

# Realistic analysis of algorithms

Cyril Nicaud

LIGM – Université Paris-Est Marne-la-Vallée & CNRS

AofA 2019, june 25

# Realistic?

- ① Which algorithms are implemented in standard libraries and why?
- ② How accurate is our model of computation? Can we improve it?

# Realistic?

- ① Which algorithms are implemented in standard libraries and why?
  - ▶ Java's Dual-Pivot Quicksort  
[C. Martínez, M. Nebel, R. Neininger, S. Wild, ...]
  - ▶ **TimSort (Python, Java, ...)**
  - ▶ ...
- ② How accurate is our model of computation? Can we improve it?

# Realistic?

- ① Which algorithms are implemented in standard libraries and why?
  - ▶ Java's Dual-Pivot Quicksort  
[C. Martínez, M. Nebel, R. Neininger, S. Wild, ...]
  - ▶ **TimSort (Python, Java, ...)**
  - ▶ ...
- ② How accurate is our model of computation? Can we improve it?
  - ▶ External memory model
  - ▶ Cache-oblivious model
  - ▶ **Branch predictions**
  - ▶ ...

# I. TimSort

*with N. Auger, V. Jugé & C. Pivoteau*

# TimSort algorithm

a	c	t	r	b	w	k	i	e	d	u	n
---	---	---	---	---	---	---	---	---	---	---	---

# TimSort algorithm


a	c	t	r	b	w	k	i	e	d	u	n
---	---	---	---	---	---	---	---	---	---	---	---

- The input is split into **runs**, which are **monotonic subsequences**

# TimSort algorithm


a	c	t	r	b	w	k	i	e	d	u	n
---	---	---	---	---	---	---	---	---	---	---	---

- The input is split into **runs**, which are **monotonic subsequences**
- Every discovered run is added to a **stack**, then some **consecutive** runs can be **merged** (as in MergeSort) ← **more details later**






## TimSort algorithm



- The input is split into **runs**, which are **monotonic subsequences**
- Every discovered run is added to a **stack**, then some **consecutive** runs can be **merged** (as in MergeSort) ← **more details later**




- When there is no more run, the runs in the stack are merged top-down

# TimSort algorithm

a	c	t	r	b	w	k	i	e	d	u	n
---	---	---	---	---	---	---	---	---	---	---	---

- The input is split into **runs**, which are **monotonic subsequences**
- Every discovered run is added to a **stack**, then some **consecutive** runs can be **merged** (as in MergeSort) ← **more details later**



- When there is no more run, the runs in the stack are merged top-down

**Remark:** TimSort also contains a lot of heuristics that we don't consider here (especially in the merge procedure)

## timsort.txt (from Tim Peters)

This describes an adaptive, stable, natural mergesort, modestly called timsort (hey, I earned it <wink>). It has supernatural performance on many kinds of partially ordered arrays (less than  $\lg(N!)$  comparisons needed, and as few as  $N-1$ ), yet as fast as Python's previous highly tuned samplesort hybrid on random arrays.

I believe that lists very often do have exploitable partial order in real life, and this is the strongest argument in favor of timsort

# Running Time

In 2003, TimSort is announced to be in  $\mathcal{O}(\log n!)$ , with no formal proof.

Theorem (Auger, Nicaud, Pivoteau 2015)

**TimSort** has a **worst-case running time** of  $\mathcal{O}(n \log n)$ .

The proof is not very difficult, but hard to read (and to teach!)

Theorem (Folklore)

The running time of any sorting by comparisons algorithm is  $\Omega(n \log n)$ .

So TimSort is optimal, as many other algorithms: it does not explain why it is **used in practice**!

# Parameterize Running Time

I believe that lists very often do have exploitable partial order in real life, and this is the strongest argument in favor of timsort

**Idea:** add a **parameter** to describe the running time

# Parameterize Running Time

I believe that lists very often do have exploitable partial order in real life, and this is the strongest argument in favor of timsort

**Idea:** add a **parameter** to describe the running time

- **First choice:** the **number of runs**  $\rho$ .

It was conjectured that TimSort runs in  $\mathcal{O}(n + n \log \rho)$

# Parameterize Running Time

I believe that lists very often do have exploitable partial order in real life, and this is the strongest argument in favor of timsort

**Idea:** add a **parameter** to describe the running time

- **First choice:** the **number of runs**  $\rho$ .

It was conjectured that TimSort runs in  $\mathcal{O}(n + n \log \rho)$

- **Better choice:** the **run lengths entropy**  $\mathcal{H}$ .

If the runs have size  $r_1, \dots, r_\rho$ , then

$$\mathcal{H} := - \sum_{i=1}^{\rho} \frac{r_i}{n} \log_2 \frac{r_i}{n}$$

# Parameterize Running Time

I believe that lists very often do have exploitable partial order in real life, and this is the strongest argument in favor of timsort

**Idea:** add a **parameter** to describe the running time

- **First choice:** the **number of runs**  $\rho$ .

It was conjectured that TimSort runs in  $\mathcal{O}(n + n \log \rho)$

- **Better choice:** the **run lengths entropy**  $\mathcal{H}$ .

If the runs have size  $r_1, \dots, r_\rho$ , then

$$\mathcal{H} := - \sum_{i=1}^{\rho} \frac{r_i}{n} \log_2 \frac{r_i}{n}$$

If the runs have sizes  $\frac{n}{11}, \dots, \frac{n}{11}$ :  $\mathcal{H} = \log_2 11 \approx 3.46$

If the runs have sizes  $\frac{90n}{100}, \frac{n}{100}, \dots, \frac{n}{100}$ :  $\mathcal{H} \approx 0.80$

If the runs have sizes  $\sqrt{n}, \dots, \sqrt{n}$ :  $\mathcal{H} = \frac{1}{2} \log_2 n$



# Our results

Theorem (Auger, Jugé, Nicaud, Pivoteau. ESA 2018)

**TimSort** has a **worst-case running time** of  $\mathcal{O}(n + n \log \rho)$ .

Theorem (Auger, Jugé, Nicaud, Pivoteau. Talk ESA 2018)

**TimSort** has a **worst-case running time** of  $\mathcal{O}(n + n\mathcal{H})$ .

We always have  $\mathcal{H} \leq \log_2 \rho \leq \log_2 n$ .

Theorem (Auger, Jugé, Nicaud, Pivoteau. Buss, Knop 2019)

**TimSort** needs  $1.5n\mathcal{H} + \mathcal{O}(n)$  comparisons in the worst case.

Theorem (Barbay, Navarro 2013)

Sorting by comparisons algorithms use more than  $n\mathcal{H} - \mathcal{O}(n)$  comparisons.

# Optimal algorithms?

Theorem (Auger, Jugé, Nicaud, Pivoteau. Buss, Knop 2019)

**TimSort** needs  $1.5n\mathcal{H} + \mathcal{O}(n)$  comparisons in the worst case.

Theorem (Barbay, Navarro 2013)

Sorting by comparisons algorithms use more than  $n\mathcal{H} - \mathcal{O}(n)$  comparisons.

There is a gap, and in fact, TimSort is **not optimal** (for this parameter  $\mathcal{H}$ ).

# Optimal algorithms?

Theorem (Auger, Jugé, Nicaud, Pivoteau. Buss, Knop 2019)

**TimSort** needs  $1.5n\mathcal{H} + \mathcal{O}(n)$  comparisons in the worst case.

Theorem (Barbay, Navarro 2013)

Sorting by comparisons algorithms use more than  $n\mathcal{H} - \mathcal{O}(n)$  comparisons.

There is a gap, and in fact, TimSort is **not optimal** (for this parameter  $\mathcal{H}$ ).

Some optimal algorithms are known: **Takaoka 2009, Barbay & Navarro 2013, Munro & Wild 2018.**

So why analyzing TimSort?

# Optimal algorithms?

Theorem (Auger, Jugé, Nicaud, Pivoteau. Buss, Knop 2019)

**TimSort** needs  $1.5n\mathcal{H} + \mathcal{O}(n)$  comparisons in the worst case.

Theorem (Barbay, Navarro 2013)

Sorting by comparisons algorithms use more than  $n\mathcal{H} - \mathcal{O}(n)$  comparisons.

There is a gap, and in fact, TimSort is **not optimal** (for this parameter  $\mathcal{H}$ ).

Some optimal algorithms are known: **Takaoka 2009, Barbay & Navarro 2013, Munro & Wild 2018.**


So why analyzing TimSort? **because it is used in Python, Java, ...**

# Back to TimSort

## Recall:

- monotonic runs are computed and added to a stack
- some merges of consecutive runs may happen when a run is added
- at the end, the remaining runs are merged top-down

## Merges:




- 1 **Run merging** algorithm: standard + many optimizations
  - ▶ time  $\mathcal{O}(k + \ell)$ , using  $k + \ell$  comparisons<sup>1</sup>
  - ▶ memory  $\mathcal{O}(\min(k, \ell))$
- 2 **Policy** for choosing runs to merge:
  - ▶ depends on **run lengths** only

Let us forget array values – **only remember run lengths!**

---

<sup>1</sup>It is  $k + \ell - 1$ , but we'll use  $k + \ell$  to simplify.

# TimSort's Merging Rules



STACK

## Notations:

- the run  $R_i$  has length  $r_i$
- the stack has height  $h$
- the topmost run is  $R_h$


## Merges after adding a new run:

- **While true**
  - ▶ **if**  $r_h > r_{h-2}$  **then** merge  $R_{h-1}$  and  $R_{h-2}$
  - ▶ **else if**  $r_h \geq r_{h-1}$  **then** merge  $R_h$  and  $R_{h-1}$
  - ▶ **else if**  $r_h + r_{h-1} \geq r_{h-2}$  **then** merge  $R_h$  and  $R_{h-1}$
  - ▶ **else** break

## Remarks:

- we only consider the three topmost runs
- we only merge  $R_h$  and  $R_{h-1}$ , or  $R_{h-1}$  and  $R_{h-2}$

# TimSort's Merging Rules



STACK

Merges after adding a new run:


- **While true**

- ▶ **if**  $r_h > r_{h-2}$  **then** merge  $R_{h-1}$  and  $R_{h-2}$
- ▶ **else if**  $r_h \geq r_{h-1}$  **then** merge  $R_h$  and  $R_{h-1}$
- ▶ **else if**  $r_h + r_{h-1} \geq r_{h-2}$  **then** merge  $R_h$  and  $R_{h-1}$
- ▶ **else** break

**timsort.txt:**

Note that, by induction, it implies the lengths of pending runs form a decreasing sequence. It implies that, reading the lengths right to left, the pending-run lengths grow at least as fast as the Fibonacci numbers. Therefore the stack can never grow larger than about  $\log_\phi(N)$  entries

# TimSort's Merging Rules



STACK


Merges after adding a new run:

- **While true**

- ▶ **if**  $r_h > r_{h-2}$  **then** merge  $R_{h-1}$  and  $R_{h-2}$
- ▶ **else if**  $r_h \geq r_{h-1}$  **then** merge  $R_h$  and  $R_{h-1}$
- ▶ **else if**  $r_h + r_{h-1} \geq r_{h-2}$  **then** merge  $R_h$  and  $R_{h-1}$
- ▶ **else** break

**timsort.txt:**

Note that, by induction,  $r_h$  implies the lengths of pending runs form a Fibonacci sequence. It implies that, reading the stack from right to left, the pending-run lengths are at least as fast as the Fibonacci numbers. Therefore the stack can never grow larger than about  $\log_\phi(N)$  entries





## An error in timsort.txt

- While true

- ▶ if  $r_h > r_{h-2}$  then merge  $R_{h-1}$  and  $R_{h-2}$
- ▶ else if  $r_h \geq r_{h-1}$  then merge  $R_h$  and  $R_{h-1}$
- ▶ else if  $r_h + r_{h-1} \geq r_{h-2}$  then merge  $R_h$  and  $R_{h-1}$
- ▶ else break

 The invariant  $r_{i+2} + r_{i+1} < r_i$  does not hold!

Discovered by de Gouw et al (2015) while trying to prove (formally) the correctness of Java's Timsort, using KeY (verification tool for Java)

## An error in timsort.txt

- **While true**

- ▶ **if**  $r_h > r_{h-2}$  **then** merge  $R_{h-1}$  and  $R_{h-2}$
- ▶ **else if**  $r_h \geq r_{h-1}$  **then** merge  $R_h$  and  $R_{h-1}$
- ▶ **else if**  $r_h + r_{h-1} \geq r_{h-2}$  **then** merge  $R_h$  and  $R_{h-1}$
- ▶ **else** break

 The invariant  $r_{i+2} + r_{i+1} < r_i$  **does not hold!**

Discovered by **de Gouw et al (2015)** while trying to prove (formally) the correctness of Java's Timsort, using KeY (verification tool for Java)

Is it a **real** problem?

- In **Python**: not really, the algorithm is still efficient and correct
- In **Java**: they use the invariant to fix the maximum size of the stack, implemented with a static array  $\Rightarrow$  **de Gouw et al (2015)** built an array that produces an error for Java's `sort()`!

# Two versions of TimSort

de Gouw et al (2015) proposed two solutions to fix the problem:

## 1. Adding a new rule (implemented in Python)

### • While true

- ▶ if  $r_h > r_{h-2}$  then merge  $R_{h-1}$  and  $R_{h-2}$
- ▶ else if  $r_h \geq r_{h-1}$  then merge  $R_h$  and  $R_{h-1}$
- ▶ else if  $r_h + r_{h-1} \geq r_{h-2}$  then merge  $R_h$  and  $R_{h-1}$
- ▶ else if  $r_{h-1} + r_{h-2} \geq r_{h-3}$  then merge  $R_h$  and  $R_{h-1}$
- ▶ else break

The invariant now holds, the algorithm is certified in KeY.

## 2. Computing correct maximal heights for the stack (implemented in Java)

### Lemma

Throughout execution of TimSort, the invariant cannot be violated at two consecutive runs in the stack.

## Running time analysis: $\mathcal{O}(n \log n)$

We focus on the main loop: other parts are done in  $\mathcal{O}(n)$  comparisons.

- **While there are remaining runs**

- (#1) Add a new run to the stack

- Repeat until stabilized**

- (#2) if  $r_h > r_{h-2}$  then merge  $R_{h-1}$  and  $R_{h-2}$

- (#3) else if  $r_h \geq r_{h-1}$  then merge  $R_h$  and  $R_{h-1}$

- (#4) else if  $r_h + r_{h-1} \geq r_{h-2}$  then merge  $R_h$  and  $R_{h-1}$

- (#5) else if  $r_{h-1} + r_{h-2} \geq r_{h-3}$  then merge  $R_h$  and  $R_{h-1}$

## Amortized analysis:

- $\diamond$ -tokens and  $\heartsuit$ -tokens are given to the elements of the input
- tokens are used to pay for comparisons
- the total number of tokens granted is our upper bound


**Tokens' rules:** an element gets two  $\diamond$  and one  $\heartsuit$

- when its run enters the stack
- when its height in the stack decreases

Running time analysis:  $\mathcal{O}(n \log n)$ , case #2

(#2) if  $r_h > r_{h-2}$  then merge  $R_{h-1}$  and  $R_{h-2}$

Every element of  $R_h$  and  $R_{h-1}$  pays one  $\diamond$ : the merge cost is  $r_{h-1} + r_{h-2} \leq r_{h-1} + r_h$ , hence it is fully paid.




The height of every element that paid one  $\diamond$  decreases by one: they all gain two  $\diamond$  and one  $\heartsuit$

Running time analysis:  $\mathcal{O}(n \log n)$ , case #3

(#3) else if  $r_h \geq r_{h-1}$  then merge  $R_h$  and  $R_{h-1}$

Every element of  $R_h$  pays two  $\diamond$ : the merge cost is  $r_h + r_{h-1} \leq 2r_h$ , hence it is fully paid.




The height of every element that paid two  $\diamond$  decreases by one: they all gain two  $\diamond$  and one  $\heartsuit$

Running time analysis:  $\mathcal{O}(n \log n)$ , case #4

(#4) **else if**  $r_h + r_{h-1} \geq r_{h-2}$  **then** merge  $R_h$  and  $R_{h-1}$

Every element of  $R_h$  pays one  $\diamond$ , every element of  $R_{h-1}$  pays one  $\heartsuit$ : the merge cost is  $r_h + r_{h-1}$ , hence it is fully paid.




The height of the elements of  $R_h$  decreases by one: **ok** for  $\diamond$

- Elements that paid one  $\heartsuit$  are now in the topmost run
- Elements in the topmost run never pay with  $\heartsuit$
- In the new stack,  $r_h \geq r_{h-1}$  so another merge is going to occur (#3)
- The height of the new topmost run is going to decrease during this new merge, its elements will get two  $\diamond$  and one  $\heartsuit$

## Running time analysis: $\mathcal{O}(n \log n)$ , case #5

(#5) **else if**  $r_{h-1} + r_{h-2} \geq r_{h-3}$  **then** merge  $R_h$  and  $R_{h-1}$

Every element of  $R_h$  pays one  $\diamond$ , every element of  $R_{h-1}$  pays one  $\heartsuit$ : the merge cost is  $r_h + r_{h-1}$ , hence it is fully paid.



The height of the elements of  $R_h$  decreases by one: **ok** for  $\diamond$

- Elements that paid one  $\heartsuit$  are now in the topmost run
- Elements in the topmost run never pay with  $\heartsuit$
- In the new stack,  $r_h + r_{h-1} \geq r_{h-2}$  so another merge is going to occur (#4)
- The height of the new topmost run is going to decrease during this new merge, its elements will get two  $\diamond$  and one  $\heartsuit$



# Running time analysis: $\mathcal{O}(n \log n)$

## Summary:

- Computing the run decomposition takes  $\mathcal{O}(n)$
- For the main loop:
  - ▶ each element gets  $2\blacklozenge$  and  $1\heartsuit$  when entering the stack
  - ▶ each merge is paid with  $\blacklozenge$  and  $\heartsuit$
  - ▶ when an element pays with  $\blacklozenge$ , it get it (them) back immediately after
  - ▶ when an element pays with  $\heartsuit$ , another merge occurs just after, during which it get it back
- The final merges are done in  $\mathcal{O}(n)$  by direct computation

## Lemma

At any moment during TimSort, the stack has height in  $\mathcal{O}(\log n)$ .

**Proof:** the invariant holds.

**Theorem** (Auger, Jugé, Nicaud, Pivoteau 2018)


The running time of TimSort is in  $\mathcal{O}(n \log n)$ .

## Running time analysis: $\mathcal{O}(n + n\mathcal{H})$

Recall: #1 is the insertion of a new run in the stack

Recall:  $\mathcal{H} = -\sum \frac{r_i}{n} \log \frac{r_i}{n}$

We use the following decomposition of the sequence of events:



Two lemmas (both consequences of the invariant):

- The total cost in ♠-tokens is **linear**
- The **height** of the stack at the **beginning of the ending sequence** after inserting a run of length  $r$  is  $\mathcal{O}(\log \frac{n}{r})$ .

Theorem (Auger, Jugé, Nicaud, Pivoteau 2018)

The running time of TimSort is in  $\mathcal{O}(n + n\mathcal{H})$ .

## Running time analysis: summary

We proved that for the “new” TimSort, we have:

Theorem (Auger, Jugé, Nicaud, Pivoteau 2018)

The running time of TimSort is in  $\mathcal{O}(n + n\mathcal{H})$ .

This can be improved to:

Theorem (Auger, Jugé, Nicaud, Pivoteau 2019)

TimSort needs at most  $1.5n\mathcal{H} + \mathcal{O}(n)$  comparisons.

## Running time analysis: summary

We proved that for the “new” TimSort, we have:

Theorem (Auger, Jugé, Nicaud, Pivoteau 2018)

The running time of TimSort is in  $\mathcal{O}(n + n\mathcal{H})$ .

This can be improved to:

Theorem (Auger, Jugé, Nicaud, Pivoteau 2019)

TimSort needs at most  $1.5n\mathcal{H} + \mathcal{O}(n)$  comparisons.

What about the **Java's version of TimSort**? it is also in  $\mathcal{O}(n + n \log \rho)$ , but it is much more complicated to establish (no nice invariant).

## Running time analysis: summary

We proved that for the “new” TimSort, we have:

Theorem (Auger, Jugé, Nicaud, Pivoteau 2018)

The running time of TimSort is in  $\mathcal{O}(n + n\mathcal{H})$ .

This can be improved to:

Theorem (Auger, Jugé, Nicaud, Pivoteau 2019)

TimSort needs at most  $1.5n\mathcal{H} + \mathcal{O}(n)$  comparisons.

What about the **Java's version of TimSort**? it is also in  $\mathcal{O}(n + n \log \rho)$ , but it is much more complicated to establish (no nice invariant).

*but wait a minute ...*

# Another bug in Java's TimSort

## Lemma

Throughout execution of TimSort, the invariant cannot be violated at two consecutive runs in the stack.

# Another bug in Java's TimSort

## Lemma

Throughout execution of TimSort, ~~the~~ invariant cannot be violated at two consecutive runs in the stack.



The lemma is incorrect!

We built an array **that produces an error to Java's (patched) TimSort!**

# Another bug in Java's TimSort

## Lemma

Throughout execution of TimSort, the invariant cannot be violated at two consecutive runs in the stack.



The lemma is incorrect!

We built an array **that produces an error to Java's (patched) TimSort!**

The screenshot shows the Oracle JIRA issue page for JDK-8203864, titled "Execution error in Java's Timsort". The issue is marked as "RESOLVED".

**Details**

- Type: Bug
- Priority: P3
- Affects Version/s: None
- Component/s: core-libs
- Labels: None
- Subcomponent: java.util.collections
- Introduced In: 6
- Version: b20
- Resolved In Build: b20

**Status:** RESOLVED

**Resolution:** Fixed

**Fix Version/s:** 11

**People**

**Assignee:** Martin Buchholz

**Reporter:** Rémi Forax

**Votes:** 0

**Watchers:** 0

**Backports**

Issue	Fix Version	Assignee	Priority	Status	Resolution	Resolved In Build
JDK-8206770 12		Doug Lea	P3	Resolved	Fixed	team
JDK-8206547 11.0.1		Doug Lea	P3	Resolved	Fixed	b01

**Description**

Carine Pivoteau wrote:  
While working on a proper complexity analysis of the algorithm, we realised that there was an error in the last paper reporting such a bug (<https://envisage-project.eu/wp-content/uploads/2015/02/sorting.pdf>). This implies that the correction implemented in the java source code (changing Timsort stack size) is wrong and that it is still possible to make it break. This is explained in full details in our analysis: <https://arxiv.org/pdf/1805.08612.pdf>

**Dates**

**Created:** 2018-05-27 23:04

**Updated:** 2018-07-07 00:07

**Resolved:** 2018-06-25 10:12



# Another bug in Java's TimSort

## Lemma

Throughout execution of TimSort,  invariant cannot be violated at two consecutive runs in the stack.

The lemma is incorrect!

We built an array **that produces an error to Java's (patched) TimSort!**



Apache Commons JIRA

Dashboards Projects Issues Search

Down for Maintenance on Friday January 25 from 3-5 PM PST

JDK / JDK-8203864

Execution error in Java's Timsort

Details

Type: Exception  
Priority: P  
Affects Version/s: None  
Component/s: core-11  
Labels: None  
Subcomponent: java.util  
Introduced In Version: 6  
Resolved In Build: b20

Backlogs

Issue	Fix Version	Assignee
JDK-8206770	12	Doug Lea
JDK-8206547	11.0.1	Doug Lea

Description

Carine Pivoteau wrote:  
While working on a proper complexity ana...  
was an error in the last paper reporting suc...  
content/uploads/2015/02/sorting.pdf. This i...  
the java source code (changing Timsort stac...  
to make it break. This is explained in full deta...  
/pdf/1805\_08612.pdf

**Hacker News** new | comments | ask | show | jobs | submit

On the Worst-Case Complexity of TimSort (pajuhh.de)  
204 points by pajuhh 4 months ago | hide | post | web | favorite | 74 comments

204 points by pajuhh 4 months ago | hide | post | web | favorite | 74 comments

The linked java test file, <http://pwn.milky.frl/~pivoteau/TimSort/Test.java>, still crashes the latest Java 10.0.2 with an "Exception in thread 'main'"

java.lang.ArrayIndexOutOfBoundsException: 49 - Amazing! I wonder if this makes some web services vulnerable... if the user can submit a just-so array of ints to be sorted? But it does seem like it would require uploading a really huge array (~4GB)

This looks like <https://bugs.openjdk.java.net/browse/JDK-8203864>, which has the following additional information:

"While working on a proper complexity analysis of the algorithm, we realised that there was an error in the last paper reporting such a bug (<http://pwn.milky.frl/~pivoteau/TimSort/Test.java>)... This implies that the correction implemented in the Java source code (changing Timsort stack size) is wrong and that it is still possible to make it break. (changing Timsort stack size) is wrong and that it is still possible to make it break. This is explained in full details in our analysis: <https://arxiv.org/pdf/1805.08612.pdf>"

That is not additional information. That bug was created by the authors of this post, and the link is a reference to the exact paper linked to in this post.

Crashing out of execution doesn't really seem like an actionable attack vector, unless that vector bubbles all the way to the application invocation.

# Another bug in Java's TimSort

## Lemma

Throughout execution of TimSort, the invariant cannot be violated at two consecutive runs in the stack.



The lemma is incorrect!

We built an array **that produces an error to Java's (patched) TimSort!**



### Details

Type: Bug

Priority: P3

Affects Version/s: 10.0.2

Component/s: core-rt

Labels: None

Subcomponent: java.util

Introduced in: 6

Version: b20

Resolved in Build: b20

### Backports

Issue

Fix Version

Assignee

JDK-8206770 12

Doug Lea

JDK-8206547 11.0.1

Doug Lea

Description

Carine Pivoteau wrote:

While working on a proper complexity ana-

was an error in the last paper reporting suc-

content/uploads/2015/02/sorting.pdf. This i-

the java source code (changing Timsort stac-

to make it break. This is explained in full deta-

/pdf/1805.08612.pdf

### Stacker News

new | comments | all

On the Worst-Case Complexity of Timsort

204 points by gordon 4 months ago | 1 vote | 1 comment

On the Worst-Case Complexity of Timsort

The linked java test file, [https://github.com/OracleOpen/java-lang-ArrayIndexOutOfBounds](#)

java-lang-ArrayIndexOutOfBounds

services vulnerable... if the v

like it would require upload

4 months ago

4 months ago

This look like it

additional info

"While working on a

error in r

code (cha

This is e

4 months ago

4 months ago

4 months ago

4 months ago

4 months ago

4 months ago

4 months ago

```
1.13
1.14
1.15
1.16 +
1.17 +
1.18 +
1.19 +
1.20 +
1.21
1.22
1.23
1.24
1.25 -
1.26 +
1.27 +
1.28
1.29
1.30 -
1.31 -
1.32 -
1.33 -
1.34 +
1.35
1.36
1.37 +
1.38
1.39
}

* This method is called each time a new run is pushed onto the stack,
* so the invariants are guaranteed to hold for i < stackSize upon
* entry to the method.

* Thanks to Stijn de Gouw, Jurriaan Rot, Frank S. de Boer,
* Richard Bubel and Reiner Hahnle, this is fixed with respect to
* the analysis in "On the Worst-Case Complexity of TimSort" by
* Nicolas Auger, Vincent Jug, Cyril Nicaud, and Carine Pivoteau.
*/

private void mergeCollapse() {
    while (stackSize > 1) {
        int n = stackSize - 2;
        if (n > 0 && runLen[n-1] <= runLen[n] + runLen[n+1]) {
            n > 1 && runLen[n-1] <= runLen[n] + runLen[n+1] ||
            if (runLen[n-2] <= runLen[n] + runLen[n+1]) {
                n--;
            } else if (runLen[n] <= runLen[n+1]) {
                mergeAt(n);
            } else {
                mergeAt(n);
            } else if (n < 0 || runLen[n] > runLen[n+1]) {
                break; // Invariant is established
            }
        }
        mergeAt(n);
    }
}
```

# Conclusion for TimSort

- TimSort is an **efficient** algorithm, in theory and in practice
- It is not **entropy-optimal**, but not far from it
- There are many optimisation to build the runs, to perform the merges,  
...
- Its  $\mathcal{O}(n \log n)$  running time was proved **more than 10 years** after it was announced
- There were **two consecutive bugs** in Java's version, due to improper analysis of the algorithm

# Conclusion for TimSort

- TimSort is an **efficient** algorithm, in theory and in practice
- It is not **entropy-optimal**, but not far from it
- There are many optimisation to build the runs, to perform the merges,  
...
- Its  $\mathcal{O}(n \log n)$  running time was proved **more than 10 years** after it was announced
- There were **two consecutive bugs** in Java's version, due to improper analysis of the algorithm

# Conclusion for TimSort

- TimSort is an **efficient** algorithm, in theory and in practice
- It is not **entropy-optimal**, but not far from it
- There are many optimisation to build the runs, to perform the merges,  
...
- Its  $\mathcal{O}(n \log n)$  running time was proved **more than 10 years** after it was announced
- There were **two consecutive bugs** in Java's version, due to improper analysis of the algorithm

Every used algorithm deserves a **fine grain analysis**

## II. Branch predictions

*with N. Auger & C. Pivoteau*

## A toy example: looking for the min and the max

We want to find the minimum and the maximum of an array  $T$  of size  $n$ .

```
min = T[n-1];  
max = T[n-1];  
for(i=0; i<n-1; i++){  
    a = T[i];  
    if (a < min) min = a;  
    if (a > max) max = a;  
}
```

### Naive solution:

foreach element  $a$  of  $T$ , if  $a$  is smaller than the current minimum, update the minimum; if it is greater than the current maximum, update the maximum.

**Fact:** the naive solution uses  $2n - 2 \sim 2n$  comparisons.

Can we do better?

## Min & Max: Optimal Algorithm

**Idea:** take the elements by pairs ( $a_1, a_2$ ), compare them, then compare the smallest to the current min and the largest to the current max

```
min = max = T[n-1];
for(i=0; i<n-1; i+=2){
    a1 = T[i];
    a2 = T[i+1];
    if (a1 < a2) {
        if (a1 < min) min = a1;
        if (a2 > max) max = a2;
    }
    else {
        if (a2 < min) min = a2;
        if (a1 > max) max = a1;
    }
}
```

Number of comparisons:

- $\sim \frac{n}{2}$  loop iterations
- 3 comparisons by iterations
- number of comparisons:  $\sim \frac{3}{2}n$

That's better!



## Min & Max: Optimal Algorithm

**Idea:** take the elements by pairs  $(a_1, a_2)$ , compare them, then compare the smallest to the current min and the largest to the current max

```
min = max = T[n-1];
for(i=0; i<n-1; i+=2){
    a1 = T[i];
    a2 = T[i+1];
    if (a1 < a2) {
        if (a1 < min) min = a1;
        if (a2 > max) max = a2;
    }
    else {
        if (a2 < min) min = a2;
        if (a1 > max) max = a1;
    }
}
```

Number of comparisons:


- $\sim \frac{n}{2}$  loop iterations
- 3 comparisons by iterations
- number of comparisons:  $\sim \frac{3}{2}n$

That's better!

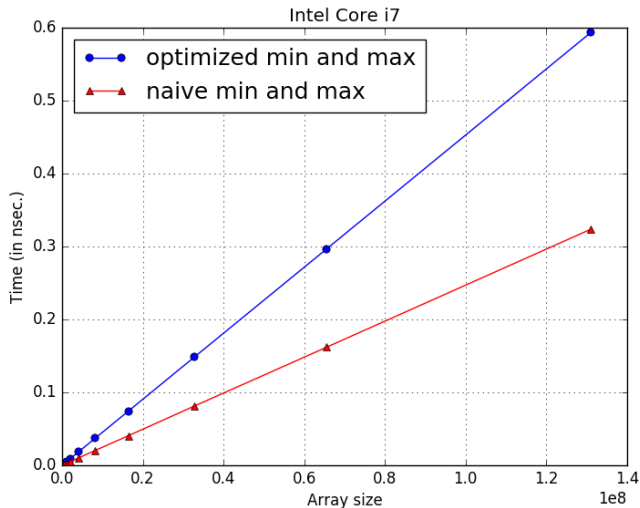
### Theorem (Folklore)

At least  $\sim \frac{3}{2}n$  comparisons are needed to compute the min and the max.

# Min & Max: experiments



## Min & Max: experiments




The **naive solution** is more efficient in practice!

# Pipeline

## Notion of pipeline:

- During the execution of a program, **instructions** are executed sequentially: `i=3`, `a<b`, `if (...)`, ...
- Instructions are divided into **several sequential steps**
- Different steps can be handle **in parallel** by the processor


## Example with 5 steps:



It can be up to **five time** as fast

# Pipeline and Branches


- A **branch** is an instruction with several possible following instructions: if, while, ...
- Branches constitute a **problem** for the pipeline:



- **We have to wait** for the completion of all the stages of **<condition>** to know whether it is followed by **A** or by **B**!

# Pipeline and Branches

- A **branch** is an instruction with several possible following instructions: if, while, ...
- Branches constitute a **problem** for the pipeline:




- **We have to wait** for the completion of all the stages of `<condition>` to know whether it is followed by `A` or by `B`!

**Solution:** try to anticipate if condition = **true** or **false**

## Branch predictions

- Branches does not fit well with the pipeline
- We try to **anticipate** whether the branch will be:
  - ▶ **Taken (T)**: when **<condition>** is **true**
  - ▶ **Not Taken (NT)**: when **<condition>** is **false**
- We push the predicted next instruction in the pipeline:
  - ▶ if the prediction is **correct**, we **gain** some time
  - ▶ if it is **incorrect**, we have to undo what we did, we **lose** some time


A simple **local predictor**, the **2-bit predictor** (one for each branch):



## Back to the toy example

```
min = T[n-1];  
max = T[n-1];  
for(i=0; i<n-1; i++){  
    a = T[i];  
    if (a < min) min = a;  
    if (a > max) max = a;  
}
```

The condition `if (a < min)` is true when there is a **min-record**, and false otherwise.



We have a pure AofA exercise:


- Start at any state, draw a uniform random permutation
- Scan it from left to right: when there is a min-record, go to the right in the automaton (if possible), otherwise go to the left
- What is the **expected number of mispredictions**?



## Back to the toy example

```
min = T[n-1];
max = T[n-1];
for(i=0; i<n-1; i++){
    a = T[i];
    if (a < min) min = a;
    if (a > max) max = a;
}
```

The condition `if (a < min)` is true when there is a **min-record**, and false otherwise.



We have a pure AofA exercise:

- Start at any state, draw a uniform random permutation
- Scan it from left to right: when there is a min-record, go to the right in the automaton (if possible), otherwise go to the left
- What is the **expected number of mispredictions**?

### Lemma

The expected number of mispredictions produced by each `if` in the naive solution is asymptotically equivalent to  $\log n$ .

## What About the Optimal Algorithm?

**Idea:** take the elements by pairs  $(a_1, a_2)$ , compare them, then compare the smallest to the current min and the largest to the current max

```
min = max = T[n-1];
for(i=0; i<n-1; i+=2){
    a1 = T[i];
    a2 = T[i+1];
    if (a1 < a2) {
        if (a1 < min) min = a1;
        if (a2 > max) max = a2;
    }
    else {
        if (a2 < min) min = a2;
        if (a1 > max) max = a1;
    }
}
```

- The first branch: if  $(a_1 < a_2)$  is **true** with probability  $\frac{1}{2}$  for uniform random permutations.
- This cannot be well predicted: there is a misprediction here with probability  $\frac{1}{2}$  for each loop iteration
- The expected number of mispredictions is asymptotically  $\frac{n}{4}$ !

## Toy example: conclusion

- We proposed two solutions to this simple problem
- For uniform random permutations, in expectation:
  - ▶ The **naive algorithm** uses  $2n$  comparisons and  $2 \log n$  mispredictions
  - ▶ The **optimal algorithm** uses  $\frac{3}{2}n$  comparisons and  $\frac{1}{4}n$  mispredictions
  - ▶ Experimentally, the **naive algorithm** is more efficient!

## Toy example: conclusion

- We proposed two solutions to this simple problem
- For uniform random permutations, in expectation:
  - ▶ The **naive algorithm** uses  $2n$  comparisons and  $2 \log n$  mispredictions
  - ▶ The **optimal algorithm** uses  $\frac{3}{2}n$  comparisons and  $\frac{1}{4}n$  mispredictions
  - ▶ Experimentally, the **naive algorithm** is more efficient!

*We have to add branch predictors to our model of computation (RAM model) to fully describe the complexity of some algorithms.*

# Some Related Works

- Brodal & Moruz, 2005 : mispredictions and (adaptive) sorting

## Tradeoffs Between Branch Mispredictions and Comparisons for Sorting Algorithms

Gerth Stølting Brodal<sup>1,\*</sup> and Gabriel Moruz<sup>1</sup>

BRICS<sup>\*\*</sup>, Department of Computer Science, University of Aarhus,  
IT Parken, Åbogade 34, DK-8200 Århus N, Denmark  
{gerth, gabi}@daimi.au.dk

**Abstract.** Branch mispredictions is an important factor affecting the running time in practice. In this paper we consider tradeoffs between the number of branch mispredictions and the number of comparisons for sorting algorithms.

sorting  
algorithm  
mispred  
by ad  
tions. F  
rithm p  
 $\Omega(n \log$   
of inver  
by Esti  
col and  
mispred

Measure	Comparisons	Branch mispredictions
Dis	$O(dn(1 + \log(1 + \text{Dis})))$	$\Omega(n \log_d(1 + \text{Dis}))$
Exc	$O(dn(1 + \text{Exc} \log(1 + \text{Exc})))$	$\Omega(n \text{Exc} \log_d(1 + \text{Exc}))$
Enc	$O(dn(1 + \log(1 + \text{Enc})))$	$\Omega(n \log_d(1 + \text{Enc}))$
Inv	$O(dn(1 + \log(1 + \text{Inv}/n)))$	$\Omega(n \log_d(1 + \text{Inv}/n))$
Max	$O(dn(1 + \log(1 + \text{Max})))$	$\Omega(n \log_d(1 + \text{Max}))$
Osc	$O(dn(1 + \log(1 + \text{Osc}/n)))$	$\Omega(n \log_d(1 + \text{Osc}/n))$
Reg	$O(dn(1 + \log(1 + \text{Reg})))$	$\Omega(n \log_d(1 + \text{Reg}))$
Rem	$O(dn(1 + \text{Rem} \log(1 + \text{Rem})))$	$\Omega(n \text{Rem} \log_d(1 + \text{Rem}))$
Runs	$O(dn(1 + \log(1 + \text{Runs})))$	$\Omega(n \log_d(1 + \text{Runs}))$
SMS	$O(dn(1 + \log(1 + \text{SMS})))$	$\Omega(n \log_d(1 + \text{SMS}))$
SUS	$O(dn(1 + \log(1 + \text{SUS})))$	$\Omega(n \log_d(1 + \text{SUS}))$

**Fig. 4.** Lower bounds on the number of branch mispredictions for deterministic comparison based adaptive sorting algorithms for different measures of presortedness, given the upper bounds on the number of comparisons

# Some Related Works

- Brodal & Moruz, 2005 : mispredictions and (adaptive) sorting
- Biggar et al, 2008 : experimental, branch prediction and sorting

## An Experimental Study of Sorting and Branch Prediction

PAUL BIGGAR<sup>1</sup>, NICHOLAS NASH<sup>1</sup>, KEVIN WILLIAMS<sup>2</sup> and DAVID GREGG<sup>1</sup>  
Trinity College Dublin

Sorting is one of the most important and well studied problems in Computer Science. Many good

algorithms are known for other factors. However, architectures that support different features, and while of general purpose properties. In this common sorting algorithm, the predictability of the branch mispredictions of an algorithm in a fashion which sort's branches may have an effect on mergesort. We examine the choice point out a simple and show also that predictability of its branch predictors is that two-level adaptive Categories and Systems Organizations  
General Terms: Algorithms  
Additional Key Words

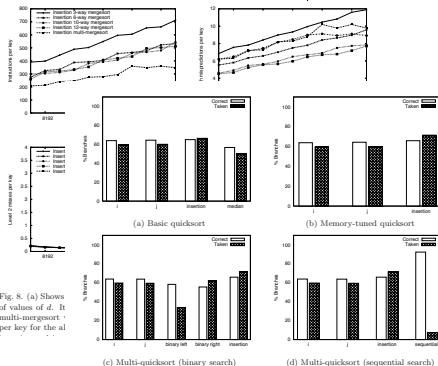


Fig. 8. (a) Shows values of  $d$ . It multi-mergesort per key for the algorithm.

Fig. 9. Overview of branch prediction behaviour in our quicksort implementations. Every figure

# Some Related Works

- Brodal & Moruz, 2005 : mispredictions and (adaptive) sorting
- Biggar et al, 2008 : experimental, branch prediction and sorting
- Sanders and Winkel, 2004 : quicksort variant without branches

## Super Scalar Sample Sort


Peter Sanders<sup>1</sup> and Sebastian Winkel<sup>2</sup>

<sup>1</sup> Max Planck Institut für Informatik  
Saarbrücken, Germany, [sanders@mpi-sb.mpg.de](mailto:sanders@mpi-sb.mpg.de)

<sup>2</sup> Chair for Prog. Lang. and Compiler Construction  
Saarland University, Saarbrücken, Germany, [sewi@cs.uni-sb.de](mailto:sewi@cs.uni-sb.de)

**Abstract.** Sample sort, a generalization of quicksort that partitions the input into many pieces, is known as the best practical comparison based sorting algorithm for distributed memory parallel computers. We show that

```
mic i t: = (sk/2, sk/4, s3k/4, sk/8, s3k/8, s5k/8, s7k/8, ...) //
condi for i := 1 to n do // locate each element
facili j := 1 // current tree node := root
final repeat log k times // will be unrolled
ber o j := 2j + (ak > tj) // left or right?
Itani j := j - k + 1 // bucket index
the (|bj|)++ // count bucket size
quick o(i) := j // remember oracle
```



**Fig. 2.** Finding buckets using implicit search trees. The picture is for  $k = 8$ . We adopt the C convention that “ $x > y$ ” is one if  $x > y$  holds, and zero else.

# Some Related Works

- Brodal & Moruz, 2005 : mispredictions and (adaptive) sorting
- Biggar [et al](#), 2008 : experimental, branch prediction and sorting
- Sanders and Winkel, 2004 : quicksort variant without branches
- Elmasry [et al](#), 2012 : mergesort variant without branches

## Branch Mispredictions Don't Affect Mergesort\*

Amr Elmasry<sup>1</sup>, Jyrki Katajainen<sup>1,2</sup>, and Max Stenmark<sup>2</sup>

<sup>1</sup> Department of Computer Science, University of Copenhagen  
Universitetsparken 1, 2100 Copenhagen East, Denmark

<sup>2</sup> Jyrki Katajainen and Company  
Thorsgade 101, 2200 Copenhagen North, Denmark

**Abstract.** In quicksort, due to branch mispredictions, a skewed pivot-selection strategy can lead to a better performance than the exact-median pivot-selection strategy, even in the worst case. This is not free. In this paper we investigate the behaviour of mergesort. By avoiding most negative branches, we can avoid most negative mispredictions. When sorting a sequence of  $n$  elements, mergesort performs  $n \log_2 n + O(n)$  comparisons and most  $O(n)$  branch mispredictions.

```
1 while (p != t1 && q != t2) {
2   if (less(*q, *p)) {
3     s = q;
4     ++q;
5   }
6   else {
7     s = p;
8     ++p;
9   }
10 }
```

```
1 test:
2   done = (q == t2);
3   if (done) goto exit;
4 entrance:
5   x = *p;
6   s = p + 1;
7   y = *q;
8   t = q + 1;
9   smaller = less(y, x);
10  if (smaller) s = t;
11  if (smaller) q = t;
12  if (! smaller) p = s;
13  if (! smaller) y = x;
14  x = *r;
15  *r = y;
16  --s;
17  *s = x;
18  ++r;
19  done = (p == t1);
20  if (! done) goto test;
21 exit:
```

Table 3. The execution time [ns], the number of conditional branches, and the number of mispredictions, each per  $n \log_2 n$ , for two in-situ variants of mergesort.

Program $n$	In-situ std::stable_sort			In-situ mergesort		
	Time Per Ares	Branches	Mispredicts	Time Per Ares	Branches	Mispredicts
$2^{10}$	49.2 29.7	9.0	2.08	7.3 5.7	1.93	0.26
$2^{15}$	57.6 35.0	11.1	2.38	7.1 5.6	1.94	0.15
$2^{20}$	62.7 38.5	12.9	2.53	7.4 5.7	1.92	0.11
$2^{25}$	68.0 41.3	14.4	2.62	7.6 5.7	1.92	0.09



# Some Related Works

- Brodal & Moruz, 2005 : mispredictions and (adaptive) sorting
- Biggar et al, 2008 : experimental, branch prediction and sorting
- Sanders and Winkel, 2004 : quicksort variant without branches
- Elmasry et al, 2012 : mergesort variant without branches
- Kaligosi and Sanders, 2006 : mispredictions and quicksort

## How Branch Mispredictions Affect Quicksort

Kanala Kaligosi<sup>1</sup> and Peter Sanders<sup>2</sup>

<sup>1</sup> Max Planck Institute  
Saarbrücken  
kaligosi@mpi-sb.de  
<sup>2</sup> Universität  
sanders@mpi-sb.de

Table 1. Number of branch mispredictions

	random pivot	$\alpha$ -skewed pivot
static predictor	$\frac{\ln 2}{2} n \lg n + \mathcal{O}(n)$ , $\frac{\ln 2}{2} \approx 0.3466$	$\frac{\pi^2}{12} n \lg n + \mathcal{O}(n)$ , $\alpha < 1/2$ $\frac{1-\alpha}{H(\alpha)} n \lg n + \mathcal{O}(n)$ , $\alpha \geq 1/2$
1-bit predictor	$\frac{2 \ln 2}{3} n \lg n + \mathcal{O}(n)$ , $\frac{2 \ln 2}{3} \approx 0.4621$	$\frac{2\alpha(1-\alpha)}{H(\alpha)} n \lg n + \mathcal{O}(n)$
2-bit predictor	$\frac{28 \ln 2}{45} n \lg n + \mathcal{O}(n)$ , $\frac{28 \ln 2}{45} \approx 0.4313$	$\frac{2\alpha^4 - 4\alpha^3 + \alpha^2 + \alpha}{(1-\alpha)(1-\alpha)H(\alpha)} n \lg n + \mathcal{O}(n)$

**Abstract.** We explain the counter-intuitive fact that “good” pivots (close to the median) do not improve the performance of quicksort in the presence of branch mispredictions. We explain the counter-intuitive fact that “good” pivots (close to the median) do not improve the performance of quicksort in the presence of branch mispredictions.


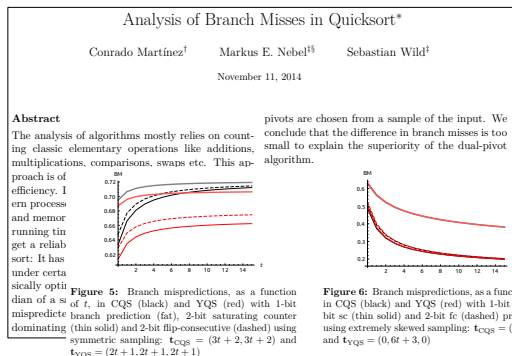


Fig. 3. Time /  $n \lg n$  for random pivot, median of 3, exact median, 1/10-skewed pivot

# Some Related Works

- Brodal & Moruz, 2005 : mispredictions and (adaptive) sorting
- Biggar et al, 2008 : experimental, branch prediction and sorting
- Sanders and Winkel, 2004 : quicksort variant without branches
- Elmasry et al, 2012 : mergesort variant without branches
- Kaligosi and Sanders, 2006 : mispredictions and quicksort
- Martínez, Nebel and Wild, 2014 : mispredictions and quicksort



# Some Related Works

- Brodal & Moruz, 2005 : mispredictions and (adaptive) sorting
- Biggar et al, 2008 : experimental, branch prediction and sorting
- Sanders and Winkel, 2004 : quicksort variant without branches
- Elmasry et al, 2012 : mergesort variant without branches
- Kaligosi and Sanders, 2006 : mispredictions and quicksort
- Martínez, Nebel and Wild, 2014 : mispredictions and quicksort
- Brodal and Moruz, 2006 : skewed binary search trees

## Skewed Binary Search Trees

Gerth Stølting Brodal<sup>1,\*</sup> and Gabriel Moruz<sup>1</sup>

BRICS<sup>\*\*</sup>, Department of Computer Science, University of Aarhus, IT Parken, Åbøgade 34, DK-8200 Århus N, Denmark. E-mail: {gerth, gabi}@dsi.imi.au.dk

**Abstract.** It is well-known that a binary search tree should be shown that a dominating factor for the number of cache faults per layout of a binary search tree by several hundred percent. branching to the left or right same cost, e.g. because of the study the class of skewed binary search tree the ratio  $t$  size of the tree is a fixed constant (trees). In this paper we present layouts of static skewed binary tree is accessed with a uniform many of the memory layouts perform better than perfect balanced search trees. The improvements in the running time are on the order of 15%.




Fig. 1. Bound on the expected cost for a random search, where the cost for visiting the left child is  $c_l = 1$  and the cost for processing the right child is  $c_r = 0, 1, 2, \dots, 28$  ( $c_r = 0$  being the lowest curve).

THE SKewed BInary SEARCH TREES (SABST) CLASS OF SEARCH TREES PERFORM BETTER THAN PERFECT BALANCED SEARCH TREES. THE IMPROVEMENTS IN THE RUNNING TIME ARE ON THE ORDER OF 15%.

# Some Related Works

- Brodal & Moruz, 2005 : mispredictions and (adaptive) sorting
- Biggar et al, 2008 : experimental, branch prediction and sorting
- Sanders and Winkel, 2004 : quicksort variant without branches
- Elmasry et al, 2012 : mergesort variant without branches
- Kaligosi and Sanders, 2006 : mispredictions and quicksort
- Martínez, Nebel and Wild, 2014 : mispredictions and quicksort
- Brodal and Moruz, 2006 : skewed binary search trees

## Skewed Binary Search Trees

Gerth Stølting Brodal<sup>1,\*</sup> and Gabriel Moruz<sup>1</sup>

BRICS<sup>\*,</sup> Department of Computer Science, University of Aarhus, IT Parken, Åbøgade 34, DK-8200 Århus N, Denmark. E-mail: {gerth, gabi}@dsi.imi.au.dk

**Abstract.** It is well-known that a binary search tree should be shown that a dominating factor for the number of cache faults per layout of a binary search tree by several hundred percent. branching to the left or right same cost, e.g. because of the study the class of skewed binary search tree the ratio  $t$  size of the tree is a fixed constant (trees). In this paper we present layouts of static skewed binary tree is accessed with a uniform many of the memory layouts perform better than perfect balanced search trees. The improvements in the running time are on the order of 15%.




Fig. 1. Bound on the expected cost for a random search, where the cost for visiting the left child is  $c_l = 1$  and the cost for processing the right child is  $c_r = 0, 1, 2, \dots, 28$  ( $c_r = 0$  being the lowest curve).

## What next?

- Branch predictors **exist** in computers
- They cannot easily be turned off
- **Classical paradigm:** ignore them, they are doing their job
- **AofA:** sometimes, it is necessary to take them into account

*What if we take them into account to **design** new algorithms?*

## Exponentiation by Squaring

We consider the classical **Exponentiation by Squaring** algorithm, and we **unroll** the main loop, to have two iterations each time.

pow(x,n)

```
r = 1;
while (n > 0) {
    // n is odd?
    if (n & 1)
        r = r * x;
    n /= 2;
    x = x * x;
}
```

$$x^n = (x^2)^{\lfloor n/2 \rfloor} x^{n_0}$$

unrolled(x,n)

```
r = 1;
while (n > 0) {
    t = x * x;
    // n0 == 1?
    if (n & 1)
        r = r * x;
    // n1 == 1?
    if (n & 2)
        r = r * t;
    n /= 4;
    x = t * t;
}
```

$$x^n = (x^4)^{\lfloor n/4 \rfloor} (x^2)^{n_1} x^{n_0}$$

## Exponentiation by Squaring

If  $n$  is taken uniformly at random in  $\{0, \dots, 4^k\}$ , then each `if` is taken with probability  $\frac{1}{2}$ : **it is difficult to predict.**

unrolled(x,n)

```
r = 1;
while (n > 0) {
    t = x * x;
    //  $n_0 == 1?$ 
    if (n & 1)  $\leftarrow \mathbb{P} = \frac{1}{2}$ 
        r = r * x;
    //  $n_1 == 1?$ 
    if (n & 2)  $\leftarrow \mathbb{P} = \frac{1}{2}$ 
        r = r * t;
    n /= 4;
    x = t * t;
}
```

# Exponentiation by Squaring

**Idea:** guide the predictors using a **unnecessary** test!

unrolled(x,n)

```
r = 1;
while (n > 0) {
    t = x * x;
    //  $n_0 == 1?$ 
    if (n & 1)  $\leftarrow \mathbb{P} = \frac{1}{2}$ 
        r = r * x;
    //  $n_1 == 1?$ 
    if (n & 2)  $\leftarrow \mathbb{P} = \frac{1}{2}$ 
        r = r * t;
    n /= 4;
    x = t * t;
}
```

guided(x,n)

```
r = 1;
while (n > 0) {
    t = x * x;
    //  $n_0 n_1 \neq 00?$ 
    if (n & 3) {  $\leftarrow \mathbb{P} = \frac{3}{4}$ 
        if (n & 1)  $\leftarrow \mathbb{P} = \frac{2}{3}$ 
            r = r * x;
        if (n & 2)  $\leftarrow \mathbb{P} = \frac{2}{3}$ 
            r = r * t;
    }
    n /= 4;
    x = t * t;
}
```

We have one more comparison by iteration, but predictions are easier.





# Exponentiation by Squaring

## Results:


- 25 % more comparisons for guided than for unrolled
- guided is 14% faster than unrolled
- yet, the number of multiplications is essentially the same.

## Analysis: Markov chains!





- The expected number of mispredictions after  $k$  steps in the Markov chain is asymptotically  $\mu(p)k$ , with  $\mu(p) = \frac{p(1-p)}{1-2p(1-p)}$ .
- The expected number of mispredictions in guided is  $\alpha \log_2 n$ , with  $\alpha = \frac{1}{2}\mu(3/4) + \frac{3}{4}\mu(2/3) = 0.45$




# Binary Search




## Binary Search



# Binary Search



# Binary Search



## Binary Search: Analysis

For arrays of size  $n$  filled with random uniform integers.  $C_n$  is the number of comparisons and  $M_n$  the number of mispredictions.

	BinarySearch	BiasedBinarySearch	SkewSearch
$\mathbb{E}[C_n]$	$\frac{\log n}{\log 2}$	$\frac{4 \log n}{4 \log 4 - 3 \log 3}$	$\frac{7 \log n}{6 \log 2}$
$\mathbb{E}[M_n]$	$\frac{\log n}{(2 \log 2)}$	$\mu(\frac{1}{4})\mathbb{E}[C_n]$	$(\frac{4}{7}\mu(\frac{1}{4}) + \frac{3}{7}\mu(\frac{1}{3}))\mathbb{E}[C_n]$

	BinarySearch	BiasedBinarySearch	SkewSearch
$\mathbb{E}[C_n]$	$1.44 \log n$	$1.78 \log n$	$1.68 \log n$
$\mathbb{E}[M_n]$	$0.72 \log n$	$0.53 \log n$	$0.58 \log n$

### Proof:

- Master Theorem gives the expected number of times each conditional is executed
- Ensure that our predictors behave *almost* like Markov chains.

# Branch Predictions: Conclusion

- Branch prediction mechanism alters the running time of algorithms
- It explains why the naive solution is better for the min/max problem
- We use it to **finely tune** classical algorithms:
  - ▶ Exponentiation by squaring is more efficient by **adding a useless test!**
  - ▶ Binary search is more efficient if we **don't cut in the middle!**
- The importance of branch mispredictions is limited to basic algorithms

*Thank you!*