Flexible Snooping:

Adaptive Forwarding and Filtering of Snoops in Embedded-Ring Multiprocessors

Karin Strauss, Xiaowei Shen*, Josep Torrellas

University of Illinois at Urbana-Champaign

*IBM Research

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

http://iacoma.cs.uiuc.edu

Motivation

- CMPs are becoming standard components
- cheaper to build medium size machines
  - 32 to 128 cores (multi-CMP)
- shared memory, cache coherent
  - easier to program, easier to manage
- supporting cache coherence is difficult

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Cache coherence solutions

strategy	ordered network?	pros	cons
snoopy broadcast bus	yes	simple	difficult to scale
directory based protocol	no	scalable	indirection, extra hardware
snoopy embedded ring	no	simple	long latencies

- other proposals (e.g. token coherence)

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Contributions

- family of adaptive coherence protocols for rings
- two were chosen as best options

compared to **fastest** state-of-the-art scheme

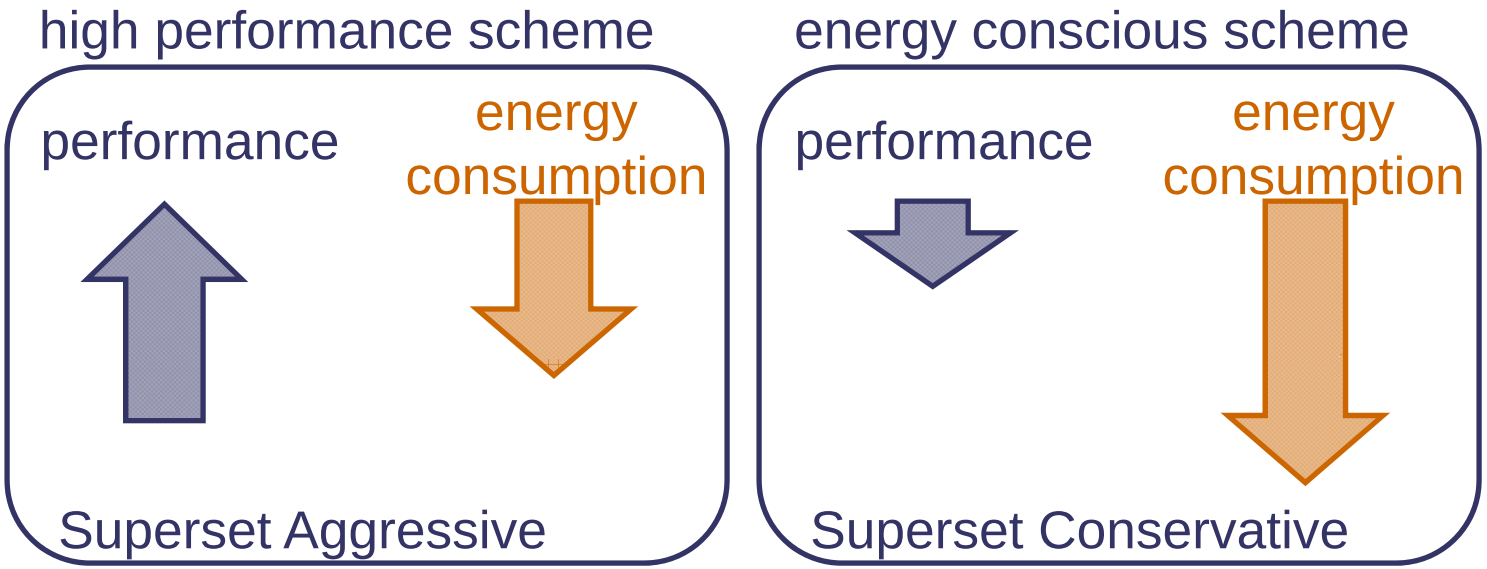

QuickTime™ and a TIFF (Uncompressed) decompressor are needed to see this picture.

Multi-CMP multiprocessor

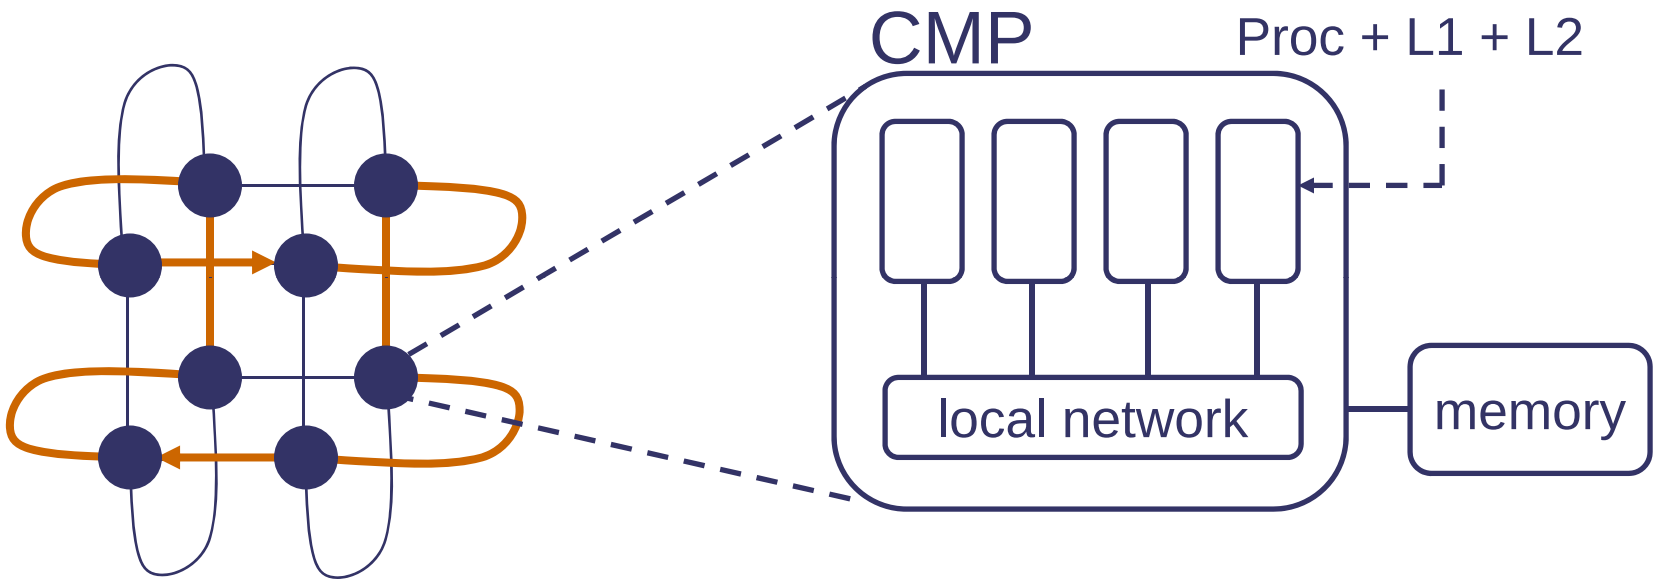

- coherence protocol used: only one supplier if line is cached

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Ring in action

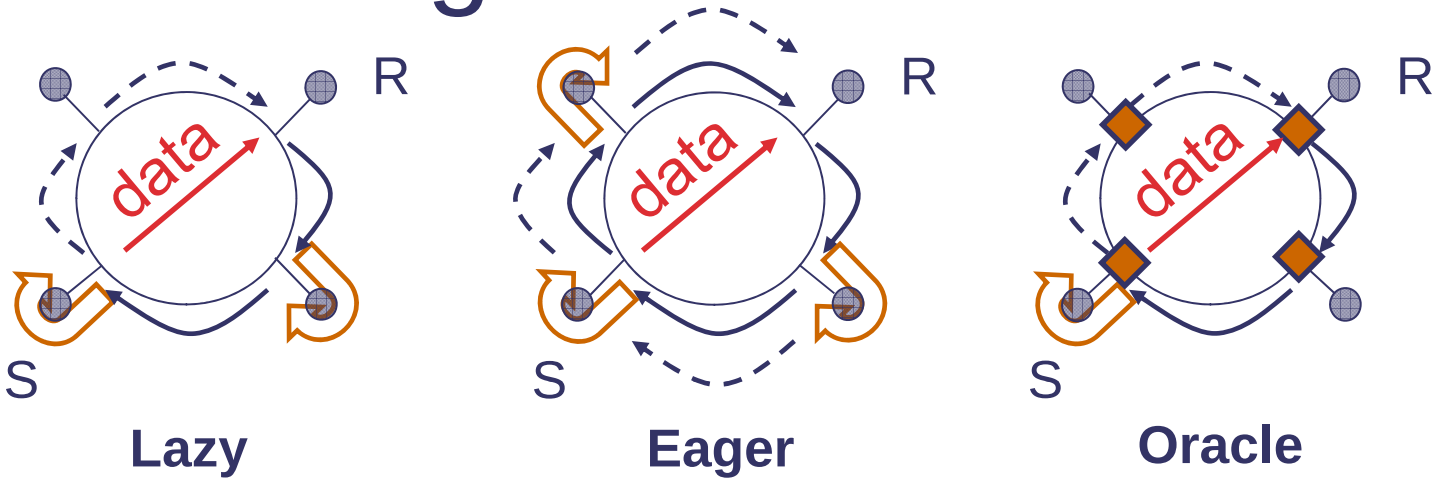

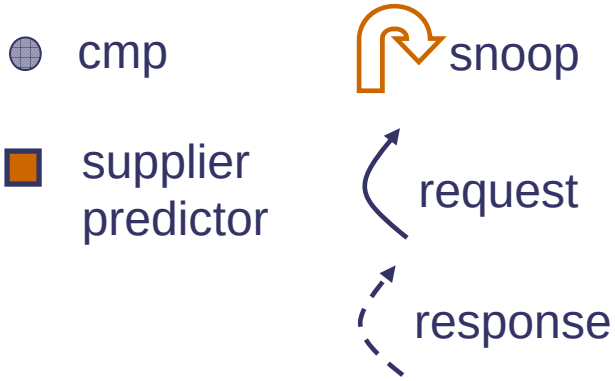

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Ring in action

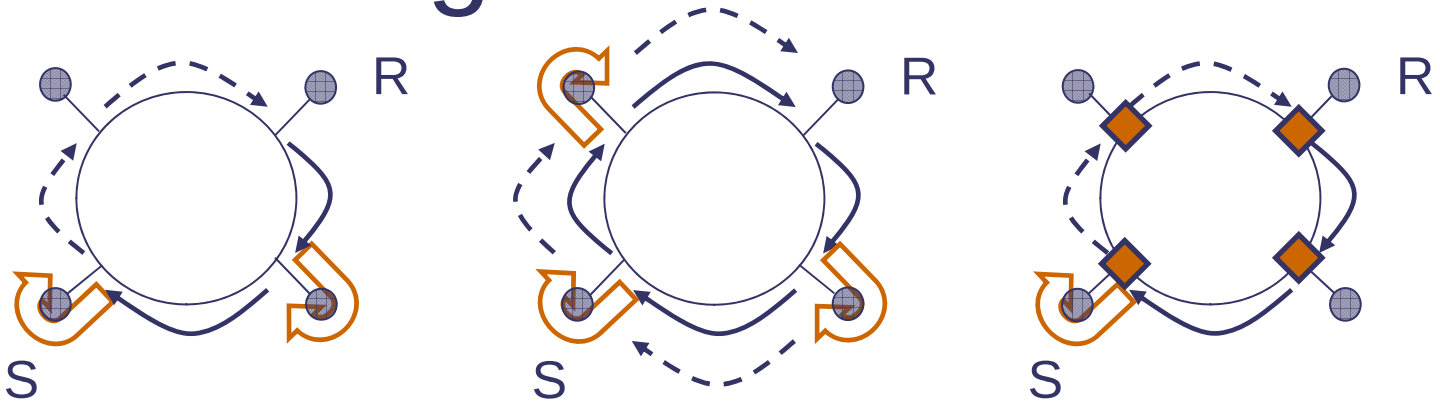

	Lazy	Eager	Oracle
latency	↑	↓	↓
snoops	↓	↑	↓ ↓
messages	↓	↑	↓

- **goal:** adaptive schemes that approximate Oracle's behavior

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Primitive snooping actions

- snoop and then forward
  - + fewer messages
- forward and then snoop
  - + shorter latency
- forward only
  - + fewer snoops
  - + shorter latency
  - – false negative predictions **not** allowed

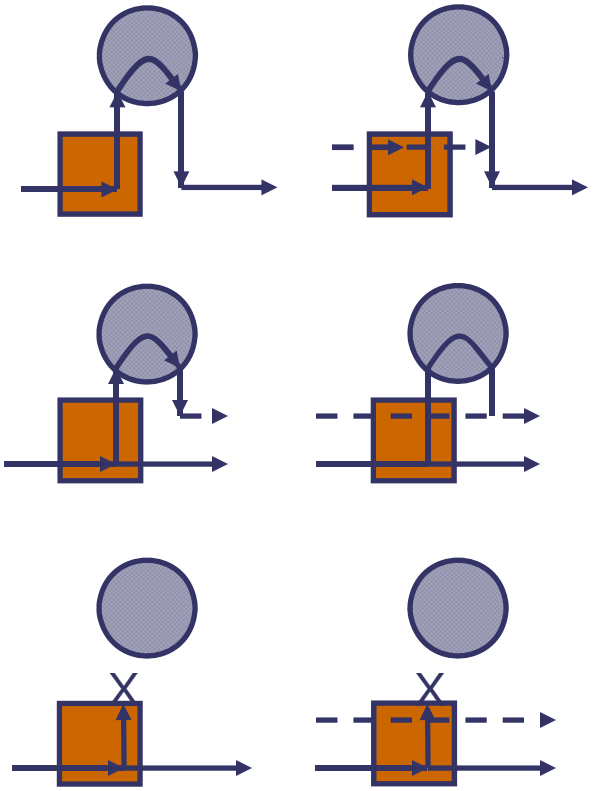

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Predictors and algorithms

set of addresses:

- node can supply
- in predictor

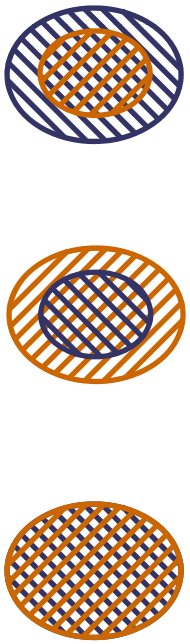

predictor / algorithm		action on negative prediction	action on positive prediction
Subset		forward then snoop	snoop
Super set	**Con**	forward	snoop then forward
	Agg		forward then snoop
Exact		forward	snoop

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Algorithms

algorithm		negative	positive
Subset		forward then snoop	snoop
Superset	**Con**	forward	snoop then forward
	Agg		forward then snoop
Exact		forward	snoop

Per miss service:

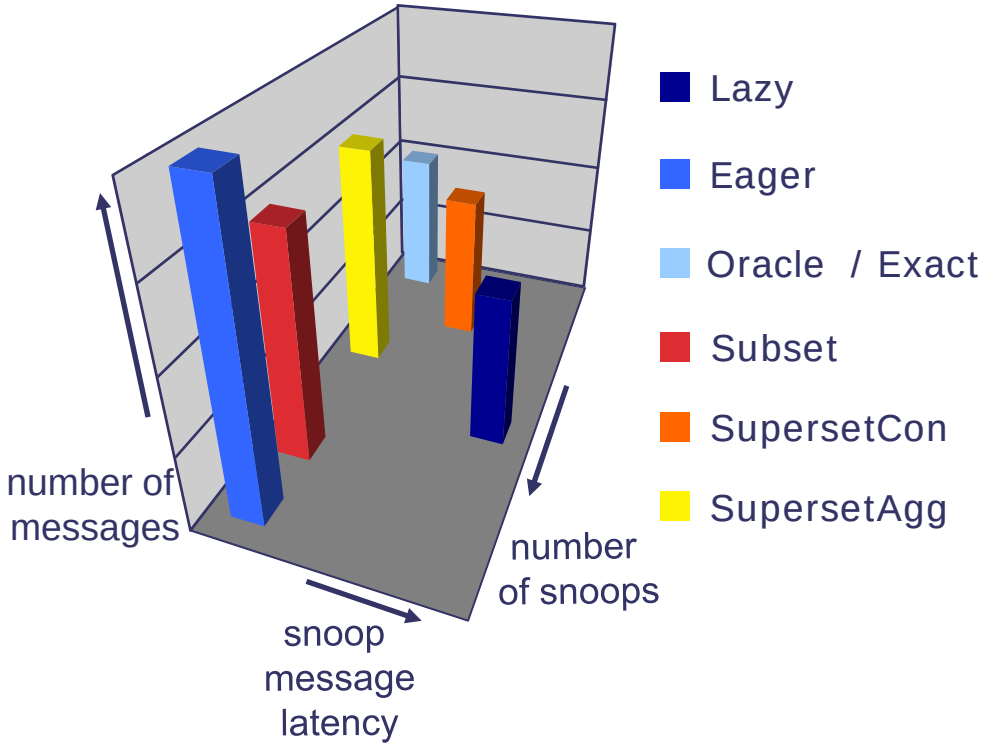

QuickTime™ and a TIFF (Uncompressed) decompressor are needed to see this picture.

Predictor implementation

- Subset
  - associative table:

    subset of addresses that can be supplied by node
- Superset
  - bloom filter: superset of addresses that can be supplied by node
  - associative table (exclude cache):

    addresses that recently suffered false positives
- Exact
  - associative table: all addresses that can be supplied by node
  - downgrading: if address has to be evicted from predictor table, corresponding line in node has to be downgraded

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Downgrading

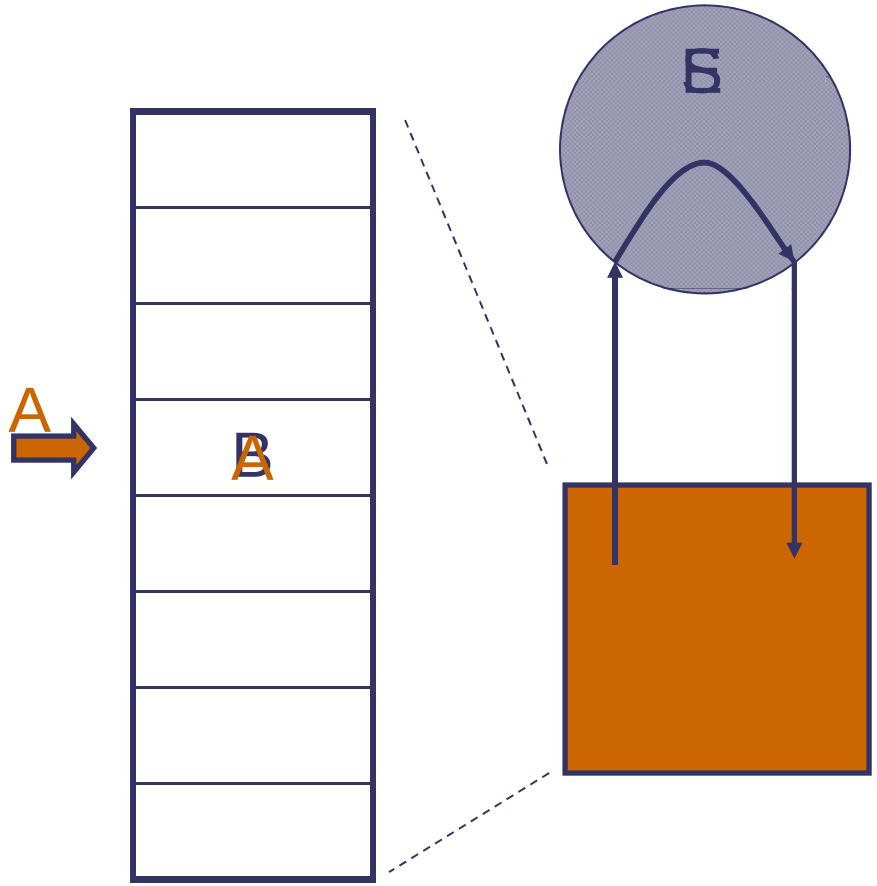

Negative effects:

- writes by this node need to snoop other nodes
- reads and writes by other nodes need to fetch line from memory

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Experiments

- 8 CMPs, 4 ooo cores each = 32 cores
  - private L2 caches
- on-chip bus interconnect
- off-chip 2D torus interconnect with embedded unidirectional ring
- per node predictors: latency of 3 processor cycles
- sesc simulator (sesc.sourceforge.net)
- SPLASH-2, SPECjbb, SPECweb

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Execution time

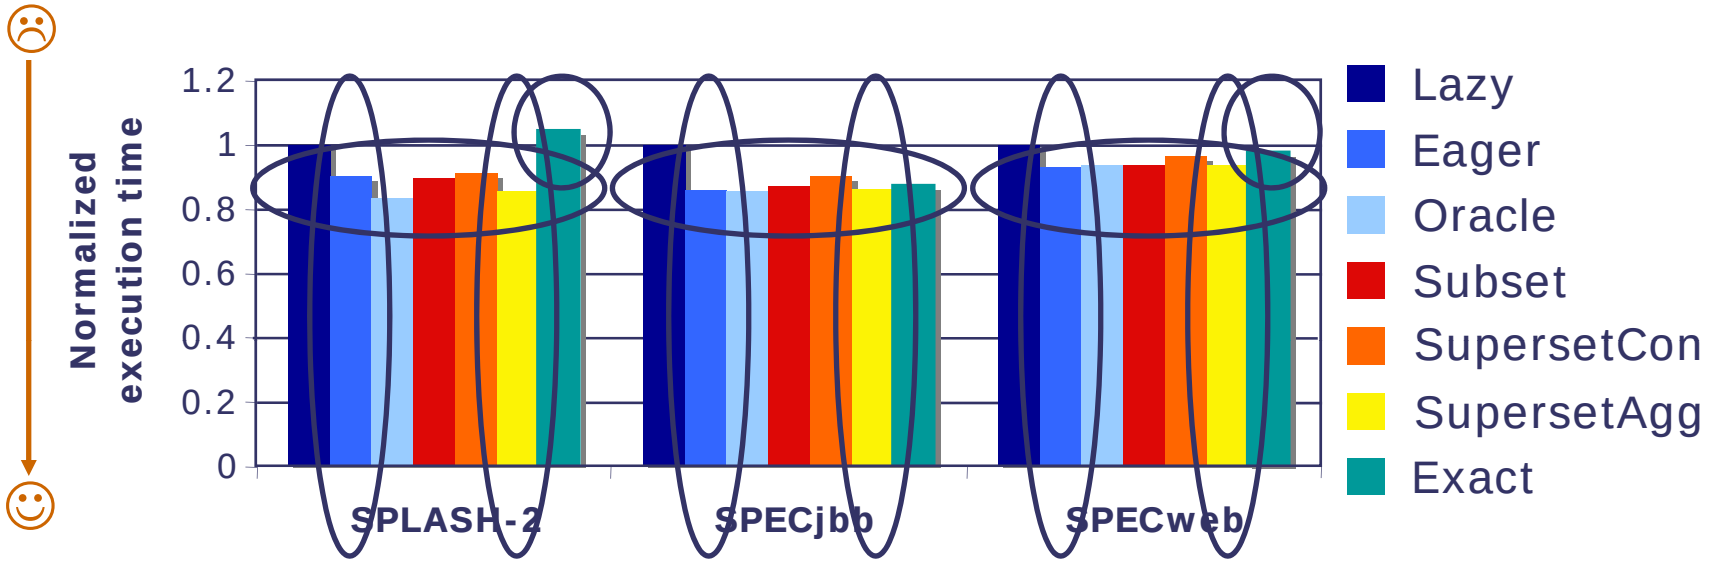

- performance of most flexible snooping algorithms is similar to Eager
- the fastest of all algorithms is SupersetAgg

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Miss service energy

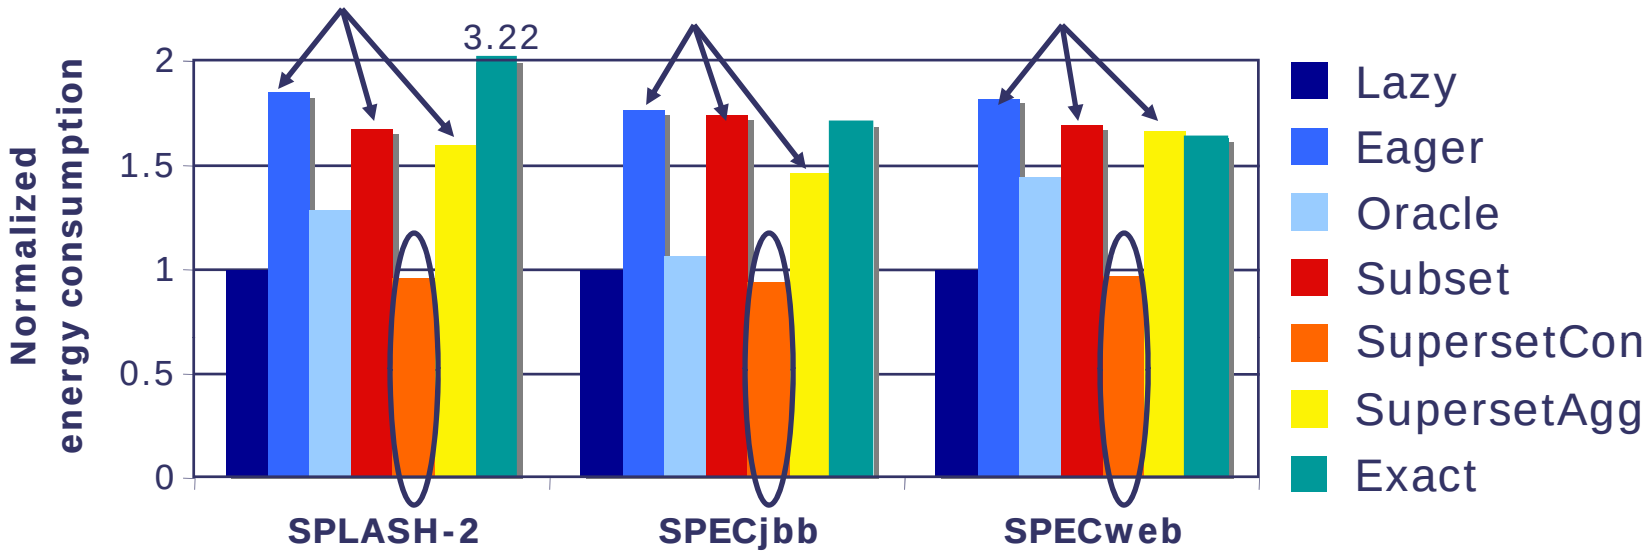

- algorithms that eagerly forward messages use more energy
- SupersetCon is least energy-hungry algorithm

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Most cost-effective algorithms

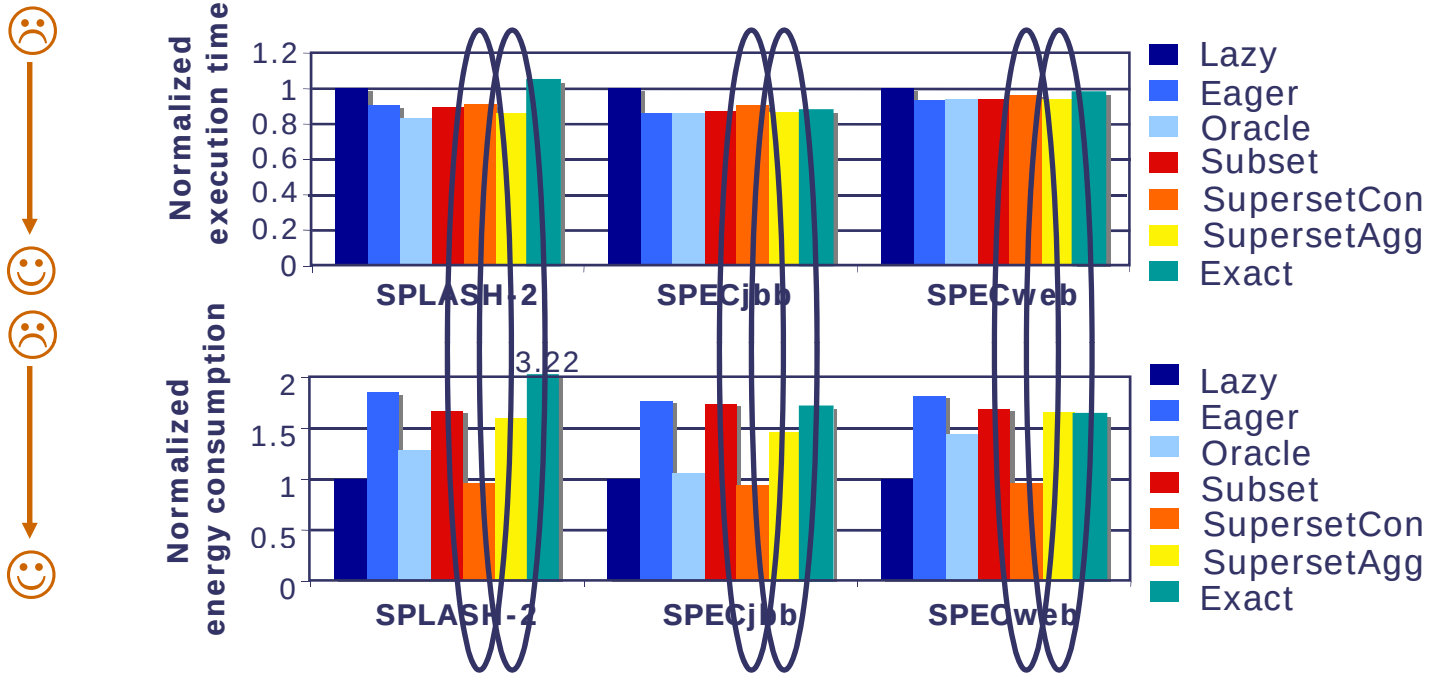

- high performance: Superset Aggressive
  - faster than Eager at lower energy consumption
- energy conscious: Superset Conservative
  - slightly slower than Eager at much lower energy consumption

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Most cost-effective algorithms

compared to **fastest** state-of-the-art scheme (Eager)

Superset Aggressive
high performance scheme

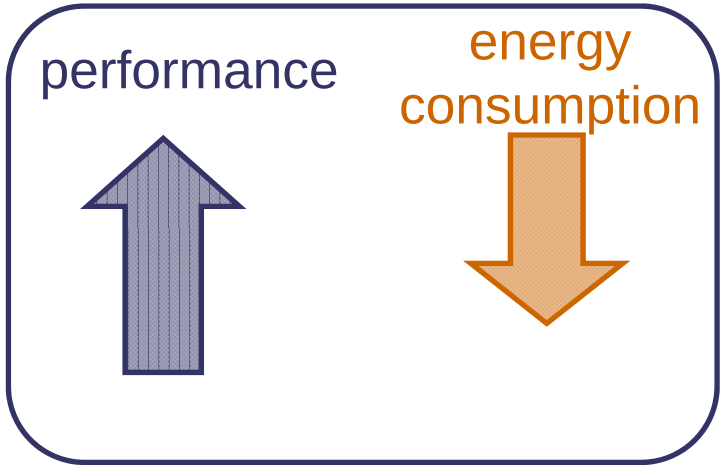

Superset Conservative
energy conscious scheme

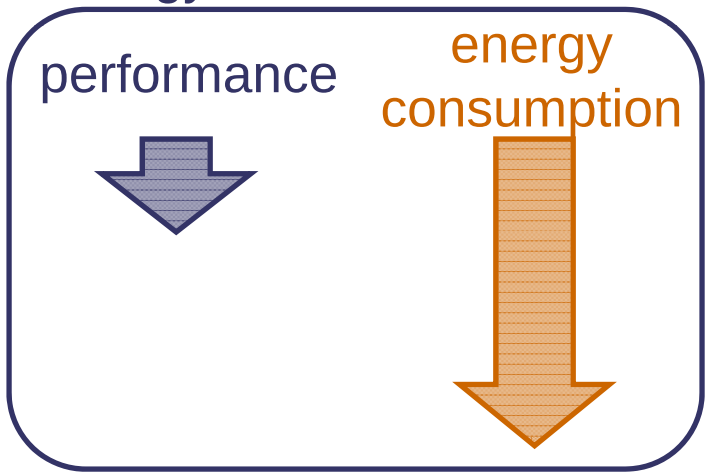

can be combined by only changing forwarding policy

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Conclusions

- proposed flexible snooping, a family of adaptive protocols for embedded rings
- two chosen protocols
  - high performance: Superset Aggressive
  - energy conservation: Superset Conservative
  - can be selected dynamically
- embedded-ring protocols more attractive

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Arch map

http://iacoma.cs.uiuc.edu/students/archmap/archmap.html

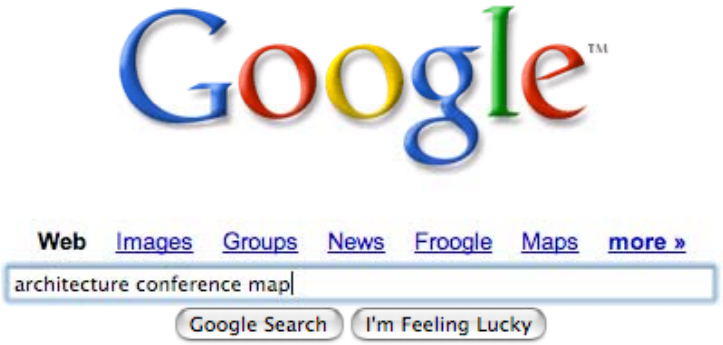

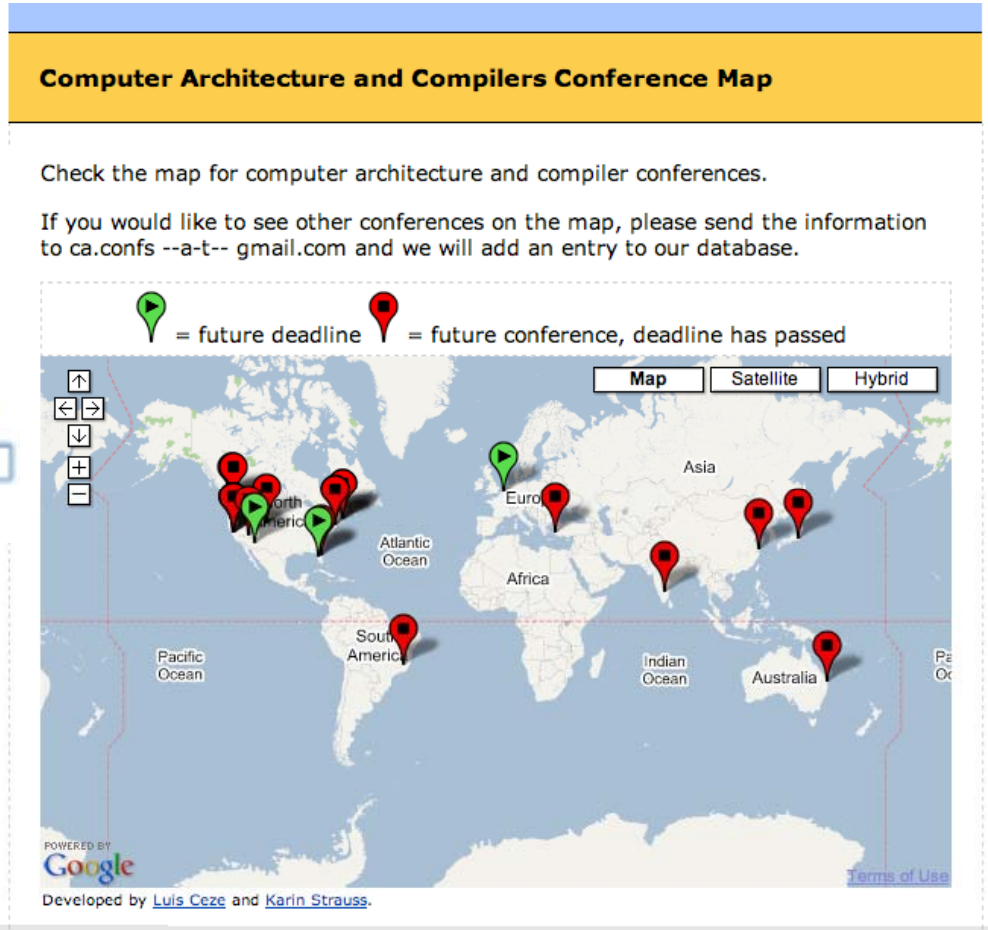

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Flexible Snooping:

Adaptive Forwarding and Filtering of Snoops in Embedded-Ring Multiprocessors

Karin Strauss, Xiaowei Shen*, Josep Torrellas

University of Illinois at Urbana-Champaign

*IBM Research

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

http://iacoma.cs.uiuc.edu