

Why Is Dual-Pivot Quicksort Fast?

Sebastian Wild^(A)

^(A)Fachbereich Informatik, TU Kaiserslautern
wild@cs.uni-kl.de

Abstract

I discuss the new dual-pivot Quicksort that is nowadays used to sort arrays of primitive types in Java. I sketch theoretical analyses of this algorithm that offer a possible, and in my opinion plausible, explanation why (a) dual-pivot Quicksort is faster than the previously used (classic) Quicksort and (b) why this improvement was not already found much earlier.

1. Introduction

Quicksort became popular shortly after its original presentation by Hoare [7] and many authors contributed variations of and implementation tricks for the basic algorithm. From a practical point of view, the most notable improvements appear in [17, 15, 2, 14].

After the improvements to Quicksort in the 1990's, almost all programming libraries used almost identical versions of the algorithm: the classic Quicksort implementation had reached calm waters.

It was not until 2009, over a decade later, that previously unknown Russian developer Vladimir Yaroslavskiy caused a sea change upon releasing the outcome of his free-time experiments to the public: a dual-pivot Quicksort implementation that clearly outperforms the classic Quicksort in Oracle's Java 6. The core innovation is the arguably natural ternary partitioning algorithm given to the right.

Yaroslavskiy's finding was so surprising that people were initially reluctant to believe him, but his Quicksort has finally been deployed to millions of devices with the release of Java 7 in 2011.

How could this substantial improvement to the well-studied Quicksort algorithm escape the eyes of researchers around the world for nearly 50 years?

```
DUALPIVOTQUICKSORT(A, left, right) // sort A[left..right]
1  if right - left ≥ 1
    // Take outermost elements as pivots (replace by sampling)
2  p := min { A[left], A[right] }
3  q := max { A[left], A[right] }
4  ℓ := left + 1; g := right - 1; k := ℓ
5  while k ≤ g
6      if A[k] < p
7          Swap A[k] and A[ℓ]; ℓ := ℓ + 1
8      else if A[k] ≥ q
9          while A[g] > q and k < g
10             g := g - 1
11         end while
12         Swap A[k] and A[g]; g := g - 1
13         if A[k] < p
14             Swap A[k] and A[ℓ]; ℓ := ℓ + 1
15         end if
16     end if
17     k := k + 1
18 end while
19 ℓ := ℓ - 1; g := g + 1
20 A[left] := A[ℓ]; A[ℓ] := p // p to final position
21 A[right] := A[g]; A[g] := q // q to final position
22 DUALPIVOTQUICKSORT(A, left, ℓ - 1)
23 DUALPIVOTQUICKSORT(A, ℓ + 1, g - 1)
24 DUALPIVOTQUICKSORT(A, g + 1, right)
25 end if
```

For programs as heavily used as library sorting methods, it is advisable to back up experimental data with mathematically proven properties. The latter consider however only a *model* of reality, which may or may not reflect accurately enough the behavior of actual machines.

The answer to above question is in part a tale of the pitfalls of theoretical models, so we start with a summary of the mathematical analysis of Quicksort and the underlying model in Section 2. We then briefly discuss the “memory wall” metaphor and its implications for Quicksort in Section 3, and finally propose an alternative model in Section 4 that offers an explanation for the superiority of Yaroslavskiy’s Quicksort.

2. Analysis of Quicksort

The classical model for the analysis of sorting algorithm considers the average number of key comparisons on random permutations. Quicksort has been extensively studied under this model, including variations like choosing the pivot as median of a sample [7, 4, 15, 10, 3]: Let c_n denote the expected number of comparisons used by classic Quicksort (as given in [16]), when each pivot is chosen as median of a sample of $2t + 1$ random elements. c_n fulfills the recurrence

$$c_n = n - 1 + \sum_{\substack{0 \leq j_1, j_2 \leq n-1 \\ j_1 + j_2 = n-1}} \frac{\binom{j_1}{t} \binom{j_2}{t}}{\binom{n}{2t+1}} (c_{j_1} + c_{j_2}) \quad (1)$$

since $n - 1$ comparisons are needed in the first partitioning step, and we have two recursive calls of random sizes, where the probability to have sizes j_1 and j_2 is given by the fraction of binomials (see [9] for details). This recurrence can be solved asymptotically [4, 15] to

$$c_n \sim \frac{1}{H_{2(t+1)} - H_{t+1}} \cdot n \ln n,$$

where $H_n = \sum_{i=1}^n 1/i$ is the n th harmonic number and $f(n) \sim g(n)$ means $\lim_{n \rightarrow \infty} f(n)/g(n) = 1$. The mathematical details are beyond the scope of this abstract, but a rather elementary derivation is possible [10]. Large values of t are impractical; a good compromise in practice is given by the “ninter”, the median of three medians, each chosen from three elements [2]. This scheme can be analyzed similarly to the above [3].

The author generalized Equation (1) to Yaroslavskiy’s Quicksort [19, 18, 9]. Note that unlike for classic Quicksort, the comparison count of Yaroslavskiy’s partitioning depends on pivot values, so its expected value has to be computed over the choices for the pivots. We obtain for tertiles-of- $(3t + 2)$

$$c_n = \left(\frac{5}{3} - \frac{1}{9t+12} \right) \cdot n + O(1) + \sum_{\substack{0 \leq j_1, j_2, j_3 \leq n-2 \\ j_1 + j_2 + j_3 = n-2}} \frac{\binom{j_1}{t} \binom{j_2}{t} \binom{j_3}{t}}{\binom{n}{3t+2}} \cdot (c_{j_1} + c_{j_2} + c_{j_3}); \quad (2)$$

with solution

$$c_n \sim \frac{\frac{5}{3} - \frac{1}{9t+12}}{H_{3(t+1)} - H_{t+1}} \cdot n \ln n.$$

Oracle’s Java runtime library previously used classic Quicksort with *n*inther, and now uses Yaroslavskiy’s Quicksort with tertiles-of-five; the average number of comparisons are asymptotically $1.5697n \ln n$ vs. $1.7043n \ln n$. According to the comparison model, Yaroslavskiy’s algorithm is significantly *worse* than classic Quicksort! Moreover, this result does not change qualitatively if we consider *all* primitive instructions of a machine instead of only comparisons [19, 9]. It is thus not surprising that researchers found the use of two pivots not promising [15, 6].

But if Yaroslavskiy’s algorithm actually uses more comparisons and instructions, how comes it is still faster in practice? And why was this discrepancy between theory and practice not noticed earlier? The reason is most likely related to a phenomenon known as the “memory wall” [20, 13] or the “von-Neumann bottleneck” [1]: Processor speed has been growing considerably faster than memory bandwidth for a long time.

3. The Memory Wall

Based on the extensive data for the STREAM benchmark [12, 11], CPU speed has increased with an average annual growth rate of 46% over the last 25 years, whereas memory bandwidth, the amount of data transferable between RAM and CPU in a given amount of time, has increased by 37% per year. Even though one should not be too strict about the exact numbers as they are averages over very different architectures, a significant increase in *imbalance* is undeniable. Figure 1 gives a direct quantitative view of this trend.

If the imbalance between CPU and memory transfer speed continues to grow exponentially, at some point in the future any further improvements of CPUs will be futile: the processor is waiting for data all the time; we hit a “memory wall”.

It is debatable if and when this extreme will be reached [5, 13], and consequences certainly depend on the application. In any case, however, the (average) relative costs of memory accesses have increased significantly over the last 25 years.

So when Quicksort with two pivots was first studied, researchers correctly concluded that it does not pay off. But computers have changed since then, and so must our models.

4. Scanned Elements

Our new cost model for sorting counts the number of “*scanned elements*”. An element scan is essentially an accesses “ $A[i]$ ” to the input array A , but we count all accesses as one that use the same index variable i and the same value for i . For example, a linear scan over A




Figure 1: Development of CPU speed against memory bandwidth over the last 25 years. Each point shows one reported result of the STREAM benchmark [12, 11], with the date on the x -axis and the machine balance (peak MFLOPS divided by Bandwidth in MW/s in the “triad” benchmark) on a logarithmic y -axis. The fat line shows the linear regression (on log-scale). Data is taken from www.cs.virginia.edu/stream/by_date/Balance.html.

entails n scanned elements, and several interleaved scans (with different index variables) cost the traveled distance, summed up over all indices, even when the scanned ranges overlap. We do not distinguish read and write accesses.

We claim that for algorithms built on potentially interleaved sequential scans, in particular for classic and dual-pivot Quicksort, the number of scanned elements is asymptotically proportional to the amount of data transferred between CPU and main memory [9].

Scanned elements are related to cache misses [8], but the latter is a machine-dependent quantity, whereas the former is a clean, abstract cost measure that is easy to analyze: One partitioning step of classic Quicksort scans A exactly once, resulting in n scanned elements. In Yaroslavskiy's partitioning, indices k and g together scan A once, but index ℓ scans the leftmost segment a second time. On average, the latter contains a third of all elements, yielding $\frac{4}{3}n$ scanned elements in total.

Using these in recurrences (1) resp. (2) yields $1.5697n \ln n$ vs. $1.4035n \ln n$ scanned elements; the Java 7 Quicksort saves 12% of the element scans over the version in Java 6, which matches the roughly 10% speedup observed in running time studies.

5. Conclusion

Memory speed has not fully kept pace with improvements in processing power. This growing imbalance forces us to economize on memory accesses in algorithms that were almost entirely CPU-bound in the past, and calls for new cost models for the analysis of algorithms. For sorting algorithms that build on sequential scans over their input, the proposed “scanned elements” counts serve as such a model and give a good indication of the amount of memory traffic caused by an algorithm. It is exactly this data traffic where dual-pivot outclasses classic Quicksort, offering a plausible explanation for its superiority in practice.

References

- [1] J. BACKUS, Can Programming Be Liberated from the Von Neumann Style? A Functional Style and Its Algebra of Programs. *Communications of the ACM* **21** (1978) 8, 613–641.
- [2] J. L. BENTLEY, M. D. MCILROY, Engineering a sort function. *Software: Practice and Experience* **23** (1993) 11, 1249–1265.
- [3] M. DURAND, Asymptotic analysis of an optimized quicksort algorithm. *Information Processing Letters* **85** (2003) 2, 73–77.
- [4] M. H. VAN EMDEN, Increasing the efficiency of quicksort. *Communications of the ACM* (1970), 563–567.
- [5] M. A. ERTL, The Memory Wall Fallacy. 2001.
<https://www.complang.tuwien.ac.at/anton/memory-wall.html>
- [6] P. HENNEQUIN, *Analyse en moyenne d'algorithmes : tri rapide et arbres de recherche*. PhD Thesis, Ecole Polytechnique, Palaiseau, 1991.
- [7] C. A. R. HOARE, Quicksort. *The Computer Journal* **5** (1962) 1, 10–16.

-
- [8] S. KUSHAGRA, A. LÓPEZ-ORTIZ, A. QIAO, J. I. MUNRO, Multi-Pivot Quicksort: Theory and Experiments. In: *ALLENEX 2014*. SIAM, 2014, 47–60.
- [9] C. MARTÍNEZ, M. E. NEBEL, S. WILD, Analysis of Pivot Sampling in Dual-Pivot Quicksort. *Algorithmica* (2015).
<http://arxiv.org/abs/1412.0193>
- [10] C. MARTÍNEZ, S. ROURA, Optimal Sampling Strategies in Quicksort and Quickselect. *SIAM Journal on Computing* **31** (2001) 3, 683–705.
- [11] J. D. MCCALPIN, *STREAM: Sustainable Memory Bandwidth in High Performance Computers*. Technical report, University of Virginia, Charlottesville, Virginia, 1991-2007. Continually updated technical report.
<http://www.cs.virginia.edu/~mccalpin/papers/bandwidth/bandwidth.html>
- [12] J. D. MCCALPIN, Memory Bandwidth and Machine Balance in Current High Performance Computers. *IEEE Computer Society Technical Committee on Computer Architecture (TCCA) Newsletter* (1995), 19–25.
<http://www.cs.virginia.edu/~mccalpin/papers/balance/index.html>
- [13] S. A. MCKEE, Reflections on the Memory Wall. In: *Proceedings of the first conference on computing frontiers*. 2004, 162–167.
- [14] D. R. MUSSER, Introspective Sorting and Selection Algorithms. *Software: Practice and Experience* **27** (1997) 8, 983–993.
- [15] R. SEDGEWICK, *Quicksort*. Ph. D. Thesis, Stanford University, 1975.
- [16] R. SEDGEWICK, Implementing Quicksort programs. *Communications of the ACM* **21** (1978) 10, 847–857.
- [17] R. C. SINGLETON, Algorithm 347: an efficient algorithm for sorting with minimal storage [M1]. *Communications of the ACM* **12** (1969) 3, 185–186.
- [18] S. WILD, *Java 7's Dual Pivot Quicksort*. Master thesis, University of Kaiserslautern, 2012.
<http://nbn-resolving.de/urn/resolver.pl?urn:nbn:de:hbz:386-kluedo-34638>
- [19] S. WILD, M. E. NEBEL, Average Case Analysis of Java 7's Dual Pivot Quicksort. In: L. EPSTEIN, P. FERRAGINA (eds.), *ESA 2012*. LNCS 7501, Springer, 2012, 825–836.
<http://arxiv.org/abs/1310.7409>
- [20] W. A. WULF, S. A. MCKEE, Hitting the Memory Wall: Implications of the Obvious. 1995.