Beyond Worst-case Analysis of Multicore Caching Strategies



Shahin Kamali



APOCS 2021



Helen Xu

QJPSE07, SRD04], but our theoretical understanding lags behind.

Caching for multicore architectures has been well-studied in practice [FSS06,



Caching for multicore architectures has been well-studied in practice [FSS06, QJPSE07, SRD04], but our theoretical understanding lags behind.

Multicore caching diverges from classical single-core caching because it has an added **scheduling aspect** due to fetch delay [LS12, KX20].

З



Shared Cache

QJPSE07, SRD04], but our theoretical understanding lags behind.

an added scheduling aspect due to fetch delay [LS12, KX20].

Furthest-In-Future (FIF) is the optimal algorithm [B66] and Least-Recently-Used (LRU) is k-competitive (where k is the size of the cache) [ST85] in single-core.

- Caching for multicore architectures has been well-studied in practice [FSS06,
- Multicore caching diverges from classical single-core caching because it has
 - **Shared Cache**



QJPSE07, SRD04], but our theoretical understanding lags behind.

an added scheduling aspect due to fetch delay [LS12, KX20].

Furthest-In-Future (FIF) is the optimal algorithm [B66] and cache) [ST85] in single-core.

in multicore caching [KX20].

- Caching for multicore architectures has been well-studied in practice [FSS06,
- Multicore caching diverges from classical single-core caching because it has
 - **Shared Cache** Least-Recently-Used (LRU) is k-competitive (where k is the size of the

A large class of "lazy algorithms" (including LRU) cannot be competitive



Relaxing Competitive Analysis

Resource augmentation [ST85] relaxes strict worst-case (competitive) analysis by giving online algorithms additional resources.





Online algorithm

Relaxing Competitive Analysis

Resource augmentation [ST85] relaxes strict worst-case (competitive) analysis by giving online algorithms additional resources.

In single-core caching, doubling the space makes LRU 2-competitive. The same result **does not hold** in multicore [LS12].





Relaxing Competitive Analysis

Resource augmentation [ST85] relaxes strict worst-case (competitive) analysis by giving online algorithms additional resources.

same result does not hold in multicore [LS12].



- In single-core caching, doubling the space makes LRU 2-competitive. The
- In multicore caching, a constant factor in resource augmentation yields an **O(log p)-competitive** algorithm (where p is the number of cores) [ABDKPS20].



Many alternative measures have been proposed for single-core caching but have not been studied in the **parallel setting** [Y94, Y98, KP00, KPR92, BLN01, BFL05, BB94].

Many alternative measures have been proposed for single-core caching but have not been studied in the **parallel setting** [Y94, Y98, KP00, KPR92, BLN01, BFL05, BB94].

Bijective analysis [ADL07, ADL08, AS09, AS13] directly compares online algorithms and has been used to separate LRU in the single-core setting from other online algorithms.

Many alternative measures have been proposed for single-core caching but have not been studied in the **parallel setting** [Y94, Y98, KP00, KPR92, BLN01, BFL05, BB94].

Bijective analysis [ADL07, ADL08, AS09, AS13] directly compares online algorithms and has been used to separate LRU in the single-core setting from other online algorithms.

Overly pessimistic

Bijective analysis

Competitive analysis



Many alternative measures have been proposed for single-core caching but have not been studied in the **parallel setting** [Y94, Y98, KP00, KPR92, BLN01, BFL05, BB94].

Bijective analysis [ADL07, ADL08, AS09, AS13] directly compares online algorithms and has been used to separate LRU in the single-core setting from other online algorithms.

Overly pessimistic Compares online algorithms directly

Bijective analysis

Competitive analysis







Many alternative measures have been proposed for single-core caching but have not been studied in the **parallel setting** [Y94, Y98, KP00, KPR92, BLN01, BFL05, BB94].

Bijective analysis [ADL07, ADL08, AS09, AS13] directly compares online algorithms and has been used to separate LRU in the single-core setting from other online algorithms.

Overly pessimistic

Bijective analysis

Competitive analysis



Compares online algorithms directly

Applicability









Cyclic Analysis for Multicore Caching

We introduce **cyclic analysis**, an alternative measure for online algorithms inspired by bijective analysis.

Cyclic Analysis for Multicore Caching

We introduce **cyclic analysis**, an alternative measure for online algorithms inspired by bijective analysis.

Bijective analysis compares inputs of the same "length", while cyclic analysis compares algorithms **across all inputs** rather than within partitions.



Bijective analysis



Cyclic analysis

Cyclic Analysis Measure

Applicability

Bijective analysis



Cyclic analysis



Cyclic Analysis Measure

Applicability

Bijective analysis



Cyclic analysis



Compares online algorithms directly





Cyclic Analysis Measure

Applicability

Bijective analysis



Cyclic analysis



Compares online algorithms directly





Used to separate LRU





Separating LRU With Cyclic Analysis

analysis on the universe of inputs I if there exists a bijection $\pi: I \leftrightarrow I$ satisfying $A(R) \leq B(\pi(R))$ for all $R \in I$. We denote this by $A \leq_c B$.

Algorithm cost

- An online algorithm A is no worse than online algorithm B under cyclic

Separating LRU With Cyclic Analysis

An online algorithm A is **no worse** than online algorithm B under **cyclic analysis** on the universe of inputs *I* if there exists a bijection $\pi: I \leftrightarrow I$ satisfying $A(R) \leq B(\pi(R))$ for all $R \in I$. We denote this by $A \leq_c B$.

Algorithm cost

Main proof



Separating LRU With Cyclic Analysis

An online algorithm A is **no worse** than online algorithm B under **cyclic analysis** on the universe of inputs *I* if there exists a bijection $\pi: I \leftrightarrow I$

Algorithm cost

- satisfying $A(R) \leq B(\pi(R))$ for all $R \in I$. We denote this by $A \leq_c B$. Main
 - Main theorem: For any lazy caching algorithm A, $LRU \leq_{c} A$ on inputs with locality of reference.



Given algorithms A, B and a set of inputs I:



Input-cost graph: modeling mappings between inputs with different costs.

Given algorithms A, B and a set of inputs I:



Goal: Define a bijection $\pi: I \leftrightarrow I$ such that for all inputs $R \in I, A(R) \leq B(\pi(R))$.

Input-cost graph: modeling mappings between inputs with different costs.

Given algorithms A, B and a set of inputs I:



Analysis technique: Partition the input-cost graph and analyze cycles in the resulting subgraphs.

Input-cost graph: modeling mappings between inputs with different costs.

Goal: Define a bijection $\pi : I \leftrightarrow I$ such that for all inputs $R \in I, A(R) \leq B(\pi(R))$.

Partitioning the input-cost graph: For problems such as single-core caching, the length of inputs is well defined.

For general online problems (e.g. multicore caching), the length is not a measure of difficulty.

Partitioning the input-cost graph: For problems such as single-core caching, the length of inputs is well defined.

For general online problems (e.g. multicore caching), the length is not a measure of difficulty.

> **Bijective analysis** partitions the input-cost graph into finite subgraphs.

Bijective analysis



Cyclic analysis

Example: Bijective vs Cyclic Analysis

Given algorithms A, B for some online problem P.

$$A(R_1) = 10 \qquad B(R_1)$$
$$A(R_2) = 40 \qquad B(R_2)$$
$$A(R_3) = 20 \qquad B(R_3)$$
$$B(R_4) = 30 \qquad B(R_4)$$



Example: Bijective vs Cyclic Analysis

Given algorithms A, B for some online problem P.

$$A(R_1) = 10 \qquad B(R_1)$$
$$A(R_2) = 40 \qquad B(R_2)$$
$$B(R_3) = 20 \qquad B(R_3)$$
$$A(R_4) = 30 \qquad B(R_3)$$



Example: Bijective vs Cyclic Analysis

Given algorithms A, B for some online problem P.

$$A(R_{1}) = 10 \longrightarrow B(R_{1})$$

$$A(R_{2}) = 40 \longrightarrow B(R_{2})$$

$$A(R_{3}) = 20 \longrightarrow B(R_{3})$$

$$A(R_{4}) = 30 \longrightarrow B(R_{4})$$

$$B(R_{4})$$



Properties of Cyclic Analysis

The cyclic analysis measure has several **natural properties**:

Transitivity ($A \leq_{c} B, B \leq_{c} C \rightarrow A \leq_{c} C$).

Directedness ($A <_{c} B$ implies $B \not\leq_{c} A$).

Properties of Cyclic Analysis

The cyclic analysis measure has several **natural properties**:

Transitivity ($A \leq_{c} B, B \leq_{c} C \rightarrow A \leq_{c} B$

Directedness ($A <_{c} B$ implies $B \not\leq_{c} A$).

Cyclic analysis can be used to show "to-be-expected" relationships between multicore caching algorithms:

All lazy algorithms are equivalent.

$$\leq_{c} C$$
).

Any lazy algorithm A is strictly better than Flush-When-Full (for p=2).

Problem setup: given a multicore processor with p cores and k pages.

page requests, where each core must serve its corresponding request sequence.

Page requests arrive at discrete timesteps.

An input *R* is formed by *p* request sequences $R = R_1, ..., R_p$ composed of

Problem setup: given a multicore processor with p cores and k pages.

page requests, where each core must serve its corresponding request sequence.

Page requests arrive at discrete timesteps.



An input *R* is formed by *p* request sequences $R = R_1, ..., R_p$ composed of

Cache: *ab*

Problem setup: given a multicore processor with p cores and k pages.

page requests, where each core must serve its corresponding request sequence.

Page requests arrive at discrete timesteps.



An input *R* is formed by *p* request sequences $R = R_1, ..., R_p$ composed of

Cache: *ab*

Problem setup: given a multicore processor with p cores and k pages.

page requests, where each core must serve its corresponding request sequence.

Page requests arrive at discrete timesteps.



An input *R* is formed by *p* request sequences $R = R_1, \ldots, R_p$ composed of

Locality of Reference

$LRU \leq_{c} A$ on inputs with locality of reference.

Intuitively: an input has locality of reference if the number of distinct pages requested is bounded [AFG02].

Main theorem: For any lazy caching algorithm A,

Locality of Reference

Intuitively: an input has locality of reference if the number of distinct pages requested is bounded [AFG02].

abcde... No locality

Main theorem: For any lazy caching algorithm A,

 $LRU \leq_{c} A$ on inputs with locality of reference.



Advantage of LRU with Locality

Goal: For any lazy algorithm *A*, define a bijection π over the space of inputs with locality *I* satisfying $LRU(R) \leq A(\pi(R))$ for all $R \in I$.

Intuition: Analyze timesteps where A and LRU differ. Suppose at time t, LRU evicts page σ_{LRU} , while A evicts σ_{NLRU} .

Advantage of LRU with Locality

with locality I satisfying $LRU(R) \leq A(\pi(R))$ for all $R \in I$.

evicts page σ_{LRU} , while A evicts σ_{NLRU} .

Case 1: Swapping σ_{LRU} , σ_{NLRU} in the continuation maintains locality. $\pi(R)$ swaps them in the continuation. $LRU(R) = A(\pi(R))$.

- **Goal:** For any lazy algorithm A, define a bijection π over the space of inputs
- Intuition: Analyze timesteps where A and LRU differ. Suppose at time t, LRU

Advantage of LRU with Locality

Goal: For any lazy algorithm *A*, define a bijection π over the space of inputs with locality *I* satisfying $LRU(R) \leq A(\pi(R))$ for all $R \in I$.

Intuition: Analyze timesteps where A and LRU differ. Suppose at time t, LRU evicts page σ_{LRU} , while A evicts σ_{NLRU} .

Case 1: Swapping σ_{LRU} , σ_{NLRU} in the continuation maintains locality. $\pi(R)$ swaps them in the continuation. $LRU(R) = A(\pi(R))$.

Case 2: Swapping σ_{LRU} , σ_{NLRU} in the continuation does not maintain locality. To show: $LRU(R) \leq A(\pi(R))$

Constructing a Mapping

Case 2: Swapping σ_{LRU} , σ_{NLRU} in the continuation does not maintain locality. To show: $LRU(R) \leq A(\pi(R))$ If there is a miss before a request to σ_{NLRU} after time *t*, set $\pi(R) = R$. Therefore, $LRU(R) = A(\pi(R))$.





$$LRU(R) = A(\pi(R))$$

The previous two cases cover of



entire codomain, but not the domain.



$$LRU(R) = A(\pi(R))$$
V
The previous two cases cover

For remaining inputs, map them arbitrarily to inputs of higher cost.



entire codomain, but not the domain.

 $LRU(R) < A(\pi(R))$



$$LRU(R) = A(\pi(R))$$
V
The previous two cases cover

For remaining inputs, map them arbitrarily to inputs of higher cost.

Equivalence with bijective mapping



entire codomain, but not the domain.

 $LRU(R) < A(\pi(R))$



$$LRU(R) = A(\pi(R))$$
V
The previous two cases cover

For remaining inputs, map them arbitrarily to inputs of higher cost.

Equivalence with bijective mapping



entire codomain, but not the domain.

 $LRU(R) < A(\pi(R))$

Cyclic analysis compares online algorithms directly on a set of inputs and establishes the **advantage of LRU** on inputs with locality.

establishes the advantage of LRU on inputs with locality.

be useful in the study of other online problems.

- Cyclic analysis compares online algorithms directly on a set of inputs and
 - Since it is more general and flexible than bijective analysis, we expect it to

establishes the advantage of LRU on inputs with locality.

be useful in the study of other online problems.

memory.

- Cyclic analysis compares online algorithms directly on a set of inputs and
 - Since it is more general and flexible than bijective analysis, we expect it to
- Multicore caching is well-studied in practice and motivated by hierarchical

establishes the advantage of LRU on inputs with locality.

be useful in the study of other online problems.

memory.

The need for theoretical understanding will only grow as multicore architectures become more prevalent.

- Cyclic analysis compares online algorithms directly on a set of inputs and
 - Since it is more general and flexible than bijective analysis, we expect it to
- Multicore caching is well-studied in practice and motivated by hierarchical