

Why is Dual-Pivot Quicksort Fast?

Sebastian Wild


wild@cs.uni-kl.de



29 September 2015

Theorietage 2015 Speyer

Sorting History



Sorting History

1961,62 Hoare: publication, first analysis

1969 Singleton: median-of-three & Insertionsort on small subarrays

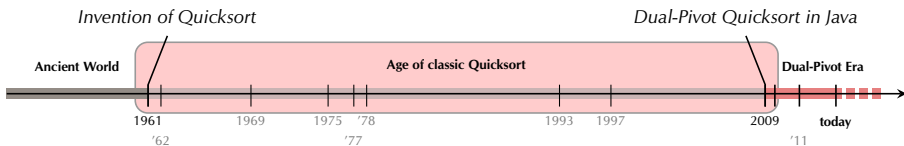
1975-78 Sedgewick: analysis of many optimizations

1993 Bentley, McIlroy: duplicate elements & “ninter”

1997 Musser: $\mathcal{O}(n \log n)$ worst case by truncating recursion

↪ Basic algorithm settled since 1961; latest tweaks from 1990's.

Since then: Almost identical in all programming libraries!



Sorting History

1961,62 Hoare: publication, first analysis

1969 Singleton: median-of-three & Insertionsort on small subarrays

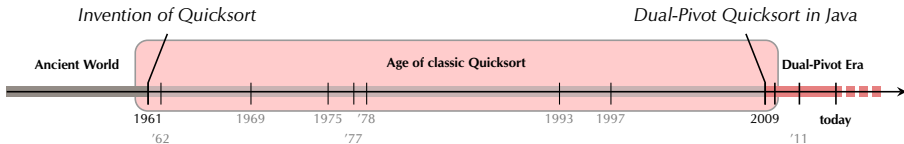
1975-78 Sedgewick: analysis of many optimizations

1993 Bentley, McIlroy: duplicate elements & “ninth”

1997 Musser: $\Theta(n \log n)$ worst case by truncating recursion

↪ Basic algorithm settled since 1961; latest tweaks from 1990's.

Since then: Almost identical in all programming libraries!



Sorting History

2008 – 2009 **Vladimir Yaroslavskiy** (developer at Sun)


experiments with Quicksort with **two pivots**

11 Sep 2009 announcement on Java core library mailing list

29 Oct 2009 **inclusion** in development version of OpenJDK

2009 – 2011 optimizations by Joshua Bloch, Jon Bentley and others

28 July 2011 **public release** of Java 7 with Yaroslavskiy's Quicksort




Dual-Pivot Quicksort

Algorithm (Conceptual View)

- 1 Choose **two pivots** $P \leq Q$
- 2 For each element x , determine its **class**
 - **small** for $x < P$
 - **medium** for $P < x < Q$
 - **large** for $Q < x$

by comparing x to **pivots** P and Q

- 3 Arrange elements according to classes:



- 4 Sort subarrays recursively.

How to implement 1 efficiently on arrays?

Dual-Pivot Quicksort

Algorithm (Conceptual View)

- 1 Choose **two pivots** $P \leq Q$
- 2 For each element x , determine its **class**
 - **small** for $x < P$
 - **medium** for $P < x < Q$
 - **large** for $Q < x$

by comparing x to **pivots** P and Q

- 3 Arrange elements according to classes:



- 4 Sort subarrays recursively.

How to implement 3 efficiently on arrays?

Dual-Pivot Quicksort

Algorithm (Conceptual View)

- 1 Choose **two pivots** $P \leq Q$
- 2 For each element x , determine its **class**
 - **small** for $x < P$
 - **medium** for $P < x < Q$
 - **large** for $Q < x$

by comparing x to **pivots** P and Q

- 3 Arrange elements according to classes:



- 4 Sort subarrays recursively.

How to implement 3 efficiently on arrays?

Dual-Pivot Quicksort

Algorithm (Conceptual View)

- 1 Choose **two pivots** $P \leq Q$
- 2 For each element x , determine its **class**
 - **small** for $x < P$
 - **medium** for $P < x < Q$
 - **large** for $Q < x$

by comparing x to **pivots** P and Q



- 3 Arrange elements according to classes:




- 4 Sort subarrays recursively.

How to implement 3 efficiently on arrays?




Yaroslavskiy's Algorithm



Invariant:






Yaroslavskiy's Algorithm




Invariant:






Yaroslavskiy's Algorithm



Invariant:





Yaroslavskiy's Algorithm



Invariant:





Yaroslavskiy's Algorithm



Invariant:





Yaroslavskiy's Algorithm



Invariant:






Yaroslavskiy's Algorithm




Invariant:






Yaroslavskiy's Algorithm



Invariant:






Yaroslavskiy's Algorithm




Invariant:






Yaroslavskiy's Algorithm



Invariant:






Yaroslavskiy's Algorithm



Invariant:






Yaroslavskiy's Algorithm



Invariant:






Yaroslavskiy's Algorithm



Invariant:






Yaroslavskiy's Algorithm



Invariant:






Yaroslavskiy's Algorithm



Invariant:






Yaroslavskiy's Algorithm




Invariant:






Yaroslavskiy's Algorithm




Invariant:






Yaroslavskiy's Algorithm




Invariant:





Yaroslavskiy's Algorithm





Invariant:





Yaroslavskiy's Algorithm



Yaroslavskiy's Algorithm




Yaroslavskiy's Algorithm




Running Time Experiments

Why switch to new, unknown algorithm?




Running Time Experiments

Why switch to new, unknown algorithm? Because it is faster!



Running Time Experiments


Why switch to new, unknown algorithm? Because it is faster!



- remains true for **basic** variants of algorithms: -●- vs. -□-!

Running Time Experiments

Why switch to new, unknown algorithm? Because it is faster!




No theoretical explanation for running time known in 2009!

● remain

Running Time Experiments

Why switch to new, unknown algorithm? Because it is faster!




No theoretical explanation for running time known in 2009!

Only lucky experiments?

● remain

Running Time Experiments

Why switch to new, unknown algorithm? Because it is faster!



No theoretical explanation for running time known in 2009!

Only lucky experiments?

Why did noone come up with this earlier?

● remain




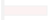




Track Record of Dual-Pivot Quicksort

- **Observation in practice:**

Yaroslavskiy's Quicksort (YQS) **faster** than classic Quicksort (CQS) ... **why?**

→ We did a mathematical analysis of YQS.

- Traditional cost measures do **not** explain observation!

	CQS	YQS	Relative
Running Time	(from various experiments)		 -10±2%
Comparisons	2	1.9	 -5%
Swaps	0.3	0.6	+80% 
Bytecode Instructions	18	21.7	+20.6% 
MMIX oops ν	11	13.1	+19.1% 
MMIX mems μ	2.6	2.8	+5% 
Scanned Elements	2	1.6	 -20%
Branch Mispredictions	0.57	0.58	+2% 

(0, 10, 20, 30, 40, 50, 60, 70, 80, 90, 100) (average case results)









Track Record of Dual-Pivot Quicksort

- **Observation in practice:**

Yaroslavskiy's Quicksort (YQS) **faster** than *classic Quicksort (CQS)* ... **why?**

↪ We did a mathematical analysis of YQS.

- Traditional cost measures do **not** explain observation!

	CQS	YQS	Relative
Running Time	(from various experiments)		 -10±2%
Comparisons	2	1.9	 -5%
Swaps	0.3	0.6	 +80%
Bytecode Instructions	18	21.7	 +20.6%
MMIX oops ν	11	13.1	 +19.1%
MMIX mems μ	2.6	2.8	 +5%
Scanned Elements	2	1.6	 -20%
Branch Mispredictions	0.57	0.58	 +2%

(0, 10, 20, 30, 40, 50, 60, 70, 80, 90, 100) average case results

Track Record of Dual-Pivot Quicksort

- **Observation in practice:**

Yaroslavskiy's Quicksort (YQS) **faster** than classic Quicksort (CQS) ... **why?**

↪ We did a mathematical analysis of YQS.

- Traditional cost measures do **not** explain observation!

	CQS	YQS	Relative
Running Time	(from various experiments)		-10±2%
Comparisons	2	1.9	-5%
Swaps	$0.\bar{3}$	0.6	+80%
Bytecode Instructions	18	21.7	+20.6%
MMIX oops ν	11	13.1	+19.1%
MMIX mems μ	$2.\bar{6}$	2.8	+5%
Scanned Elements	2	1.6	-20%
Branch Mispredictions	0.57	0.58	+2%

$\cdot n \ln n + O(n)$, average case results

Only plausible explanation for running time: 20% less memory transfers in YQS.









Track Record of Dual-Pivot Quicksort

- **Observation in practice:**

Yaroslavskiy's Quicksort (YQS) **faster** than classic Quicksort (CQS) ... **why?**

↪ We did a mathematical analysis of YQS.

- Traditional cost measures do **not** explain observation!

	CQS	YQS	Relative
Running Time	(from various experiments)		 $-10 \pm 2\%$
Comparisons	2	1.9	 -5%
Swaps	$0.\bar{3}$	0.6	$+80\%$ 
Bytecode Instructions	18	21.7	$+20.6\%$ 
MMIX ops ν	11	13.1	$+19.1\%$ 
MMIX mems μ	$2.\bar{6}$	2.8	$+5\%$ 
Scanned Elements	2	1.6	 -20%
Branch Mispredictions	0.57	0.58	$+2\%$ 

$\cdot n \ln n + O(n)$, average case results

Only plausible explanation for running time: 20% less memory transfers in YQS.









Track Record of Dual-Pivot Quicksort

- **Observation in practice:**

Yaroslavskiy's Quicksort (YQS) **faster** than *classic Quicksort (CQS)* ... **why?**

↪ We did a mathematical analysis of YQS.

- **Traditional** cost measures do **not** explain observation!

	CQS	YQS	Relative
Running Time	(from various experiments)		 $-10 \pm 2\%$
Comparisons	2	1.9	 -5%
Swaps	$0.\bar{3}$	0.6	$+80\%$ 
Bytecode Instructions	18	21.7	$+20.6\%$ 
MMIX oops ν	11	13.1	$+19.1\%$ 
MMIX mems μ	$2.\bar{6}$	2.8	$+5\%$ 
Scanned Elements	2	1.6	 -20%
Branch Mispredictions	0.57	0.58	$+2\%$ 

$\cdot n \ln n + O(n)$, average case results

Only plausible explanation for running time: 20% less memory transfers in YQS.


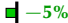



Track Record of Dual-Pivot Quicksort

- **Observation in practice:**

Yaroslavskiy's Quicksort (YQS) **faster** than *classic Quicksort (CQS)* ... **why?**

↪ We did a mathematical analysis of YQS.


- **Traditional** cost measures do **not** explain observation!

	CQS	YQS	Relative
Running Time	(from various experiments)		 $-10 \pm 2\%$
Comparisons	2	1.9	 -5%
Swaps	$0.\bar{3}$	0.6	$+80\%$ 
Bytecode Instructions	18	21.7	$+20.6\%$ 
MMIX oops ν	11	13.1	$+19.1\%$ 

What happened?!

Only plausible explanation for running time: 20% less memory transfers in YQS.

The “Memory Wall”



—□— CPU speed (MFLOPS)


STREAM benchmark data with linear regressions

www.cs.virginia.edu/stream/by_date/Balance.html

- 20 years since Bentley and McIlroy developed the classic Quicksort implementation

... *this most likely changes the game for sorting!*


The “Memory Wall”



- 20 years since Bentley and McIlroy developed the classic Quicksort implementation

... *this most likely changes the game for sorting!*

The “Memory Wall”



—□— CPU speed (MFLOPS)


Averaged annual growth rates:
46% CPU speed

STREAM benchmark data with linear regressions
www.cs.virginia.edu/stream/by_date/Balance.html

- 20 years since Bentley and McIlroy developed the classic Quicksort implementation

... this most likely changes the game for sorting!


The “Memory Wall”



- 20 years since Bentley and McIlroy developed the classic Quicksort implementation

... *this most likely changes the game for sorting!*


The “Memory Wall”



- 20 years since Bentley and McIlroy developed the classic Quicksort implementation

... *this most likely changes the game for sorting!*

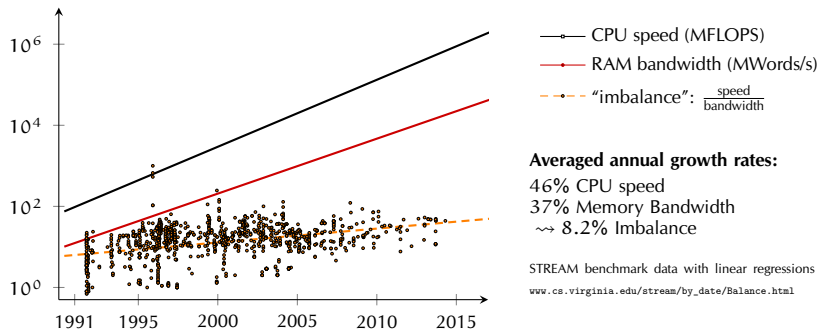
The “Memory Wall”



- 20 years since Bentley and McIlroy developed the classic Quicksort implementation

... this most likely changes the game for sorting!


The “Memory Wall”



- 20 years since Bentley and McIlroy developed the classic Quicksort implementation

... this most likely changes the game for sorting!

The “Memory Wall”



—○— CPU speed (MFLOPS)
—○— RAM bandwidth (MWords/s)
- - -○- - - “imbalance”: $\frac{\text{speed}}{\text{bandwidth}}$

Averaged annual growth rates:

46% CPU speed
37% Memory Bandwidth
~ 8.2% Imbalance

STREAM benchmark data with linear regressions


www.cs.virginia.edu/stream/by_date/Balance.html

- 20 years since Bentley and McIlroy developed the classic Quicksort implementation

→ relative cost of RAM accesses today 5 times as big

... this most likely changes the game for sorting!

The “Memory Wall”




- 20 years since Bentley and McIlroy developed the classic Quicksort implementation

↪ relative cost of RAM accesses today 5 times as big

... this most likely changes the game for sorting!

The “Memory Wall”

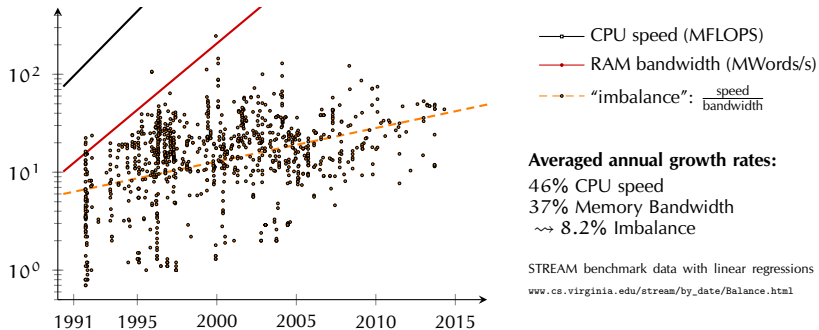


- 20 years since Bentley and McIlroy developed the classic Quicksort implementation

~> relative cost of RAM accesses today **5 times as big**

... *this most likely changes the game for sorting!*

The “Memory Wall”



- 20 years since Bentley and McIlroy developed the classic Quicksort implementation

~> relative cost of RAM accesses today **5 times as big**

... *this most likely changes the game for sorting!*

Analyzing Memory Transfers

Need an **abstract** and **simple** cost model to approximate **memory transfer**.

- abstract \rightarrow machine-independent results
- simple \rightarrow easy to analyze

- avoid any kind of memory accesses that are probably not needed
- \rightarrow not the same as cache misses and cache misses are not measurable

My proposal: number of *“scanned elements”*

- Machine model: Access to array only through iterators
- Cache size: 1
- Cache line: 1
- Cache hit or miss (no conditional)
- Cache miss: Access to next position
- Cache hit: Access to next position
- Cache miss: Cache miss + next position
- Cache hit: Cache hit + next position
- Cache miss: Cache miss + next position
- Cache hit: Cache hit + next position



... let's compare scanned elements for our Quicksorts!

Analyzing Memory Transfers

Need an **abstract** and **simple** cost model to approximate **memory transfer**.

- abstract \rightsquigarrow machine-independent results
 - simple \rightsquigarrow easy to analyze
 - should **only** count memory accesses that are probably **not** cached
- \rightsquigarrow neither swaps nor comparisons are suitable measures!

My proposal: number of *“scanned elements”*

Machine model: Access to array only through pointers

Access to memory

Access to array (left or right from *current*)

Access to *next* element

Access to *next* element with *current* pointer

Access to *next* element of *swaps*



... let's compare scanned elements for our Quicksorts!

Analyzing Memory Transfers

Need an **abstract** and **simple** cost model to approximate **memory transfer**.

- abstract \rightsquigarrow machine-independent results
 - simple \rightsquigarrow easy to analyze
 - should **only** count memory accesses that are probably **not** cached
- \rightsquigarrow neither swaps nor comparisons are suitable measures!

My proposal: number of *“scanned elements”*

● Machine model: Access to array only through pointers

● Cache model:

● Cache size: 2^k elements (one-dimensional)

● Cache associativity: 2^l elements

● Cache replacement: LRU (least recently used)

● Cache number of misses: $2^k - 2^l$



... let's compare scanned elements for our Quicksorts!

Analyzing Memory Transfers

Need an **abstract** and **simple** cost model to approximate **memory transfer**.

- abstract \rightsquigarrow machine-independent results
 - simple \rightsquigarrow easy to analyze
 - should **only** count memory accesses that are probably **not cached**
- \rightsquigarrow neither swaps nor comparisons are suitable measures!

My proposal: number of “**scanned elements**”

... means that I have to iterate only through elements

... that are not in the current partition

... that are not in the current subarray

... that are not in the current subarray

... that are not in the current subarray

... let's compare scanned elements for our Quicksorts!

Analyzing Memory Transfers

Need an **abstract** and **simple** cost model to approximate **memory transfer**.

- abstract \rightsquigarrow machine-independent results
- simple \rightsquigarrow easy to analyze
- should **only** count memory accesses that are probably **not cached**
- \rightsquigarrow neither swaps nor comparisons are suitable measures!

My proposal: number of “*scanned elements*”

- Machine model: Access to array only through iterators

... *scanned elements* for the *array*

... *scanned elements* for the *partition*

... *scanned elements* for the *partition*

... *scanned elements* for the *partition*

... *let's compare scanned elements for our Quicksorts!*

Analyzing Memory Transfers

Need an **abstract** and **simple** cost model to approximate **memory transfer**.

- abstract \rightsquigarrow machine-independent results
- simple \rightsquigarrow easy to analyze
- should **only** count memory accesses that are probably **not cached**
 \rightsquigarrow neither swaps nor comparisons are suitable measures!

My proposal: number of **“scanned elements”**

- Machine model: Access to array only through **iterators**
- Iterators can
 - head left or right (one-directional!)
 - **advance** to next position
 - **read** and **write** current position
- Cost: number of advances



... let's compare scanned elements for our Quicksorts!

Analyzing Memory Transfers

Need an **abstract** and **simple** cost model to approximate **memory transfer**.

- abstract \rightsquigarrow machine-independent results
- simple \rightsquigarrow easy to analyze
- should **only** count memory accesses that are probably **not cached**
 \rightsquigarrow neither swaps nor comparisons are suitable measures!

My proposal: number of **“scanned elements”**

- Machine model: Access to array only through **iterators**
- Iterators can
 - head left or right (one-directional!)
 - **advance** to next position
 - **read** and **write** current position
- Cost: number of advances



... let's compare scanned elements for our Quicksorts!


Analyzing Memory Transfers

Need an **abstract** and **simple** cost model to approximate **memory transfer**.

- abstract \rightsquigarrow machine-independent results
- simple \rightsquigarrow easy to analyze
- should **only** count memory accesses that are probably **not cached**
 \rightsquigarrow neither swaps nor comparisons are suitable measures!

My proposal: number of **“scanned elements”**

- Machine model: Access to array only through **iterators**
- Iterators can
 - head left or right (one-directional!)
 - **advance** to next position
 - **read** and **write** current position
- Cost: number of advances



... let's compare scanned elements for our Quicksorts!


Analyzing Memory Transfers

Need an **abstract** and **simple** cost model to approximate **memory transfer**.

- abstract \rightsquigarrow machine-independent results
- simple \rightsquigarrow easy to analyze
- should **only** count memory accesses that are probably **not cached**
 \rightsquigarrow neither swaps nor comparisons are suitable measures!

My proposal: number of **“scanned elements”**

- Machine model: Access to array only through **iterators**
- Iterators can
 - head left or right (one-directional!)
 - **advance** to next position
 - **read** and **write** current position
- Cost: number of advances



... let's compare scanned elements for our Quicksorts!


Analyzing Memory Transfers

Need an **abstract** and **simple** cost model to approximate **memory transfer**.

- abstract \rightsquigarrow machine-independent results
- simple \rightsquigarrow easy to analyze
- should **only** count memory accesses that are probably **not cached**
 \rightsquigarrow neither swaps nor comparisons are suitable measures!

My proposal: number of **“scanned elements”**

- Machine model: Access to array only through **iterators**
- Iterators can
 - head left or right (one-directional!)
 - **advance** to next position
 - **read** and **write** current position
- Cost: number of advances



... *let's compare scanned elements for our Quicksorts!*

Scanned Elements in CQS and YQS

How many scanned elements (SE) do we need for partitioning?

Classic Quicksort



- array scanned exactly once
- no scanned elements

Yaroslavskiy's Quicksort



- 1/3n SE on average

How does this translate to sorting costs?

Classic Quicksort



Recursion tree: $2n$ nodes

Yaroslavskiy's Quicksort




Recursion tree: n nodes

Scanned Elements in CQS and YQS

How many scanned elements (SE) do we need for partitioning?

Classic Quicksort



- array scanned exactly once
- ↪ n scanned elements

Yaroslavskiy's Quicksort



↪ $2n$ scanned elements

How does this translate to sorting costs?

Classic Quicksort



↪ $\log_2 n$ levels

Yaroslavskiy's Quicksort




↪ $\log_2 n$ levels

Scanned Elements in CQS and YQS

How many scanned elements (SE) do we need for partitioning?

Classic Quicksort



- array scanned exactly once

↔ n scanned elements

Yaroslavskiy's Quicksort



↔ $2n$ scanned elements

How does this translate to sorting costs?

Classic Quicksort



↔ $\log_2 n$ levels

Yaroslavskiy's Quicksort




↔ $\log_2 n$ levels

Scanned Elements in CQS and YQS

How many scanned elements (SE) do we need for partitioning?

Classic Quicksort



- array scanned exactly once
- ↪ n scanned elements

Yaroslavskiy's Quicksort



↪ $1.5n$ SE on average

How does this translate to sorting costs?

Classic Quicksort



Worst Case

Yaroslavskiy's Quicksort



Average Case

Scanned Elements in CQS and YQS

How many scanned elements (SE) do we need for partitioning?


Classic Quicksort



● array scanned exactly once

↪ n scanned elements

Yaroslavskiy's Quicksort



↪ $1.3n$ SE on average
worse than CQS?

How does this translate to sorting costs?

Classic Quicksort




Yaroslavskiy's Quicksort



Scanned Elements in CQS and YQS

How many scanned elements (SE) do we need for partitioning?


Classic Quicksort



• array scanned exactly once

↪ n scanned elements

Yaroslavskiy's Quicksort



↪ $1.3\bar{n}$ SE on average
worse than CQS!

How does this translate to sorting costs?

Classic Quicksort




Yaroslavskiy's Quicksort



Scanned Elements in CQS and YQS

How many scanned elements (SE) do we need for partitioning?


Classic Quicksort



• array scanned exactly once

↪ n scanned elements

Yaroslavskiy's Quicksort



↪ $1.3\bar{n}$ SE on average
worse than CQS!

How does this translate to sorting costs?

Classic Quicksort




Yaroslavskiy's Quicksort



Scanned Elements in CQS and YQS

How many scanned elements (SE) do we need for partitioning?


Classic Quicksort



● array scanned exactly once

↪ n scanned elements

Yaroslavskiy's Quicksort



↪ $1.3\bar{n}$ SE on average
worse than CQS!

How does this translate to sorting costs?

Classic Quicksort



↪ $2n \ln n$ SE overall

Yaroslavskiy's Quicksort



↪ $1.3\bar{n} \ln n$ SE overall

Scanned Elements in CQS and YQS

How many scanned elements (SE) do we need for partitioning?


Classic Quicksort



• array scanned exactly once

↪ n scanned elements


Yaroslavskiy's Quicksort



↪ $1.3\bar{n}$ SE on average
worse than CQS!


How does this translate to sorting costs?

Classic Quicksort



↪ $2 n \ln n$ SE overall

Yaroslavskiy's Quicksort




↪ $1.6 n \ln n$ SE overall

Scanned Elements in CQS and YQS

How many scanned elements (SE) do we need for partitioning?


Classic Quicksort



• array scanned exactly once

↪ n scanned elements


Yaroslavskiy's Quicksort



↪ $1.3\bar{n}$ SE on average
worse than CQS!


How does this translate to sorting costs?

Classic Quicksort



↪ $2n \ln n$ SE overall

Yaroslavskiy's Quicksort




↪ $1.6n \ln n$ SE overall

Scanned Elements in CQS and YQS

How many scanned elements (SE) do we need for partitioning?


Classic Quicksort



• array scanned exactly once

↪ n scanned elements


Yaroslavskiy's Quicksort



↪ $1.3\bar{n}$ SE on average
worse than CQS!


How does this translate to sorting costs?

Classic Quicksort



↪ $2n \ln n$ SE overall

Yaroslavskiy's Quicksort



↪ $1.6n \ln n$ SE overall









Track Record of Dual-Pivot Quicksort

- **Observation in practice:**

Yaroslavskiy's Quicksort (YQS) **faster** than *classic Quicksort (CQS)* ... **why?**

↪ We did a mathematical analysis of YQS.

- **Traditional** cost measures do **not** explain observation!

	CQS	YQS	Relative
Running Time	(from various experiments)		 $-10 \pm 2\%$
Comparisons	2	1.9	 -5%
Swaps	$0.\bar{3}$	0.6	$+80\%$ 
Bytecode Instructions	18	21.7	$+20.6\%$ 
MMIX oops ν	11	13.1	$+19.1\%$ 
MMIX mems μ	$2.\bar{6}$	2.8	$+5\%$ 
Scanned Elements	2	1.6	 -20%
Branch Mispredictions	0.57	0.58	$+2\%$ 

$\cdot n \ln n + O(n)$, average case results

Only plausible explanation for running time: 20% less memory transfers in YQS.









Track Record of Dual-Pivot Quicksort

- **Observation in practice:**

Yaroslavskiy's Quicksort (YQS) **faster** than *classic Quicksort (CQS)* ... **why?**

↪ We did a mathematical analysis of YQS.

- **Traditional** cost measures do **not** explain observation!

	CQS	YQS	Relative
Running Time	(from various experiments)		 $-10 \pm 2\%$
Comparisons	2	1.9	 -5%
Swaps	$0.\bar{3}$	0.6	$+80\%$ 
Bytecode Instructions	18	21.7	$+20.6\%$ 
MMIX oops ν	11	13.1	$+19.1\%$ 
MMIX mems μ	$2.\bar{6}$	2.8	$+5\%$ 
Scanned Elements	2	1.6	 -20%
Branch Mispredictions	0.57	0.58	$+2\%$ 

$\cdot n \ln n + O(n)$, average case results

Only plausible explanation for running time: 20% less memory transfers in YQS.









Track Record of Dual-Pivot Quicksort

- **Observation in practice:**

Yaroslavskiy's Quicksort (YQS) **faster** than *classic Quicksort (CQS)* ... **why?**

↪ We did a mathematical analysis of YQS.

- **Traditional** cost measures do **not** explain observation!

	CQS	YQS	Relative
Running Time	(from various experiments)		 -10±2%
Comparisons	2	1.9	 -5%
Swaps	$0.\bar{3}$	0.6	+80% 
Bytecode Instructions	18	21.7	+20.6% 
MMIX oops ν	11	13.1	+19.1% 
MMIX mems μ	$2.\bar{6}$	2.8	+5% 
Scanned Elements	2	1.6	 -20%
Branch Mispredictions	0.57	0.58	+2% 

$\cdot n \ln n + O(n)$, average case results

Only plausible explanation for running time: 20% less memory transfers in YQS.

- 1 Dual-Pivot Quicksort most likely faster because of fewer **memory references**.
- 2 The “memory wall” calls for new view on classic algorithms.
How about others?

3 Don't stop looking for the next algorithm
4 All the good ones are already implemented

Conclusion

- 1 Dual-Pivot Quicksort most likely faster because of fewer **memory references**.
- 2 The “memory wall” calls for new view on classic algorithms.
How about others?
- 3 Don't stop looking for theoretical explanations,
but do question models and assumptions!

Conclusion

- 1 Dual-Pivot Quicksort most likely faster because of fewer **memory references**.
- 2 The “memory wall” calls for new view on classic algorithms.
How about others?
- 3 Don't stop looking for theoretical explanations,
but do **question models** and assumptions!

Conclusion

- 1 Dual-Pivot Quicksort most likely faster because of fewer **memory references**.
- 2 The “memory wall” calls for new view on classic algorithms.
How about others?
- 3 Don't stop looking for theoretical explanations,
but do **question models** and assumptions!