

Cache Misses Prediction Using Stack Distances

Calin Cascaval and David A. Padua

Department of computer science

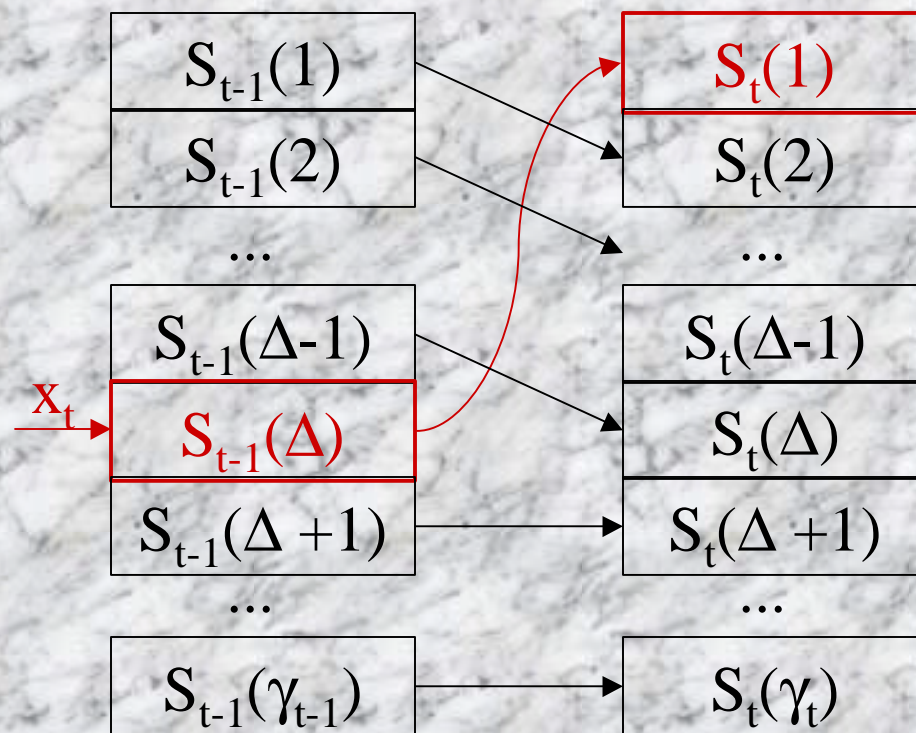
University of Illinois at Urbana-Champaign

cascaval,padua@cs.uiuc.edu

Outline

- Stack Distances as a Metric for Locality
- Run-time Instrumentation
- Compile-time Analysis
- Future Work

Stack Algorithms

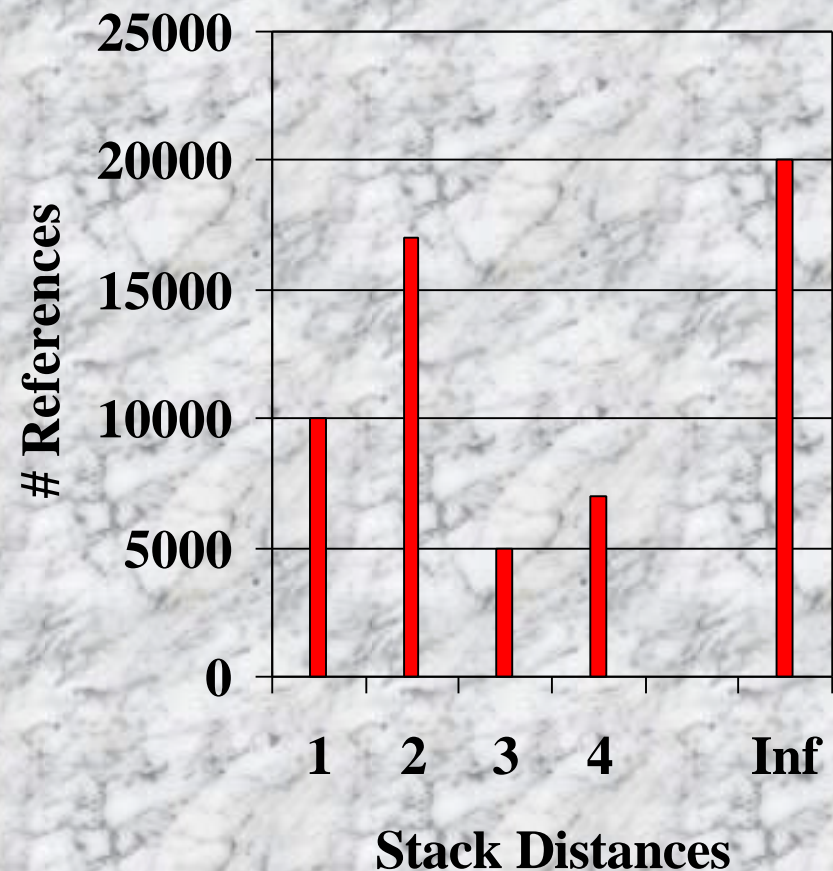


Stack Distances As Cache Misses

- compute the number of cache hits and misses as follows:

$$\text{hits}(C) = \sum_{\Delta=1}^C s(\Delta)$$

$$\text{misses}(C) = \sum_{\Delta=C+1}^{\text{Inf}} s(\Delta)$$




Metric Validation

- Stack algorithm implemented as a library
- Polaris instrumented codes from the SPEC95 and Perfect Club benchmarks
- measured actual number of cache misses using the hardware counters on the R10K processors (for both L1 and L2 caches)

Experiments Serial

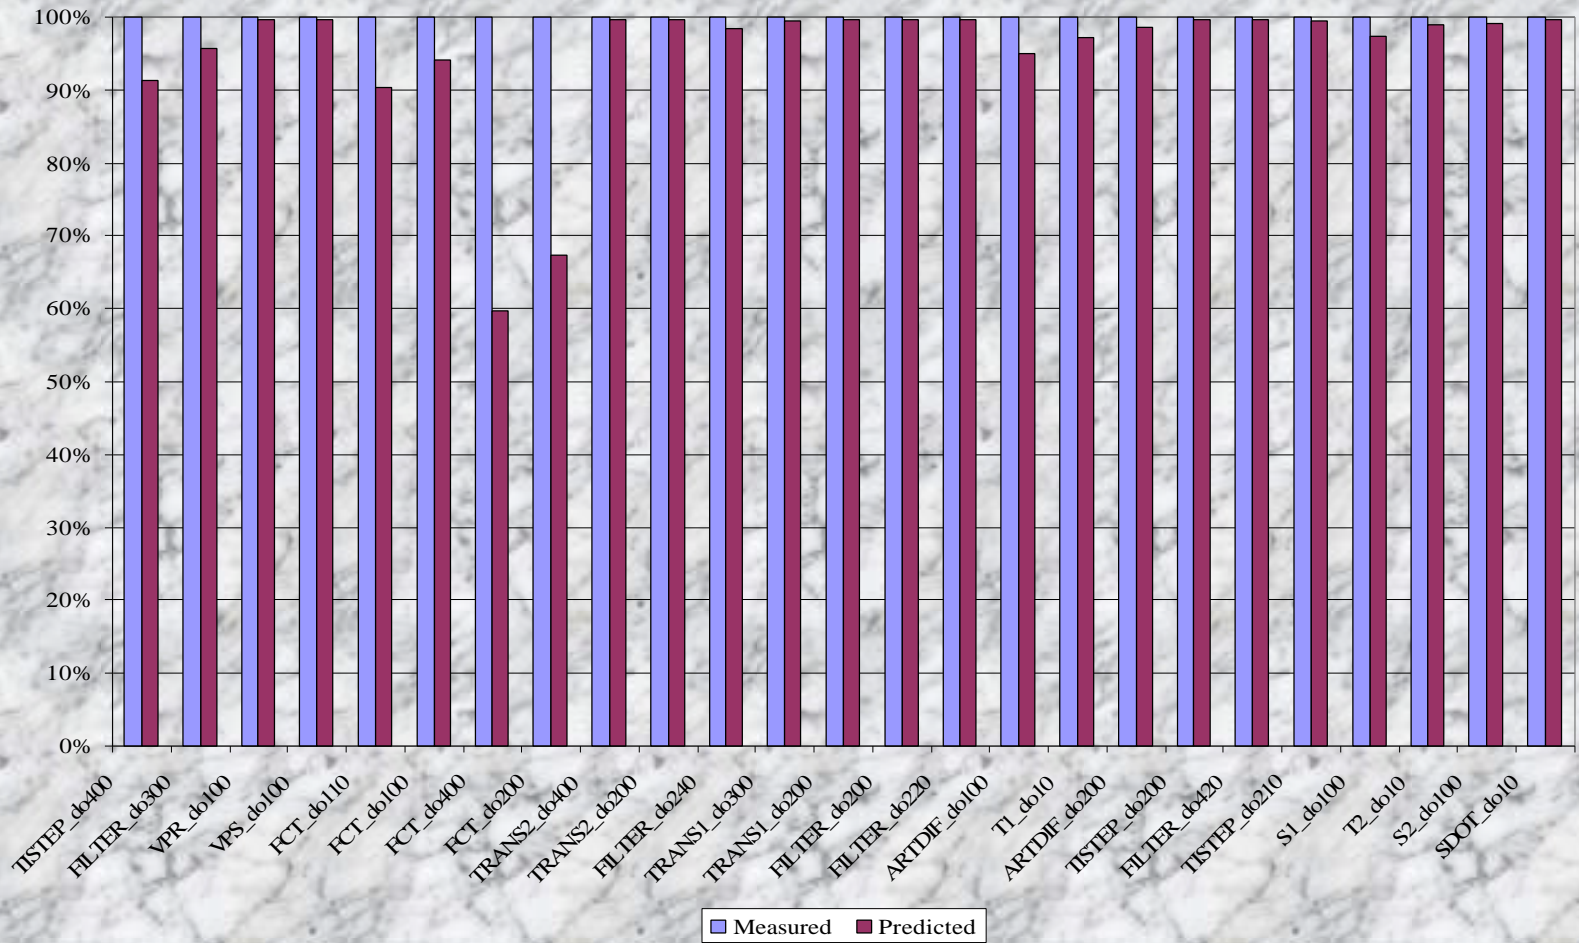
Prediction Accuracy



■ < 75% ■ 75%-90% ■ 90%-100% ■ 100%-110% ■ > 110%

HYDRO2D

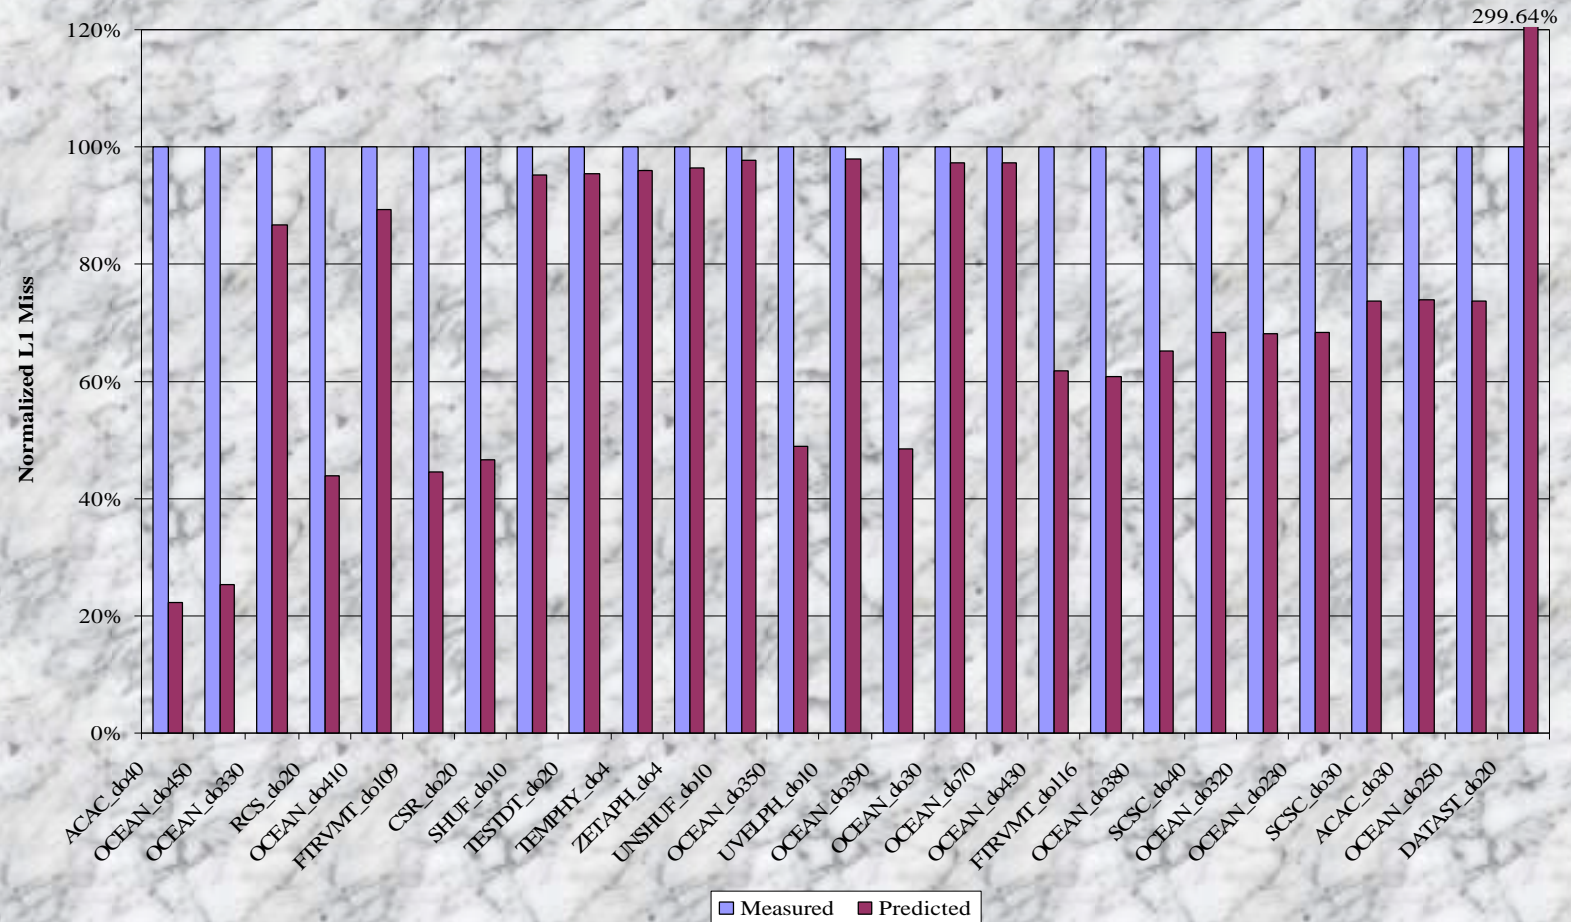
Normalized L1 Misses



August 27, 1998

Univ. of Illinois at Urbana-Champaign

Ocean



Solution Advantages

- Accurate in most cases.
- One pass through the trace estimates misses for any cache size
- Architecture independent (except for the cache line size)
- Same model applicable to independent loops as well as entire programs
- Easily applicable to parallel programs

Run-time Solution Advantages

- Works in all cases

Run-time Solution Disadvantages

- It is run-time, therefore consumes CPU time
- Cannot easily identify references with bad locality
- Needs separate runs for different cache sizes/processors
- Assumes a fully associative cache (to reduce overhead)

Compile-time Solution

- Integrated within the Polaris parallelizing compiler
- Algorithm based on data dependence distance vectors to compute the stack distances for loops
- Computes symbolic expressions (based on loop bounds) for stack distances and number of references

Compile-time Solution (cont.)

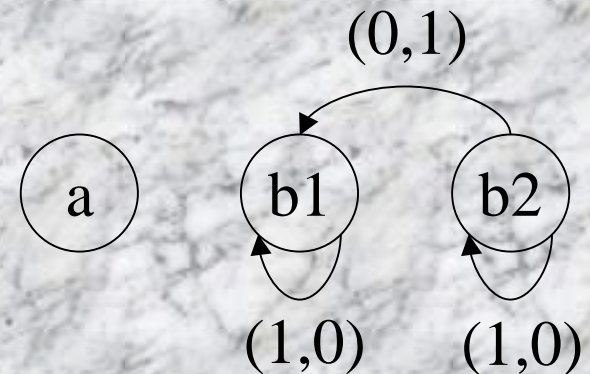
- Preserves the advantages given by the run-time solution, adding
 - misbehaving reference identification
 - array size knowledge may improve accuracy
 - flexibility and locality information readily available in the compiler
- May need a run-time pass for unknown loop bounds or data depending on the input

Example

```
do j = 1, n
  do i = 1, n
    a(i,j) = b(i,1) + b(i+1,1)
```

Distance	References
$ \delta R_i^1 + d - s $	$N-1$
Inf	$3N-(N-1)$

Distance	References
$ \delta R_i^1 + d - s $	$N(N-1)$
$ \delta R_j^1 + d - s $	$(N-1)(N+1)$
Inf	N^2+N+1



Future work

- Integrating the compile-time solution with the run-time instrumentation
- Solve the false sharing integration within the compiler approach
- Address some of the limitations: cache associativity, multi-word cache lines