Verification of Real-Time Systems
Cache Persistence Analysis Beyond LRU

Jan Reineke

Department of Computer Science
Saarland University
Saarbrücken, Germany

Advanced Lecture, Summer 2015

Notion of Persistence

SAARLAND UNIVERSITY
COMPUTER SCIENCE

- Intuition:
  "Block *b* is *persistent* if it can only cause one cache miss in any execution."
- What is an appropriate *concrete semantics* that captures this property?

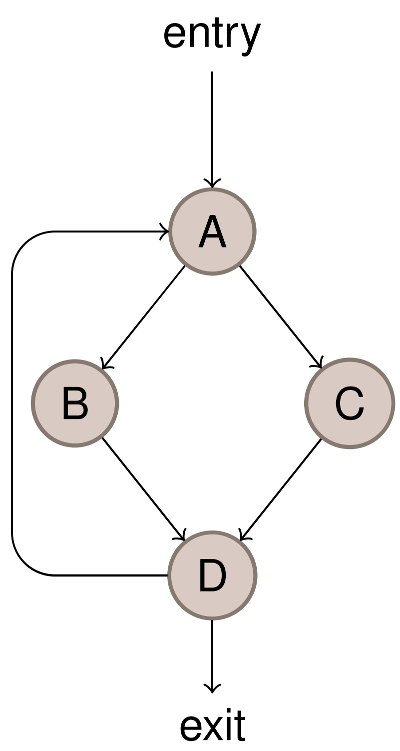

What about FIFO, MRU, PLRU, etc.?

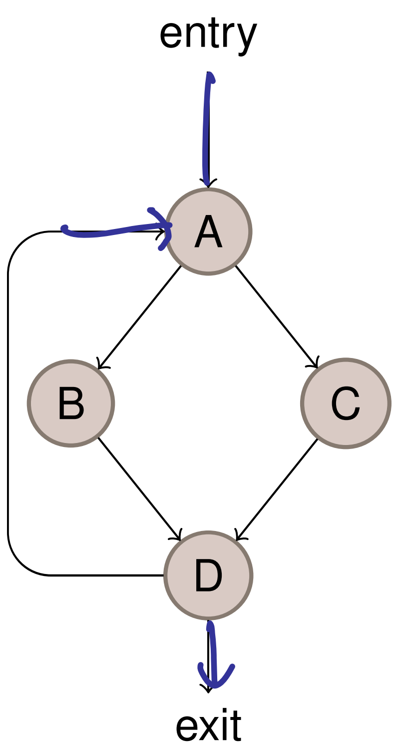

Are any of these blocks persistent under FIFO, MRU or PLRU?

For associativity 3, 4, 5...?

Block-Level Relative Competitiveness

- Competitiveness (Sleator and Tarjan, 1985):
  worst-case performance of an online policy *relative to the optimal offline policy*
  - used to evaluate online policies

- Relative competitiveness (Reineke and Grund, 2008):
  worst-case performance of an online policy *relative to another online policy*
  - used to derive local and global cache analyses

- Block-Level Relative competitiveness:
  worst-case performance of an online policy *relative to another online policy* concerning accesses to a single memory block
  - used to transfer persistence analysis results

Definition – Relative Miss-Competitiveness

Notation

$$m_{\mathbf{P}}(p, s) \quad = \quad \textit{number of misses that policy } \mathbf{P} \textit{ incurs on access sequence } s \in M^* \textit{ starting in state } p \in C^{\mathbf{P}}$$

Definition – Relative Miss-Competitiveness

Notation

$m_{\mathbf{P}}(p, s)$ = *number of misses that policy* **P** *incurs on access sequence* $s \in M^*$ *starting in state* $p \in C^{\mathbf{P}}$

Definition (Relative miss competitiveness)

Policy **P** is (k, c)-miss-competitive relative to policy **Q** if

$$m_{\mathbf{P}}(p, s) \leq k \cdot m_{\mathbf{Q}}(q, s) + c$$

for all access sequences $s \in M^*$ and cache-set states $p \in C^{\mathbf{P}}, q \in C^{\mathbf{Q}}$ that are compatible $p\ q$.

Definition – Relative Miss-Competitiveness

Notation

$m_{\mathbf{P}}(p, s) =$ *number of misses that policy* **P** *incurs on access sequence* $s \in M^*$ *starting in state* $p \in C^{\mathbf{P}}$

Definition (Relative miss competitiveness)

Policy **P** is (k, c)-miss-competitive relative to policy **Q** if

$$m_{\mathbf{P}}(p, s) \leq k \cdot m_{\mathbf{Q}}(q, s) + c$$

for all access sequences $s \in M^*$ and cache-set states $p \in C^{\mathbf{P}}, q \in C^{\mathbf{Q}}$ that are compatible p q.

Definition (Competitive miss ratio of **P** relative to **Q**)

The smallest k, s.t. **P** is (k, c)-miss-competitive rel. to **Q** for some c.

Definition – Block-Level Relative Miss-Competitiveness

Notation

$m_{\mathbf{P}}(p, s, b)$ $=$ *number of misses that policy* **P** *incurs on access sequence* $s \in M^*$ *starting in state* $p \in C^{\mathbf{P}}$ *for accesses to memory block* b

a b c d a b c

Definition – Block-Level Relative Miss-Competitiveness

Notation

$m_{\mathbf{P}}(p, s, b) \quad = \quad$ *number of misses that policy* **P** *incurs on access sequence* $s \in M^*$ *starting in state* $p \in C^{\mathbf{P}}$ *for accesses to memory block b*

Definition (Block-Level relative miss-competitiveness)

Policy **P** is (k, c)-block-level-miss-competitive relative to policy **Q** if

$$m_{\mathbf{P}}(p, s, b) \leq k \cdot m_{\mathbf{Q}}(q, s, b) + c$$

for all access sequences $s \in M^*$, cache-set states $p \in C^{\mathbf{P}}, q \in C^{\mathbf{Q}}$ that are compatible $p \sim q$, and memory blocks $b \in M$.

Consequences of Block-Level Relative Competitiveness

Allows to transfer persistence results:

- Let memory block b be persistent under LRU, and
- let policy P be (k, c)-block-level-miss-competitive relative to LRU.

$\rightarrow$ Accesses to b may cause at most $k + c$ misses under policy P.

Evaluation of Policies

SAARLAND UNIVERSITY COMPUTER SCIENCE

Computation of block-level relative competitiveness can be automated similarly to regular relative competitiveness.

Manual proofs are required to achieve parameterized results, i.e., results in terms of the associativities of the policies.

Evaluation of Policies: Results

SAARLAND UNIVERSITY
COMPUTER SCIENCE

P	Q	$c^{b,m}_{P(k),Q(l)}$	$c^{m}_{P(k),Q(l)}$

Evaluation of Policies: Results

P	Q	$c^{b,m}_{P(k),Q(l)}$	$c^{m}_{P(k),Q(l)}$
FIFO	*LRU*	∞	$\frac{k}{k-l+1}$
LRU	*FIFO*	$\left\lceil \frac{l}{k-l+1} \right\rceil$	$\left\lceil \frac{l}{k-l+1} \right\rceil$

Not useful for persistence analysis of FIFO.

Evaluation of Policies: Results

P	Q	$c^{b,m}_{P(k),Q(l)}$	$c^{m}_{P(k),Q(l)}$
MRU	LRU	$\begin{cases} 1 & : l = 2 \\ l & : l > 2 \end{cases}$ $k \geq l$	$\frac{k-1}{k-l+1}$
LRU	MRU	$\begin{cases} 1 & : k \geq 2l-2 \\ \infty & : k < 2l-2 \end{cases}$	$\max\{1, \frac{l-1}{k-l+1}\}$

A memory block that is persistent in LRU(l) will cause at most l misses in MRU(k), where $k \geq l$.

Evaluation of Policies: Results

P	Q	$c^{b,m}_{P(k),Q(l)}$	$c^{m}_{P(k),Q(l)}$
PLRU	*LRU*	$\begin{cases} 1 & : k \geq 2^{l-1} \\ \infty & : k < 2^{l-1} \end{cases}$	$\begin{cases} 1 & : k \geq 2^{l-1} \\ \infty & : k < 2^{l-1} \end{cases}$

A memory block that is persistent in LRU($\log_2 k + 1$) will also be persistent in PLRU(k).

Evaluation of Policies: Results

SAARLAND UNIVERSITY
COMPUTER SCIENCE

P	Q	$c^{b,m}_{P(k),Q(l)}$	$c^{m}_{P(k),Q(l)}$
FIFO	*LRU*	∞	$\frac{k}{k-l+1}$
LRU	*FIFO*	$\left\lceil \frac{l}{k-l+1} \right\rceil$	$\left\lceil \frac{l}{k-l+1} \right\rceil$
MRU	*LRU*	$\begin{cases} 1 & : l = 2 \\ l & : l > 2 \end{cases}$	$\frac{k-1}{k-l+1}$
LRU	*MRU*	$\begin{cases} 1 & : k \geq 2l-2 \\ \infty & : k < 2l-2 \end{cases}$	$\max\{1, \frac{l-1}{k-l+1}\}$
PLRU	*LRU*	$\begin{cases} 1 & : k \geq 2^{l-1} \\ \infty & : k < 2^{l-1} \end{cases}$	$\begin{cases} 1 & : k \geq 2^{l-1} \\ \infty & : k < 2^{l-1} \end{cases}$

Can Nothing Be Done for FIFO?

Intuitively, a large FIFO cache will hit *almost always* if a small LRU cache hits almost always.

Definition – Block-Level Relative Hit-Competitiveness

Notation

$$h_{\mathbf{P}}(p, s, b) = \text{number of hits that policy } \mathbf{P} \text{ incurs on access sequence } s \in M^* \text{ starting in state } p \in C^{\mathbf{P}} \text{ for accesses to memory block } b$$

Definition – Block-Level Relative Hit-Competitiveness

SAARLAND UNIVERSITY
COMPUTER SCIENCE

Notation

$h_{\mathbf{P}}(p, s, b) \quad = \quad$ *number of hits that policy* **P** *incurs on access sequence* $s \in M^*$ *starting in state* $p \in C^{\mathbf{P}}$ *for accesses to memory block b*

Definition (Block-Level relative hit-competitiveness)

Policy **P** is (k, c)-block-level-hit-competitive relative to policy **Q** if

$$h_{\mathbf{P}}(p, s, b) \geq k \cdot h_{\mathbf{Q}}(q, s, b) - c$$

for all access sequences $s \in M^*$, cache-set states $p \in C^{\mathbf{P}}, q \in C^{\mathbf{Q}}$ that are compatible $p\ q$, and memory blocks $b \in M$.

Evaluation of Policies: Results

SAARLAND UNIVERSITY
COMPUTER SCIENCE

$k = 4l-4$ $\qquad \left\lceil \frac{4l-3}{l-1} \right\rceil = 5$

P	Q	$c^{b,h}_{P(k),Q(l)}$	$c^{h}_{P(k),Q(l)}$
FIFO	*LRU*	$\frac{\left\lceil \frac{k}{l-1} \right\rceil - 1}{\left\lceil \frac{k}{l-1} \right\rceil}$	$\frac{\left\lceil \frac{k}{l-1} \right\rceil - 1}{\left\lceil \frac{k}{l-1} \right\rceil}$
LRU	*FIFO*	$\begin{cases} 1 & : k \geq 2l-1 \\ 0 & : k < 2l-1 \end{cases}$	$\begin{cases} 1 & : k \geq 2l-1 \\ 0 & : k < 2l-1 \end{cases}$

E.g. 4/5 = 80% of all hits of LRU(k) must also be hits of FIFO(4k-3)

Summary

Block-level Relative Competitiveness

- Refinement of relative competitiveness.
- Allows to transfer persistence results into bounds on the number of hits/misses for other policies.
  Useful for MRU, FIFO, and PLRU.