

A Local Limit Theorem for QuickSort Key Comparisons

Jim Fill

Department of Applied Mathematics and Statistics
The Johns Hopkins University

June 20, 2017

Analysis of Algorithms at Princeton University

This is joint work with Béla Bollobás and Oliver Riordan.

QuickSort: The Algorithm (Hoare, early 1960s)

- Assume distinct keys (numbers—or perhaps symbol strings—to be sorted).
- Choose “pivot” key uniformly at random.
- Use pivot to partition into two subsets: smaller and larger.
- Apply QuickSort recursively to subsets. Measure runtime by

$$\begin{aligned} K_n &:= \text{number of key comparisons, with } K_0 = 0 \\ &\stackrel{\mathcal{L}}{=} K_{U_n-1} + K_{n-U_n}^* + n - 1, \quad n \geq 1, \end{aligned}$$

where on the RHS:

$$U_n \sim \text{unif}\{1, \dots, n\}$$

and

$$U_n; K_0, \dots, K_{n-1}; K_0^*, \dots, K_{n-1}^*$$

are all *independent*.

- Applied focus: the law $\mathcal{L}(K_n)$ and asymptotics as $n \rightarrow \infty$

The importance of QuickSort

- In a special issue of *Computing in Science & Engineering* (2000), guest editors Jack Dongarra and Francis Sullivan chose QuickSort as one of the ten algorithms “with the greatest influence on the development and practice of science and engineering in the 20th century.”
- QuickSort is the standard sorting procedure in Unix systems.
- QuickSort is among “some of the most basic algorithms—the ones that deserve deep investigation.” — Ph. Flajolet (1999)
- “probably most widely used sorting algorithm” — U. Rösler (1991)
- “one of the fastest, the best-known, the most generalized, the most completely analyzed, and the most widely used algorithms for sorting” — W. F. Eddy and M. J. Schervish (1995) (But *much* more analysis has followed!)

1: Basics of QuickSort analysis

Conditioning on the pivot index U_n and using the law of total expectation, the distributional recurrence

$$K_n \stackrel{\mathcal{L}}{=} K_{U_n-1} + K_{n-U_n}^* + n - 1, \quad n \geq 1$$

implies a simple **divide and conquer** recurrence relation for expected values with explicit solution

$$\kappa_n := \mathbf{E}K_n = 2(n+1)H_n - 4n, \quad n \geq 0.$$

We have

$$\kappa_n \sim 2n \ln n.$$

The law of total variance gives another recurrence relation, with solution

$$\begin{aligned} \mathbf{Var} K_n &= 7n^2 - 4(n+1)^2 H_n^{(2)} - 2(n+1)H_n + 13n \\ &\sim \left(7 - \frac{2}{3}\pi^2\right)n^2 =: \sigma^2 n^2 \end{aligned}$$

[e.g., Exercise 6.2.2-8 in **D. E. Knuth (1973, vol. 3)**].

- Higher-order cumulants: **Pascal Hennequin (1991 PhD diss.)**

2: Convergence in distribution: heuristic derivation

Is there convergence in distribution? **YES!:**

$$Y_n \xrightarrow{\mathcal{L}} Y \stackrel{\mathcal{L}}{=} UY + (1 - U)Y^* + 2U \ln U + 2(1 - U) \ln(1 - U) + 1$$

with

$$\mathbf{E}Y = 0 \quad \text{and} \quad \mathbf{Var} Y = \sigma^2 = 7 - \frac{2}{3}\pi^2 < \infty.$$

This is **TRUE!**, and at least three methods of proof are possible:

Method 1: $Y_n \xrightarrow{\mathcal{L}} Y$ by *method of moments* [Hennequin (1991)]

Method 2: $Y_n \rightarrow Y$ a.s. and in $L^p \forall 0 < p < \infty$ by *martingale* arguments [M. Régnier (1989)]

Method 3: $Y_n \rightarrow Y$ in *Mallows* $d_p \quad \forall p = 1, 2, \dots$ by *contraction method* [Rösler (1991)]

OPEN PROBLEM 6.2 in F and Janson (2002):

Is there a *local* limit theorem for distns.?

YES!: this talk, based on Bollobás, F, and Riordan (2017)

3: Selected contributions to analysis of QS (2002-)

The following contributions are highlighted because they are relevant to establishment of a LLT!

- K. H. Tan and P. Hadjicostas (1995) — limit distribution is absolutely continuous; density is positive a.e.
- F and S. Janson (*Mathematics and Computer Science: Algorithms, Trees, Combinatorics, and Probabilities*, 2000) — smoothness and decay properties of the limiting density and its derivatives; in particular (Theorems 3.1 and 3.3 and Corollary 4.2 there), there is an infinitely smooth (analytic?) everywhere positive limiting density f satisfying

$$\max_x f(x) < 16 \quad \text{and} \quad \max_x f'(x) < 2466$$

(with the true bounds probably closer to 1 and 2, resp.).

Aside: These bounds were one ingredient used by L. Devroye, F, and R. Neininger (2000) to devise an algorithm for perfect simulation from the limiting distribution F for QuickSort.

Plot of the (smooth) density function f for F

Kok Hooi Tan, P. Hadjicostas / Statistics & Probability Letters 25 (1995) 87–94




Fig. 4. Density function $f_X(x)$ of limiting distribution.

Thinness of the tails of F (rigorous)

- Invariably, someone asks: How thin are the tails of F ?
- A **rigorously** established answer, sharp to lead order in log-probability, was provided by **S. Janson (Electron. J. Probab., 2015)**, improving significantly on results of **F and S. Janson (RSA, 2000)**.
- Let $Y \sim F$. The left tail is doubly-exponentially thin, and the right tail is Poisson(1)-thin:

Theorem (**Janson, 2015**)

(a) Let $\gamma := (2 - \frac{1}{\ln 2})^{-1}$. As $x \rightarrow \infty$ we have

$$\exp \left[-e^{\gamma x + \ln \ln x + O(1)} \right] \leq \mathbf{P}(Y \leq -x) \leq \exp \left[-e^{\gamma x + O(1)} \right].$$

(b) As $x \rightarrow \infty$ we have

$$\exp \left[-x \ln x - x \ln \ln x + O(x) \right] \leq \mathbf{P}(Y \geq x) \leq \exp \left[-x \ln x + O(x) \right].$$

Thinness of the tails of F (non-rigorous)

- Using the non-rigorous WKB method, C. Knessl and W. Szpankowski (*Discrete Math. Theor. Comput. Sci.*, 1999) found very sharp log-asymptotics for both tails:

(a) As $x \rightarrow \infty$, for some constant c we have

$$P(Y \leq -x) = \exp \left[-e^{\gamma x + c + o(1)} \right].$$

(b) As $x \rightarrow \infty$ we have

$$P(Y \geq x) = \exp \left[-x \ln x - x \ln \ln x + (1 + \ln 2)x + o(x) \right].$$

OPEN PROBLEMS: Prove that f is unimodal. Is f in fact *strongly* unimodal? What can one say about changes of signs of the derivatives of f ? Is F infinitely divisible?

3: Selected contributions to analysis of QS (2002-)

Let F_n denote the distribution function of $Y_n = (K_n - \kappa_n)/n$. The following two results about convergence of F_n to its limit F are from **F and S. Janson** (*J. of Algorithms*, 2002), and the upper bounds are proved by first treating d_p -distances inductively using the *contraction method* à la **Rösler's** (1991) proof of convergence and then relating *Kolmogorov-Smirnov distance* to d_p -distances:

- We have $F_n \rightarrow F$ at rate $O\left(n^{-(\frac{1}{2}-\epsilon)}\right)$ in the K-S distance (for every $\epsilon > 0$); but our lower bound is only $\Omega(n^{-1})$ (and **improving either bound** in this global theorem remains an **open problem**).
- There is a constant C such that, for any x and any $n \geq 1$,

$$\left| \frac{F_n(x + \frac{\delta_n}{2}) - F_n(x - \frac{\delta_n}{2})}{\delta_n} - f(x) \right| \leq Cn^{-1/6},$$

where $\delta_n = 2Cn^{-1/6}$ (w/ no claim of sharpness in the bound).

4: From a semi-local LT to a local LT

Recall that F_n denotes the distribution function of

$$Y_n = (K_n - \kappa_n)/n$$

and that we have the following **semi-local limit theorem**:

- There is a constant C such that for any x and any $n \geq 1$ we have

$$\left| \frac{F_n(x + \frac{\delta_n}{2}) - F_n(x - \frac{\delta_n}{2})}{\delta_n} - f(x) \right| \leq Cn^{-1/6},$$

where $\delta_n = 2Cn^{-1/6}$.

Our **new result** is that δ_n can be decreased to n^{-1} , with the error bound increasing from $O(n^{-1/6})$ to a little more than $O(n^{-1/18})$. Expressed in terms of the distribution of the unnormalized comparisons-count K_n , the result is as follows (**again with no claim of sharpness in the bound**).

Main theorem: A LLT for QuickSort

The following **local limit theorem** is our main result and gives a positive answer to **Open Problem 6.2** in **F and S. Janson (J. of Algorithms, 2002)**:

Theorem (Local Limit Theorem for QuickSort)

Let K_n denote the (random) number of key comparisons required by **QuickSort** to sort a file of n distinct keys, let $\kappa_n = \mathbf{E}K_n$, and let f denote the continuous density of the limiting distribution for the normalized random variable $Y_n = (K_n - \kappa_n)/n$. Then there is a constant C such that for any integer k and any $n \geq 1$ we have

$$|\mathbf{P}(K_n = k) - n^{-1}f((k - \kappa_n)/n)| \leq n^{-1} \times Cn^{-1/18} \log n.$$

We obtain the LLT from the semi-LLT by multiple rounds of a **smoothing argument**. The proof is too complex for a 20-minute talk, but I will give some ideas.

5: Obtaining the LLT from the semi-LLT: main ideas

We obtain the LLT from the semi-LLT by multiple rounds of a **smoothing argument**. The basic idea of strengthening a distributional (often normal) limit theorem to a local one by smoothing is by now quite old, but we find it necessary (and sufficient) to do this smoothing in multiple rounds. [Multi-round smoothing has recently been used independently by **Diaconis and Hough (2015)** in a different context.]

- We know that there is global (and, indeed, semi-local) convergence to a well-behaved distribution. (GLT)
- To deduce a LLT from the GLT it would suffice to show that “nearby” values for K_n have probabilities that are close, namely: If $k_n, k'_n = \kappa_n + O(n)$ and $|k_n - k'_n| = o(n)$, then

$$\mathbf{P}(K_n = k_n) = \mathbf{P}(K_n = k'_n) + o(n^{-1}). \quad (1)$$

- For this, in turn, we might (as in **D. R. McDonald (1979)**) try to find a “smooth part” within the distribution of K_n . More precisely, ...

Obtaining the LLT from the semi-LLT: main ideas

- We might try to find a “smooth part” within the distribution of K_n . More precisely, we might try to write

$$K_n = A_n + B_n$$

where, for some σ -field \mathcal{F}_n , we have that A_n is \mathcal{F}_n -measurable and the conditional distribution of B_n given \mathcal{F}_n (wvhp) obeys a “closeness” relation corresponding to (1).

- Then it follows easily (by first considering conditional probabilities given \mathcal{F}_n) that (1) holds.
- One idea is to choose \mathcal{F}_n so that B_n has a very well understood conditional distribution, such as binomial.
- In some contexts, this approach works directly. Here it does not. We can get such a decomp. with (conditionally) $B_n \sim \text{Bin}(\Theta(n), 2/3)$. But then the variance of B_n grows linearly, while the variance of K_n grows quadratically. Roughly speaking, this allows us to prove (1) when $|k_n - k'_n| = o(\sqrt{n})$, but we need (1) for $|k_n - k'_n| = o(n)$ [or, in light of the semi-LLT, at least for $|k_n - k'_n| = O(n^{5/6})$].

6: The tree-exploration lemma

The key idea, then, is not to try to jump straight from the GLT (or from the semi-LLT) to the LLT, but to proceed in stages (rounds). Rather than outline the entire argument (which makes use of the CLT for sums of independent random variables after suitable exponential tilting, among many other things), I will discuss a **tree-exploration** lemma (along with its proof using **martingale arguments**) that lies at the heart of the argument. Let $c := 1/100$.

Lemma (tree-exploration lemma)

Let $r \geq 2$ be even, and assume that $r \leq n/75$. Then setting $s = \lceil cn/r \rceil$, we may write $K_n = A + B$ where, for some σ -field \mathcal{F} :

- *A is \mathcal{F} -measurable, and*
- *with probability at least $1 - e^{-s}$, the conditional distribution of B given \mathcal{F} is the sum of s independent random variables B_1, \dots, B_s with each B_i having the distribution K_{r_i} for some r_i with $r/2 \leq r_i \leq r$.*

7: Proof of the tree-exploration lemma

The remainder of the talk is devoted to the proof of the **tree-exploration lemma**. Before any **claims**, much preliminary discussion:

- T_n = random binary search tree for n nodes labeled $1, \dots, n$.
- $T(v)$ denotes the subtree of T_n consisting of a given node v and its descendants (the “**fringe**” **subtree** of T_n rooted at v). We consider two counts, X_n (primary) and Y_n (auxiliary):
- $X_n :=$ number of nodes v such that (i) $r/2 \leq |T(v)| \leq r$ and (ii) either v is the root or $|T(\text{parent of } v)| > r$. Call such a node **special**.
- Observe: If v and w are distinct special nodes, then $T(v)$ and $T(w)$ are disjoint.
- We will show that, with probability at least $1 - e^{-s}$, there are at least s special nodes in T_n , i.e.: $\mathbf{P}(X_n \geq s) \geq 1 - e^{-s}$.
- $Y_n :=$ number of fringe trees with at least $r + 1$ nodes.

Proof of the tree-exploration lemma (cont.): exploration

For convenience in presenting the arguments that follow, attach (unlabeled) external nodes as necessary to each node of T_n so that as a result each of the original nodes of T_n has two children; these external nodes do not have labels and do not contribute to the sizes of subtrees.

Consider the following discrete-time procedure for **exploring** the original nodes of T_n in order to learn their search-tree labels.

- At time 0 nothing is known.
- At time 1, the label of the root is revealed, thus also revealing the sizes of the left and right subtrees.

Now let $2 \leq t \leq n$

Proof of the tree-exploration lemma: exploration (cont.)

Now let $2 \leq t \leq n$. At time $t - 1$, it will be true inductively that nodes with revealed labels will have fringe subtree size at least $r + 1$ and the subtree sizes of their two children will be known.

- At time t , choose an unrevealed child v of a revealed node such that $|T(v)| \geq r + 1$ and reveal its label, if such a v exists.
 - For definiteness, among such nodes with smallest level, choose the leftmost one.
 - If there are no such nodes v , then nothing is revealed at time t (nor at later times).
 - At any time $t \in [1, n]$, call the unrevealed children (including all external nodes) of revealed nodes “leaves”.
- Note that each time the label of a node is revealed, two new leaves (namely, the children of that node) are created; for later technical reasons, we consider the left child to be created as a leaf before the right child.

Proof of the tree-exploration lemma (cont.): martingales

- $\tau := Y_n$ is the random time at which the procedure ends, i.e., the first time at which all of the leaves have subtree sizes $\leq r$.
- $\mathcal{F}_t :=$ the σ -field corresponding to the labels that have been revealed through time t .
- **Observe!:** For any random variable W with finite expectation, the stochastic process $(\mathbf{E}[W|\mathcal{F}_t])_{0 \leq t \leq n}$ is a (Doob's) **martingale** with values $\mathbf{E}W, W$ at times $t = 0, n$, resp.
- We will apply this observation taking W to be X_n [with martingale (M_t)] and Y_n [with martingale (N_t)].
Clearly $N_\tau = Y_n = \tau$.
- The respective means $\xi_n := \mathbf{E}X_n$ and $\eta_n := \mathbf{E}Y_n$ of the martingales (M_t) and (N_t) can be computed by solving standard **divide-and-conquer recurrence relations**.
- From those computations it is simple to deduce that

$$-2 < N_t - N_{t-1} \leq 1, \quad -\frac{1}{2} < M_t - M_{t-1} < 1$$

for every t .

Proof of the tree-exploration lemma (cont.): bound on τ

First claim: $\mathbf{P}(\tau < \frac{3n}{r+1}) \geq 1 - \exp\left[-\frac{2}{27} \left(\frac{n}{r+1}\right)\right]$.

Proof: Recalling that $\tau = N_\tau$, we have

$$\mathbf{P}\left(\tau \geq \frac{3n}{r+1}\right) \leq \mathbf{P}\left(N_t = t \text{ for some } t \geq \frac{3n}{r+1}\right).$$

By the **generalized Azuma inequality** we have

$$\begin{aligned} & \mathbf{P}\left(\tau \geq \frac{3n}{r+1}\right) \\ & \leq \mathbf{P}\left(N_t - \eta_n = t - \left\lfloor \frac{2(n+1)}{r+2} - 1 \right\rfloor \text{ for some } t \geq \frac{3n}{r+1}\right) \\ & \leq \mathbf{P}\left(N_t - \eta_n \geq \frac{1}{3}t \text{ for some } t \geq \frac{3n}{r+1}\right) \\ & \leq \exp\left[-\frac{2}{81} \left\lfloor \frac{3n}{r+1} \right\rfloor\right] \\ & \leq \exp\left[-\frac{2}{27} \left(\frac{n}{r+1}\right)\right] =: \varepsilon. \quad \square \end{aligned}$$

Pf. of T-E L (cont.): whp, X_n is not too small rel. to n/r

First claim: $\mathbf{P}(\tau < \frac{3n}{r+1}) \geq 1 - \exp \left[-\frac{2}{27} \left(\frac{n}{r+1} \right) \right]$. □

Second claim: $\mathbf{P} \left(X_n \geq \frac{(1/100)n}{r} \right) \geq 1 - \exp \left[- \left\lceil \frac{1}{100} \left(\frac{n}{r} \right) \right\rceil \right]$.

Proof: By the **first claim**, because $t \geq \tau$ implies $M_t = X_n$, we have

$$\mathbf{P} \left(X_n < \frac{(1/2)(n+1)}{r+1} \right) \leq \mathbf{P} \left(\tau \geq \frac{3n}{r+1} \right) + \mathbf{P} \left(M_{\lceil 3n/(r+1) \rceil} < \frac{(1/2)(n+1)}{r+1} \right).$$

Further, by **Azuma's inequality** we have

$$\begin{aligned} & \mathbf{P} \left(X_n < \frac{(1/100)n}{r} \right) \\ & \leq \mathbf{P} \left(X_n < \frac{(1/2)(n+1)}{r+1} \right) \leq \varepsilon + \mathbf{P} \left(M_{\lceil 3n/(r+1) \rceil} < \frac{(1/2)(n+1)}{r+1} \right) \\ & = \varepsilon + \mathbf{P} \left(M_{\lceil 3n/(r+1) \rceil} - \xi_n < -\frac{(1/2)(n+1)}{r+1} \right) \\ & \leq \varepsilon + \exp \left[-\frac{8}{9} \left[\frac{(1/2)(n+1)}{r+1} \right]^2 / \lceil 3n/(r+1) \rceil \right] \leq \varepsilon + \exp \left[-\frac{1}{18} \left(\frac{n}{r+1} \right) \right] \\ & \leq \exp \left[-\frac{4}{81} \left(\frac{n}{r} \right) \right] + \exp \left[-\frac{1}{27} \left(\frac{n}{r} \right) \right] \leq \exp \left[-\left\lceil \frac{1}{100} \left(\frac{n}{r} \right) \right\rceil \right], \text{ where} \end{aligned}$$

Pf. of T-E L: whp, X_n is not too small rel. to n/r (cont.)

- the orange ineq. holds because $\lceil 3n/(r+1) \rceil \leq 4n/(r+1)$,
- the blue inequality holds because $r+1 \leq (3/2)r$, and
- the green inequality holds because

$$e^{-4x/81} + e^{-x/27} \leq e^{-(x/100)+1} \leq e^{-\lceil x/100 \rceil} \text{ for } x \geq 75.$$

We have proven the second claim that the number X_n of special nodes is at least $s = \lceil (1/100)n/r \rceil$ wp at least $1 - e^{-s}$.

The proof of the tree-exploration lemma, and hence of the LLT, concludes on the next slide.

Proof of the tree-exploration lemma (conclusion)

- Let $\sigma \leq \tau$ denote
 - the stopping time at which the exploration process has discovered either s or (by means of having simultaneously discovered two special-leaf children with a single reveal) $s + 1$ special leaves, in the event E that such a time exists;
 - otherwise set $\sigma = n$.
- If the event E occurs, fully explore the subtrees rooted at all leaves except the first s special leaves that have been discovered by time σ .
- Take the desired σ -field \mathcal{F} to be (informally stated) the σ -field corresponding to all the information uncovered as we have just described.
- Taking B_i , $i = 1, \dots, s$, to be (when E occurs) the total (internal) path length (i.e., number of key comparisons for **QuickSort**) of the i th special leaf discovered by time σ , the **tree-exploration lemma** follows. □

I hope you enjoyed attending
this talk!