A best-offset prefetcher

Pierre Michaud

2nd data prefetching championship, june 2015

DPC2 rules

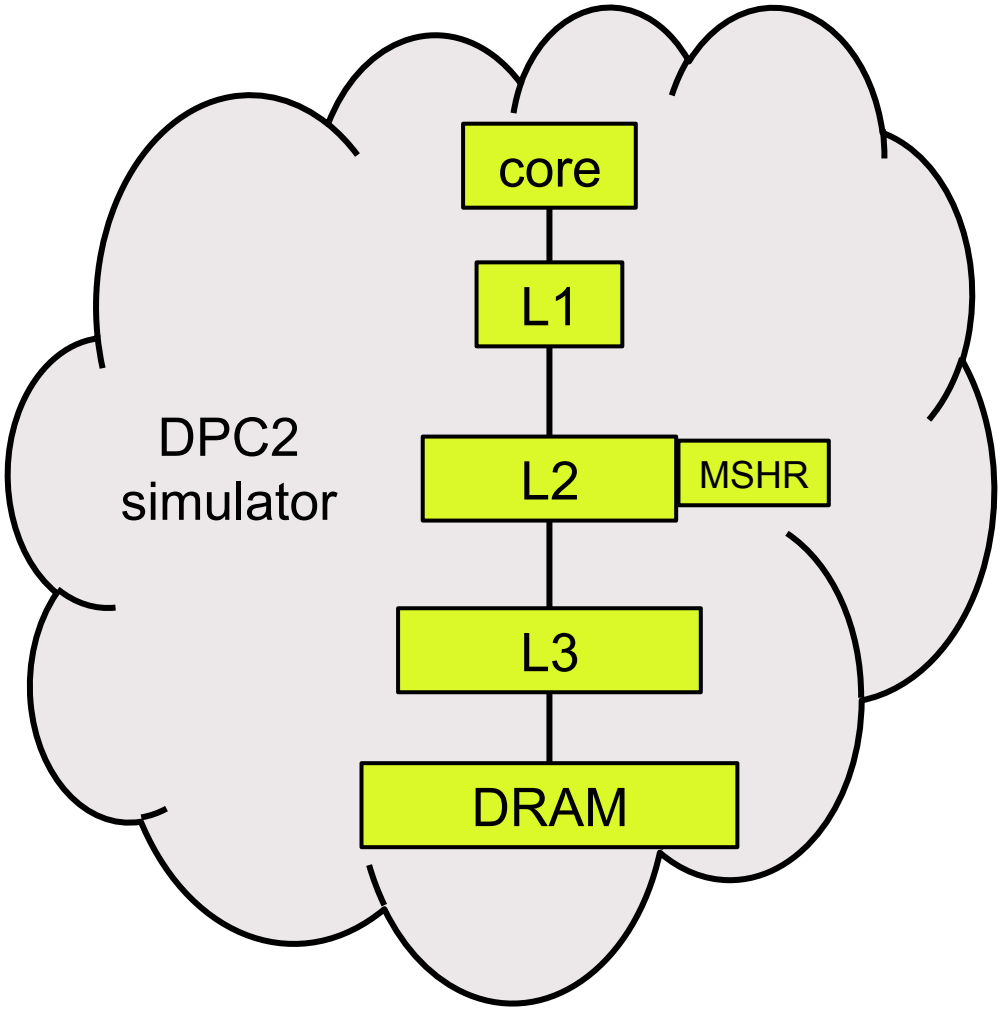

prefetcher

DPC2 rules

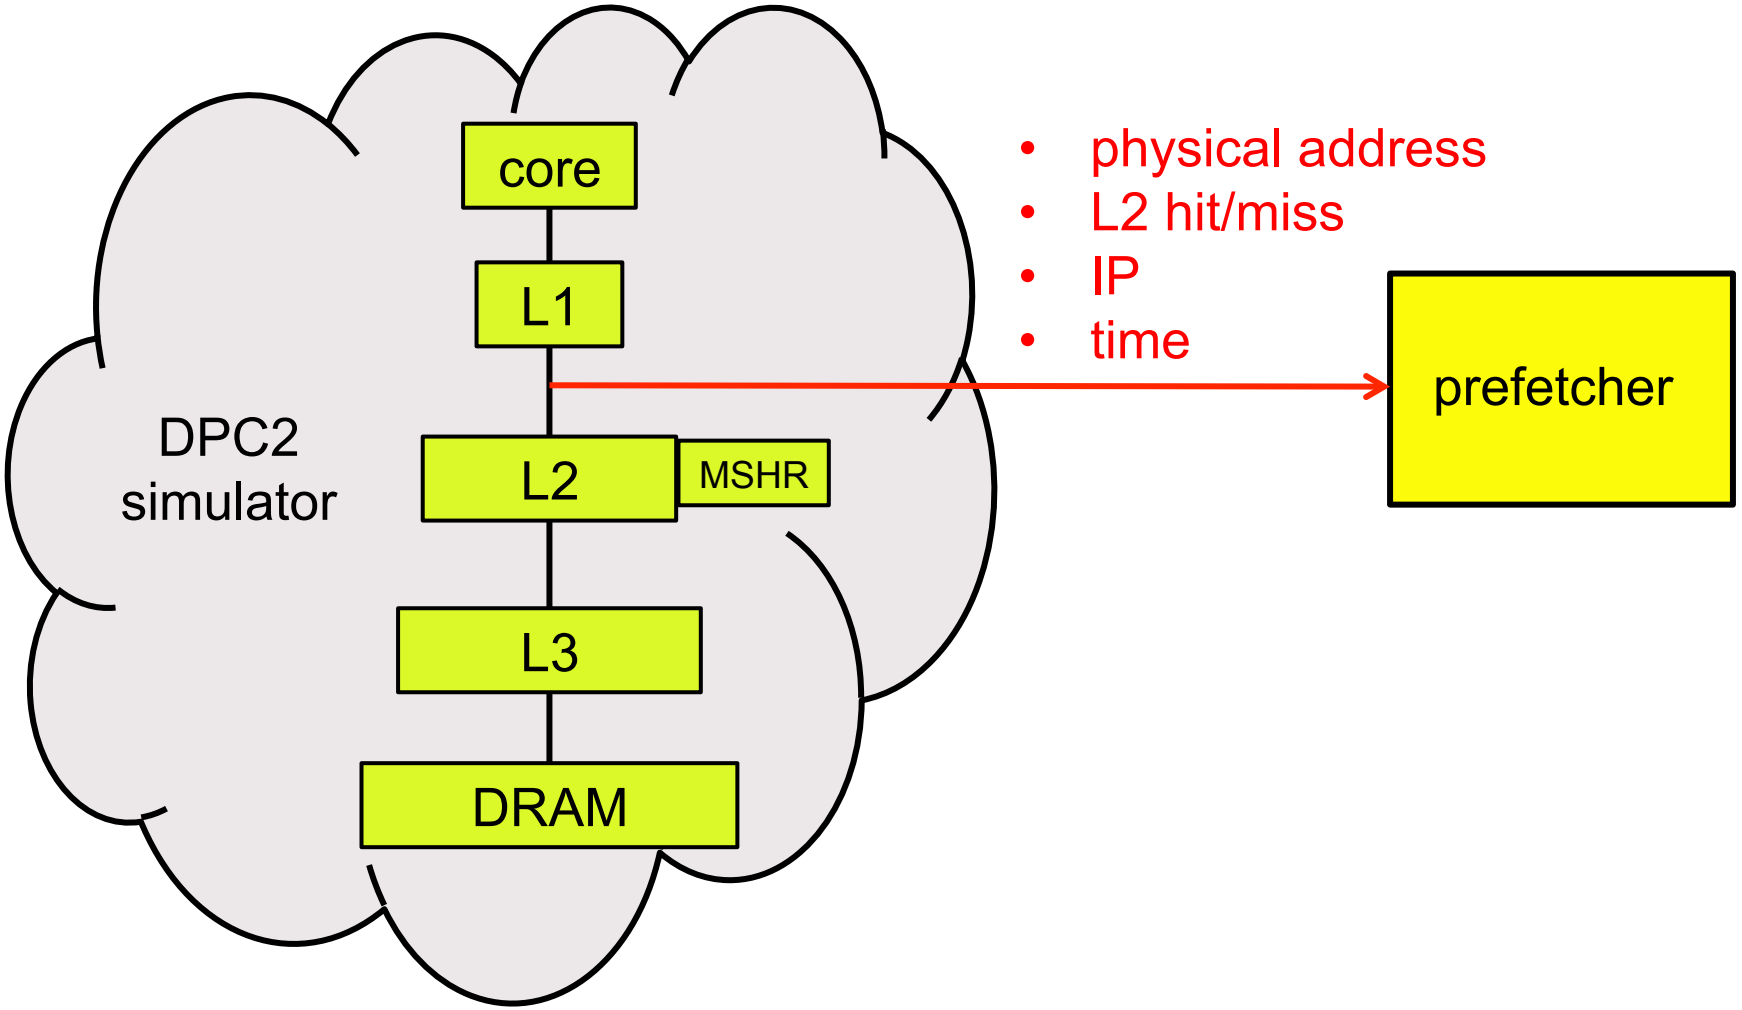
core
L1
DPC2
simulator
L2
MSHR
L3
DRAM
physical address
L2 hit/miss
IP
time
prefetcher

DPC2 rules

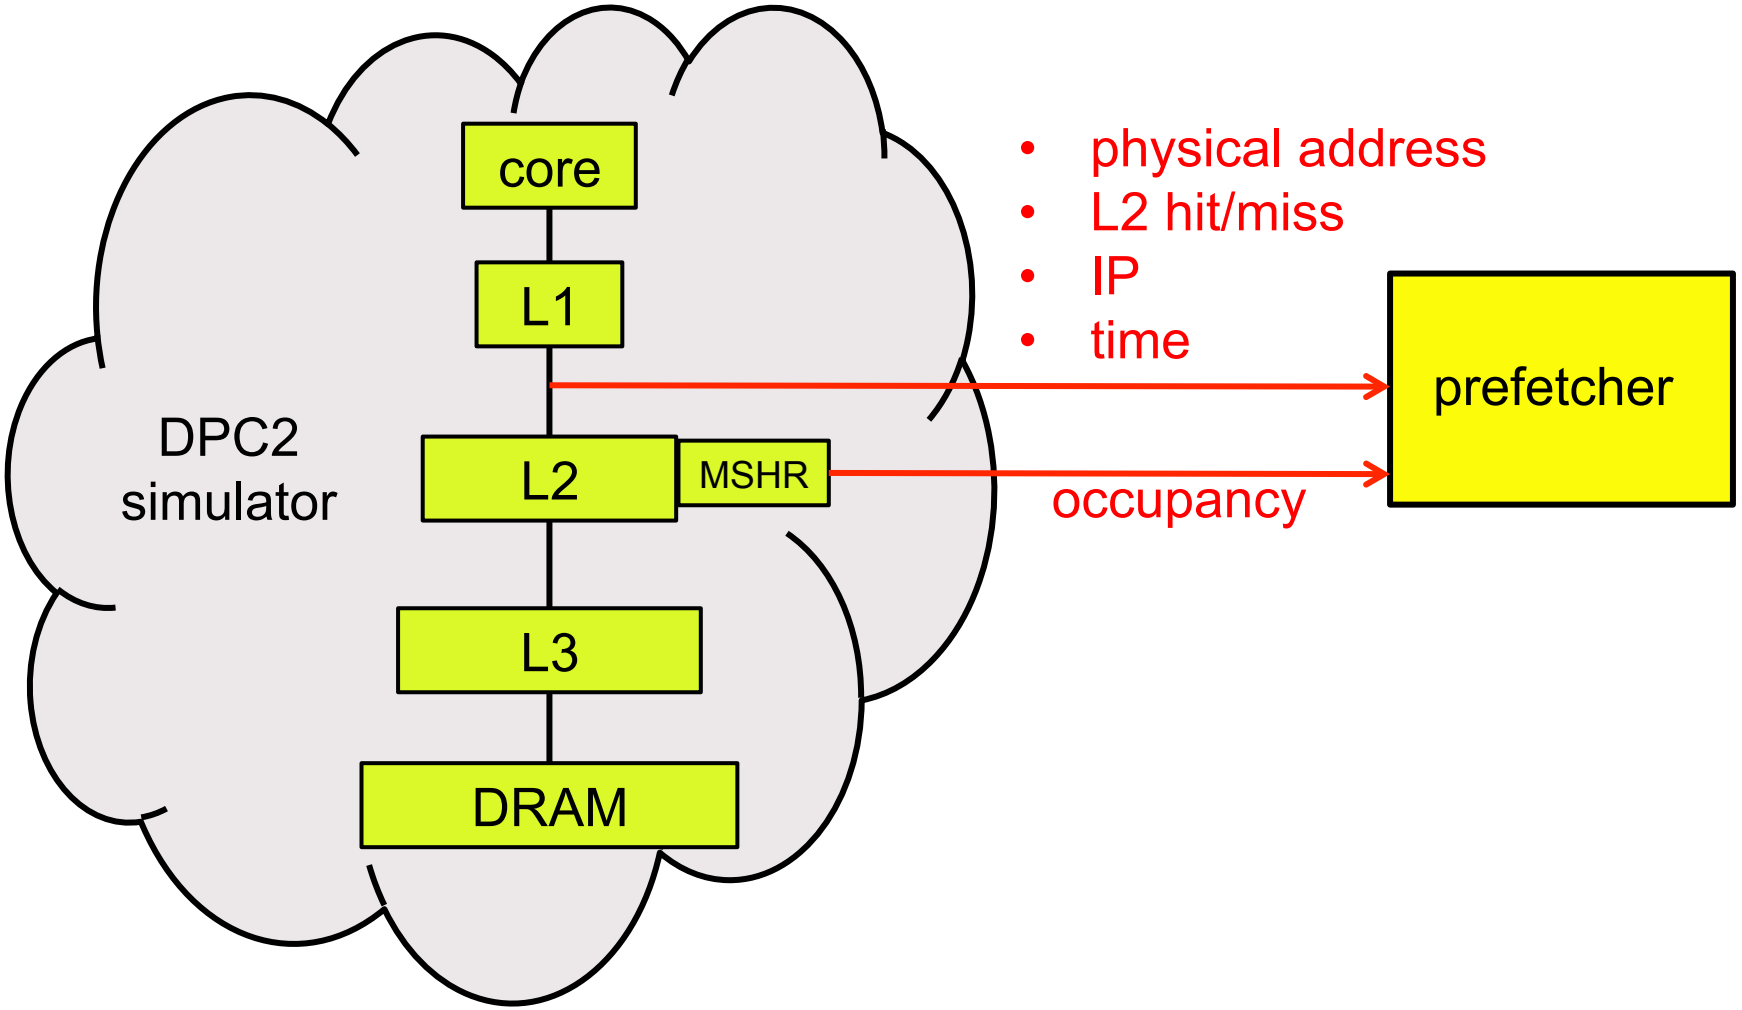
core
L1
L2
MSHR
L3
DRAM
DPC2
simulator
physical address
L2 hit/miss
IP
time
occupancy
prefetcher

DPC2 rules

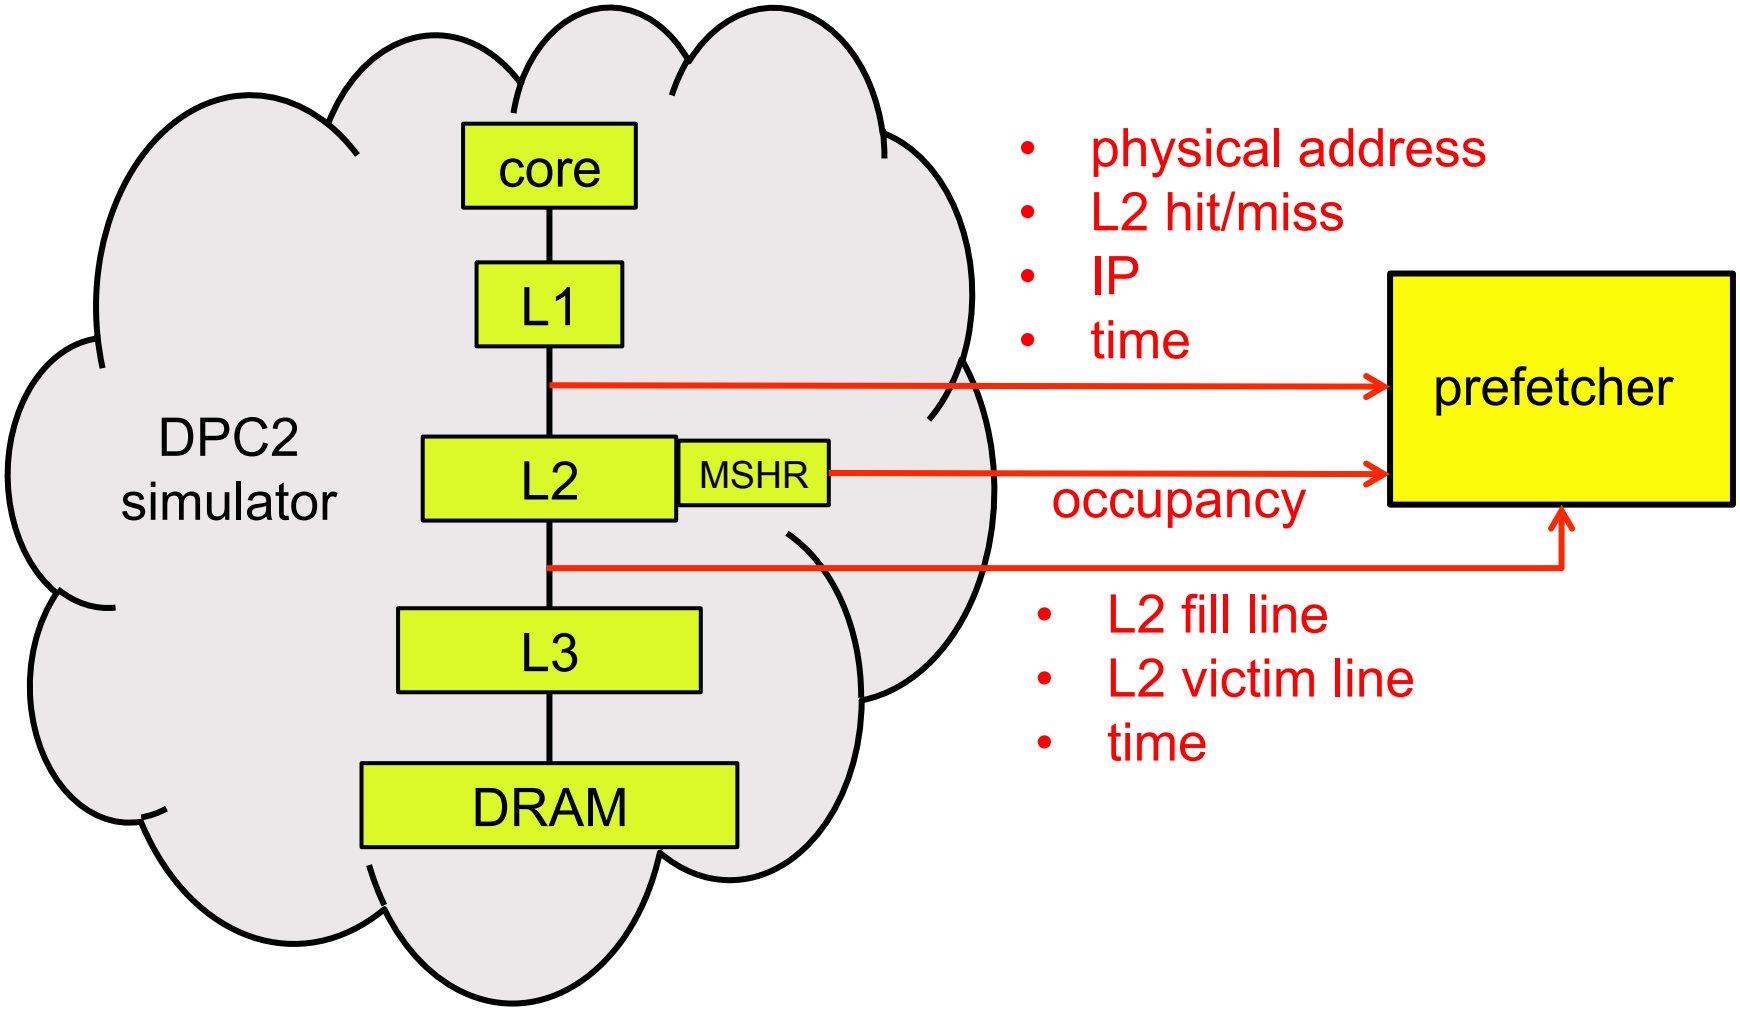
core
L1
L2
MSHR
L3
DRAM
DPC2
simulator
prefetcher
• physical address
• L2 hit/miss
• IP
• time
occupancy
• L2 fill line
• L2 victim line
• time

DPC2 rules

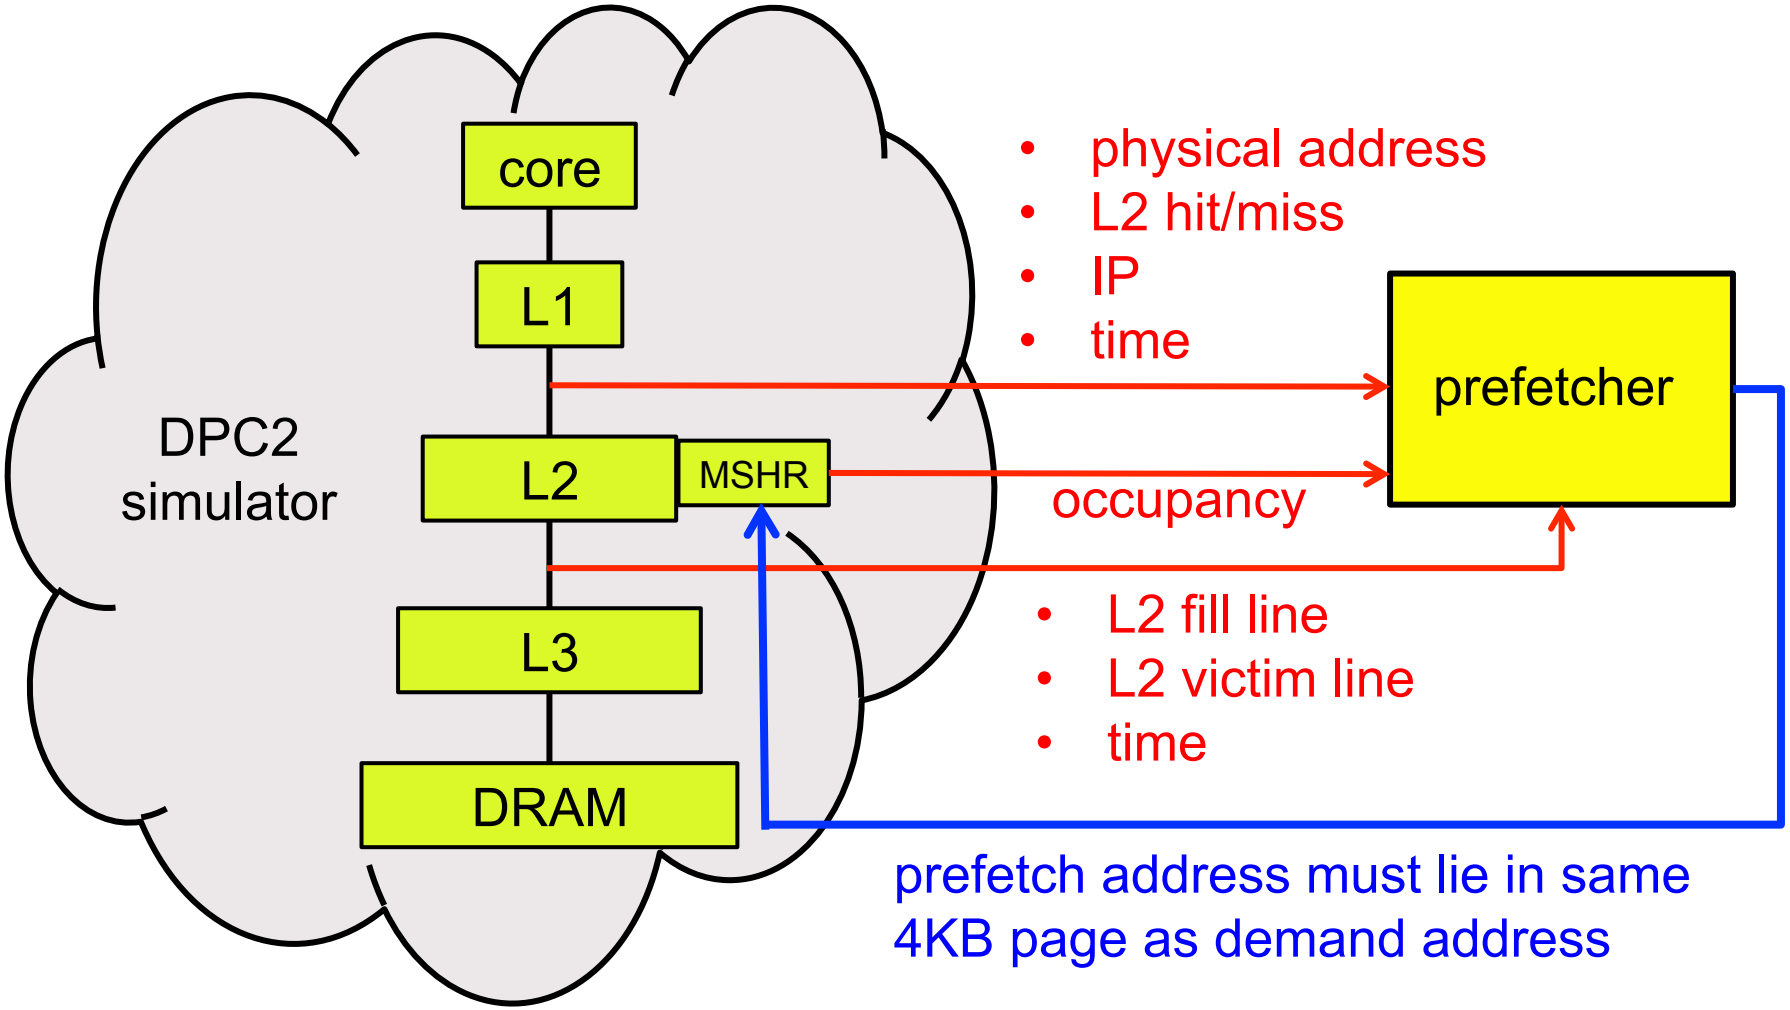
core
L1
L2
MSHR
L3
DRAM
DPC2
simulator
prefetcher
• physical address
• L2 hit/miss
• IP
• time
occupancy
• L2 fill line
• L2 victim line
• time
prefetch address must lie in same
4KB page as demand address

Offset prefetching

- Next-line prefetching ➔ O=1
- Full-fledged offset prefetcher ➔ varying offset
- Sandbox prefetcher (Pugsley et al., HPCA 2014)

Proposed Best-Offset (BO) prefetcher

- New method for setting the offset automatically
  - different from Sandbox
  - first implementation in an in-house simulator in 2011

- Bandwidth & cache pollution ➔ prefetch throttling method
  - somewhat specific to DPC2
  - DPC2 rules limit what can be done

Sequential stream

(neglect page boundary effect)

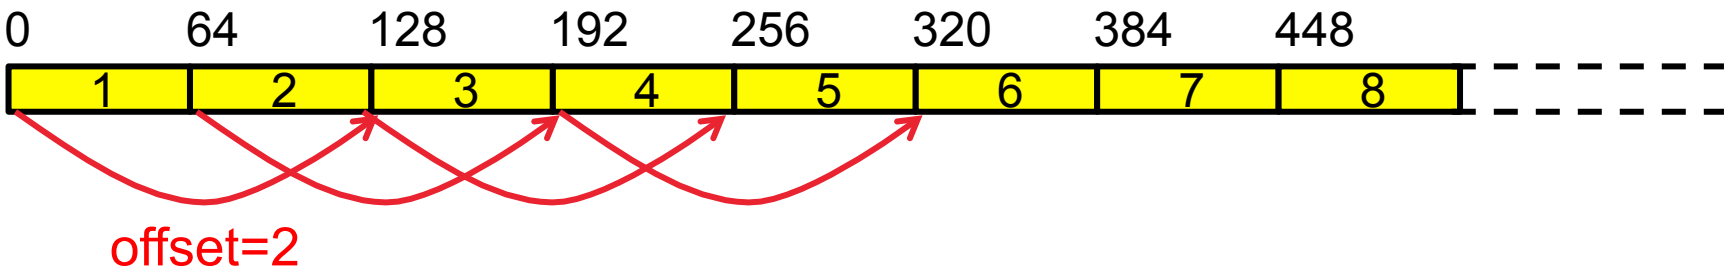

- if the offset is too small, prefetches may not be timely

Strided stream

example: stride=+96 bytes

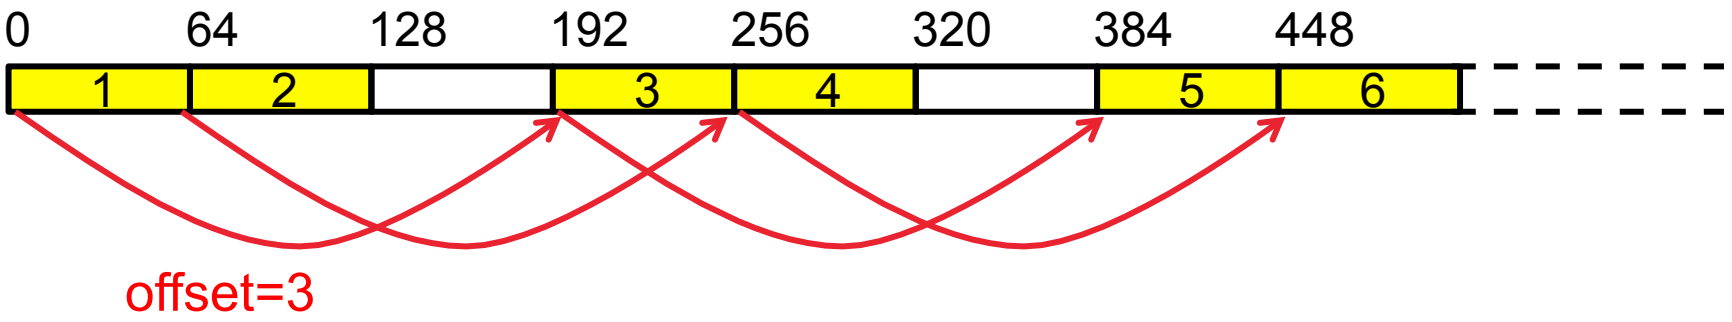

- constant byte-stride ➔ periodic sequence of line-strides (1,2,1,2,...)
- offset = sum of line-strides in a period (offset=1+2=3)
- ...or multiple of that sum (6,9,...)

Interleaved streams

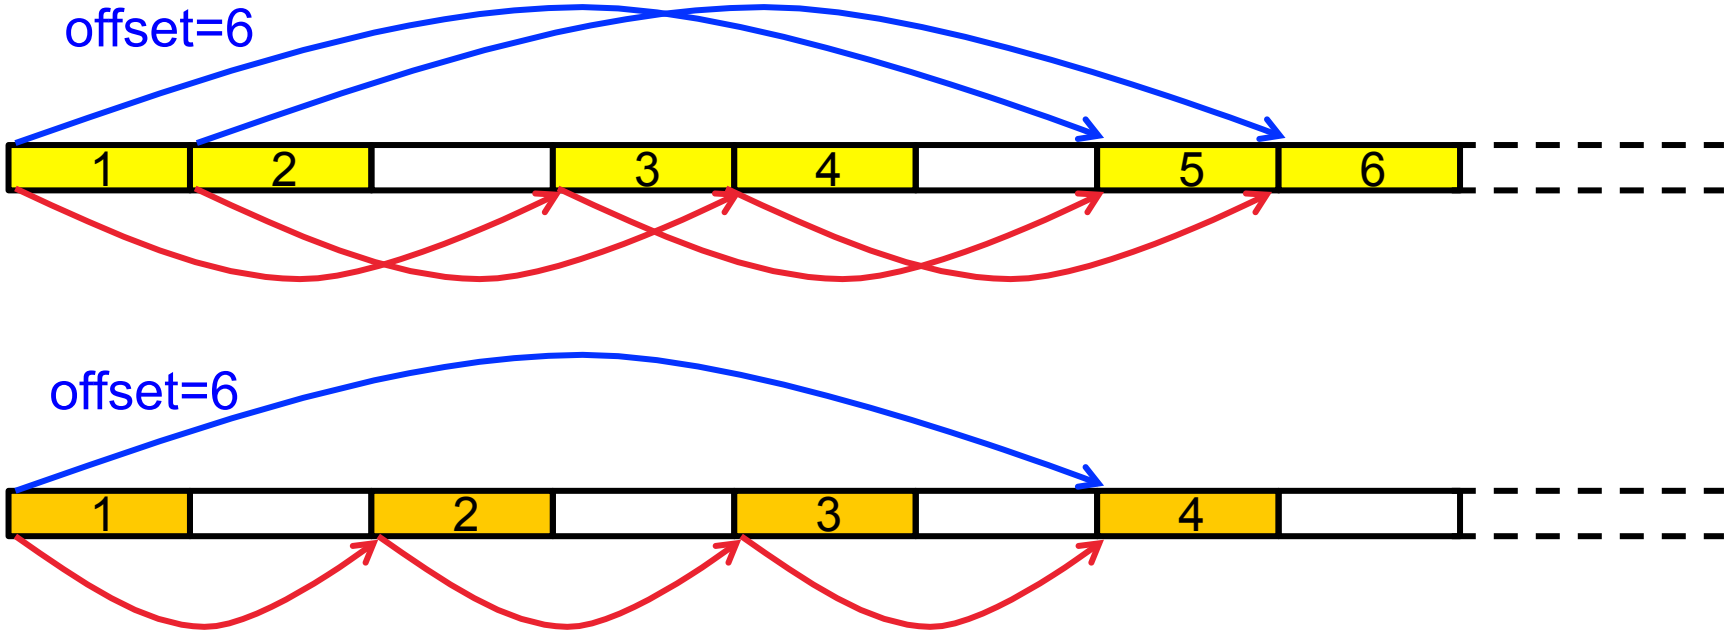

- 1st stream alone ➔ offset = multiple of 3
- 2nd stream alone ➔ offset = multiple of 2
- Both streams ➔ offset = multiple of 6

BO prefetcher: main idea

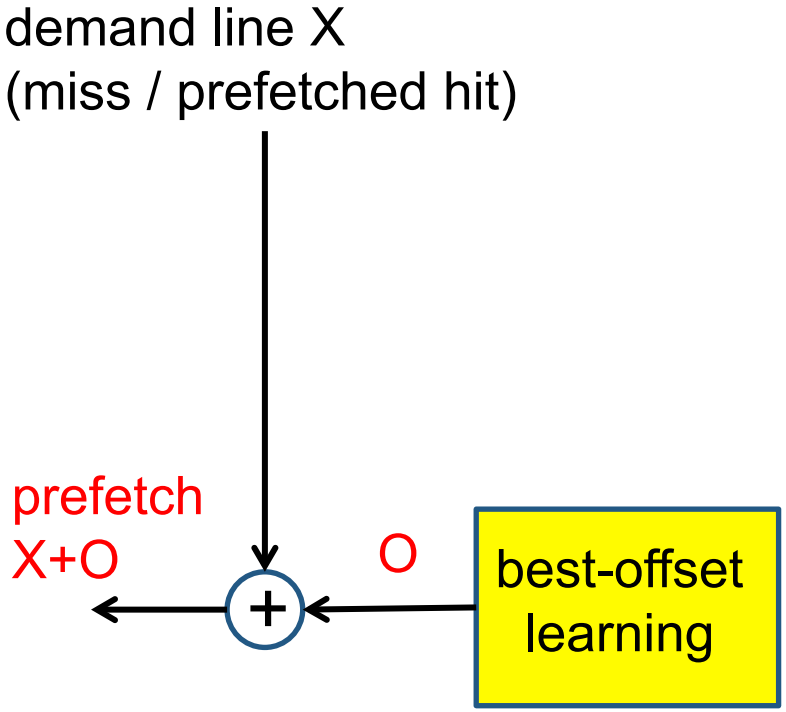

BO prefetcher: main idea

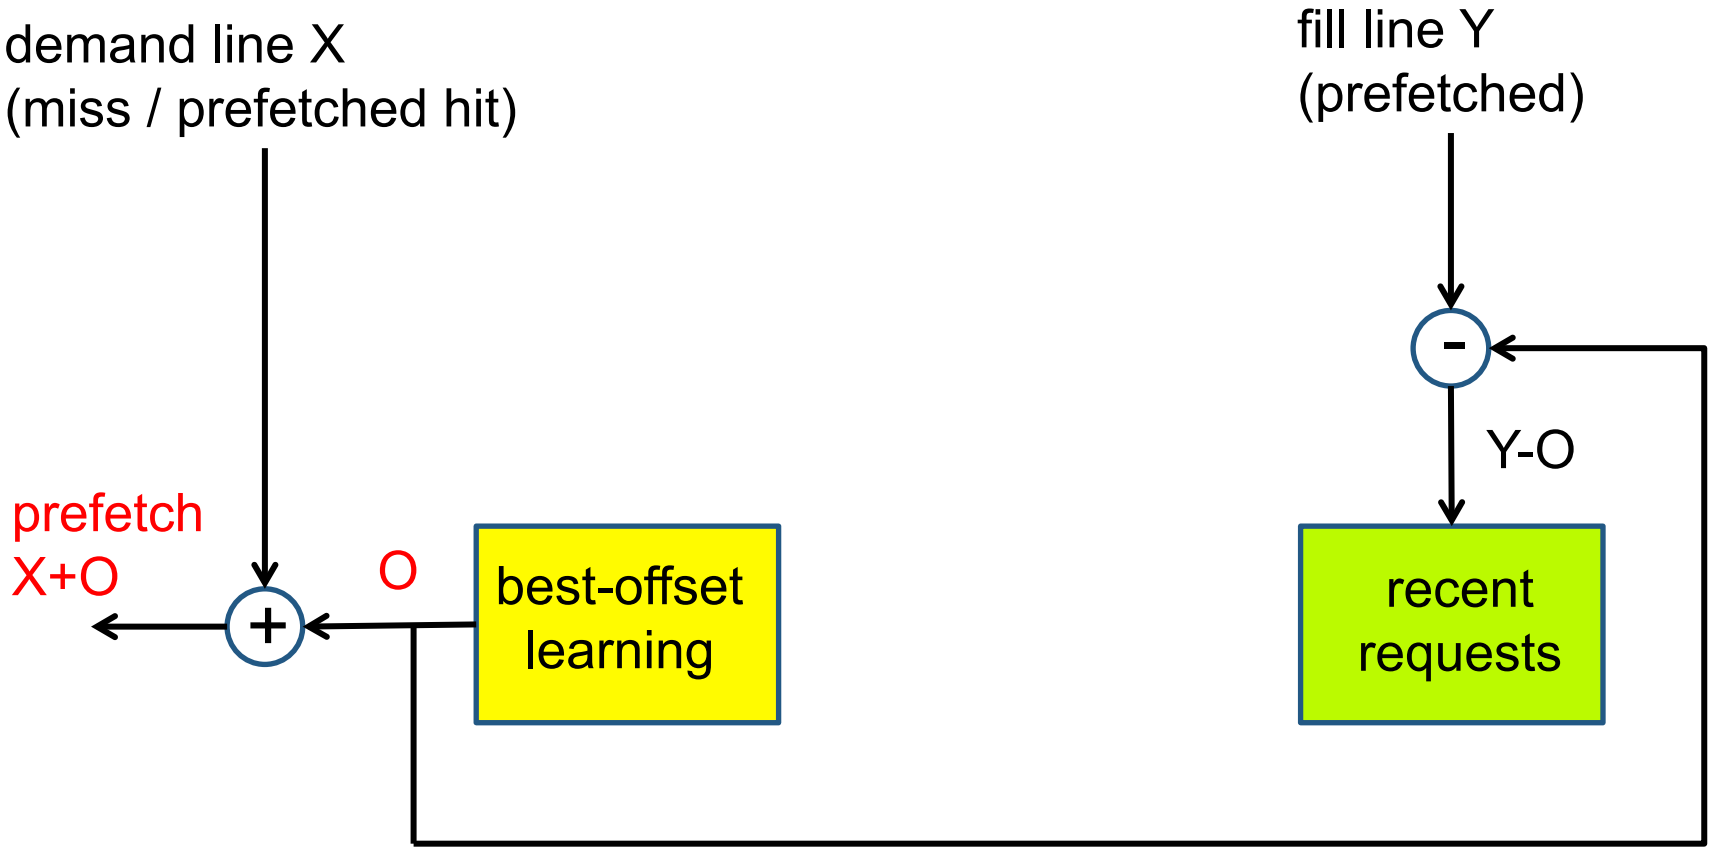

BO prefetcher: main idea

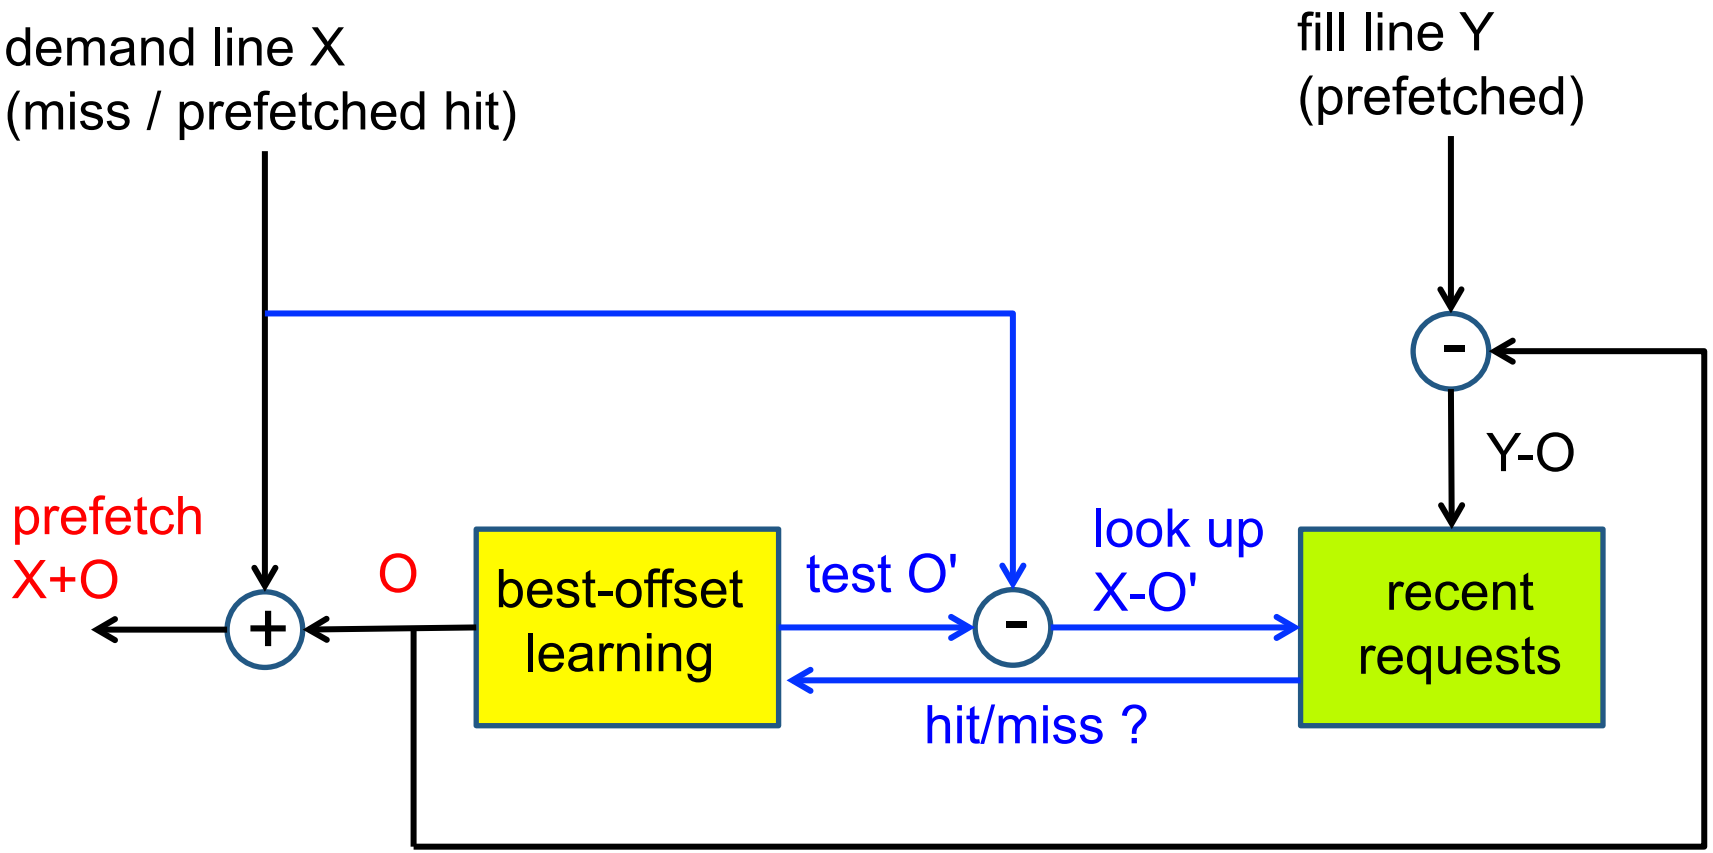

Recent Requests (RR) Table

- in 2011: 64-entry fully-associative FIFO
- for DPC2: two direct-mapped banks with different hashing
  - resembles 2-way skewed-associative
  - 2 x 64 x 12-bit tags ➔ 1536 bits
- Write same tag redundantly in both banks

Learning the best offset

- 46 different offsets evaluated
  - 23 positive + 23 negative
  - 1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,18,20,24,30,32,36,40
- Each offset has a 5-bit score
  - 46 x 5 ➔ 230 bits
- Test the 46 offsets successively (46 L2 accesses) = one round
  - if hit in RR table for an offset, increment its score
- Learning phase finishes after 100 rounds, or if one of the scores reaches 31
  - select the offset with the greatest score ➔ this is the new prefetch offset
  - new learning phase starts ➔ reset scores

Prefetch timeliness vs. prefetch accuracy

- BO prefetcher tries to do timely prefetches
- However...
- Sometimes, better to choose a smaller offset, even if it generates late prefetches
  - Example: short sequential streams
- Imperfect solution: delay queue

BO prefetcher with a delay queue

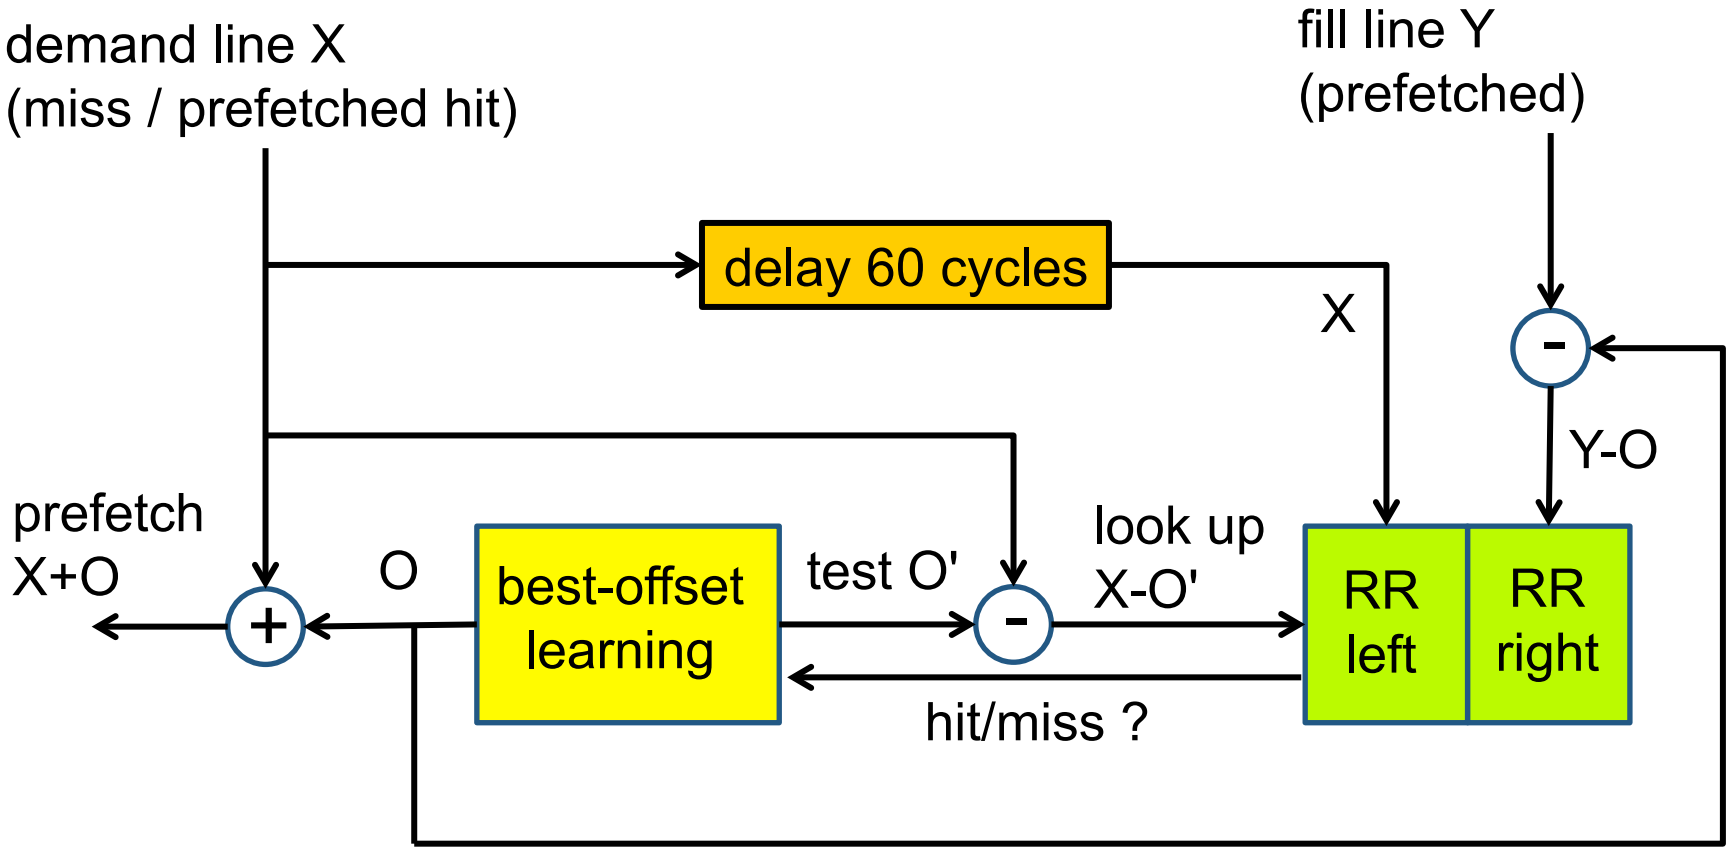

Prefetch throttling (DPC2)

- Turn prefetch on only if BO score > BADSCORE
  - DPC2➔ BADSCORE=1 (10 for small L3 config)
  - best-offset learning continues while prefetch is off

- Drop prefetch request if MSHR occupancy is above a threshold
  - Vary MSHR threshold depending on BO score and L3 access rate

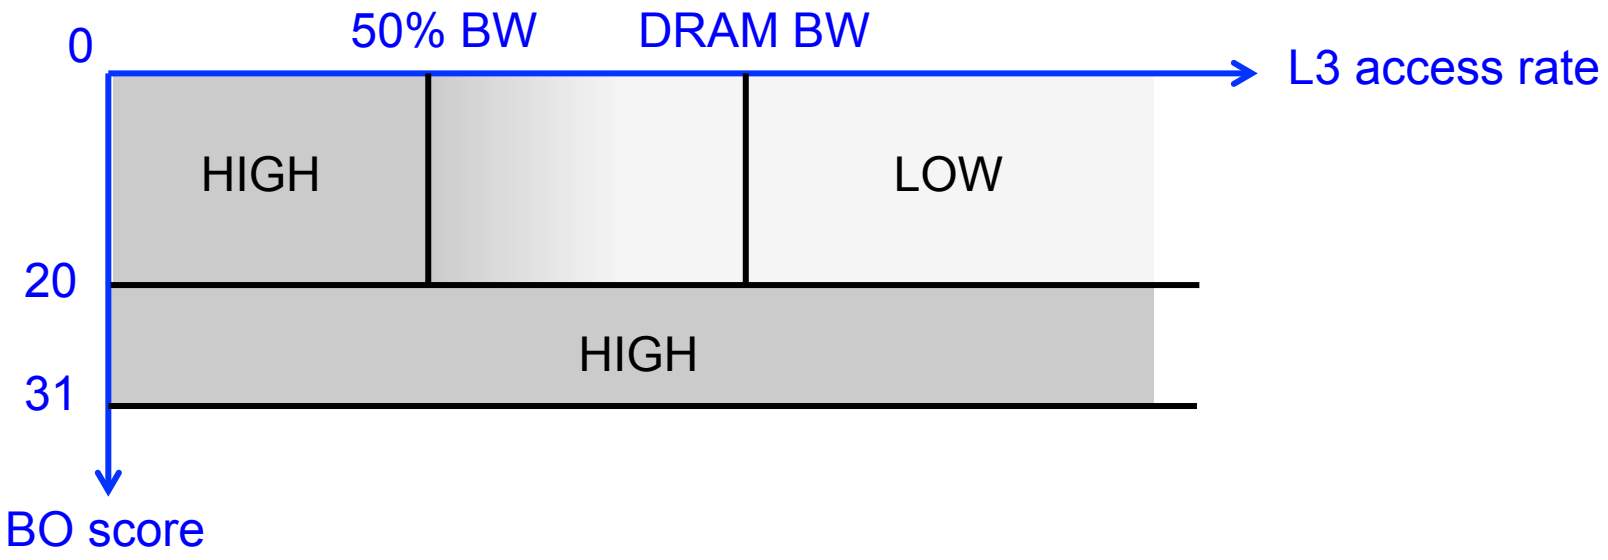

State (number of bits)

	bits
prefetch bits (1 bit per L2 line)	2048
recent requests (2x64x12)	1536
scores (46x5)	230
delay queue (15 slots)	473
miscellaneous	74
TOTAL	4361

fixed vs. adaptive offset (437.leslie3d)

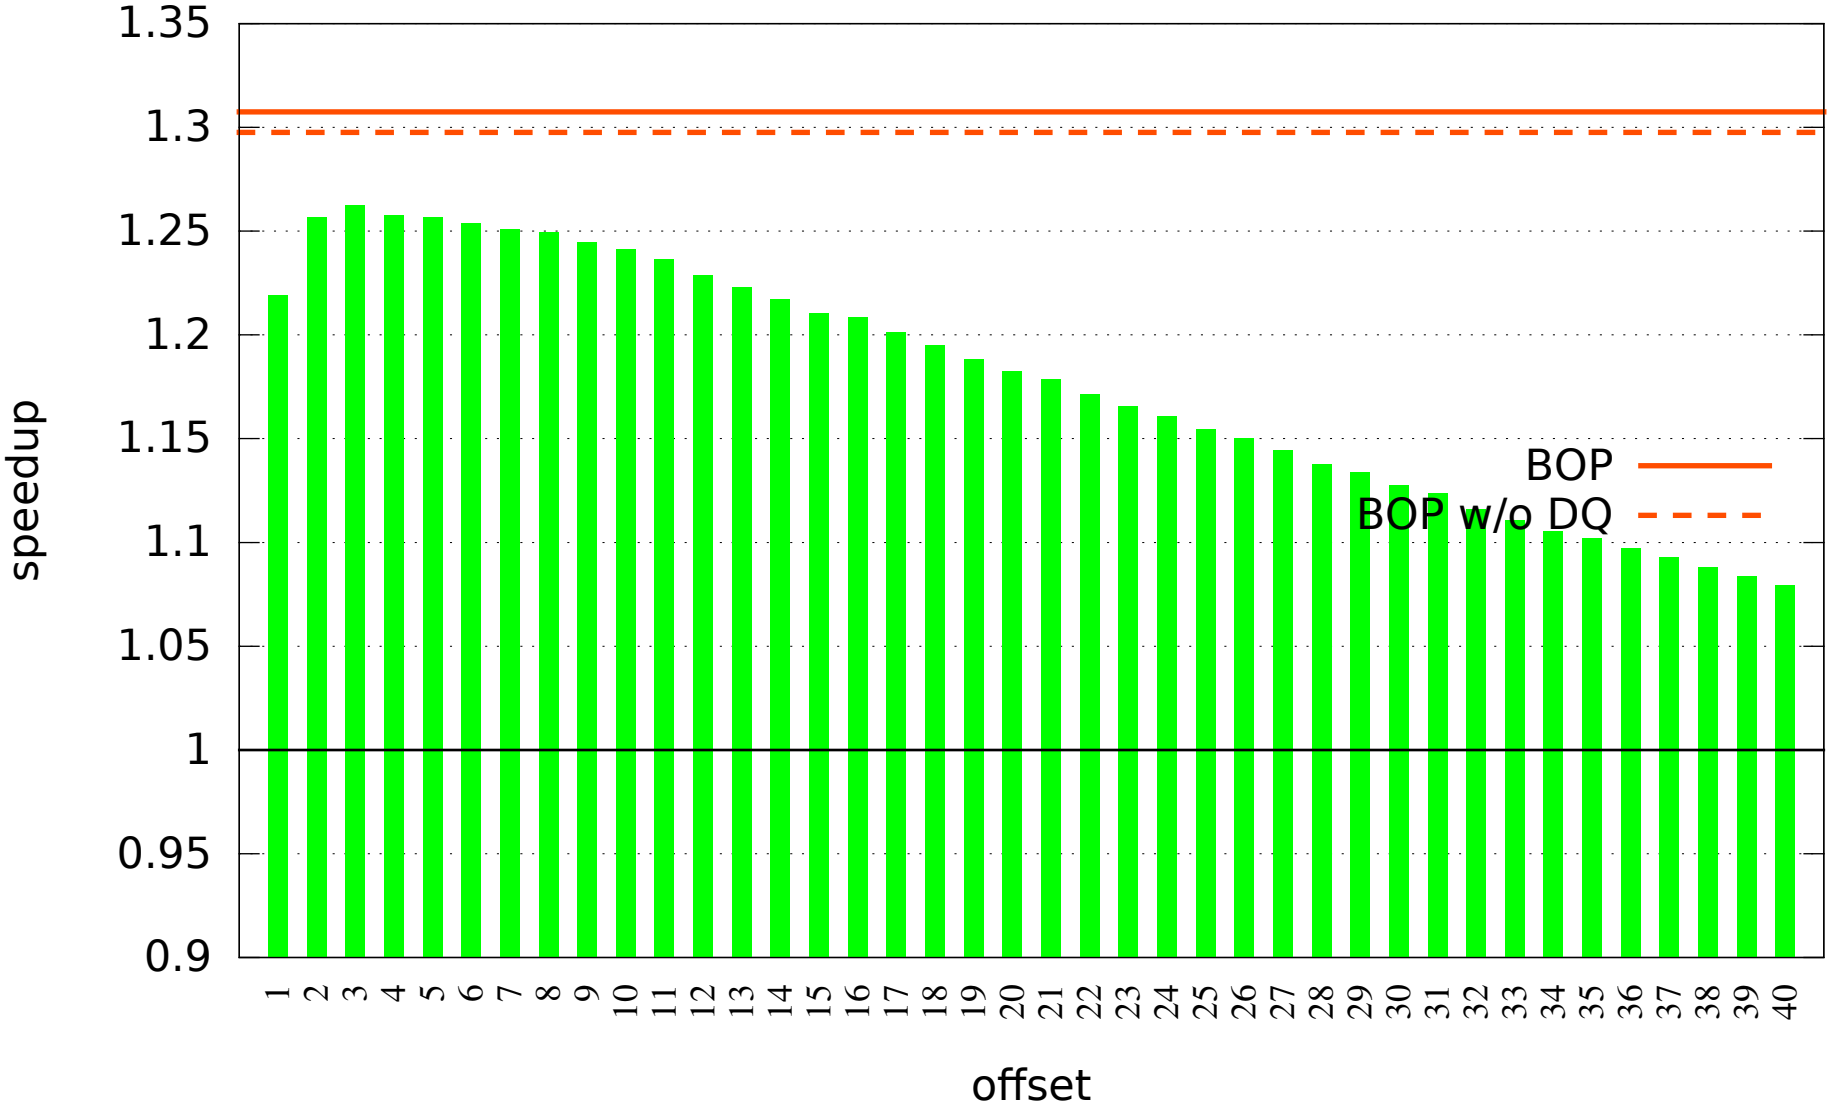

Fixed vs. adaptive offset (433.milc)

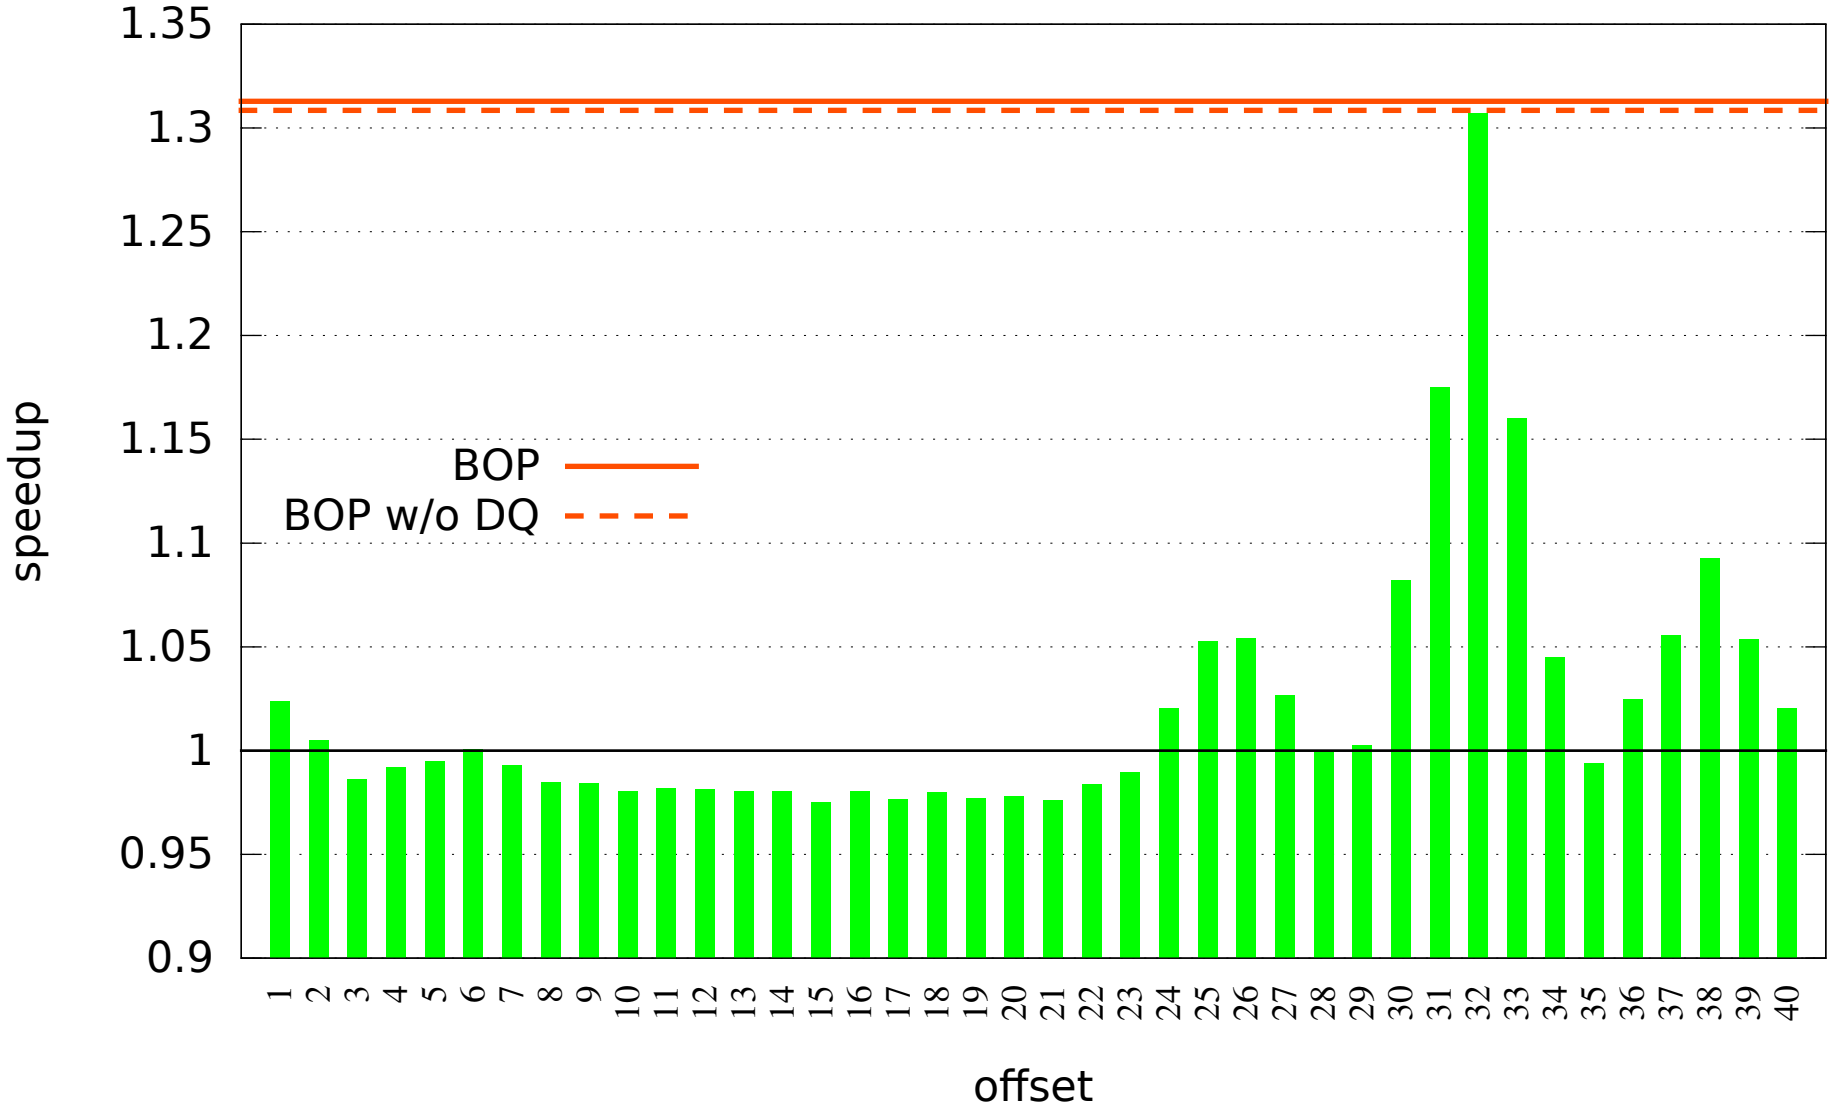

Fixed vs. adaptive offset (434.zeusmp)

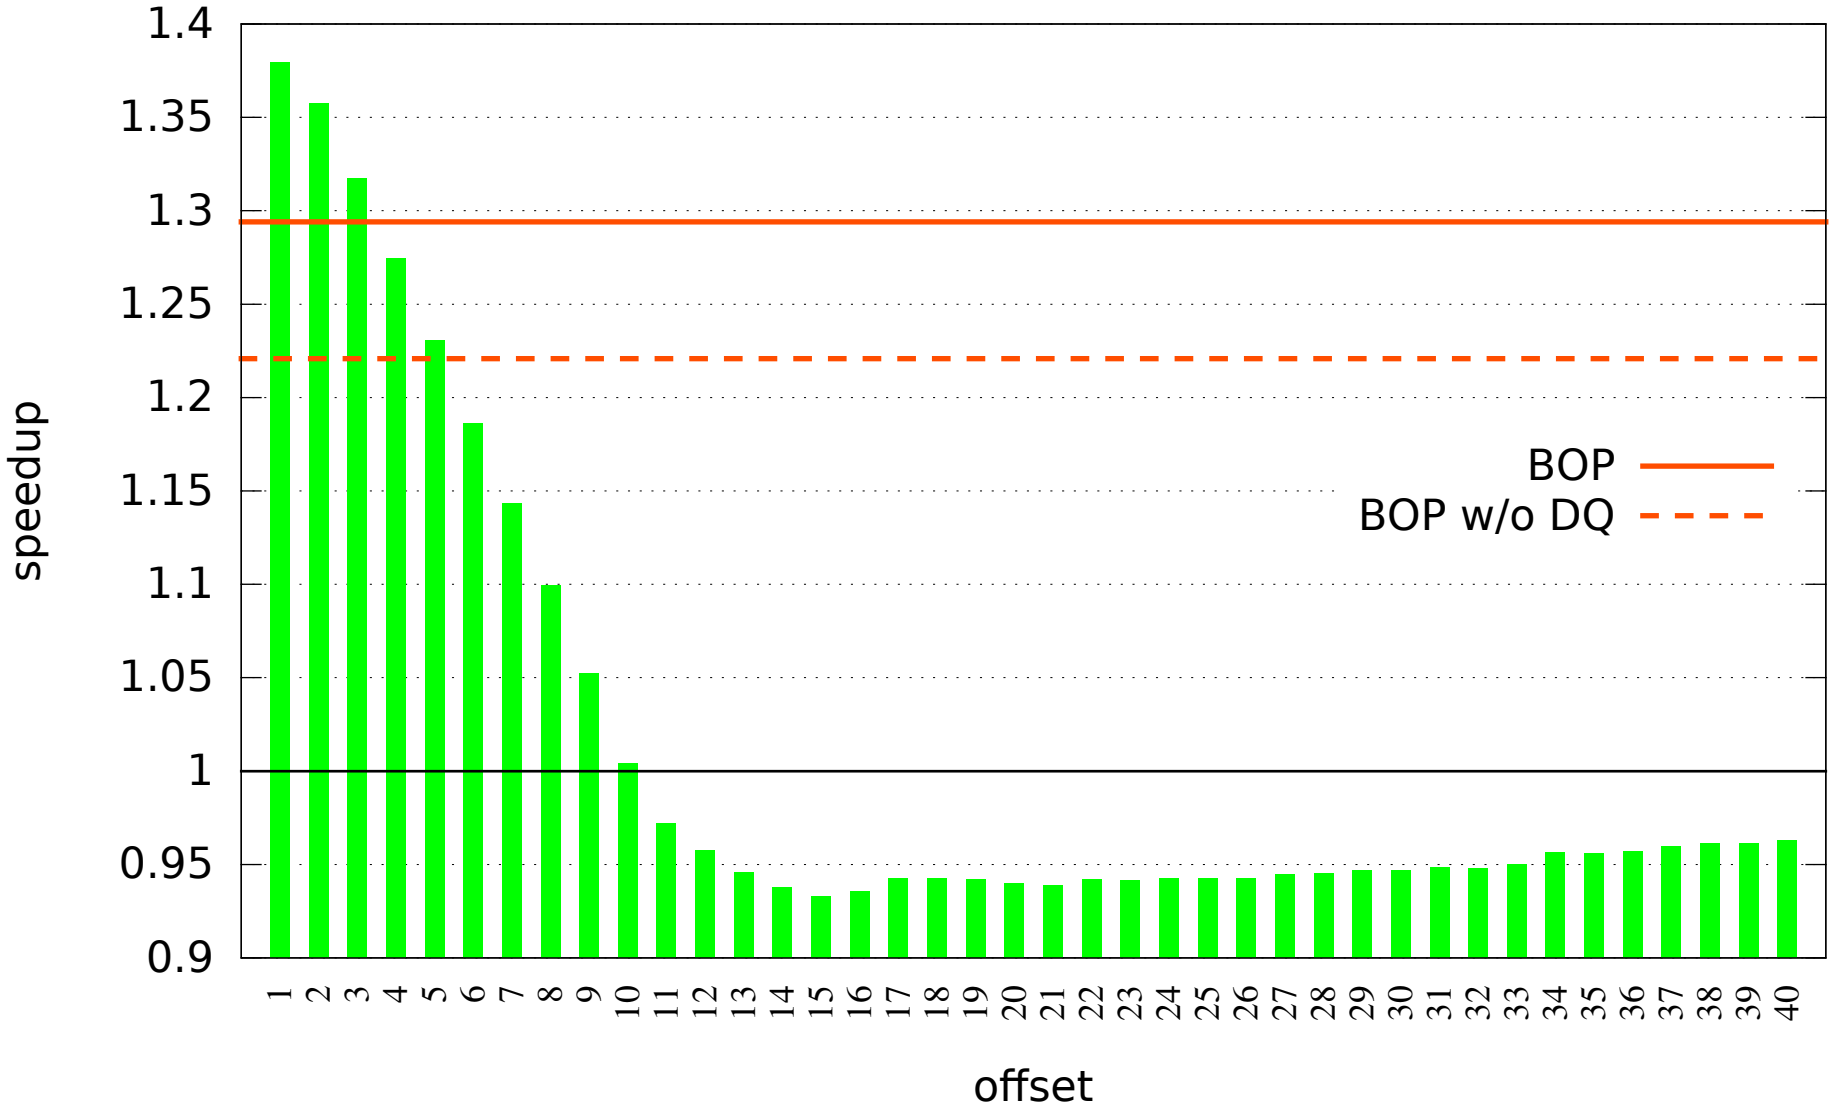

BO prefetcher vs. Sandbox prefetcher

- Sandbox prefetcher (Pugsley et al., HPCA 2014)
  - first published full-fledged offset prefetcher
  - fake prefetches ➔ evaluate an offset by setting bits in a Bloom filter
  - if demand access hits in Bloom filter ➔ fake prefetch successful
  - prefetch timeliness not considered
  - Sandbox method is orthogonal to offset prefetching

- BO prefetcher
  - no fake prefetches
  - strive for prefetch timeliness

FIN