

# Neighborhood Prefetching

David M. Koppelman

Electrical & Computer Engineering Department  
Louisiana State University  
Baton Rouge, LA U.S.A.

Prefetch:

Fetching a cache line in anticipation of its use.

Motivation:

Cache misses becoming larger component of execution time.

Multiprocessors add interconnect and protocol delays.

Challenges:

Prefetching addresses that will be used. (Easy)

Avoiding addresses that won't be used. (Hard)

## Sequential

Infinite sequential sequence.

Sequential sequence indicated in **bold**.

100, 3000, **101**, 2000, 3000, **102**, **103**, 6, ...

## Stride

Infinite stride sequence.

Stride 10 sequence indicated in **bold**.

100, 3000, **110**, 2000, 3000, **120**, **130**, 6, ...

Real address reference sequences are finite, guaranteeing useless prefetches.

### Unneeded Prefetch Problems in All Systems

Resource Consumption. (Cache ports, interconnect bandwidth, etc.)

Eviction of useful lines.

### Unneeded Prefetch Problems in Multiprocessors

Adds to false sharing.

### Advantages

Larger variety of address sequences than stride.

Less tendency to prefetch unneeded lines.

Execution of following code fragment ...

```
! Iteration 1: r2 = 0x2000, Line Size 0x100 Characters
!  
A: load r1, [r2]  
...  
B: load r3, [r2+0x500]  
...  
C: load r4, [r2-0x700]
```

... generates ...

Address Seq.: 0x2000, ..., 0x2500, ..., 0x1900, ...

## Neighborhood Prefetch Terminology

<i>Term:</i>	<i>Example</i>
Base: Address used as reference.	0x2000
Initiator: Instruction accessing base address.	A: load r1, [r2]
Offset: Distance from base, in lines.	0, ..., +5, ... -7
Neighborhood: Set of small-magnitude offsets.	{0, +5, -7}

Consider ...

```
! Iteration 1: r2 = 0x2000 (Base), Line Size 0x100 Characters
!  
A: load r1, [r2] ! (Initiator)  
...  
B: load r3, [r2+0x500]  
...  
C: load r4, [r2-0x700]
```

... generates ...

Address Seq.: 0x2000, ..., 0x2500, ..., 0x1900, ...

Neighborhood: {0, +5, -7}

Notice that the neighborhood ...



Consider ... and

```
! Iteration 2: r2 = 0x8000 (Base), Line Size 0x100 Characters
```

```
!
```

```
A: load r1, [r2] ! (Initiator)
```

```
...
```

```
B: load r3, [r2+0x500]
```

```
...
```

```
C: load r4, [r2-0x700]
```

... generates ...

Address Seq.: 0x8000, ..., 0x8500, ..., 0x7900, ...

Neighborhood: {0, +5, -7}

Notice that the neighborhood ... is the same in both cases.

Base address and neighborhood used to construct prefetch addresses.

Hardware monitors addresses that **miss** the cache.

Hardware determines neighborhoods and stores them in PC-indexed table.

Table checked on miss; if entry found used for prefetching.

Superset (loosely) of existing schemes.

Can detect sequential patterns:  $\{0, 1, 2, 3, \dots\}$ .

Can detect stride patterns:  $\{0, 10, 20, 30, \dots\}$ .

Can detect arbitrary patterns:  $\{0, 10, 20, 50\}$ .

Better at avoiding unneeded and harmful prefetches.

Prefetch candidates based on past usage by instruction.

Other schemes inevitably fetch unneeded, possibly harmful, lines.

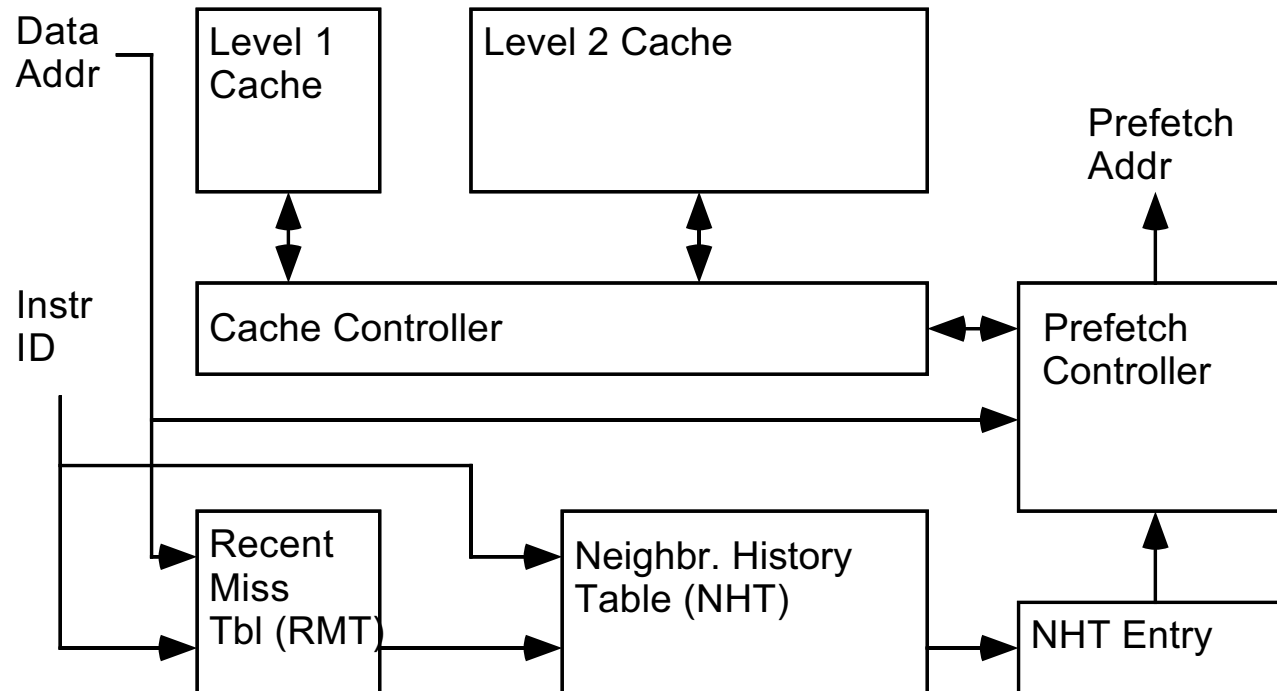
Particularly important in multiprocessors: less false sharing.

## Main Parts

*Recent Miss Table (RMT)*—for determining neighborhoods.

*Neighborhood History Table (NHT)*—stores neighborhoods.

Prefetch Controller—handles prefetches.



## Recent Miss Table (RMT)

Entries for recently encountered neighborhoods.

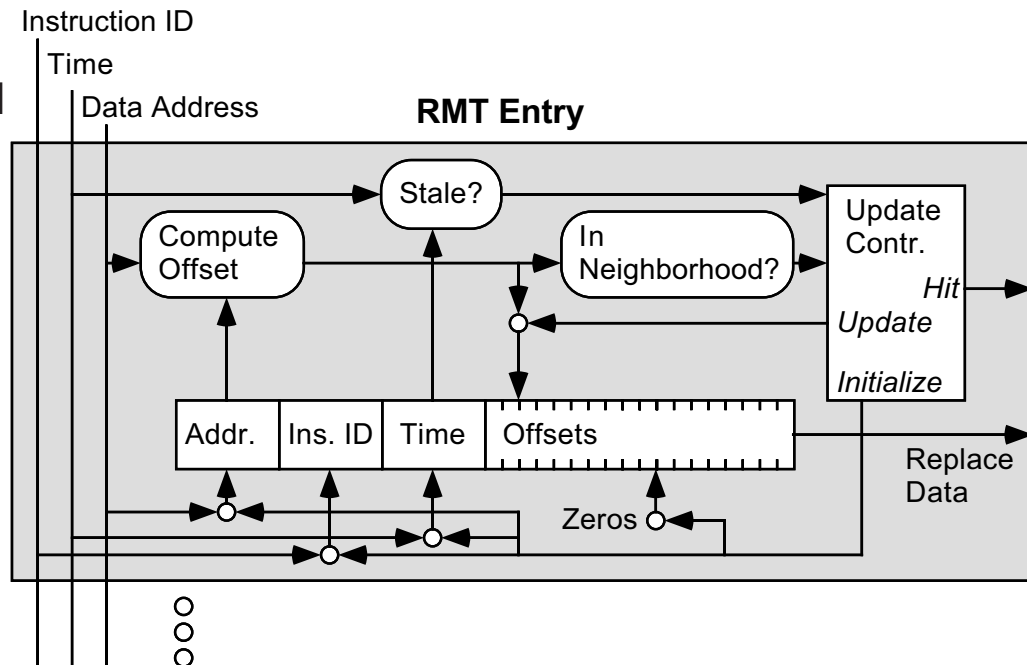
Each entry holds:

Base Address.

Offsets.

Initiating instruction.

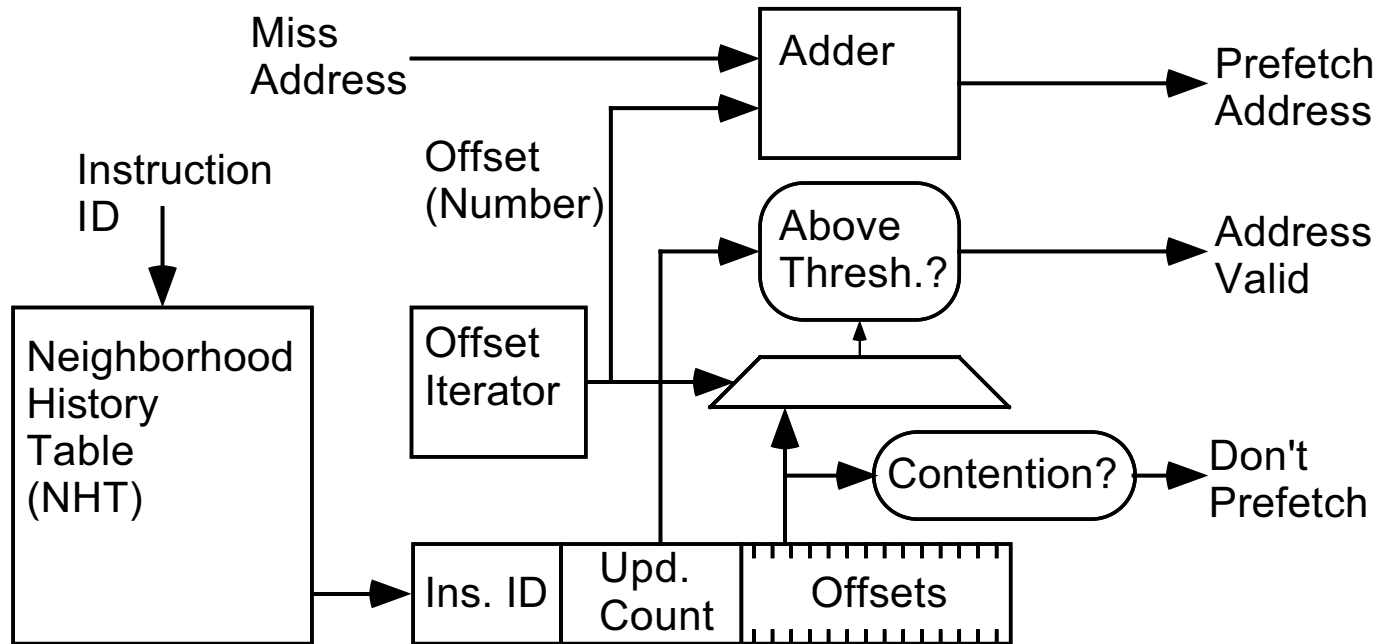
And other data.



On miss update existing entry or create new ones.

Entries occasionally moved to neighborhood history table.

## Neighborhood History Table



### Neighborhood History Table

Entry for memory access instructions.

Retrieved on miss using address of missing instruction.

Used to construct prefetch addresses.

Evaluated by execution driven simulation using Proteus.

Machine model:

Sixteen-node multiprocessor.

Full-map directory-based cache coherence.

In-order execution, four-way issue, SPARC ISA.

Virtual memory, nonblocking stores, blocking loads.

Mesh topology.

Execution Bottlenecks

Cache and Network Latency (all accesses)

Hot Spots at Memories (application dependent)


## Base Configuration Parameters

Simulation Parameter	Value
System Size	16 processors
Network Topology	4 × 4 mesh
VM Page Size	2 <sup>12</sup> bytes
TLB Capacity	64 entries
TLB Replacement	LRU, fully assoc.
Cache Size	2 <sup>7</sup> sets
Cache Associativity	8, LRU Repl.
Cache Line Size	64 bytes
L1 Cache Hit Latency	1 cycle
L2 Cache Hit Latency	7 cycles
Total L2 Miss Latency	50 (min), 135 (typ)
Directory Size	full map
Completion Buffer	5 stores
Raw Memory Latency	10 cycles
Protocol Message Size	8 bytes (plus data)
Network Interface Width	4 bytes
Network Link Width	4 bytes
Hop Latency	20 cycles (plus waiting)
Neighborhood Size	16 offsets
NHT Size	256 entries
RMT Size	8 entries




# SPLASH-2 Benchmarks Used


Cholesky - Processor States




FFT: Processor States (65536 elements)




LU: Processor States (128 x 128)




Processor States, Radix 262,144 Keys




Processor States: Barnes, 16384 Particles




Processors States: FMM, 16384 Particles



Processor States: Ocean, 258 x 258



Processor States: Water N Squared, 512 Molec



Ran SPLASH-2 Benchmarks

Compared:

Neighborhood Prefetch (Cache Size  $\frac{7}{8}$  Other Systems)

Adaptive Sequential Prefetch (Preupgrades, Tuned for System)

Conventional System

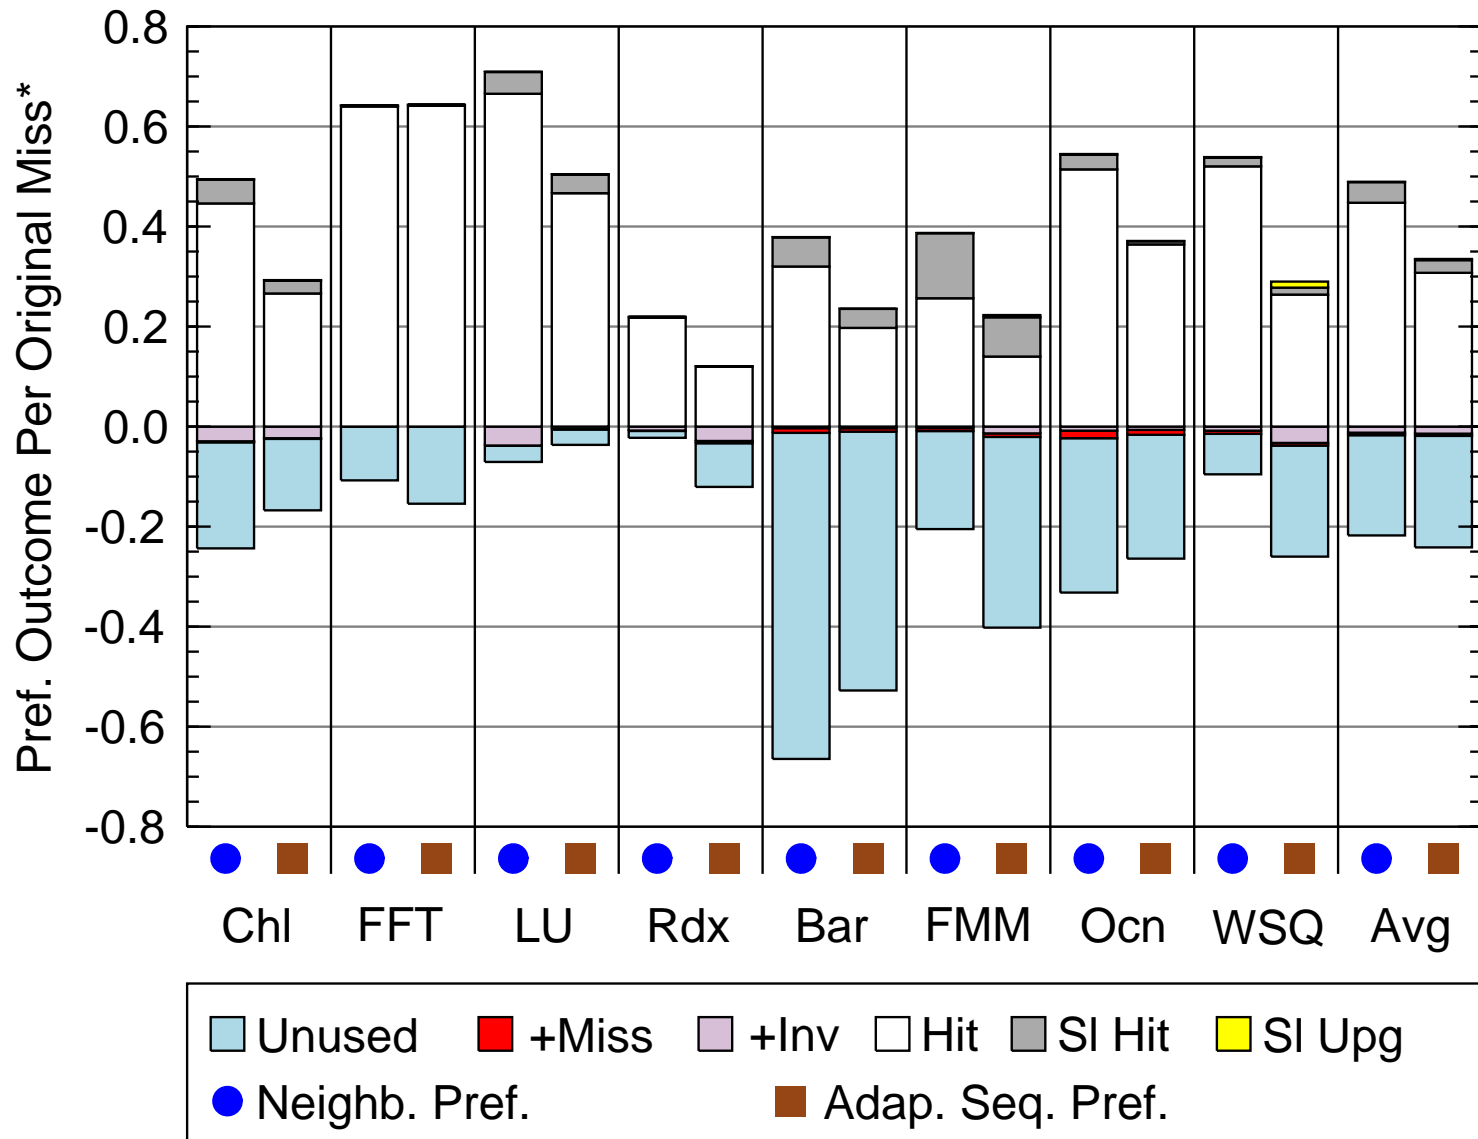
Experiments Presented

Base: 64KB: 8 way assoc.  $\times$  64 byte lines  $\times$  128 sets.


Cache Size: 8KB ( $2^4$  sets) to 1MB ( $2^{11}$  sets).

Line Size: 8 bytes to 1024 bytes.

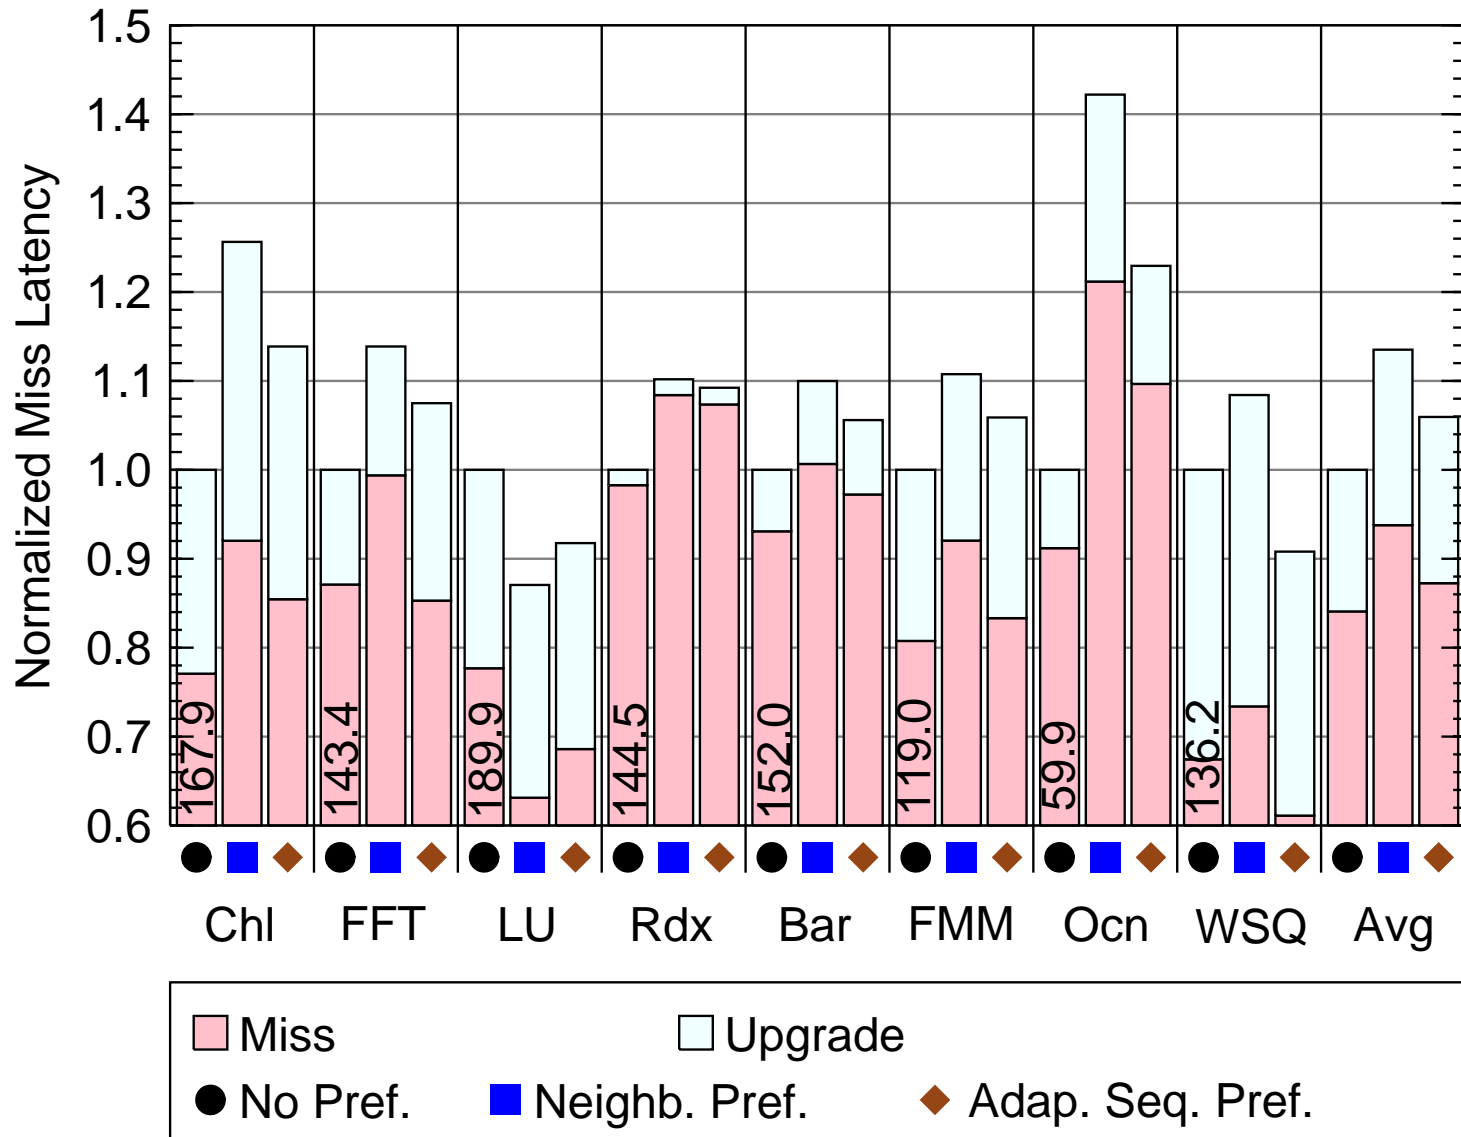
### Base: Prefetch Outcome



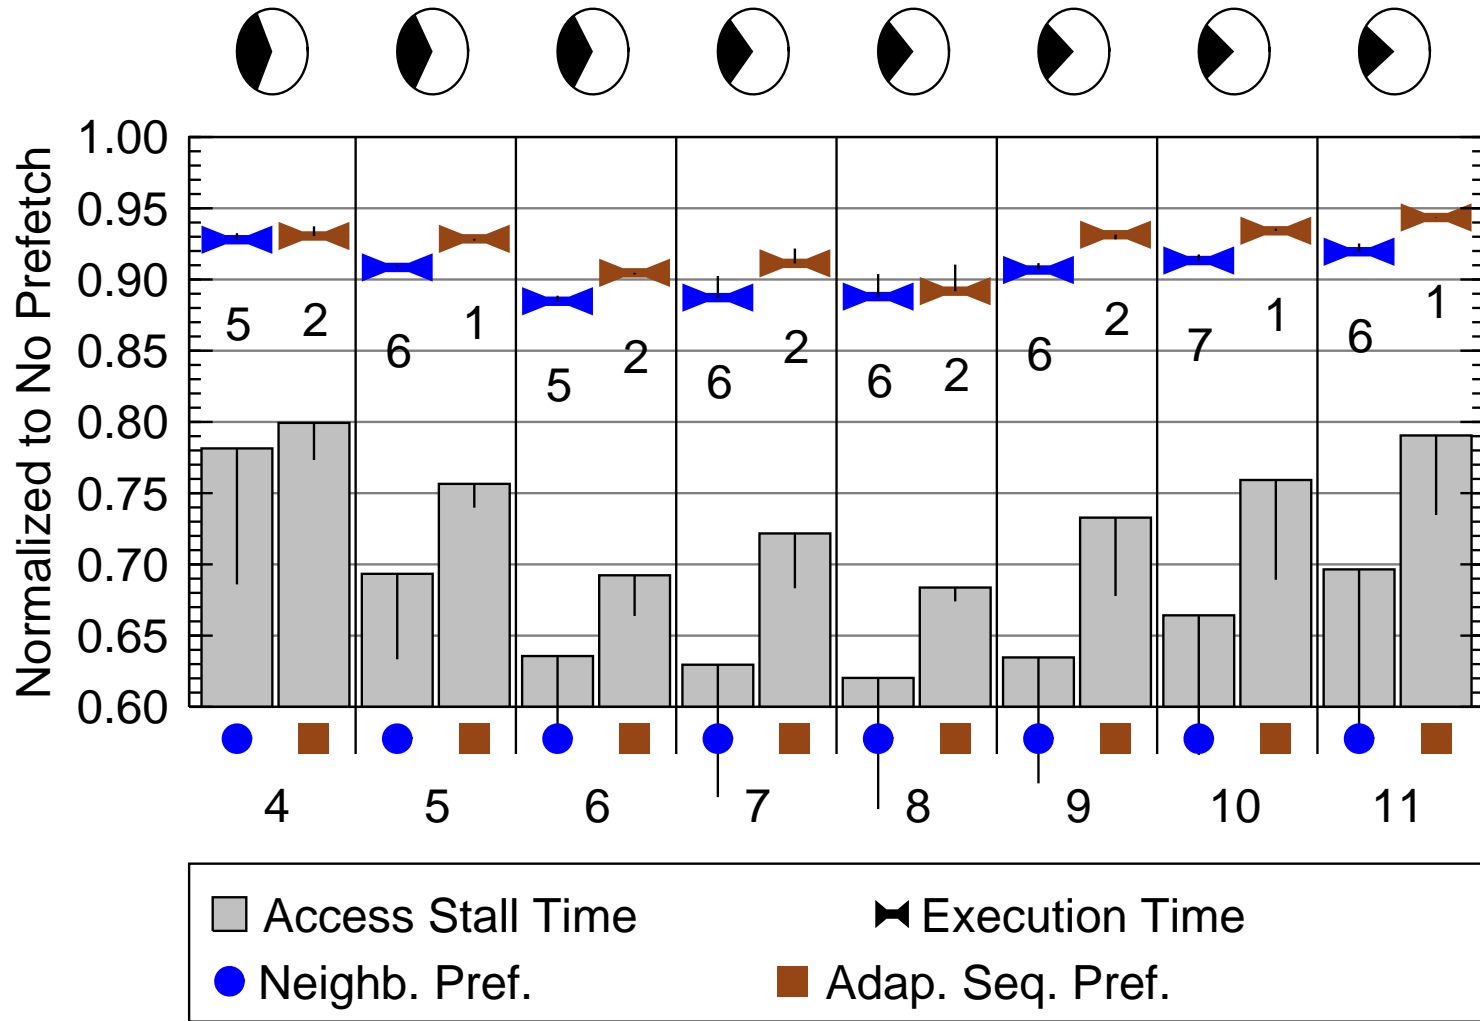
## Base: Normalized Stall and Execution Time



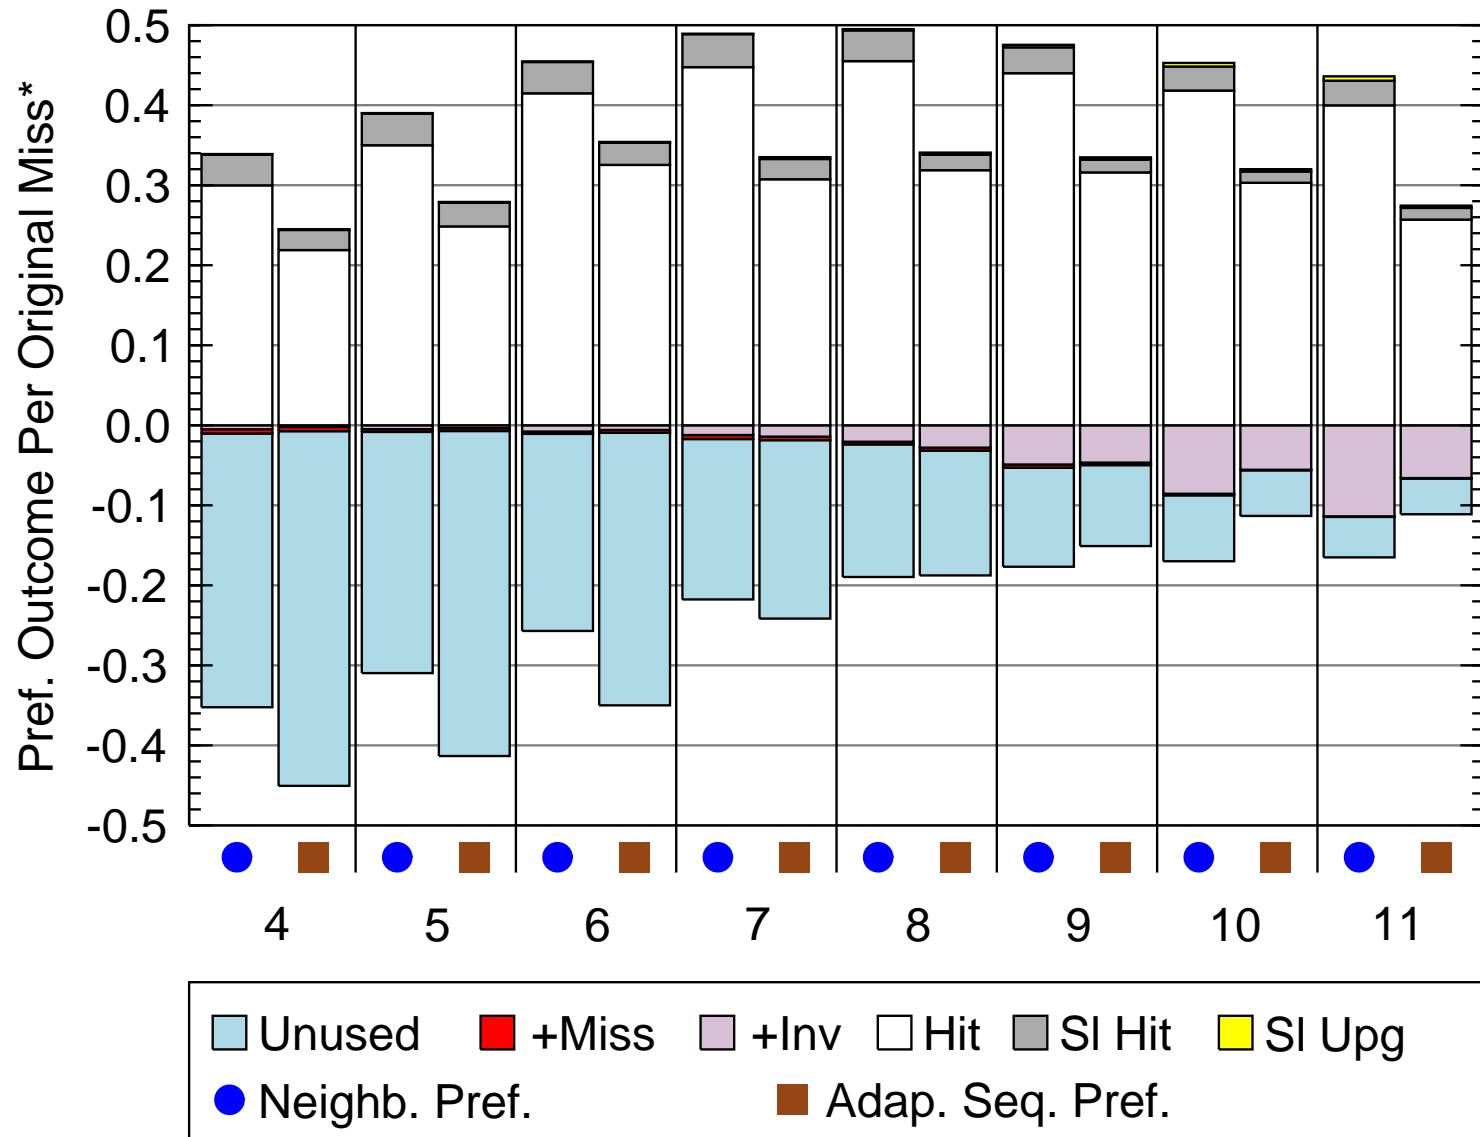
### Base: Miss Latency



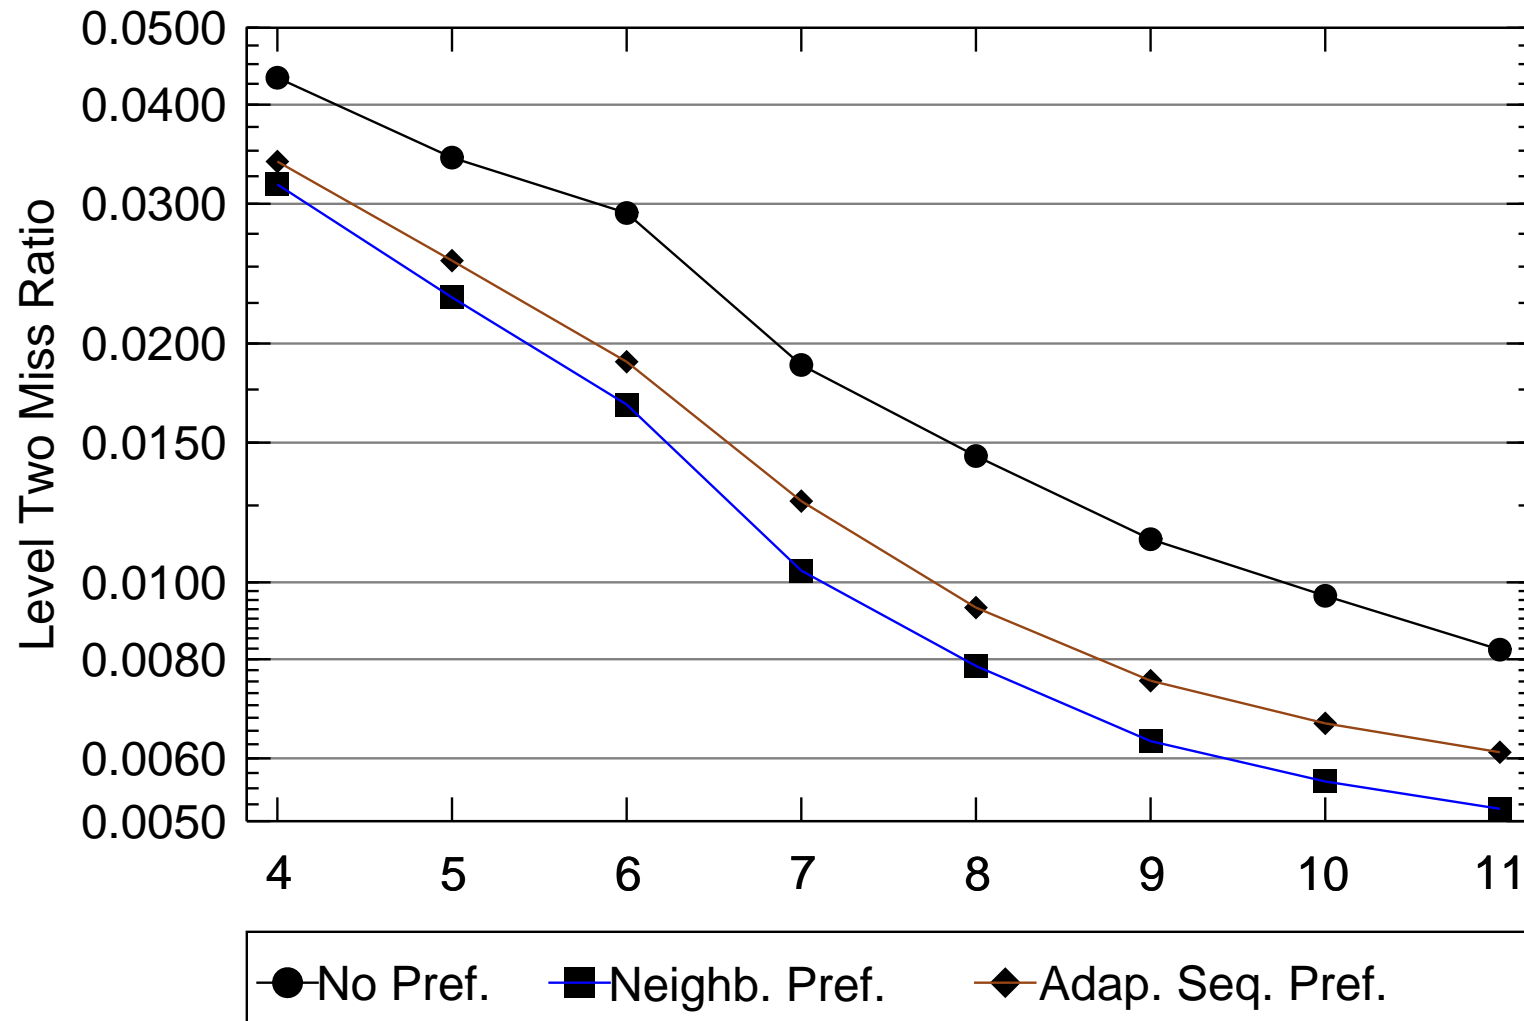
# Cache Size: Normalized Stall and Execution Time



## Cache Size: Prefetch Outcome




## Cache Size: Miss Ratio

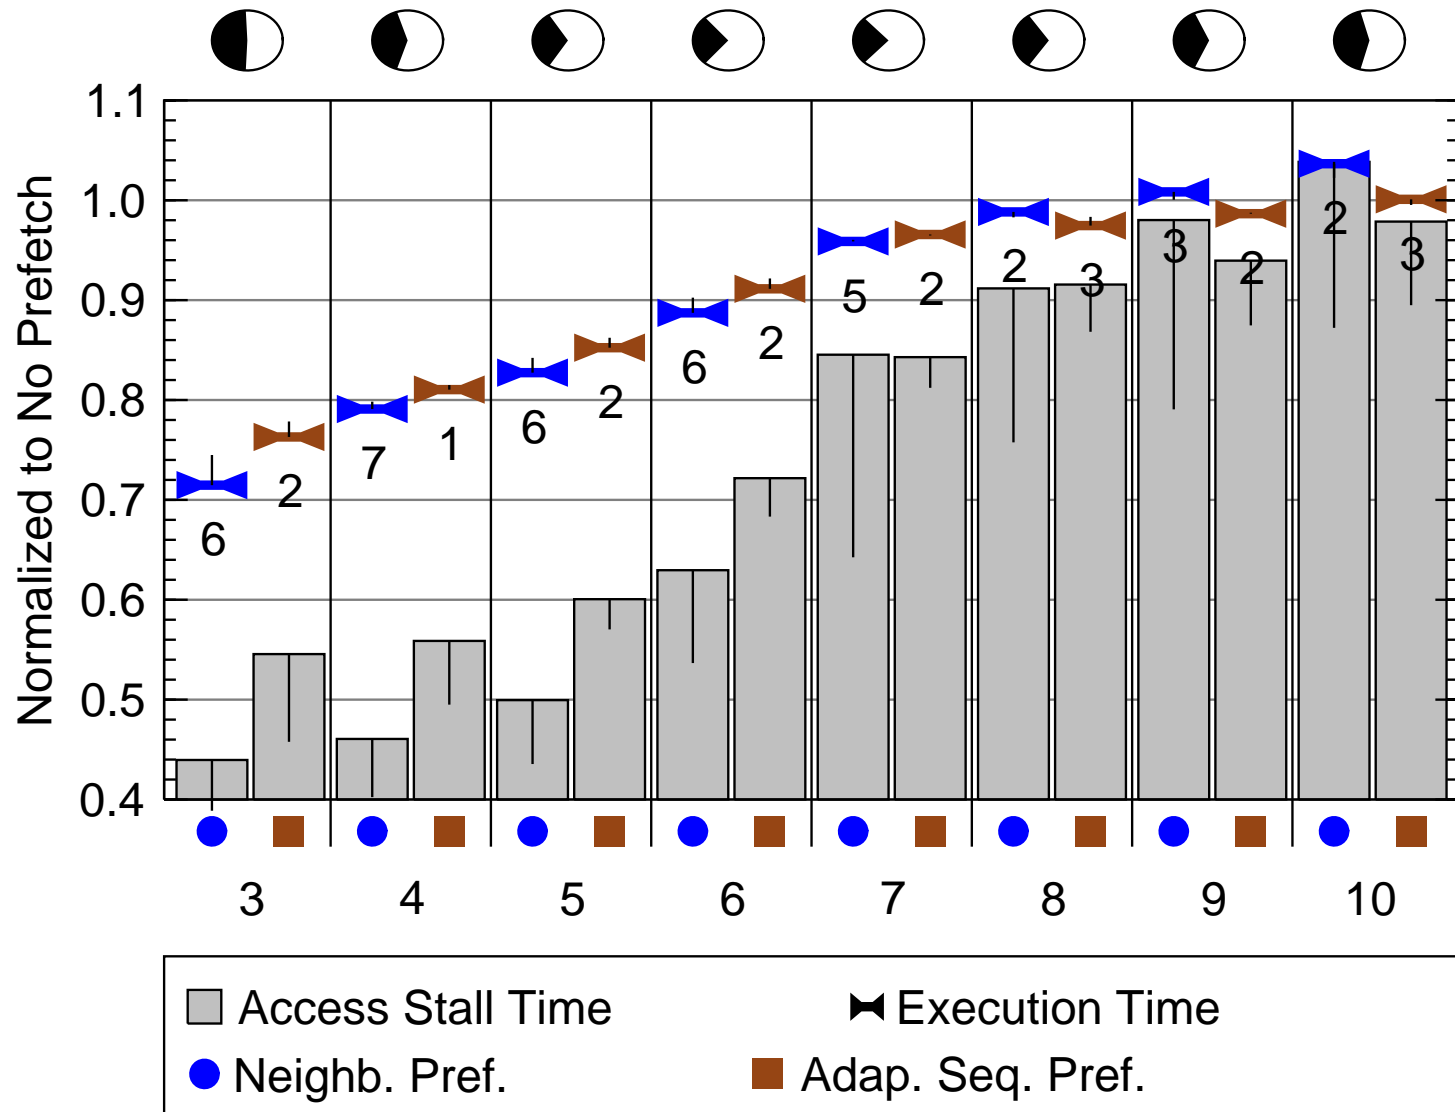





### Line Size: Execution Time



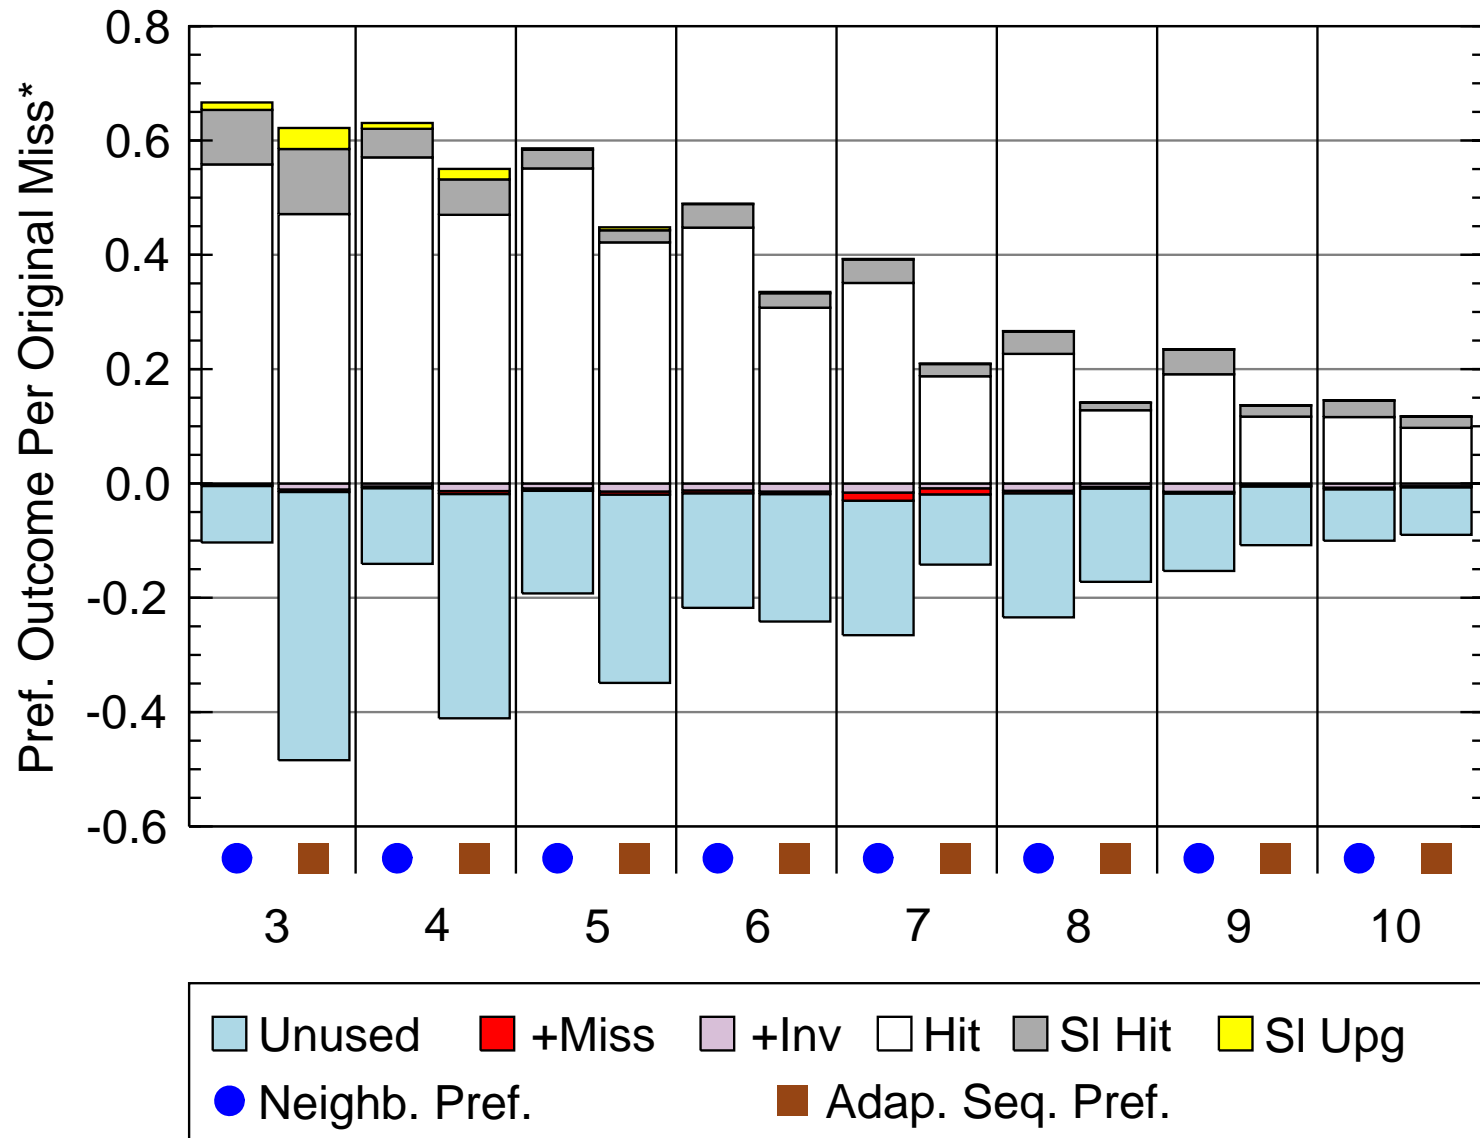
# Line Size: Normalized Stall and Execution Time




# Line Size: Miss Ratio



## Line Size: Prefetch Outcome





## Conclusions

Neighborhood Prefetch handles wider variety of reference patterns.

Less prefetching of unneeded lines.

Reduces miss ratio by nearly 50%.

Reduces execution time by 10% or more.

Performance better than adaptive sequential prefetch . . .  
 . . . though implementation considerably more complex.

## Future Work: Prefetch on Hit

Wallpaper prefetch? (Repeat neighborhoods.)