

Multigrain Coherence Directories

Dr. Jason Zebchuk

Principal Engineer, Cavium

Prof. Andreas Moshovos, *University of Toronto*

Prof. Babak Falsafi, *EcoCloud, EPFL*

Multigrain Coherence Directories

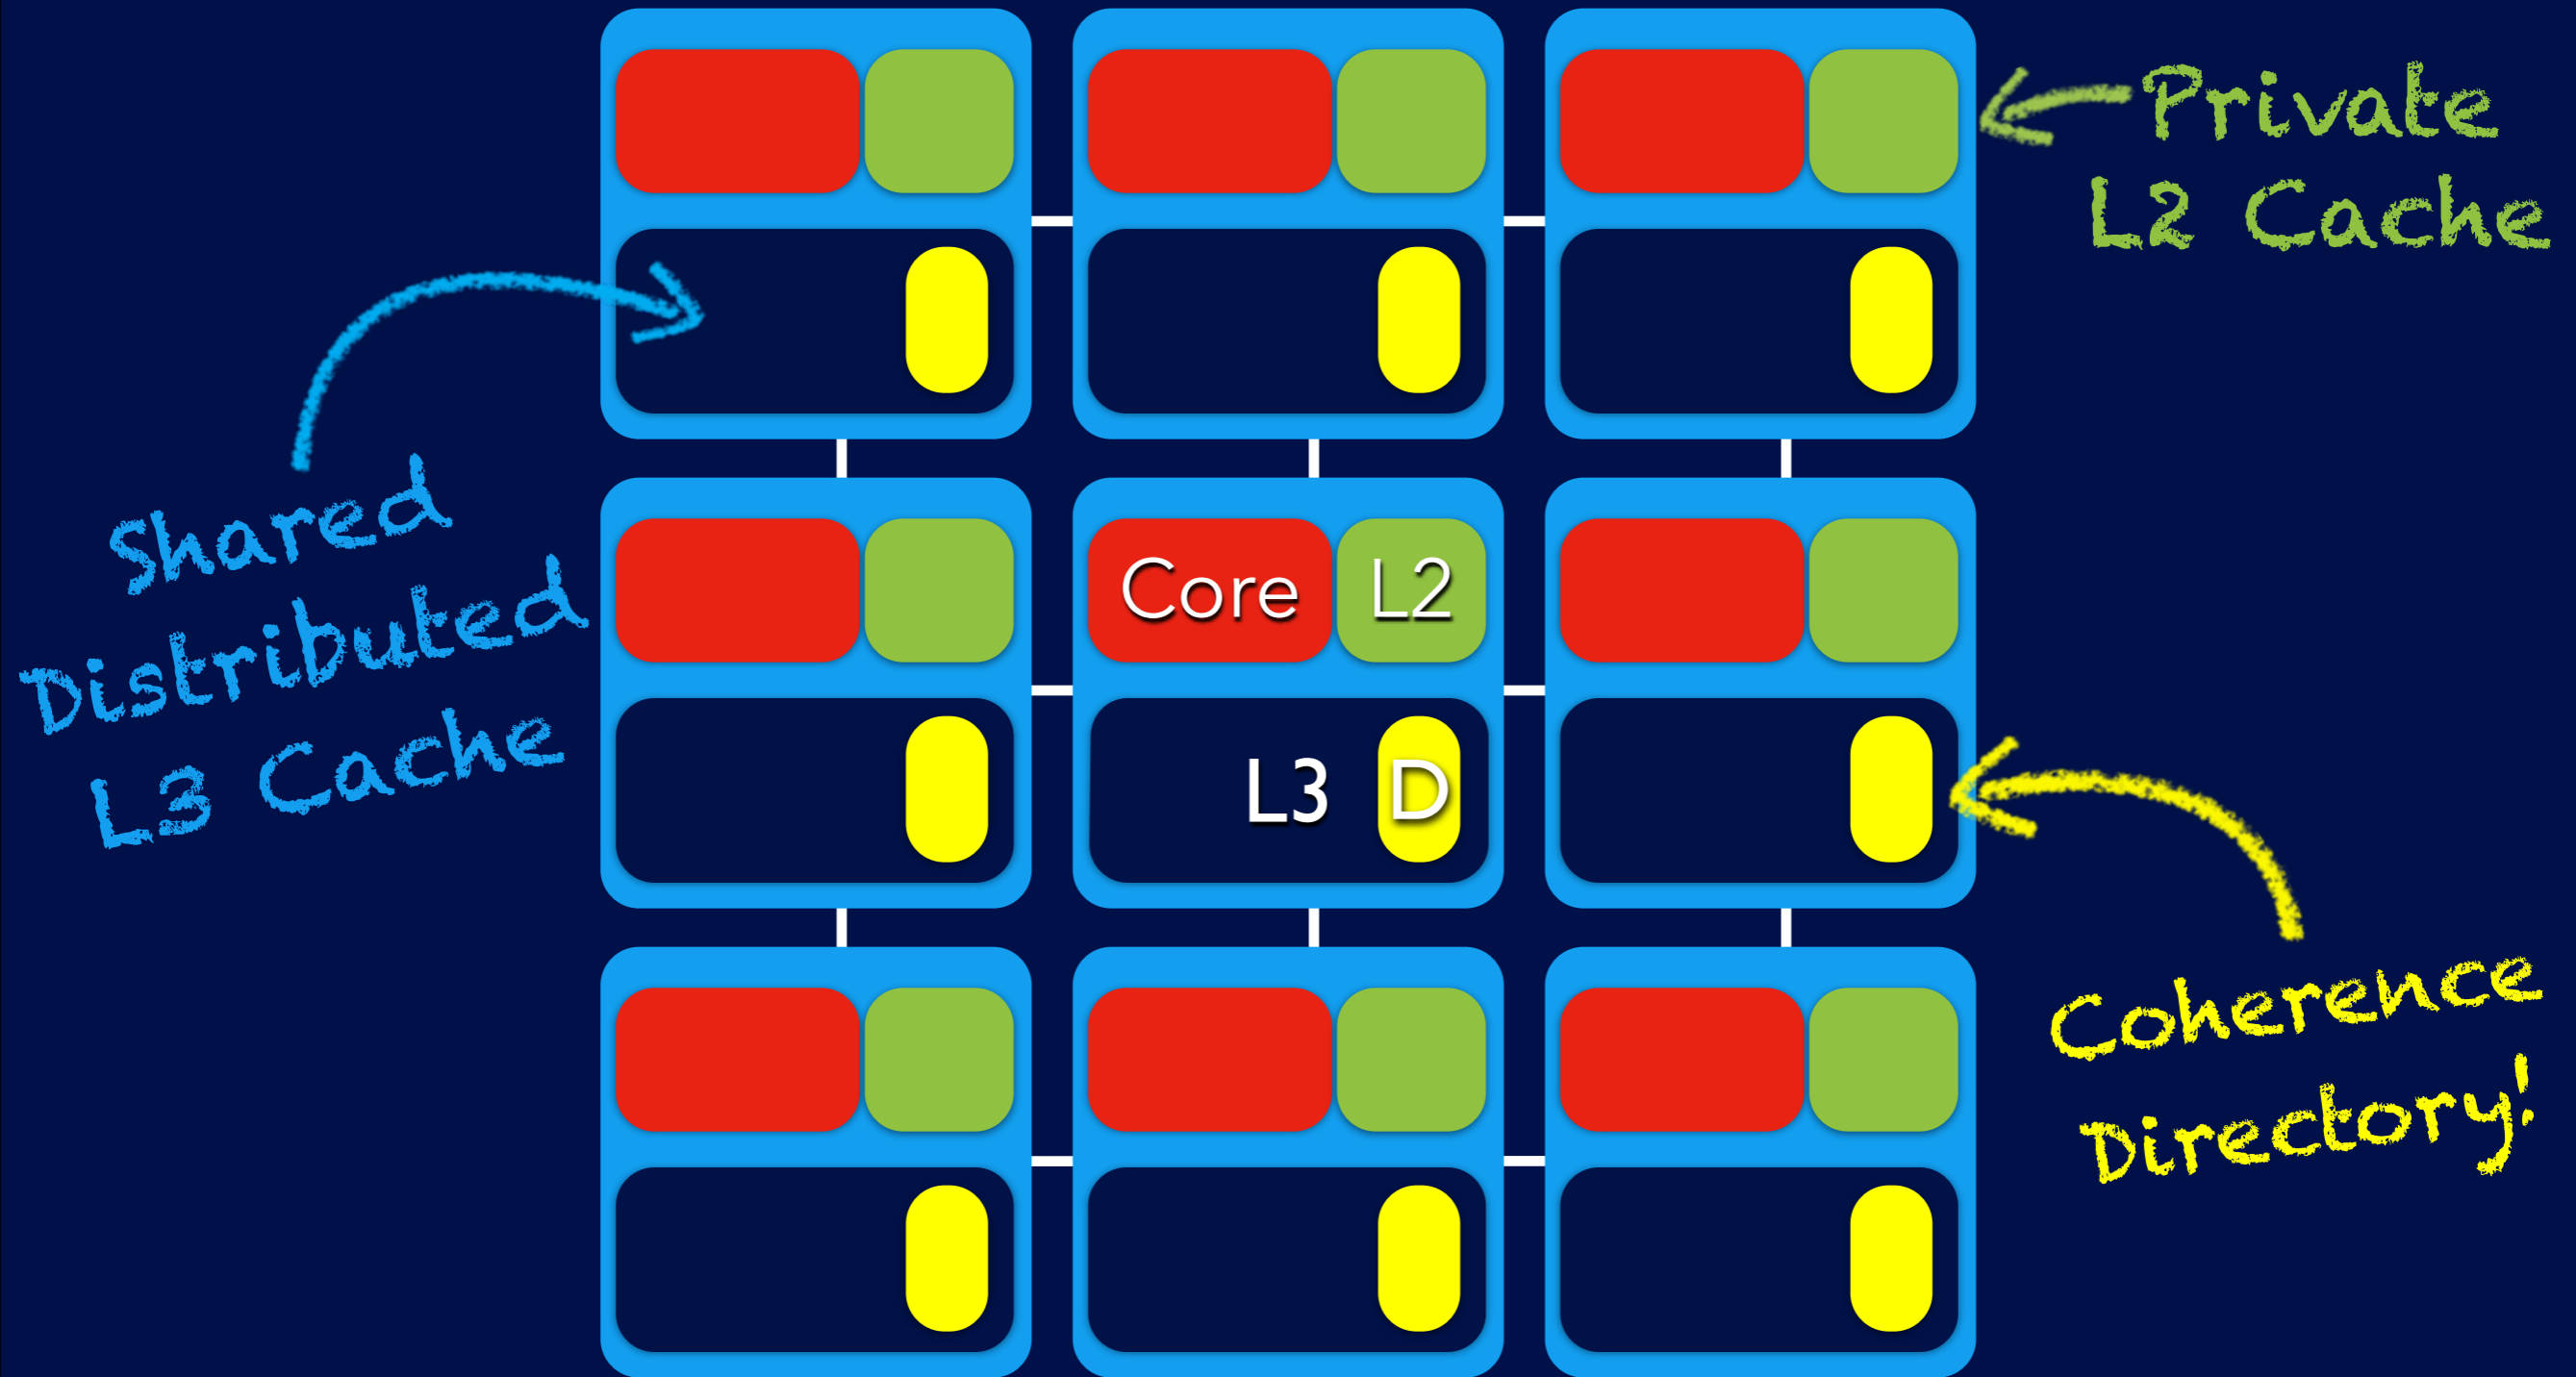
Dr. Jason Zebchuk

Principal Engineer, Cavium

Prof. Andreas Moshovos, *University of Toronto*

Prof. Babak Falsafi, *EcoCloud, EPFL*

Many-Core Processor




Memory

Core 1

Core 2

Core 3




Core 1


Core 2

Core 3


Memory



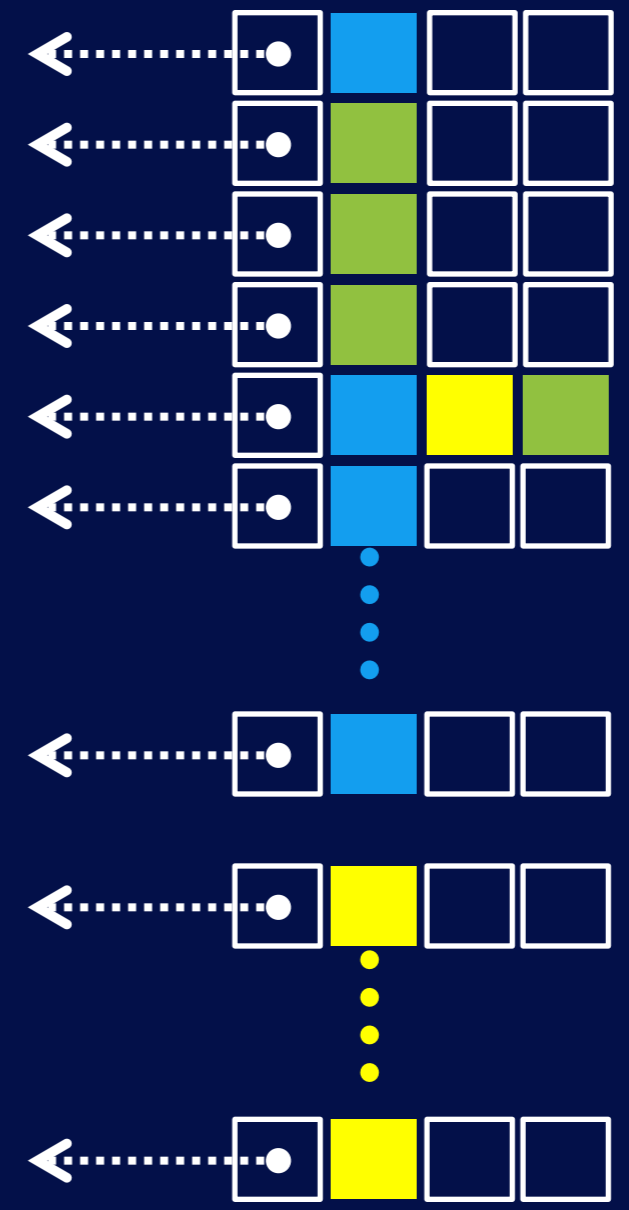
Coherence Directory

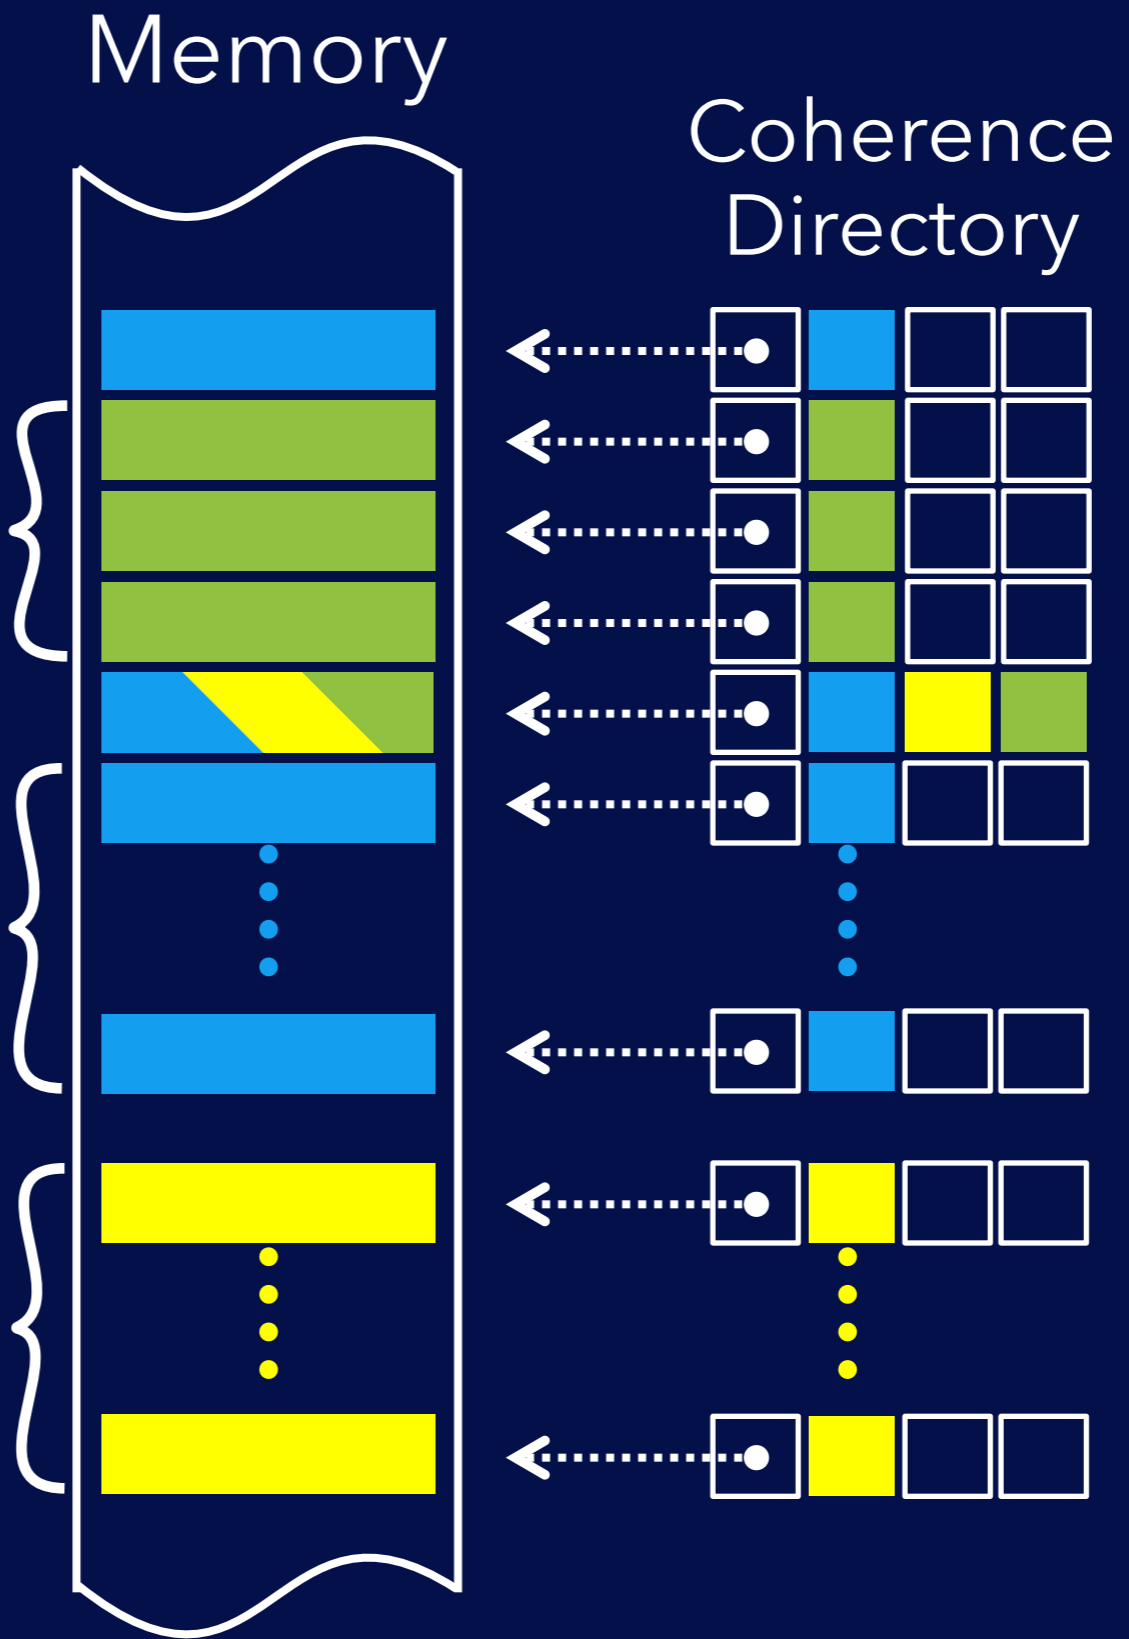


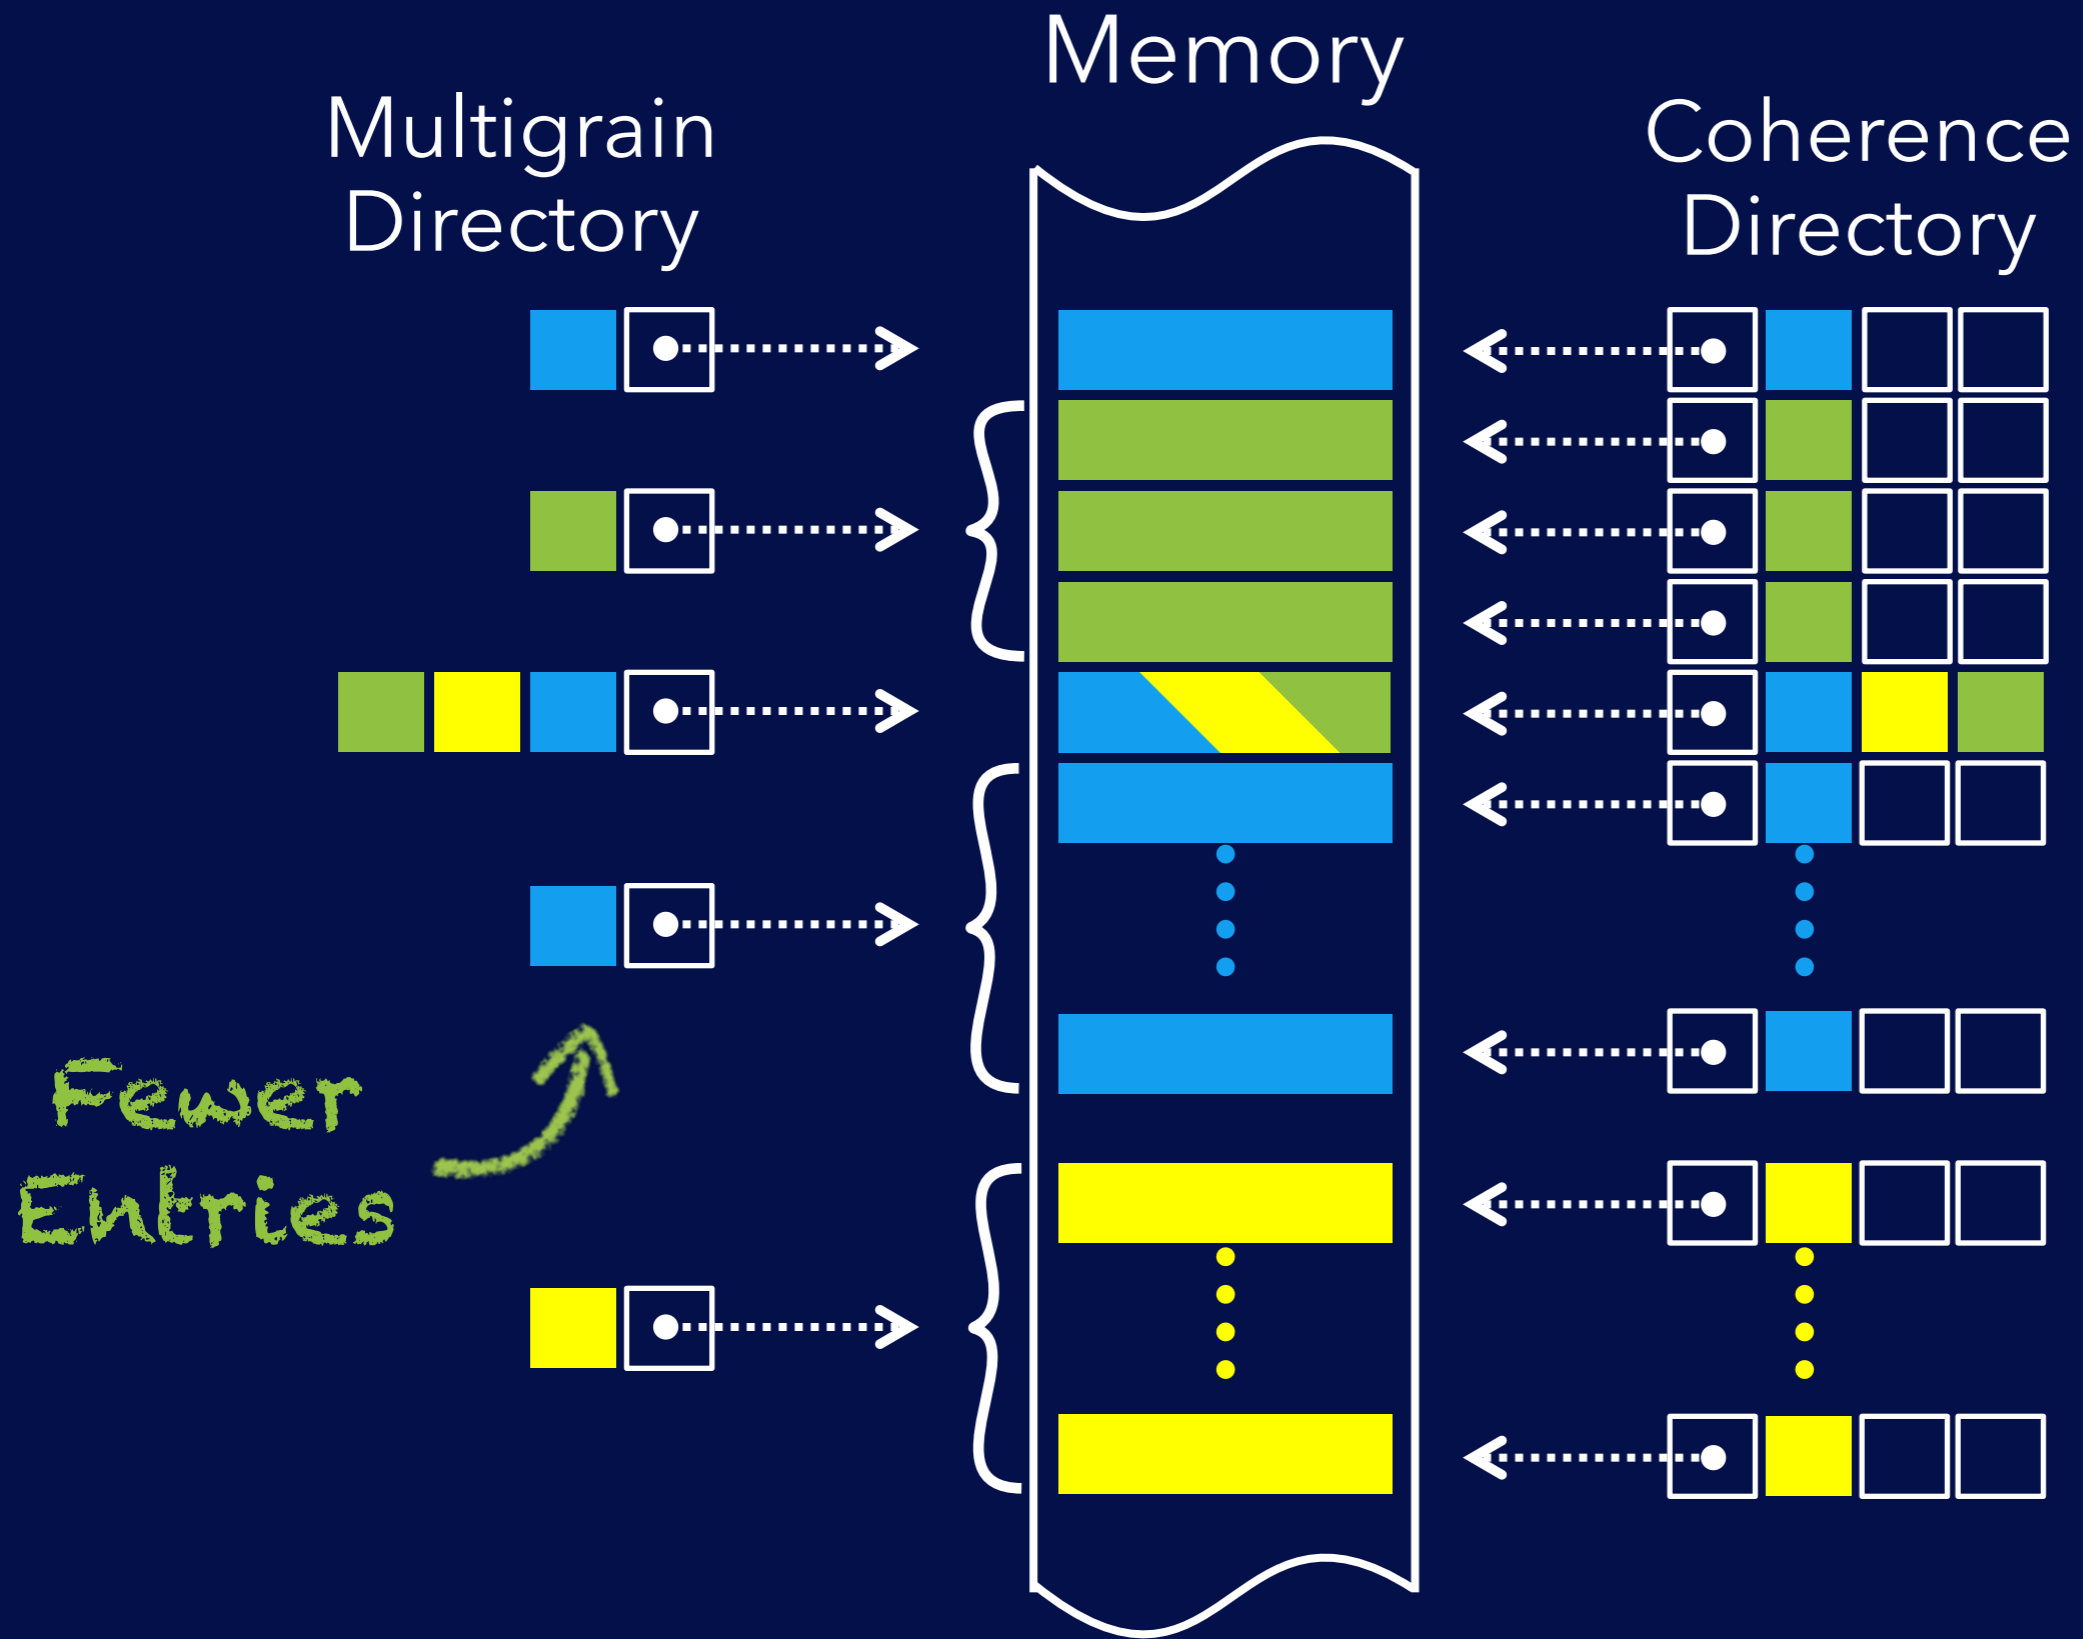
Memory




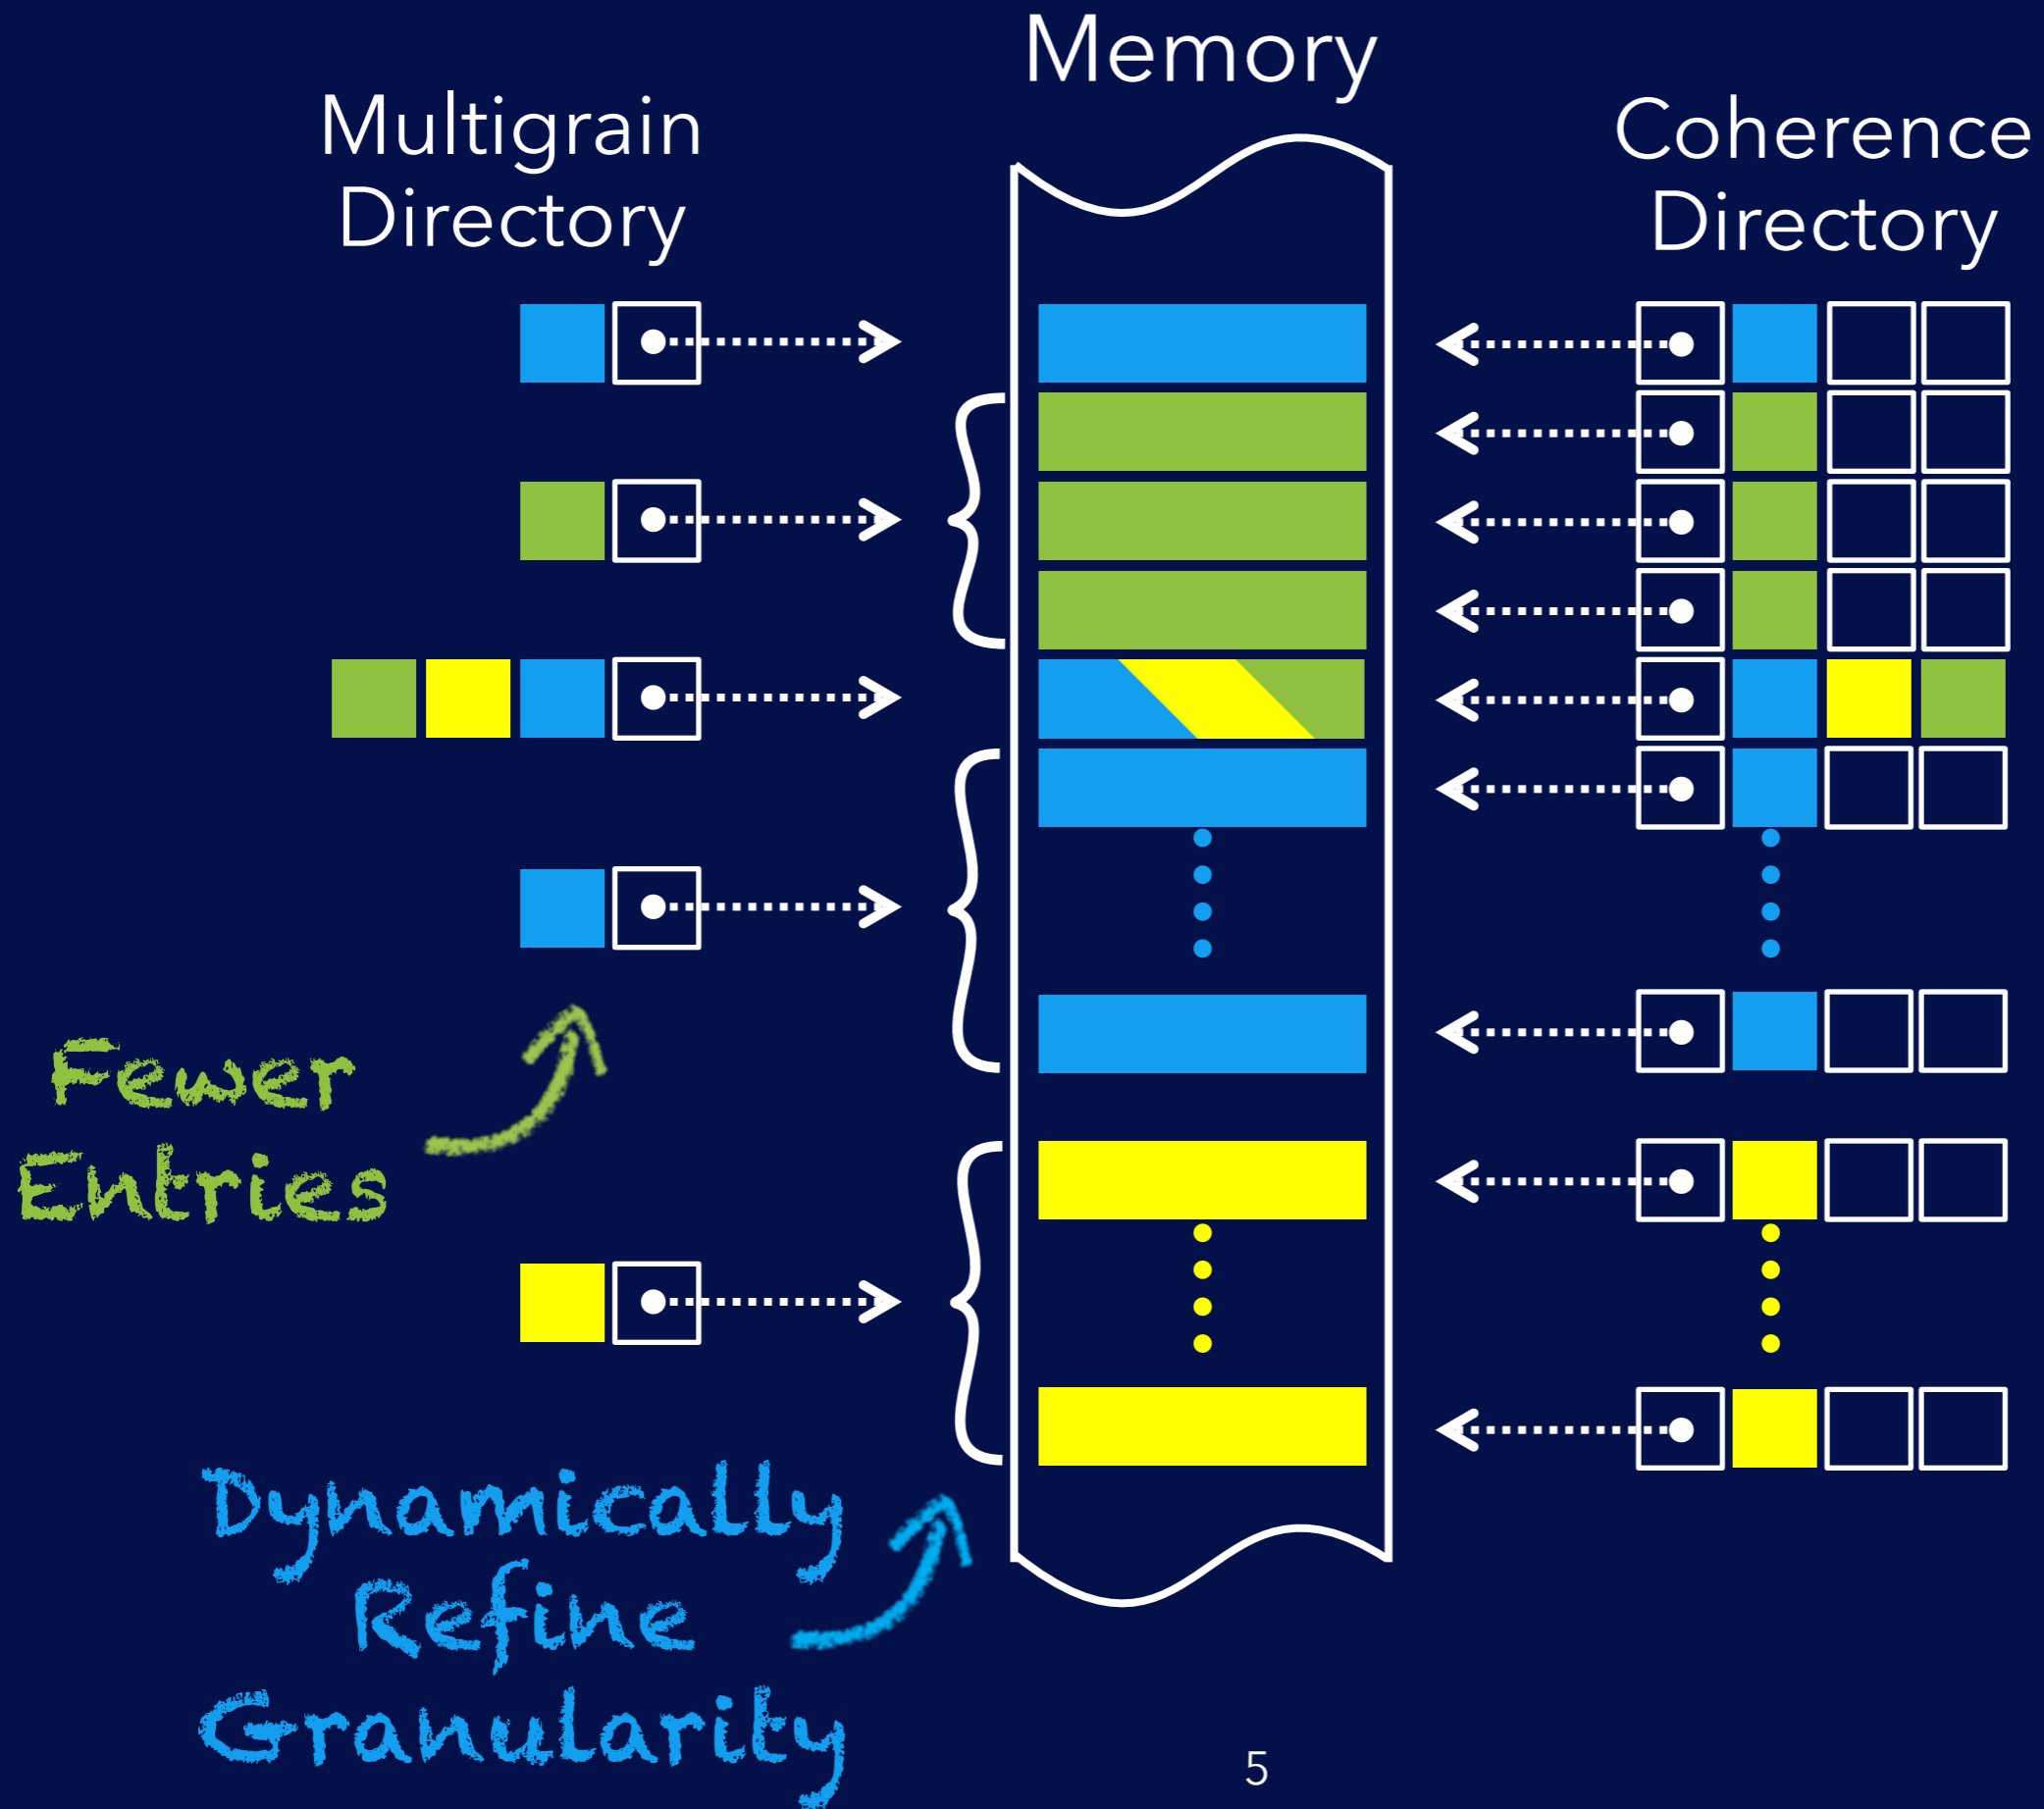
Coherence Directory











Multi-Grain Coherence Directory (MGD)

Conceptual MGD Directory:

- ✓ Dynamically refine granularity of entries
- ✓ 78% fewer directory entries (on average)

Practical MGD Directory:

- ✓ Limited number of fixed granularities
- ✓ 41% less area
- ✓ No coherence protocol changes
- ✓ Robust performance


- ▶ **Original MGD concept**
- ▶ Potential benefits of MGD
- ▶ Making a practical design
- ▶ A recent competing design
- ▶ Some nice graphs
- ▶ Final Thoughts

Core 1


Core 2

Core 3

Memory



Each entry tracks a private region of memory




Core 1

Core 2

Core 3

Memory





*Dynamically adapt
region size for
each entry*

Core 1


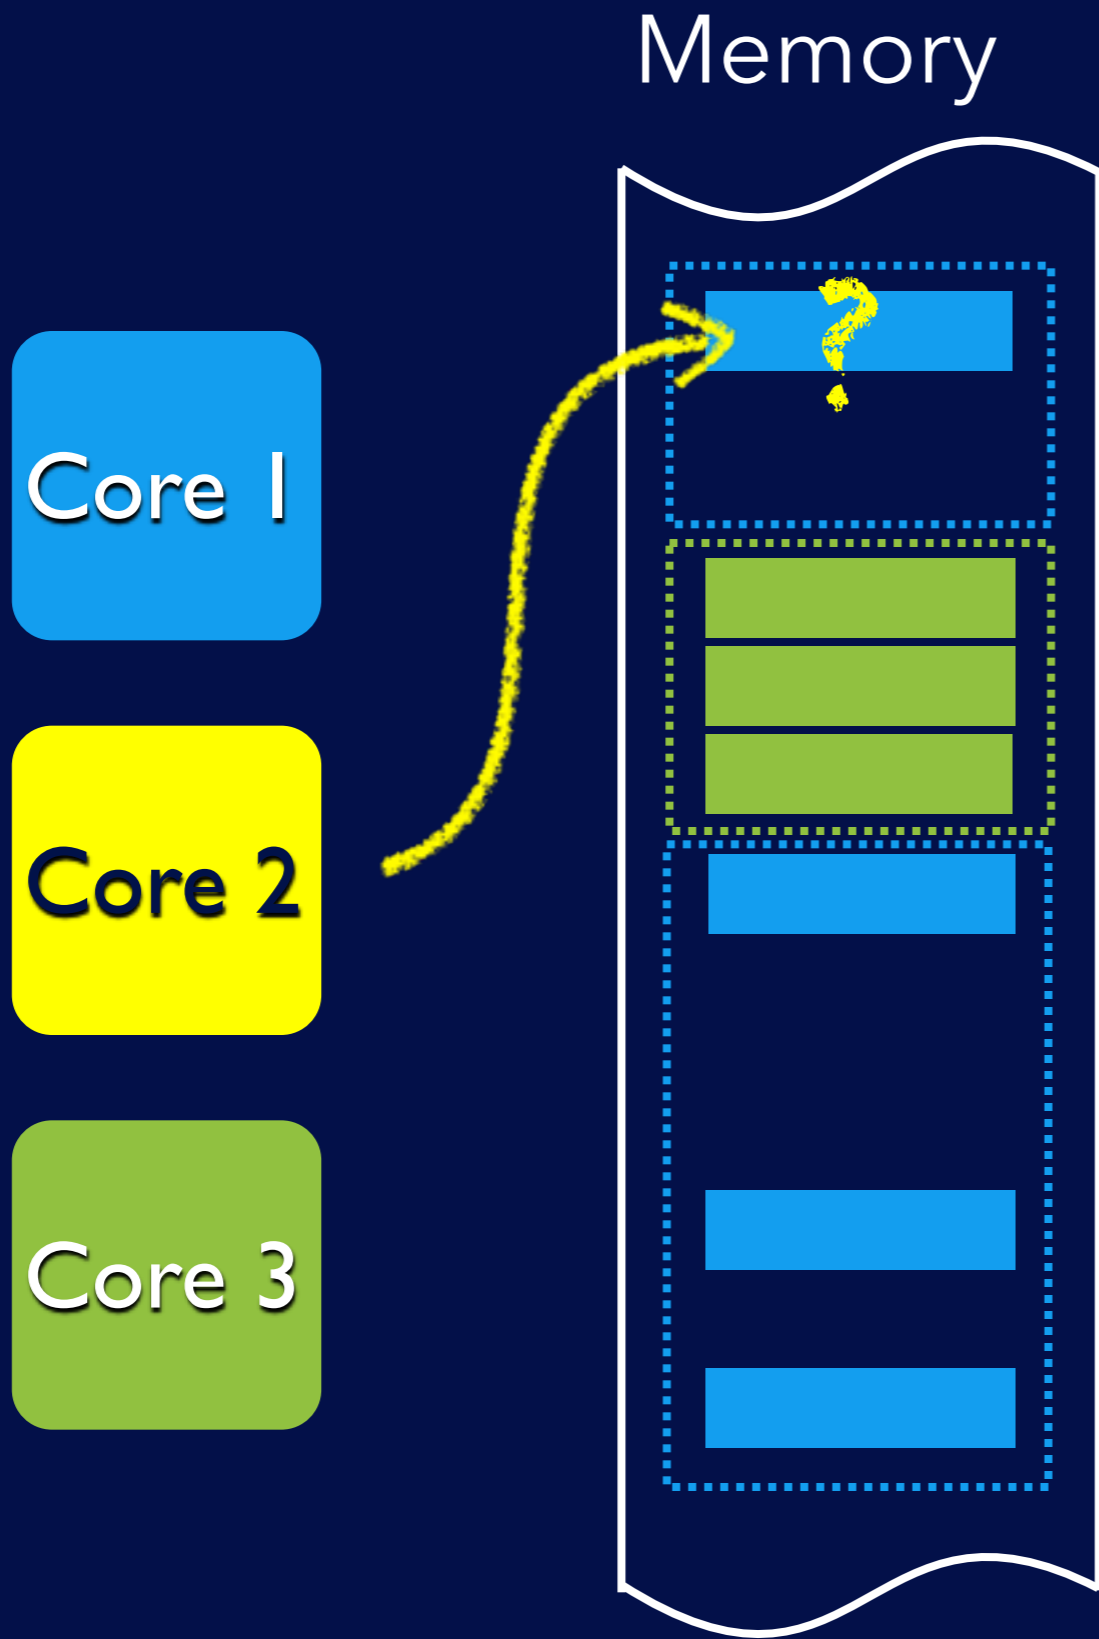
Core 2

Core 3

Memory



*Dynamically adapt
region size for
each entry*




*Dynamically adapt
region size for
each entry*

Core 1

Core 2

Core 3

Memory




single entry for each shared block

Conceptual Multi-Grain Directory

- ❖ **Thought experiment, not a real design**
- ❖ Entries track *instantaneous* private regions
- ❖ Dynamically refine region granularity
 - ❖ *Maximize size of each private region*
 - ❖ *Limit to powers of two, properly aligned*
- ❖ Shared blocks assigned individual entries

- ▶ Original MGD concept
- ▶ **Potential benefits of MGD**
- ▶ Making a practical design
- ▶ A recent competing design
- ▶ Some nice graphs
- ▶ Final Thoughts

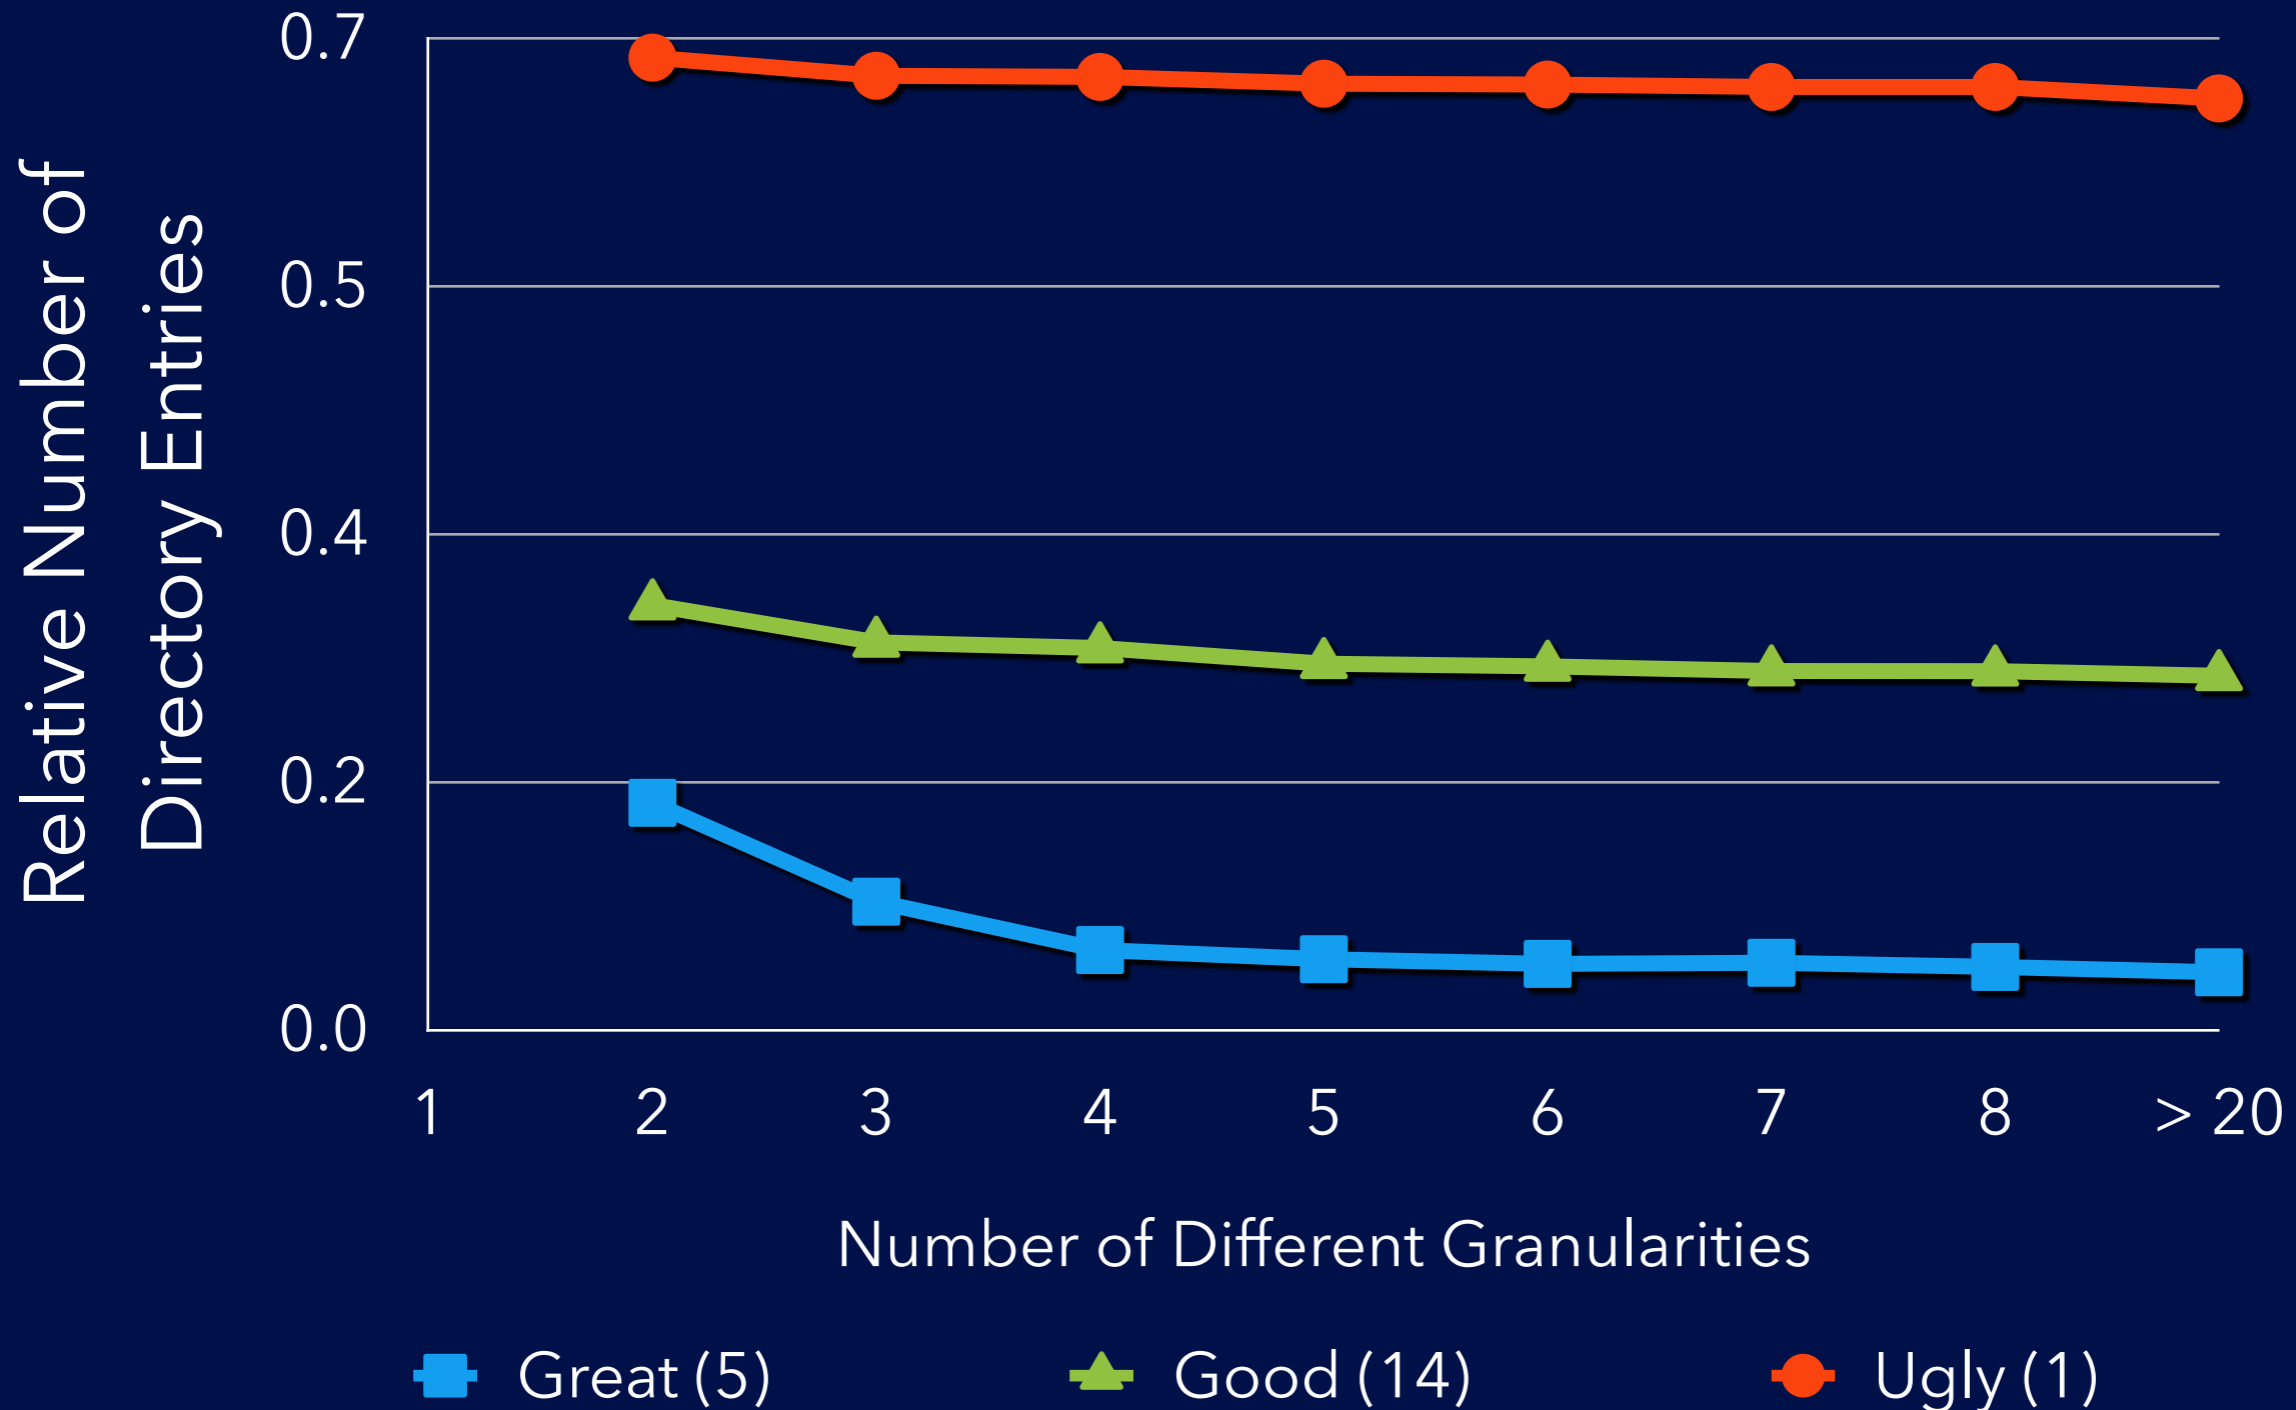
Conceptual MGD Potential



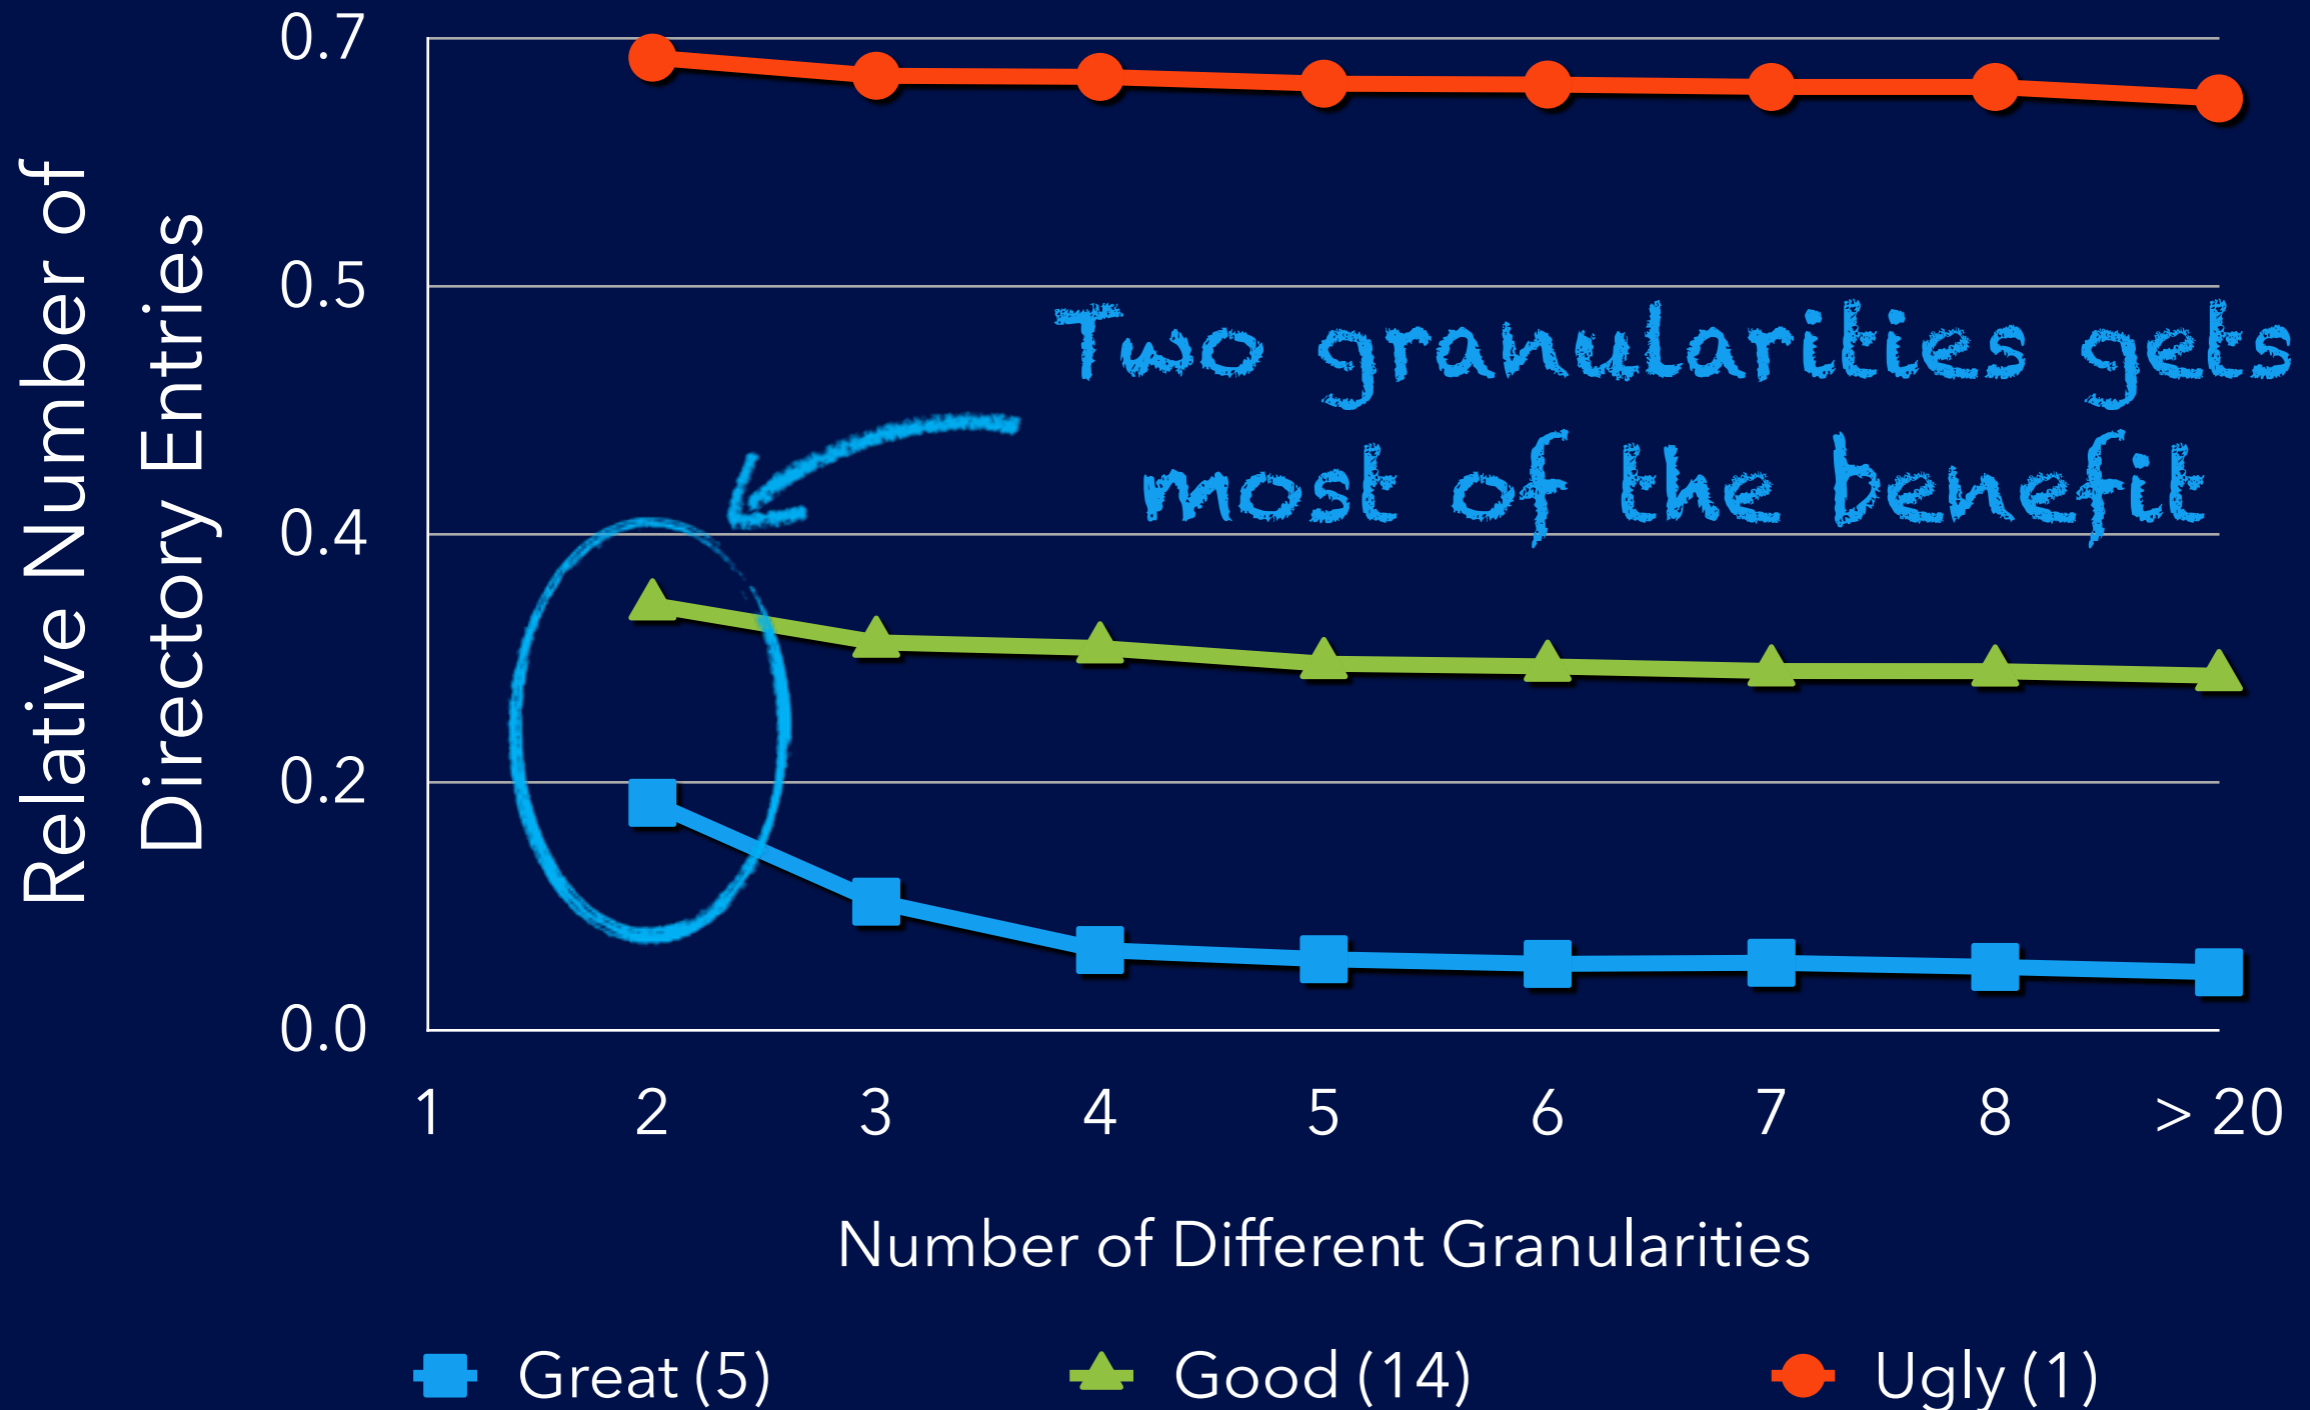
Conceptual MGD Potential




Conceptual MGD Potential




Conceptual MGD Potential




Conceptual MGD Potential



Conceptual MGD Potential

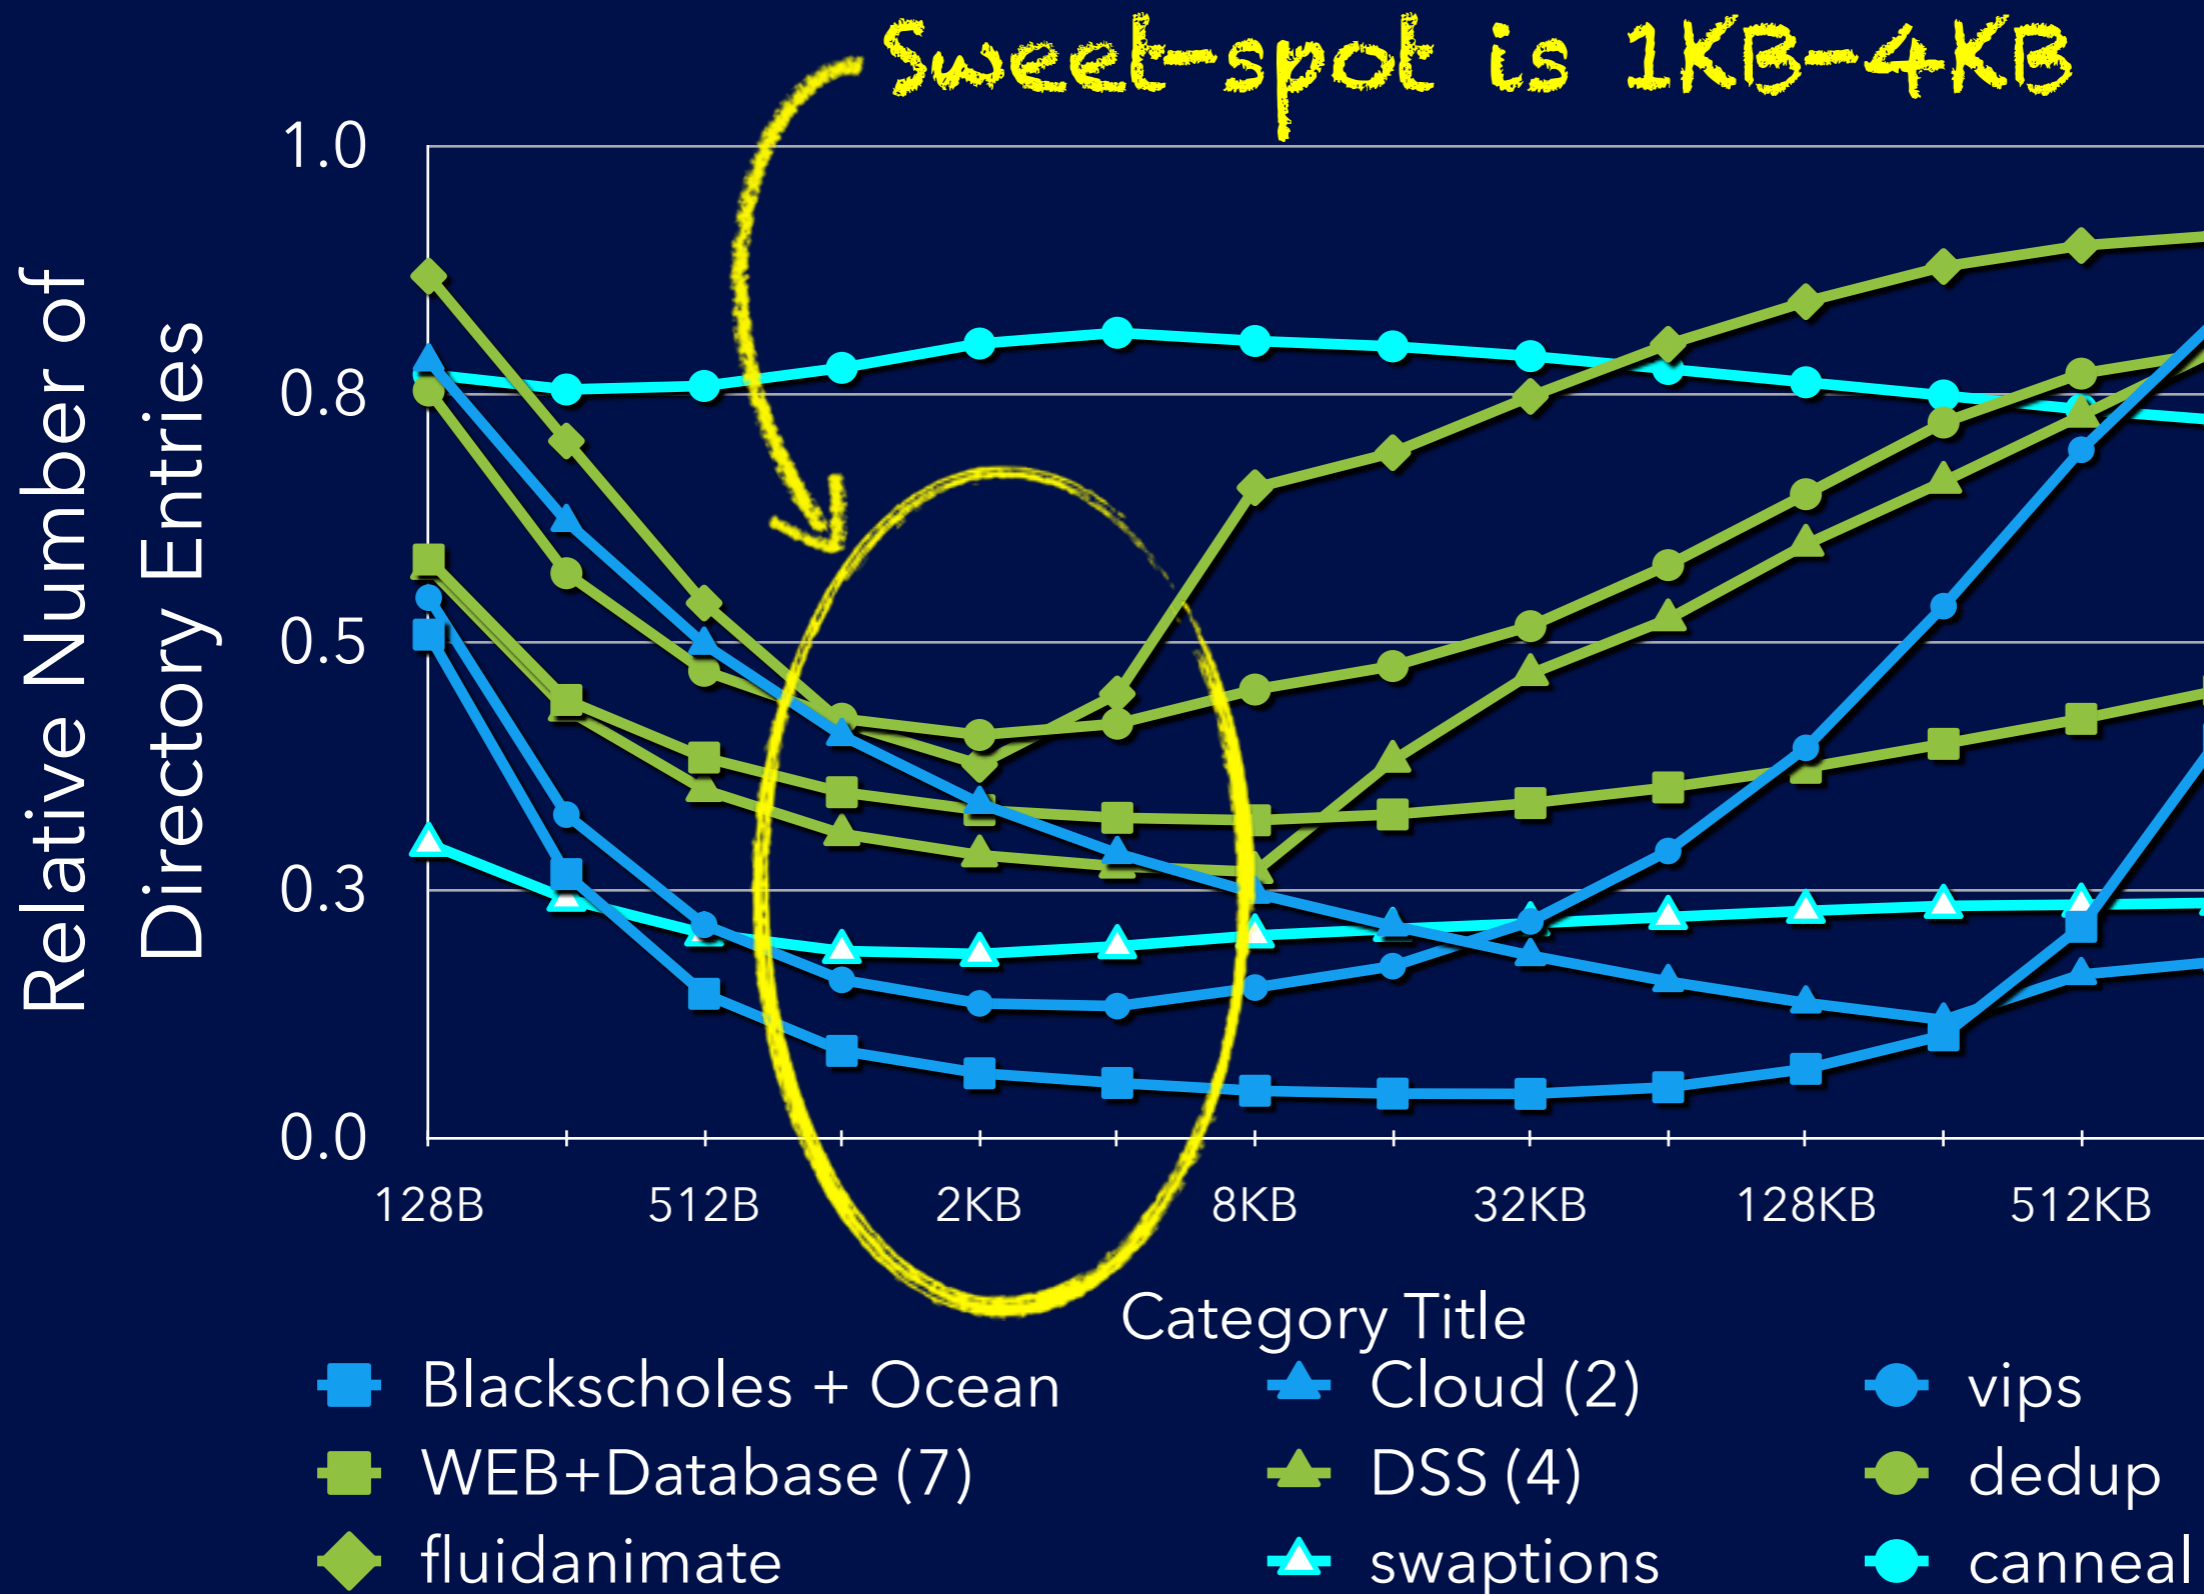


Region Size for Dual-Grain MGD



- Category Title
- Blackscholes + Ocean
 - WEB+Database (7)
 - fluidanimate
 - Cloud (2)
 - DSS (4)
 - swaptions
 - vips
 - dedup
 - canneal

Region Size for Dual-Grain MGD



- ▶ Original MGD concept
- ▶ Potential benefits of MGD
- ▶ **Making a practical design**
- ▶ A recent competing design
- ▶ Some nice graphs
- ▶ Final Thoughts

Practical Multi-Grain Directory

A Dual-Grain Directory (DGD)

- ❖ **Real design, not just a thought experiment**
- ❖ Track instantaneous private regions
- ❖ Use two fixed granularities
 - ✦ *individual cache blocks + regions of 1KB - 4KB*
- ❖ ***How will this work?***

What do Entries Look Like?



What do Entries Look Like?



Bit vector of blocks
cached by region owner

What do Entries Look Like?

Both entry types are the same size



← 32 × 16 →

Block Entry:



Region Entry:



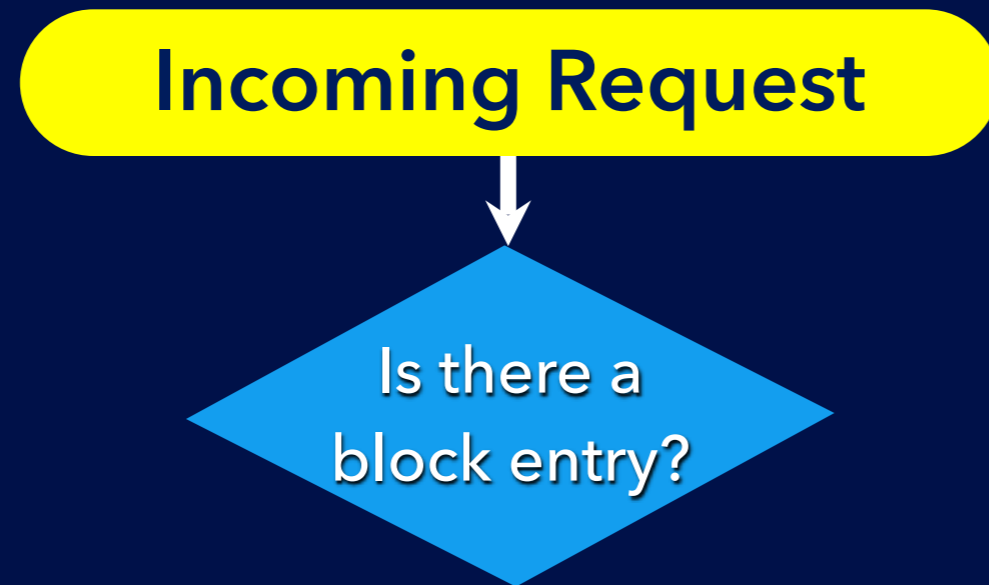
← 28 × 4 × 16 →

Bit vector of blocks
cached by region owner

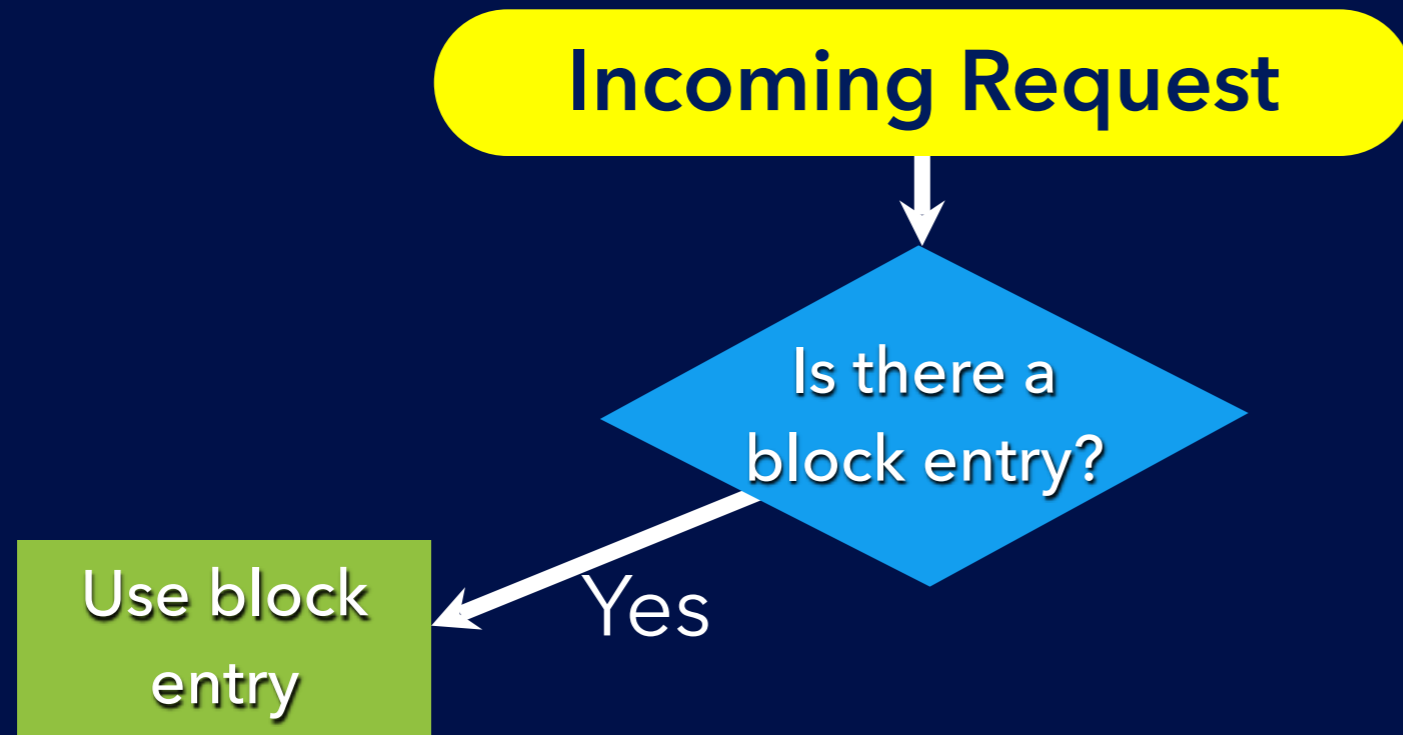
Dual-Grain Directory (DGD) Function

Incoming Request

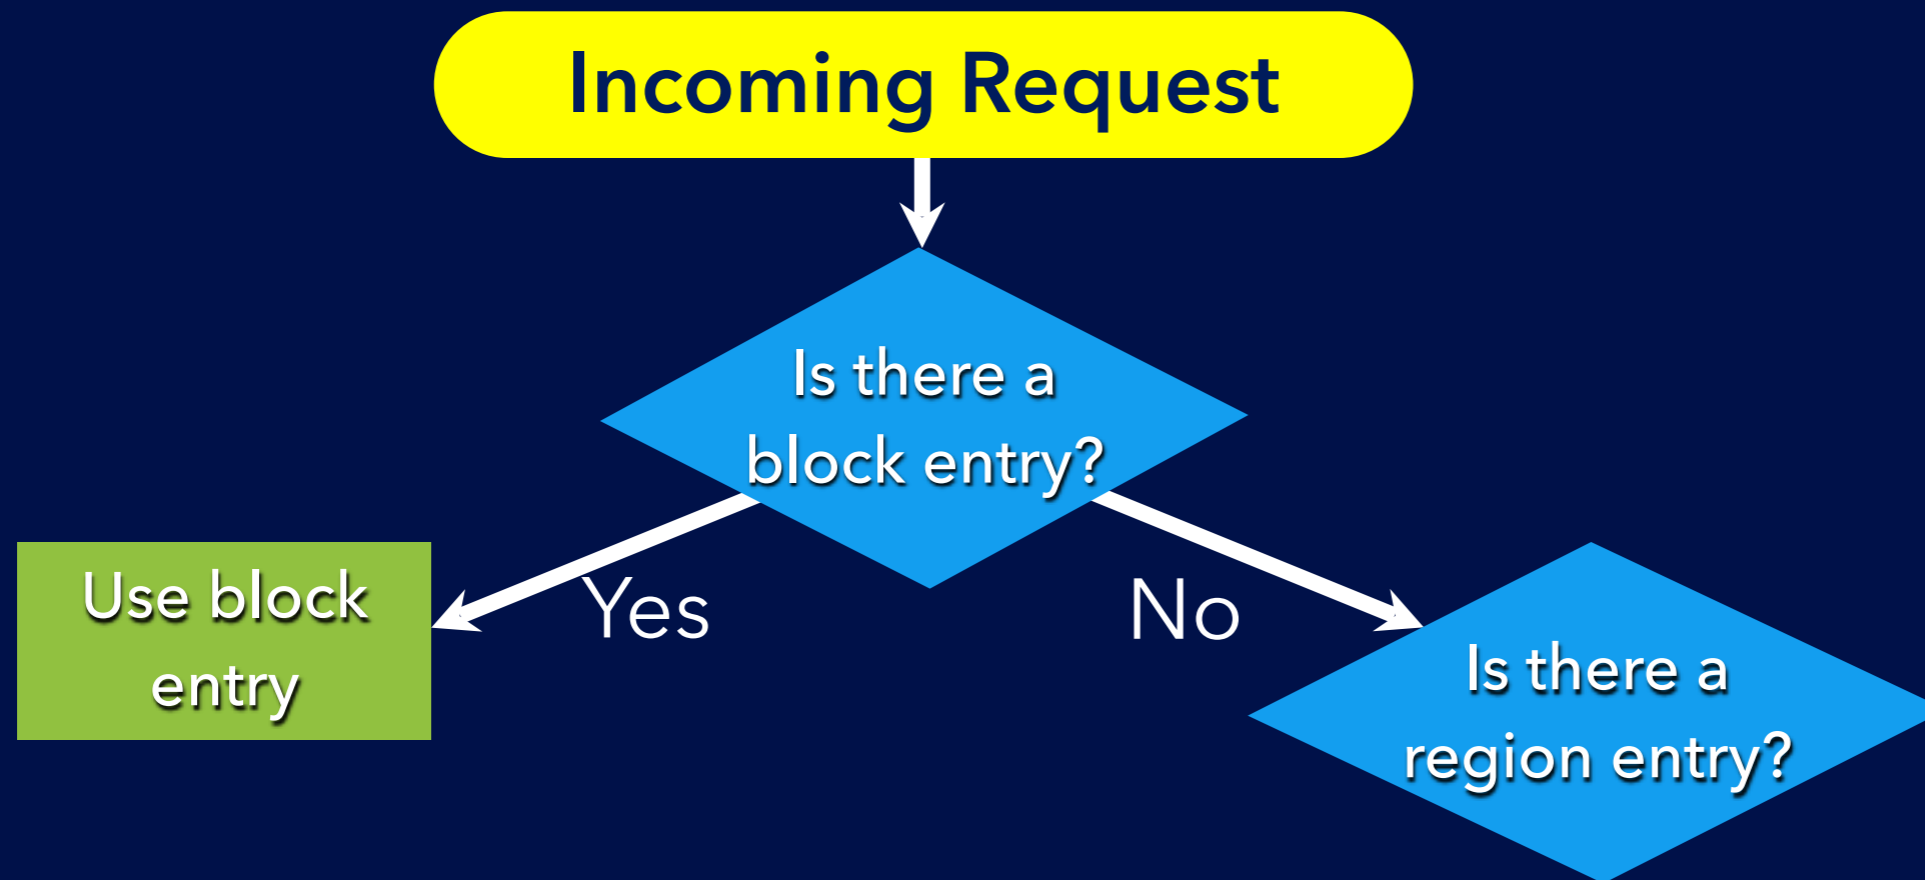
Dual-Grain Directory (DGD) Function



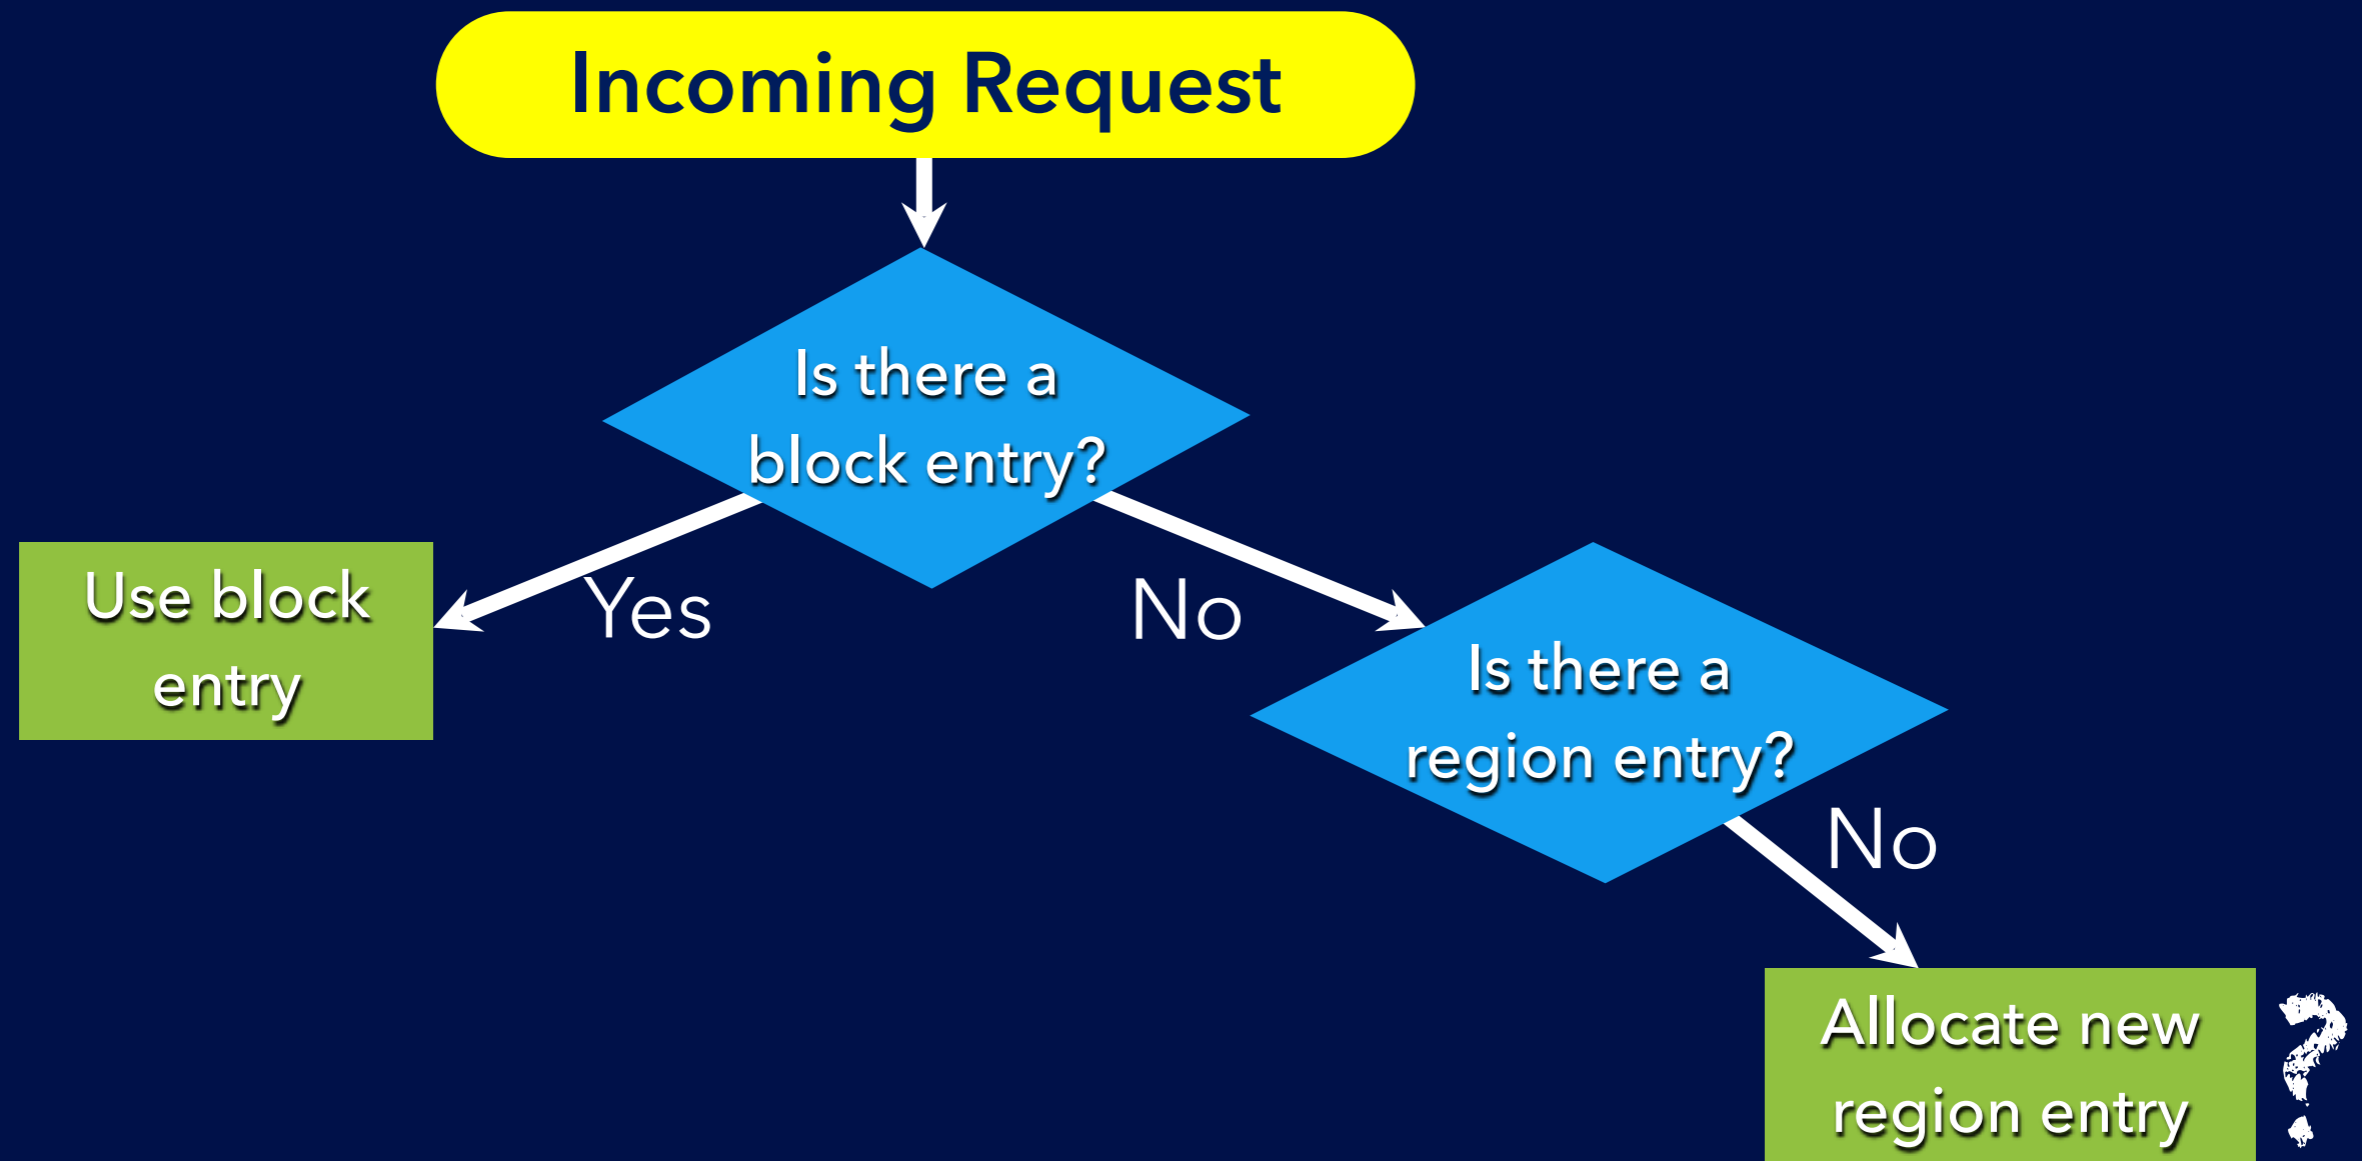
Dual-Grain Directory (DGD) Function



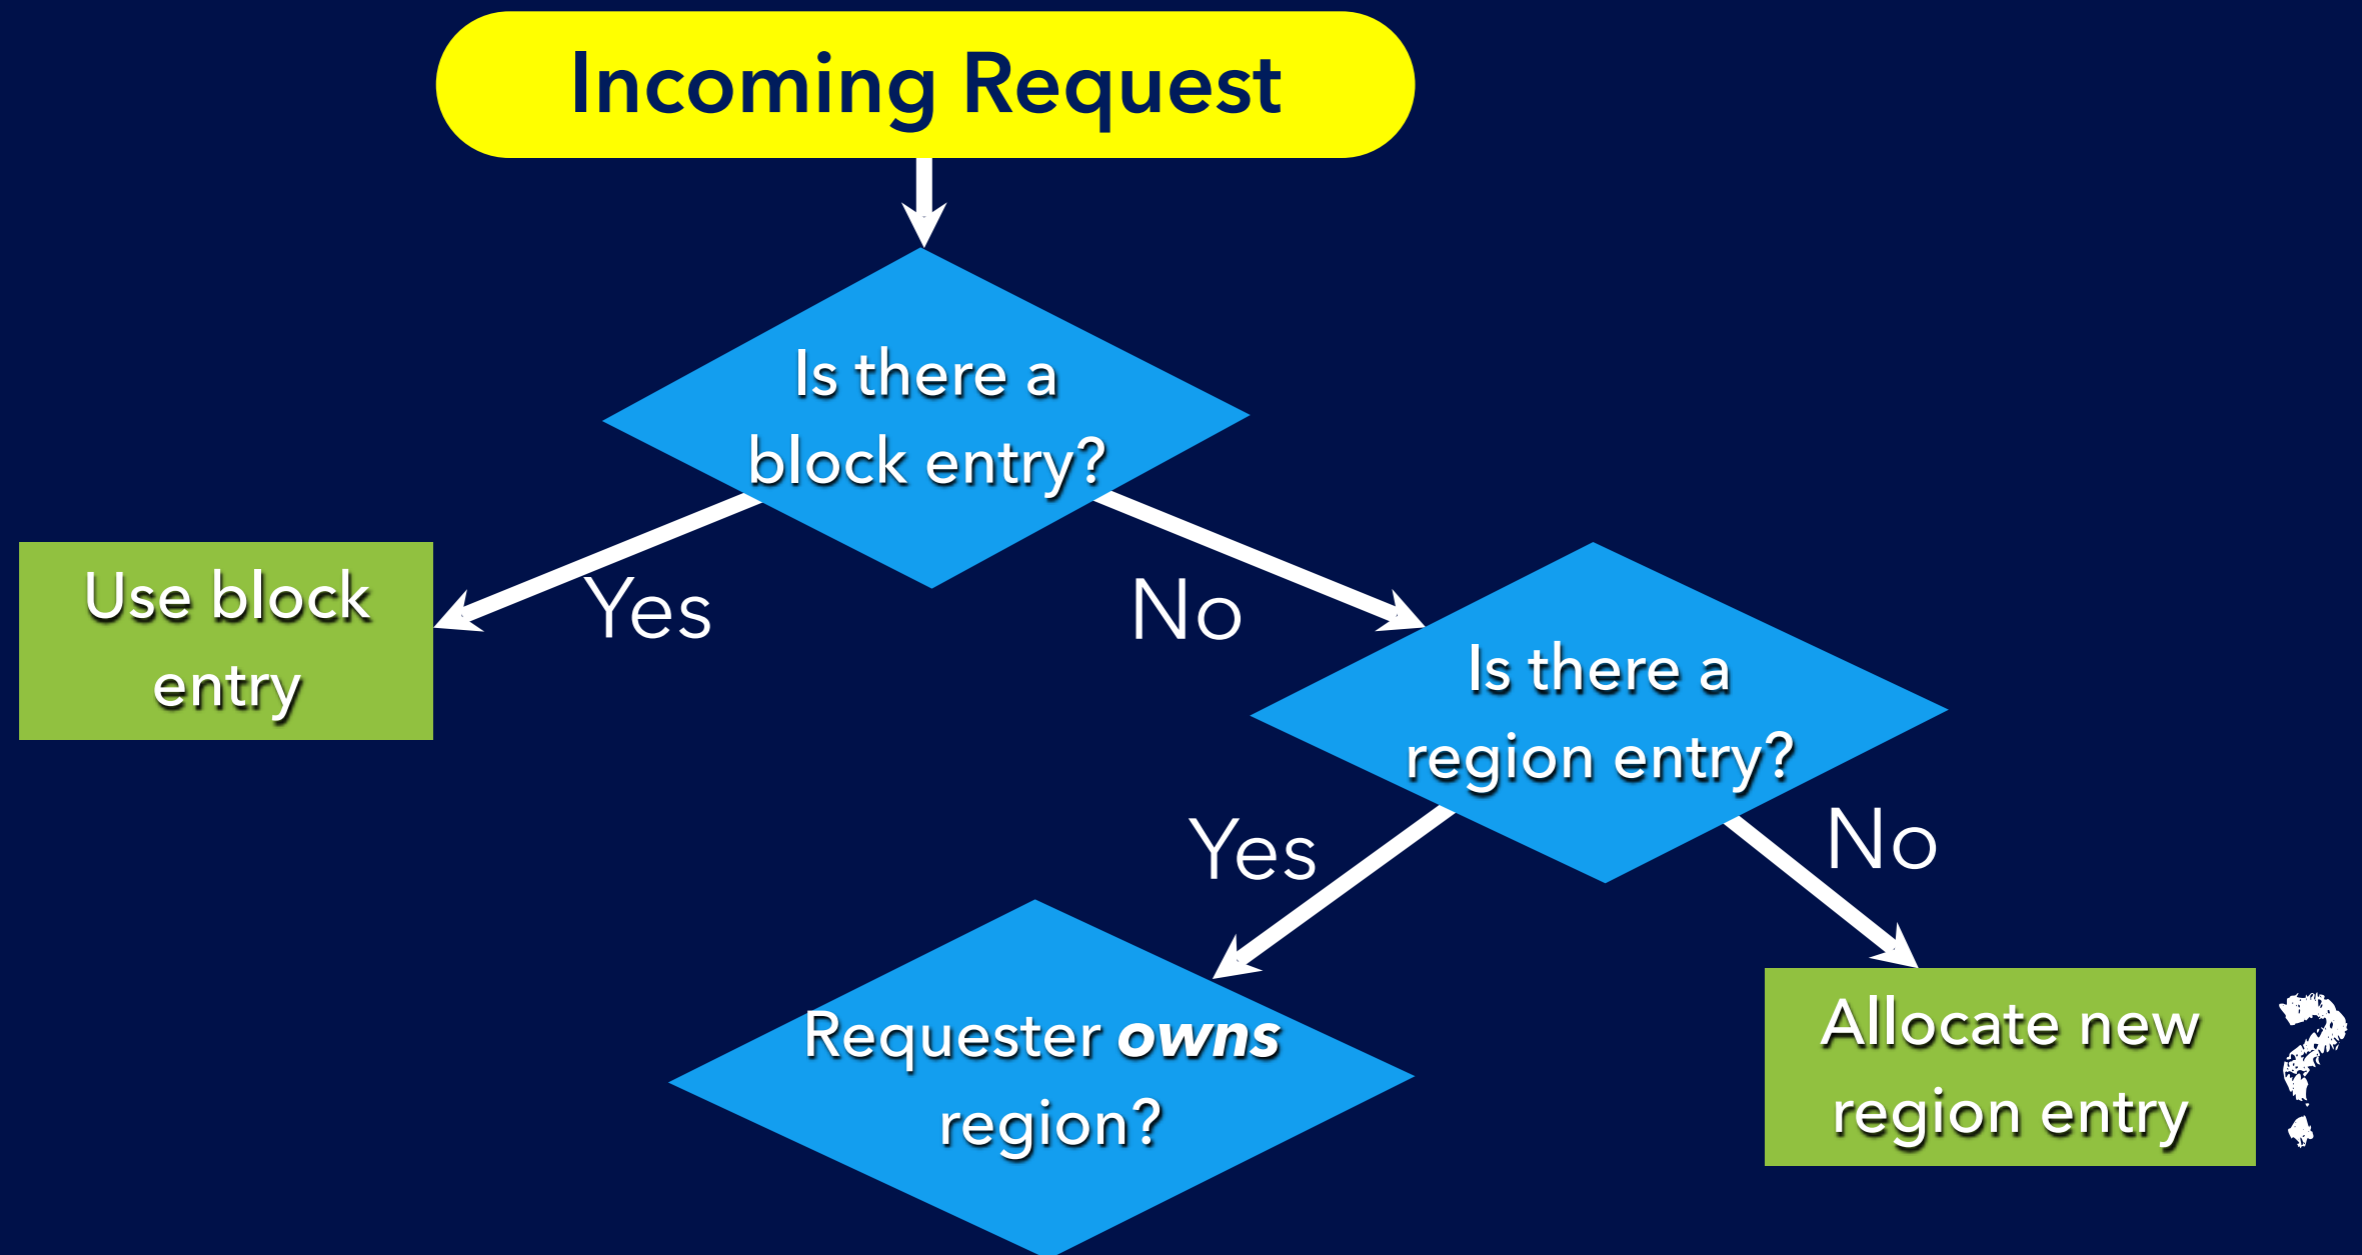
Dual-Grain Directory (DGD) Function



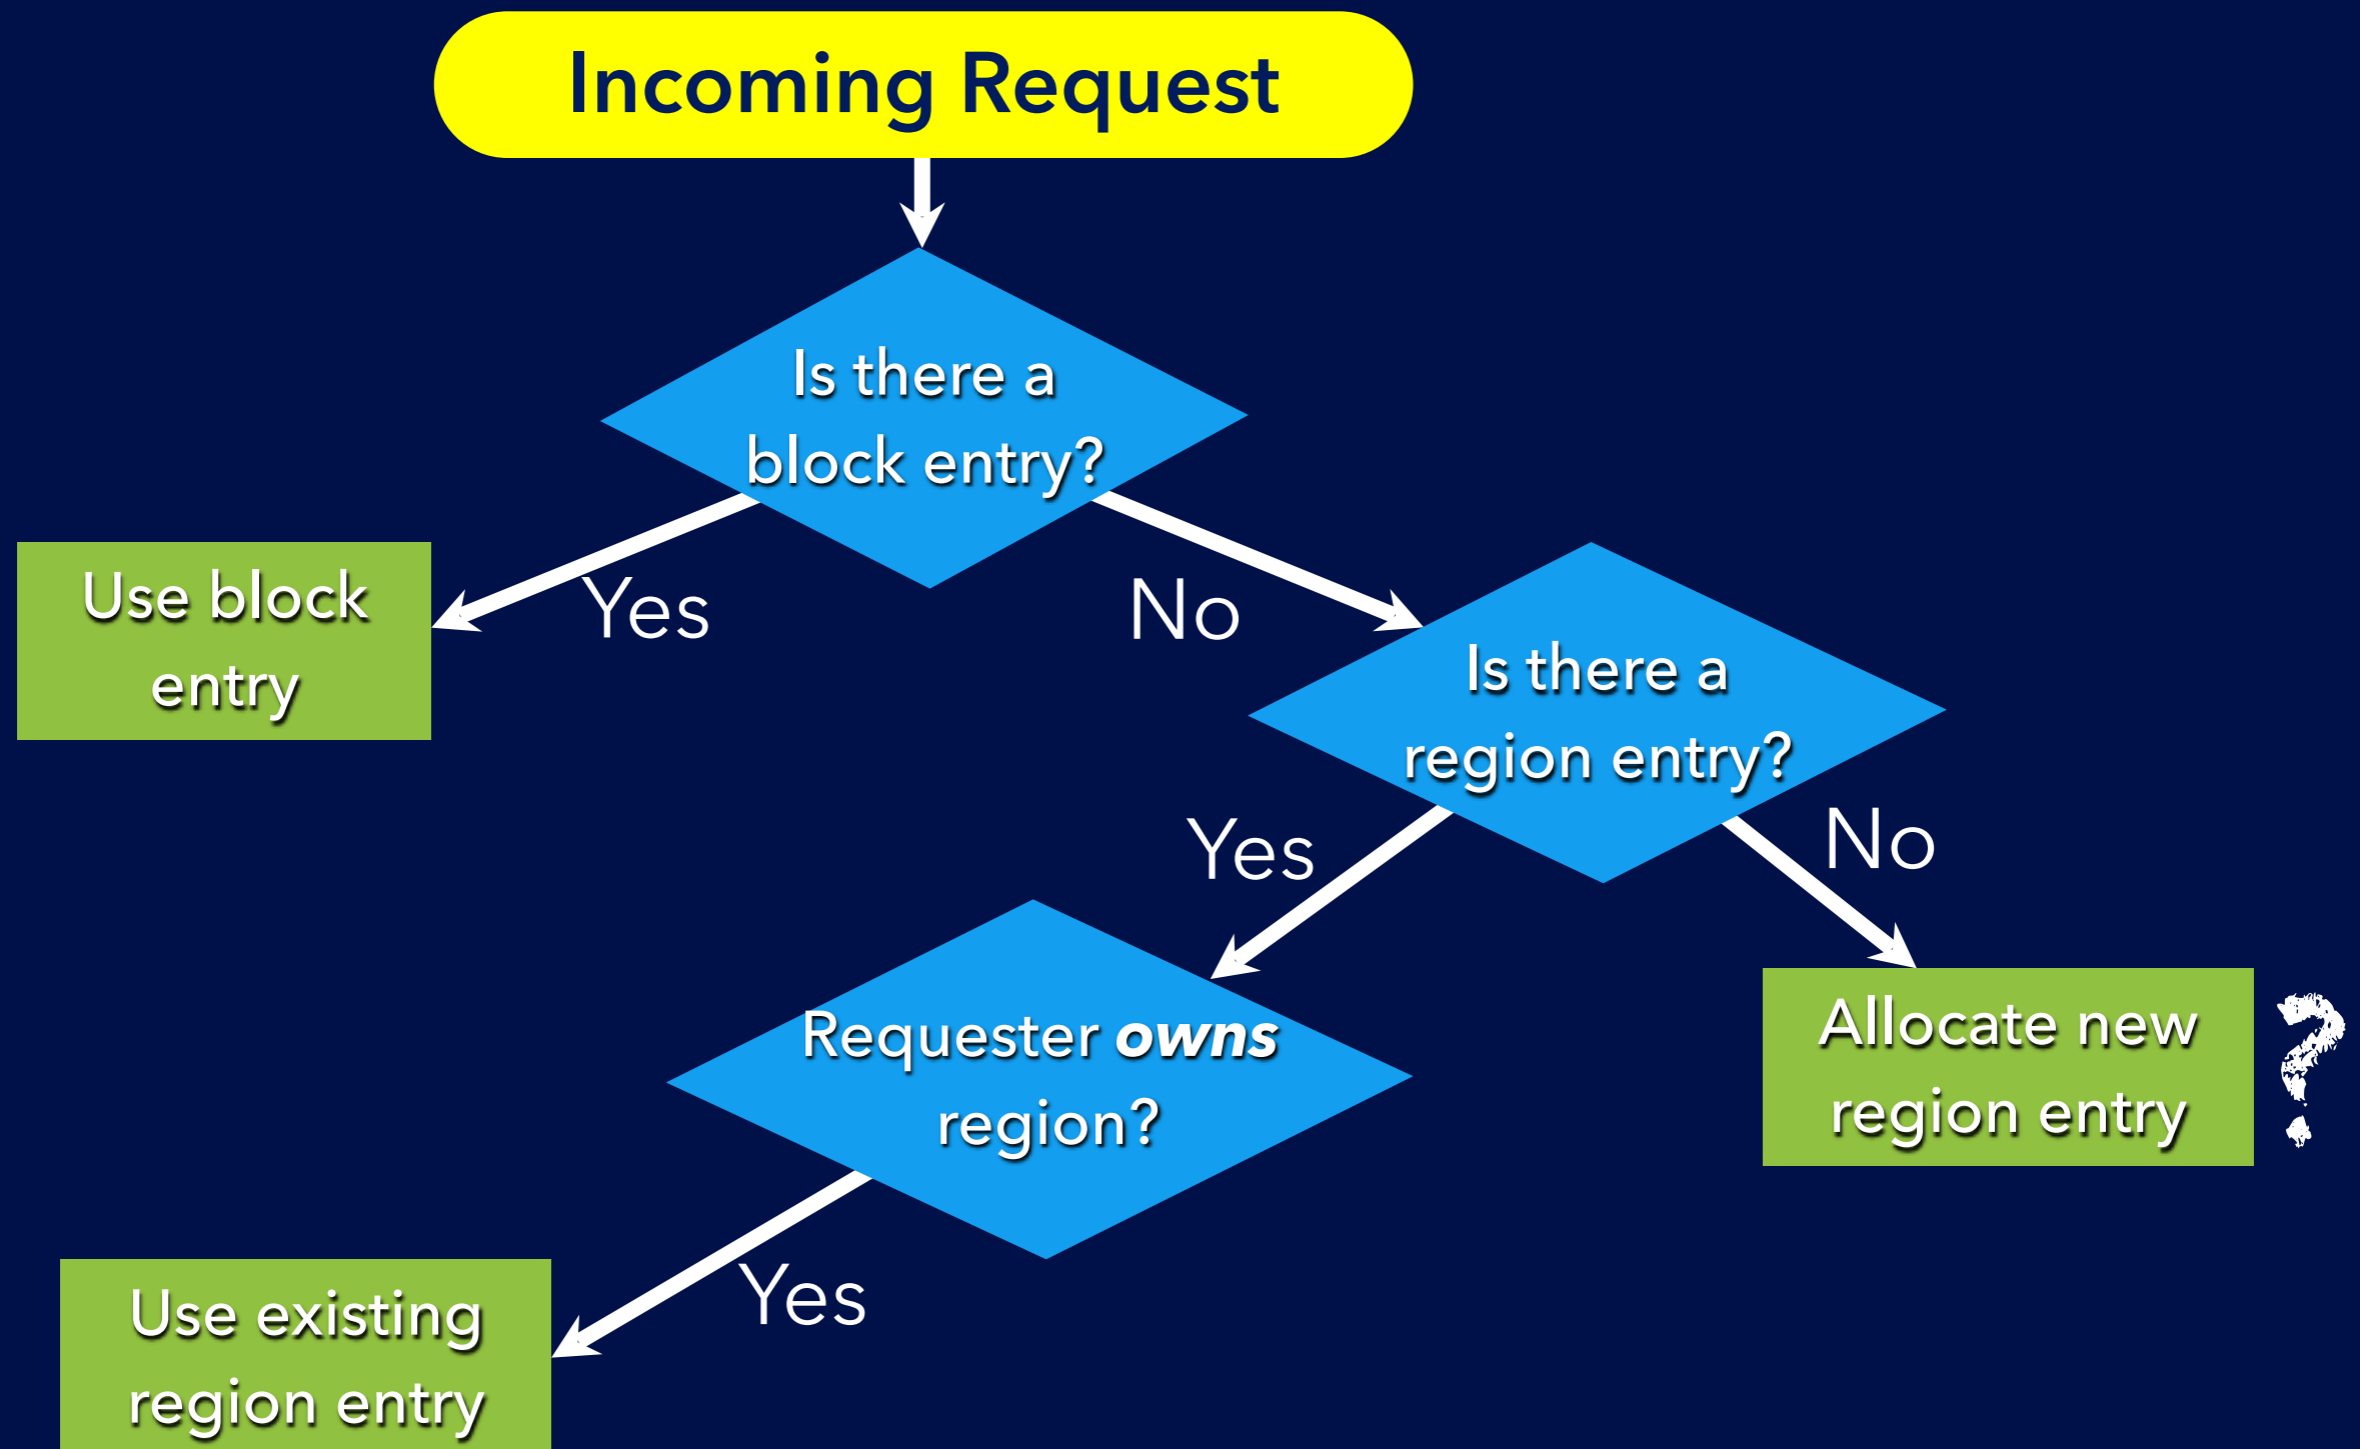
Dual-Grain Directory (DGD) Function



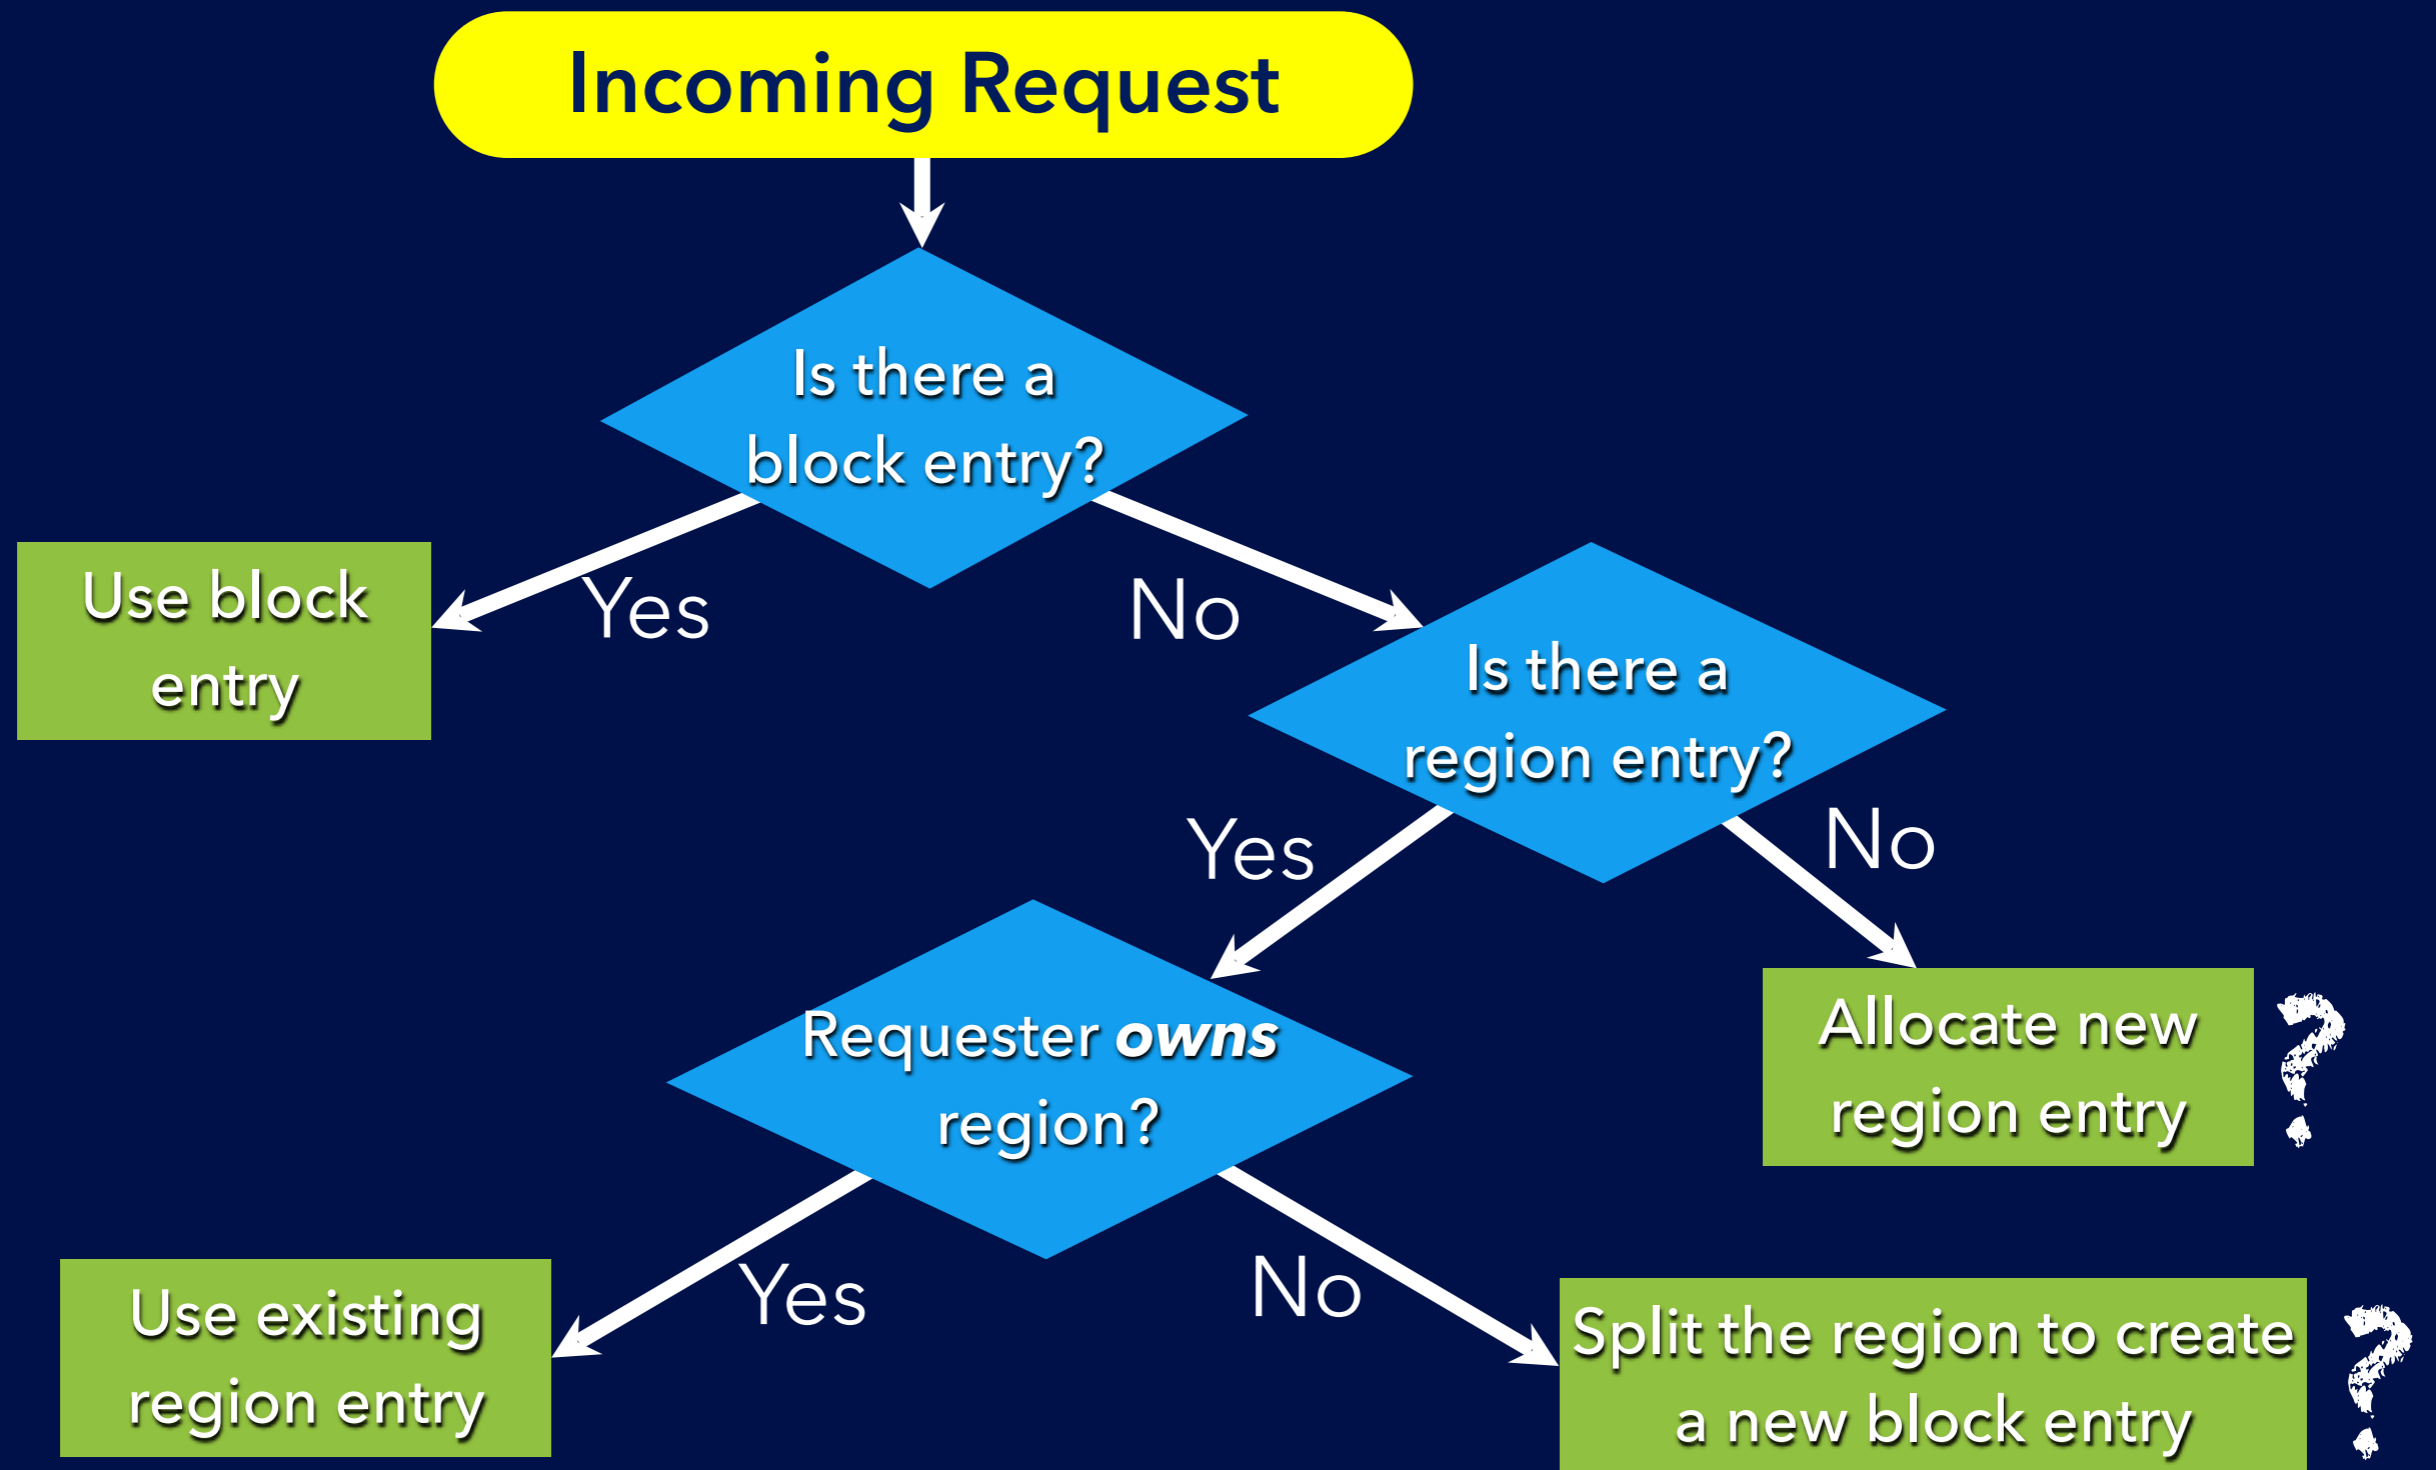
Dual-Grain Directory (DGD) Function

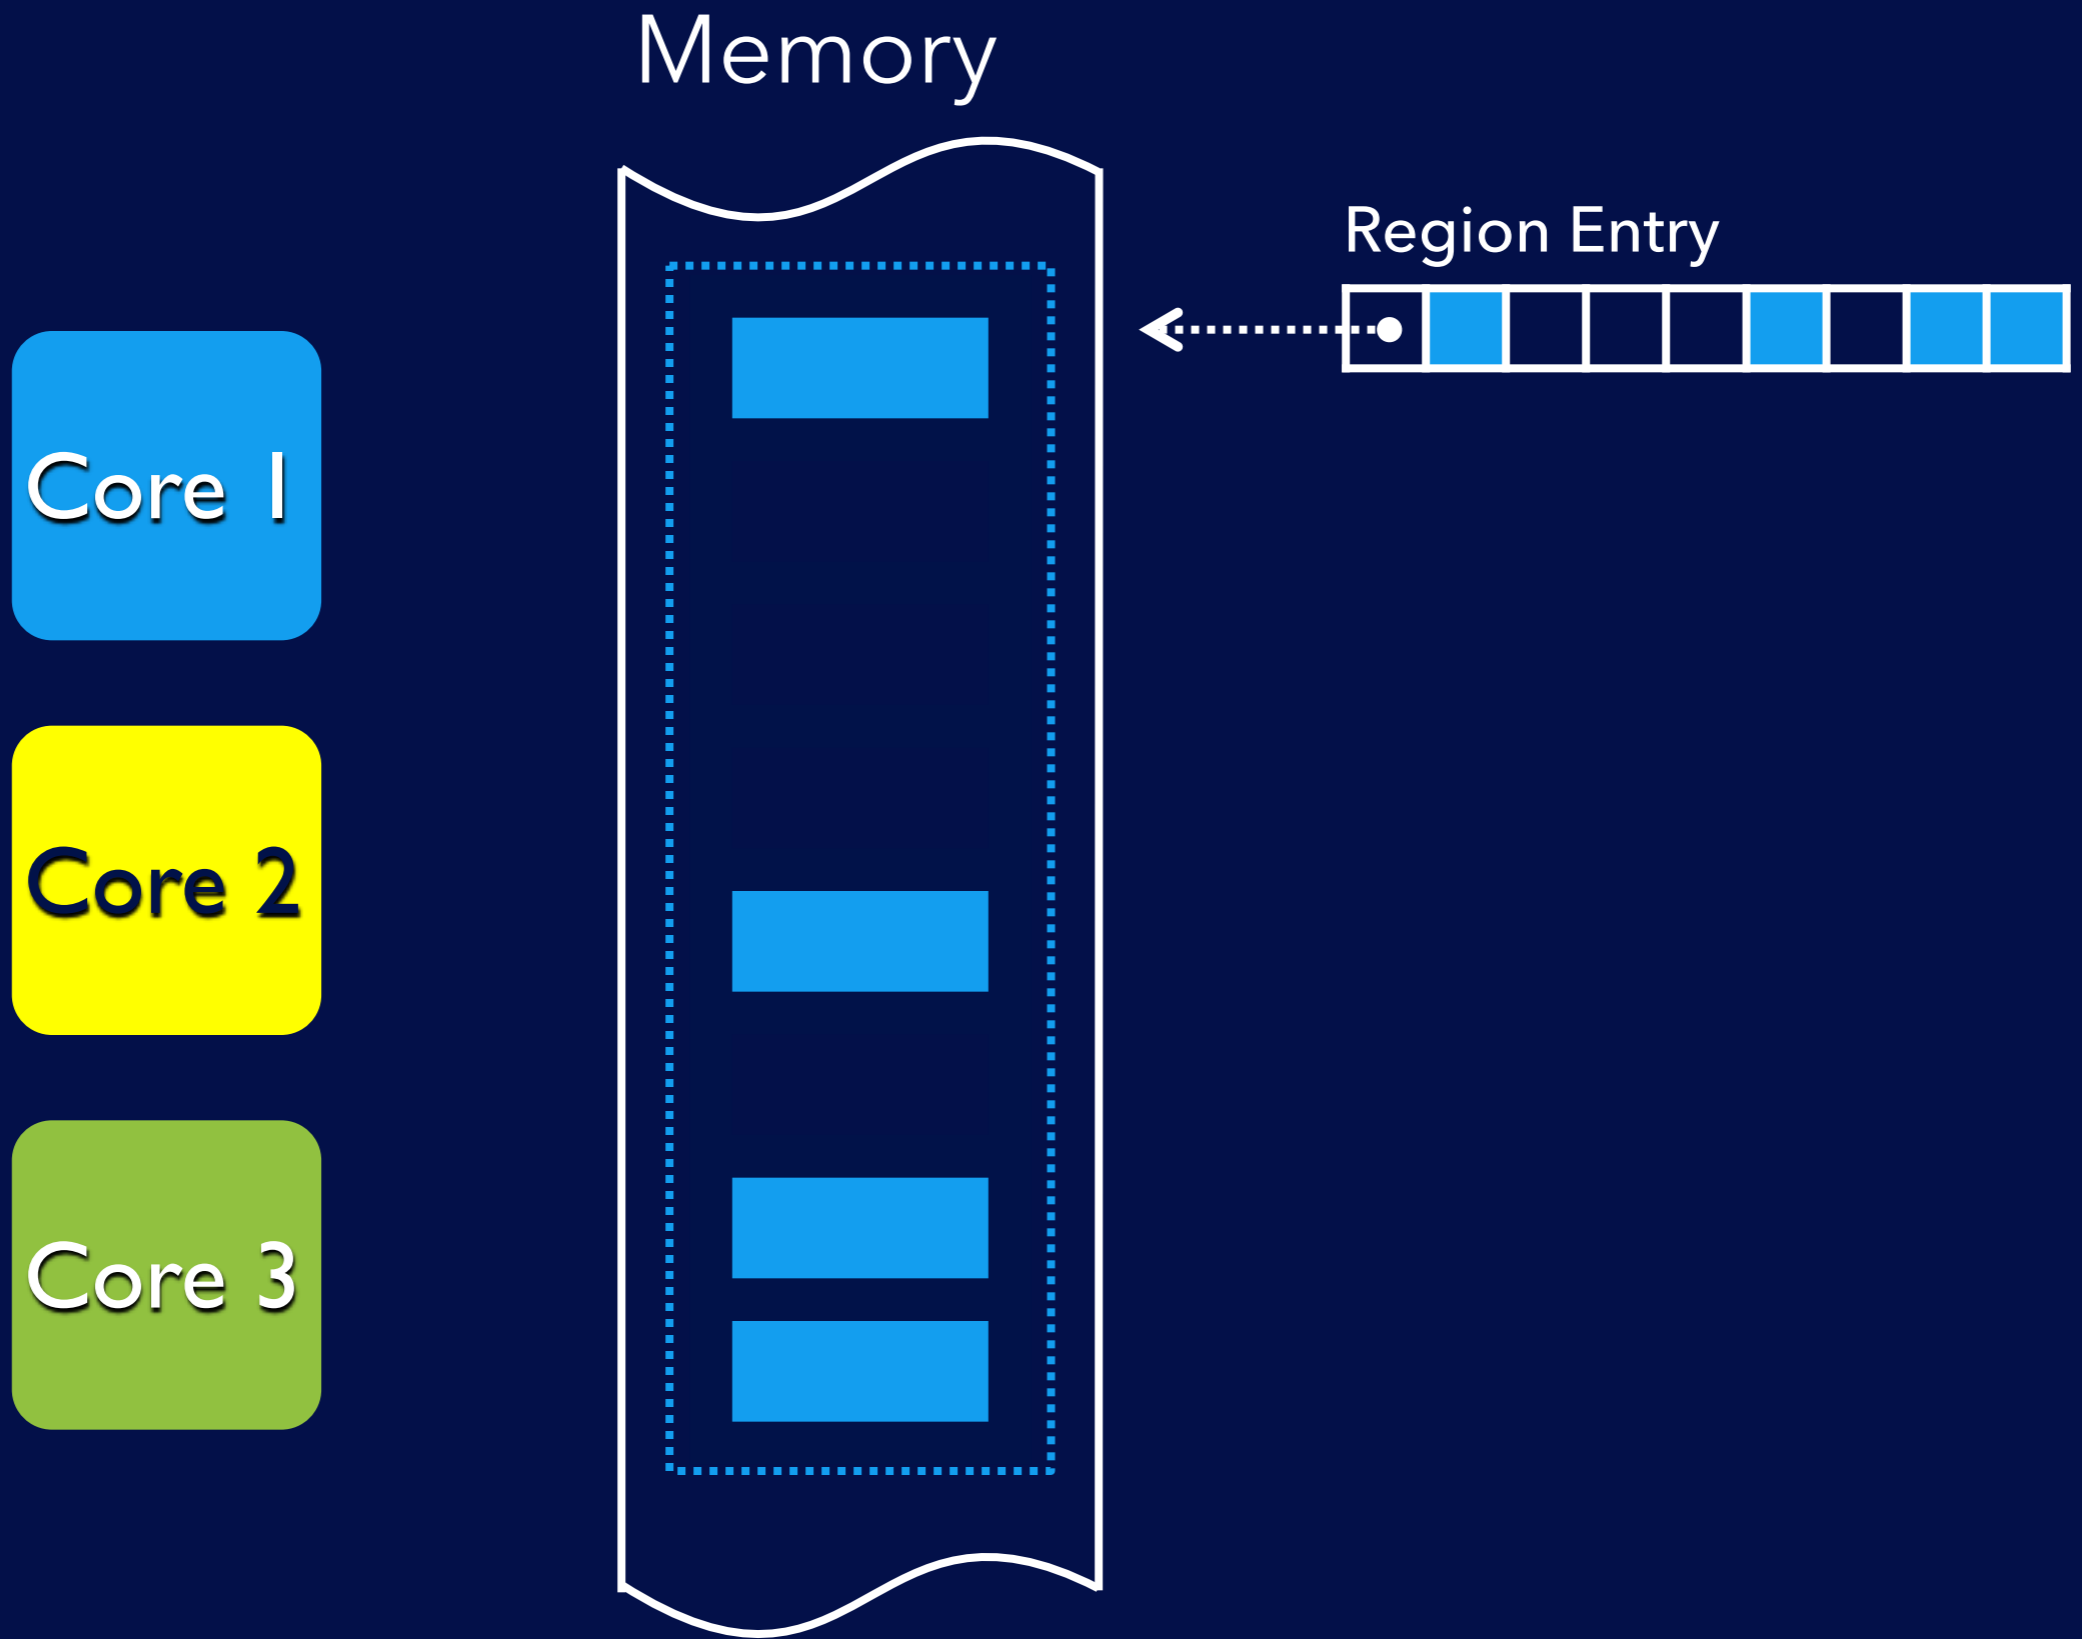


Dual-Grain Directory (DGD) Function



Dual-Grain Directory (DGD) Function






Memory

Core 1

Core 2

Core 3



Region Entry




How do we split a region entry?

Core 1

Core 2

Core 3

Memory



Region Entry

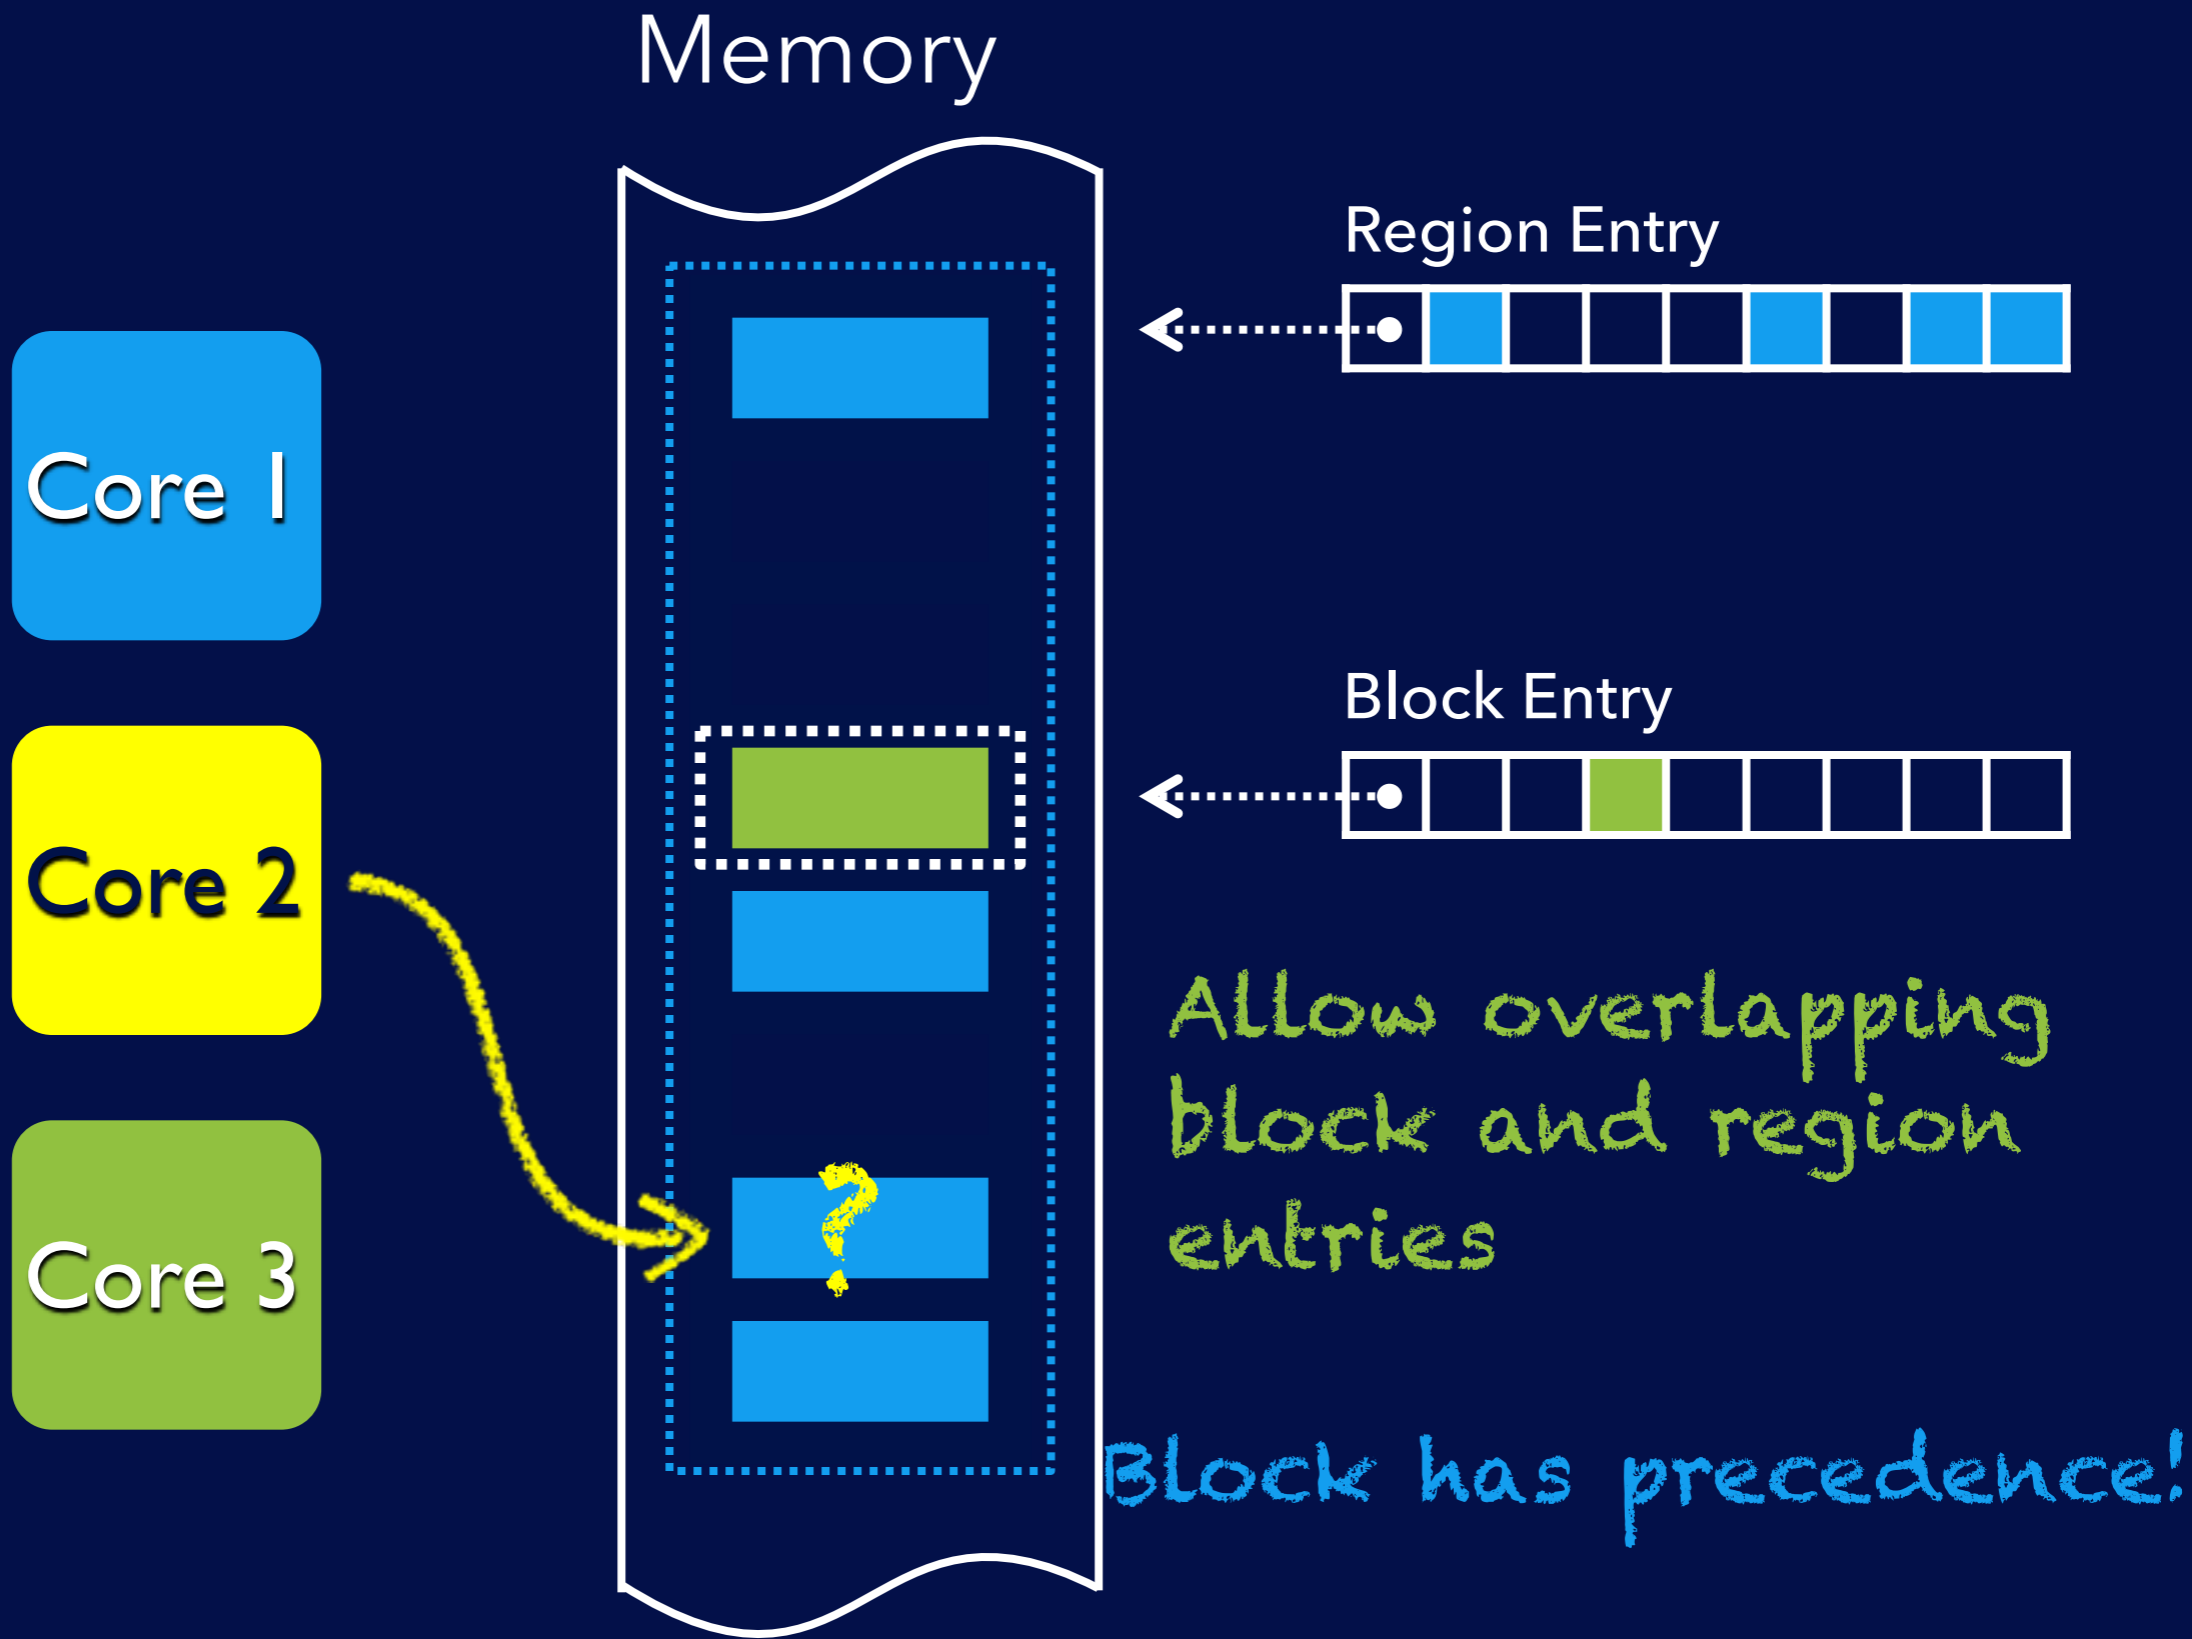


Block Entry



Allow overlapping
block and region
entries

Block has precedence!




Core 1

Core 2

Core 3

Memory



Shared blocks not tracked in region

Region Entry



Block Entry



Block Entry




Memory

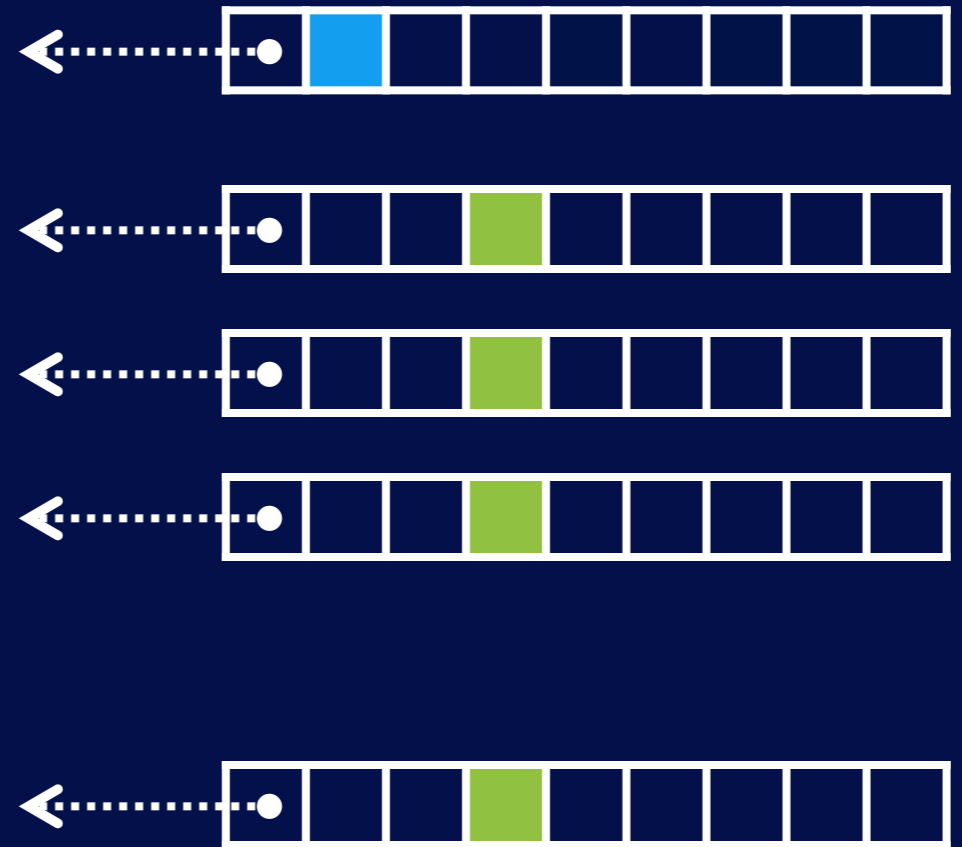
Core 1

Core 2



Core 3





Region Entry




After eviction?



After eviction?



Just invalidate region




Core 1

Core 2


Core 3

Memory



Re-allocation

Region Entry



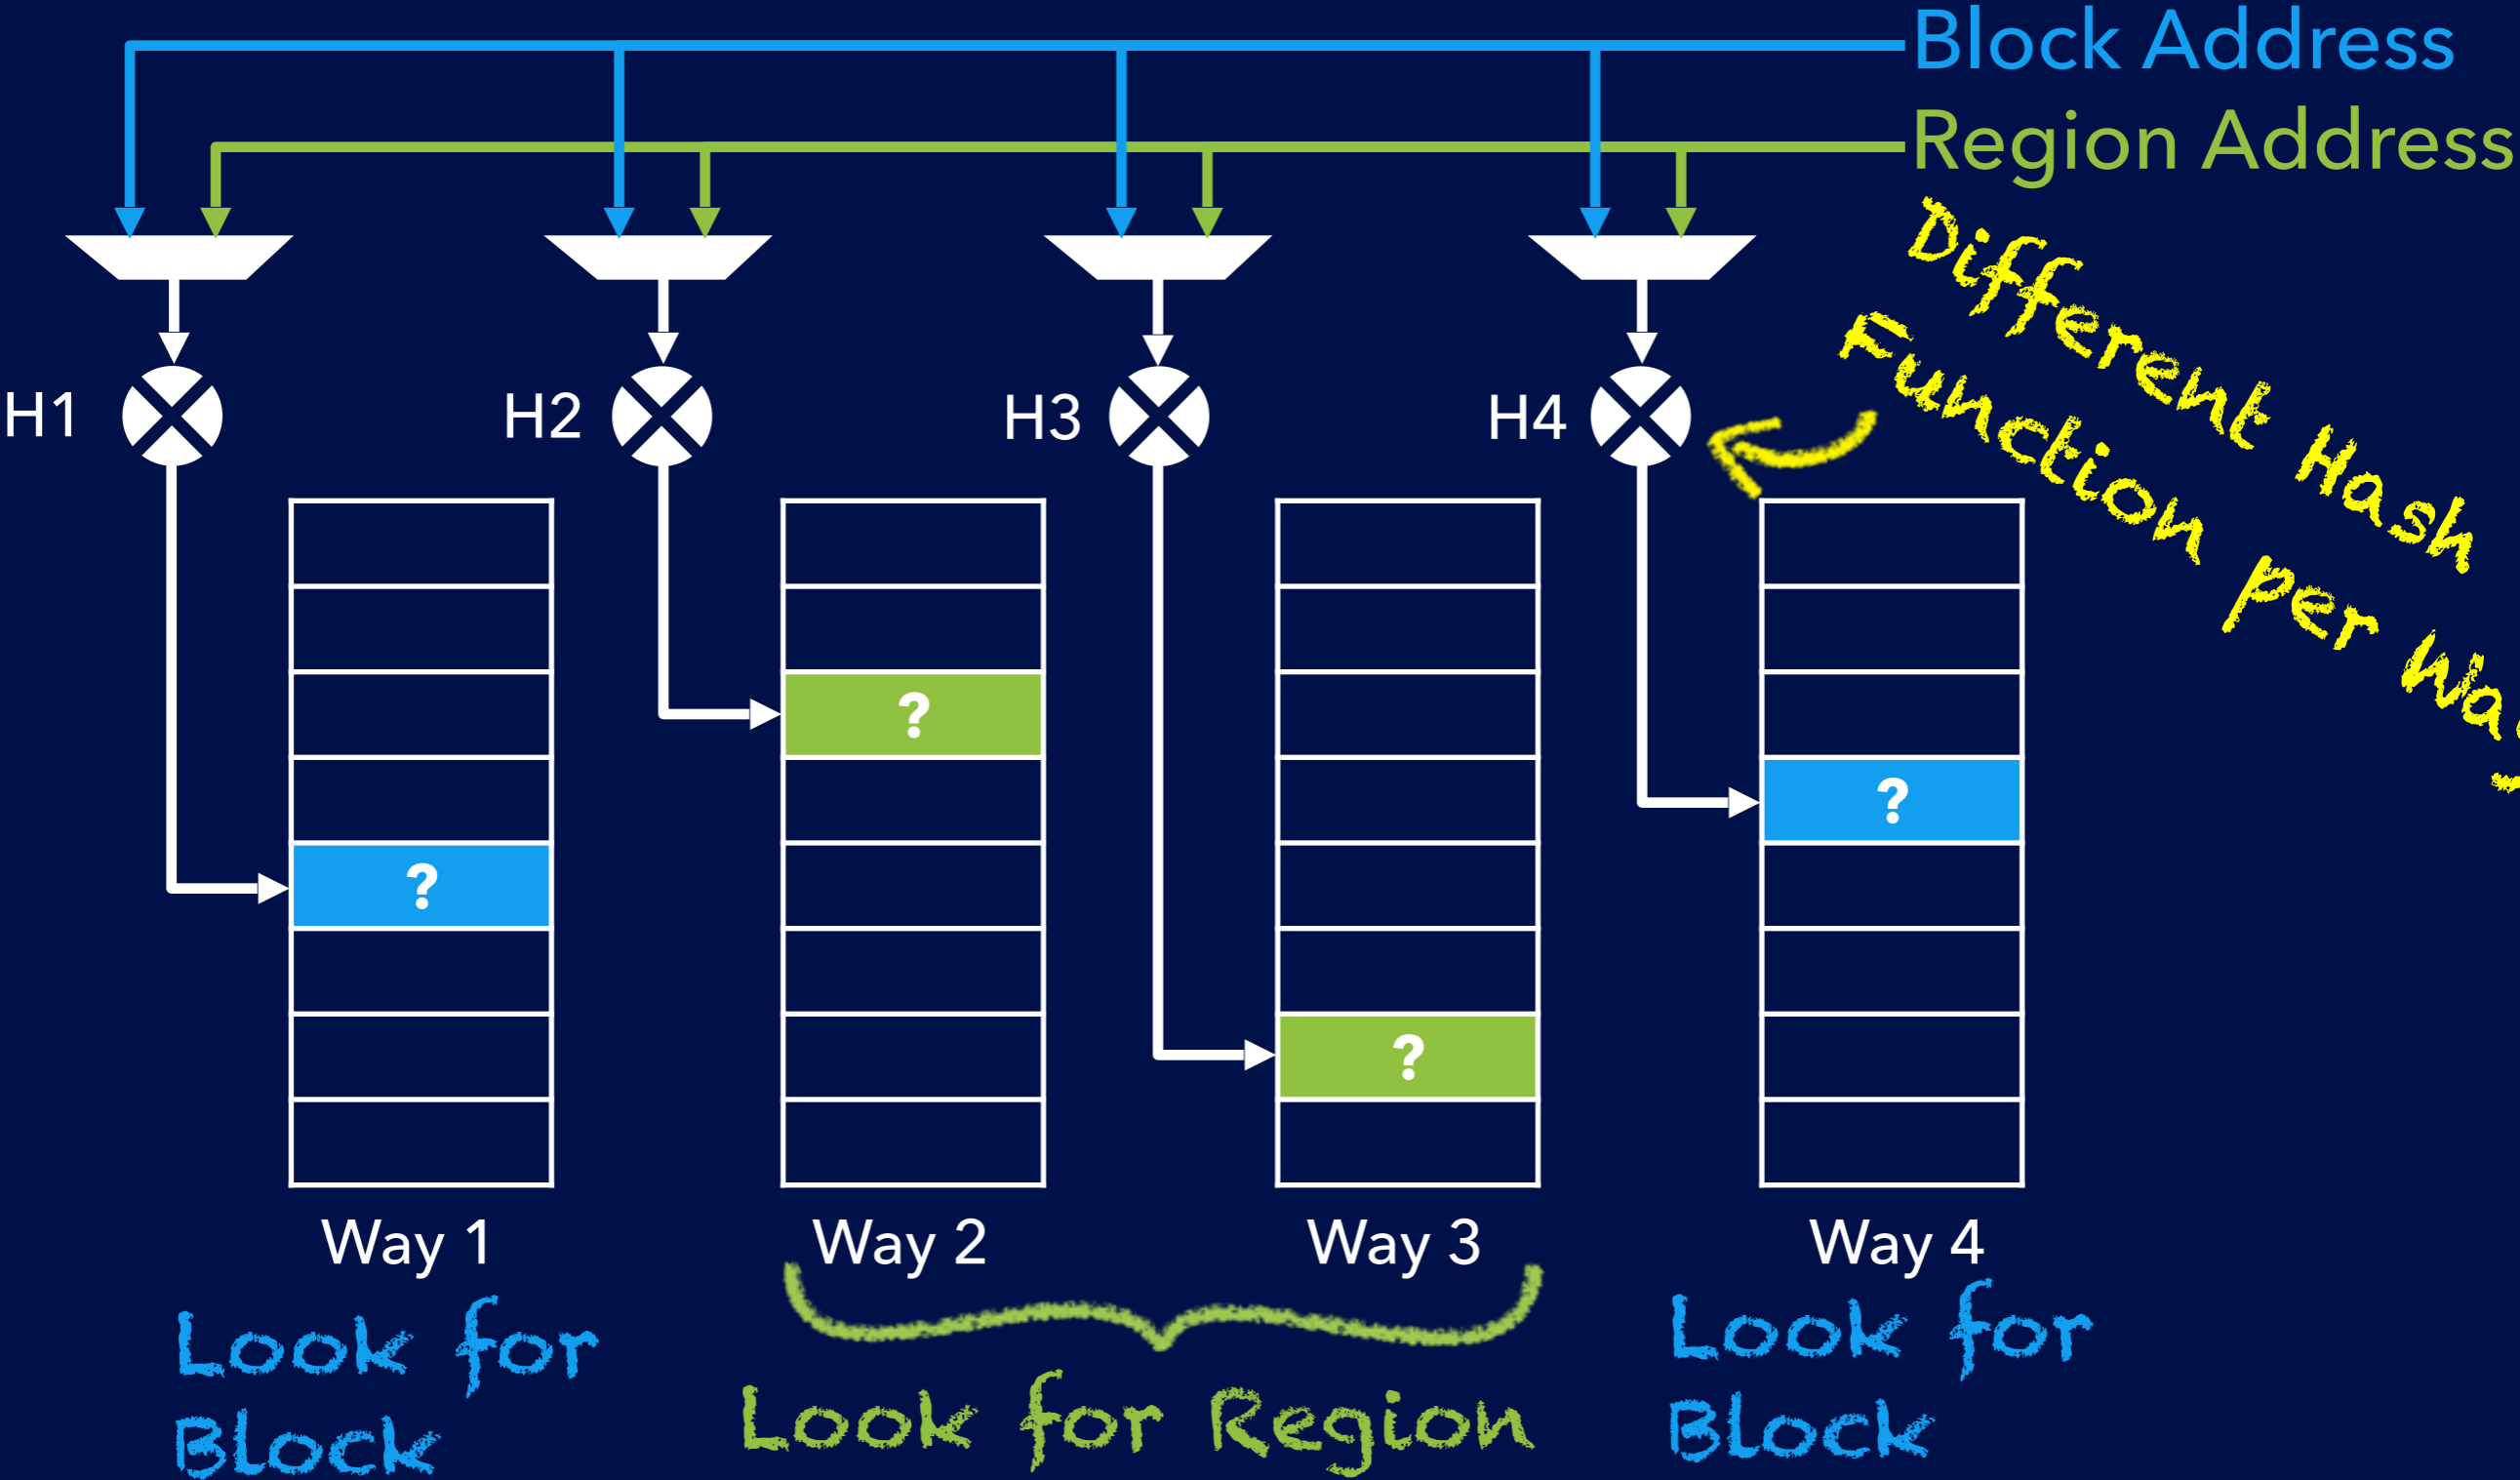
Lazy Merging!
Attempt merge if block
entry gets evicted from
directory.

DGD Directory Structure

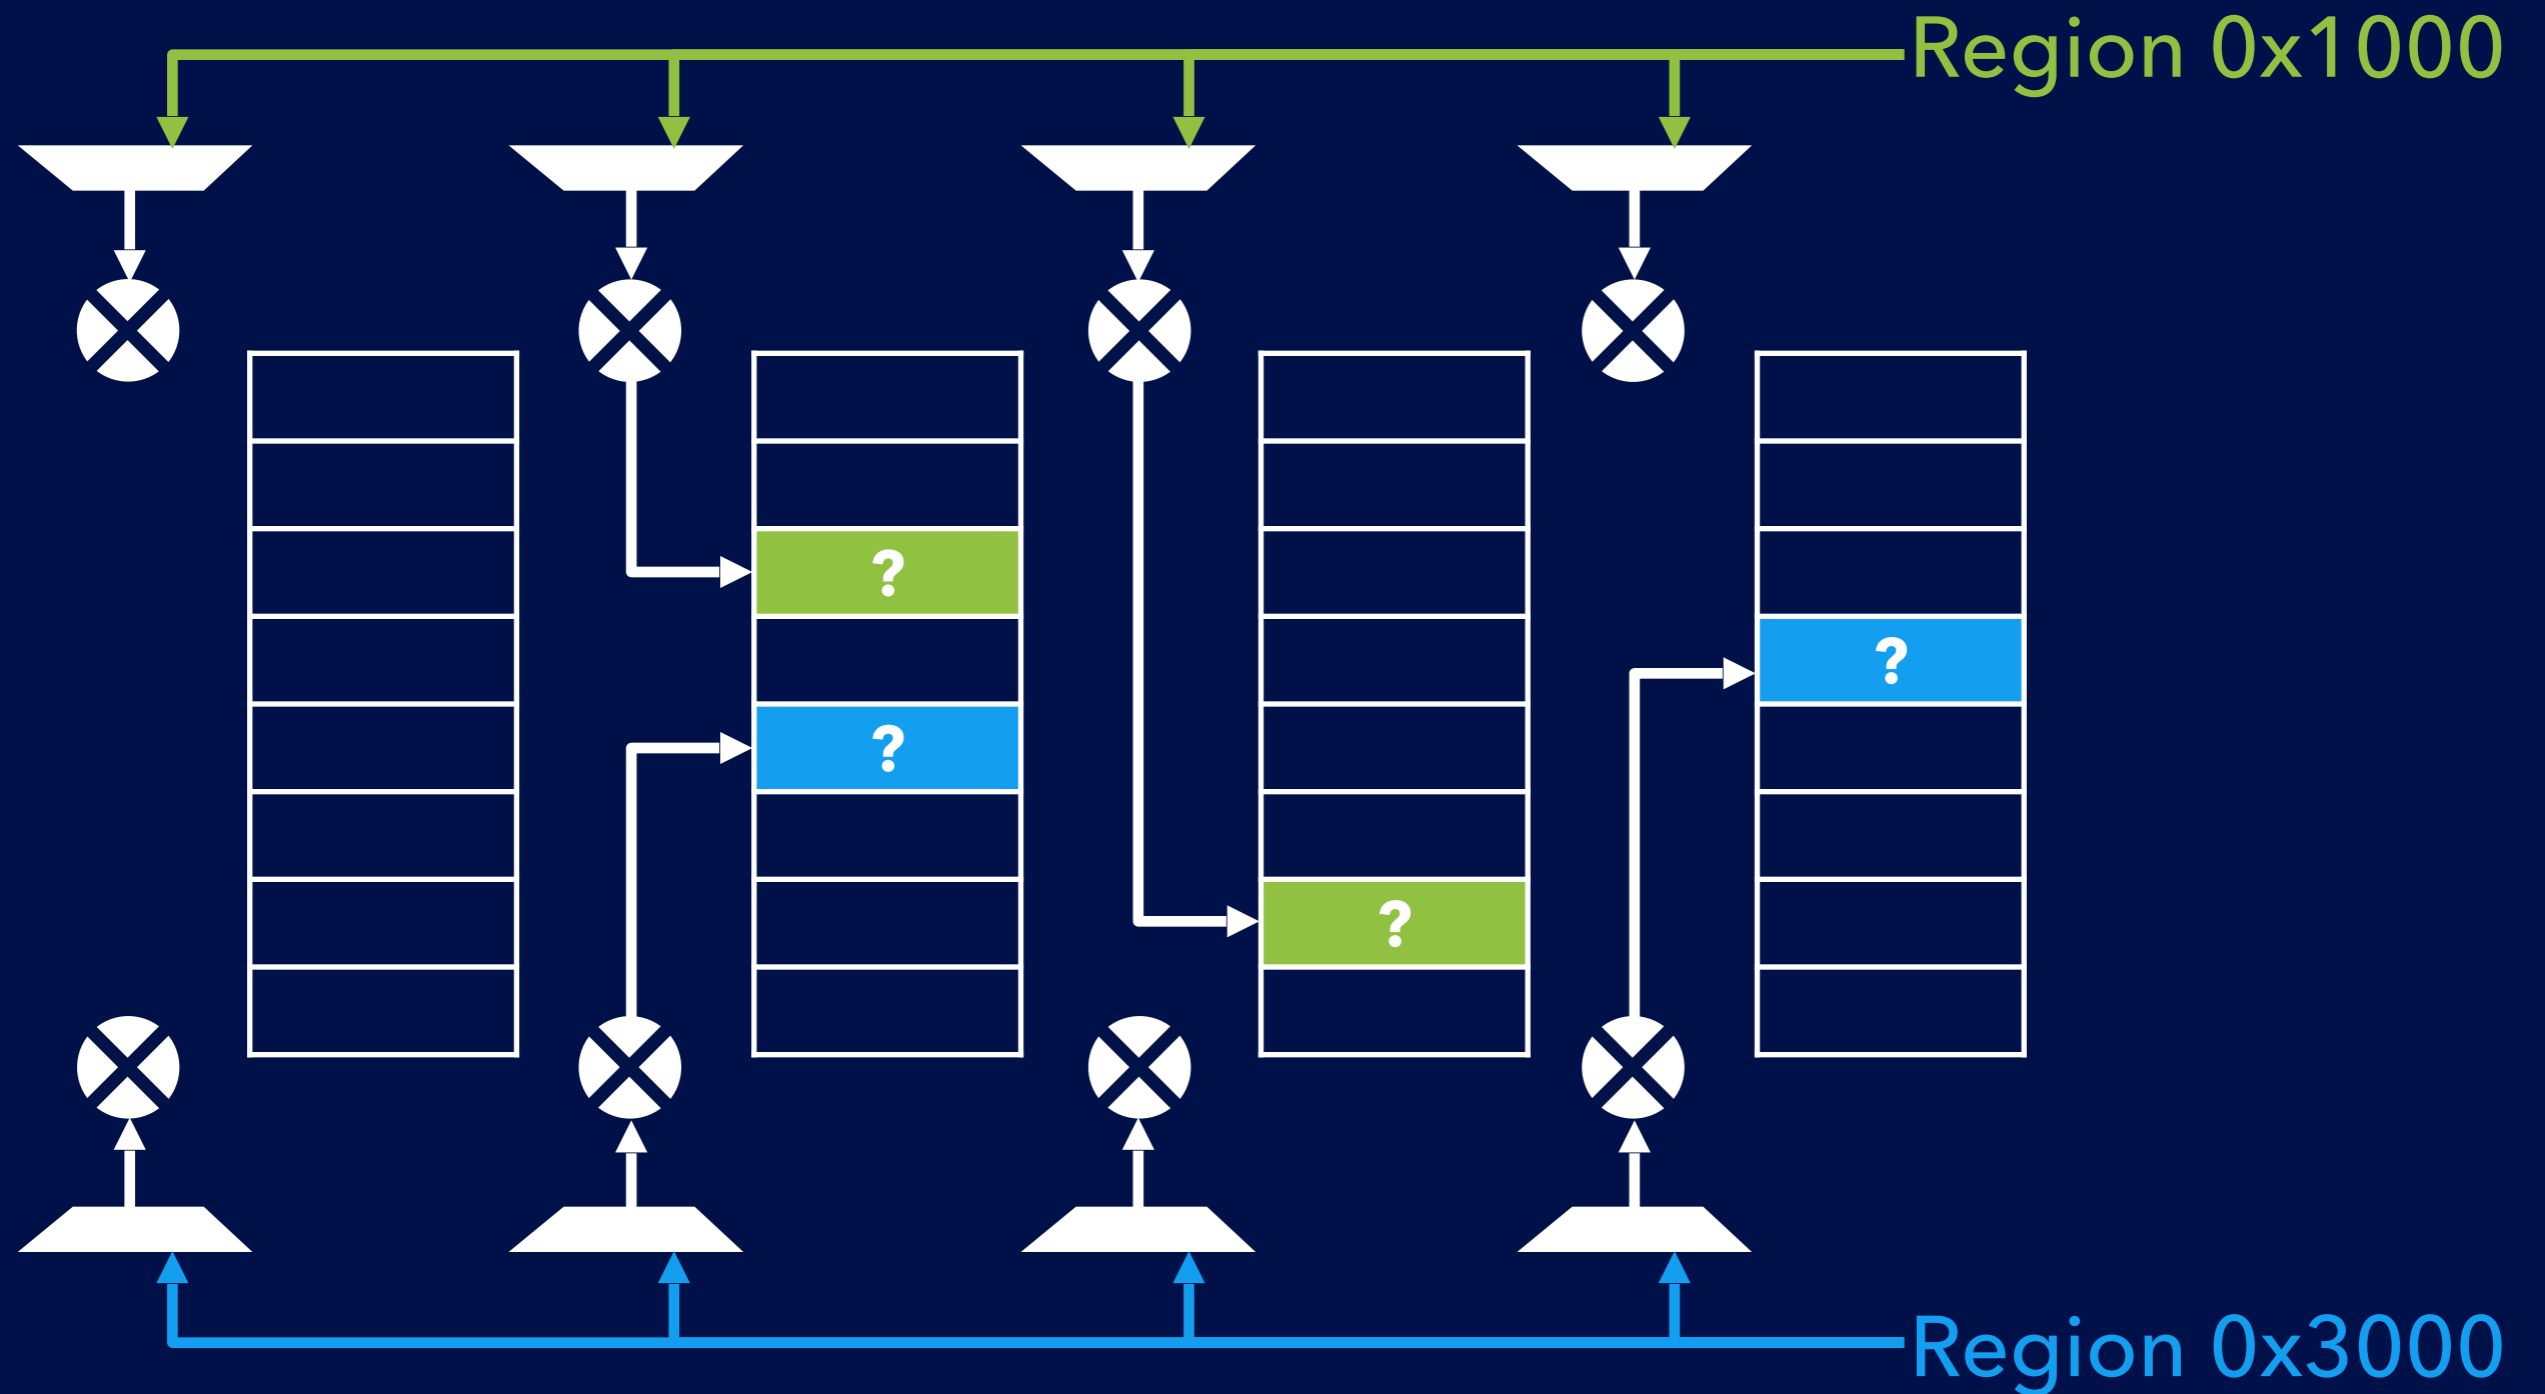
Objectives

- ❖ Single structure for block and region entries
- ❖ Find block and region entries with one lookup
 - *avoid multiple sequential lookups on critical path*
- ❖ Make lookups efficient
 - *allocation and replacement less critical*

DGD Structure

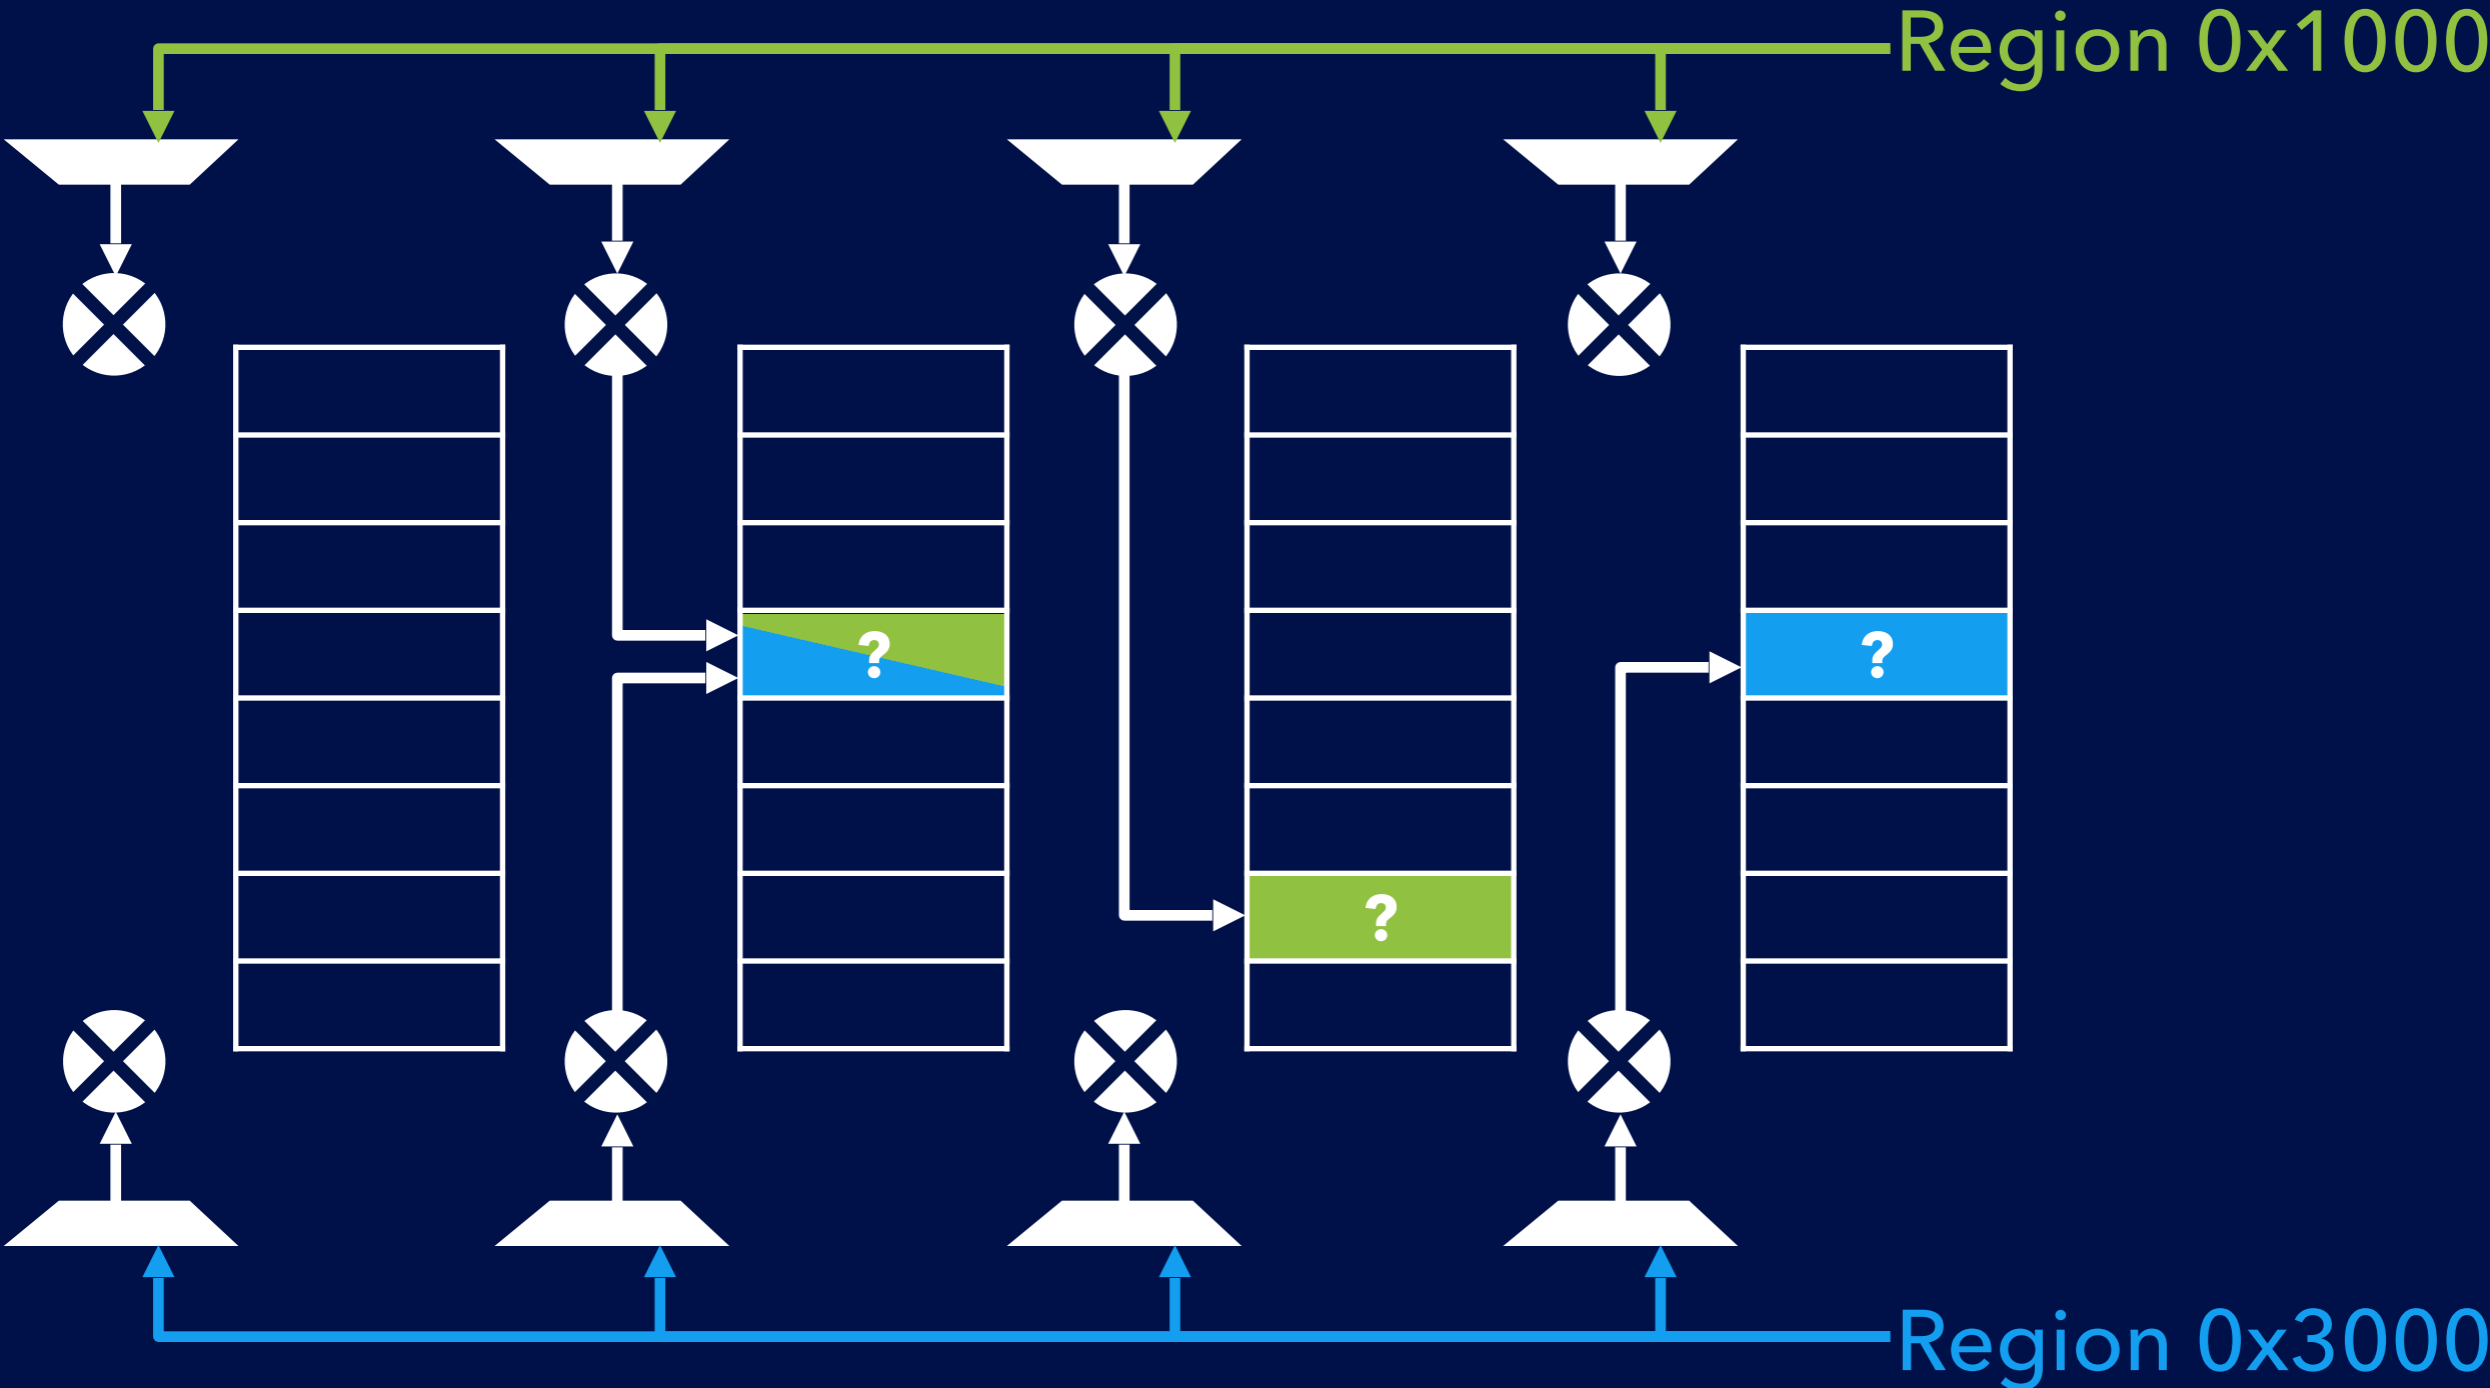


DGD Structure



Regions and blocks spread across all ways

DGD Structure



Complex options for resolving conflicts

Replacement Policy

- ❖ Skew-Associative with ZCache-style replacement
[Sanchez 2010]
 - *likely to find an unused entry*
 - *replacement can require multiple lookups*
- ❖ Equal preference to both block and region entries
 - *Invalidate cached blocks on directory eviction*
 - *Invalidated blocks allocated in shared cache if necessary*
- ❖ Not on the critical path
 - *Insertion/Eviction queue to delay actions*

- ▶ Original MGD concept
- ▶ Potential benefits of MGD
- ▶ Making a practical design
- ▶ **A recent competing design**
- ▶ Some nice graphs
- ▶ Final Thoughts

Spatiotemporal Coherence Tracking (SCT)

Mohammad Alisafae, MICRO 2012

Block Entry:

0	Block Tag	Sharers	V	LRU
---	-----------	---------	---	-----

Region Entry:

1	Region Tag	Owner	PBC	SBC		V	LRU
---	------------	-------	-----	-----	--	---	-----

	<u>DGD</u>	<u>SCT</u>
Consider more than 2 granularities	✓	✗
No Protocol Changes	✓	✗
No New Race Conditions	✓	✗
Single Lookup for Each Access	✓	✗

- ▶ Original MGD concept
- ▶ Potential benefits of MGD
- ▶ Making a practical design
- ▶ A recent competing design
- ▶ **Some nice graphs**
- ▶ Final Thoughts


Simulation Details

- ❖ 16-core chip-multiprocessor
 - ❖ 64 KB L1 Instruction and Data caches
 - ❖ 256 KB private L2 cache
 - ❖ 16 MB distributed shared L3 cache
 - ❖ 8-way set-associative directory
-
- ❖ Simics + Flexus full-system simulation
 - ❖ SPARCv9 ISA, Solaris OS (version 8 or 10)
 - ❖ Detailed timing of out-of-order cores and memory hierarchy
-
- ❖ Selection of 18 diverse applications
 - ❖ SpecWeb, Parsec, TPC-C, TPC-H, CloudSuite

Dual-Grain Directory Performance

50% Entry Reduction vs. Baseline Directory


% Speedup vs. Sparse Directory



Dual-Grain Directory Performance

50% Entry Reduction vs. Baseline Directory

% Speedup vs. Sparse Directory




Dual-Grain Directory Performance

50% Entry Reduction vs. Baseline Directory

Most robust performance

% Speedup vs. Sparse Directory



- ▶ Original MGD concept
- ▶ Potential benefits of MGD
- ▶ Making a practical design
- ▶ A recent competing design
- ▶ Some nice graphs
- ▶ **Final Thoughts**

Multi-Grain Coherence Directory

Ideal MGD Directory:

- ✓ Conceptual exploration
- ✓ Majority of benefit from just 2 granularities

Practical Dual-Grain Directory:

- ✓ Track private regions and shared blocks
- ✓ No coherence protocol changes
- ✓ **50% fewer entries**
- ✓ Robust performance, minimal losses

