



# **CS698Y: Modern Memory Systems**

## **Lecture-11 (Hardware Prefetching)**

---

**Biswabandan Panda**

**[biswap@cse.iitk.ac.in](mailto:biswap@cse.iitk.ac.in)**

**<https://www.cse.iitk.ac.in/users/biswap/CS698Y.html>**

# Flow of the Module

**Data Prefetching Techniques**

**Metrics Related to Prefetching**

**Interaction with Cache Replacement**

**Instruction Prefetching**

**But, Why Prefetching? Remember Memory Wall: It is still hurting**

# Hardware Prefetching

*What?*

**Latency-hiding technique** - Fetches data before the core demands.

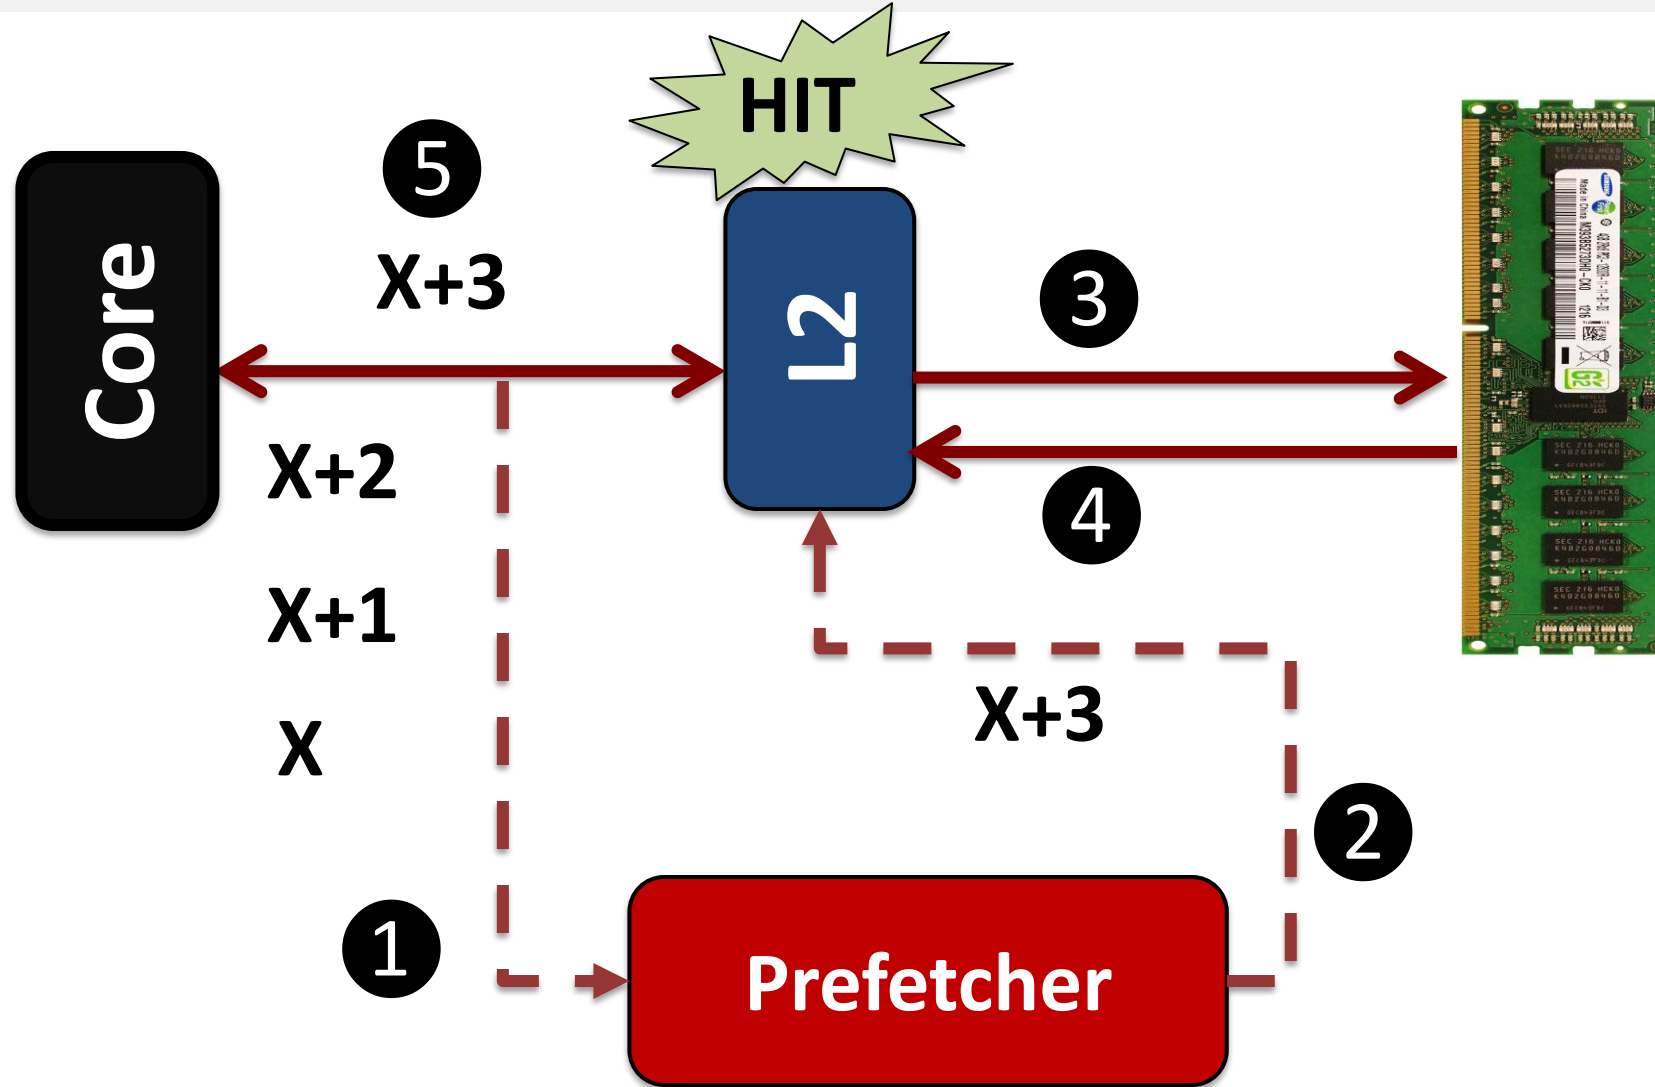
*Why?*

Off-chip DRAM latency has grown up to 400 to 800 cycles.

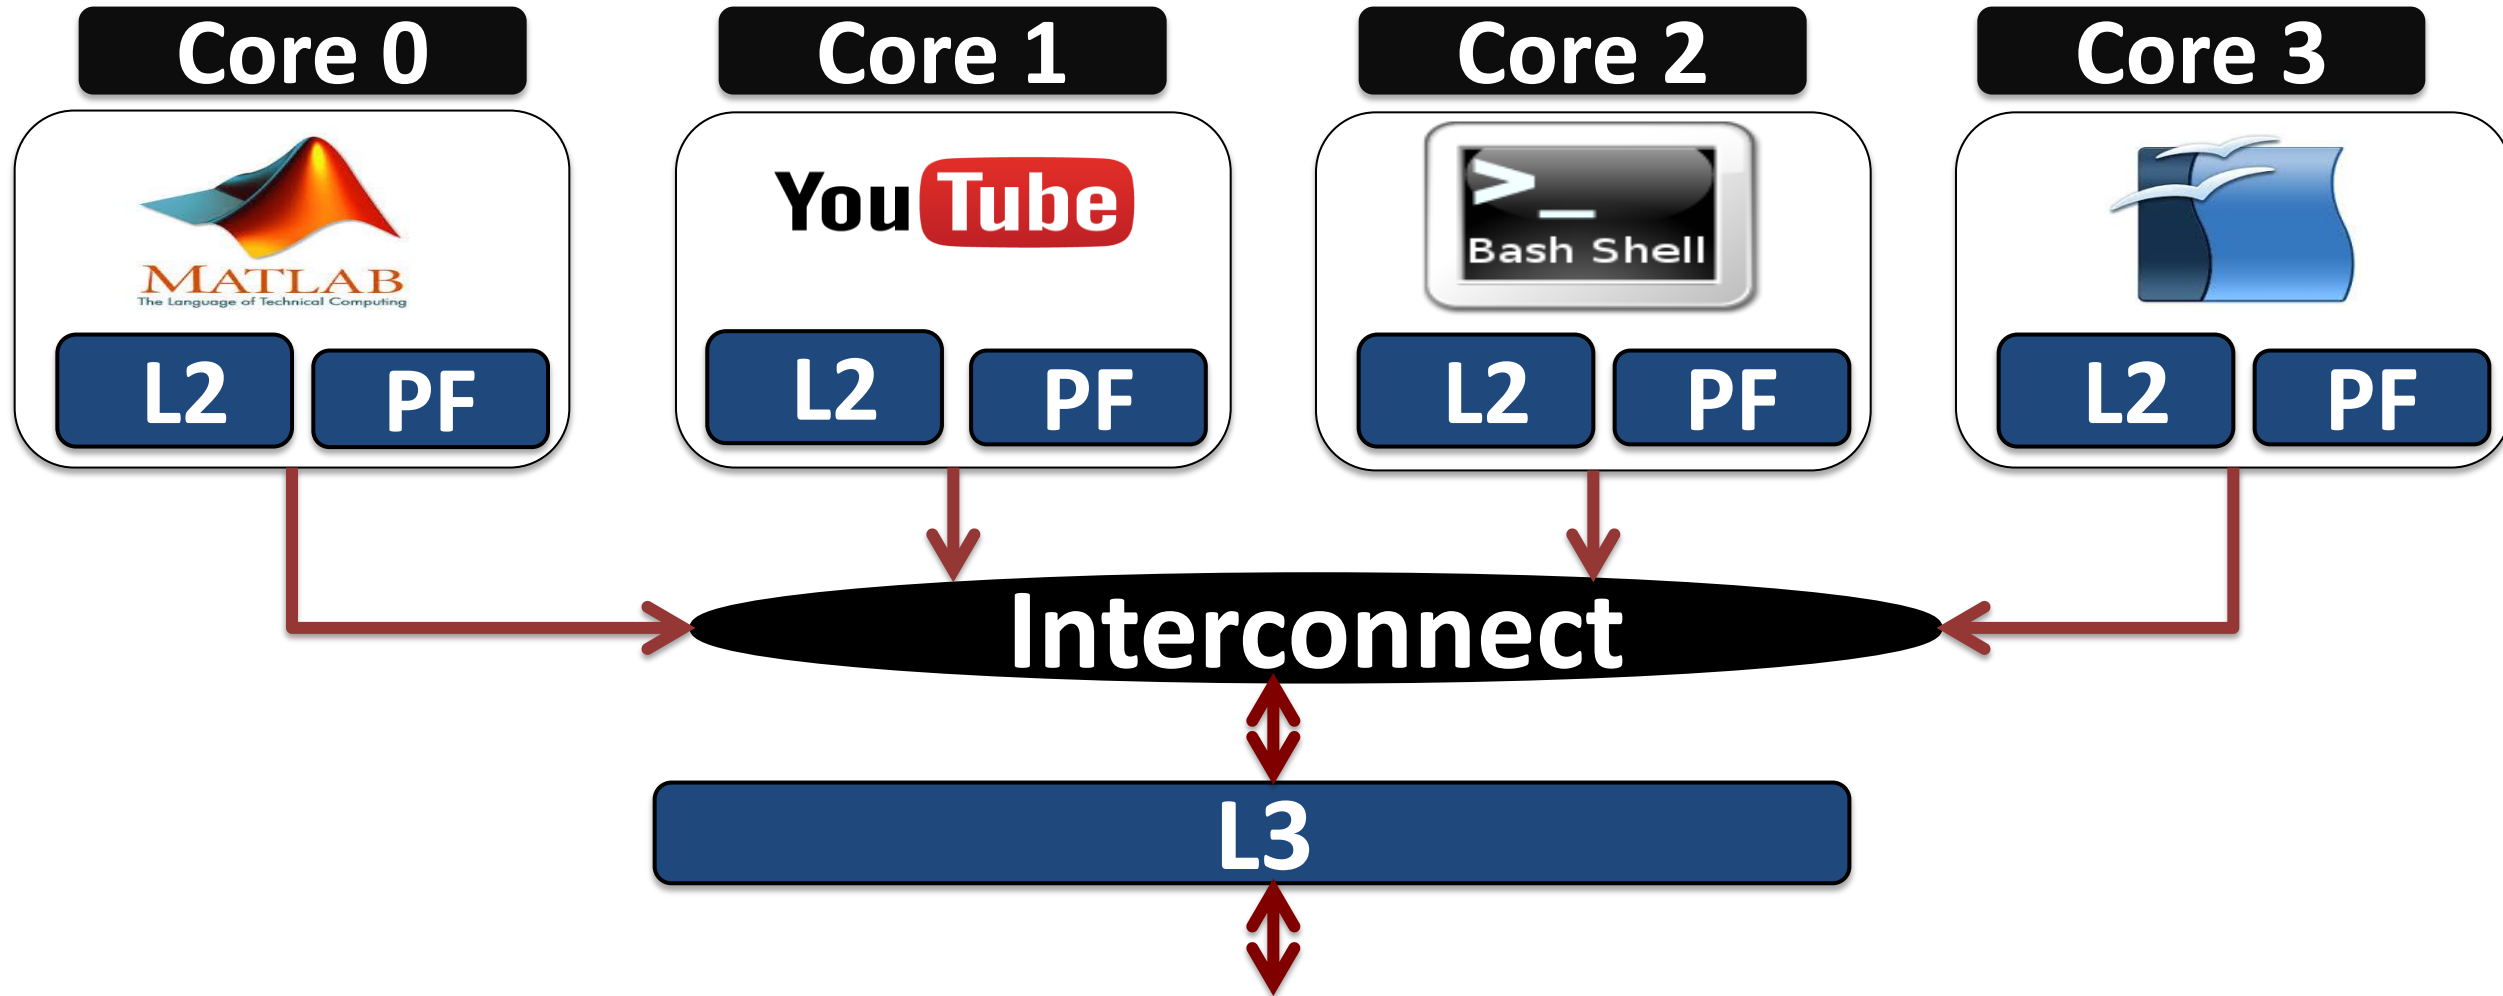
*How?*

By observing/predicting the demand access (LOAD/STORE) patterns.

# Hardware Prefetch Engine

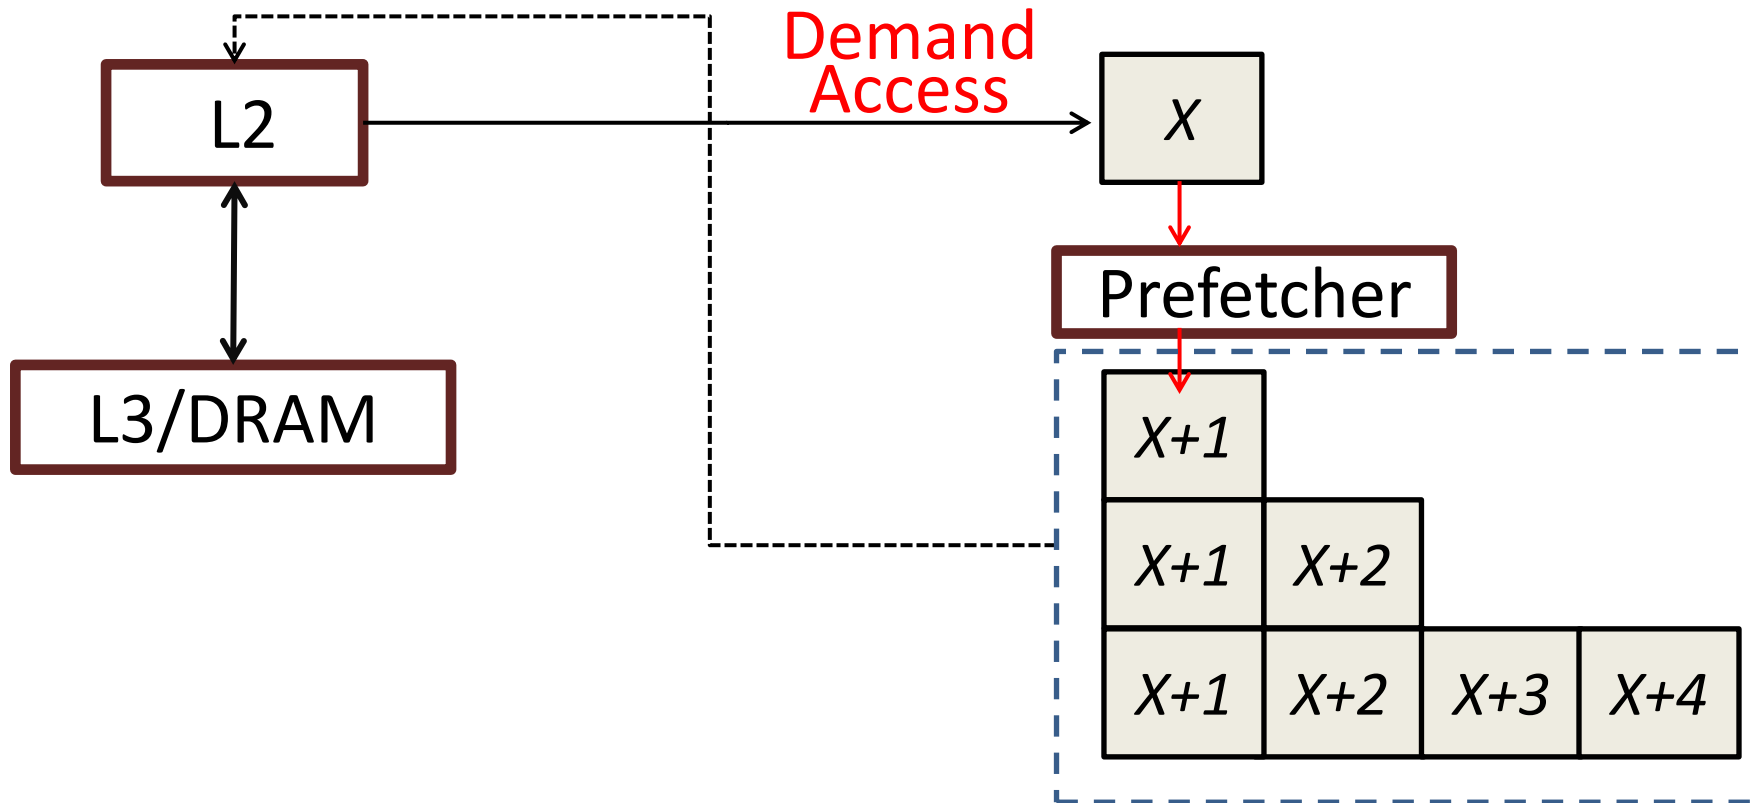


# Prefetchers in Multicore - 101




# Prefetching Knobs

**Prefetch Degree:** Number of prefetch requests to issue at a given time.



# Prefetching Knobs

**Prefetch Distance:** How far ahead of the demand access stream are the prefetch requests issued?



*Prefetch-distance*


$$Y = X + 4$$

$$Y = X + 8$$


$$Y = X + 16$$

# Aggressiveness [degree, distance]

**Prefetch degree:** #Prefetch requests issued on a miss



**Prefetch distance:** How far ahead (in terms of # blocks) of the demand miss ?





# The Simplest Prefetcher


**Next Line:** Miss to cache block  $X$ , prefetch  $X+1$ . Degree=1, Distance=1

Works well for L1 Icache and L1 Dcache.

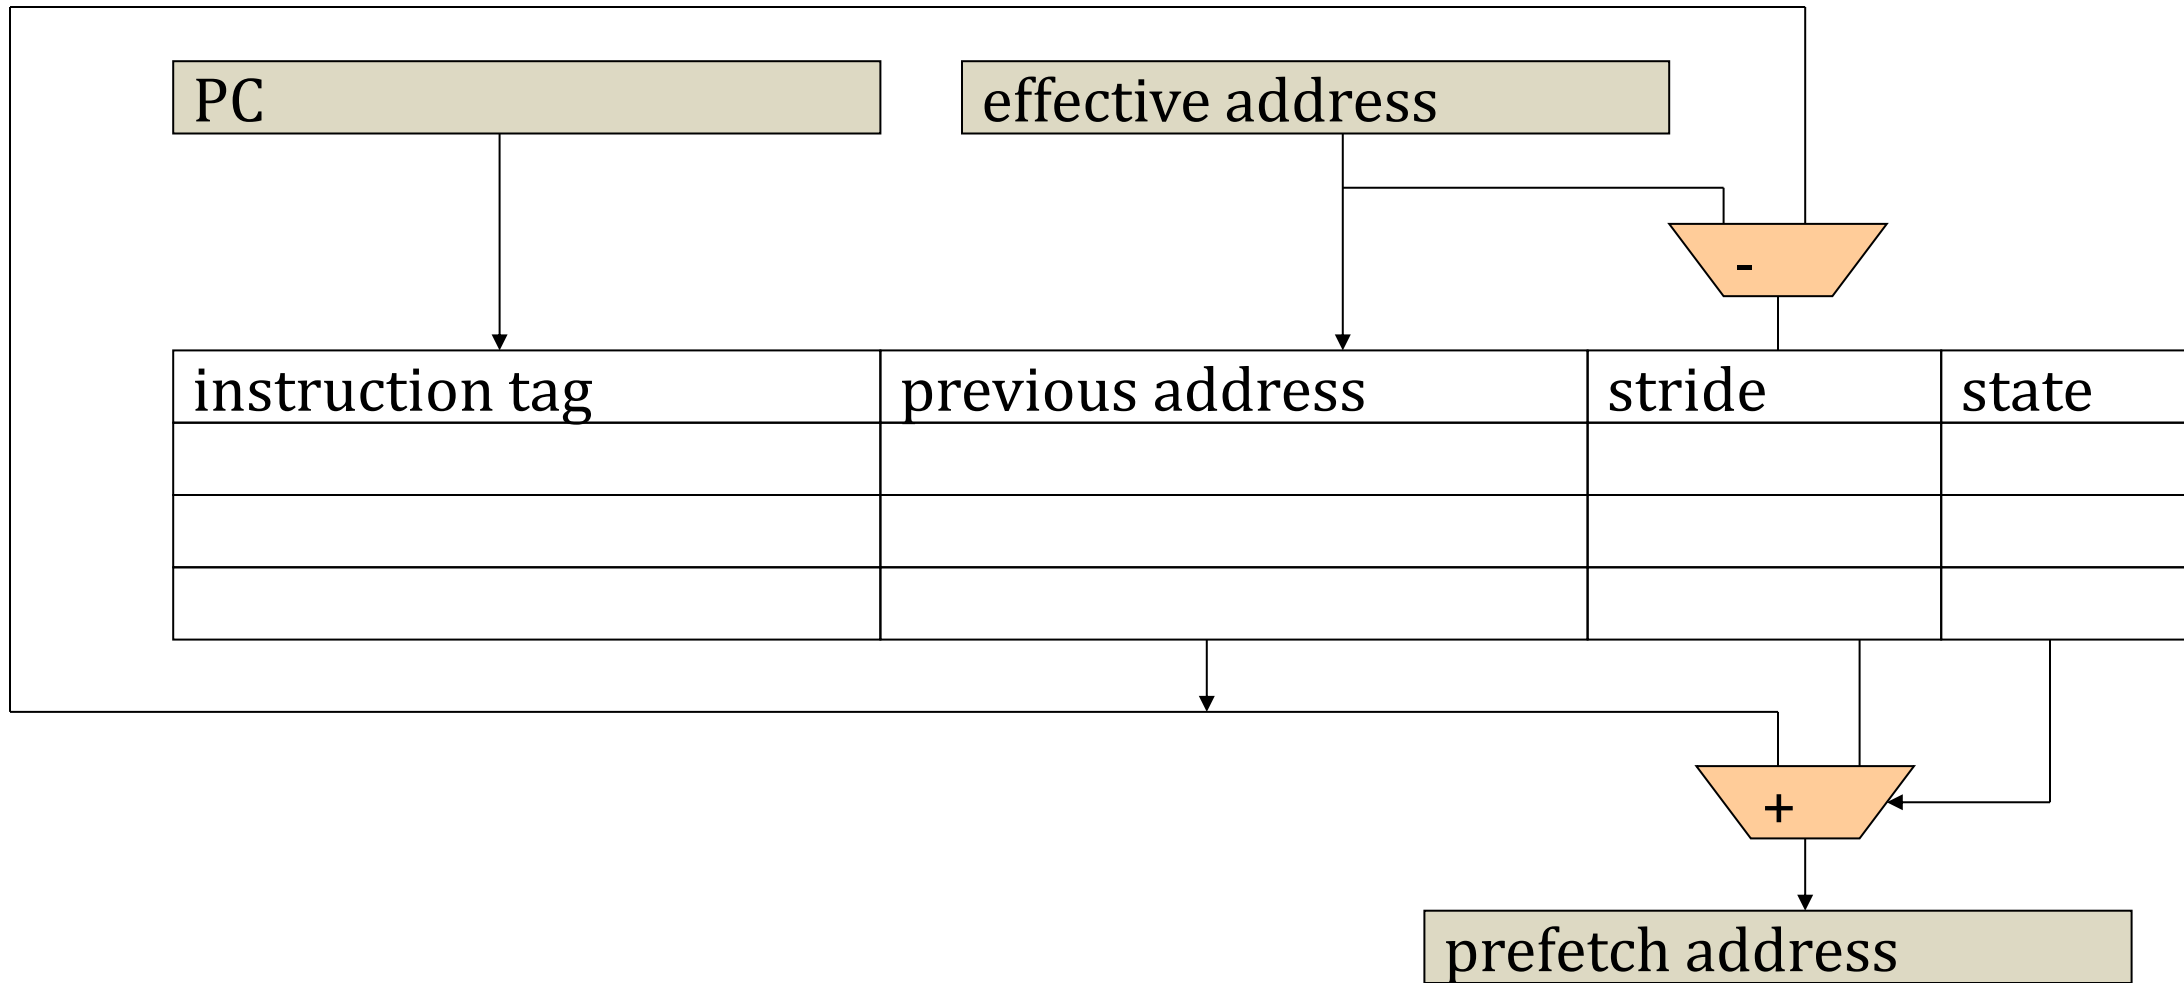
**Next N Line:** Miss to cache block  $X$ , prefetch  $X+1, X+2, \dots, X+N$ ,  
Degree= $N$ , Distance= min. 1 and max.  $N$

# What about this?

Load R1 = [R2]  
Load R3 = [R4]  
Add R5, R1, R3  
Store [R6] = R5



# Stride Prefetching



# An Example



```
float a[100][100], b[100][100], c[100][100];
```

```
...
```

```
for (i = 0; i < 100; i++)  
  for (j = 0; j < 100; j++)  
    for (k = 0; k < 100; k++)  
      a[i][j] += b[i][k] * c[k][j];
```

instruction tag	previous address	stride	state
ld b[i][k]	50008	4	steady
ld c[k][j]	90800	800	steady
ld a[i][j]	10000	0	steady

# Pointer Chasers



# Stream Prefetching [DPC1]

miss  
sequence →

100	102	104	503	504	501	499
-----	-----	-----	-----	-----	-----	-----


100	102	104	504	501	499
-----	-----	-----	-----	-----	-----

1<sup>st</sup> 2<sup>nd</sup> 3<sup>rd</sup> 1<sup>st</sup> 2<sup>nd</sup> 3<sup>rd</sup>  
miss miss miss miss miss miss

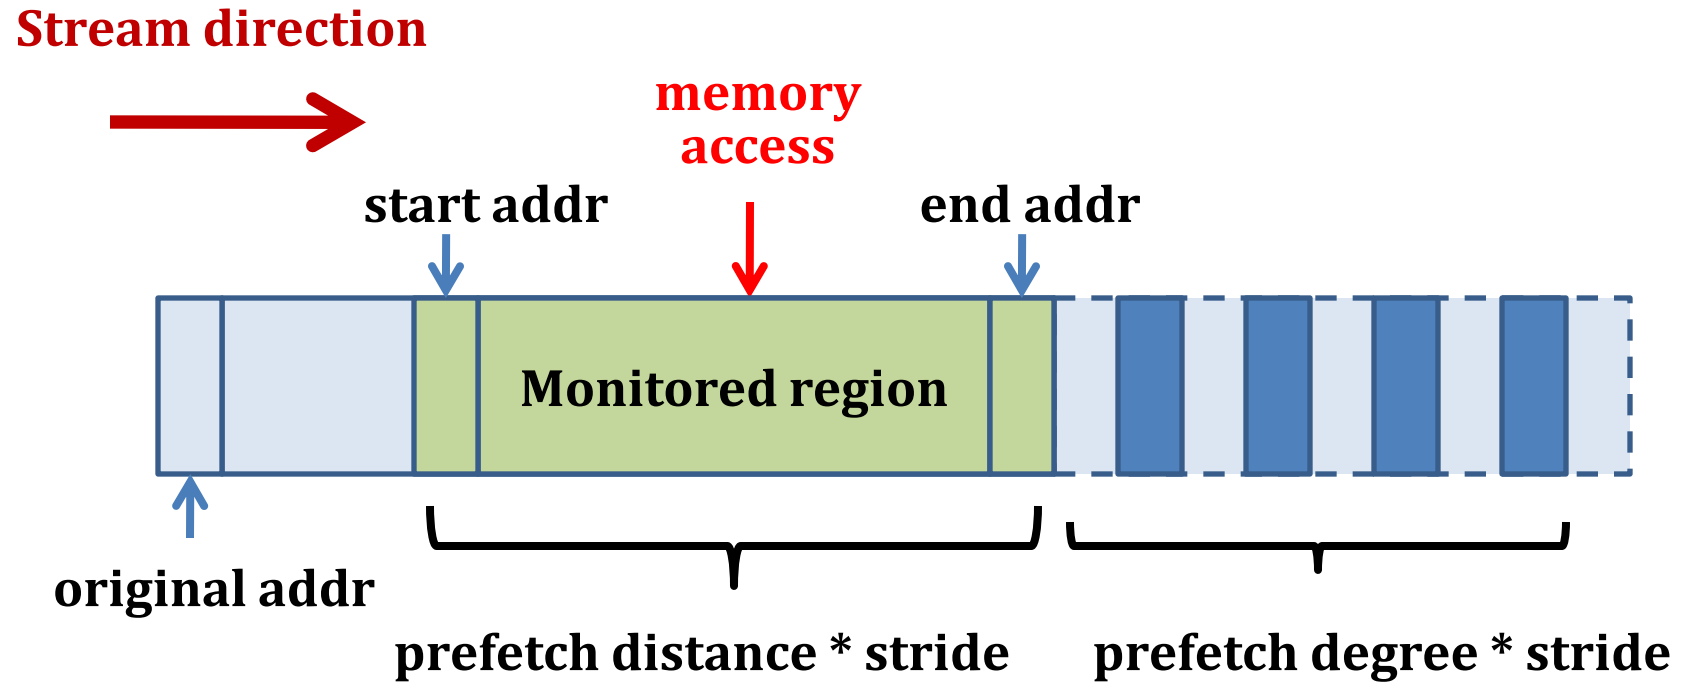
← **Trained!**

**Training: Consecutive misses in the same direction.**

# Stream Prefetcher in Action



# Stream + Stride





# Quantifying Prefetchers

$$\text{PrefetchAccuracy}(i) = \frac{\text{Prefetch}_{\text{hits}}(i)}{\text{Prefetch}_{\text{issued}}(i)}$$

Prefetched Block in the Cache.

$$\text{Lateness}(i) = \frac{\text{Prefetch}_{\text{late}}(i)}{\text{Prefetch}_{\text{hits}}(i)}$$

Prefetched Block Still on its way

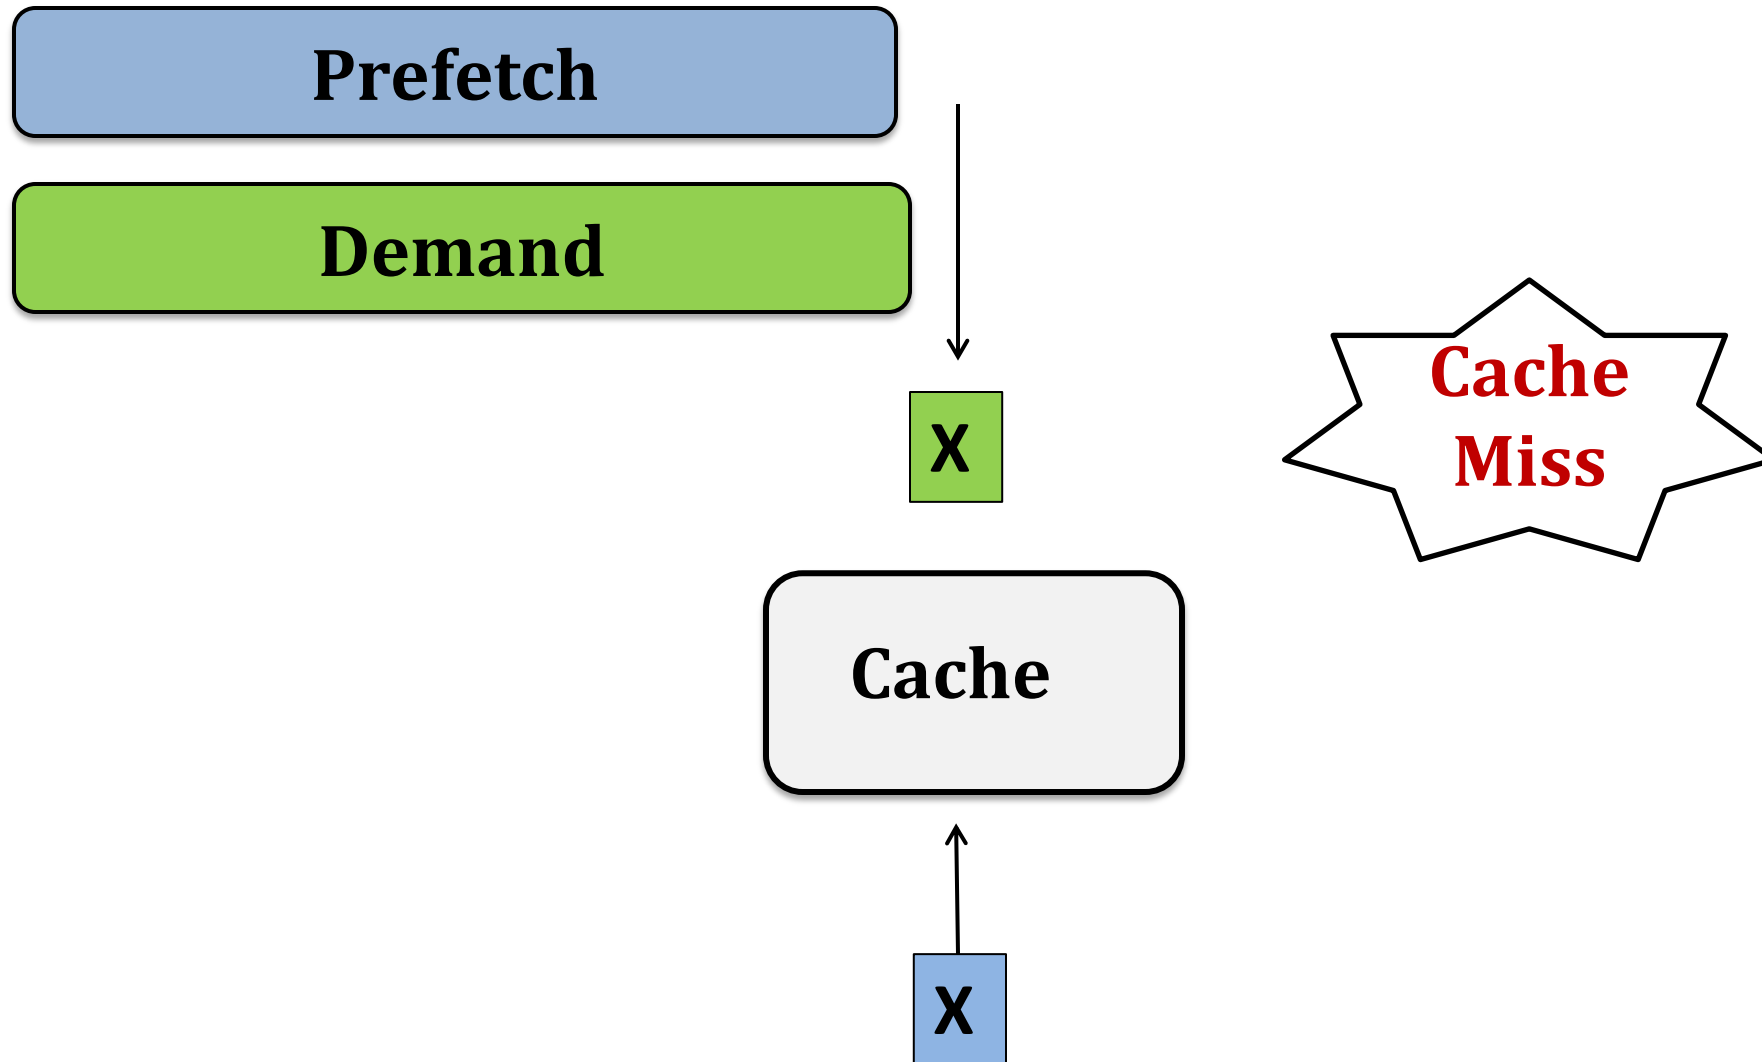
$$\text{Pollution}(i) = \frac{\text{LLC Poll}(i)}{\text{Demand}_{\text{misses}}(i)}$$

Prefetched Block evicted a demand block that will be reused

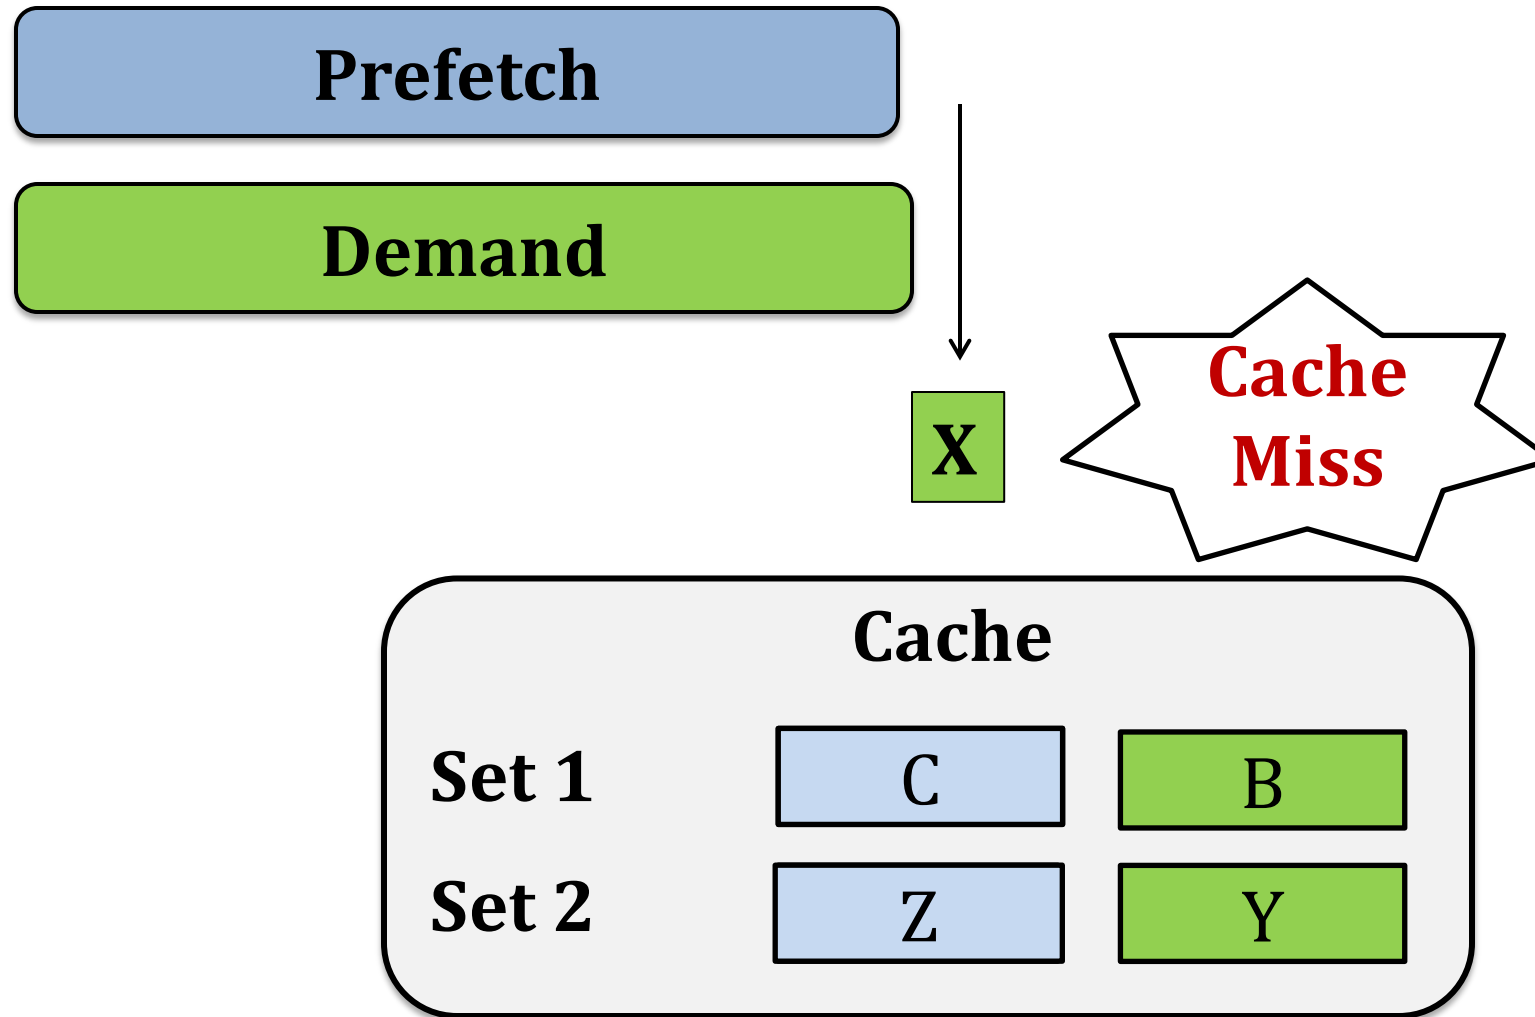
$$\text{Coverage}(i) = \frac{\text{Prefetch Hits}(i)}{\text{Prefetch Hits}(i) + \text{Demand}_{\text{misses}}(i)}$$

Fraction of misses avoided

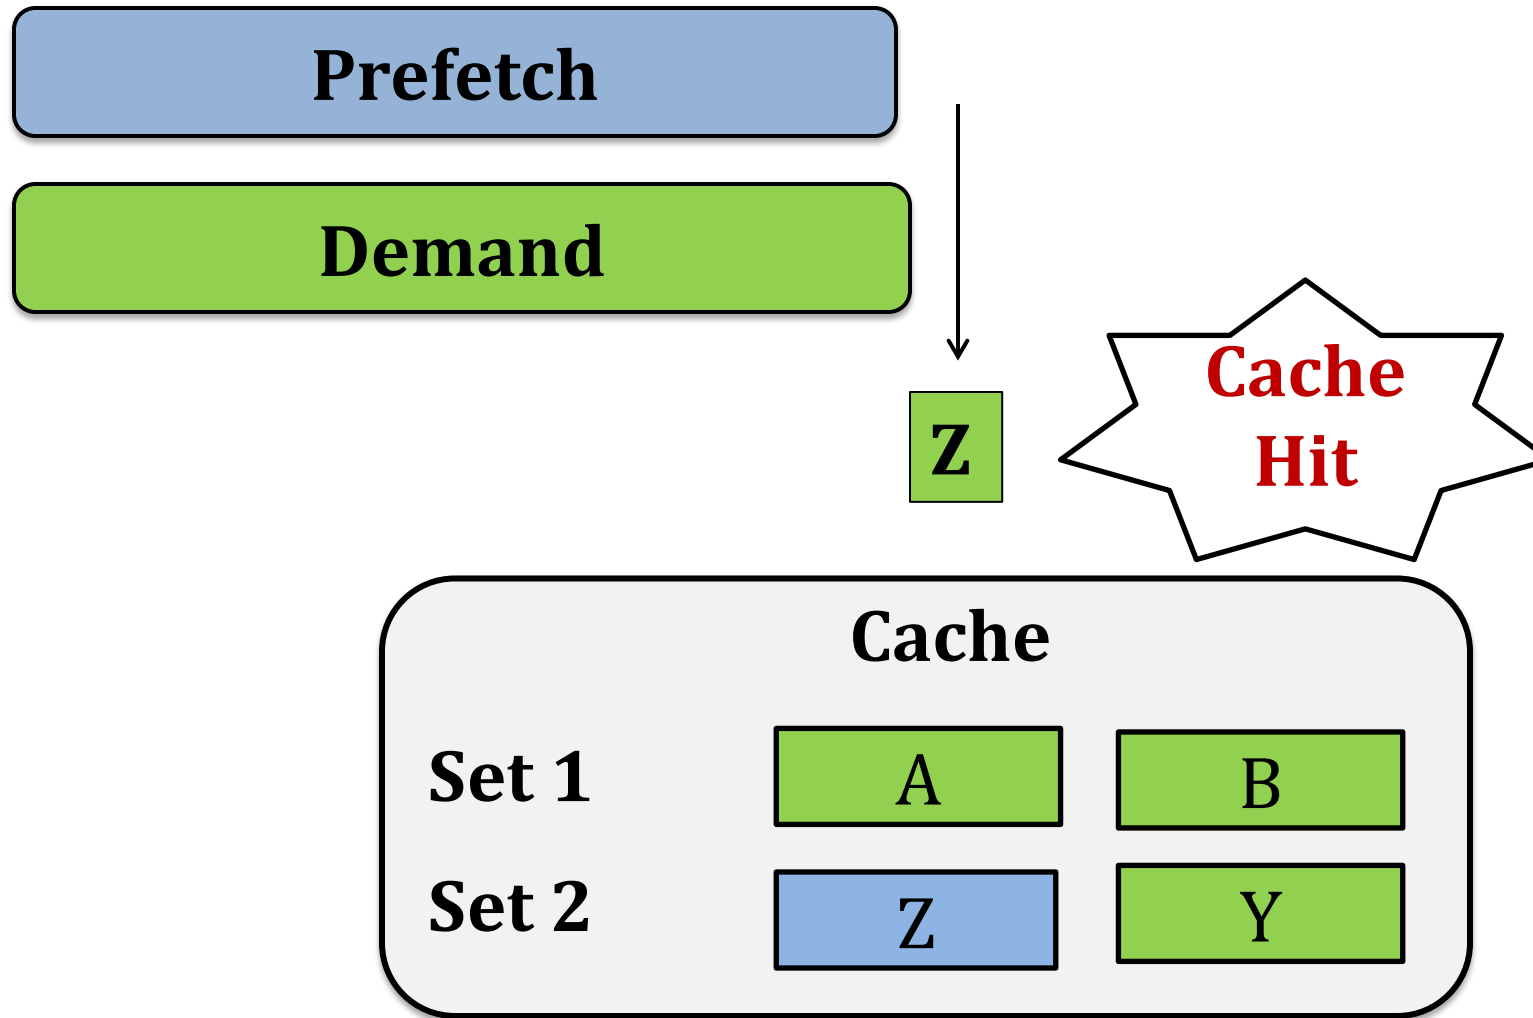
# Prefetch Lateness



# Cache Pollution



# Cache Hits & Accuracy



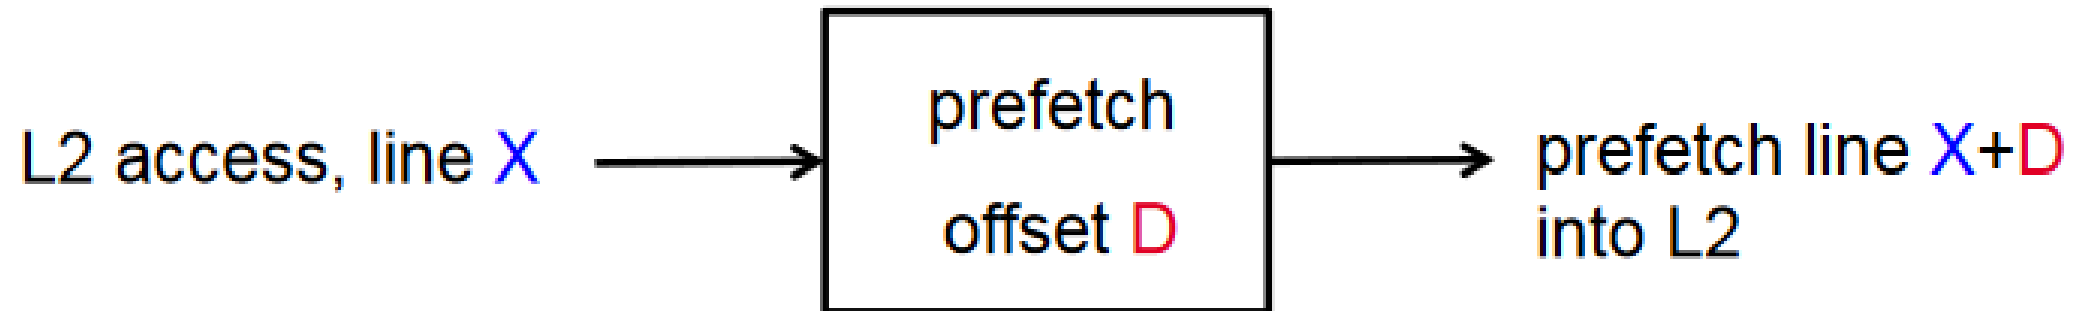
# Where to Put These Prefetchers?

L1? Next-line, PC-localized stride predictors

L2? Stream + Stride, Other variants


L1 instruction cache ? Predict the future PC

# State-of-the-art Prefetchers




next-line prefetching → offset = 1

# Perfect Timing




# Delayed Prefetching






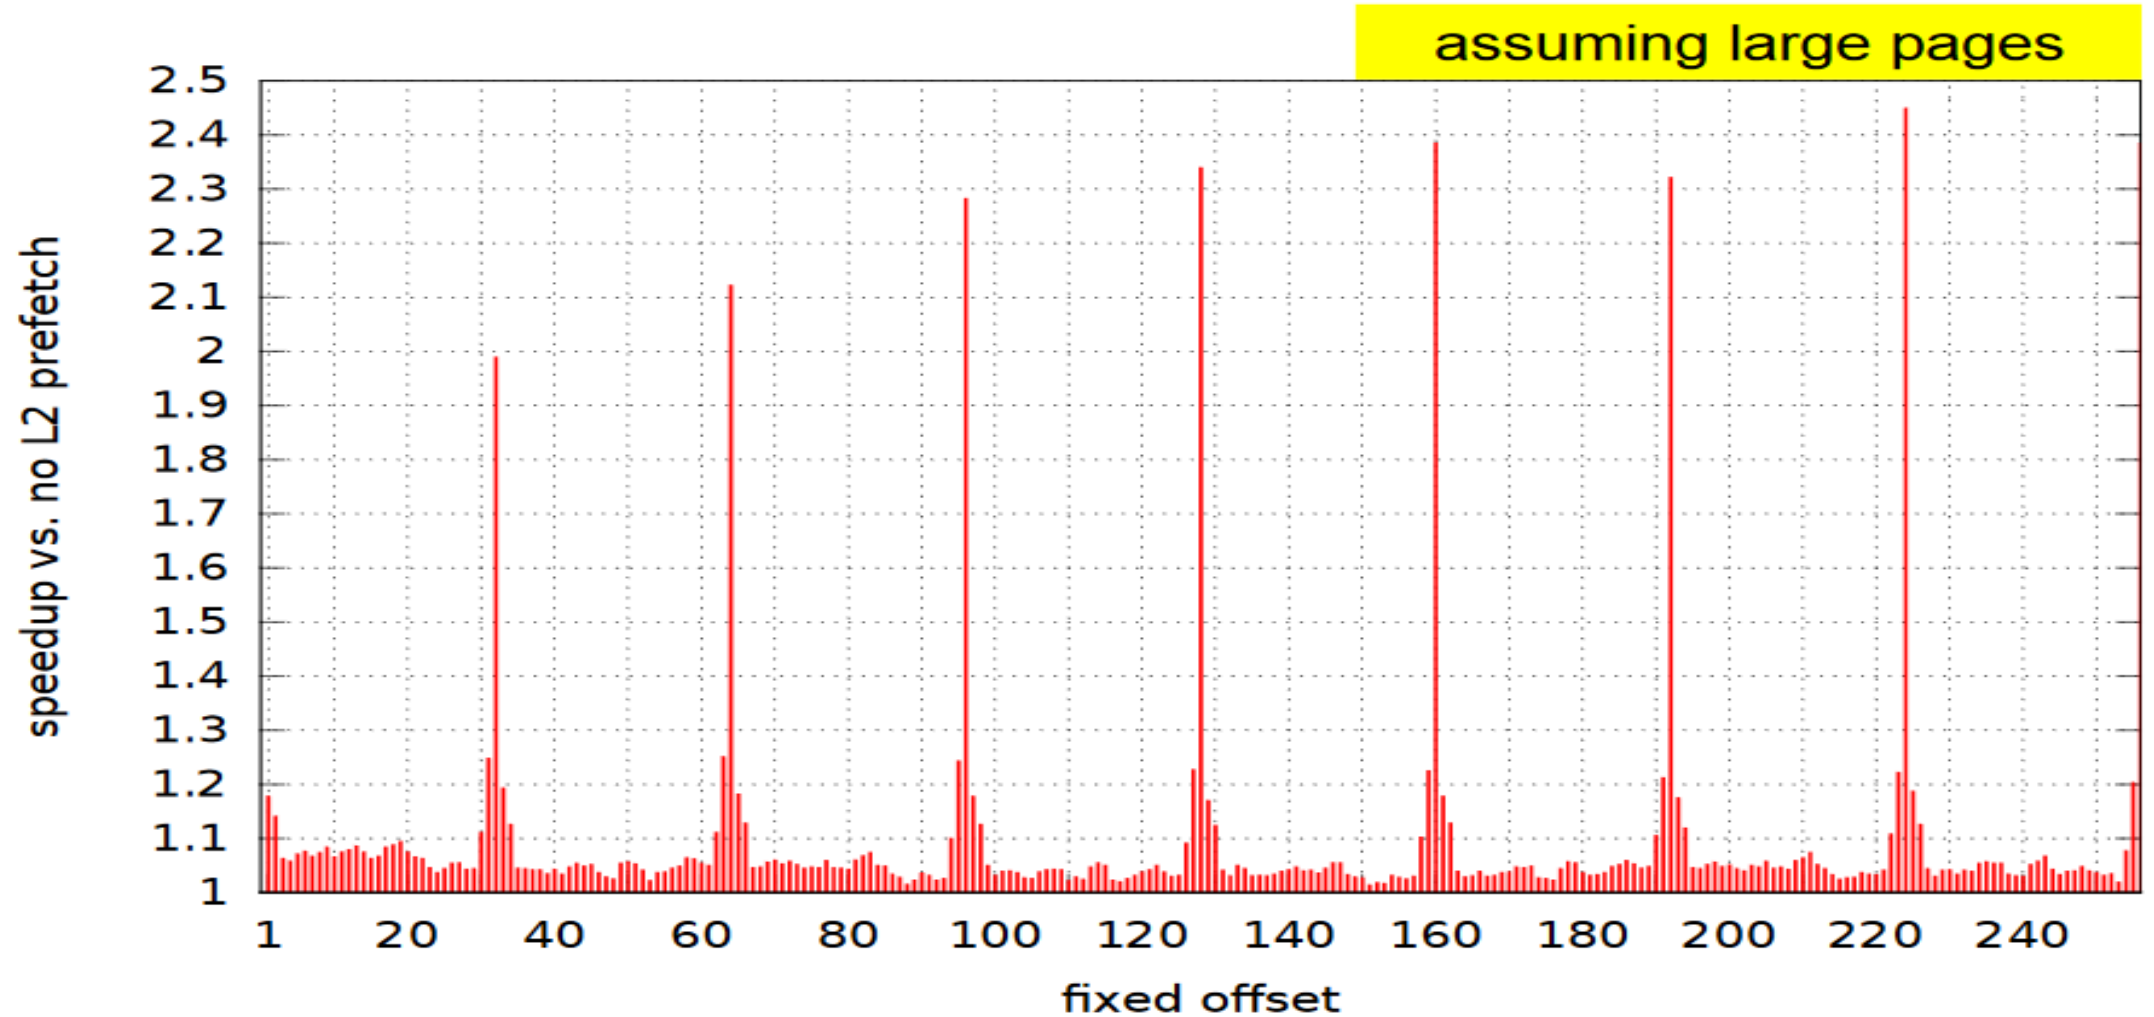
# Offset




# Offset = Sum of strides



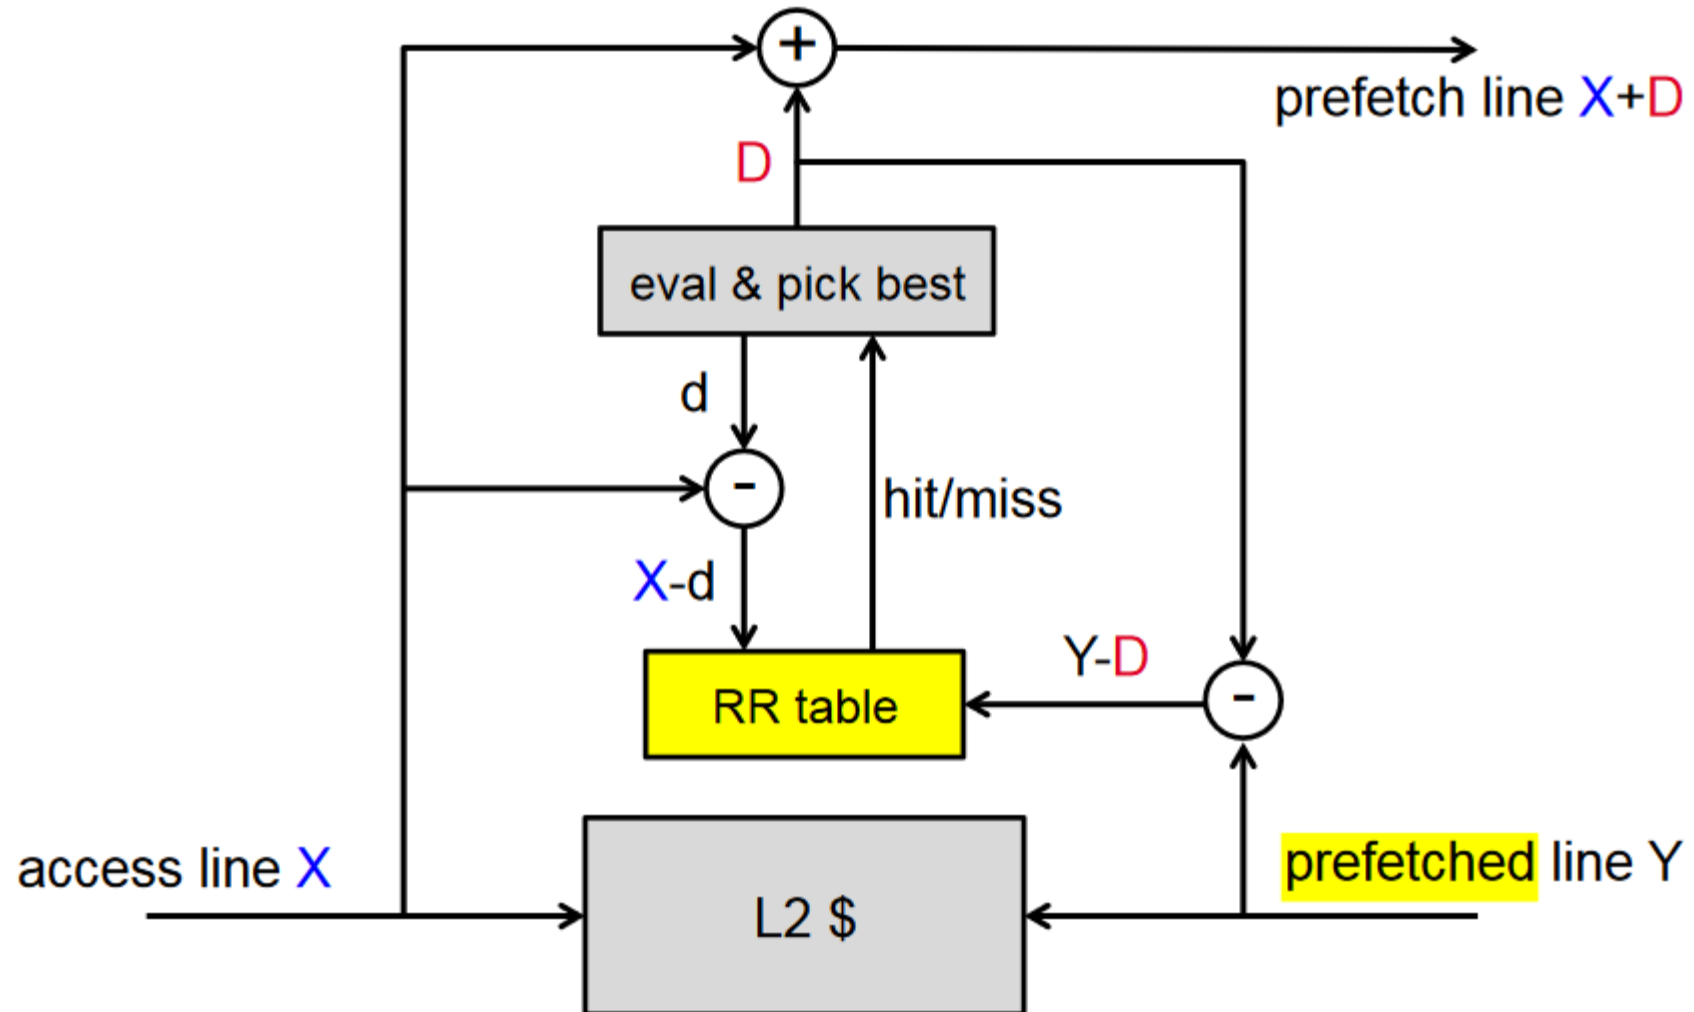
# milc: Offset



# GemsFDTD: Offset



# Best-offset Prefetcher [HPCA '16]



# Specialized Streams

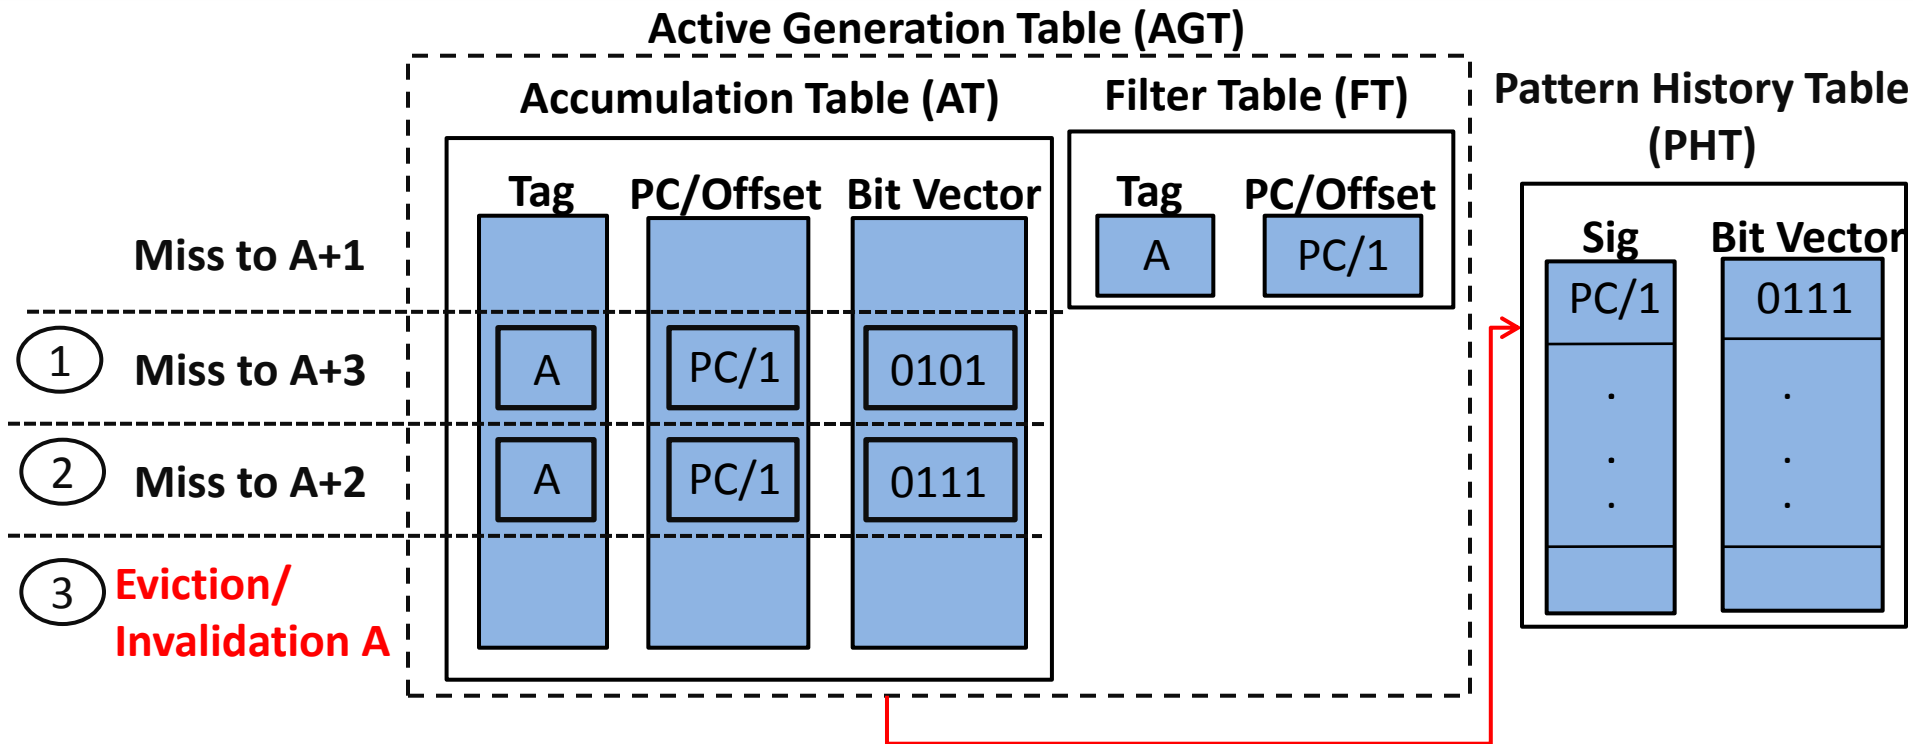
**Temporal Streams – Sequences of temporally correlated addresses, exploited by TMS [ISCA '05].**

**Spatial Streams – Streams, which are correlated in space, exploited by SMS [ISCA '06].**

**SpatioTemporal Streams – Temporal correlation among the spatial regions, and spatial correlation within a region, exploited by STeMS [ISCA '09].**

# Spatial Memory Streaming (SMS)

- Divides the memory space into fixed size regions, indexed by a **signature (PC/offset)**.
- Each signature contains a **bit vector**.
- Each bit in the bit vector corresponds to a cache line.



# Reading Assignment-1

Proactive Instruction Fetch [MICRO '11]

Indirect Memory Prefetcher [MICRO '15]

Deadline: October 7, 2017, 17:00 hrs through Canvas

More details through Piazza



# Programming Assignment-2

Will be released on Sept 11, 2017

Based on Hardware Prefetchers

This time: You have to code (no analysis)

# PA1 Presentation


12<sup>th</sup> Sept, 15:00 hrs IST, KD-103

7+1 min presentation


Do not put MPKI and IPC numbers

You will evaluate your peers

# What About Irregular Applications? [MICRO '13]



# PC Localization




**Physical Address Space**



**Indirection**

# Structural Address Space



## PC-Localized Stream : A

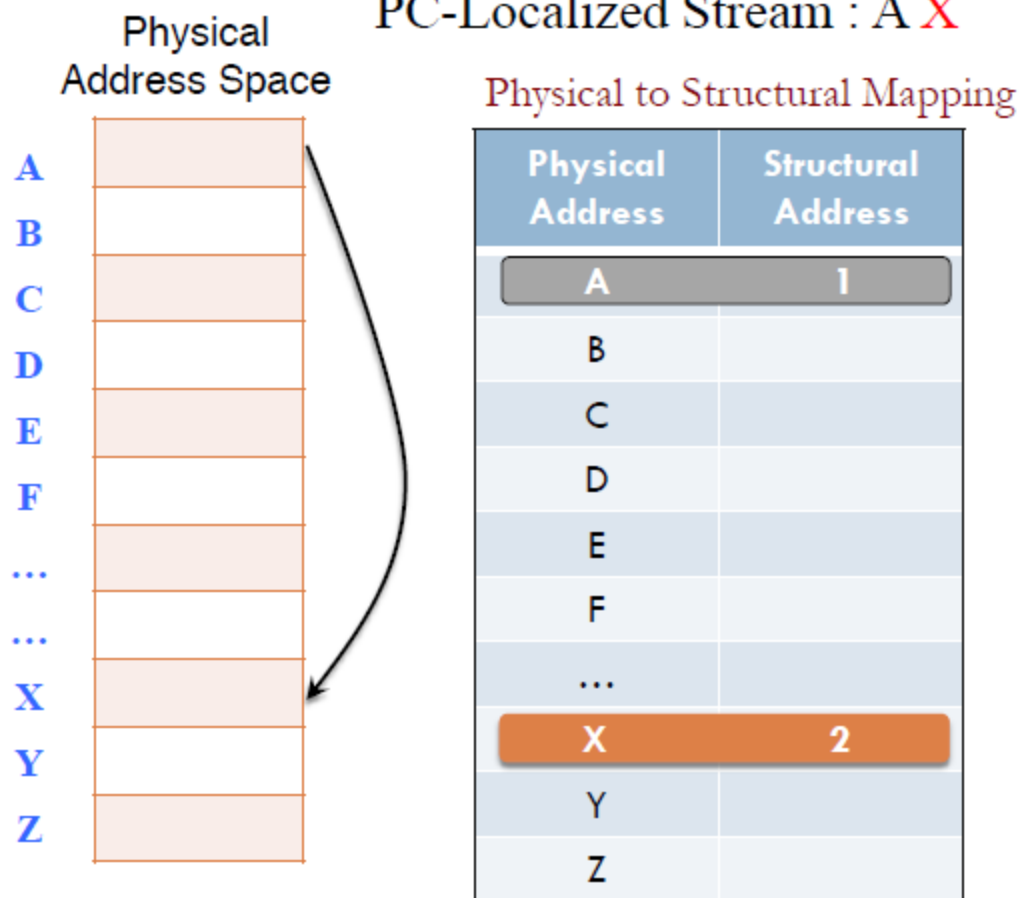
Physical  
Address Space

A	
B	
C	
D	
E	
F	
...	
...	
X	
Y	
Z	

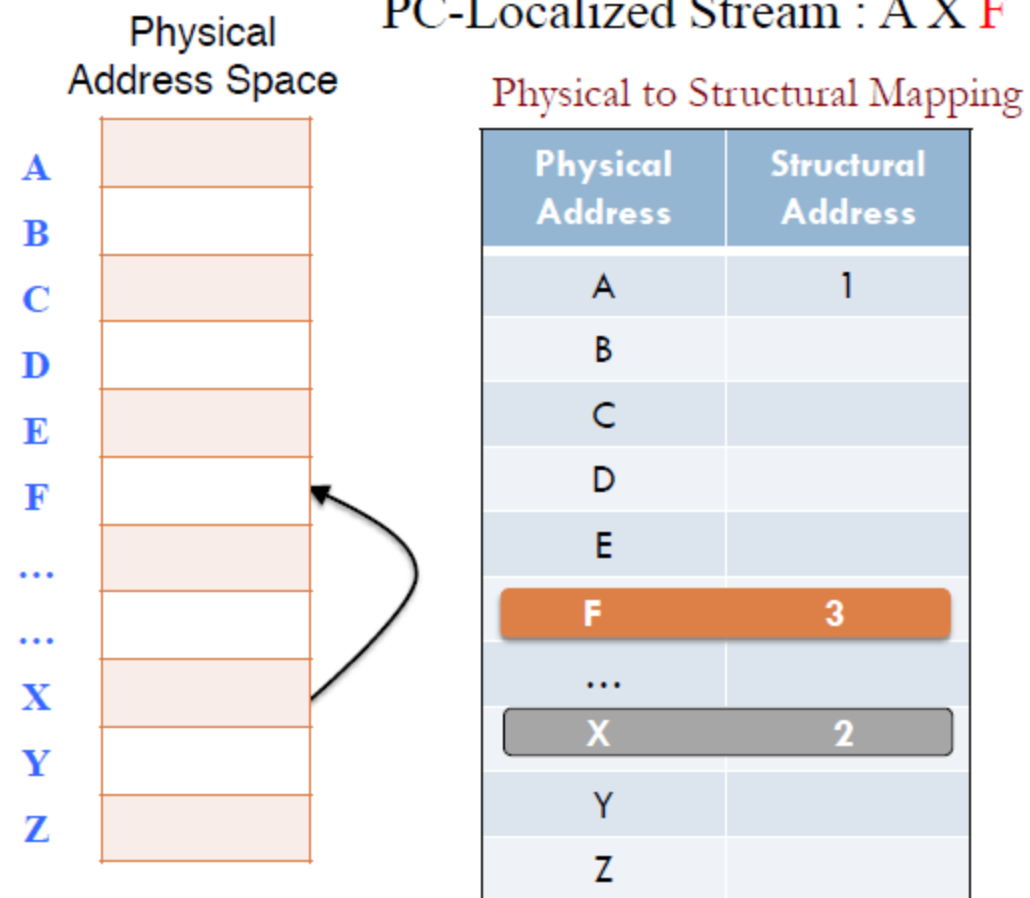
Physical to Structural Mapping

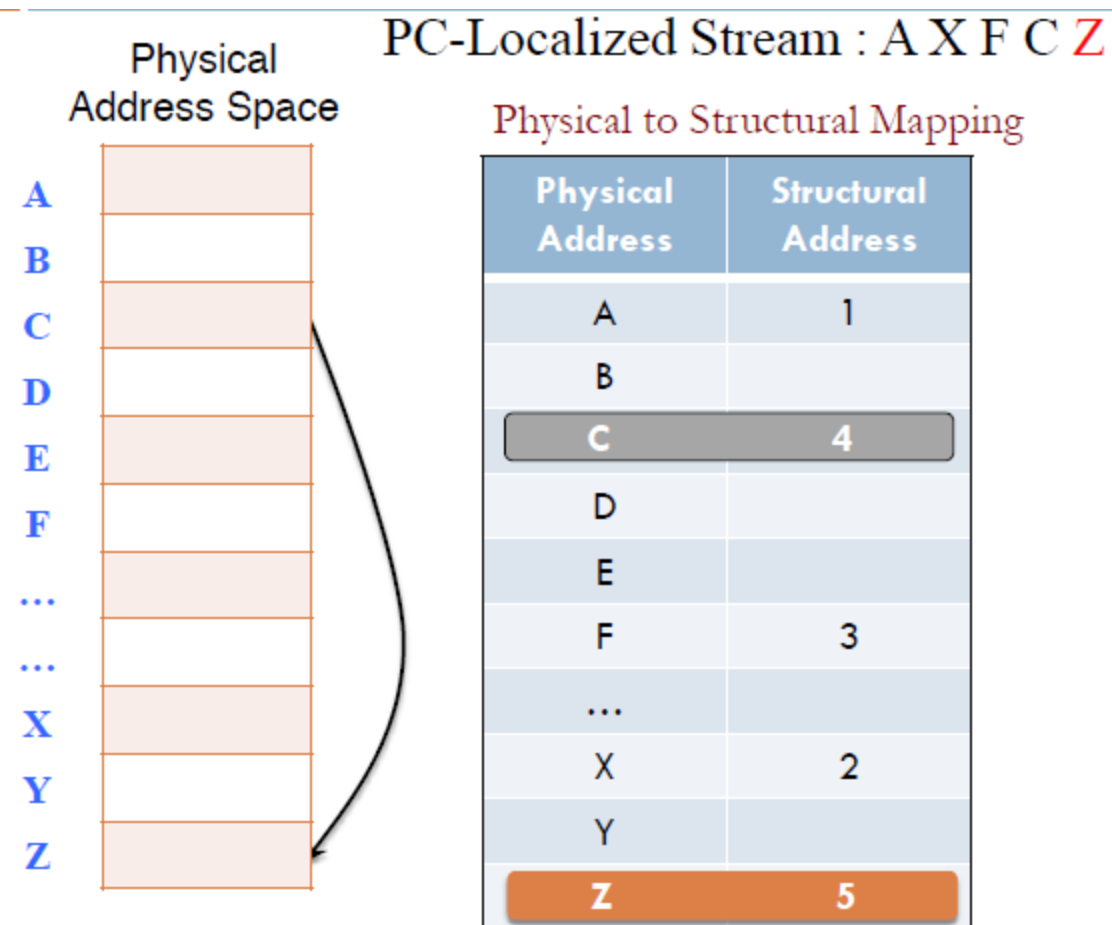
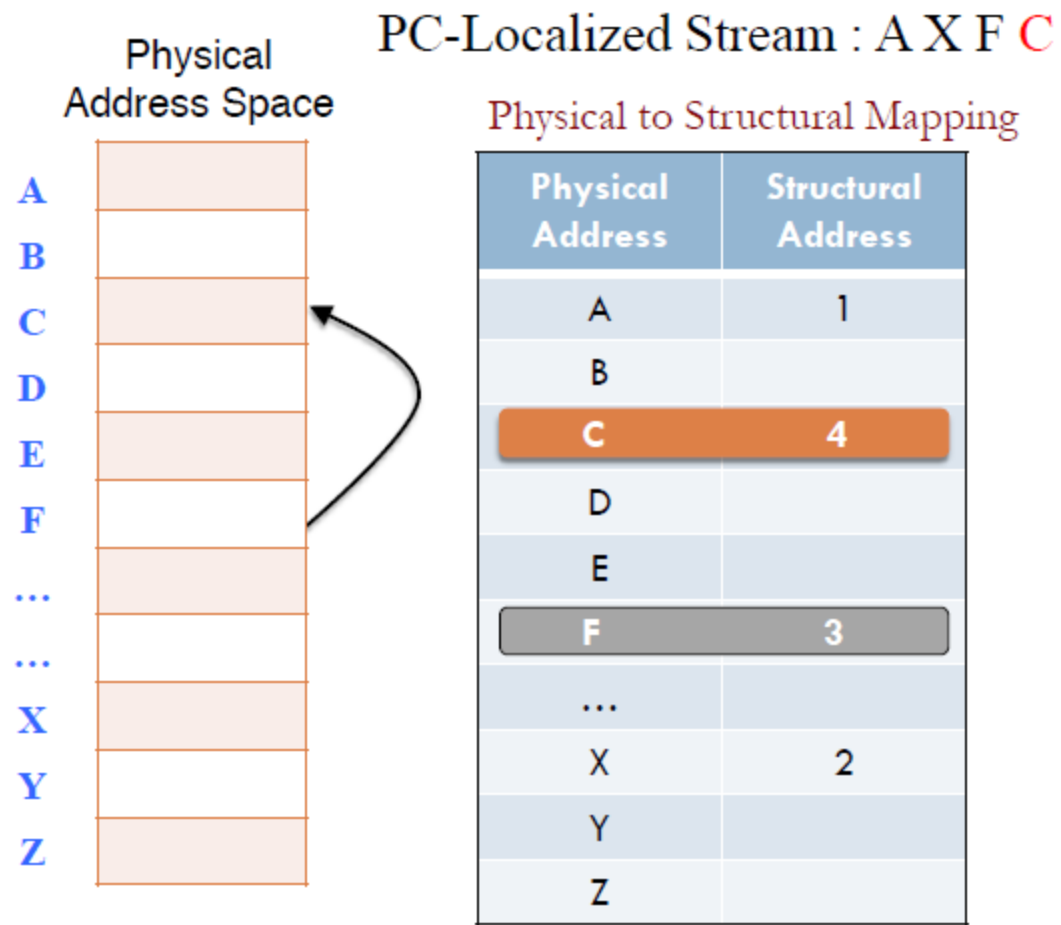
Physical Address	Structural Address
A	1
B	
C	
D	
E	
F	
...	
X	
Y	
Z	

## PC-Localized Stream : A X



## PC-Localized Stream : A X F







Trigger Address: **X**

Physical to Structural Mapping

Physical Address	Structural Address
A	1
B	
C	4
D	
E	
F	3
...	
<b>X</b>	<b>2</b>

Structural to Physical Map

Structural Address	Physical Address
1	A
2	X
3	F
4	C
5	Z
6	E

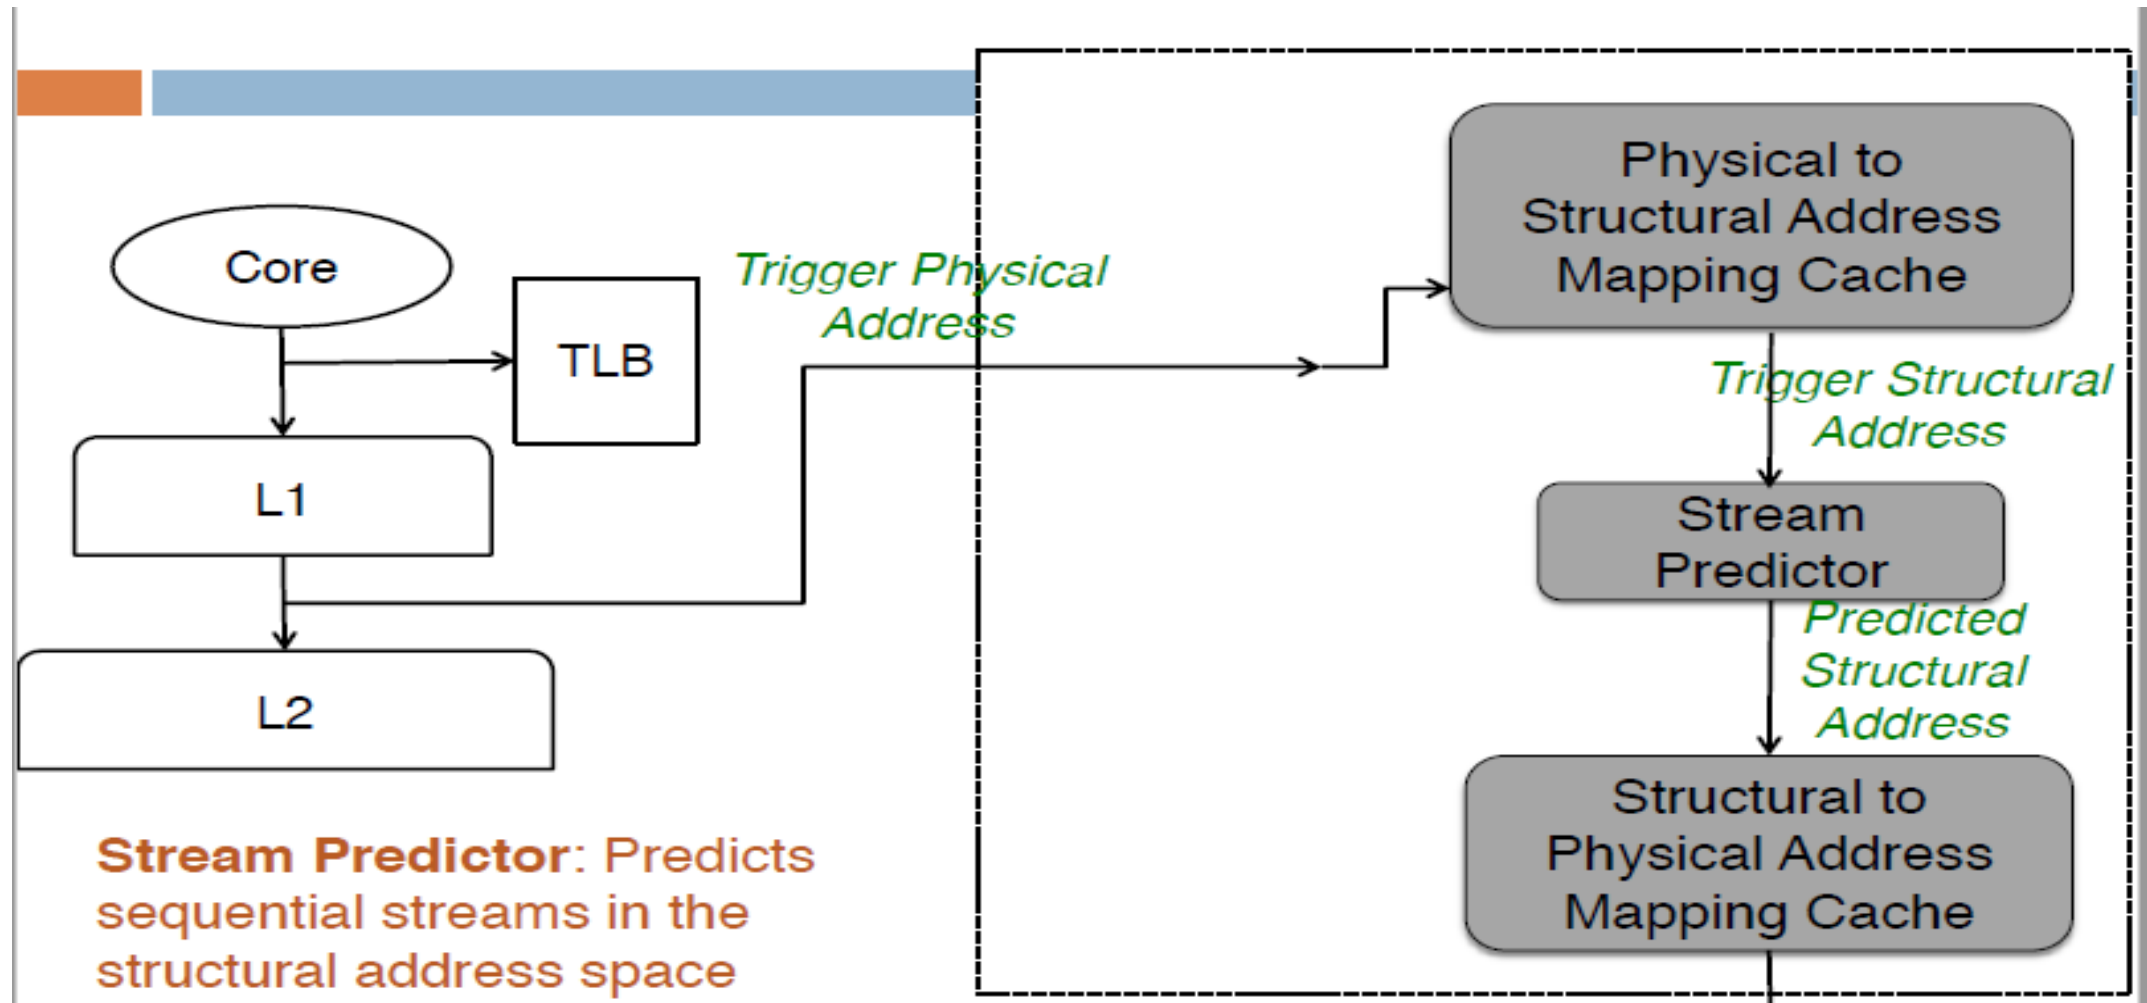
# Irregular Stream Buffer [MICRO '13]

Physical to Structural Mapping

Physical Address	Structural Address
A	1
B	
C	4
D	
E	
F	3
...	
X	2
Y	
Z	5

Structural to Physical Mapping

Structural Address	Physical Address
1	A
2	X
3	F
4	C
5	Z
6	E



# Interaction with Cache Replacement

**Read PACMan [MICRO '11]**

**Crux: Prefetched blocks are not reused after their first-use. So insert them with lowest priority.**