ISCA' 21

IIT Kanpur

Don't Forget the I/O When Allocating Your LLC

Yifan Yuan, Mohammad Alian, Yipeng Wang, Ren Wang, Ilia Kurakin, Charlie Tai, Nam Sung Kim

Shiv Bhushan Tripathi

CDOS Talk
04th July 2022

Resource Sharing in Cloud

Resource Sharing in Cloud

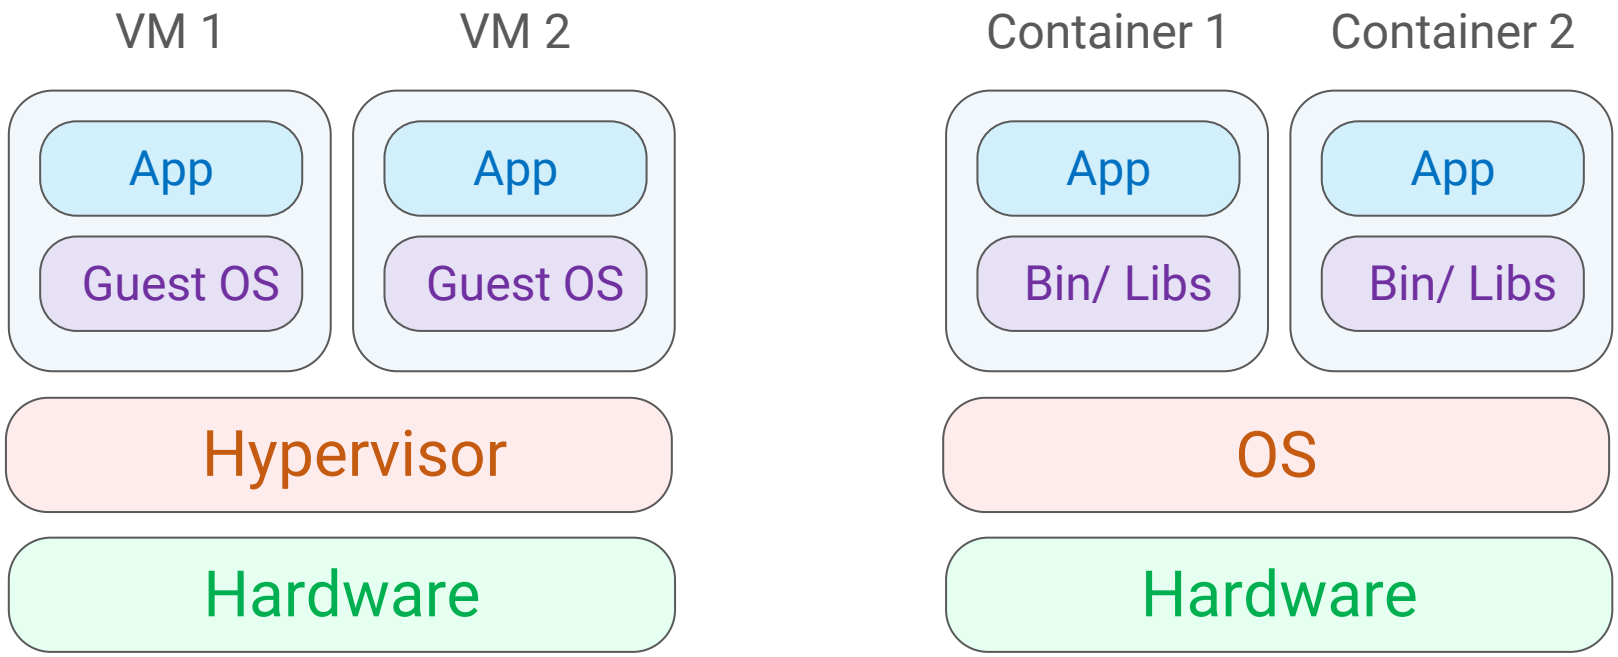

Resource Sharing in Cloud

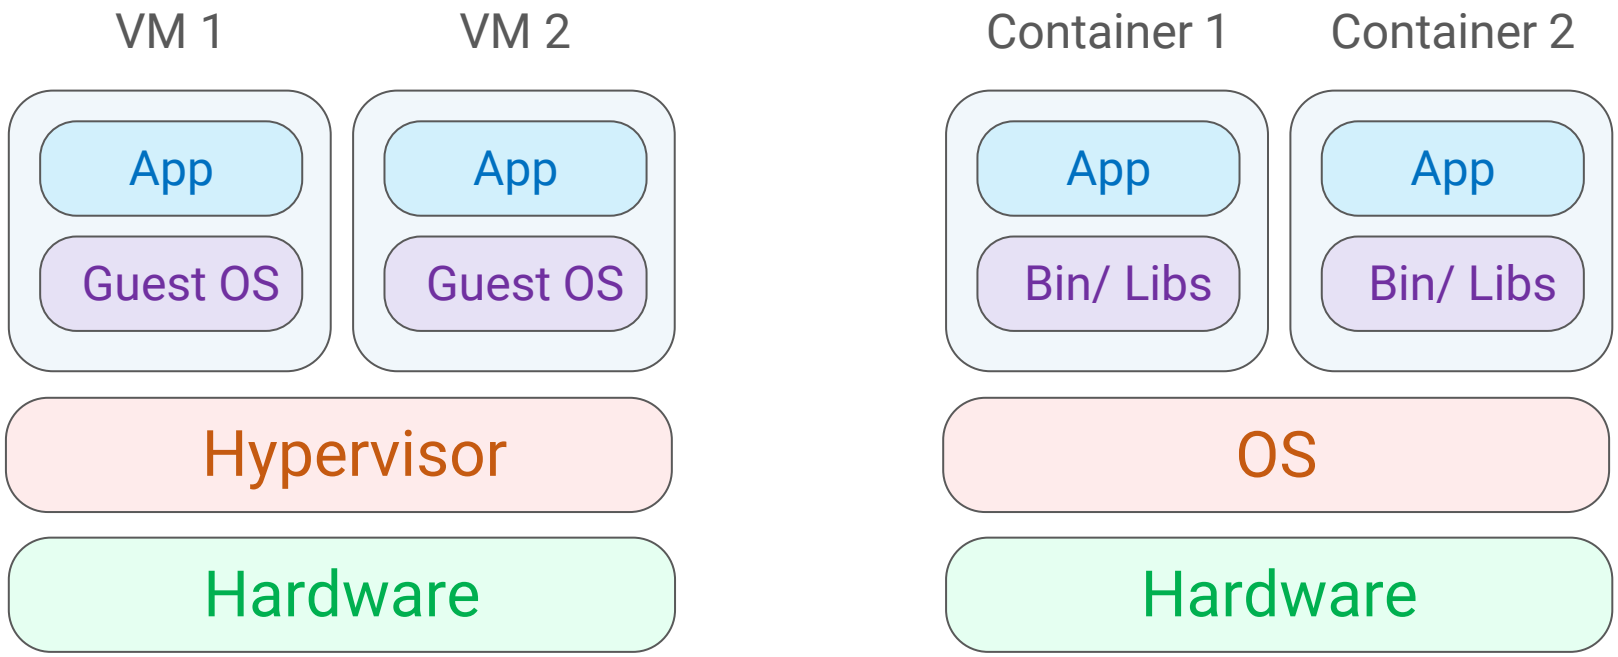

- Shared hardware resources for better utilization, e.g. Core, Memory, LLC, etc.

Resource Sharing in Cloud

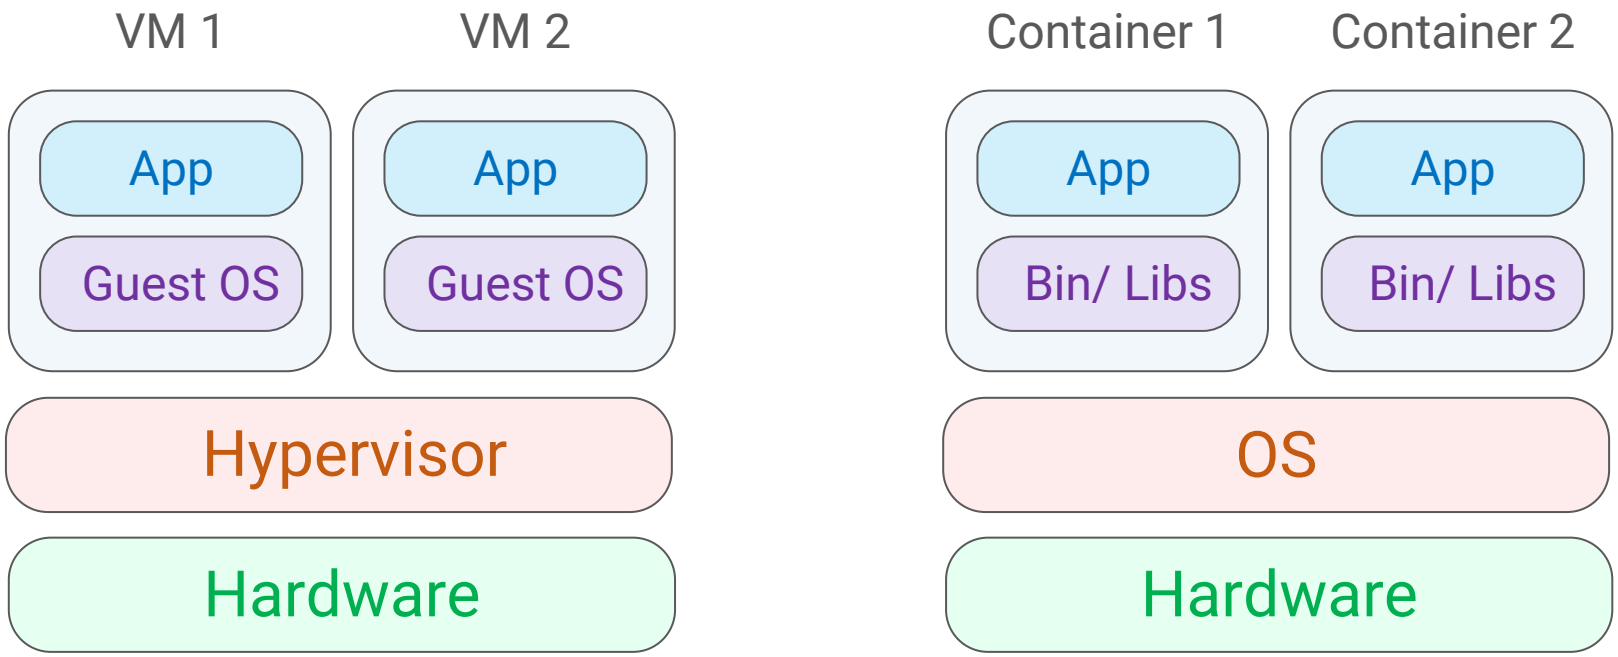

- Shared hardware resources for better utilization, e.g. Core, Memory, LLC, etc.
- Leads to interference: Causing performance degradations for tenants

Resource Sharing in Cloud

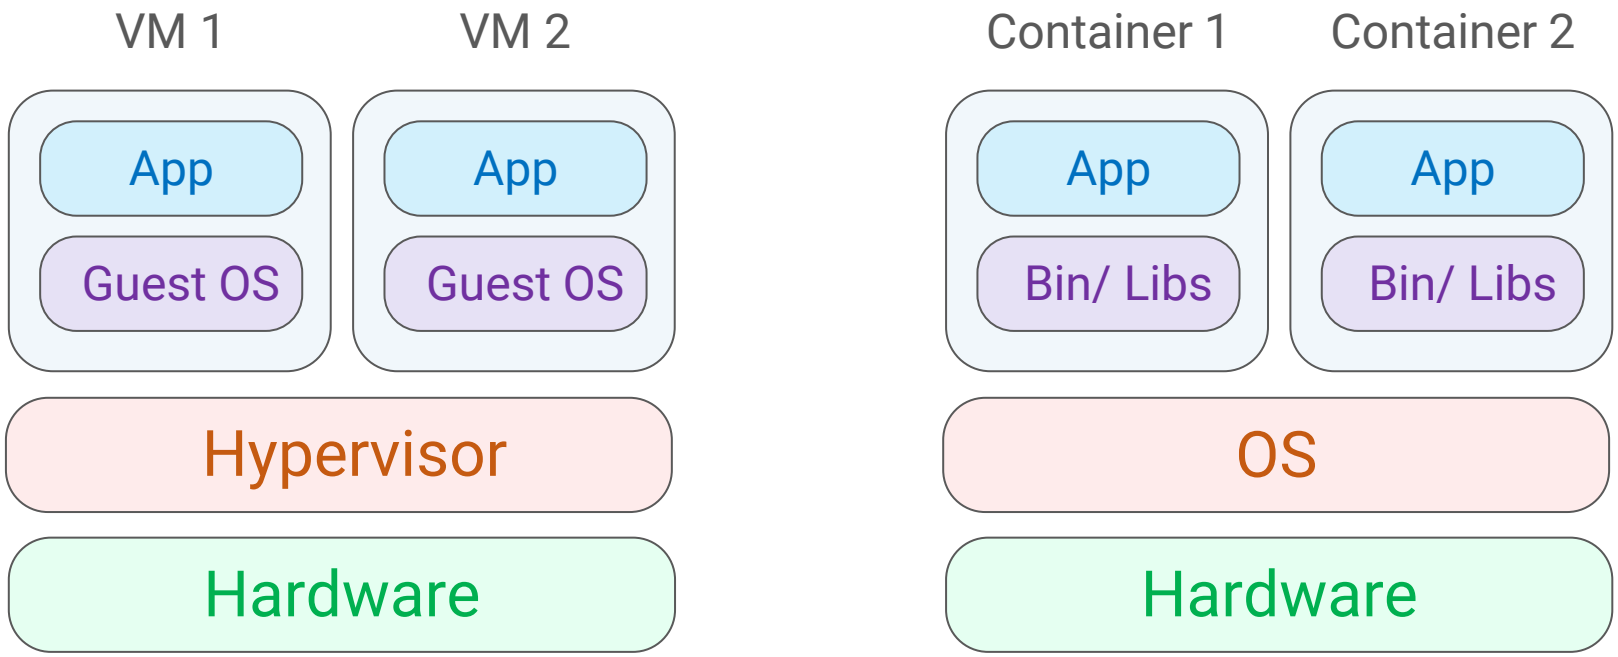

- Shared hardware resources for better utilization, e.g. Core, Memory, LLC, etc.
- Leads to interference: Causing performance degradations for tenants
- Need to allocate and isolate for different tenants: the focus is LLC

Resource Sharing in Cloud

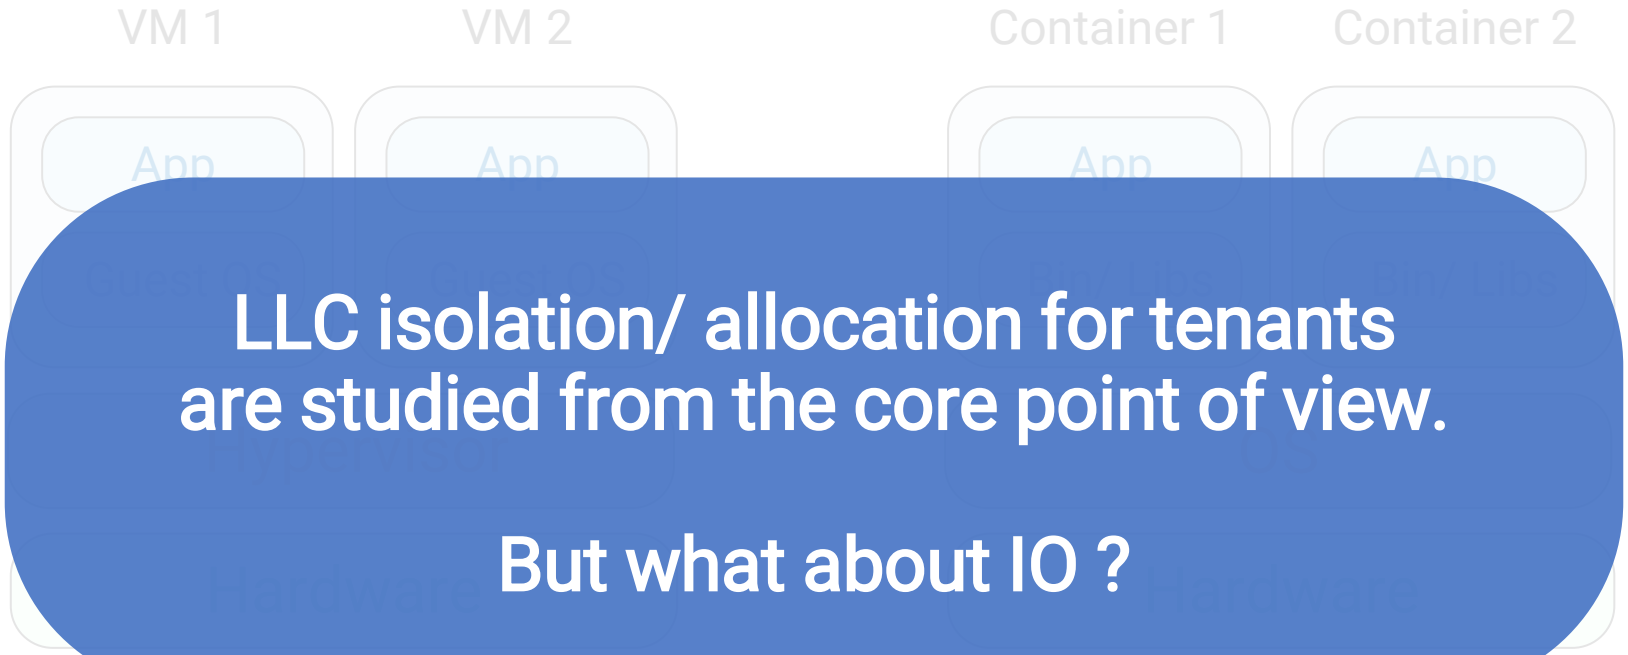

LLC isolation/ allocation for tenants are studied from the core point of view.

But what about IO ?

- Shared hardware resources for better utilization, e.g. Core, Memory, LLC, etc.
- Leads to interference: Causing performance degradations for tenants
- Need to allocate and isolate for different tenants: the focus is LLC

LLC Management and IO

LLC Management and IO

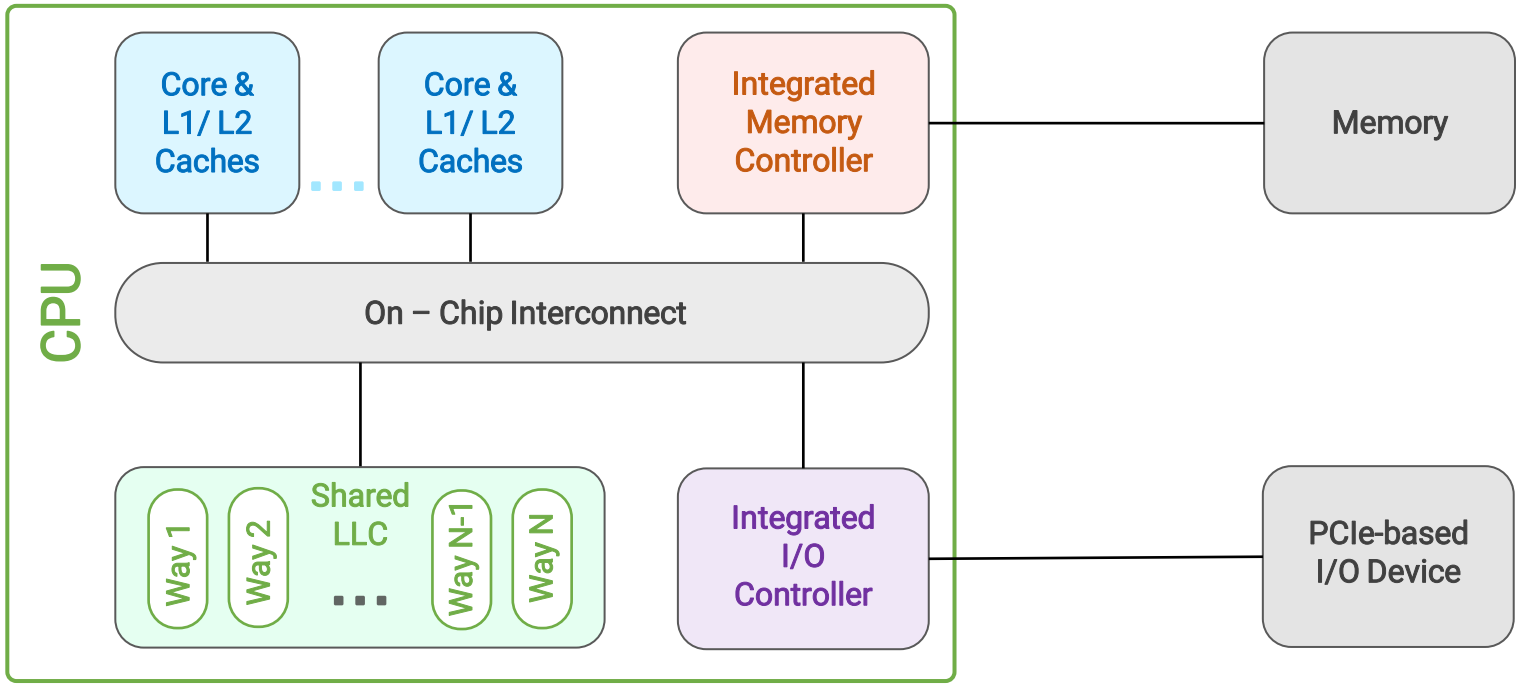

LLC Management and IO

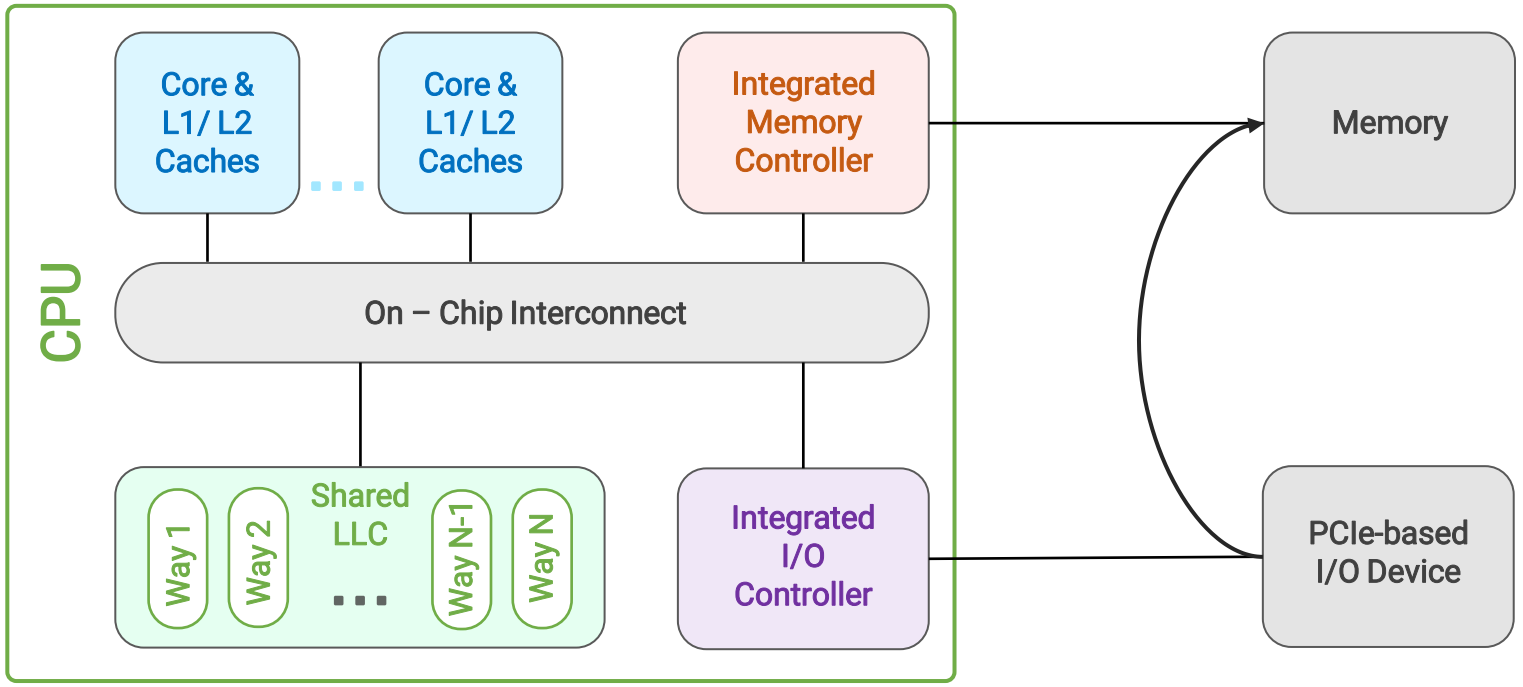

- Memory involvement for IO (DMA)

LLC Management and IO

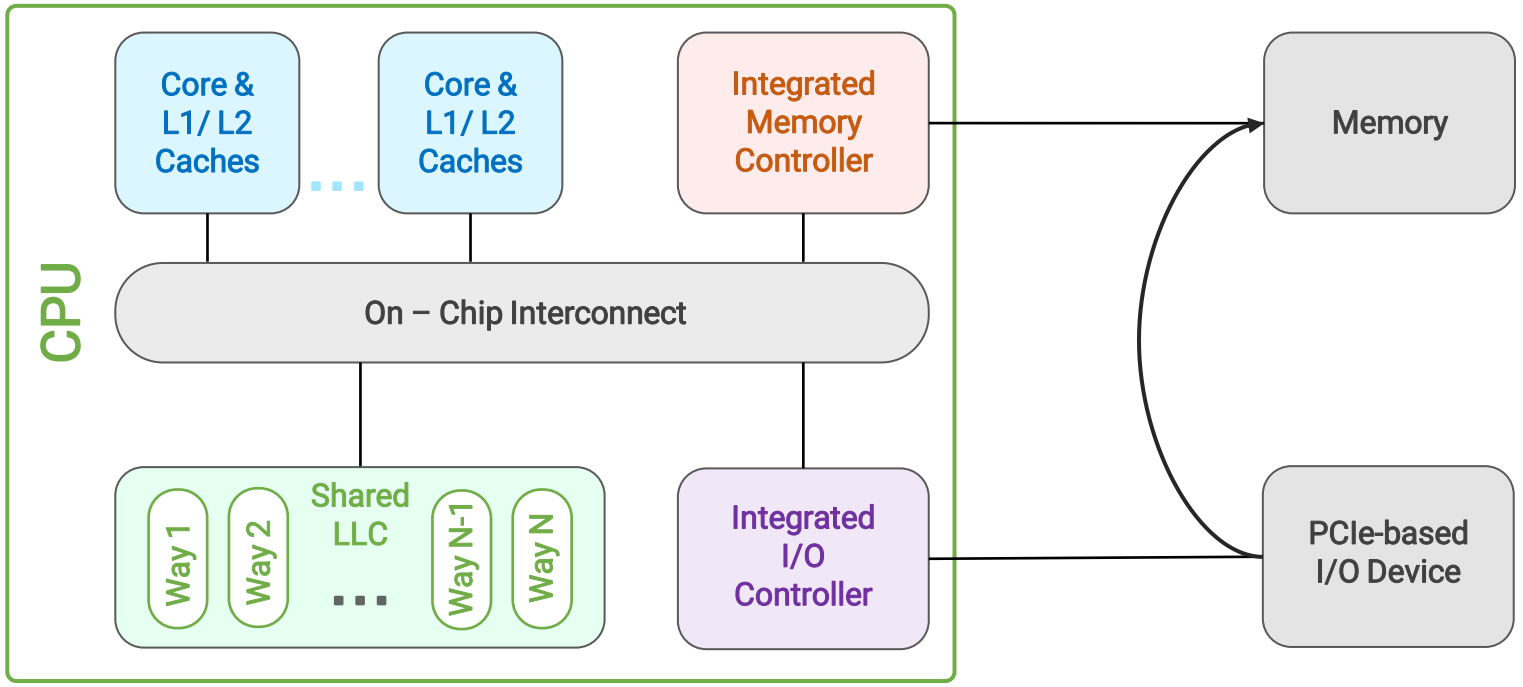

- Memory involvement for IO (DMA)
- Problems: High Speed Devices (NIC, SSDs)

LLC Management and IO

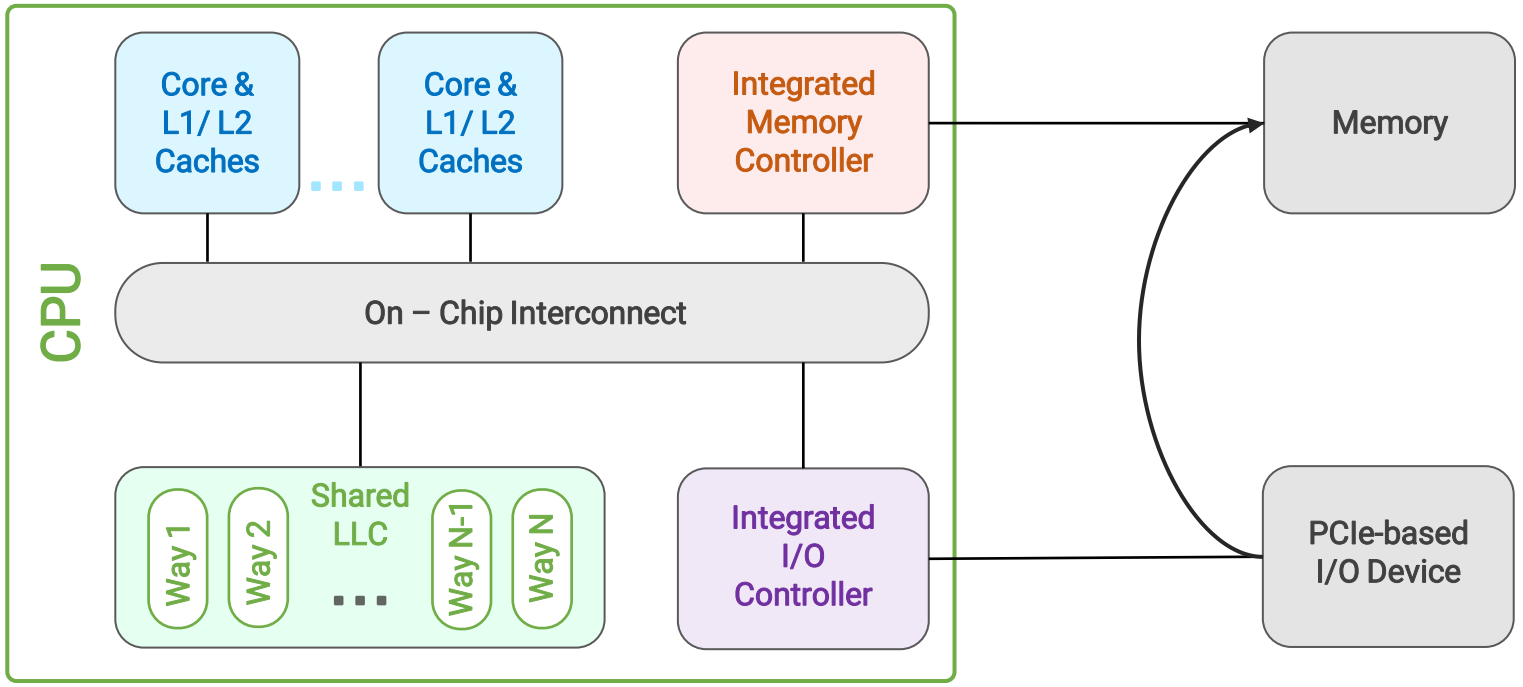

- Memory involvement for IO (DMA)
- Problems: High Speed Devices (NIC, SSDs), Poor memory bandwidth utilization

LLC Management and IO

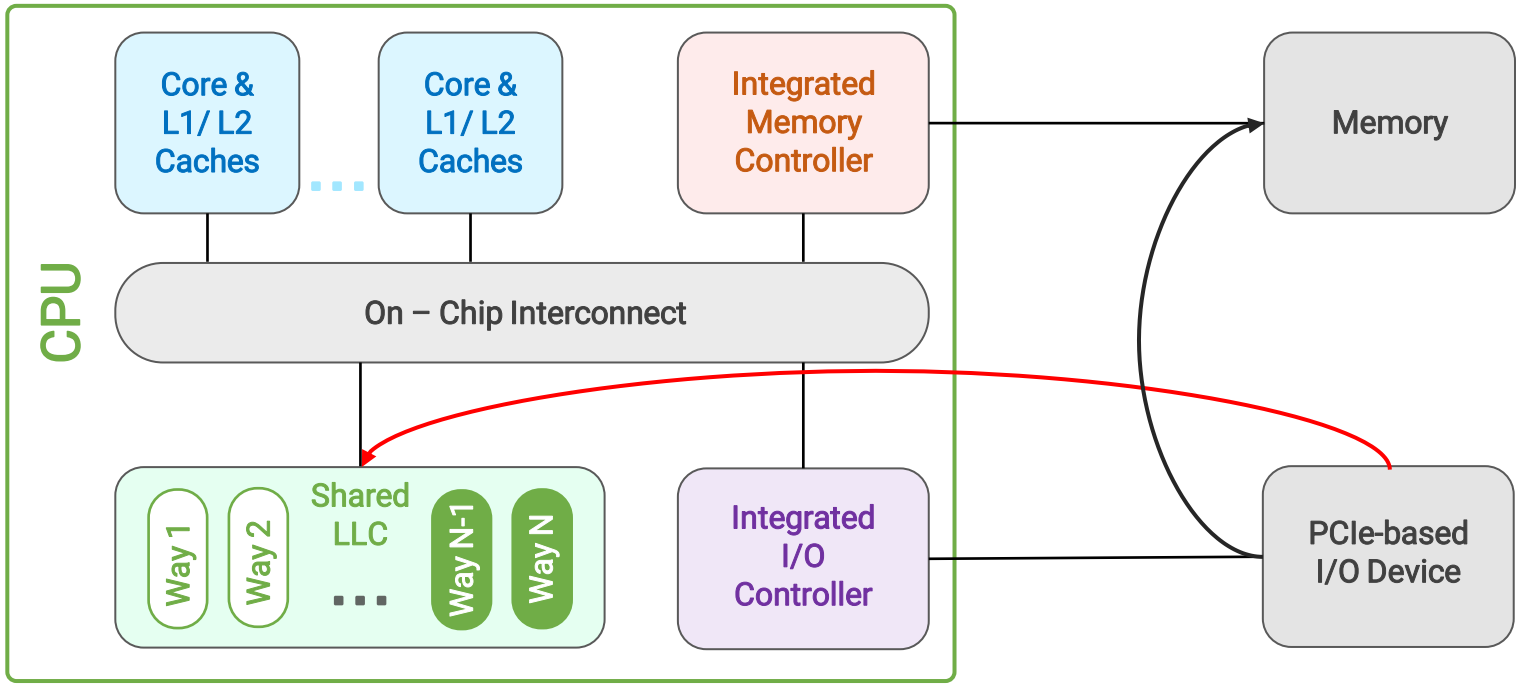

- Memory involvement for IO (DMA)
- Problems: High Speed Devices (NIC, SSDs), Poor memory bandwidth utilization
- DDIO (2 LLCs ways)

LLC Management and IO

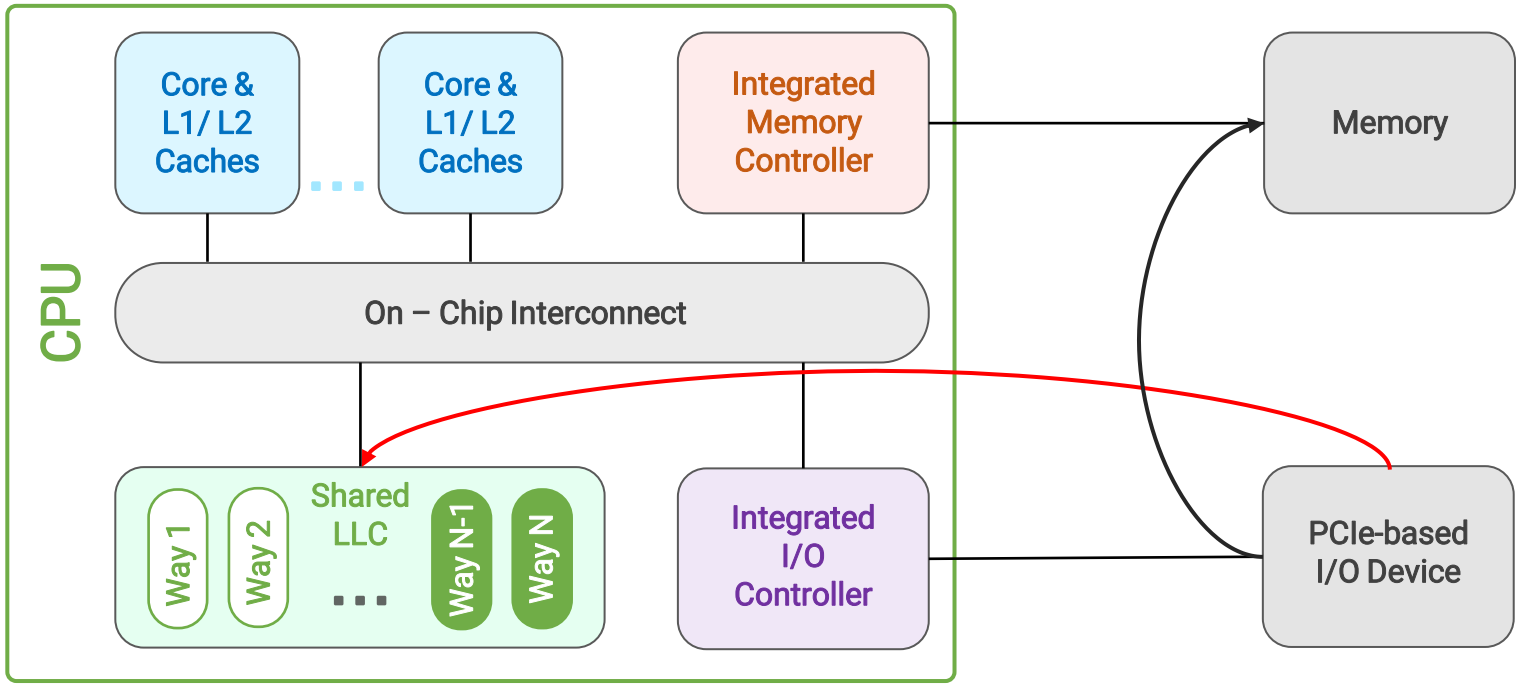

- Memory involvement for IO (DMA)
- Problems: High Speed Devices (NIC, SSDs), Poor memory bandwidth utilization
- DDIO (2 LLCs ways): introduces new challenges

Challenge 1: Leaky DMA

Challenge 1: Leaky DMA

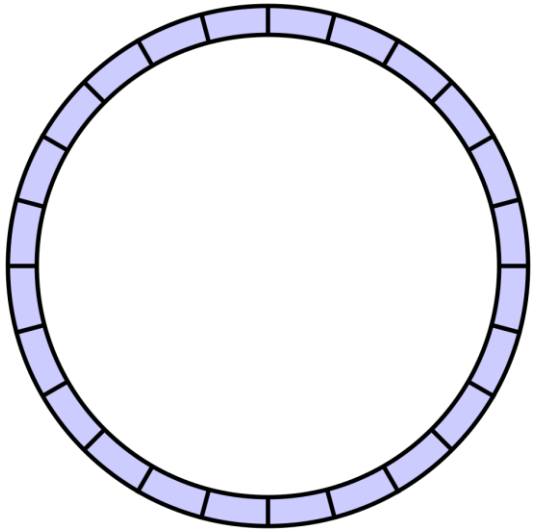

Rx Ring Buffer

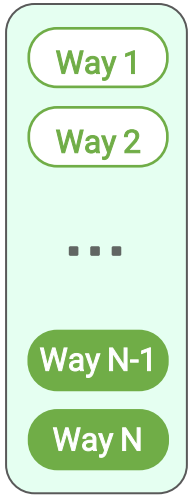

Shared LLC

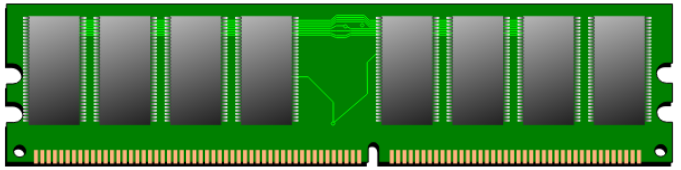

Memory

Challenge 1: Leaky DMA

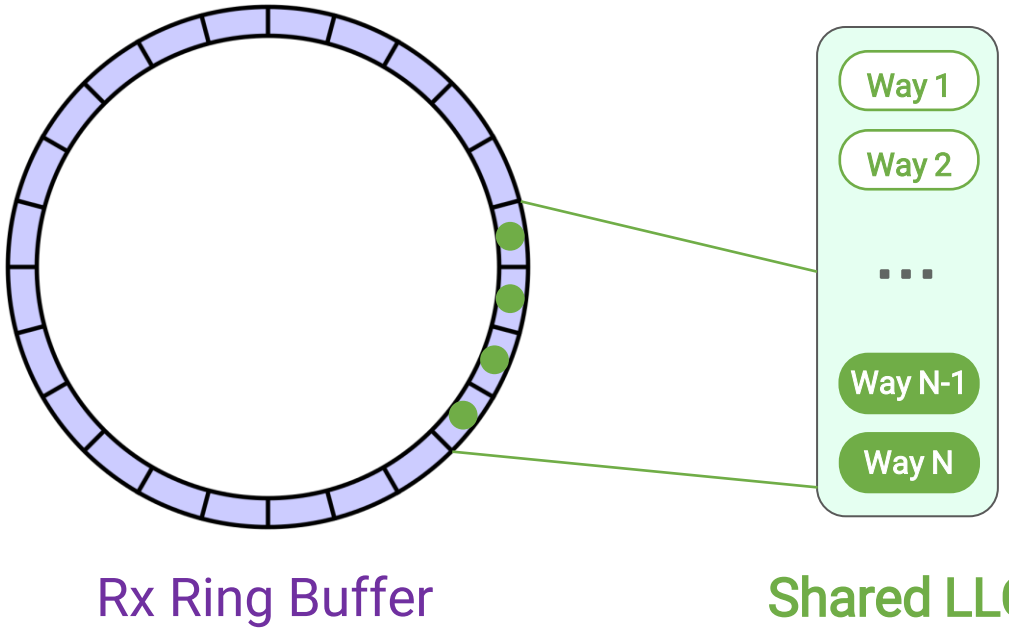

Rx Ring Buffer

Shared LLC

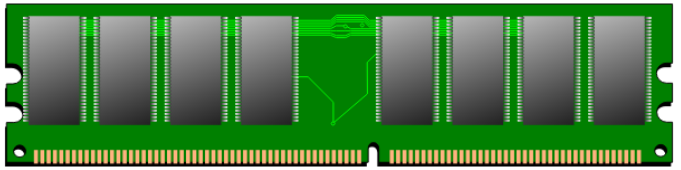

Memory

Challenge 1: Leaky DMA

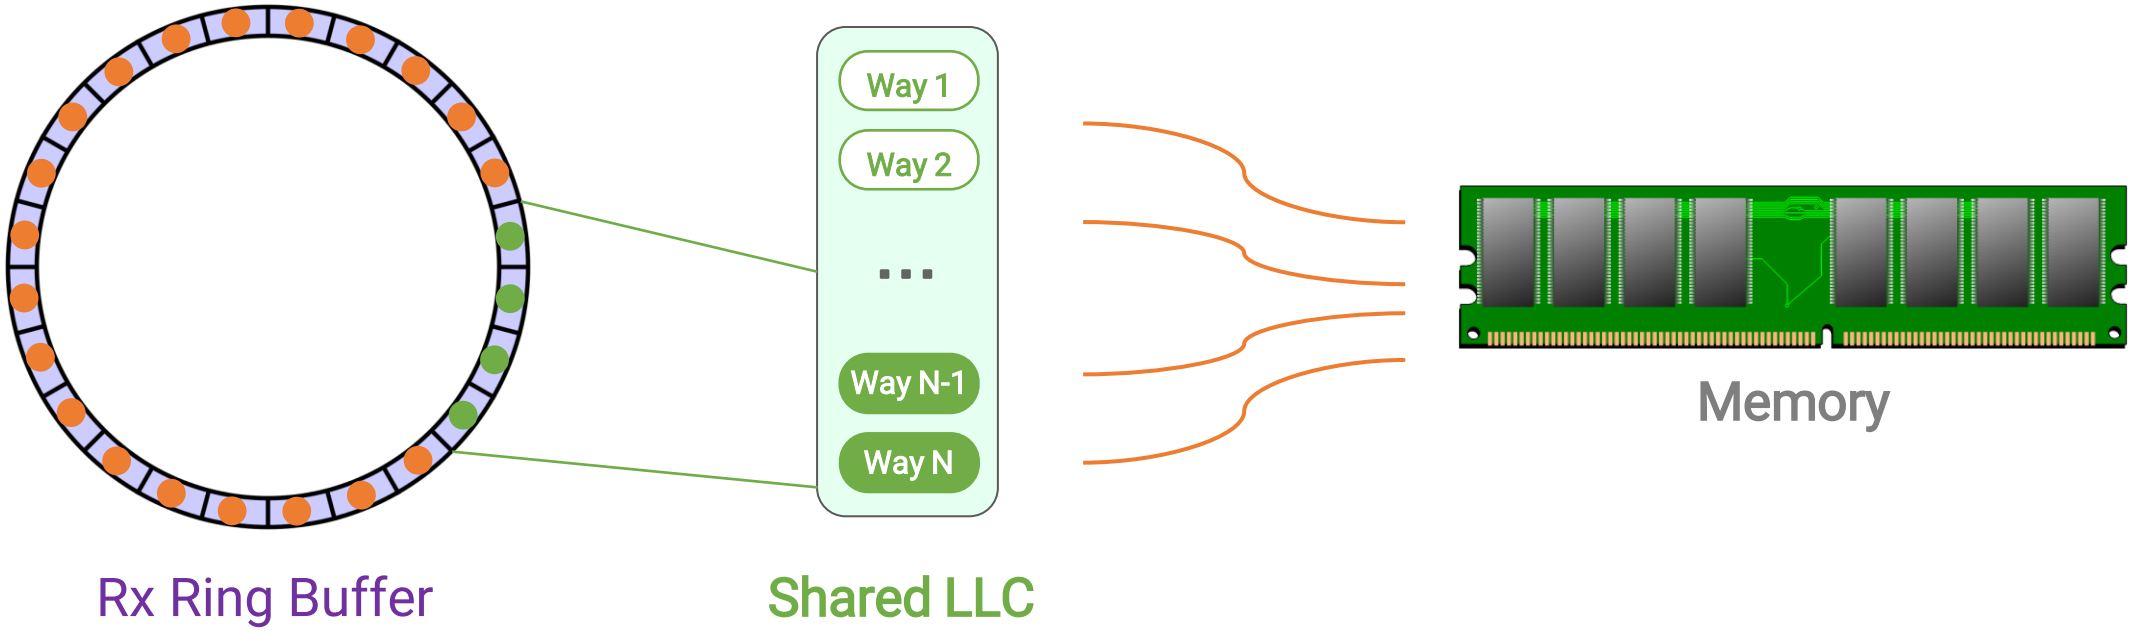

Challenge 1: Leaky DMA

- ResQ proposes to reduce Rx buffer size

Challenge 1: Leaky DMA

- ResQ proposes to reduce Rx buffer size:
  - Very less buffer size for multiple tenants

Challenge 1: Leaky DMA

- ResQ proposes to reduce Rx buffer size:
  - Very less buffer size for multiple tenants
  - Throughput drop

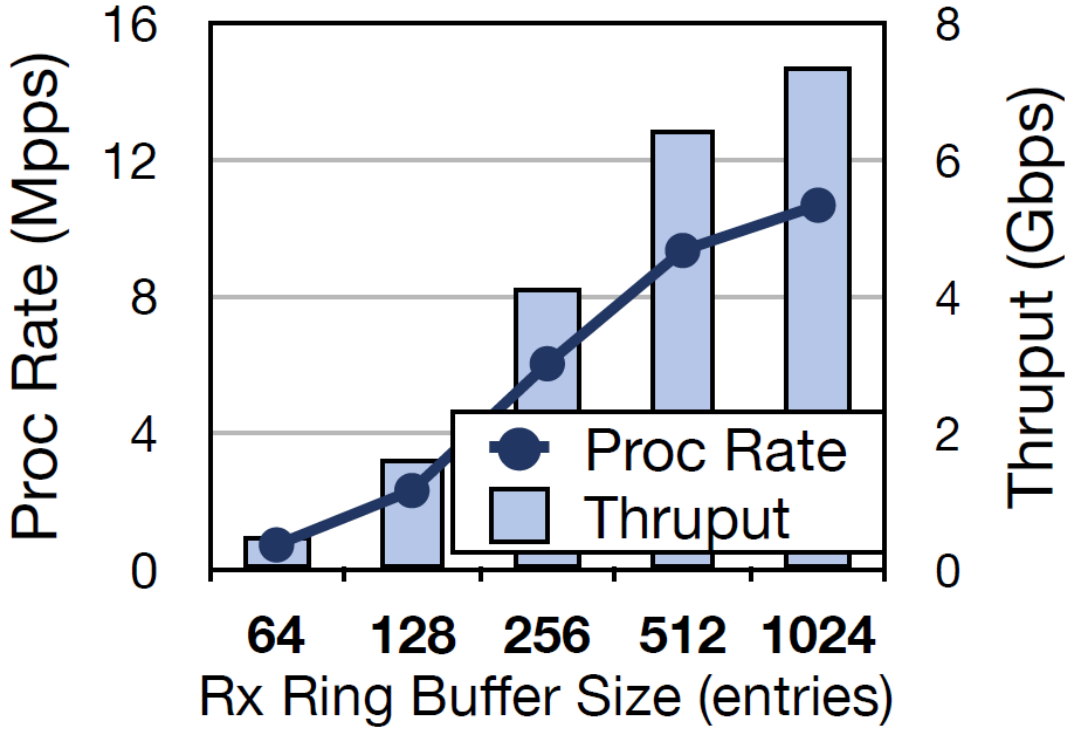

(a) 64B small packet.

Challenge 1: Leaky DMA

- ResQ proposes to reduce Rx buffer size:
  - Very less buffer size for multiple tenants
  - Throughput drop

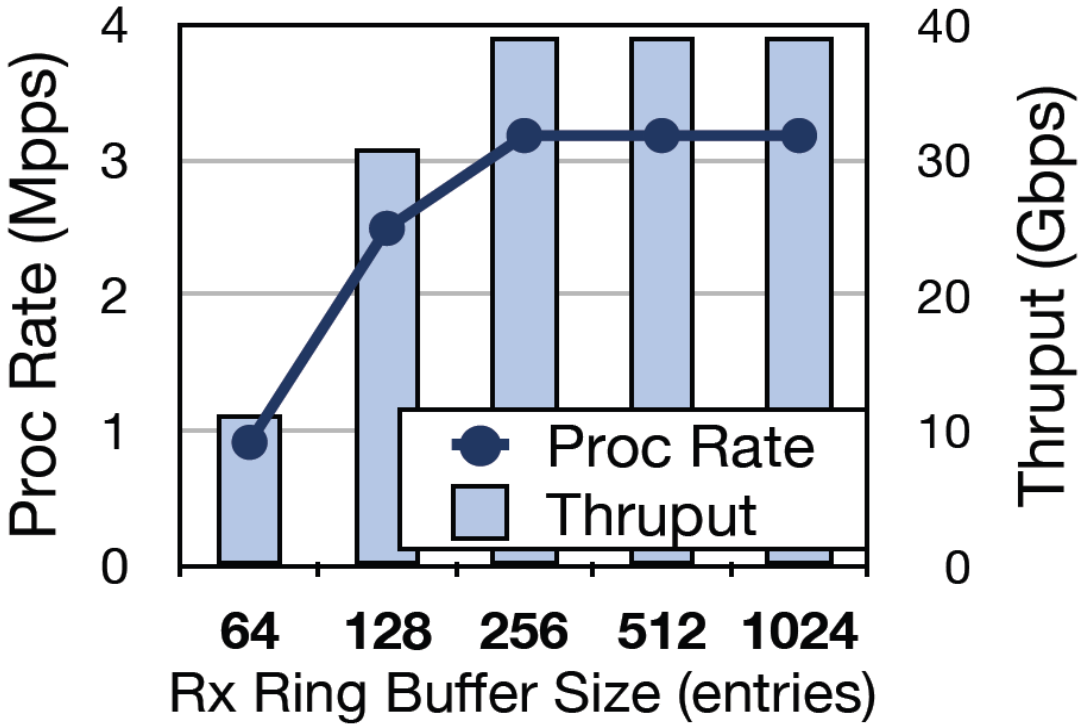

(b) 1.5KB large packet.

Challenge 1: Leaky DMA

- ResQ proposes to reduce Rx buffer size:
  - Very less buffer size for multiple tenants
  - Throughput drop
- **How to adaptively change DDIO ways ?**

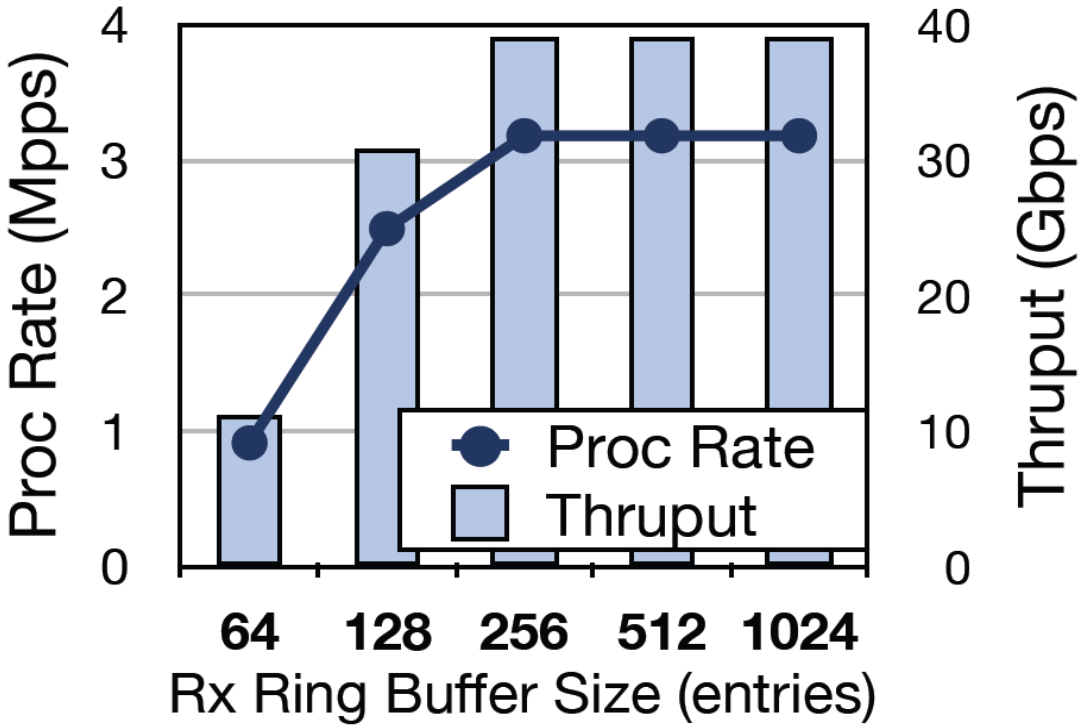

(b) 1.5KB large packet.

Challenge 2: Latent Contender

Challenge 2: Latent Contender

- Overlapping of DDIO ways with Core affect performance

Challenge 2: Latent Contender

- Overlapping of DDIO ways with Core affect performance

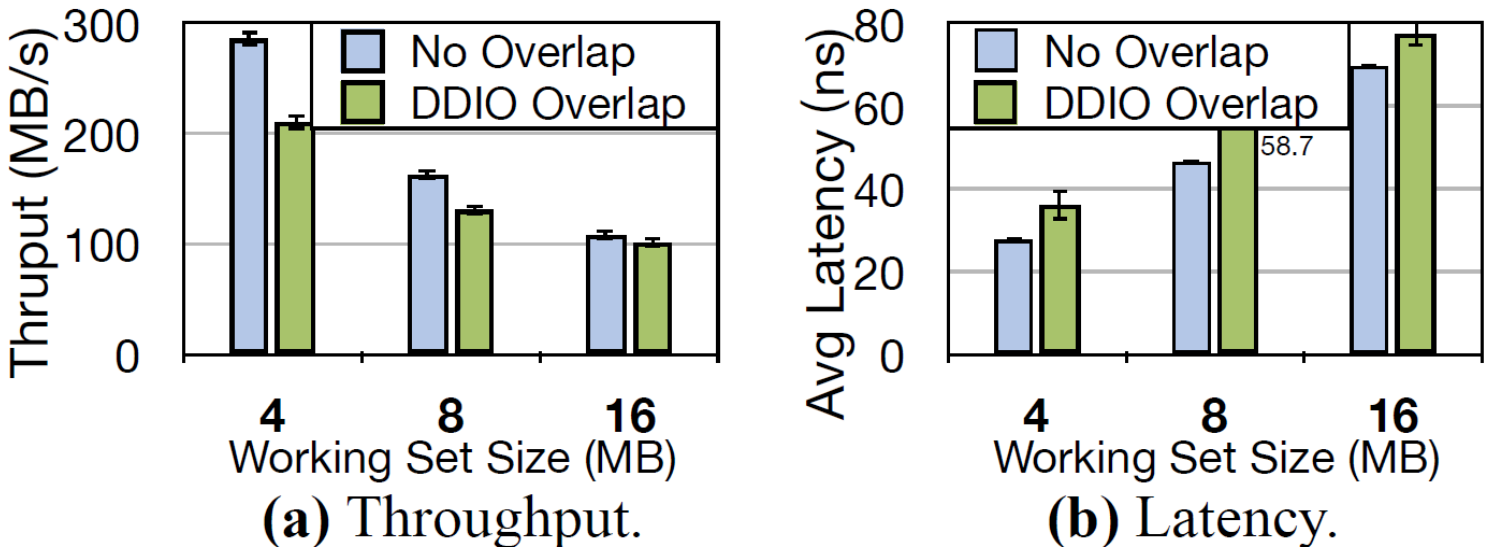

- DDIO overlap worsen the X-Mem's (random read pattern) throughput and avg latency

Challenge 2: Latent Contender

- Overlapping of DDIO ways with Core affect performance

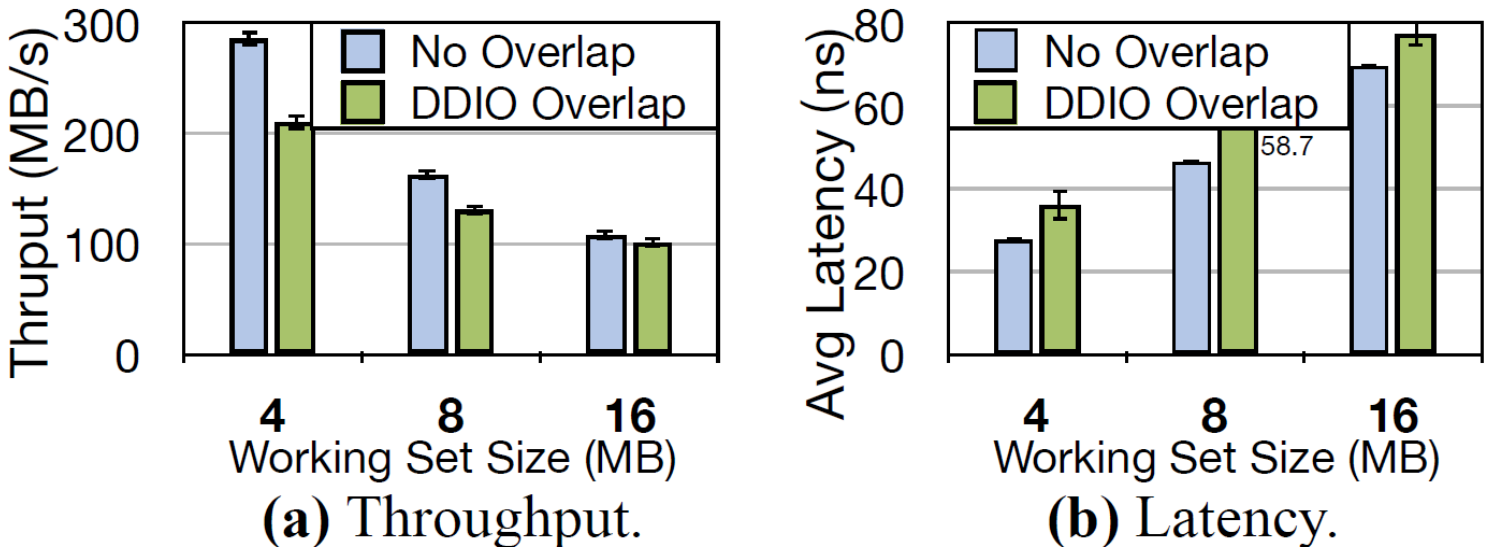

- DDIO overlap worsen the X-Mem's (random read pattern) throughput and avg latency
- **How to share LLC ways with DDIO ?**

IAT: Design

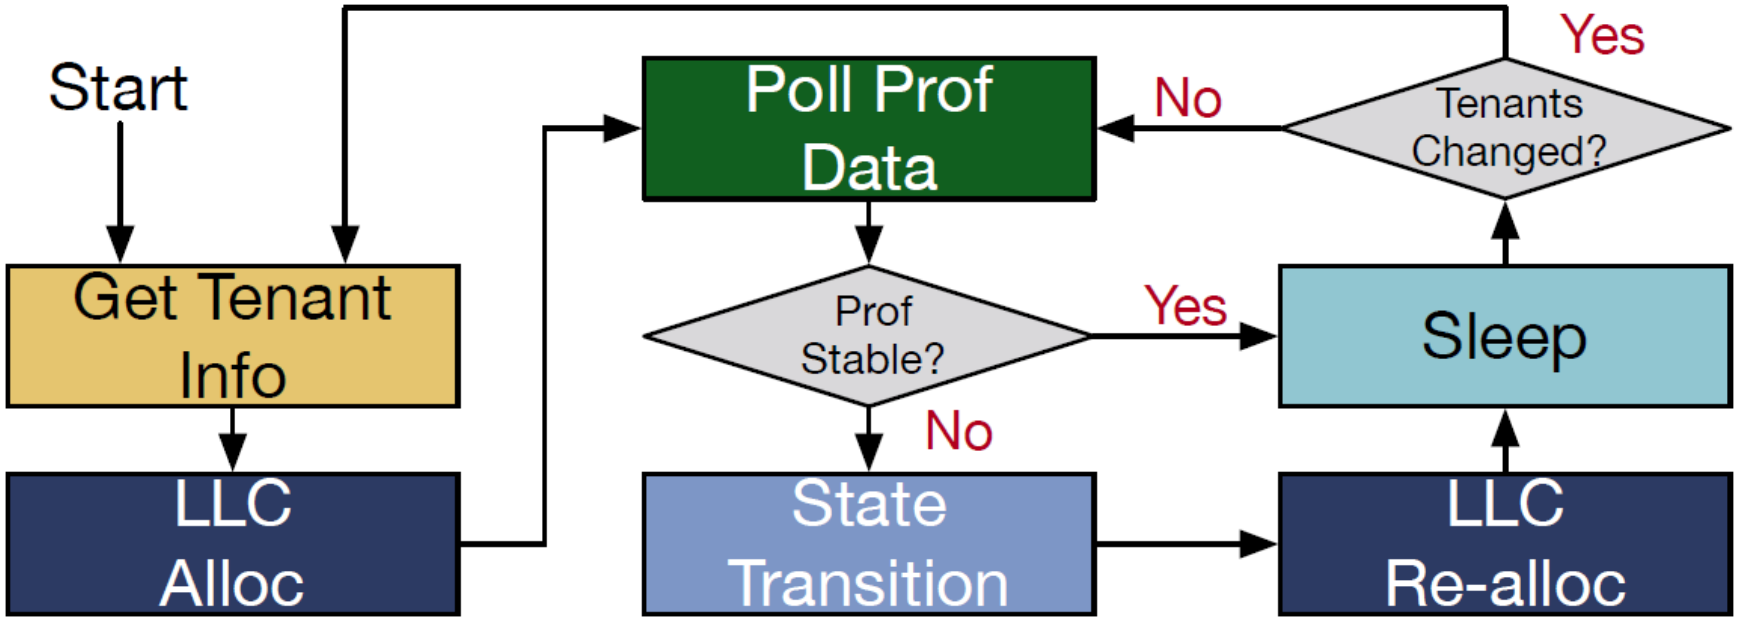
Start
Get Tenant Info
LLC Alloc
Poll Prof Data
Prof Stable?
No
Yes
State Transition
Sleep
Tenants Changed?
No
Yes
LLC Re-alloc

IAT: Design

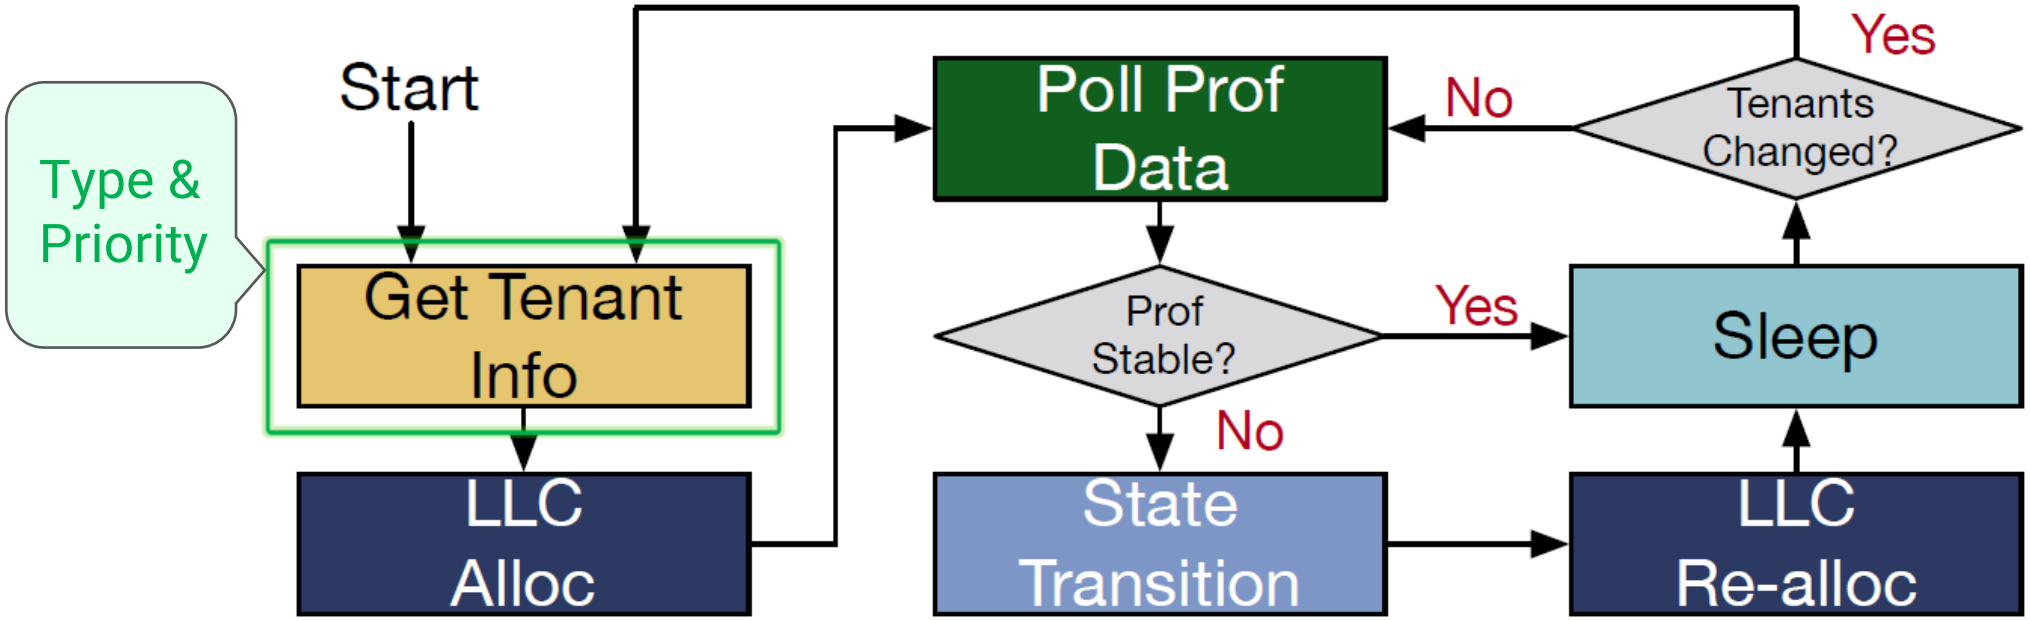
Start
Type & Priority
Get Tenant Info
LLC Alloc
Poll Prof Data
Prof Stable?
Yes
No
State Transition
Sleep
Tenants Changed?
No
Yes
LLC Re-alloc

IAT: Design

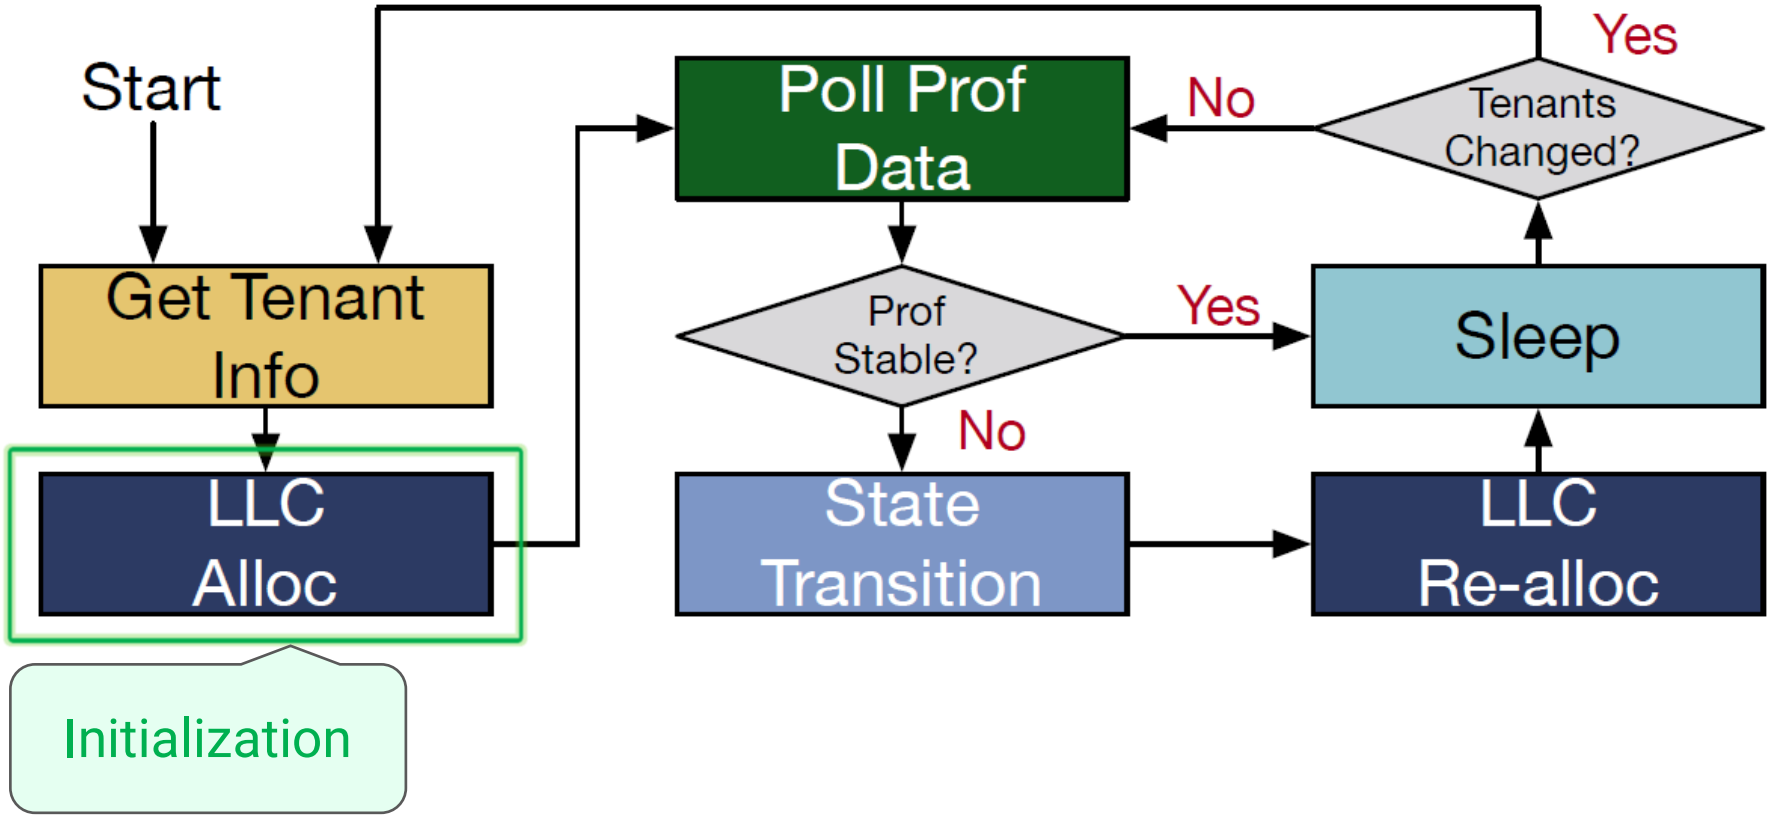
Start
Get Tenant Info
LLC Alloc
Initialization
Poll Prof Data
Prof Stable?
No
Yes
State Transition
Tenants Changed?
No
Yes
Sleep
LLC Re-alloc

IAT: Design

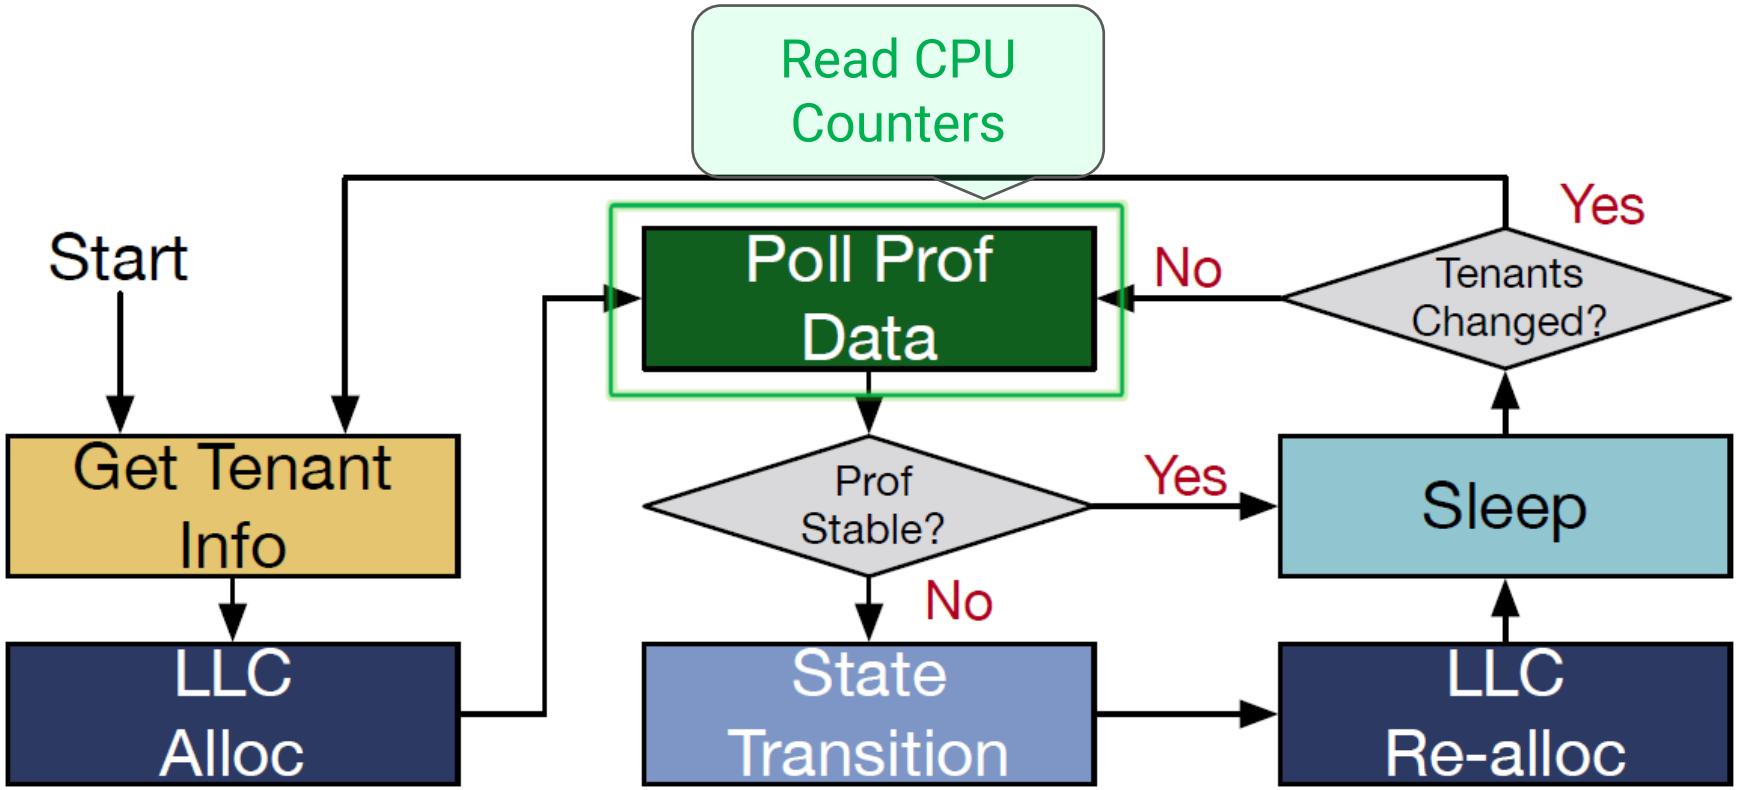
Read CPU Counters
Start
Poll Prof Data
No
Tenants Changed?
Yes
Get Tenant Info
Prof Stable?
Yes
Sleep
No
LLC Alloc
State Transition
LLC Re-alloc

IAT: Design

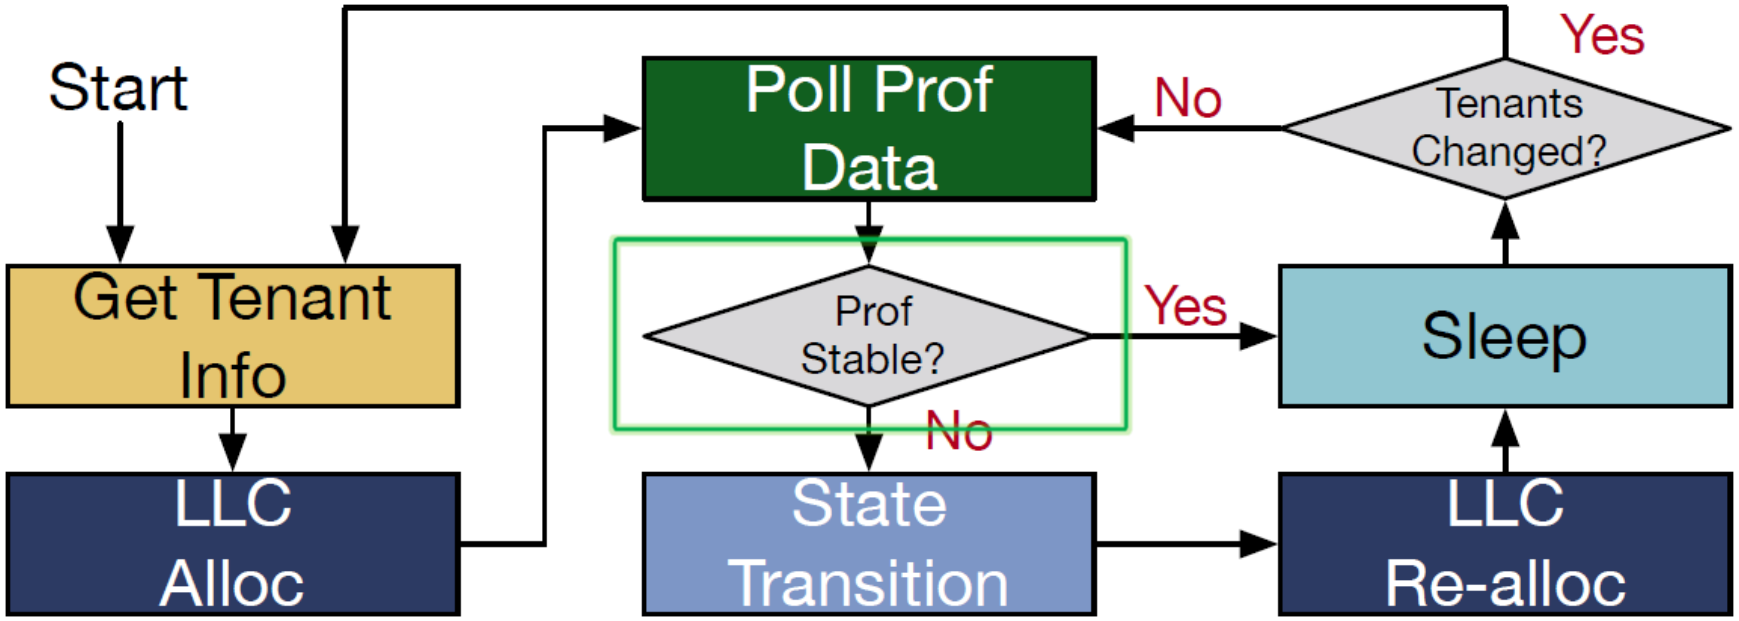
Start
Get Tenant Info
LLC Alloc
Poll Prof Data
Prof Stable?
Yes
No
State Transition
Sleep
Tenants Changed?
No
Yes
LLC Re-alloc

IAT: Design

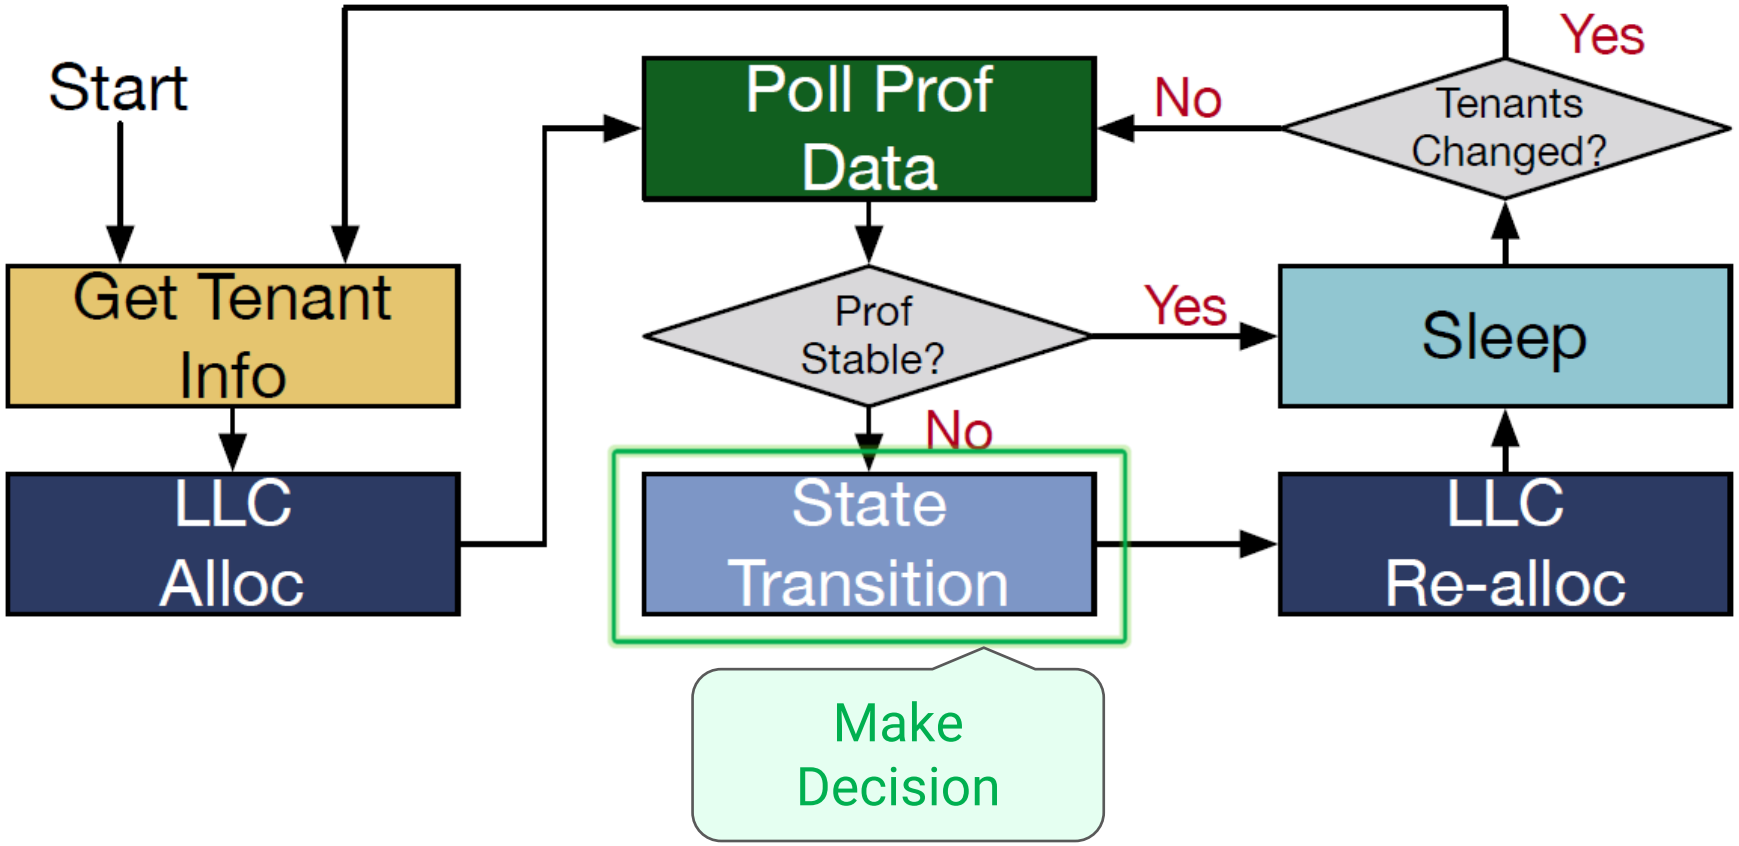
Start
Get Tenant Info
LLC Alloc
Poll Prof Data
Prof Stable?
Yes
No
State Transition
Make Decision
Tenants Changed?
No
Yes
Sleep
LLC Re-alloc

IAT: Design

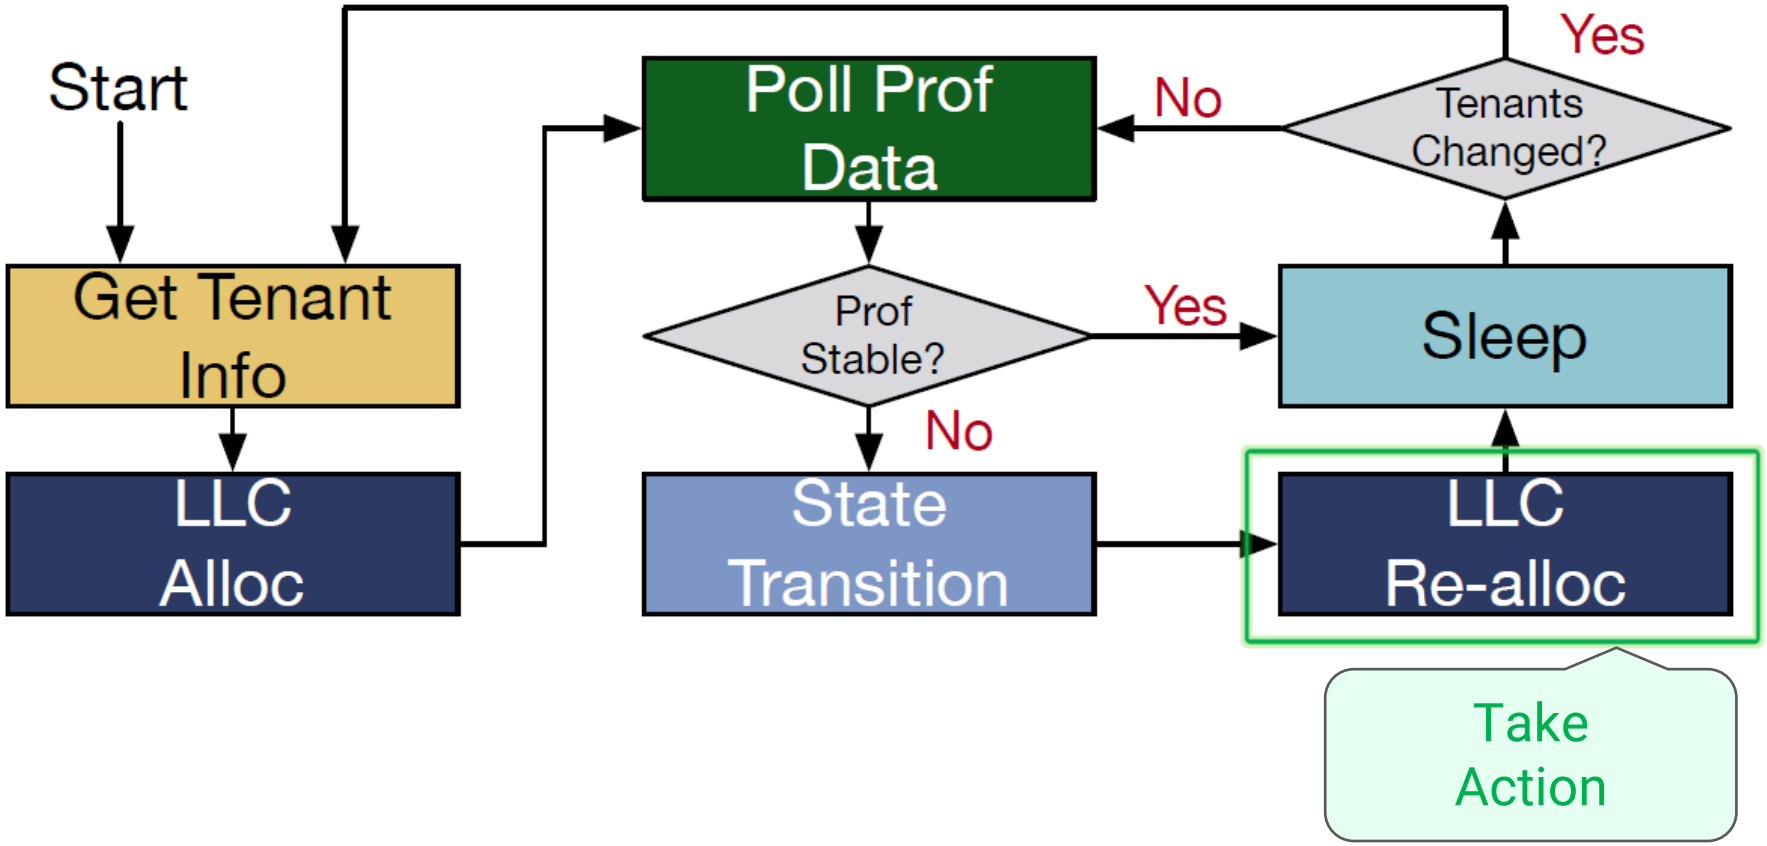
Start
Get Tenant Info
LLC Alloc
Poll Prof Data
Prof Stable?
No
Yes
State Transition
Tenants Changed?
No
Yes
Sleep
LLC Re-alloc
Take Action

IAT: Design

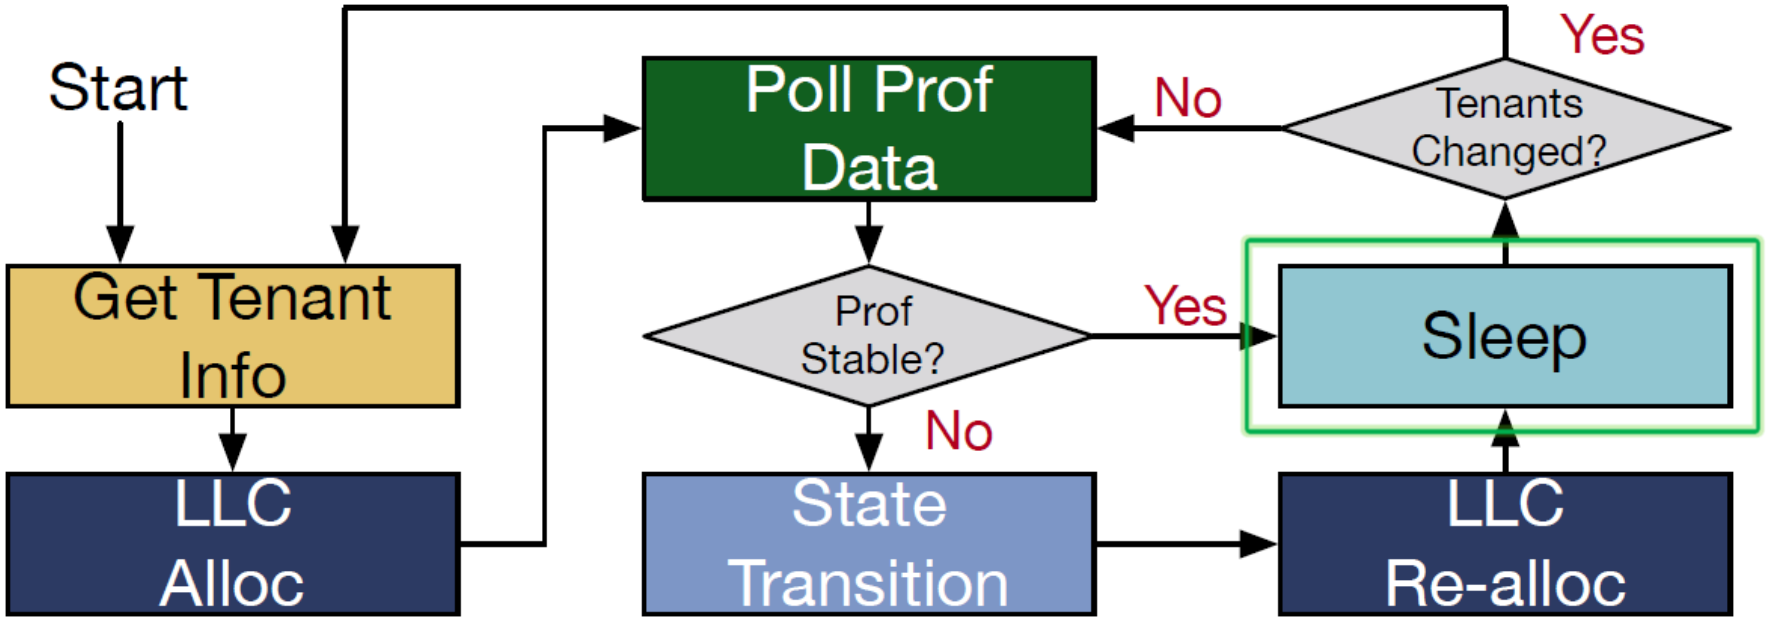
Start
Get Tenant Info
LLC Alloc
Poll Prof Data
Prof Stable?
No
Yes
State Transition
Tenants Changed?
No
Yes
Sleep
LLC Re-alloc

IAT: Poll Prof Data

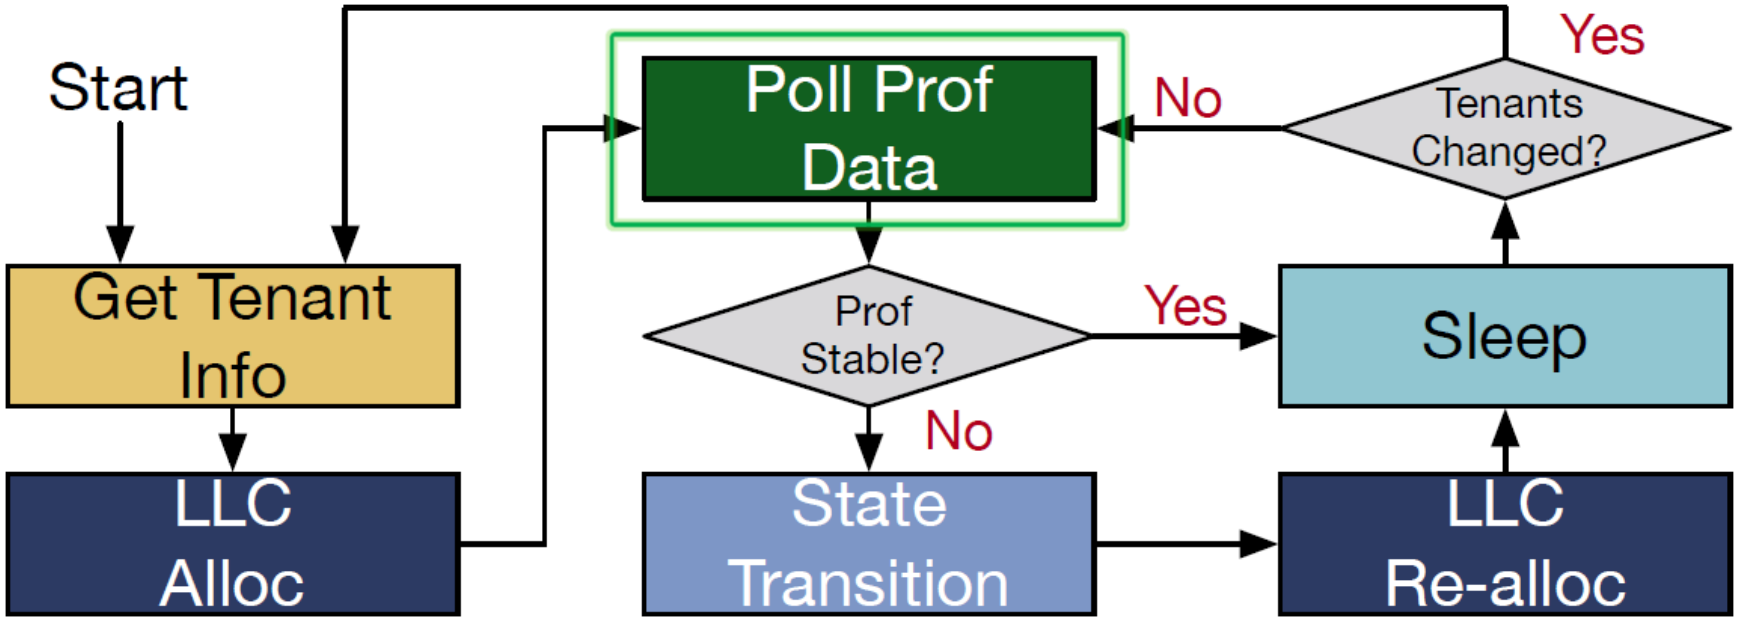

IAT: Poll Prof Data

- IPC

IAT: Poll Prof Data

- IPC: performance of an application

IAT: Poll Prof Data

- IPC: performance of an application
- LLC ref and miss

IAT: Poll Prof Data

- IPC: performance of an application
- LLC ref and miss: memory/ cache access characteristics of a workload

IAT: Poll Prof Data

- IPC: performance of an application
- LLC ref and miss: memory/ cache access characteristics of a workload
- DDIO hit and miss (chip wide metrics)

IAT: Poll Prof Data

- IPC: performance of an application
- LLC ref and miss: memory/ cache access characteristics of a workload
- DDIO hit and miss (chip wide metrics):
  - Hit when targeted cache-line in LLC

IAT: Poll Prof Data

- IPC: performance of an application
- LLC ref and miss: memory/ cache access characteristics of a workload
- DDIO hit and miss (chip wide metrics):
  - Hit when targeted cache-line in LLC
  - Miss when a victim has to be evicted for DDIO

IAT: Poll Prof Data

- IPC: performance of an application
- LLC ref and miss: memory/ cache access characteristics of a workload
- DDIO hit and miss (chip wide metrics):
  - Hit when targeted cache-line in LLC
  - Miss when a victim has to be evicted for DDIO
  - Tells about IO pressure on LLC

IAT: State Transition

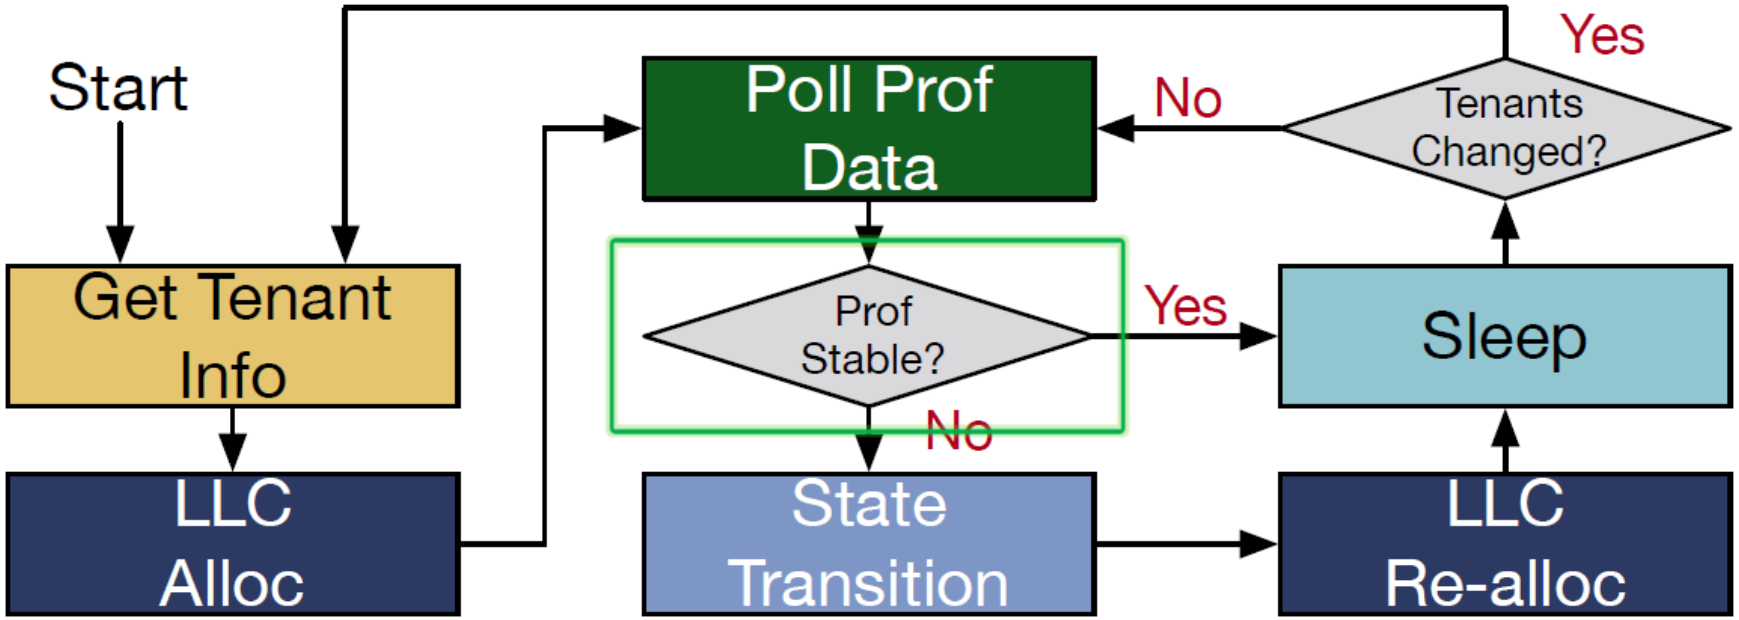

IAT: State Transition

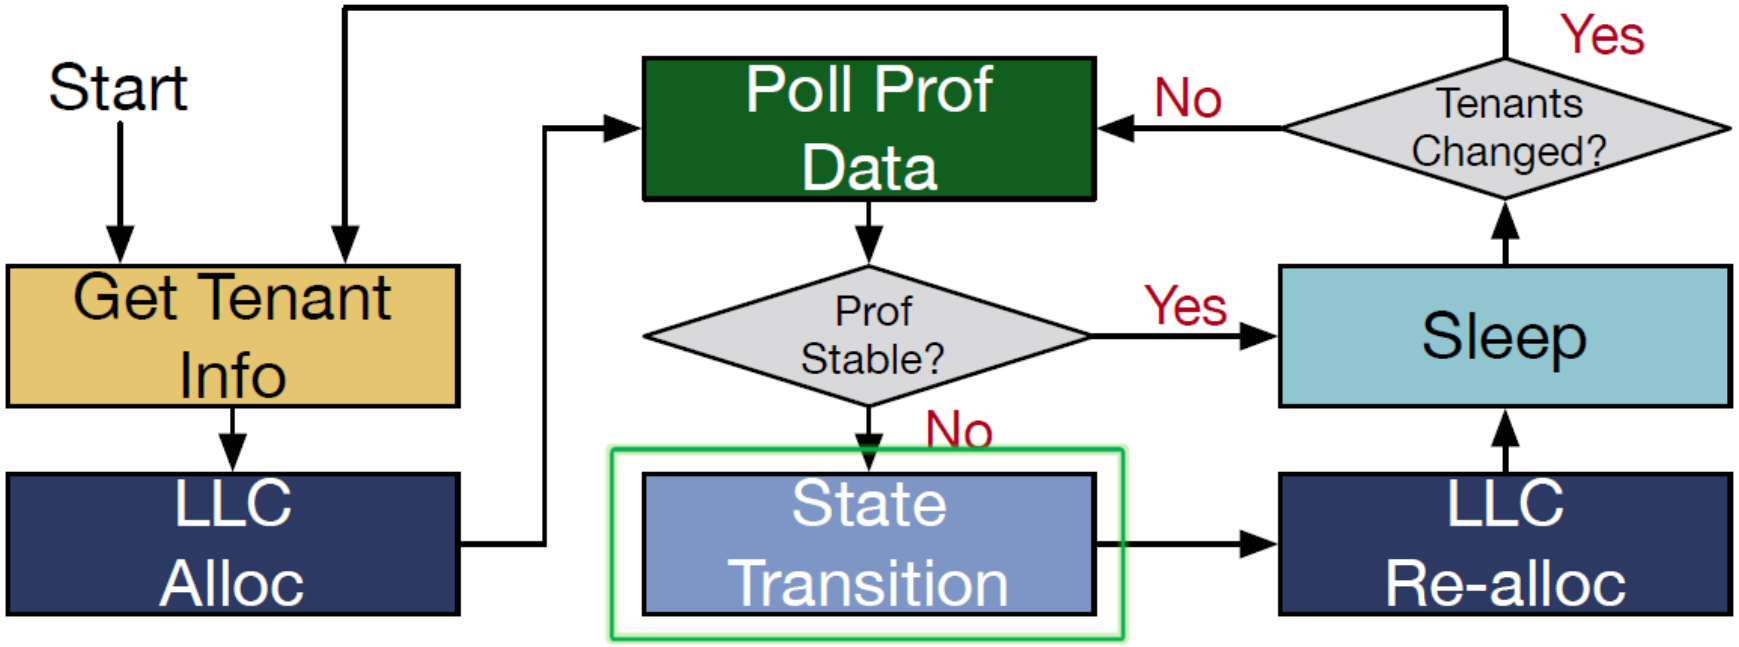

IAT: State Transition

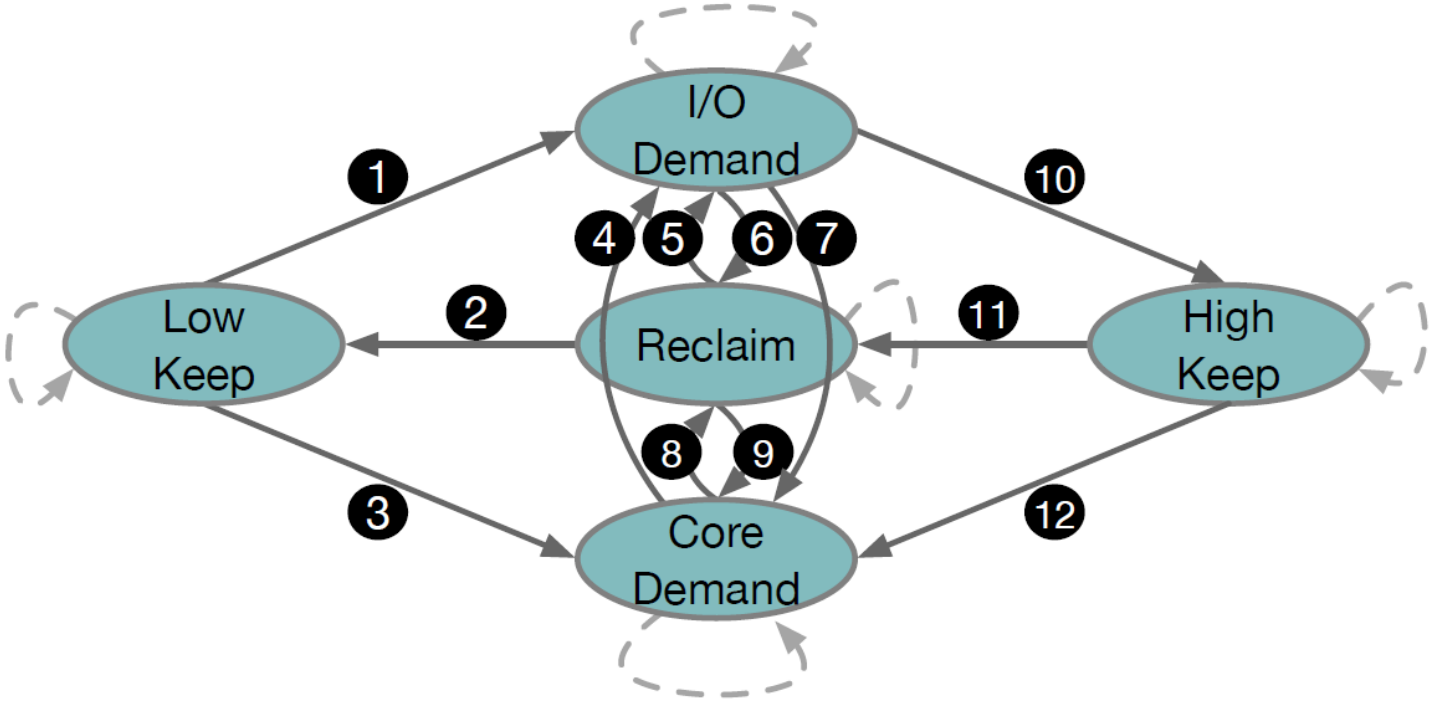

IAT: State Transition

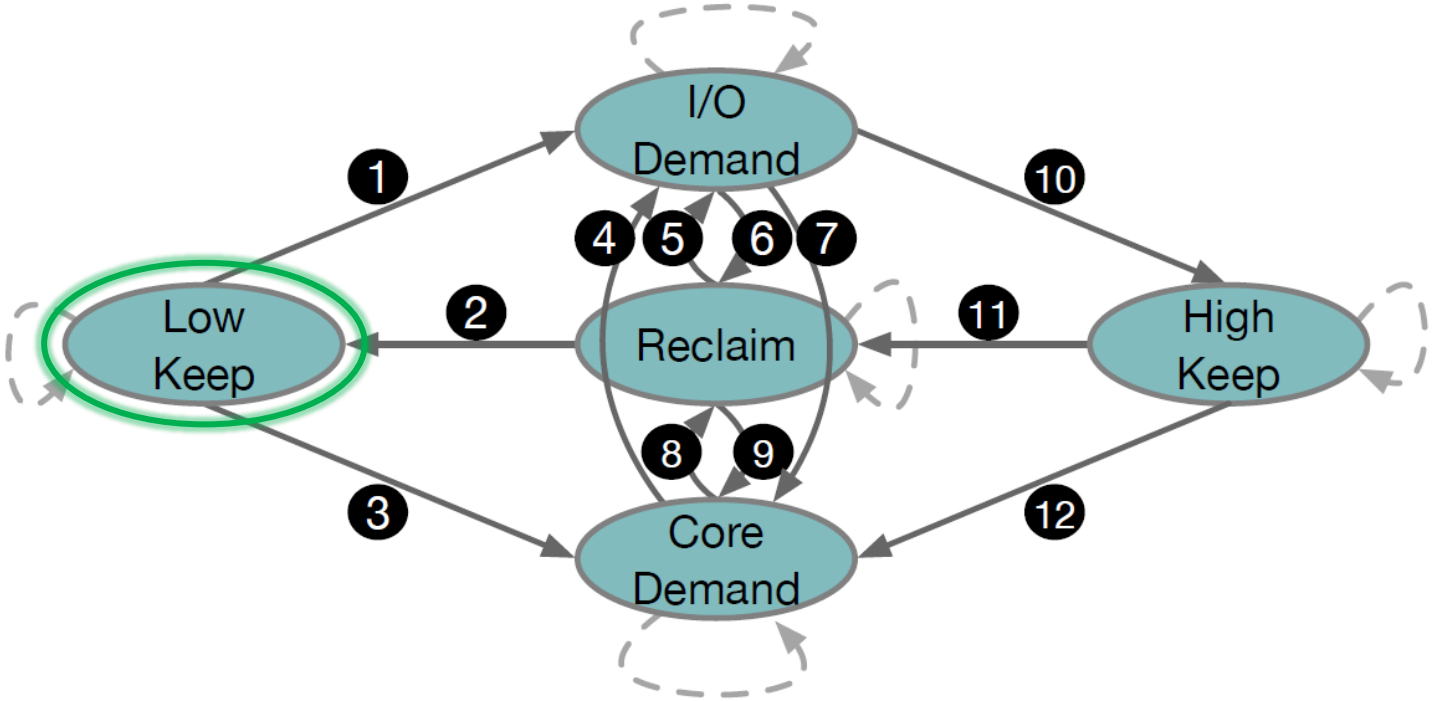

IAT: State Transition

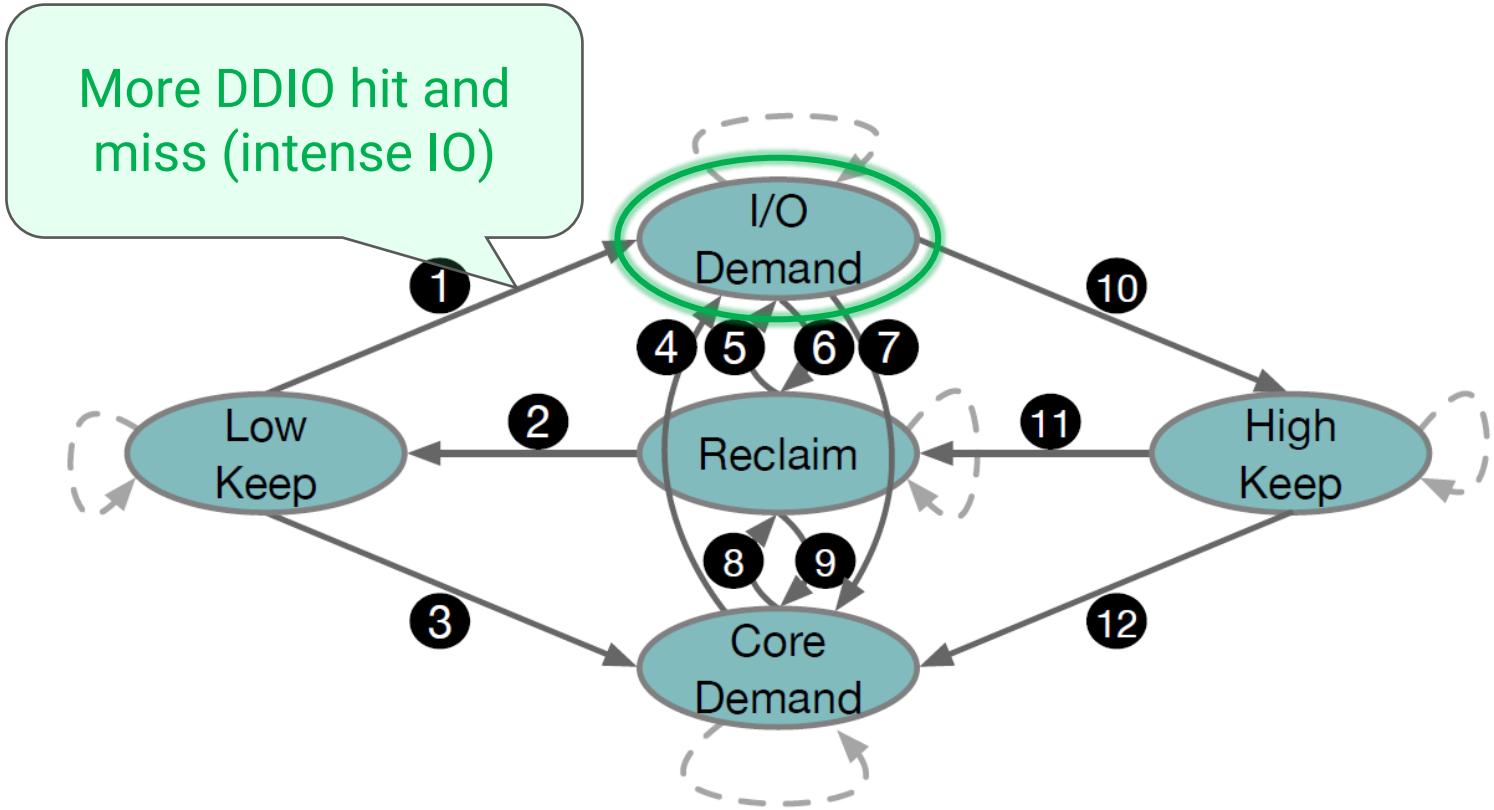

IAT: State Transition

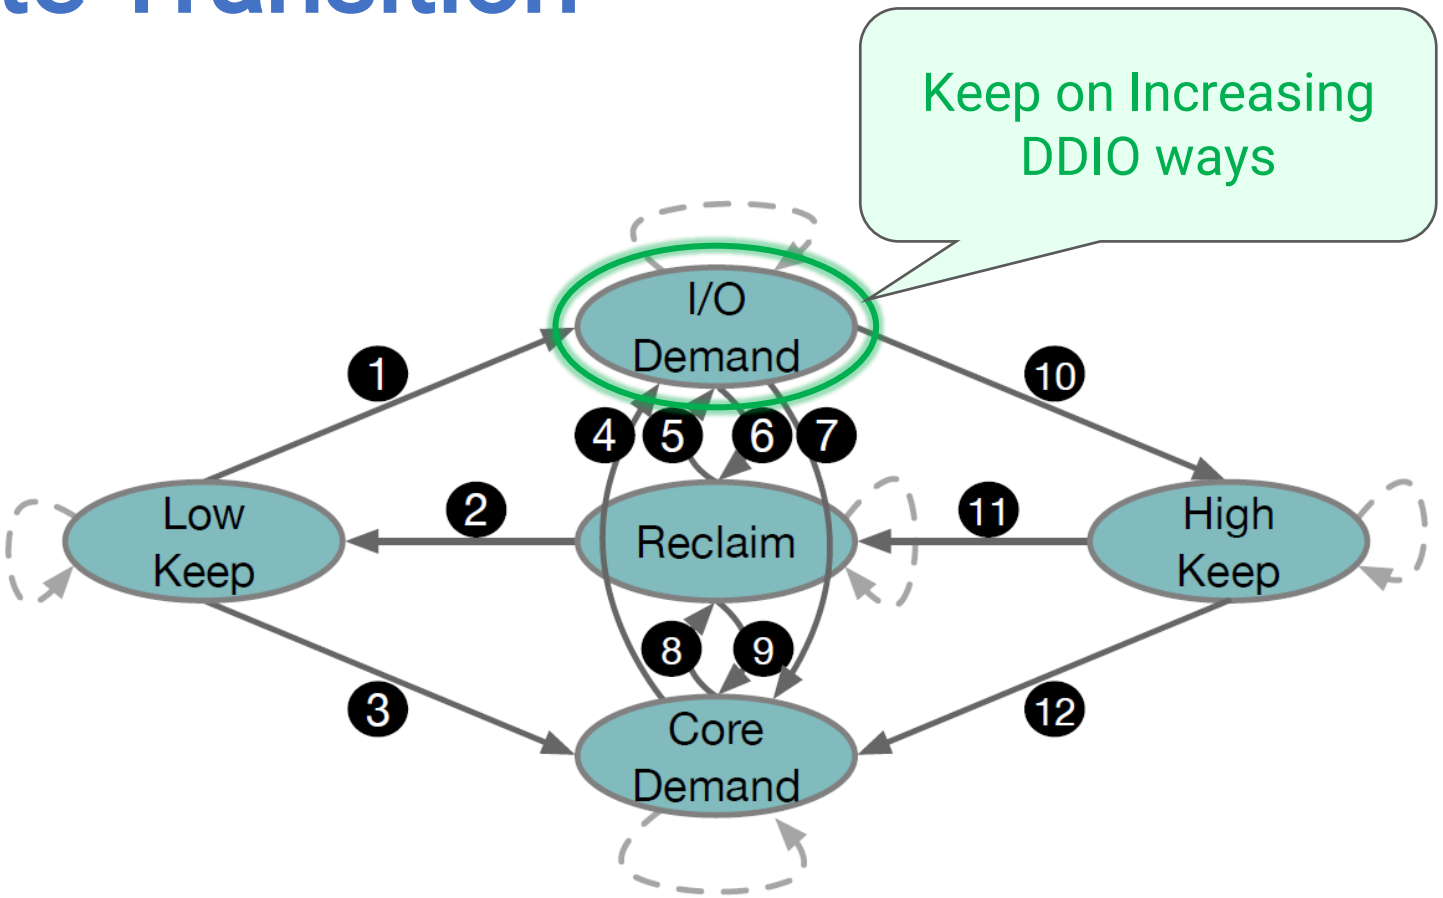

IAT: State Transition

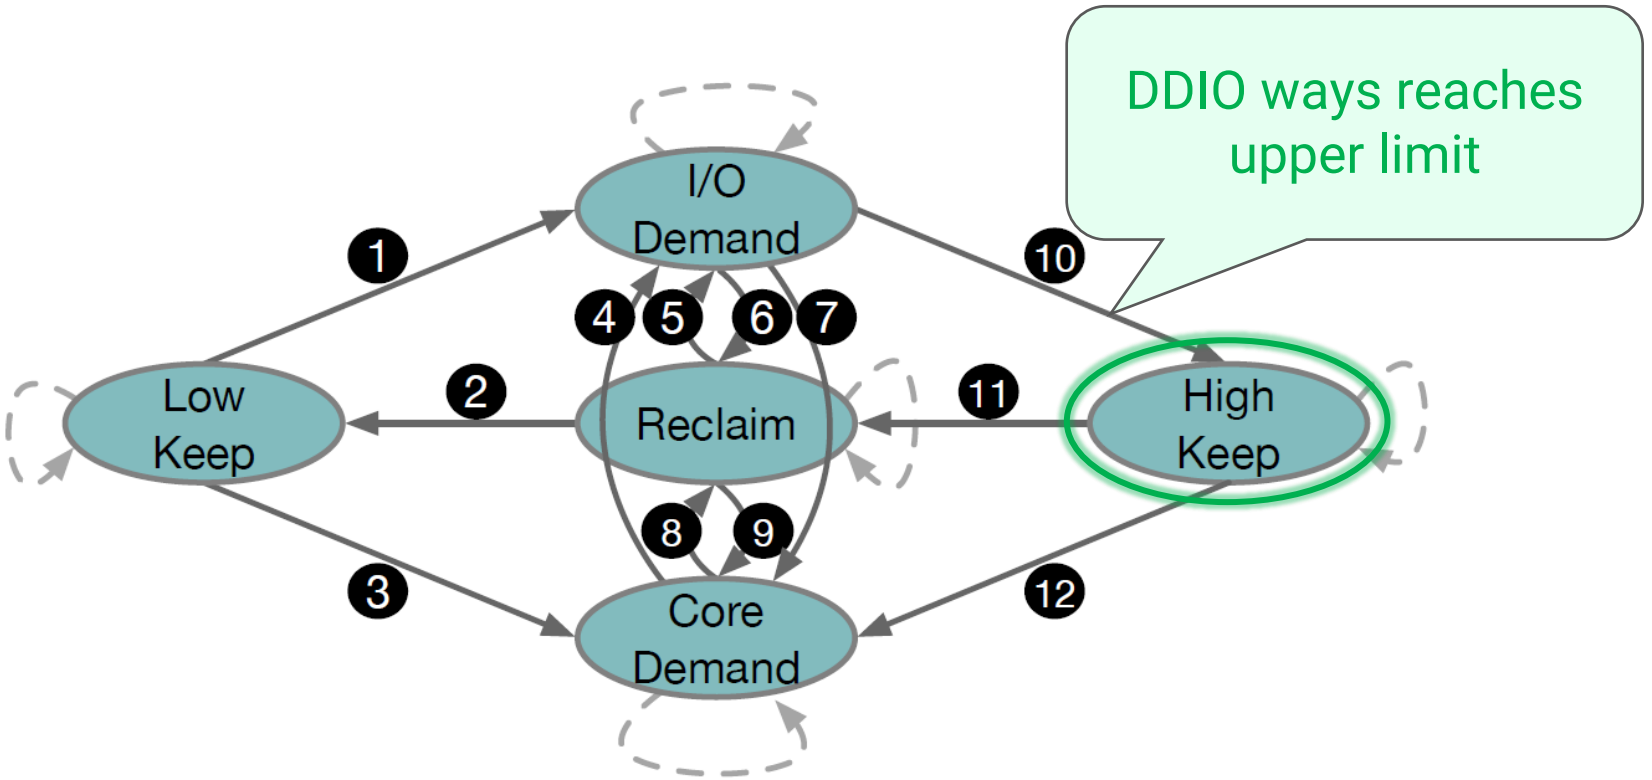

IAT: State Transition

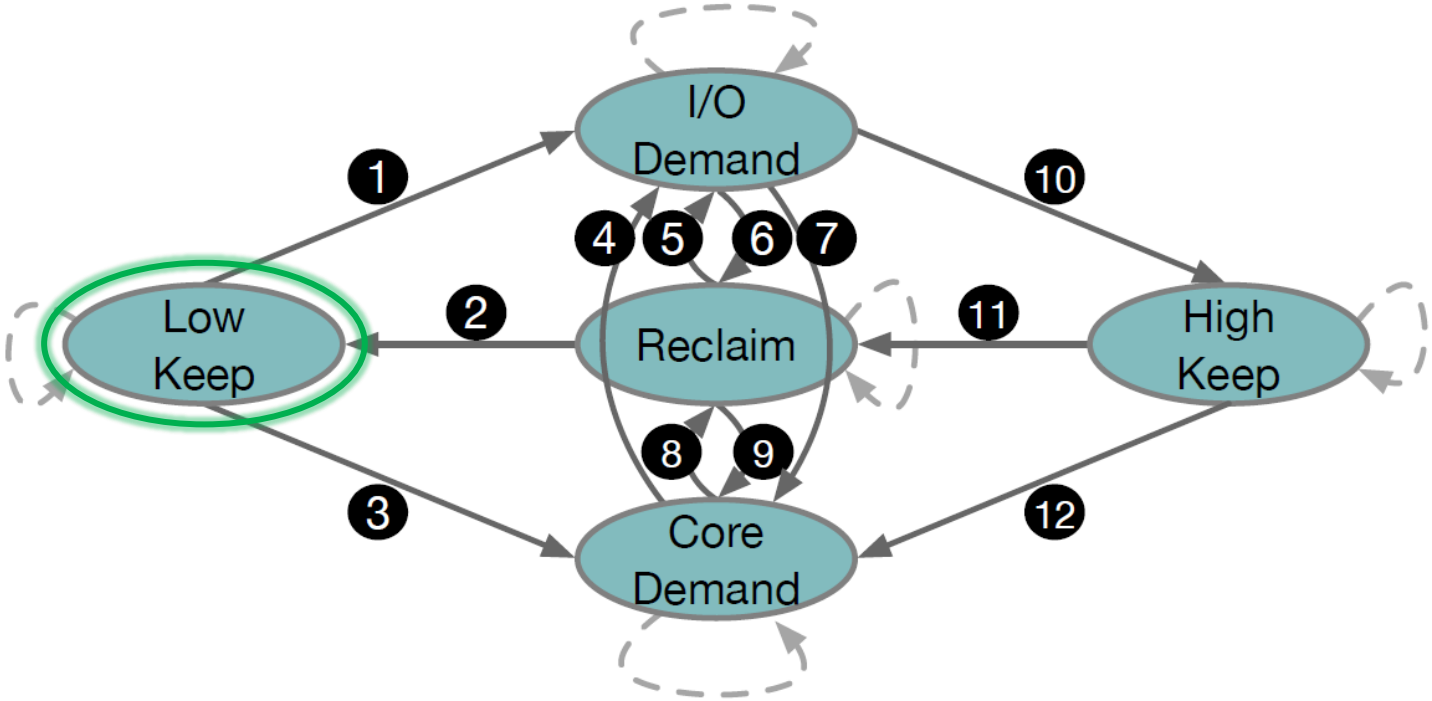

IAT: State Transition

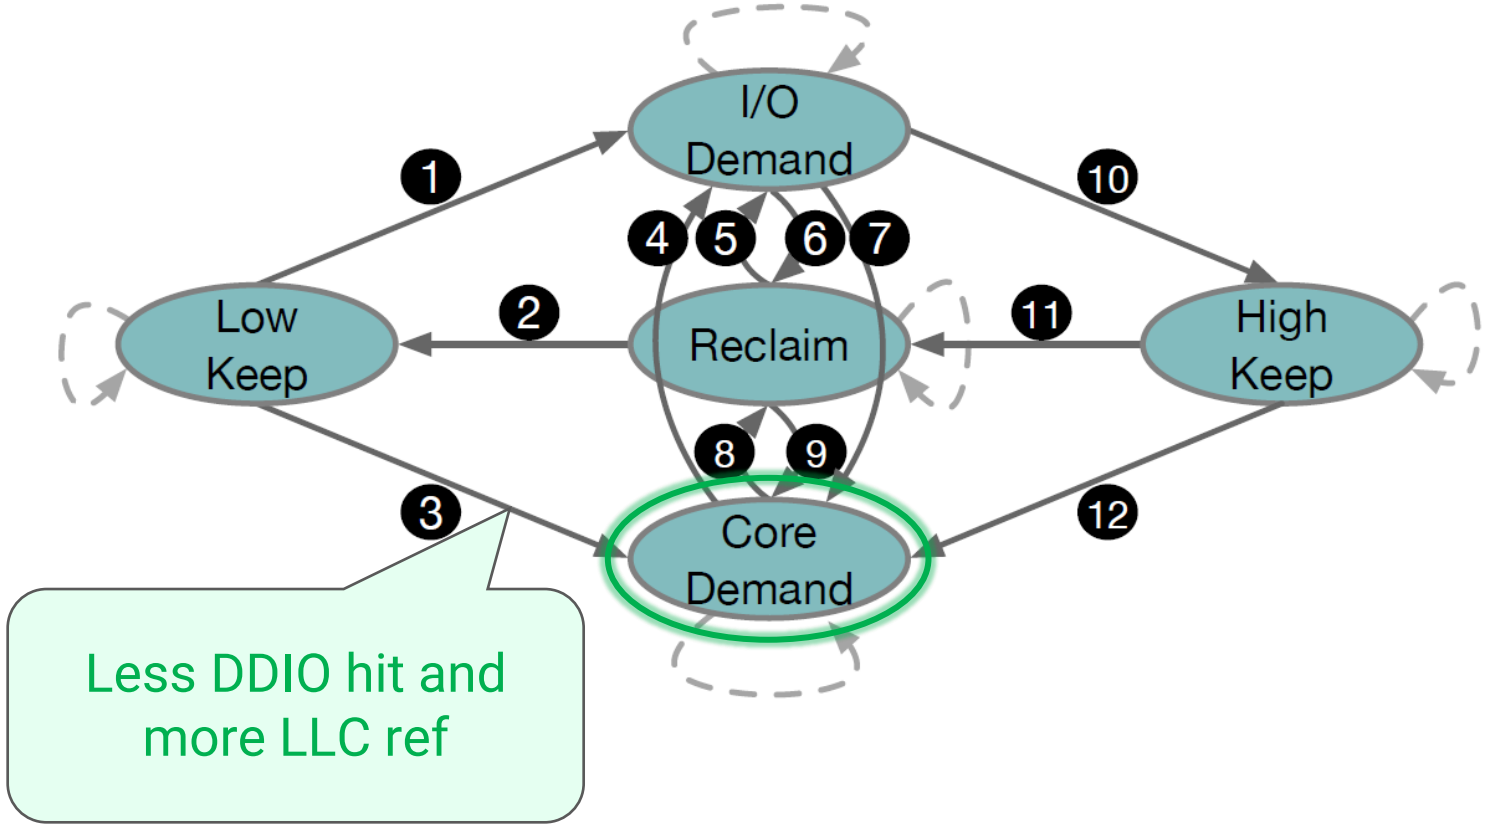

IAT: State Transition

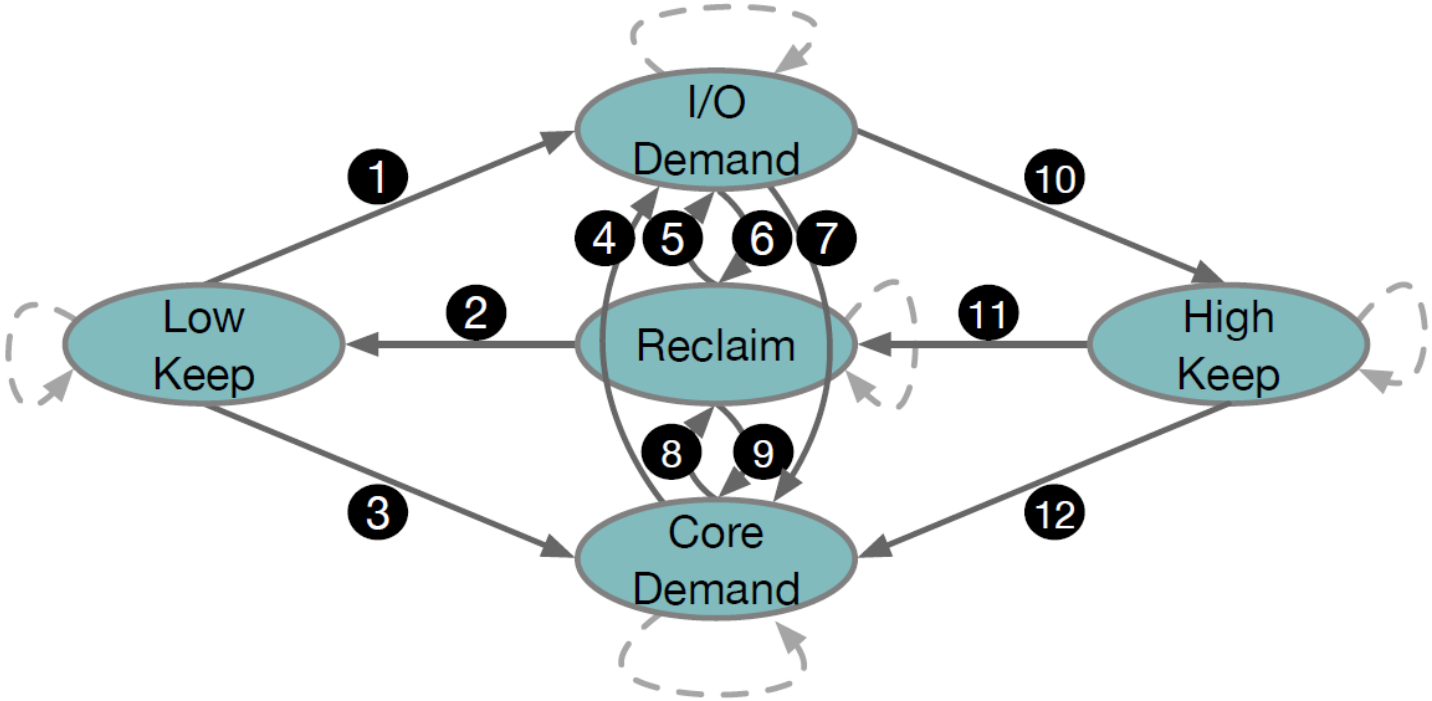

IAT: LLC Re – alloc

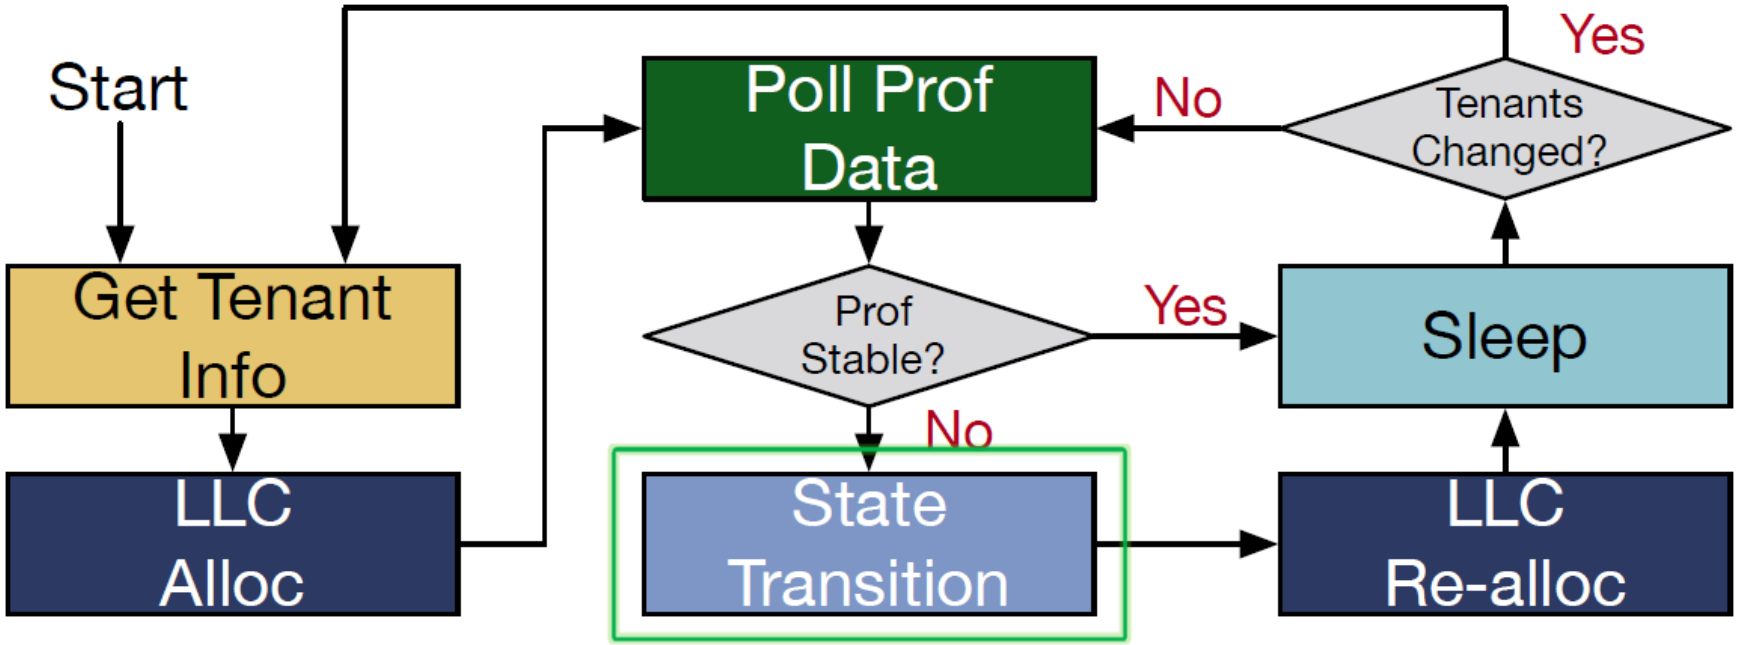

IAT: LLC Re – alloc

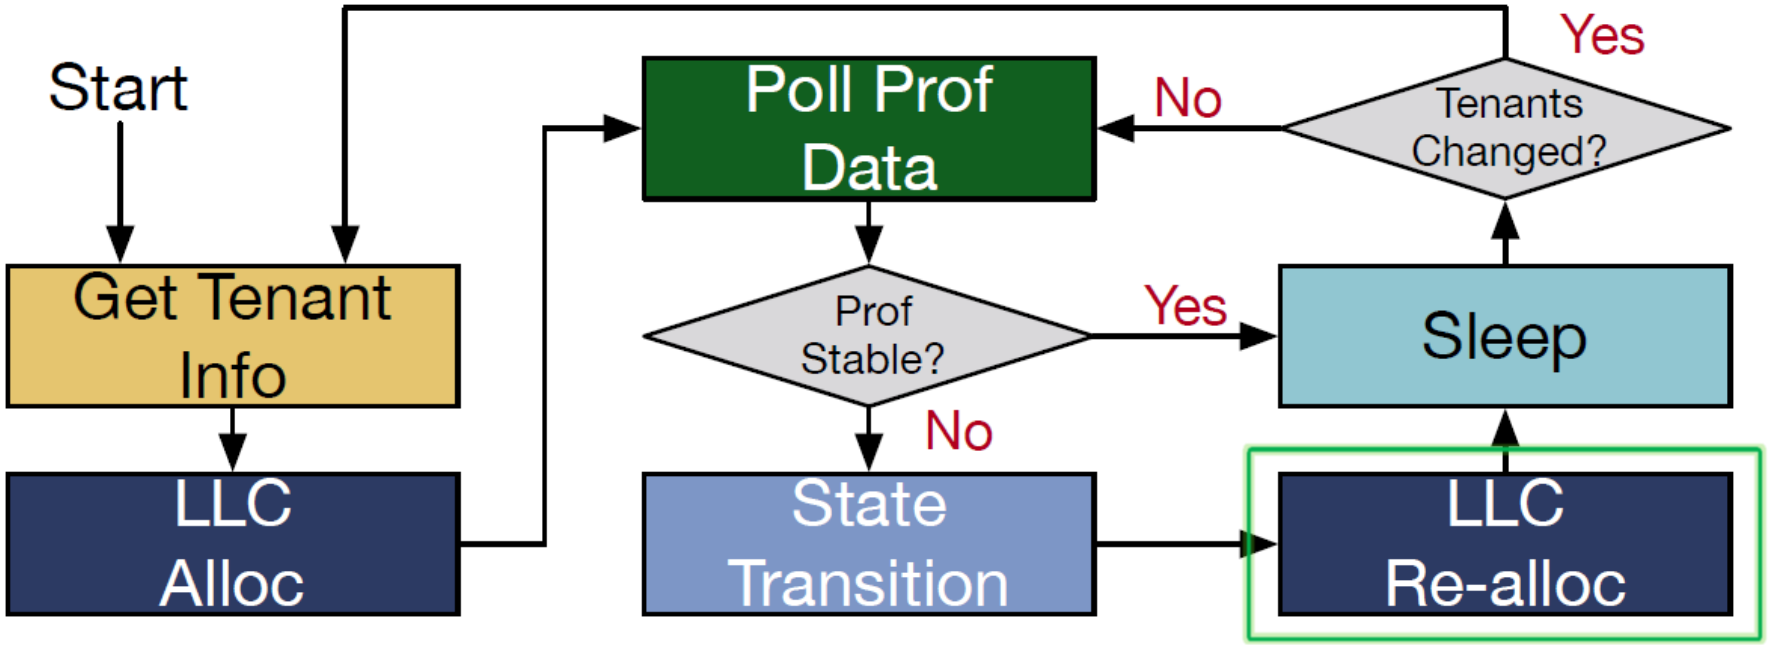

IAT: LLC Re – alloc (Example 1)

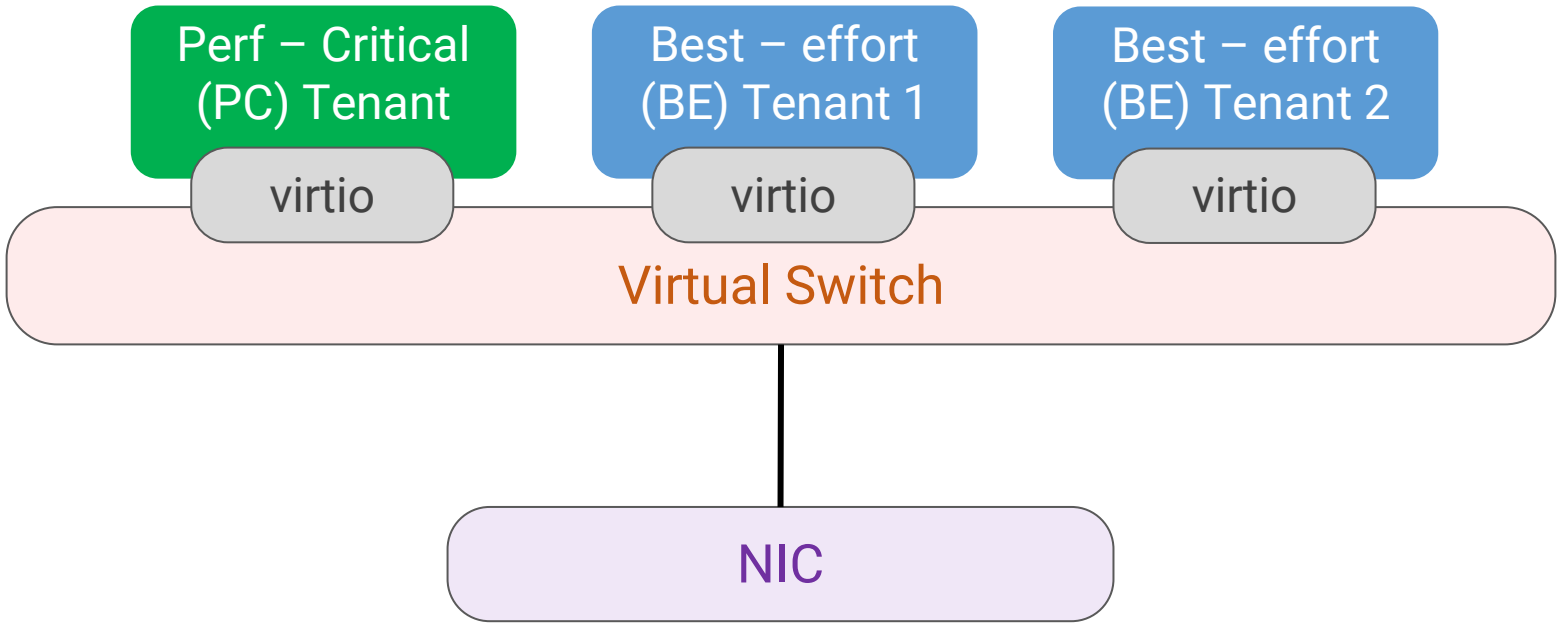

Example 1: Aggregation Model

IAT: LLC Re – alloc (Example 1)

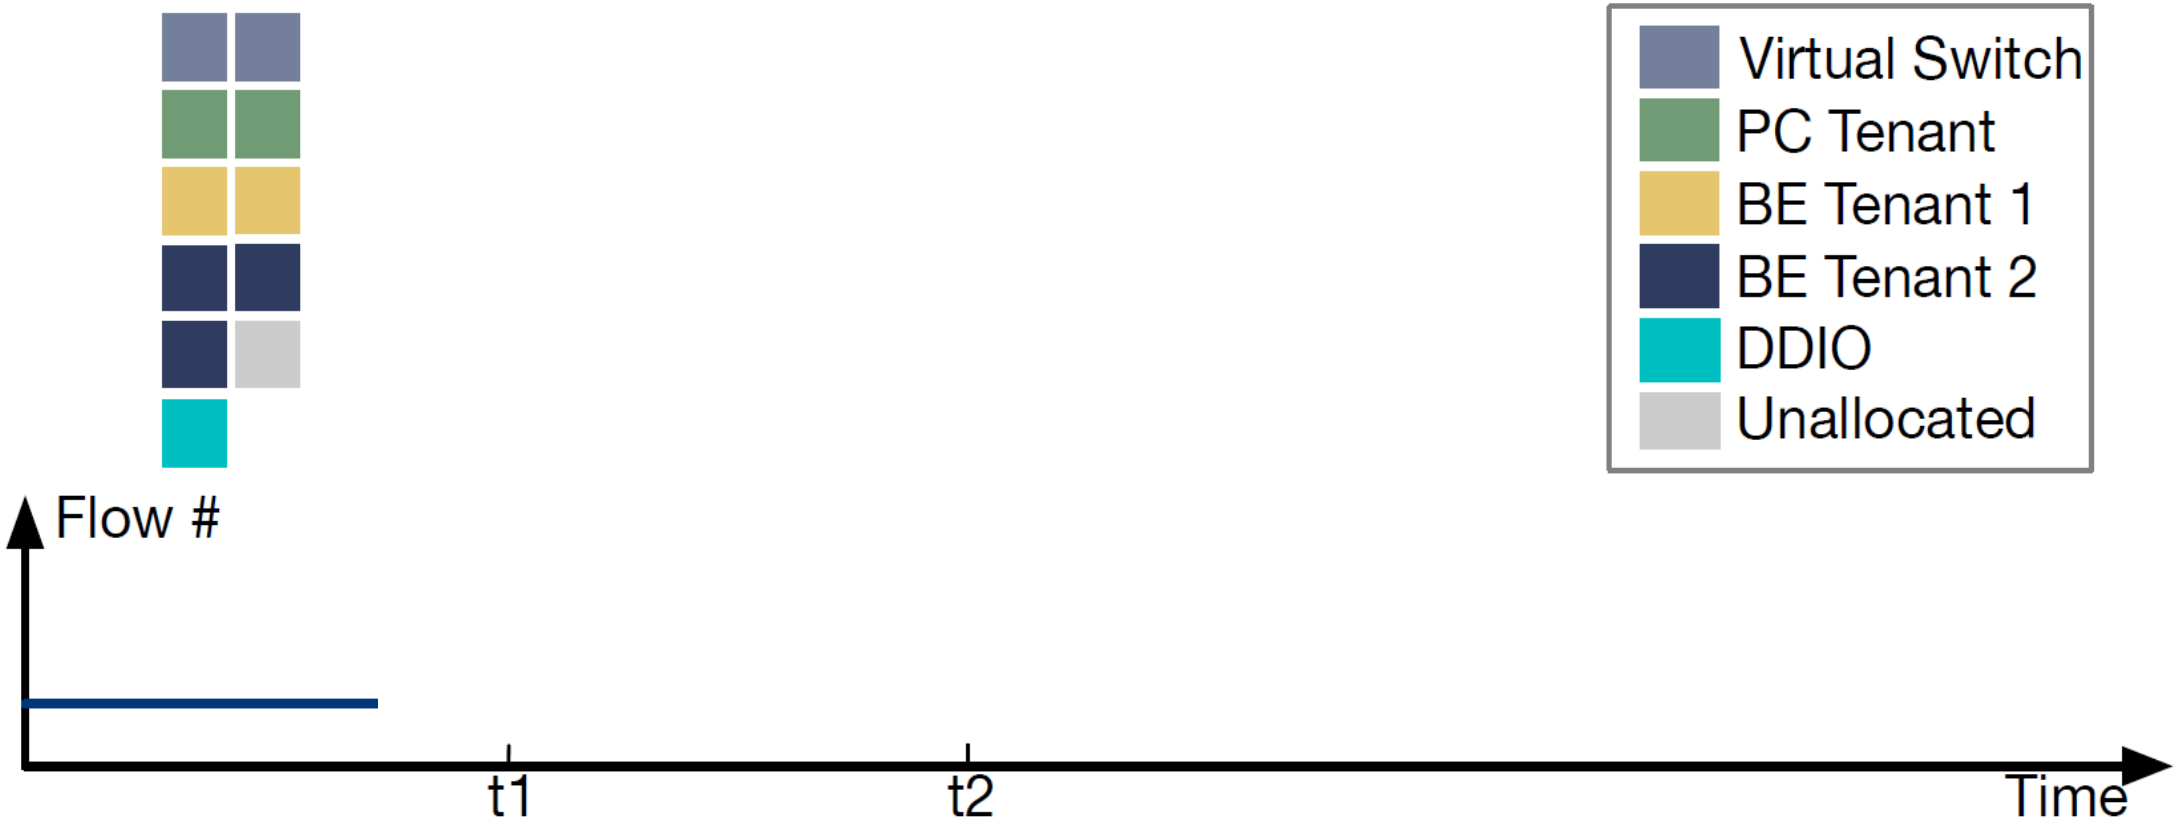

IAT: LLC Re – alloc (Example 1)

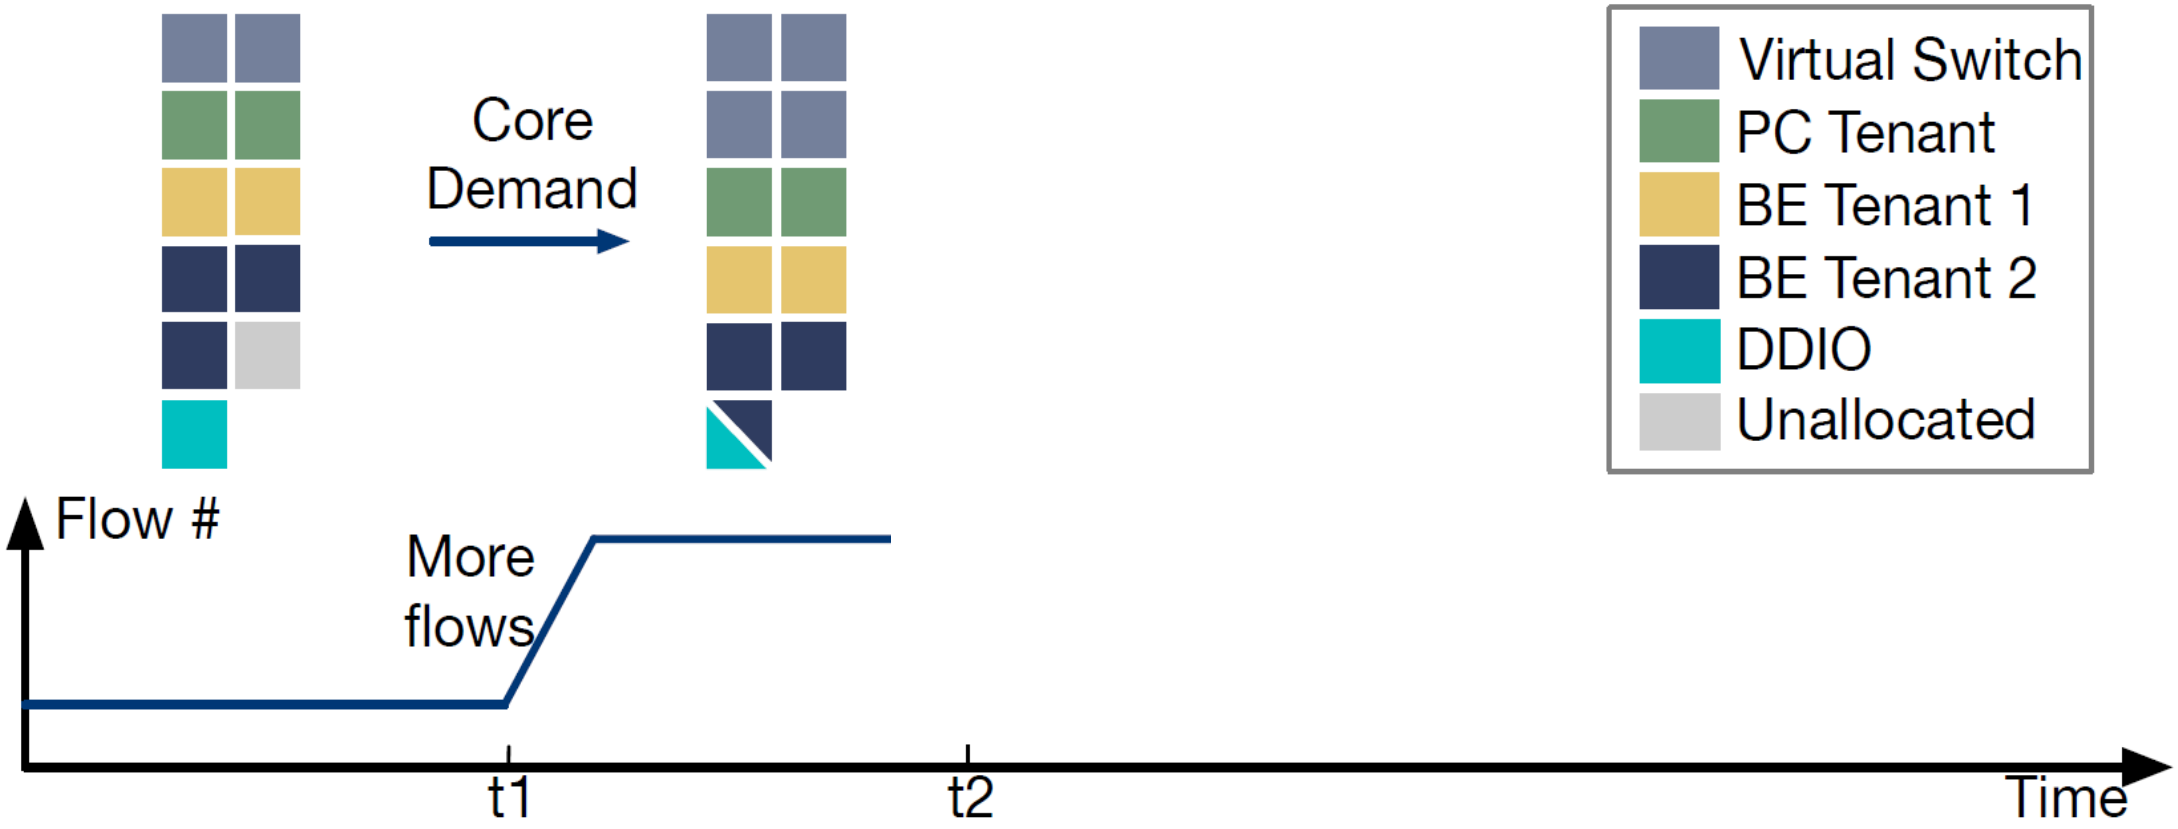

IAT: LLC Re – alloc (Example 1)

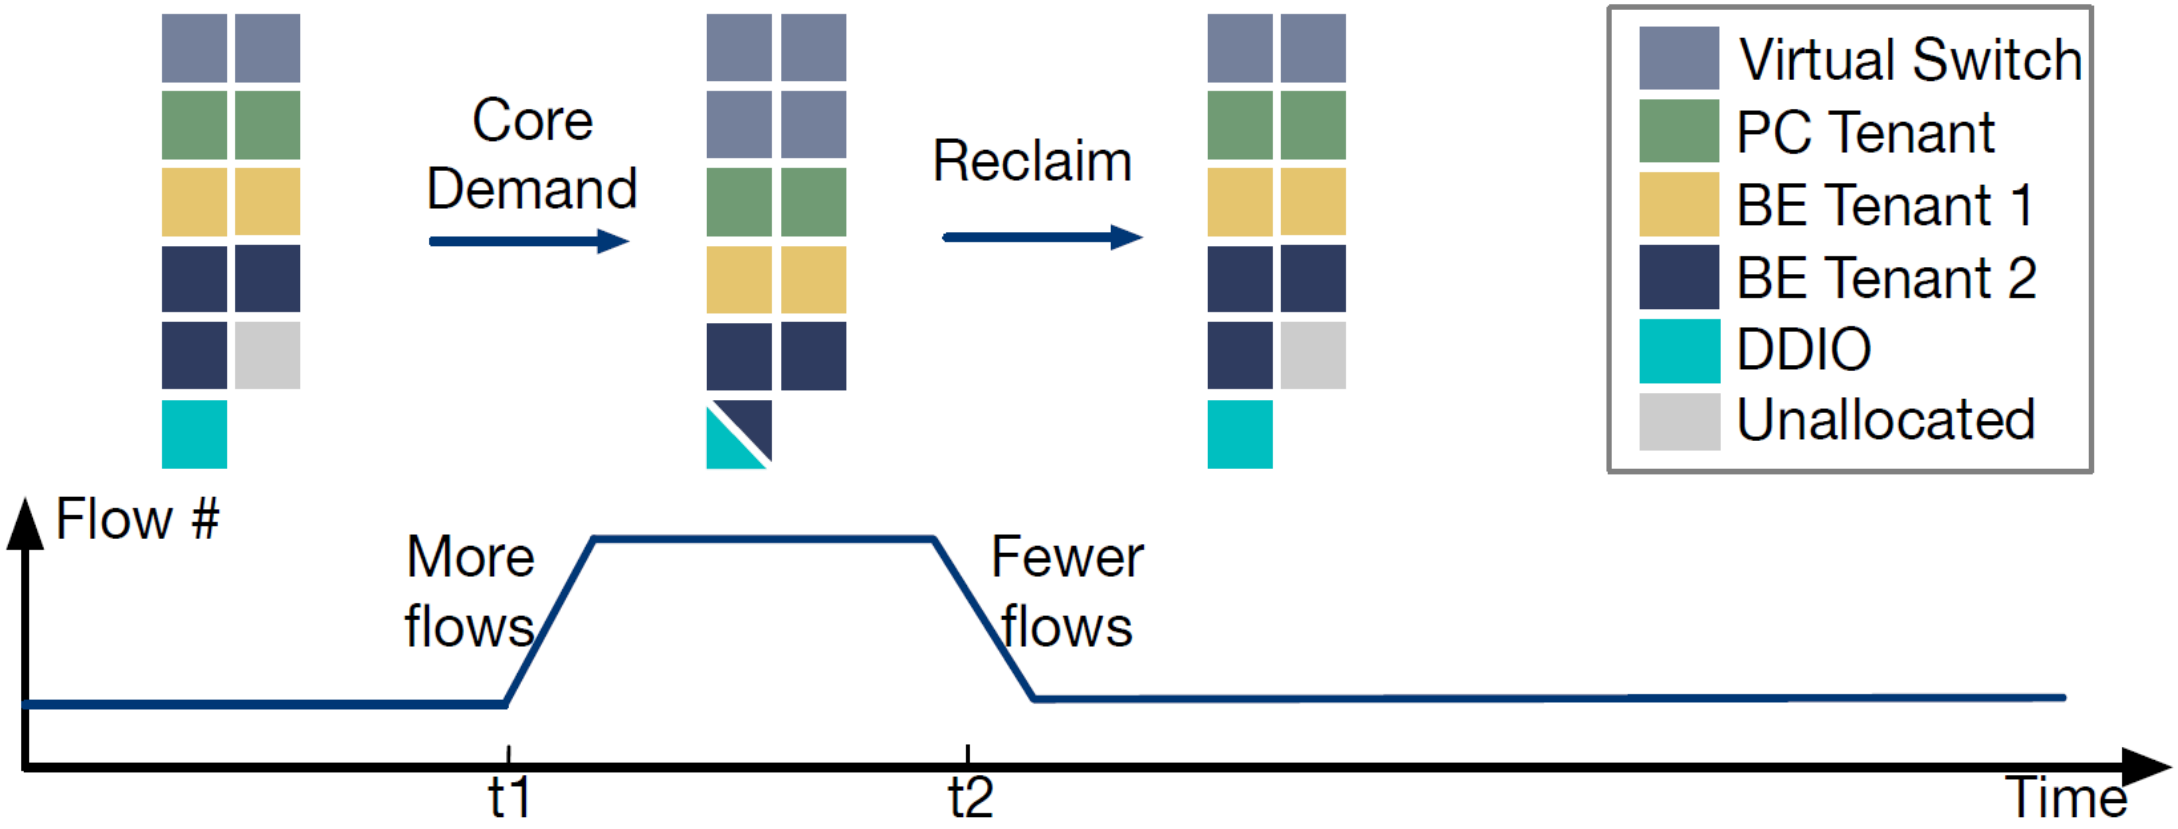

IAT: LLC Re – alloc (Example 2)

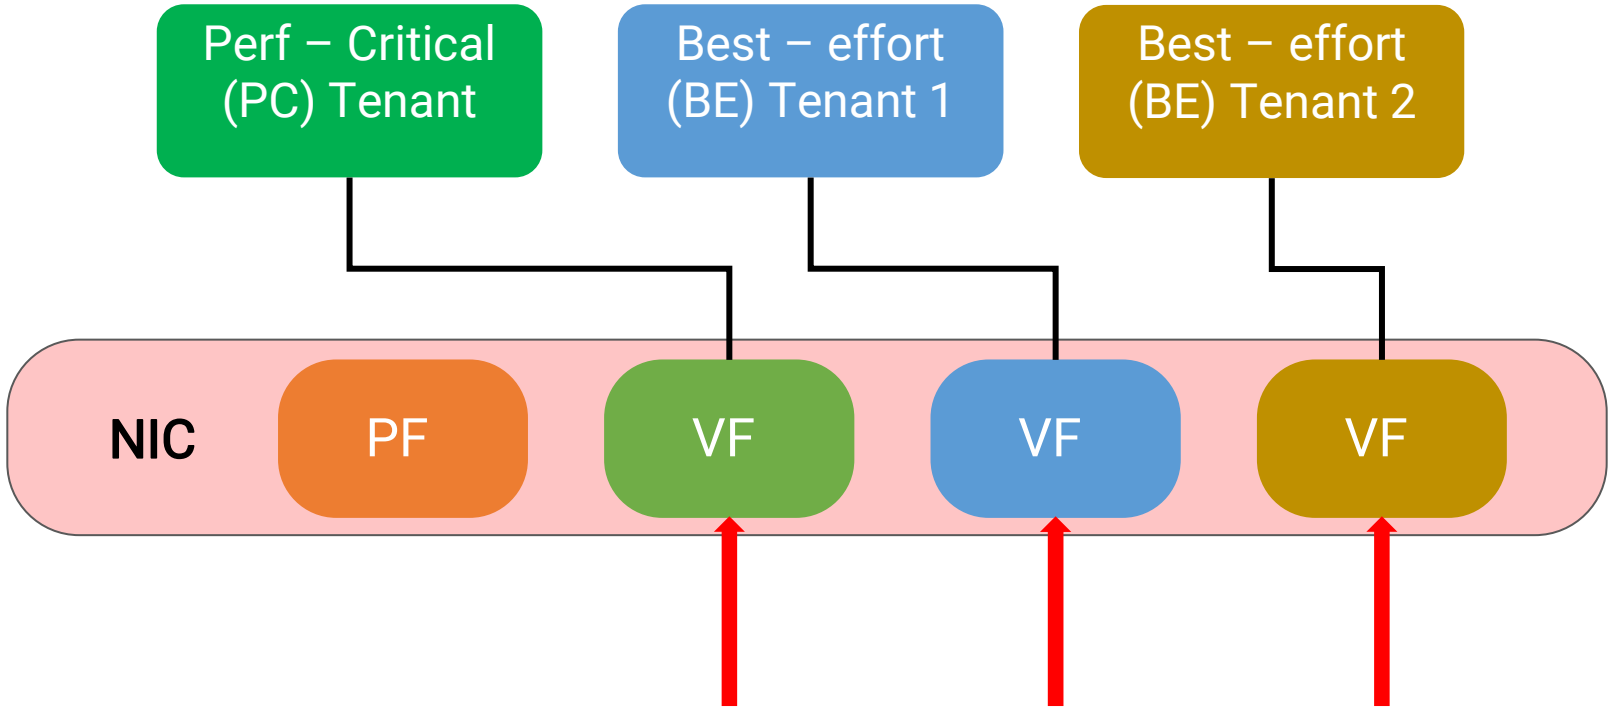

Example 2: Slicing Model

IAT: LLC Re – alloc (Example 2)

IAT: LLC Re – alloc (Example 2)

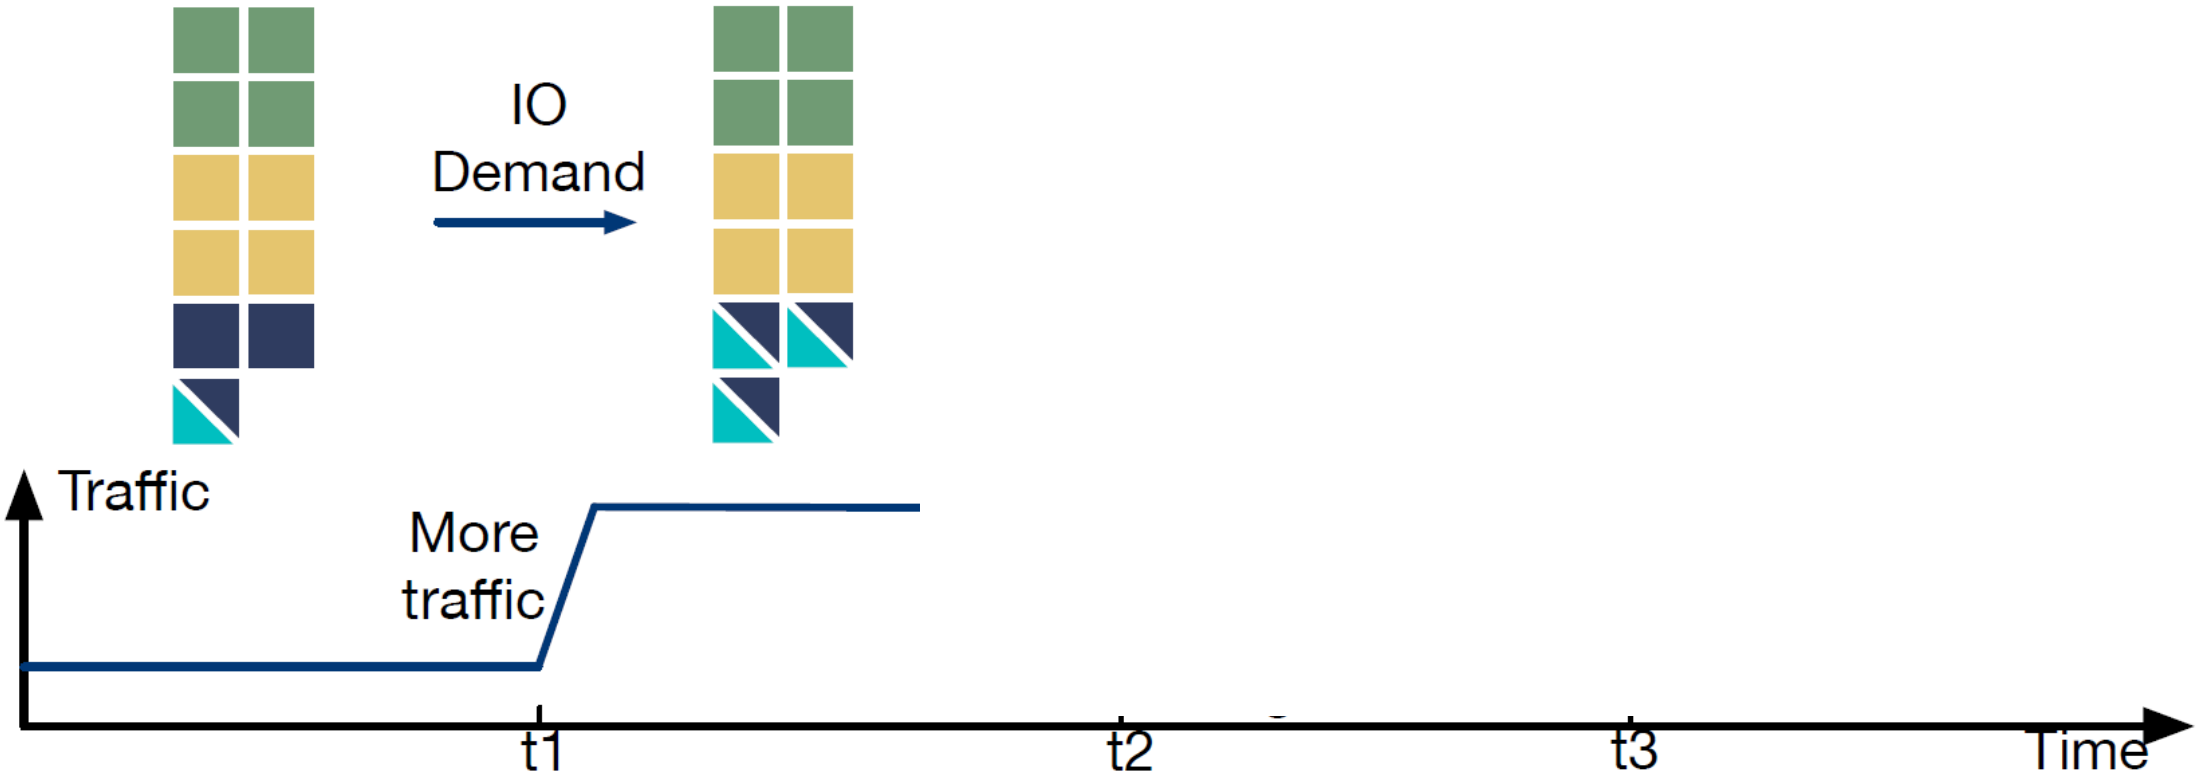

IAT: LLC Re – alloc (Example 2)

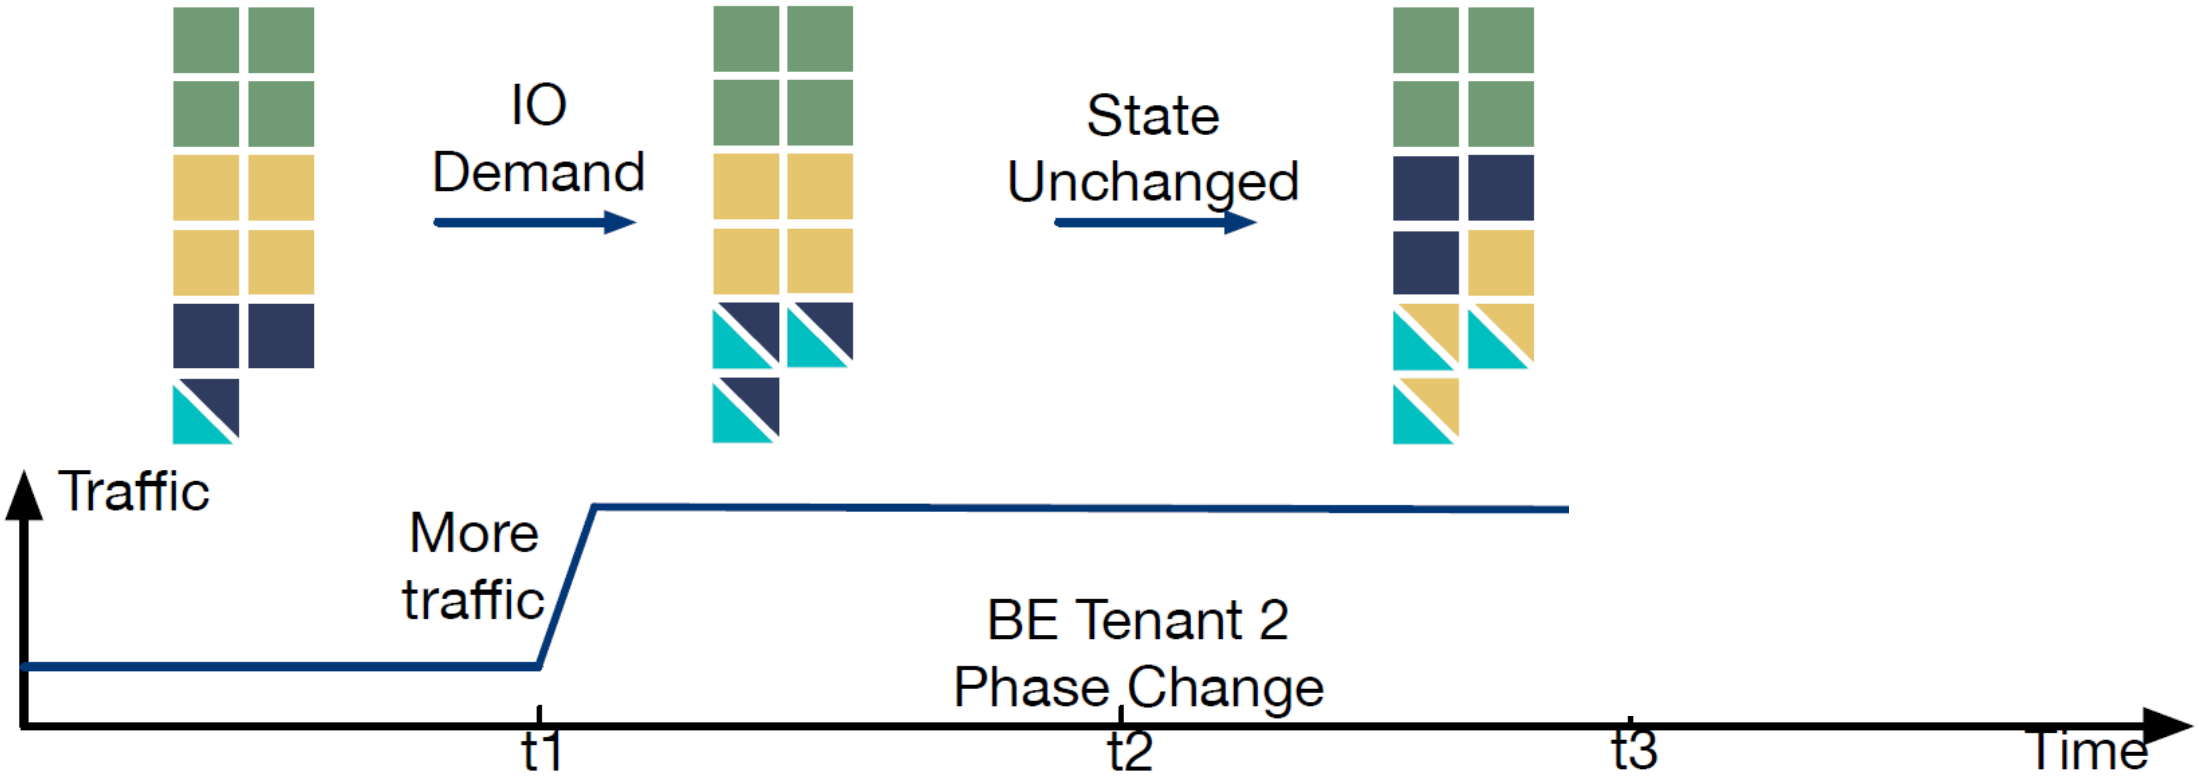

IAT: LLC Re – alloc (Example 2)

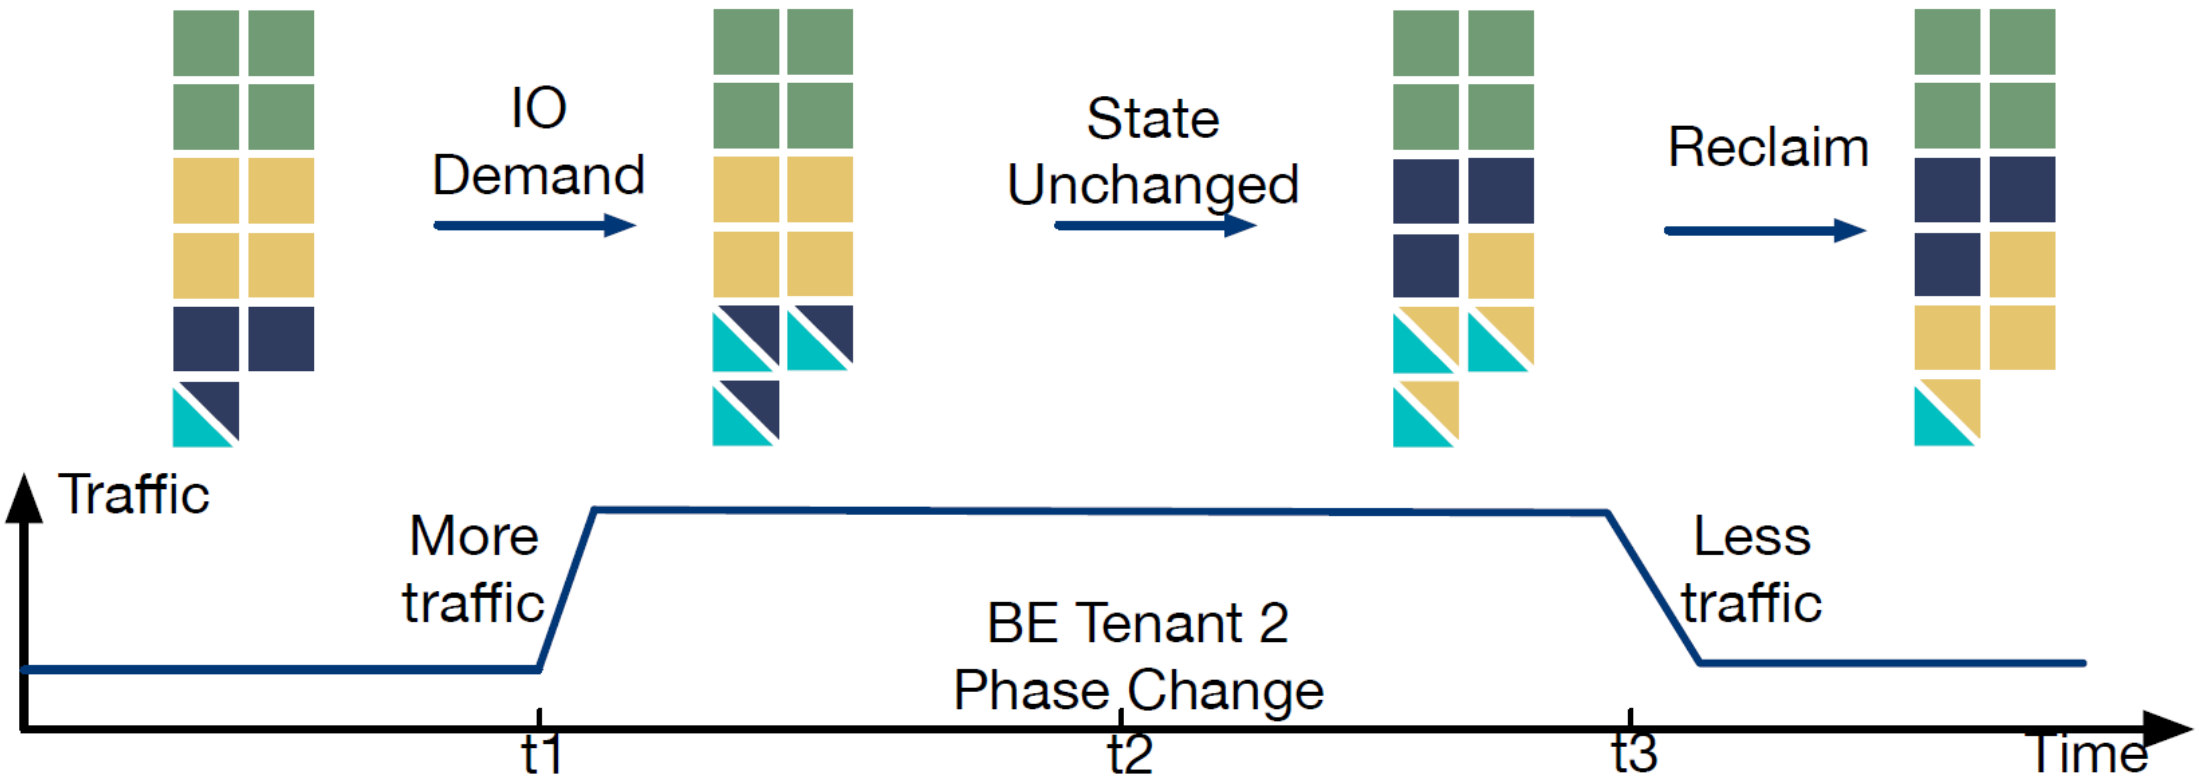

IAT: Solving Leaky DMA Problem

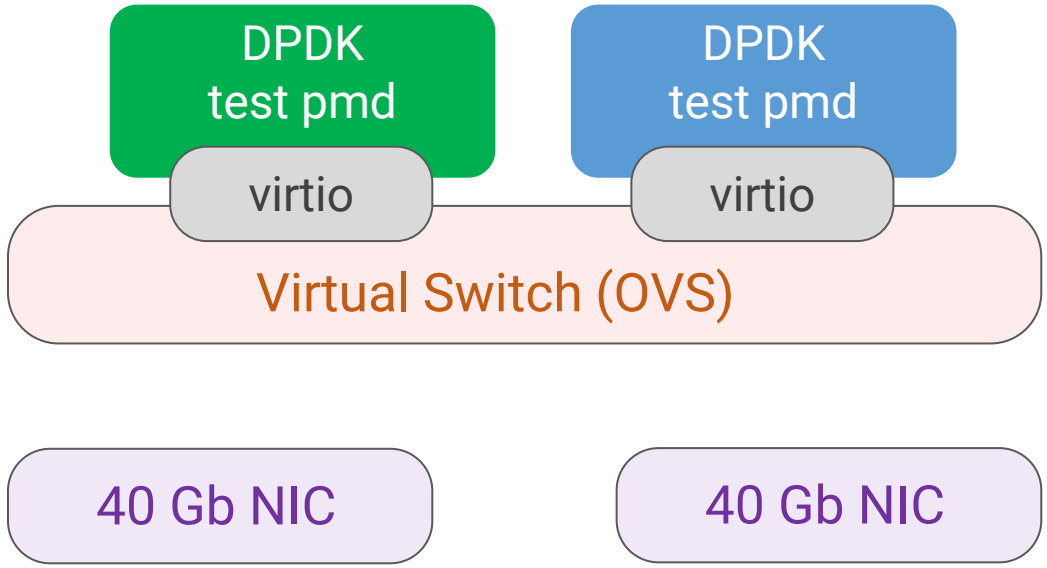

IAT: Solving Leaky DMA Problem

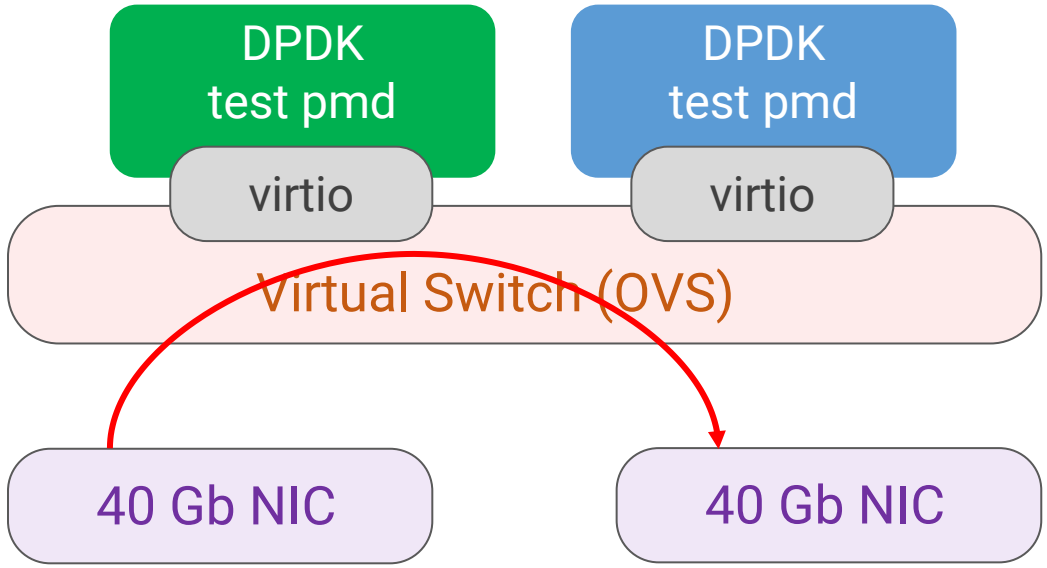

IAT: Solving Leaky DMA Problem

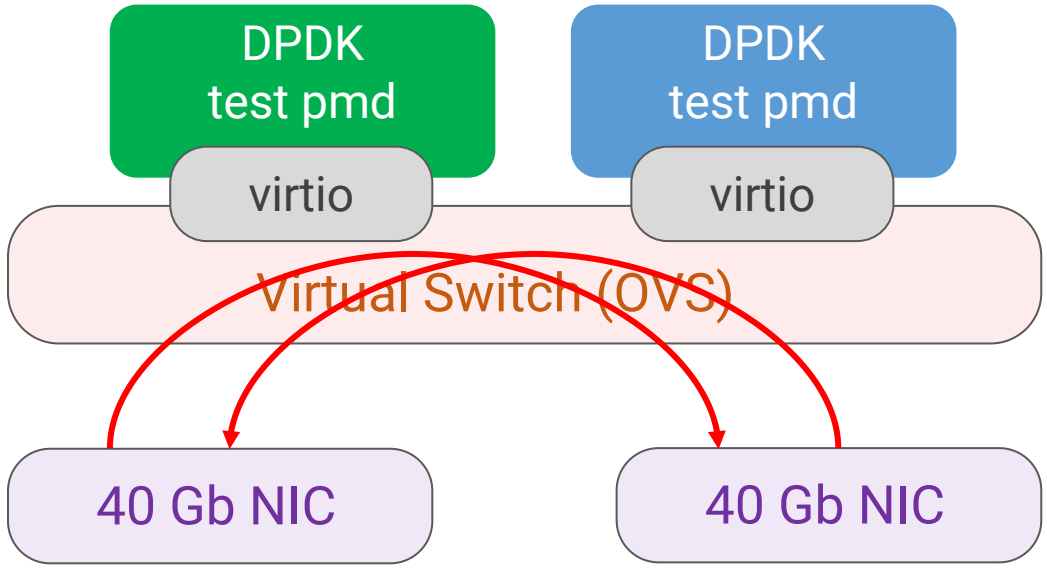

IAT: Solving Leaky DMA Problem

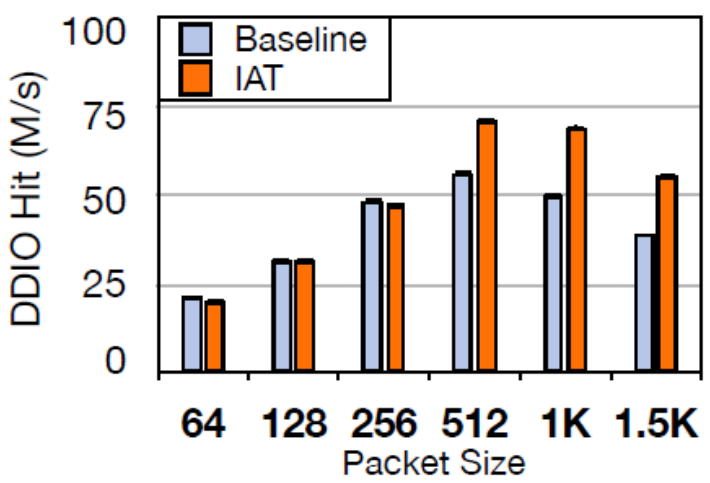

(a) DDIO hit count per second.

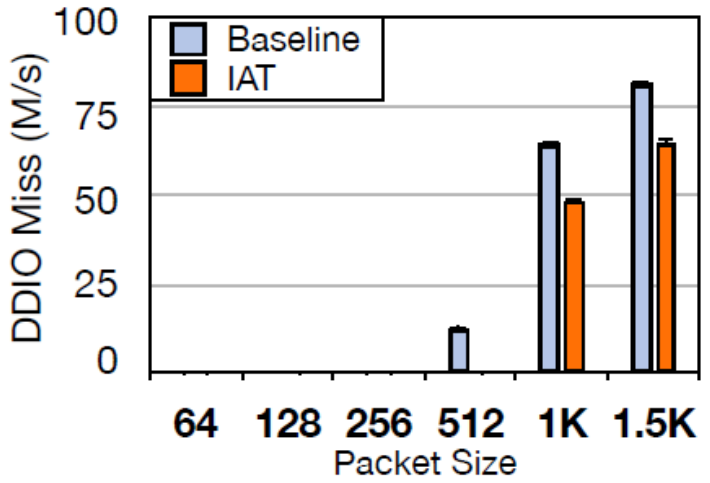

(b) DDIO miss count per second.

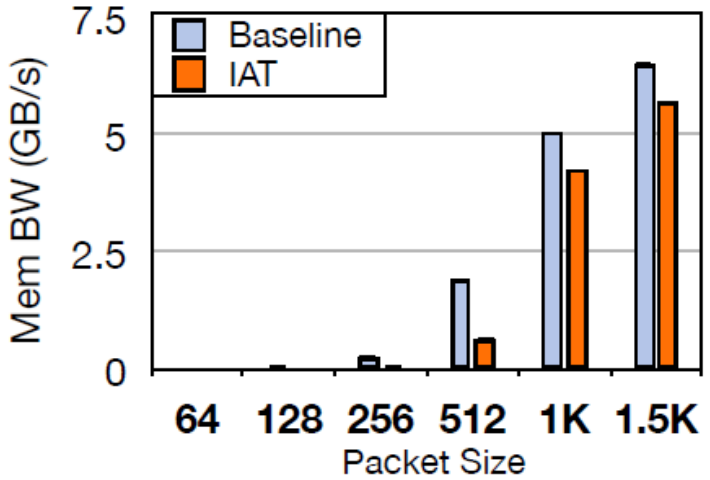

(c) Memory bandwidth consumption.

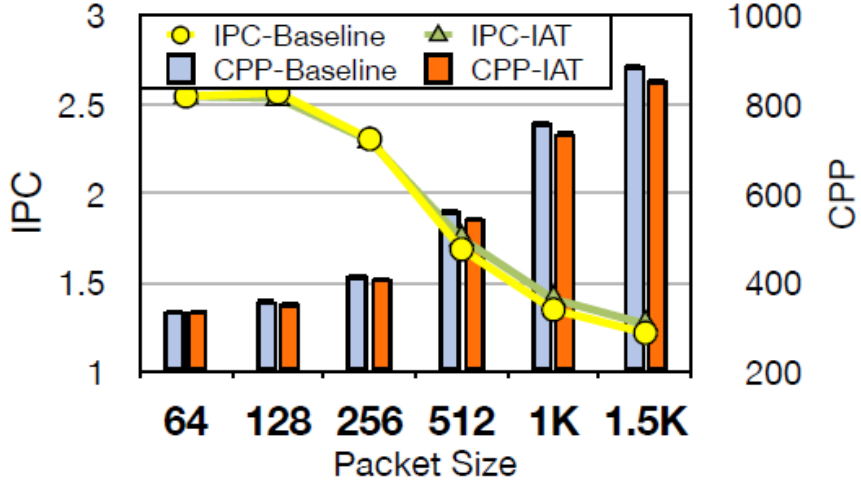

(d) OVS IPC and CPP.

IAT: Solving Latent Contender Problem

IAT: Solving Latent Contender Problem

- Change in X-Mem working set size and DDIO ways at different time interval

IAT: Solving Latent Contender Problem

- Change in X-Mem working set size and DDIO ways at different time interval

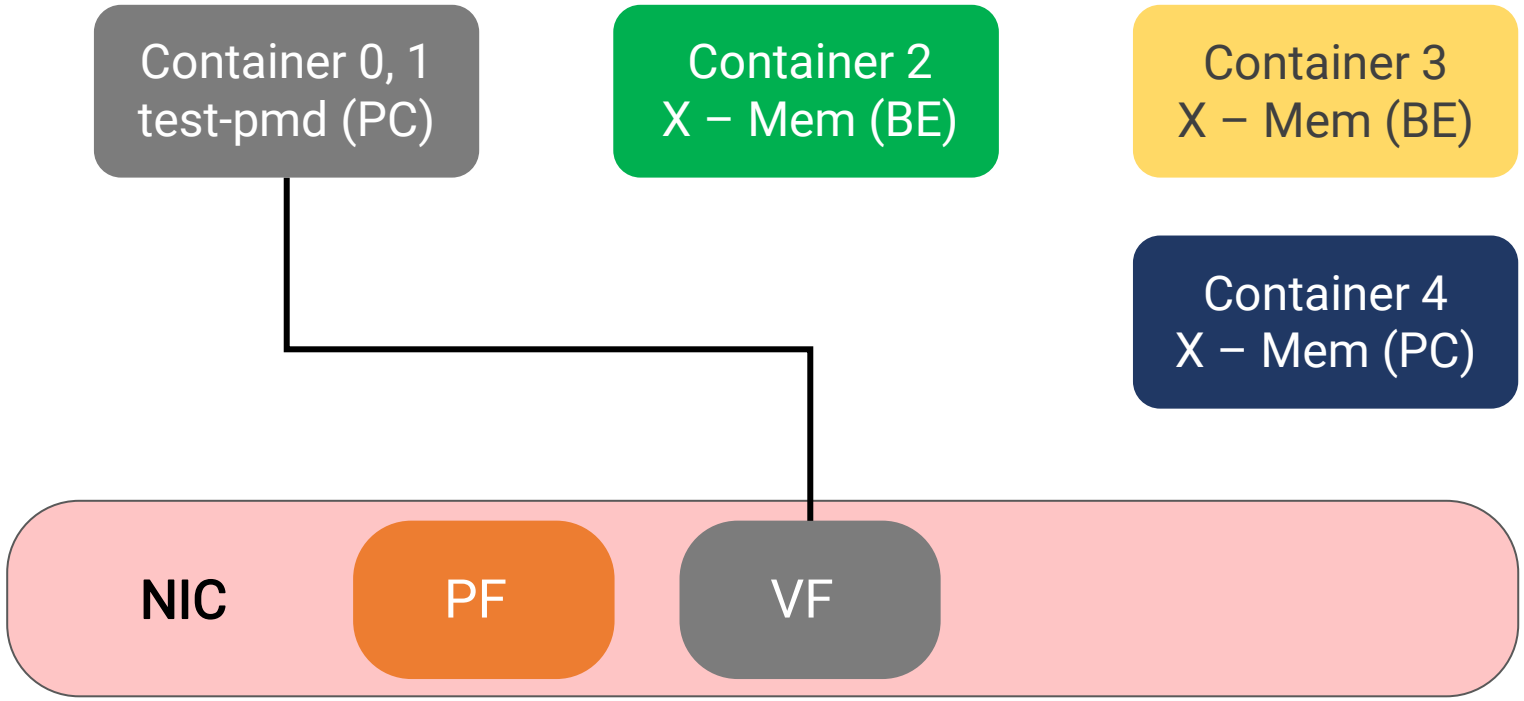

IAT: Solving Latent Contender Problem

- Change in X-Mem working set size and DDIO ways at different time interval: **Container 4**

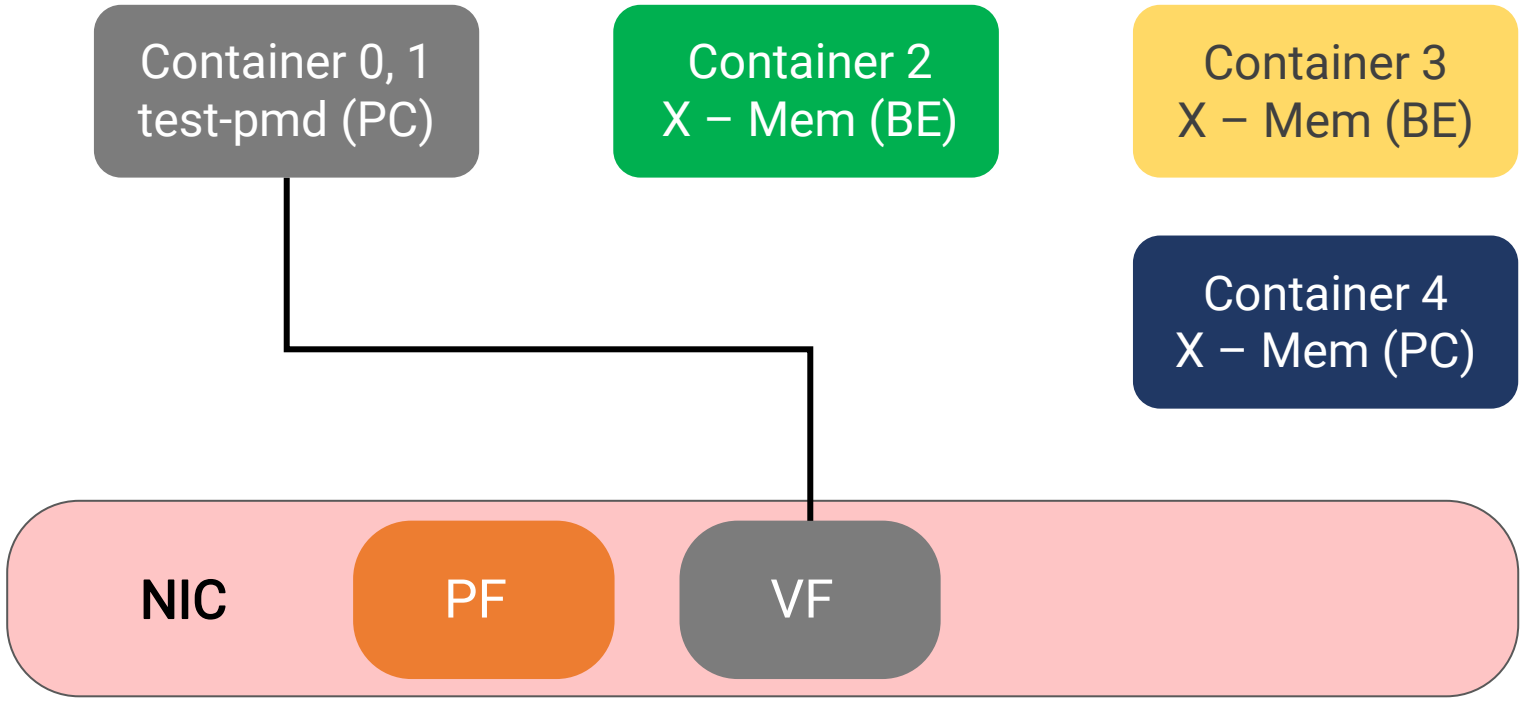

IAT: Solving Latent Contender Problem

- Change in X-Mem working set size and DDIO ways at different time interval: **Container 4**

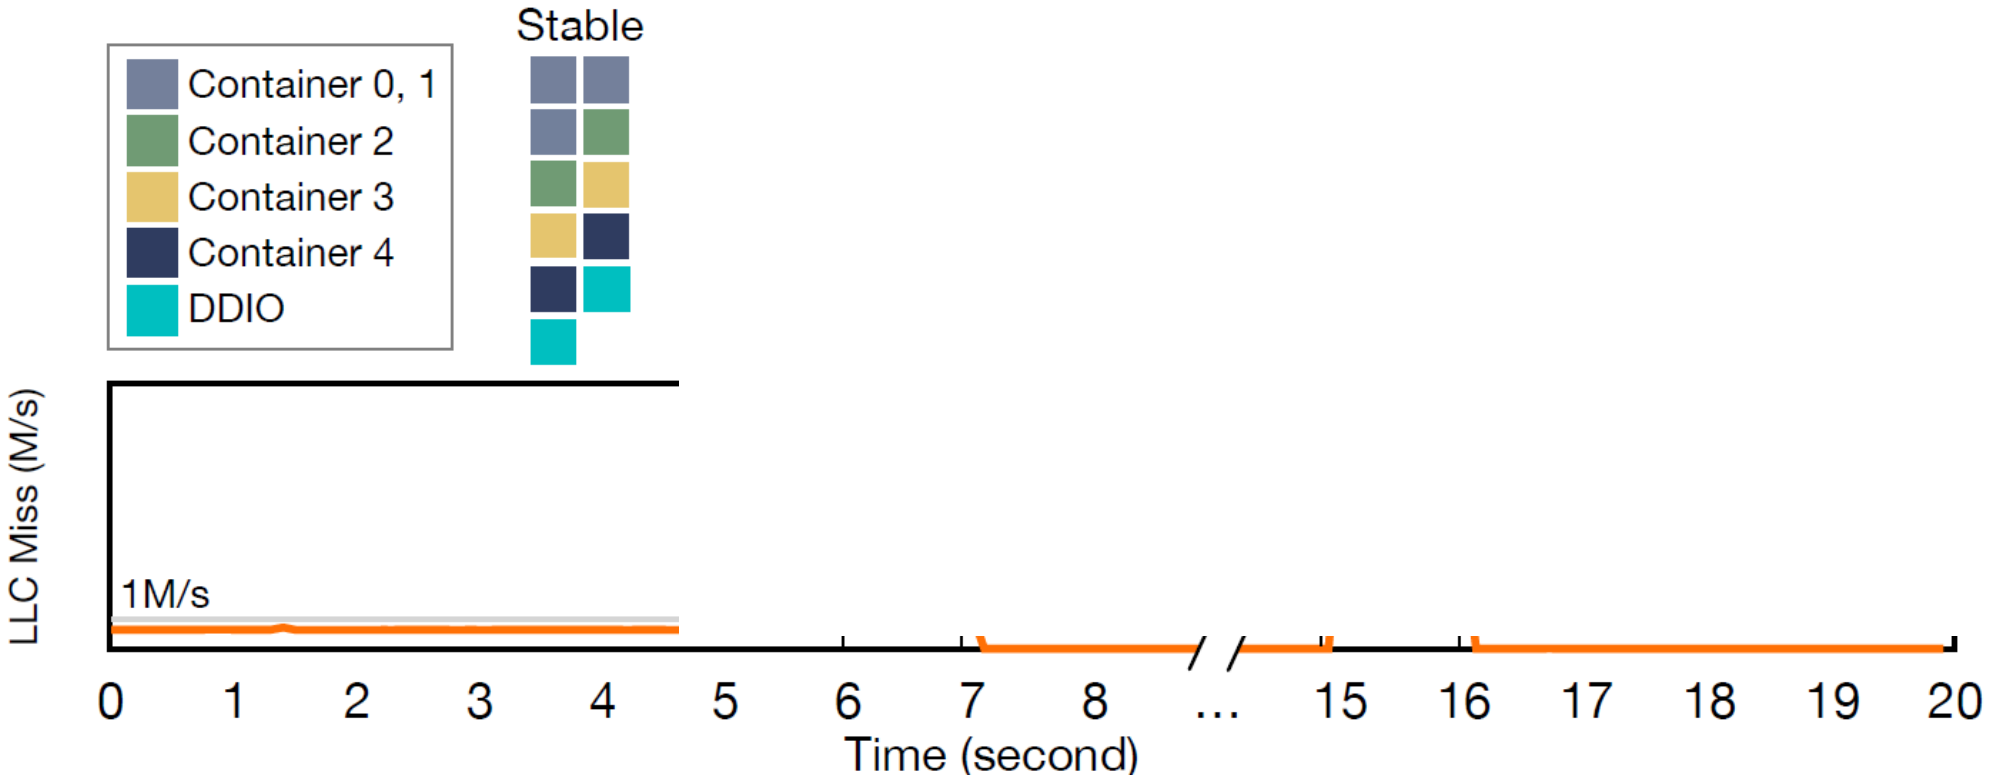

IAT: Solving Latent Contender Problem

- Change in X-Mem working set size and DDIO ways at different time interval: **Container 4**

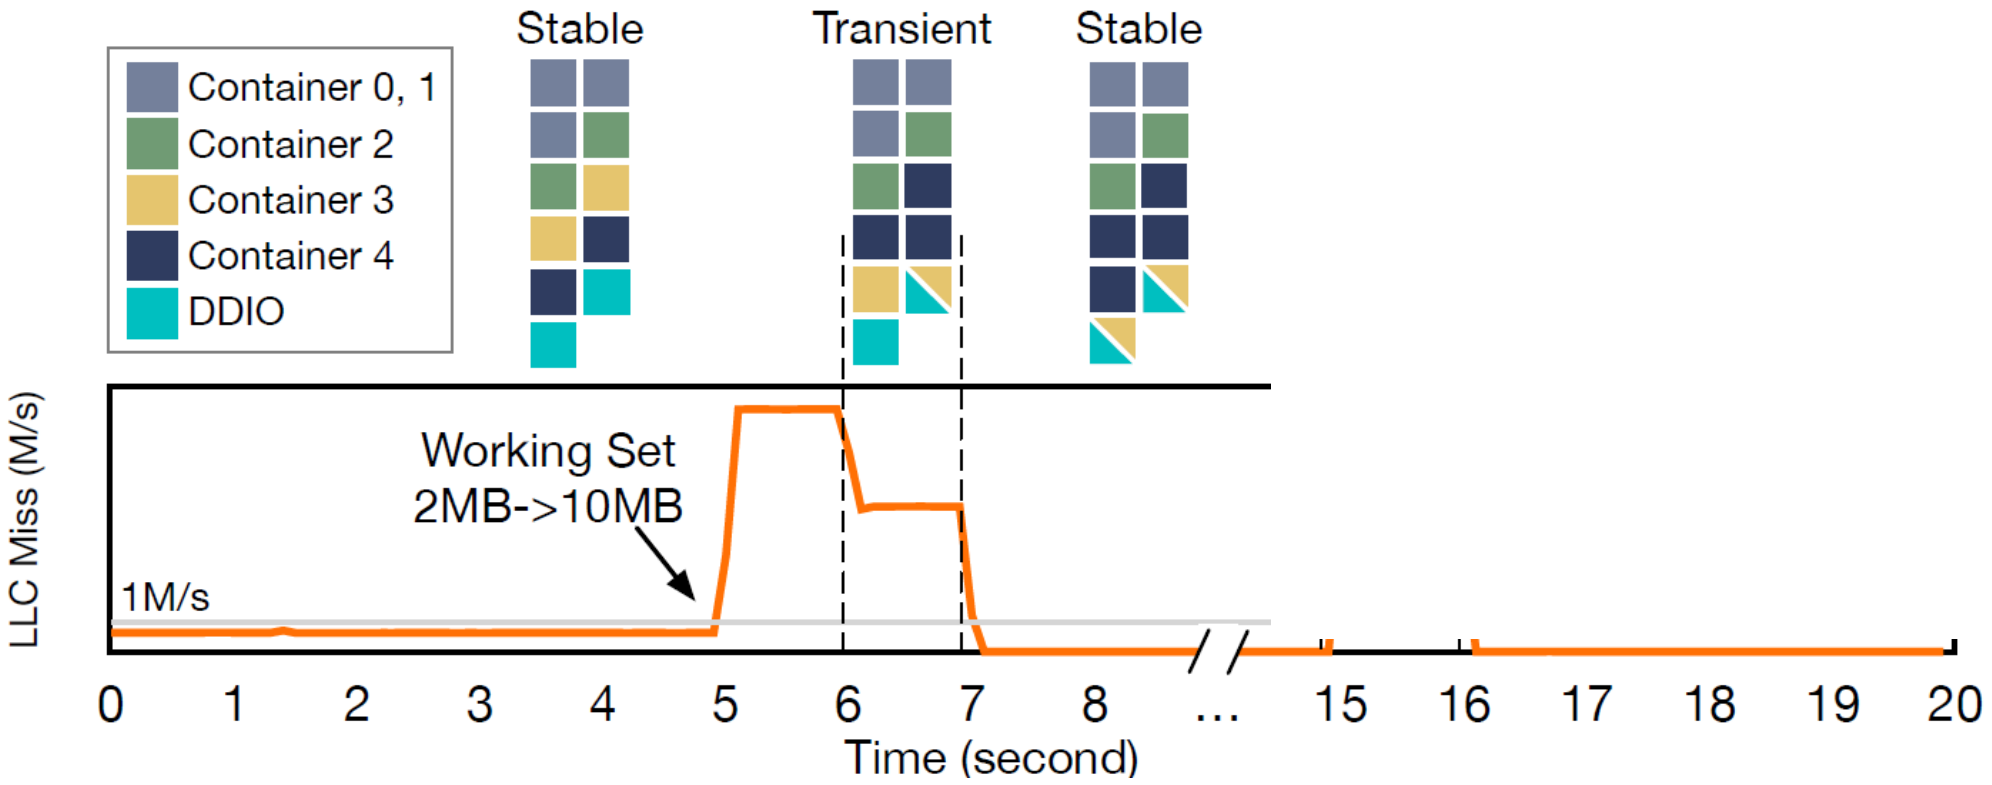

IAT: Solving Latent Contender Problem

- Change in X-Mem working set size and DDIO ways at different time interval: **Container 4**

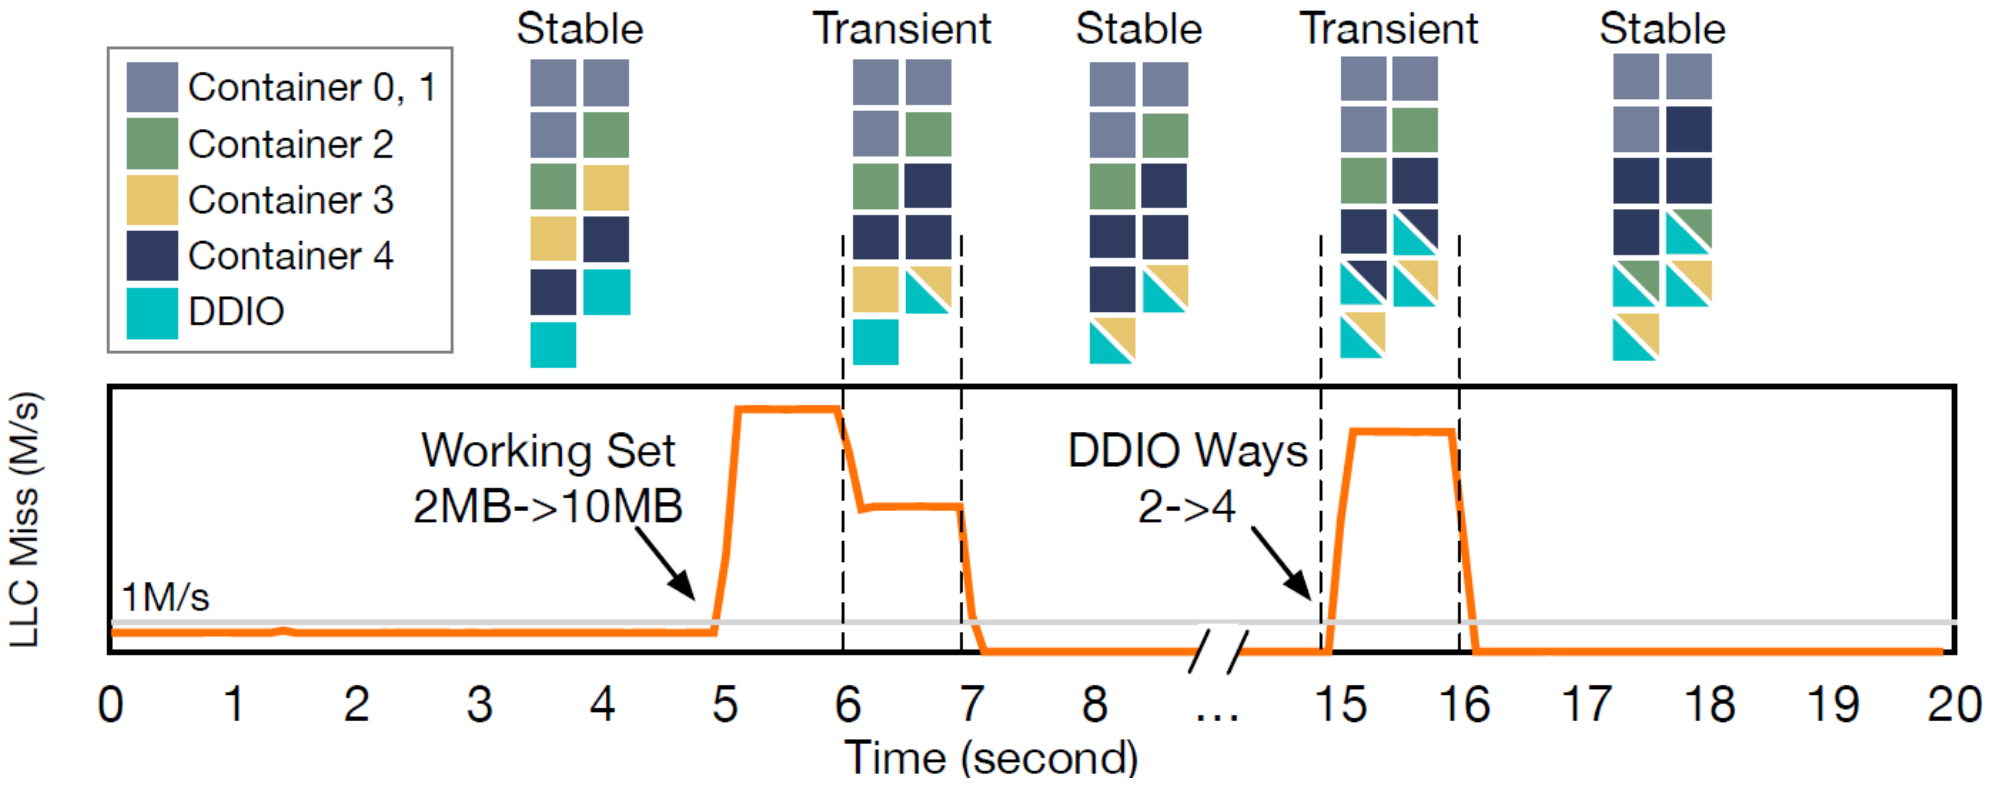

IAT: Solving Latent Contender Problem

- Change in X-Mem working set size and DDIO ways at different time interval: **Container 4**

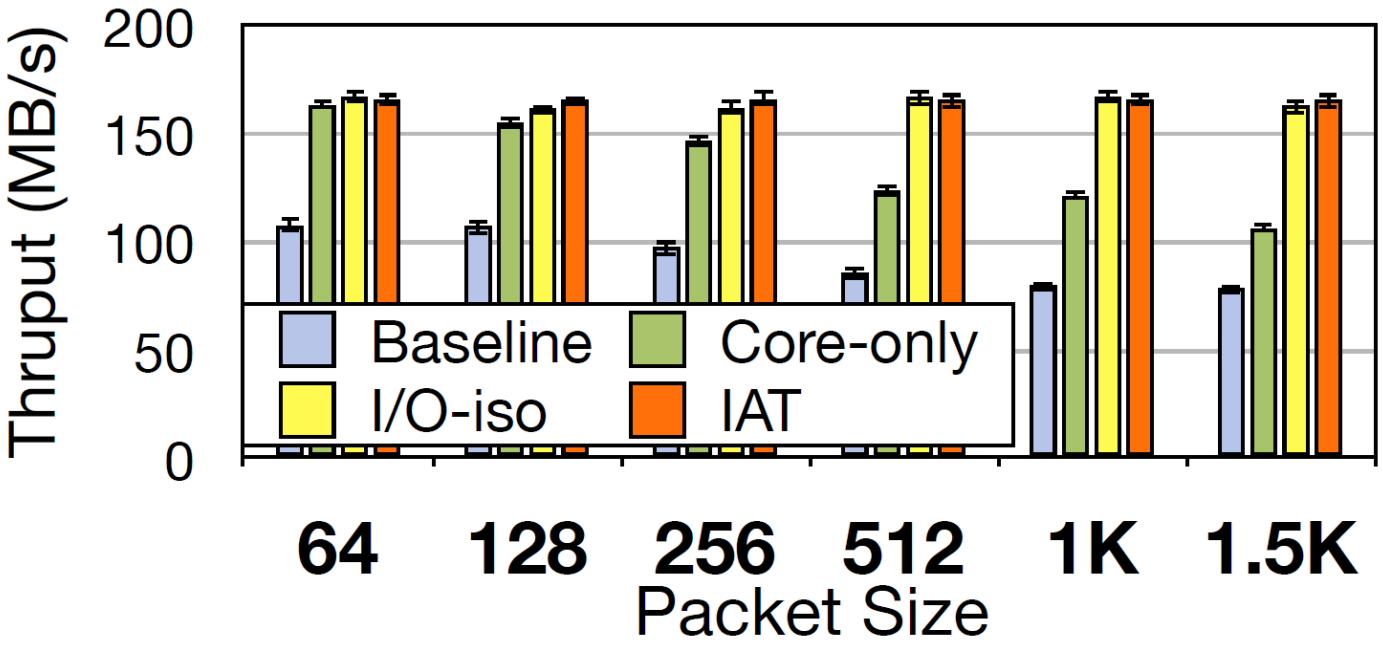

(a) Throughput between 5s and 15s.

IAT: Solving Latent Contender Problem

- Change in X-Mem working set size and DDIO ways at different time interval: **Container 4**

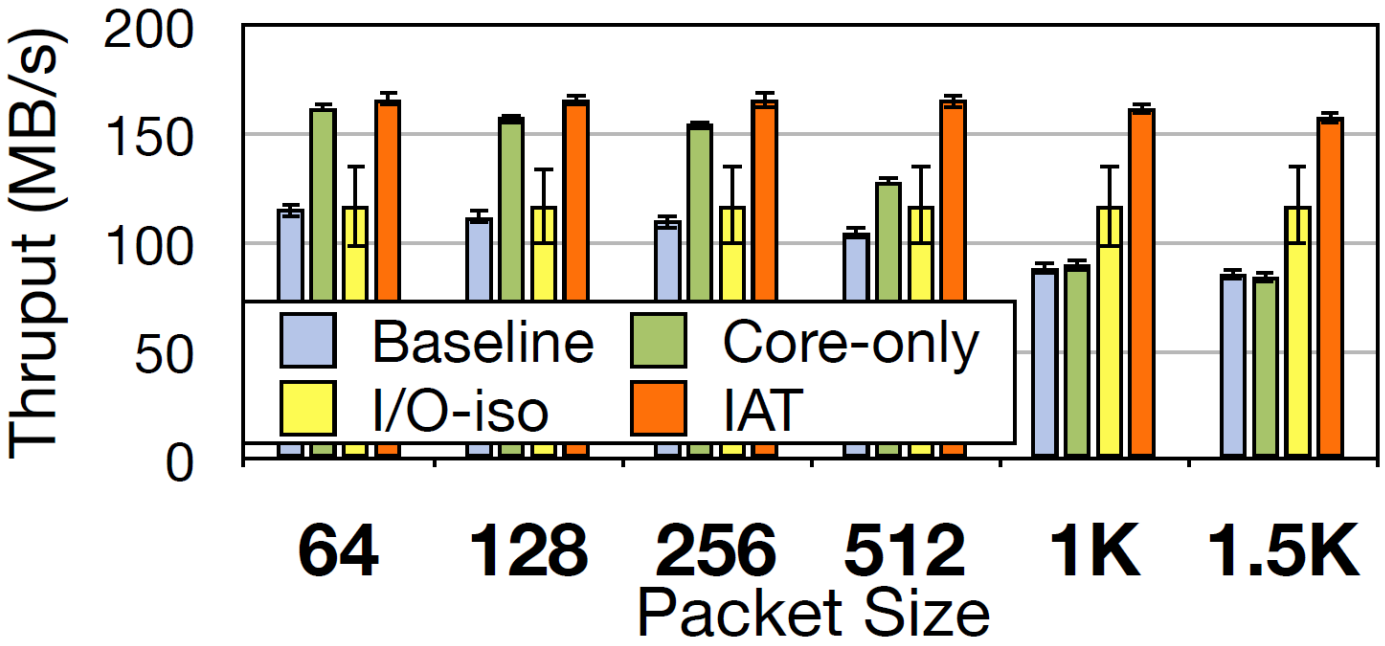

(c) Throughput after 15s.

IAT: Overhead

IAT: Overhead

- Implemented as user – space daemon in Ubuntu

IAT: Overhead

- Implemented as user – space daemon in Ubuntu
- IAT execution time is measured in following cases:
  - One vs two cores per tenant
  - Stable: no need to re – allocate LLC (only `Poll Prof Data` time)
  - Unstable: need to take action (`PPD + ST + LLC Re – alloc` time)

IAT: Overhead

- Implemented as user – space daemon in Ubuntu
- IAT execution time is measured in following cases:
  - One vs two cores per tenant
  - Stable: no need to re – allocate LLC (only `Poll Prof Data` time)
  - Unstable: need to take action (`PPD + ST + LLC Re – alloc` time)

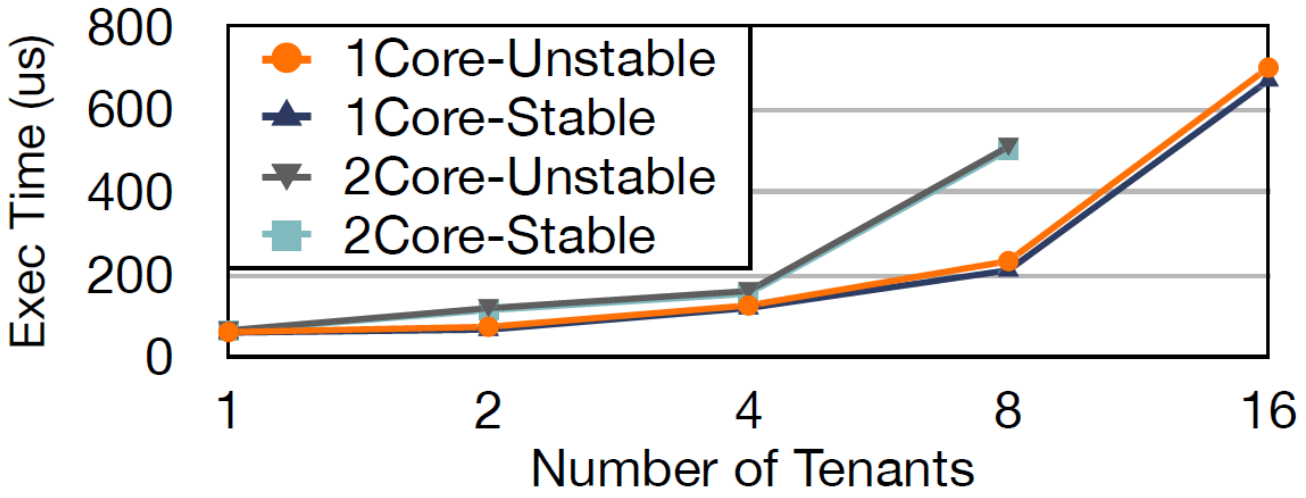

IAT: Overhead

- Implemented as user – space daemon in Ubuntu
- IAT execution time is measured in following cases:
  - One vs two cores per tenant
  - Stable: no need to re – allocate LLC (only `Poll Prof Data` time)
  - Unstable: need to take action (`PPD + ST + LLC Re – alloc` time)

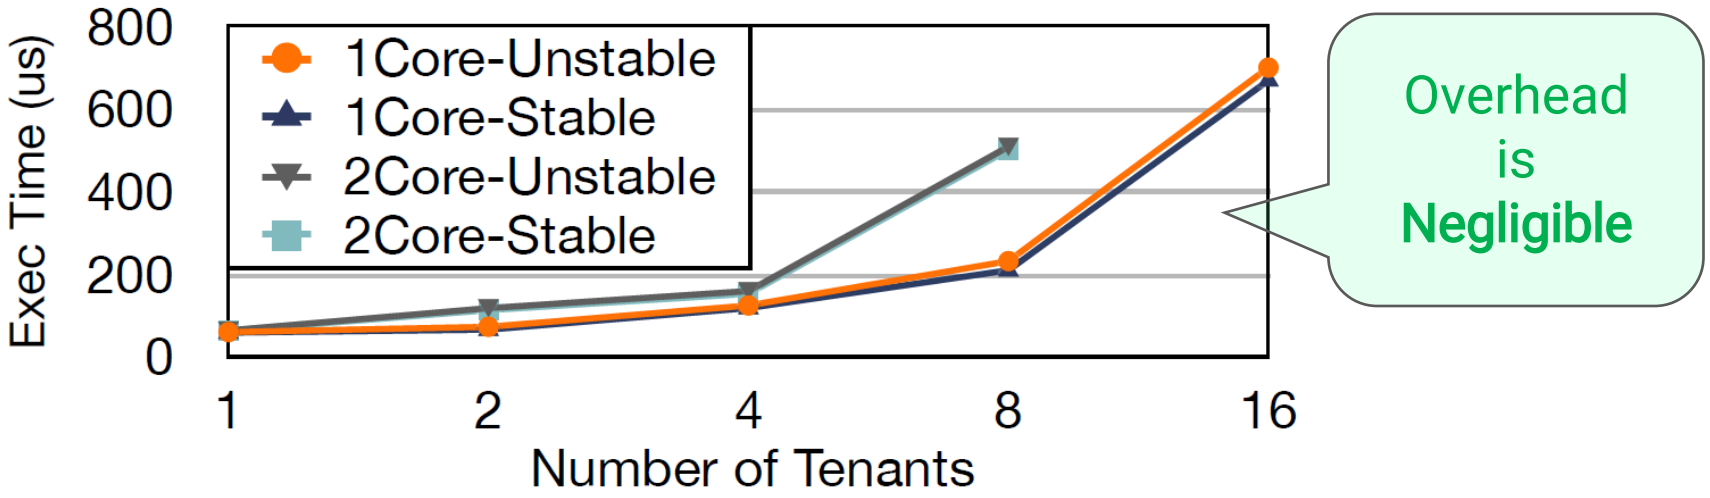

Conclusion

- DDIO introduces new challenges for better LLC management (IO aware)

Conclusion

- DDIO introduces new challenges for better LLC management (IO aware)
- Both the problems, Leaky DMA and Latent Contender, are solved by IAT

Conclusion

- DDIO introduces new challenges for better LLC management (IO aware)
- Both the problems, Leaky DMA and Latent Contender, are solved by IAT
- IAT reduces interference caused by DDIO with negligible overhead

Questions ?

Thank You