On the Use of Trace Sampling for Architectural Studies of Desktop Applications

SIGMETRICS 1999
Atlanta, GA

Patrick Crowley & Jean-Loup Baer
Department of Computer Science & Engineering
University of Washington
{pcrowley, baer}@cs.washington.edu

Motivation

- Trace-based simulation is expensive
- Desktop applications are large and can execute billions of instructions
- Often, approximate results or trends are OK

Question: *When can sampling be used to efficiently (i.e., ~10x reduction) drive architectural studies?*

Application Traces

- 5 interactive desktop applications

Program	*Application*	*Description*
acrord32	Adobe Acrobat Reader 3.0	PDF document viewer
netscape	Netscape Navigator 3.1	Web browser
photoshp	Adobe Photoshop 4.0	Image editing package
powerpnt	MS PowerPoint 7.0	Presentation software
winword	MS Word 7.0	Word processor

- Binary instrumentation used to record memory references*
- Full traces total ~2GB in size
- sampled traces ~200MB

(*OS kernel activity not included)

Trace Sampling

- "sample" the trace

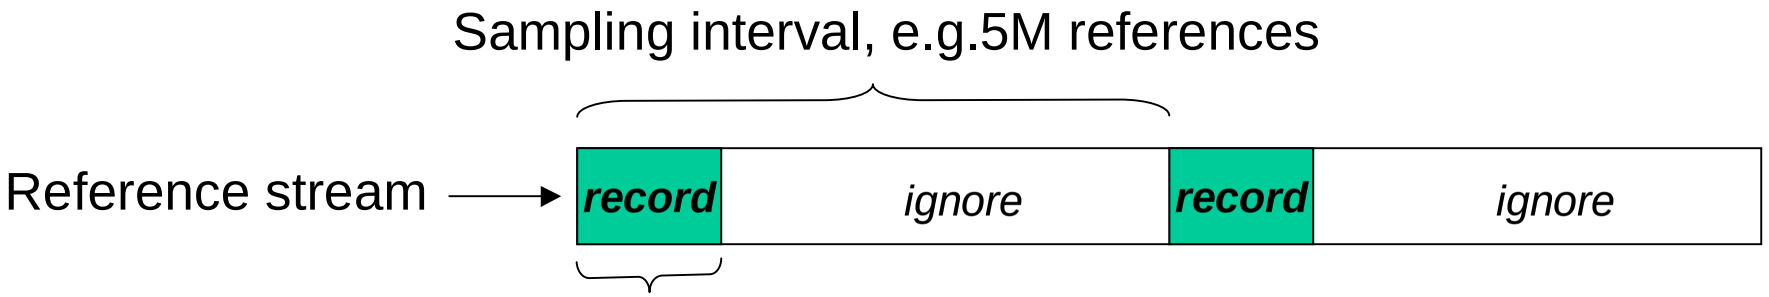

Sampling ratio = Sample size/Sampling interval =0.1

- Two types of errors:
  - Sampling Errors
    - samples are not always representative of population
    - Result of parameter choice (i.e., sample size & sampling ratio)
  - Errors due to **unknown references**
    - cache state is unknown at start of sample
    - is the first reference to a line a hit or a miss?

Trace Sampling Techniques

- Techniques for mitigating error due to unknown references

Technique	*Description*
cold	Assumes that each unknown reference misses (i.e., an effectively empty cache).
half	Uses the first half of each sample to "prime" the cache
stitch	uses the end state of the previous sample as the start state for the current sample
INITMR	estimates the miss ratio of unknown references [Wood et al. 91, Kessler et al. 94]
true-sample	Starts each sample with the correct cache state

Miss Ratio Determination

Miss Rates for acrord32
(direct-mapped, instruction cache)

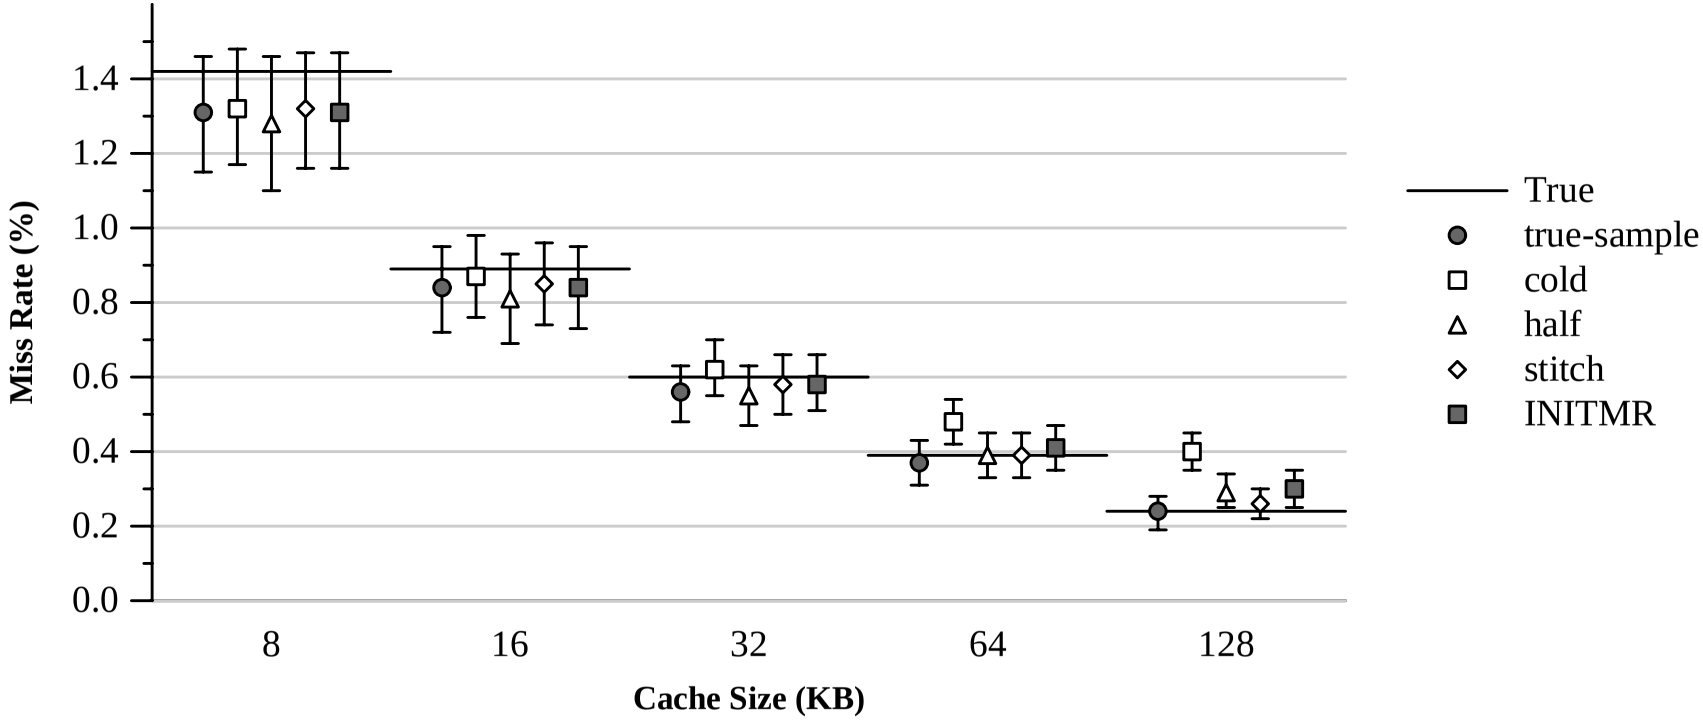

– All techniques work up to 32KB (stitch is best)

– stitch and INITMR are reliable up to 64KB

Architectural Study: Victim Cache Trends

stitch results for acrord32
(direct-mapped, data cache)

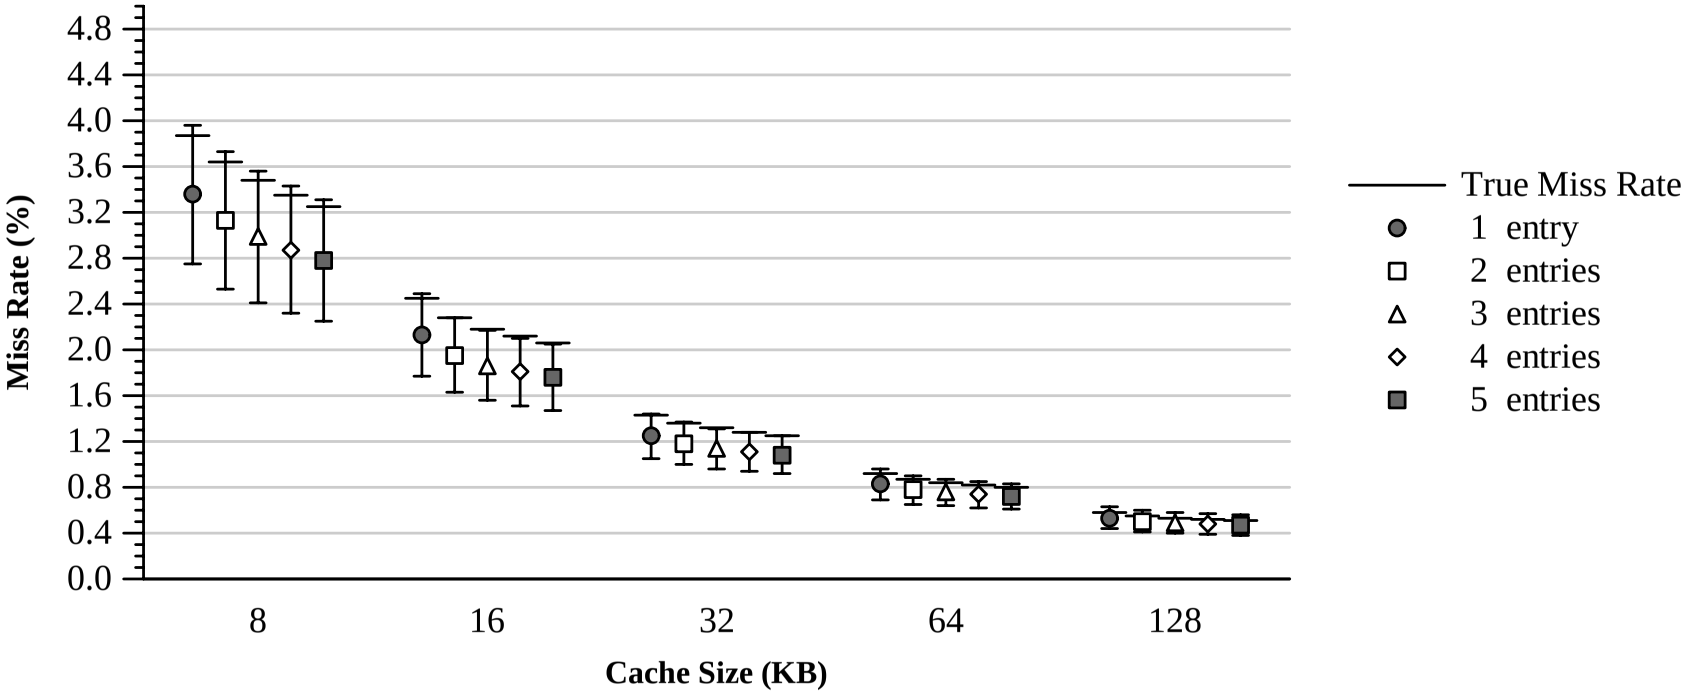

- Results from stitch are good
- Trends clearly demonstrated

See the paper for ...

- More sampling techniques
- Trace sampling for branch prediction
- Trace sampling for analytical cache models
- Large set of results

Conclusions

- **stitch** & **INITMR** are reliable for caches up to 64KB in size
- Sampling *can* be used effectively to drive architectural studies