image from pachd.com

RADISH:

Sound and Complete Race Detection in Software and Hardware

Joseph Devietti[*]
Benjamin P. Wood[*]
Karin Strauss[*^]
Luis Ceze[*]
Dan Grossman[*]
Shaz Qadeer[^]

[*]University of Washington
[^]Microsoft Research

Uses of Race Detection

- multithreaded record+replay
- simplifying consistency models
- atomicity checking
- atomicity enforcement
- testing & verification
- determinism checking
- determinism enforcement
- concurrency bug detection

many uses require
sound+complete detection

software

Flanagan and Freund, PLDI 2009

hardware

Min and Choi, ASPLOS 1991
Muzahid et al., ISCA 2009
Prvulovic, HPCA 2006

complementary strengths

	software	hardware
+	sound+complete static analysis	$ coherence event-based
−	slow polling-based	unsound/incomplete $ evictions descheduled threads $-line granularity

RADISH overview

sound+complete race detection in sw+hw

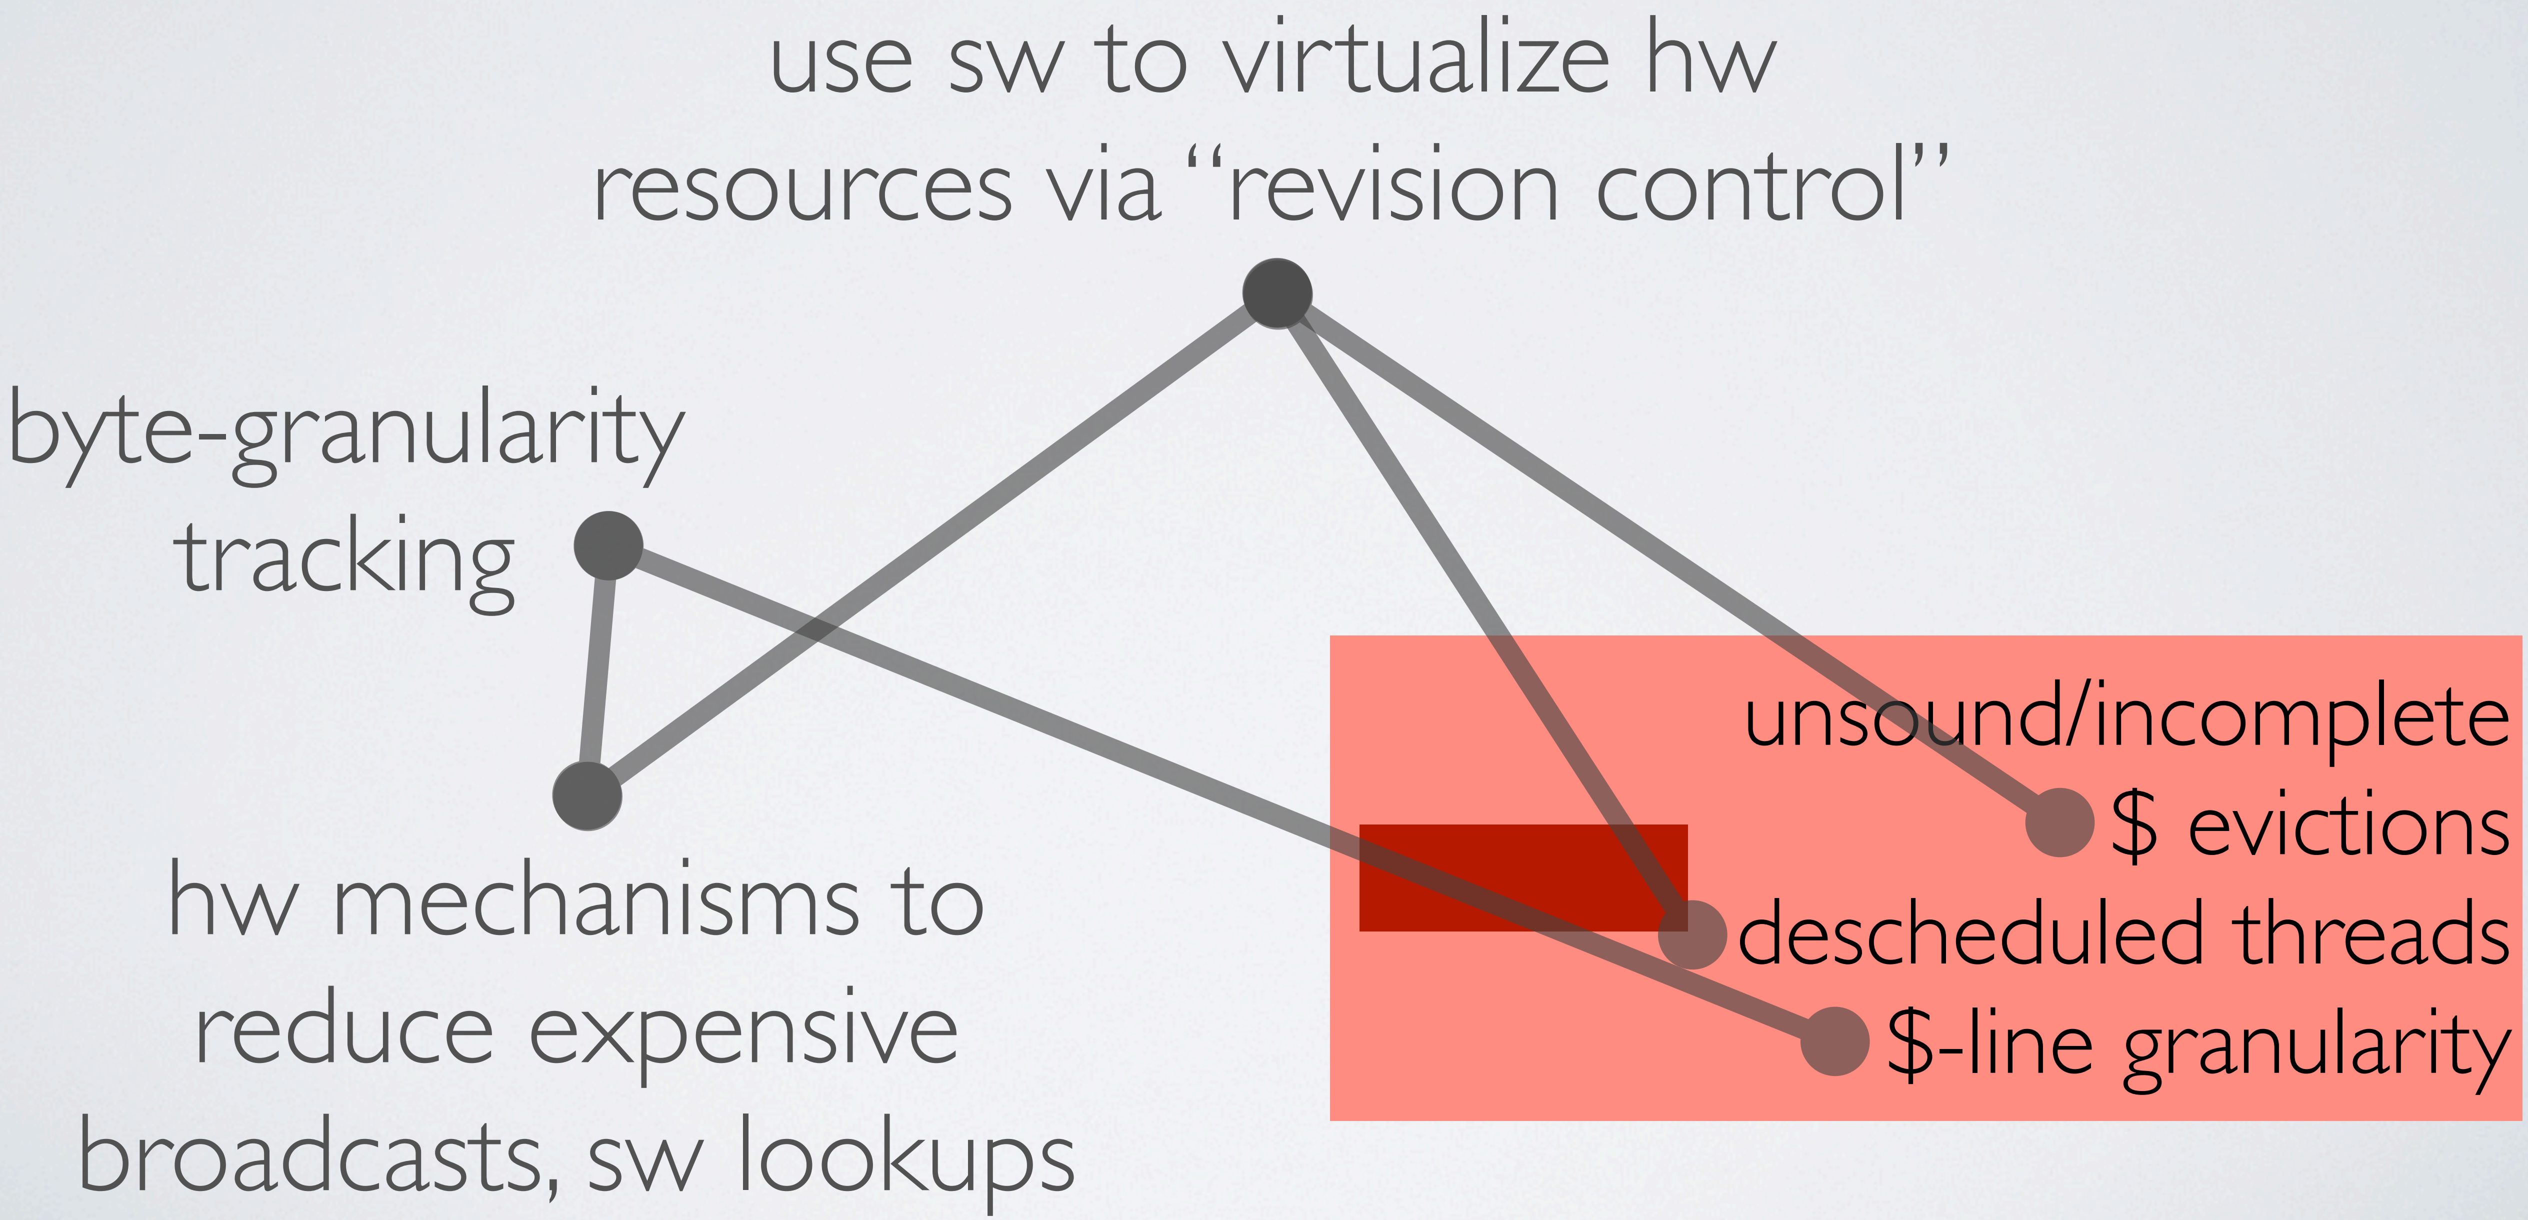

image from pachd.com

outline

data races

Lamport, CACM 1978

2 **concurrent** accesses to the same memory location,
≥ 1 of which is a write

unordered wrt the **happens-before relation**

transitive closure of
program order + synchronization order

data races

Lamport, CACM 1978

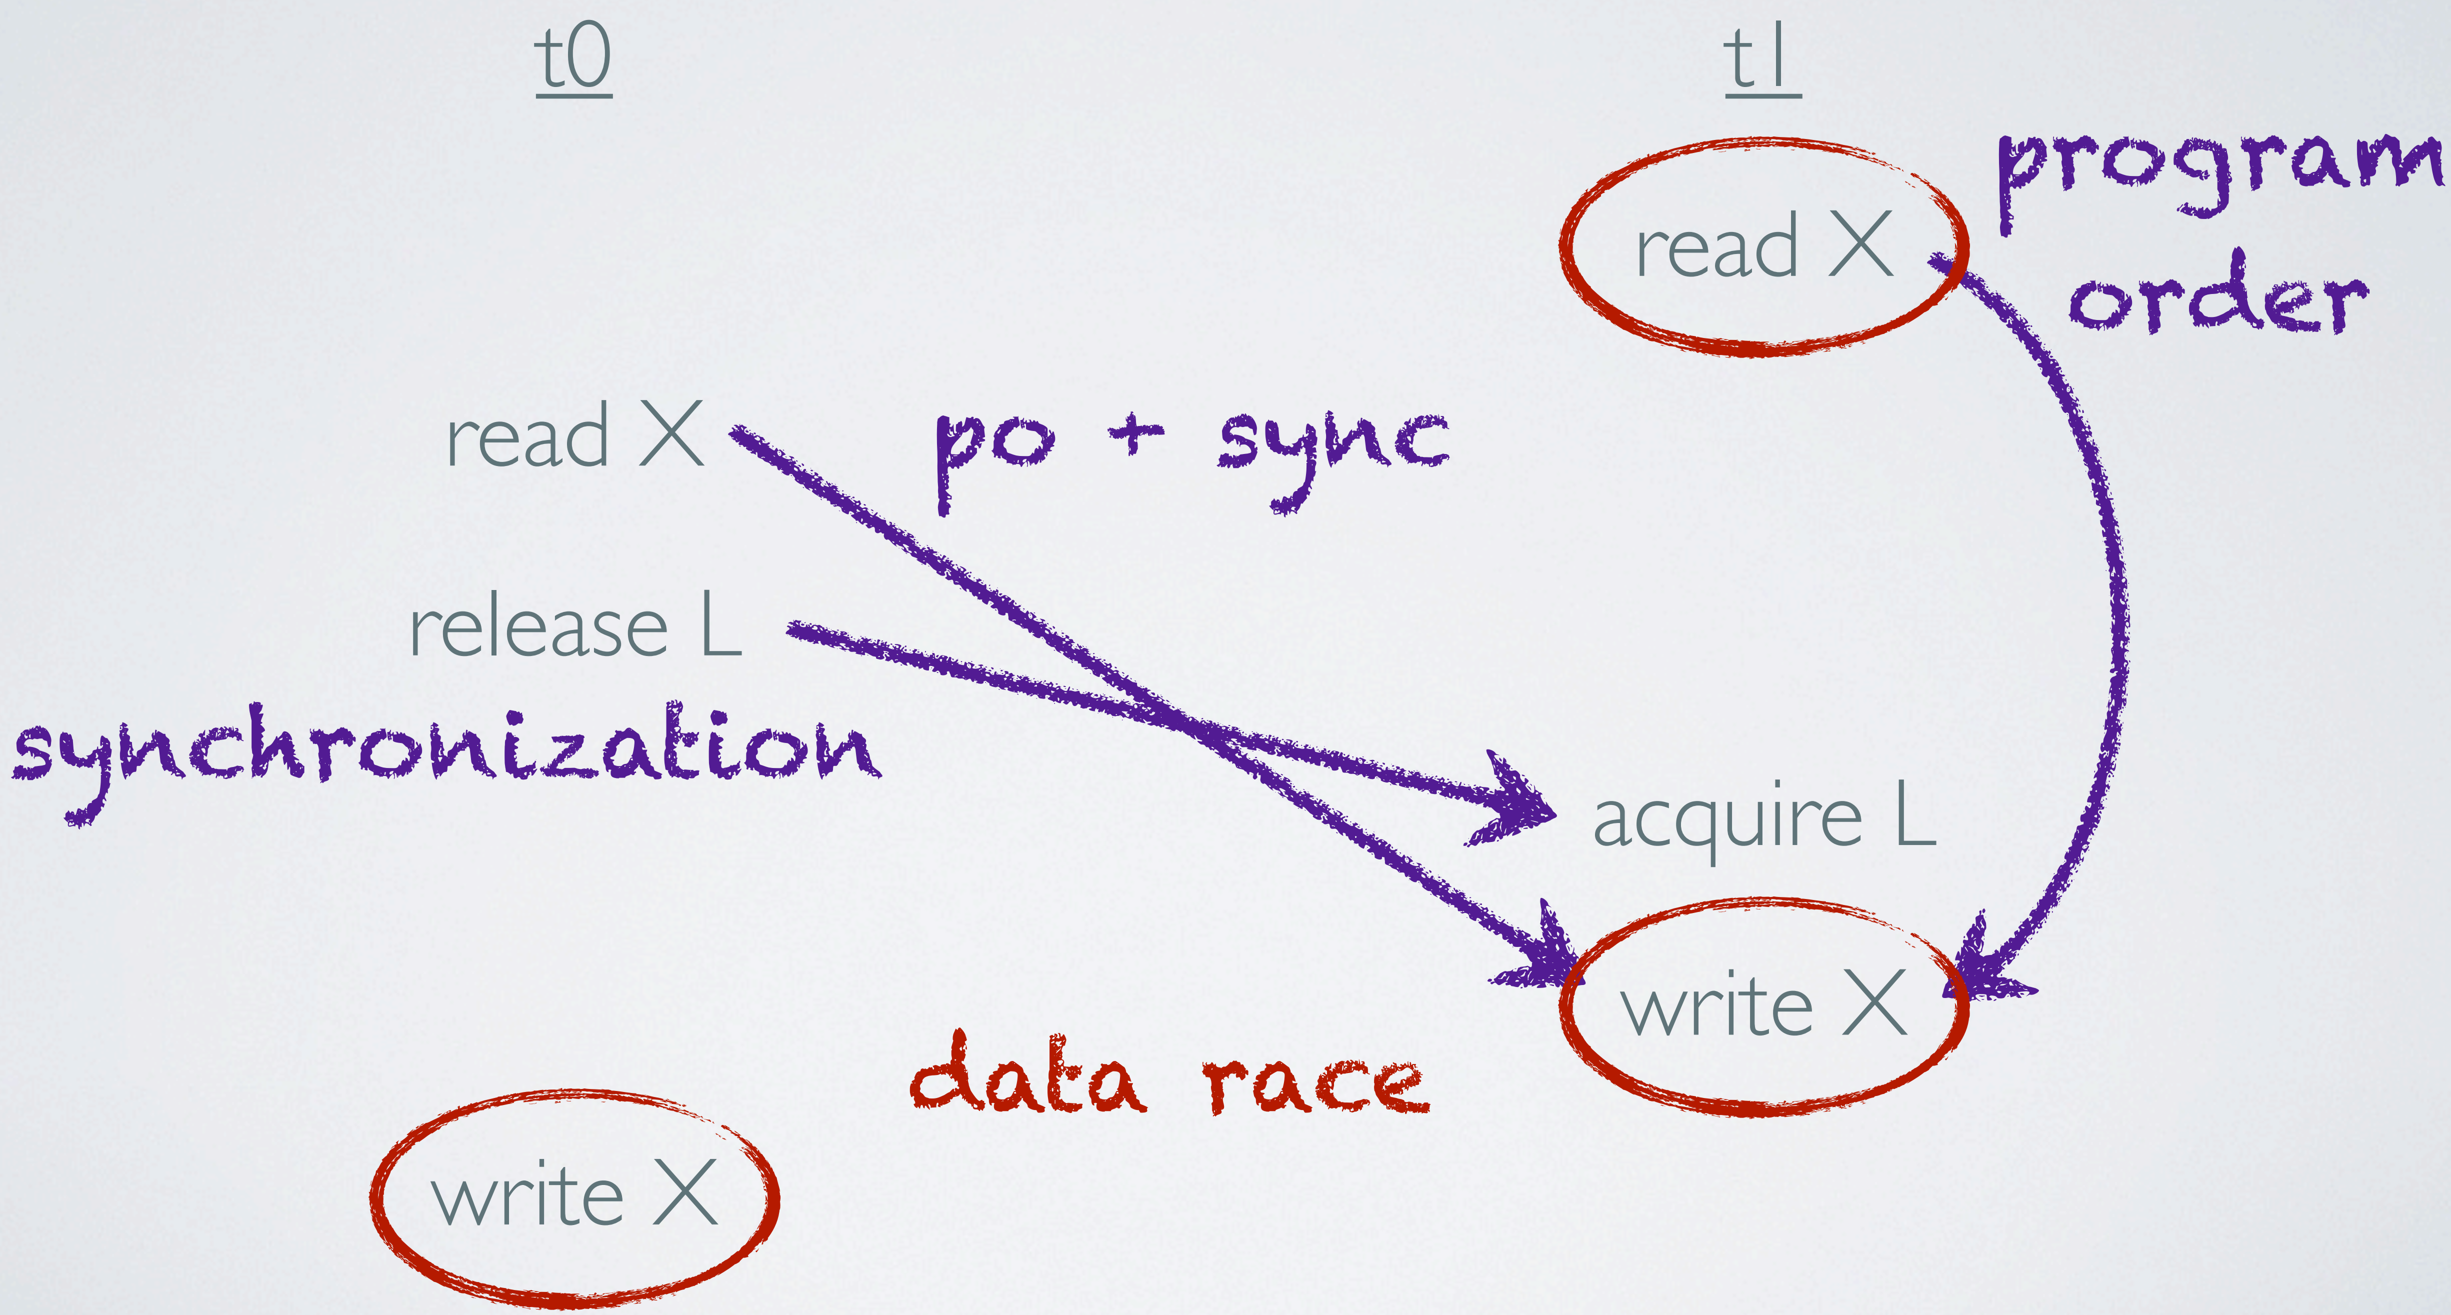

happens-before race detection

canonical sound+complete approach

Fidge, Computer 1991   Mattern, IWPDA 1989

per-**thread** metadata
read ordered with last write

thread	last synchronized with
t0	t1@*T*, t2@*U*

write ordered with last write
and all last reads
per-**location** metadata

location	last write	last reads
X	t2@*T*	t0@*U*, t1@-, t2@*W*

mapping to hardware

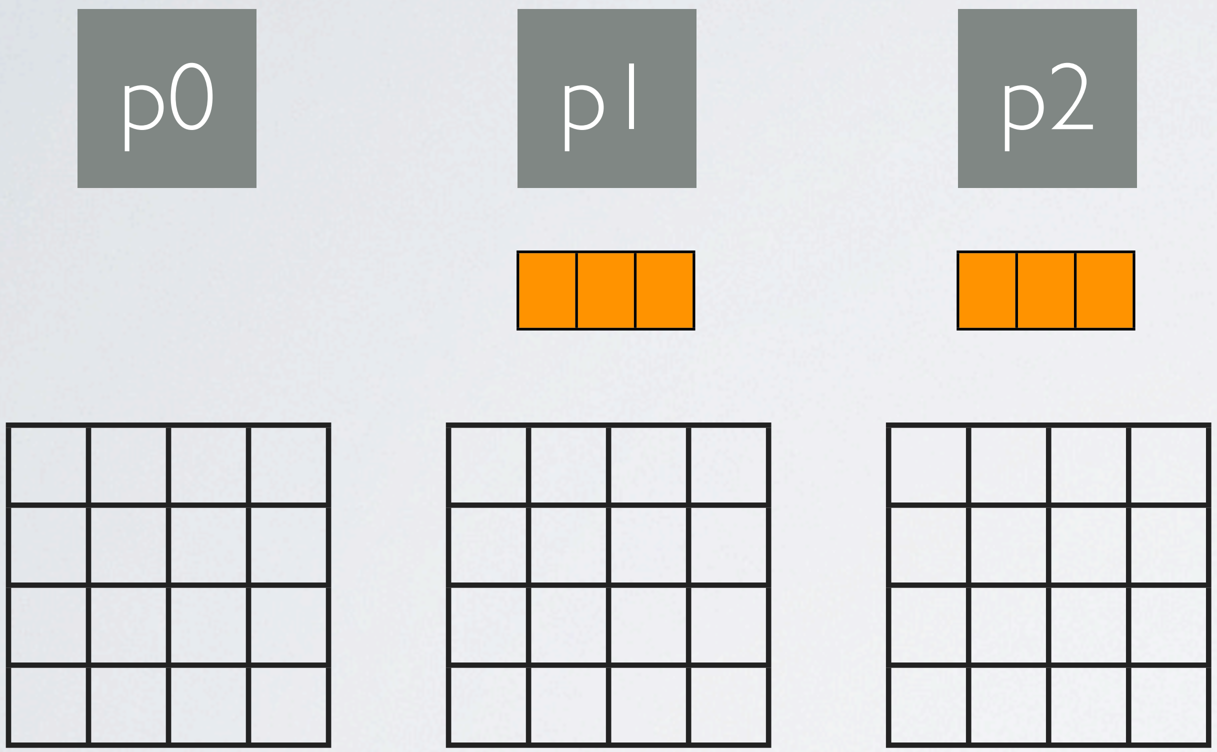

thread	last synchronized with
t0	t1@[illegible]@U

location	last write	last reads
X	t2@T	t0[illegible]U, t1[illegible]–, t2[illegible]W

unbounded # threads

unbounded # locations

image from pachd.com

outline

happens-before
data race detection

in-$ RADISH

full RADISH

results

conclusions

strawman

~~all metadata is in hw,~~
~~so broadcast on every access?~~

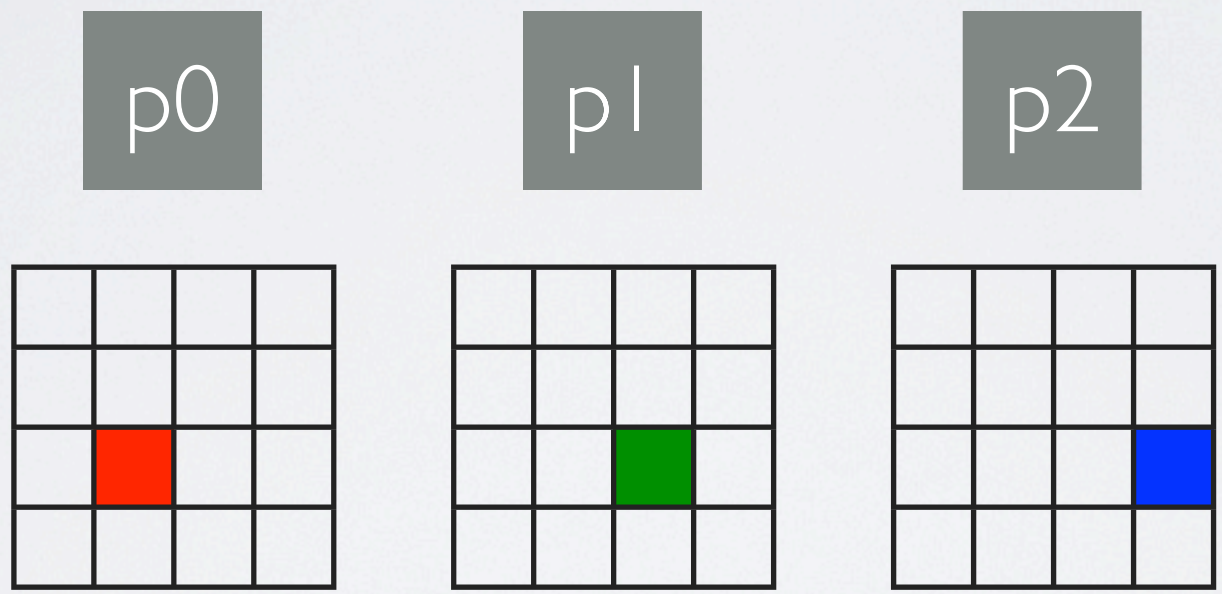

p0

write X

write X

local permissions cache what can be done without communication

local permissions

READ, WRITE or NONE permissions
to each byte in a $ line

updated only on permissions violations
and coherence events

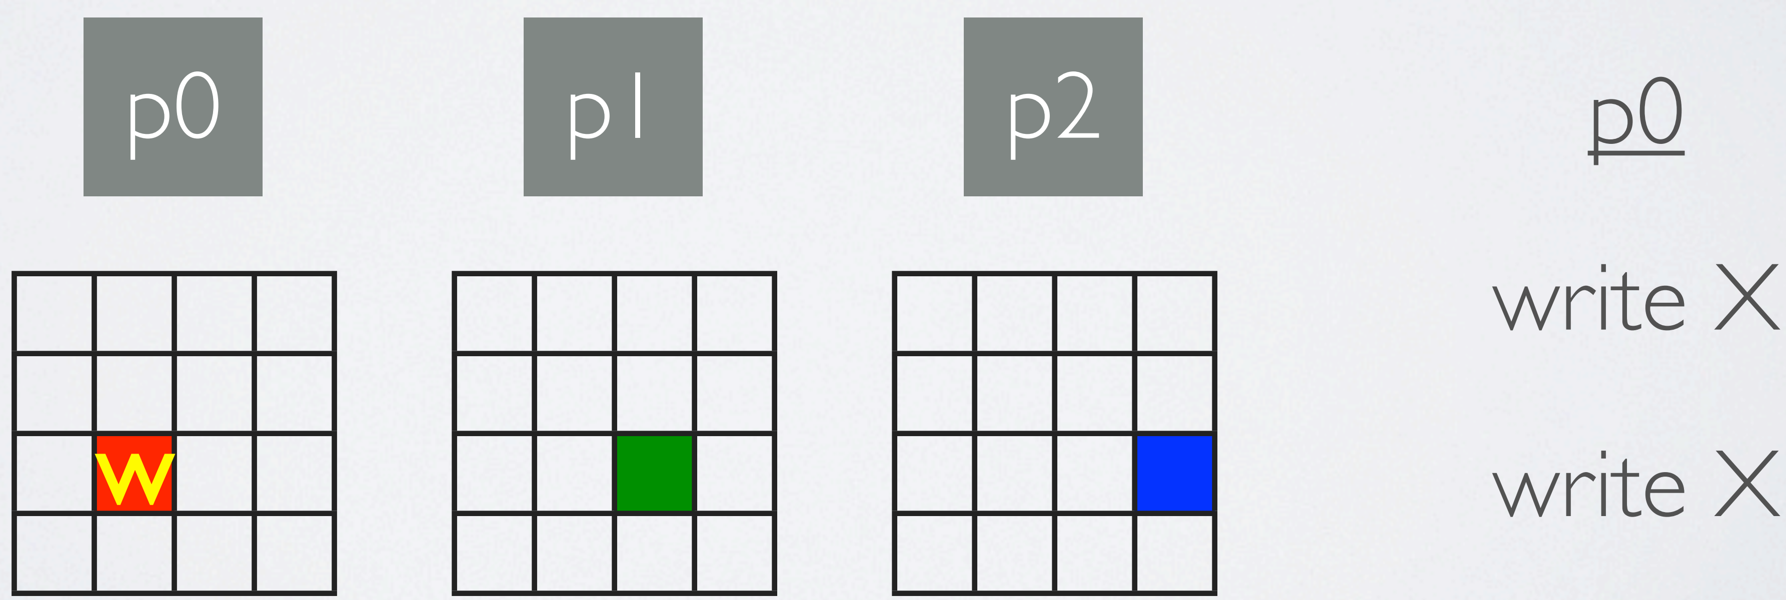

image from pachd.com

outline

strawman

~~metadata can be in hw or sw~~
~~so check sw on every access?~~

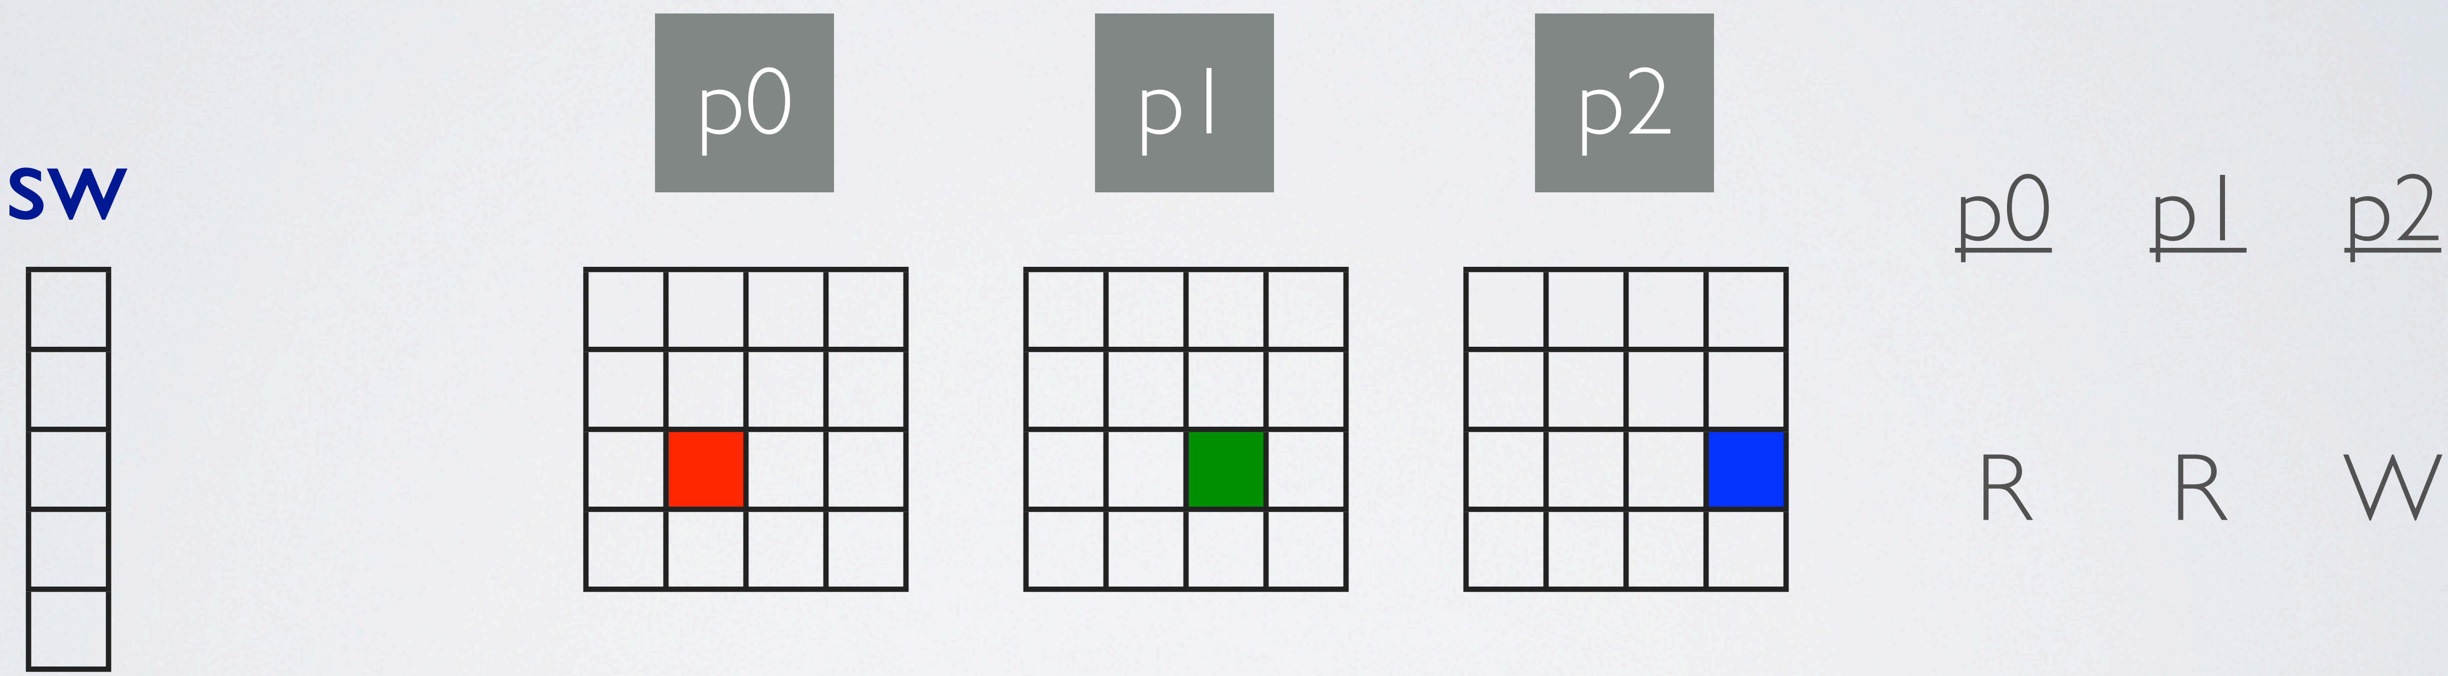

in-hardware status summarizes
what metadata resides in hw

in-hardware status

what can we figure out without going to sw?

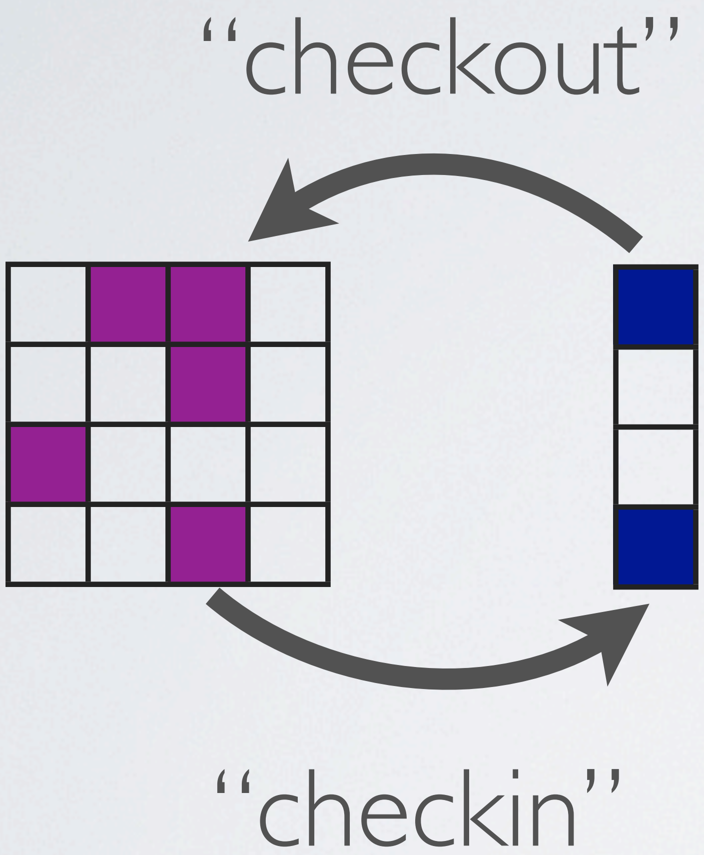

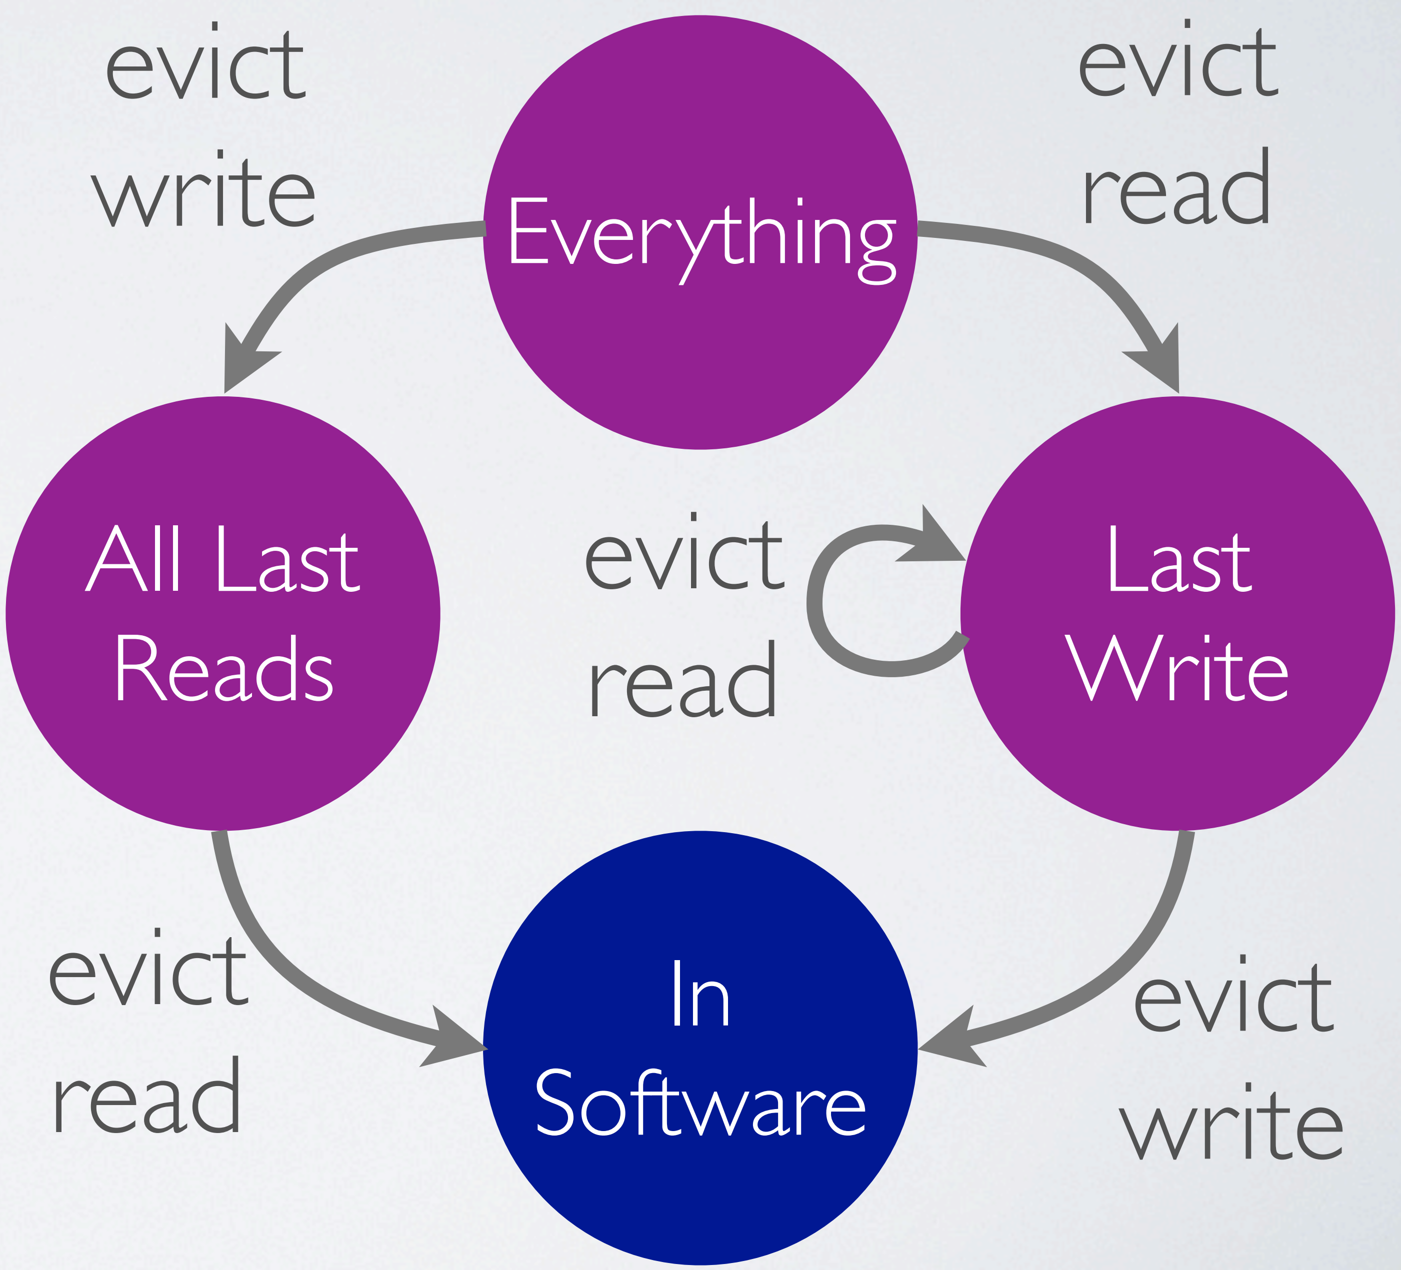

set IHS on checkout, degrade on $ evictions

also in the paper

leveraging type-safe languages to reduce metadata space overheads

asynchronous software lookups to reduce overheads

image from pachd.com

outline

- happens-before data race detection
- in-$ RADISH
- full RADISH
- results
- conclusions

simulation methodology

- Pin-based simulator
- 8 cores, MESI coherence
- 8-way 64KB L1, 8-way 256KB private L2, 16-way 16MB L3
- PARSEC 2.1
- compare with FastTrack [Flanagan and Freund, PLDI 2009]

runtime compared to native

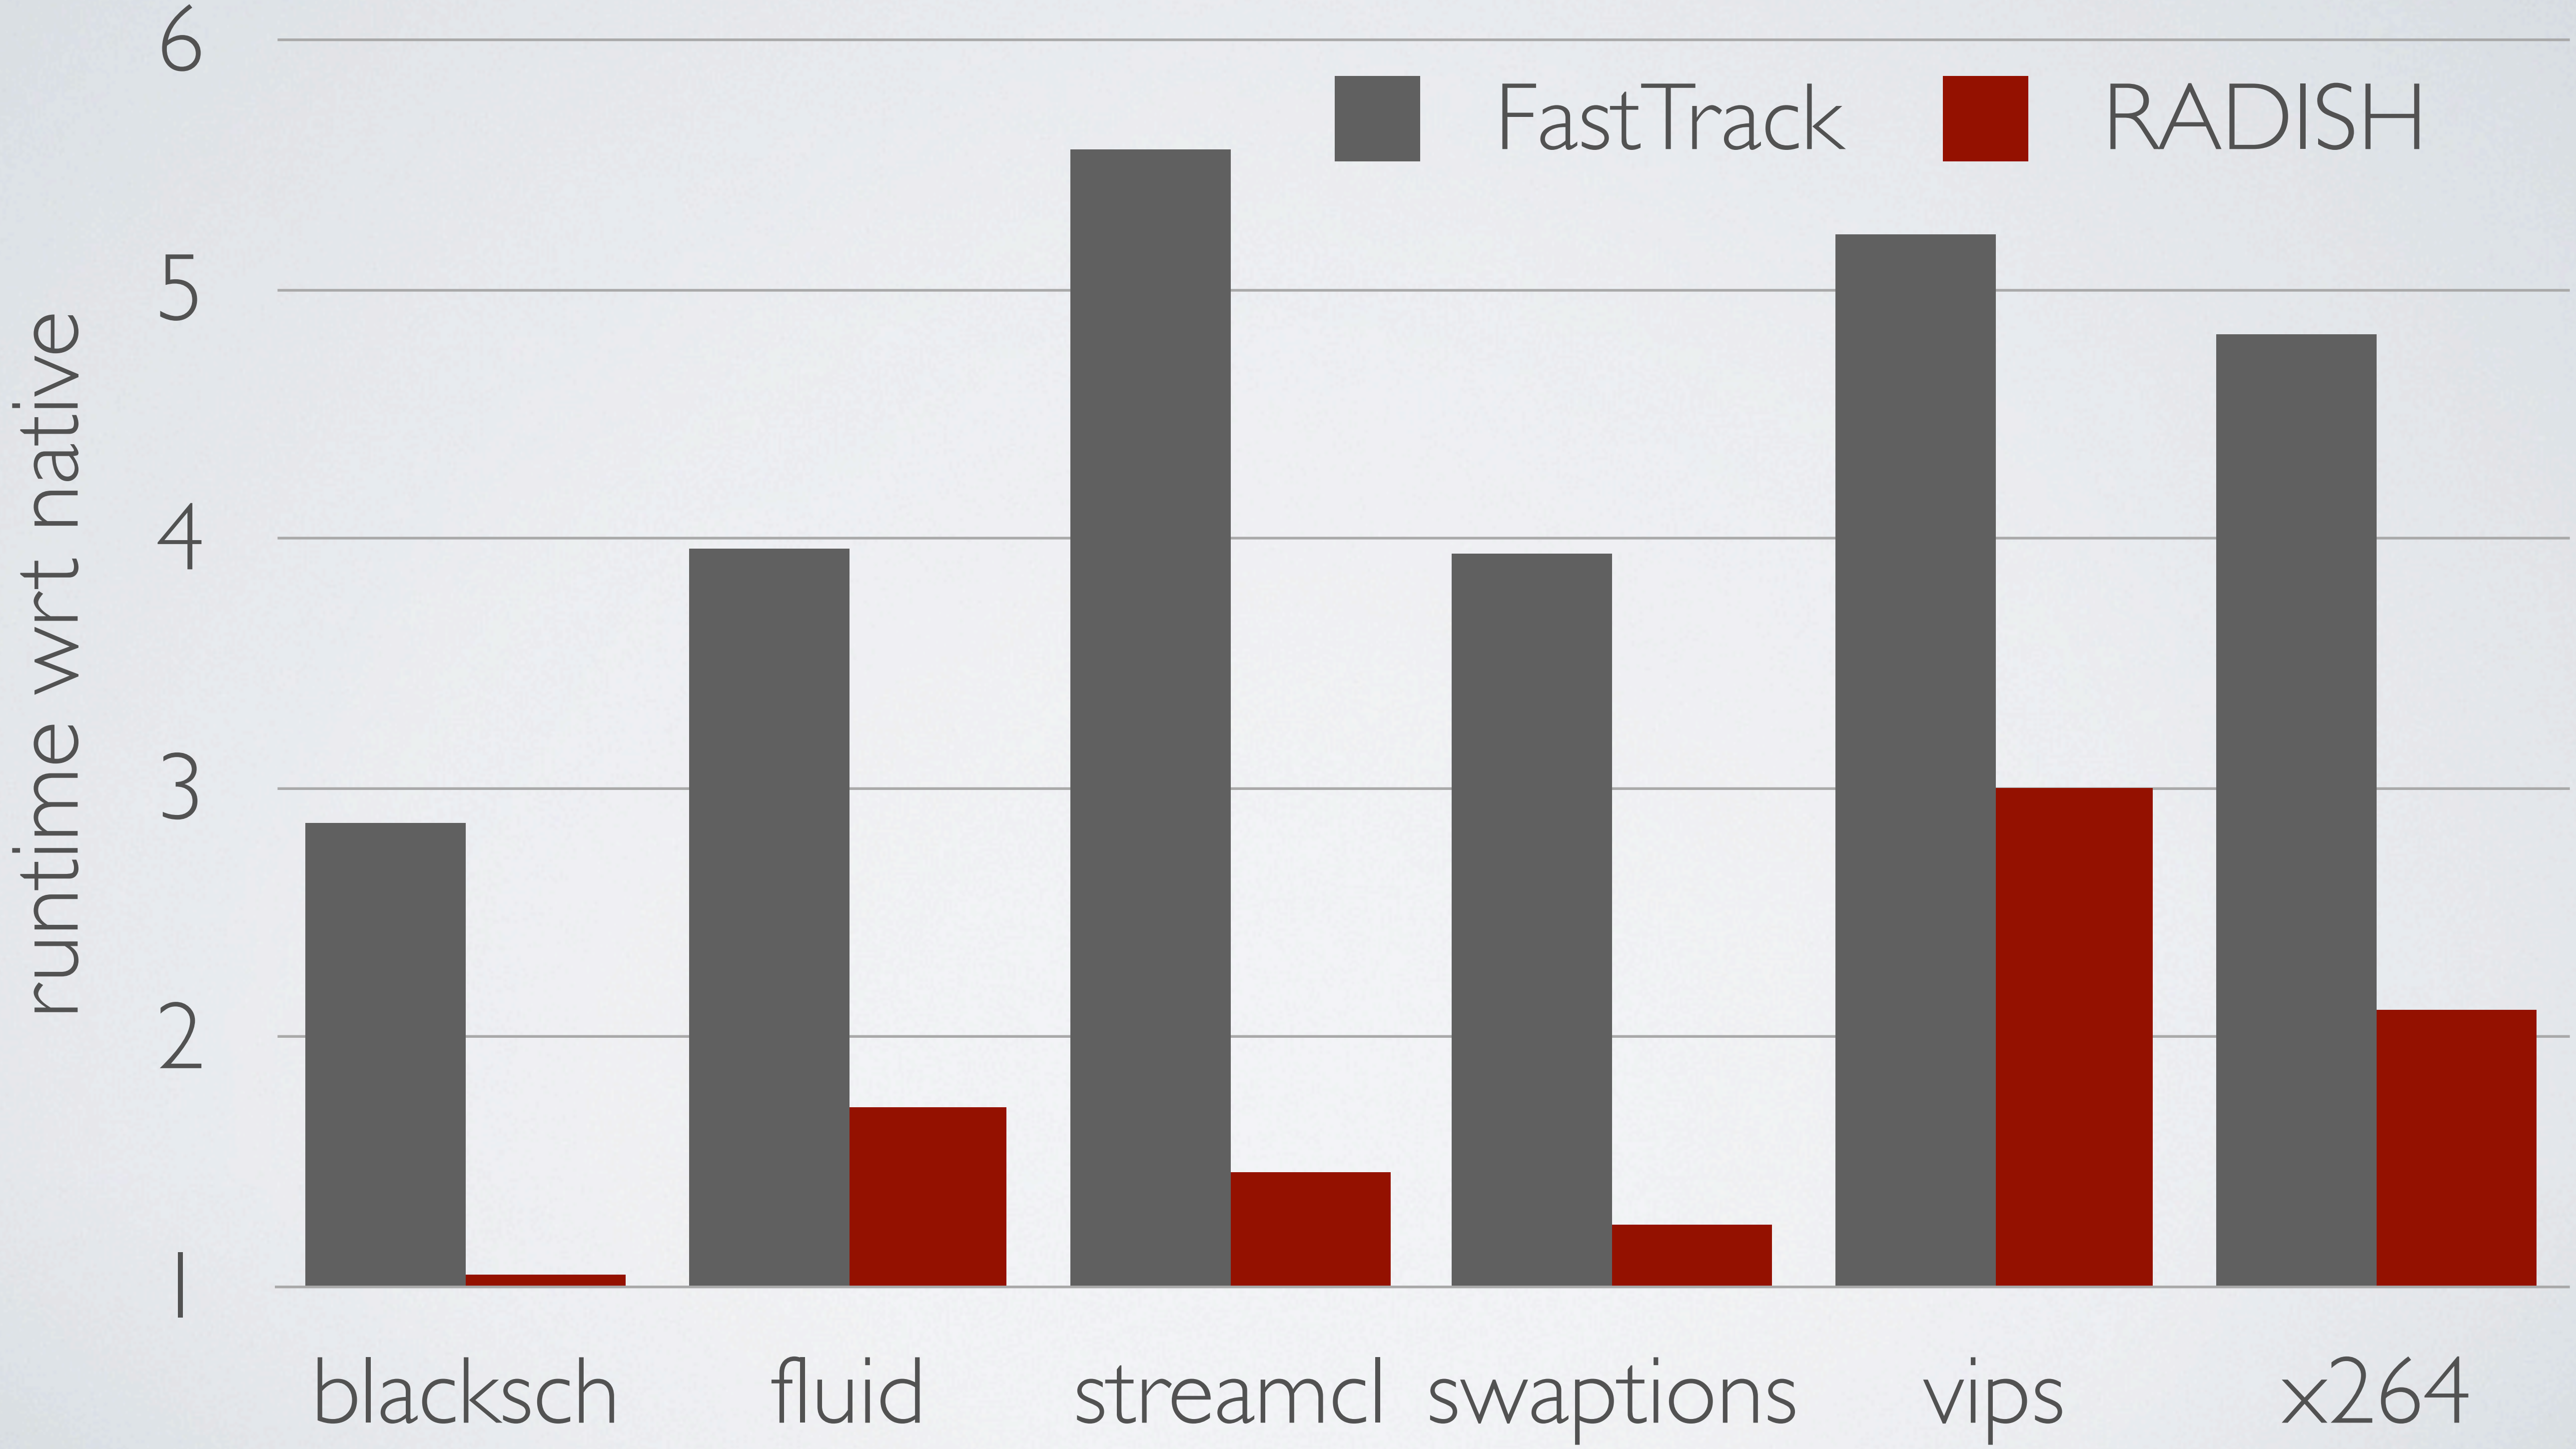

runtime compared to FastTrack

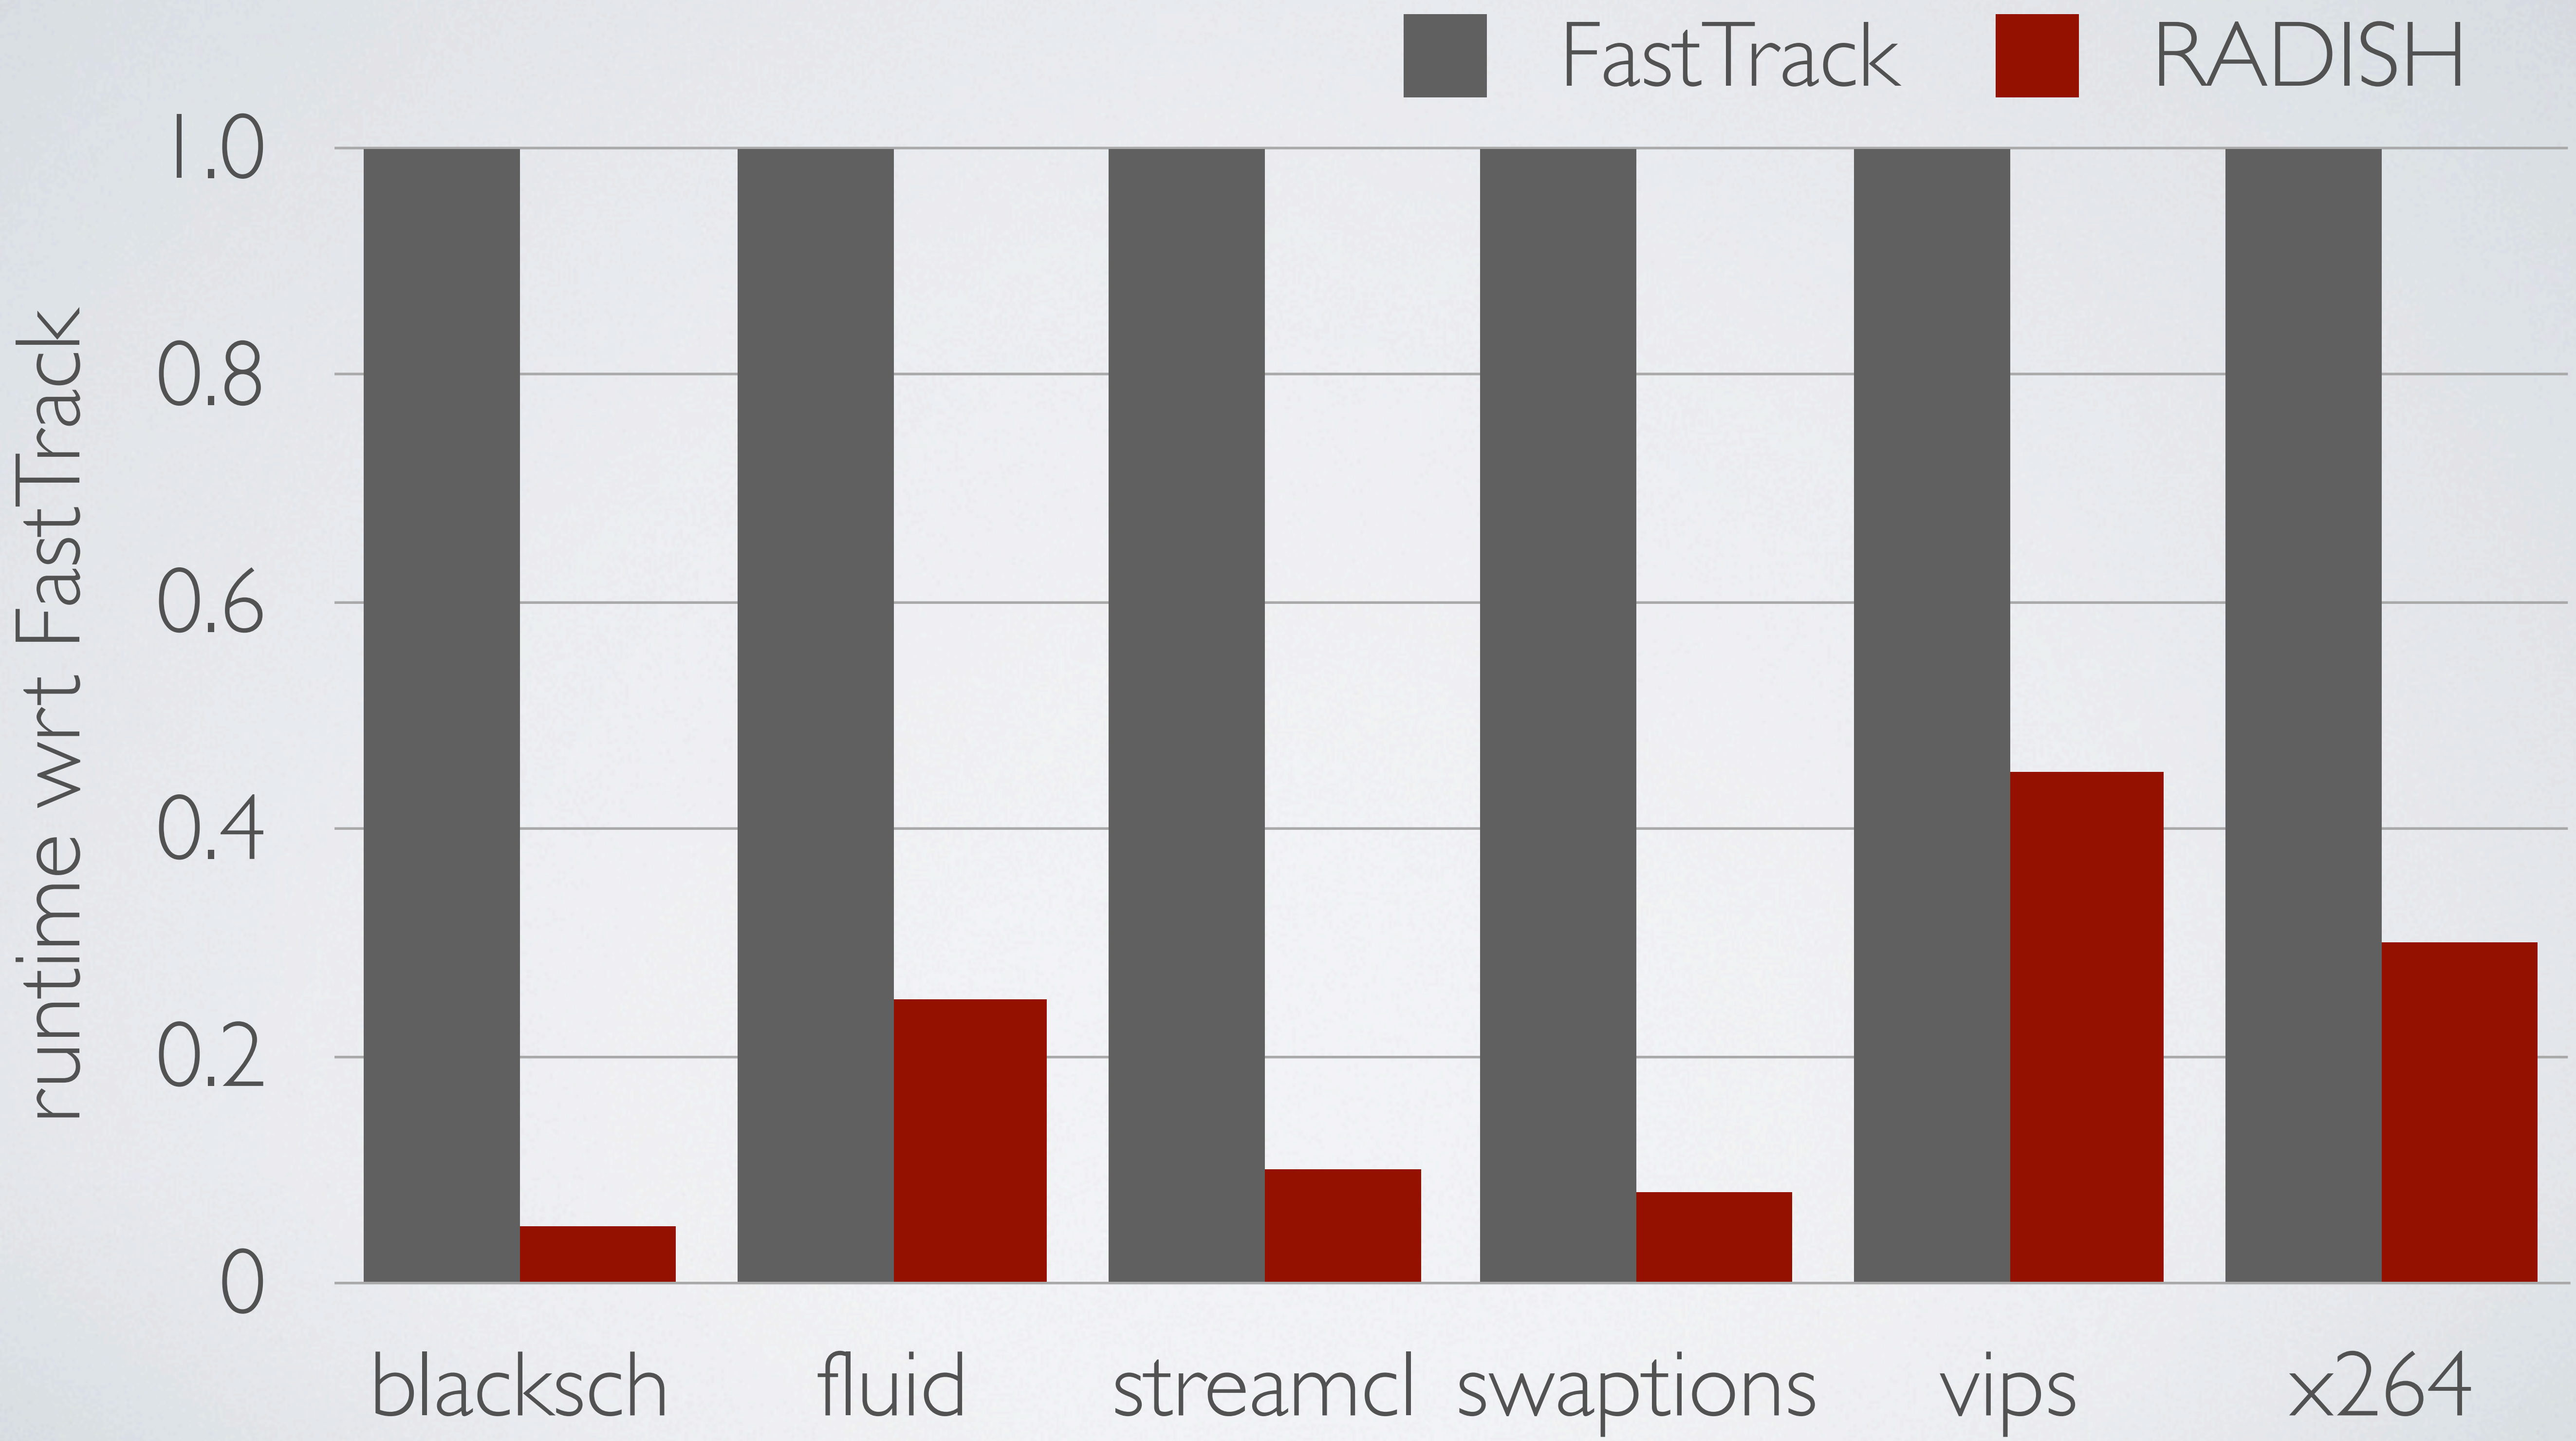

conclusions

sound+complete race detection in hw+sw

unmodified cache design

much faster than software-only race detection

image from pachd.com

thanks!

W

University of Washington Computer Science & Engineering

sampa

Microsoft Research