CASPER

CS305: Computer Architecture

Hardware Prefetching

https://www.cse.iitb.ac.in/~biswa/courses/CS305/main.html

https://www.cse.iitb.ac.in/~biswa/

Hardware Prefetching

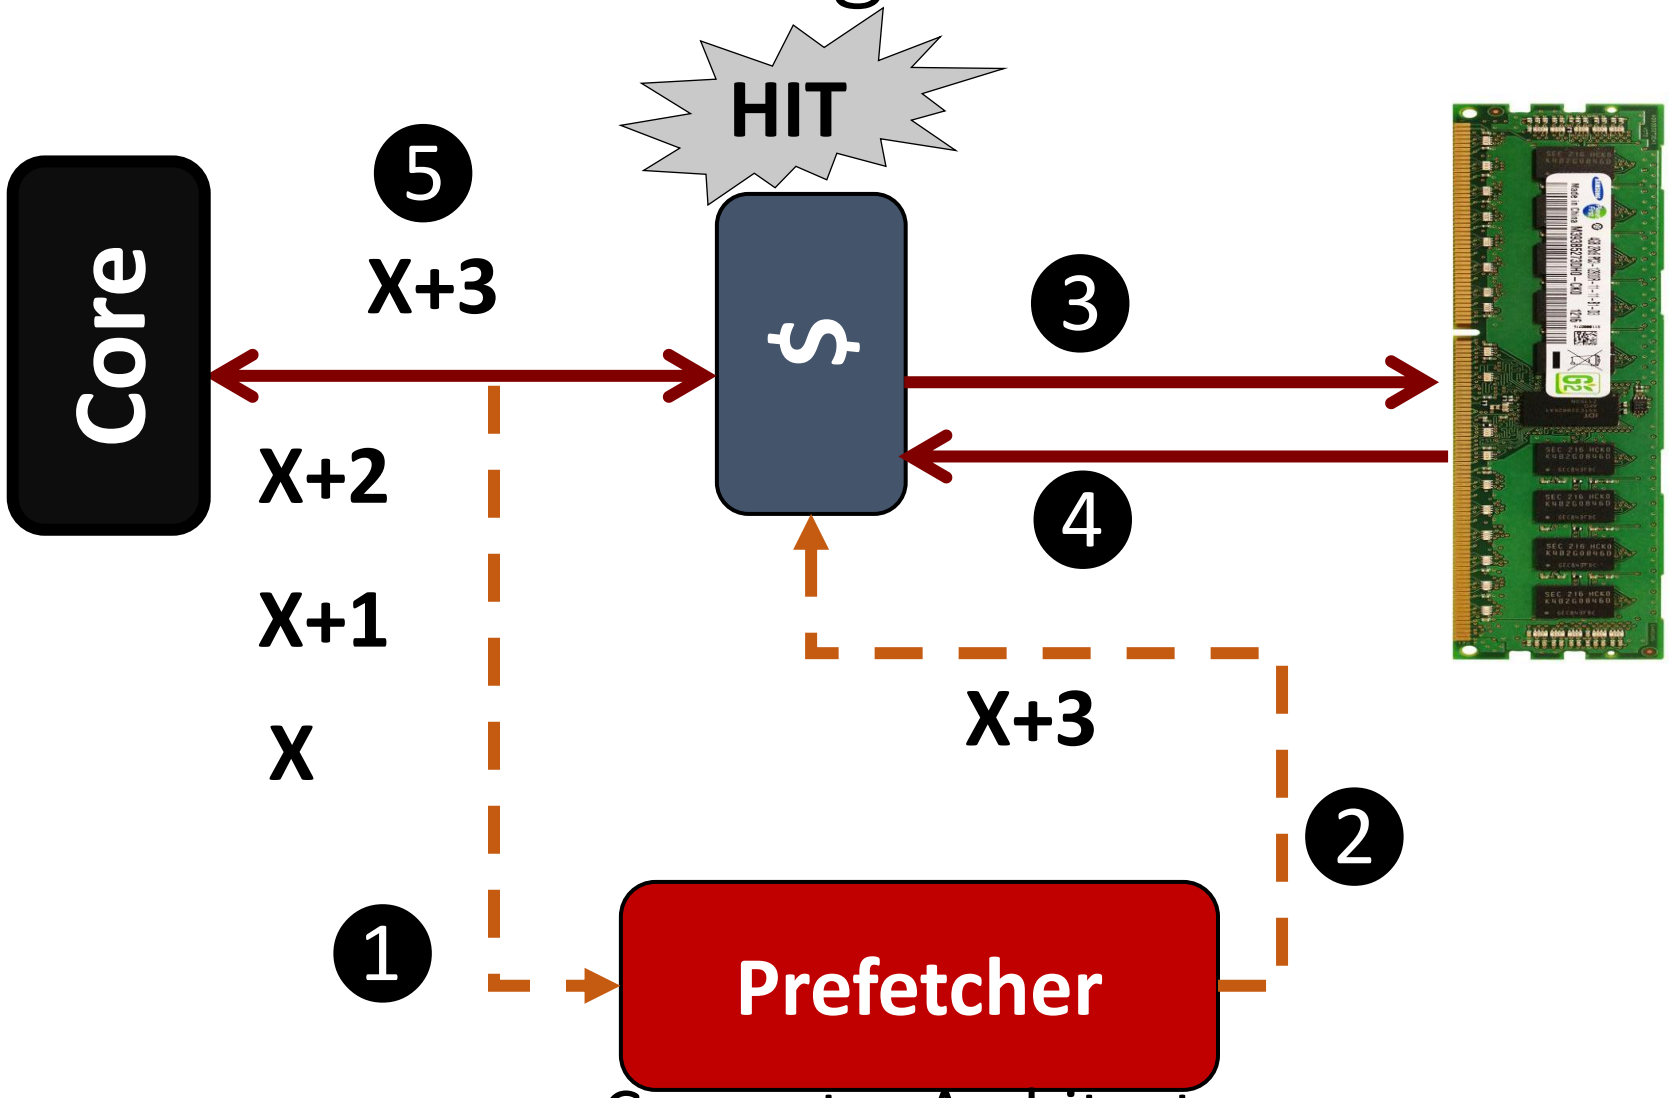

10K Feet View

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.

Prefetch Degree

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

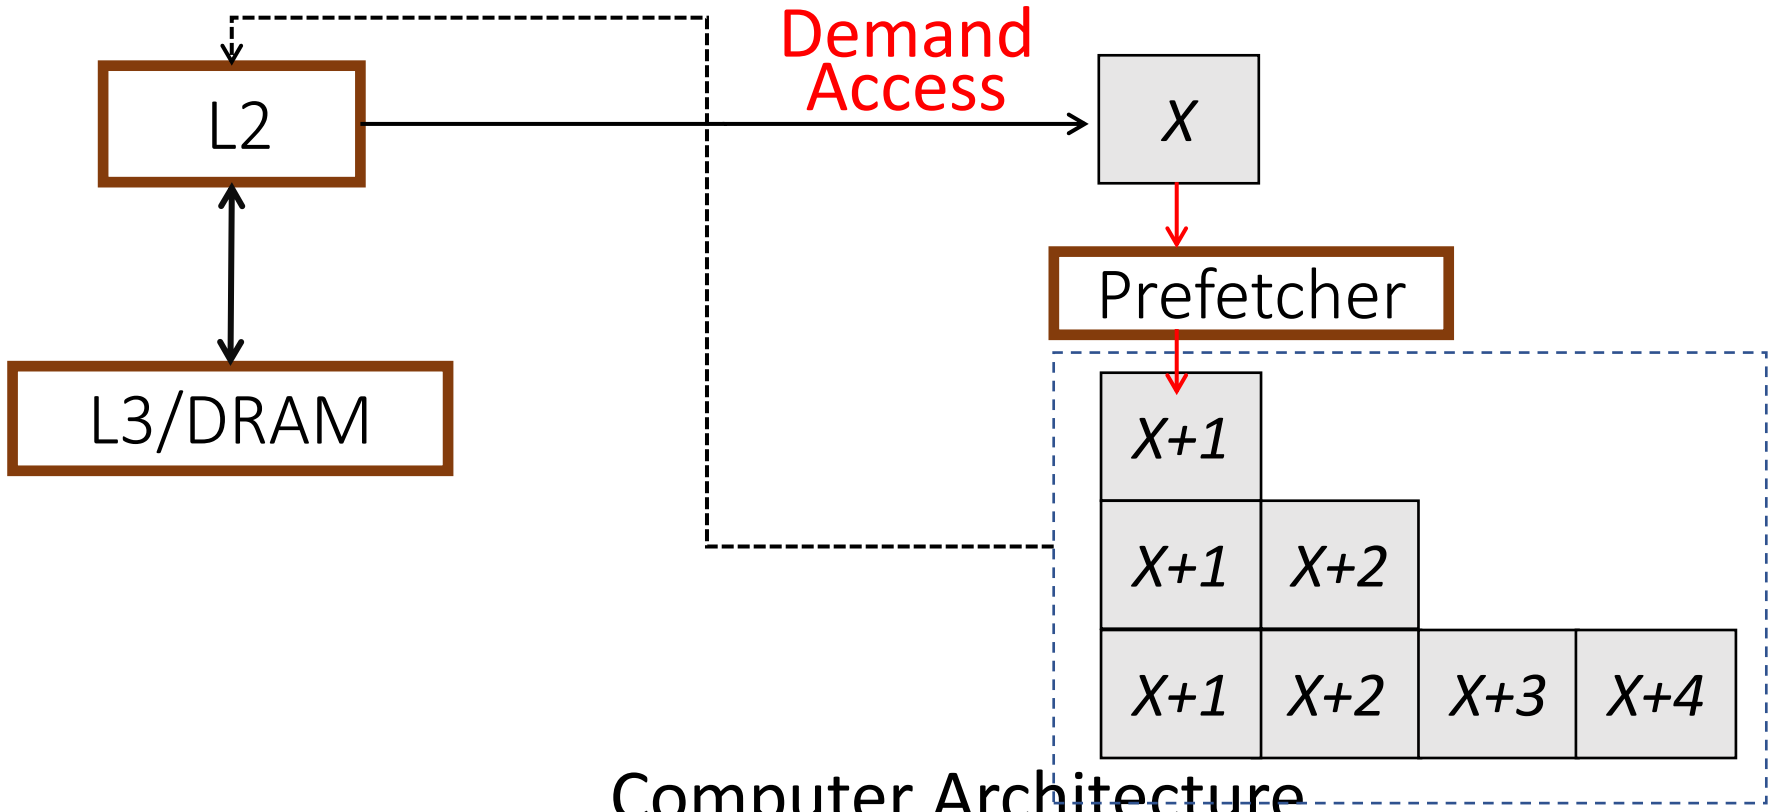

Prefetch Distance

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

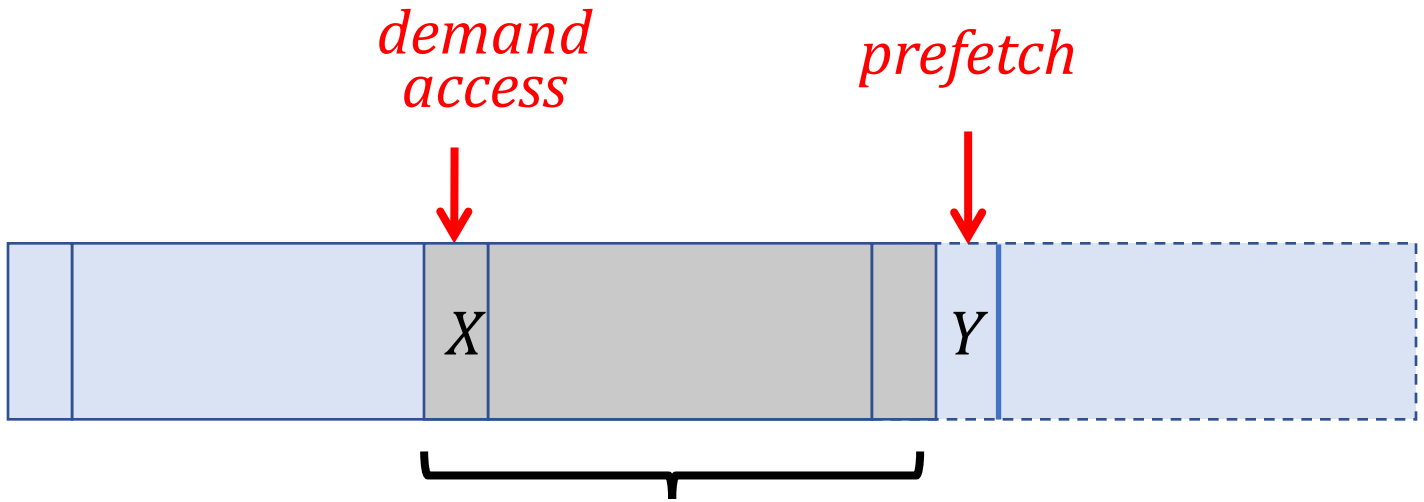

$Y = X + 4$

$Y = X + 8$

$Y = X + 16$

Next-line prefetcher

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

Works well for L1 Icache and L1 Dcache.

What About this?

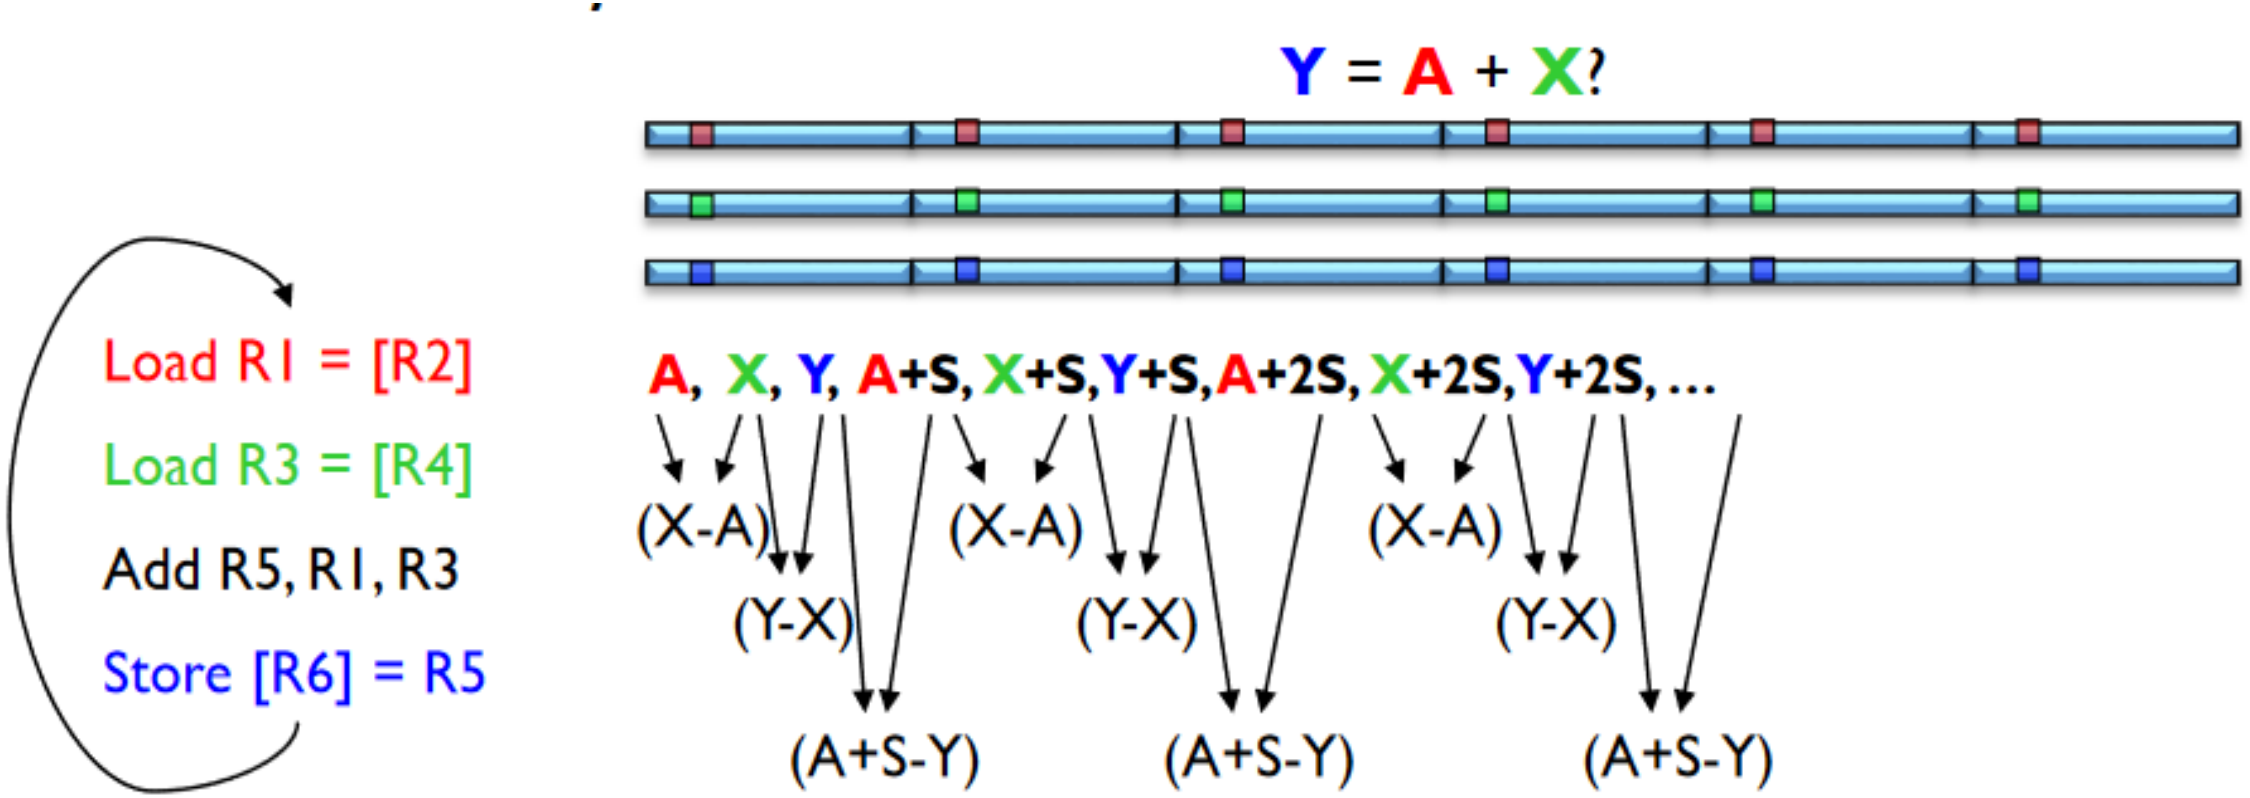

IP-stride prefetcher

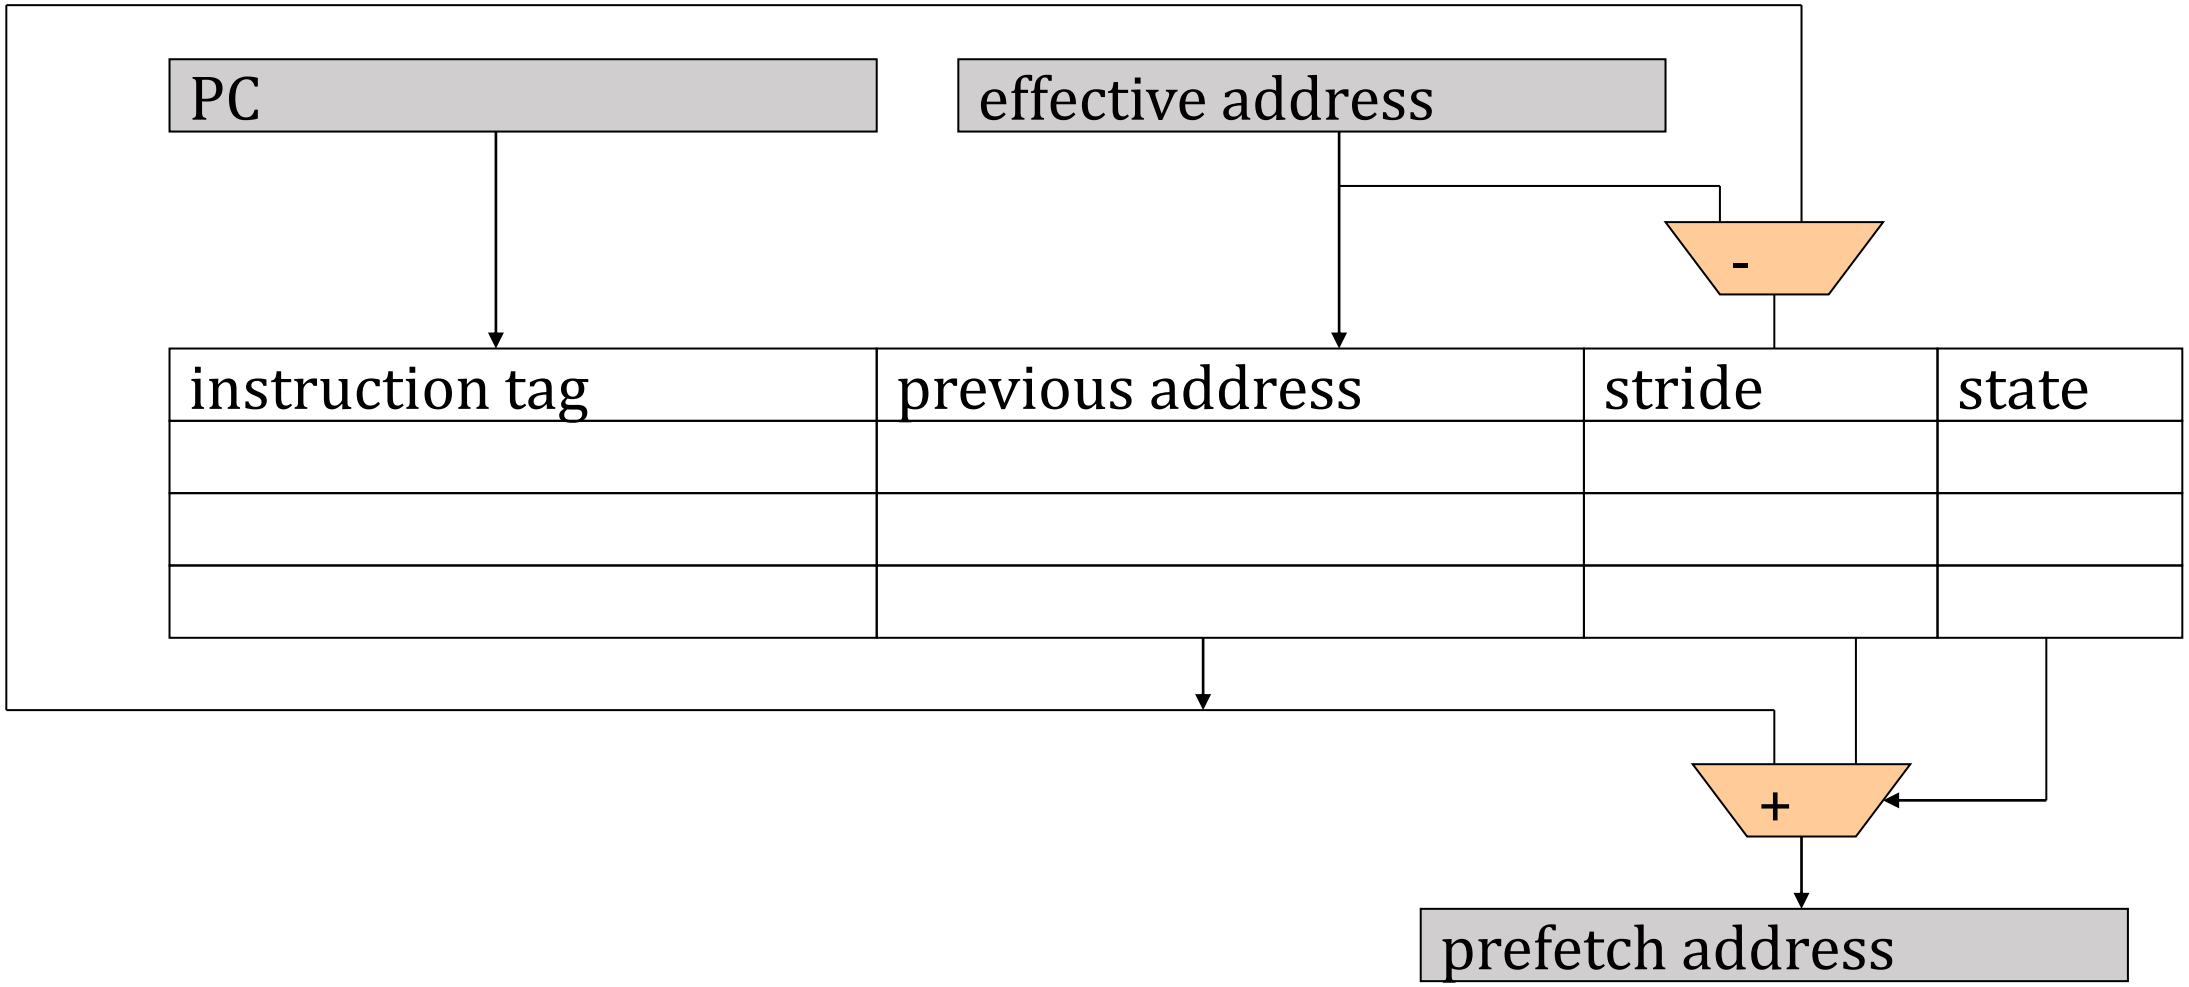

Metrics of interest

Accuracy

Coverage

Timeliness

If interested, have a read

Bouquet of Instruction Pointers: Instruction Pointer Classifier-based Hardware Prefetching

DPC3@ISCA '19

ISCA '20

https://www.cse.iitb.ac.in/~biswa/IPCP_ISCA20.pdf
https://biswabandan.medium.com/from-cricket-to-winning-the-data-prefetching-championship-at-isca-2019-7ffe4bf5a710

Choukran