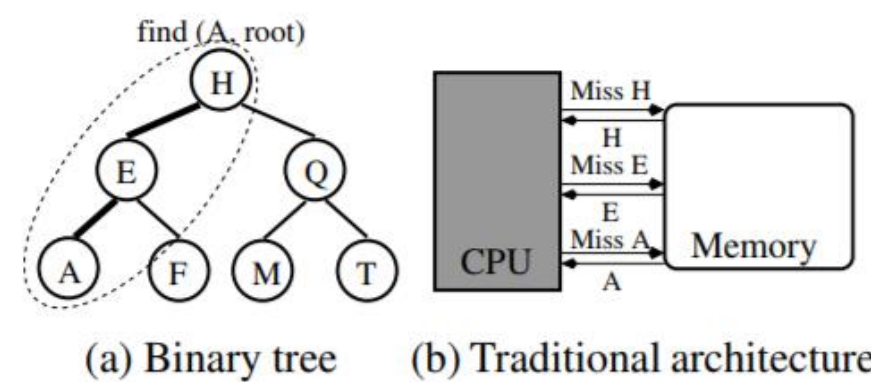


## Problem Statement

Accessing pointer-based linked data structures (LDS) such as linked lists, hash tables, B-trees exhibit highly irregular memory accesses.

Problem with pointer-chasing programs:

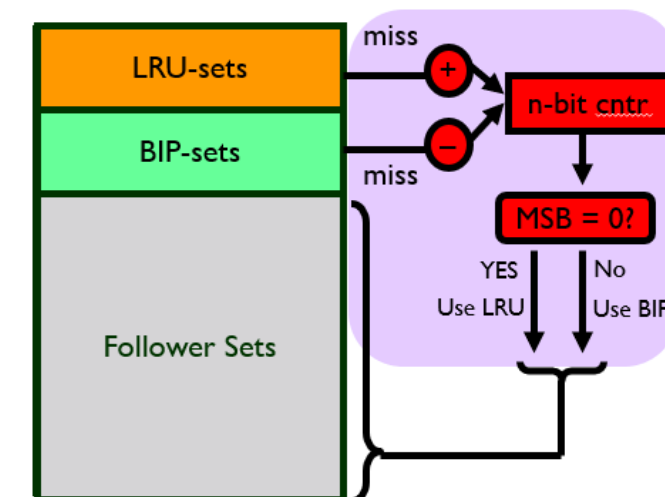
- Poor locality of data
- More cache misses
- Memory latency
- Pre-fetching is not helpful
- Limits parallelism



## Existing Policies

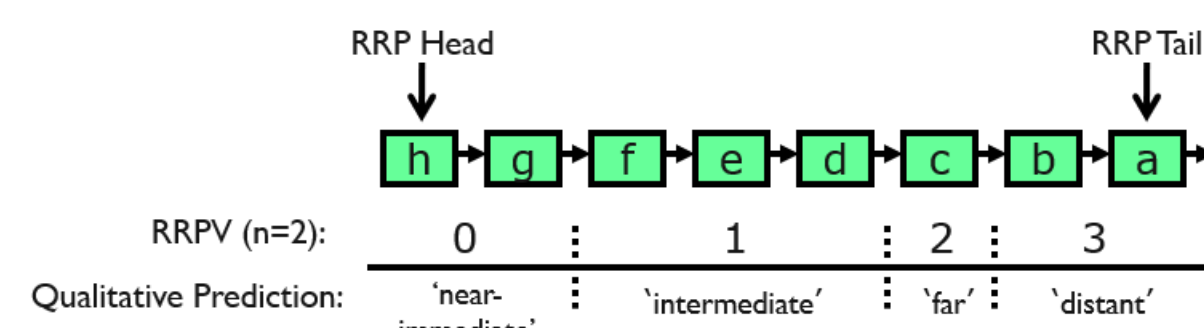
### DIP:

- BIP: Insertion at LRU position and MRU with low probability (1/32)
- Set-Dueling between LRU and BIP
- Follower sets work based on pSel counter
- Works well for thrashing workloads, but not good for LDS.



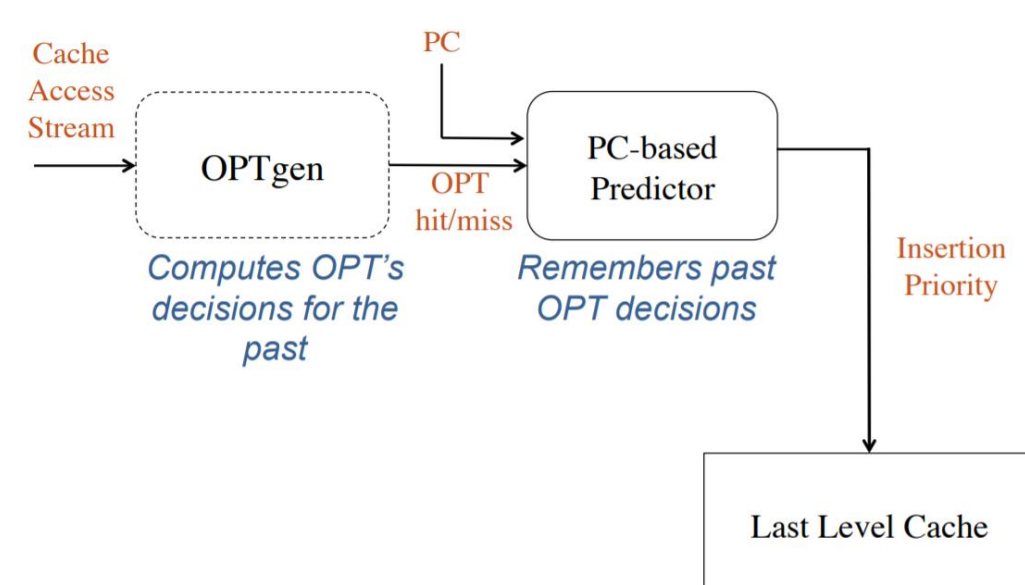
### DRRIP:

- Tries to predict the future reference interval
- Inserts lines with RRPV, higher value means farther in future reference.
- Works well with scan and thrashing workload, but not with LDS.



### Hawkeye:

- Train predictor per-PC based on past OPT behavior
- Predictor decides between cache friendly or cache averse data, and accordingly applies RRPV.



## Proposed Policy

- Pointer Chasing workloads are in iterative or recursive forms.
- Typically small set of instructions with non-uniform memory access pattern.

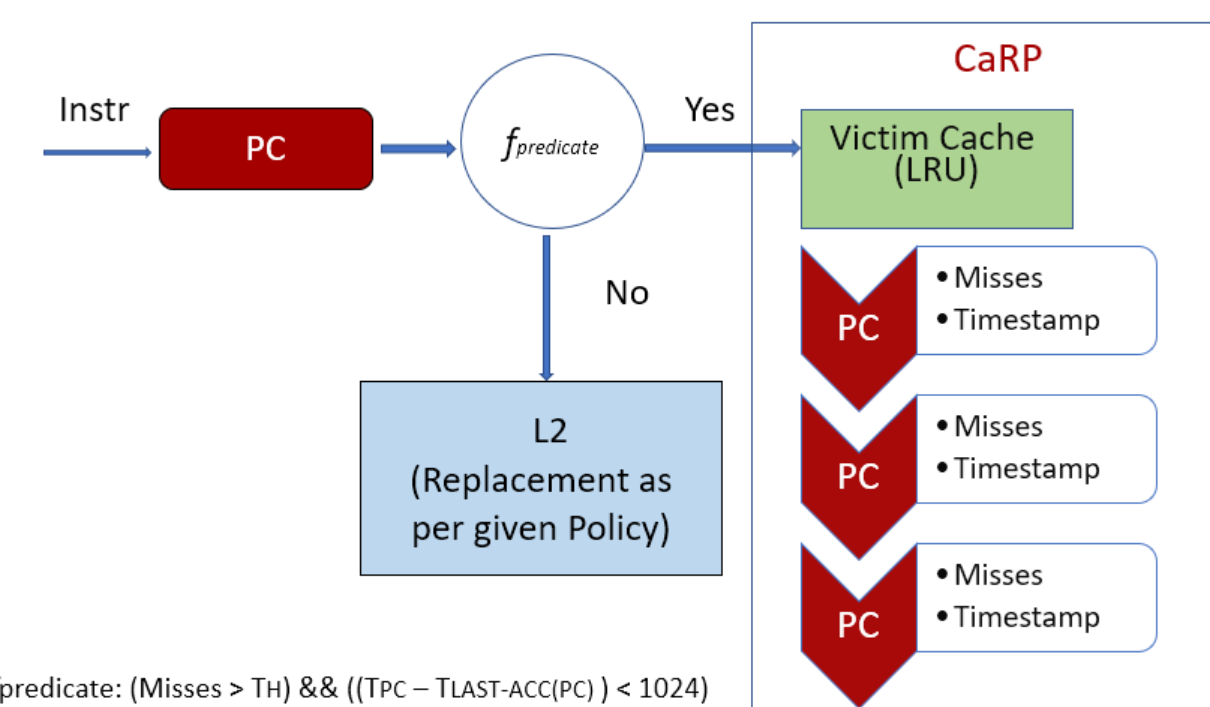
```

/* Given a binary tree, print its nodes in inorder*/
void printInorder(struct node* node)
{
    if (node == NULL)
        return;
    /* first recur on left child */
    printInorder(node->left);
    /* then print the data of node */
    int val = node->data;
    /* now recur on right child */
}

int main()
{
    struct node* root = (struct node*) malloc(sizeof(struct node));
    root->data = 1;
    root->left = (struct node*) malloc(sizeof(struct node));
    root->right = (struct node*) malloc(sizeof(struct node));
    root->left->data = 2;
    root->right->data = 3;
    printInorder(root);
    return 0;
}

```

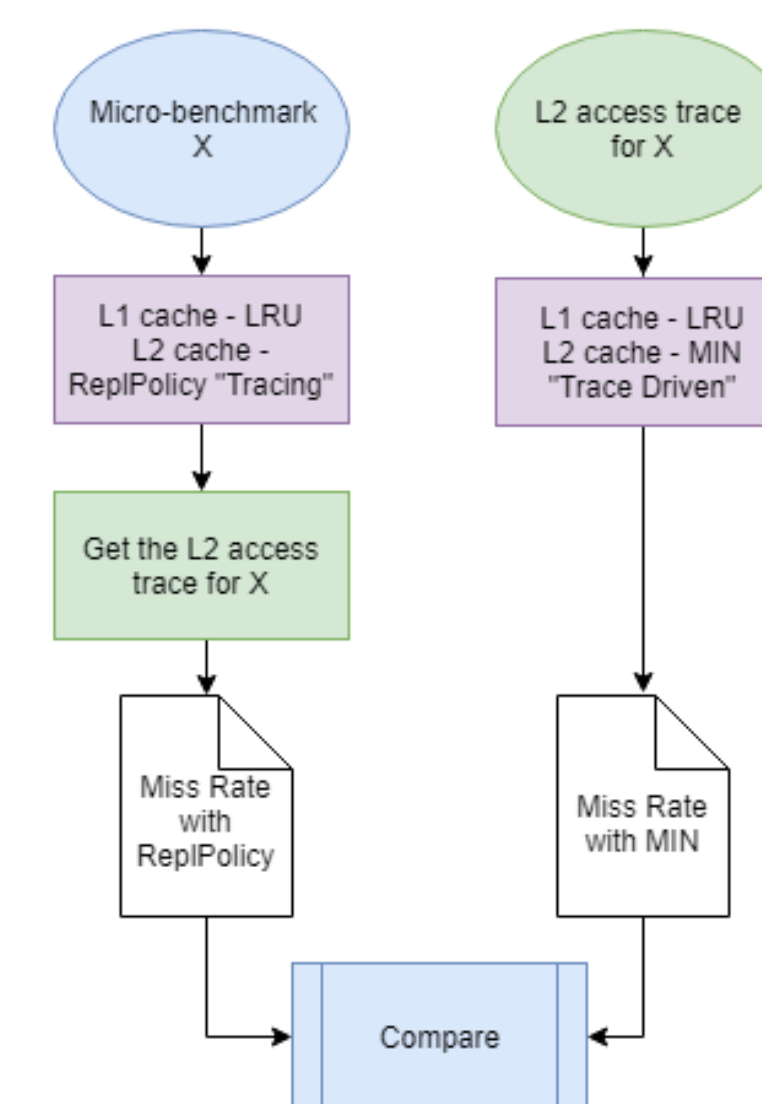
- Track last 1024 PCs for consecutive misses on that PC.
- Apply heuristic to select the Victim cache for PC if misses are above  $Th$ .
- 64 lines Victim Cache and  $Th = 4$ .



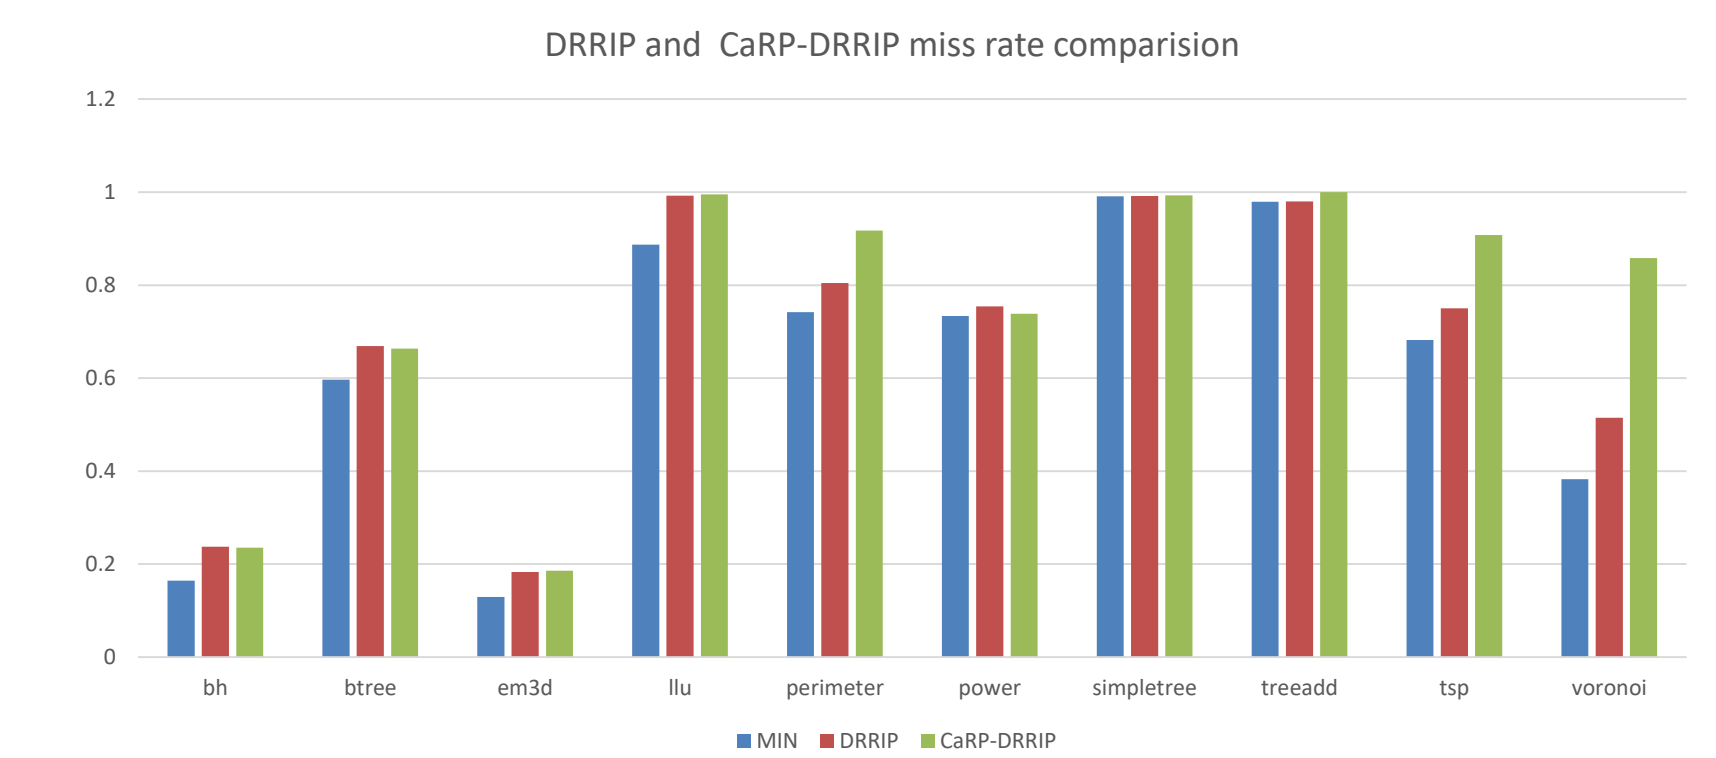
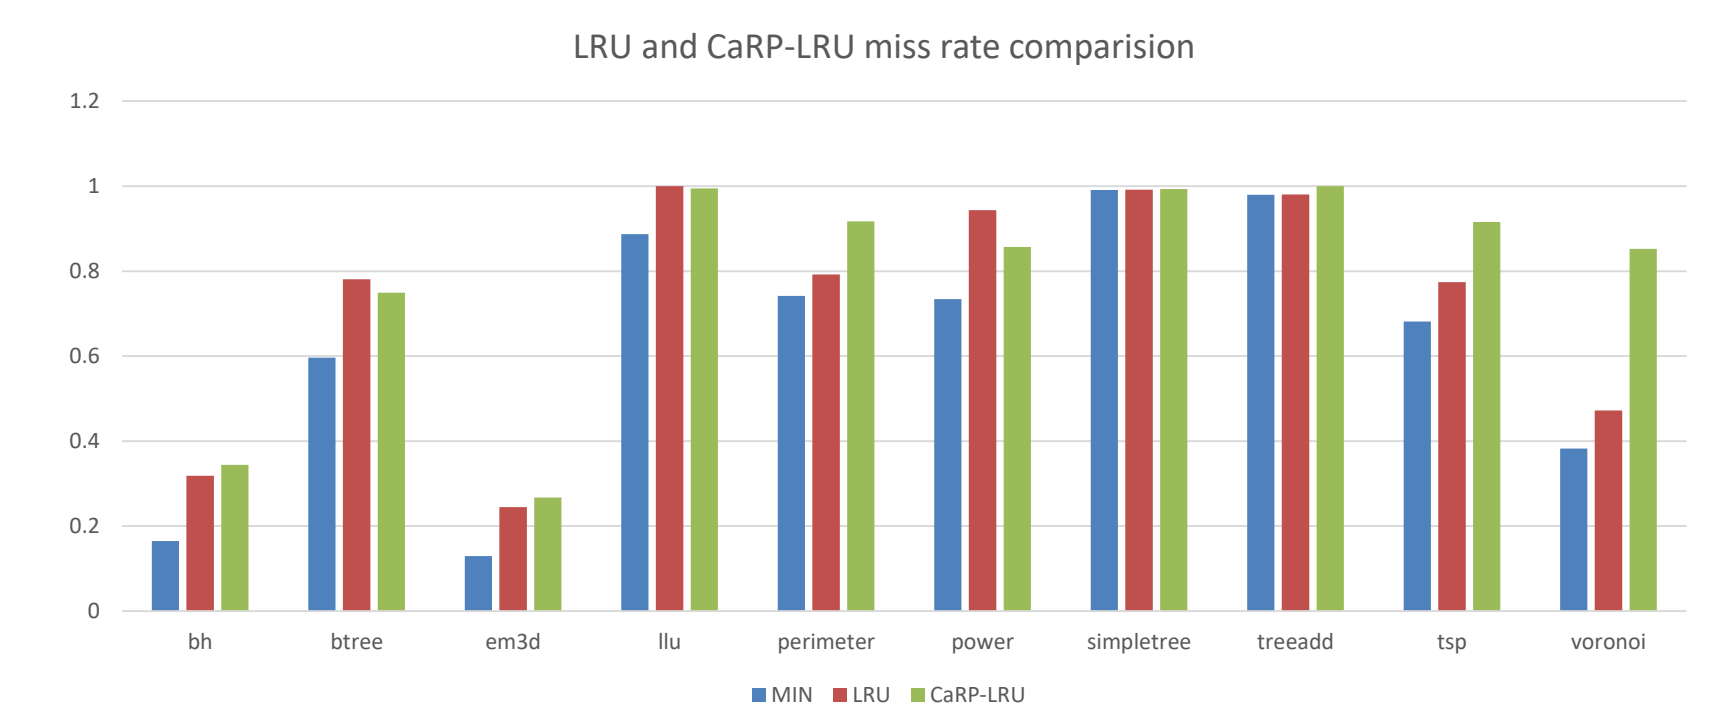
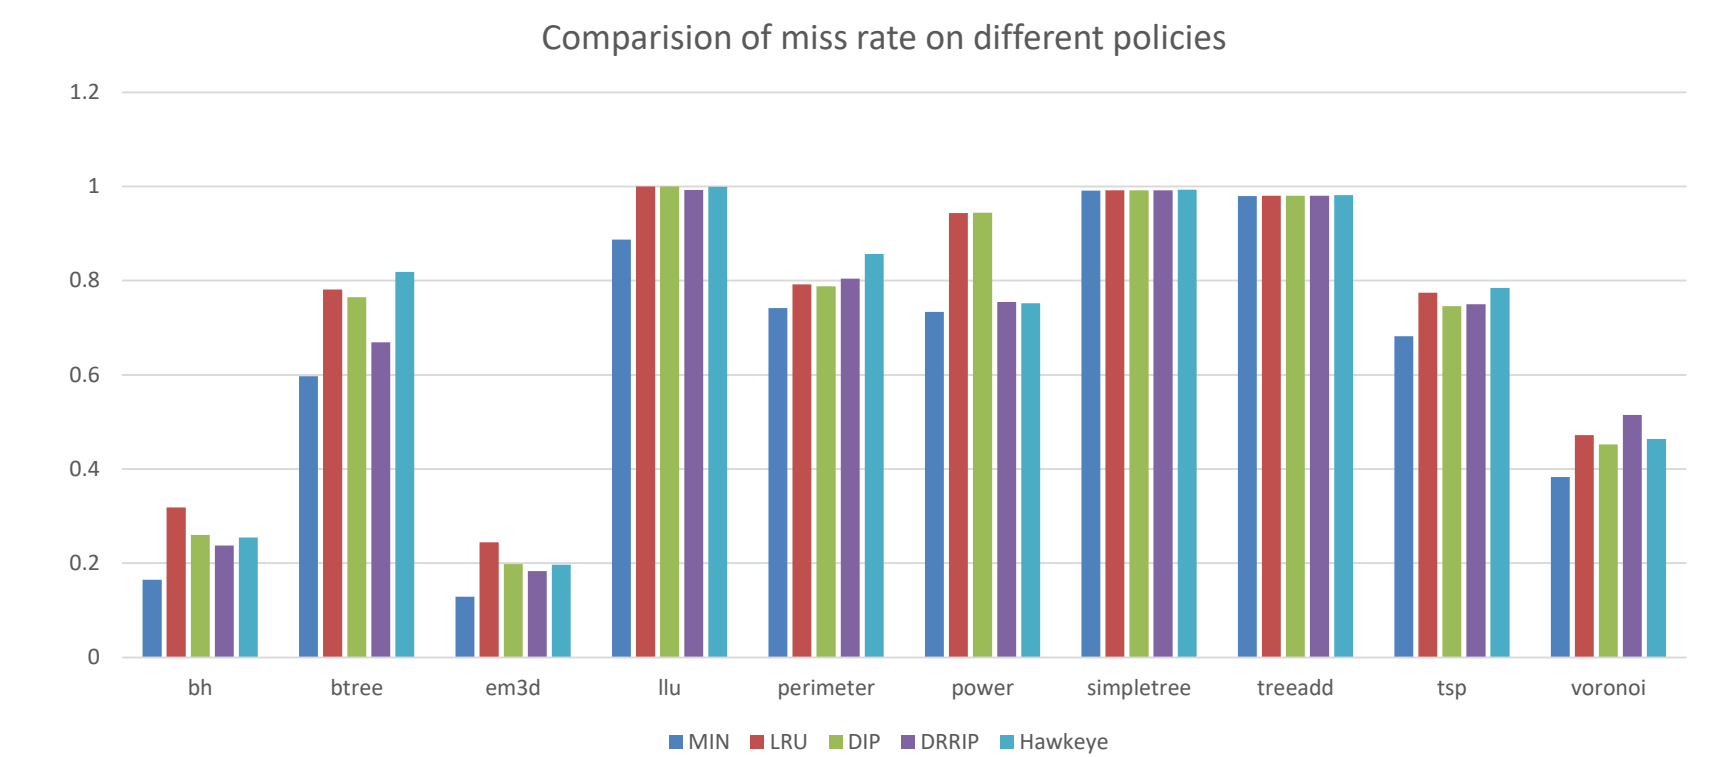
## Experimental Evaluation

- Simulator : Zsim
- Machine : Ubuntu 12.04 with Pintool 2.14
- Memory hierarchy:
  - L1 - 32 kB (4-way, 5 cycle latency)
  - L2 - 256 kB (8-way, 12 cycle latency)
  - L3 - 2 MB (16-way, 35 cycle latency)
- Max simulation time = 720 sec

Benchmark	Data Organisation	Parameters
BH	Heterogenous OcTree	100000 bodies, 32 nodes
Em3D	Single Linked Lists	2000 H 500 E Nodes
Perimeter	Quad-Tree	11 levels
Power	N-way Tree, single-linkd Lists	N/A
TreeAdd	Binary-Tree	25 levels
Tsp	Balanced binary-tree	2 Million Nodes
Voronoi	Balanced binary-tree	1 Million Nodes
Simple Tree	Binary Tree	100 levels
B-Tree	B+ Tree	3 Million Nodes



## Results



## Conclusions

- Given workloads have little scope of improvement over MIN. (Max 8-10 %)
- CaRP with DRRIP achieves behavior equivalent to MIN on *power* benchmark.
- CaRP with LRU performs better than LRU on *power*, *llu*, and *btree*.

## Future Work

- Select optimal Victim Cache size, hardware overhead
- Minimize the overhead of keeping PC history, buffering
- Try to incorporate memory accesses along with PC history to reduce the false positives