

Fairness via Source Throttling:

A configurable and high-performance fairness
substrate for multi-core memory systems

Eiman Ebrahimi*

Chang Joo Lee*


Onur Mutlu[‡]

Yale N. Patt*

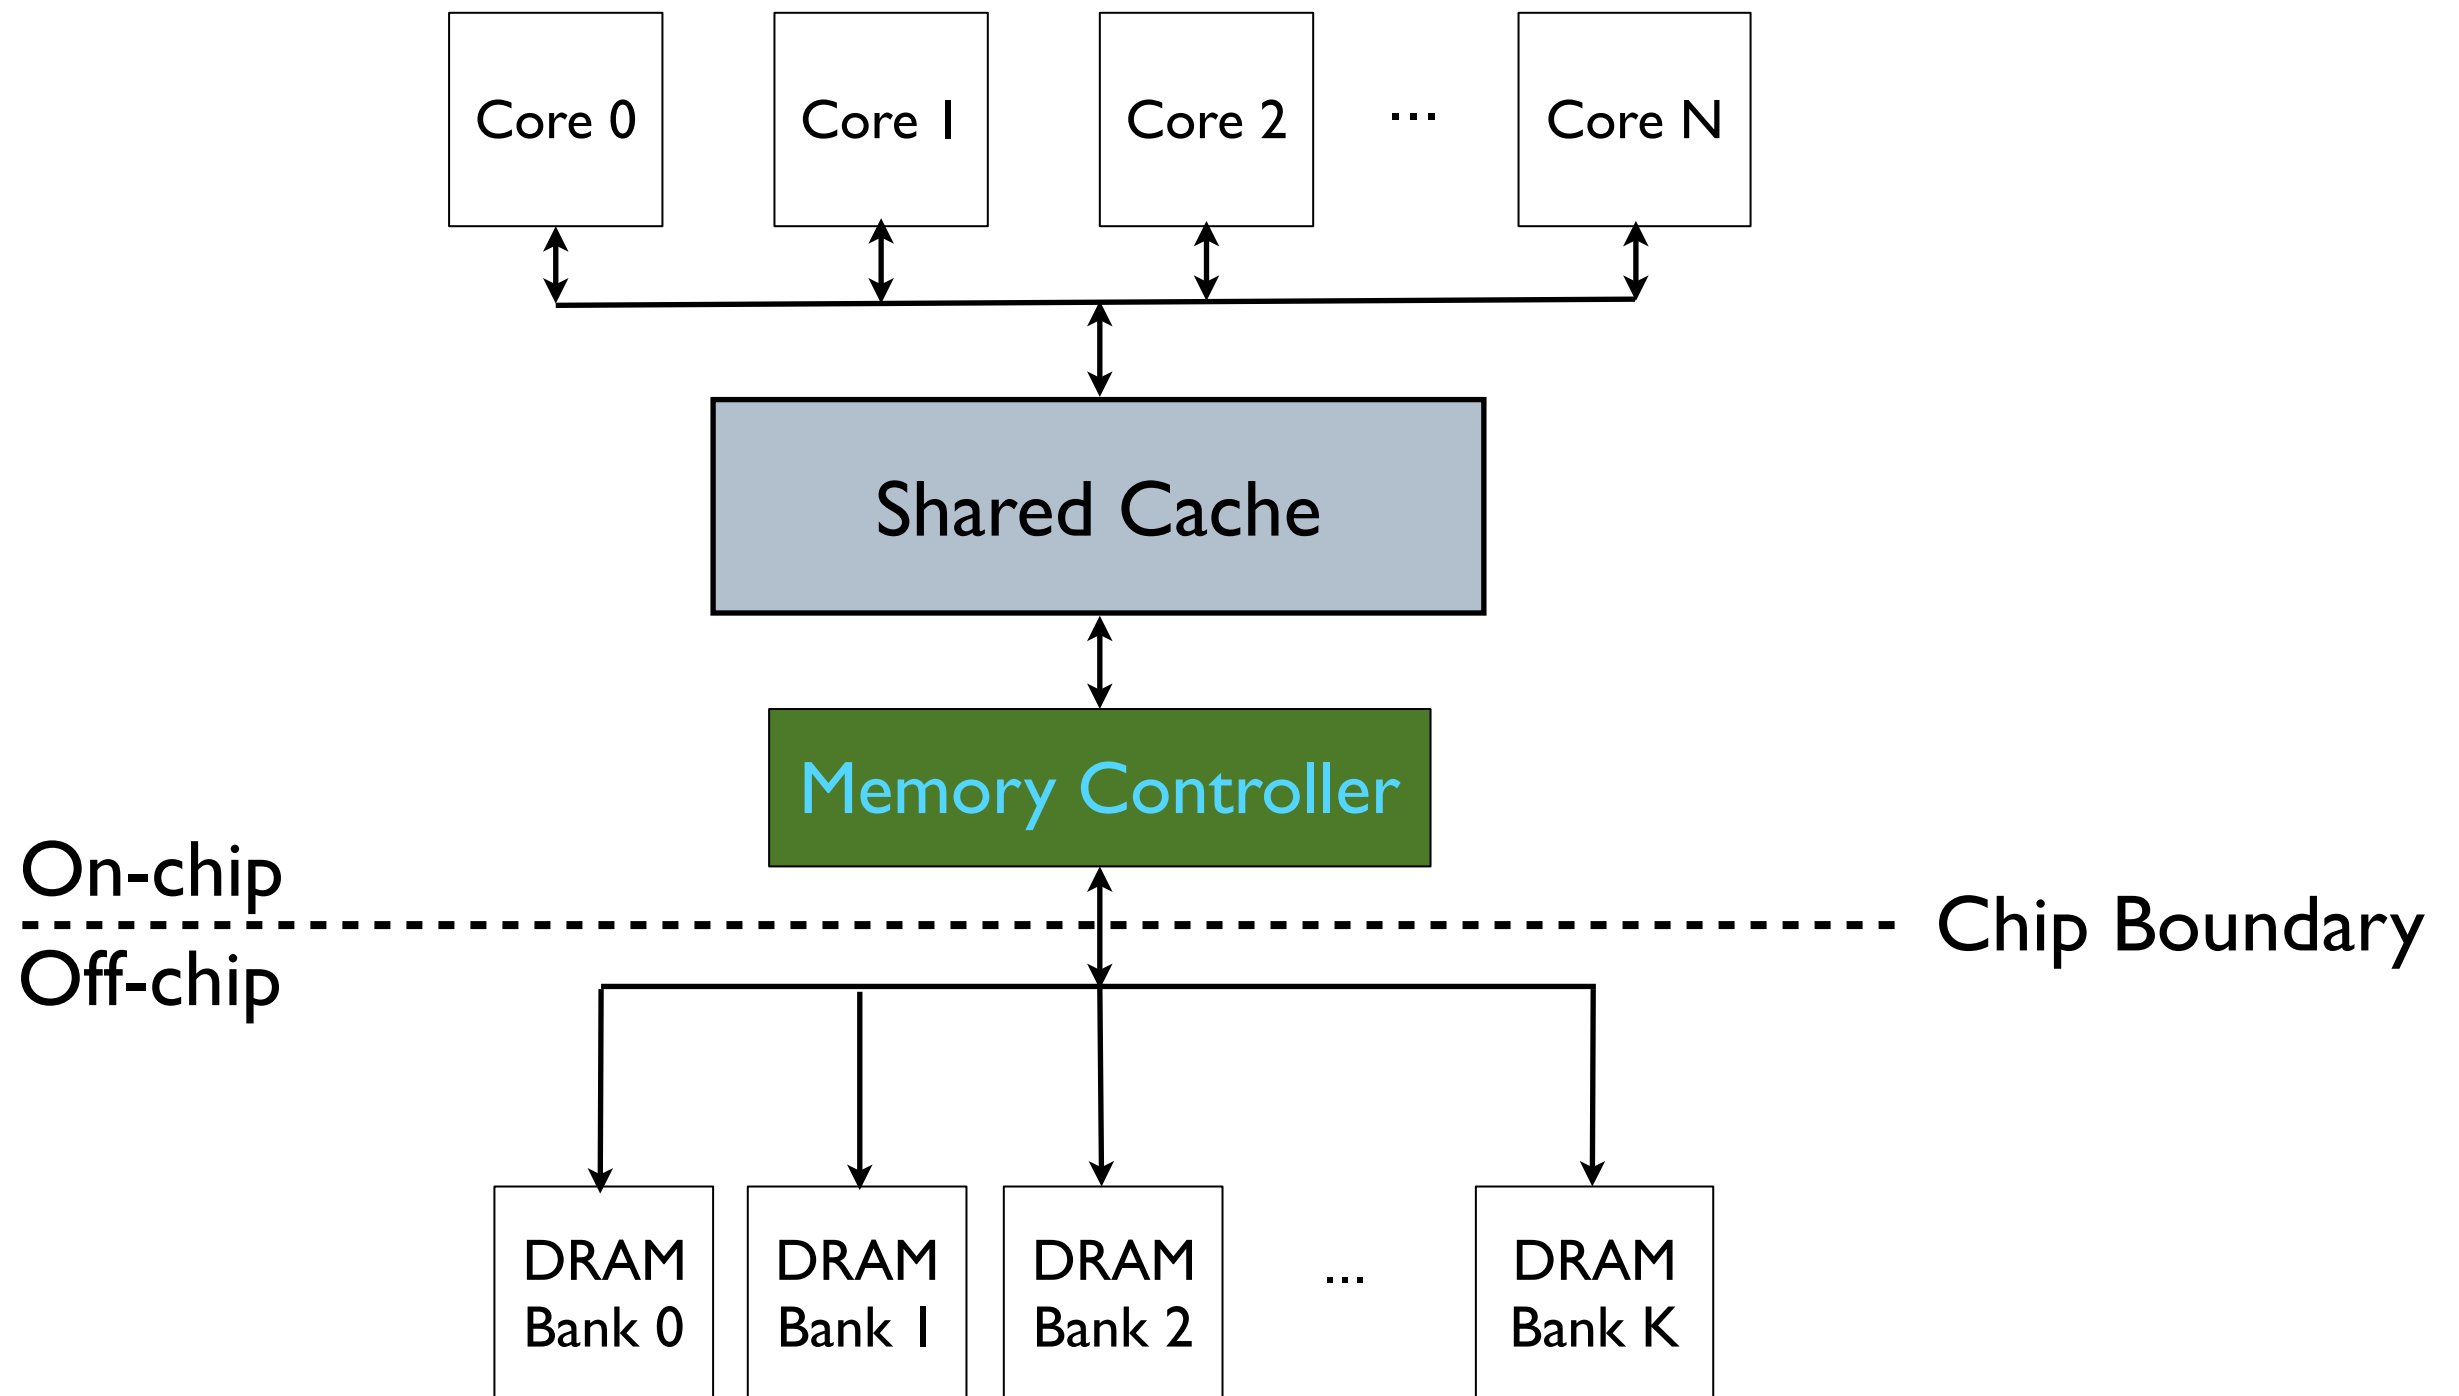
*** HPS Research Group
The University of Texas at Austin**

**‡ Computer Architecture Laboratory
Carnegie Mellon University**

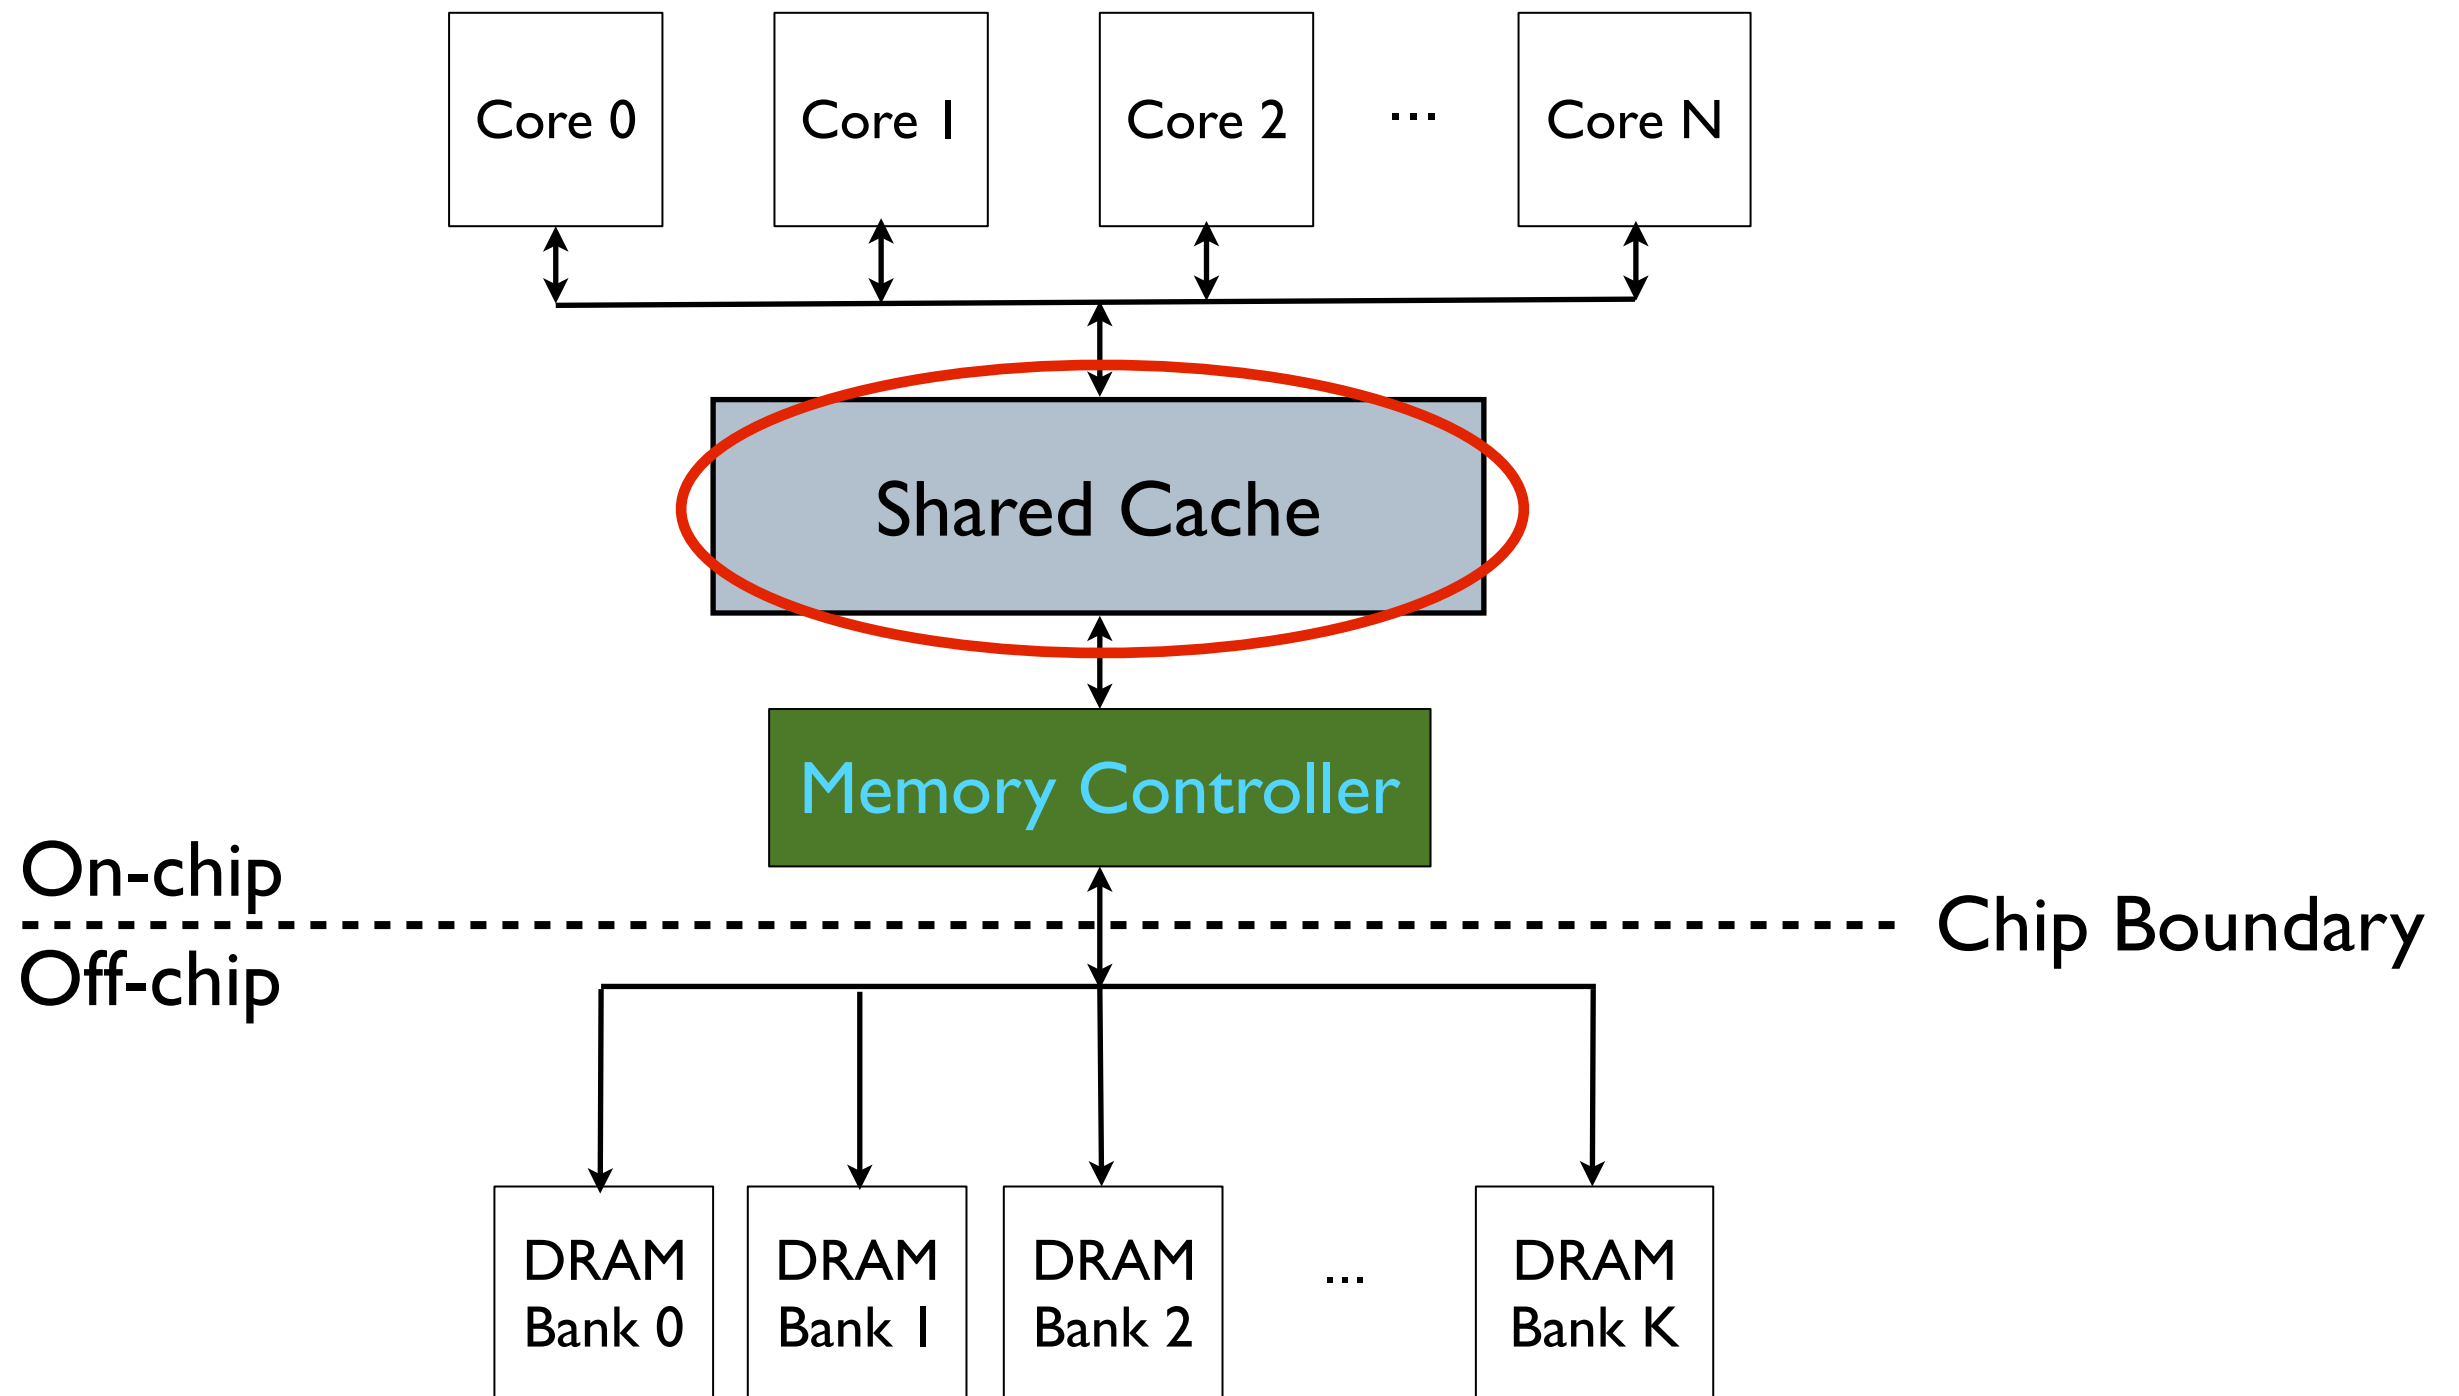
Background and Problem



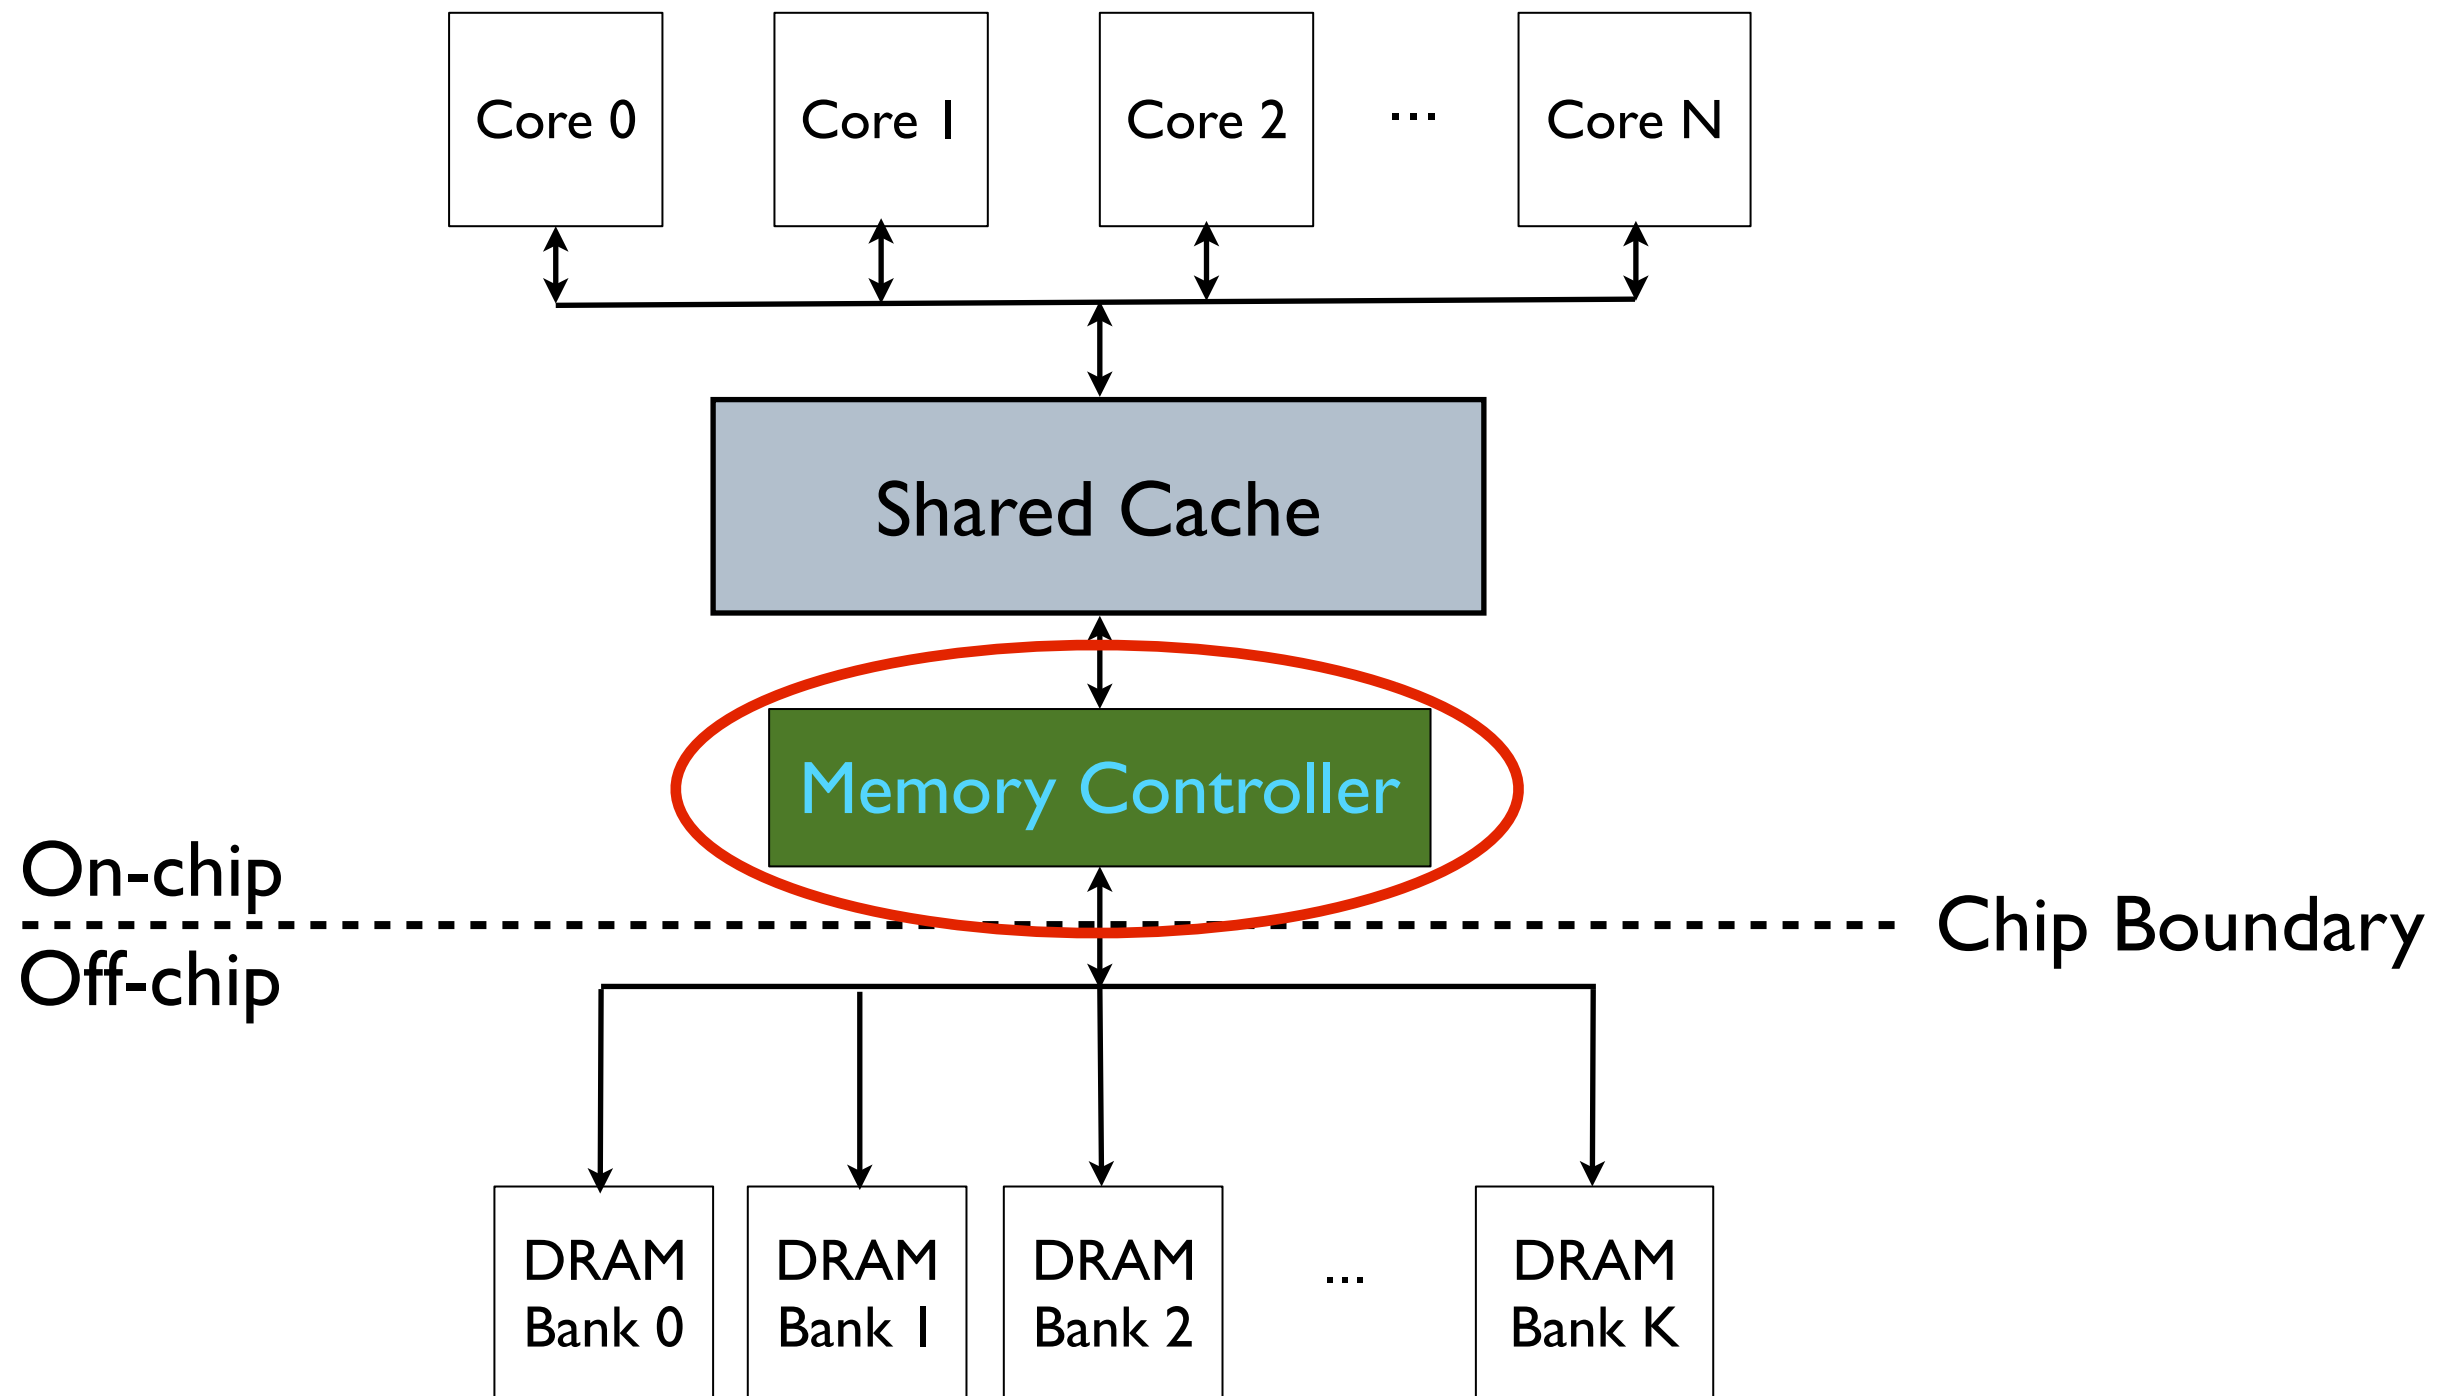
Background and Problem



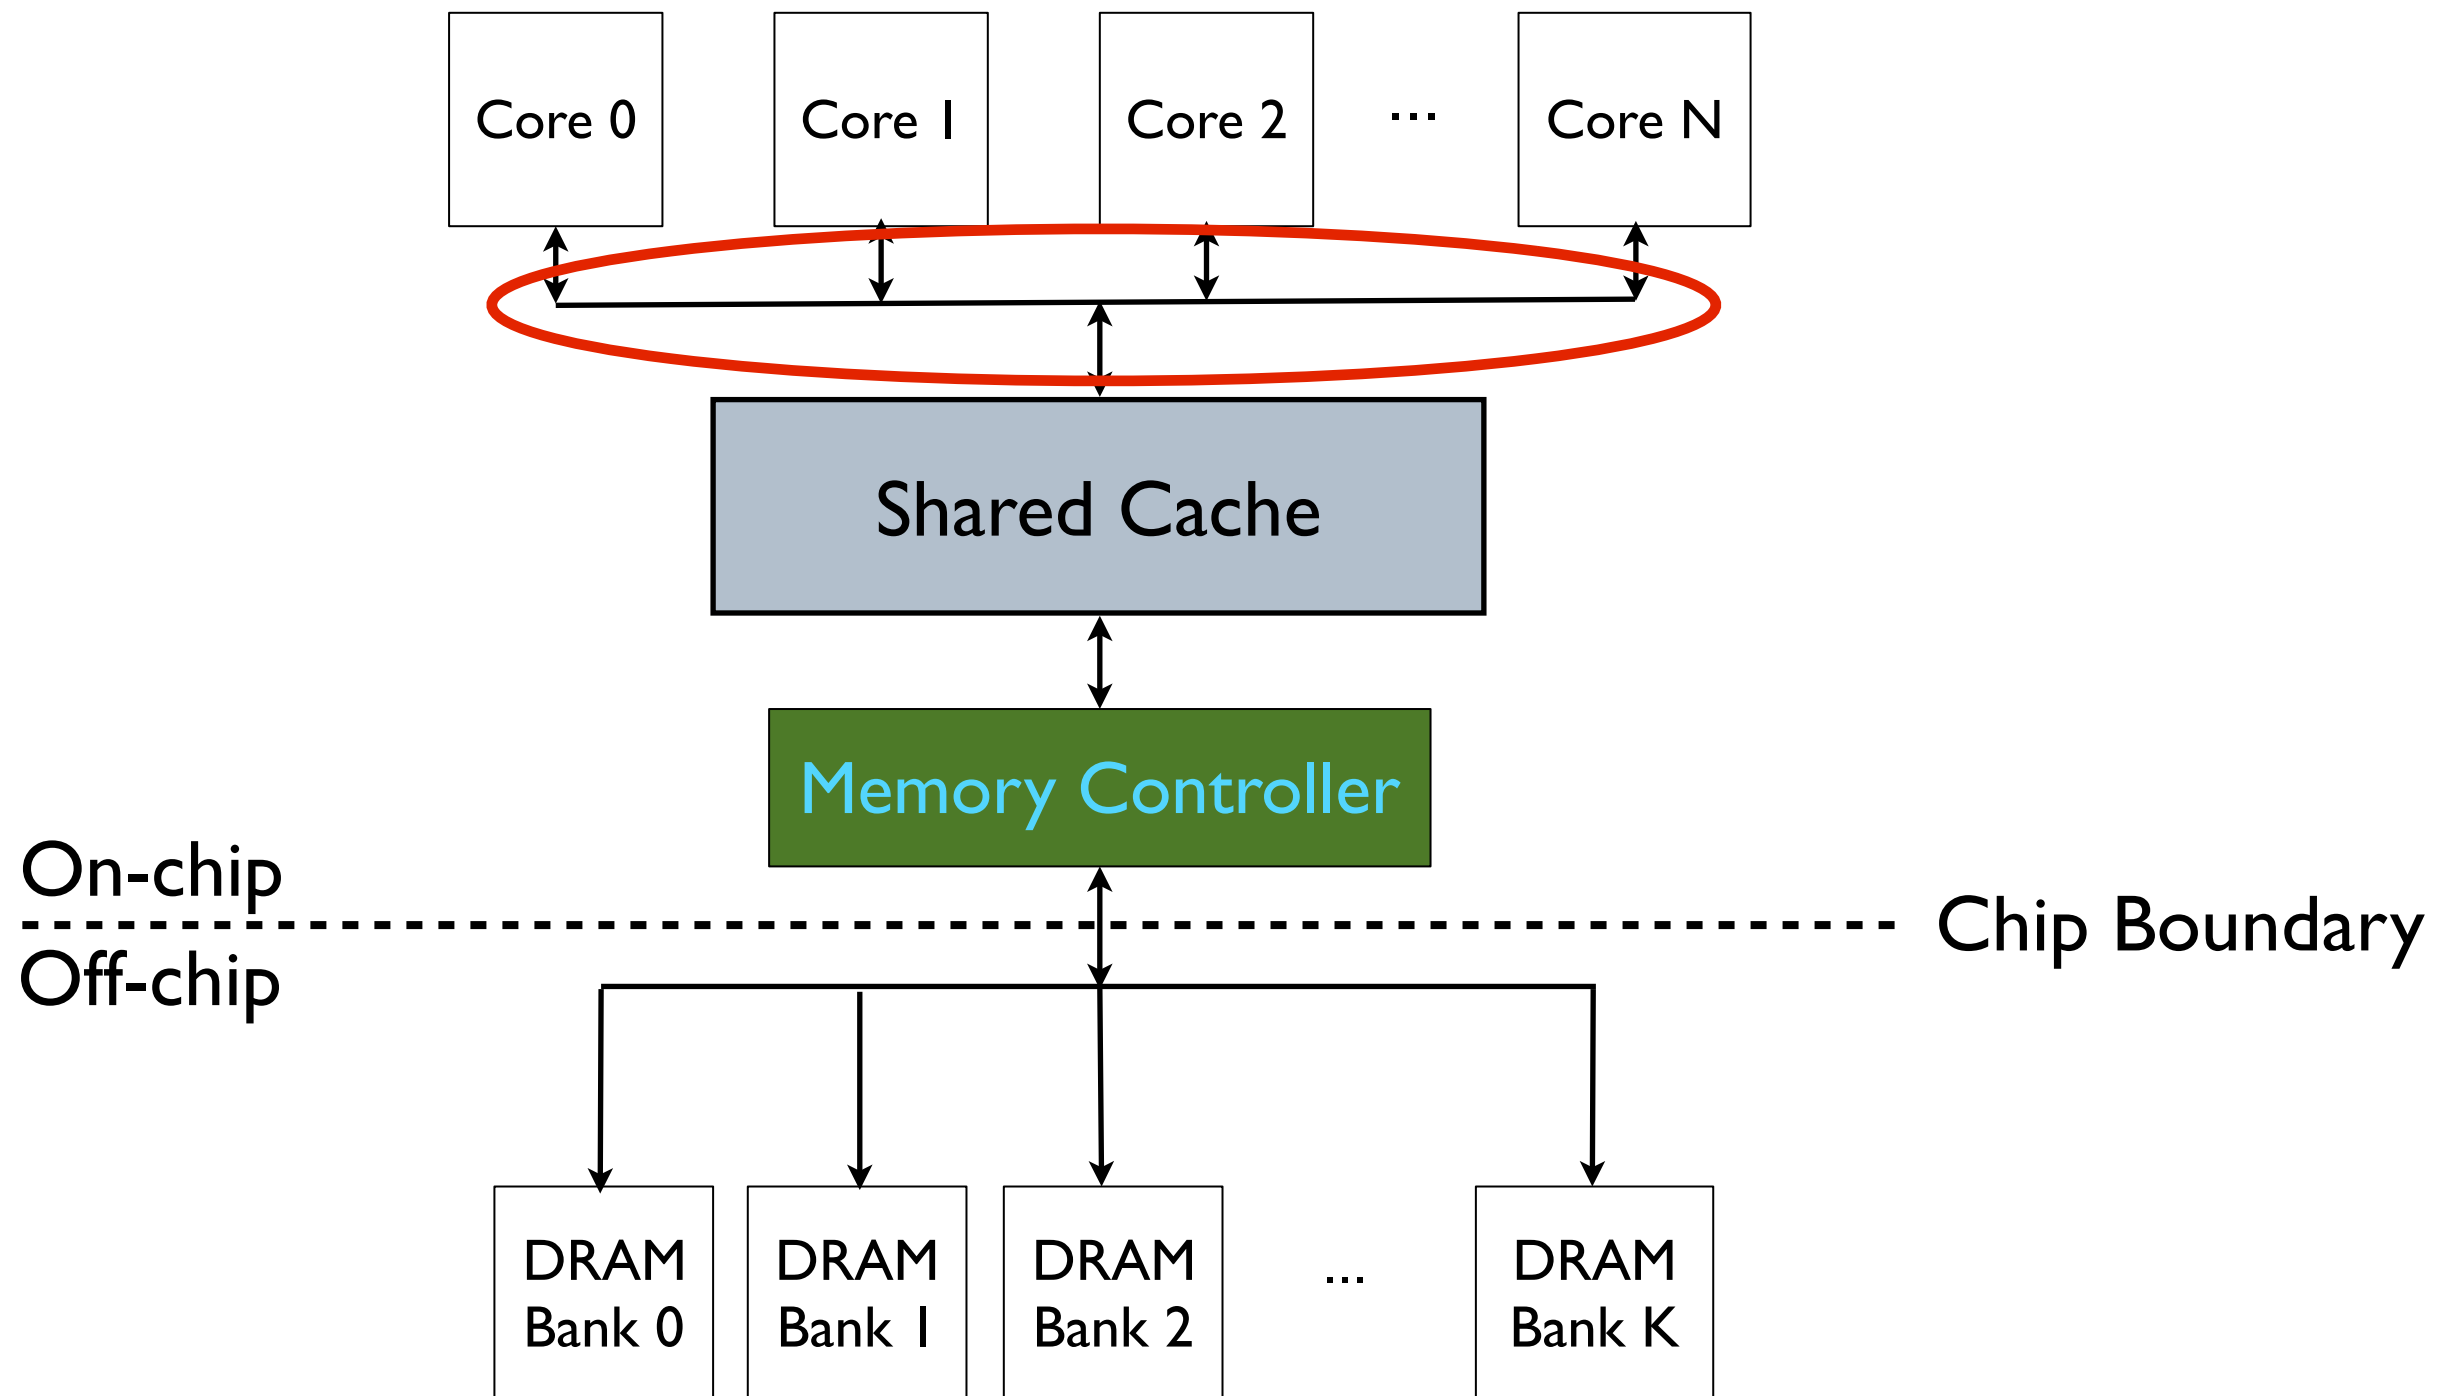
Background and Problem



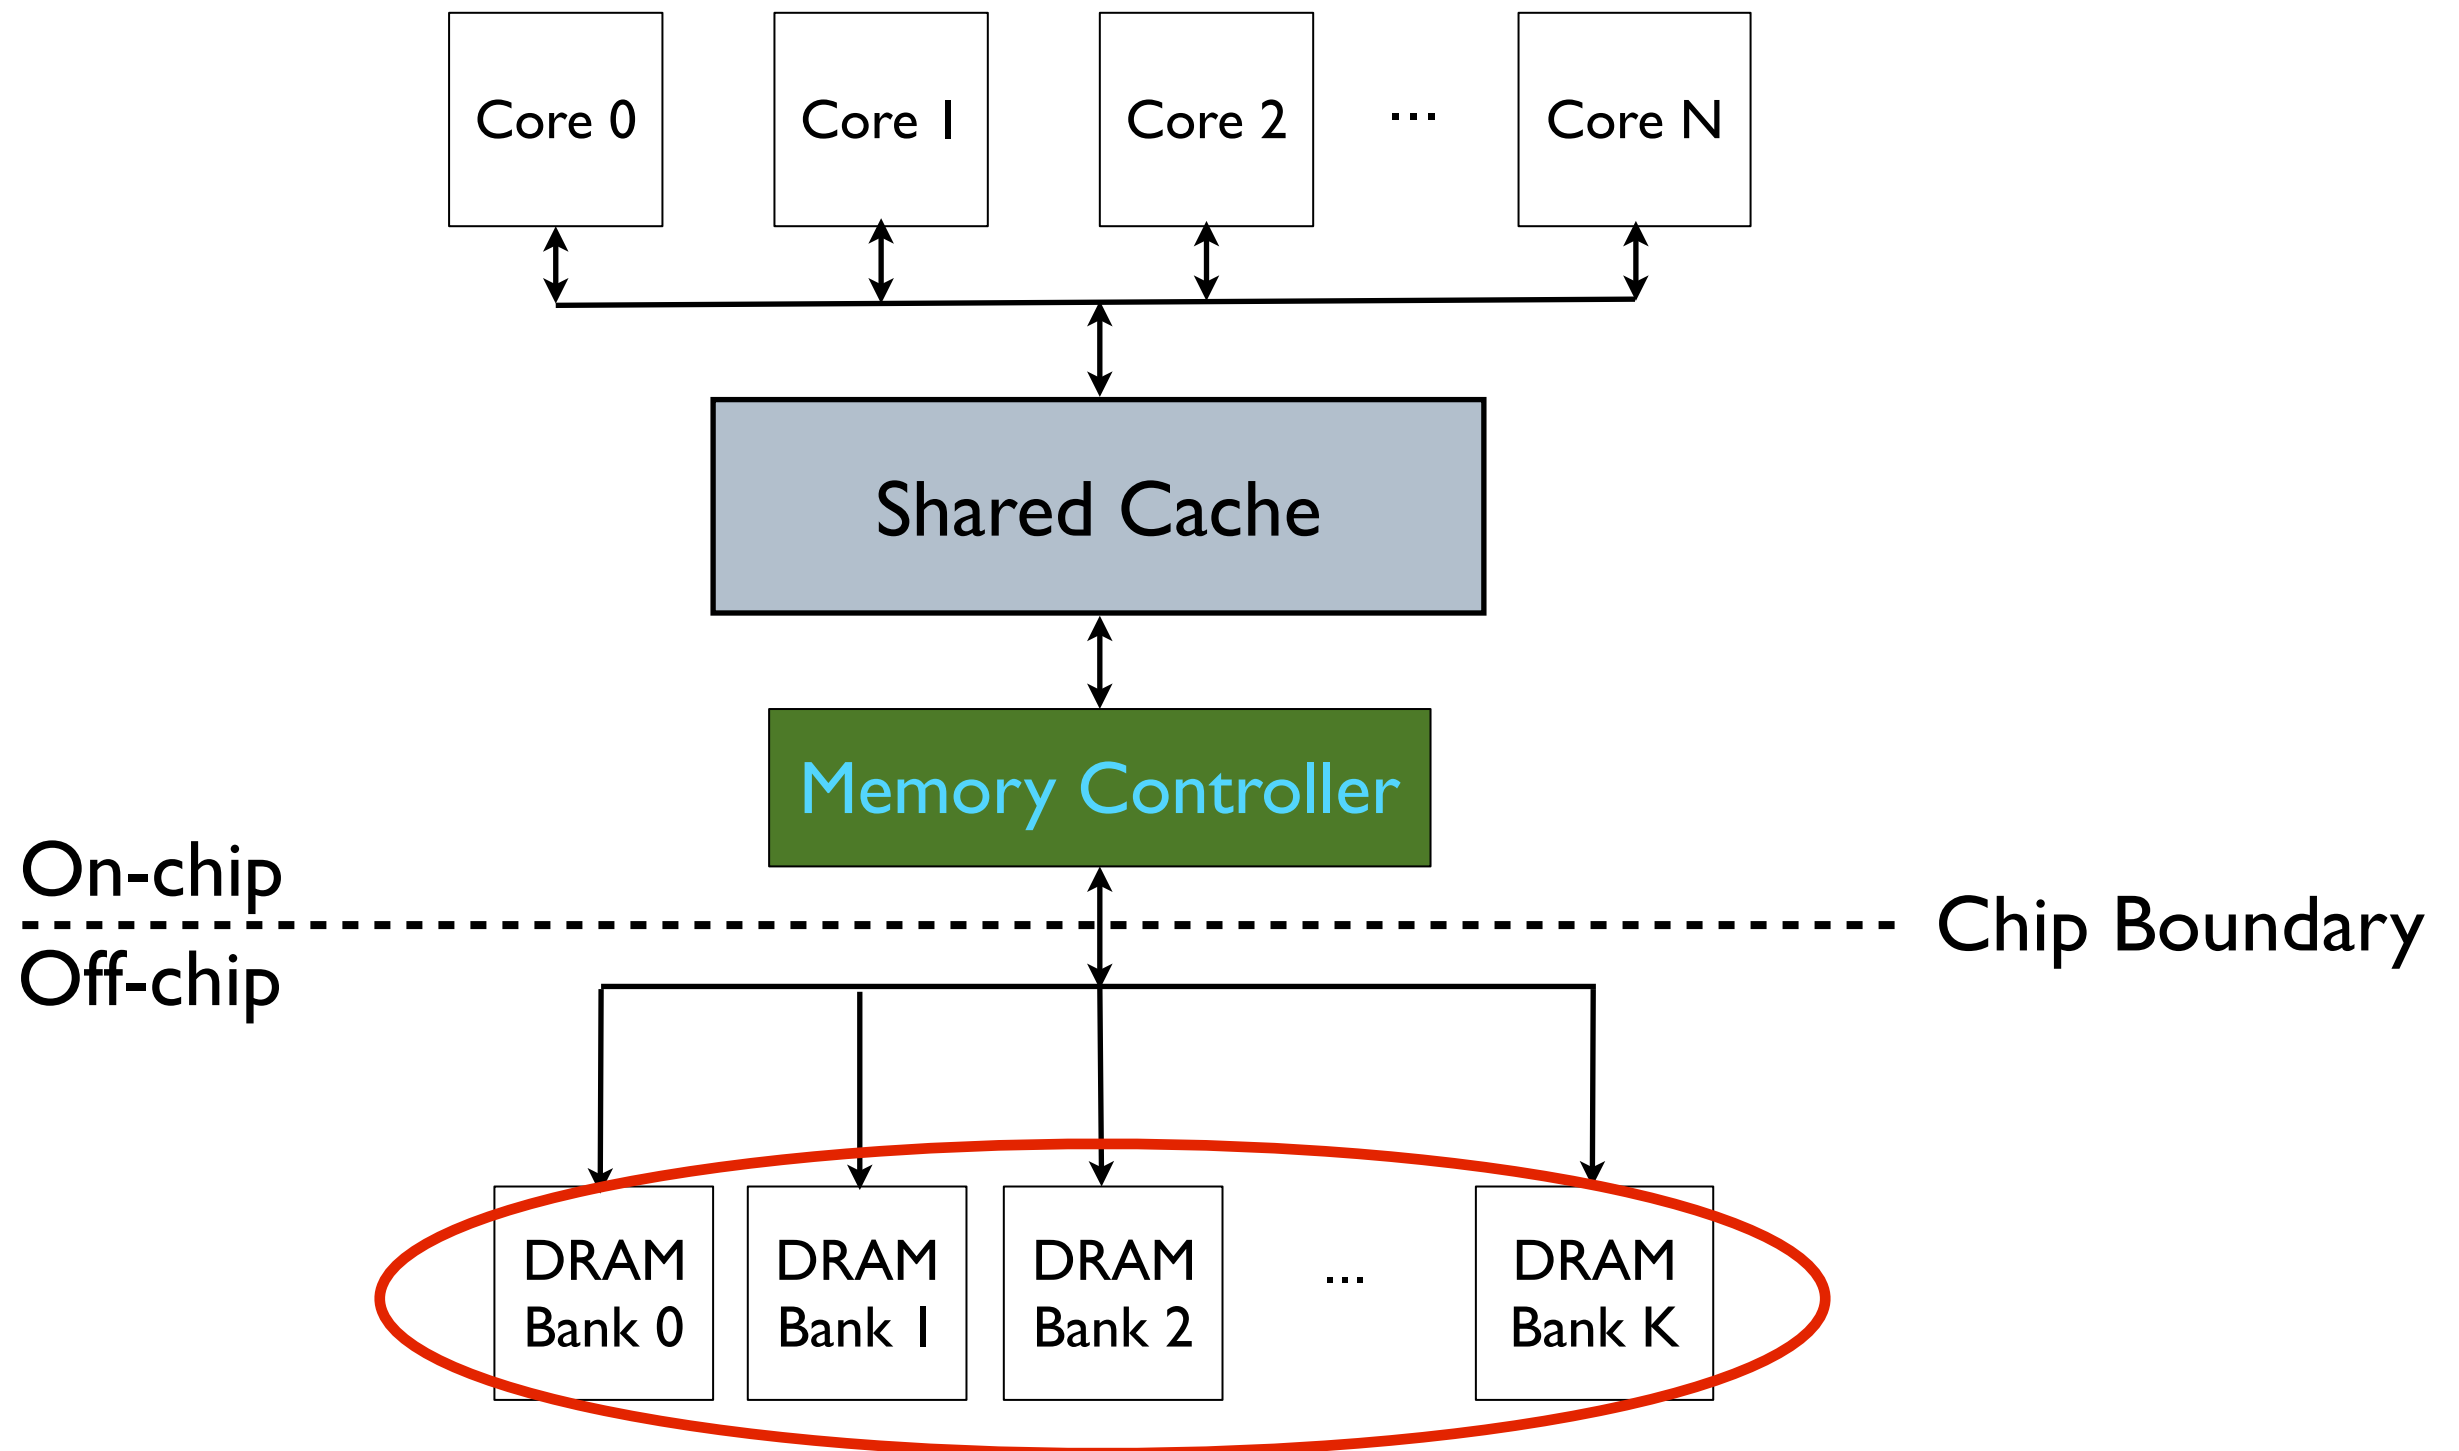
Background and Problem



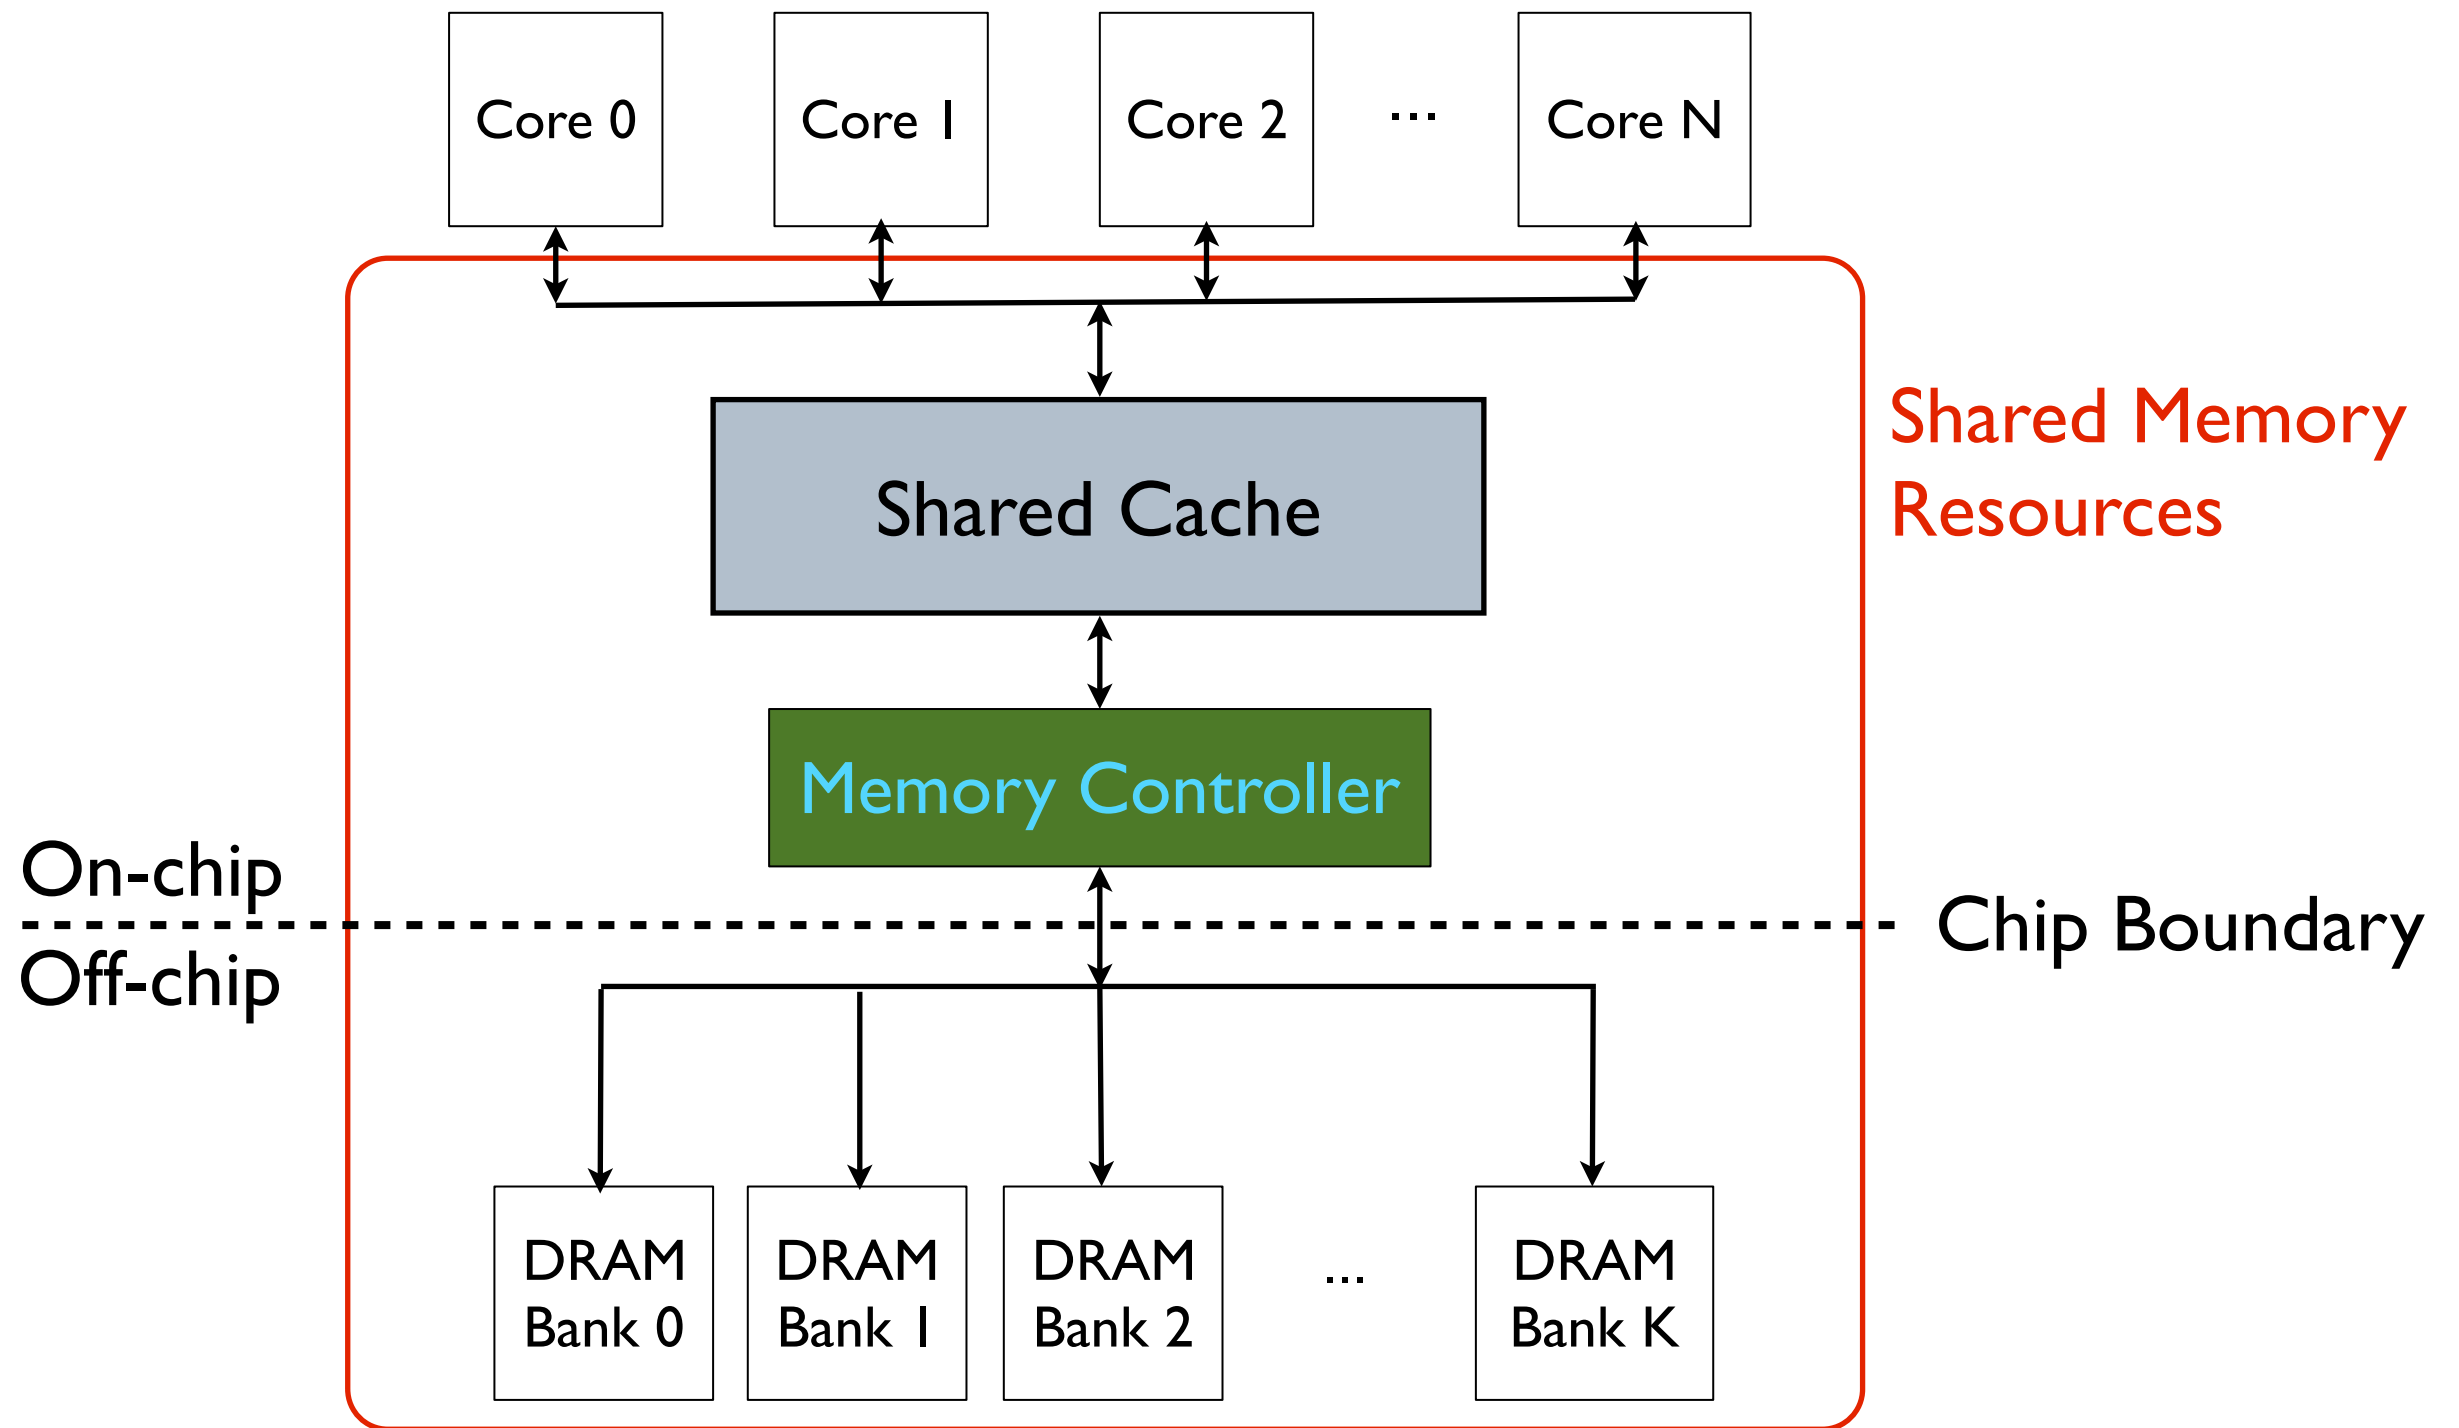
Background and Problem



Background and Problem



Background and Problem



Background and Problem

- Applications slow down due to **interference** from memory requests of other applications

Background and Problem

- Applications slow down due to **interference** from memory requests of other applications
- A memory system is **fair** if slowdowns of **same-priority** applications are **equal**
(MICRO '06, MICRO '07, ISCA '08)

Background and Problem

- Applications slow down due to **interference** from memory requests of other applications
- A memory system is **fair** if slowdowns of **same-priority** applications are **equal**
(MICRO '06, MICRO '07, ISCA '08)

- Slowdown of application $i = \frac{T_i^{\text{Shared}}}{T_i^{\text{Alone}}}$

Background and Problem

- Applications slow down due to **interference** from memory requests of other applications
- A memory system is **fair** if slowdowns of **same-priority** applications are **equal**
(MICRO '06, MICRO '07, ISCA '08)

- Slowdown of application $i = \frac{T_i^{\text{Shared}}}{T_i^{\text{Alone}}}$

- Unfairness = $\frac{\text{Max}\{\text{Slowdown } i\} \text{ over all applications } i}{\text{Min}\{\text{Slowdown } i\} \text{ over all applications } i}$
(MICRO '07)


Background and Problem

Background and Problem

- Magnitude of each application's slowdown depends on concurrently running applications' memory behavior



Background and Problem

- Magnitude of each application's slowdown depends on concurrently running applications' memory behavior




Background and Problem

- Magnitude of each application's slowdown depends on concurrently running applications' memory behavior



Background and Problem


- Magnitude of each application's slowdown depends on concurrently running applications' memory behavior



- Large disparities in slowdowns are unacceptable

Background and Problem


- Magnitude of each application's slowdown depends on concurrently running applications' memory behavior



- Large disparities in slowdowns are unacceptable
 - Low system performance

Background and Problem


- Magnitude of each application's slowdown depends on concurrently running applications' memory behavior



- Large disparities in slowdowns are unacceptable
 - Low system performance
 - Vulnerability to denial of service attacks

Background and Problem

- Magnitude of each application's slowdown depends on concurrently running applications' memory behavior



- Large disparities in slowdowns are unacceptable
 - Low system performance
 - Vulnerability to denial of service attacks
 - Difficult for system software to enforce priorities

Outline

- Background and Problem
- Motivation for Source Throttling
- Fairness via Source Throttling (FST)
- Evaluation
- Conclusion

Prior Approaches

Prior Approaches

- Primarily manage inter-application interference in only one particular resource
 - Shared Cache, Memory Controller, Interconnect, etc.

Prior Approaches

- Primarily manage inter-application interference in only one particular resource
 - Shared Cache, Memory Controller, Interconnect, etc.
 - Combining techniques for the different resources can result in negative interaction

Prior Approaches

- Primarily manage inter-application interference in only one particular resource
 - Shared Cache, Memory Controller, Interconnect, etc.
 - Combining techniques for the different resources can result in negative interaction
- Approaches that coordinate interaction among techniques for different resources require complex implementations

Prior Approaches

- Primarily manage inter-application interference in only one particular resource
 - Shared Cache, Memory Controller, Interconnect, etc.
 - Combining techniques for the different resources can result in negative interaction
- Approaches that coordinate interaction among techniques for different resources require complex implementations

Our Goal: Enable fair sharing of the **entire memory system** by dynamically detecting and controlling interference in a **coordinated manner**

Our Approach

Our Approach

- Manage inter-application interference at the **cores**, **not** at the **shared resources**

Our Approach

- Manage inter-application interference at the **cores**, **not** at the **shared resources**
- **Dynamically estimate unfairness** in the memory system

Our Approach

- Manage inter-application interference at the **cores**, **not** at the **shared resources**
- **Dynamically estimate unfairness** in the memory system
- If unfairness $>$ system-software-specified target then
throttle down core causing unfairness &
throttle up core that was unfairly treated

Unmanaged
Interference

A:

B:

Fair Source
Throttling

A:

B:

queue of requests to
shared resources

Unmanaged
Interference

A:

B:

Oldest →

Shared Memory
Resources

Fair Source
Throttling

A:

B:

queue of requests to
shared resources

Unmanaged
Interference

A: Compute

B: Compute

Oldest →

Shared Memory
Resources

Fair Source
Throttling

A:

B:

queue of requests to
shared resources

Request Generation Order:
A1,A2,A3,A4, B1

Unmanaged
Interference

A: Compute

B: Compute

Oldest →

Shared Memory
Resources

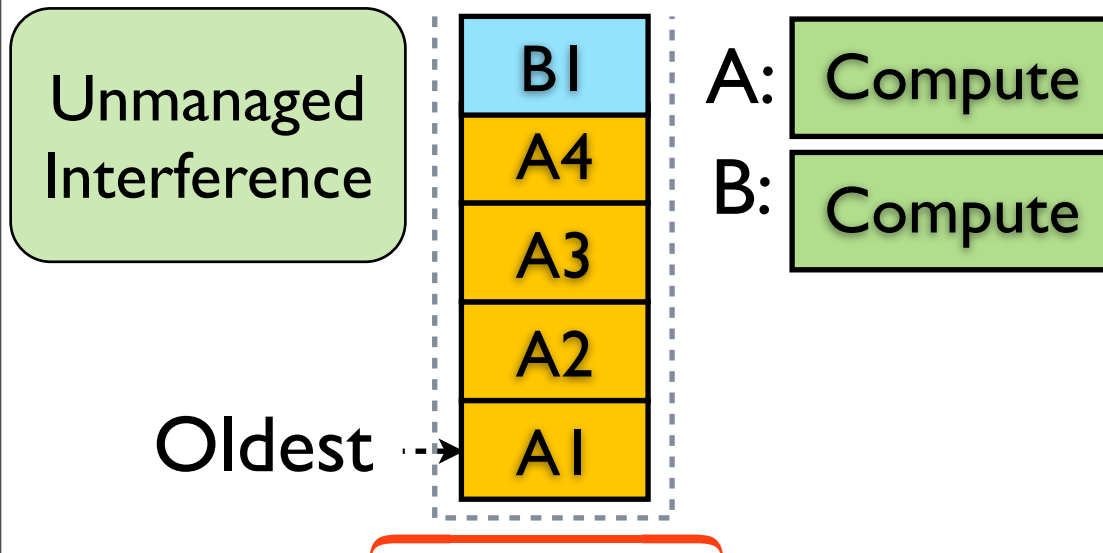
Fair Source
Throttling

A:


B:

queue of requests to
shared resources

Request Generation Order:
A1,A2,A3,A4, B1


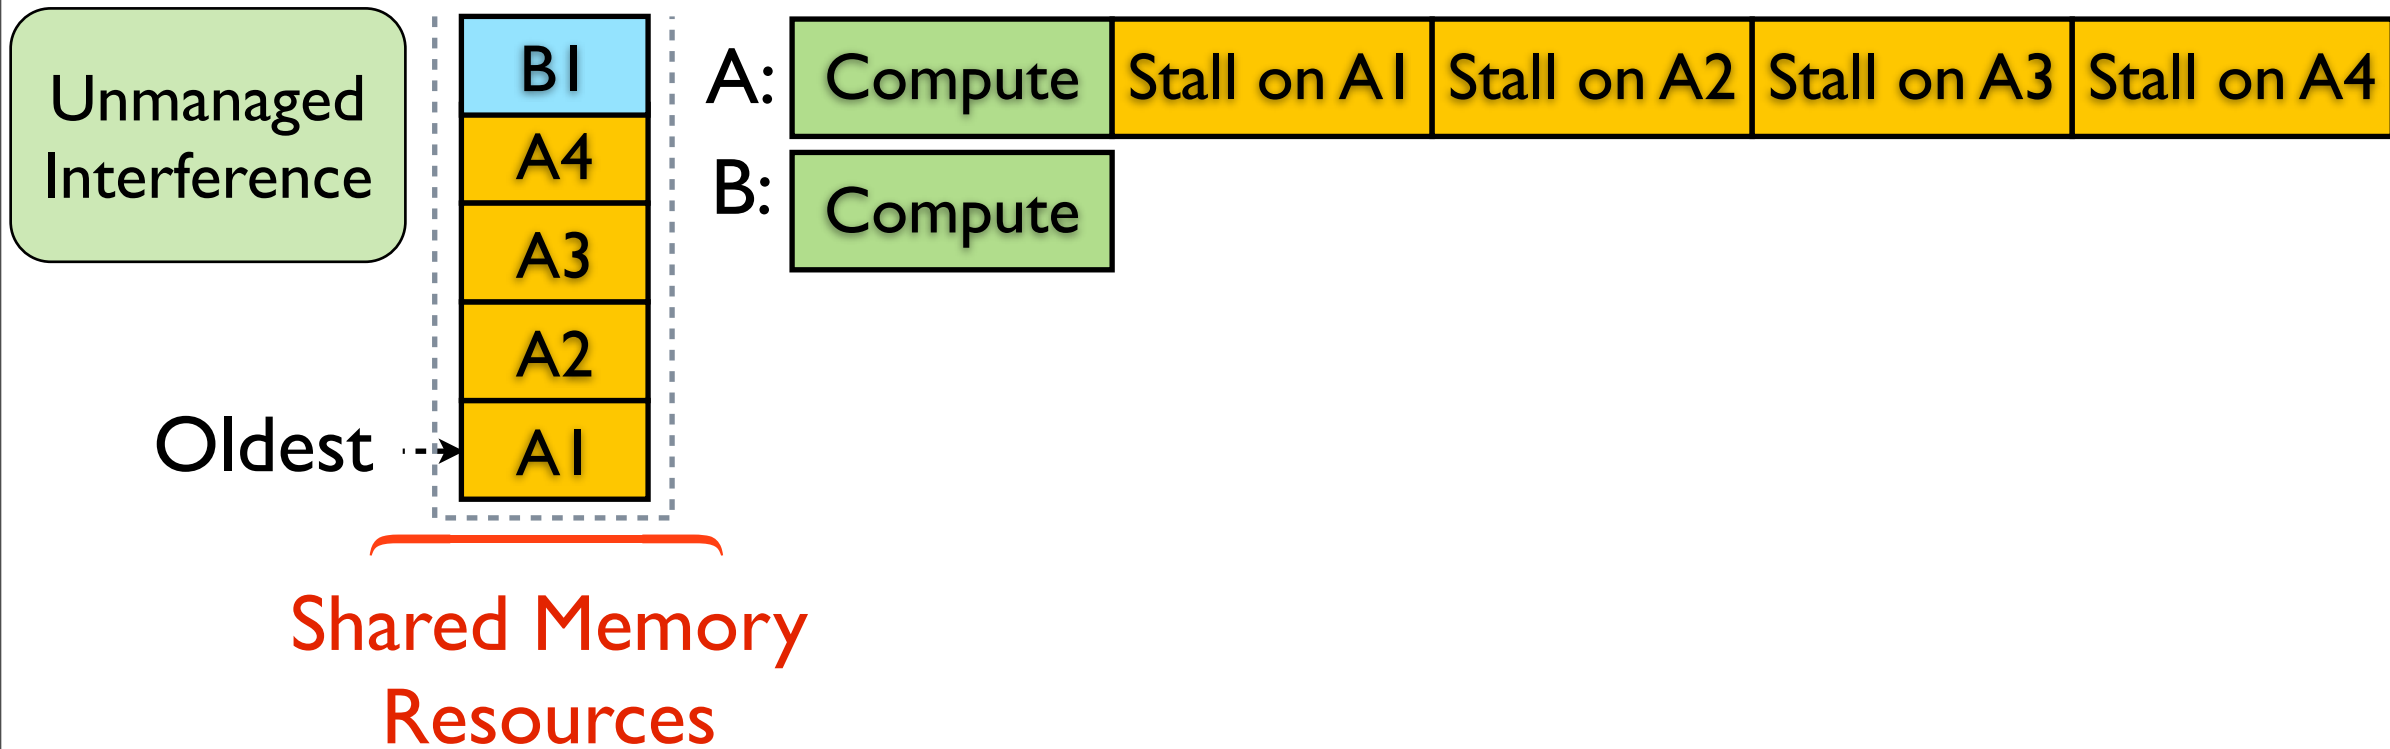


Shared Memory
Resources




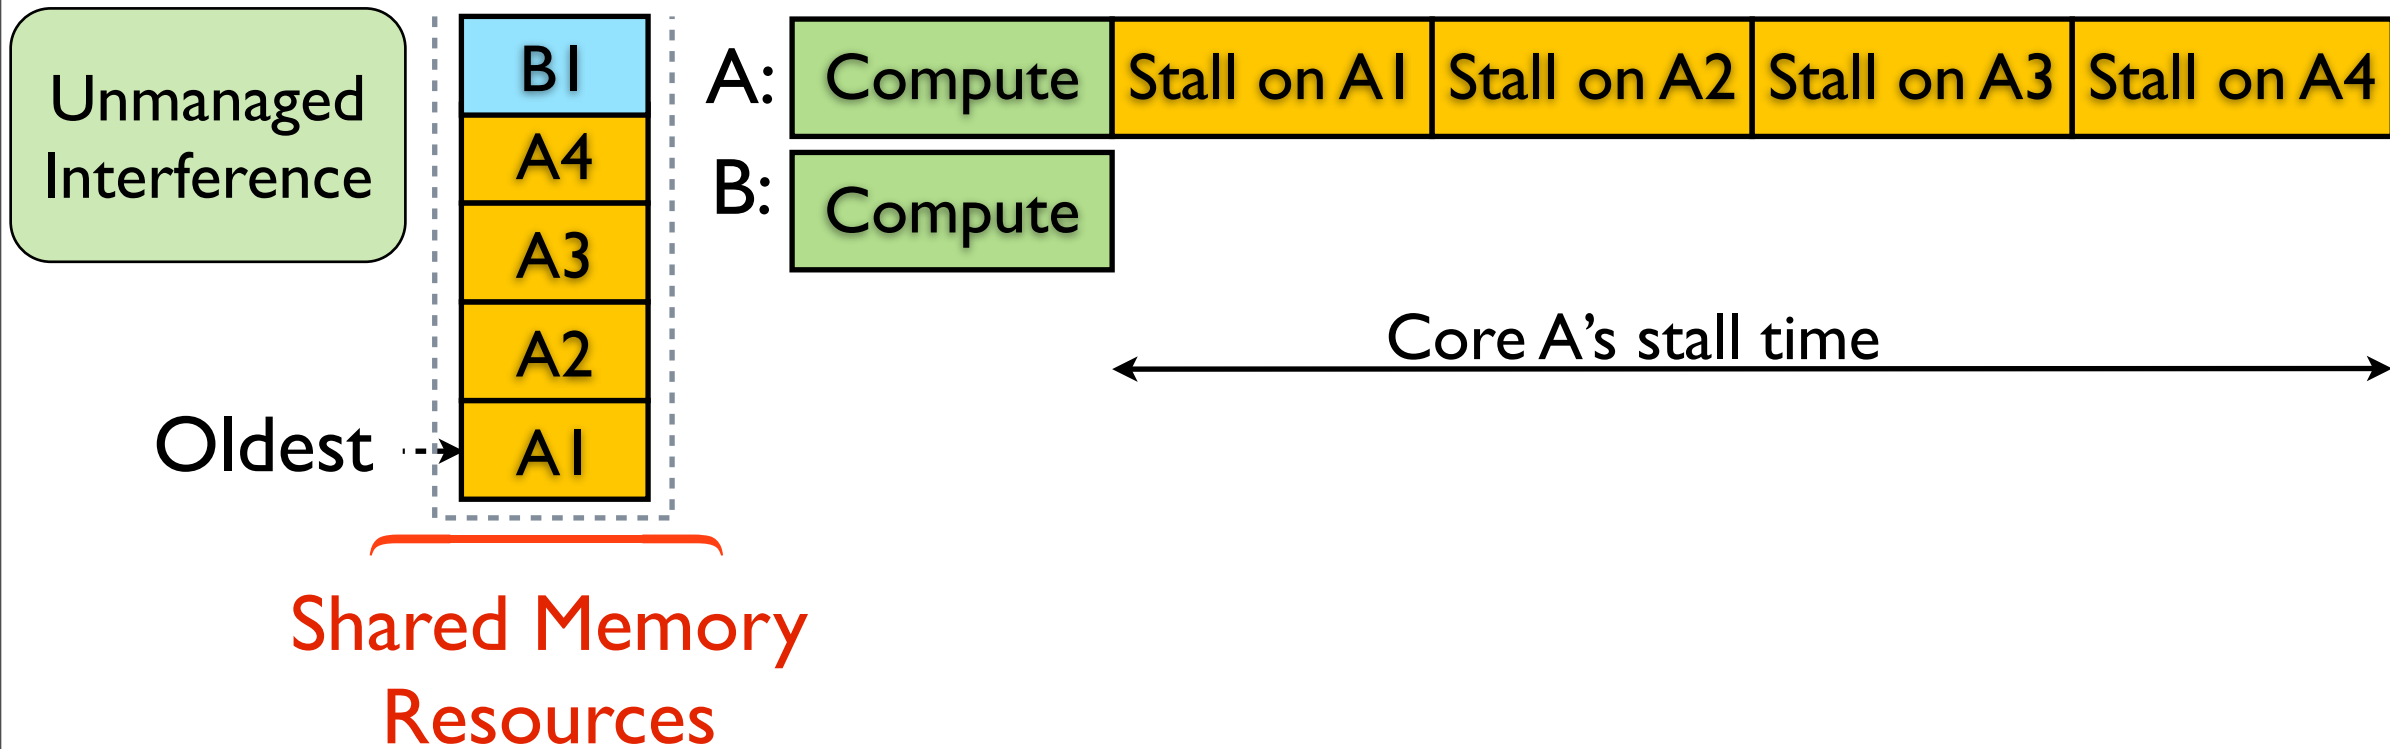
queue of requests to
shared resources

Request Generation Order:
A1,A2,A3,A4, B1




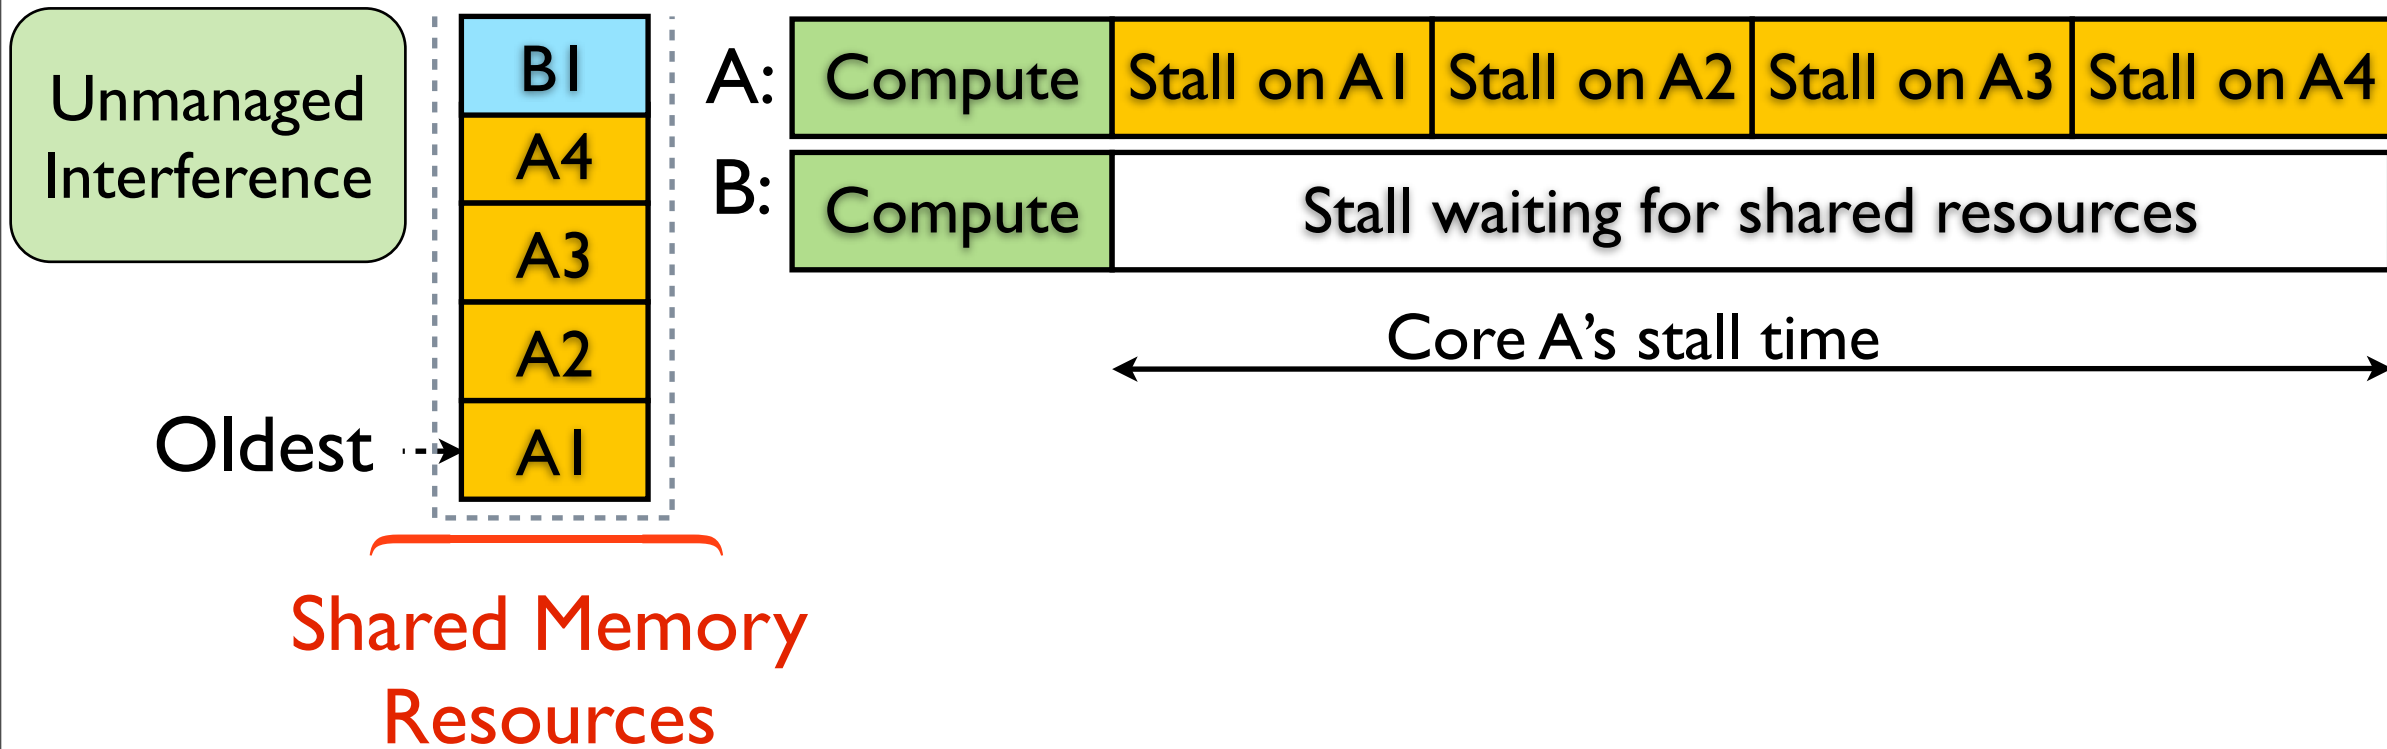
queue of requests to
shared resources

Request Generation Order:
A1,A2,A3,A4, B1



queue of requests to
shared resources


Request Generation Order:
A1,A2,A3,A4, B1



queue of requests to
shared resources

Request Generation Order:
A1,A2,A3,A4, B1

Unmanaged
Interference



Shared Memory
Resources



Core A's stall time

Fair Source
Throttling


A:

B:

queue of requests to
shared resources

Request Generation Order:
A1, A2, A3, A4, B1

Unmanaged
Interference



A:



B:



Core A's stall time

Core B's stall time

Fair Source
Throttling


A:

B:

queue of requests to
shared resources

Request Generation Order:
A1,A2,A3,A4, B1

Unmanaged
Interference



Shared Memory
Resources

A: Compute Stall on A1 Stall on A2 Stall on A3 Stall on A4

B: Compute Stall waiting for shared resources Stall on B1

Core A's stall time

Core B's stall time

Intensive application A generates many requests and causes long stall times for less intensive application B

Fair Source
Throttling


A:

B:

queue of requests to
shared resources

Request Generation Order:
A1,A2,A3,A4, B1

Unmanaged
Interference



A: Compute Stall on A1 Stall on A2 Stall on A3 Stall on A4

B: Compute Stall waiting for shared resources Stall on B1

Core A's stall time

Core B's stall time

Intensive application A generates many requests and causes long stall times for less intensive application B

Request Generation Order
A1,A2,A3,A4,B1

Fair Source
Throttling


A: Compute

B: Compute

queue of requests to
shared resources

Request Generation Order:
A1,A2,A3,A4, B1

Unmanaged
Interference



Oldest

Shared Memory
Resources



Core A's stall time

Core B's stall time

Intensive application A generates many requests and causes long stall times for less intensive application B

Request Generation Order
A1,A2,A3,A4,B1


Fair Source
Throttling



queue of requests to
shared resources

Request Generation Order:
A1,A2,A3,A4, B1

Unmanaged
Interference



Core A's stall time

Core B's stall time

Intensive application A generates many requests and causes long stall times for less intensive application B

Request Generation Order
A1, B1, A2,A3,A4

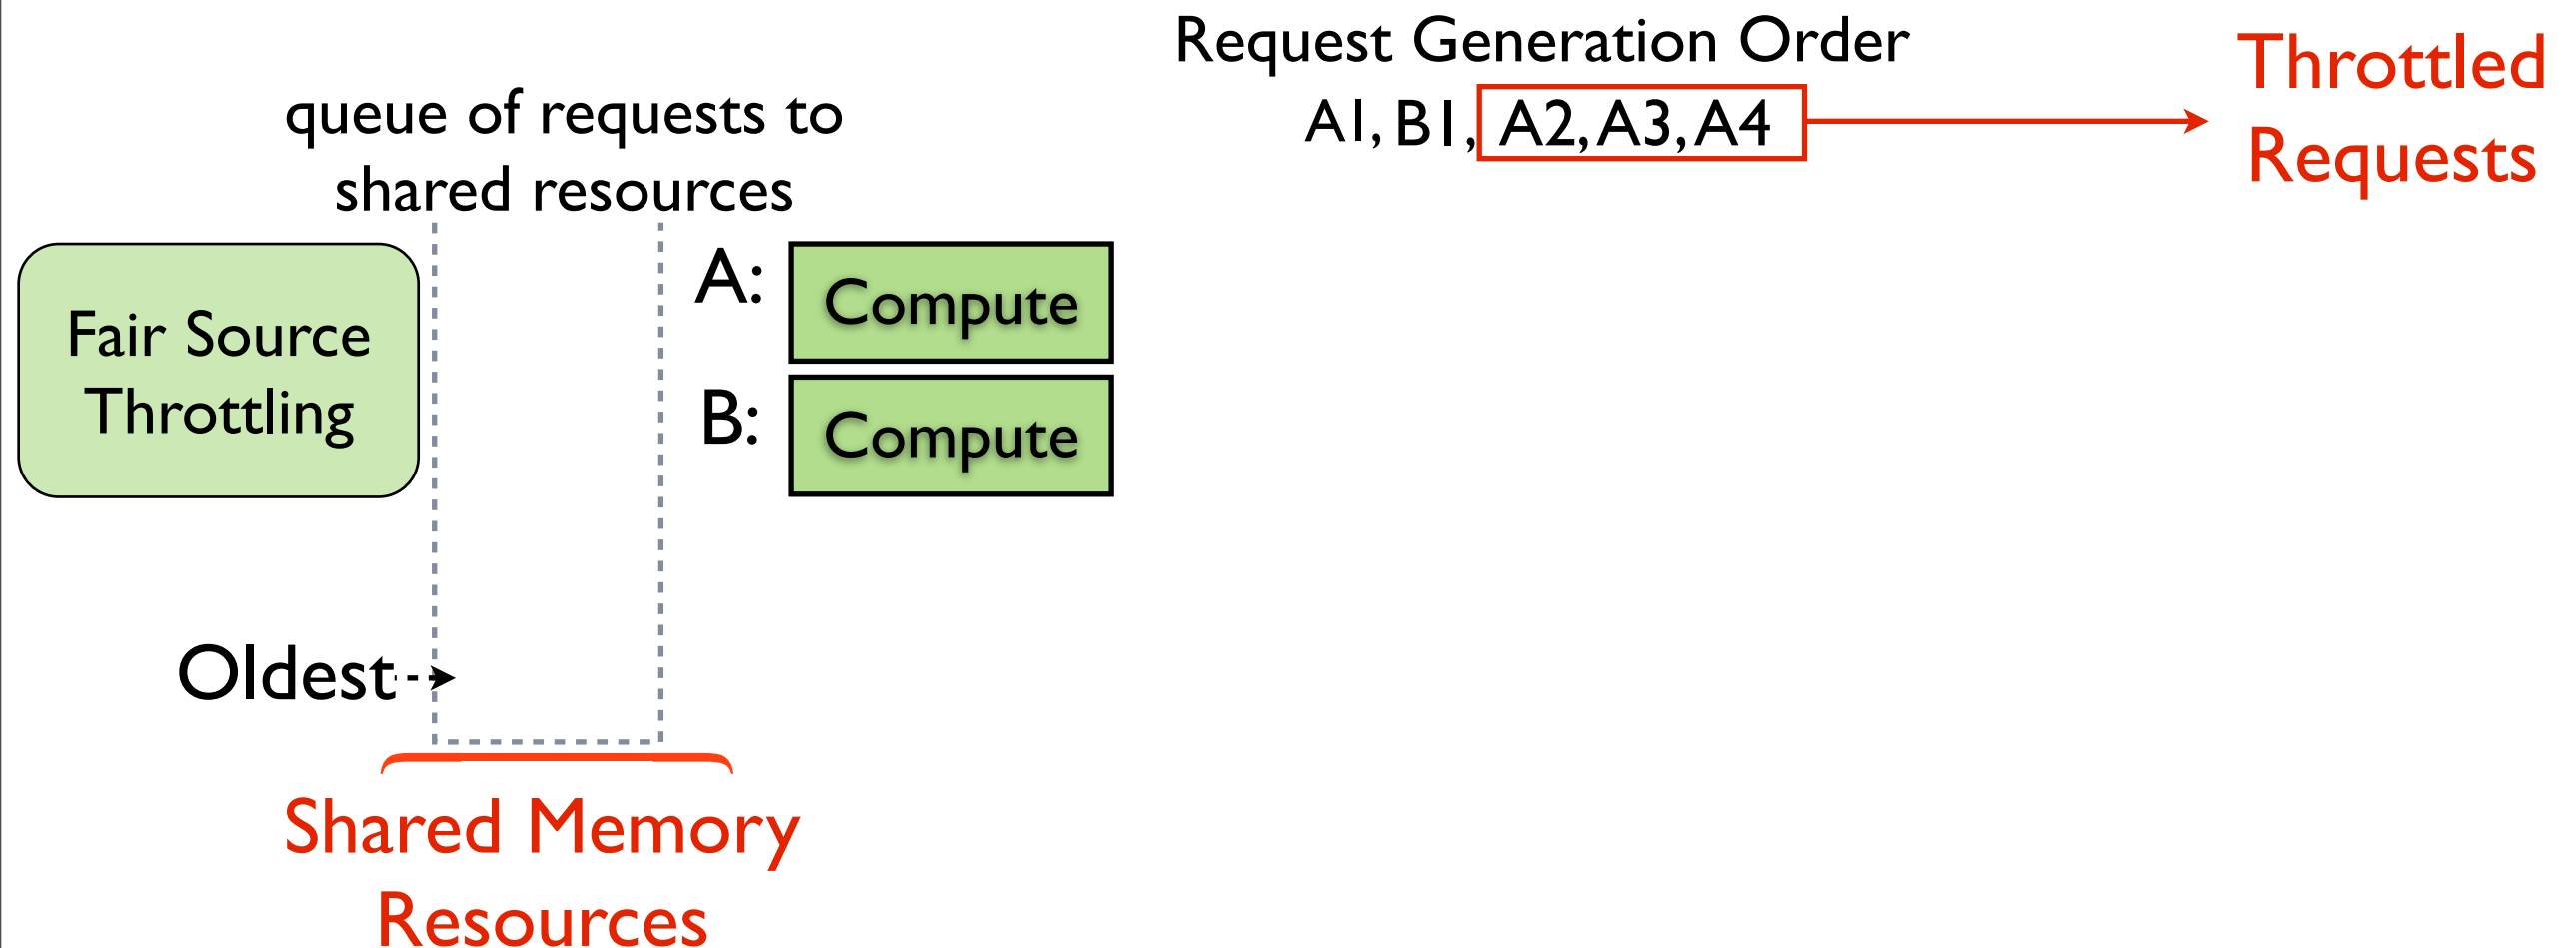
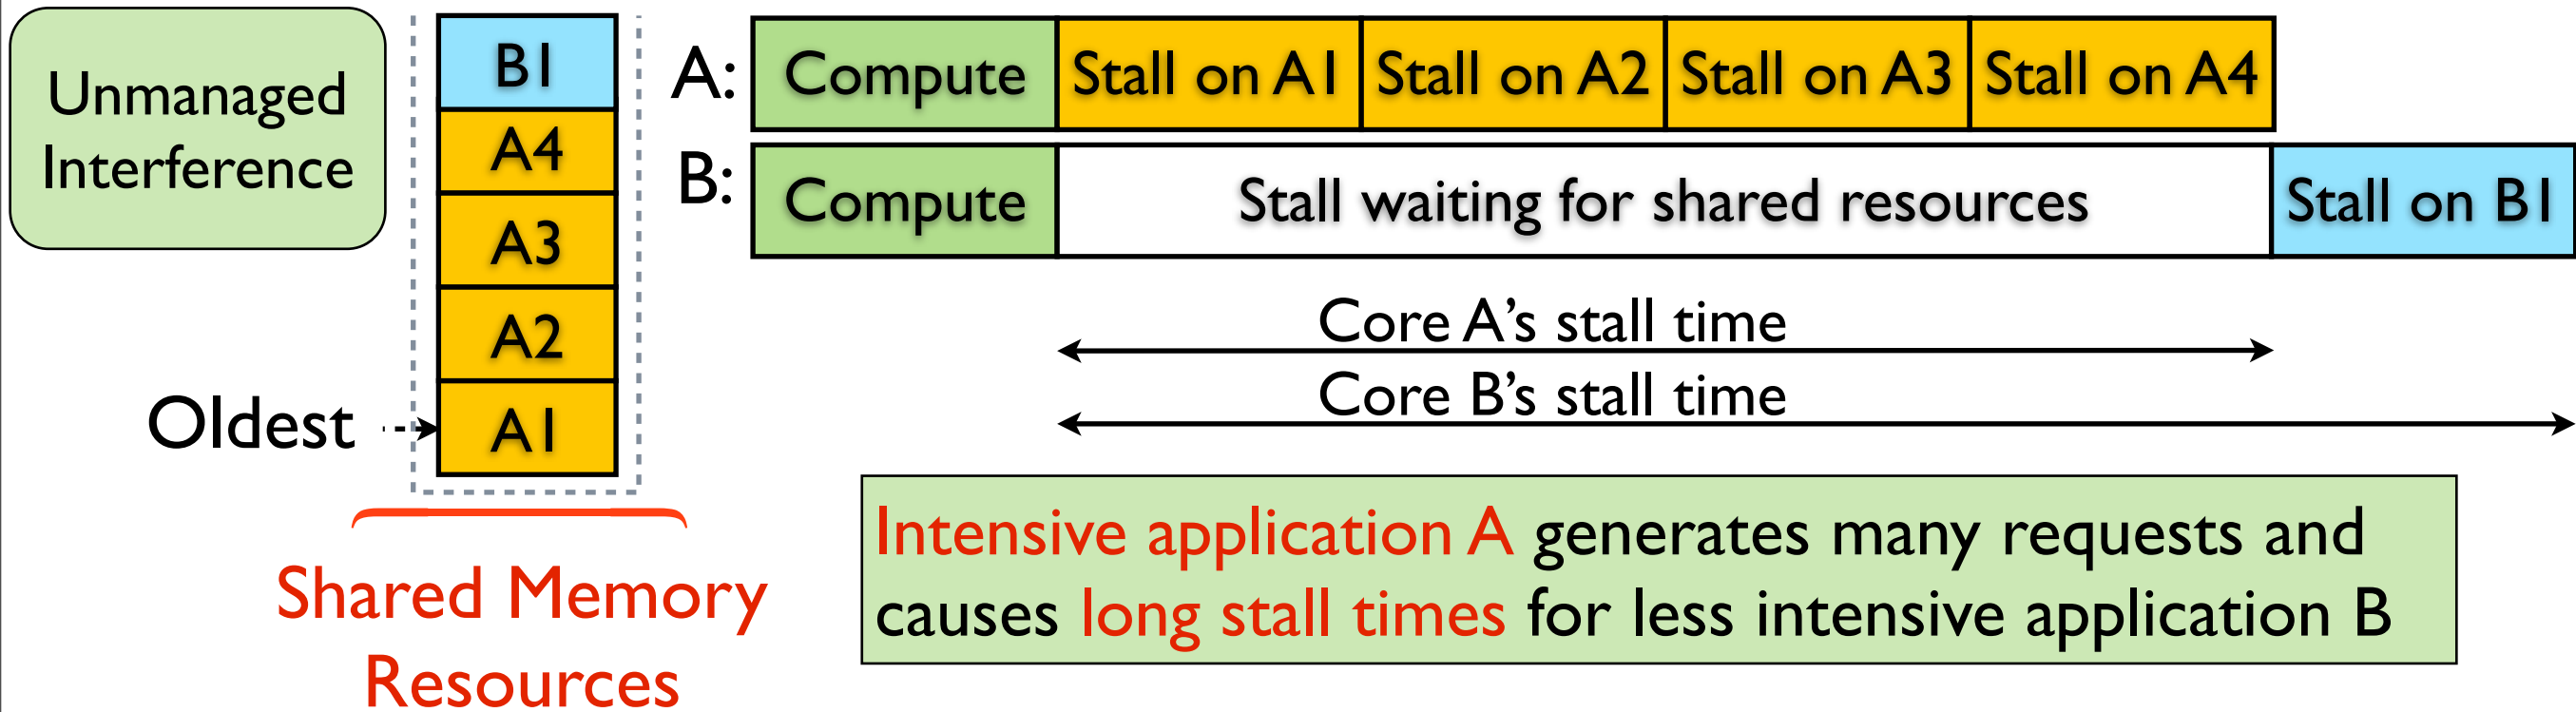
Throttled
Requests

Fair Source
Throttling



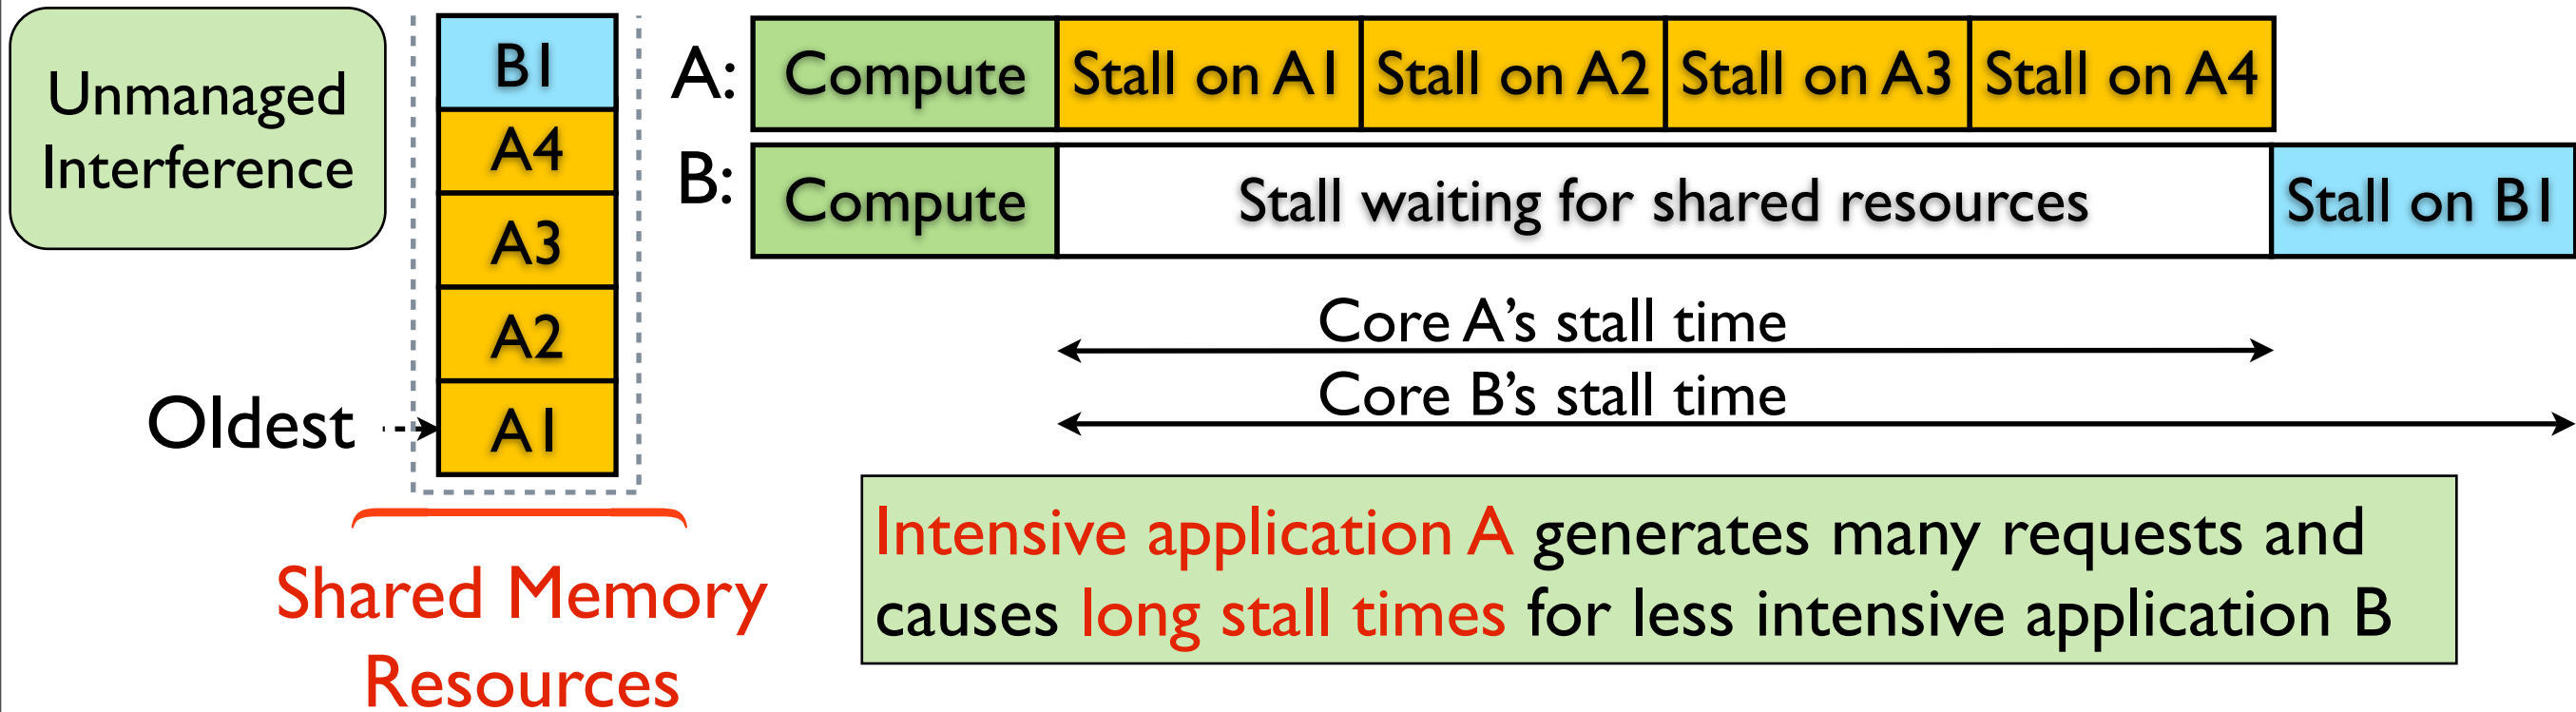
queue of requests to shared resources

Request Generation Order:
A1, A2, A3, A4, B1



queue of requests to shared resources

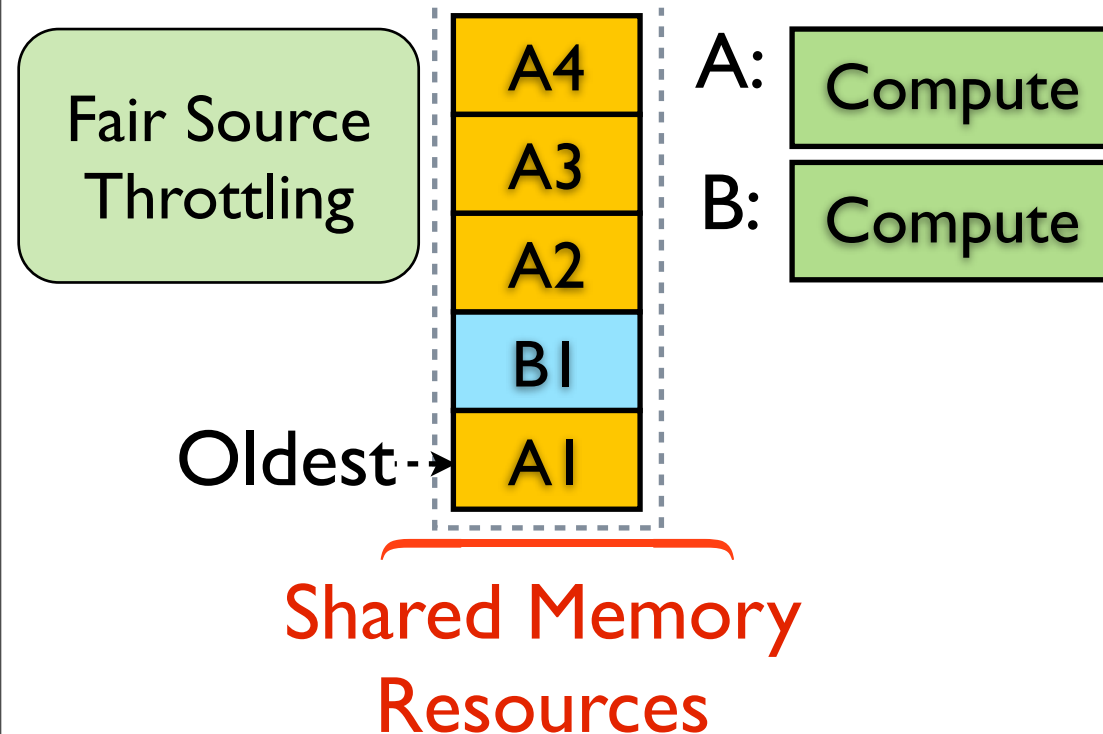
Request Generation Order:
A1, A2, A3, A4, B1



queue of requests to shared resources

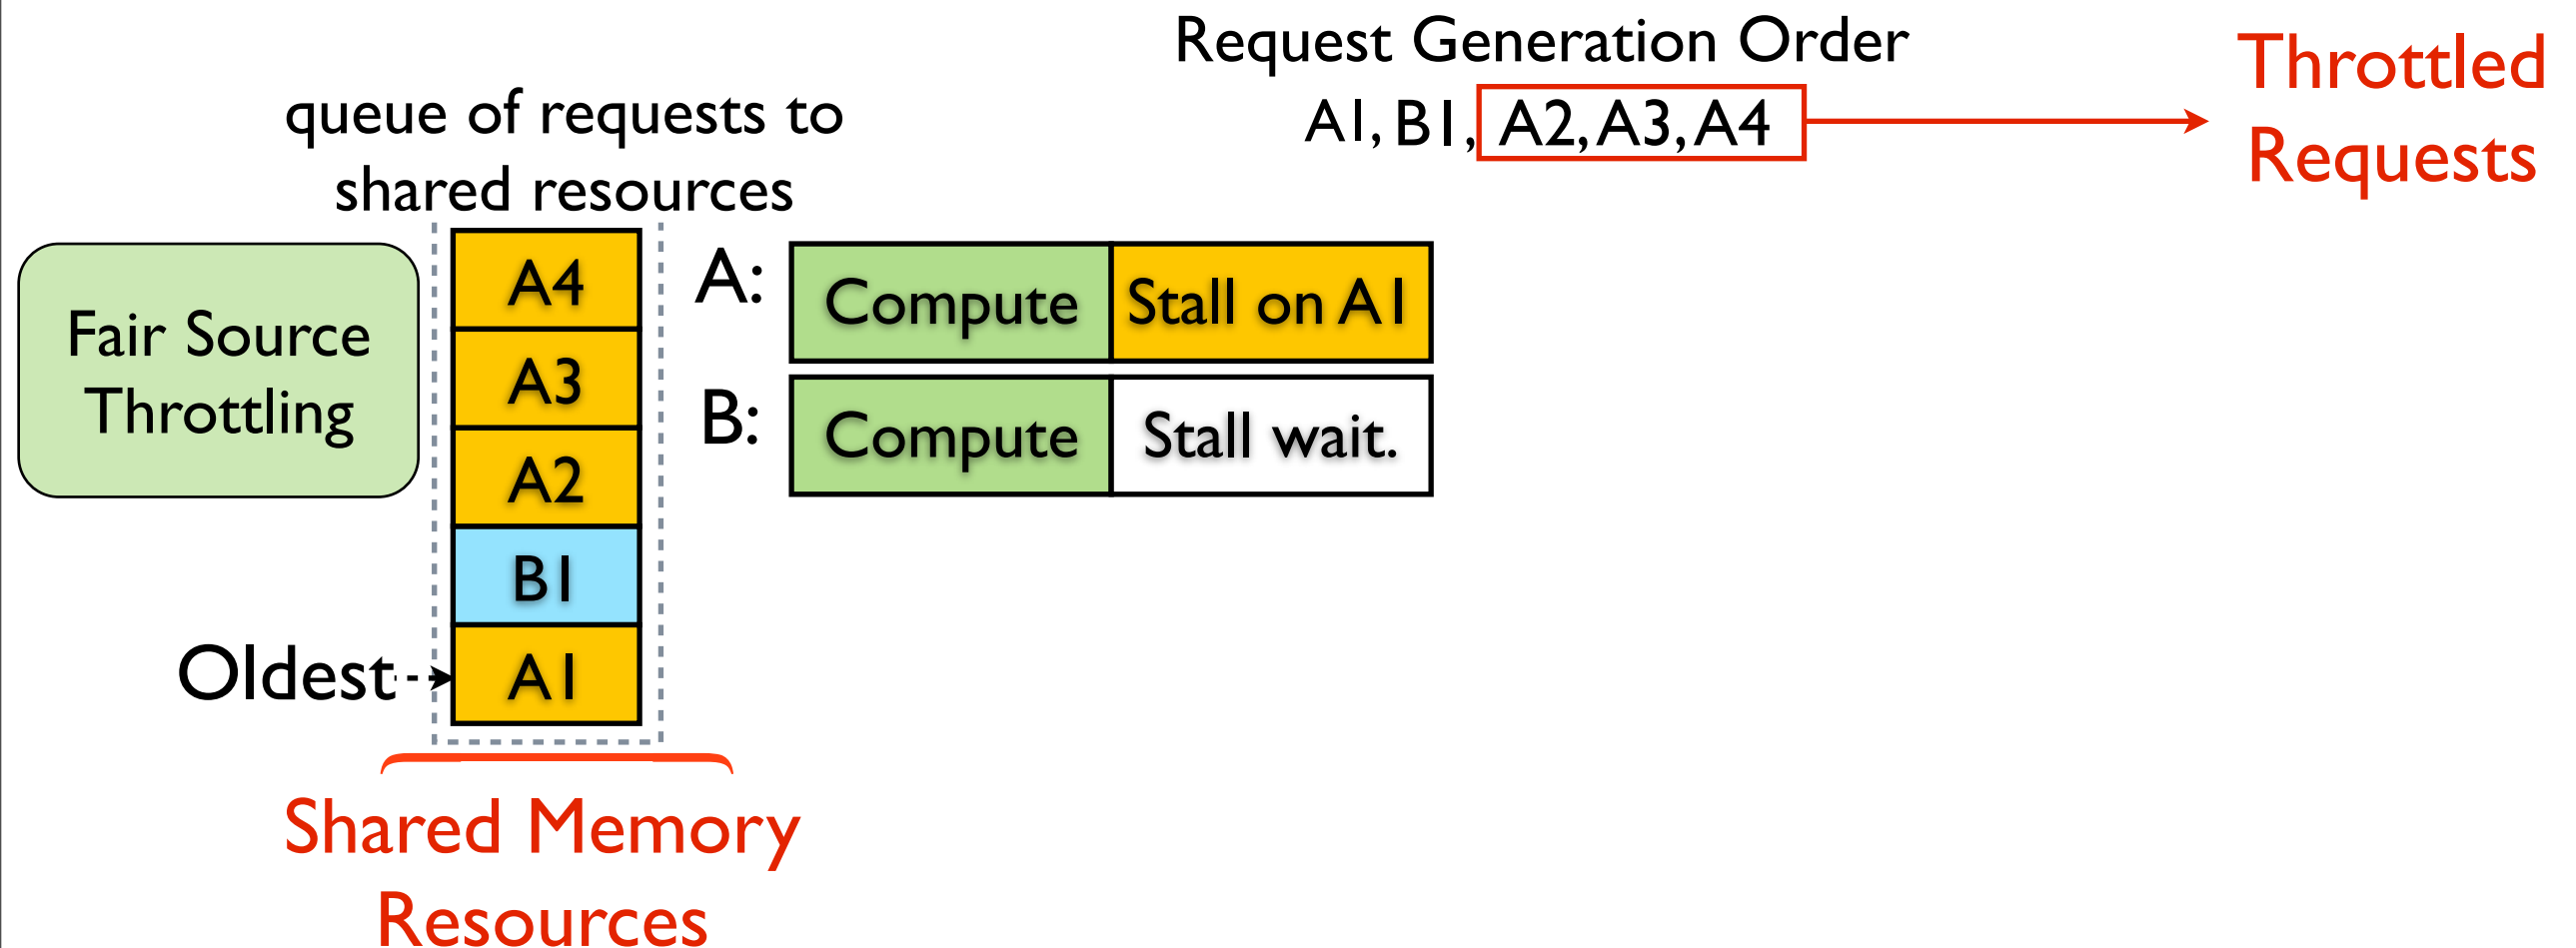
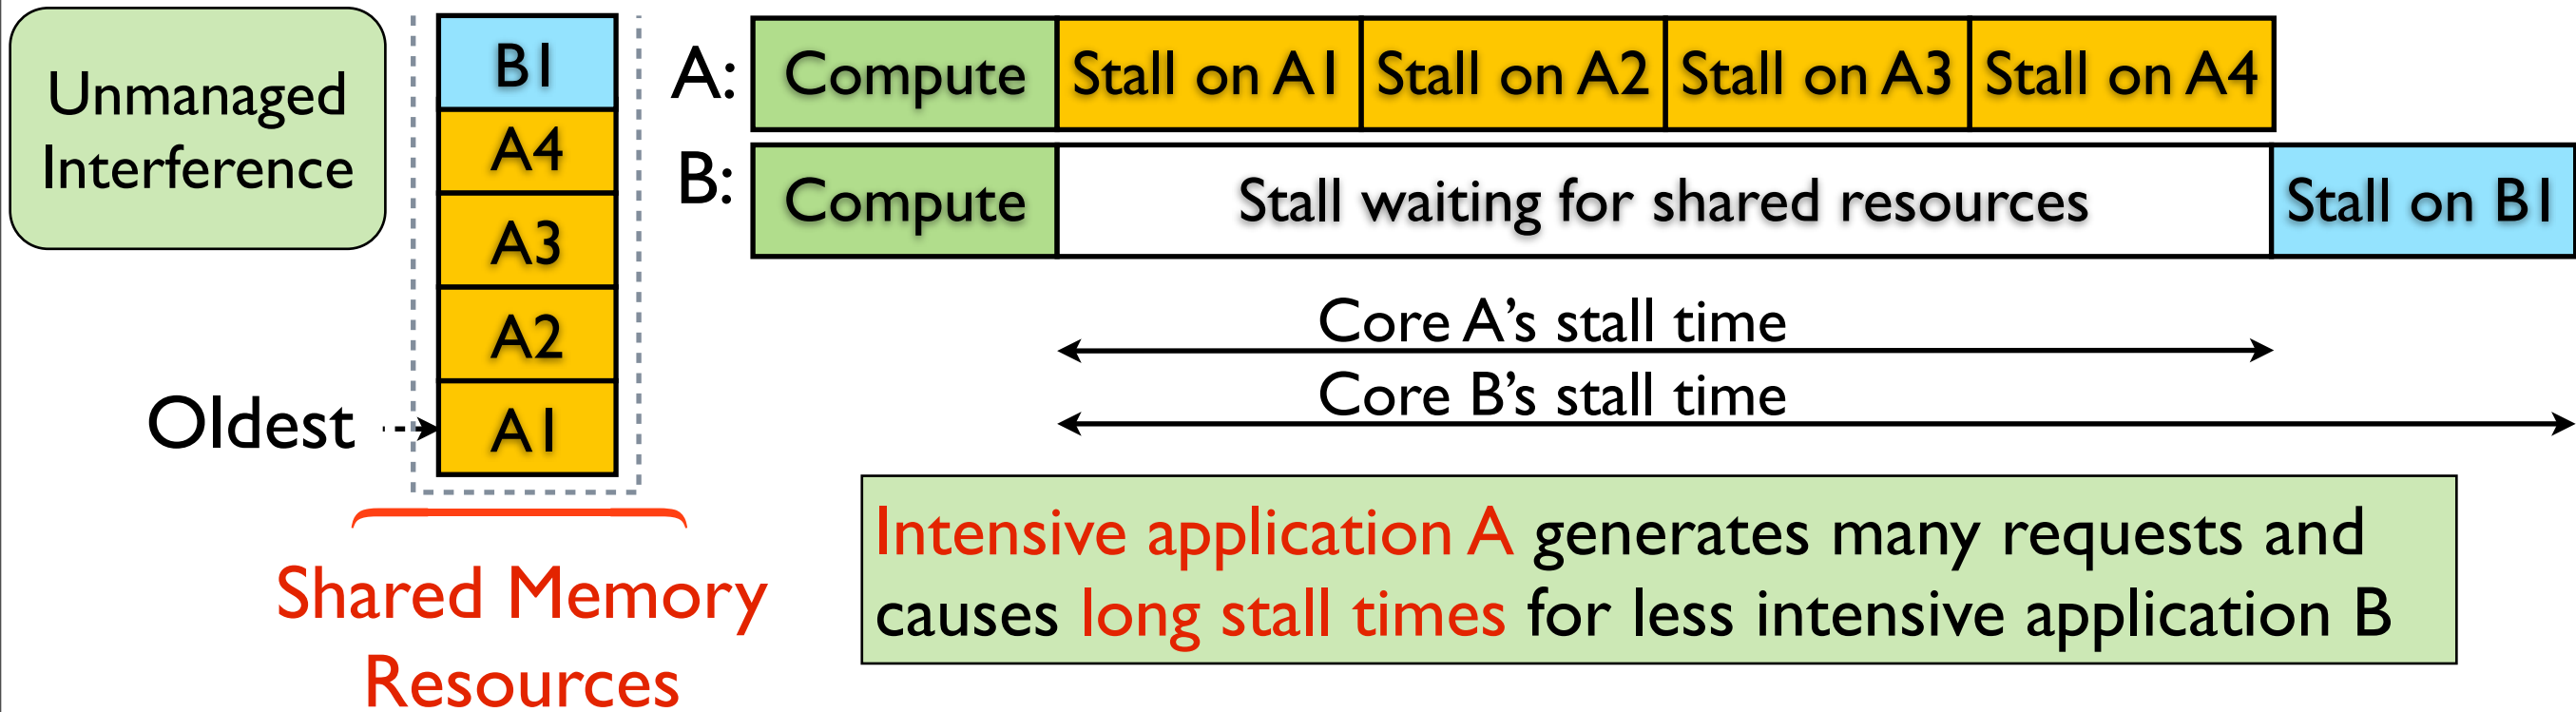
Request Generation Order
A1, B1, A2, A3, A4

Throttled Requests



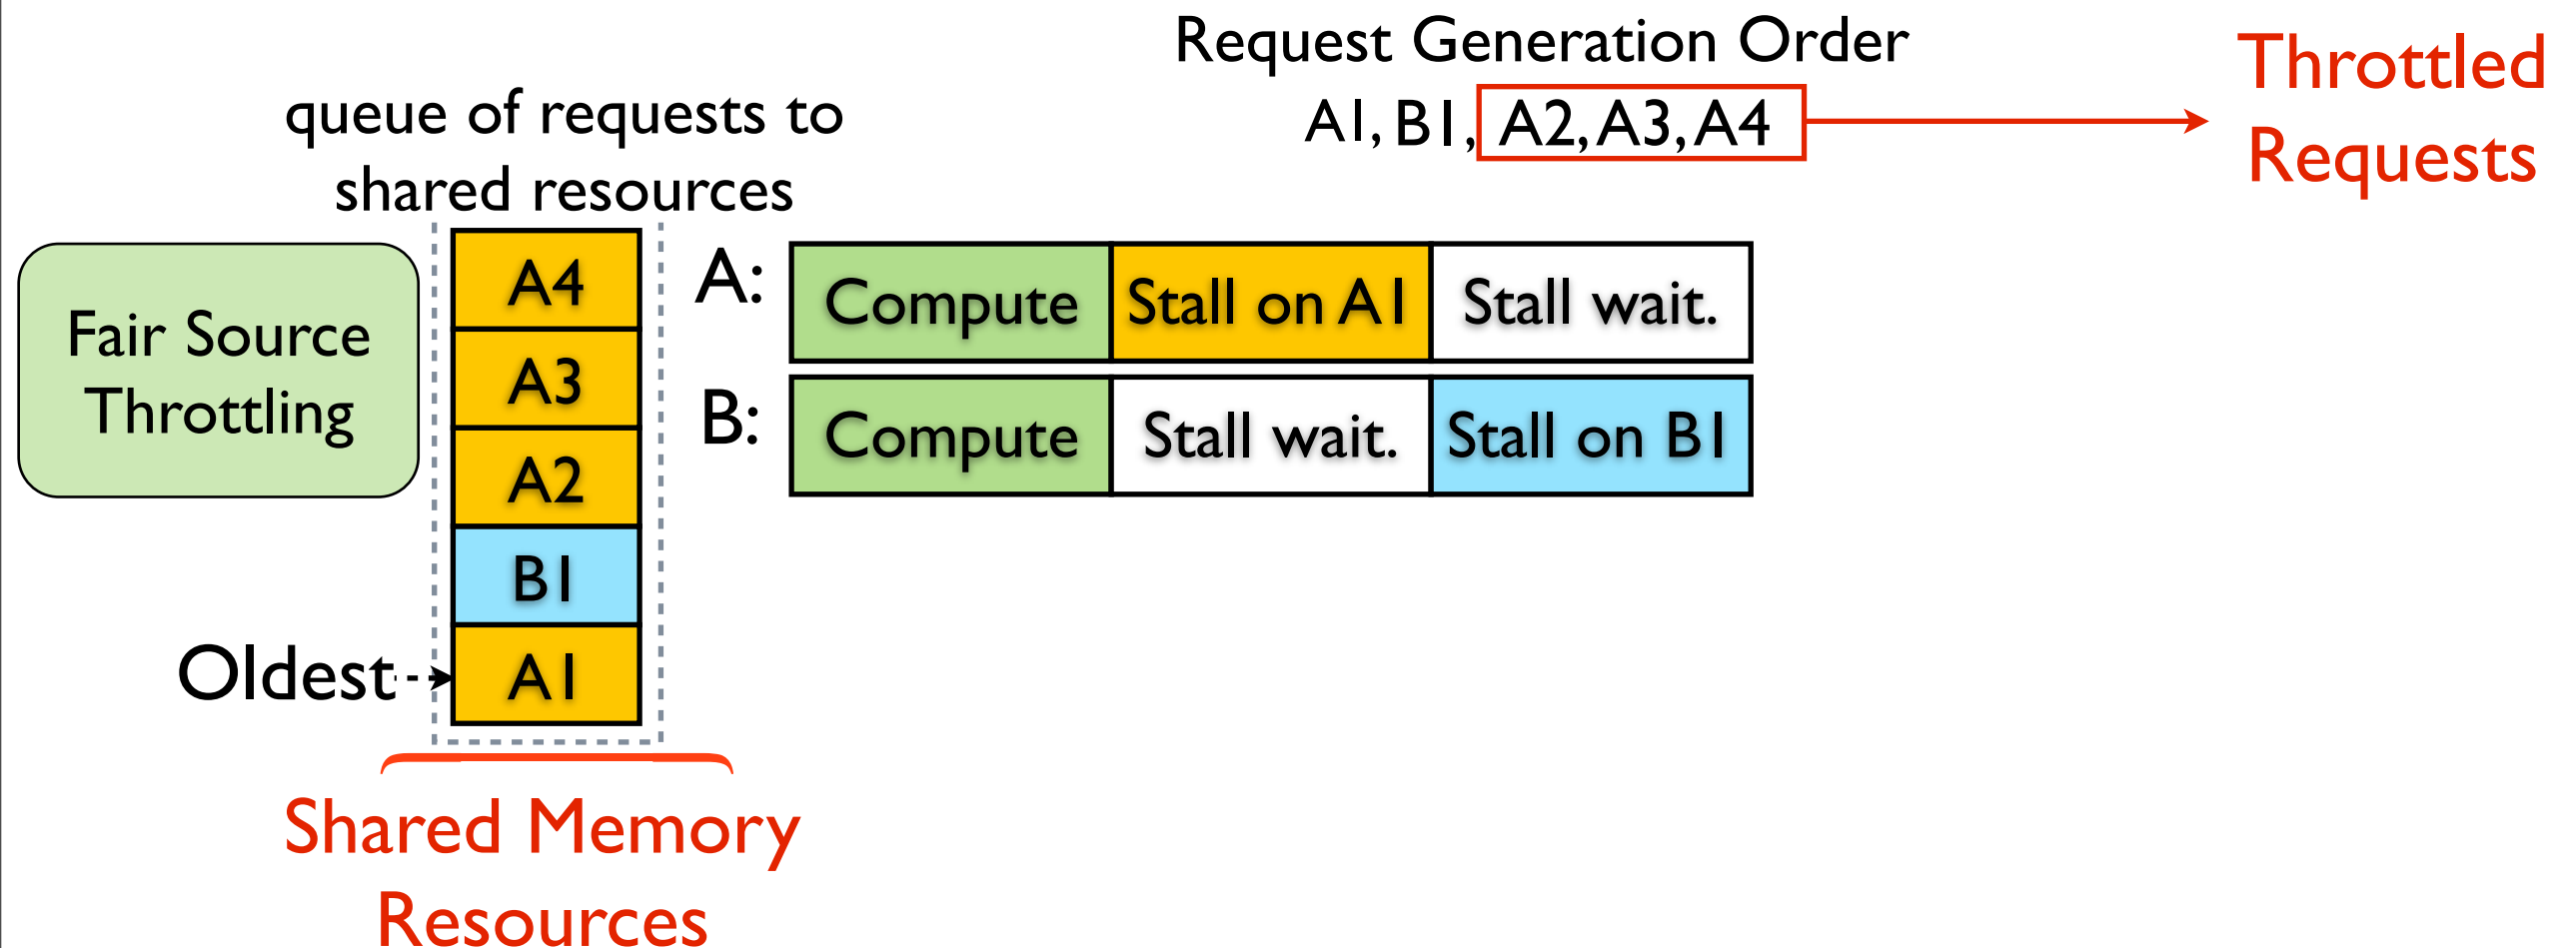
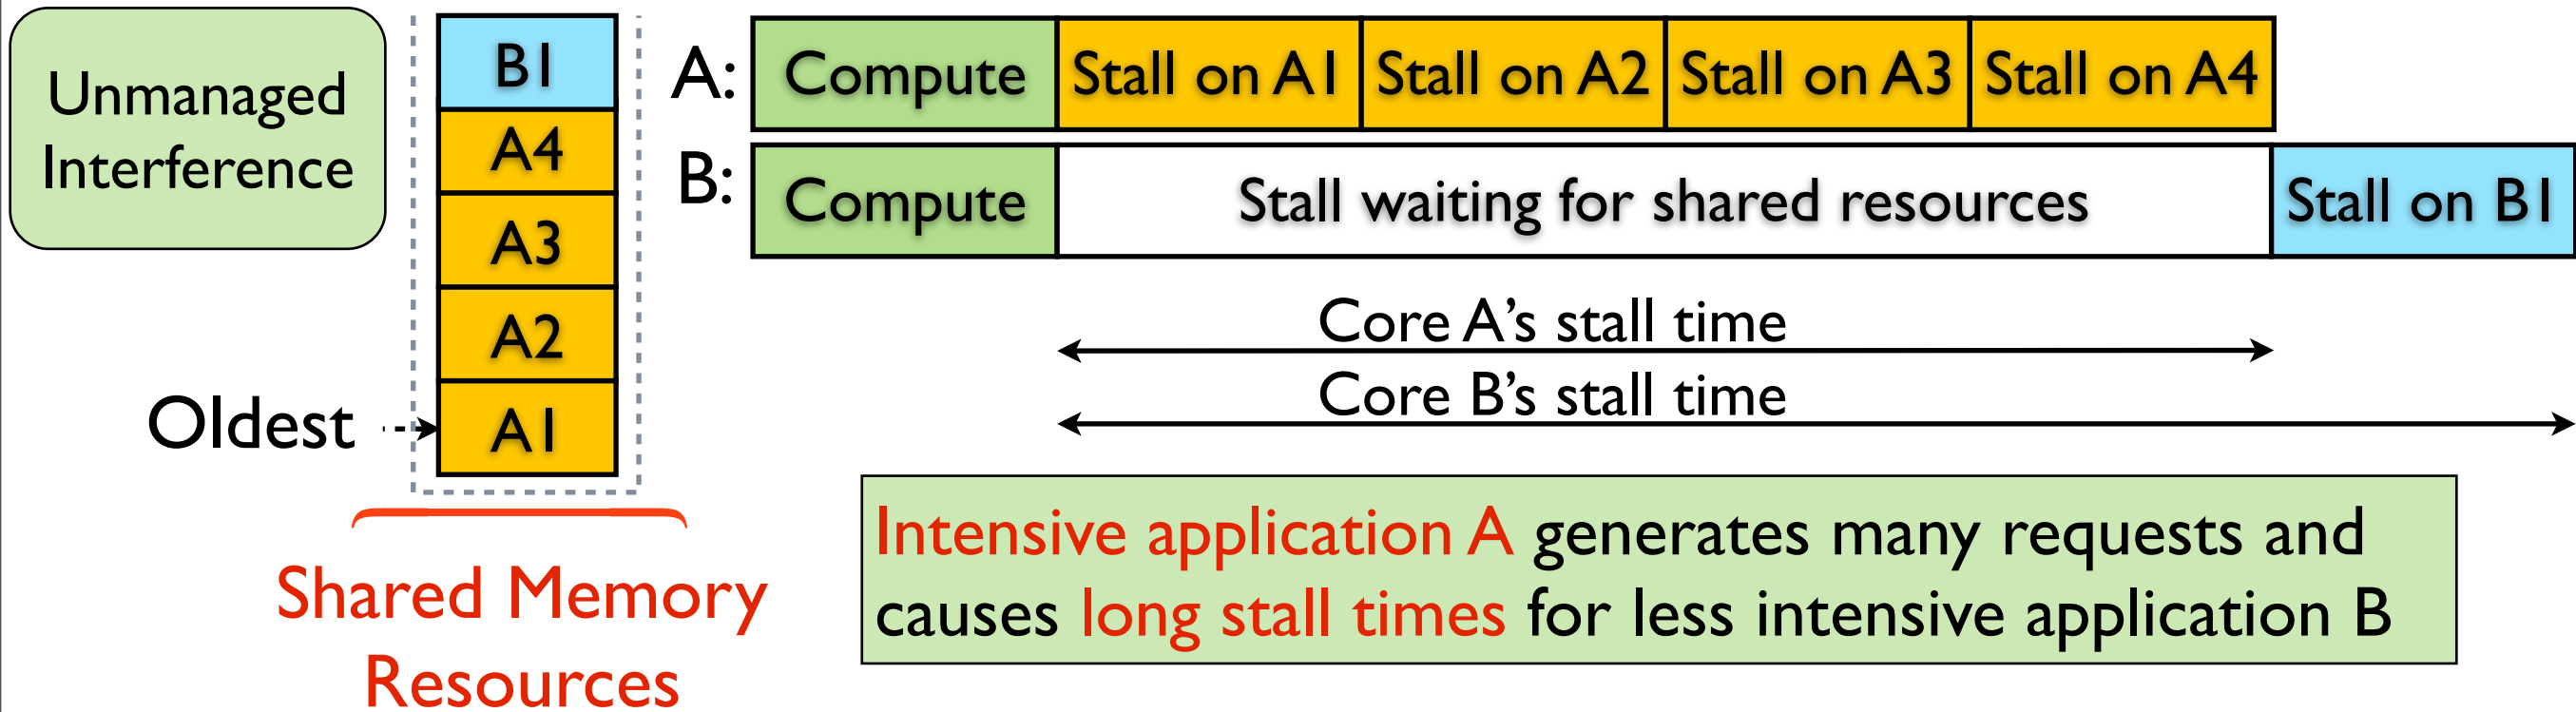
queue of requests to shared resources

Request Generation Order:
A1, A2, A3, A4, B1



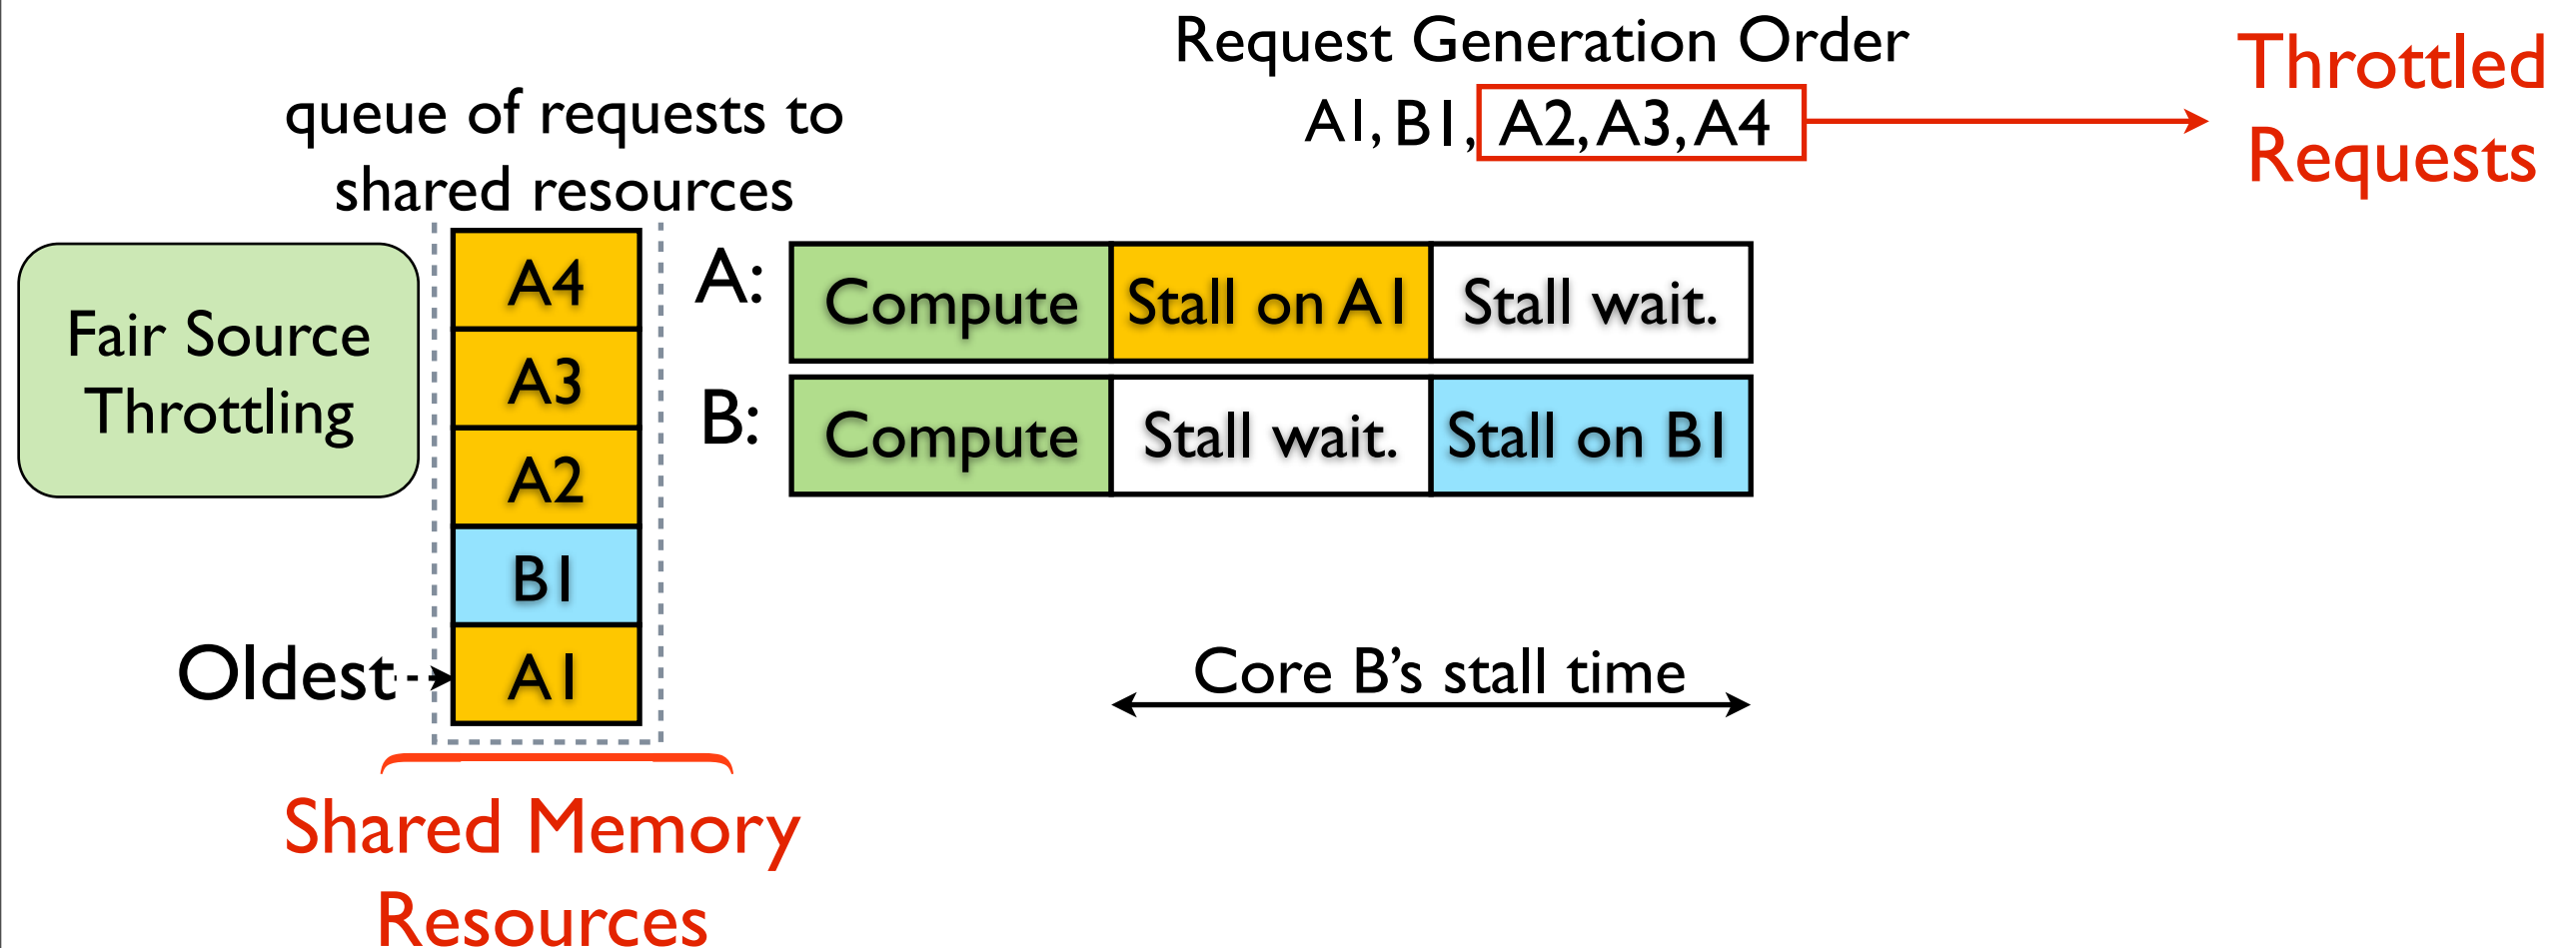
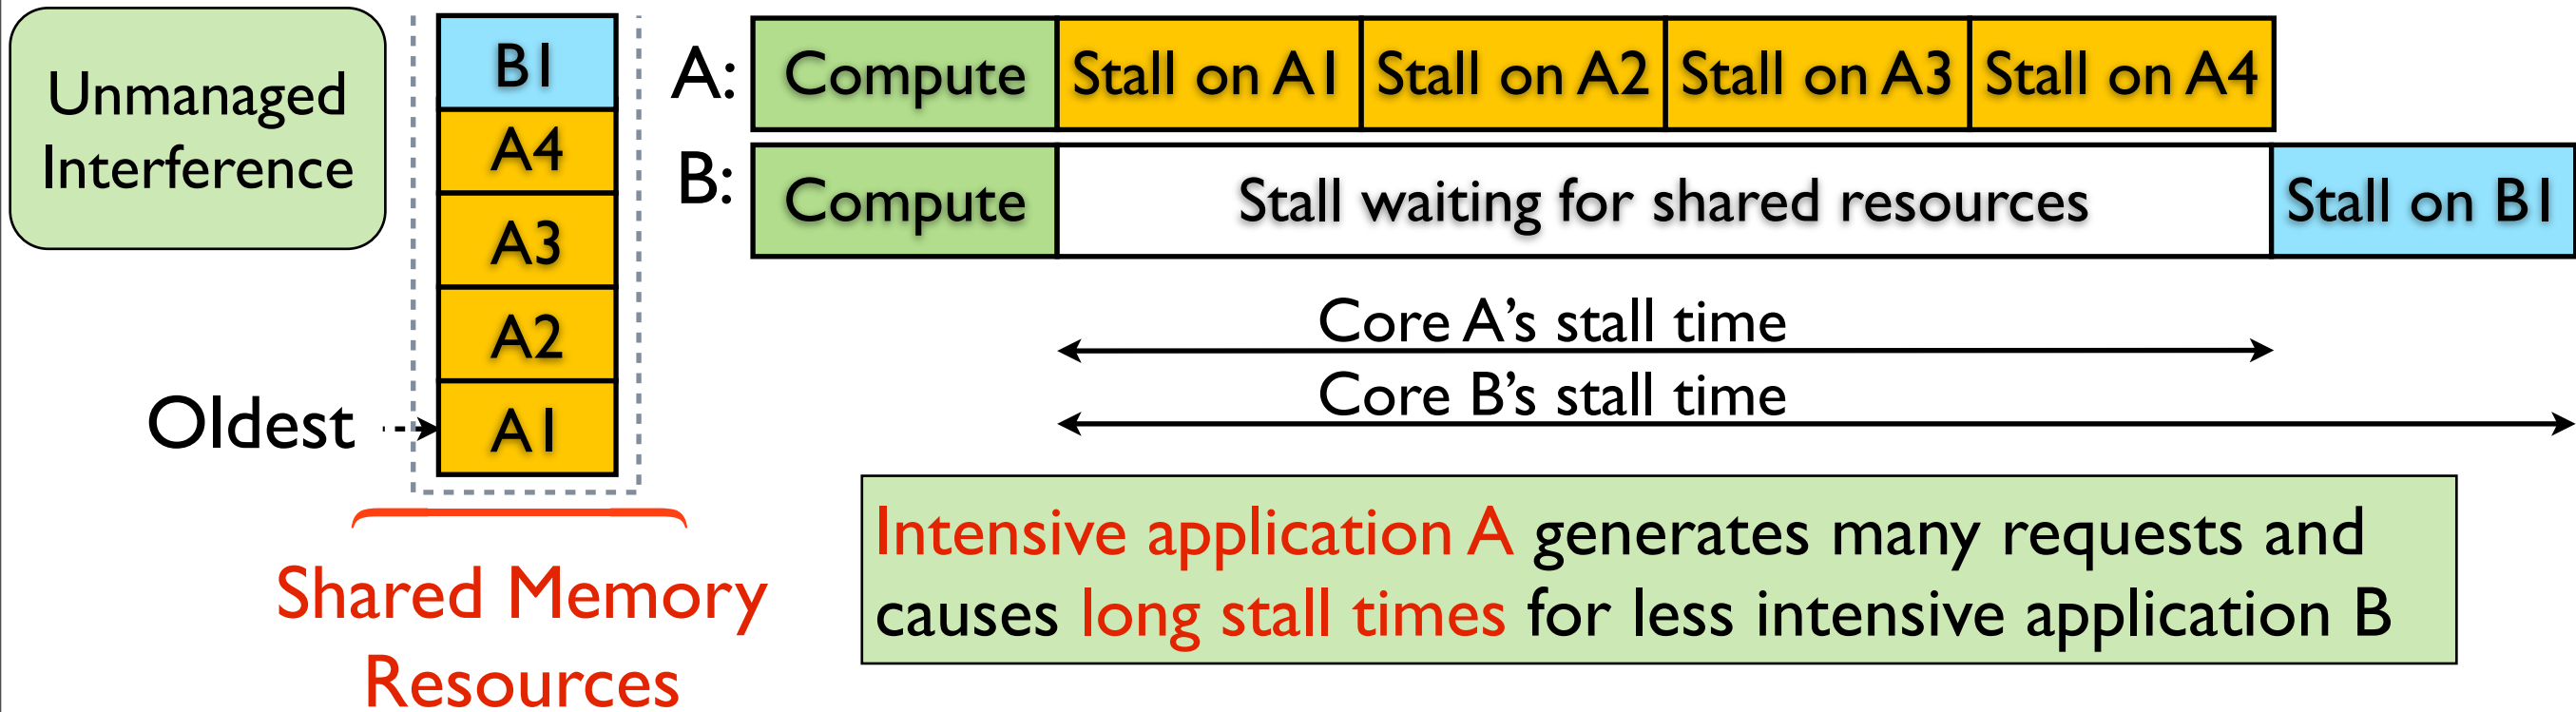
queue of requests to shared resources

Request Generation Order:
A1, A2, A3, A4, B1



queue of requests to shared resources


Request Generation Order:
A1, A2, A3, A4, B1



queue of requests to
shared resources

Request Generation Order:
A1,A2,A3,A4, B1

Unmanaged
Interference



Core A's stall time

Core B's stall time

Oldest

Shared Memory
Resources


Intensive application A generates many requests and causes long stall times for less intensive application B

queue of requests to
shared resources

Request Generation Order
A1, B1, A2,A3,A4

Throttled
Requests

Fair Source
Throttling



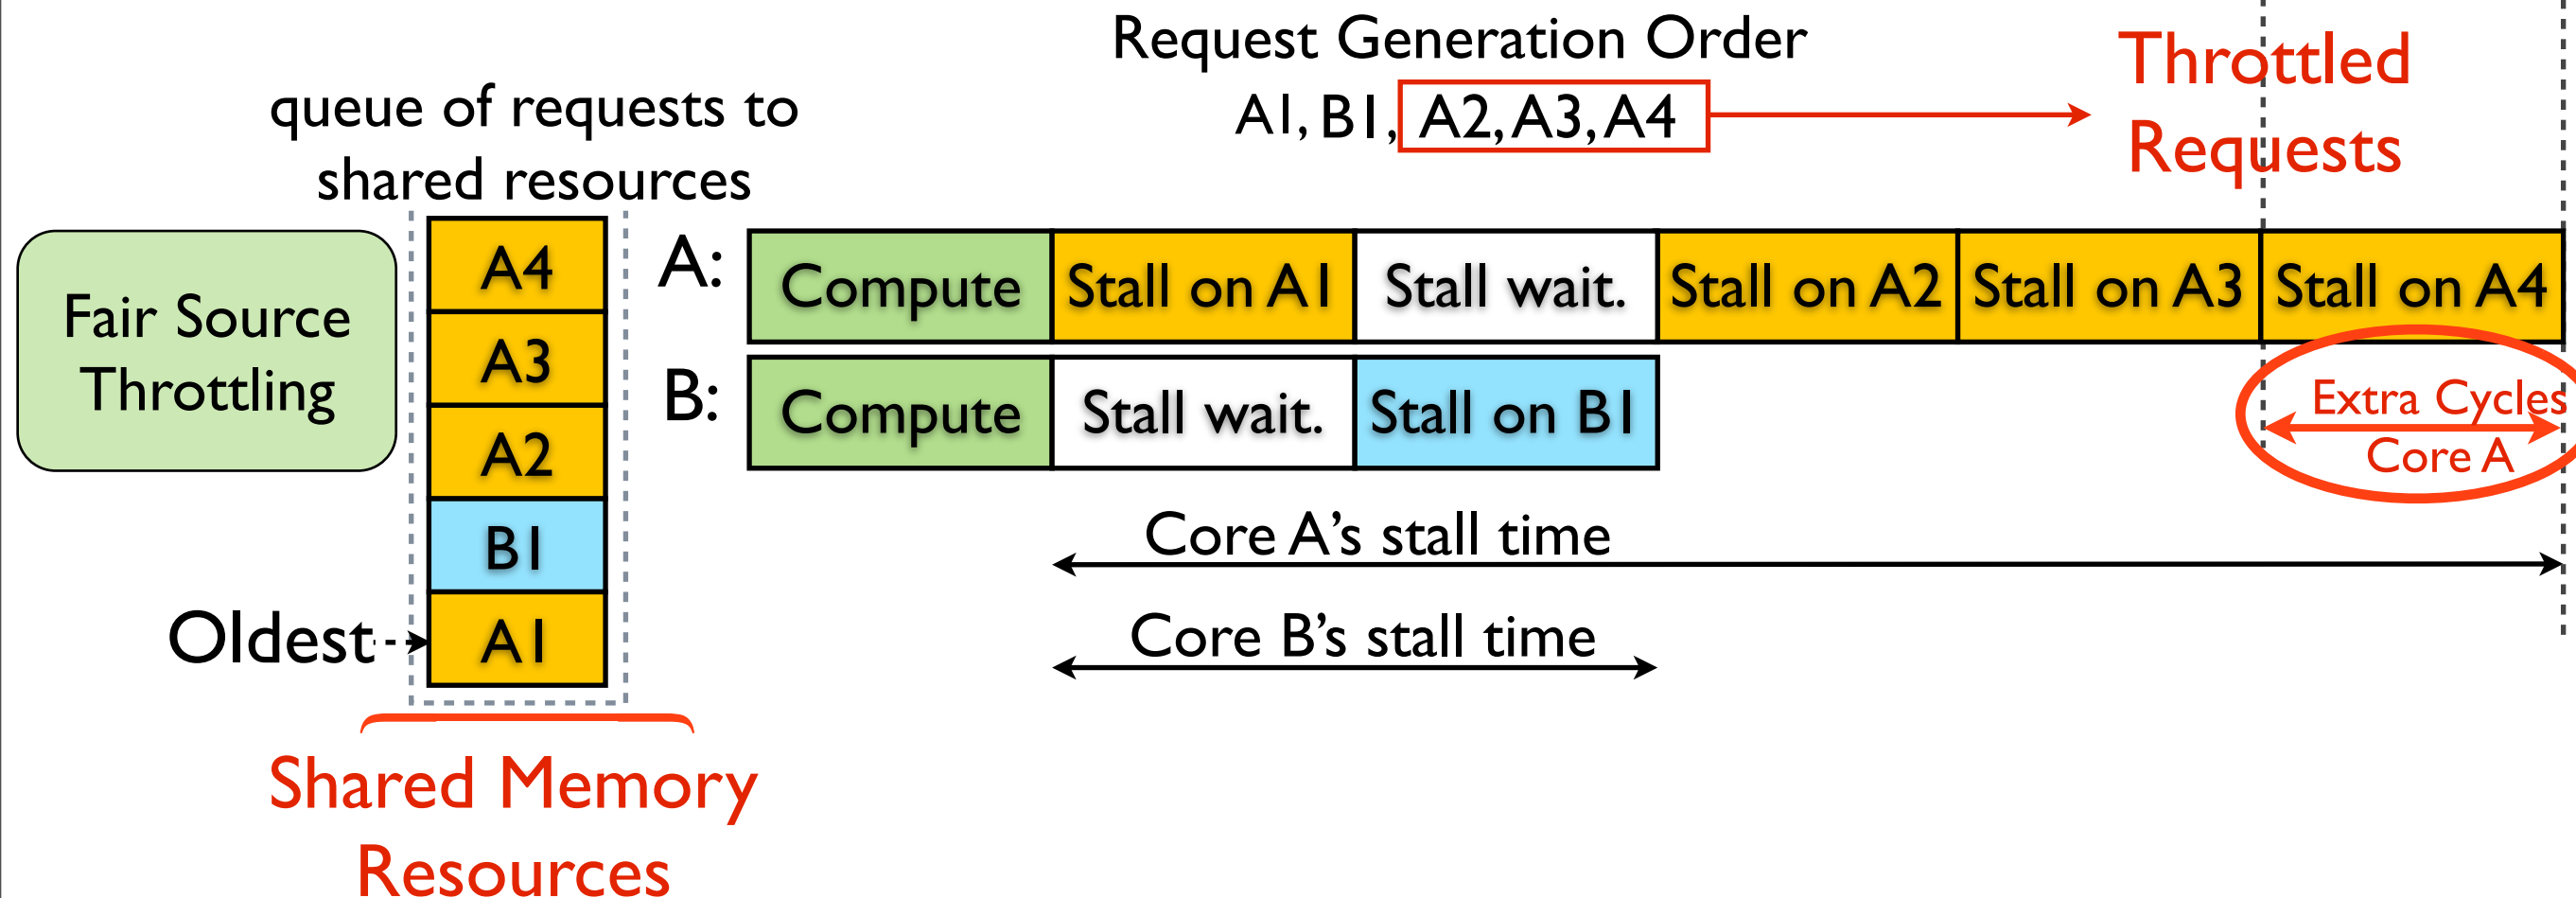
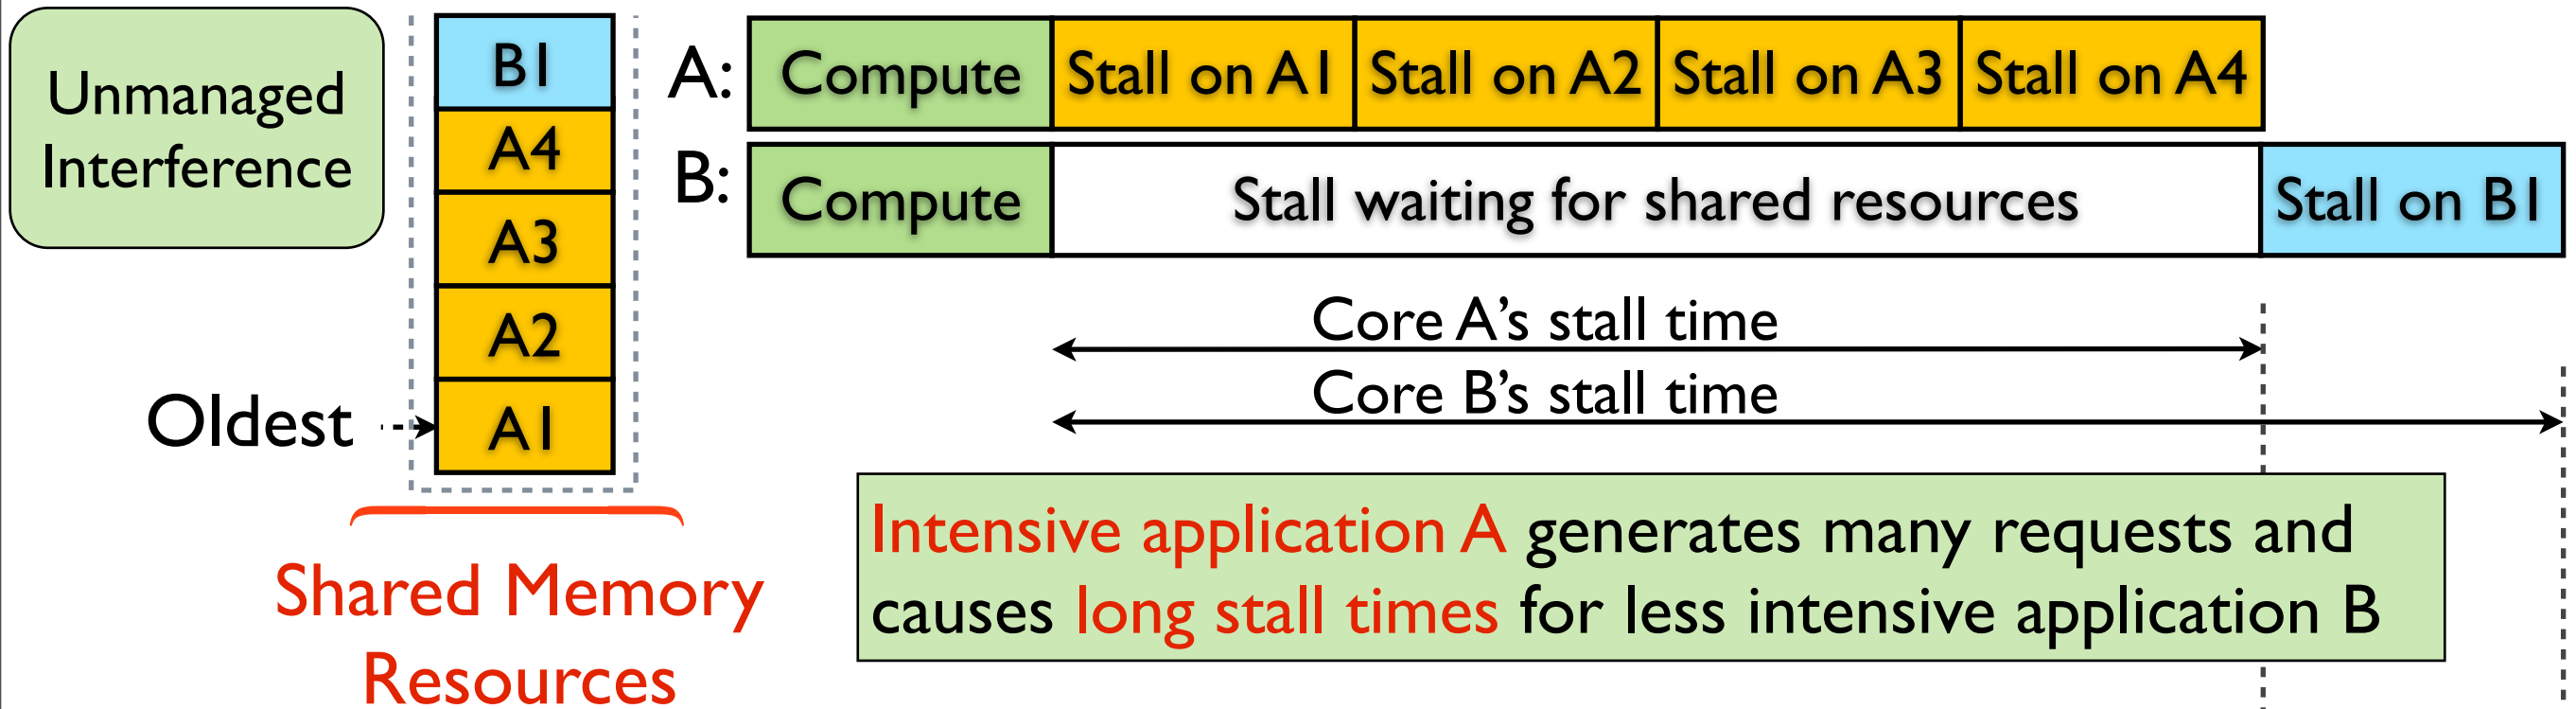
Core B's stall time

Oldest


Shared Memory
Resources

queue of requests to shared resources

Request Generation Order:
A1, A2, A3, A4, B1



Request Generation Order:
A1, A2, A3, A4, B1



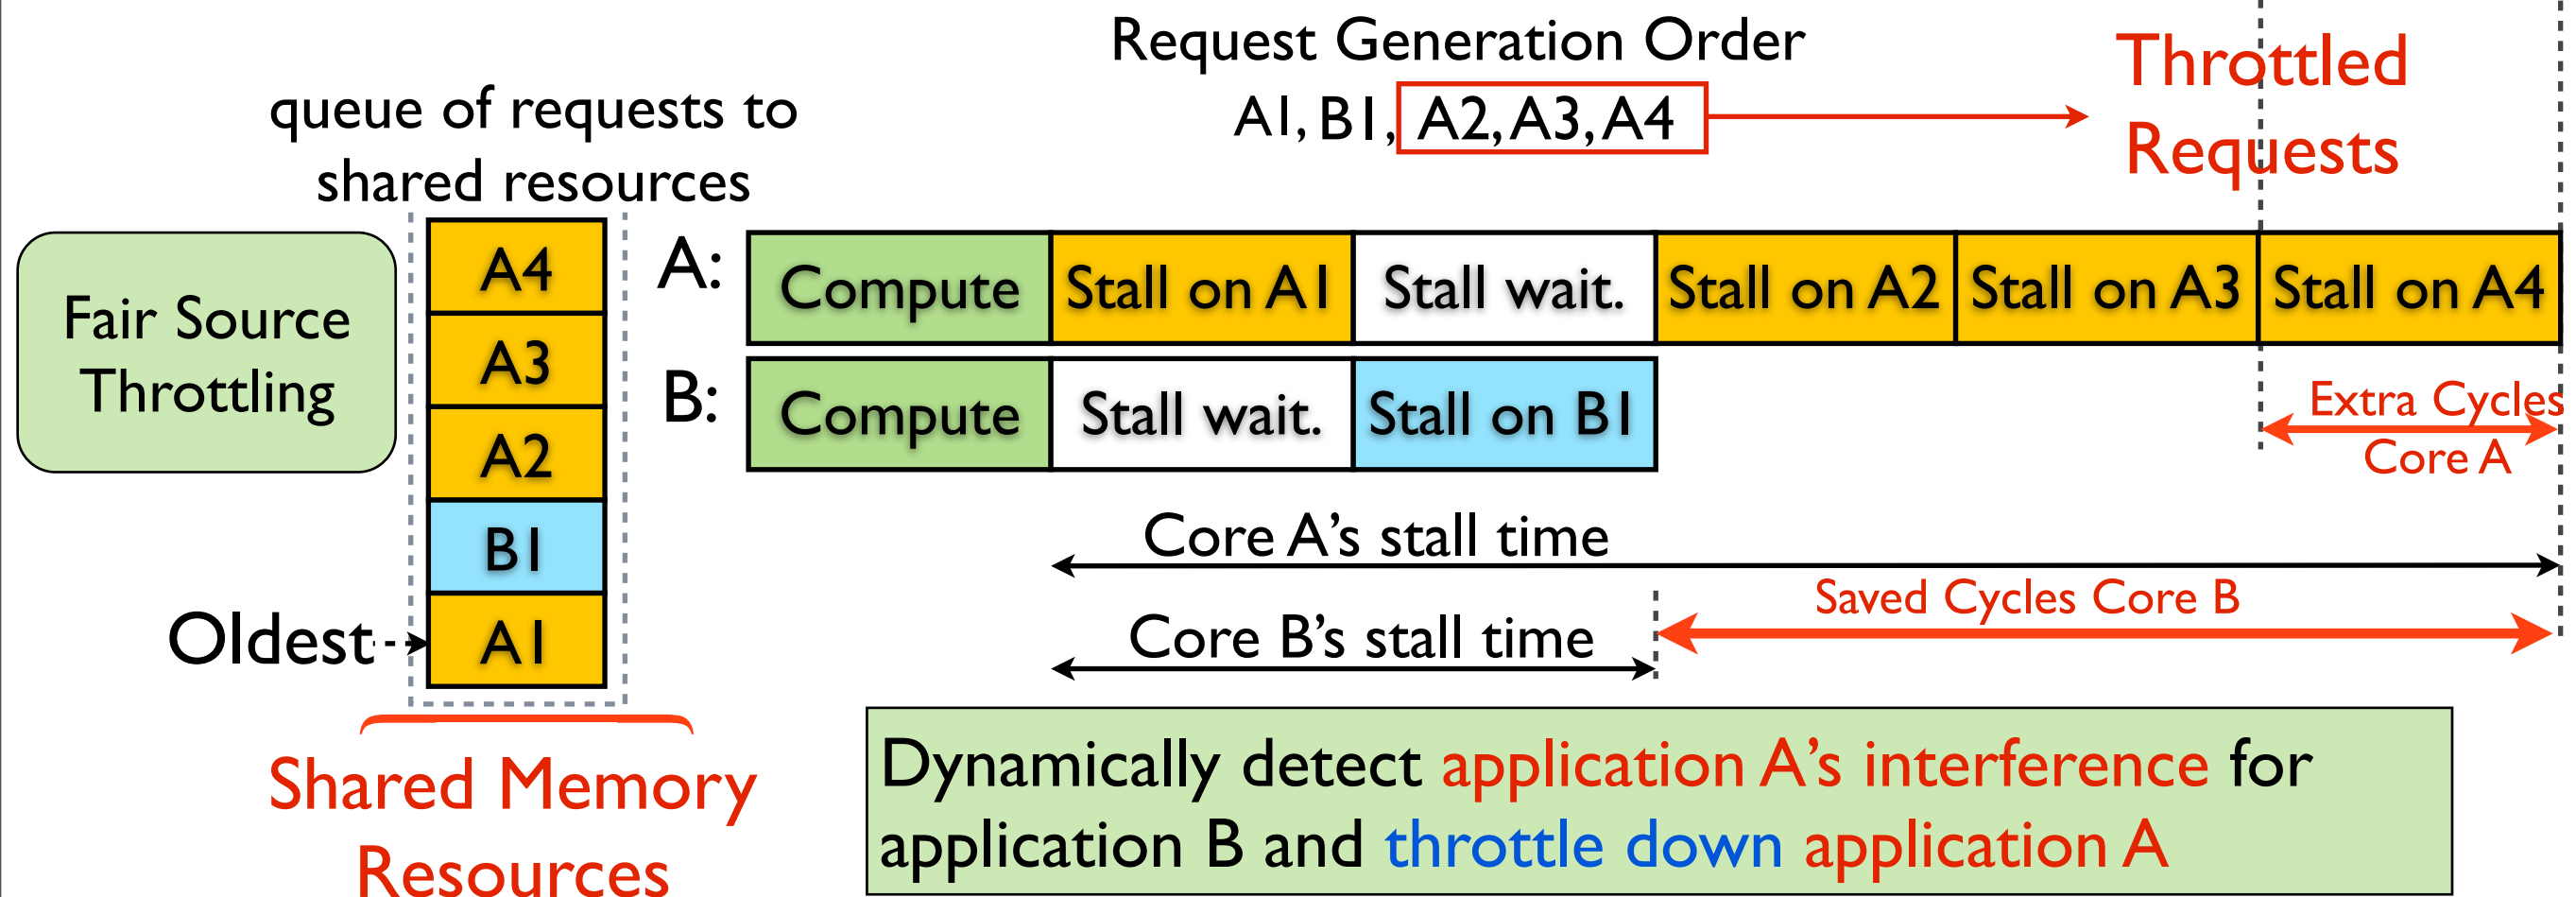
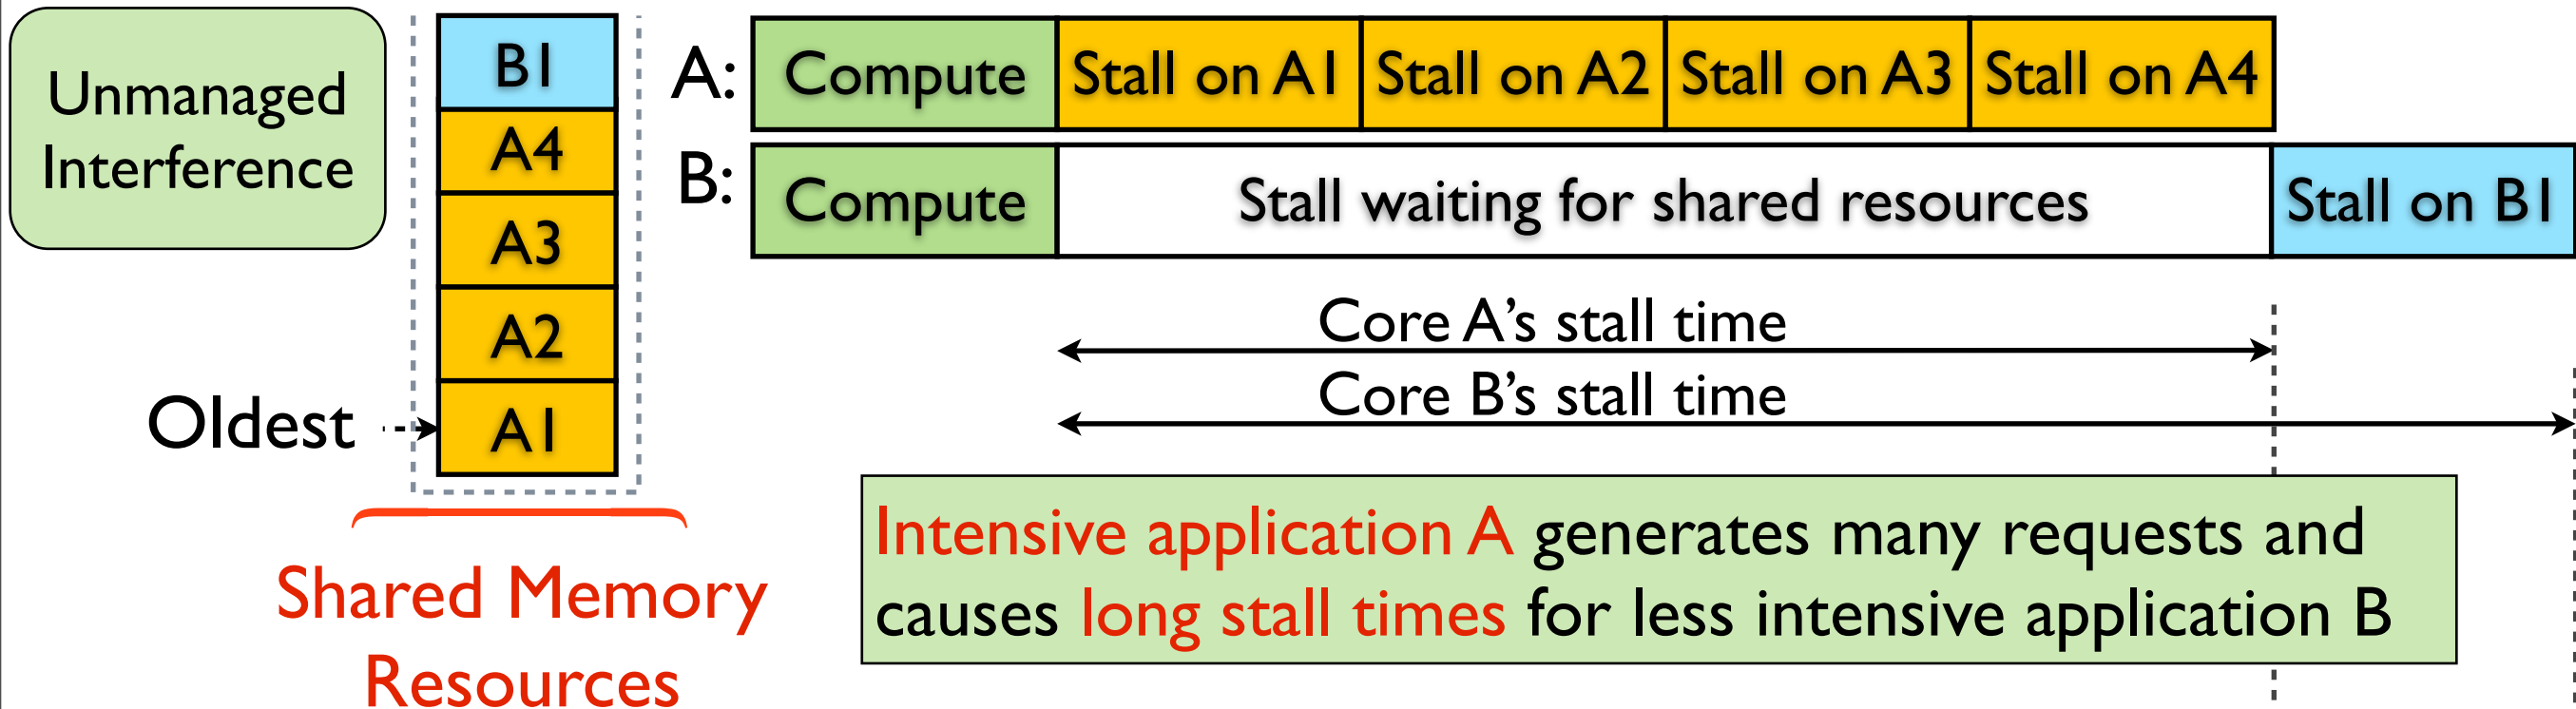
Request Generation Order
A1, B1, A2, A3, A4

A2, A3, A4



queue of requests to shared resources

Request Generation Order:
A1, A2, A3, A4, B1



Outline

- Background and Problem
- Motivation for Source Throttling
- Fairness via Source Throttling (FST)
- Evaluation
- Conclusion

Fairness via Source Throttling (FST)


Fairness via Source Throttling (FST)

- Runtime Unfairness Evaluation
 - Dynamically estimates the unfairness in the memory system

Fairness via Source Throttling (FST)

- Runtime Unfairness Evaluation
 - Dynamically estimates the unfairness in the memory system
- Dynamic Request Throttling
 - Adjusts how aggressively each core makes requests to the shared resources

Fairness via Source Throttling (FST)




FST

Runtime Unfairness
Evaluation

Dynamic
Request Throttling

Fairness via Source Throttling (FST)




FST

Runtime Unfairness
Evaluation

Dynamic
Request Throttling

Fairness via Source Throttling (FST)




FST

Runtime Unfairness
Evaluation

Dynamic
Request Throttling

1- Estimating system unfairness

Fairness via Source Throttling (FST)




FST

Runtime Unfairness
Evaluation

Dynamic
Request Throttling

- 1- Estimating system unfairness
- 2- Find app. with the highest slowdown (App-slowest)

Fairness via Source Throttling (FST)




FST

Runtime Unfairness
Evaluation


Dynamic
Request Throttling

- 1- Estimating system unfairness
- 2- Find app. with the highest slowdown (App-slowest)
- 3- Find app. causing most interference for App-slowest (App-interfering)

Fairness via Source Throttling (FST)




FST




- 1- Estimating system unfairness
- 2- Find app. with the highest slowdown (App-slowest)
- 3- Find app. causing most interference for App-slowest (App-interfering)

Fairness via Source Throttling (FST)




FST



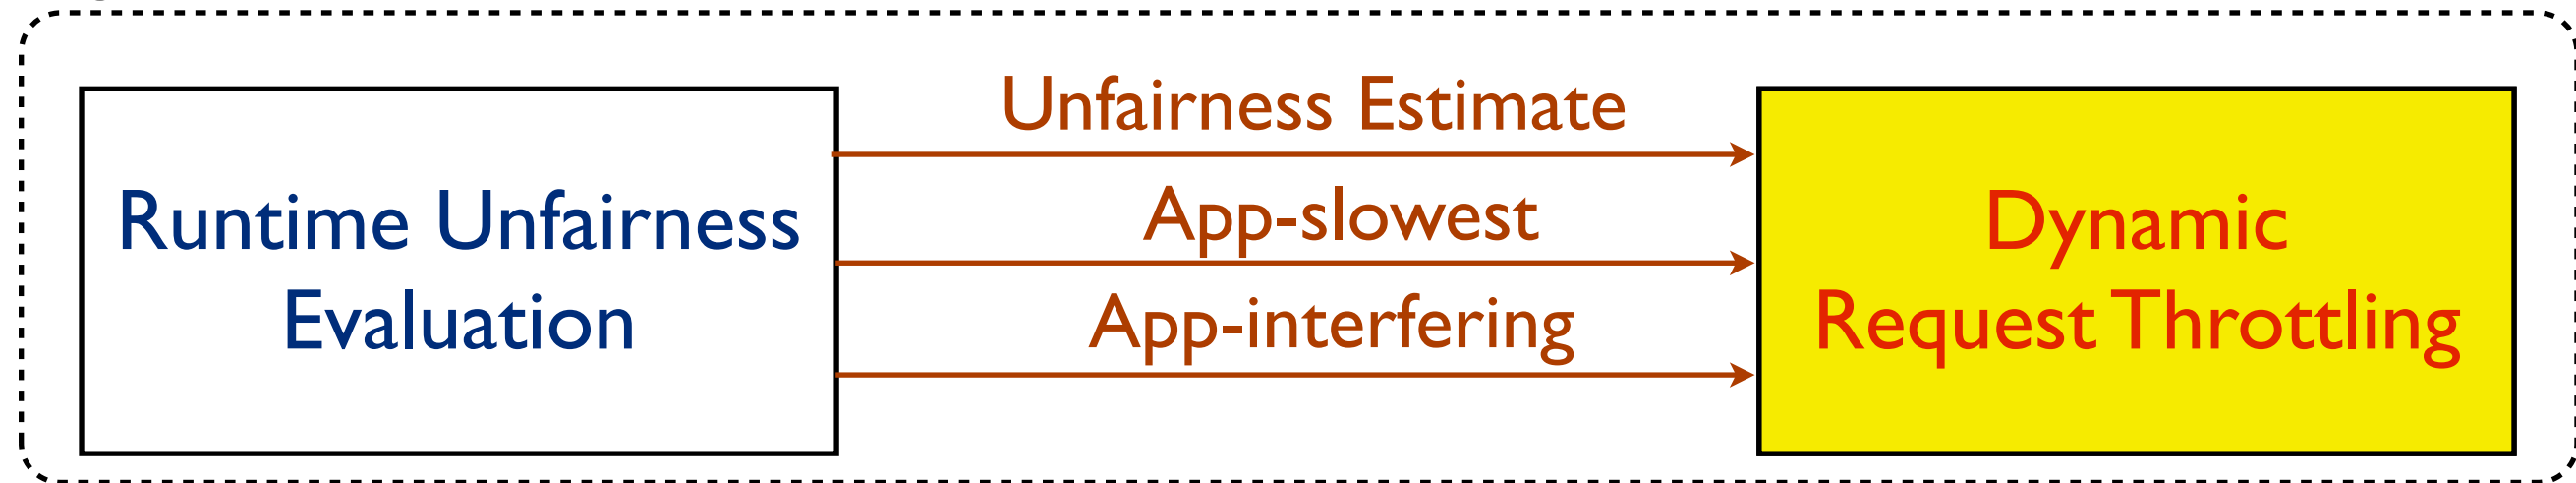
- 1- Estimating system unfairness
- 2- Find app. with the highest slowdown (App-slowest)
- 3- Find app. causing most interference for App-slowest (App-interfering)

```
if (Unfairness Estimate > Target)
{
```

Fairness via Source Throttling (FST)




FST



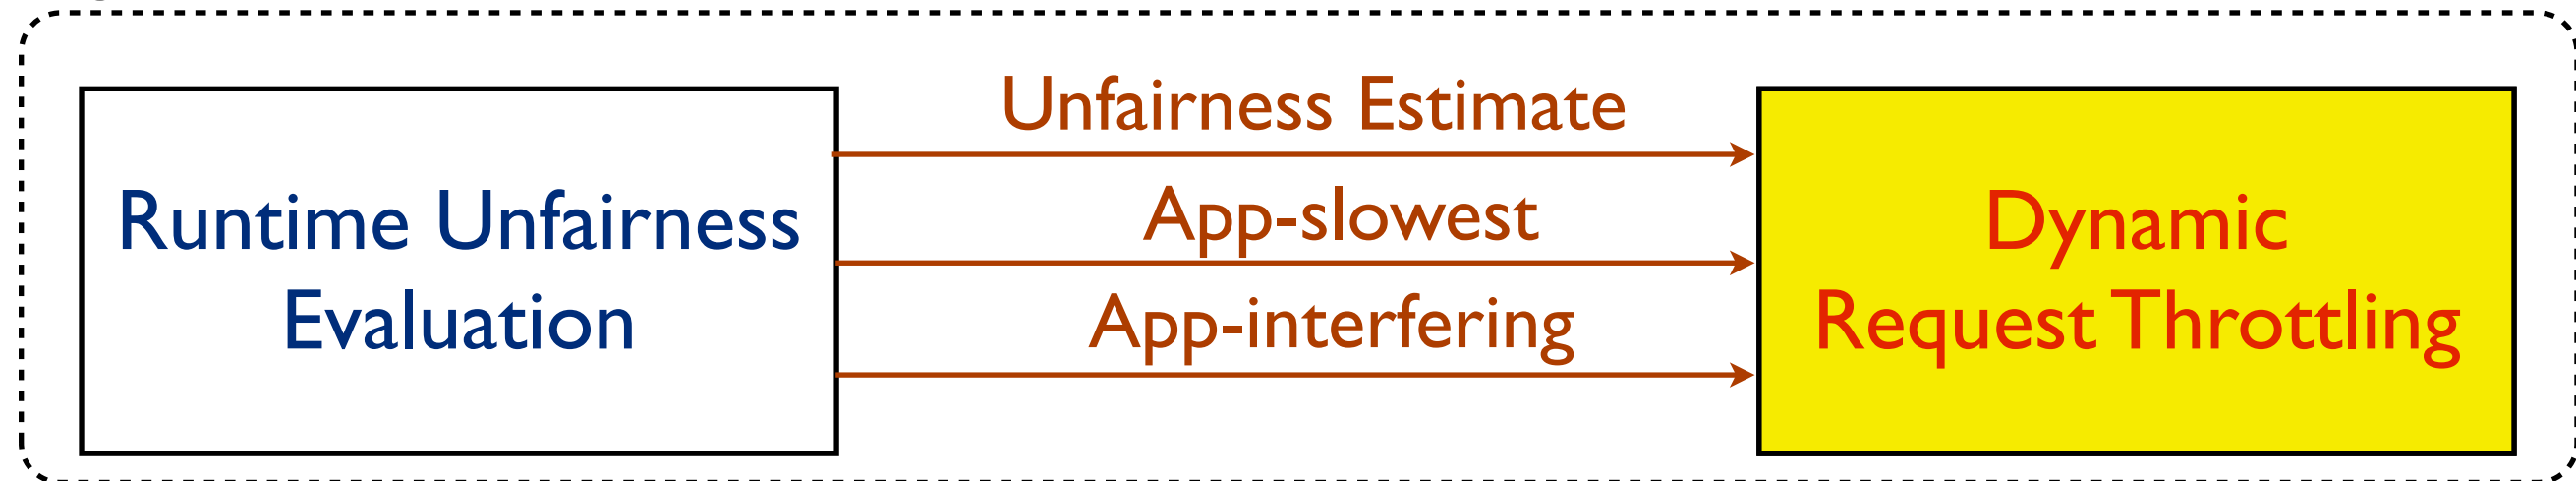
- 1- Estimating system unfairness
- 2- Find app. with the highest slowdown (App-slowest)
- 3- Find app. causing most interference for App-slowest (App-interfering)

```
if (Unfairness Estimate > Target)
{
  1-Throttle down App-interfering
}
```

Fairness via Source Throttling (FST)



FST

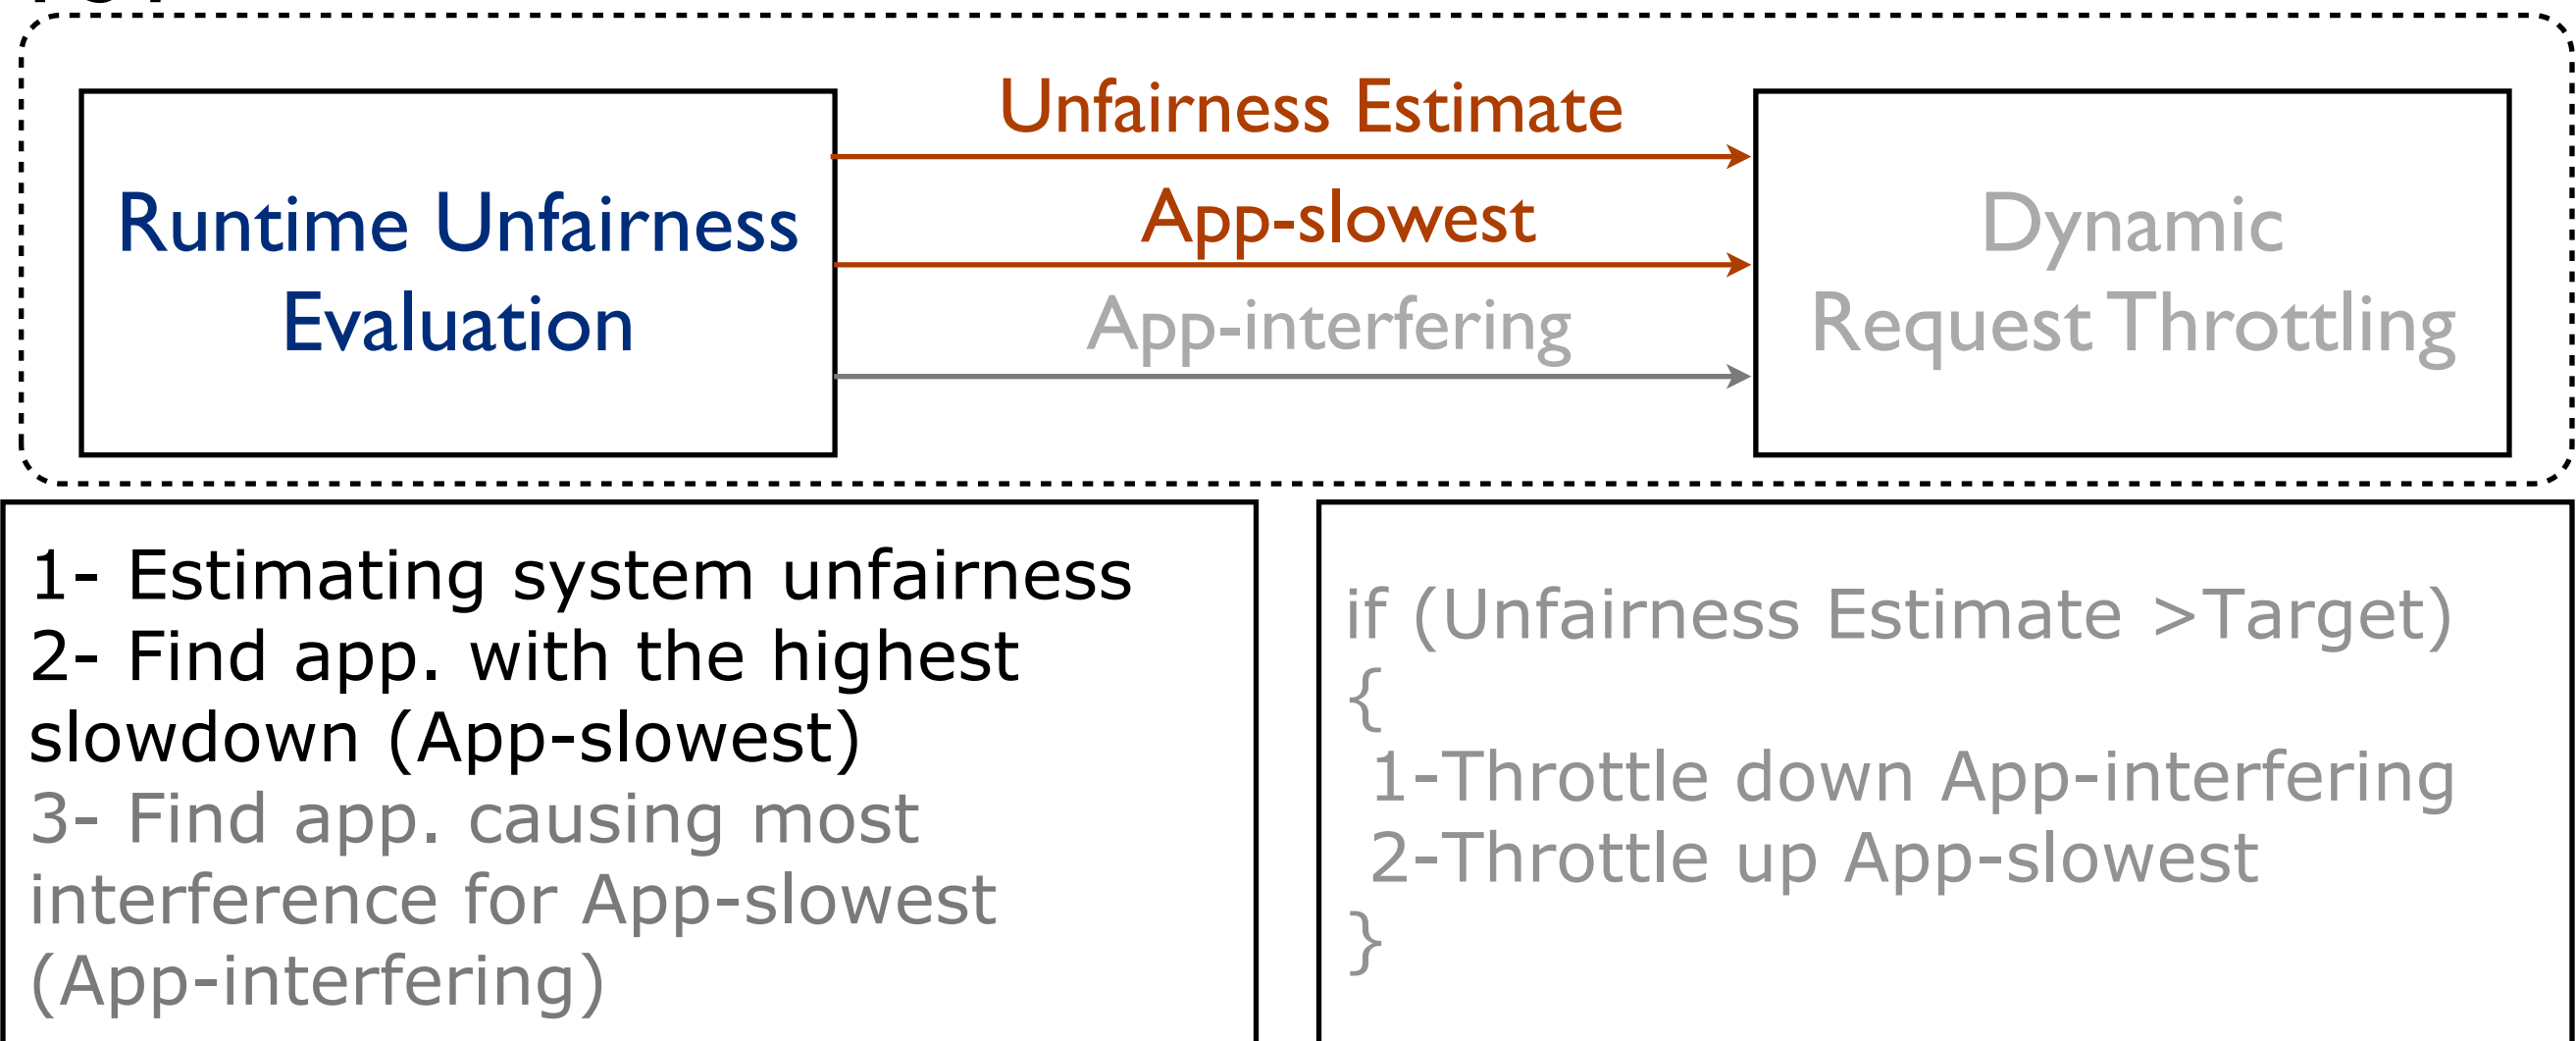


- 1- Estimating system unfairness
- 2- Find app. with the highest slowdown (App-slowest)
- 3- Find app. causing most interference for App-slowest (App-interfering)

```
if (Unfairness Estimate > Target)
{
  1-Throttle down App-interfering
  2-Throttle up App-slowest
}
```

Fairness via Source Throttling (FST)

FST



Estimating System Unfairness

- Unfairness = $\frac{\text{Max}\{\text{Slowdown } i\} \text{ over all applications } i}{\text{Min}\{\text{Slowdown } i\} \text{ over all applications } i}$
- Slowdown of application $i = \frac{T_i^{\text{Shared}}}{T_i^{\text{Alone}}}$

Estimating System Unfairness

- Unfairness = $\frac{\text{Max}\{\text{Slowdown } i\} \text{ over all applications } i}{\text{Min}\{\text{Slowdown } i\} \text{ over all applications } i}$
- Slowdown of application $i = \frac{T_i^{\text{Shared}}}{T_i^{\text{Alone}}}$
- How can T_i^{Alone} be estimated in shared mode?

Estimating System Unfairness


- Unfairness = $\frac{\text{Max}\{\text{Slowdown } i\} \text{ over all applications } i}{\text{Min}\{\text{Slowdown } i\} \text{ over all applications } i}$
- Slowdown of application $i = \frac{T_i^{\text{Shared}}}{T_i^{\text{Alone}}}$
- How can T_i^{Alone} be estimated in shared mode?
- T_i^{Excess} is the number of **extra cycles** it takes application i to execute **due to interference**

Estimating System Unfairness


- Unfairness = $\frac{\text{Max}\{\text{Slowdown } i\} \text{ over all applications } i}{\text{Min}\{\text{Slowdown } i\} \text{ over all applications } i}$
- Slowdown of application $i = \frac{T_i^{\text{Shared}}}{T_i^{\text{Alone}}}$
- How can T_i^{Alone} be estimated in shared mode?
- T_i^{Excess} is the number of **extra cycles** it takes application i to execute **due to interference**
- $T_i^{\text{Alone}} = T_i^{\text{Shared}} - T_i^{\text{Excess}}$

Tracking Inter-Core Interference

Tracking Inter-Core Interference




Tracking Inter-Core Interference




Three interference sources:

Tracking Inter-Core Interference



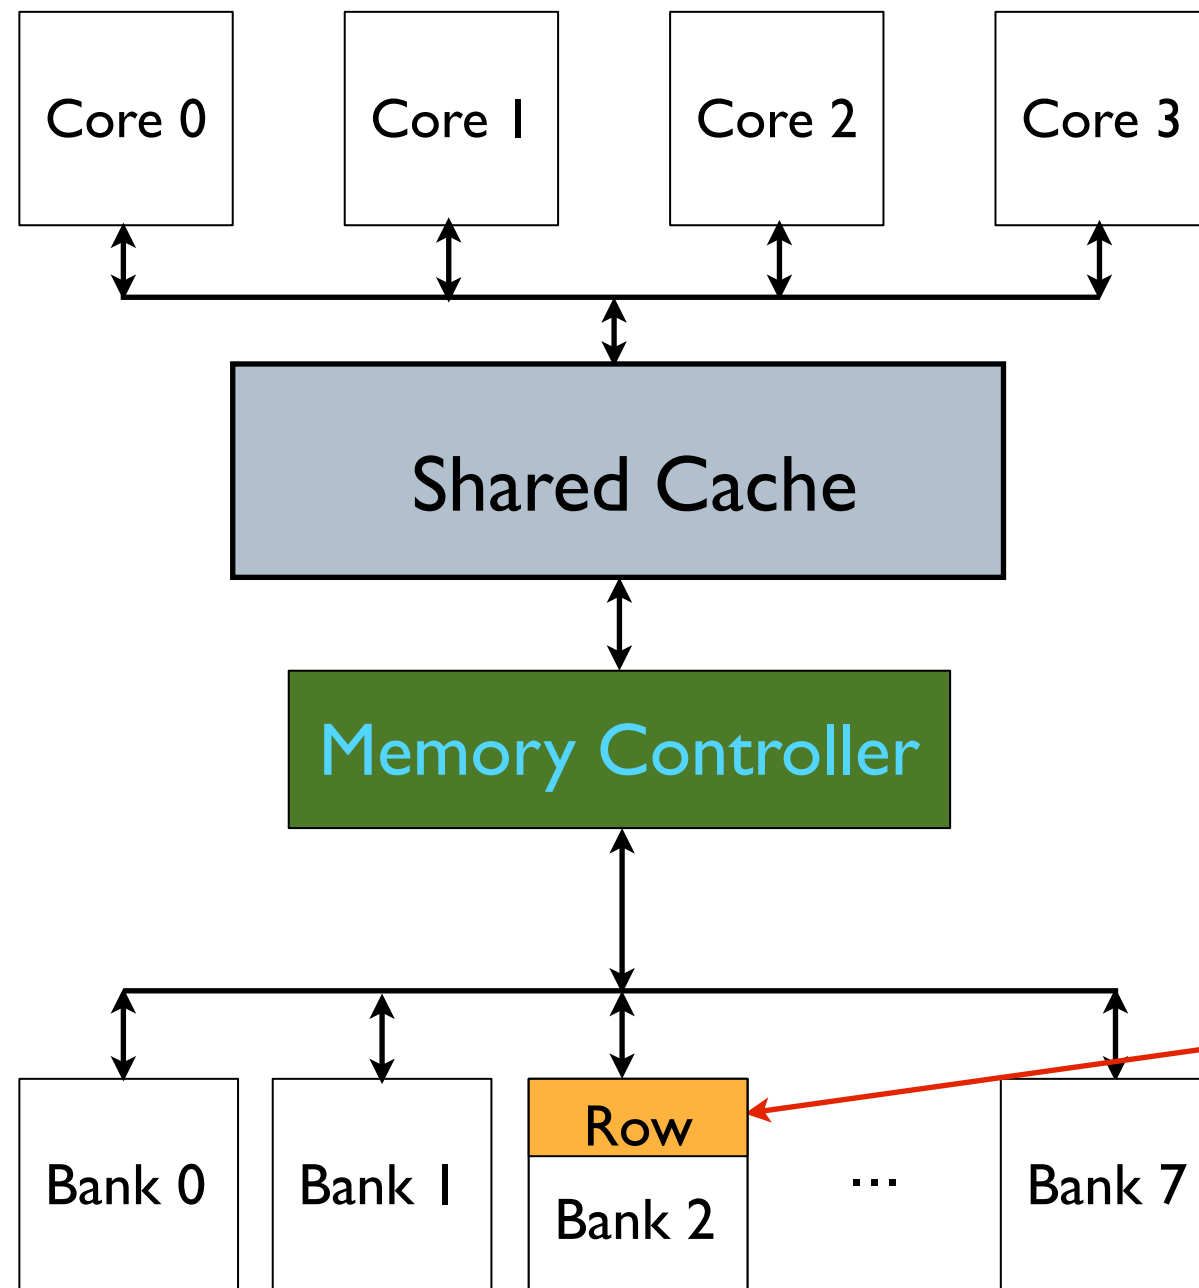
Three interference sources:
1. Shared Cache

Tracking Inter-Core Interference




Three interference sources:
1. Shared Cache
2. DRAM bus and bank

Tracking Inter-Core Interference



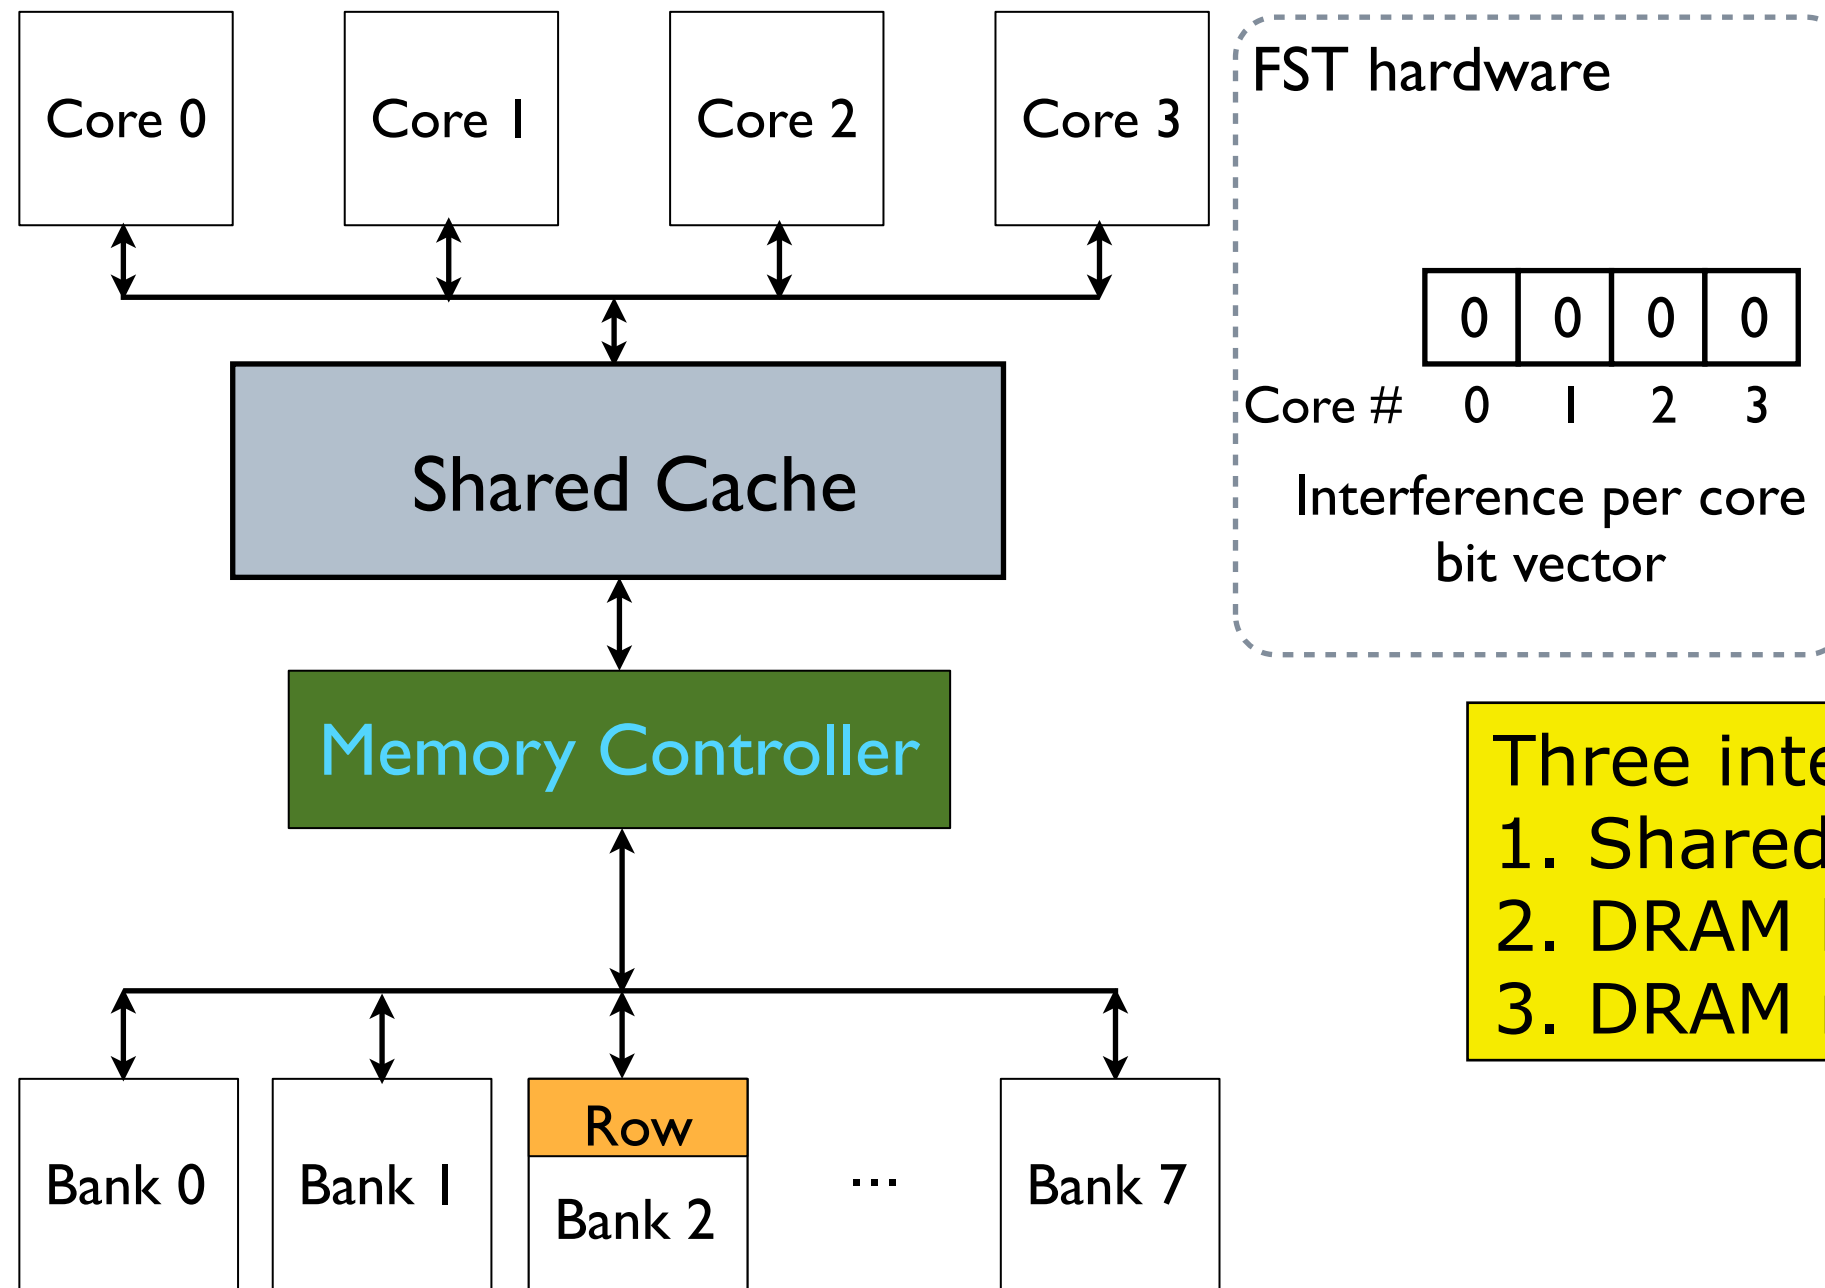
Three interference sources:
1. Shared Cache
2. DRAM bus and bank
3. DRAM row-buffers

Tracking Inter-Core Interference



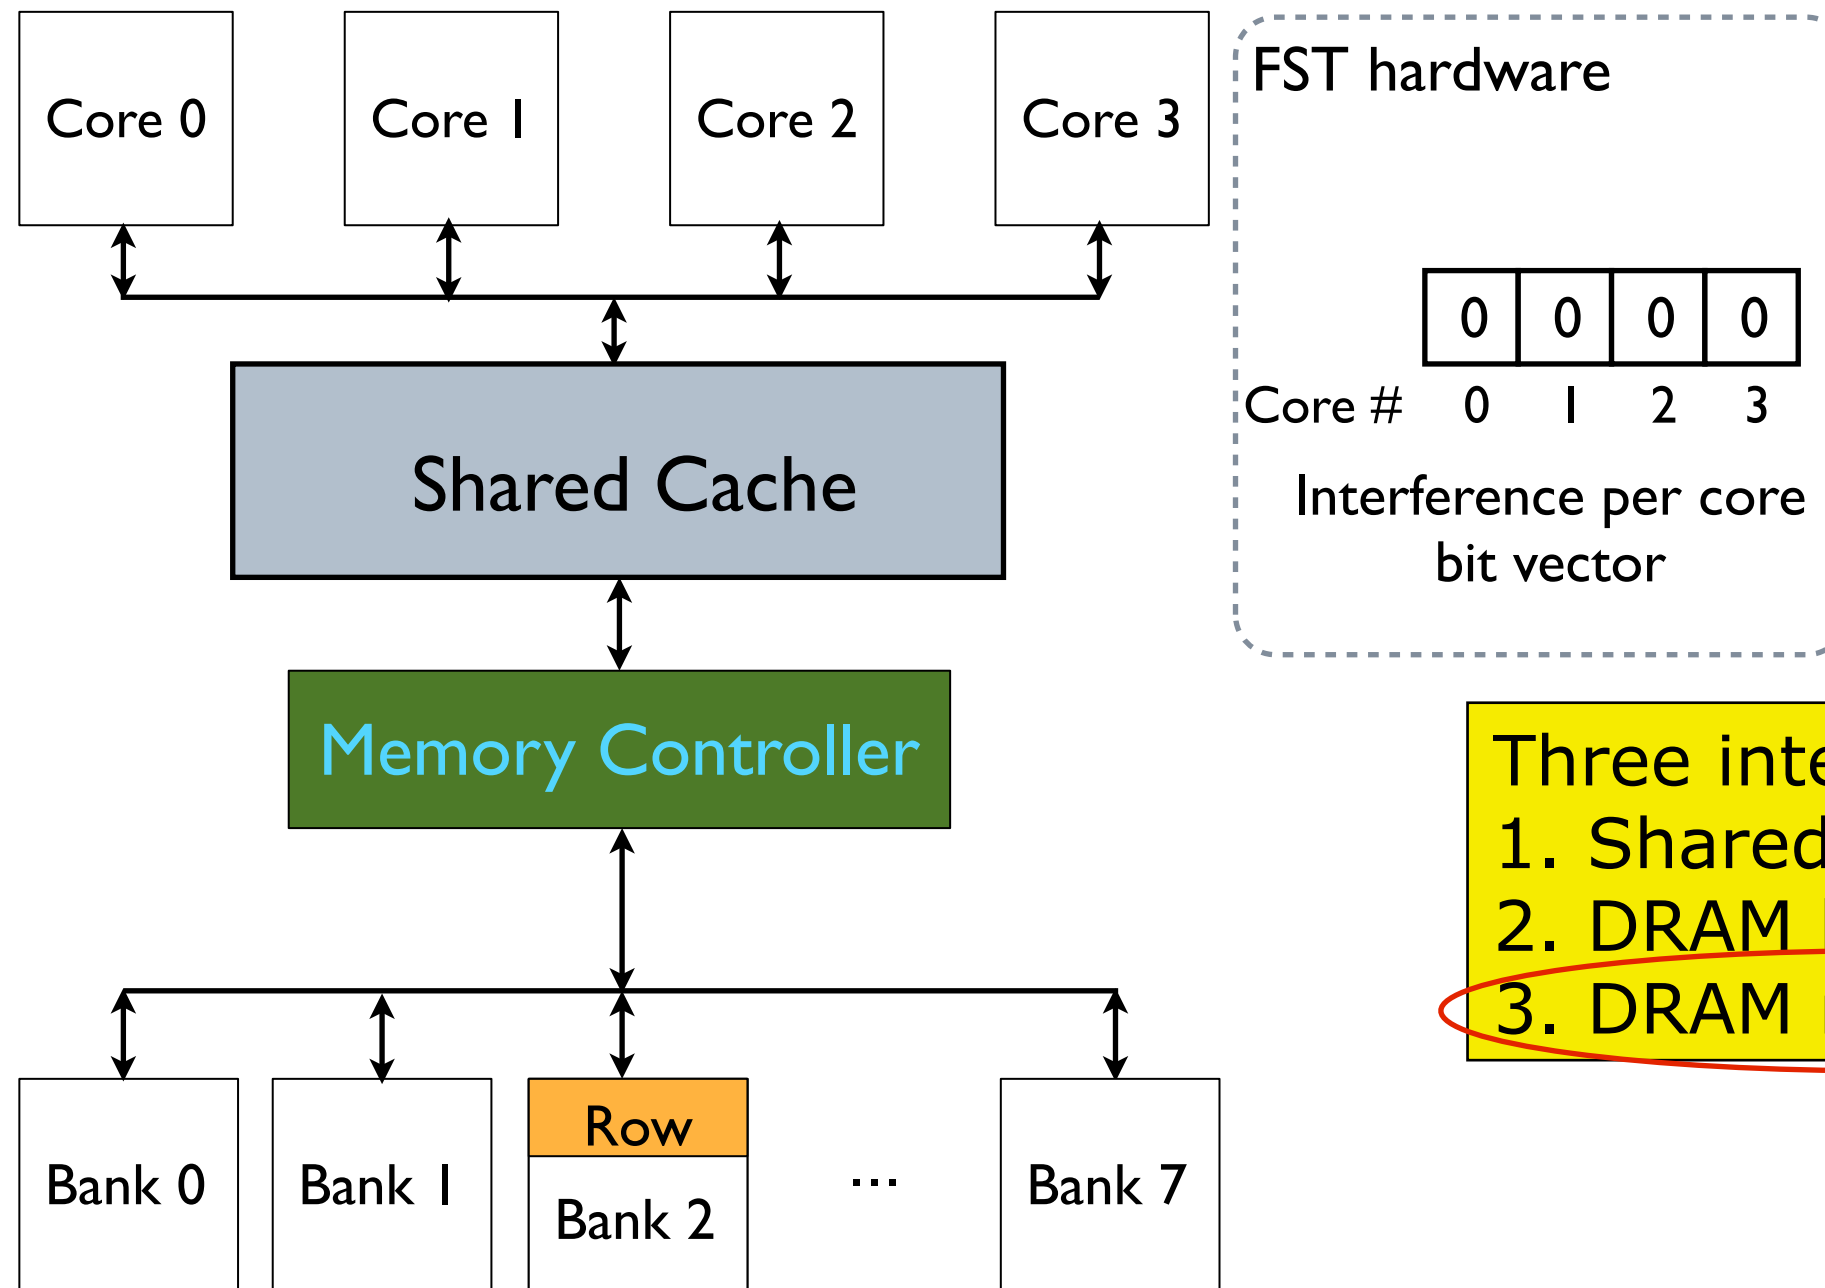
Three interference sources:
1. Shared Cache
2. DRAM bus and bank
3. DRAM row-buffers

Tracking Inter-Core Interference





Three interference sources:
1. Shared Cache
2. DRAM bus and bank
3. DRAM row-buffers

Tracking Inter-Core Interference

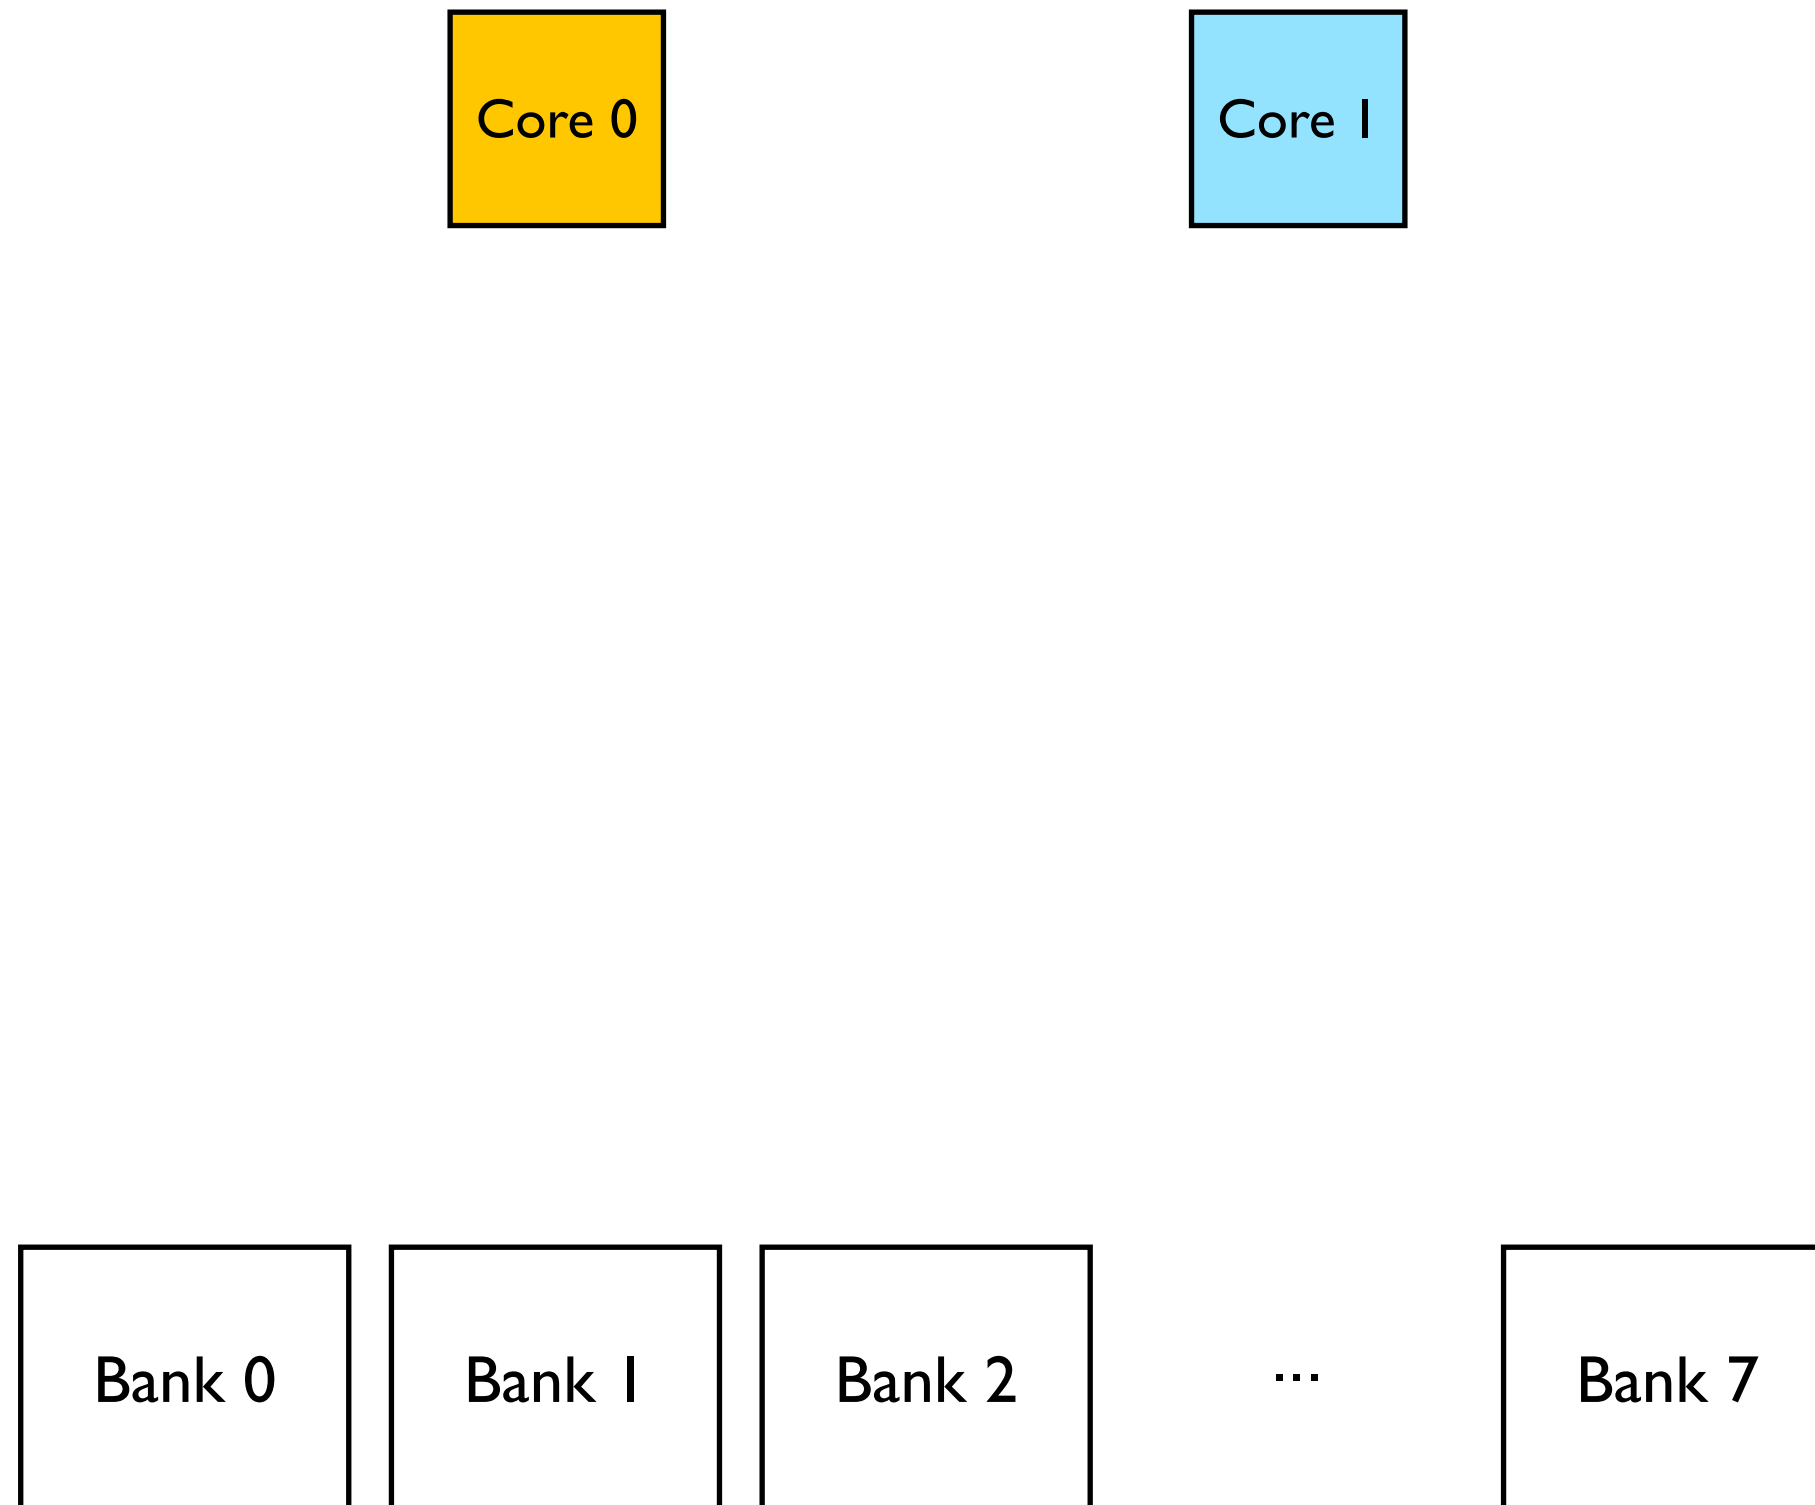


Three interference sources:
1. Shared Cache
2. DRAM bus and bank
3. DRAM row-buffers

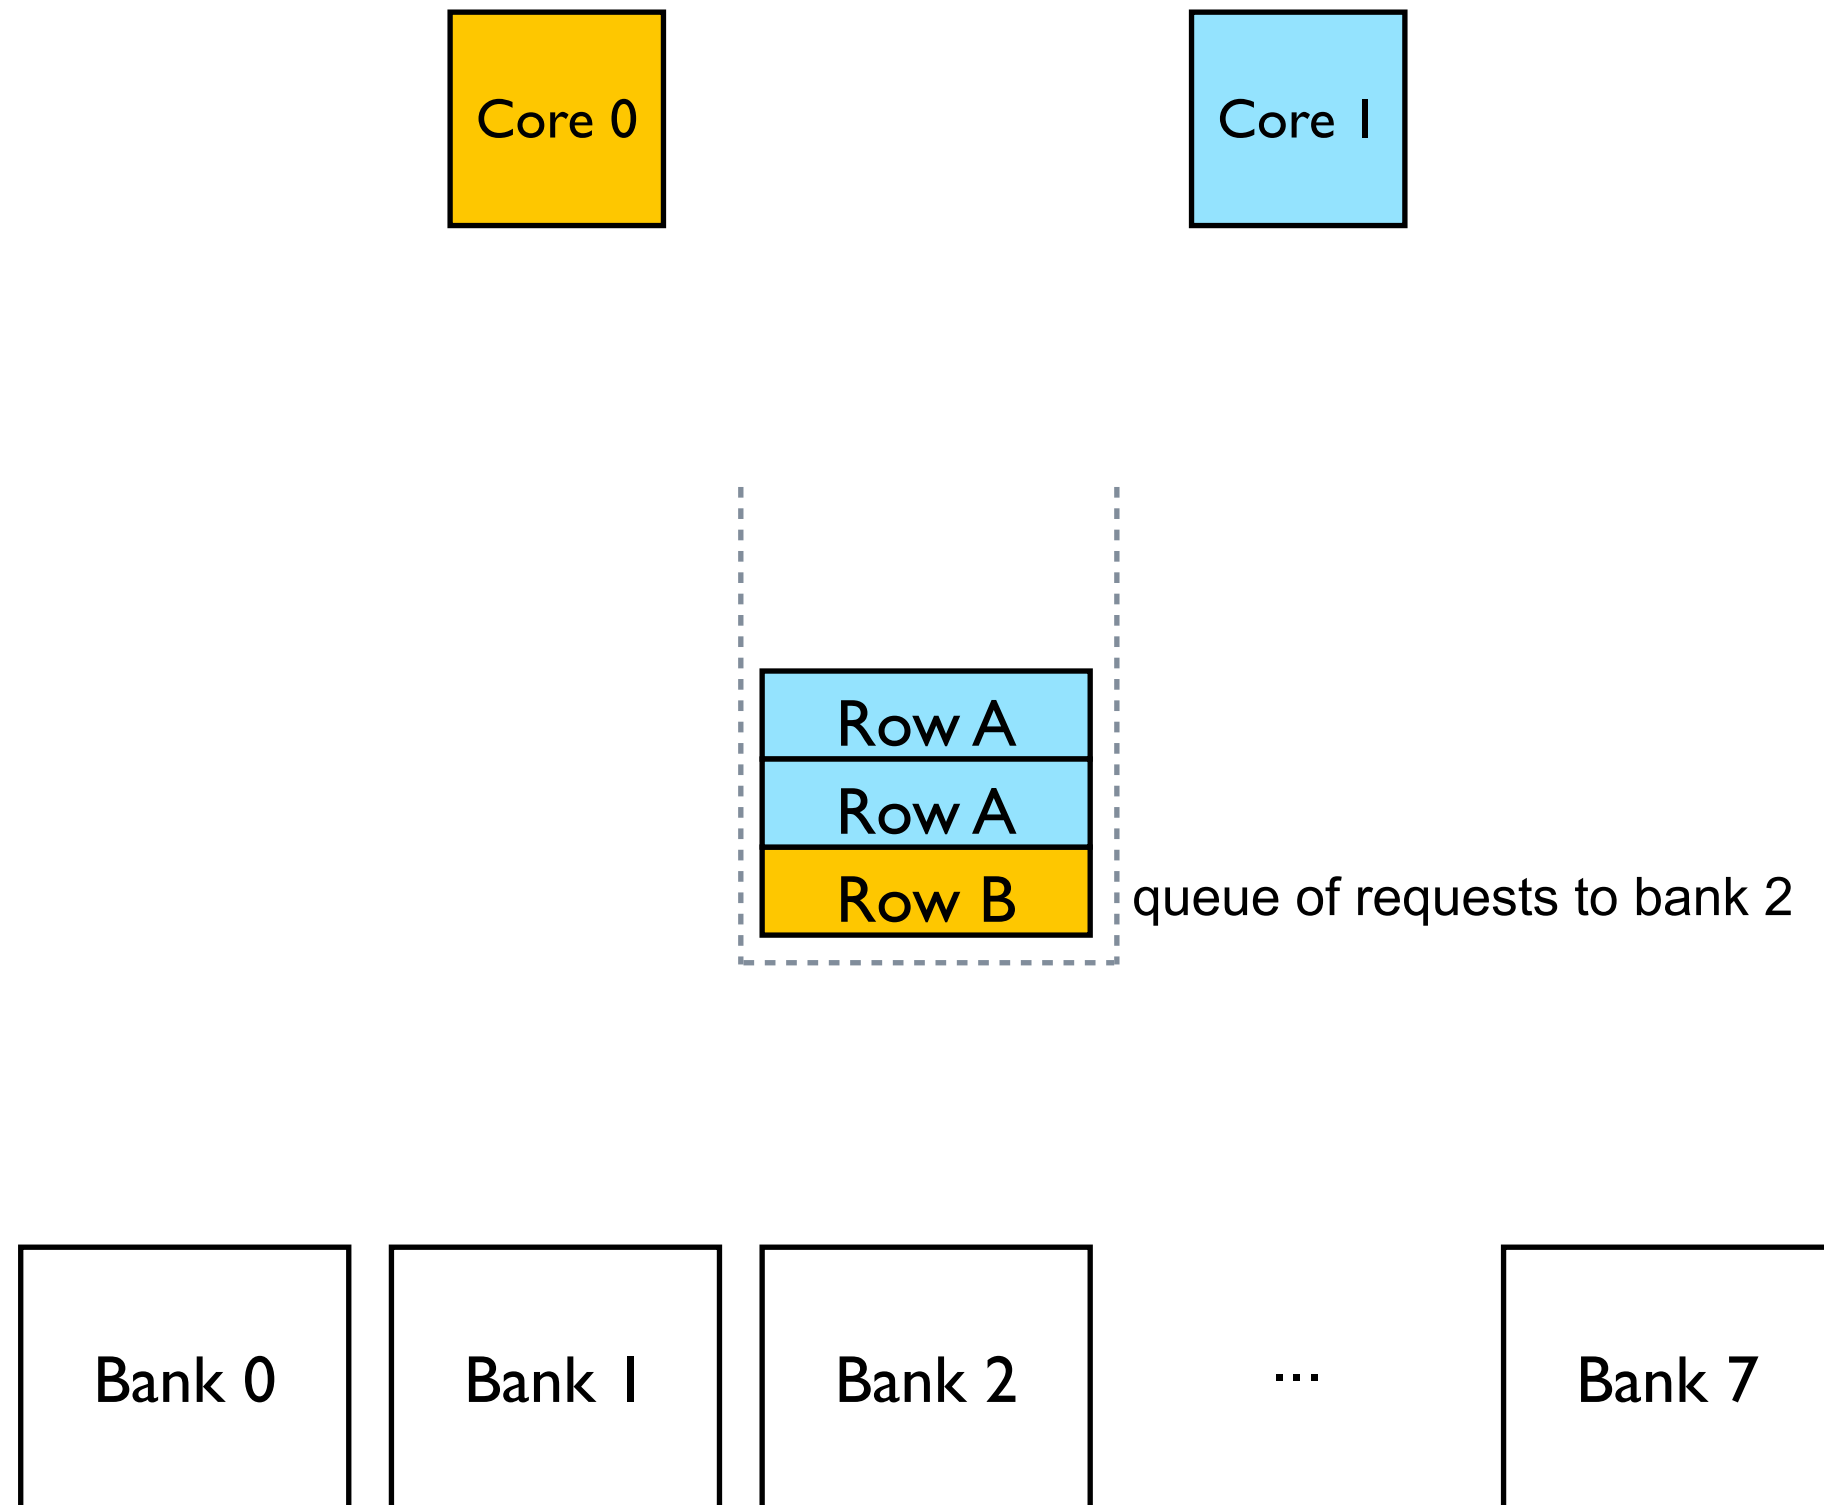
Tracking DRAM Row-Buffer Interference



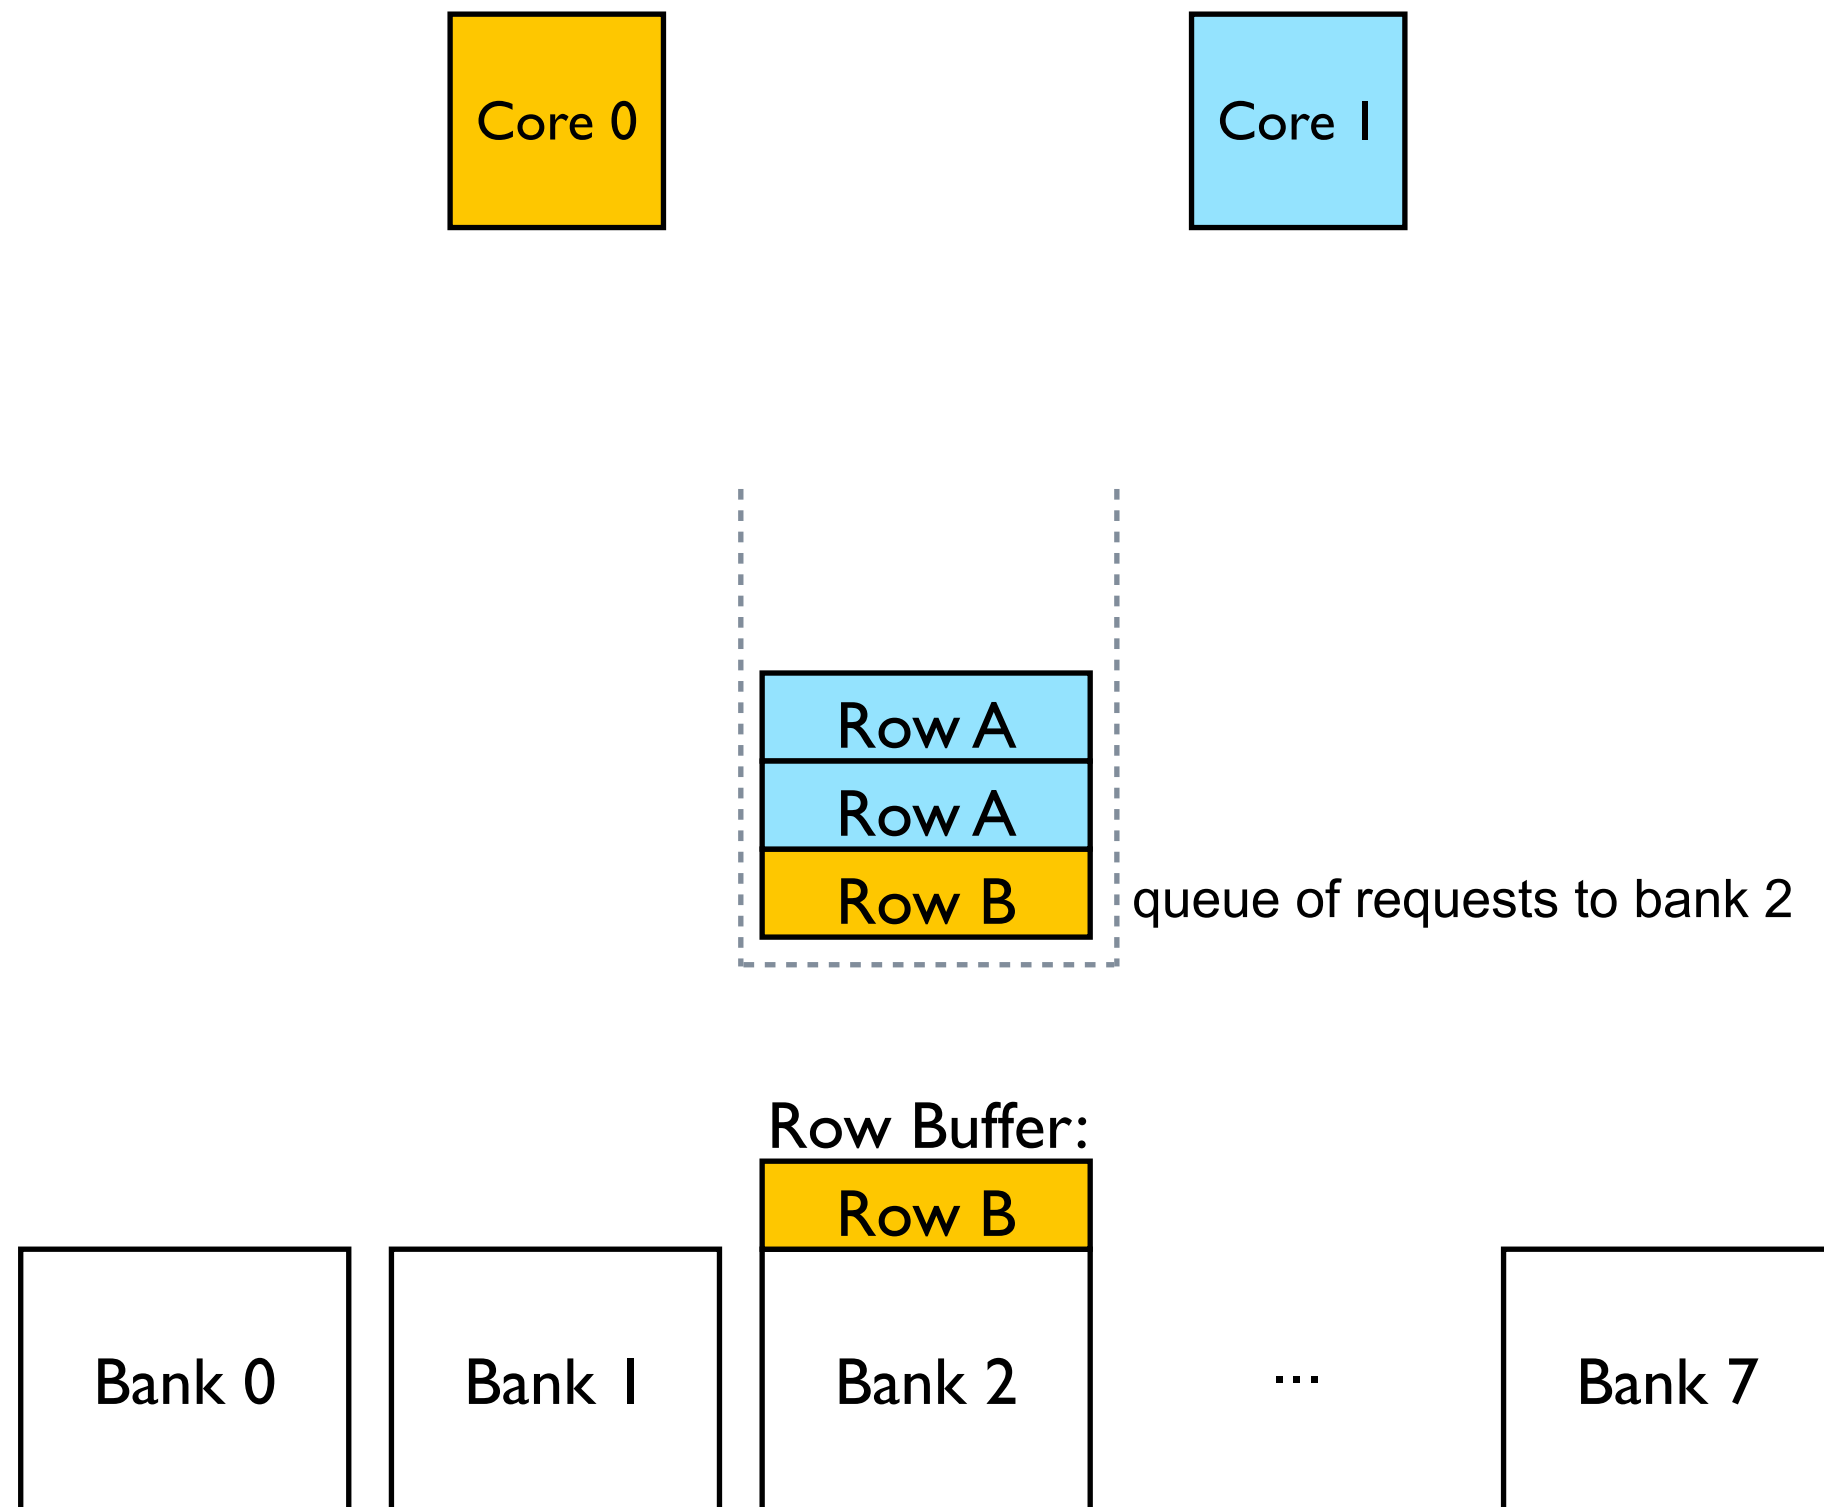
Tracking DRAM Row-Buffer Interference



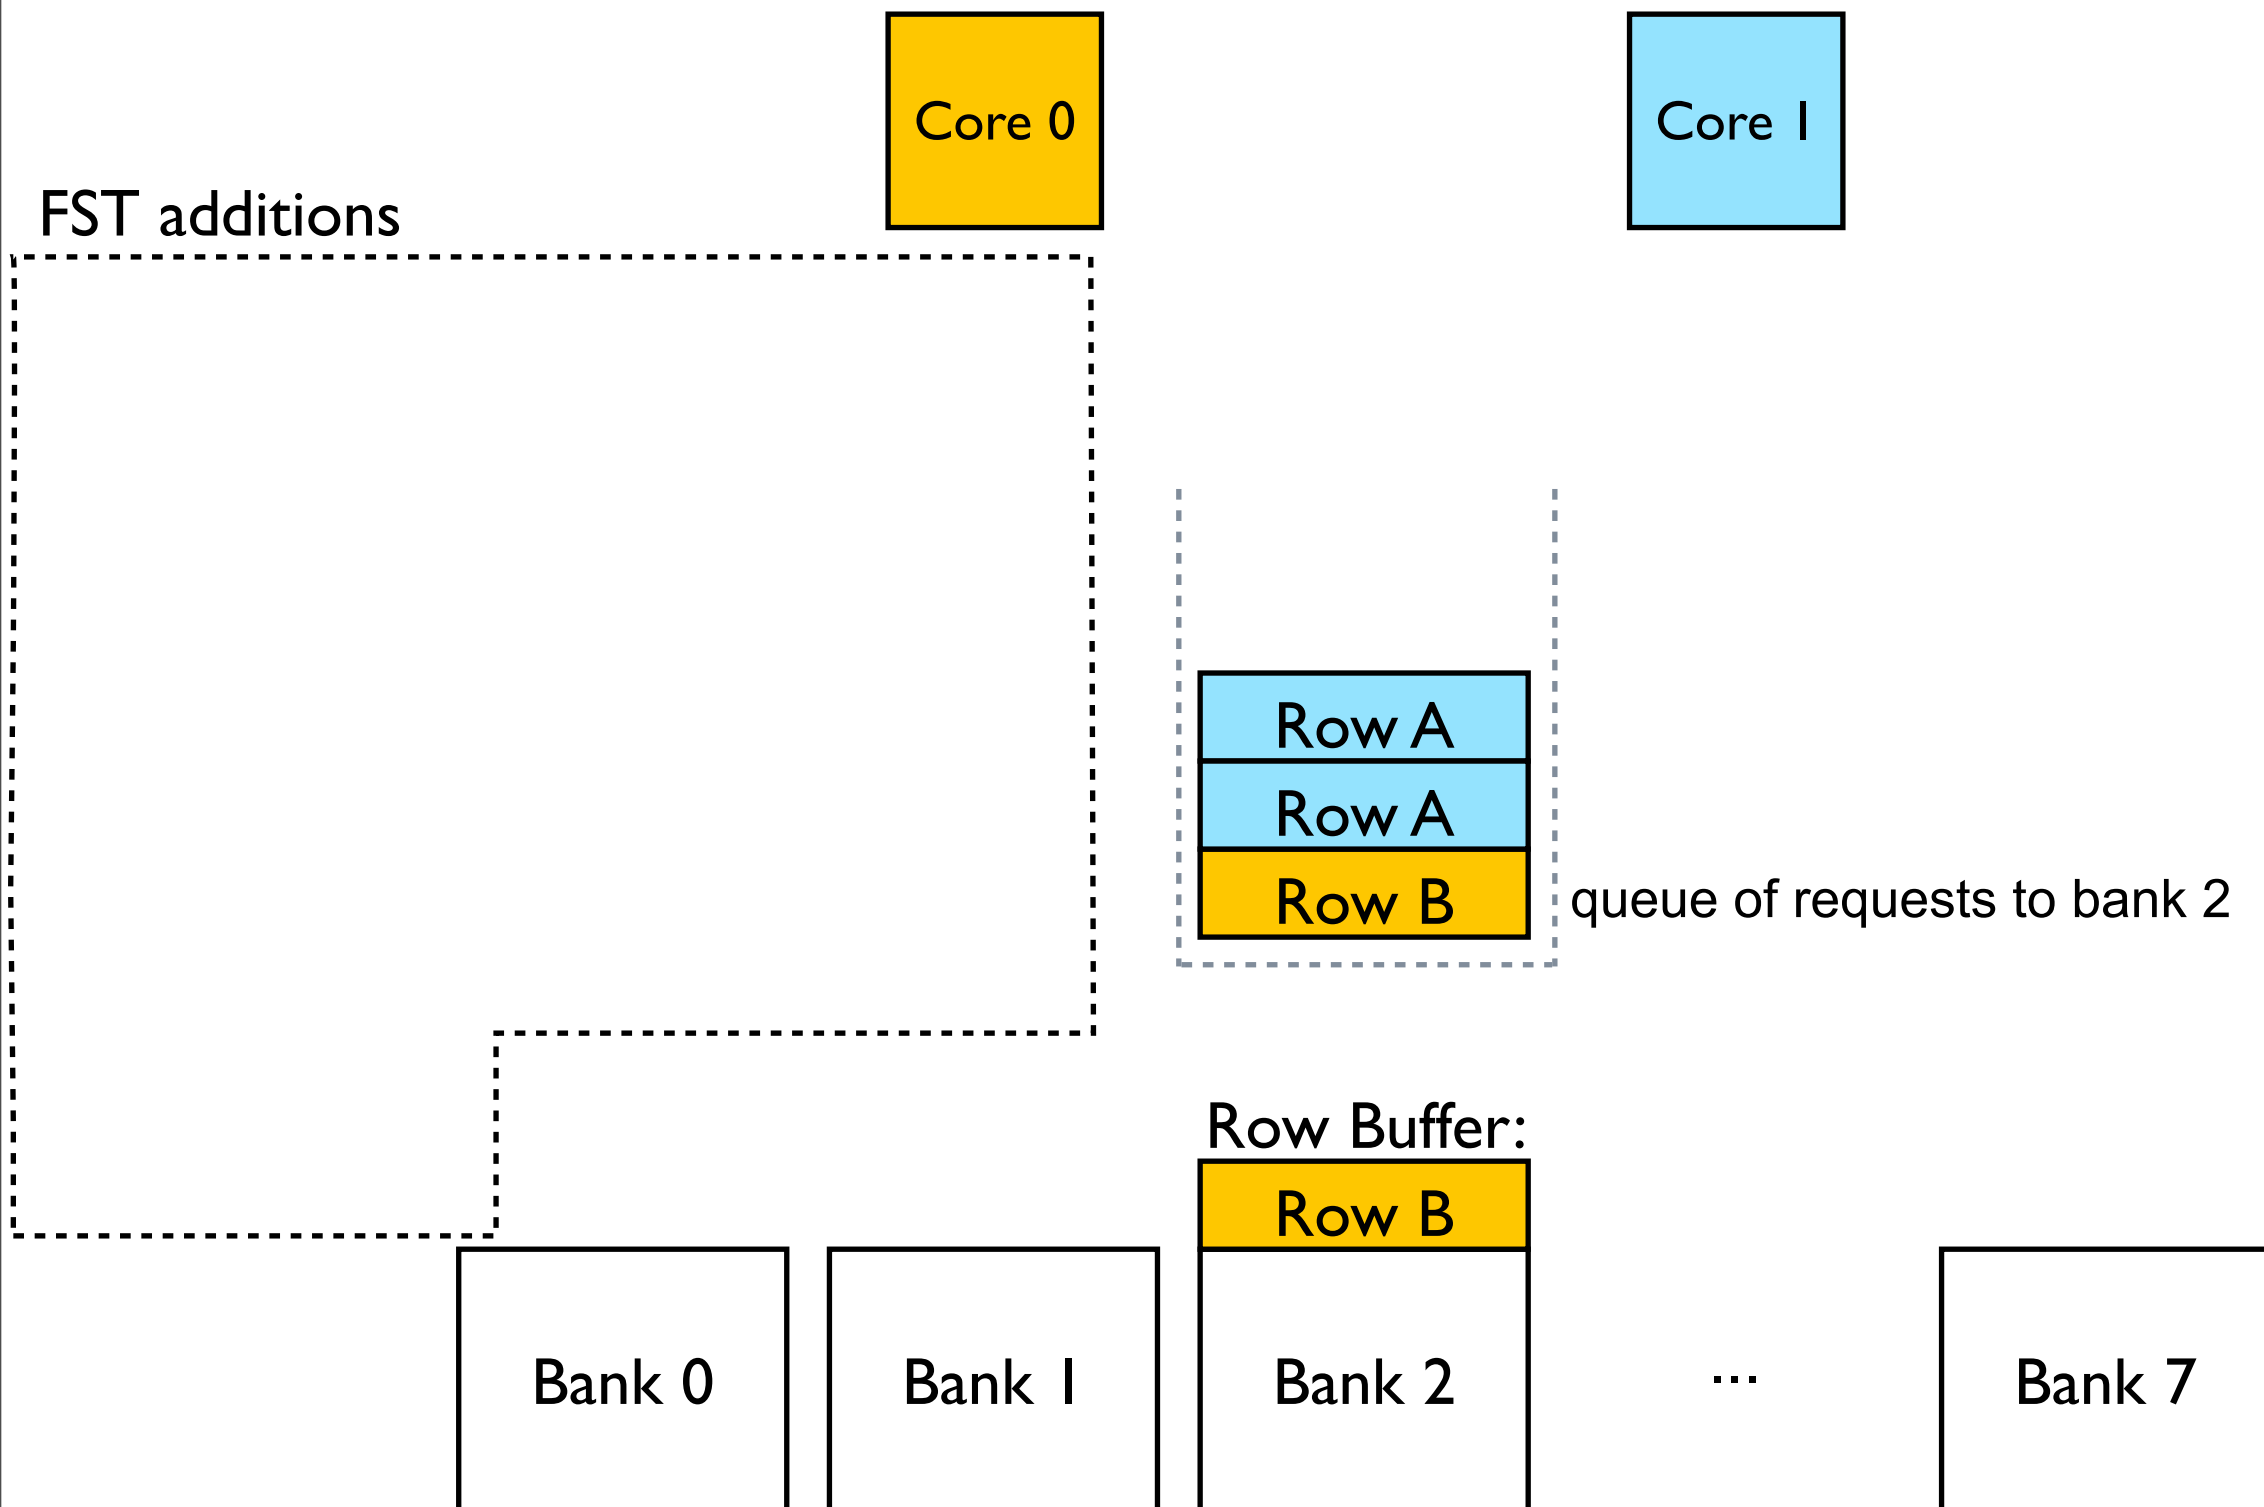
Tracking DRAM Row-Buffer Interference



Tracking DRAM Row-Buffer Interference

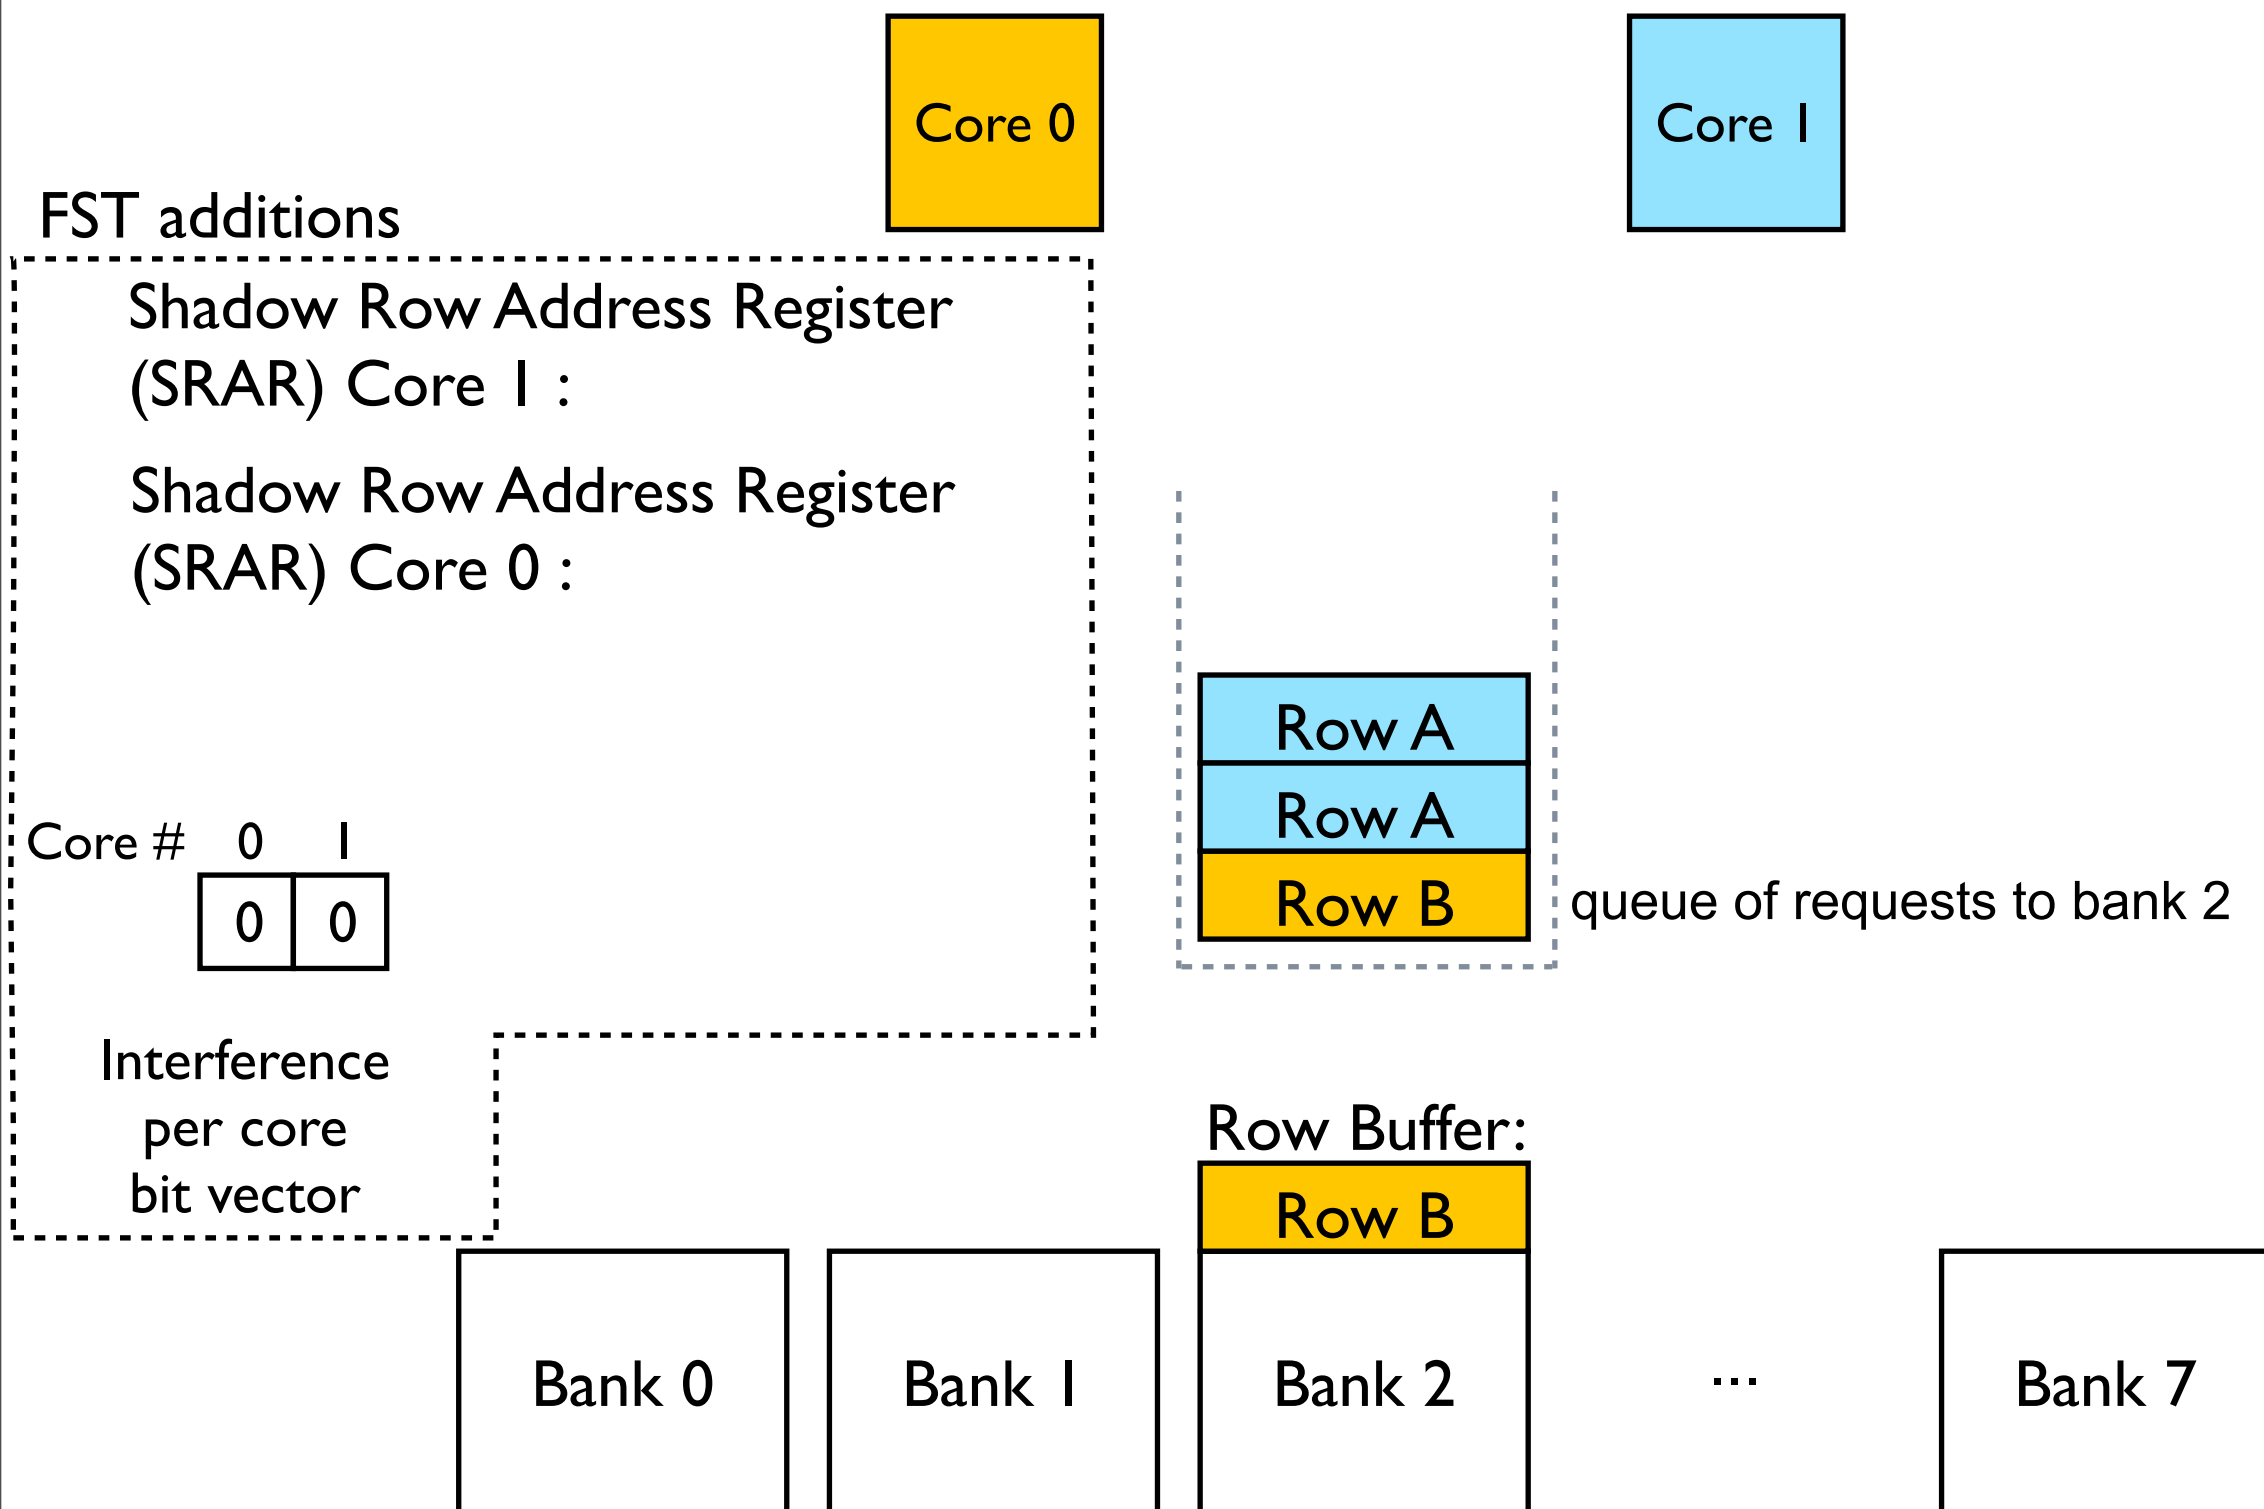


Tracking DRAM Row-Buffer Interference

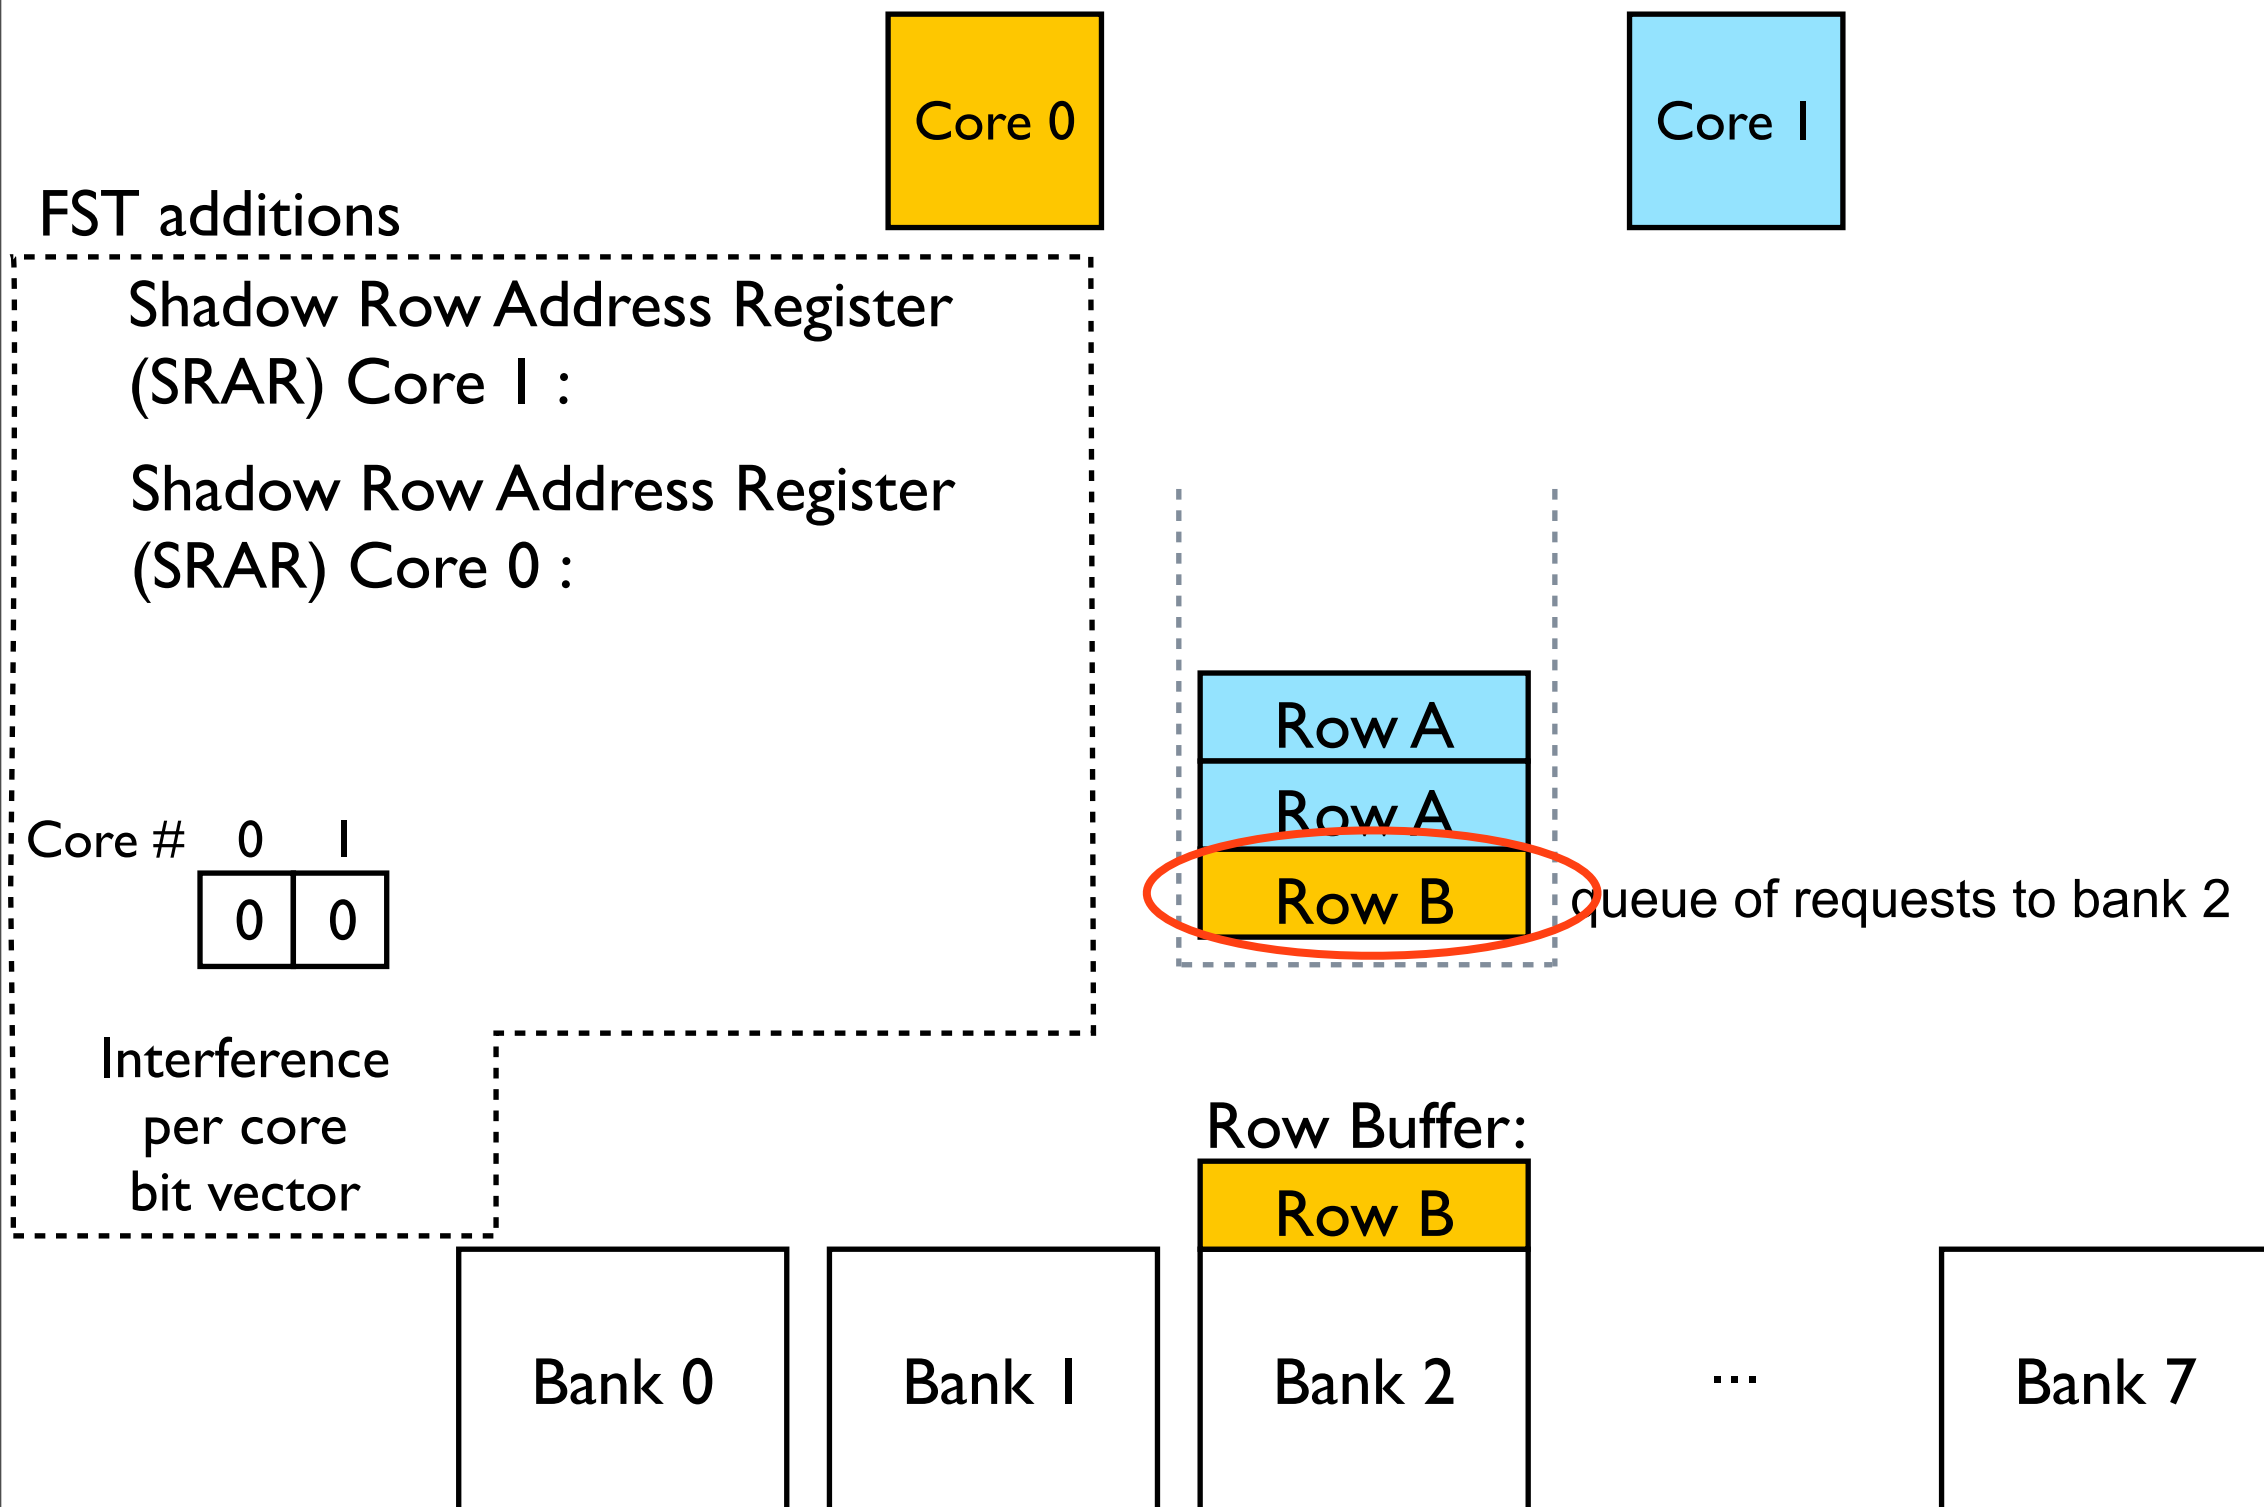




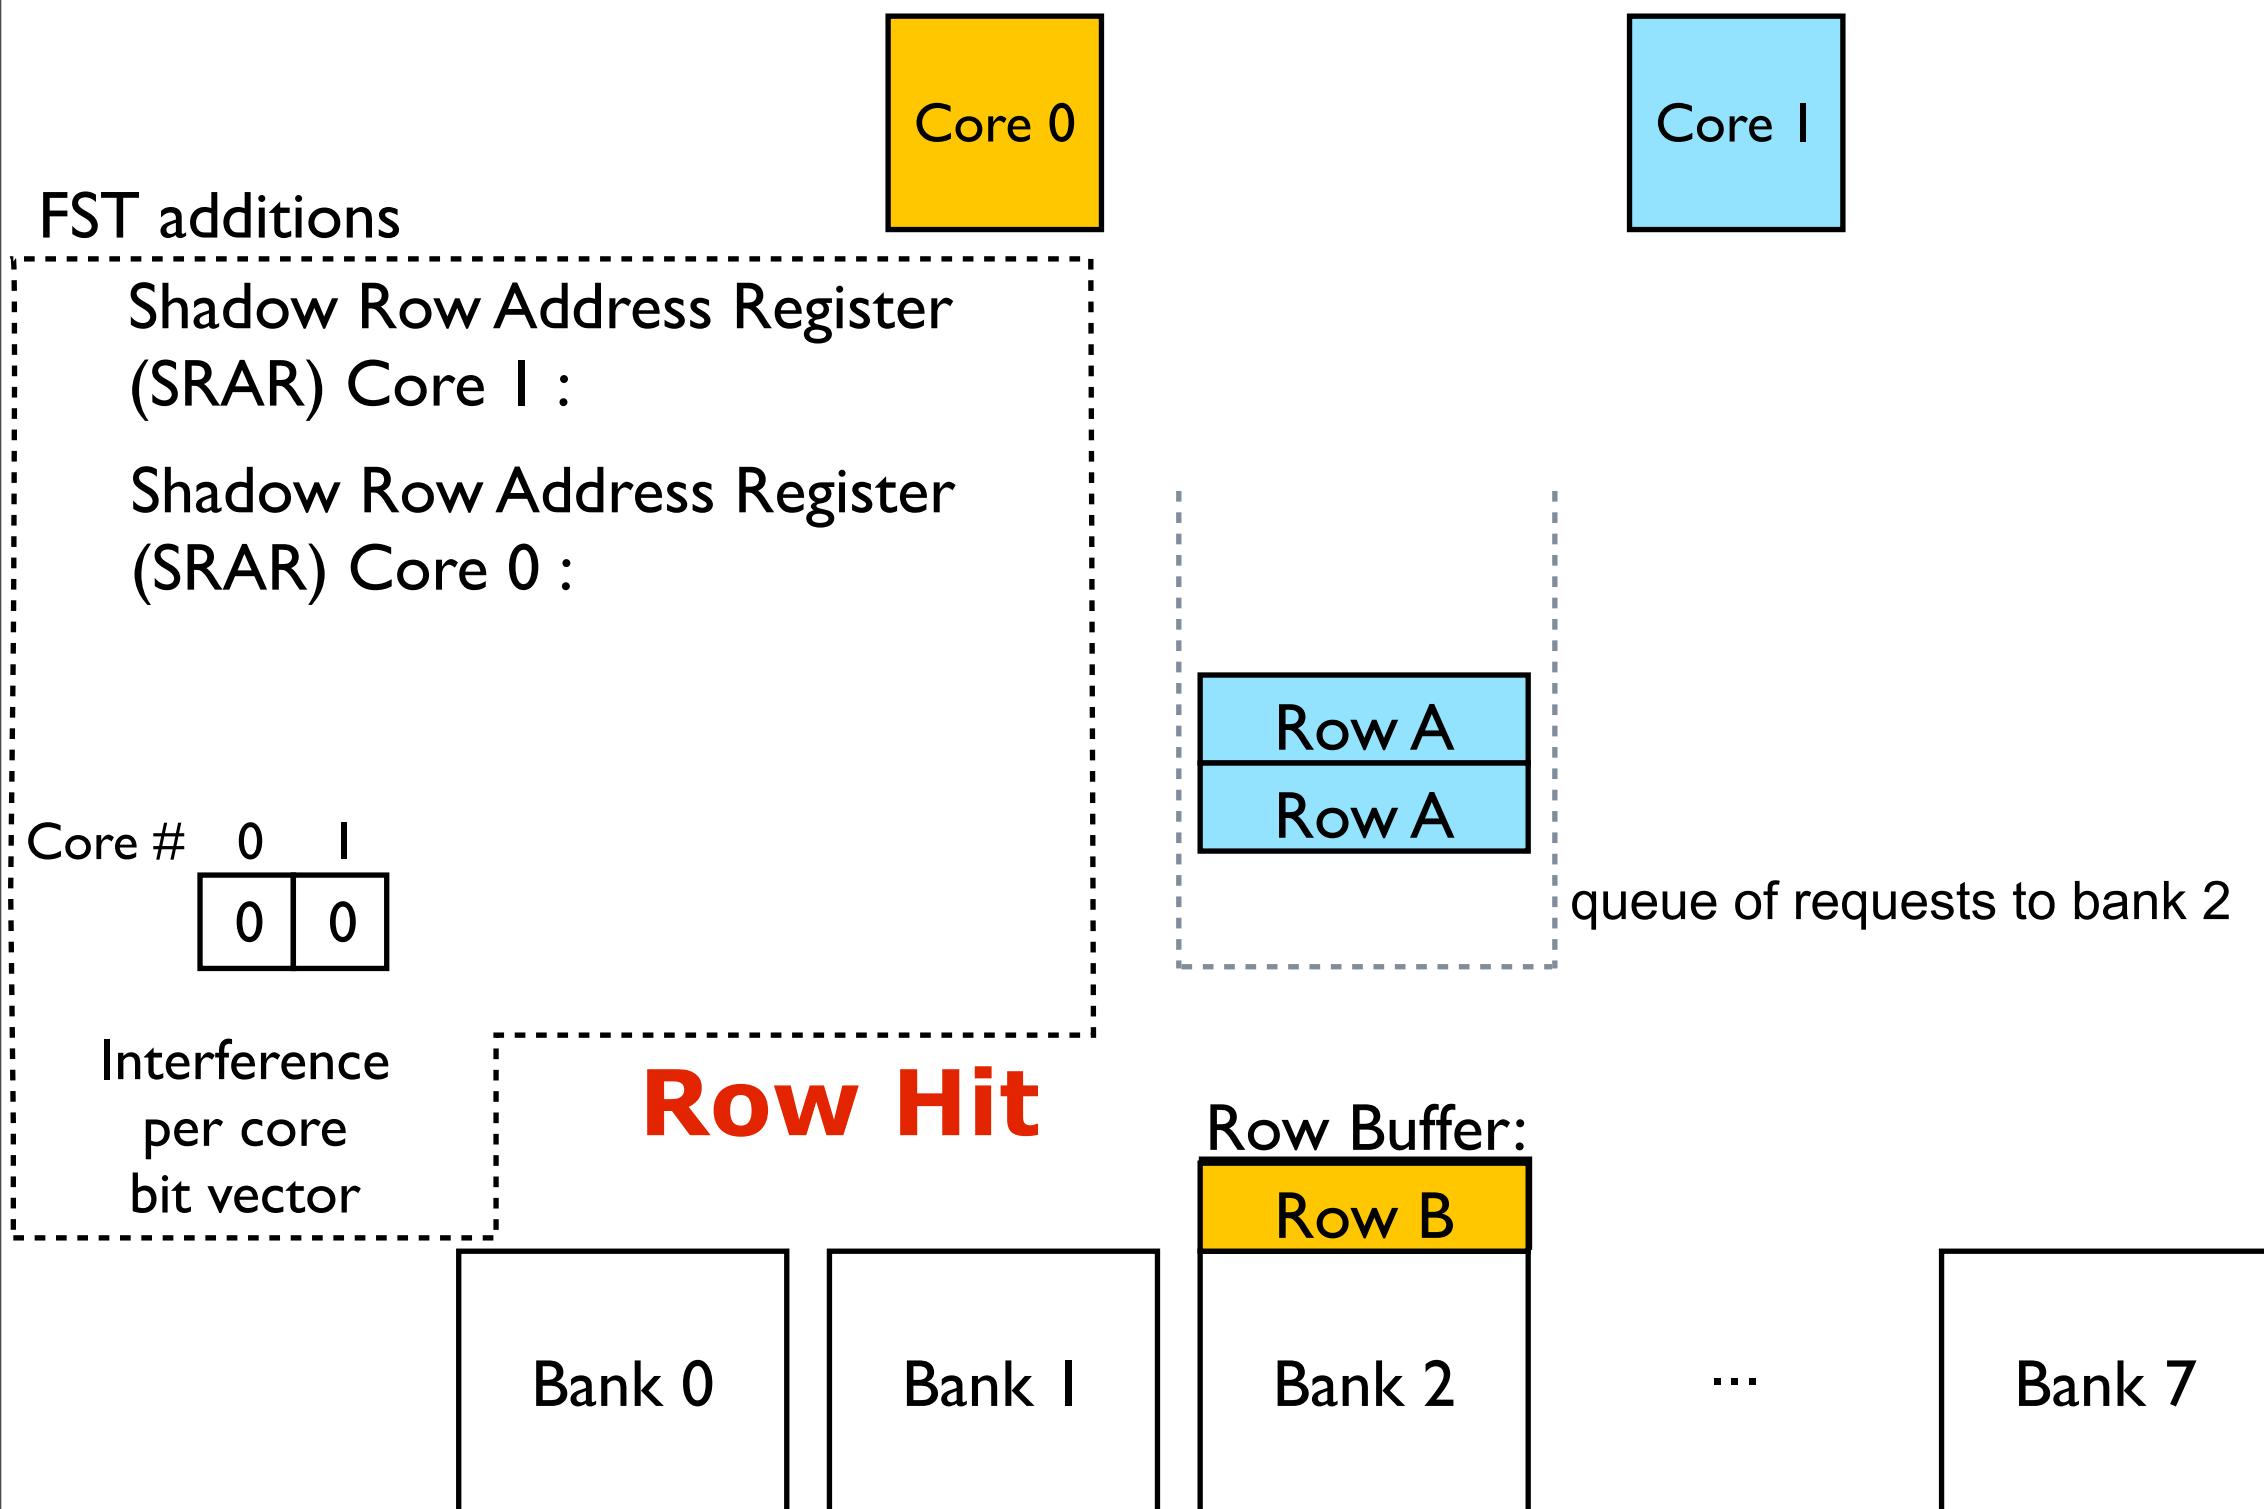
Tracking DRAM Row-Buffer Interference



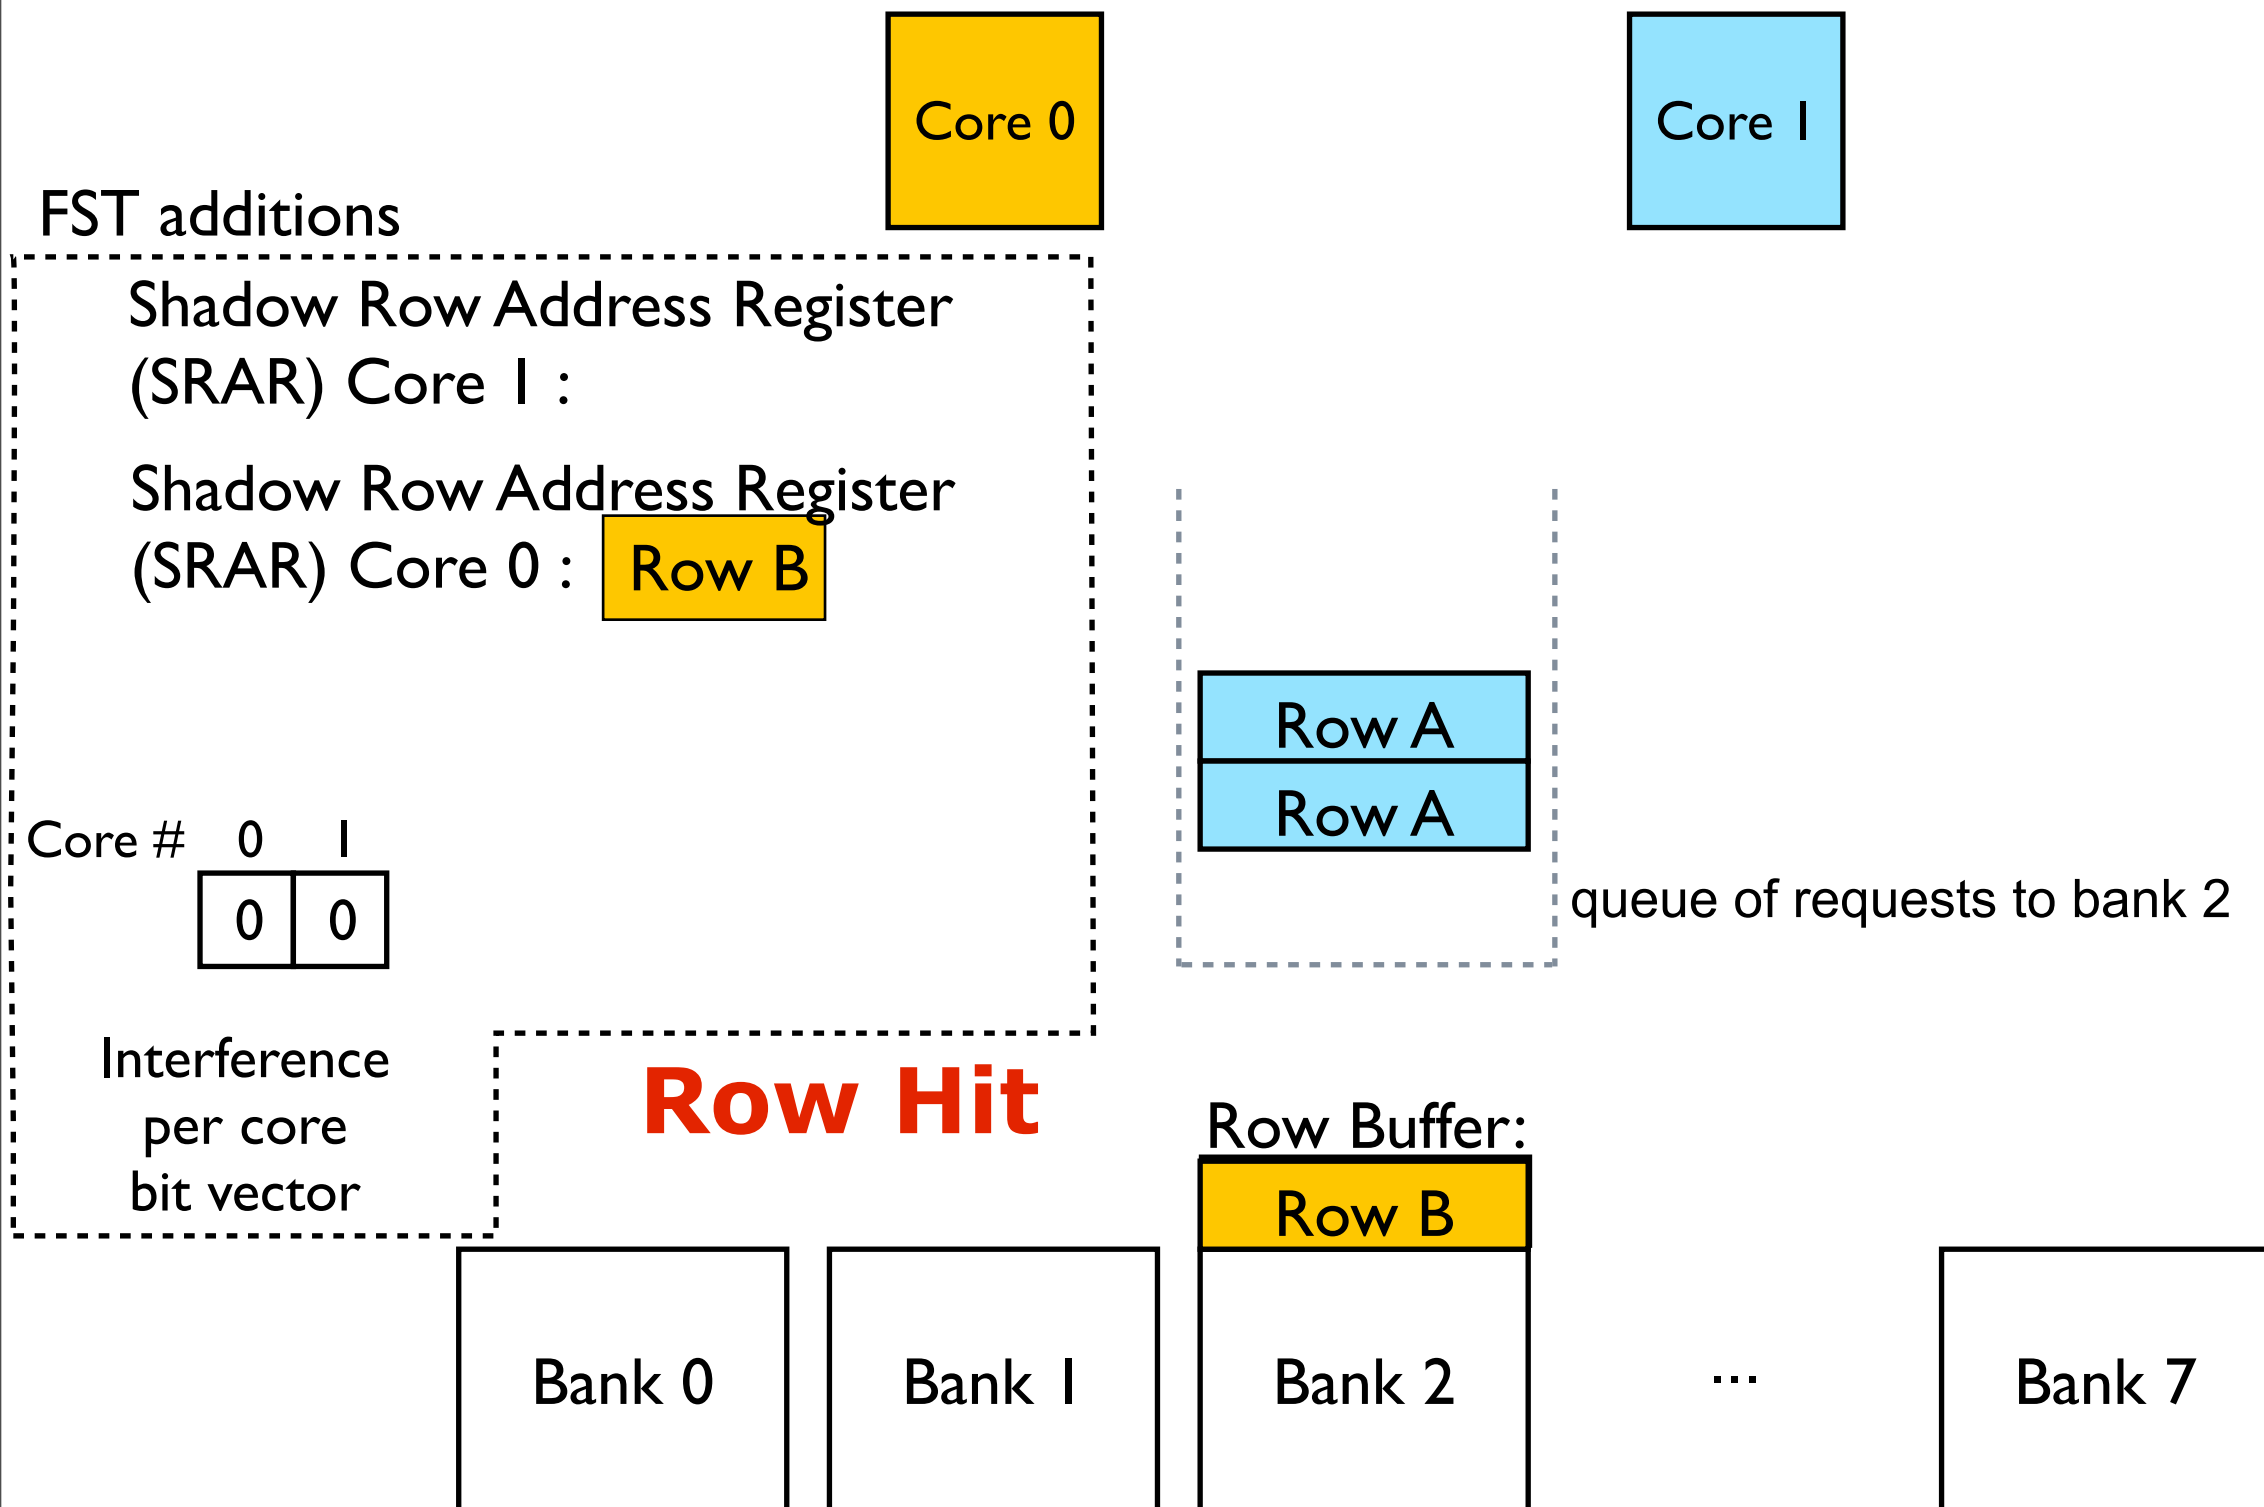
Tracking DRAM Row-Buffer Interference



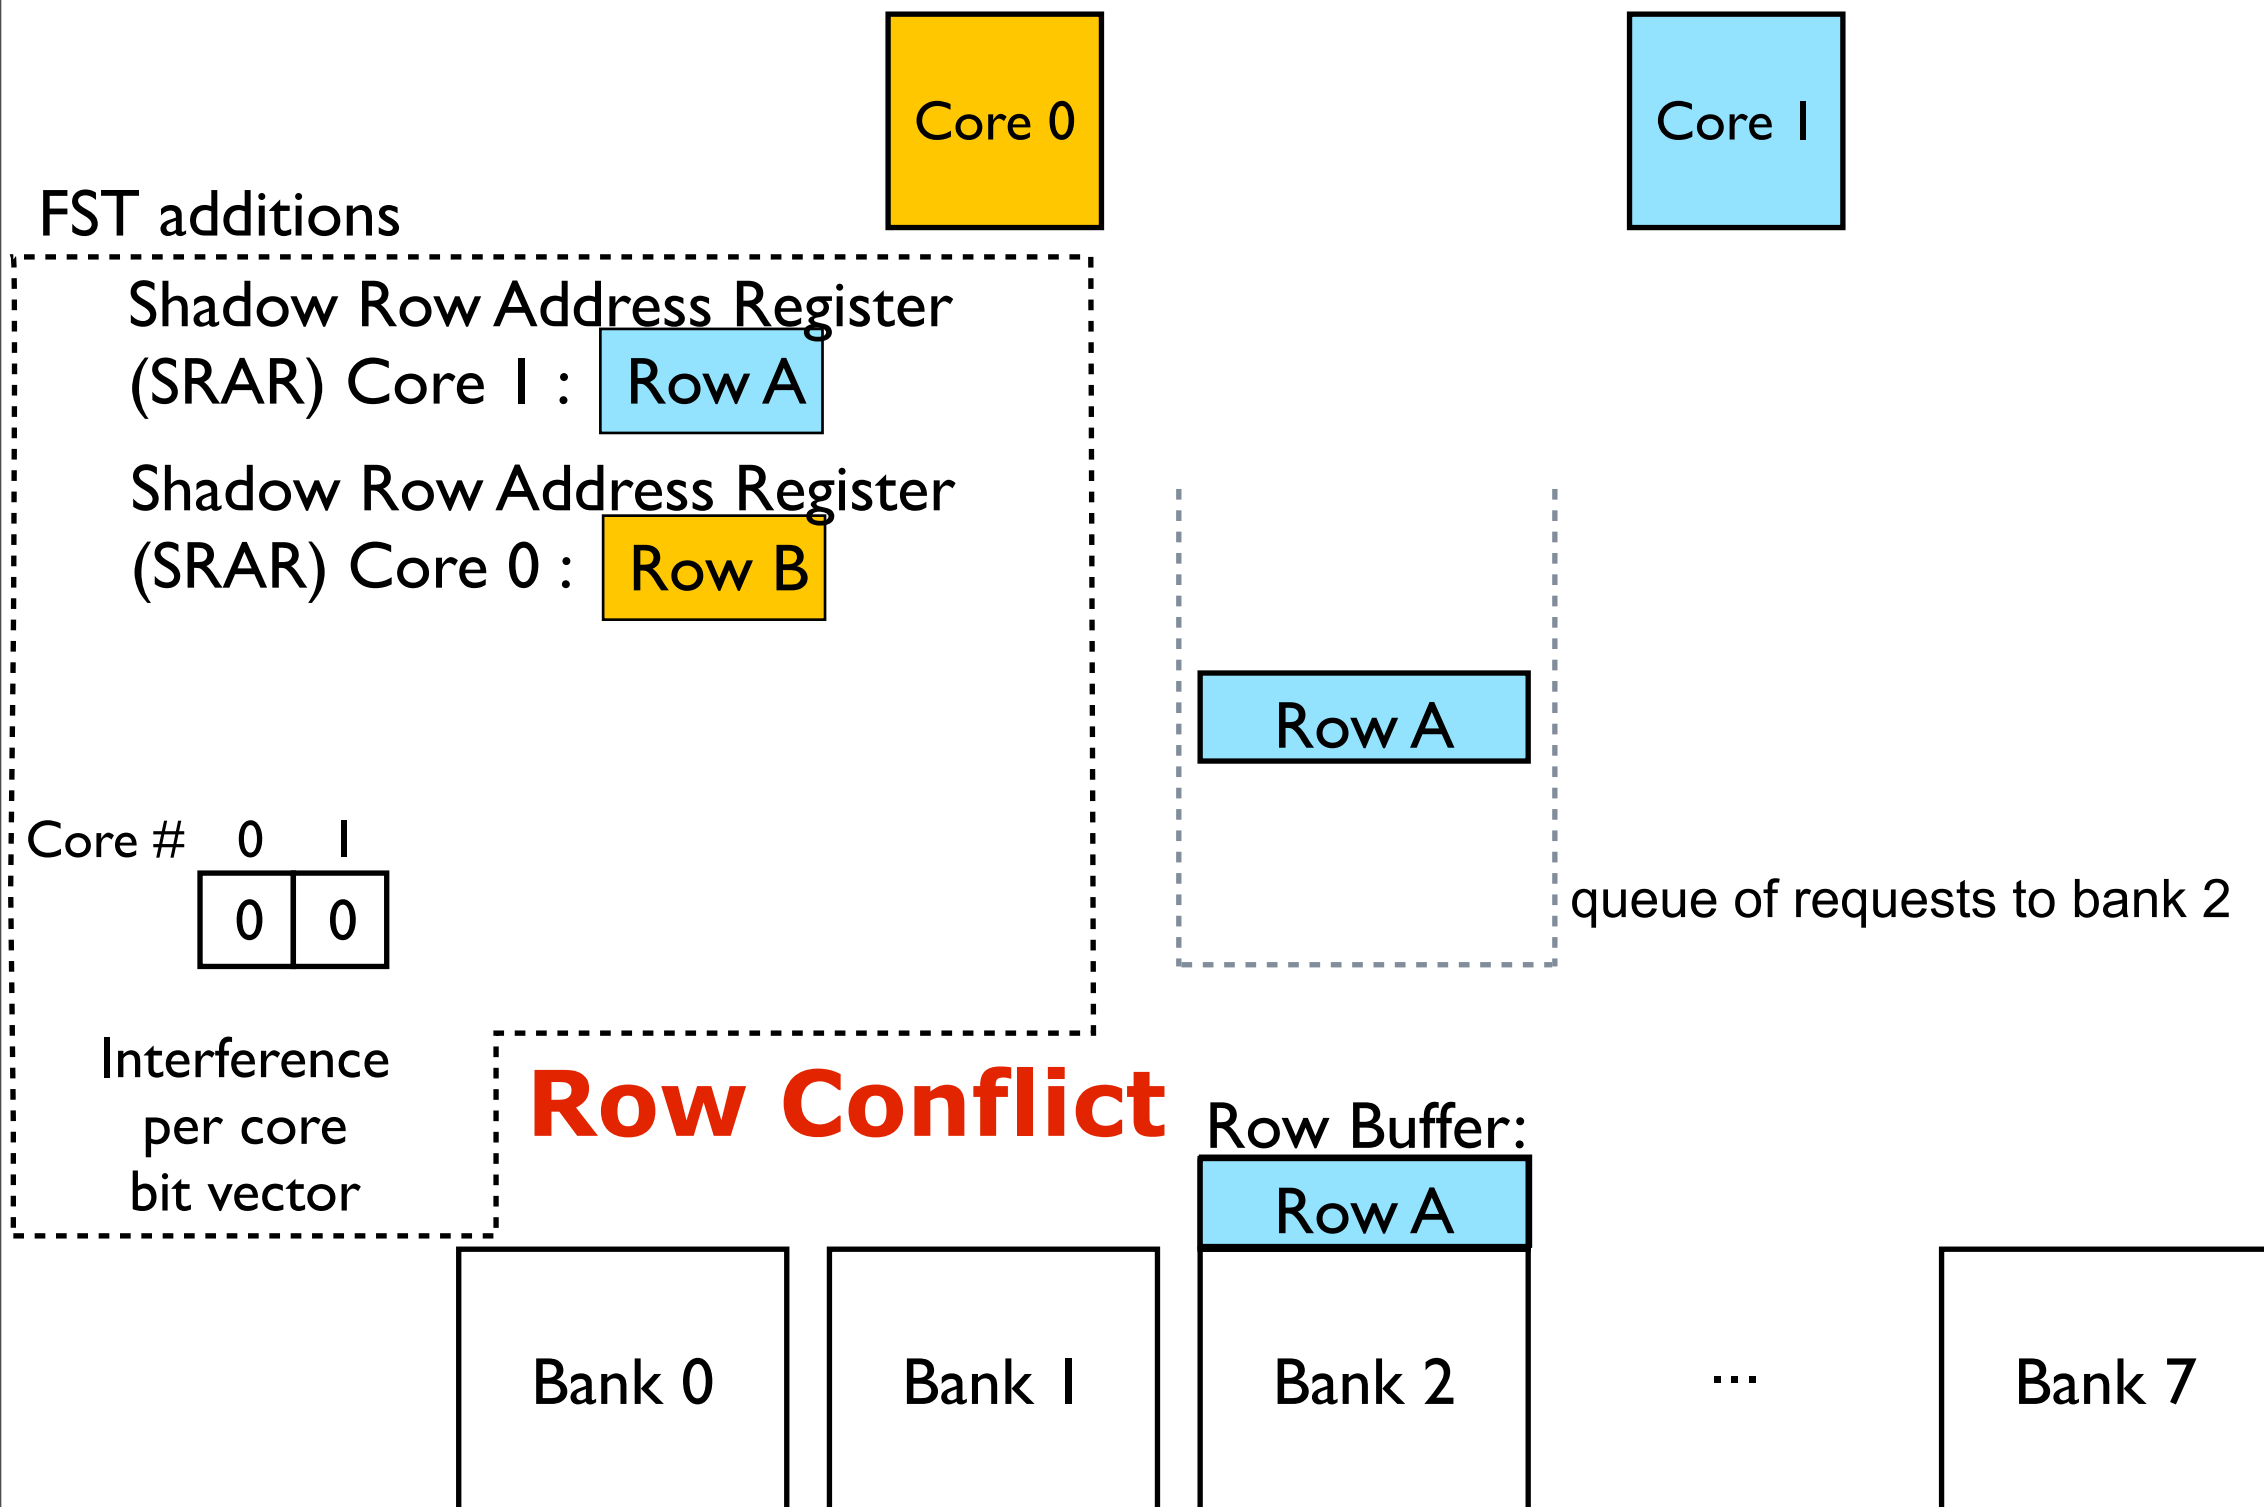
Tracking DRAM Row-Buffer Interference



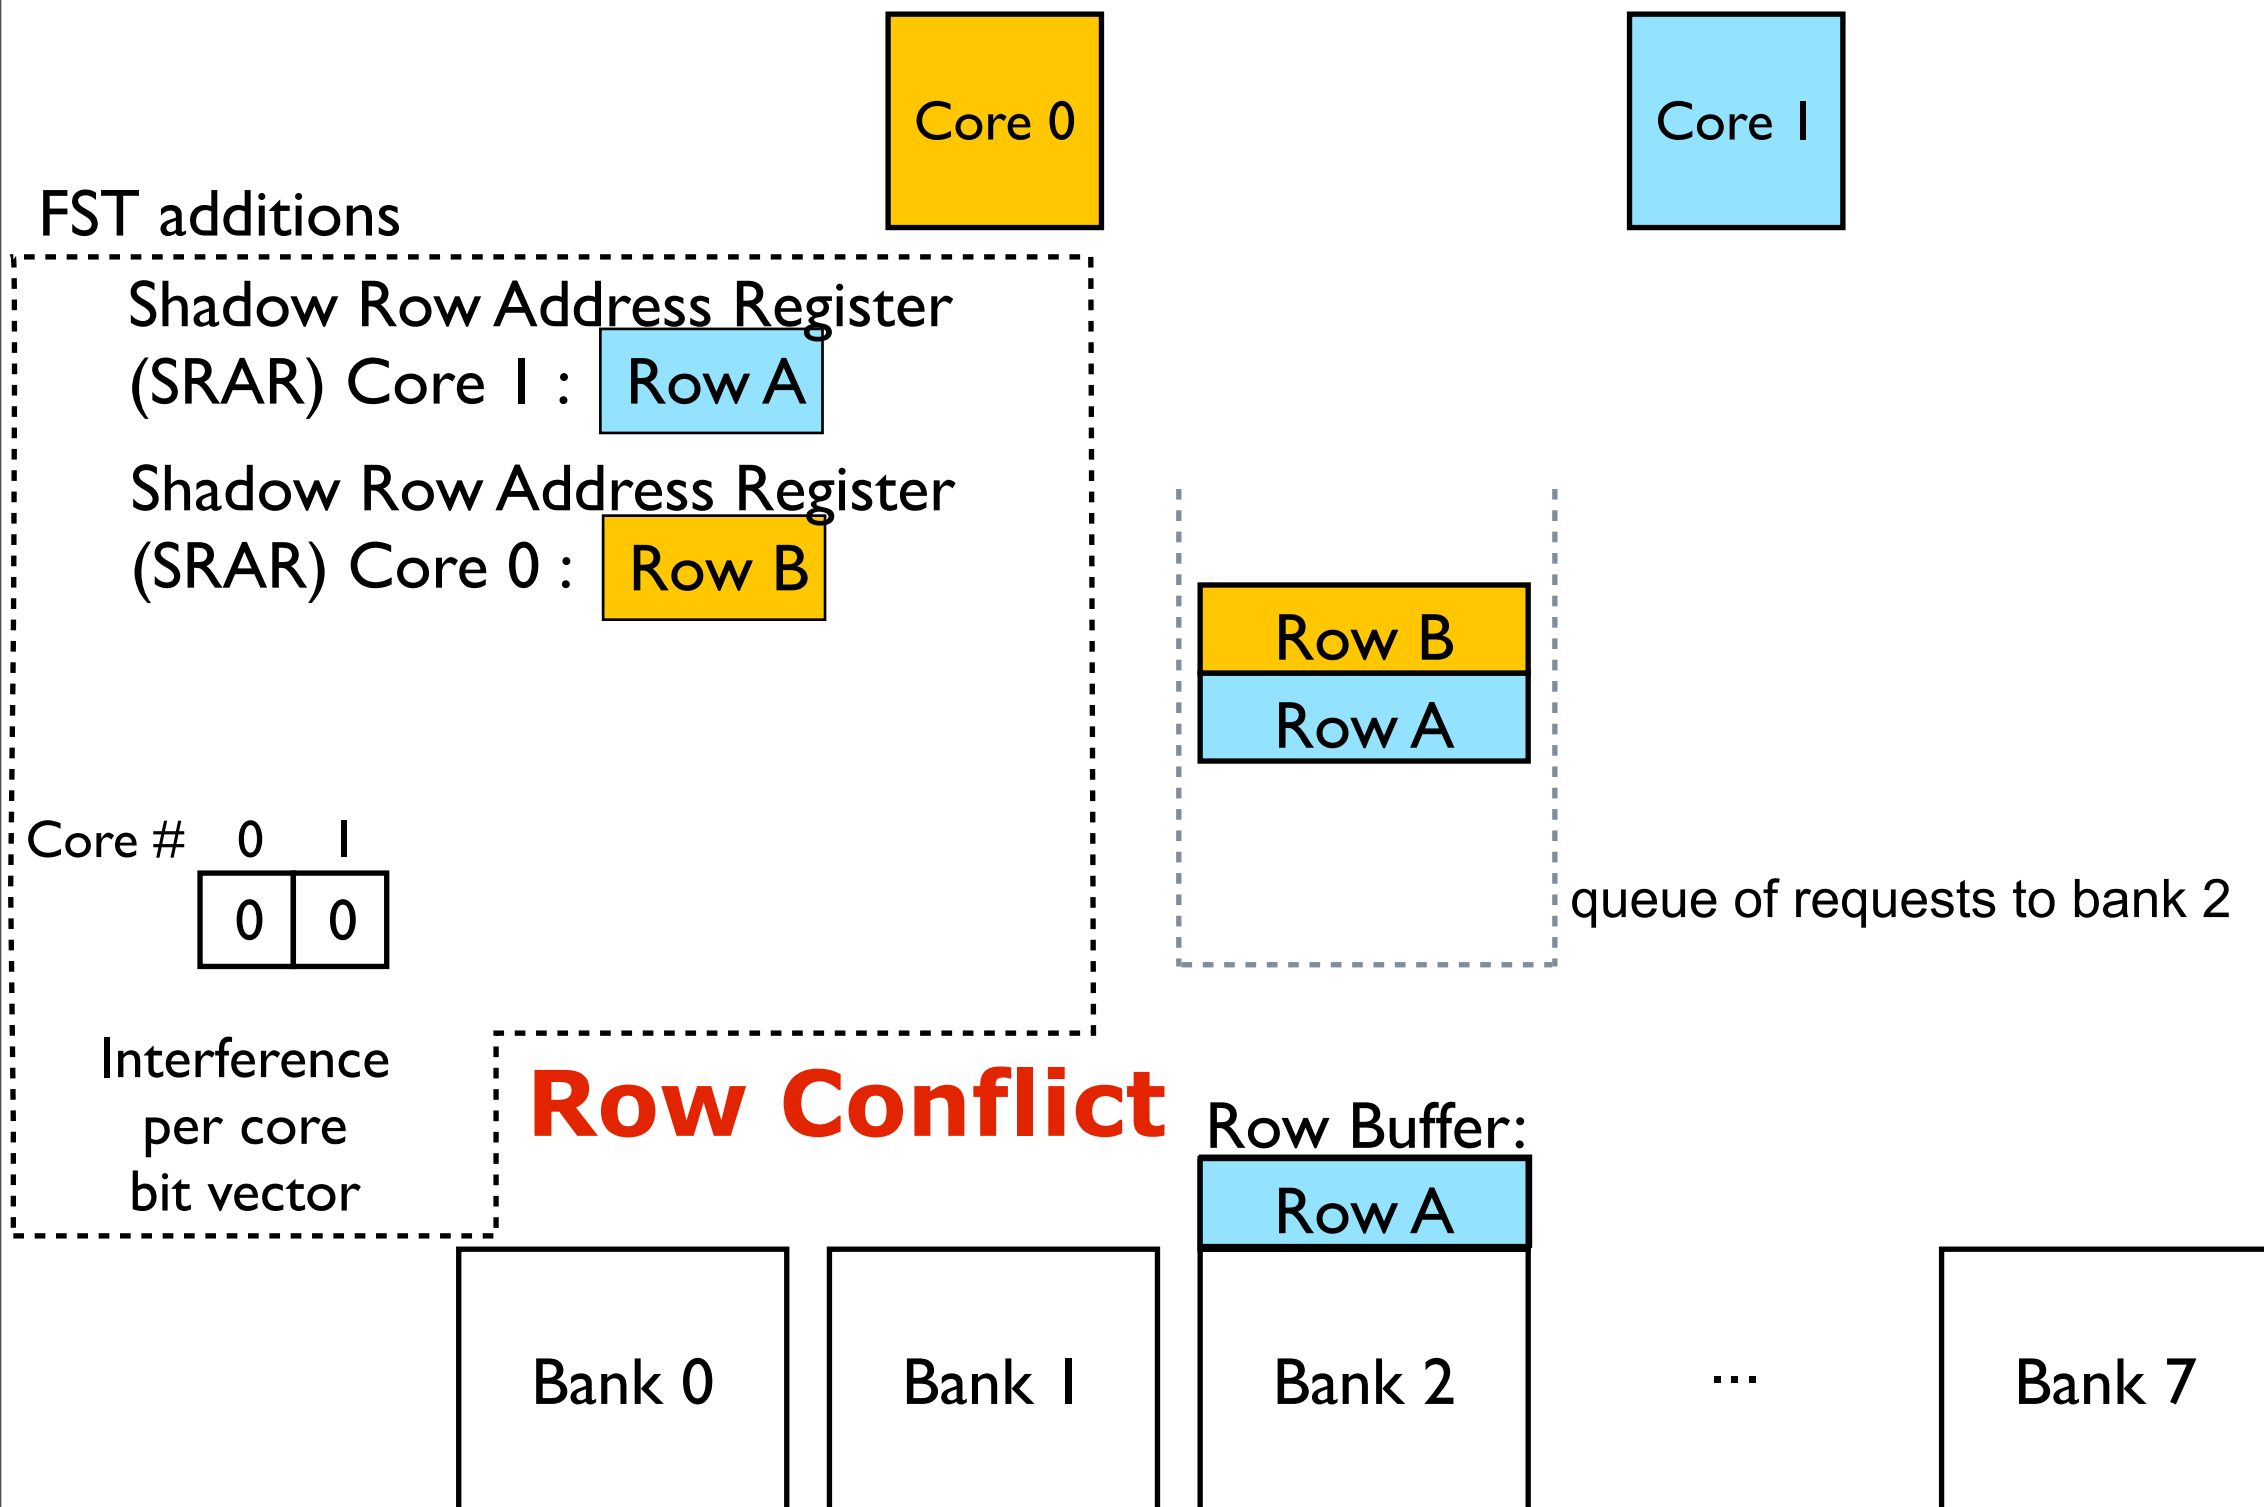
Tracking DRAM Row-Buffer Interference



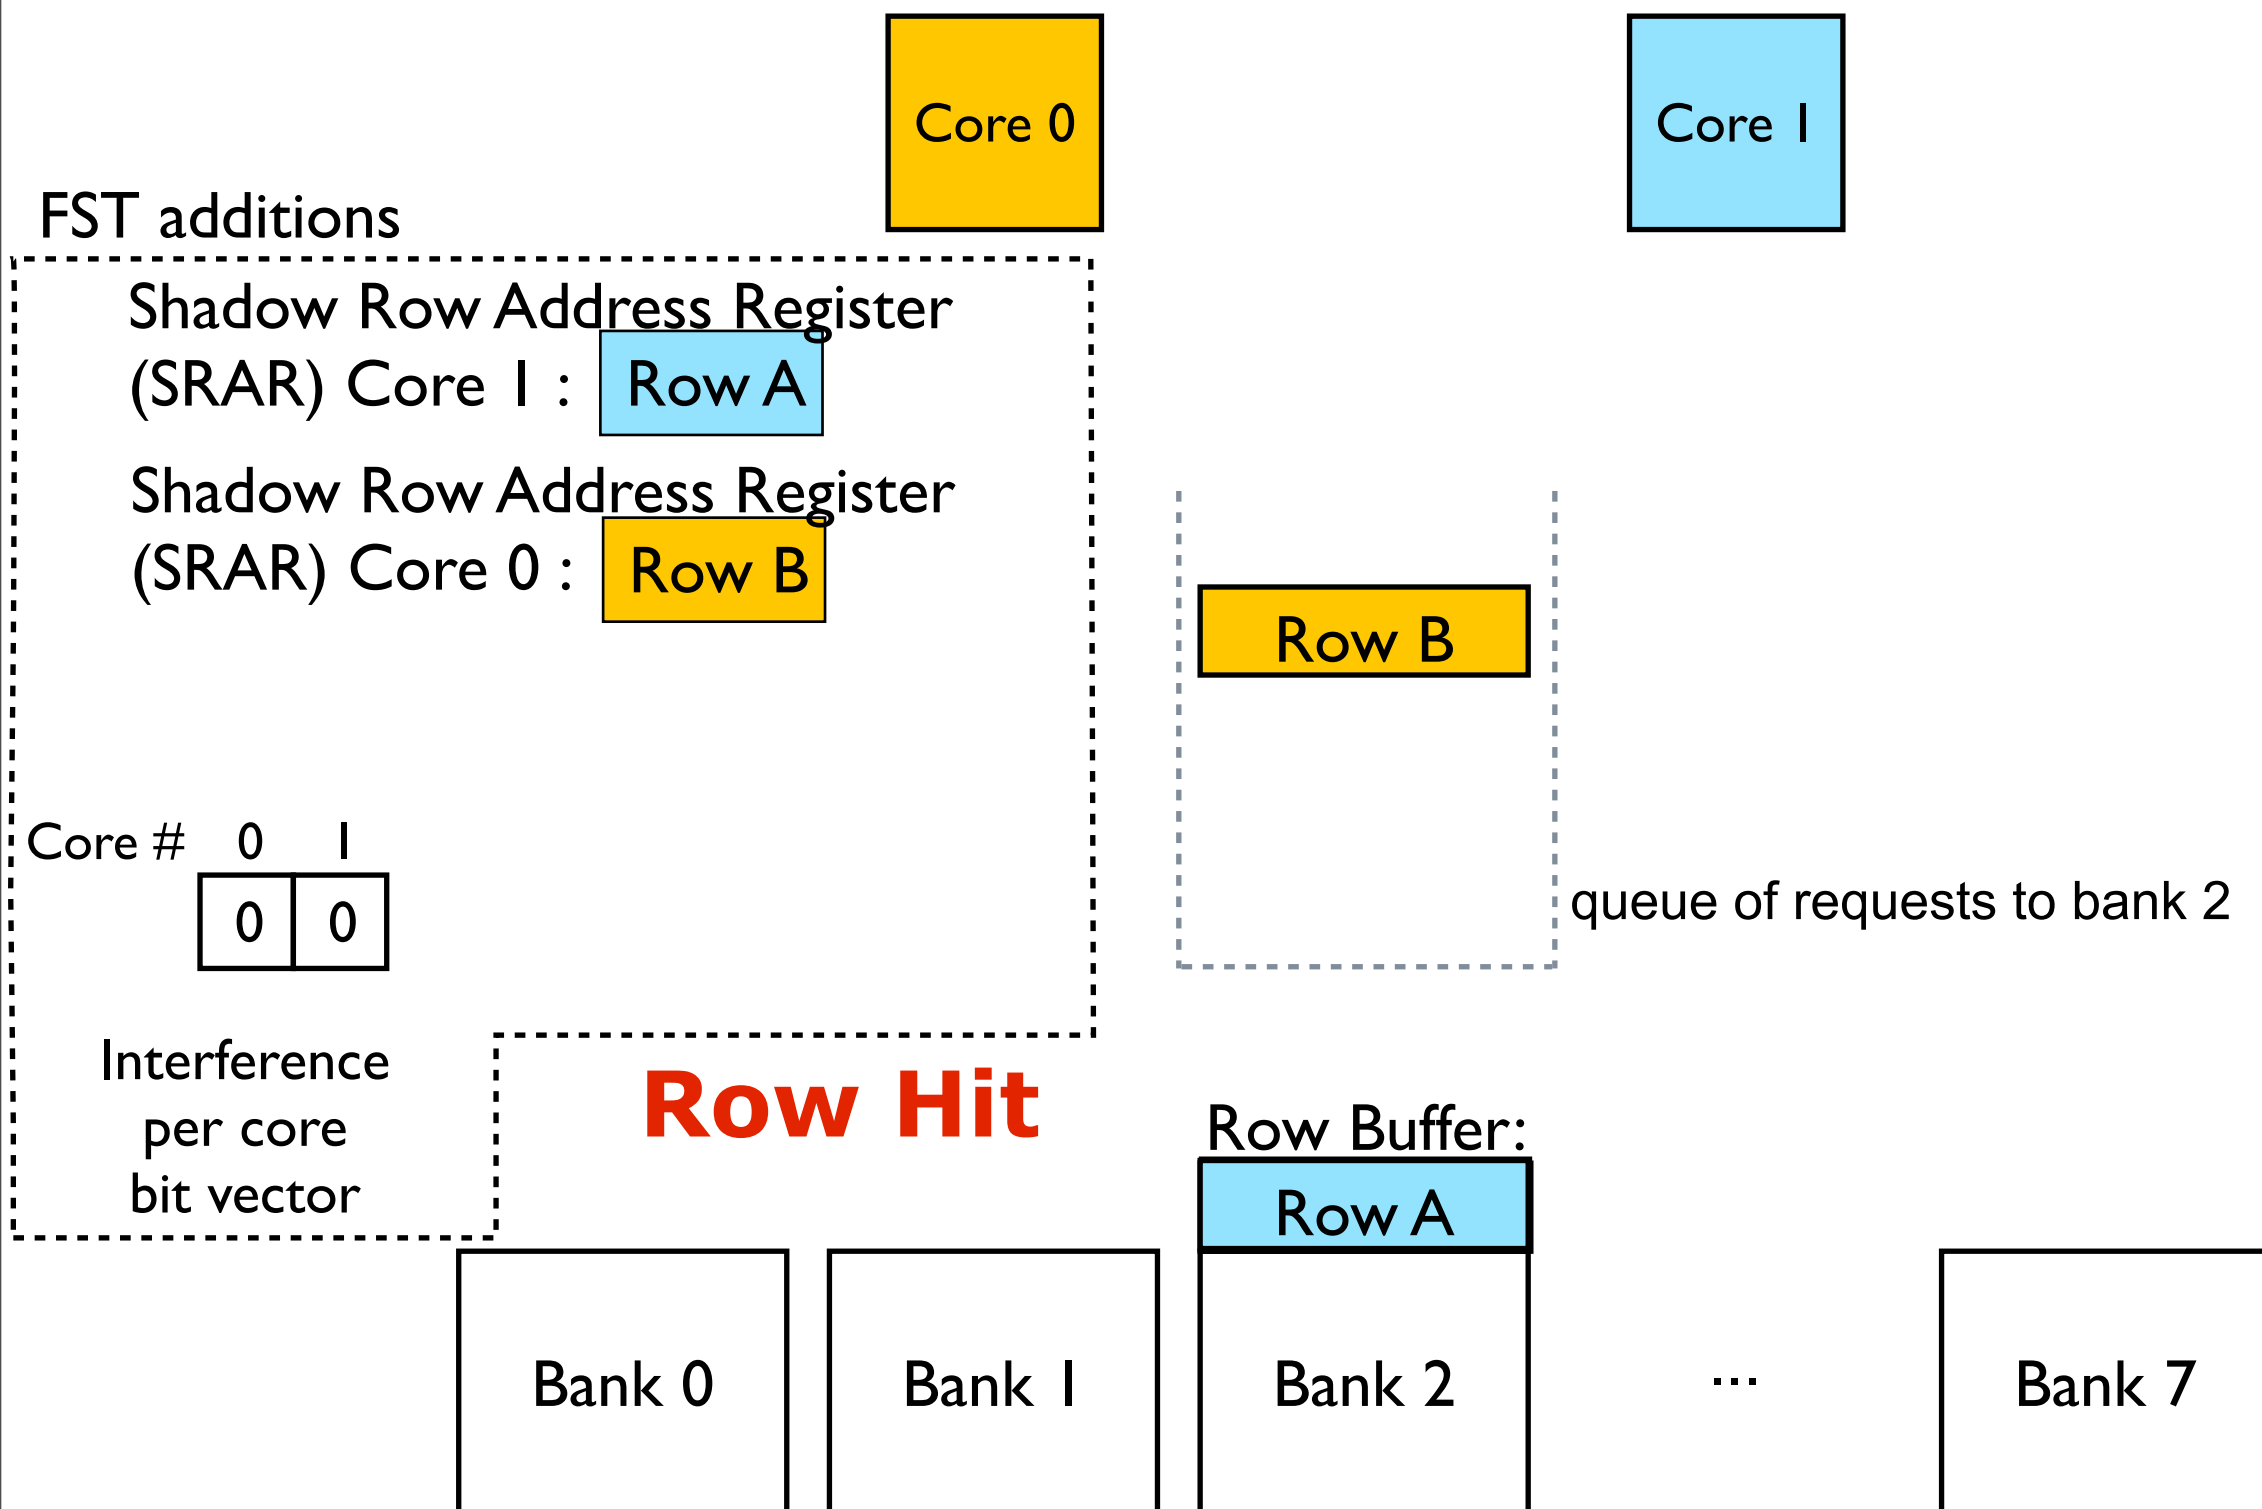
Tracking DRAM Row-Buffer Interference



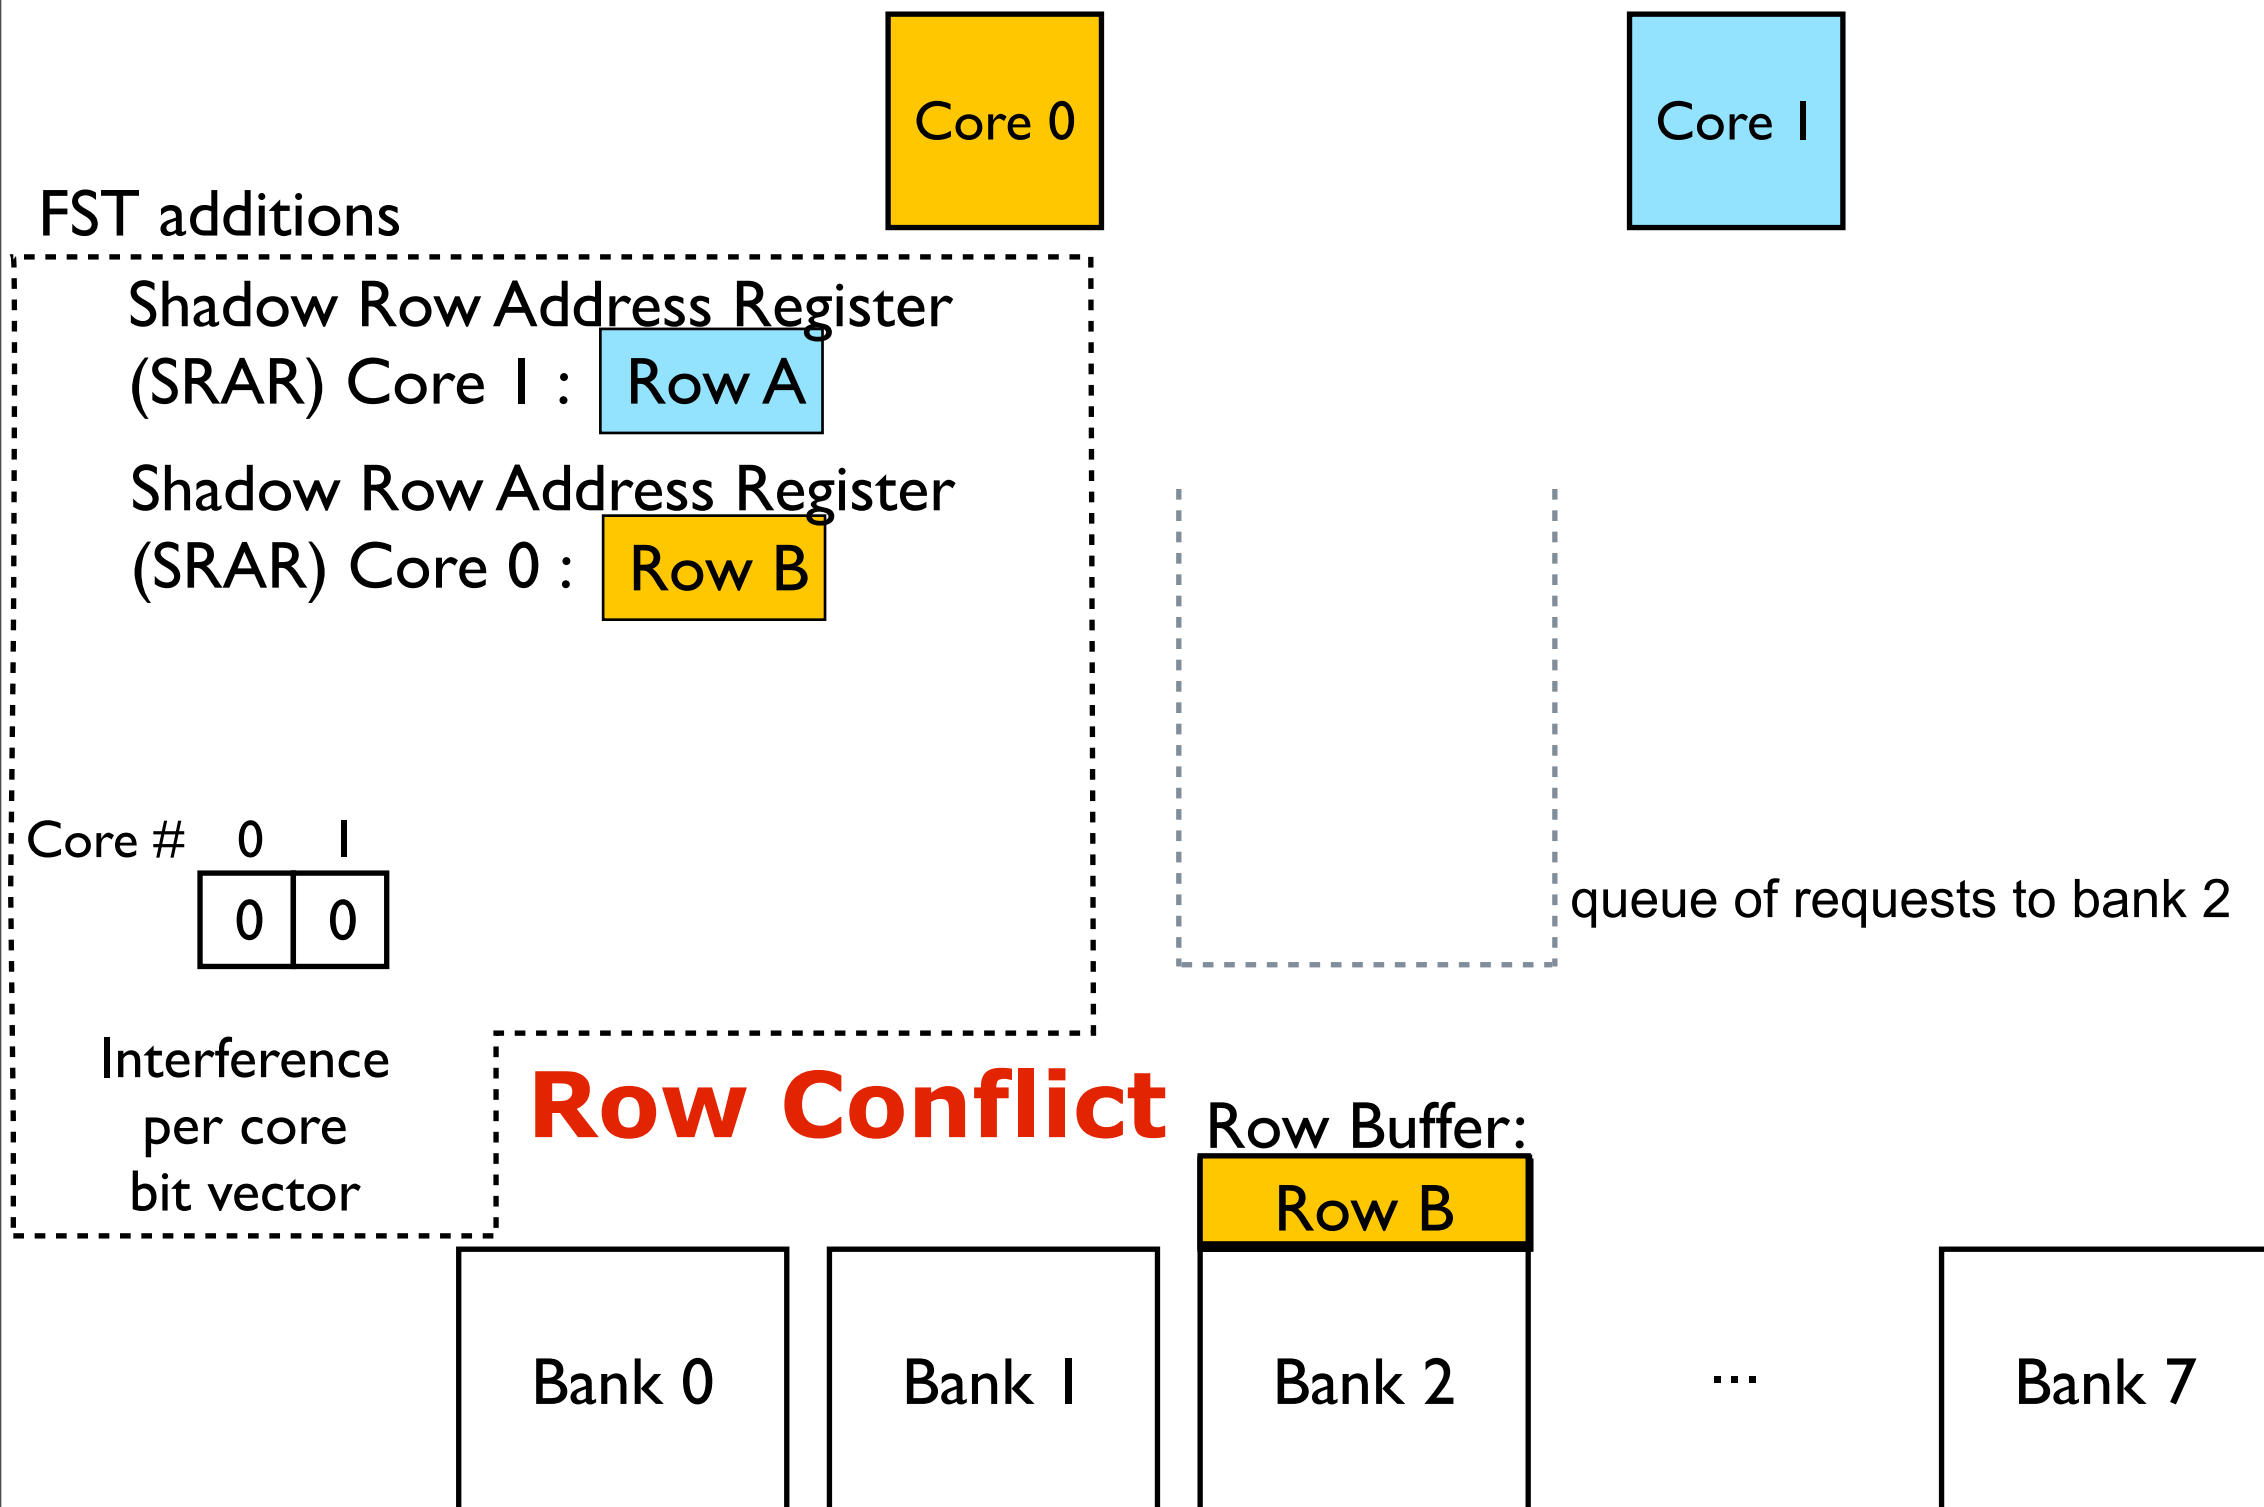
Tracking DRAM Row-Buffer Interference



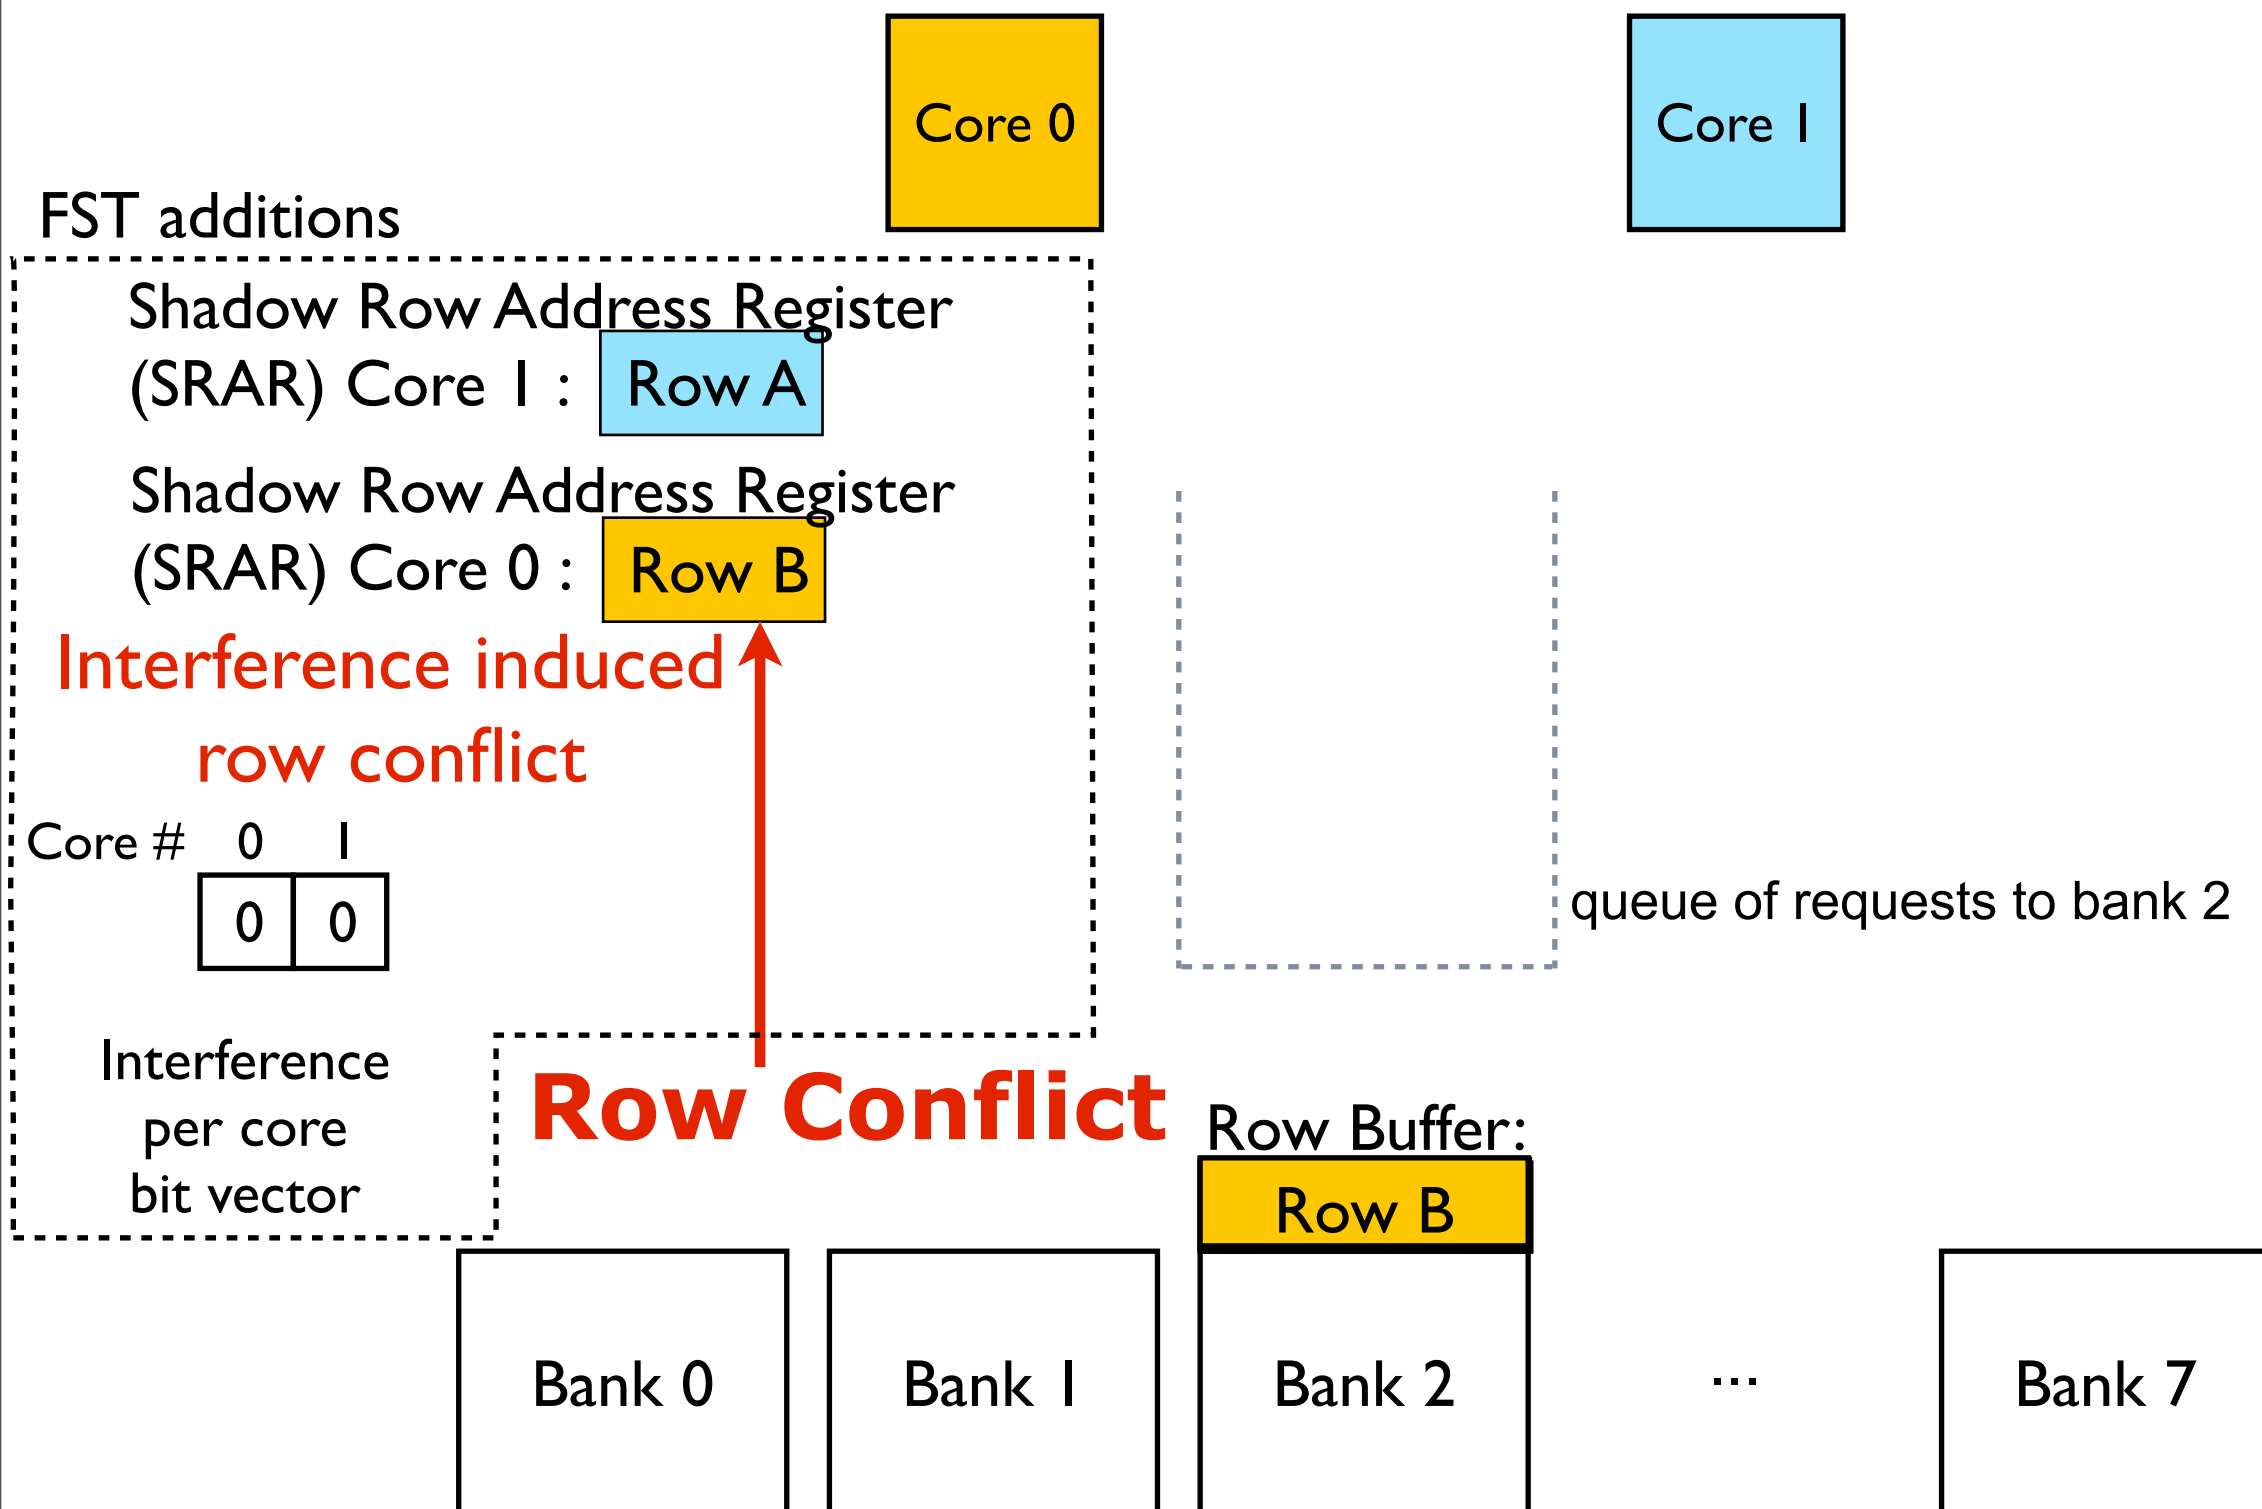
Tracking DRAM Row-Buffer Interference



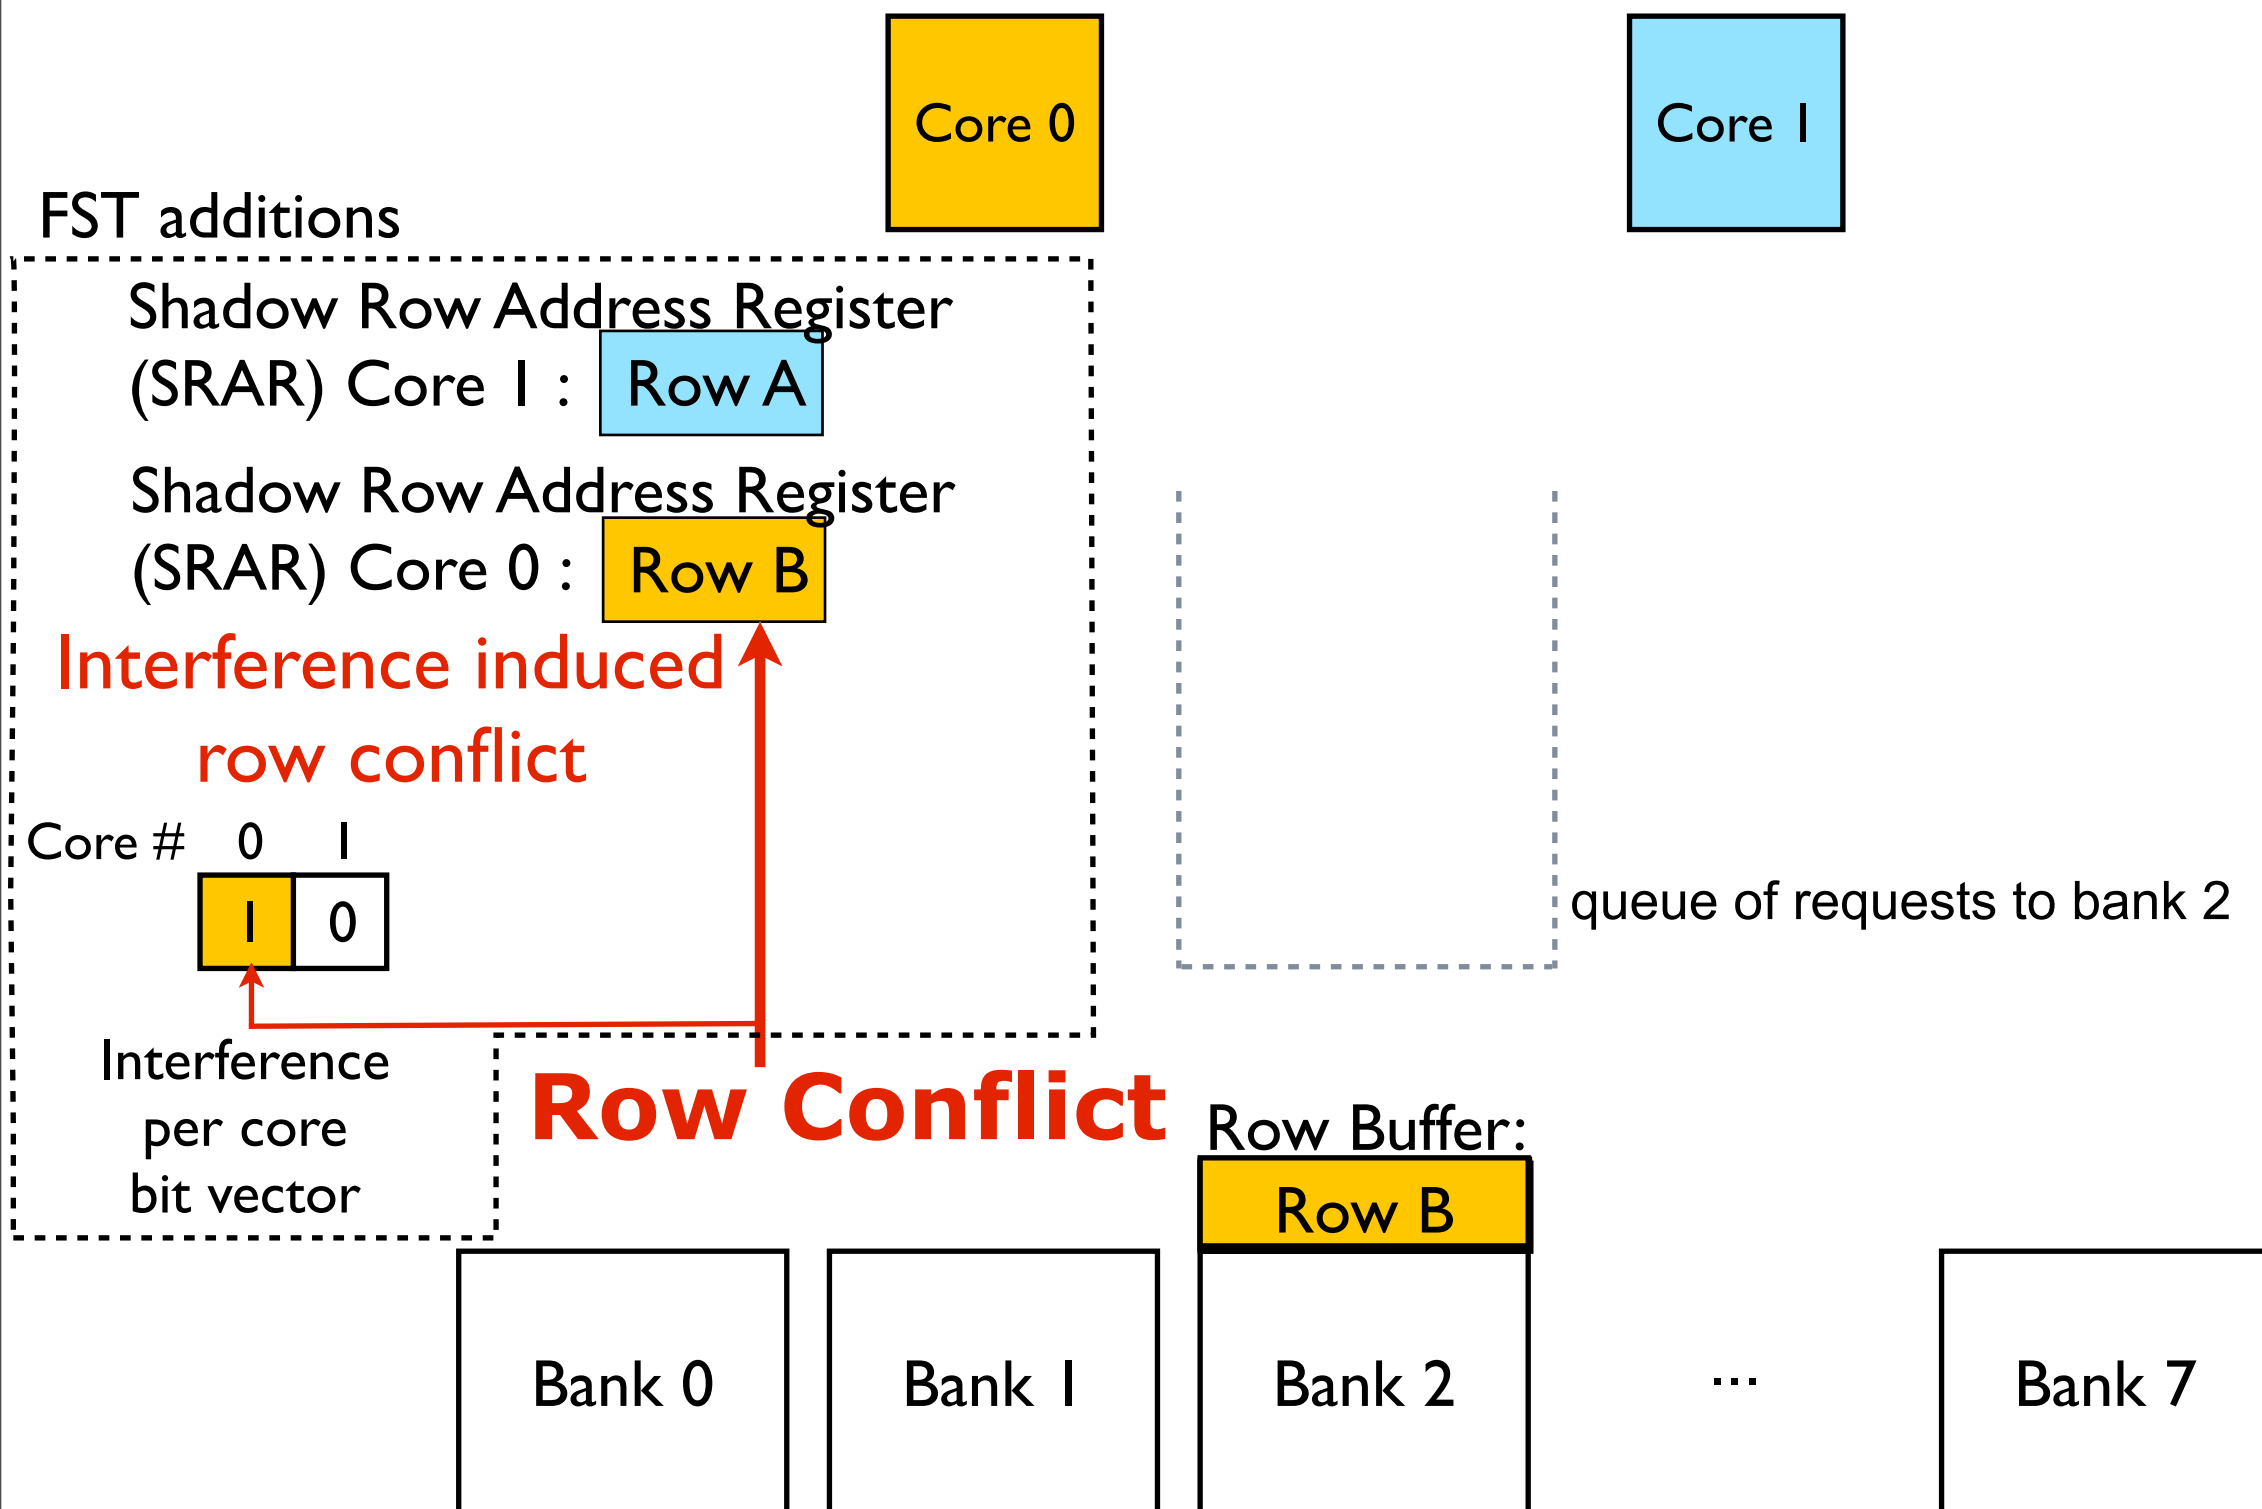
Tracking DRAM Row-Buffer Interference



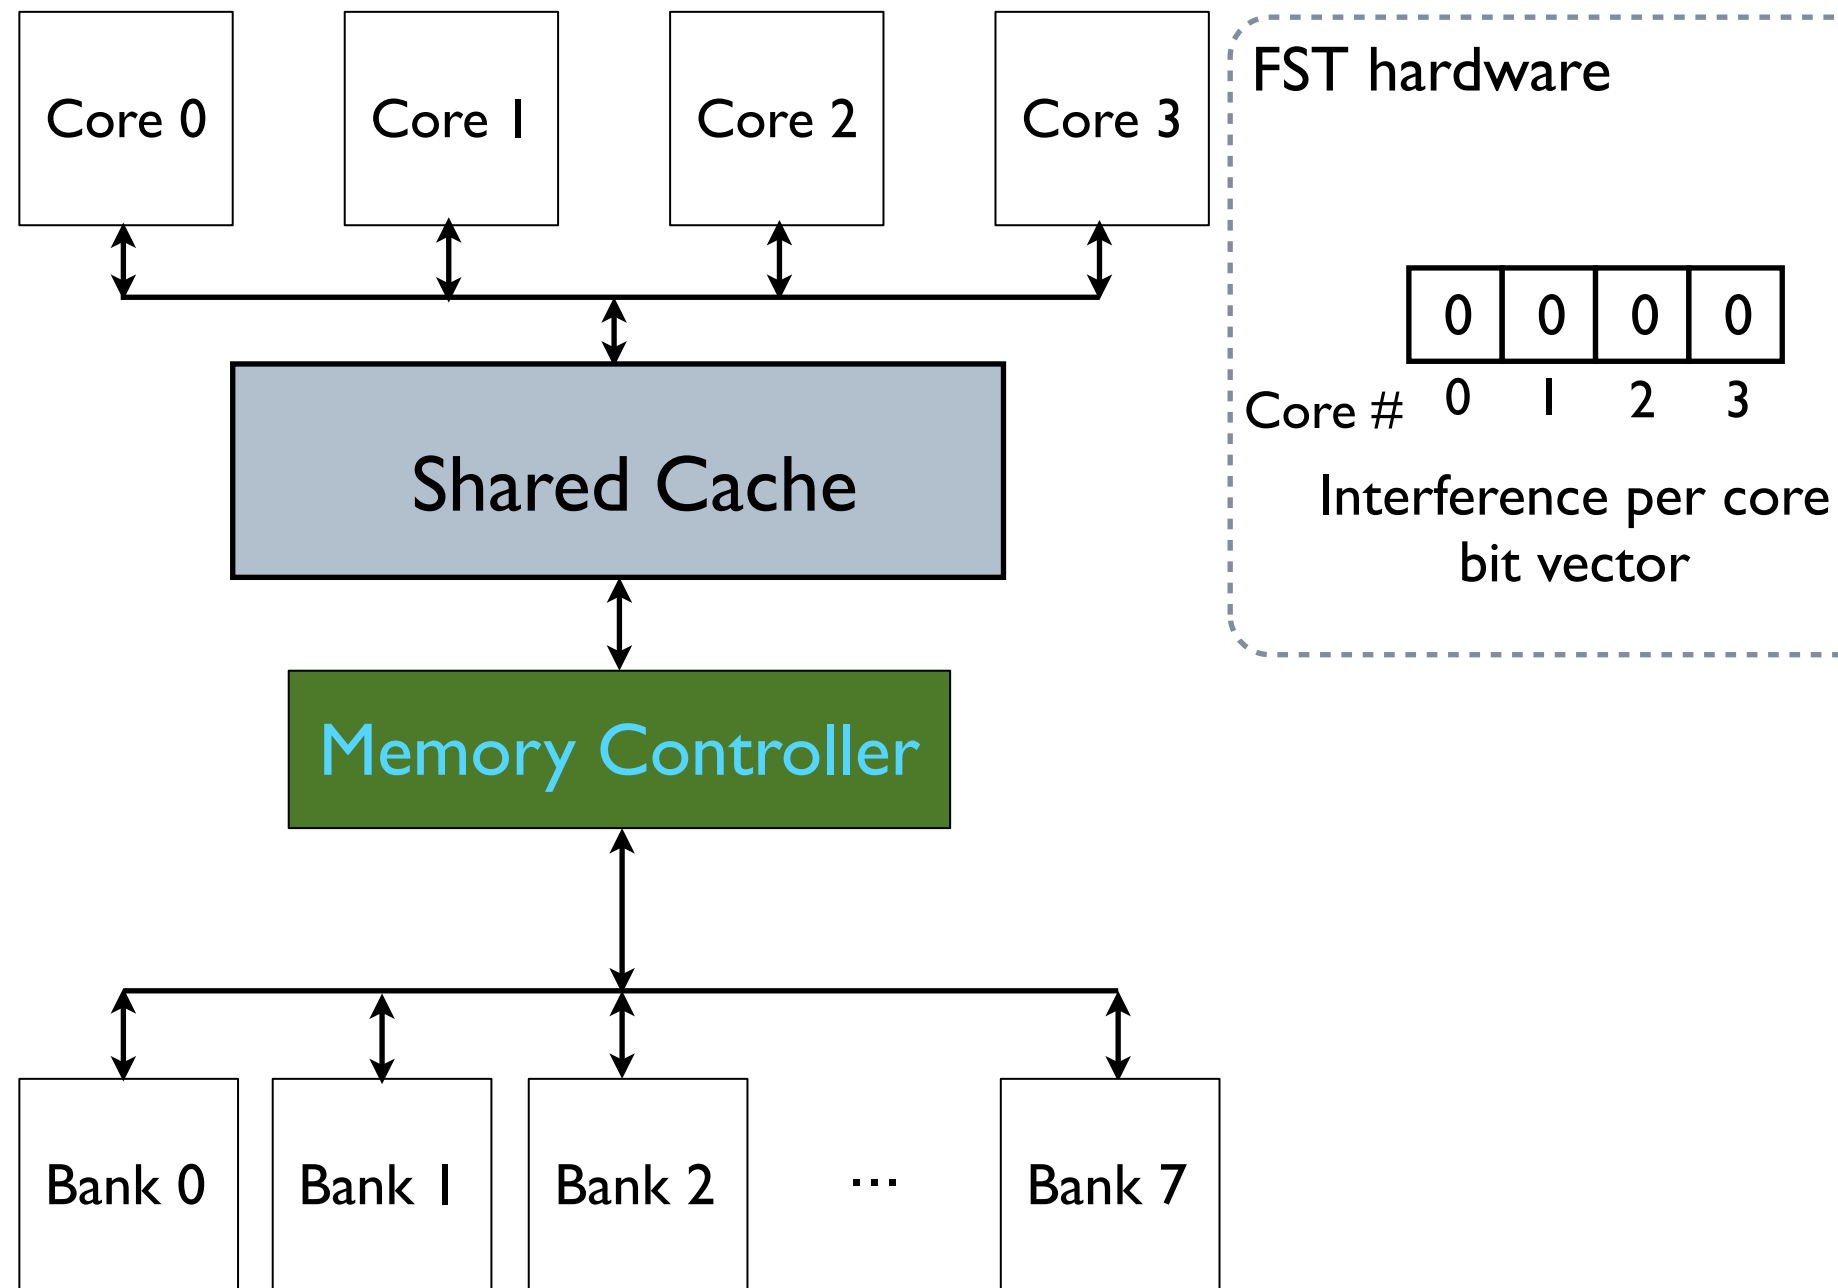
Tracking DRAM Row-Buffer Interference



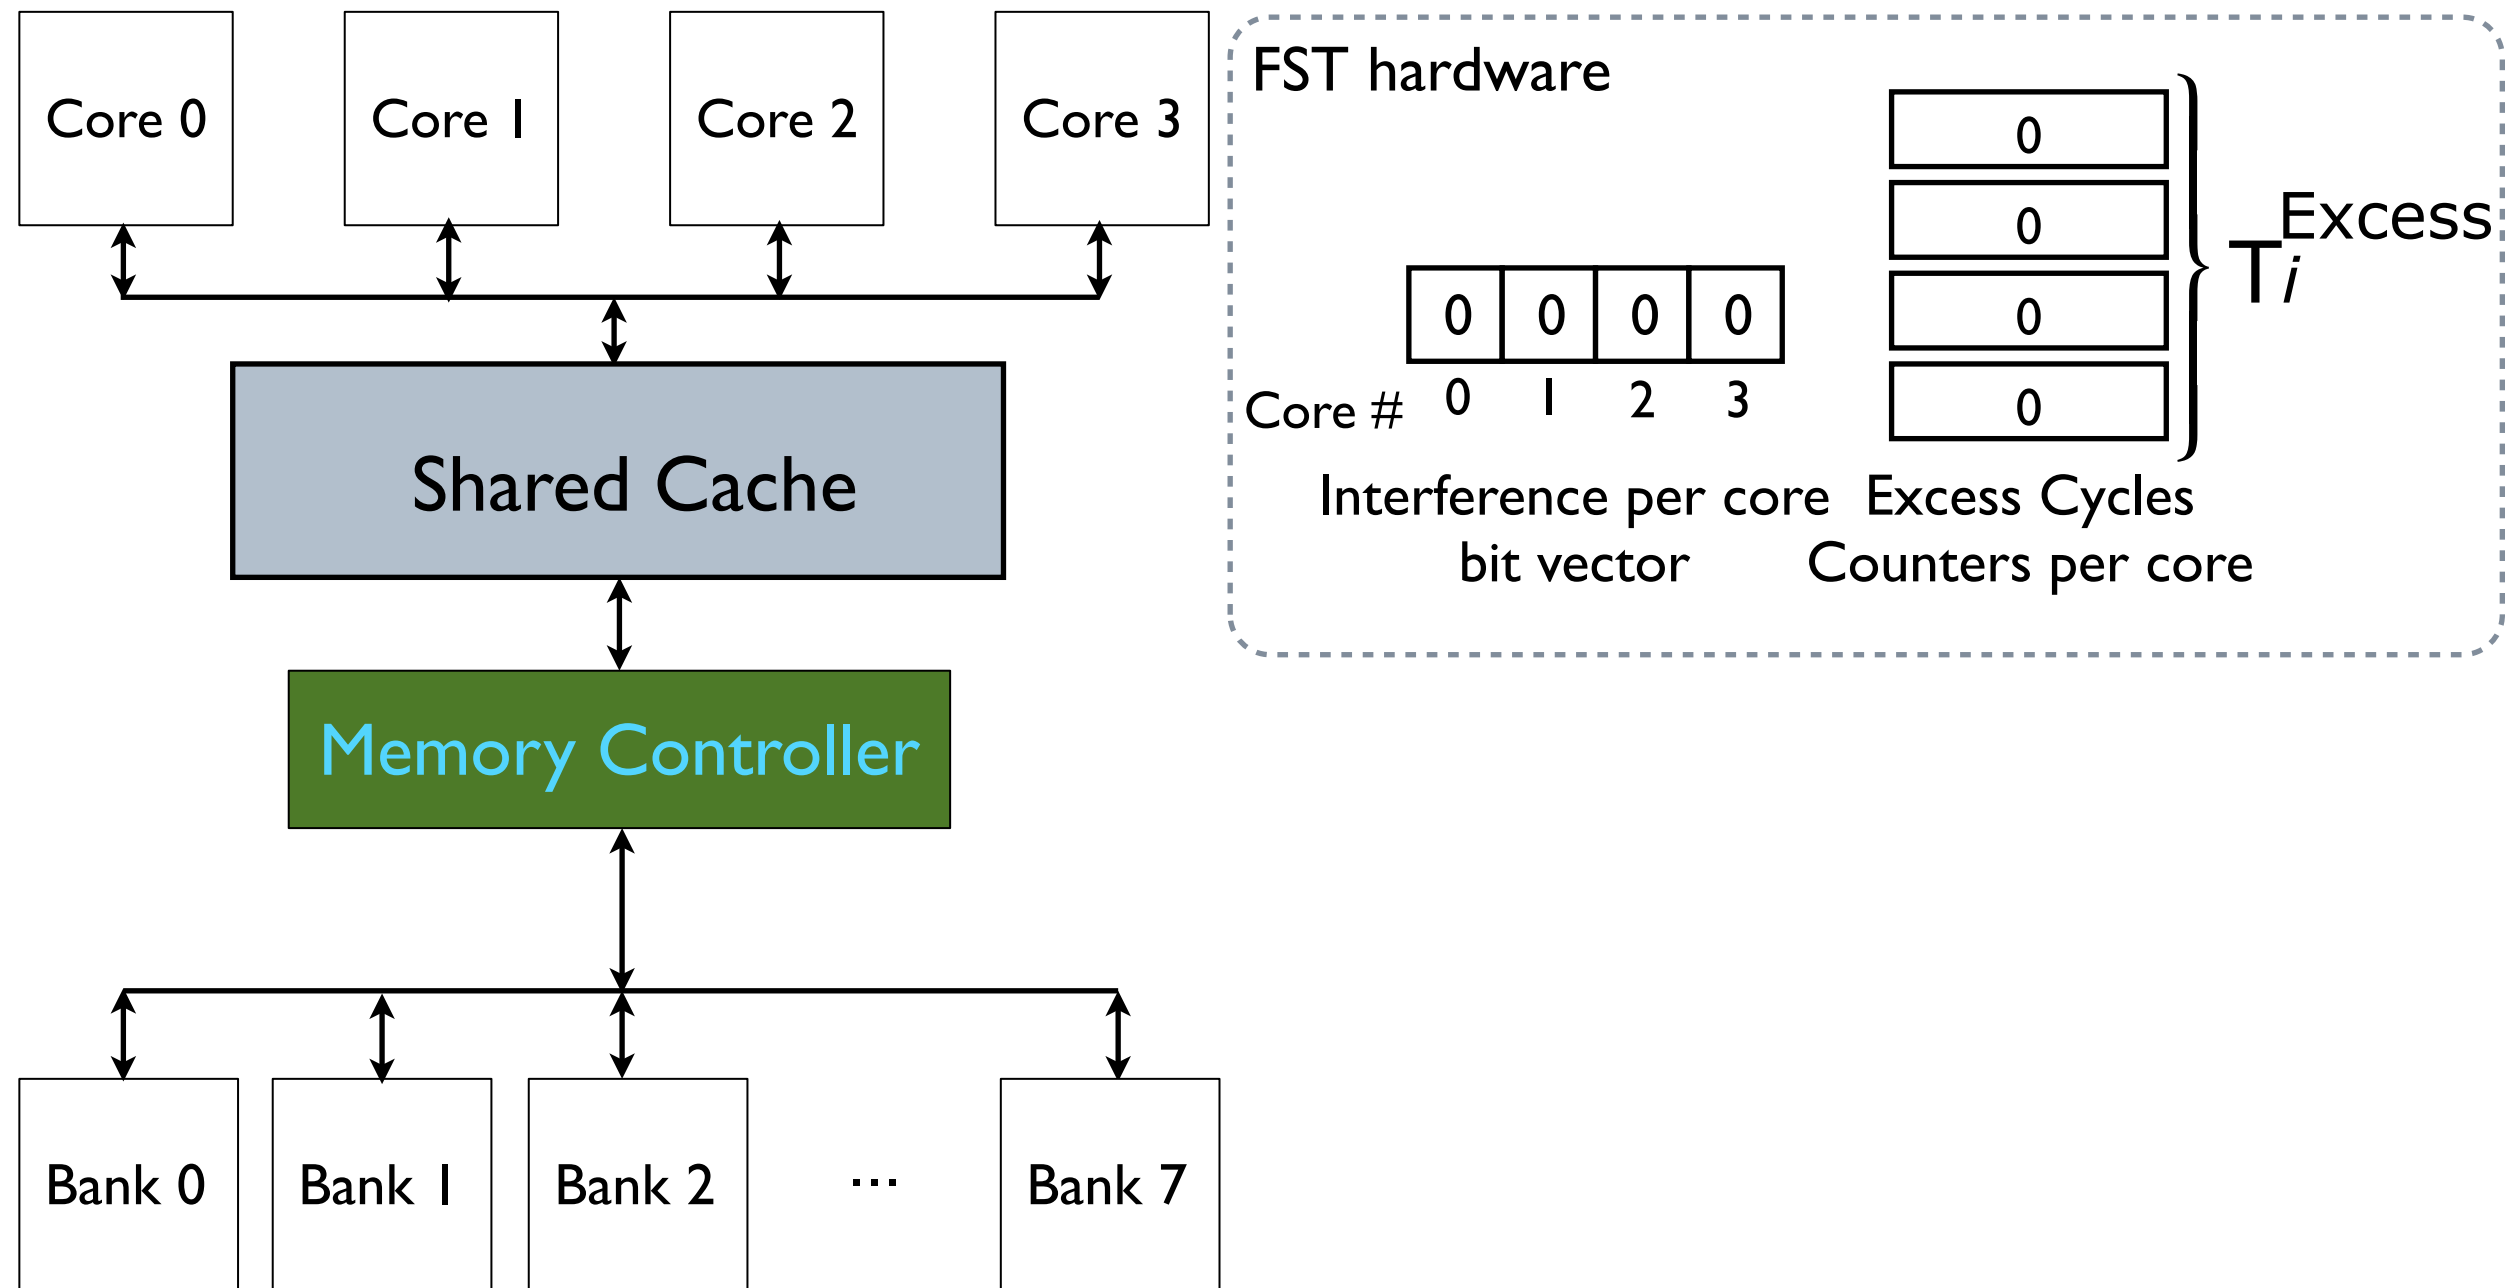
Tracking DRAM Row-Buffer Interference



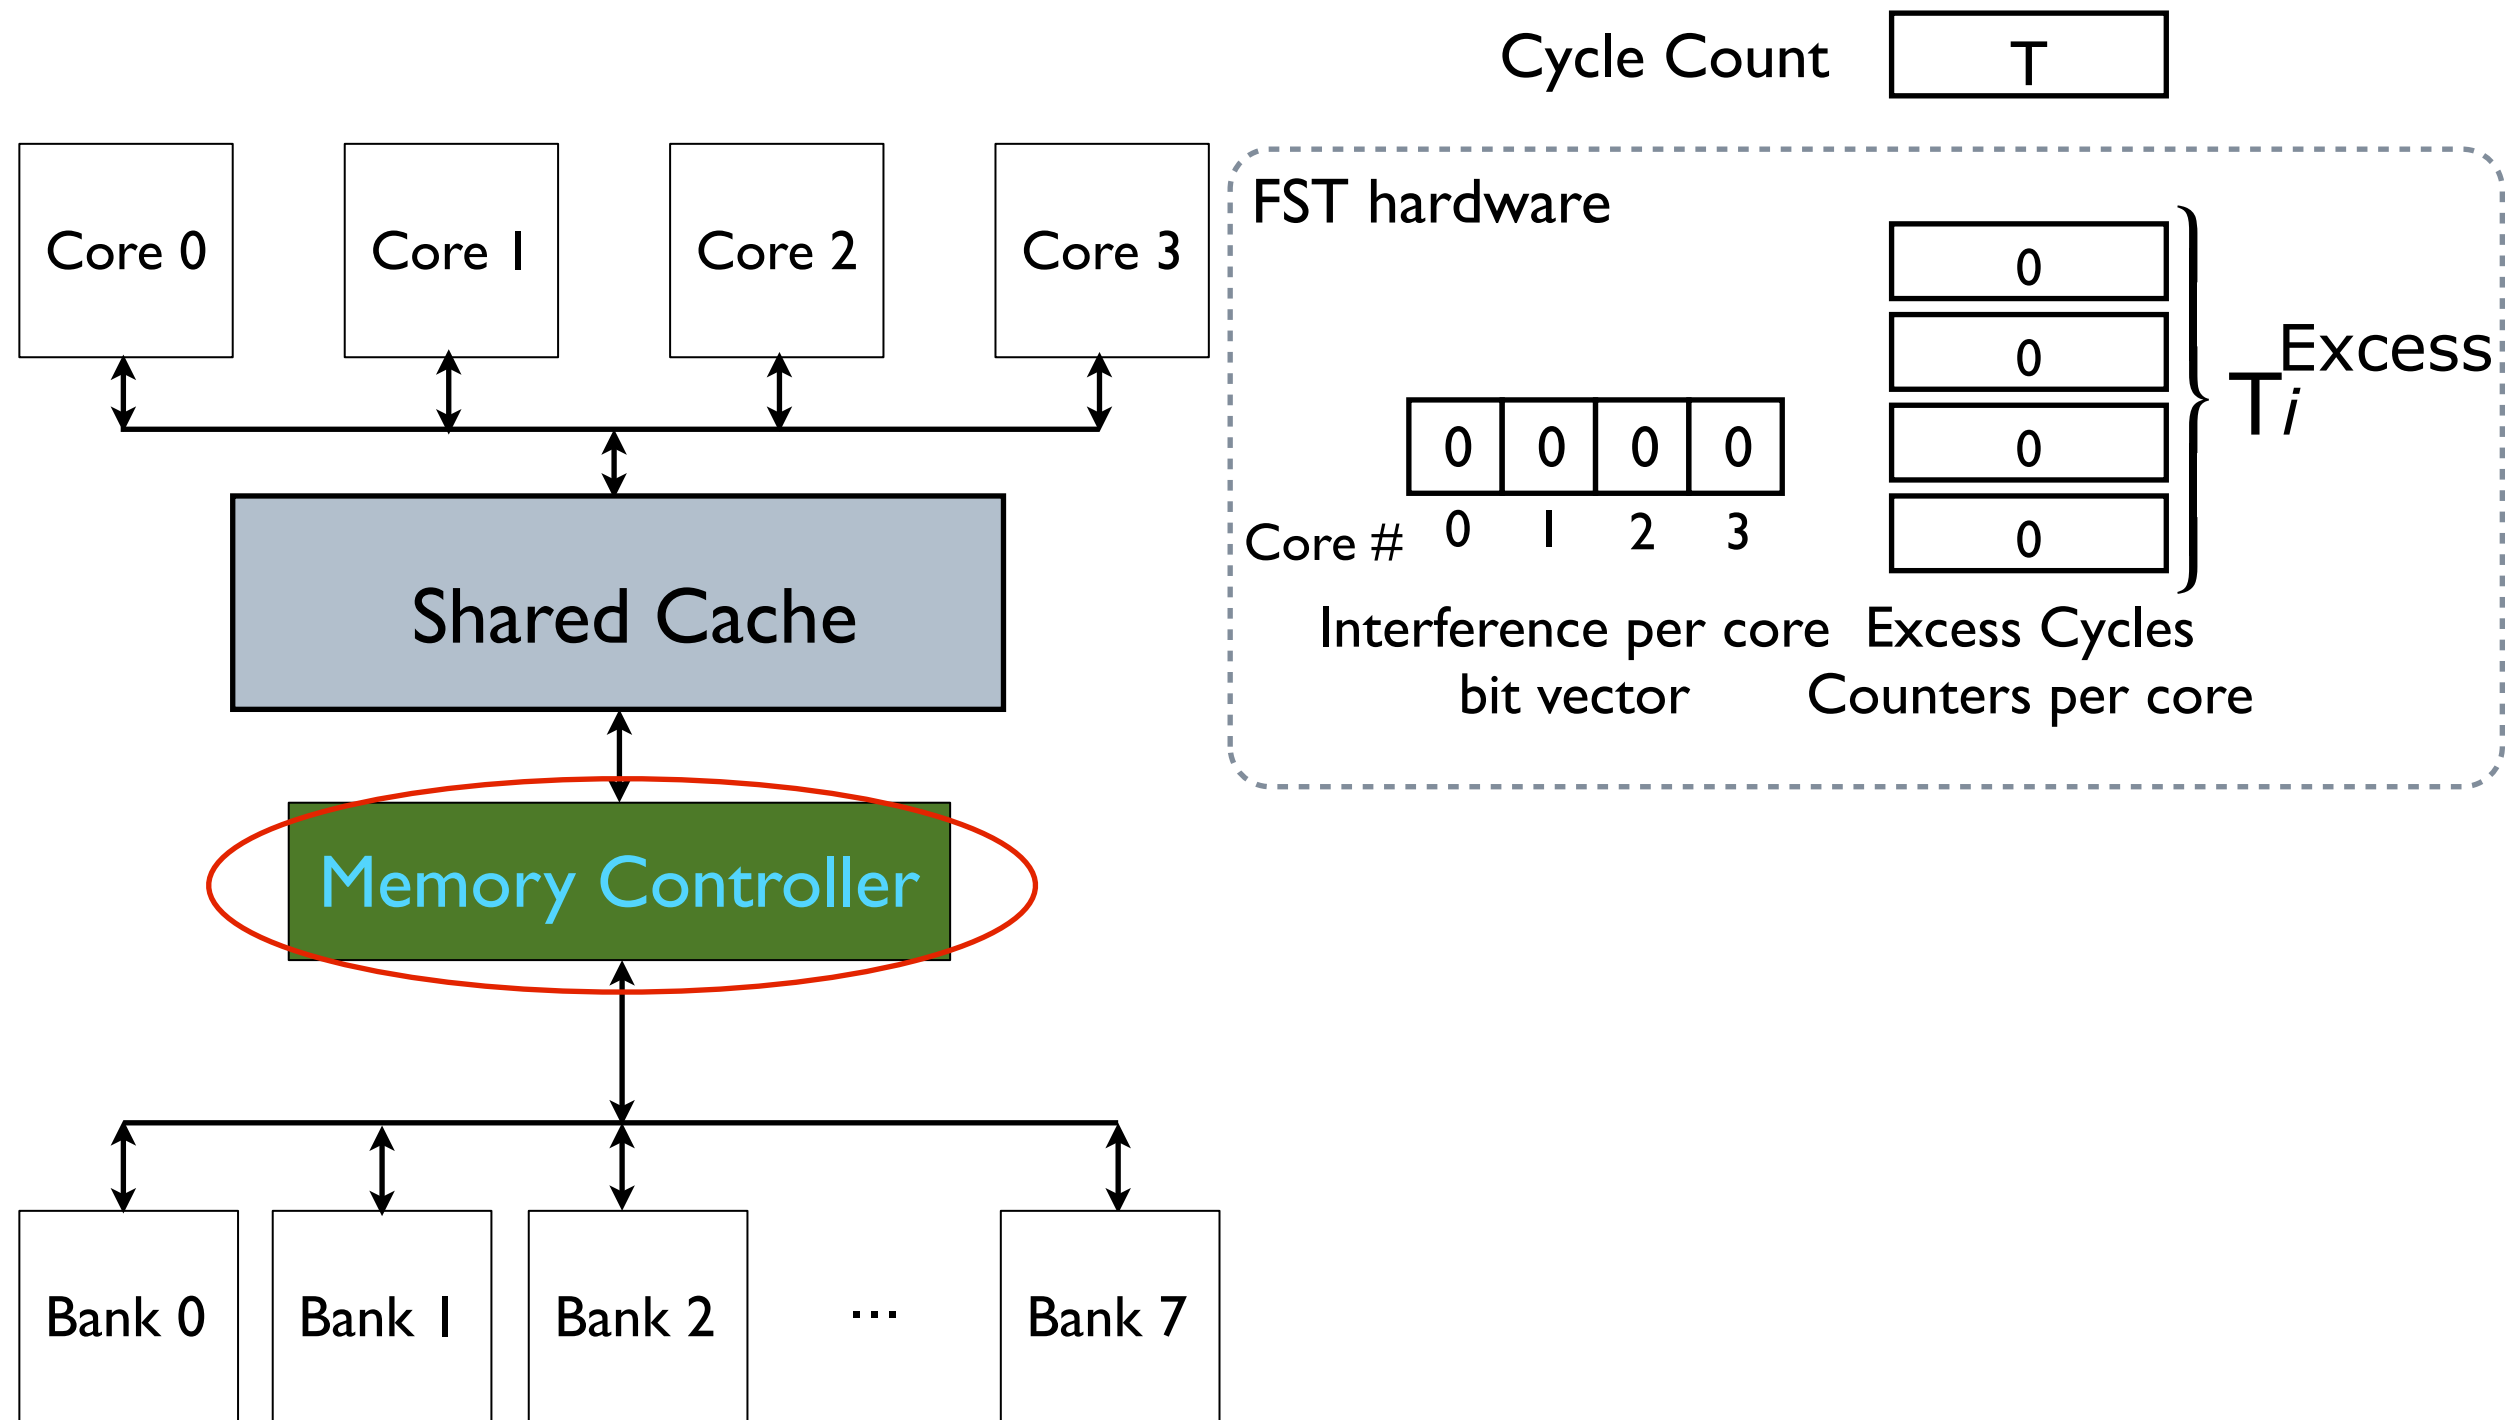
Tracking Inter-Core Interference



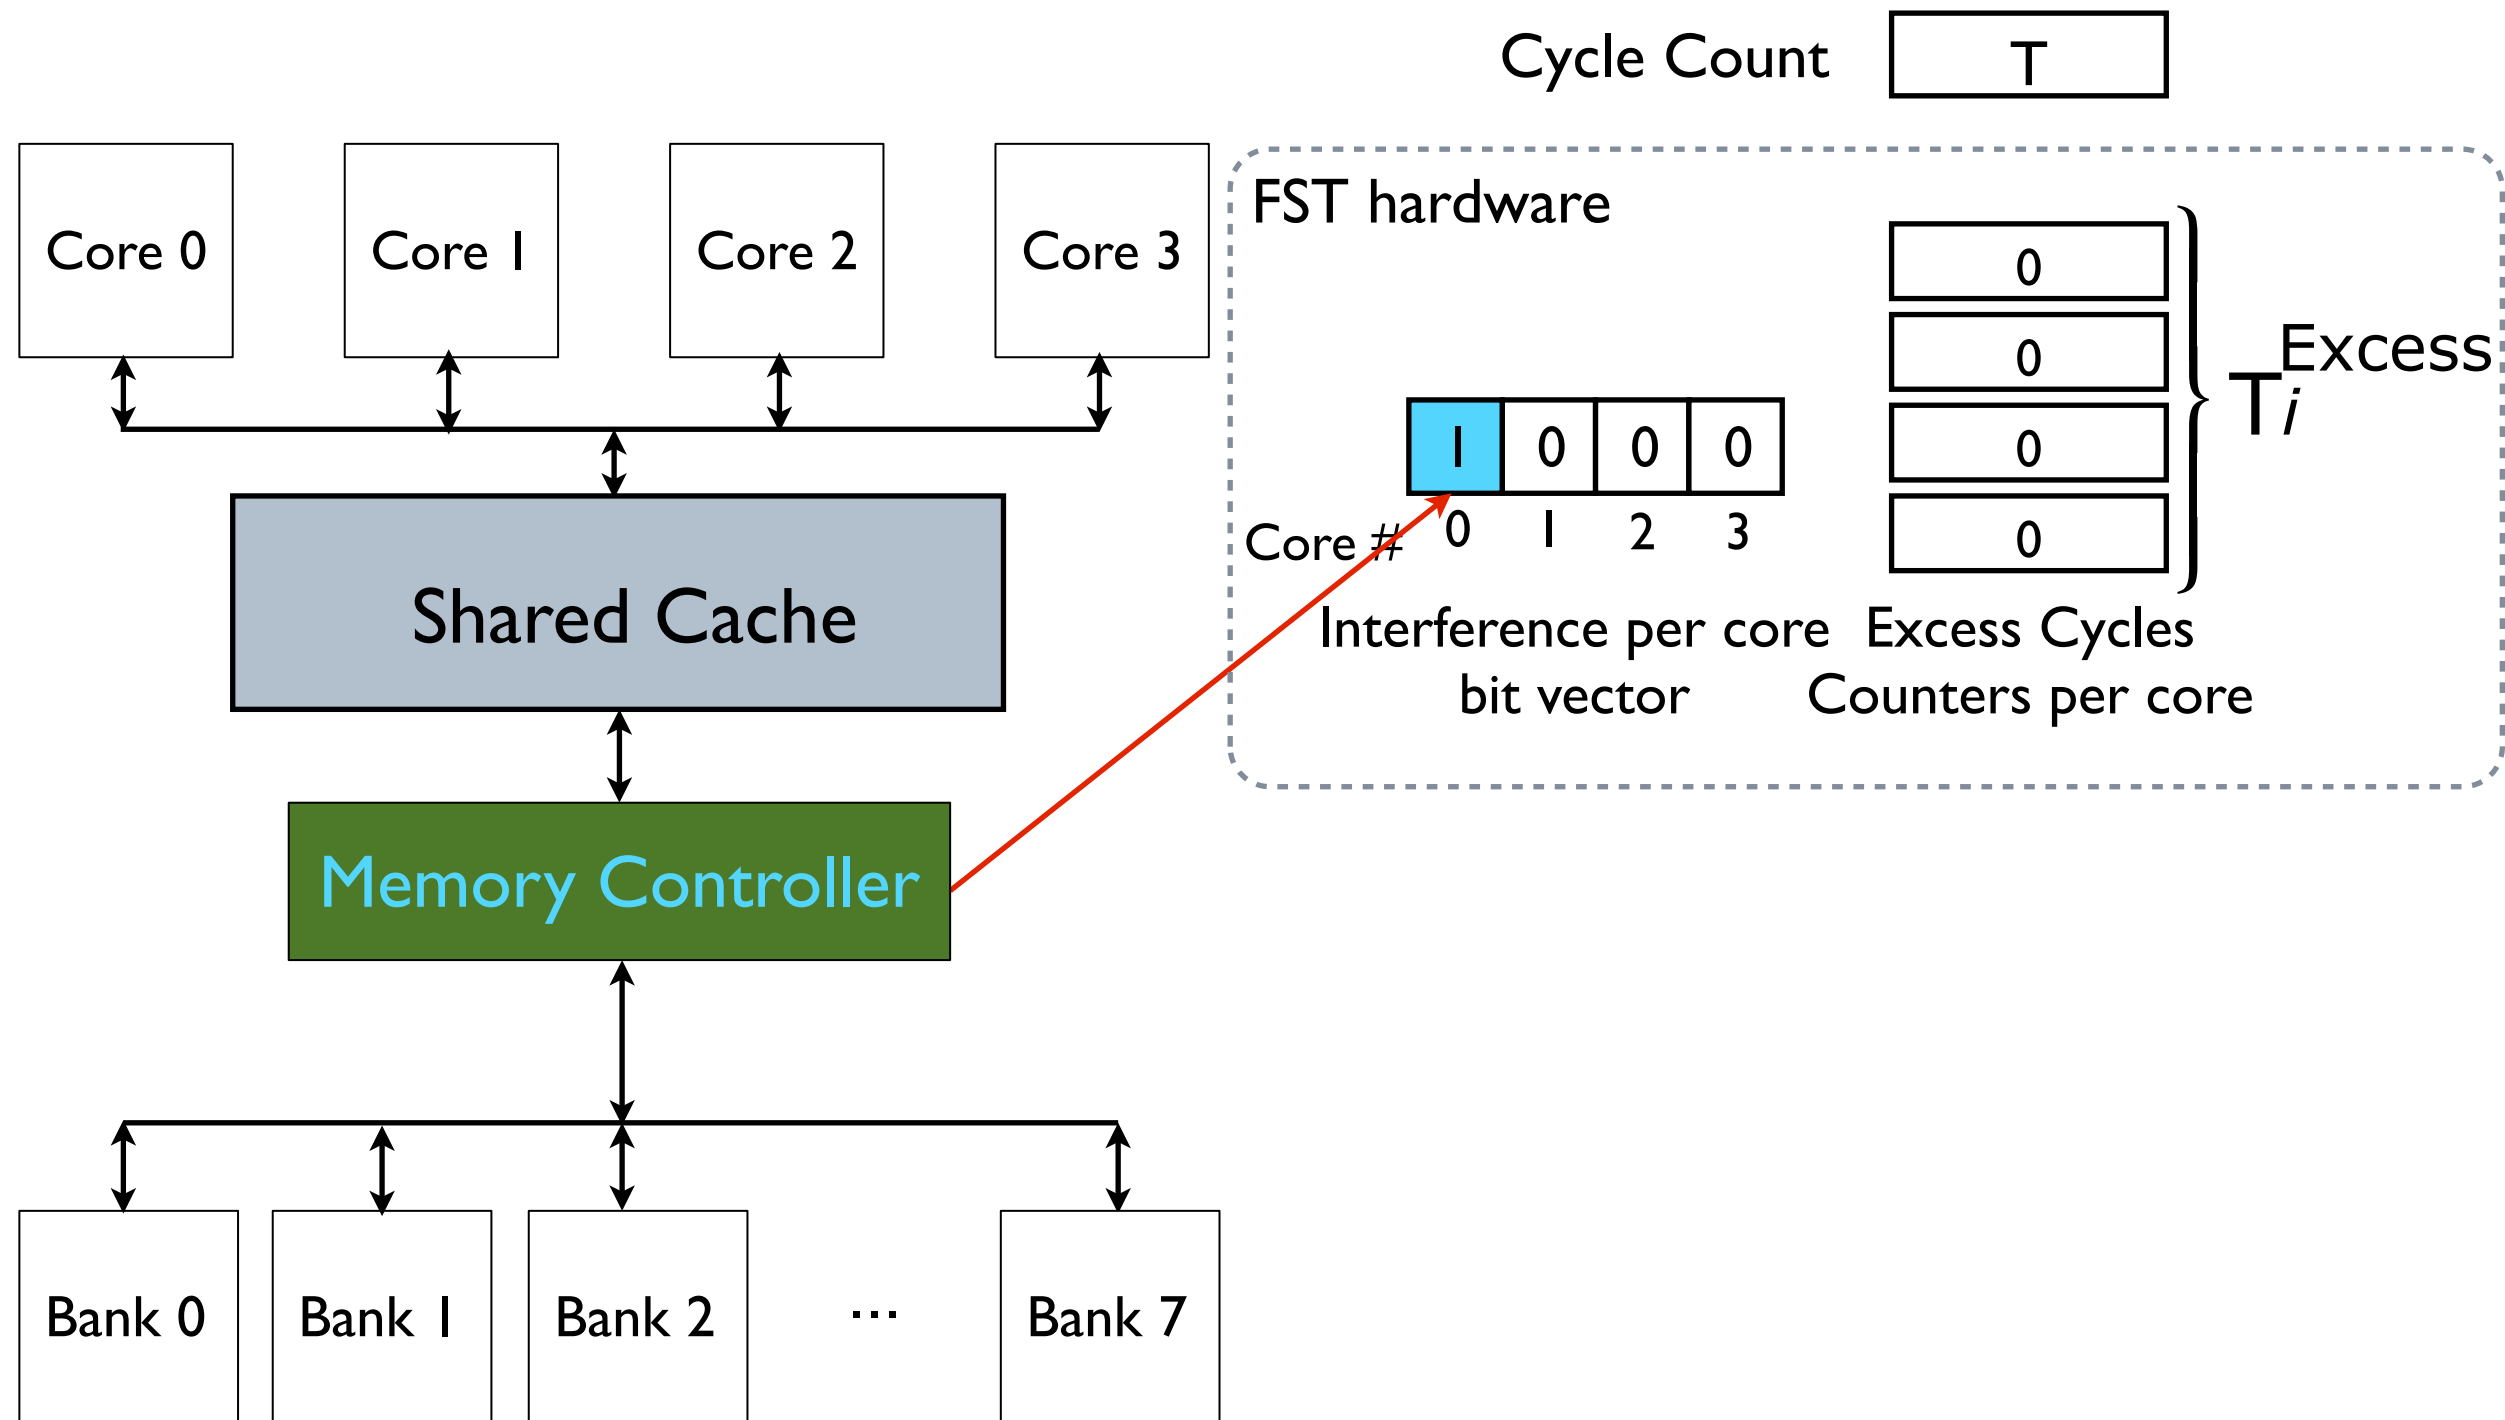
Tracking Inter-Core Interference



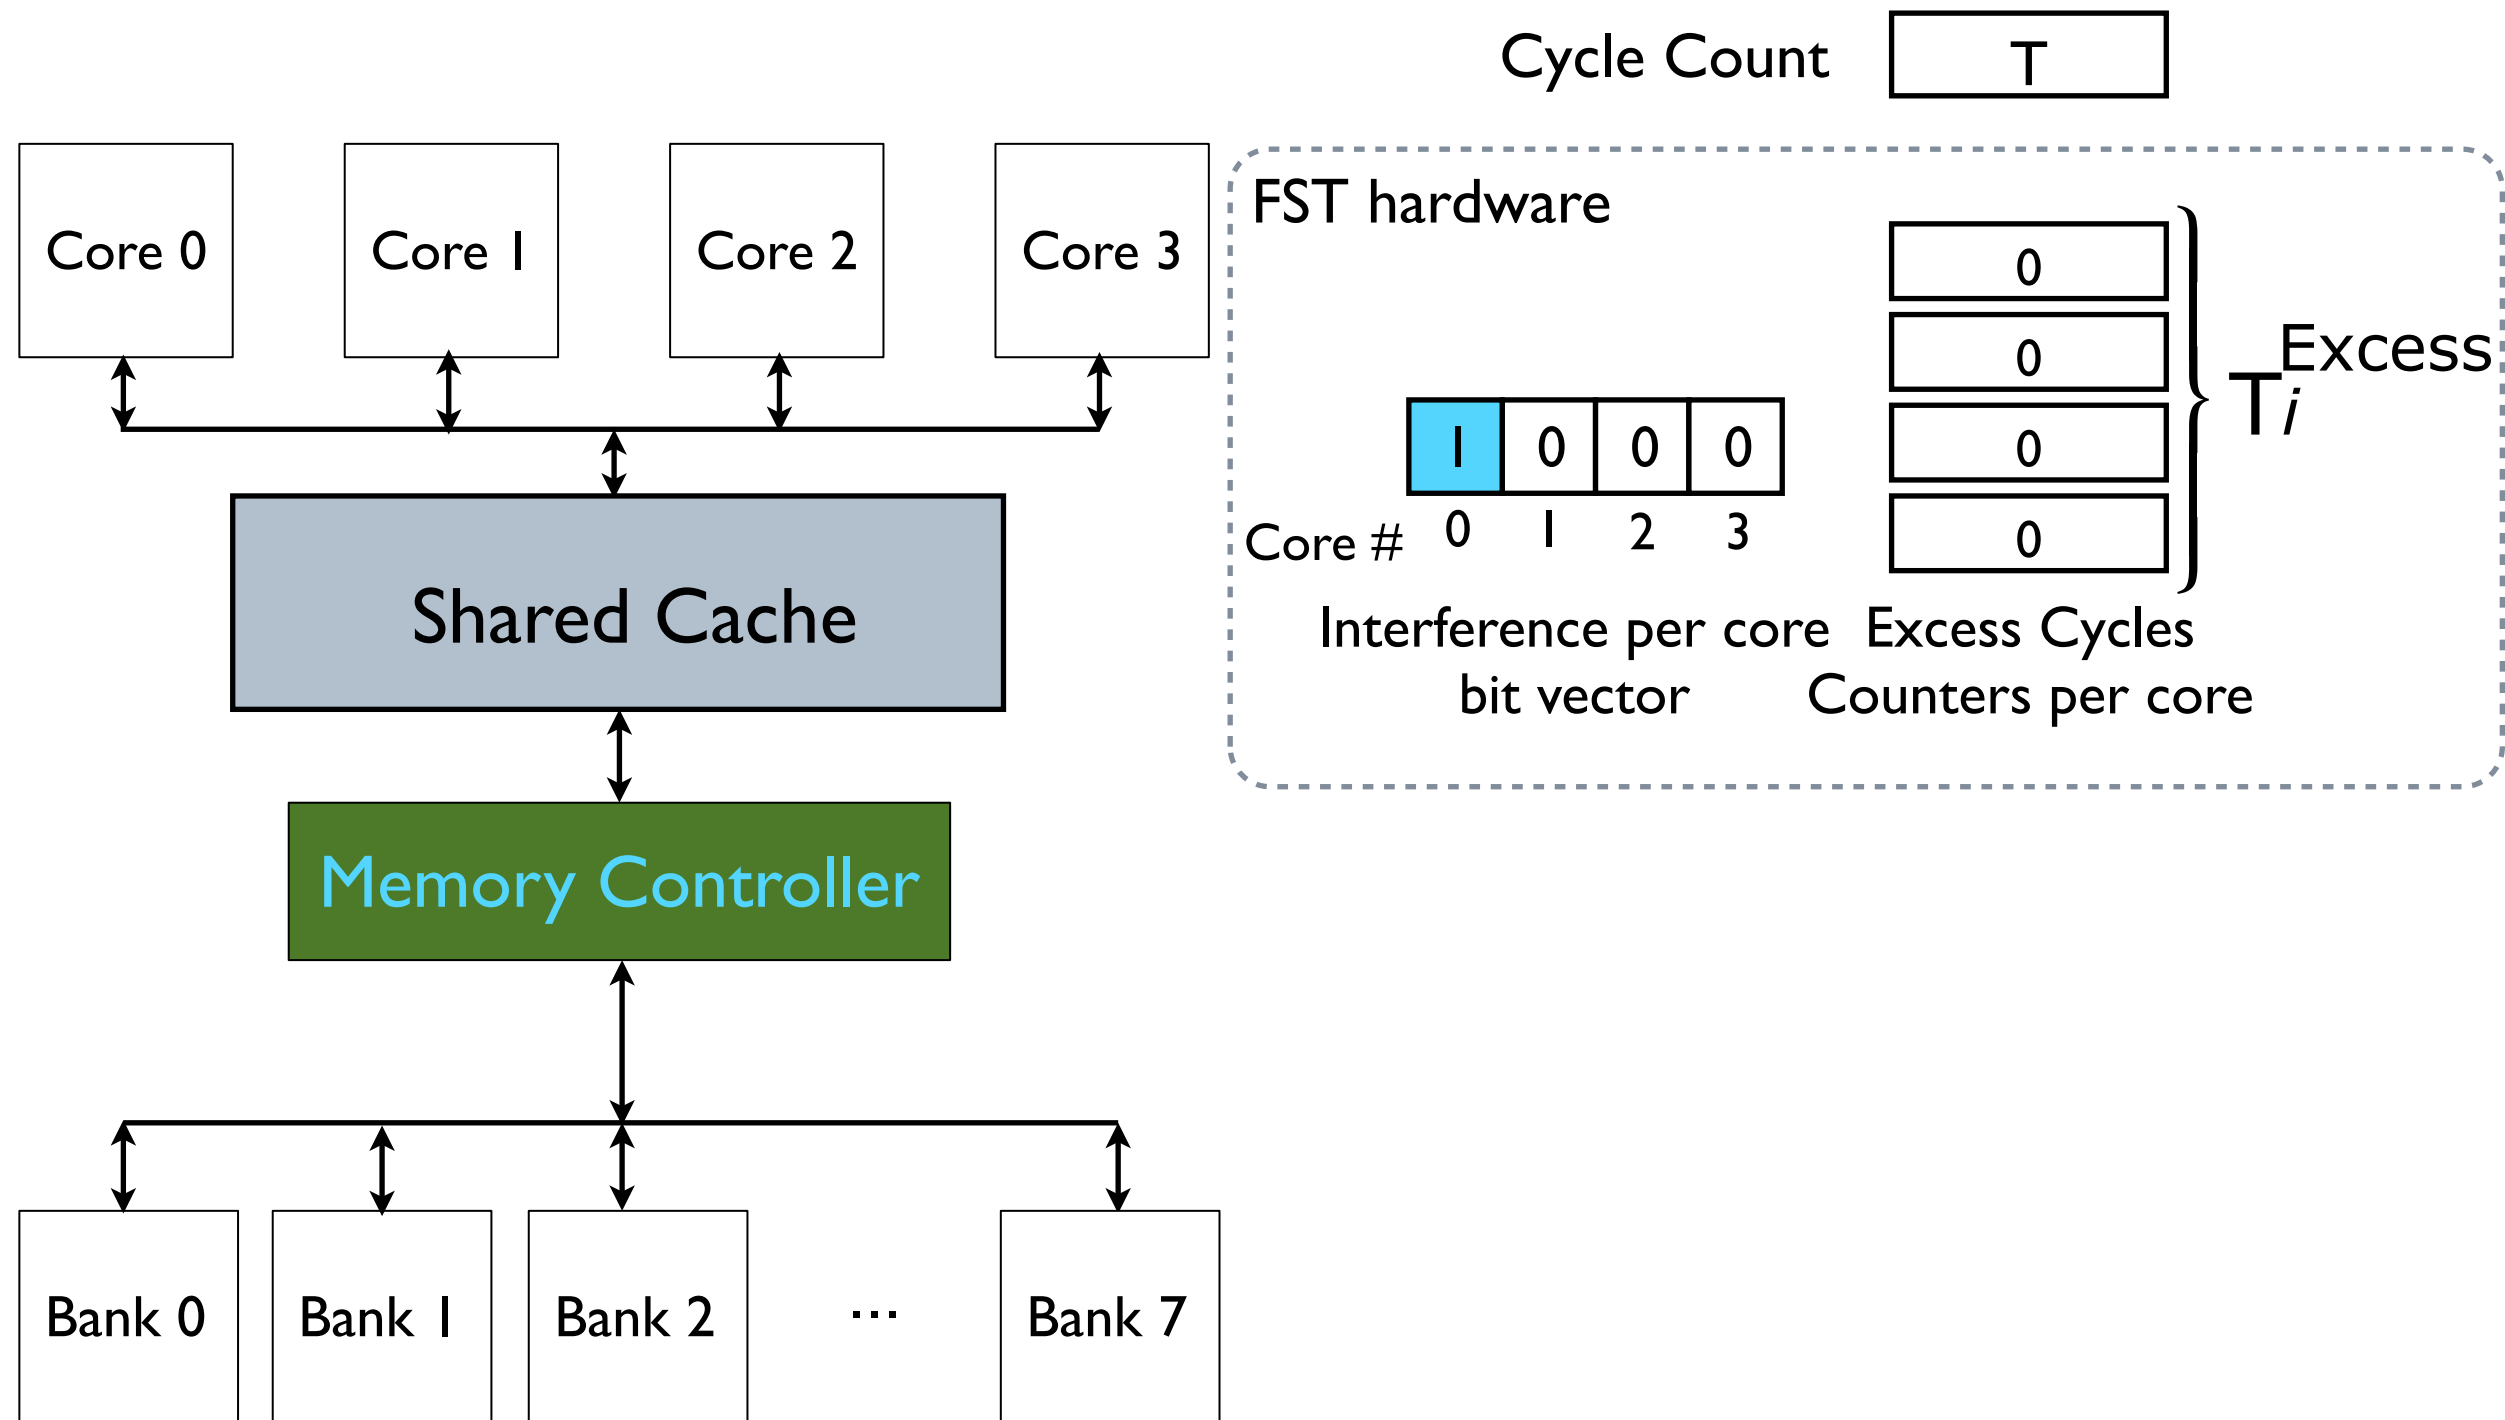
Tracking Inter-Core Interference



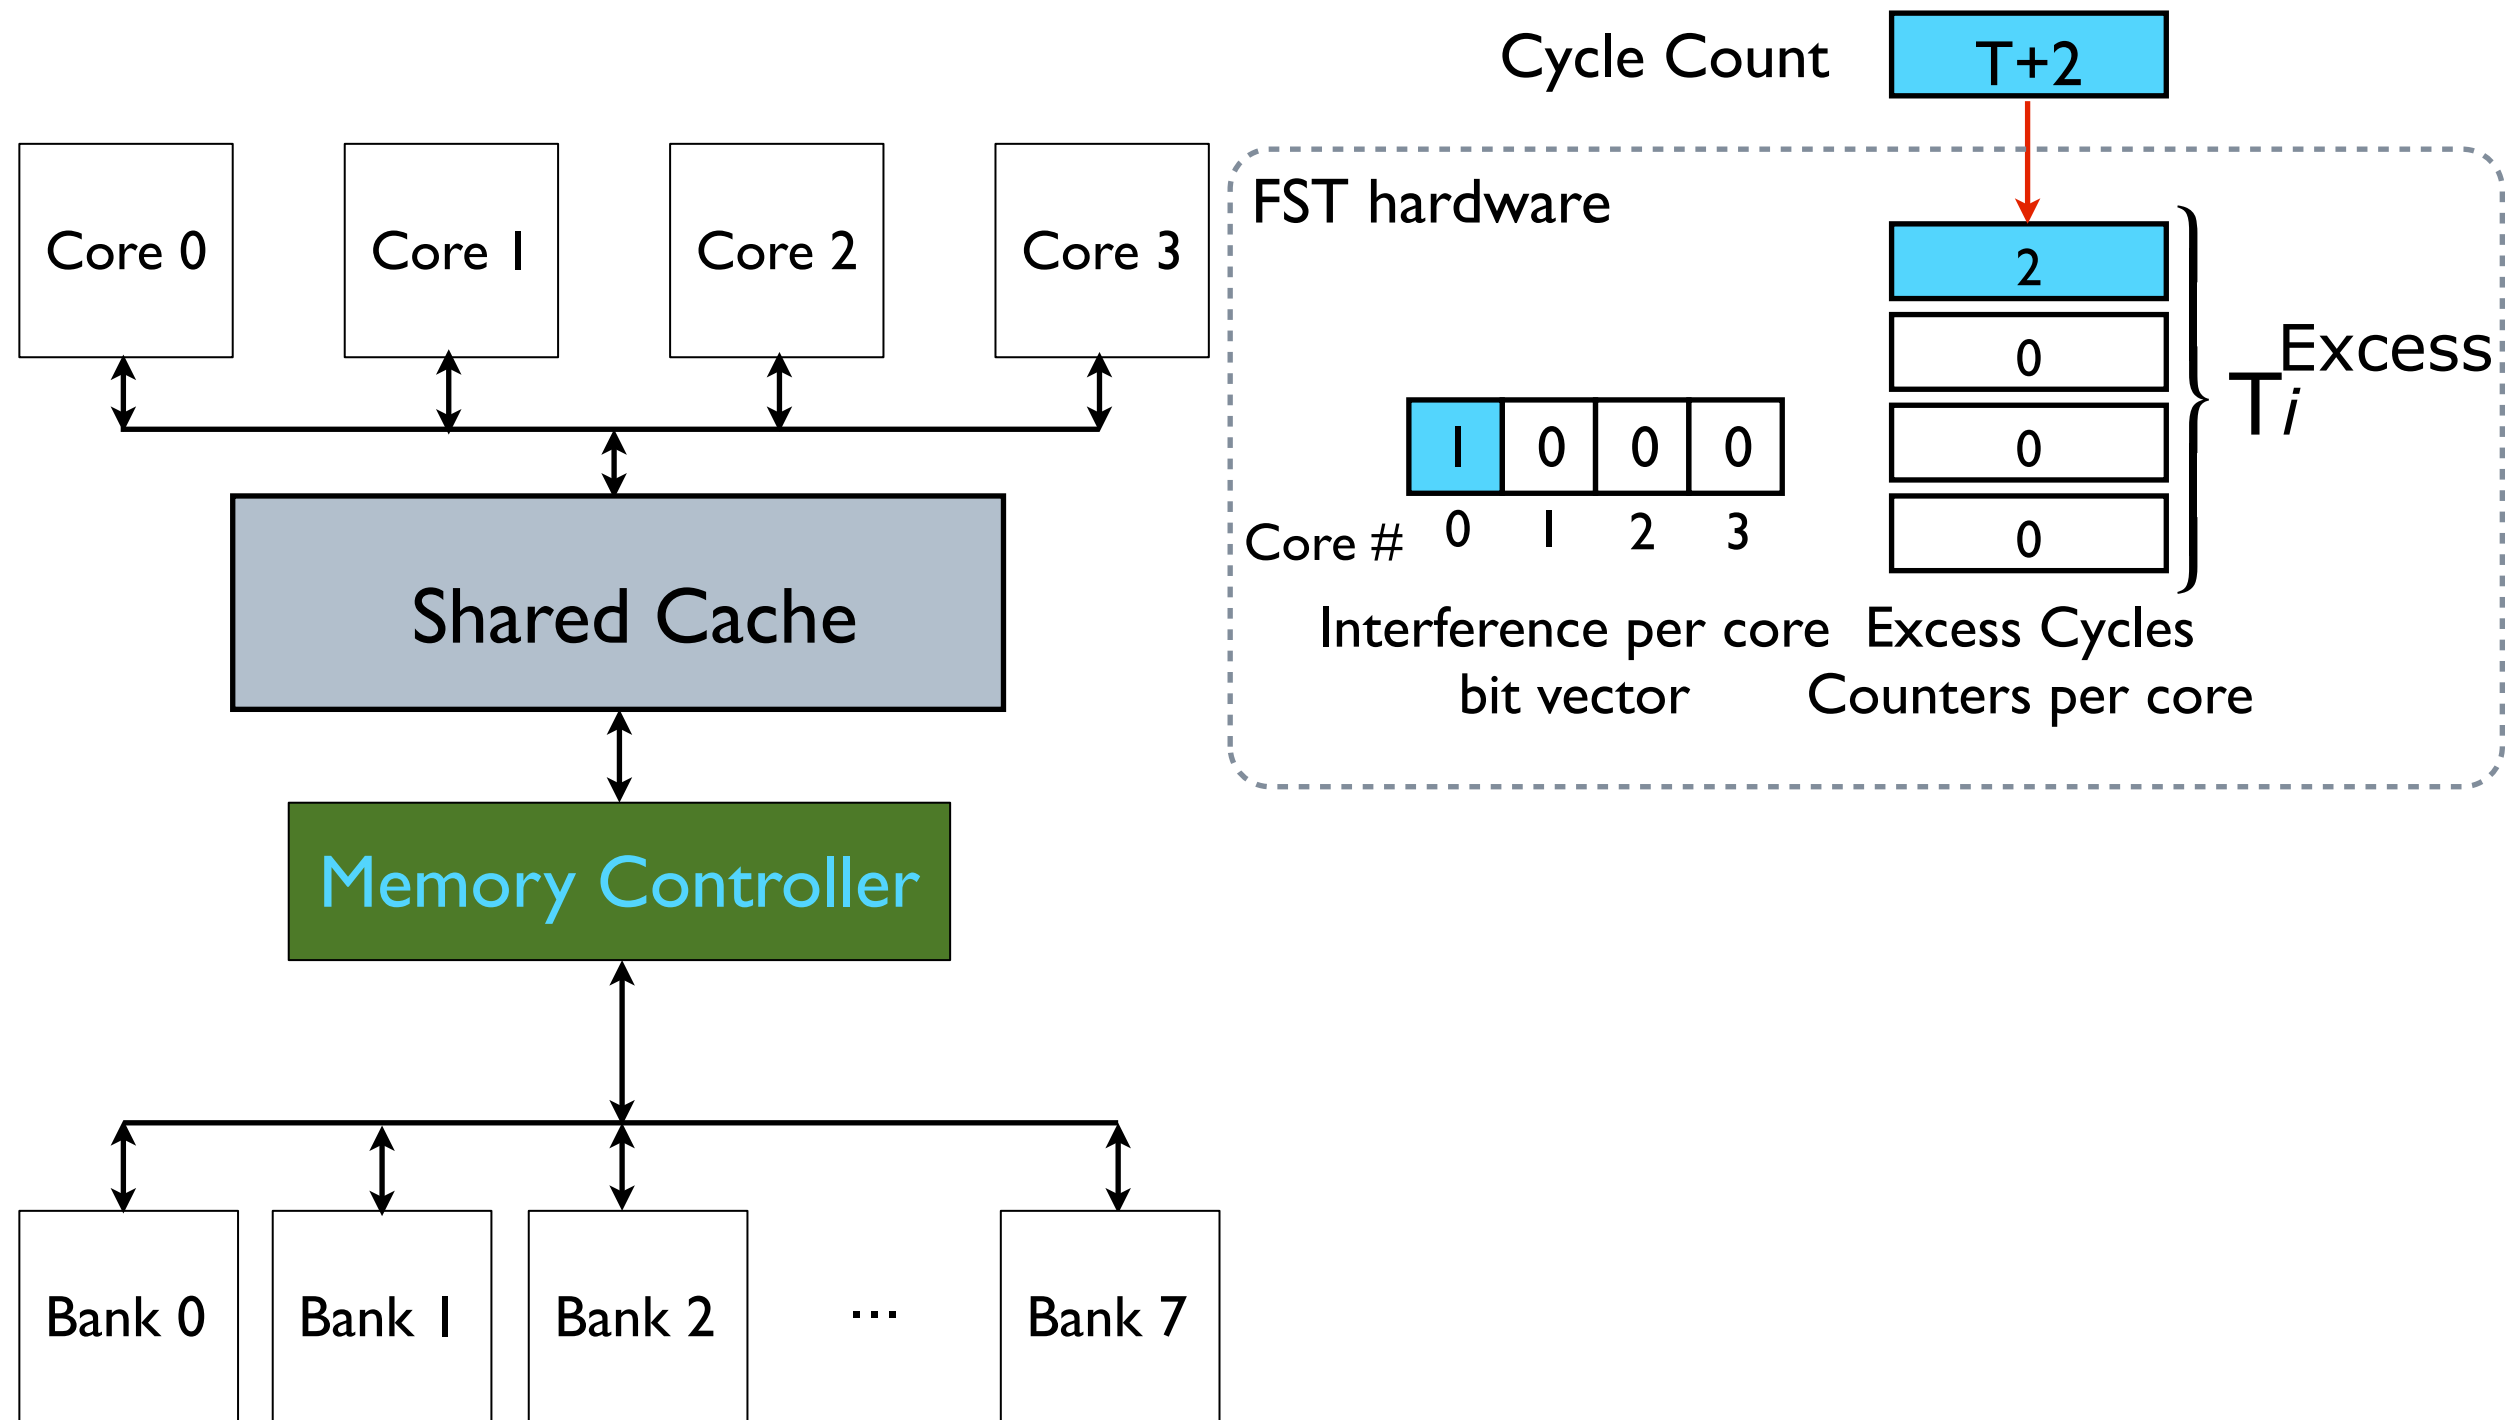
Tracking Inter-Core Interference



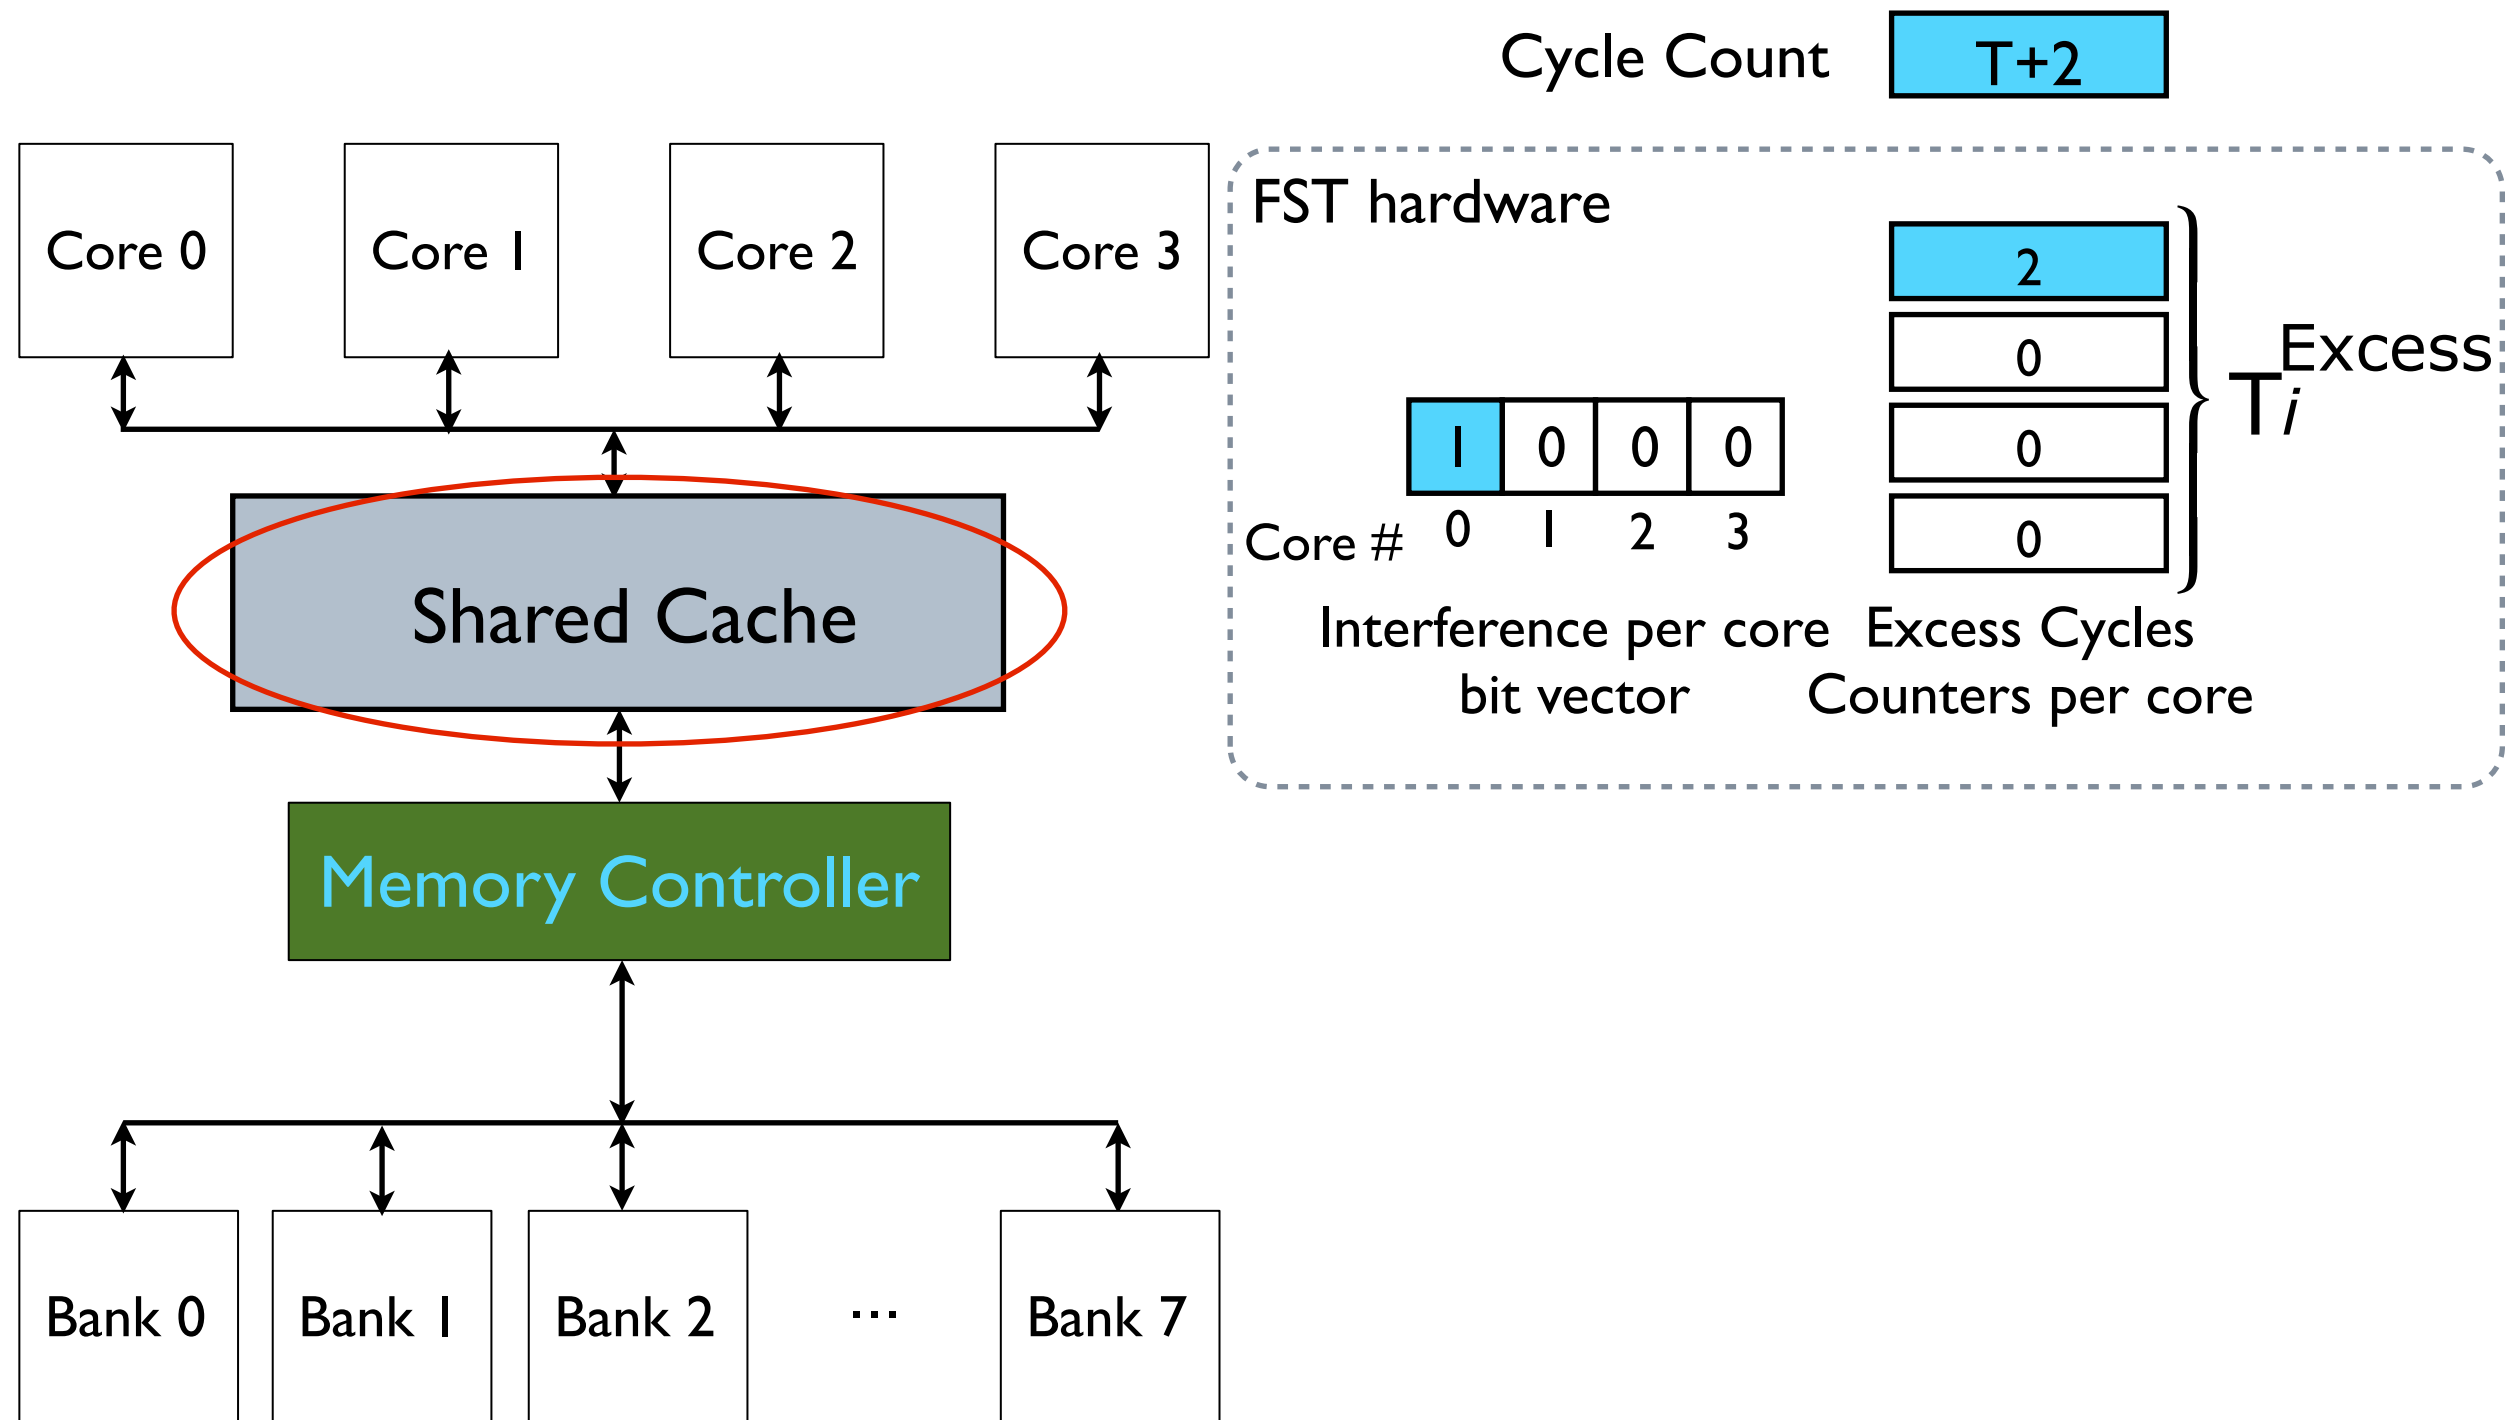
Tracking Inter-Core Interference



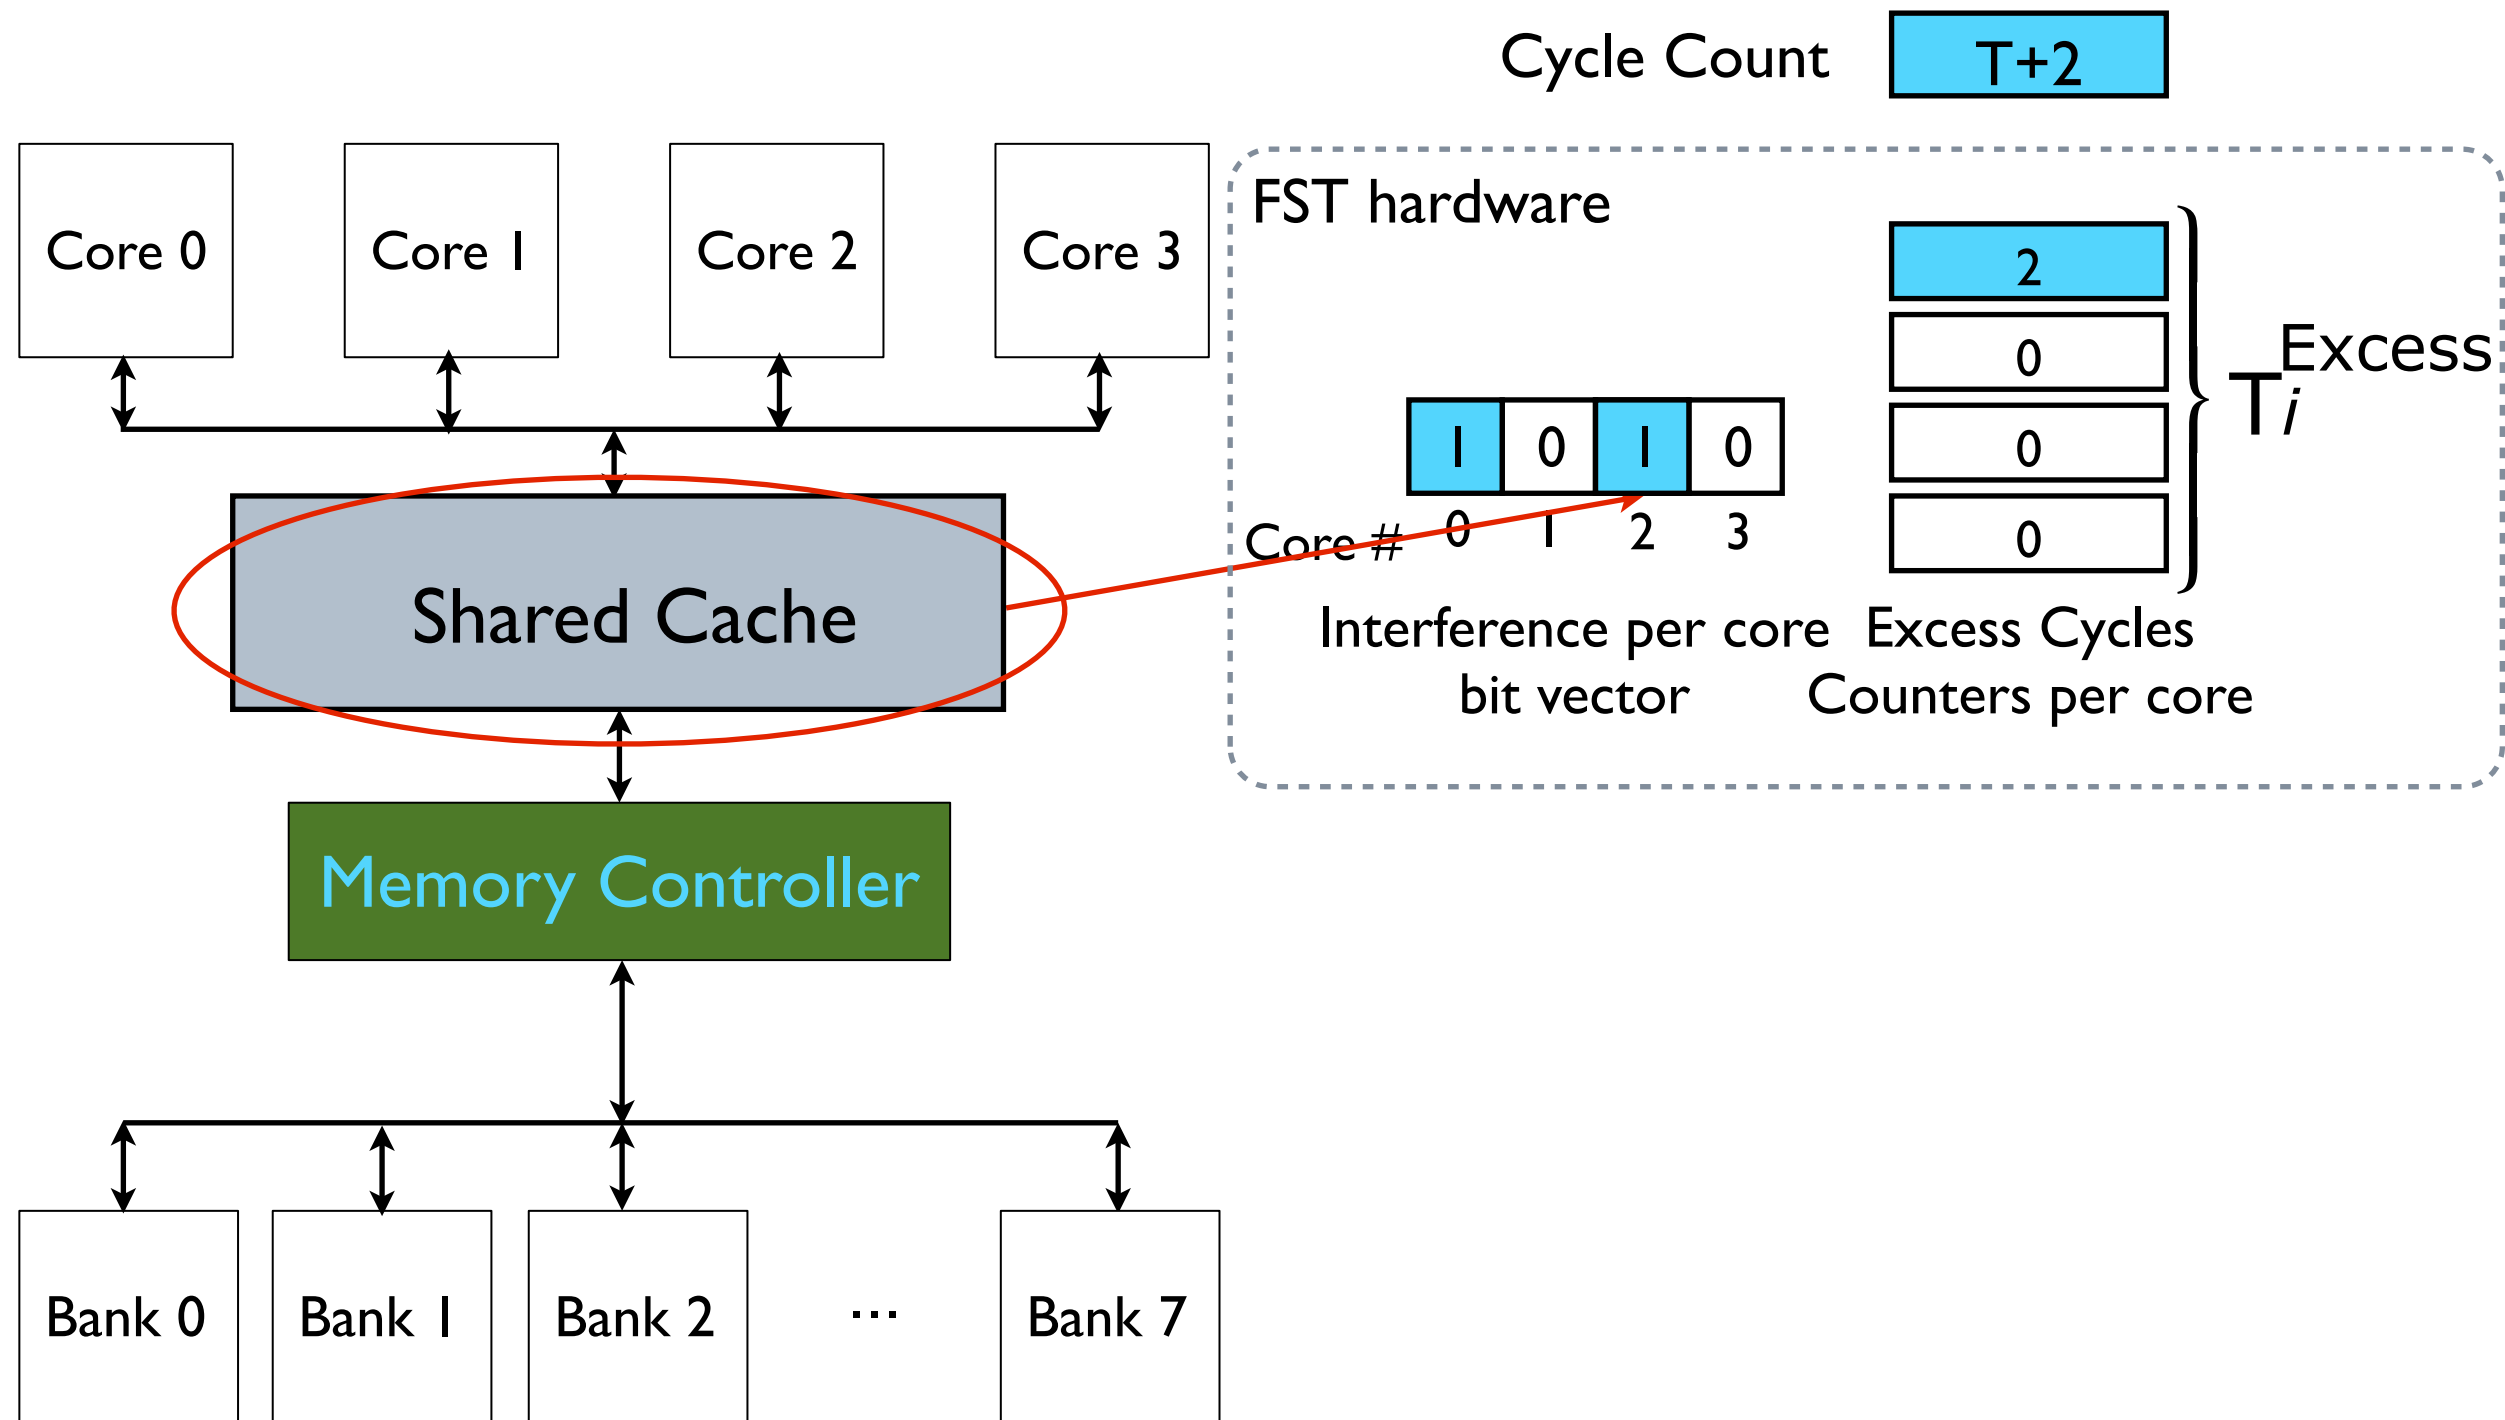
Tracking Inter-Core Interference



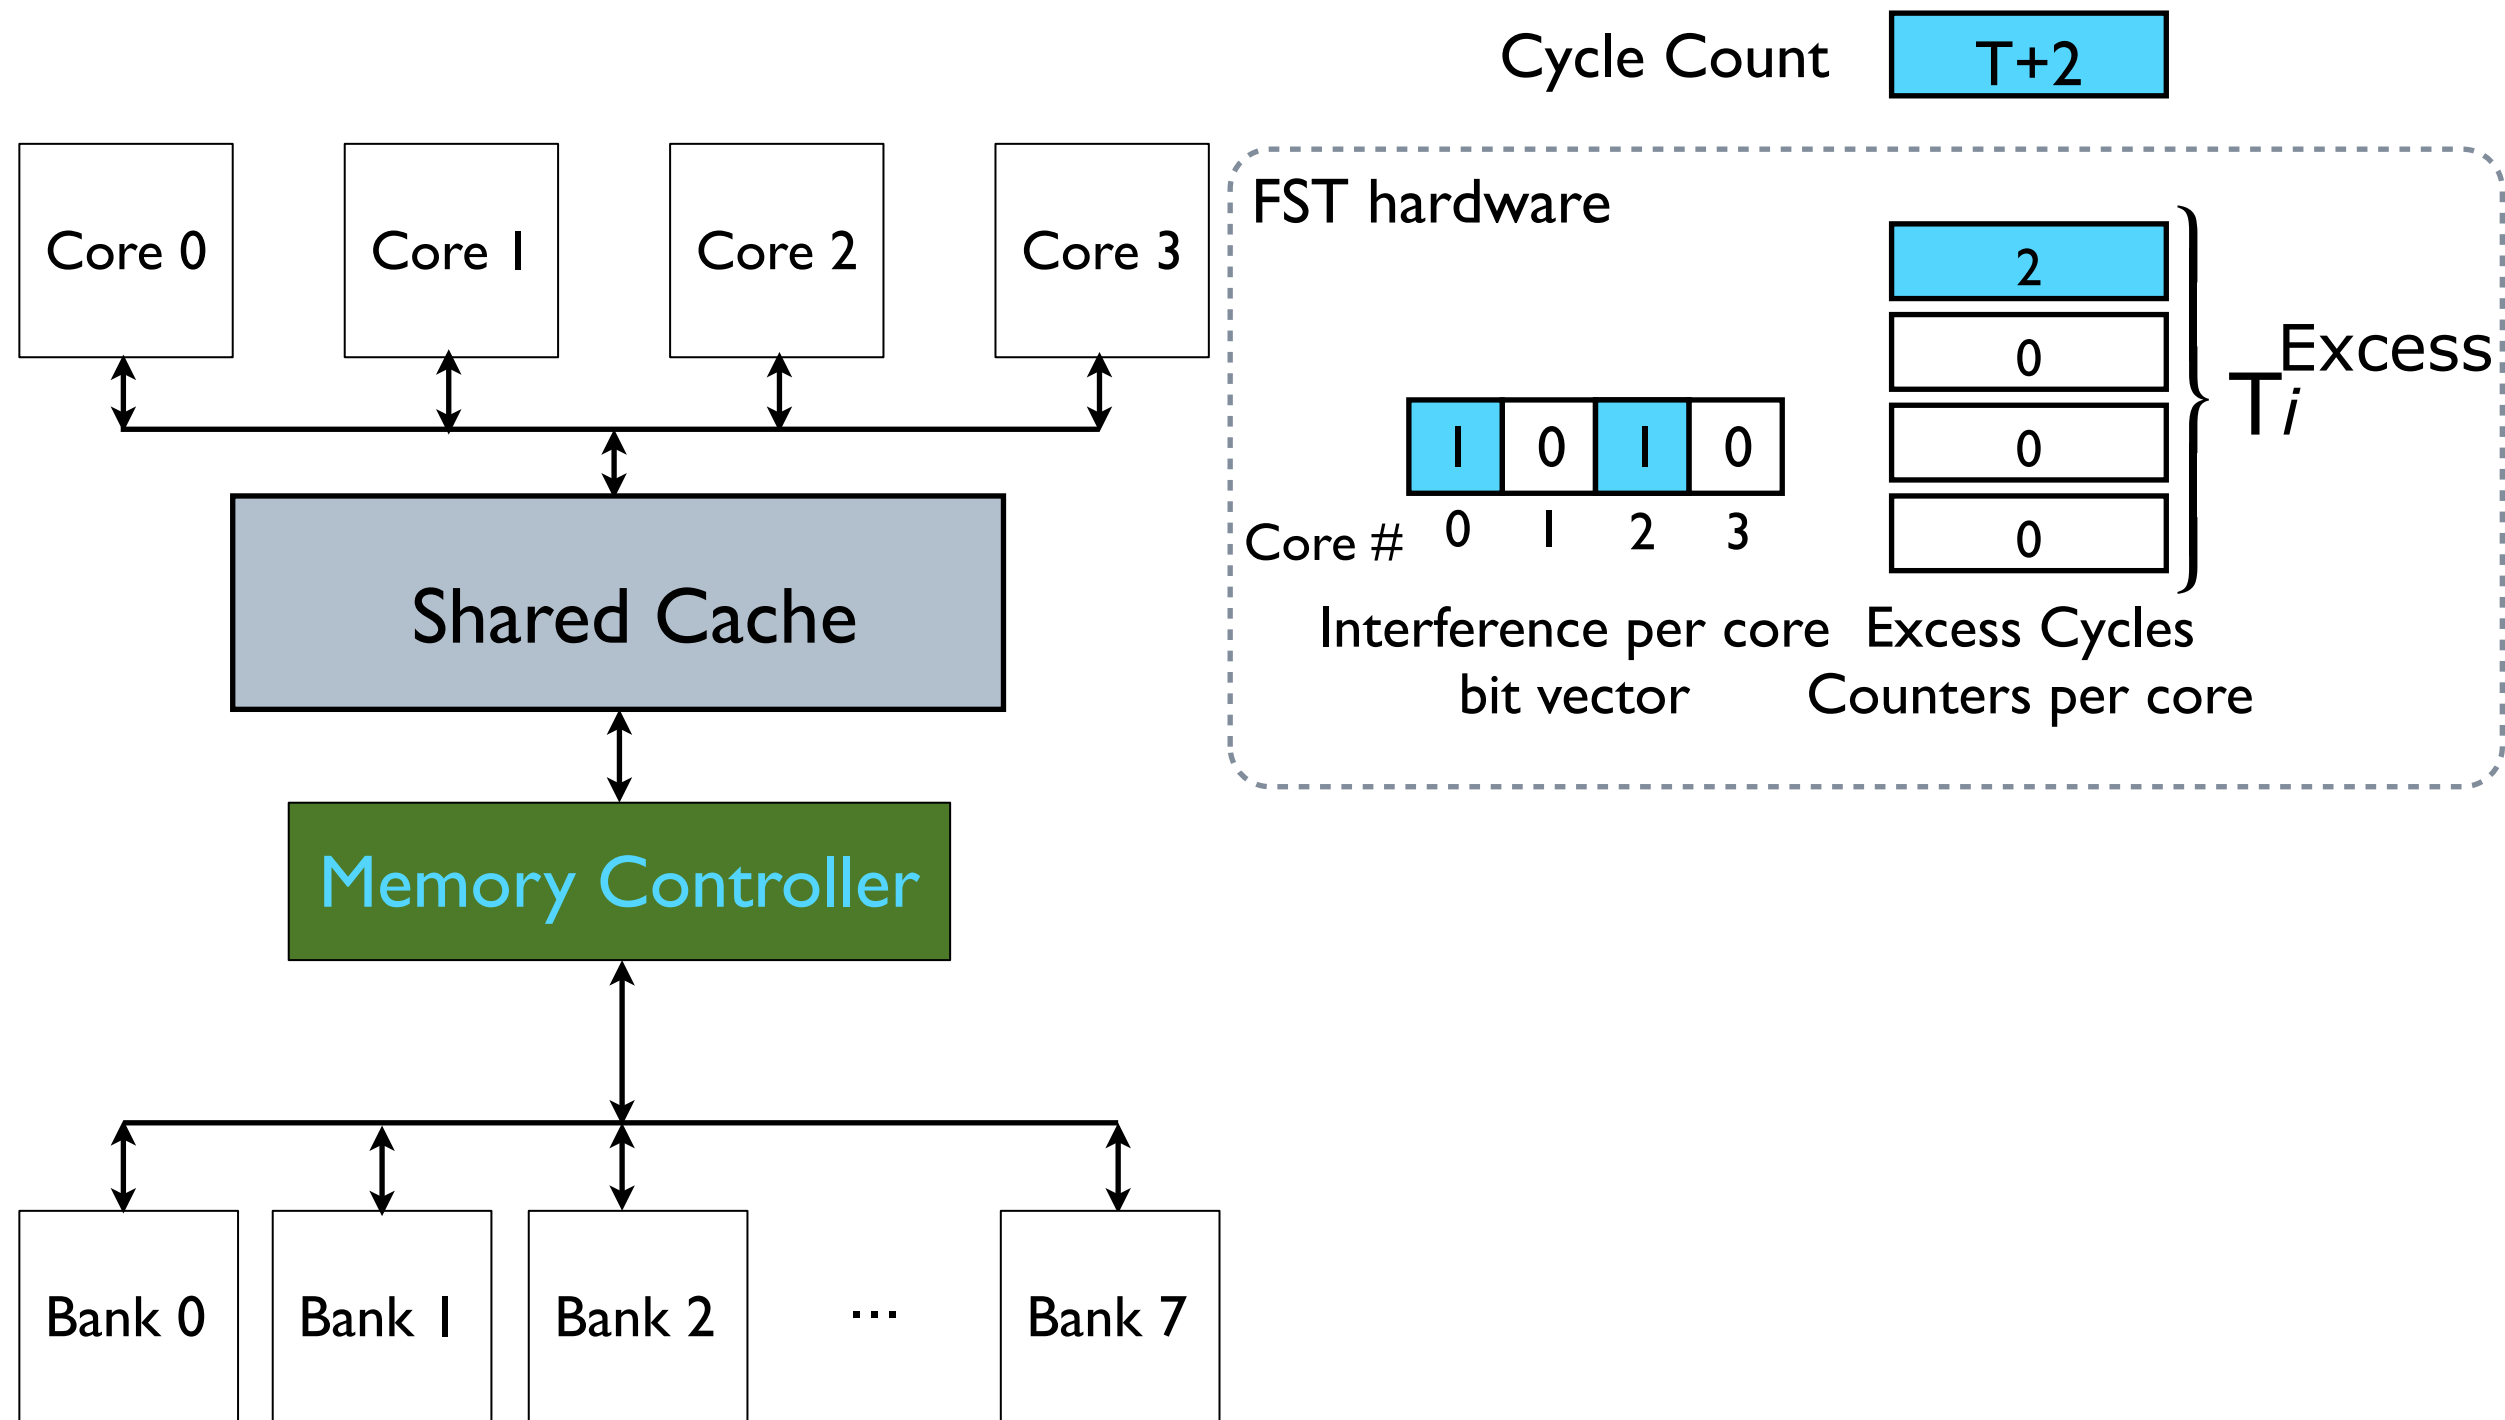
Tracking Inter-Core Interference



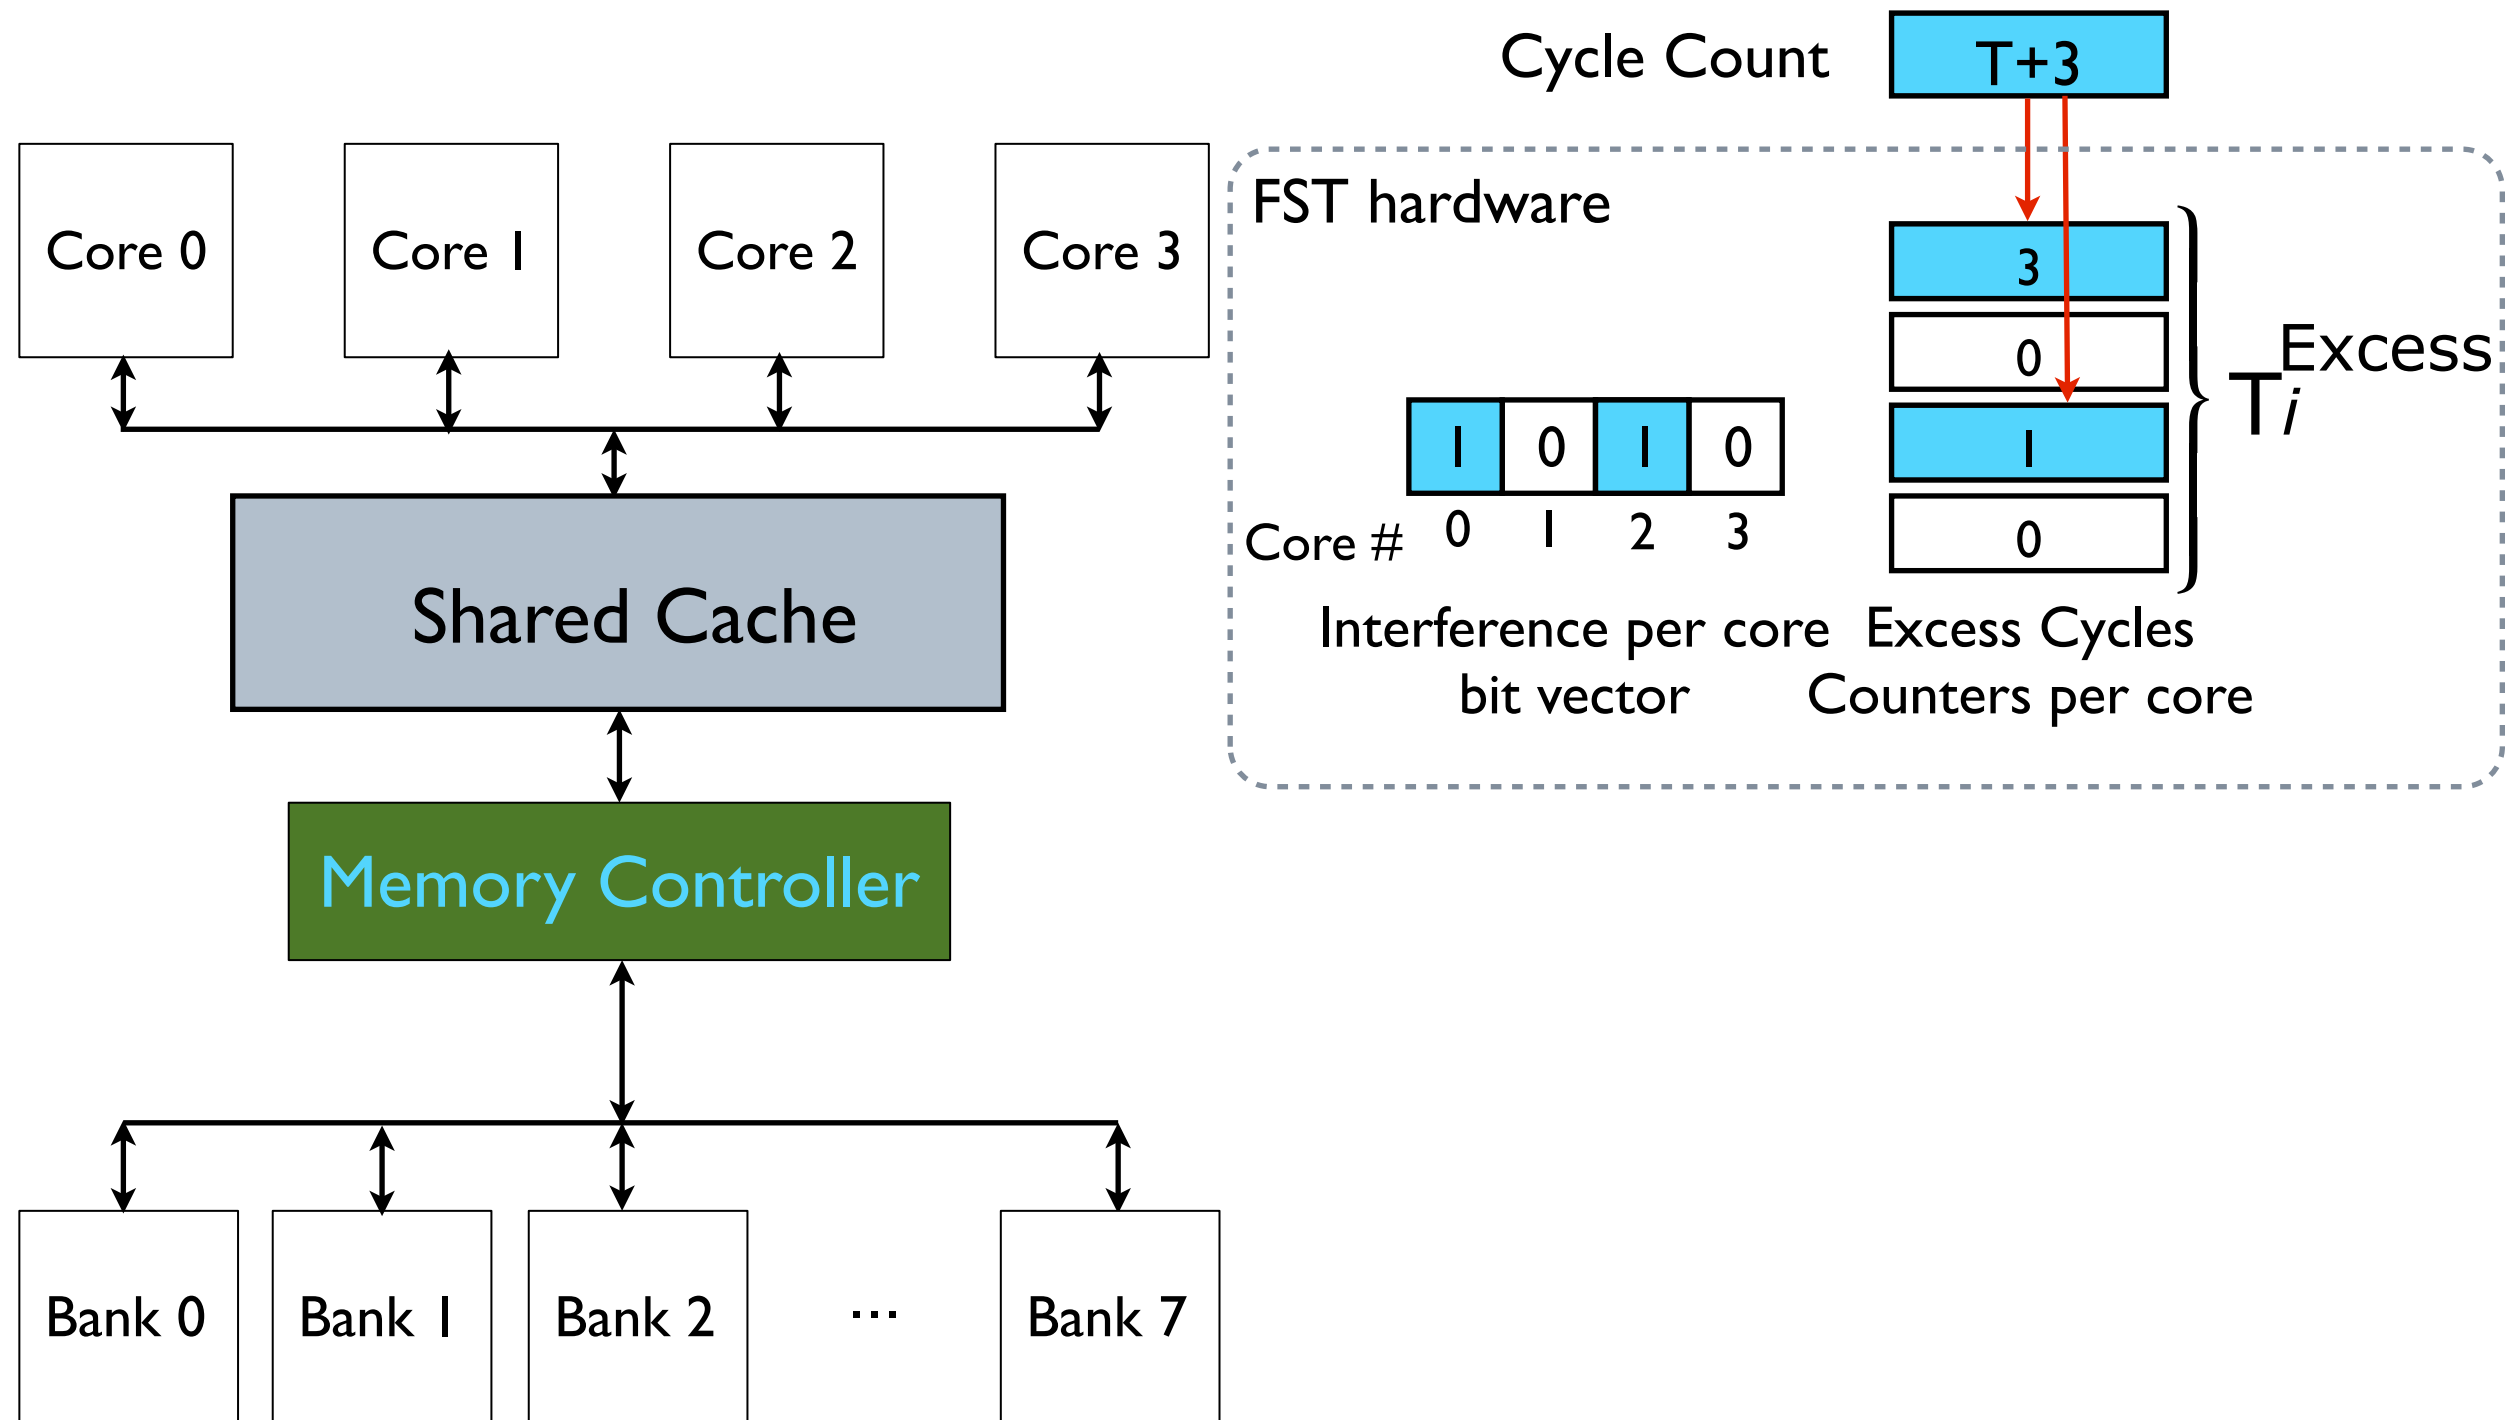
Tracking Inter-Core Interference



Tracking Inter-Core Interference

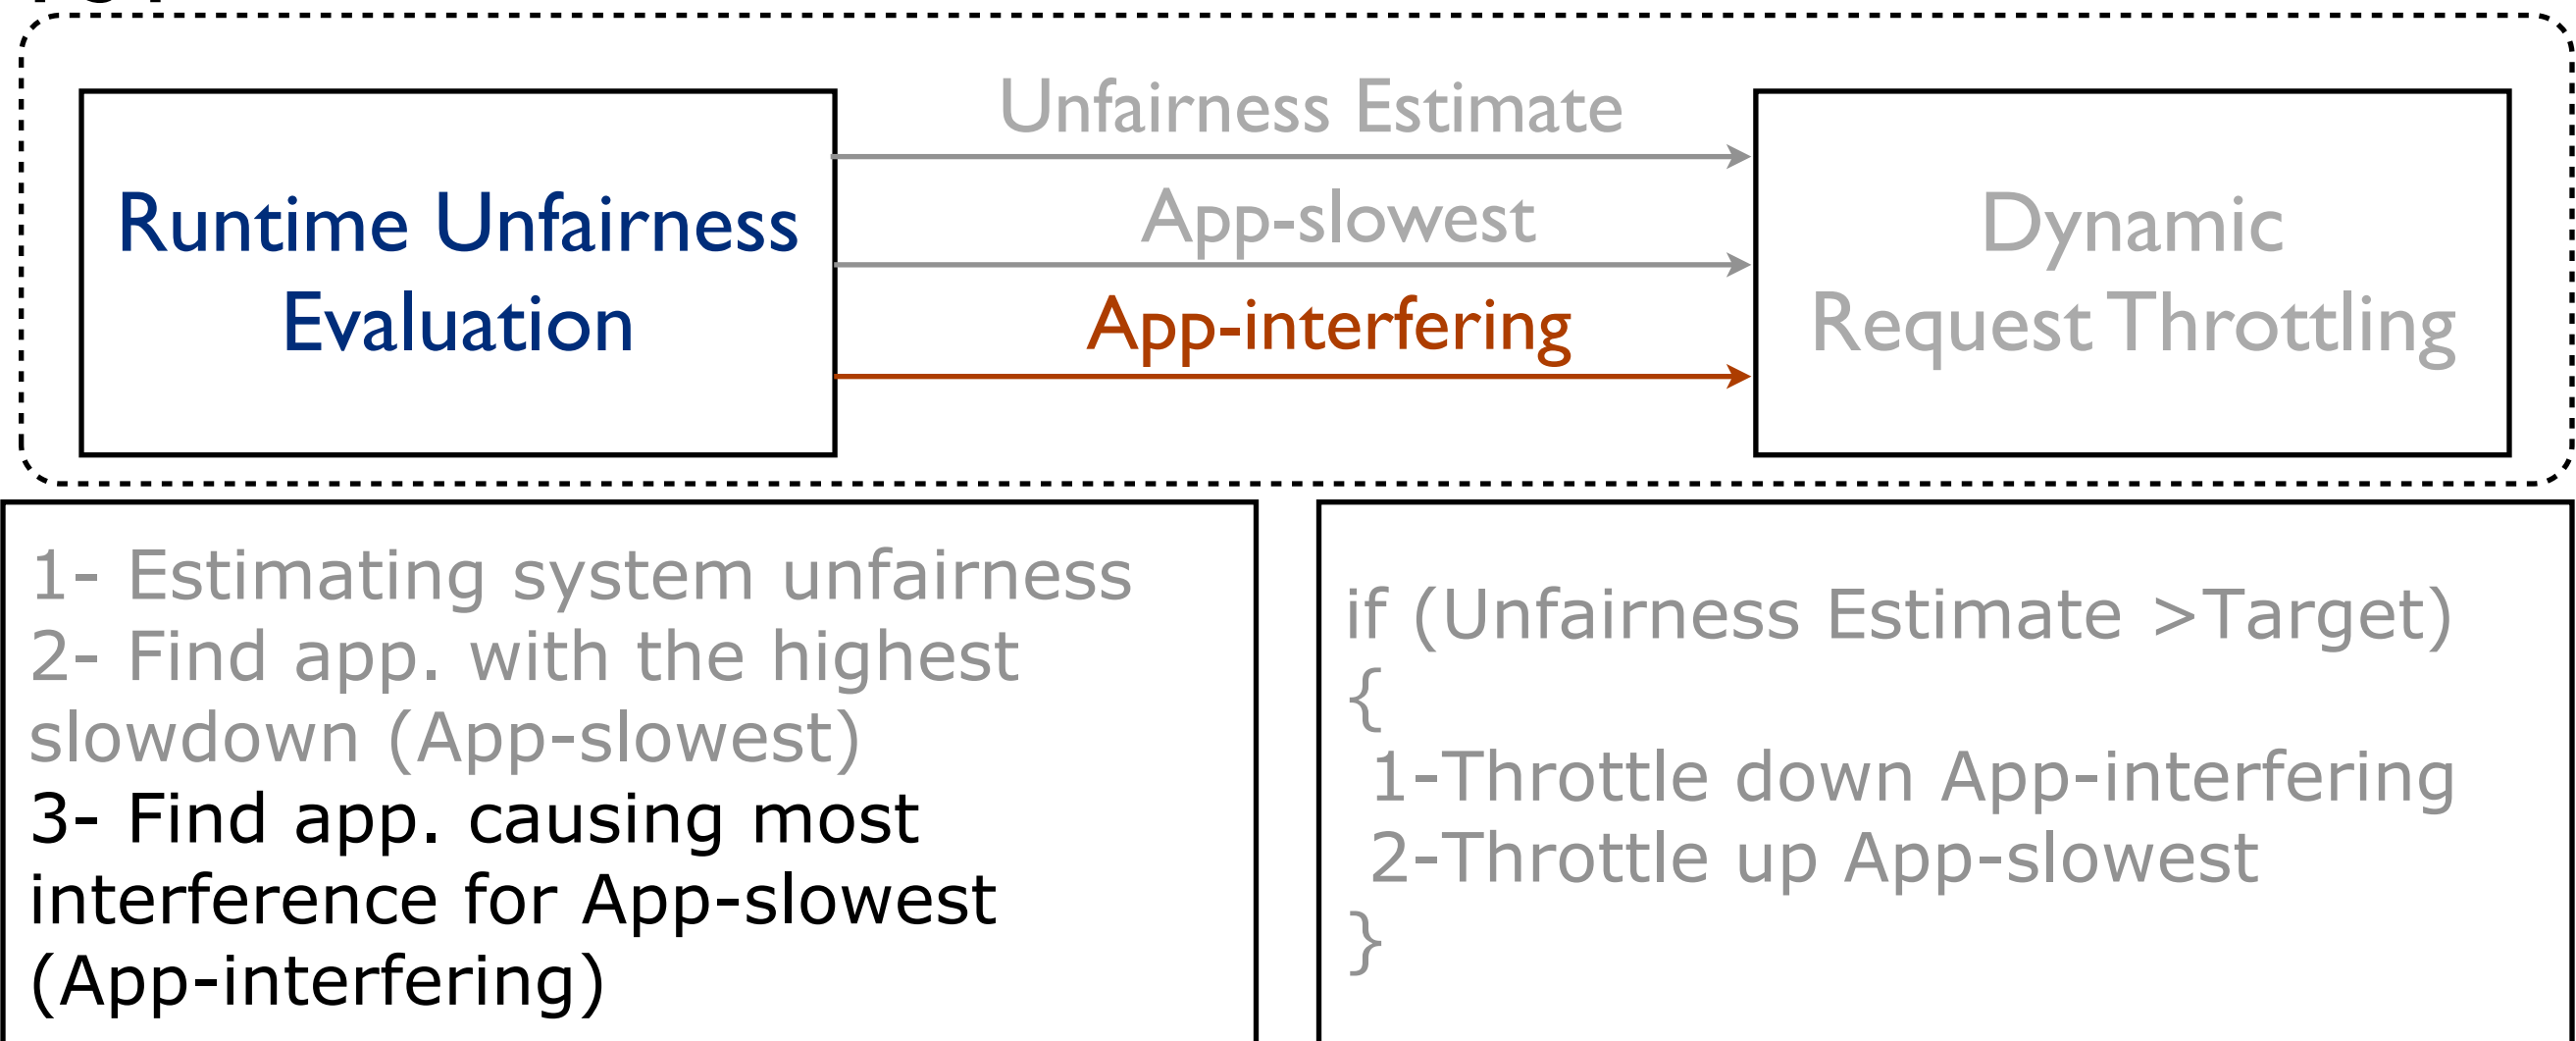


Tracking Inter-Core Interference



Fairness via Source Throttling (FST)

FST



Tracking Inter-Core Interference

Tracking Inter-Core Interference

- To identify App-interfering, for each core i

Tracking Inter-Core Interference

- To identify App-interfering, for each core i
 - FST separately tracks interference caused by each core j ($j \neq i$)

Tracking Inter-Core Interference

- To identify App-interfering, for each core i
 - FST separately tracks interference caused by each core j ($j \neq i$)

Interference per core
bit vector

Core #	0	1	2	3
	0	0	0	0

Tracking Inter-Core Interference

- To identify App-interfering, for each core i
 - FST separately tracks interference caused by each core j ($j \neq i$)

Interference per core
bit vector

Core #	0	1	2	3
0	-	0	0	0
1	0	-	0	0
2	0	0	-	0
3	0	0	0	-

Tracking Inter-Core Interference

- To identify App-interfering, for each core i
 - FST separately tracks interference caused by each core j ($j \neq i$)

Interference per core
bit vector

Interfered with core

Core #		0	1	2	3
Interfering core	0	-	0	0	0
	1	0	-	0	0
	2	0	0	-	0
	3	0	0	0	-

Tracking Inter-Core Interference

- To identify App-interfering, for each core i
 - FST separately tracks interference caused by each core j ($j \neq i$)

Interference per core
bit vector

Interfered with core

Core #		0	1	2	3
Interfering core	0	-	0	0	0
	1	0	-	0	0
	2	0	0	-	0
	3	0	0	0	-

Excess Cycles
Counters per core

Cnt 0	Cnt 1	Cnt 2	Cnt 3
-------	-------	-------	-------

Tracking Inter-Core Interference

- To identify App-interfering, for each core i
 - FST separately tracks interference caused by each core j ($j \neq i$)

Interference per core
bit vector

Interfered with core

Core #		0	1	2	3
Interfering core	0	-	0	0	0
	1	0	-	0	0
	2	0	0	-	0
	3	0	0	0	-

Excess Cycles
Counters per core

-	Cnt 0,1	Cnt 0,2	Cnt 0,3
Cnt 1,0	-	Cnt 1,2	Cnt 1,3
Cnt 2,0	Cnt 2,1	-	Cnt 2,3
Cnt 3,0	Cnt 3,1	Cnt 3,2	-

Tracking Inter-Core Interference

- To identify App-interfering, for each core i
 - FST separately tracks interference caused by each core j ($j \neq i$)

Interference per core
bit vector

Interfered with core

Core #	0	1	2	3
0	-	0	0	0
1	0	-	0	0
2	0	1	-	0
3	0	0	0	-

Interfering core

Excess Cycles
Counters per core

-	Cnt 0,1	Cnt 0,2	Cnt 0,3
Cnt 1,0	-	Cnt 1,2	Cnt 1,3
Cnt 2,0	Cnt 2,1	-	Cnt 2,3
Cnt 3,0	Cnt 3,1	Cnt 3,2	-

Tracking Inter-Core Interference

- To identify App-interfering, for each core i
 - FST separately tracks interference caused by each core j ($j \neq i$)

Interference per core
bit vector

Excess Cycles
Counters per core

Interfered with core

Core # 0 1 2 3

Interfering core {

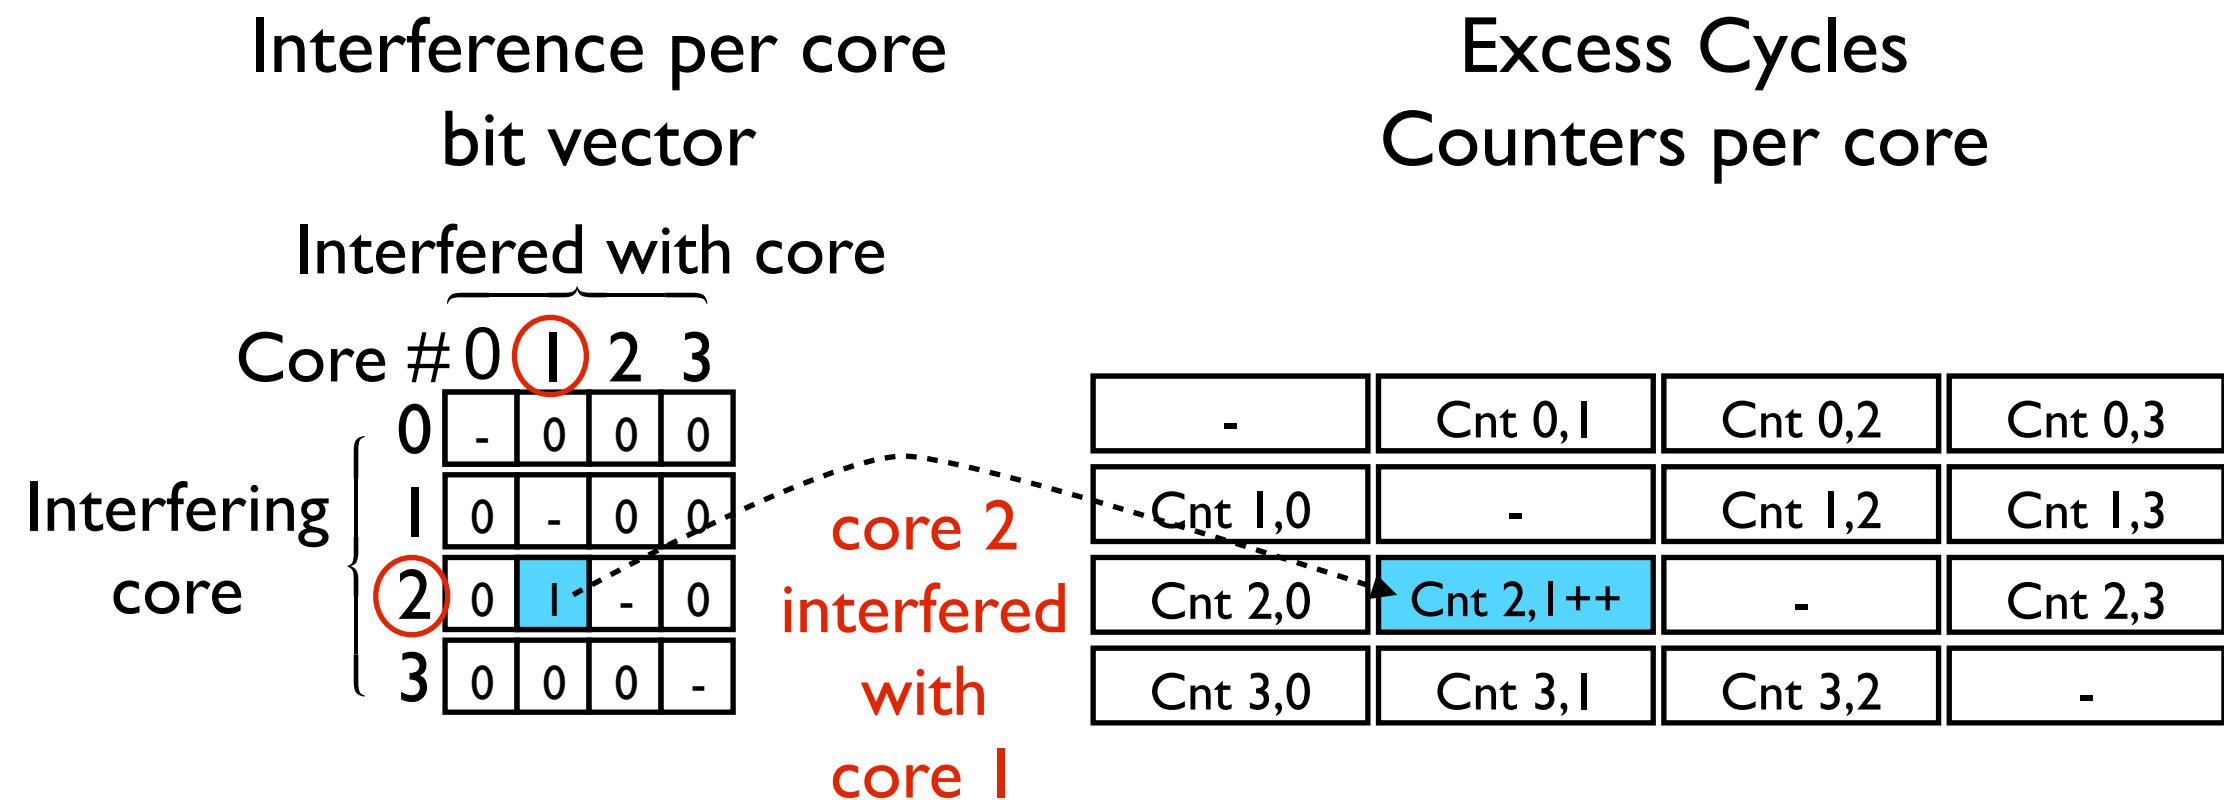
0	-	0	0	0
1	0	-	0	0
2	0	1	-	0
3	0	0	0	-

core 2 interfered with core 1

-	Cnt 0,1	Cnt 0,2	Cnt 0,3
Cnt 1,0	-	Cnt 1,2	Cnt 1,3
Cnt 2,0	Cnt 2,1	-	Cnt 2,3
Cnt 3,0	Cnt 3,1	Cnt 3,2	-

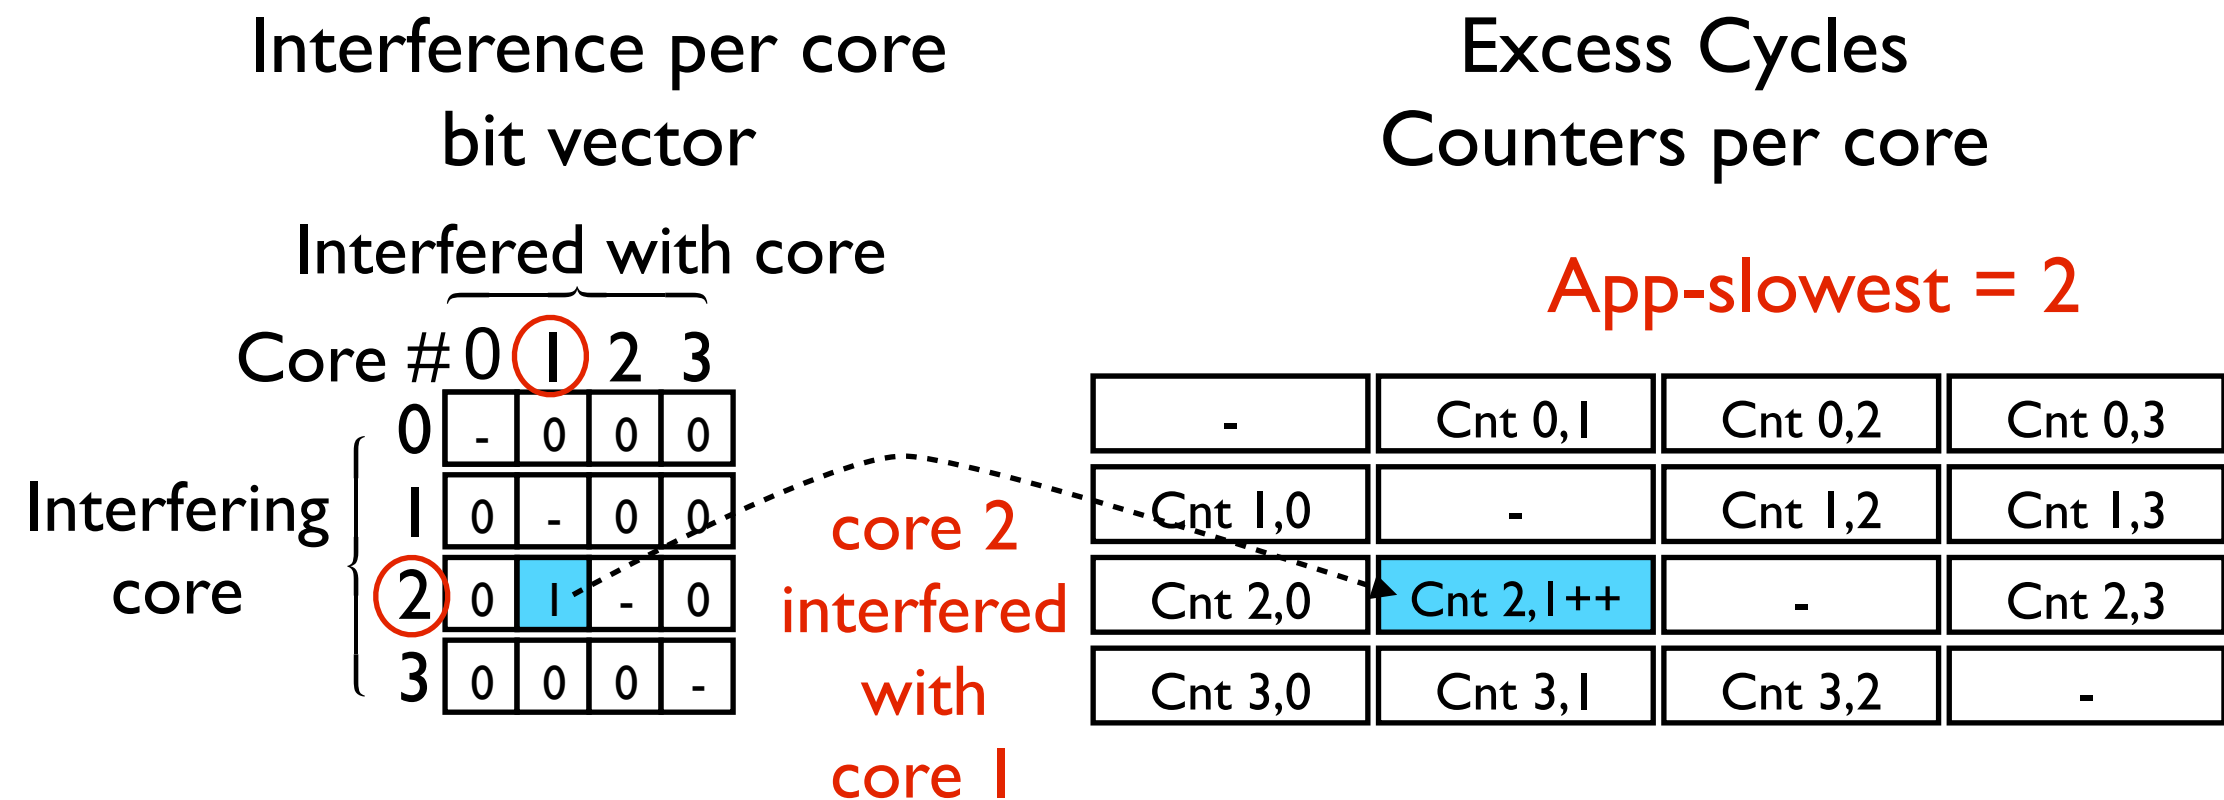
Tracking Inter-Core Interference

- To identify App-interfering, for each core i
 - FST separately tracks interference caused by each core j ($j \neq i$)



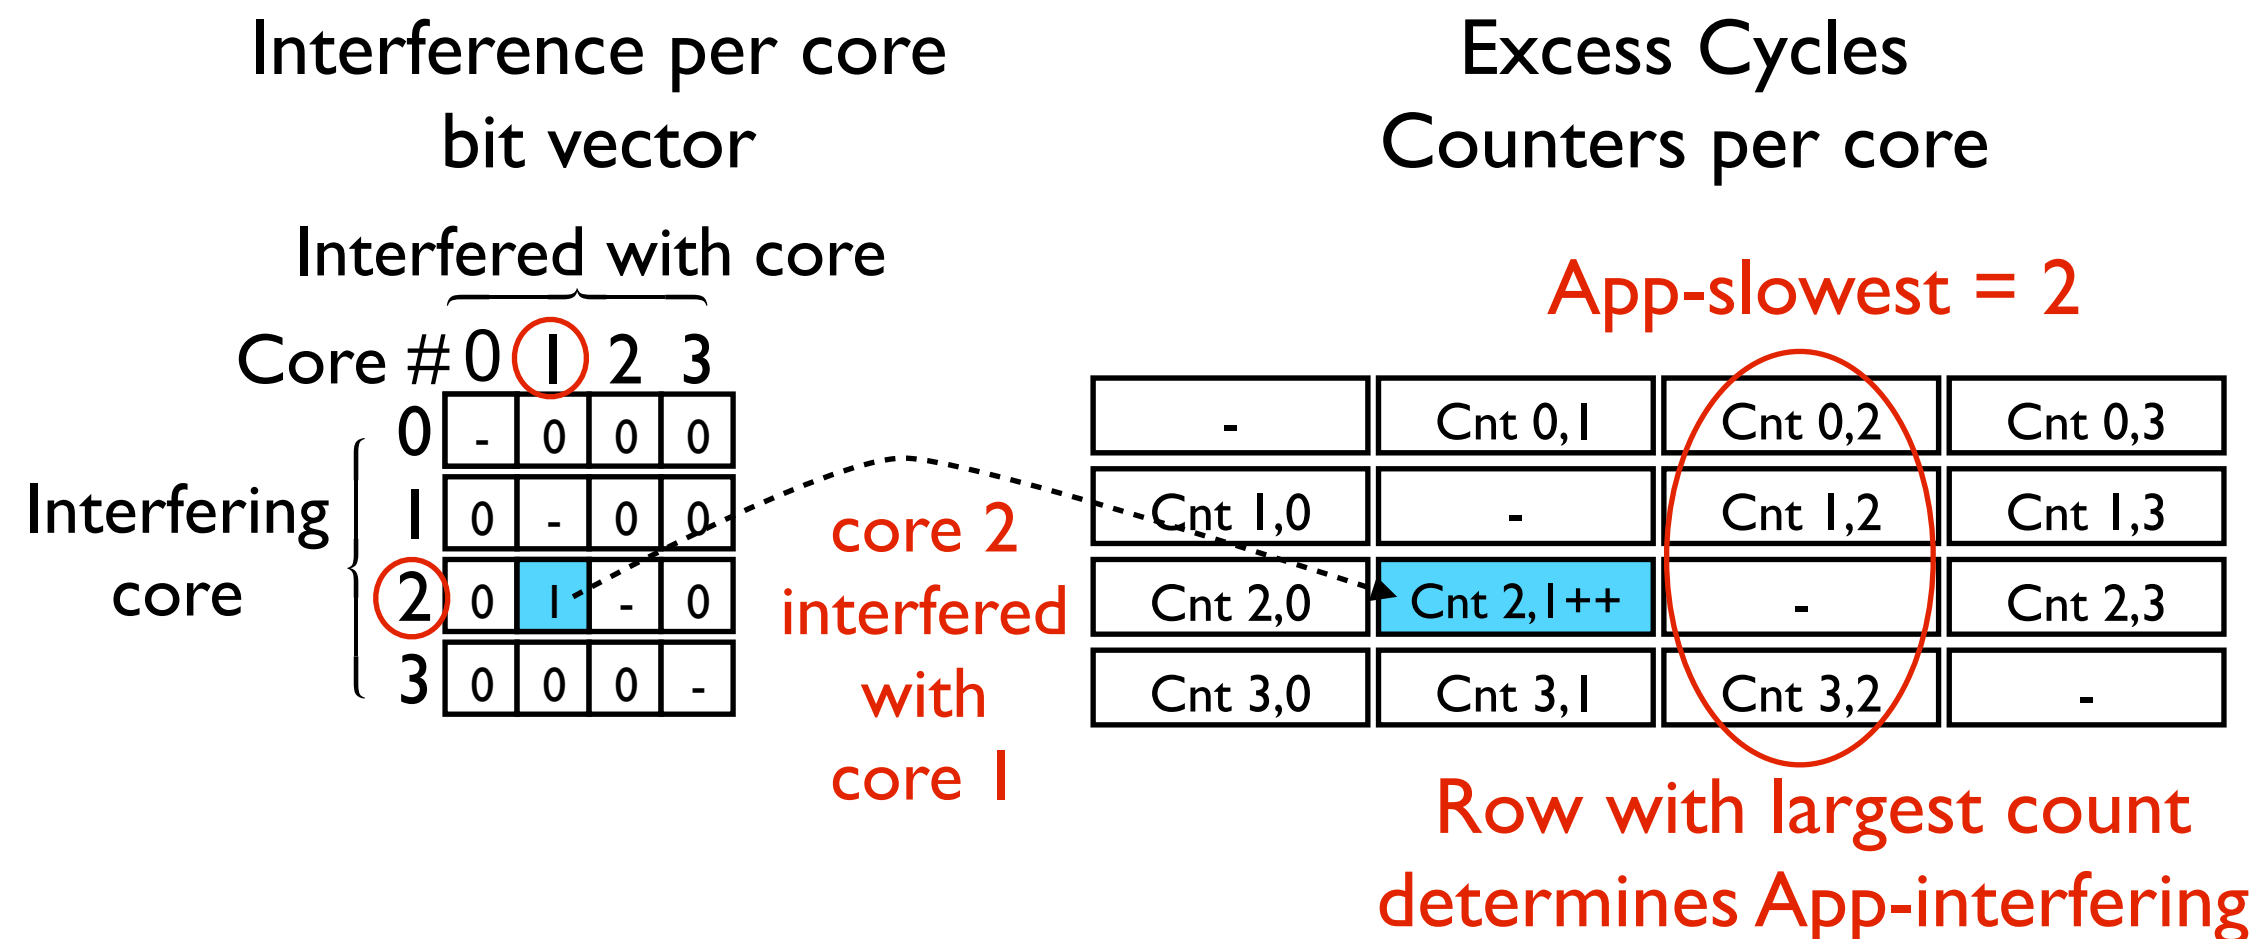
Tracking Inter-Core Interference

- To identify App-interfering, for each core i
 - FST separately tracks interference caused by each core j ($j \neq i$)



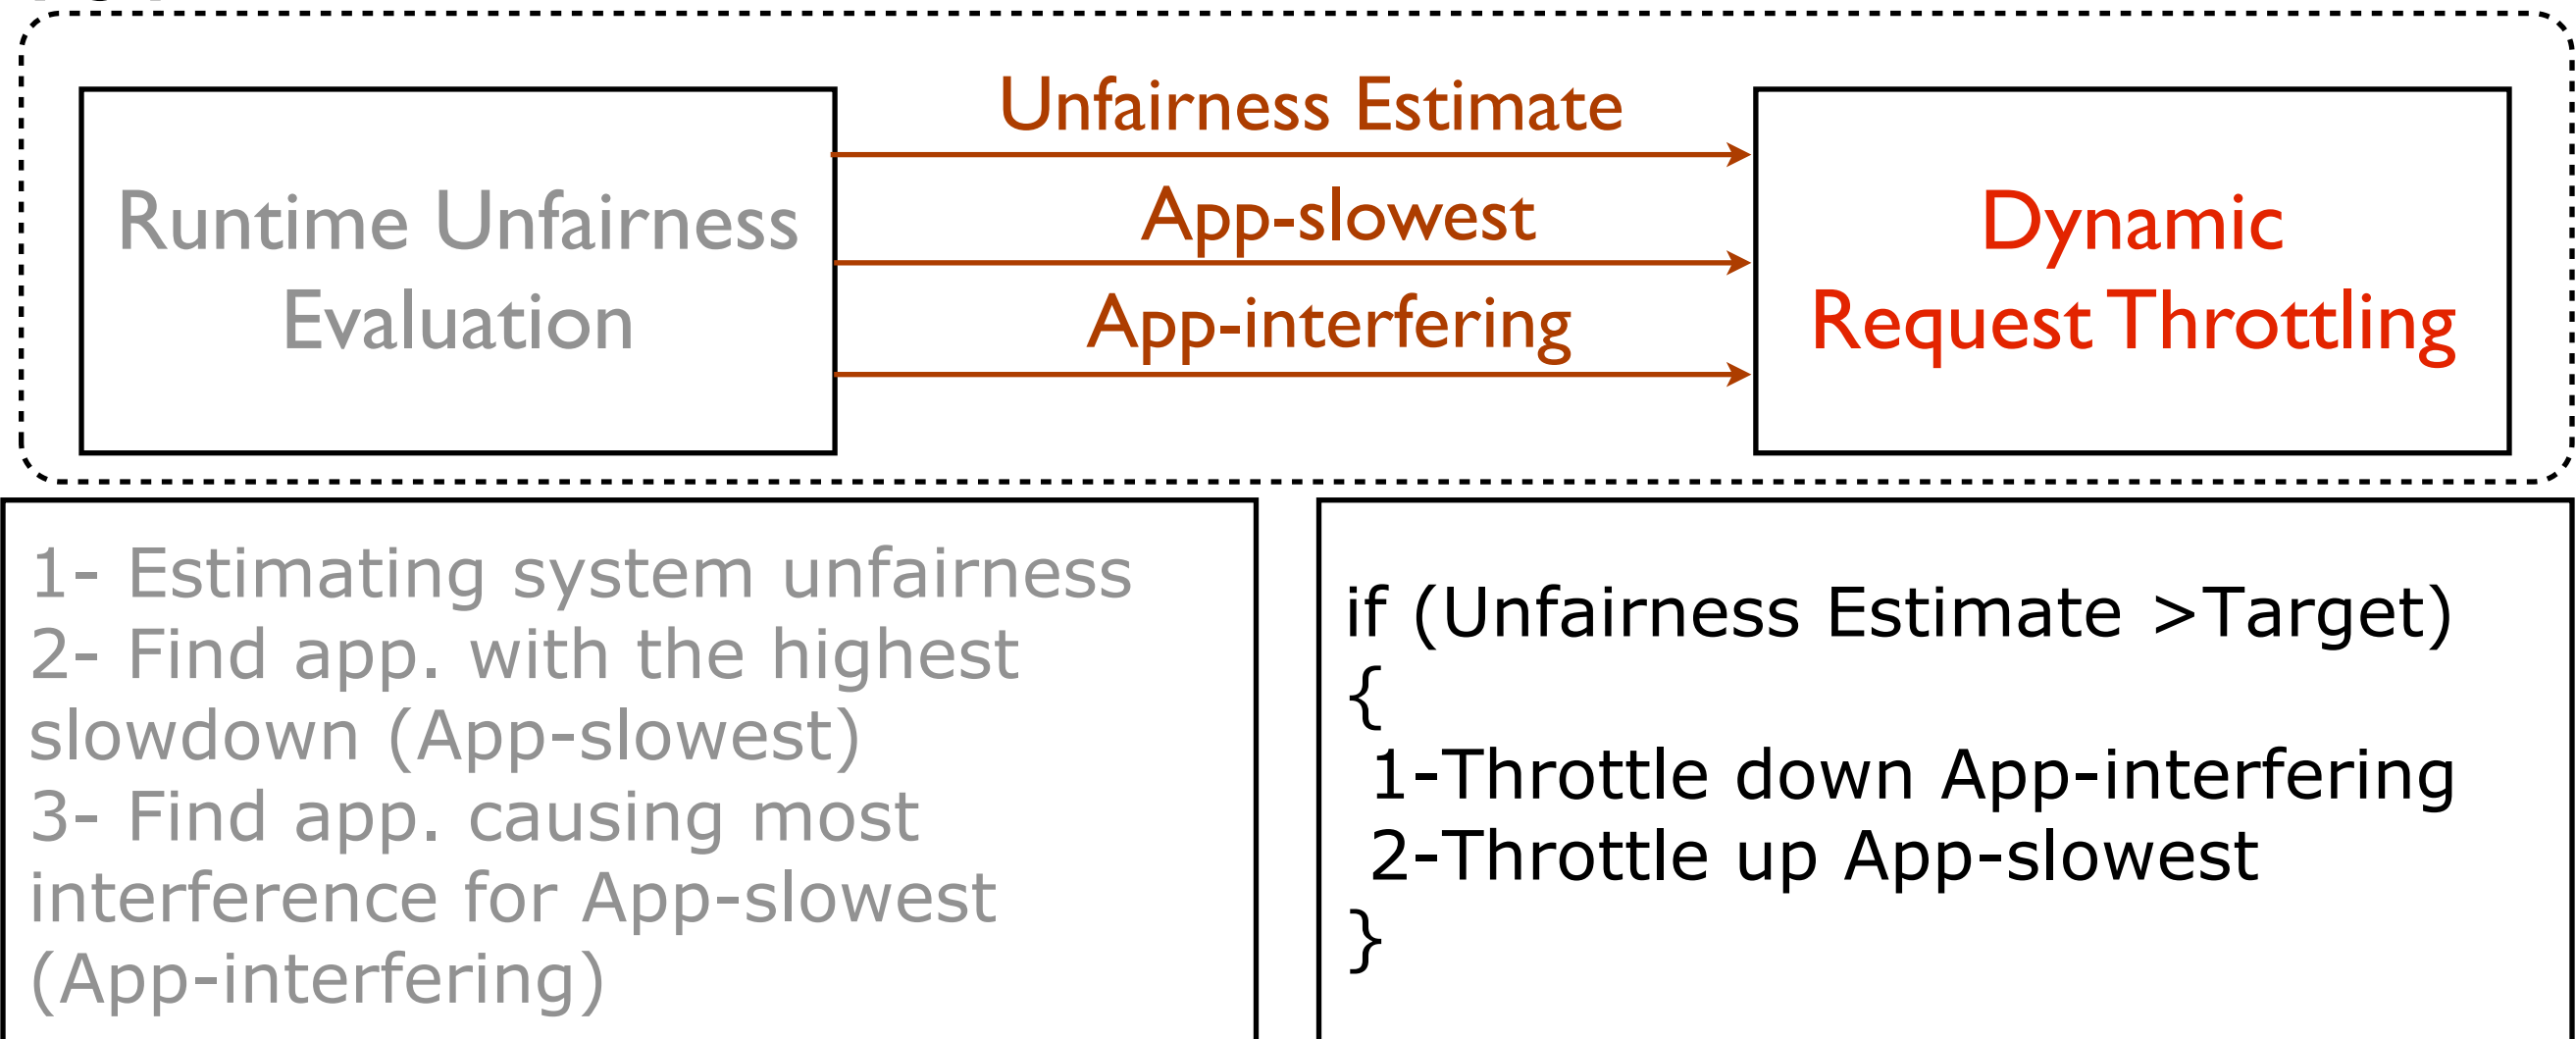
Tracking Inter-Core Interference

- To identify App-interfering, for each core i
 - FST separately tracks interference caused by each core j ($j \neq i$)



Fairness via Source Throttling (FST)

FST



Dynamic Request Throttling

Dynamic Request Throttling

- Goal: Adjust **how aggressively** each core makes requests to the shared resources

Dynamic Request Throttling

- Goal: Adjust **how aggressively** each core makes requests to the shared resources
- Mechanisms:
 - Miss Status Holding Register (MSHR) quota

Dynamic Request Throttling

- Goal: Adjust **how aggressively** each core makes requests to the shared resources
- Mechanisms:
 - Miss Status Holding Register (MSHR) quota
 - Controls the **number of concurrent requests** accessing shared resources from each application

Dynamic Request Throttling

- Goal: Adjust **how aggressively** each core makes requests to the shared resources
- Mechanisms:
 - Miss Status Holding Register (MSHR) quota
 - Controls the **number of concurrent requests** accessing shared resources from each application
 - Request injection frequency

Dynamic Request Throttling

- Goal: Adjust **how aggressively** each core makes requests to the shared resources
- Mechanisms:
 - Miss Status Holding Register (MSHR) quota
 - Controls the **number of concurrent requests** accessing shared resources from each application
 - Request injection frequency
 - Controls **how often memory requests are issued** to the last level cache from the MSHRs

Dynamic Request Throttling

- **Throttling level** assigned to each core determines both **MSHR quota** and **request injection rate**

Dynamic Request Throttling

- **Throttling level** assigned to each core determines both **MSHR quota** and **request injection rate**

Throttling level	MSHR quota	Request Injection Rate
100%	128	Every cycle
50%	64	Every other cycle
25%	32	Once every 4 cycles
10%	12	Once every 10 cycles
5%	6	Once every 20 cycles
4%	5	Once every 25 cycles
3%	3	Once every 30 cycles
2%	2	Once every 50 cycles

Total # of
MSHRs: 128

Dynamic Request Throttling

- **Throttling level** assigned to each core determines both **MSHR quota** and **request injection rate**


Throttling level	MSHR quota	Request Injection Rate
100%	128	Every cycle
50%	64	Every other cycle
25%	32	Once every 4 cycles
10%	12	Once every 10 cycles
5%	6	Once every 20 cycles
4%	5	Once every 25 cycles
3%	3	Once every 30 cycles
2%	2	Once every 50 cycles

Total # of
MSHRs: 128

FST at Work




FST




	Core 0	Core 1	Core 2	Core 3
Interval i				
Interval $i + 1$				
Interval $i + 2$				

Throttling Levels

FST at Work




FST



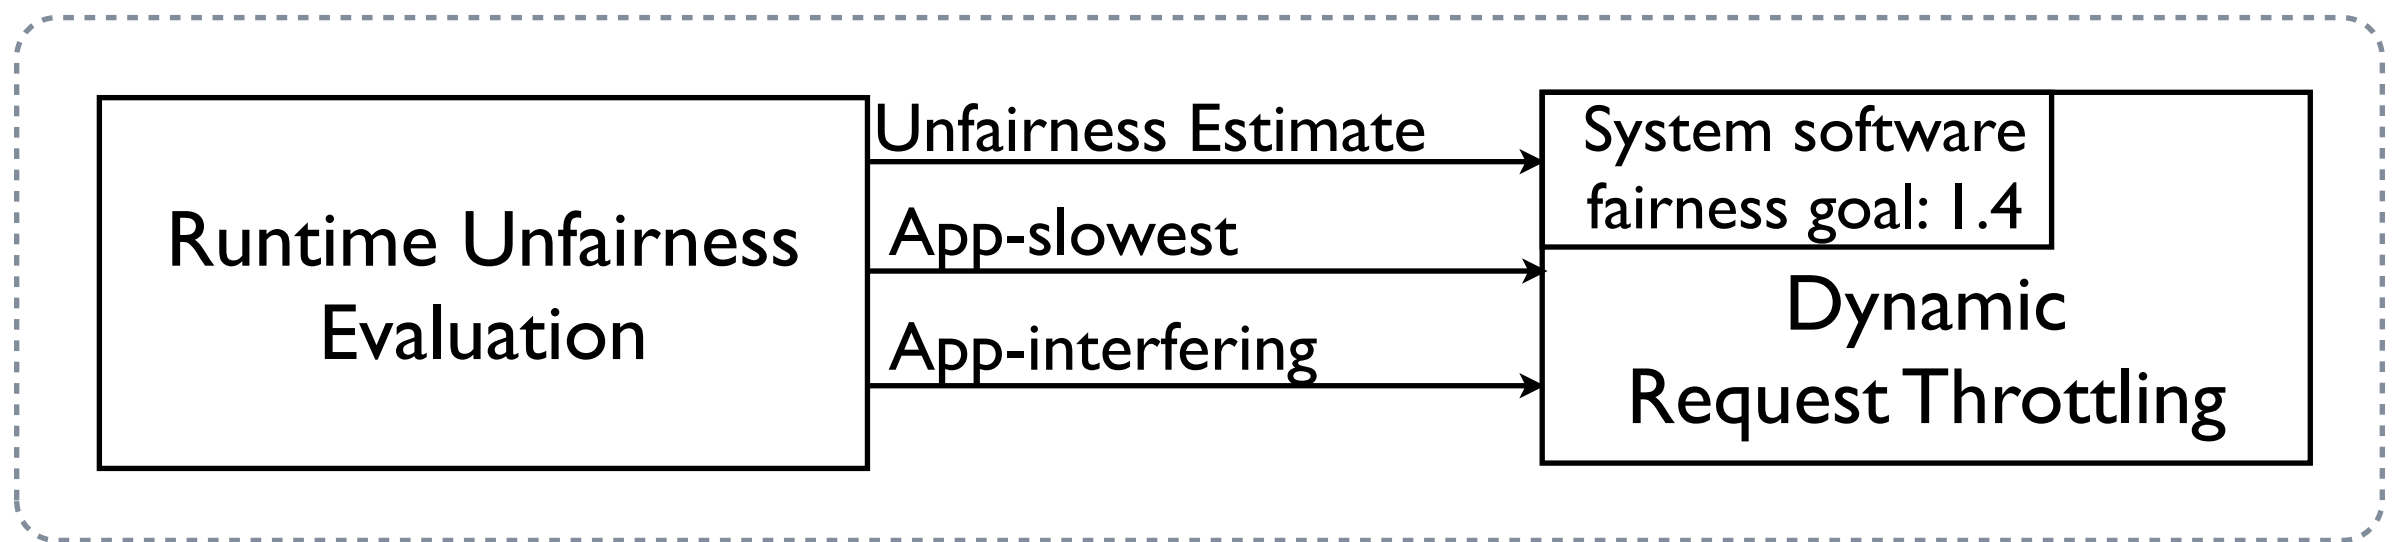
	Core 0	Core 1	Core 2	Core 3
Interval i	50%	100%	10%	100%
Interval $i + 1$				
Interval $i + 2$				

Throttling Levels

FST at Work




FST



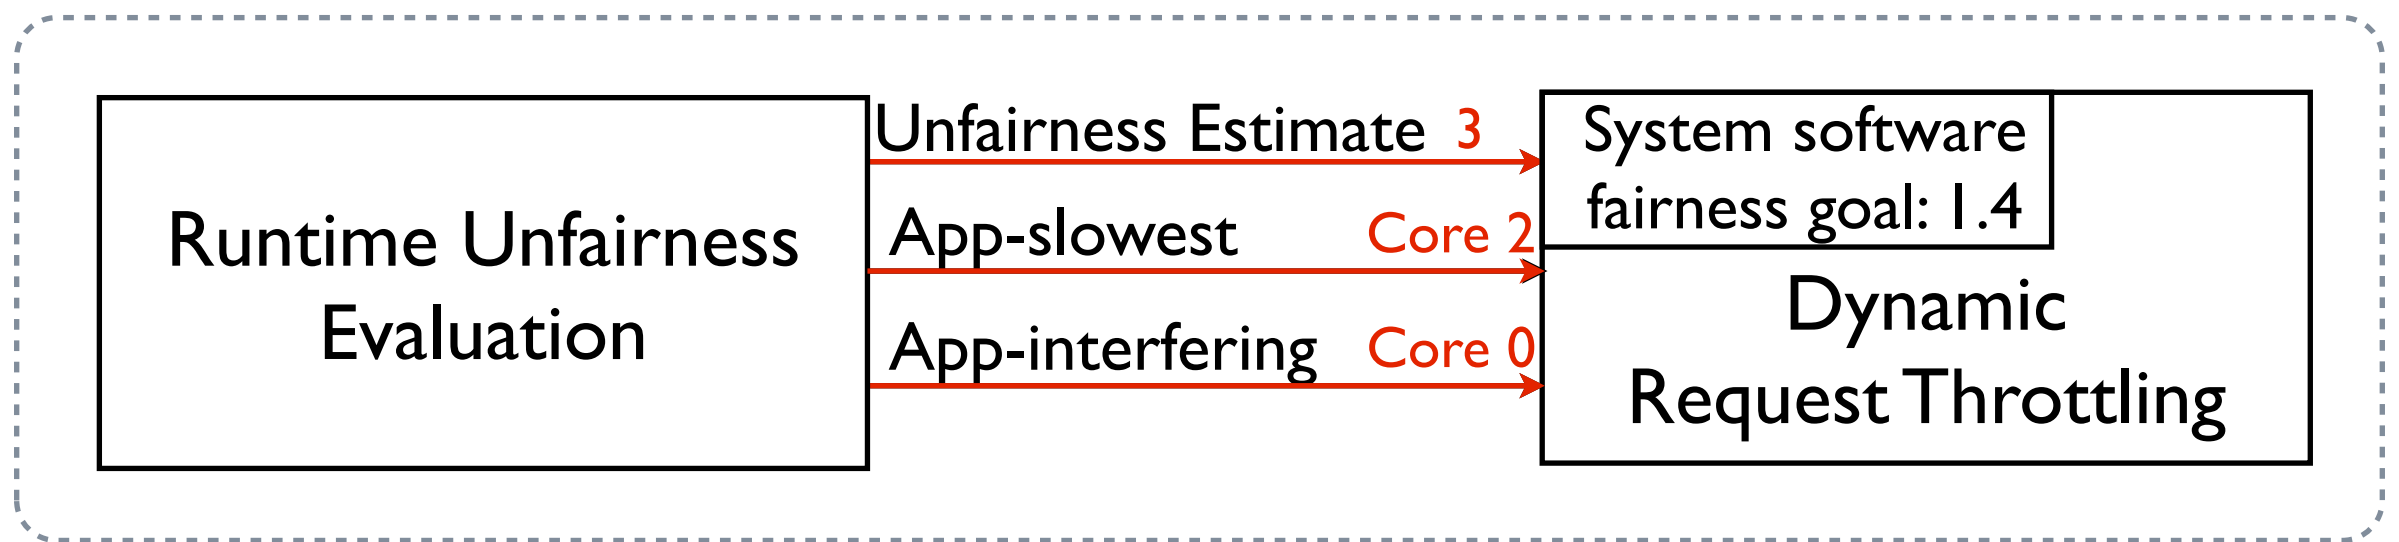
	Core 0	Core 1	Core 2	Core 3
Interval i	50%	100%	10%	100%
Interval $i + 1$				
Interval $i + 2$				

Throttling Levels

FST at Work




FST



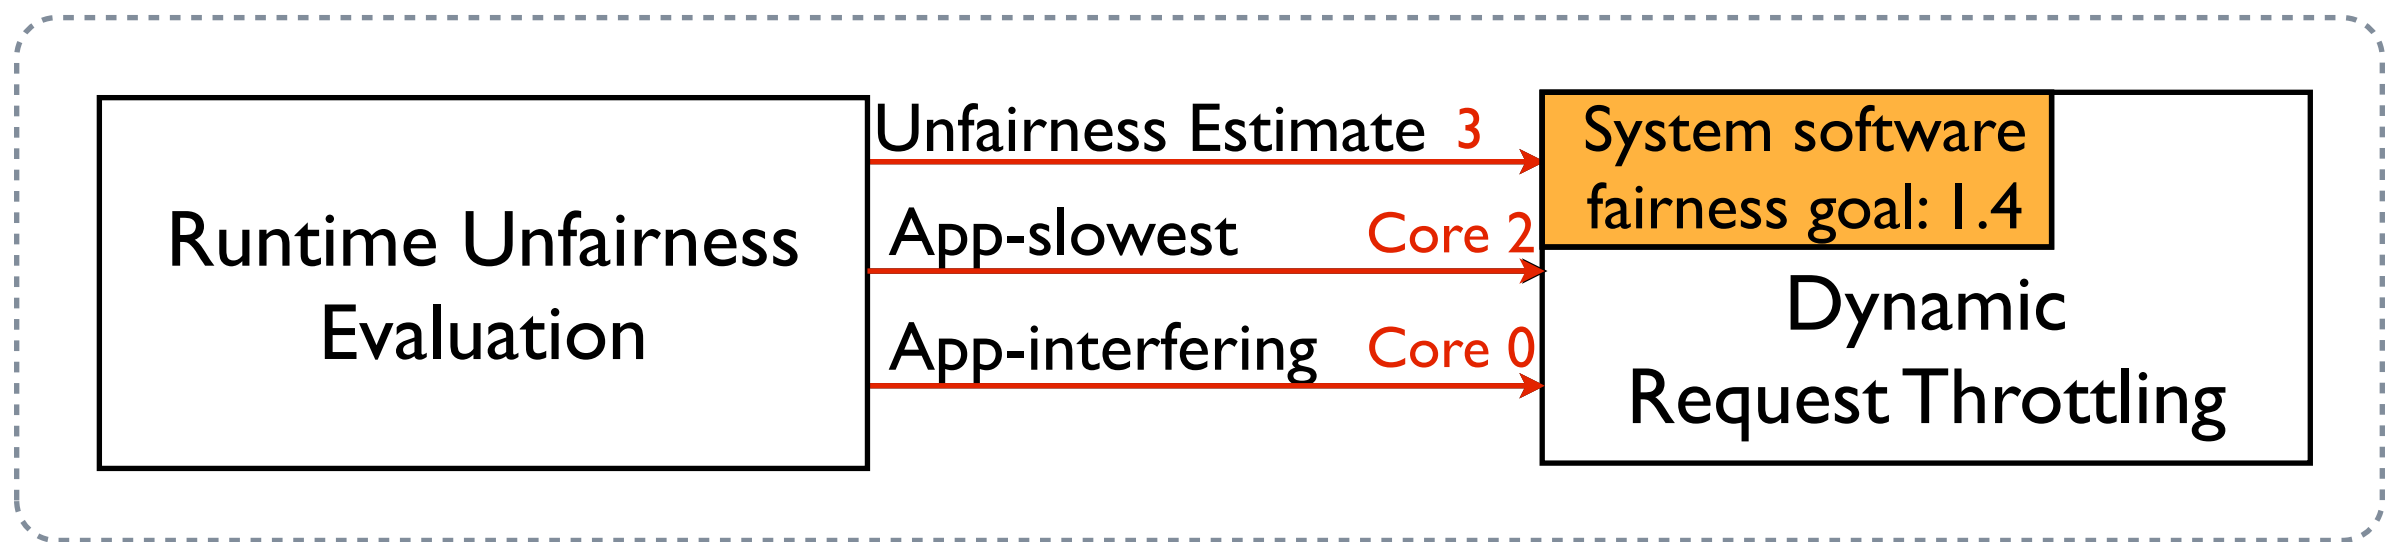
	Core 0	Core 1	Core 2	Core 3
Interval i	50%	100%	10%	100%
Interval $i + 1$				
Interval $i + 2$				

Throttling Levels

FST at Work




FST



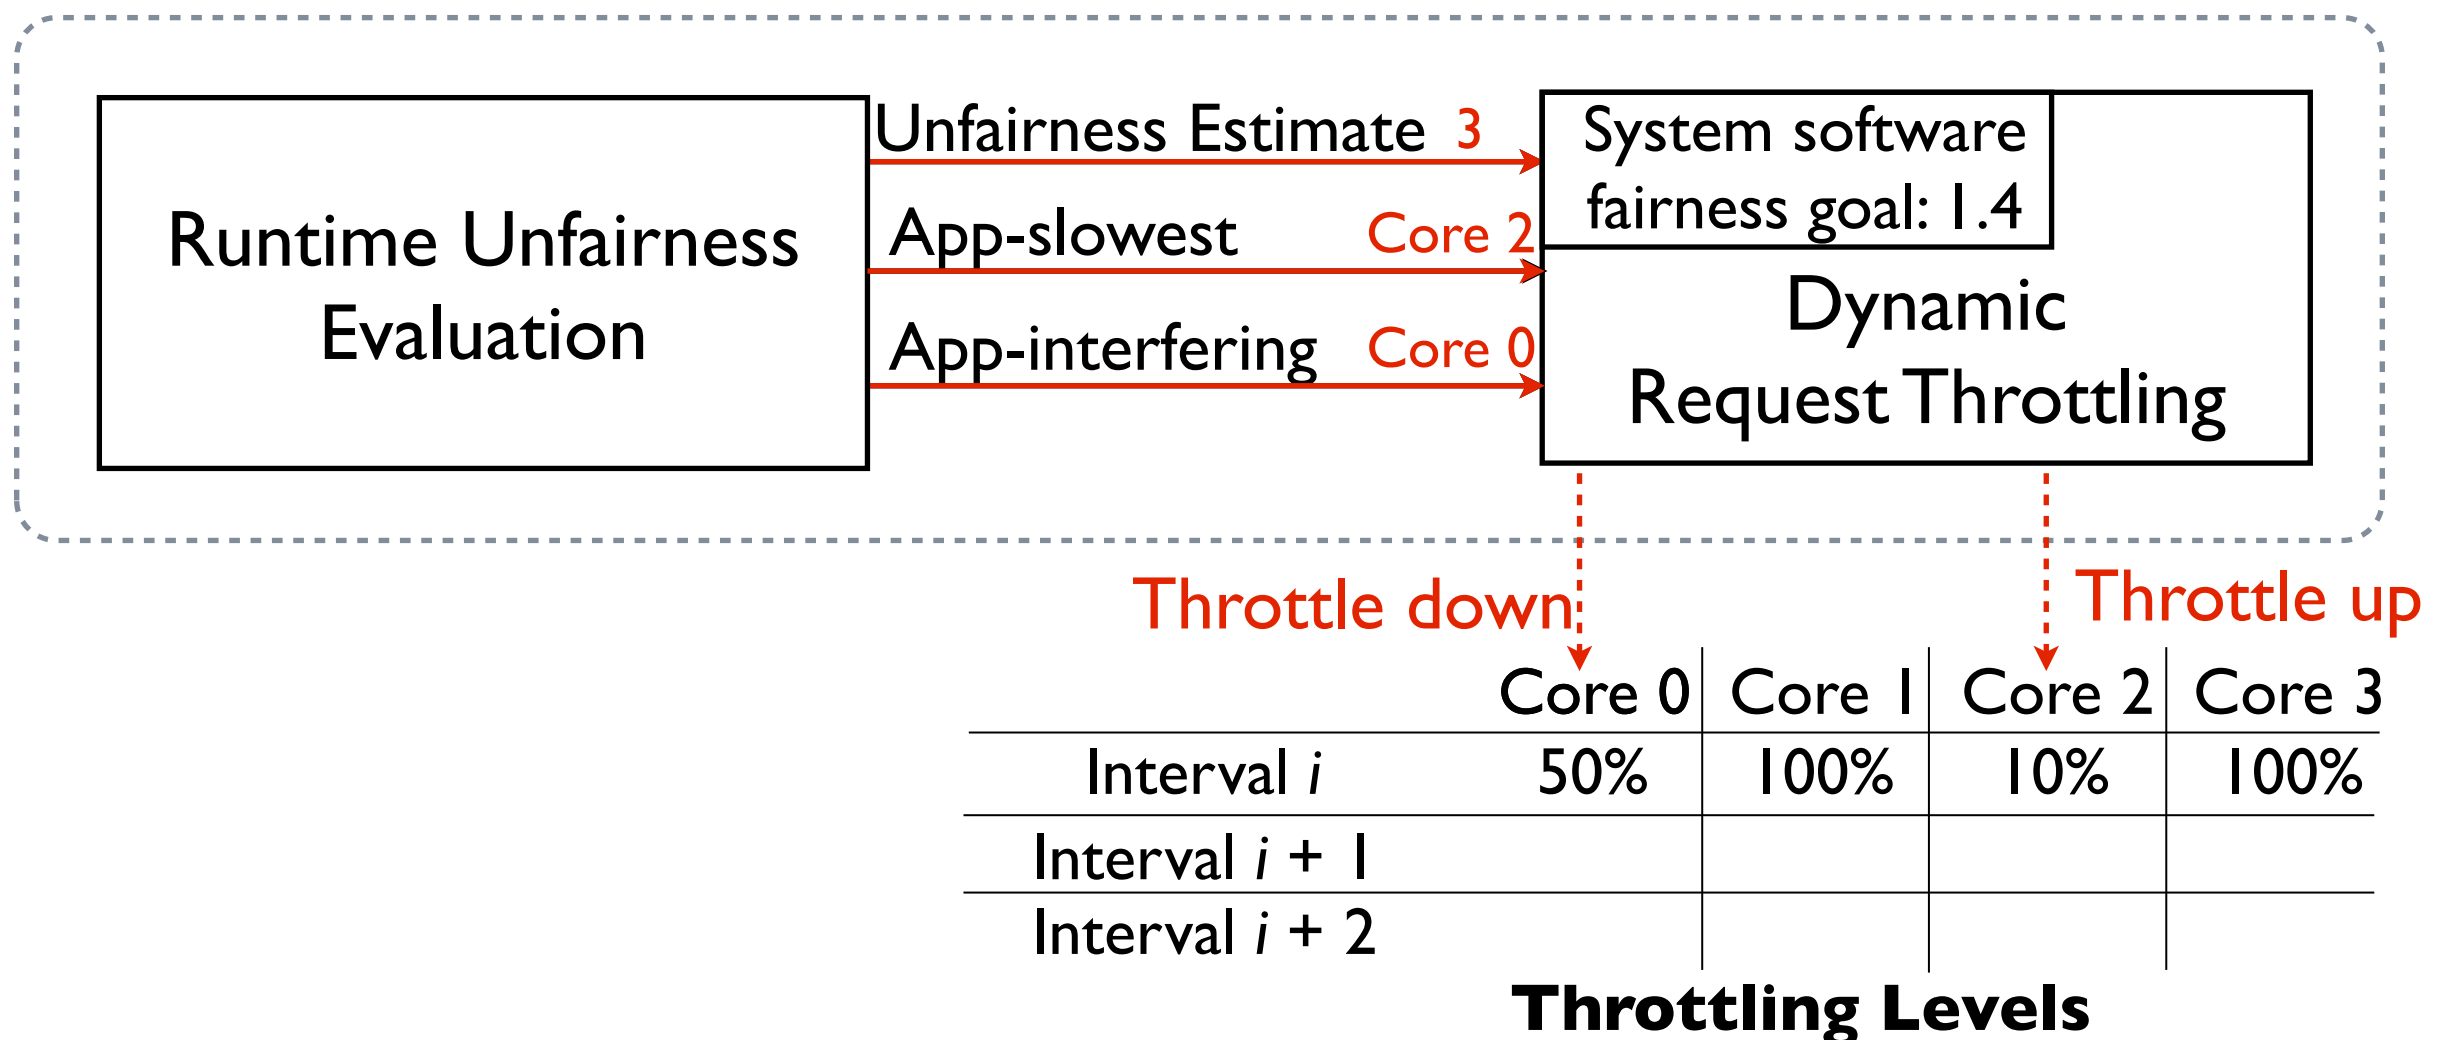
	Core 0	Core 1	Core 2	Core 3
Interval i	50%	100%	10%	100%
Interval $i + 1$				
Interval $i + 2$				

Throttling Levels


FST at Work



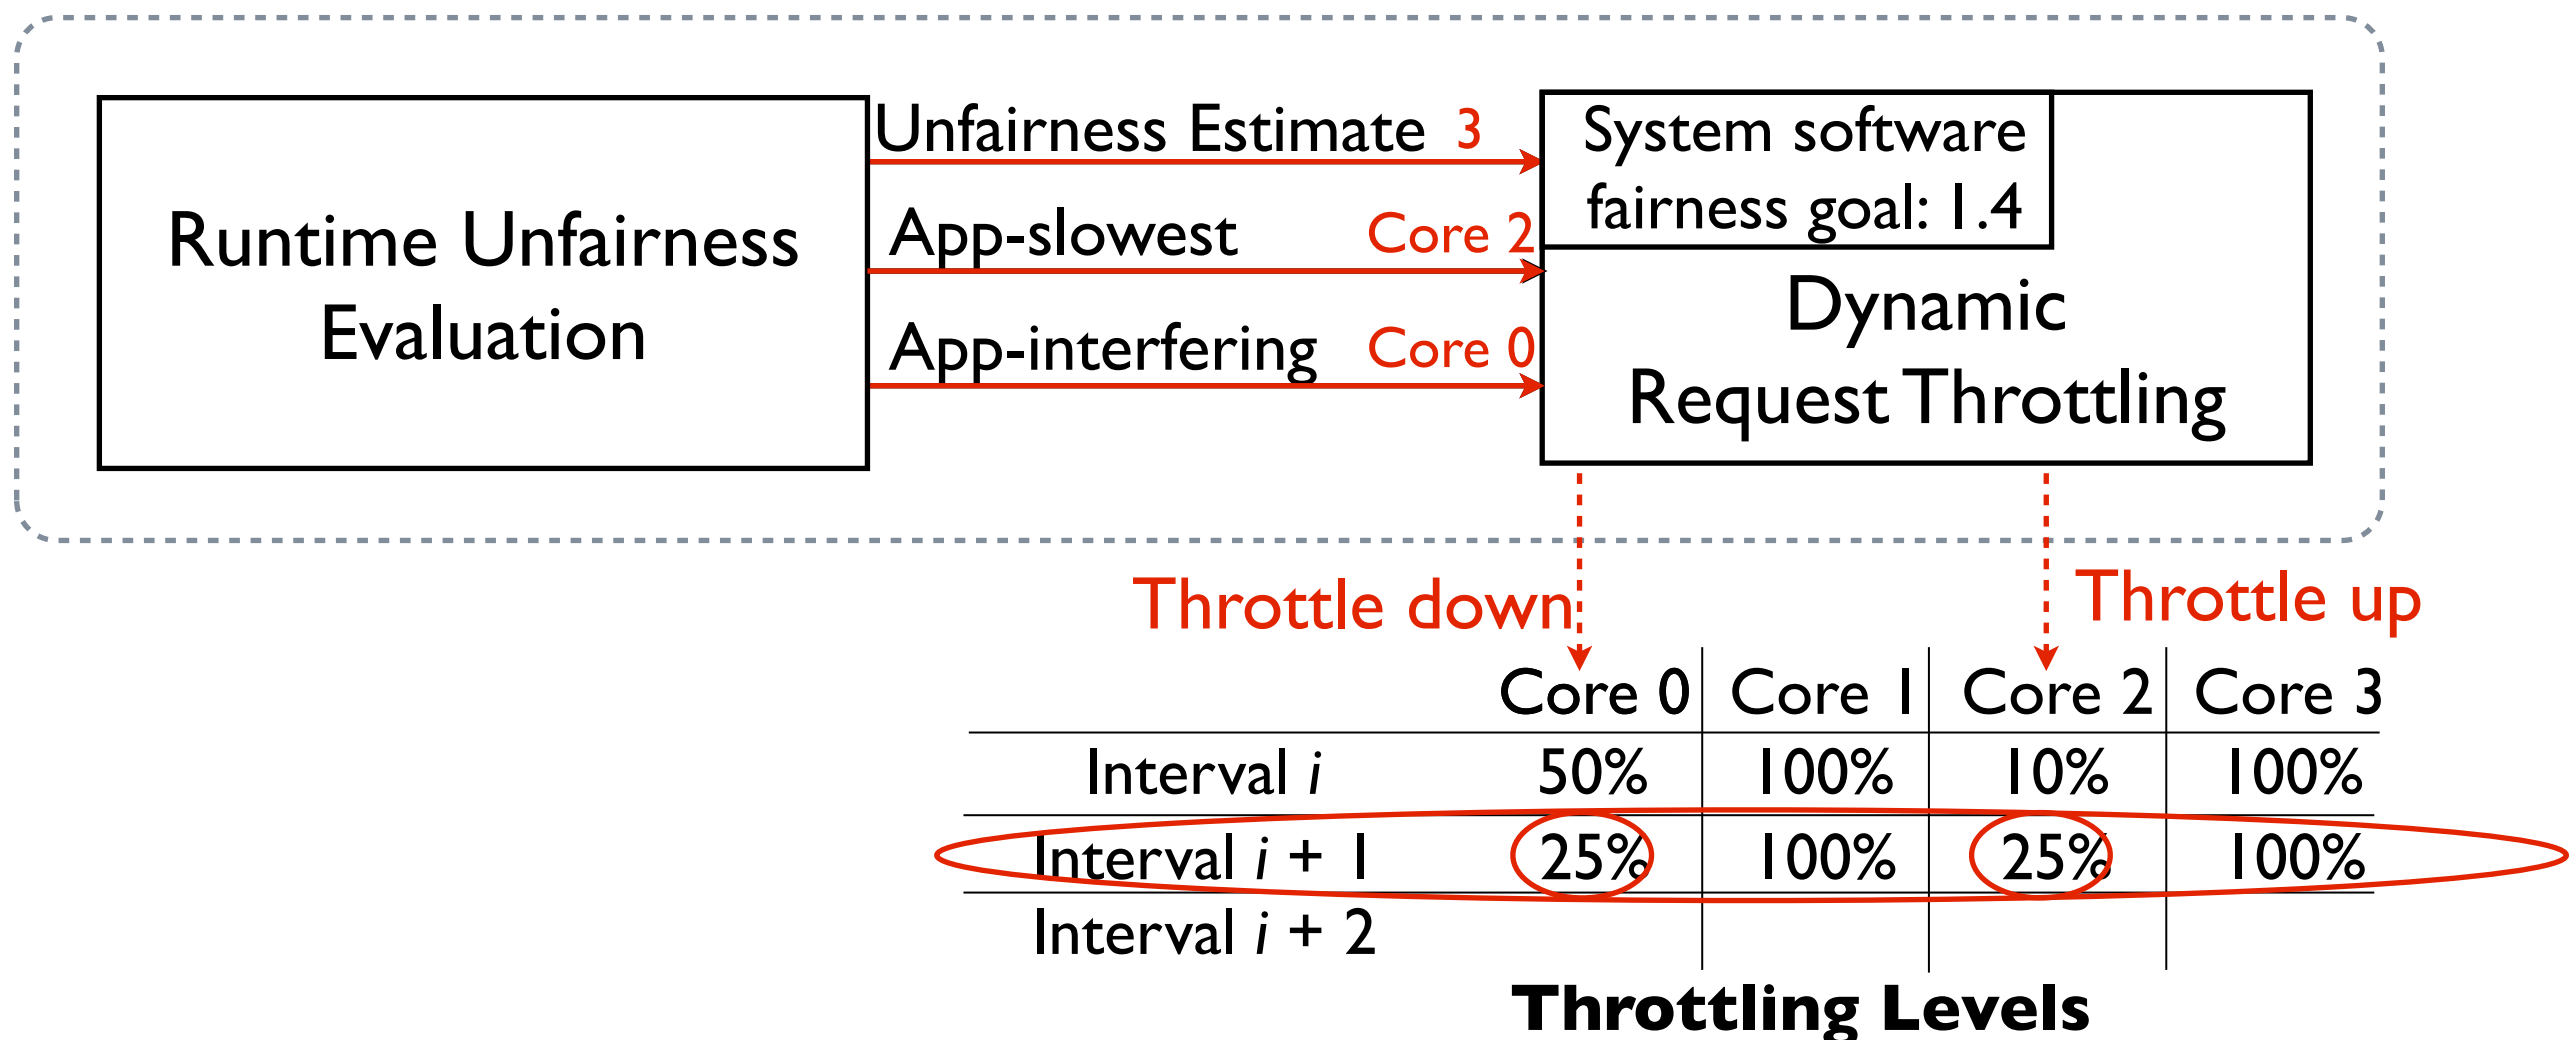
FST




FST at Work



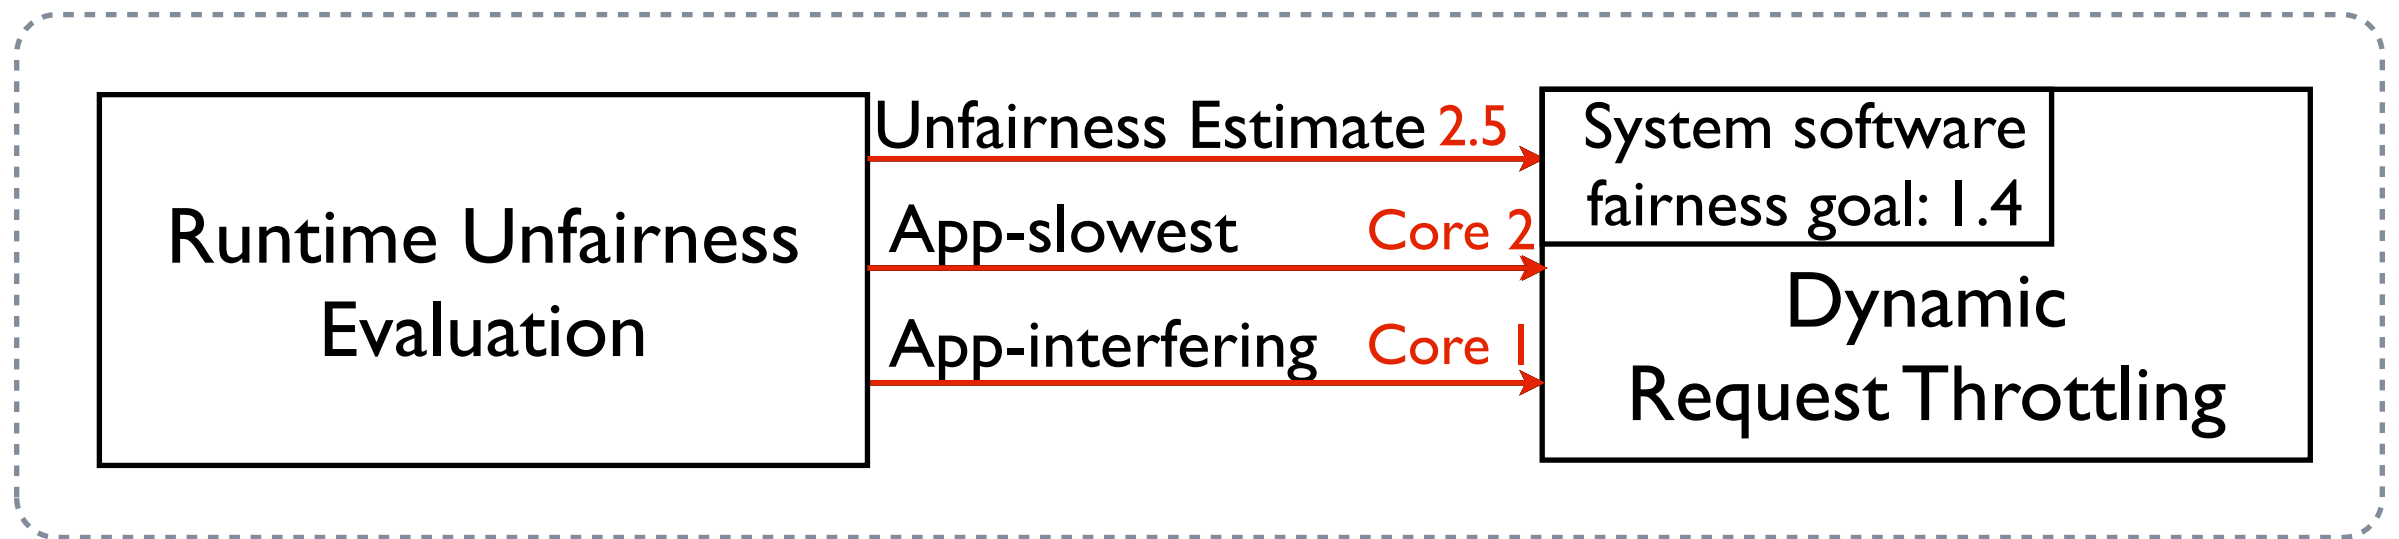
FST



FST at Work



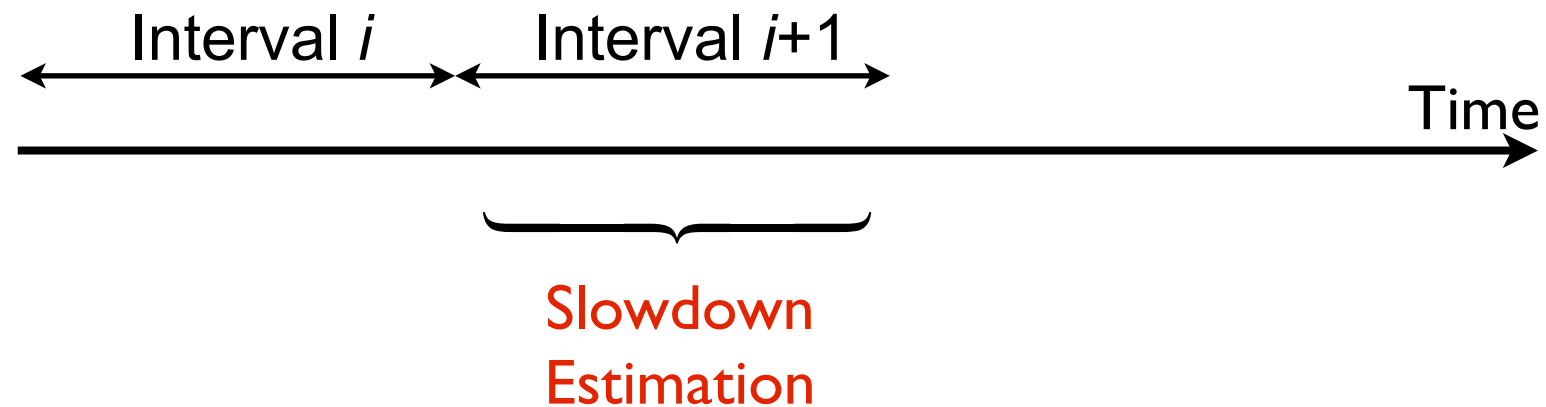
FST



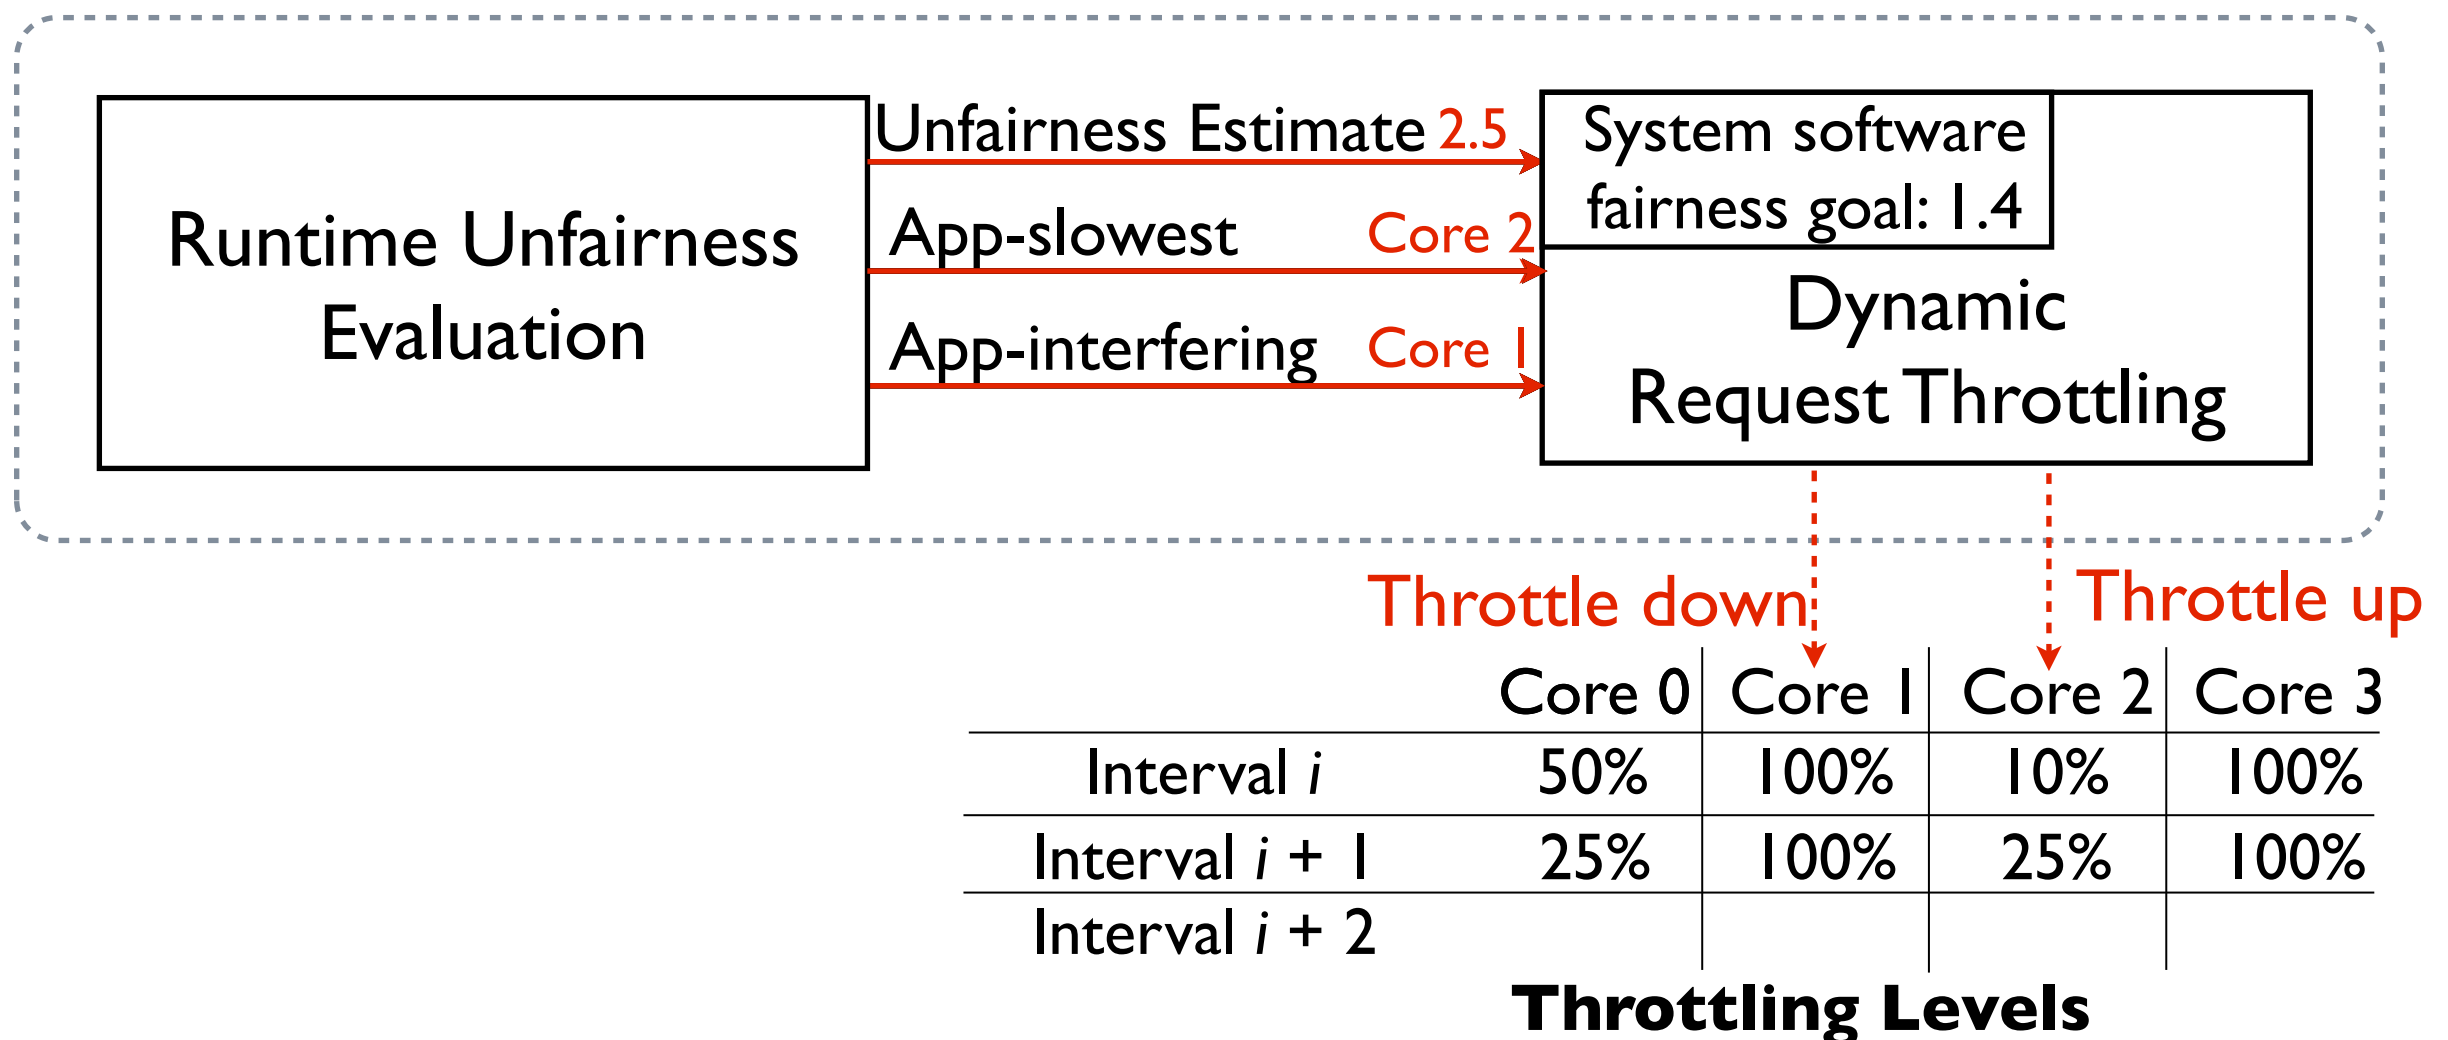
	Core 0	Core 1	Core 2	Core 3
Interval i	50%	100%	10%	100%
Interval $i + 1$	25%	100%	25%	100%
Interval $i + 2$				

Throttling Levels

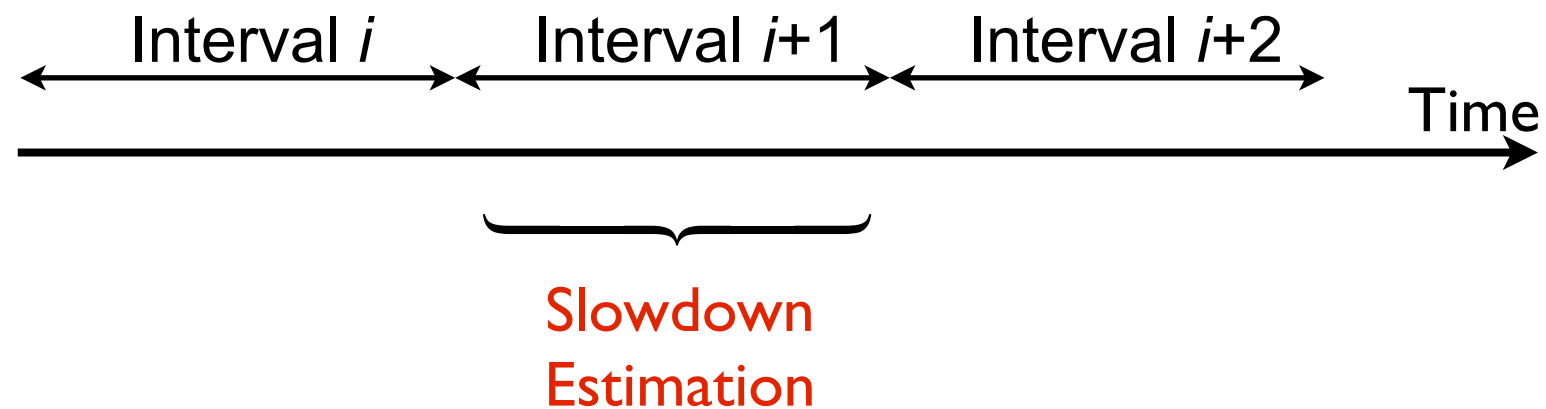
FST at Work



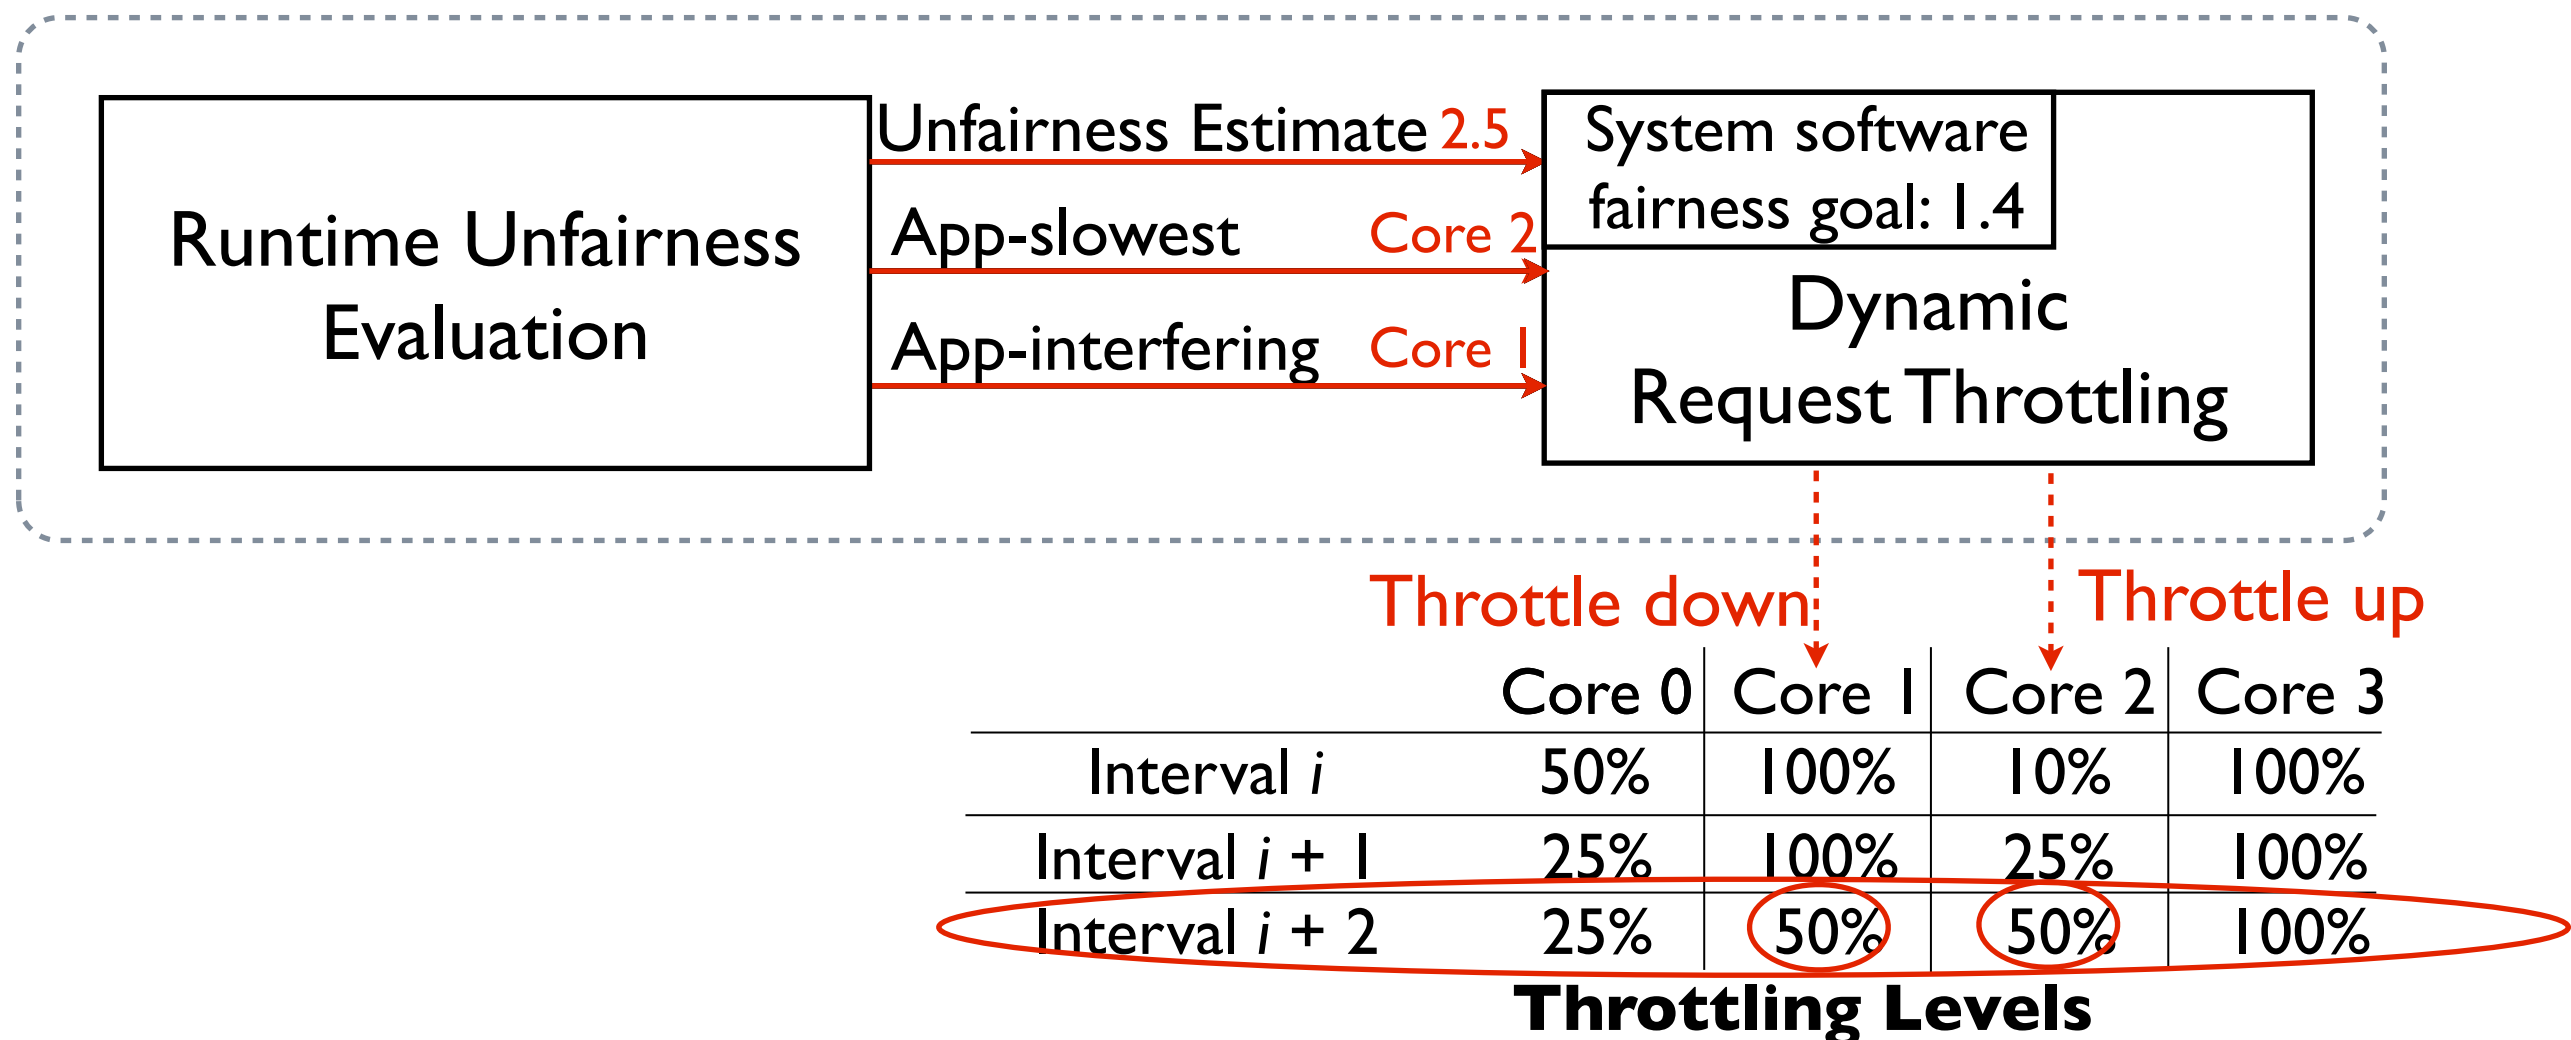
FST



FST at Work



FST



System Software Support

System Software Support

- Different fairness objectives can be configured by **system software**

System Software Support

- Different fairness objectives can be configured by **system software**
 - *Estimated Unfairness > Target Unfairness*

System Software Support

- Different fairness objectives can be configured by **system software**
 - *Estimated Unfairness* > *Target Unfairness*
 - *Estimated **Max Slowdown*** > *Target **Max Slowdown***

System Software Support

- Different fairness objectives can be configured by **system software**
 - *Estimated Unfairness* > *Target Unfairness*
 - *Estimated Max Slowdown* > *Target Max Slowdown*
 - *Estimated Slowdown(i)* > *Target Slowdown(i)*

System Software Support

- Different fairness objectives can be configured by **system software**
 - *Estimated Unfairness* > *Target Unfairness*
 - *Estimated **Max Slowdown*** > *Target **Max Slowdown***
 - *Estimated **Slowdown(i)*** > *Target **Slowdown(i)***
- Support for **thread priorities**

System Software Support

- Different fairness objectives can be configured by **system software**
 - *Estimated Unfairness* > *Target Unfairness*
 - *Estimated Max Slowdown* > *Target Max Slowdown*
 - *Estimated Slowdown(i)* > *Target Slowdown(i)*
- Support for **thread priorities**
 - *Weighted Slowdown(i)* =
$$\text{Estimated Slowdown}(i) \times \text{Weight}(i)$$

Hardware Cost

- Total storage cost required for 4 cores is ~ 12KB
- FST does not require any structures or logic that are on the processor's critical path


Outline

- Background and Problem
- Motivation for Source Throttling
- Fairness via Source Throttling (FST)
- Evaluation
- Conclusion


Evaluation Methodology

- x86 cycle accurate simulator
- Baseline processor configuration
 - Per-core
 - 4-wide issue, out-of-order, 256 entry ROB
- Shared (4-core system)
 - 128 MSHRs
 - 2 MB, 16-way L2 cache
- Main Memory
 - DDR3 1333 MHz
 - Latency of 15ns per command (t_{RP} , t_{RCD} , CL)
 - 8B wide core to memory bus


System Unfairness Results




System Unfairness Results




System Unfairness Results




System Unfairness Results




System Unfairness Results




System Unfairness Results




System Unfairness Results




System Unfairness Results




System Performance Results




System Performance Results



System Performance Results



System Performance Results



Conclusion

- Fairness via Source Throttling (FST) is a new fair and high-performance shared resource management approach for CMPs
- Dynamically monitors unfairness and throttles down sources of interfering memory requests
- Eliminates the need for and complexity of multiple per-resource fairness techniques
- Improves both system fairness and performance
- Incorporates thread weights and enables different fairness objectives

Fairness via Source Throttling:

A configurable and high-performance fairness
substrate for multi-core memory systems

Eiman Ebrahimi*

Chang Joo Lee*

Onur Mutlu[‡]

Yale N. Patt*

*** HPS Research Group
The University of Texas at Austin**

**‡ Computer Architecture Laboratory
Carnegie Mellon University**

Backups

Other Source-Based Techniques

- Herdrich et. al. ICS '09
Rate-based QoS techniques for cache/memory in CMP platforms
- Zhang et. al., USENIX '09
Hardware execution throttling for multi-core resource management
- Jahre and Natvig, Computing Frontiers '09
A light-weight fairness mechanism for chip multiprocessor memory systems

Interference-Aware Thread Scheduling

- Zhuravlev et. al. ASPLOS '10
Addressing Shared Resource Contention in Multicore Processors Via Scheduling
- Schedules applications which interfere less with each other as best as possible
- Advantages of FST:
 - The mix of apps may force co-scheduling of intensive applications
 - FST can make scheduling decisions easier for system software
- Advantages of using thread scheduling:
 - Does not require hardware support
- Approaches are complementary

Tracking Cache Interference

Shared Cache

Hash
Function

Core 1's pollution filter

Pollution
bit

Core id

0	0
0	0
⋮	⋮
0	0
0	0

Interfered with core

Core #	0	1	2	3
0	-	0	0	0
1	0	-	0	0
2	0	0	-	0
3	0	0	0	-

Interfering
core

Tracking Cache Interference

Shared Cache

Hash
Function

Core 1's pollution filter

Pollution bit
Core id

0	0
0	0
⋮	⋮
0	0
0	0

Interfered with core

Core #	0	1	2	3
0	-	0	0	0
1	0	-	0	0
2	0	0	-	0
3	0	0	0	-

Interfering
core

Tracking Cache Interference

Shared Cache

Hash
Function

Core 1's pollution filter

Pollution
bit

Core id

0	0
0	0
⋮	⋮
0	0
0	0

Interfered with core

Core #	0	1	2	3
0	-	0	0	0
1	0	-	0	0
2	0	0	-	0
3	0	0	0	-

Interfering
core

Tracking Cache Interference

Shared Cache

Hash
Function

Core 1's pollution filter

Pollution
bit

Core id

0	0
0	0
⋮	⋮
0	0
0	0

Interfered with core

Core #	0	1	2	3
0	-	0	0	0
1	0	-	0	0
2	0	0	-	0
3	0	0	0	-

Interfering
core

Tracking Cache Interference

Shared Cache

Core 1's pollution filter

Pollution bit | Core id

0	0
0	0
⋮	⋮
0	0
0	0

Interfered with core

Core #	0	1	2	3
0	-	0	0	0
1	0	-	0	0
2	0	0	-	0
3	0	0	0	-

Interfering core

Core 2's memory request evicts core 1's cache line from the shared cache

Hash Function

Tracking Cache Interference

Shared Cache

Core 2's memory request evicts core 1's cache line from the shared cache

Evicted line's address
from core 1

Hash
Function

Core 1's pollution filter

Pollution
bit

Core id

0	0
0	0
⋮	⋮
0	0
0	0

Interfered with core

Core #	0	1	2	3
0	-	0	0	0
1	0	-	0	0
2	0	0	-	0
3	0	0	0	-

Interfering
core

Tracking Cache Interference



Core 1's pollution filter

Pollution bit | Core id

0	0
0	0
⋮	⋮
0	0
0	0

Interfered with core

Core #	0	1	2	3
0	-	0	0	0
1	0	-	0	0
2	0	0	-	0
3	0	0	0	-

Interfering core

Core 2's memory request evicts core 1's cache line from the shared cache

Evicted line's address from core 1



Tracking Cache Interference



Core 1's pollution filter

Pollution bit | Core id

0	0
0	0
⋮	⋮
1	0
0	0

Interfered with core

Core #	0	1	2	3
0	-	0	0	0
1	0	-	0	0
2	0	0	-	0
3	0	0	0	-

Interfering core

Core 2's memory request evicts core 1's cache line from the shared cache

Evicted line's address from core 1



Tracking Cache Interference



Core 1's pollution filter

Pollution bit | Core id

0	0
0	0
⋮	⋮
1	1
0	0

Interfered with core

Core #	0	1	2	3
0	-	0	0	0
1	0	-	0	0
2	0	0	-	0
3	0	0	0	-

Interfering core

0
1
2
3

Core 2's memory request evicts core 1's cache line from the shared cache

Evicted line's address from core 1



Tracking Cache Interference

Shared Cache

Hash
Function

Core 1's pollution filter

Pollution
bit

Core id

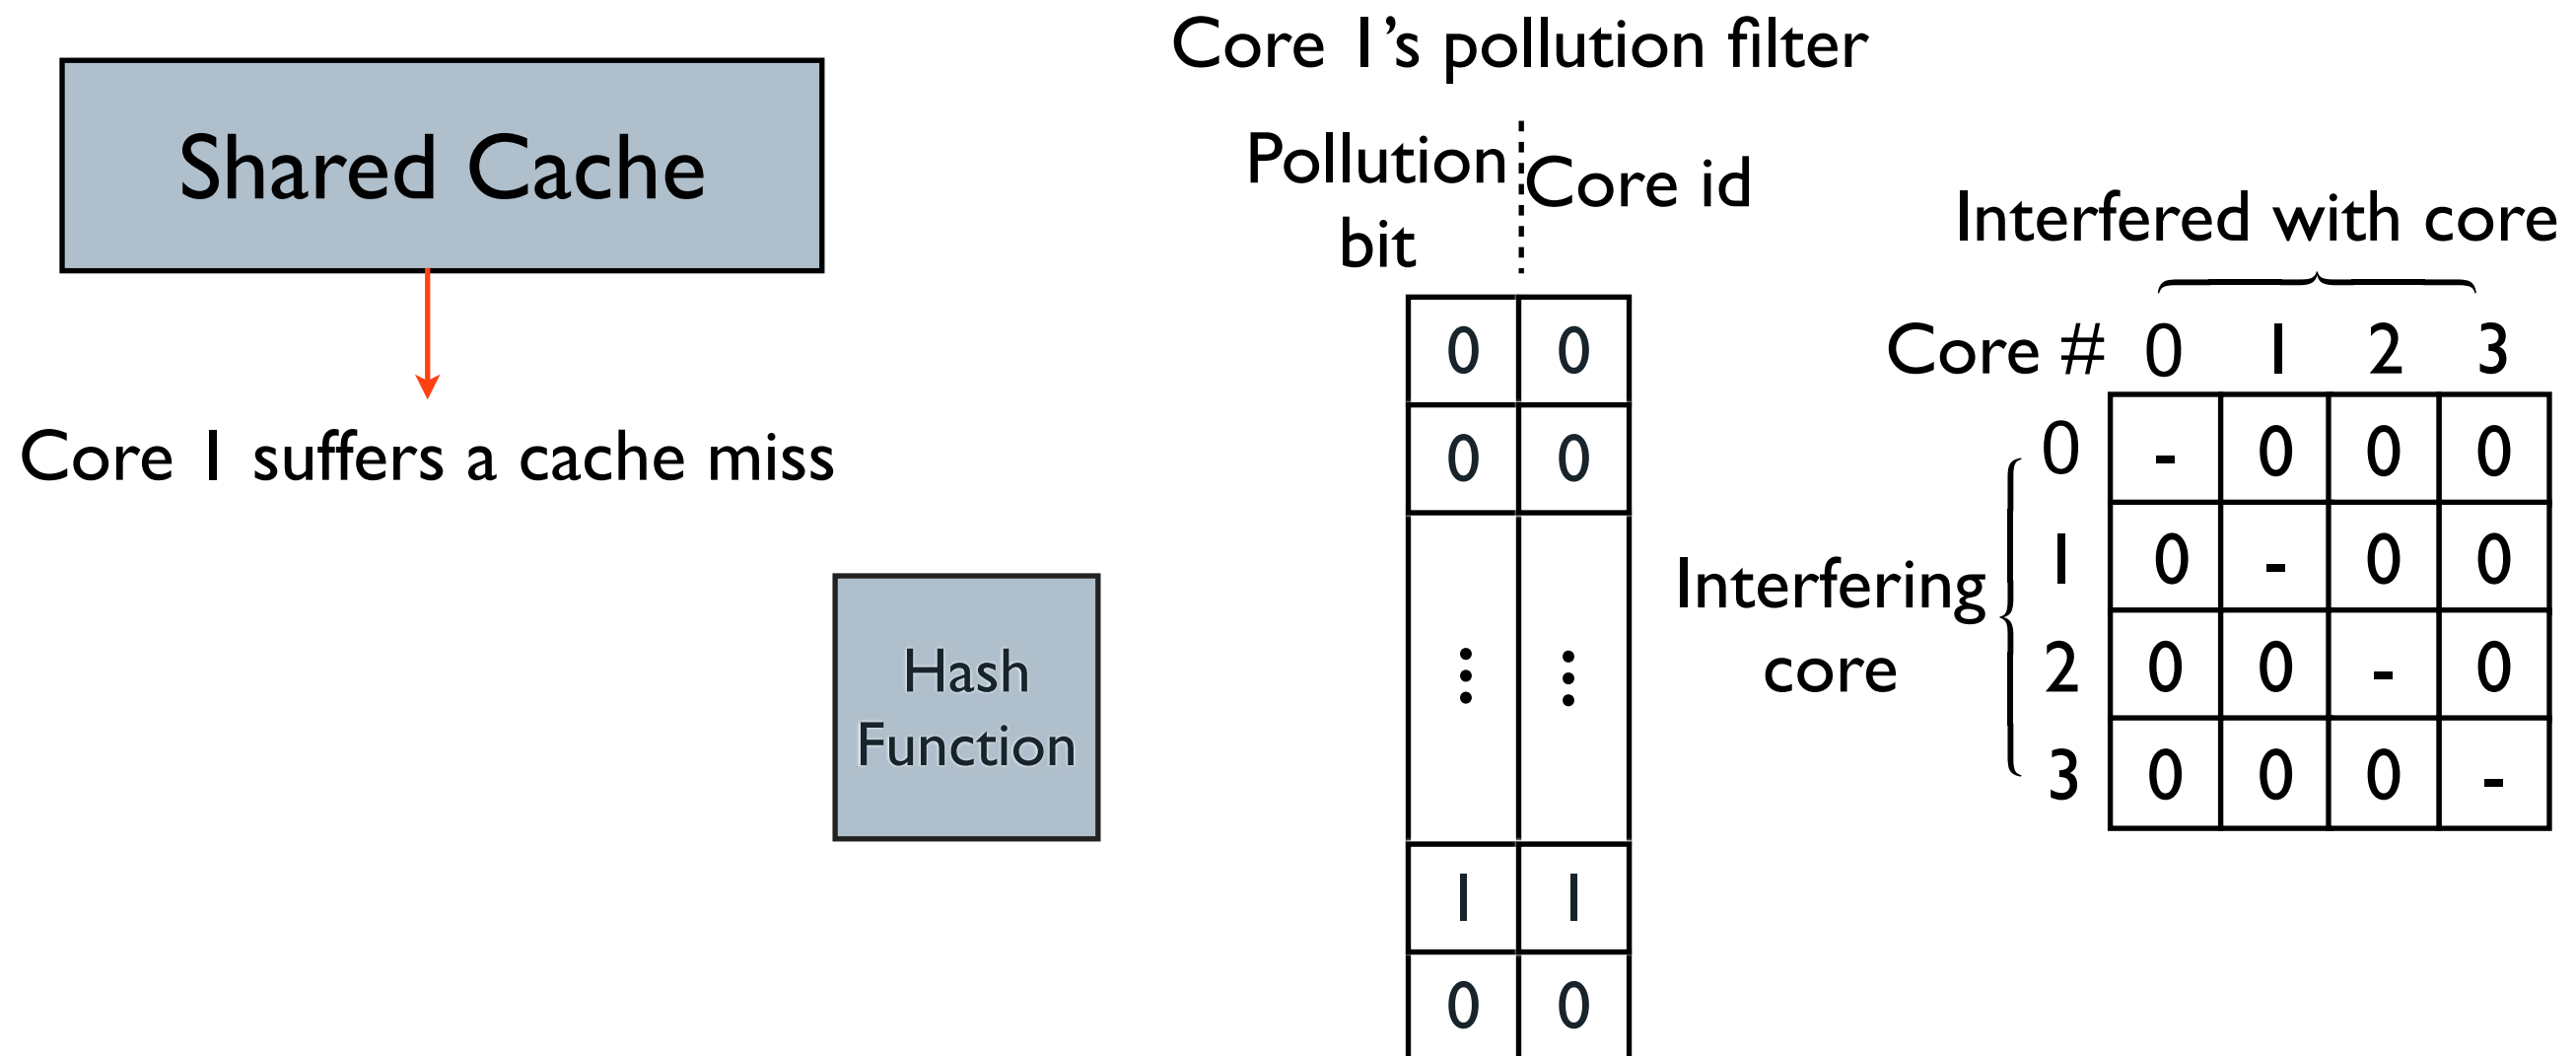
0	0
0	0
⋮	⋮
1	1
0	0

Interfered with core

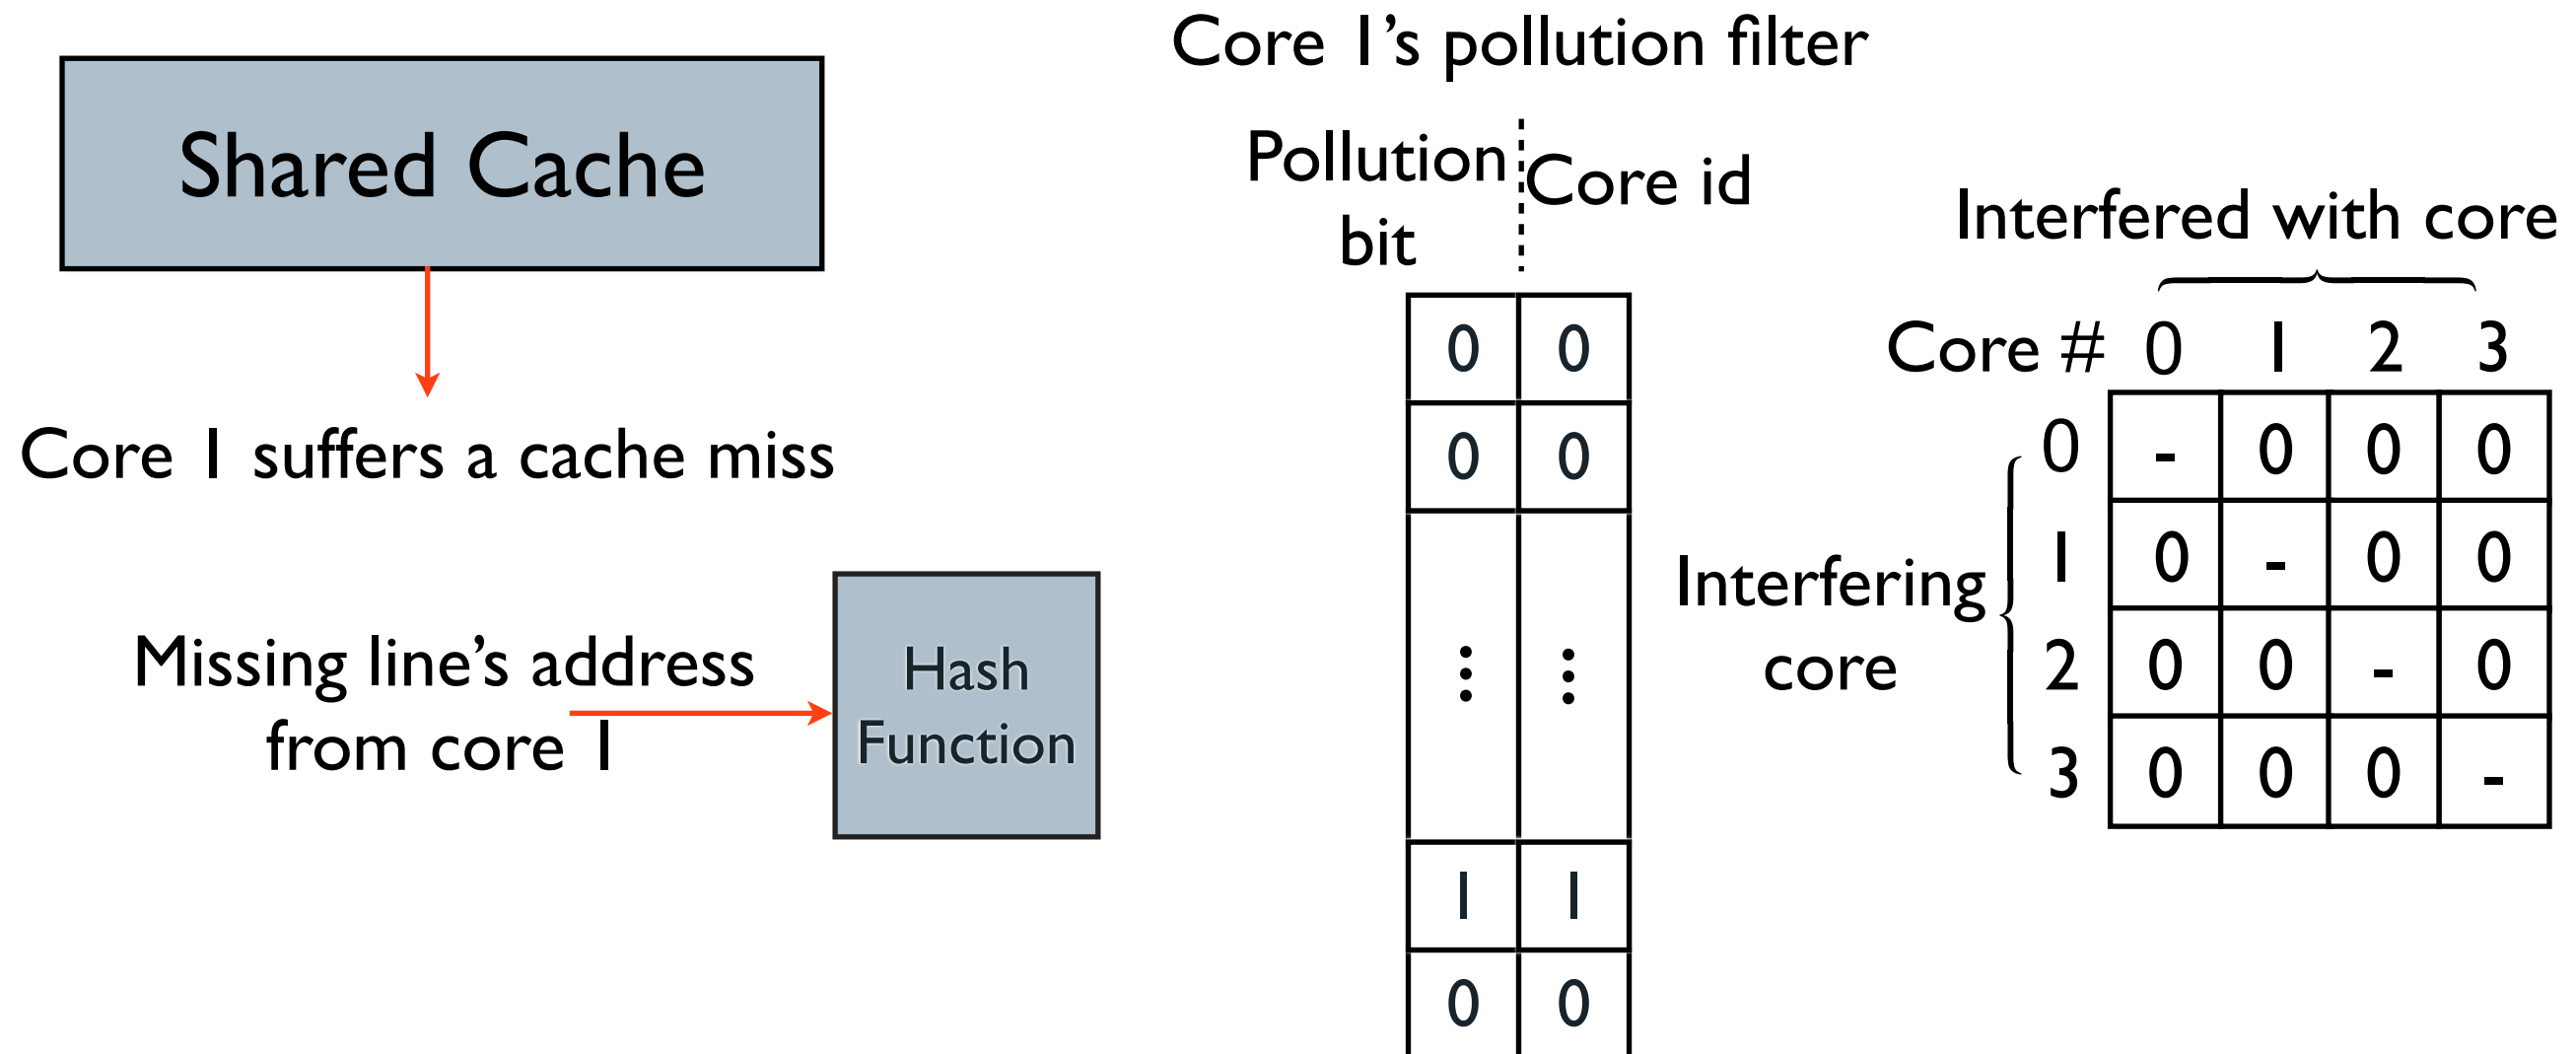
Core #	0	1	2	3
0	-	0	0	0
1	0	-	0	0
2	0	0	-	0
3	0	0	0	-

Interfering
core

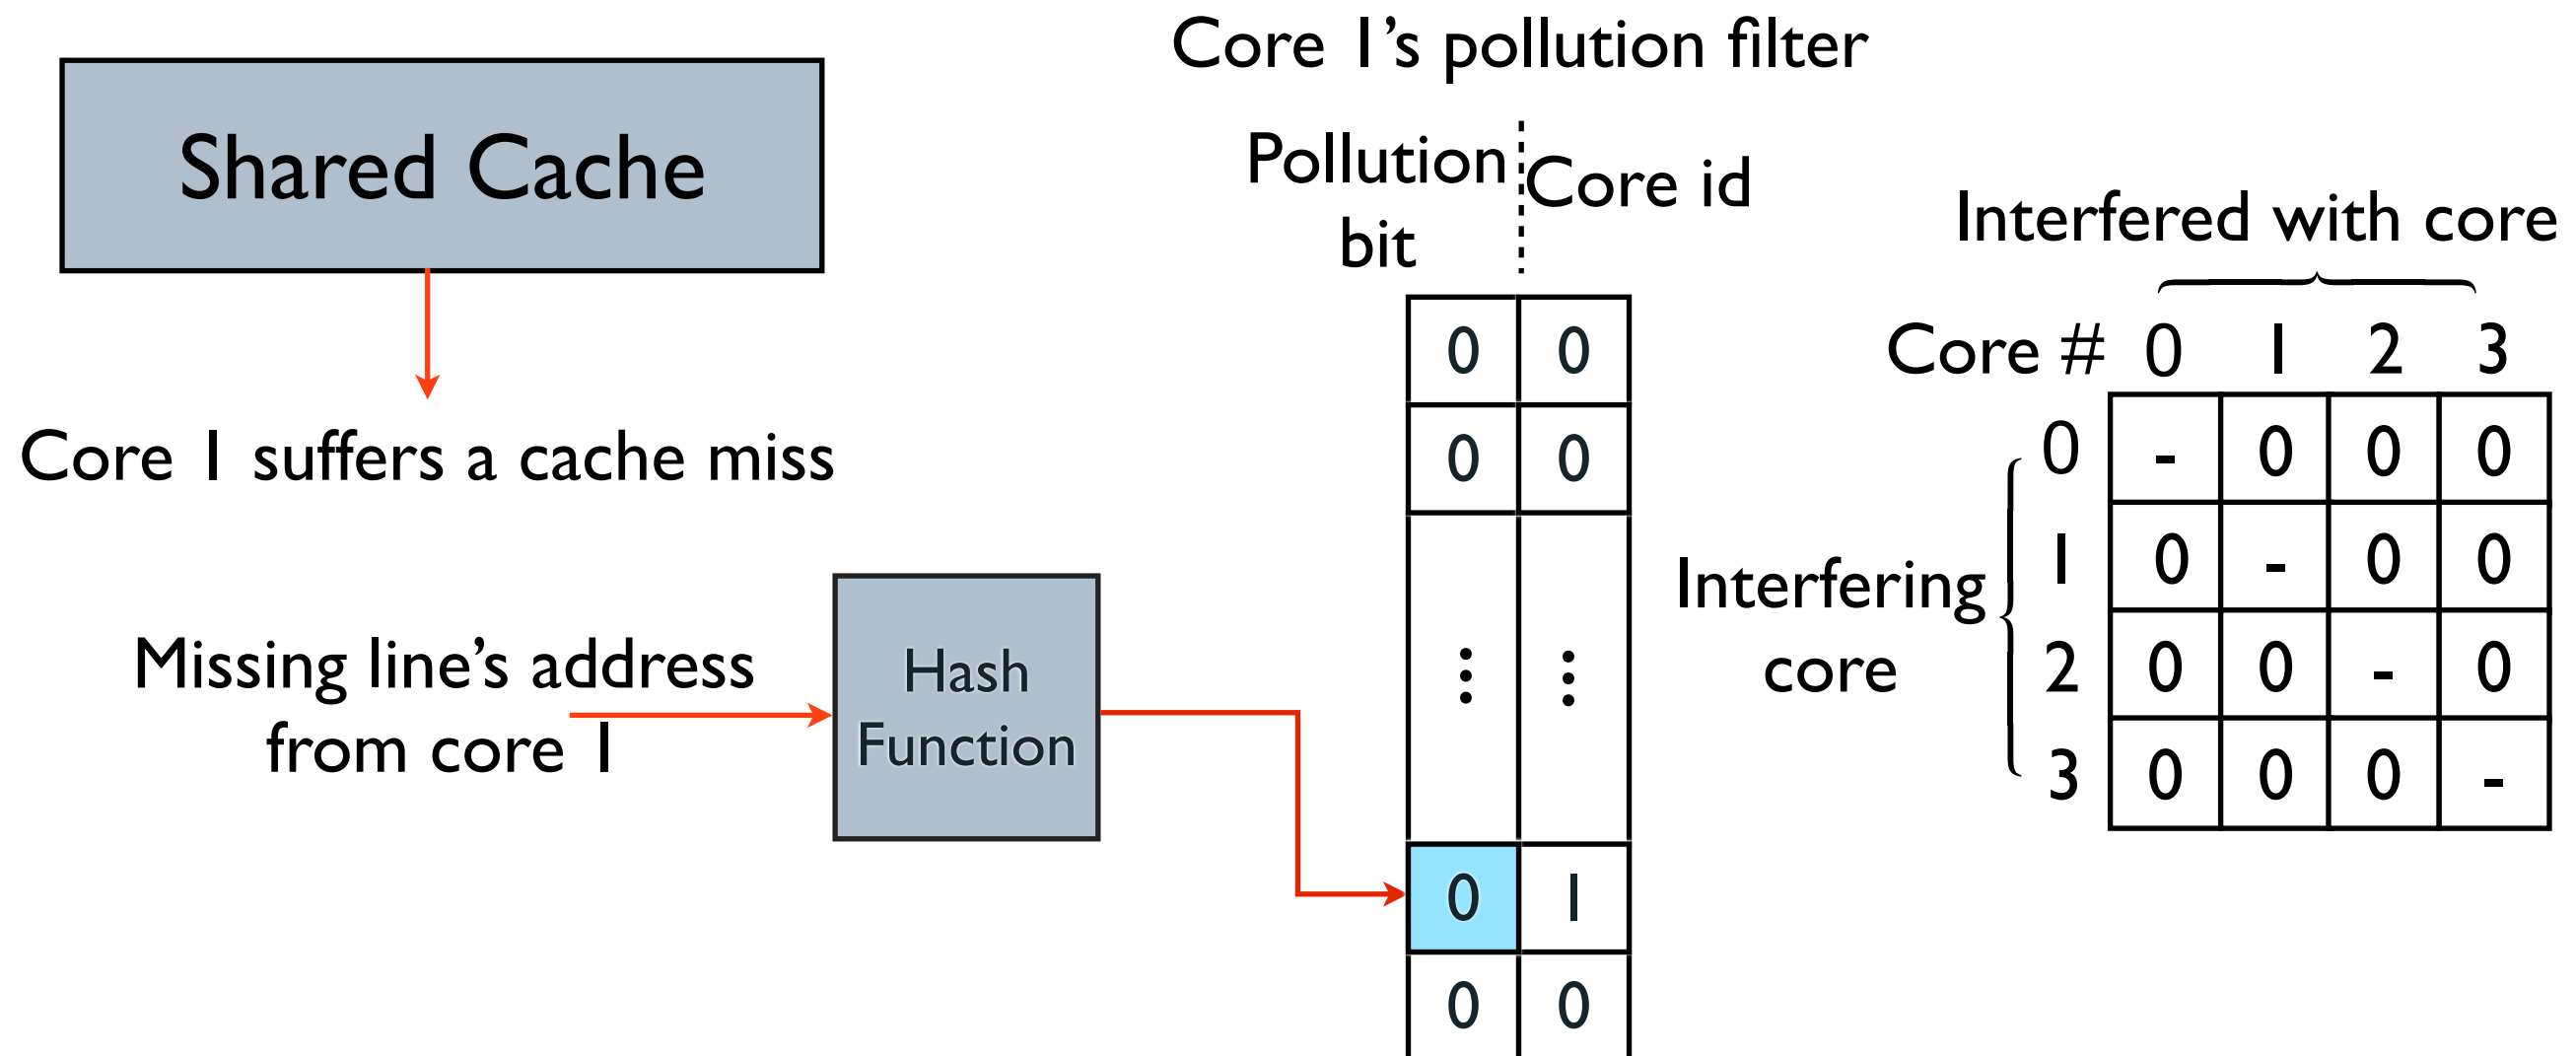
Tracking Cache Interference



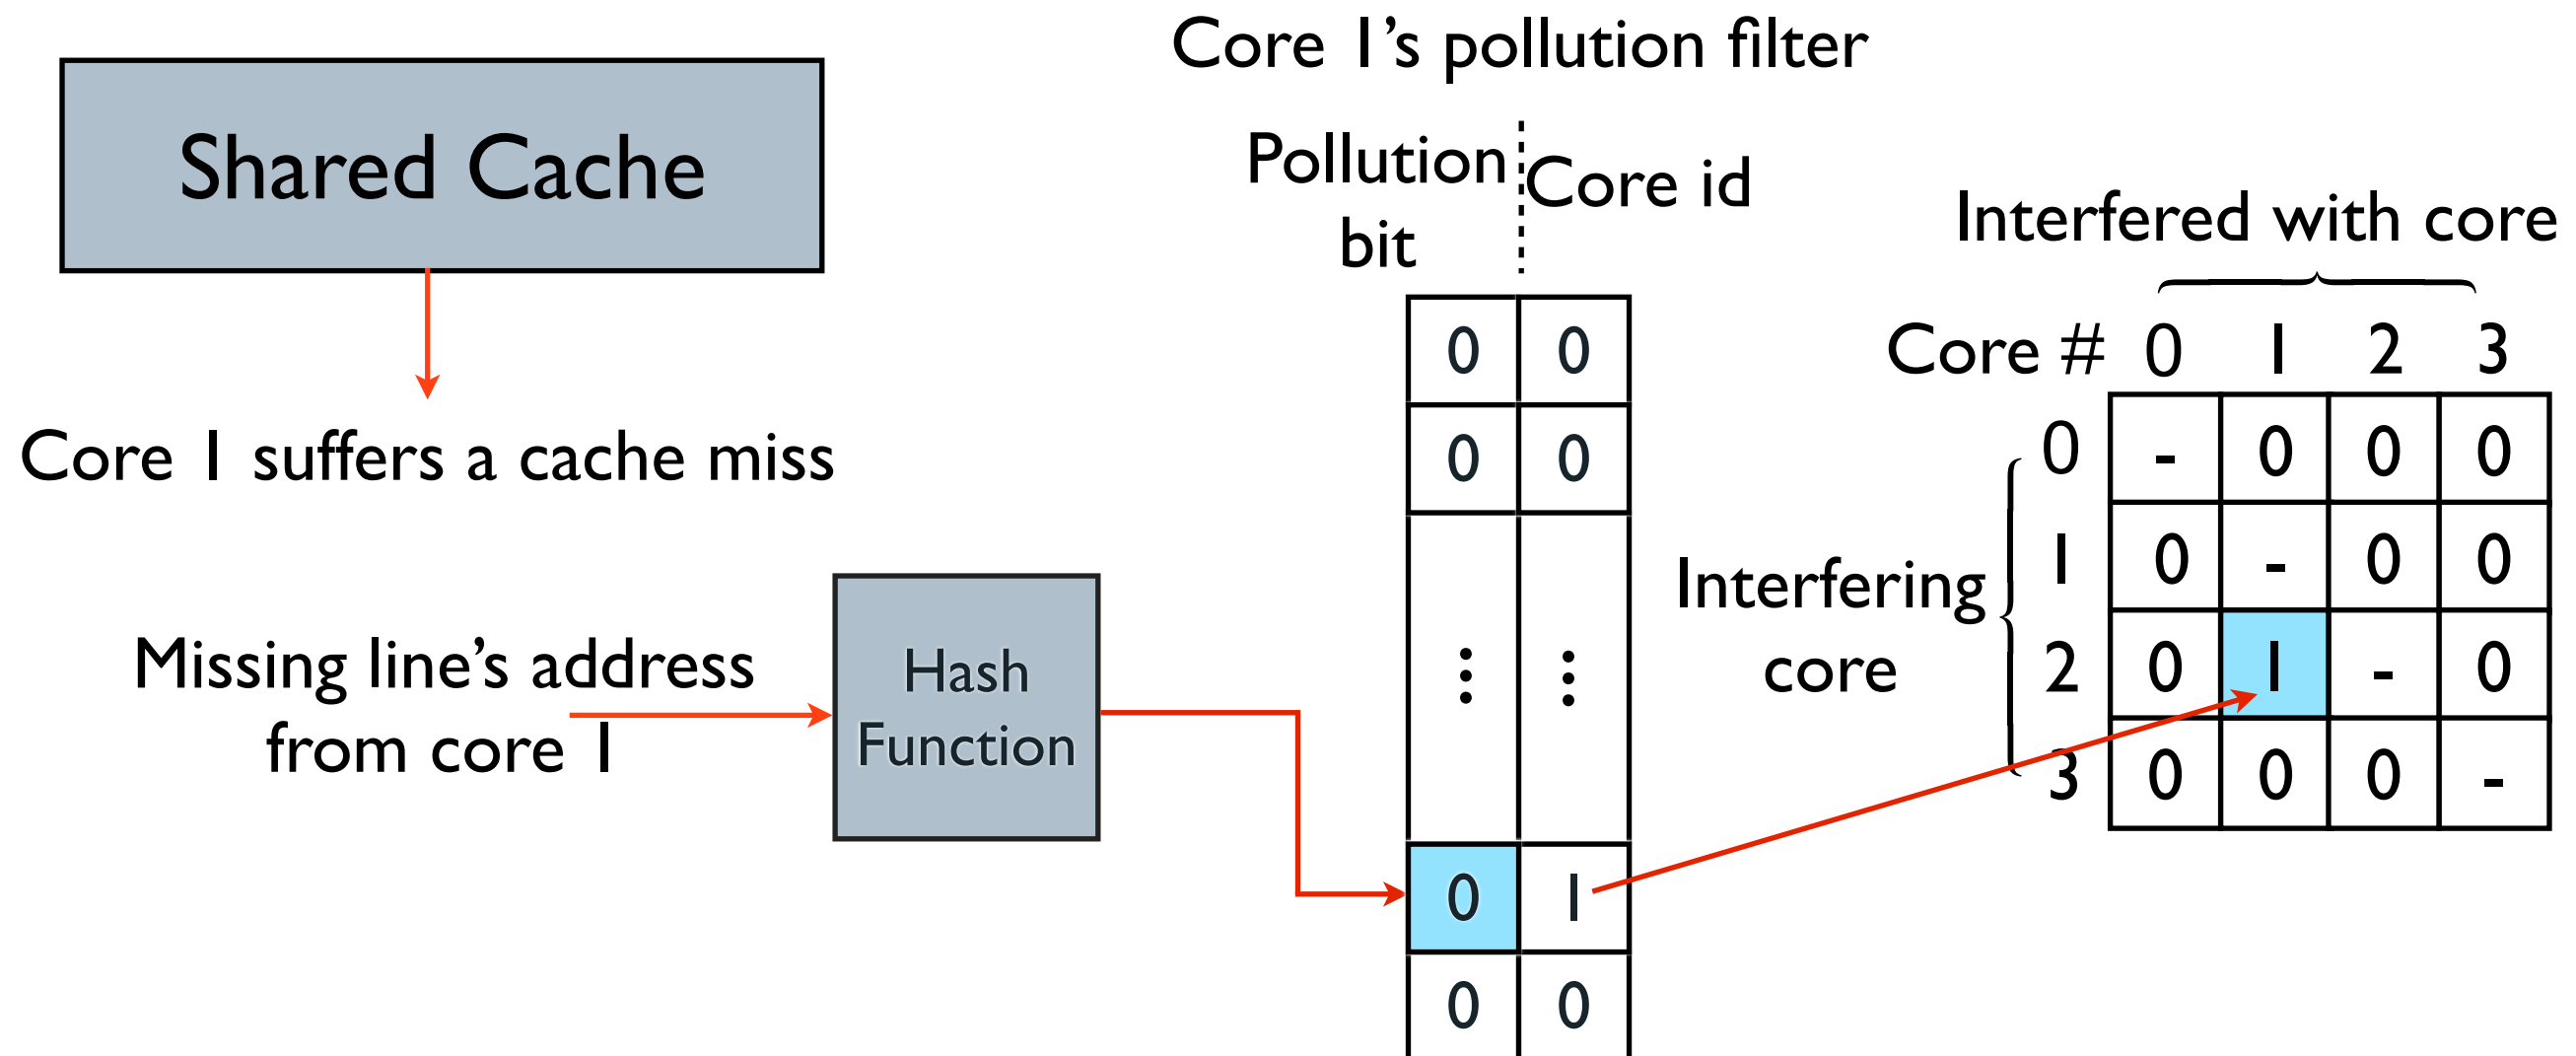
Tracking Cache Interference



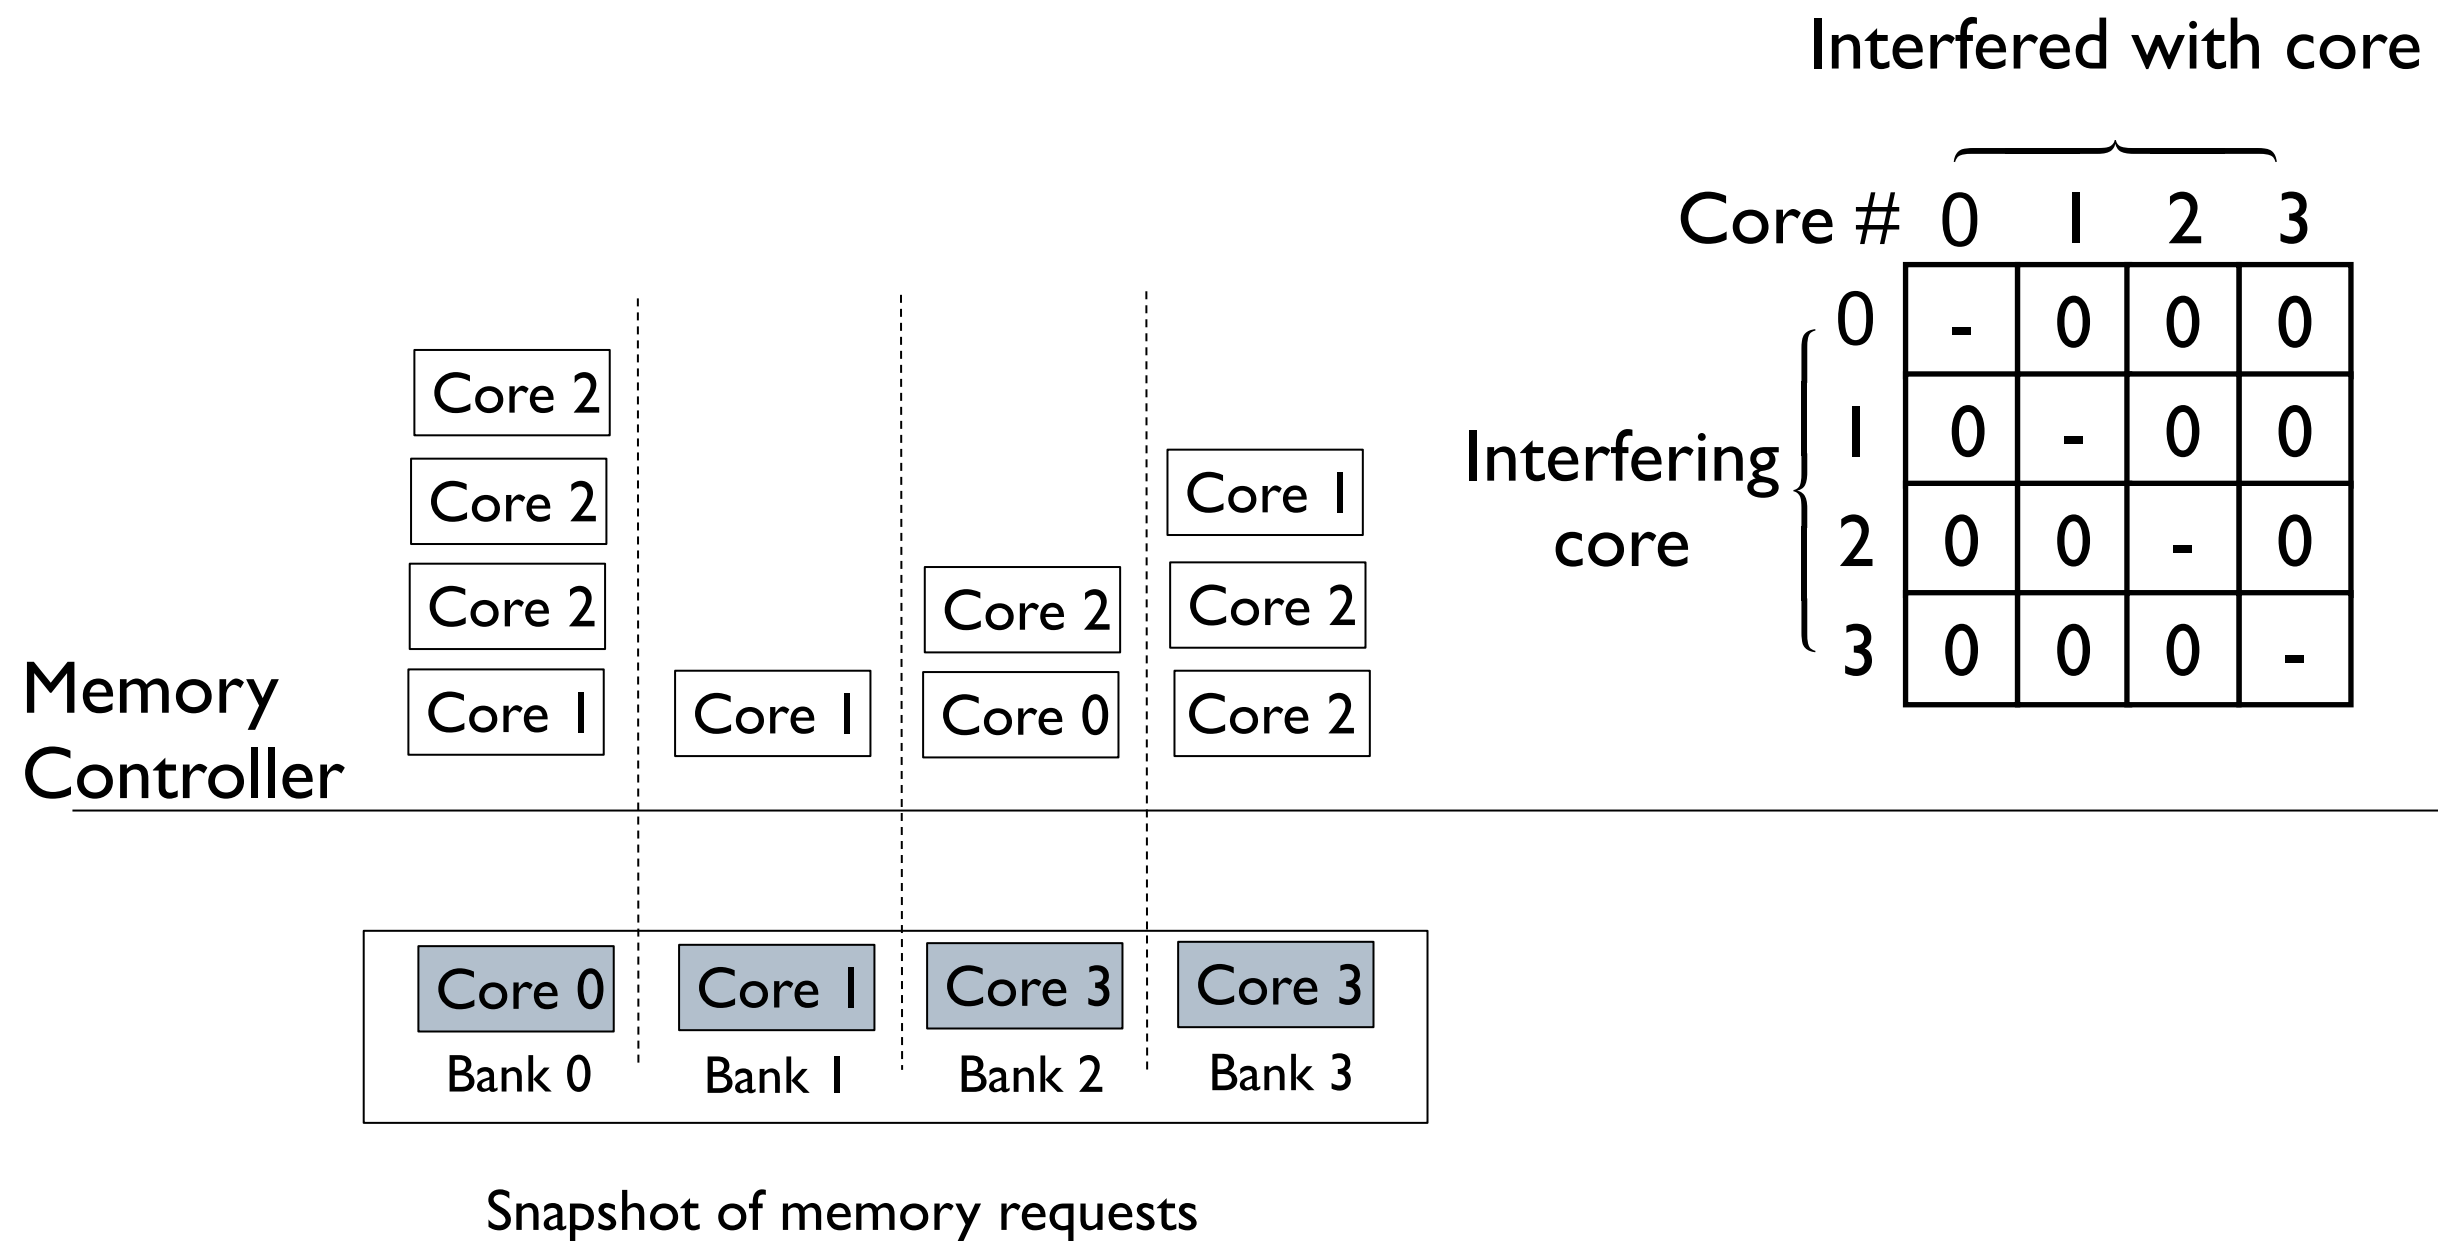
Tracking Cache Interference



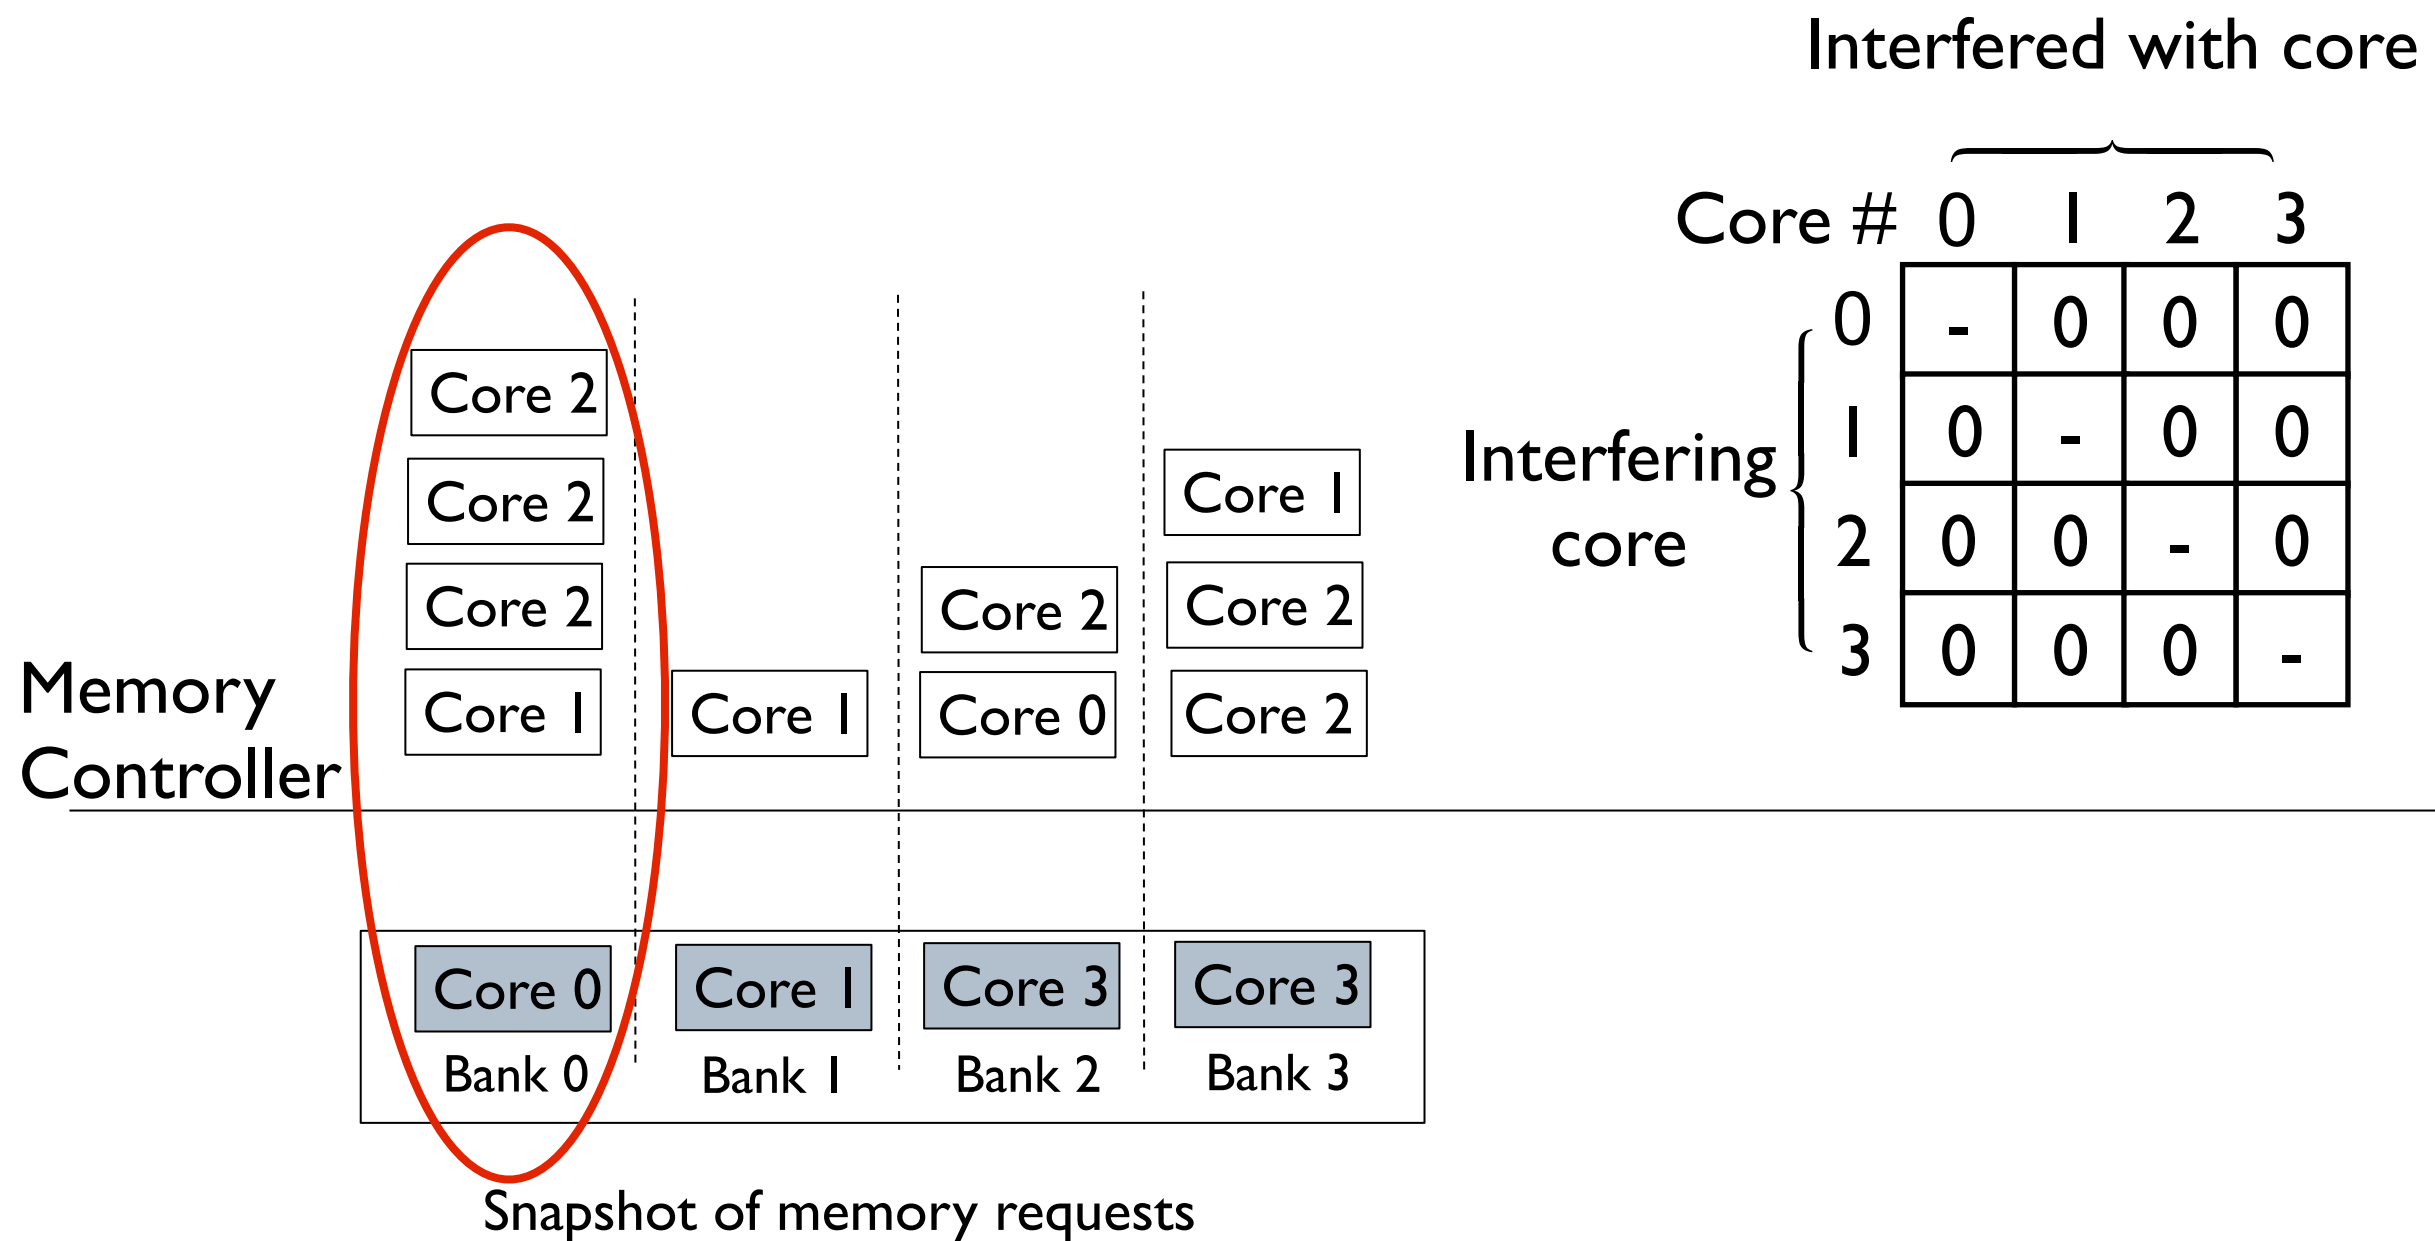
Tracking Cache Interference



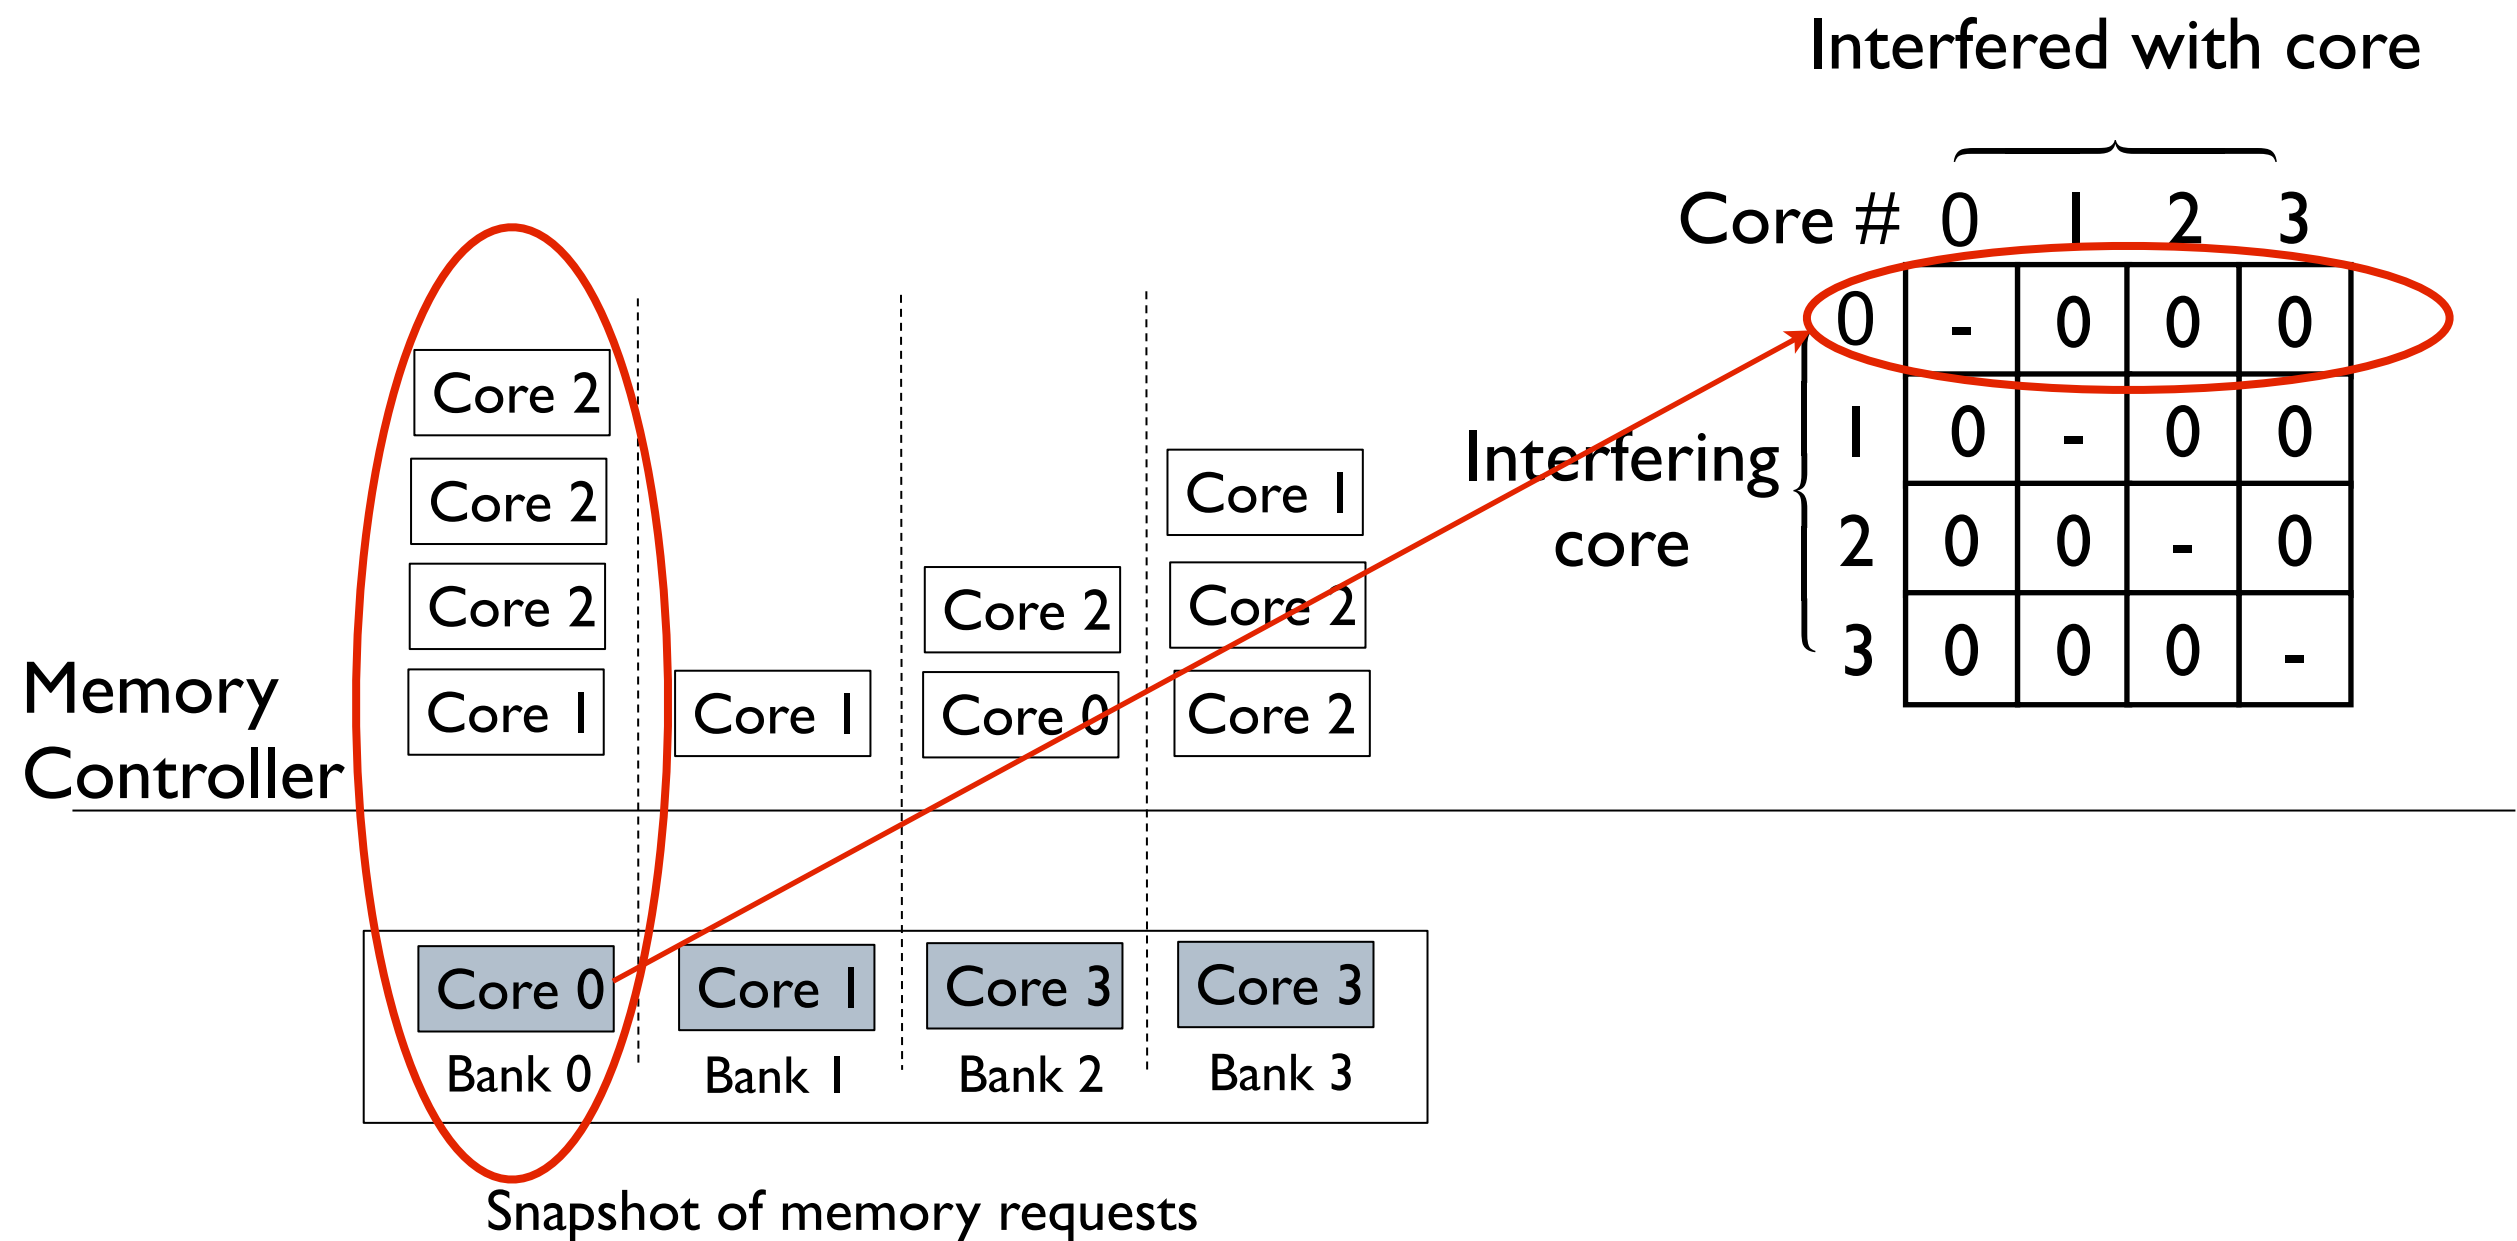
Tracking DRAM Bank Interference



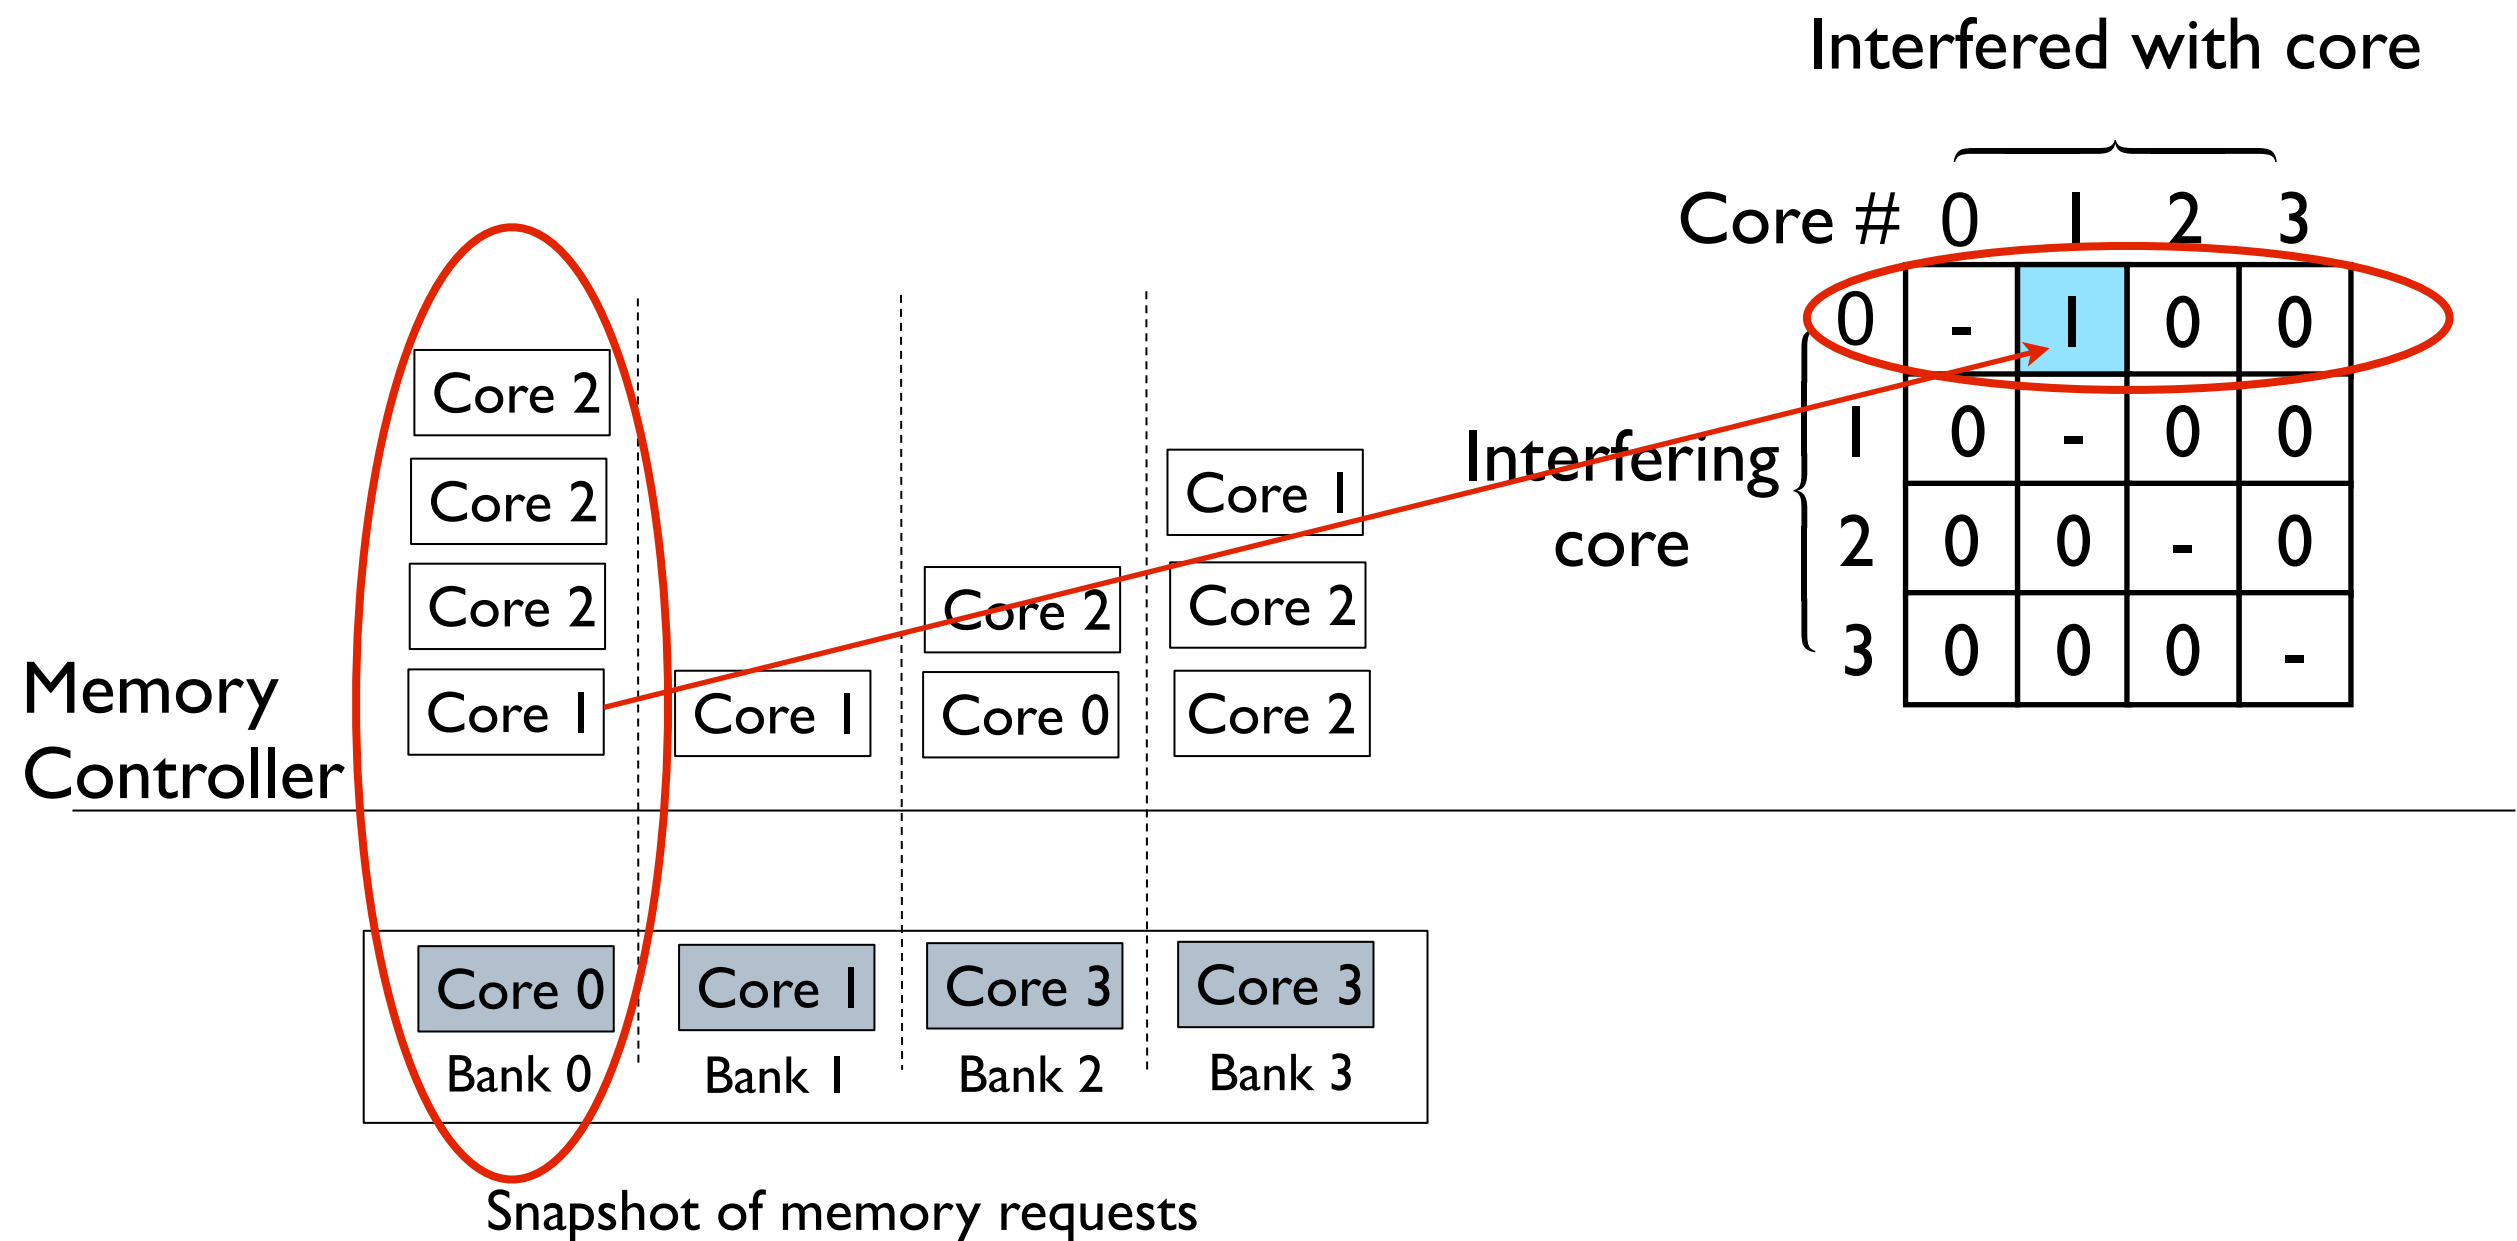
Tracking DRAM Bank Interference



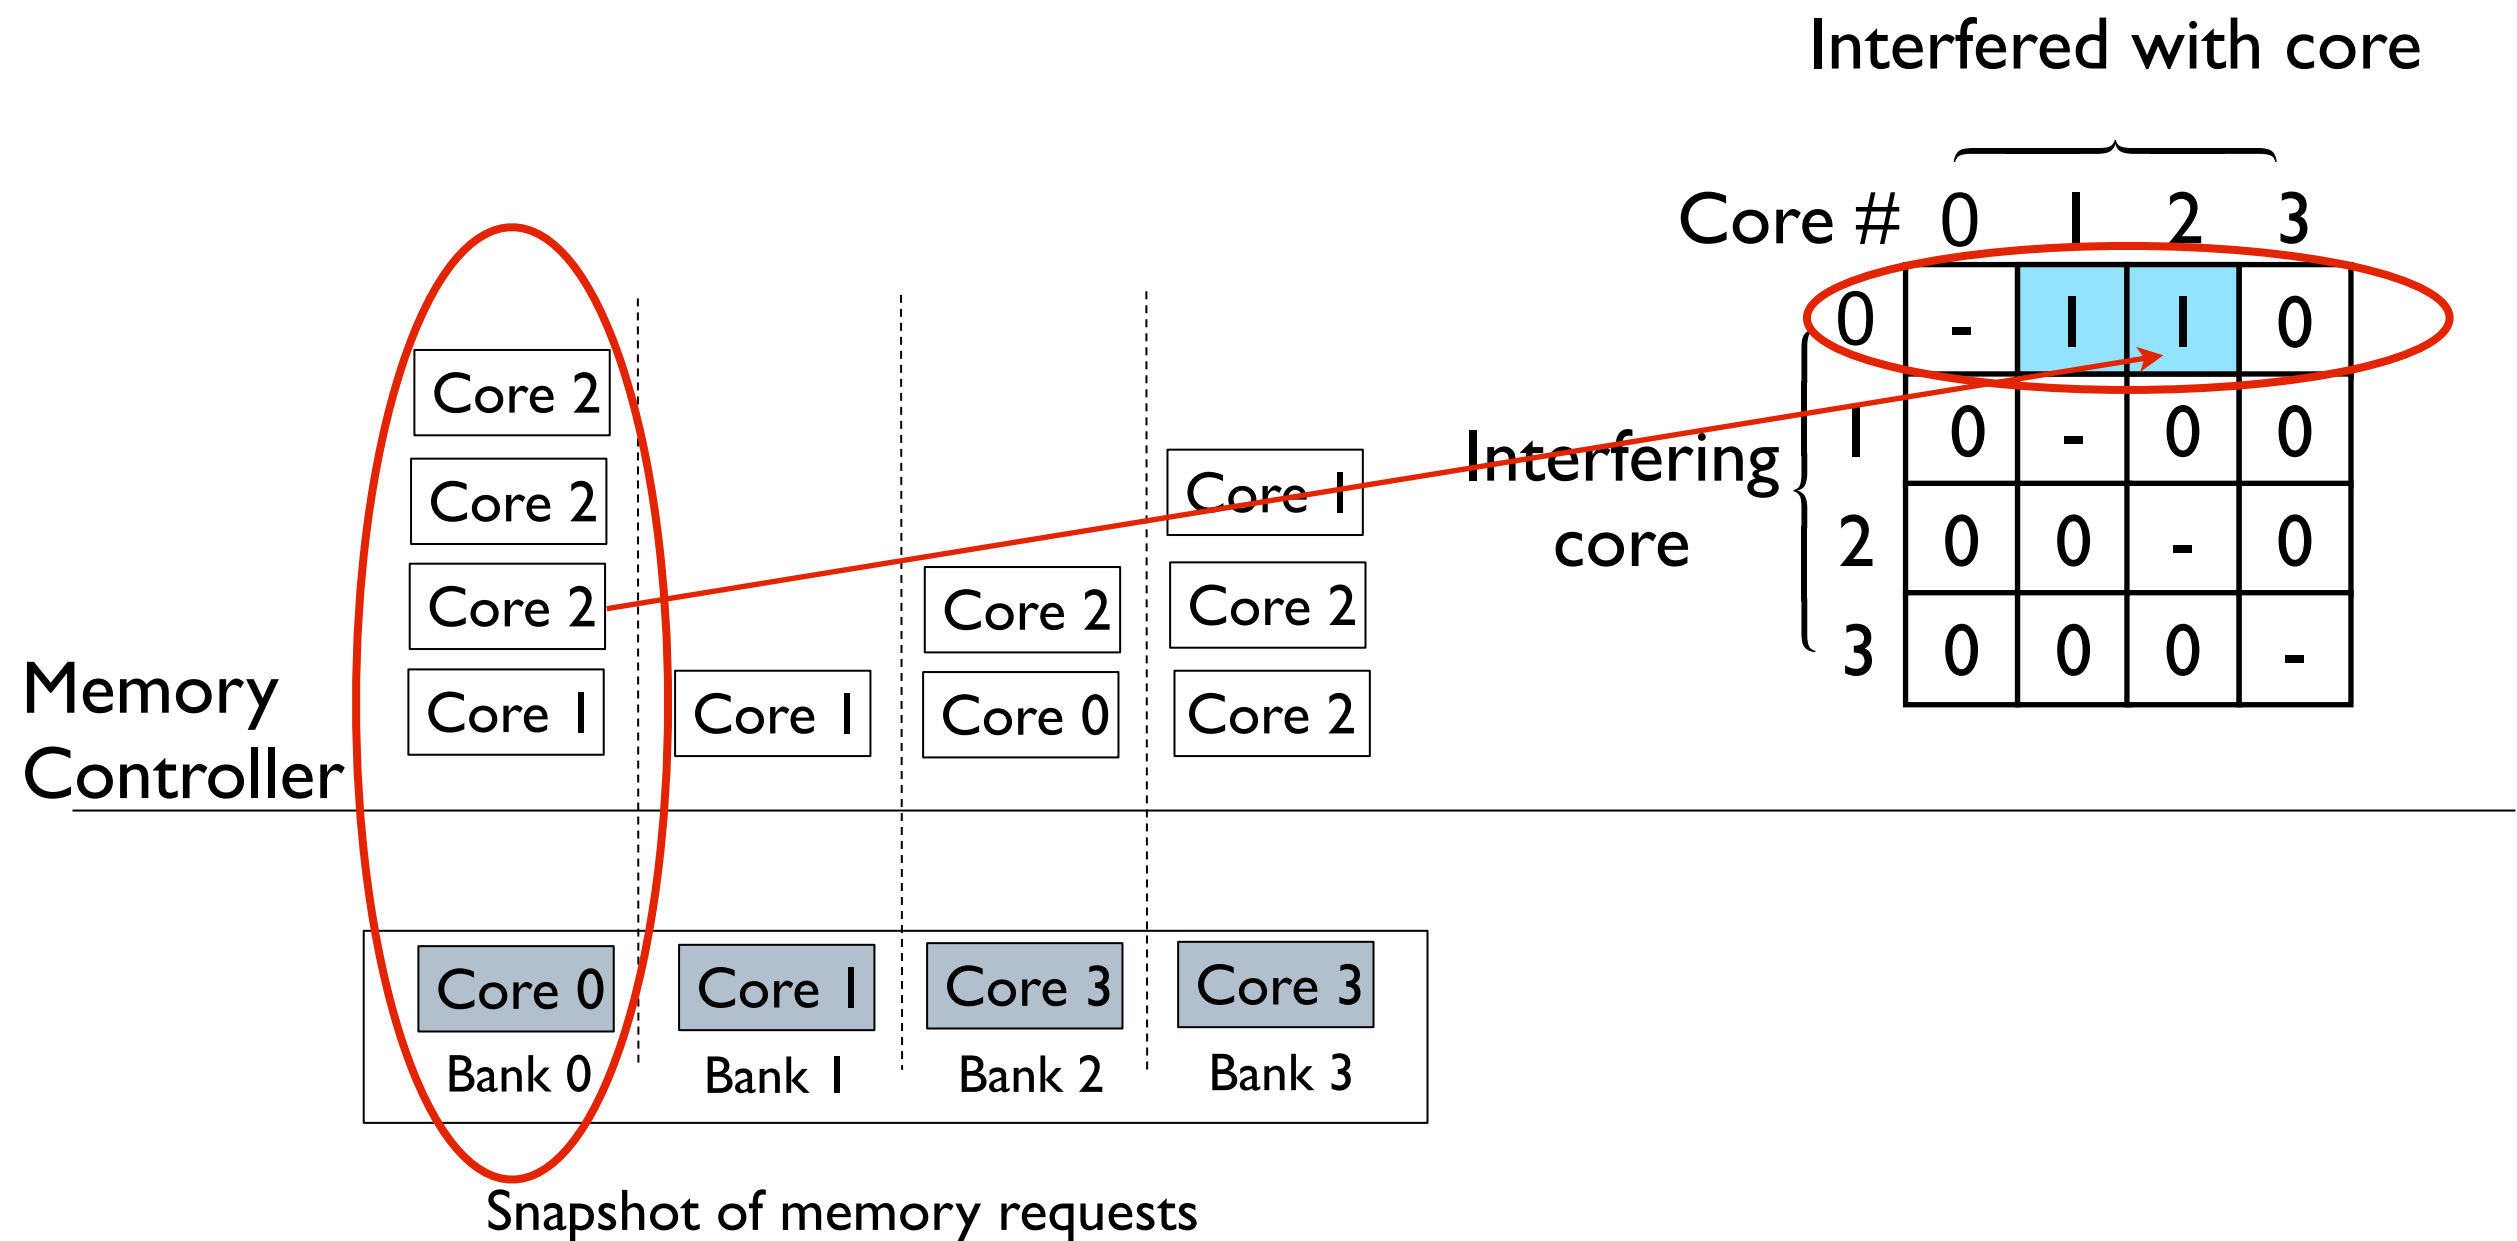
Tracking DRAM Bank Interference



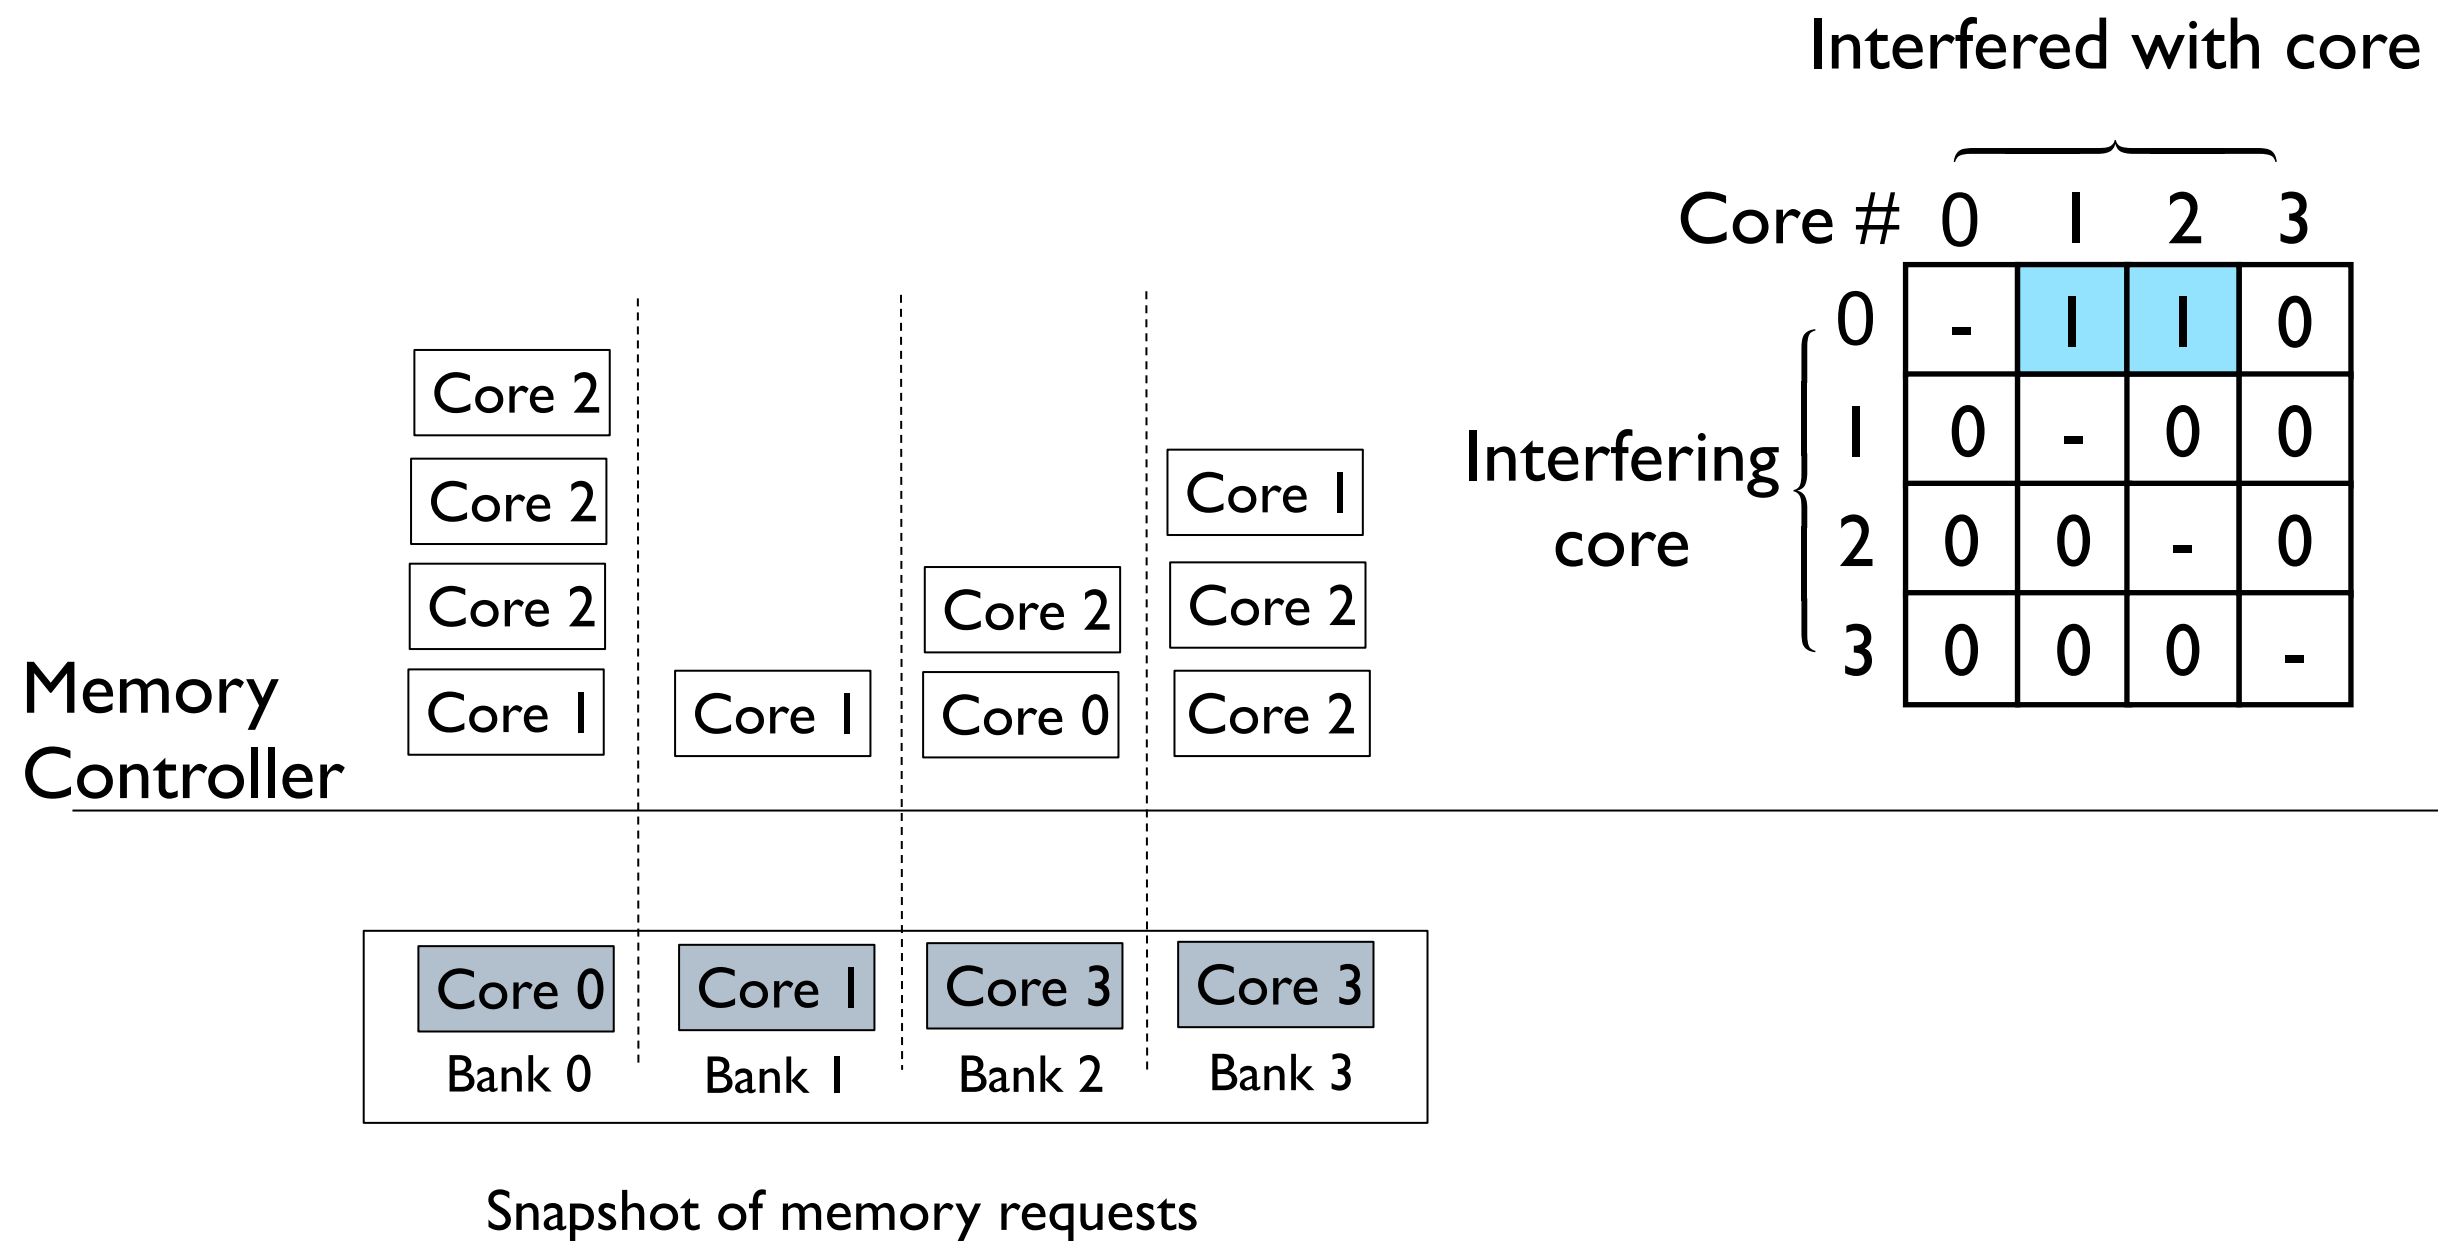
Tracking DRAM Bank Interference



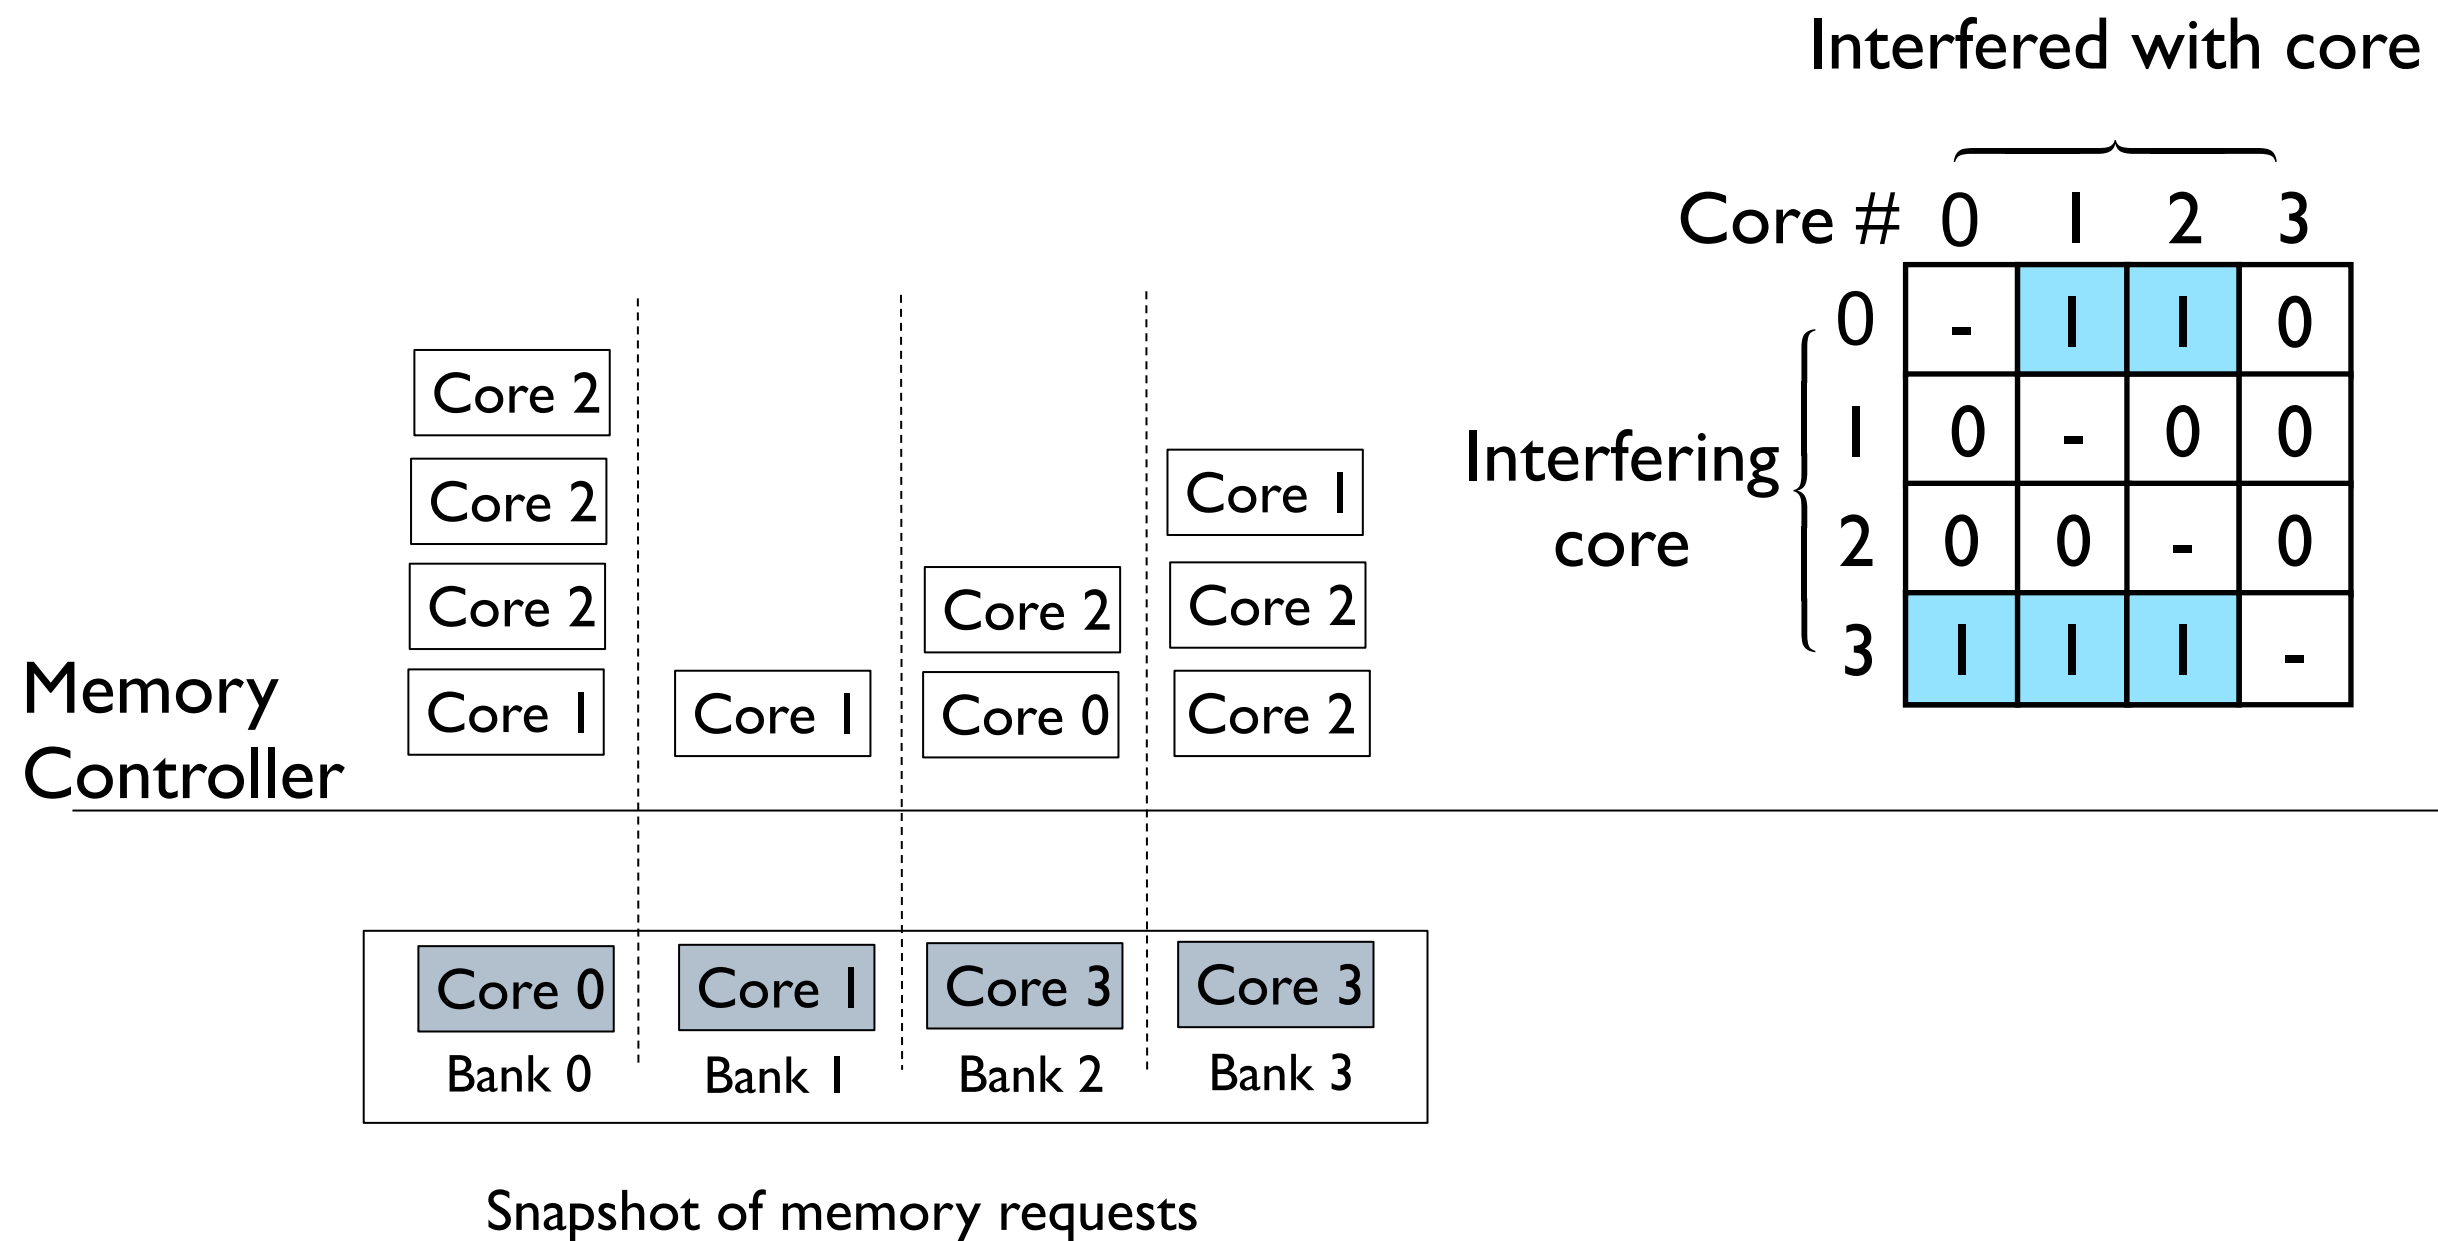
Tracking DRAM Bank Interference



Tracking DRAM Bank Interference




Tracking DRAM Bank Interference




Results For Alternative System Unfairness Metric


Deviation of Variation Normalized to No Fairness




Results of Different Throttling Mechanisms




Results With 2 Memory Channels



Support For Constraining Max Slowdown



Support For Thread Priorities



Case Study

