

Online Phase-Adaptive Data Layout Selection

Chengliang Zhang

Microsoft
(former IBM intern)


Martin Hirzel

IBM



ECOOP, 10 July 2008

Problem Statement


No
training
run.




Online Phase-Adaptive Data Layout Selection



Data Layouts from Copying Garbage Collection




(a) Depth-first




(b) Breadth-first


BF



(c) Preserving




(d) Sorting




(e) Size

HI

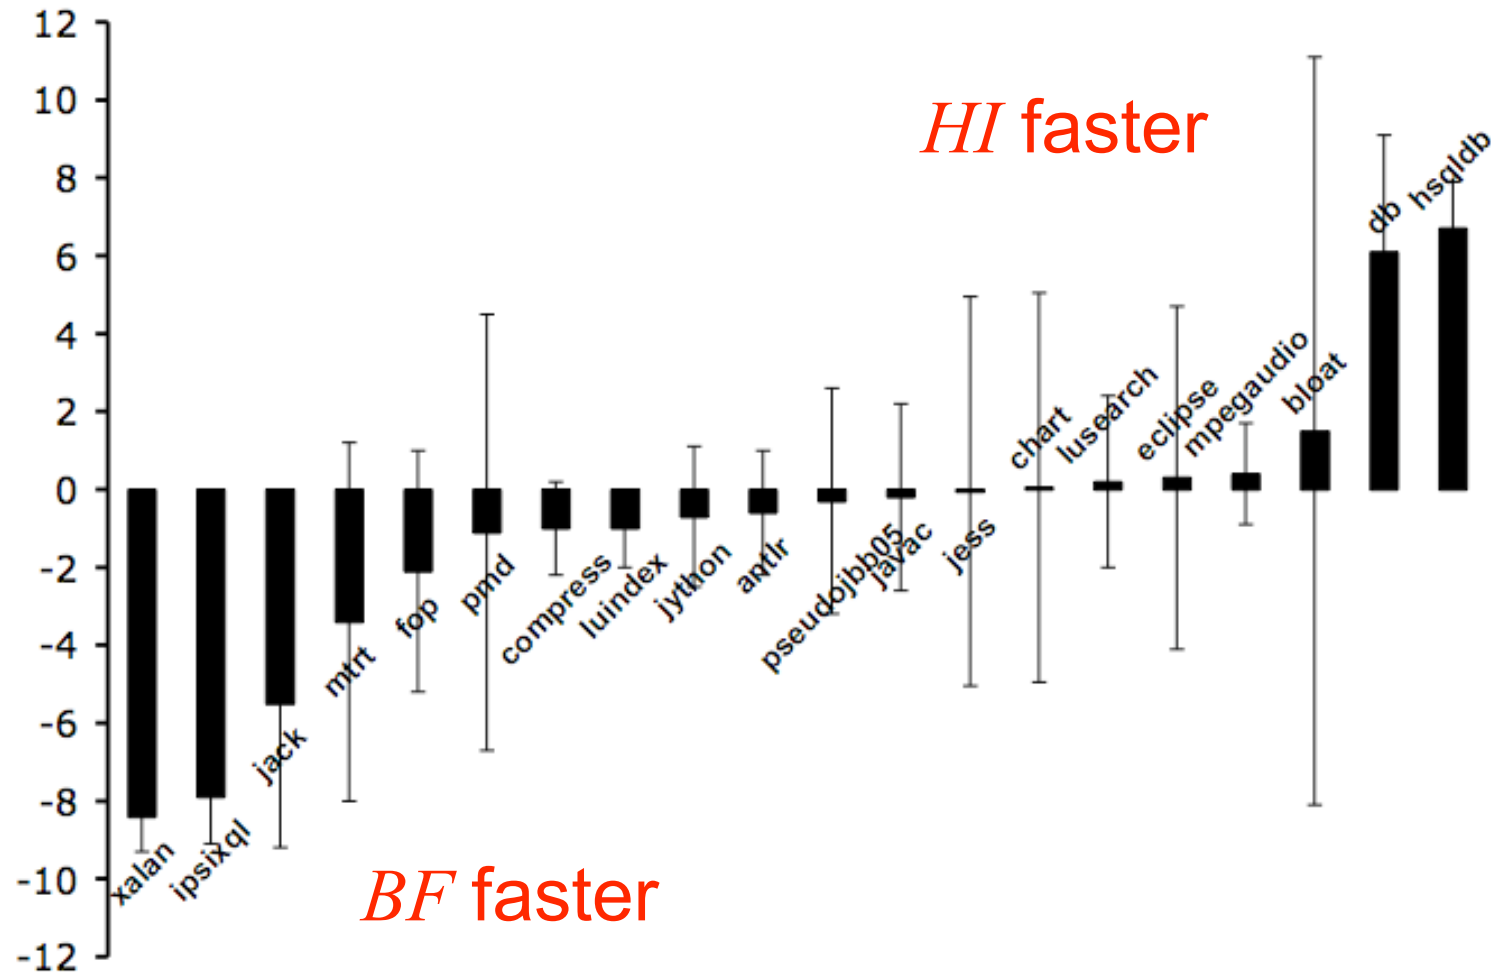


(f) Hierarchical



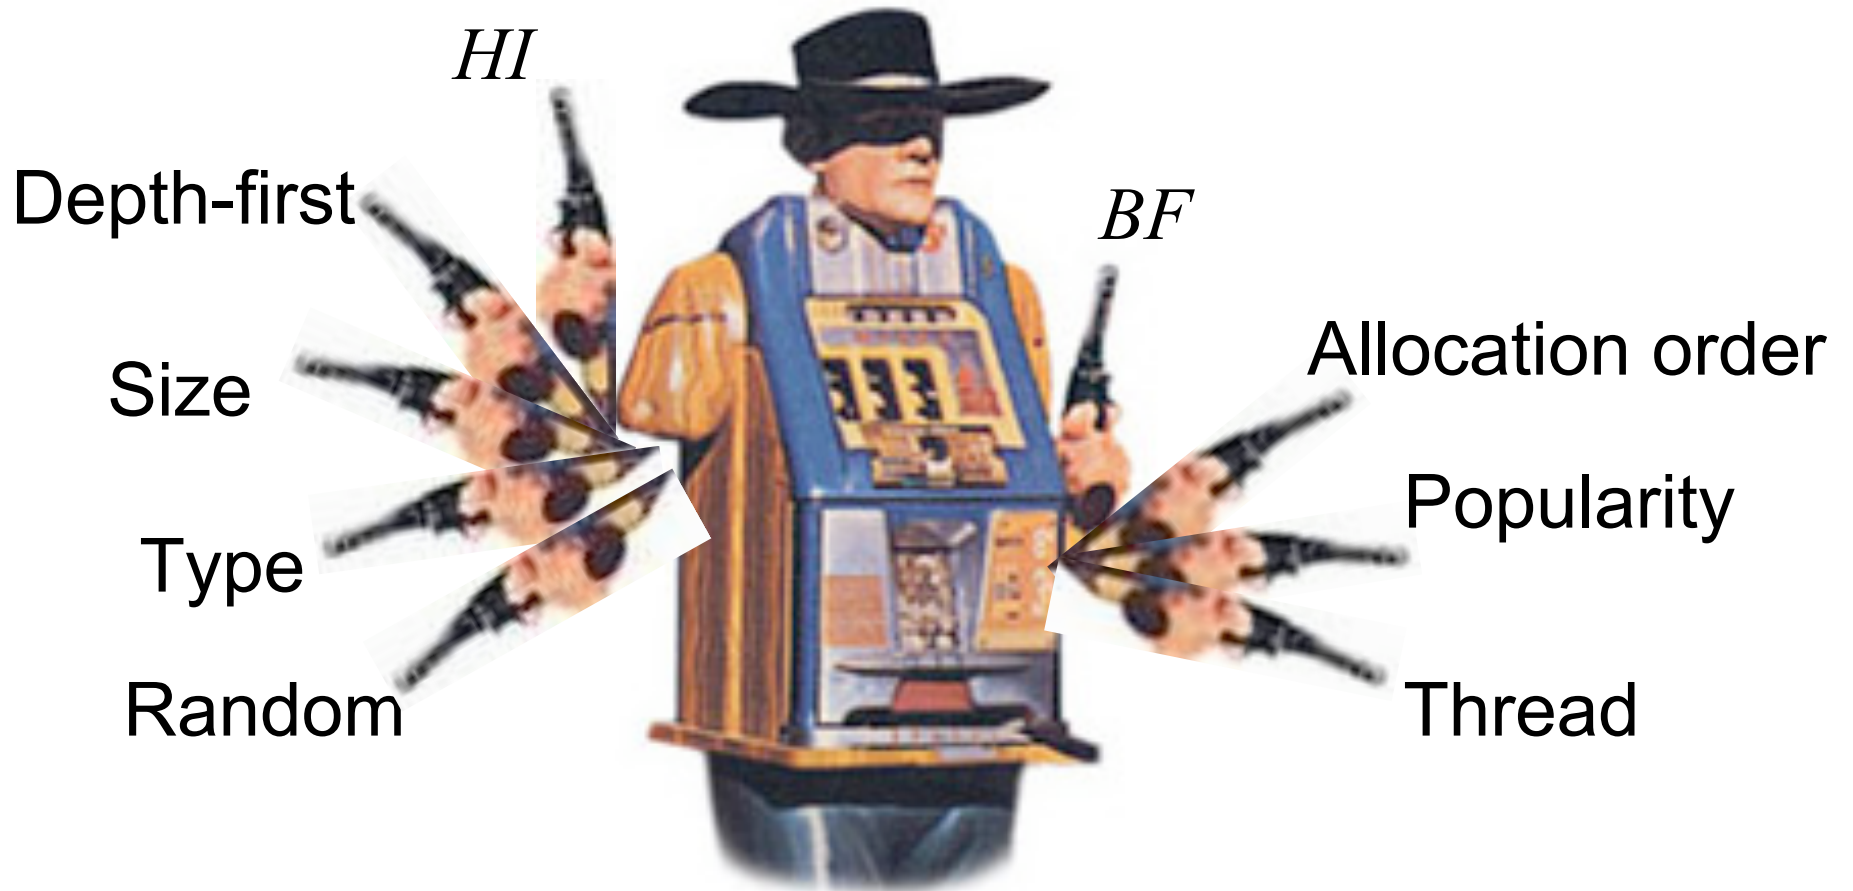
(g) Thread

Layout Performance Comparison

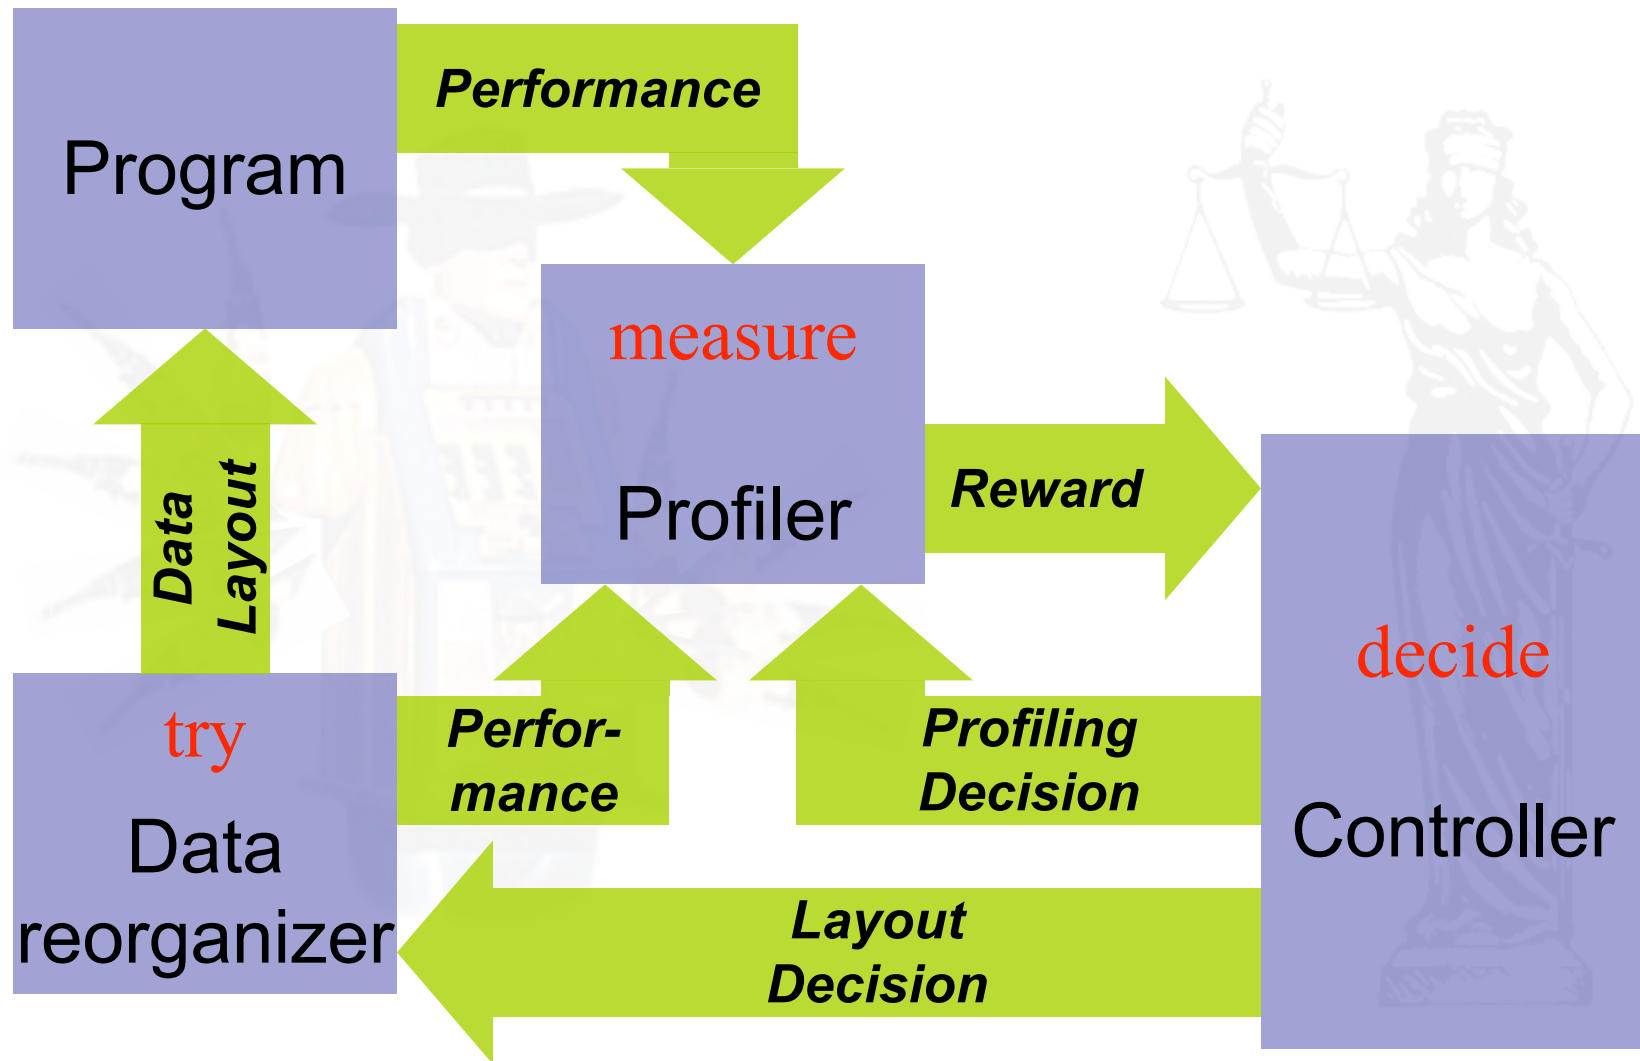


8-processor AMD

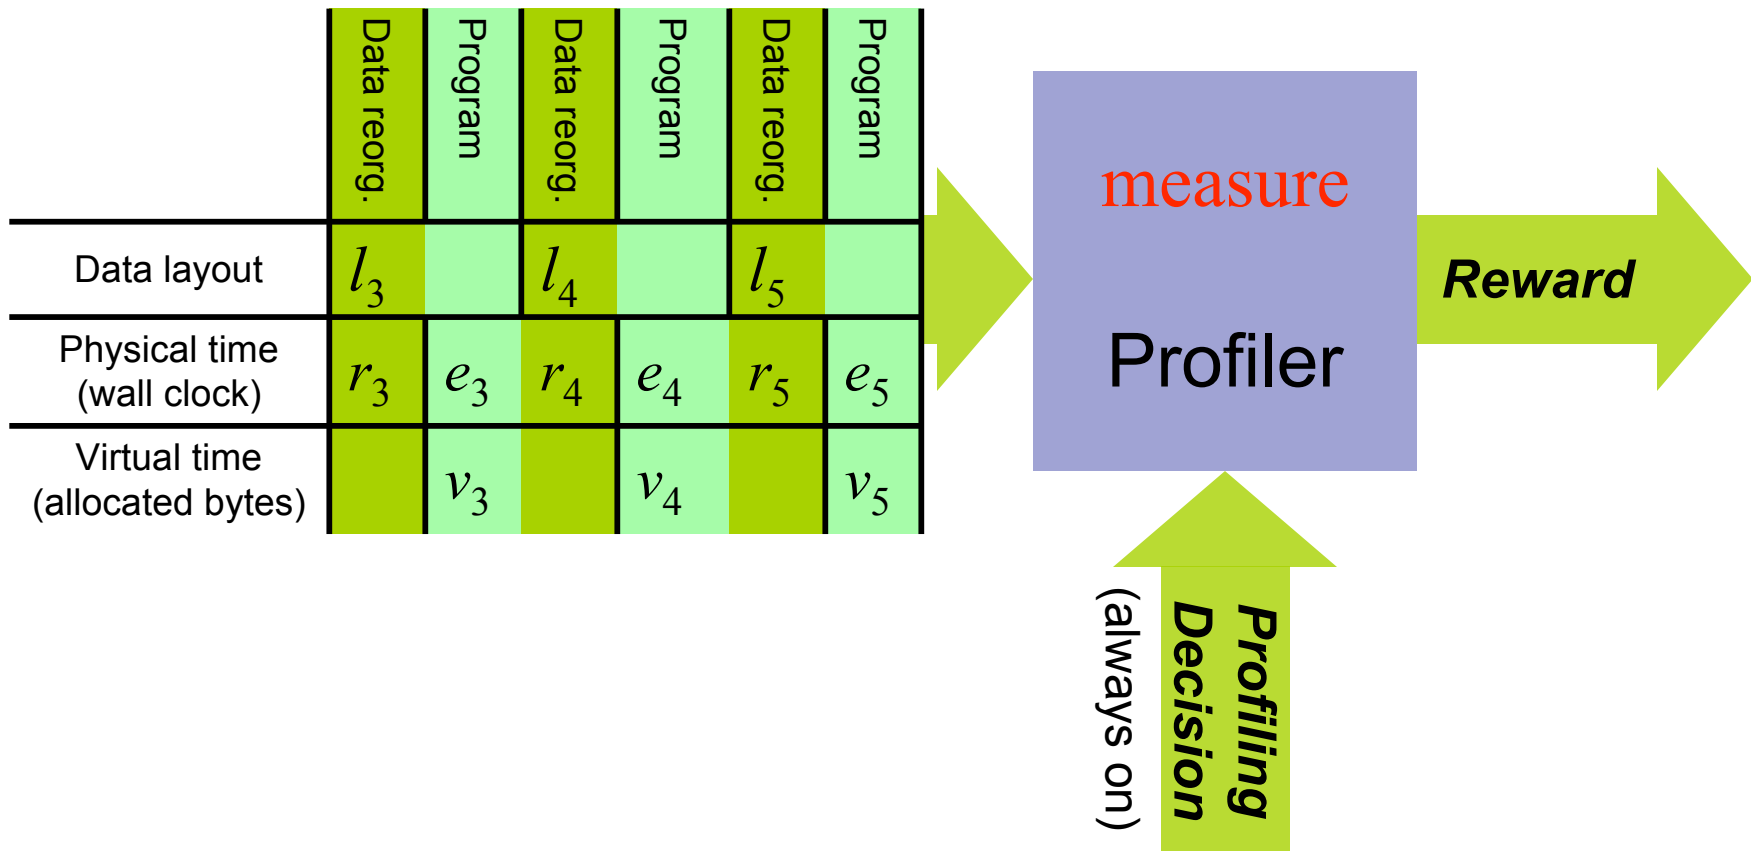
Multi-Armed Bandit Problem



Layout Auditing



Profiler



Reward for layout l_i uses historical average of:

- Virtual time v_i / program execution time e_i
- Virtual time v_{i-1} / reorganizer time r_i

Controller: Blind Justice



Goals

- Match performance of best layout
- Online

Challenges

- Confidence vs. Curiosity
- Phase changes vs. Noise

Confidence vs. Curiosity

Pick layout l if either:

- High **confidence** that l gives best reward
- High **curiosity** about l 's reward

	Confidence	Curiosity
Never tried layout		
Few samples / High variance		
Many samples / Low variance		

⇒ **use simulated annealing**

Phase Changes vs. Noise

Phase Adaptivity


- When layout performance changes, learn new best layout

⇒ **Forget** historical rewards

Noise Tolerance

- Perturbation from extraneous causes

⇒ **Remember** historical rewards



⇒ **use exponential decay**

SASO Properties of Control Systems


- **S**tability
- **A**ccuracy
- **S**ettling
- ~~Overshoot~~
- Phase adaptivity
- **O**verhead

Methodology


20 Java programs (DaCapo suite, SPECjvm98 suite, and a few more)											
J9 = IBM's product Java VM											
<i>HI</i>				<i>BF</i>				<i>LA</i>			
hierarchical				breadth-first				layout auditing			
4 Hardware Platforms											
Intel-2			AMD-2			AMD-4			AMD-8		

Accuracy and Overhead


Average %
slowdown
vs. best



Number of
programs
not optimal




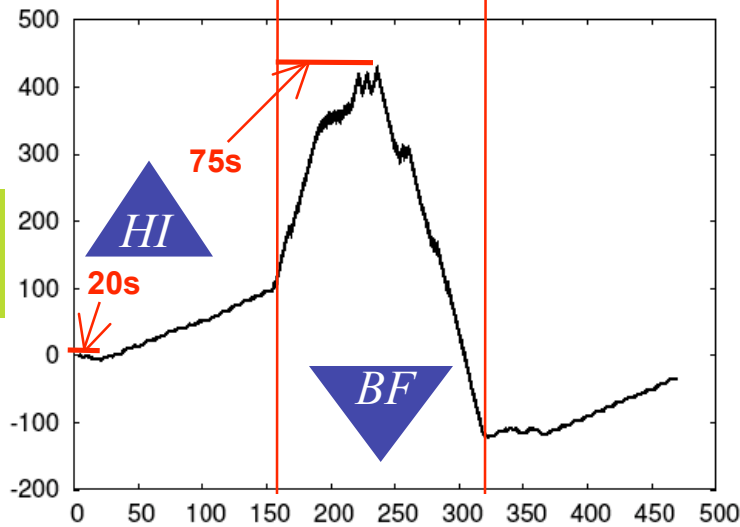
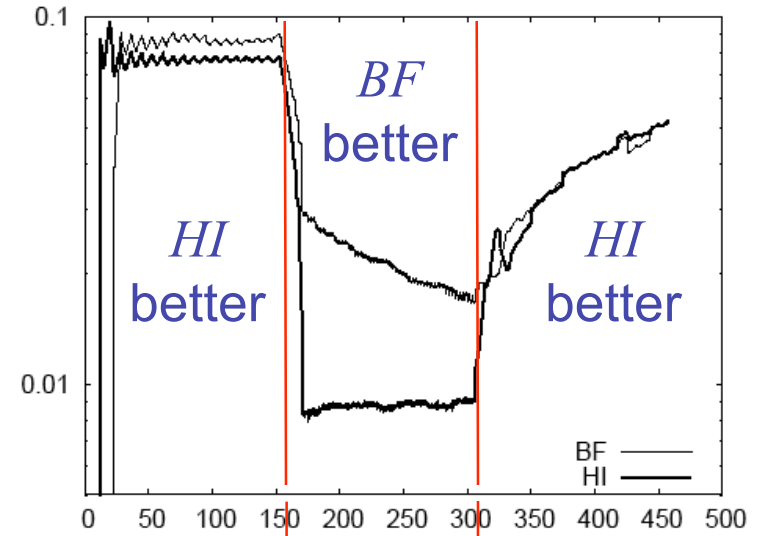

Worst %
slowdown
vs. best



Stability and Settling

Decay = 0.9



No decay (Decay=1.0)



Related Work

- Lau/Arnold/Hind/Calder PLDI'06: performance auditing for JIT optimization
- Soman/Krintz/Bacon ISMM'04: switch copy vs. mark-sweep, generations or not
- Chen/Bhansali/Chilimbi/Gao/Chuang PLDI'06: throttle unless miss rate reduced
- Saavedra/Park PACT'96: adapt prefetch distance based on cancellation & latency

Conclusions



- Accurate
- Phase adaptive (good settling/stability)
- Negligible overhead profiling
- Online, hardware independent

Clustering Layouts by Performance [SIGMETRICS 2007]

