

CLOCK-Pro+: Improving CLOCK-Pro Cache Replacement with Utility-Driven Adaptation

Cong Li, Intel Corporation


Outline

- **Introduction: Cache & Page Replacement**
- **Background: CLOCK-Pro & CLOCK for Adaptive Replacement**
- **The New Policy w/ Utility-Driven Adaptation: CLOCK-Pro+**
- **Experimental Results**
- **Conclusion**

Introduction

- **Buffer Cache Replacement**
 - Determine the victim to be replaced given a new data block to be loaded
 - Many policies proposed, e.g., LRU, ARC, LIRS, etc.
- **CLOCK**
 - Data manipulation w/ a hit → lock contention problem in low hit latency scenario
 - ✓ Page replacement in virtual memory management

CLOCK



✓ Referenced

CLOCK-Pro


- **Reuse Distance**

- Distance of a referenced page away from the top
- Page w/ a low reuse distance → more likely to be accessed in the future

- **CLOCK-Pro**


- Efficiently discriminate hot pages (low reuse distances) from cold pages (high reuse distances)
 - ✓ Approximating LIRS policy
 - ✓ Adapting to LRU-friendly workloads

CLOCK-Pro




- Hot page
- Resident cold page
- Non-resident cold page
- Referenced

CLOCK-Pro




CLOCK-Pro



Cold page promotion &
hot page demotion

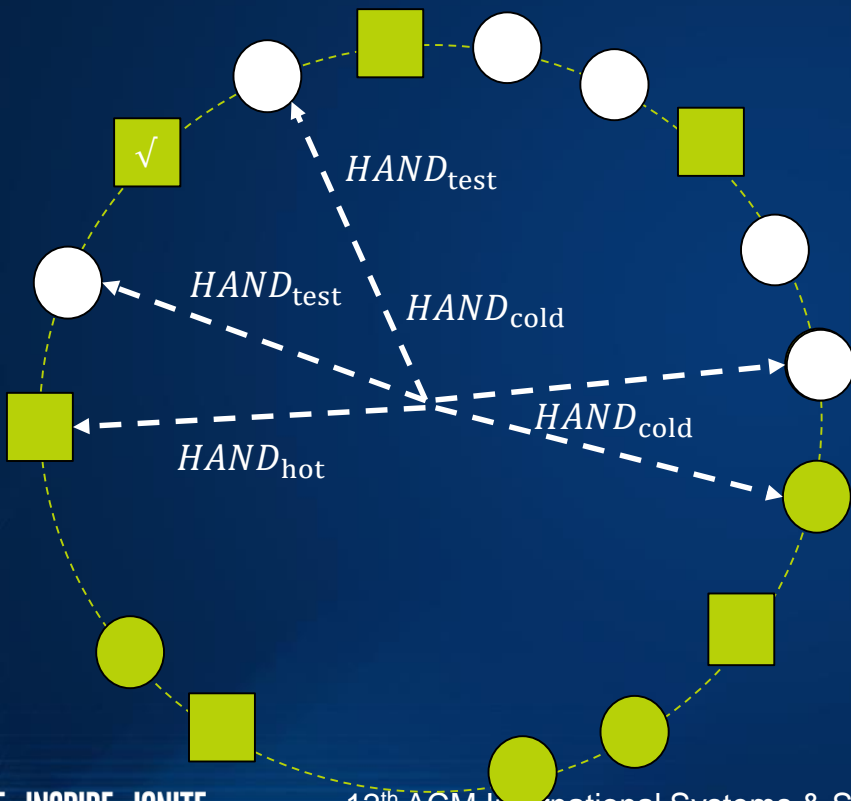
CLOCK-Pro



$HAND_{hot}$ & $HAND_{test}$ move

Test period terminates & non-resident page discarded

CLOCK-Pro



Many new pages come

Limit clock size by terminating
test pages with $HAND_{test}$

Weakness w/o Adaptation

- **Static Cache Space Allocation**
 - Small number of resident cold pages close to head position
 - Non-resident cold pages interleaved w/ hot pages
- **When Reuse Distance Is not a Good Predictor (or does not Exist)**
 - Frequent accesses to close-to-head non-resident cold pages result in misses
 - ✓ Can be captured with a basic CLOCK policy
 - ✓ Example: stack depth distribution (SDD) workload

CLOCK-Pro w/o adaptation is not good enough

CLOCK-Pro w/ Adaptation

- **Idea**
 - Cold page access → LRU friendly
 - Test period expiration → need more hot pages to extend test period
- **Issue**
 - Simple heuristics w/o utility analysis, e.g.,
 - ✓ Resident cold page accesses → not necessary to increase cold page number
 - ✓ Many test pages expire → more hot pages may not help

CLOCK-Pro w/ adaptation is still not good enough

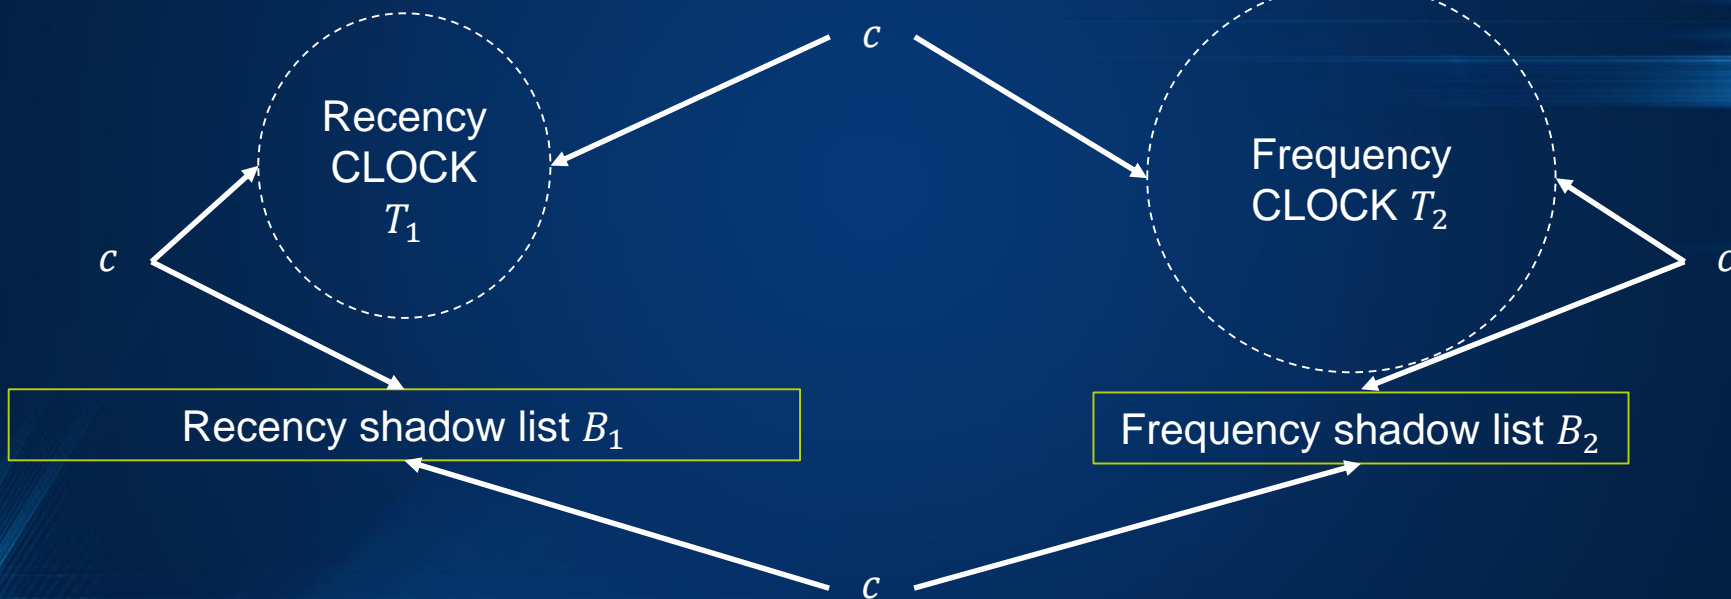
CLOCK w/ Adaptive Replacement (CAR)

- **Recency vs. Frequency**
 - Varying & requiring dynamic adaptation
- **CAR (Approximation of ARC)**
 - Maintain 2 different CLOCKS & 2 different shadow lists
 - ✓ 1 CLOCK & 1 shadow list for recency (1 recent access)
 - ✓ 1 CLOCK & 1 shadow list for frequency (at least 2 recent accesses)
 - Utility-driven adaptation to dynamically adjust the 2 CLOCKS

CAR

Recency pages: pages w/ 1 recent accesses only

Frequency pages: pages w/ at least 2 recent accesses



CAR

Recency pages: pages w/ 1 recent accesses only




Recency shadow list B_1



Access recency shadow list \rightarrow growing T_1 Access frequency shadow list \rightarrow growing T_2
Incremental utility quantified as $P_1 = 1/|B_1|$ Incremental utility quantified as $P_2 = 1/|B_2|$

Frequency pages: pages w/ at least 2 recent accesses




Frequency shadow list B_2



CAR

Recency pages: pages w/ 1 recent accesses only




Recency shadow list B_1

Adjustment given a B_1 access:

$$|T_1| \leftarrow |T_1| + \max\{1, P_1 / P_2\}$$

Frequency pages: pages w/ at least 2 recent accesses



Frequency shadow list B_2

Adjustment given a B_2 access:

$$|T_2| \leftarrow |T_2| + \max\{1, P_2 / P_1\}$$

CAR (cont.)

- **Frequency CLOCK & Shadow List**
 - Contain less granular information
- **Without a Fine-Grained Metric like Reuse Distance**
 - Less capable in capturing repeated accesses w/ relatively long temporal distances (weak locality)

CAR is not good enough as well

CLOCK-Pro vs CAR (a Glance)

Trace (cache size)	CLOCK-Pro	CAR
WebSearch1 (131072)	13.10%	8.32%
WebSearch1 (262144)	24.91%	14.90%
WebSearch1 (524288)	40.36%	32.78%
WebSearch2 (262144)	29.80%	26.94%
WebSearch2 (524288)	48.35%	41.72%
WebSearch3 (262144)	29.66%	26.68%
WebSearch3 (524288)	48.21%	41.40%
Financial1 (512)	17.78%	23.17%
Financial1 (1024)	20.62%	26.02%
Financial1 (2048)	24.16%	29.38%
Financial1 (4096)	27.58%	32.61%
Financial1 (8192)	31.31%	35.72%
No consistent winner		
SDD (256)	17.10%	20.40%
SDD (512)	31.60%	36.75%

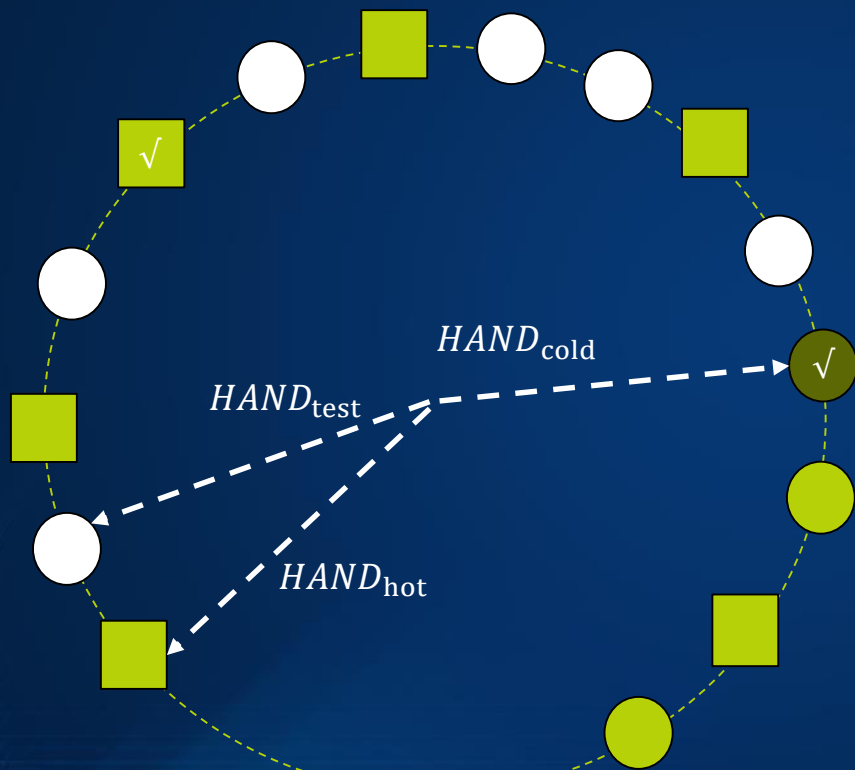
CLOCK-Pro
outperforms CAR

CAR
outperforms
CLOCK-Pro

Idea of CLOCK-Pro+

- **Idea Inspired by CAR**
 - Dynamic adaptation in CLOCK-Pro using a CAR-style utility evaluation
 - ✓ When reuse distance is a good predictor, more space allocated to hot pages
 - ✓ When reuse distance is not a good predictor, more space allocated to cold pages
- **Determining Predictor Goodness**
 - Accessing non-resident cold pages
 - Inappropriately demoting hot pages (hit shortly after demotion)

Adaptation in CLOCK-Pro+




● Resident cold pages demoted from hot pages

C_n : current number of non-resident pages


C_d : current number of resident cold pages demoted from hot pages

Adaptation in CLOCK-Pro+



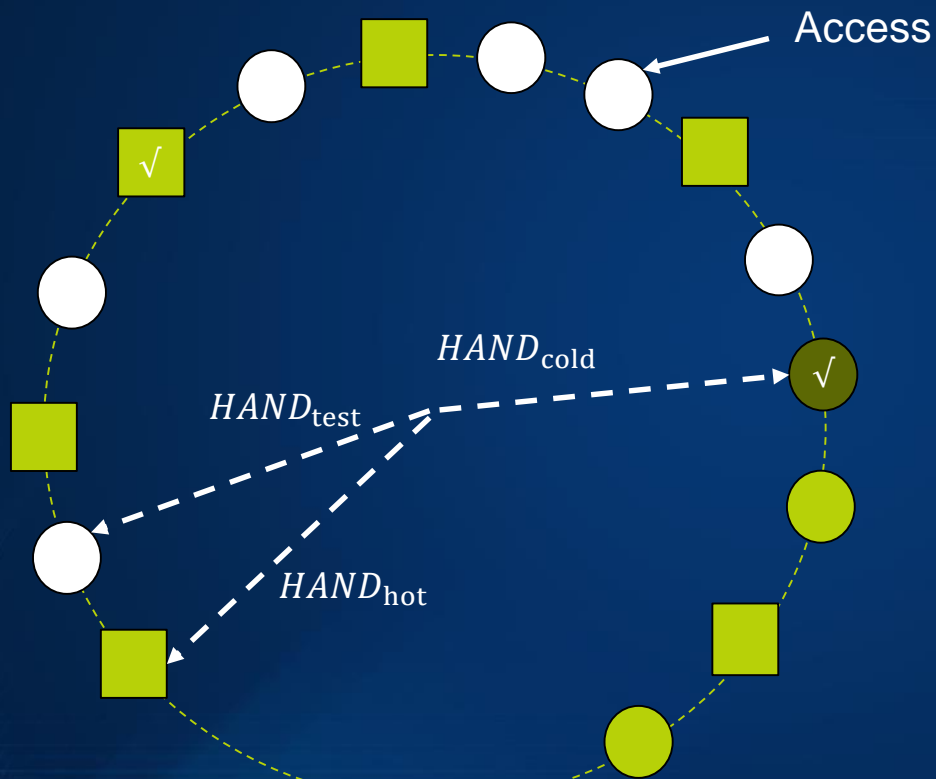
Grow resident cold page size
Utility quantified as $P_{\bar{n}} = 1/C_n$

Adaptation in CLOCK-Pro+




Grow hot page size
Utility quantified as $P_{\bar{d}} = 1/C_d$

Adaptation in CLOCK-Pro+



Grow resident cold page size by $\max\{1, P_{\bar{n}}/P_{\bar{d}}\}$

Adaptation in CLOCK-Pro+



Grow hot page size by
 $\max\{1, P_{\bar{a}}/P_{\bar{n}}\}$

Experimental settings

- **Trace-Driven Simulation**
 - I/O traces from UMass Trace Repository
 - Synthetic trace drawn from a stack depth distribution
 - Cache size varies, & shadow entry number = cache entry number
- **Comparative Study on Hit Ratio**
 - CLOCK-Pro
 - CAR
 - CLOCK-Pro+

Experimental results

Retain CLOCK-Pro's
strength

Trace (cache size)	CLOCK-Pro	CAR	CLOCK-Pro+
WebSearch1 (131072)	13.10%	8.32%	12.96%
WebSearch1 (262144)	24.91%	14.90%	24.80%
WebSearch1 (524288)	40.36%	32.78%	41.66%
WebSearch2 (262144)	29.80%	26.94%	29.64%
WebSearch2 (524288)	48.35%	41.72%	48.50%
WebSearch3 (262144)	29.66%	26.68%	29.52%
WebSearch3 (524288)	48.21%	41.40%	48.41%
Financial1 (512)	17.78%	23.17%	22.69%
Financial1 (1024)	20.62%	26.02%	25.77%
Financial1 (2048)	24.16%	29.38%	29.15%
Financial1 (4096)	27.58%	32.61%	32.35%
Financial1 (8192)	31.31%	35.72%	35.65%
Financial1 (16384)	34.33%	38.35%	38.31%
SDD (256)	17.10%	20.40%	19.34%
SDD (512)	31.60%	36.75%	35.06%

Experimental results

Overcome CLOCK-Pro's weaknesses, bringing its performance close to CAR

Trace (cache size)	CLOCK-Pro	CAR	CLOCK-Pro+
WebSearch1 (131072)	13.10%	8.32%	12.96%
WebSearch1 (262144)	24.91%	14.90%	24.80%
WebSearch1 (524288)	40.36%	32.78%	41.66%
WebSearch2 (262144)	29.80%	26.94%	29.64%
WebSearch2 (524288)	48.35%	41.72%	48.50%
WebSearch3 (262144)	29.66%	26.68%	29.52%
WebSearch3 (524288)	48.21%	41.40%	48.41%
Financial1 (512)	17.78%	23.17%	22.69%
Financial1 (1024)	20.62%	26.02%	25.77%
Financial1 (2048)	24.16%	29.38%	29.15%
Financial1 (4096)	27.58%	32.61%	32.35%
Financial1 (8192)	31.31%	35.72%	35.65%


CLOCK-Pro+ performs close to the winner between the two

SDD (256)	17.10%	20.40%	19.34%
SDD (512)	31.60%	36.75%	35.06%

Conclusion

- **Novel Improvement to CLOCK-Pro's Adaptation**
 - Borrowing idea from CAR
 - Utility-driven adaptation of cache space allocation
- **CLOCK-Pro+**
 - Enjoy the strengths of CLOCK-Pro & CAR
 - Overcome the weaknesses of CLOCK-Pro & CAR
 - Perform consistently close to the winner between the two

Q & A



Ablation Study

Sometimes CLOCK-Pro improves the performance

Trace (cache size)	CLOCK-LIRS ¹	CLOCK-Pro	CLOCK-Pro+
Financial1 (512)	15.80%	17.78%	22.69%
Financial1 (1024)	19.42%	20.62%	25.77%
Financial1 (2048)	25.36%	24.16%	29.15%
Financial1 (4096)	30.51%	27.58%	32.35%
Financial1 (8192)	34.24%	31.31%	35.65%
Financial1 (16384)	37.08%	34.33%	38.31%
SDD (256)	17.00%	17.10%	19.34%
SDD (512)	30.95%	31.60%	35.06%
SDD (1024)	51.55%	58.08%	58.07%

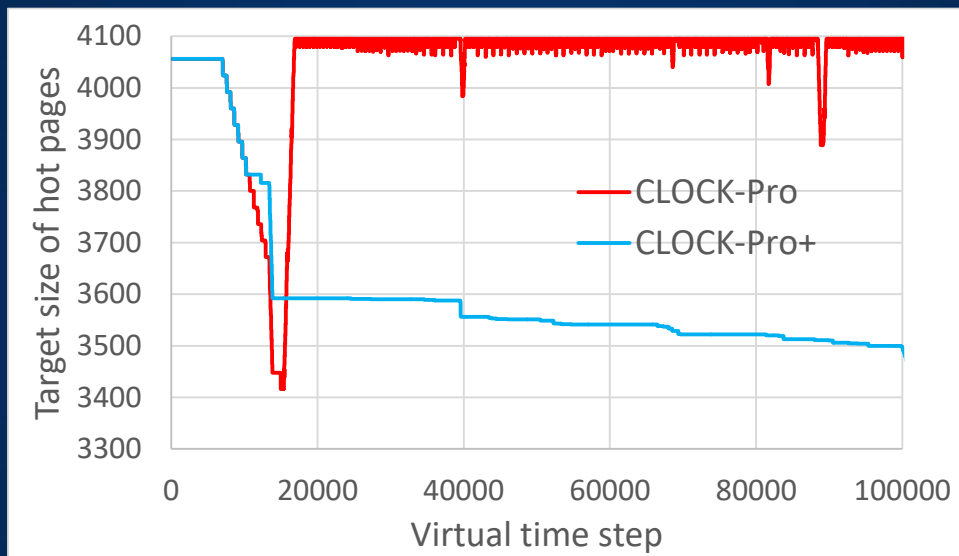
¹ CLOCK-Pro w/o adaptation

Sometimes it does not

CLOCK-Pro+ consistently improves the performance

CLOCK-Pro performs unstably but CLOCK-Pro+ performs consistently

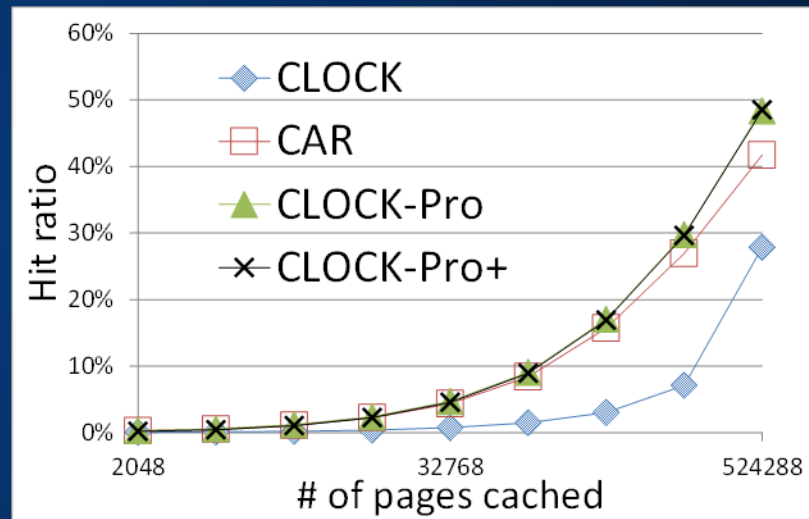
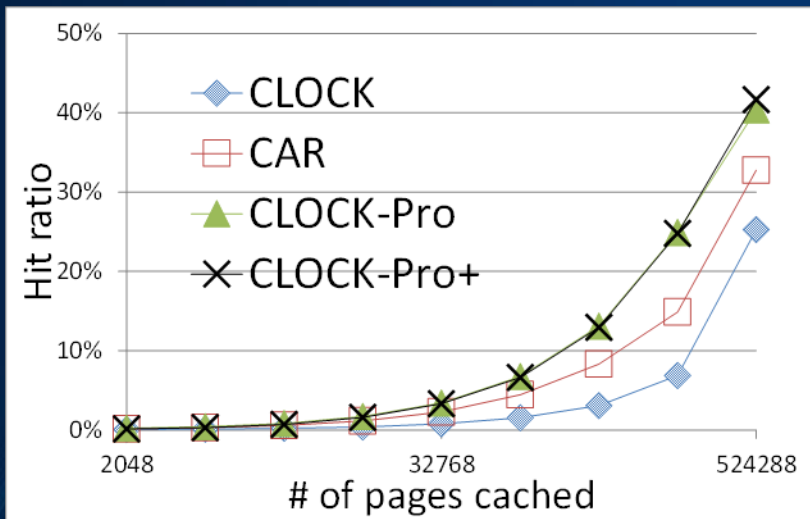
Case Study: Financial1 (4096)





CLOCK-Pro: 382,543 non-resident cold page accesses, 111,244 resident cold page hits tracked, but **3,143,452** test pages expired;

CLOCK-Pro+: 102,804 non-resident cold page accesses & 3,780 demoted page hits

Full Results: WebSearch1 & Webserach2



Full Results: WebSearch3 & Financial1



Full Results: Financial2 & SDD

