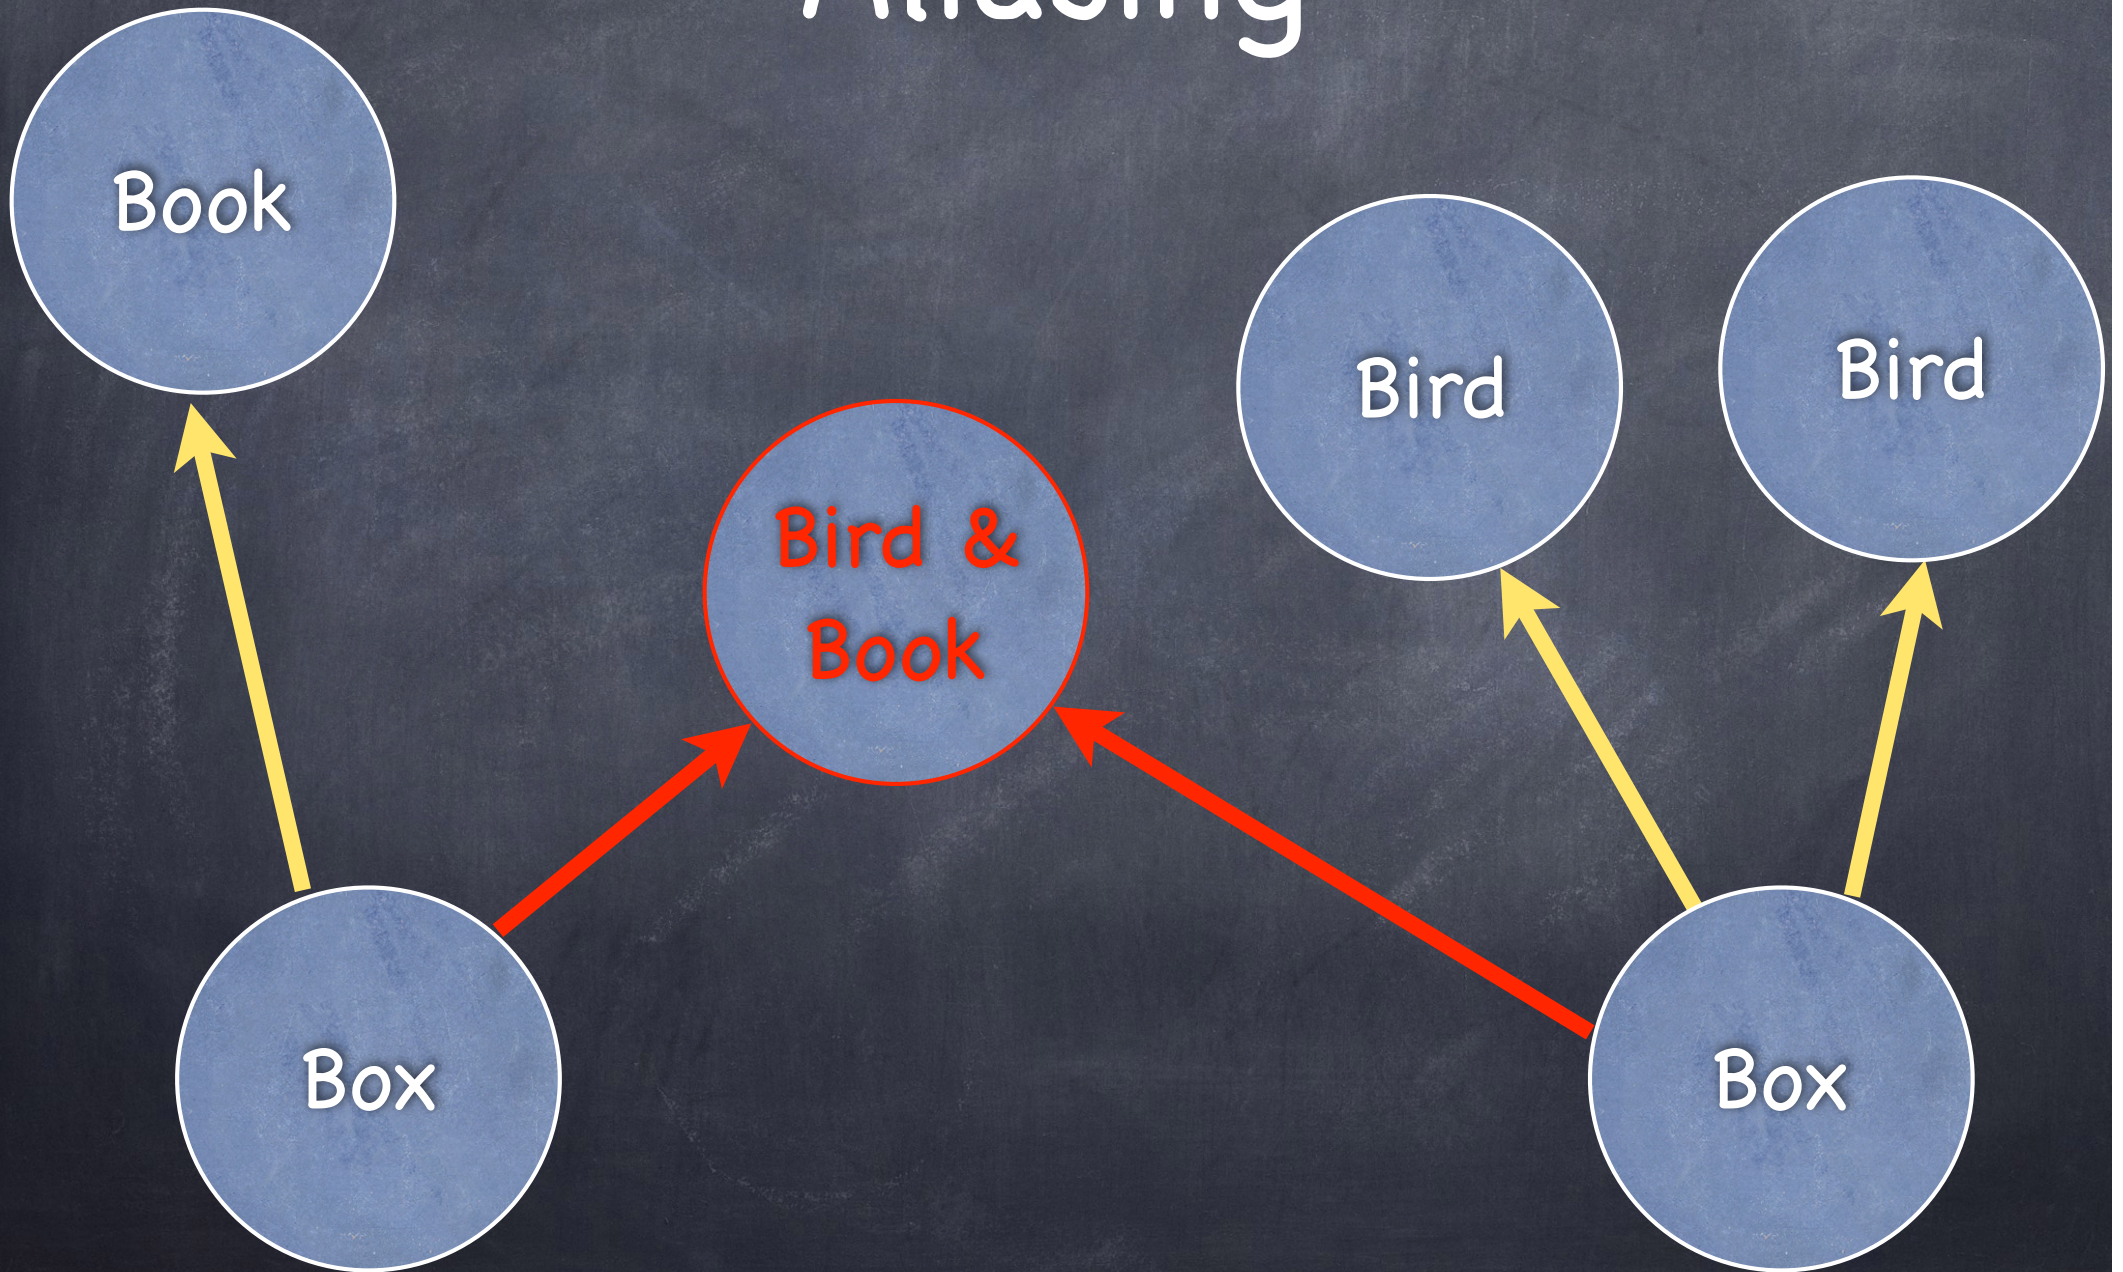
# Object-Oriented Programming



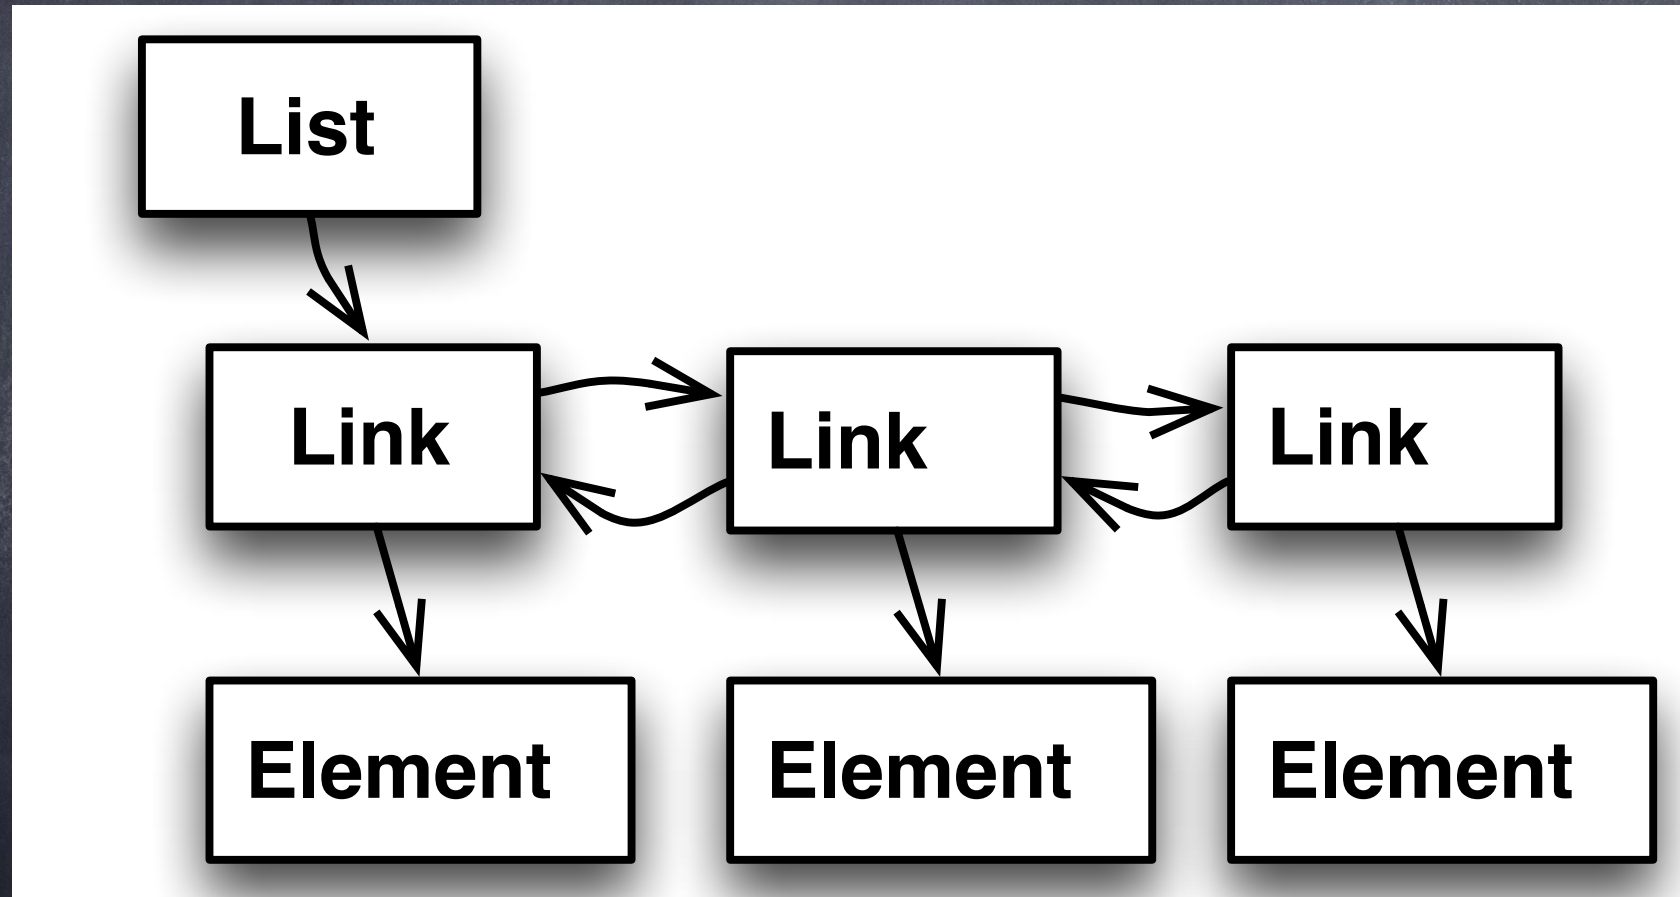
# Object-Oriented Programming



# Aliasing




# Aliasing (The Good)



# Aliasing (The Bad)

```
class Rectangle {  
    private Point topLeft;  
    private int w, h;  
    public Rectangle (Point topLeft, int w, int h)  
        { this.topLeft = topLeft; this.w = w; this.h = h; }  
}
```



...

```
Point p = new Point(100, 50);
```

```
Rectangle r = new Rectangle(p, 300, 200);
```

...

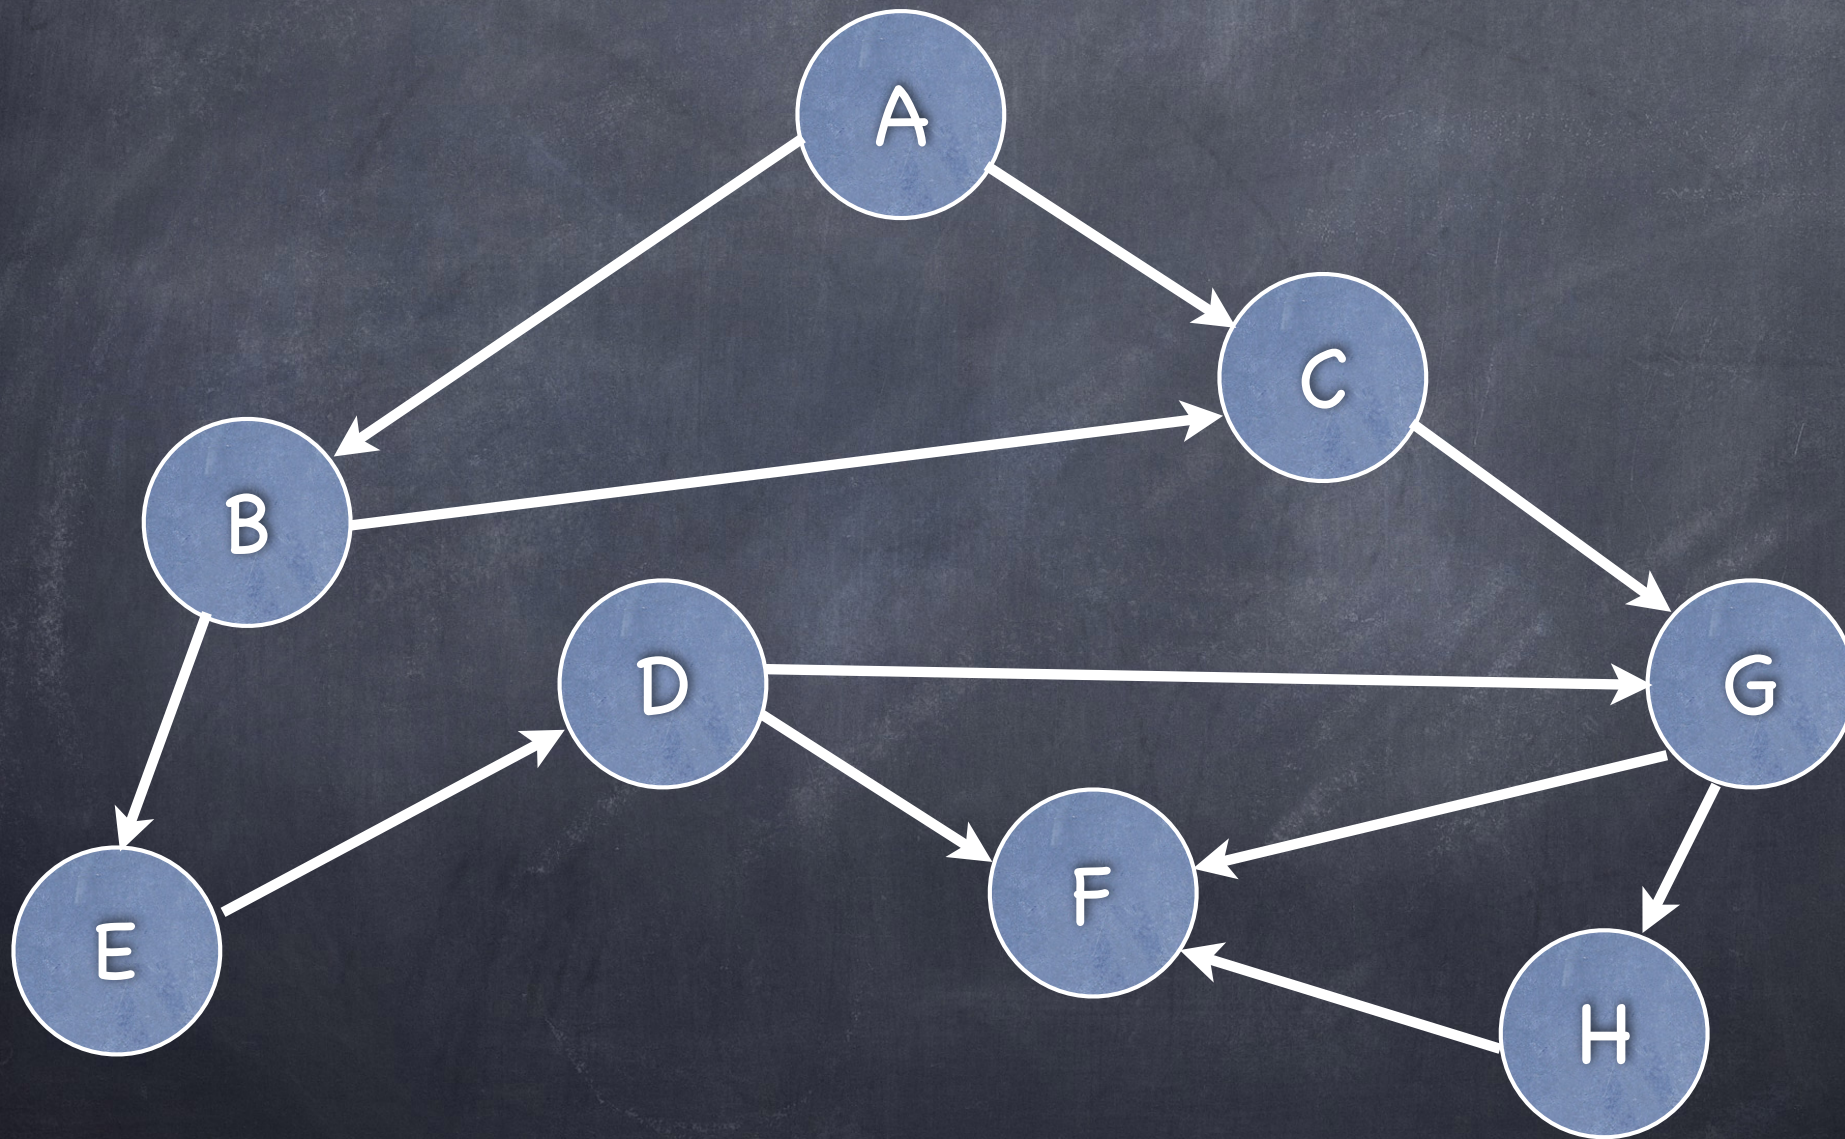
```
p.setX(400);
```



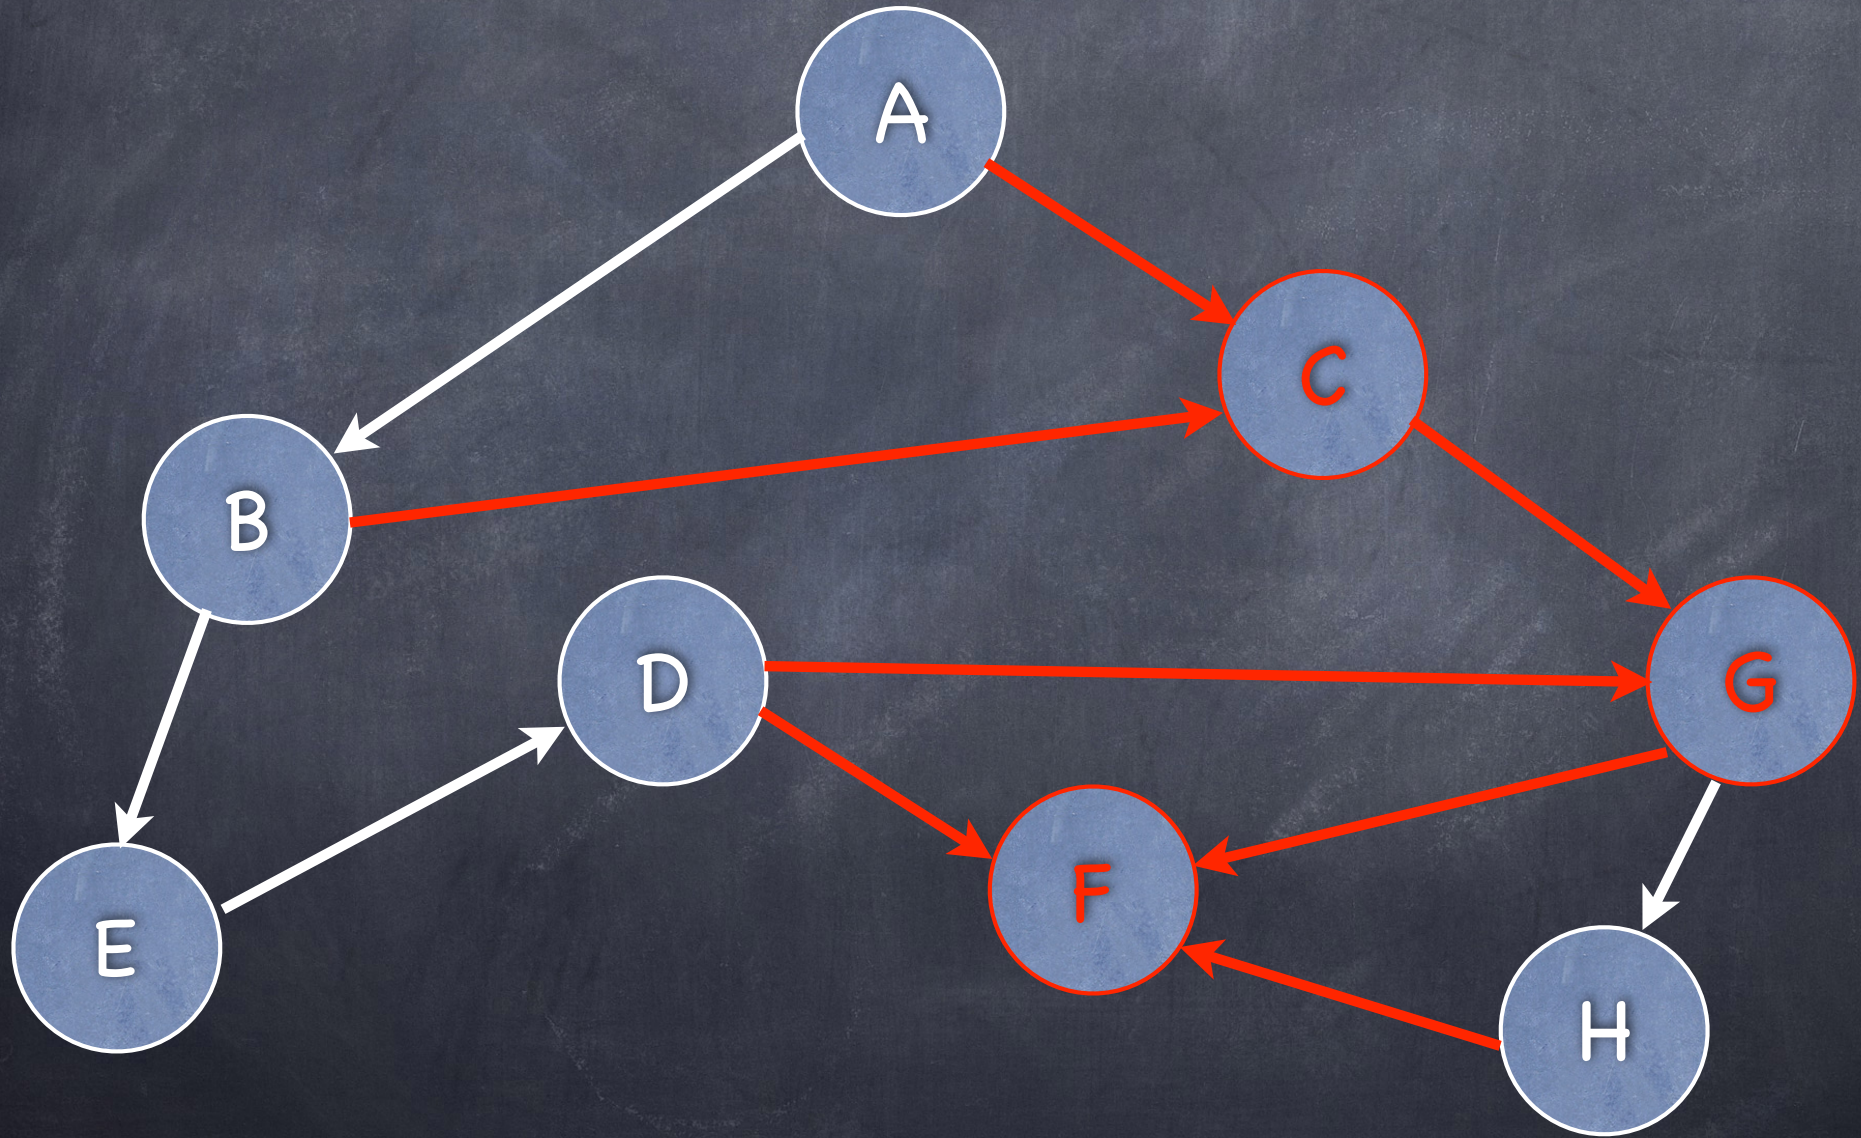
# Aliasing (The Ugly)

- Bugs due to unintentional aliasing are hard to track down (e.g. applets breaking out of JDK v1.1.1 sandbox)
- Typically enforced by coding style only (e.g. CMU SEI CERT OBJ05-J advises to “defensively copy private mutable class members before returning their references”)
- Still no mainstream language support for enforcing per-object encapsulation after 20+ years... (Rust?)

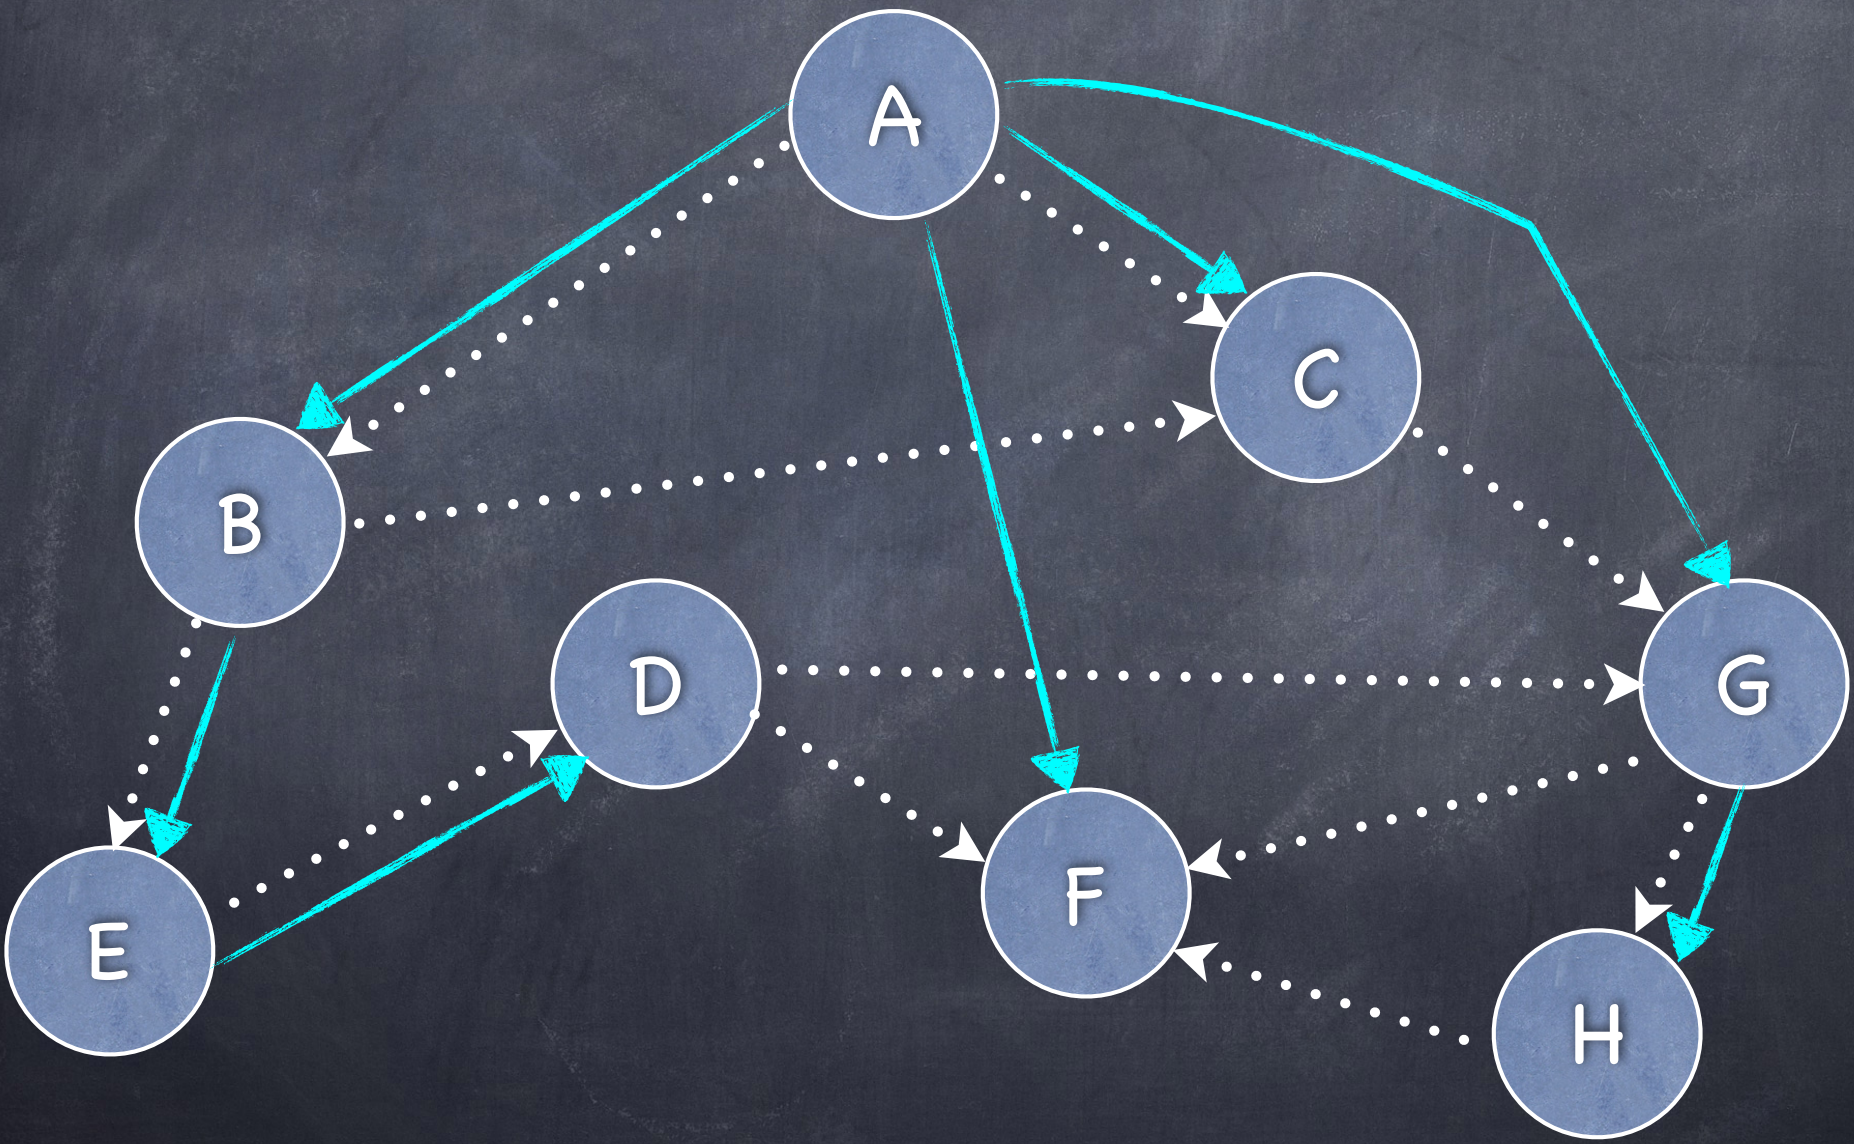
# Object Graph




# Object Graph (Aliasing)



# Ownership Tree




# Ownership Tree: HelloWorld






# Back to LinkedList




# Encapsulated LinkedList






# LinkedList with Iterator



# Our Questions

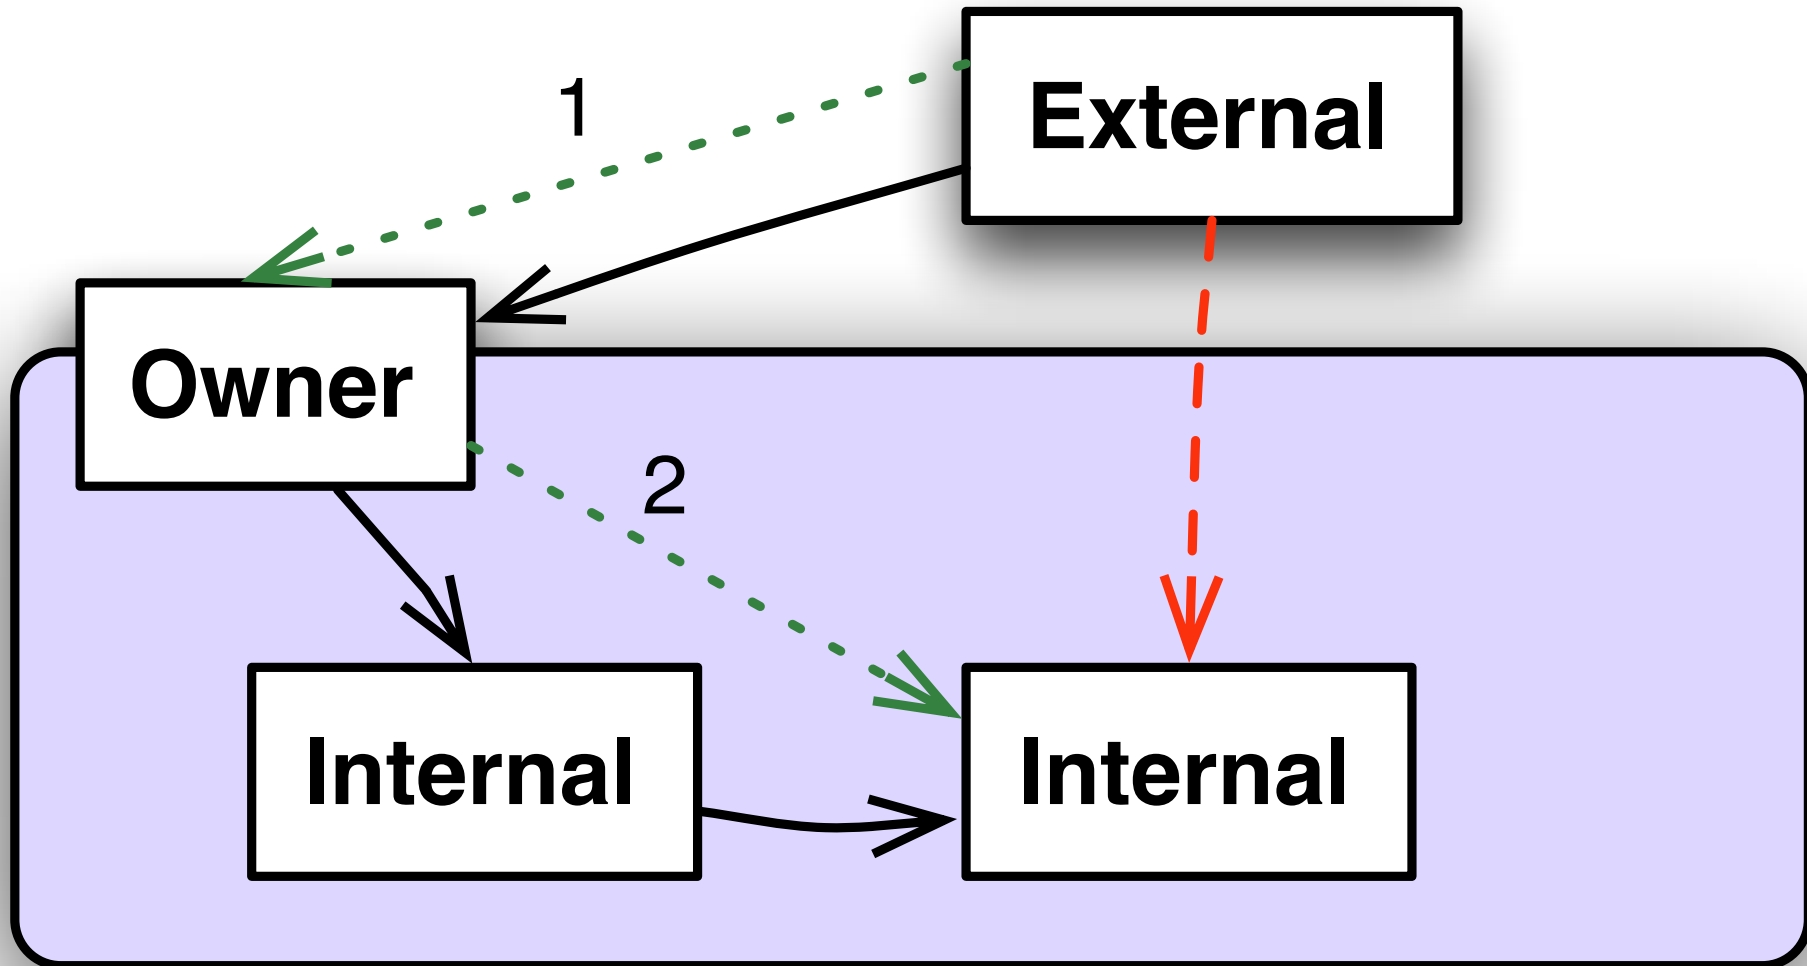
- How must designs change to respect encapsulation?
- What performance cost do these changes impose?
- How does this impact programs' performance?

# More on Ownership



# Owners as Accessors

(IWACO 2014)



# Owner-as-Accessor

## Example

```
class Internal { String s = "abc"; }
class Owner {
    Internal i = new Internal();
    void modifyI() { i.s = "def"; } }
class External {
    Internal i; Owner o = new Owner();
    void doIt() {
        this.i = o.i; // OK. Stores external ref.
        o.modifyI(); // OK. Modifies Internal.
        // this.i.s = "ghi"; // WRONG. Not owner!
    } }
}
```

# Java Collections Framework v1.5.0


		Implementations				
		Hash Table	Resizable Array	Balanced Tree	Linked List	Hash Table + Linked List
Interfaces	Set	<a href="#">HashSet</a>		<a href="#">TreeSet</a>		<a href="#">LinkedHashSet</a>
	List		<a href="#">ArrayList</a>		<a href="#">LinkedList</a>	
	Map	<a href="#">HashMap</a>		<a href="#">TreeMap</a>		<a href="#">LinkedHashMap</a>

- Effectively: Lists and Maps (implemented using Hash, Array, List or Tree).
- Our Implementations Available Online:

<http://homepages.ecs.vuw.ac.nz/~alex/software/>


# Lists

# ArrayList (and Vector)





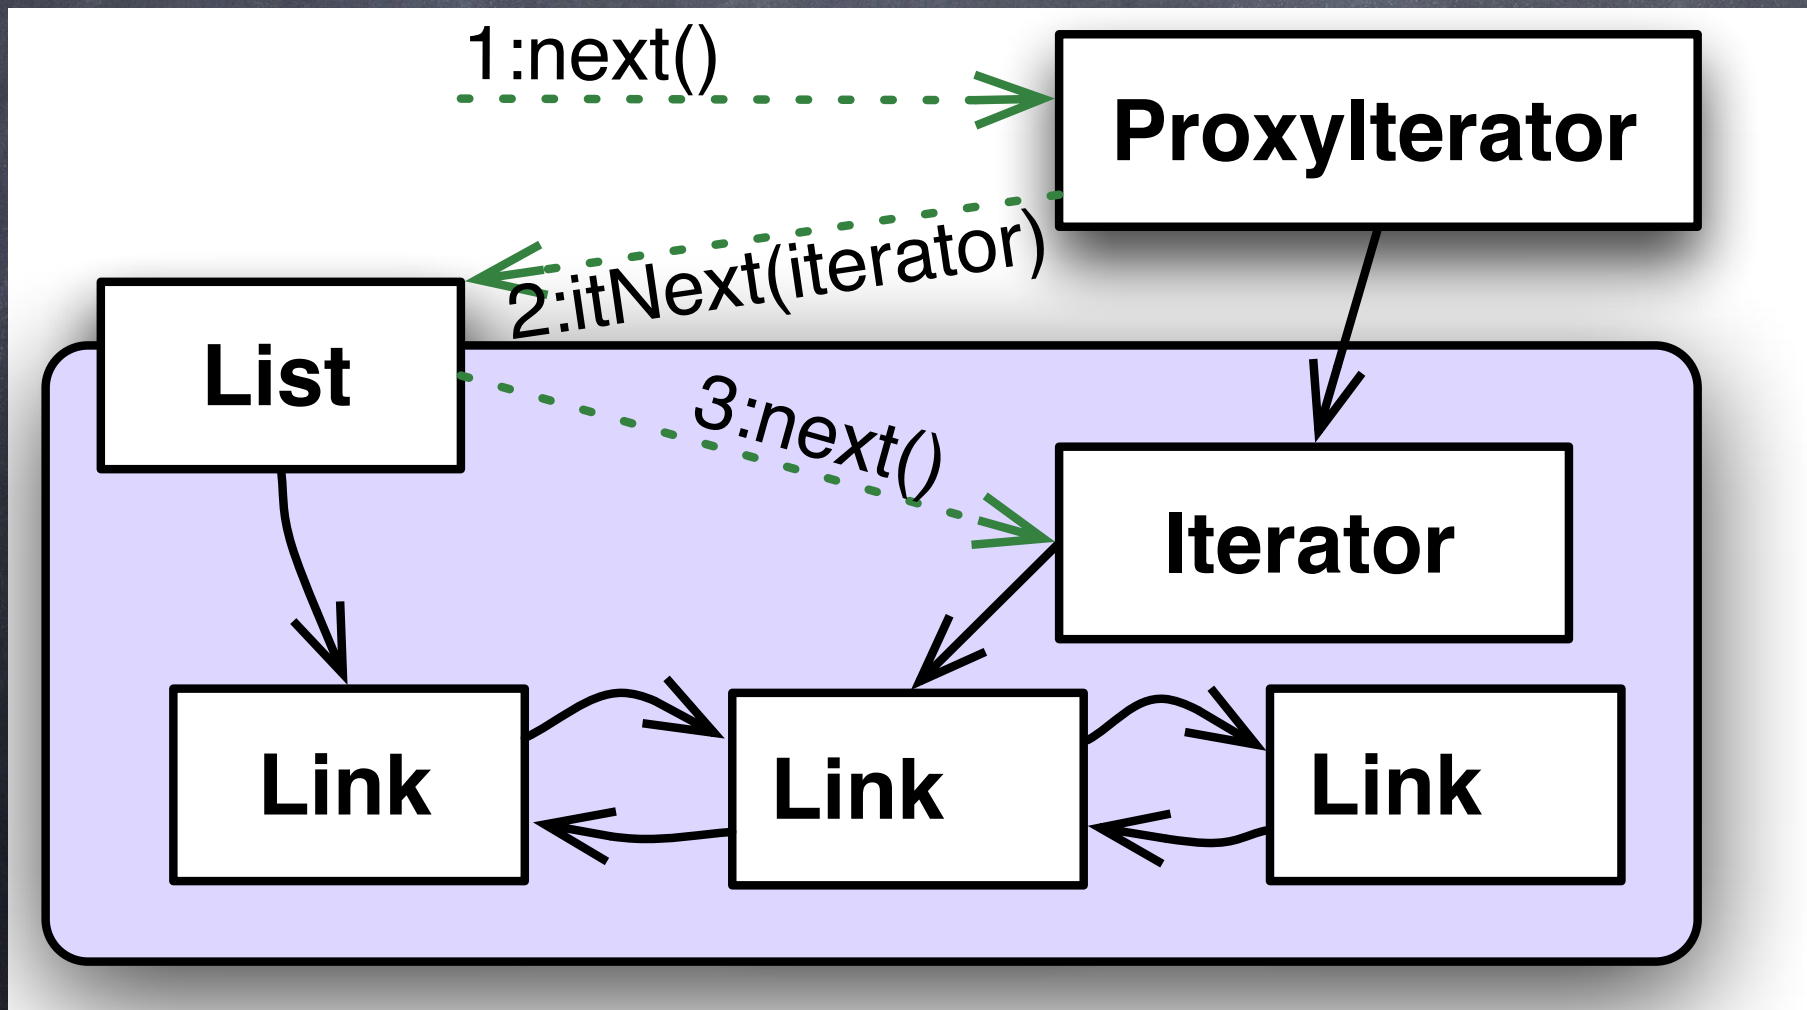
# What About LinkedList?



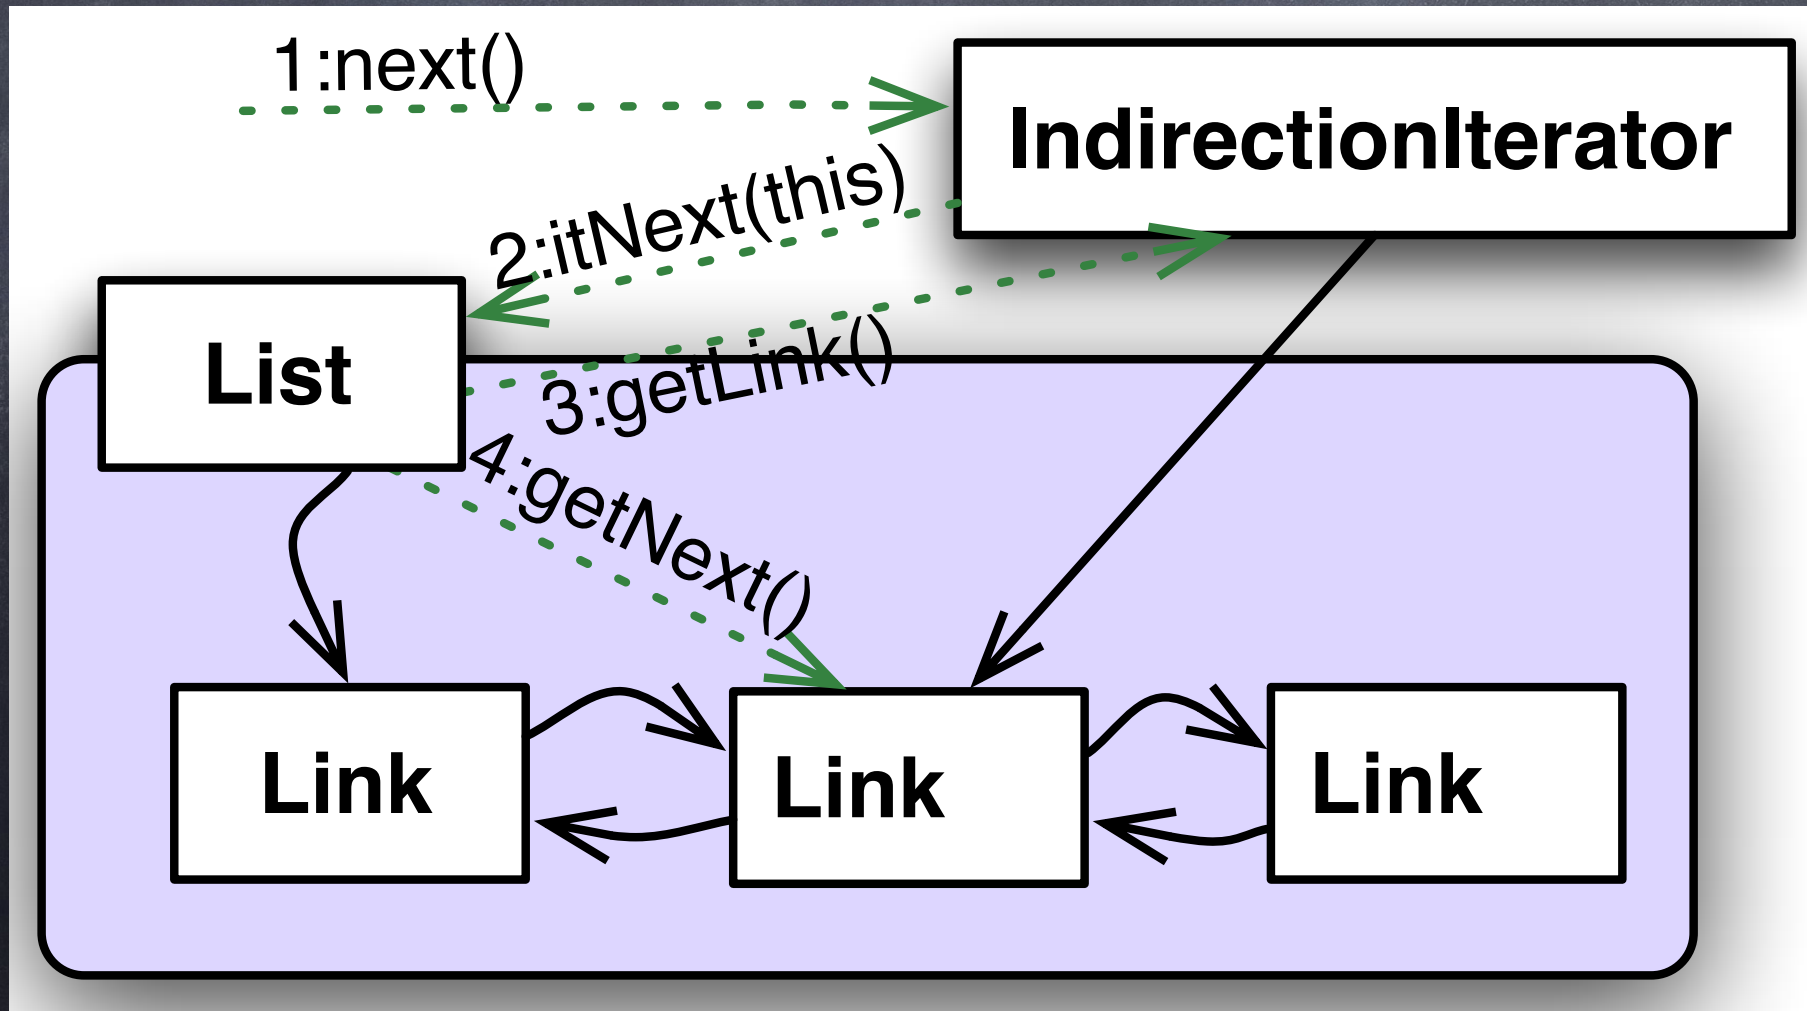
# Owner-as-Dominator

- Naive implementation that uses just the interface gives  $O(N^2)$  iteration of the whole list (duh!)
- Single Place Cache (caching last accessed item and its index) improves it back to  $O(N)$  for non-random iteration – complies with specification

# Owner-as-Accessor (Proxy Iterator)



# Owner-as-Accessor (Indirection Iterator)

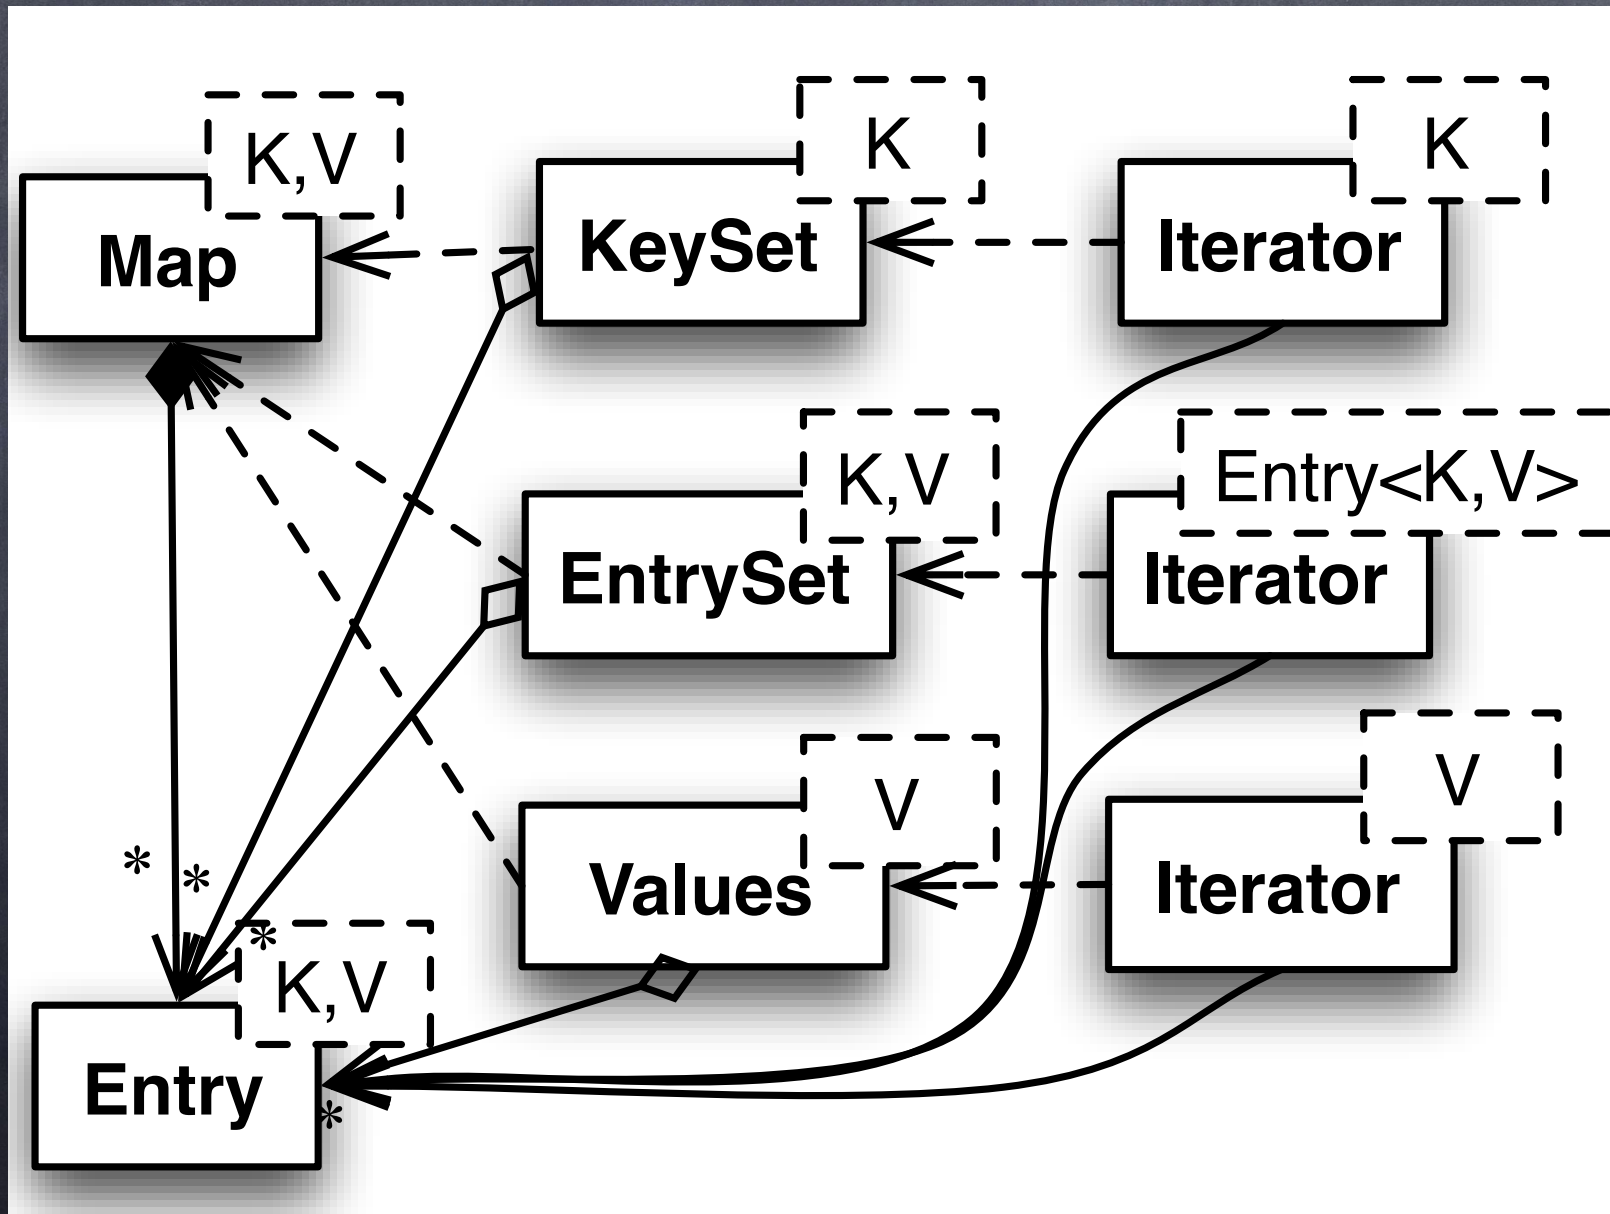


# Other Possible Iterators

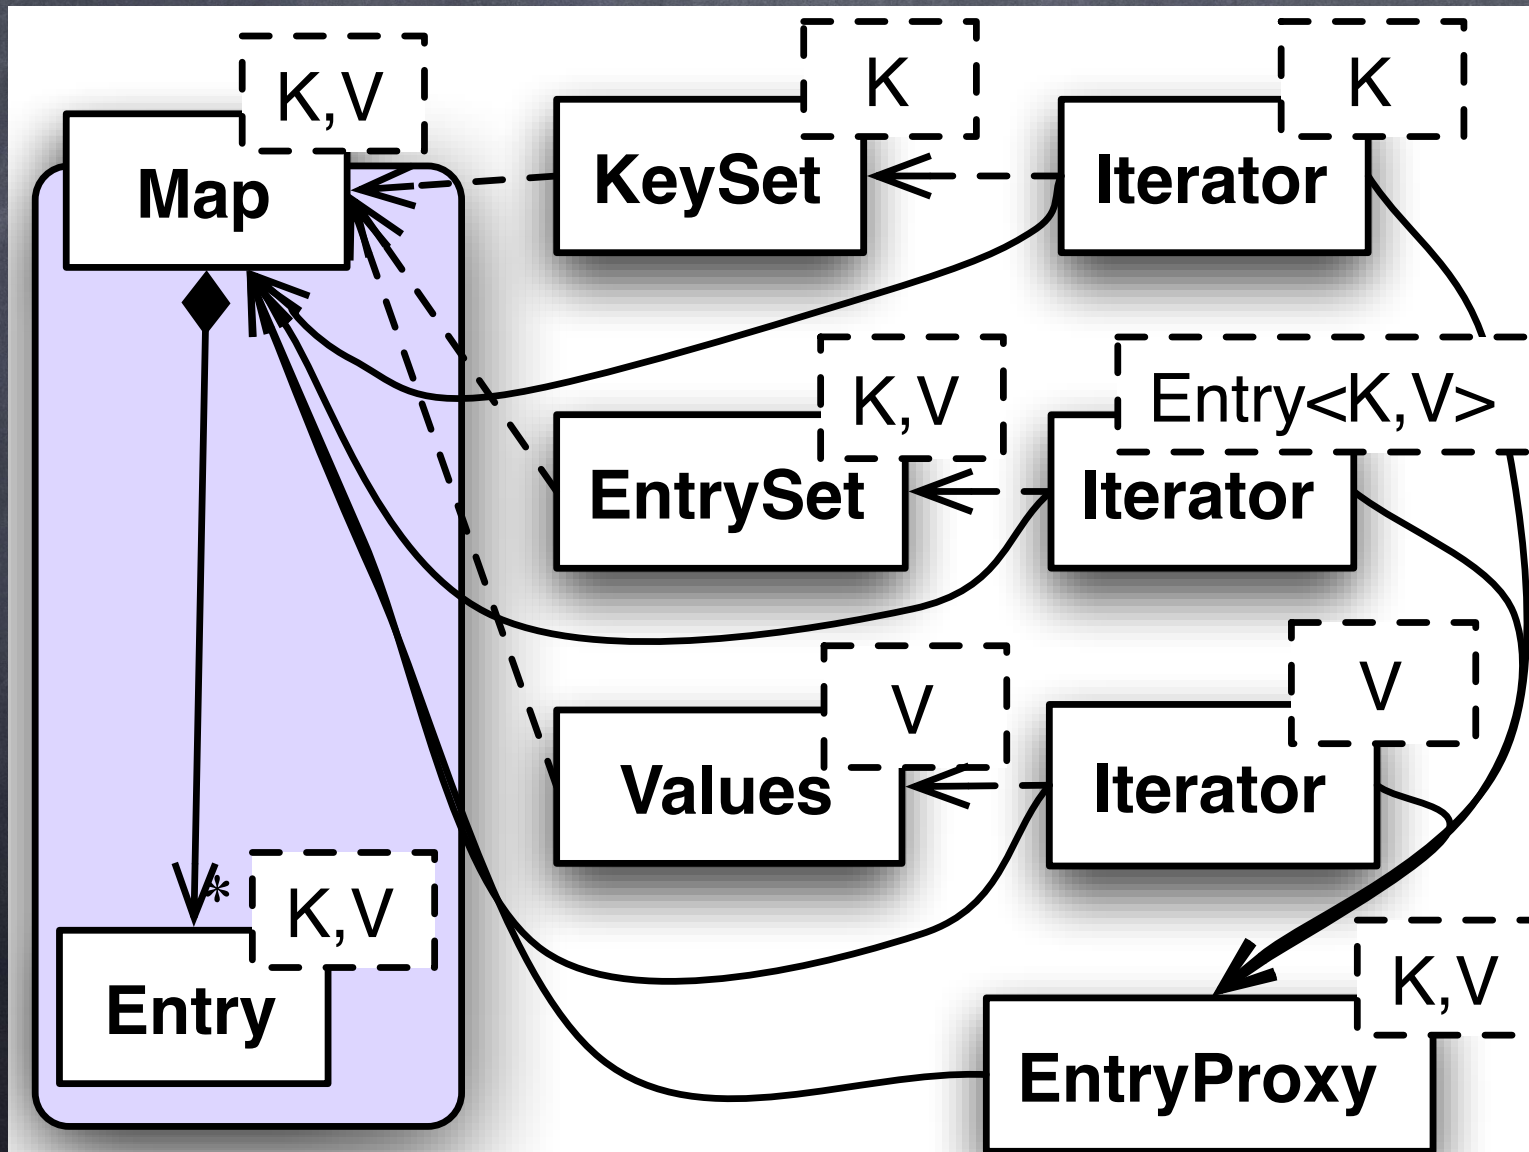
- “Magic Cookie” that uses a unique ID for each Iterator and stores its internal state inside a list in a hash map
- See: “Iterators and Encapsulation” by James Noble in TOOLS2000

Maps

# Map Interfaces



# [Linked]HashMap and Hashtable





# Extended Map Interface

```
/* *****  
/* OWNERSHIP additions for map external iterators      */  
/* *****  
int getModificationCount();  
K getFirstKey();  
K getNextKey(K key);  
boolean hasNextKey(K key);  
Map.Entry<K,V> entryProxyForKey(K key);  
void forall(Procedure<K,V> proc);
```

# Summary of Map Changes

- HashMap's HashIterator was changed to utilise the extended Map interface
- LinkedHashMap was easier as the "getNextKey" could make use of linked list woven through the map's entries
- Hashtable (like Vector) was fundamentally the same but used different interfaces and was fully synchronised
- Finally, we provided 4 more owner-as-accessor refactorings following the ones described in LinkedList

# TreeMap

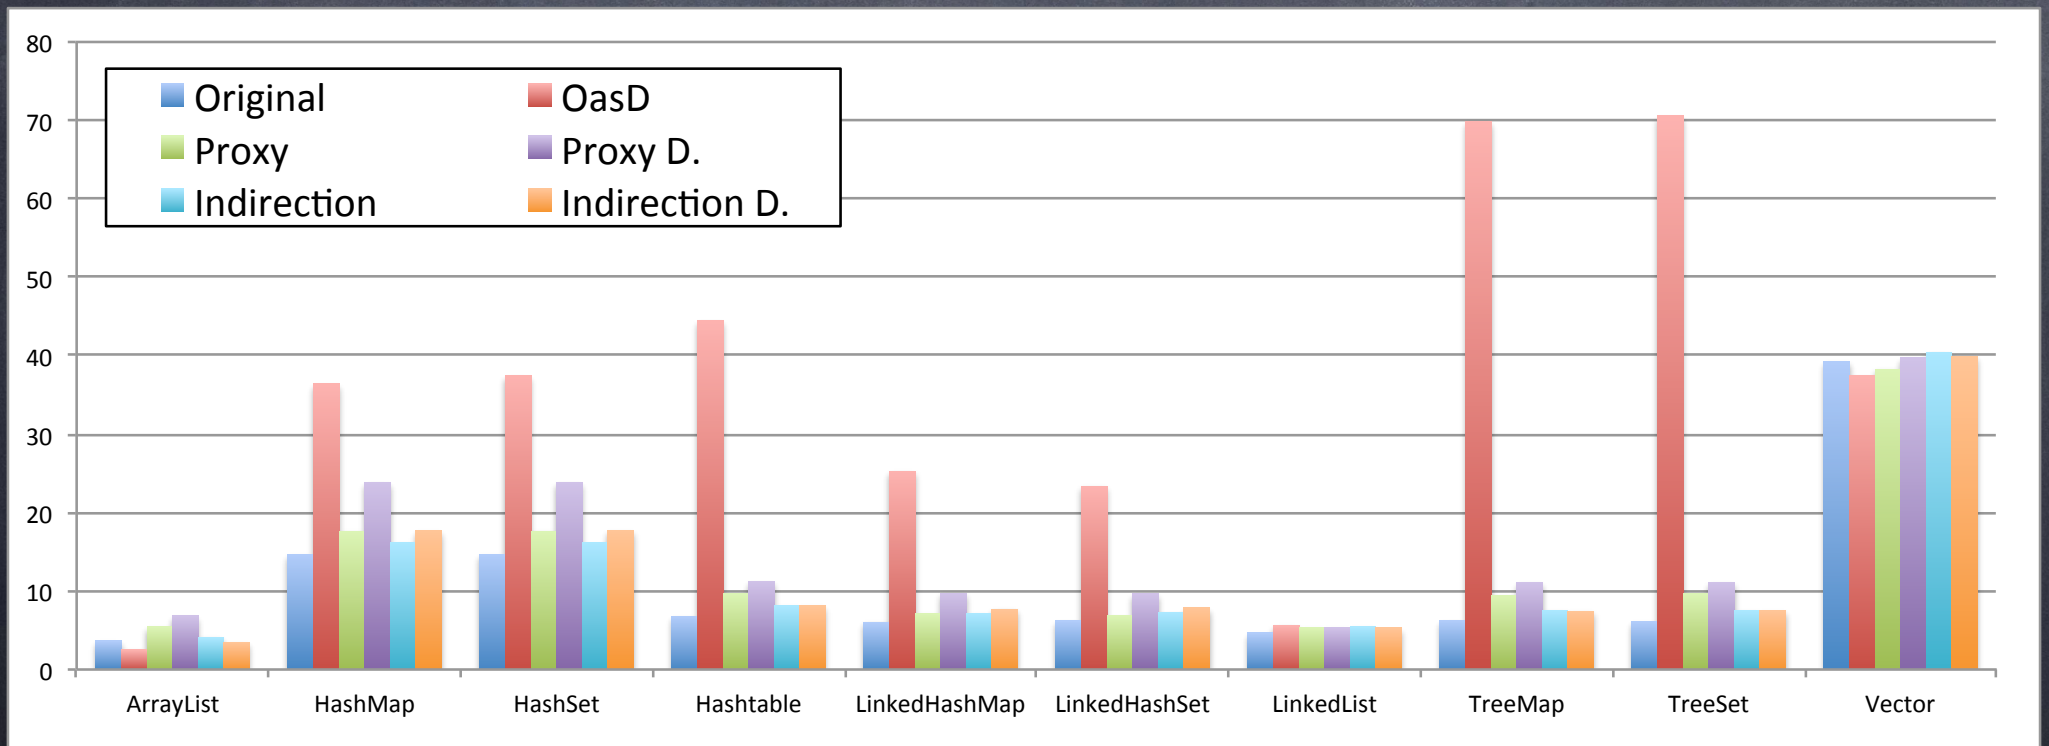
- Originally planned to refactor separately...
- Turns out we could re-use all the iterator, view, and entry objects from HashMap as they were just using Map's public interface!
- Calling getNextKey(currentKey) results in tracing down from the root of Red-Black tree (more costly and complex operation), but improvements with caching are possible

# Measurements

# Three Microbenchmarks

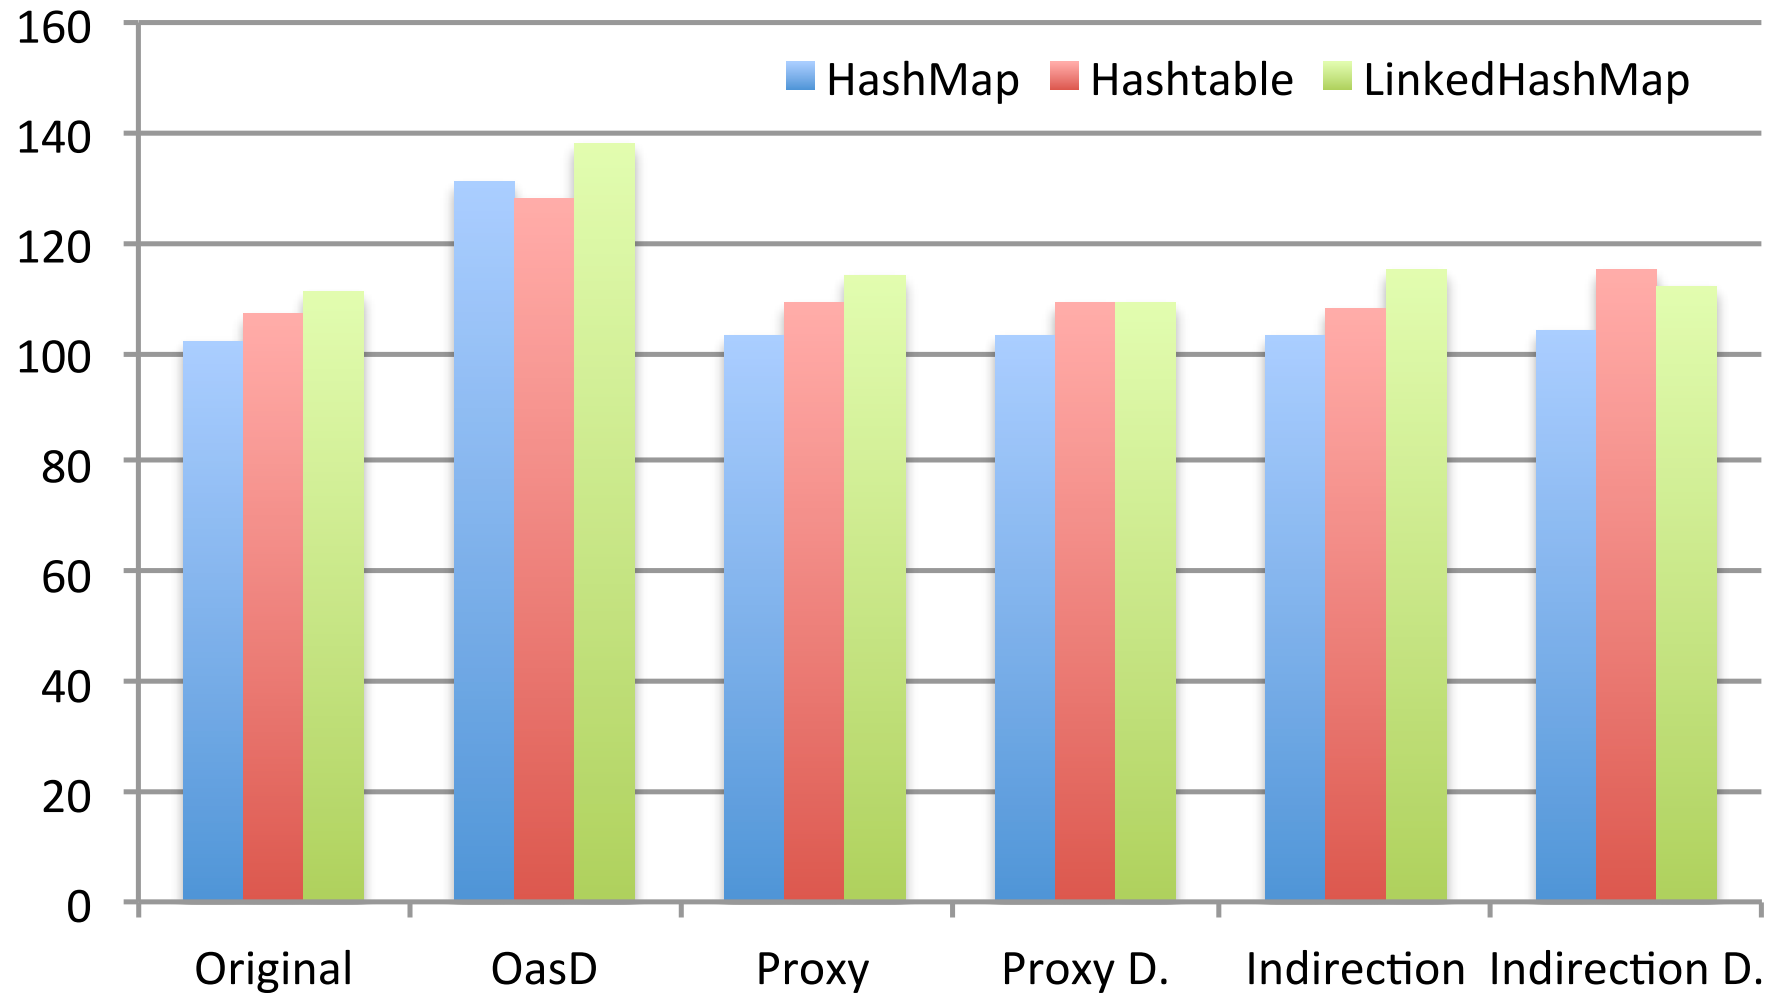
- Doug Lea's IteratorLoops from JSR166
- LinkedList Iteration: forward/backward/  
disruptive
- Doug Lea's MapMicroBenchmark from JSR166
- NB! We ensured that GC and JIT did not interfere with our tests and ran each 25 times (see Georges et al OOPSLA2007)

# IteratorLoops





# MapMicroBenchmark





# Three Macrobenchmarks

- DaCapo (Typical Usage, Improves on SPEC)
- SPECjbb2005
- SPECjvm2008

TABLE I

FULL RESULTS TABLE (FORMAT: MEAN|SD; STATISTICALLY SIGNIFICANTLY DIFFERENT VALUES SHOWN IN BOLD)

Benchmark	Original		OasD		Proxy		Proxy D.		Indirection		Indirection D.		Number	Percent
DaCapo (Time in <i>ms</i> ; lower is better)														
avro	23003	300	<b>22781</b>	<b>415</b>	<b>22816</b>	<b>236</b>	22867	179	22948	389	22873	323	219049309	78.08%
batik	2516	34	2517	19	2520	25	2519	19	2519	25	<b>2528</b>	<b>30</b>	26507124	31.37%
eclipse	53793	1031	53480	936	53716	1010	53812	1329	<b>53554</b>	<b>1406</b>	53608	738	355465429	33.63%
fop	393	27	397	29	394	20	397	23	396	23	399	20	1874892	18.07%
h2	24133	580	24238	593	24141	517	24188	380	23967	320	23934	375	90175446	8.04%
jython	15041	215	<b>15476</b>	<b>100</b>	<b>15725</b>	<b>107</b>	<b>15719</b>	<b>53</b>	<b>15837</b>	<b>110</b>	<b>17050</b>	<b>145</b>	159700109	7.49%
luindex	705	18	<b>687</b>	<b>23</b>	714	22	713	19	710	40	718	51	327466	36.66%
lusearch	7251	184	7334	105	<b>7181</b>	<b>189</b>	<b>7120</b>	<b>198</b>	7226	299	7325	78	11979688	5.45%
pmd	3944	47	<b>3992</b>	<b>39</b>	<b>4046</b>	<b>59</b>	<b>4065</b>	<b>72</b>	<b>4005</b>	<b>64</b>	<b>4054</b>	<b>67</b>	10712544	36.55%
sunflow	22656	518	22560	365	<b>23365</b>	<b>126</b>	22970	711	22851	523	<b>22331</b>	<b>207</b>	171198077	<0.01%
tomcat	7576	108	7641	135	<b>7687</b>	<b>134</b>	<b>7736</b>	<b>111</b>	<b>7733</b>	<b>88</b>	<b>7661</b>	<b>118</b>	16726923	13.95%
tradebeans	27952	409	<b>27556</b>	<b>494</b>	<b>28258</b>	<b>506</b>	28142	275	27998	328	28020	340	1621619	33.00%
tradesoap	64476	1251	65193	1549	65042	1712	65119	1463	64390	1378	65111	1499	1631193	32.82%
xalan	26604	384	26692	318	<b>26383</b>	<b>247</b>	<b>26173</b>	<b>258</b>	<b>26125</b>	<b>251</b>	<b>26310</b>	<b>291</b>	61153799	13.23%
SPECjbb2005 (Throughput; higher is better)														
SPECjbb2005	29598	405	<b>14062</b>	<b>181</b>	<b>28825</b>	<b>860</b>	<b>28540</b>	<b>619</b>	<b>28959</b>	<b>764</b>	<b>28394</b>	<b>641</b>	35542855	4.87%
SPECjvm2008 (Time in <i>ms</i> ; lower is better)														
compress	46.56	0.81	46.77	0.59	46.71	0.59	46.82	0.42	46.83	0.42	46.84	0.57	199478	20.21%
crypto.aes	18.49	0.18	18.40	0.13	18.46	0.10	18.48	0.18	<b>18.42</b>	<b>0.10</b>	18.48	0.19	254853	22.00%
crypto.rsa	35.04	0.25	34.89	0.28	34.92	0.31	34.95	0.36	34.94	0.25	34.95	0.27	6535358	11.18%
crypto.signverify	53.06	0.27	52.97	0.31	53.11	0.29	53.16	0.36	53.09	0.38	53.12	0.29	3290783	2.13%
derby	21.55	0.48	21.63	0.40	21.56	0.50	21.73	0.43	21.59	0.38	21.61	0.33	89061937	3.04%
mpegaudio	15.08	0.05	15.10	0.05	15.08	0.06	15.11	0.05	15.10	0.06	15.07	0.06	209089	20.32%
fft.large	23.61	0.27	23.53	0.30	23.49	0.27	<b>23.43</b>	<b>0.28</b>	23.55	0.24	23.51	0.38	168290	23.78%
fft.small	82.75	4.49	84.24	5.21	84.98	3.94	83.35	4.01	84.46	4.70	83.57	4.26	439646	9.22%
lu.large	6.98	1.44	6.69	1.24	6.72	1.22	7.05	1.39	6.34	0.96	7.02	1.42	166728	23.92%
lu.small	107.98	0.87	107.61	0.92	107.48	0.75	107.58	0.84	107.94	0.70	107.82	0.85	642713	6.39%
monte_carlo	15.33	1.49	15.33	1.48	15.65	0.10	15.67	0.11	15.63	0.03	15.29	1.47	179253	22.58%
sor.large	13.01	0.02	13.01	0.02	13.01	0.01	13.02	0.02	<b>12.99</b>	<b>0.05</b>	13.01	0.01	167449	23.92%
sor.small	57.47	0.11	57.47	0.10	57.49	0.10	57.47	0.12	57.43	0.07	57.44	0.10	193425	21.16%
sparse.large	11.51	0.23	<b>11.97</b>	<b>1.22</b>	11.70	0.28	11.82	0.80	11.71	0.34	11.77	0.36	166691	23.98%
sparse.small	43.87	0.11	43.79	0.12	43.80	0.18	43.83	0.11	43.85	0.18	43.80	0.15	182676	22.17%
serial	33.22	0.98	<b>32.46</b>	<b>0.96</b>	33.27	0.81	33.04	0.92	32.88	1.31	33.31	0.89	51123977	5.68%
sunflow	20.51	0.50	20.48	0.50	20.55	0.32	20.42	0.52	<b>20.22</b>	<b>0.63</b>	20.51	0.47	42506559	0.12%
xml.validation	63.38	1.09	63.01	0.96	63.36	1.03	63.19	1.42	63.75	1.39	63.23	1.03	10318760	3.33%

# Macrobenchmarks

- Full table available in our ICSE2013 paper and an accompanying Technical Report ECSTR12-22:

<http://ecs.victoria.ac.nz/Main/TechnicalReportSeries>

- In only 40 out of 165 refactored benchmarks have we detected a statistically significant slow down!
- SPECjbb2005 is the heaviest user of collections (around 8% of its running time spent in java.util.\* methods according to option -Xrunhprof:cpu-times) and as a result slowed down the most


# Conclusions

- Encapsulation reduces performance by factors of 2 to 8 (in particular on microbenchmarks)
- Owner-as-dominator is the worst (as expected)
- Owner-as-accessor (even with dynamic checking) only produces less than 3% slowdown on macrobenchmarks
- We hope these results may encourage object-oriented designers to consider object encapsulation more carefully when designing their programs - especially their use of incoming aliases to circumvent encapsulation - and to ask themselves: are their incoming aliases really necessary?

# Since 2013...

- State of the macro benchmarks is still “to be improved”: some recent 2019 developments on DeCapo update \*as well as\* a competing performance corpora by Oracle Labs just released at PLDI 2019 (Renaissance - controversial)
- New common data structure libraries but still not necessarily treating aliasing more than an algorithmic side effect.

**Are Your  
Incoming  
Aliases  
Really  
Necessary?**





**KEEP  
CALM  
AND  
CARRY  
ON**

# Are your Incoming Aliases Really Necessary?

Counting the Cost  
of Object Ownership

Alex Potanin,  
Monique Damitio,  
James Noble

**Fri 24 May Session 2.3**