

DISTRIBUTION OF A CLASS OF DIVIDE AND CONQUER RECURRENCES ARISING FROM THE COMPUTATION OF THE WALSH-HADAMARD TRANSFORM

PAWEŁ HITCZENKO

DEPT. OF MATHEMATICS, DREXEL UNIVERSITY, PHILADELPHIA, PA 19104

JEREMY JOHNSON AND HUNG-JEN HUANG

DEPT. COMPUTER SCIENCE, DREXEL UNIVERSITY, PHILADELPHIA, PA 19104

ABSTRACT. This paper explores the performance of a family of algorithms for computing the Walsh-Hadamard transform, a useful computation in signal and image processing. Recent empirical work has shown that the family of algorithms exhibit a wide range of performance and that it is non-trivial to determine which algorithm is optimal on a given computer. This paper provides a theoretical basis for the performance distribution. Performance is modelled by a family of recurrence relations that determine the number of instructions required to execute a given algorithm, and the recurrence relations can be used to explore the performance of the space of algorithms. The recurrence relations are related to standard divide and conquer recurrences, however, there are a variable number of recursive parts which can grow to infinity as the input size increases. Thus standard approaches to solving such recurrences can not be used and new techniques must be developed. In this paper, the minimum, maximum, expected values, and variances are calculated and the limiting distribution is obtained.

1. INTRODUCTION

This work is motivated by the relatively new field of “Automated Performance Tuning” [1], where various techniques are used to automatically implement and optimize an algorithm on a specific computer platform. For many important algorithms, there is a large number of variations and implementation choices which can greatly affect performance. Moreover, the optimal choice among the variations and implementations is highly dependent on the underlying computer architecture and the compiler and compiler flags used to create the executable program. A particularly useful strategy for automated performance tuning, called generate-and-test, creates alternative implementation choices and uses intelligent search to find the best implementation on a given platform. This approach has been used effectively in the implementation and optimization of linear algebra [26, 4] and signal processing [9, 20, 21, 15] kernels.

The work of Paweł Hitczenko was supported in part by NSA Grant MSPF-02G-043. Part of the research of this author was carried out while he was visiting the University of Jyväskylä, Finland in the Fall of 2002. He would like to thank the Department of Mathematics and Statistics for the invitation, and Christel and Stefan Geiss for their hospitality. The work of Jeremy Johnson was supported by DARPA through research grant DABT63-98-1-0004 administered by the Army Directorate of Contracting.

Since this approach utilizes empirical run times, it is not necessary to have a model of the underlying computer architecture nor is it necessary to understand why a particular choice leads to good performance and another choice leads to poor performance. However, search based on empirical run times, can take a significant amount of time, and the lack of an analytic model leaves the significance of the optimization problem unclear. With better understanding, the search might prove unnecessary. The goal of this paper is to provide the first step towards a performance model and an analytic understanding of the search space for one particular problem where there is a large space of alternative algorithm variations and generate-and-test has proved successful.

The problem addressed in this paper is the computation of the Walsh-Hadamard transform (WHT). The WHT is a transform used in signal and image processing and coding theory [2, 6, 18]. The WHT is similar to the discrete Fourier transform (DFT) and can be computed with a divide and conquer algorithm similar to the well-known fast Fourier transform (FFT). There is a large family of fast, i.e. $\Theta(N \lg(N))$, algorithms for computing the WHT, which despite having the same arithmetic complexity, can exhibit widely varying performance. The package described in [15] can be used to explore the performance of these algorithms and can automatically select an algorithm, using generate-and-test, with good performance on a given platform. Empirically it was shown that there is a wide range of performance, and the best algorithms are rare. Various comments were given, having to do with code structure and the underlying architecture, to suggest explanations for the distribution of run times and the selection of the best algorithms on different architectures.

In this paper a performance model is used to explore the space of WHT algorithms. The performance model, which uses instruction count, does not account for important features of modern computer architectures such as cache, pipelining, and instruction level parallelism, and thus can not be used to accurately actual performance. Nonetheless, the model does provide important insight into the performance of WHT algorithms. More importantly the model can be analyzed analytically and there is a theoretical basis for the observed distributions of run times. Using the instruction count model, an analytic solution is provided to the optimization problem, the average instruction count is calculated, and the limiting distribution is determined. These results are obtained through the study of a class of divide and conquer recurrences.

The divide and conquer algorithms for computing the WHT arise from arbitrary compositions (i.e. ordered partitions) of the exponent n of the size $N = 2^n$ of the transform. Typically, divide and conquer approach refers to decomposing a problem of a given size n into a small (typically 2), deterministic number of subproblems of smaller sizes. Our partitioning strategies, allowing search over *all* compositions (i.e. ordered partitions) of n are a sharp departure from a standard scheme: not only is the number of subproblems random, it is also allowed to grow to infinity with n (and typically does). This leads, in principle, to more complex recurrences. While there is a vast literature (see e.g. [19], [13] and references therein) on traditional divide and conquer recurrences (often referred to as quicksort type recurrences), the more general scheme has been addressed only sporadically. Certain cases are considered in [8, Section 2] under the name stochastic divide and conquer recurrences while Roura [24, Section 3] uses the name continuous divide and conquer recurrences.

These are generally recurrences of the form

$$f_n = \sum_{k=1}^{n-1} \xi_{n,k} f_k + t_n,$$

where the coefficients $(\xi_{n,k})$ and a sequence (t_n) are known. Flajolet et al. considered the case $\xi_{n,k} = (n-k)/(n(n+1))$. Roura treated a more general case of (essentially polynomial) coefficients under some regularity assumptions. Our consideration of expected values in Section 5.3 (see Equation 15) lead to the same type of recurrence with $\xi_{n,k} = (n+3-k)2^{-k-1}$, up to a constant only dependent on n , which we solve by elementary means. In view of these facts, however, it seems that would be worthwhile to try to develop a fairly general theory of such recurrences, with as mild assumptions on $\xi_{n,k}$'s as possible.

As for the issue of limiting distribution, Régnier [22] proved by martingale methods that the normalized number of comparisons in quicksort algorithm converges in distribution and Rösler [23] used Banach's fixed point theorem to characterize the limiting distribution of quicksort as the *unique* solution of an underlying distributional equation. Rösler's approach, now referred to as contraction method proved very useful in analyzing other divide and conquer recurrences (see, e.g. examples and references in [13]) as well as in studying more subtle properties of the limiting distribution of quicksort (see , for example, [7]). Of crucial importance, however, is the fact the each recursive subdivision is into two (or a number bounded independently of the size n) subproblems, an assumption that fails in our case. Consequently, the methods developed for studying quicksort type recurrences are of little use in this context. Our approach is via martingale techniques (albeit quite different than Régnier's); the main tool is the central limit theorem for martingales.

The paper is organized as follows. Section 2 reviews the Walsh-Hadamard transform and defines the space of algorithms for computing it that will be considered. Section 3 presents the instruction count model, and Section 4 summarizes empirical data that illustrates the instruction count model and properties of the space of WHT algorithms. Section 5 analyzes a family of divide and conquer recurrences that arise when the performance model is applied to the family of WHT algorithms. The minimum, maximum, expected values, and variances are calculated and the limiting distribution is obtained. These results, when restricted to binary splits, contain various known results as special cases. Section 6 compares the special cases of the general result to these known results.

2. THE WALSH-HADAMARD TRANSFORM

The Walsh-Hadamard transform of a signal x , of size $N = 2^n$, is the matrix-vector product $\mathbf{WHT}_N \cdot x$, where

$$\mathbf{WHT}_N = \bigotimes_{i=1}^n \mathbf{DFT}_2 = \overbrace{\mathbf{DFT}_2 \otimes \cdots \otimes \mathbf{DFT}_2}^n.$$

The matrix

$$\mathbf{DFT}_2 = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$$

is the 2-point DFT matrix, and \otimes denotes the tensor or Kronecker product. The tensor product of two matrices is obtained by replacing each entry of the first matrix

by that element multiplied by the second matrix. Thus, for example,

$$\begin{aligned} \mathbf{WHT}_4 &= \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \otimes \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \\ &= \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{bmatrix}. \end{aligned}$$

Algorithms for computing the WHT can be derived using properties of the tensor product [25, 16]. More precisely, algorithms for computing the WHT can be represented by structured factorizations of the WHT matrix, and such factorizations can be derived using the multiplicative property of tensor products: $(A \otimes B)(C \otimes D) = AC \otimes BD$. For example, a recursive algorithm for the WHT is obtained from the factorization

$$(1) \quad \mathbf{WHT}_{2^n} = (\mathbf{WHT}_2 \otimes I_{2^{n-1}})(I_2 \otimes \mathbf{WHT}_{2^{n-1}}),$$

where I_m is the $m \times m$ identity matrix. This factorization is obtained by using the multiplicative property with $A = \mathbf{WHT}_2$, $B = I_{2^{n-1}}$, $C = I_2$, and $D = \mathbf{WHT}_{2^{n-1}}$. An algorithm to compute $y = \mathbf{WHT}_{2^n}x$, corresponding to this factorization, is obtained by first computing $t = (I_2 \otimes \mathbf{WHT}_{2^{n-1}})x$, which involves two recursive calls to compute $W_{2^{n-1}}$, and then computing $y = (\mathbf{WHT}_2 \otimes I_{2^{n-1}})t$ which is computed by applying W_2 to subvectors of t containing the pair of elements $(t_i, t_{i+2^{n-1}})$, for $i = 0, \dots, 2^{n-1} - 1$. Note that this factorization does not specify how the recursive calls are computed. This would be determined by a recursive factorization of $\mathbf{WHT}_{2^{n-1}}$.

An iterative algorithm for computing the WHT is obtained from the factorization

$$(2) \quad \mathbf{WHT}_{2^n} = \prod_{i=1}^n (I_{2^{i-1}} \otimes \mathbf{WHT}_2 \otimes I_{2^{n-i}}),$$

which can be proven (see [16]) using induction and the multiplicative property along with the identity $I_m \otimes I_n = I_{mn}$. More generally, let $n = n_1 + \dots + n_t$, then

$$(3) \quad \mathbf{WHT}_{2^n} = \prod_{i=1}^t (I_{2^{n_1+\dots+n_{i-1}}} \otimes \mathbf{WHT}_{2^{n_i}} \otimes I_{2^{n_{i+1}+\dots+n_t}})$$

This equation encompasses both the iterative and recursive algorithm and provides a mechanism for exploring different breakdown strategies and combinations of recursion and iteration.

Let $N = N_1 \cdots N_t$, where $N_i = 2^{n_i}$, and let $x_{b,s}^M$ denote the vector $(x(b), x(b+s), \dots, x(b+(M-1)s))$. Then evaluation of $x = \mathbf{WHT}_N \cdot x$ using Equation 3 is performed using

$$\begin{aligned} R &= N; \quad S = 1; \\ \text{for } i &= 1, \dots, t \\ R &= R/N_i; \\ \text{for } j &= 0, \dots, R-1 \\ \text{for } k &= 0, \dots, S-1 \\ x_{jN_i S+k, S}^{N_i} &= \mathbf{WHT}_{N_i} \cdot x_{jN_i S+k, S}^{N_i}; \\ S &= S * N_i; \end{aligned}$$

The computation of \mathbf{WHT}_{N_i} is computed recursively in a similar fashion until a base case of the recursion is encountered. Observe that \mathbf{WHT}_{N_i} is called N/N_i times. Small WHT transforms are computed using the same approach; however, the code is unrolled in order to avoid the overhead of loops or recursion. This scheme assumes that the algorithm works in-place and is able to accept stride parameters.

Alternative algorithms are obtained through different sequences of the application of Equation 3. Each algorithm obtained this way can be represented by a tree, called a partition tree. The root of the partition tree corresponding to an algorithm for computing \mathbf{WHT}_N , where $N = 2^n$ is labeled with n . Each application of Equation 3 corresponds to an expansion of a node into children whose sum equals the node. Leaf nodes in the tree correspond to the base case of the recursion. A node labeled with m corresponds to the computation of \mathbf{WHT}_{2^m} . If the node is in a tree rooted with n , the computation of \mathbf{WHT}_{2^m} is performed 2^{n-m} times. Figure 1 shows the trees for a recursive and iterative algorithm for computing \mathbf{WHT}_{16} .




FIGURE 1. Partition Trees for Iterative and Recursive WHT Algorithms

In this paper the performance of WHT algorithms corresponding to all possible partition trees and various subsets of partition trees is explored. The total number of partition trees of size n is given by the recurrence

$$(4) \quad T_n = 1 + \sum_{n_1 + \dots + n_k = n} T_{n_1} \dots T_{n_k},$$

and the subset of fully expanded partition trees (i.e. all leaves equal to 1) satisfies the recurrence

$$(5) \quad \tilde{T}_n = \sum_{n_1 + \dots + n_k = n} \tilde{T}_{n_1} \dots \tilde{T}_{n_k}.$$

The number of binary and fully expanded binary partition trees satisfy the recurrences

$$(6) \quad B_n = 1 + \sum_{n_1 + n_2 = n} B_{n_1} \cdot B_{n_2},$$

and

$$(7) \quad \tilde{B}_n = \sum_{n_1 + n_2 = n} \tilde{B}_{n_1} \cdot \tilde{B}_{n_2}.$$

Table 1 lists the first few values of T_n , \tilde{T}_n , B_n , and \tilde{B}_n .

The generating function, $T(z)$, for T_n satisfies the functional equation

$$(8) \quad T(z) = z/(1 - z) + T(z)^2/(1 - T(z)),$$

n	1	2	3	4	5	6	7	8
T_n	1	2	6	24	112	568	3032	16768
\tilde{T}_n	1	1	3	11	45	197	903	4279
B_n	1	1	2	5	15	51	188	731
\tilde{B}_n	1	1	2	5	14	42	132	429


TABLE 1. Number of Partition Trees for \mathbf{WHT}_{2^n} 

FIGURE 2. Performance histogram on the Pentium III

and consequently $T_n = \Theta(\alpha^n/n^{3/2})$, where $\alpha = 4 + 2\sqrt{2} \approx 6.828427120$. The other recurrences satisfy similar functional equations: $\tilde{T}(z) = z + \tilde{T}(z)^2/(1 - \tilde{T}(z))$, $B(z) = z/(1 - z) + B(z)^2$, and $\tilde{B}(z) = z + \tilde{B}(z)^2$, and consequently the number of fully expanded trees is $\Theta(\beta^n/n^{3/2})$, where $\beta = 3 + 2\sqrt{2} \approx 5.828427120$, the number of binary trees is $\Theta(5^n/n^{3/2})$, and the number of fully expanded binary trees is $\Theta(4^n/n^{3/2})$.

3. PERFORMANCE MODEL FOR THE WHT

In the previous section it was shown that there is a large family of WHT algorithms which have varying degrees of recursion, iteration, and straight-line code. A natural question is to determine which algorithm leads to the best performance. The histogram in Figure 2 shows that there is in fact a wide range in performance. The times in Figure 2 were obtained using the WHT package from [15], and were obtained on a Pentium III with 128MB of memory running at 550MHz. The histogram shows the runtimes for 10,000 randomly generated WHT algorithms of size 2^{16} .

The wide range of times in Figure 2 are not due to the number of arithmetic operations. The following theorem shows that all algorithms have exactly the same number of floating point operations (flops).

Theorem 1. *Let \mathcal{W}_N be a fully expanded WHT algorithm for computing \mathbf{WHT}_N . Then $\text{flops}(\mathcal{W}_N) = N \lg(N)$.*

The proof is by induction on $N = 2^n$. The base case for $n = 1$ is clearly true. In general, assume that \mathcal{W}_N uses the factorization

$$\prod_{i=1}^t (\mathbf{I}_{2^{n_1+\dots+n_{i-1}}} \otimes \mathcal{W}_{2^{n_i}} \otimes \mathbf{I}_{2^{n_{i+1}+\dots+n_t}}),$$

where $\mathcal{W}_{2^{n_i}}$ is an algorithm to compute $\mathbf{WHT}_{2^{n_i}}$. Since $\mathcal{W}_{2^{n_i}}$ is called 2^{n-n_i} times,

$$\text{flops}(\mathcal{W}_N) = \sum_{i=1}^t 2^{n-n_i} \text{flops}(\mathcal{W}_{2^{n_i}}),$$

which by induction is equal to

$$\sum_{i=1}^t 2^{n-n_i} 2^{n_i} n_i = N \sum_{i=1}^t n_i = Nn = N \lg(N).$$

Since arithmetic operations can not distinguish algorithms, the distribution of runtimes must be due to other factors. In [12], other performance metrics, such as instruction count, memory accesses, and cache misses, were gathered and their influence on runtime was investigated. Different algorithms can have vastly different instruction counts due to the varying amounts of control overhead from recursion, iteration, and straight-line code. Different algorithms access data in different patterns, with varying amounts of locality. Consequently different algorithms can have different amounts of instruction and data cache misses.

In this paper, the focus is on an instruction count model and the mathematical techniques required to analyze the number of instructions for different WHT algorithms. While instruction count does not accurately model performance, it is an important aspect of performance and the model provides insight into many aspects of the performance tradeoff of different WHT algorithms. Furthermore, the model is sufficiently complex to exhibit non-trivial behavior yet is amenable to analytic techniques.

The number of instructions required by an algorithm depends on the way the algorithm is implemented, the compiler that translates the program implementing the algorithm into machine instructions, and the machine on which the algorithm is executed. It is possible to derive a set of parameterized recurrence relations for the number of machine instructions used by an arbitrary WHT algorithm. The parameters depend on the program, compiler and machine used.

The WHT package executes code similar to the pseudo-code in the previous section each time the algorithm is called recursively. When a leaf node is encountered a procedure to compute the corresponding WHT, implemented using straight-line code, is called. A library of small WHT procedures is generated by a family of code generators. In the WHT package straight-line code is considered only for sizes 2^n for $n = 1, \dots, 8$, since it has been determined that straight-line code for larger sizes is not beneficial on current computers.

In order to count the number of instructions, it is necessary to know how many times the recursive WHT algorithm is called, how many times each loop body is executed, and how many times each small WHT procedure is called. Given this information and the number of instructions for the straight-line code and the basic blocks in the WHT procedure, it is possible to determine a formula for the number of instructions.

α	α_1	α_2	α_3	β_1	β_2	β_3
27	12	34	106	18	18	20

TABLE 2. Instruction constants for the Pentium III using gcc version 2.91.66

Let \mathcal{W}_{2^n} be a WHT algorithm, and let $A(n)$ be the number of times the recursive WHT procedure is called, $A_l(n)$ the number of times the straight-line code for **WHT**_{2^l} is called, $L_1(n)$ the number of times the outer loop is executed, $L_2(n)$ the number of times the middle loop is executed, $L_3(n)$ the number of times the innermost loop is executed. Then the number of instructions required to execute \mathcal{W}_{2^n} is equal to

$$\alpha A(n) + \sum_{l=1}^8 \alpha_l A_l(n) + \sum_{i=1}^3 \beta_i L_i(n),$$

where α is the number of instructions for the code in the compiled WHT procedure executed outside the loops, α_i , $i = 1, \dots, 8$ is the number of instructions in the compiled straight-line code implementations of small WHT's of size one through eight, and β_i , $i = 1, 2, 3$ is the number of instructions executed in the outer-most, middle, and inner-most loops in the compiled WHT procedure.

The α and β constants are determined by examining the assembly code produced by compiling the WHT package. Table 2 shows the values obtained for the Pentium III using the gcc compiler version 2.91.66 with flags set to “-O6 -fomit-frame-pointer -pedantic -malign-double -Wall”.

The functions $A(n)$, $A_l(n)$, $L_i(n)$ satisfy recurrences that can be derived from the pseudo code in Section 3.

The WHT procedure is called once plus the number of calls in $\mathcal{W}_{2^{n_i}}$ for $i = 1, \dots, t$, and since $\mathcal{W}_{2^{n_i}}$ is called 2^{n-n_i} times, the number of recursive calls to the WHT procedure in \mathcal{W}_{2^n} is equal to

$$(9) \quad A(n) = \begin{cases} 1 + \sum_{i=1}^t 2^{n-n_i} A(n_i) & \text{if } n = n_1 + \dots + n_t \\ 0 & \text{if } n \text{ is a leaf.} \end{cases}$$

Similarly the number of calls to the straight-line code for **WHT**_{2^l} in the algorithm \mathcal{W}_{2^n} is equal to

$$(10) \quad A_l(n) = \begin{cases} \sum_{i=1}^t 2^{n-n_i} A_l(n_i) & \text{if } n = n_1 + \dots + n_t \\ 1 & \text{if } n = l \text{ is a leaf.} \end{cases}$$

It is easy to show that $A_l(n) = \nu_l 2^{n-l}$, where ν_l is the number of leaf nodes with value l in the tree corresponding to \mathcal{W}_n .

Since the outermost loop is executed t times, the middle loop is executed $2^{n_1+\dots+n_{i-1}}$ times for the i -th execution of the outermost loop, and the innermost loop is executed 2^{n-n_i} times for each iteration of the middle loop in the i -th iteration of the

Size	Right Recursive/Iterative	Left Recursive/Iterative	Balanced/Iterative
2	1.00	1.00	1.00
3	1.31	1.37	1.37
4	1.41	1.54	1.63
5	1.42	1.61	1.75
6	1.40	1.64	1.82
7	1.37	1.64	1.91
8	1.34	1.64	1.97
9	1.31	1.64	1.99
10	1.28	1.64	2.00
11	1.26	1.63	2.01
12	1.24	1.63	2.02
13	1.22	1.62	2.04
14	1.21	1.62	2.07
15	1.20	1.61	2.09
16	1.18	1.61	2.10

TABLE 3. Ratio of Instruction Counts Recursive, Balanced, and Iterative Algorithms

outer loop,

$$(11) L_1(n) = \begin{cases} t + \sum_{i=1}^t 2^{n-n_i} L_1(n_i) & \text{if } n = n_1 + \dots + n_t \\ 0 & \text{if } n \text{ is a leaf} \end{cases}$$

$$(12) L_2(n) = \begin{cases} \sum_{i=1}^t \{2^{n-n_i} L_2(n_i) + 2^{n_1+\dots+n_{i-1}}\} & \text{if } n = n_1 + \dots + n_t \\ 0 & \text{if } n \text{ is a leaf} \end{cases}$$

$$(13) L_3(n) = \begin{cases} \sum_{i=1}^t \{2^{n-n_i} L_3(n_i) + 2^{n-n_i}\} & \text{if } n = n_1 + \dots + n_t \\ 0 & \text{if } n \text{ is a leaf} \end{cases}$$

4. EMPIRICAL OBSERVATIONS OF THE PERFORMANCE MODEL

This section presents empirical data using the instruction count model with constants set to the values in Table 2. Observations from the data provide some insight into the behavior of the family of WHT algorithms presented in Section 2. Furthermore, the data suggests theorems which will be proved in the following sections.

Table 3 compares the number of instructions used by four families of WHT algorithms. Counts are reported as a ratio of the number of instructions used by a given algorithm of specified size to the number of instructions used by the iterative algorithm of the same size. The iterative algorithm is obtained by setting $n = 1 + \dots + 1$ in Equation 3. The *left recursive*, *right recursive*, and *balanced* algorithms are obtained by recursively splitting n using $n = 1 + (n - 1)$, $n = (n - 1) + 1$, and $n = \lceil n/2 \rceil + \lfloor n/2 \rfloor$ in Equation 3 respectively.

In all cases the iterative algorithm has the fewest number of instructions. Note that the number of instructions used by the right and left recursive algorithms differs. This difference is due solely to the L_2 component corresponding to the middle loop, as all other cost functions are invariant under permutations of the children. Moreover, the L_2 recurrence is lowest if the larger children are to the right.

The data in Table 3 suggest that there are limiting ratios. This can be verified by specializing recurrences 9, 10, and 11–13 to the algorithms in the table. For all of the algorithms $A_1(n) = n2^{n-1}$. For the iterative algorithm $A(n) = 1$, $L_1(n) = n$, $L_2(n) = 2^n - 1$, and $L_3(n) = n2^{n-1}$. For the right recursive algorithm, $A(n) = 2^n - 1$, $L_1(n) = 2(2^n - 1)$, $L_2(n) = 3(2^n - 1)$, and $L_3(n) = n2^{n-1} + 2^{n+1} - 2$. The left recursive algorithm has the same solutions to the recurrence relations except $L_2(n) = n2^{n-1} + 2^n - 1$. A simple closed form solution to the recurrences for the balanced algorithm was not found; however, $A(n)$ and $L_i(n)$ for $i = 1, 2, 3$ are all $\Theta(n2^n)$ and the limiting constants can be determined numerically: $A(n) = 0.1411142349n2^n$, $L_1(n) = 0.2822284699n2^n$, $L_2(n) = 0.4616713524n2^n$, and $L_3(n) = 0.6411142349n2^n$. Plugging in the constants for the Pentium III the limiting ratios in Table 3 are 1, 25/16, and 2.251410365 respectively.

The data in Table 3 and the ensuing discussion suggests that the iterative algorithm is optimal and that some combination of a balanced tree favoring larger children to the right is the worst case. Moreover it appears, for fully expanded trees using the Pentium III constants, that there is a factor of about two between the best and worst performance in terms of instruction counts.

The trees with the minimum and maximum instruction counts can be found using dynamic programming, since using the instruction count model, the optimal tree of a given size is independent of the context in which it is called (i.e. where in the tree it is located). Dynamic programming is applied by generating all possible splits of a node of size n and computing the max or min value using the max/min values for each of the children. The values for all smaller sizes must be computed and once they are computed the values are obtained by table lookup. To implement this procedure it is necessary to generate all possible splits. This is done using a one-to-one mapping between compositions of n and $(n - 1)$ -bit numbers. The mapping is obtained by considering a string of n ones and placing the bits of the $(n - 1)$ -bit binary number between adjacent ones. The string of ones are summed until a 1-bit is encountered (assume an implicit 1-bit at the end). For example the partition generated from the number $j = 0100110$ is [2,3,1,2].

This procedure was implemented using Maple version 8 and the Pentium III constants. The optimal and worst case trees were obtained for size $n = 2, \dots, 16$. For all sizes, the iterative algorithm using leaf nodes of size 3 (2 was also used when n is not a multiple of 3). More precisely, when $n \equiv 0 \pmod{3}$, the iterative tree with $n/3$ children of size 3 is optimal, when $n \equiv 1 \pmod{3}$, the iterative tree with $\lfloor n/3 \rfloor - 1$ children of size 3 and two leftmost children of size 2 is optimal, and when $n \equiv 2 \pmod{3}$, the iterative tree $\lfloor n/3 \rfloor$ children of size 3 and the left child of size 2 is optimal. The function $A_l(n)$ which counts the number of instructions due to leaf nodes can be used to suggest which leaves will occur in the optimal tree. First note that all trees with the same set of leaf nodes have the same value of $A_l(n)$. When comparing two iterative trees the contribution due to all leaf nodes with the same value can be ignored. Thus, for example, in the case when $n \equiv 1 \pmod{3}$, it is required only to compare $2^{n-2}\alpha_2 + 2^{n-2}\alpha_2$ with $2^{n-1}\alpha_1 + 2^{n-3}\alpha_3$. In this case, the tree with two leaves of size 2 will be chosen over the tree with one leaf of size 1 and one leaf of size 3 when $4\alpha_2 \leq 4\alpha_1 + \alpha_3$, as is the case for the pentium III.

Table 4 compares the best WHT algorithm with the worst. It shows that a factor of 6 to 7 is available by choosing the appropriate algorithm. Figures 3 and 4 show the trees that lead to the maximum instruction counts for size 13 and 16. The tree of size 13 is an example of a balanced power of two tree [14]. A balanced power of




FIGURE 4. Worst case tree of size 16 on the Pentium III

n	2	3	4	5	6	7	8	9
ratio bin/gen	1.000	1.000	1.000	1.000	1.000	1.065	1.031	1.031
n	10	11	12	13	14	15	16	
ratio bin/gen	1.032	1.028	1.027	1.024	1.023	1.022	1.019	

TABLE 5. Ratio of the best binary and best general WHT algorithms using Pentium III instruction counts

10,000 random WHT trees. A random tree of size n is obtained by choosing a random composition of $n = n_1 + \dots + n_t$ (the trivial composition with n by itself is not allowed) and then recursively choosing random trees with roots equal to n_1, \dots, n_t . If $n_i \leq \text{CUTOFF}$, then a leaf of size n_i is returned. In this experiment CUTOFF was set to 3, since those are the values for which code is available and useful on the Pentium III. Note that this method of generating random trees is not uniform in the space of trees. Short trees with lots of small nodes are favored.

Figure 5 shows a histogram, generated by Maple, of the instruction counts for the random sample of 10,000 trees. The mean is approximately 1.201×10^7 and the variance is approximately 4.635×10^{12} . The minimum number of instructions, determined by dynamic programming, is 6066945 and the maximum number of instructions is 37817249. The minimum number of instructions obtained in the sample is 6088617, while the maximum is 20616167.

If the randomly generated trees are restricted to fully expanded trees, then the distribution more closely resembles a normal distribution (see Figure 6). The mean is approximately 2.862×10^7 and the variance is approximately 6.147×10^{12} . The minimum number of instructions obtained in the sample is 18827489, while the maximum is 36226849.

5. ANALYSIS OF WHT RECURRENCES

In this section we establish theoretical results concerning Recurrences 9 and 11–13 derived in Section 3 and we will see that they agree with empirical observations discussed in Section 4.




FIGURE 5. Instruction count histogram on the Pentium III




FIGURE 6. Instruction count histogram on the Pentium III

5.1. Some Reductions. For the purpose of mathematical analysis we concentrate on fully expanded partition trees (i.e. all leaves equal to 1). This assumption does not change the nature of general phenomenon; it simplifies analysis a little bit and, most importantly, makes it cleaner. Of course, as illustrated in Section 4 from practical point of view it is important to consider leaves of larger sizes as well.

Consider first Recurrence 9. After dividing by 2^n and letting $L_0(k) = A(k)/2^k$, for $k \geq 2$ it becomes

$$L_0(n) = \frac{1}{2^n} + \sum_{i=1}^t L_0(n_i).$$

For future considerations, it is convenient to look at the slightly more general setting. Namely, let $T_0(n)$, $n \geq 1$ be a sequence of random variables, which we will

refer to as a toll function. Then, the above relation may be viewed as

$$A_0(n) = \sum_{i=1}^t A_0(n_i) + T_0(n),$$

where $T_0(n) \equiv 1/2^n$. As for the other recurrences, for $L_3(n)$ add 1 to both sides and then divide through by 2^n to see that the quantity $(L_3(n) + 1)/2^n$ satisfies exactly the same recurrence as $L_0(n)$ and thus will not be of interest anymore. For the remaining two, divide through by 2^n . The corresponding toll $T_1(n)$ is

$$T_1(n) = \frac{t}{2^n}.$$

Finally, for $L_2(n)$ we obtain

$$\frac{L_2(n)}{2^n} = \sum_{i=1}^t \left\{ \frac{L_2(n_i)}{2^{n_i}} + \frac{1}{2^{n_i+n_{i+1}+\dots+n_t}} \right\},$$

i.e., the toll function in this case is

$$T_2(n) = \sum_{i=1}^t \frac{1}{2^{n_i+n_{i+1}+\dots+n_t}}.$$

Letting $F_k(n) = L_k(n)/2^n$ for $k = 0, 1, 2$, we see that all F 's satisfy the same recurrence with different toll functions. Thus, we will consider a generic recurrence of the form

$$(14) \quad F(n) = \sum_{i=1}^t F(n_i) + T(n),$$

only occasionally referring to the specific toll function $T_k(n)$ or the recurrence $F_k(n)$.

We wish to analyze the limiting distribution of the random variable $F(n)$ under the assumption that a composition is chosen uniformly at random from the set Ω_n of all $2^{n-1} - 1$ compositions of n into at least two parts, and that $F(n_i)$'s depend on a particular composition only through the sizes of parts n_i . That is, $F(n_i)$'s are conditionally independent once the composition is chosen.

We will assume that the initial value $F(1)$ is a given nonrandom number. Since the calls to leaf nodes are treated separately by the recurrence (10), the initial condition for each of the recurrences is $F(1) = 0$, but mathematically it makes no difference if the initial value is set to be any other number.

Let us begin by establishing deterministic bounds on these recurrences.

5.2. Deterministic bounds. In order to state our bounds we need a bit more notation. For a binary partition tree T with internal nodes $I(T)$ we let

$$w(T) = \sum_{x \in I(T)} \frac{1}{2^x}.$$

Let us consider the so-called balanced power of two tree (see [14]) that is a partition tree that can be defined as follows: given a positive integer n , consider its binary expansion

$$n = 2^{k_1} + 2^{k_2} + \dots + 2^{k_j}, \quad k_1 > k_2 > \dots > k_j.$$

The root n is split as

$$n = \begin{cases} 2^{k_1} + (n - 2^{k_1}), & \text{if } k_2 = k_1 - 1; \\ 2^{k_1-1} + (n - 2^{k_1-1}), & \text{if } k_2 < k_1 - 1. \end{cases}$$

The same rule is then applied recursively. We let $w_n = w(T_b(n))$, where $T_b(n)$ is the balanced power of two partition tree for n . We have

Proposition 1. *The following are tight bounds on the given recurrences:*

- (i) $nF_0(1) + \frac{1}{2^n} \leq F_0(n) \leq nF_0(1) + w_n$
- (ii) $n(F_1(1) + \frac{1}{2^n}) \leq F_1(n) \leq nF_1(1) + 2w_n$
- (iii) $nF_2(1) + 1 - \frac{1}{2^n} \leq F_2(n) \leq n(F_2(1) + 1) - \frac{1}{2^n}$

Before proving this proposition let us remark that lower bounds are exactly the values of the respective recurrences when an iterative partition tree is used and that the upper bound for F_2 is the value attained at the left-recursive partition tree. The upper bounds for F_0 , and F_1 are the values obtained on $T_b(n)$. This value can be obtained numerically; for example, when $n = 2^m$ is a perfect power of 2 then,

$$w_n = \sum_{j=0}^{m-1} 2^j 2^{-2^{m-j}} = 2^m \sum_{j=1}^m 2^{-j} 2^{-2^j} \leq n \sum_{j=1}^{\infty} 2^{-j} 2^{-2^j} = nw,$$

where $w \sim .1411142349\dots$

Proof: It is an easy inductive proof to show that the given expressions are lower bounds. For example,

$$\begin{aligned} F_0(n) &= \sum_{i=1}^t F_0(n_i) + \frac{1}{2^n} = \sum_{i: n_i \geq 2} F_0(n_i) + \sum_{i: n_i=1} F_0(n_i) + \frac{1}{2^n} \\ &\geq \sum_{i: n_i \geq 2} \left(n_i F_0(1) + \frac{1}{2^{n_i}} \right) + \sum_{i: n_i=1} F_0(1) + \frac{1}{2^n} \\ &= nF_0(1) + \frac{1}{2^n} + \sum_{i: n_i \geq 2} \frac{1}{2^{n_i}} \geq nF_0(1) + \frac{1}{2^n}. \end{aligned}$$

For (ii) we obtain

$$F_1(n) \geq nF_1(1) + \frac{t}{2^n} + \sum_{i: n_i \geq 2} \frac{n_i}{2^{n_i}} \geq n + \frac{\#\{i: n_i = 1\}}{2^n} + \sum_{i: n_i \geq 2} \frac{n_i}{2^{n_i}} = n + \frac{n}{2^n}.$$

For (iii) observe that once a composition $n = n_1 + \dots + n_t$ is chosen, the n_i 's should be arranged in a non-decreasing order. That is because, the sum of $F_2(n_i)$ remains the same but the sum of $1/2^{n_i + \dots + n_t}$ is minimized if the n_i 's are non-decreasing. If now this is the case, and if there is an n_i larger than 1, then, in particular, $n_t \geq 2$, and, by inductive hypothesis, $F_2(n_t) \geq n_t F_2(1) + 1 - 1/2^{n_t}$, and also $F_2(m) \geq mF_2(1)$ for $1 \leq m \leq n-1$. Thus,

$$\begin{aligned} F_2(n) &= \sum_{i=1}^{t-1} \left(F_2(n_i) + \frac{1}{2^{n_i + \dots + n_t}} \right) + \left(\frac{1}{2^{n_t}} + F_2(n_t) \right) \\ &\geq \sum_{i=1}^{t-1} n_i F_2(1) + \frac{1}{2^{n_t}} + n_t F_2(1) + 1 - \frac{1}{2^{n_t}} \\ &= nF_2(1) + 1 \geq nF_2(1) + 1 - \frac{1}{2^n}. \end{aligned}$$

We now turn to the upper bounds. First, notice that for each recurrence, binary splits give the worst behavior. The reason is that if there are more than two parts, then merging two of them together would increase the value. For the first two any two can be merged and for the third, merging the last two parts would increase the value:

$$\begin{aligned}
F_2(n) &= \sum_{i=1}^{t-2} \left(F_2(n_i) + \frac{1}{2^{n_i + \dots + n_t}} \right) + F_2(n_{t-1}) + \frac{1}{2^{n_{t-1}}} + F_2(n_t) + \frac{1}{2^{n_t}} \\
&\leq \sum_{i=1}^{t-2} \left(F_2(n_i) + \frac{1}{2^{n_i + \dots + n_t}} \right) + F_2(n_{t-1} + n_t) + \frac{1}{2^{n_{t-1}}} \\
&\leq \sum_{i=1}^{t-2} \left(F_2(n_i) + \frac{1}{2^{n_i + \dots + n_t}} \right) + F_2(n_{t-1} + n_t) + \frac{1}{2^{n_{t-1} + n_t}} \\
&= \sum_{i=1}^{t-1} \left(F_2(n_i^*) + \frac{1}{2^{n_i^* + \dots + n_t^*}} \right),
\end{aligned}$$

where

$$n_j^* = \begin{cases} n_j, & \text{if } j = 1, \dots, t-2; \\ n_{t-1} + n_t, & \text{if } j = t-1. \end{cases}$$

Now, the claimed upper bound in (iii) can be established again by a straightforward induction and is omitted. The remaining two recurrences are the same, up to a factor of 2 in a toll function, so let us consider the first one. Suppose that the root n is split as $n = n_1 + n_2$ and, inductively, that subsequent splittings of n_1 and n_2 follow the balanced power of 2 pattern. Assume without loss that $n_1 \geq n_2$ and let

$$n_1 = 2^{k_1} + 2^{k_2} + \dots + 2^{k_i}, \quad \text{and} \quad n_2 = 2^{\ell_1} + 2^{\ell_2} + \dots + 2^{\ell_j},$$

be the respective binary representations (note that $n_1 \geq n_2$ implies $k_1 \geq \ell_1$). Thus, n_1 and n_2 are split as

$$n_1 = 2^k + m_1, \quad n_2 = 2^\ell + m_2,$$

where k is either $k_1 - 1$ or k_1 and ℓ is either ℓ_1 or $\ell_1 - 1$. We have

$$2^{k-1} \leq m_1 < 2^{k+1}, \quad 2^{\ell-1} \leq m_2 < 2^{\ell+1}.$$

The rest of the argument consists on considering various cases and reshuffling the nodes correspondingly. For example, if both m_1 and m_2 are less than 2^k we replace the node $n_1 = 2^k + m_1$ by 2^k and the node n_2 by $n_2 + m_1$ (subsequently split into n_2 and m_1). If all other splittings are kept the same, this operation will increase the value since the term $1/2^{2^k + m_1}$ will be replaced by a larger value $1/2^{2^k}$, and that could be further increased by applying the balanced power of 2 rule to the new node $n_2 + m_1$. Other cases are handled similarly, and we omit the rest of the details.

5.3. Expected value. In order to analyze a normalized distribution of $F(n)$ we first need to find, asymptotically at least, its expected value and the variance. Given a composition of n let $m_k^{(n)}$ be the multiplicity of a part size k and let $\mu_k^{(n)}$ be the expected multiplicity of that size.

Let $f(n)$ and $v(n)$ denote the expected value and the variance of $F(n)$, respectively, and set $t(n) = \mathbf{E}T(n)$. Grouping together the terms in the recurrence relation

(14) by their sizes, using linearity of expectation and the assumption that, given a composition, the distributions of $F(n_i)$ depend only on the size n_i , we have

$$\begin{aligned}
f(n) &= \mathbb{E}F(n) = \mathbb{E} \left(T(n) + \sum_{i=1}^t F(n_i) \right) \\
&= t(n) + \mathbb{E} \left(\sum_{j=1}^{n-1} \sum_{i=1}^t I_{n_i=j} F(n_i) \right) \\
&= t(n) + \sum_{j=1}^{n-1} \mathbb{E} \left(\sum_{i=1}^t I_{n_i=j} F(n_i) \right) = t(n) + \sum_{j=1}^{n-1} \mu_j^{(n)} \mathbb{E}F(j) \\
&= t(n) + \sum_{j=1}^{n-1} \mu_j^{(n)} f(j).
\end{aligned}$$

Thus we obtain the following recurrence for $f(n)$:

$$(15) \quad f(n) = \sum_{j=1}^{n-1} \mu_j^{(n)} f(j) + t(n),$$

with $f(1)$ given (for WHT computations we set $f(1) = 0$). In order to proceed, we need some information about the quantities involved. Most of these have been studied for random compositions. Since in our model we disallow one trivial composition of n we need to adjust these results.

Lemma 2. *With the above notation, the following are true:*

$$\begin{aligned}
\text{(i)} \quad \mu_k^{(n)} &= \frac{2^{n-1}}{2^{n-1} - 1} \frac{n + 3 - k}{2^{k+1}}, \\
\text{(ii)} \quad t_1(n) &= \frac{(n+1)2^{n-2} - 1}{2^n(2^{n-1} - 1)}, \\
\text{(iii)} \quad t_2(n) &= \frac{1}{2} + \frac{1}{2^n}.
\end{aligned}$$

Proof: Let κ_0 denote the trivial composition of n with one part. We denote the probability and the expectation over the set of all compositions by \mathbb{P}_0 and \mathbb{E}_0 , respectively. The relationship between \mathbb{E} and \mathbb{E}_0 is, of course, $\mathbb{E}(\cdot) = \mathbb{E}_0(\cdot | \kappa \neq \kappa_0)$. We begin with (iii); we need to find the expected value of

$$T_2(n) = \sum_{i=1}^t \frac{1}{2^{n_i+n_{i+1}+\dots+n_t}}.$$

To this end, first consider the expectation over all 2^{n-1} compositions of n . Denoting it by g_n and conditioning on the value of n_t , which is k with probability $1/2^k$, for

$1 \leq k < n$, and n with probability $1/2^{n-1}$, we can write

$$\begin{aligned}
g_n &= \mathbf{E}_0 \frac{1}{2^{n_t}} \left(1 + \cdots + \frac{1}{2^{n-n_t}} \right) \\
&= \sum_{k=1}^n \mathbf{P}_0(n_t = k) \mathbf{E}_0 \left(\frac{1}{2^{n_t}} \left(1 + \cdots + \frac{1}{2^{n-n_t}} \right) \middle| n_t = k \right) \\
&= \frac{1}{2^{n-1}} \frac{1}{2^n} + \frac{1}{2^k} \sum_{k=1}^{n-1} \mathbf{E}_0 \left(\frac{1}{2^k} \left(1 + \cdots + \frac{1}{2^{n-k}} \right) \middle| n_t = k \right) \\
&= \frac{1}{2^{n-1}} \frac{1}{2^n} + \frac{1}{2^k} \frac{1}{2^k} \sum_{k=1}^{n-1} (1 + g_{n-k}) = \frac{1}{2^{2n-1}} + \sum_{k=1}^{n-1} \frac{1}{4^k} + \sum_{k=1}^{n-1} \frac{g_{n-k}}{4^k} \\
&= \frac{1}{3} + \frac{2}{3 \cdot 4^n} + \sum_{k=1}^{n-1} \frac{g_k}{4^{n-k}},
\end{aligned}$$

from which follows that $g_n = 1/2$ is a solution. Discarding one composition amounts to computing the conditional expectation

$$\begin{aligned}
\mathbf{E}T_2(n) &= \mathbf{E}_0(T_2(n) | \kappa \neq \kappa_0) = \mathbf{E}_0(T_2(n) | n_t < n) = \frac{1}{\mathbf{P}_0(n_t < n)} \mathbf{E}_0 T_2(n) I_{n_t < n} \\
&= \frac{2^{n-1}}{2^{n-1} - 1} \mathbf{E}_0 T_2(n) I_{n_t < n}.
\end{aligned}$$

But

$$\mathbf{E}_0 T_2(n) I_{n_t < n} = \mathbf{E}_0 T_2(n) - \mathbf{E}_0 T_2(n) I_{n_t = n} = \frac{1}{2} - \frac{1}{2^{n-1}} \frac{1}{2^n},$$

which proves (iii). As for the proofs of the first two assertions we use the facts that if all 2^{n-1} compositions of n are considered then the expected multiplicity of a part size k is $(n+3-k)/2^{k+1}$ (see [17]) and that the number of parts t is equidistributed with $1 + \text{Bin}(n-1, 1/2)$ random variable [11]. Thus, its expected value is $(n+1)/2$. Since the one composition we disallow has one part of size n and multiplicity one and all other sizes have multiplicity zero, after adjusting in the same manner as above we obtain

$$\mu_k^{(n)} = \frac{2^{n-1}}{2^{n-1} - 1} \frac{n+3-k}{2^{k+1}},$$

which proves (i), and furthermore

$$\begin{aligned}
\frac{n+1}{2} &= \mathbf{E}_0 t = \mathbf{E}_0 t I(\kappa = \kappa_0) + \mathbf{E}_0 t I(\kappa \neq \kappa_0) = \frac{1}{2^{n-1}} + \mathbf{P}_0(\kappa \neq \kappa_0) \mathbf{E}_0(t | \kappa \neq \kappa_0) \\
&= \frac{1}{2^{n-1}} + (1 - \frac{1}{2^{n-1}}) \mathbf{E}t,
\end{aligned}$$

from which (ii) follows.

Having computed the coefficients and tolls, we now turn to solving (15). This can be accomplished by elementary means. Rewriting

$$f(n) = \frac{2^{n-1}}{2^{n-1} - 1} \sum_{k=1}^{n-1} f(k) \frac{n+3-k}{2^{k+1}} + t(n),$$

as

$$\frac{2^{n-1} - 1}{2^{n-1}} f(n) = \sum_{k=1}^{n-1} f(k) \frac{n+3-k}{2^{k+1}} + t(n) \frac{2^{n-1} - 1}{2^{n-1}},$$

writing a similar expression replacing n by $n + 1$ and subtracting the former from the latter we obtain:

$$\begin{aligned} & \frac{2^n - 1}{2^n} f(n + 1) - \frac{2^{n-1} - 1}{2^{n-1}} f(n) \\ &= \sum_{k=1}^n f(k) \frac{n+4-k}{2^{k+1}} - \sum_{k=1}^{n-1} f(k) \frac{n+3-k}{2^{k+1}} + t(n+1) \frac{2^n - 1}{2^n} - t(n) \frac{2^{n-1} - 1}{2^{n-1}} \\ &= \frac{f(n)}{2^{n-1}} + \sum_{k=1}^{n-1} \frac{f(k)}{2^{k+1}} + t(n+1) \frac{2^n - 1}{2^n} - t(n) \frac{2^{n-1} - 1}{2^{n-1}}, \end{aligned}$$

which yields

$$\frac{2^n - 1}{2^n} f(n + 1) = f(n) + \sum_{k=1}^{n-1} \frac{f(k)}{2^{k+1}} + t(n+1) \frac{2^n - 1}{2^n} - t(n) \frac{2^{n-1} - 1}{2^{n-1}}.$$

Once again, writing a similar expression replacing $n + 1$ by n and subtracting we get

$$\begin{aligned} & \frac{2^n - 1}{2^n} f(n + 1) - \frac{2^{n-1} - 1}{2^{n-1}} f(n) \\ &= f(n) - f(n-1) + \frac{f(n-1)}{2^n} + t(n+1) \frac{2^n - 1}{2^n} - 2t(n) \frac{2^{n-1} - 1}{2^{n-1}} + t(n-1) \frac{2^{n-2} - 1}{2^{n-2}}, \end{aligned}$$

which, after solving for $f(n + 1)$, gives

$$f(n + 1) = 2f(n) - f(n-1) + t(n+1) - 4t(n) \frac{2^{n-1} - 1}{2^n - 1} + 4t(n-1) \frac{2^{n-2} - 1}{2^n - 1}.$$

Finally, letting for $n \geq 2$, $\Delta_{n+1} = f_{n+1} - f_n$ and $\gamma_{n+1} = t_{n+1} - 4t_n \frac{2^{n-1} - 1}{2^n - 1} + 4t_{n-1} \frac{2^{n-2} - 1}{2^n - 1}$, we obtain a very simple recurrence:

$$\Delta_{n+1} = \Delta_n + \gamma_{n+1}, \quad n \geq 2;$$

with the initial condition $\Delta_2 = f(1) + t(2)$. Iterating we obtain

$$\Delta_{n+1} = \Delta_2 + \sum_{k=2}^n \gamma_{k+1}.$$

Hence,

$$\sum_{j=2}^n \Delta_{j+1} = (n-1)\Delta_2 + \sum_{j=2}^n \sum_{k=2}^j \gamma_{k+1},$$

and since the left hand side is telescoping, letting $\gamma_2 = t(2)$ and changing the order of summation yields

$$f(n + 1) = (n + 1)f(1) + \sum_{j=1}^n (n + 1 - j)\gamma_{j+1}.$$

Recapitulating, we have obtained:

Theorem 3. *The solution of a recurrence (15) with the initial value $f(1)$ and toll function $(t(n))$ is given by*

$$f(n) = nf(1) + \sum_{j=1}^{n-1} (n - j)\gamma_{j+1},$$

where $\gamma_2 = t(2)$, and for $j \geq 2$,

$$\gamma_{j+1} = t(j+1) - \frac{4}{2^n - 1} ((2^{j-1} - 1)t(j) - (2^{j-2} - 1)t(j-1)).$$

In particular, for $i = 0, 1, 2$ there exist constants ϕ_i such that we have

$$f_i(n) \sim (f_i(1) + \phi_i)n.$$

Numerically, $\phi_0 = .073\dots$, $\phi_1 = .152\dots$, $\phi_2 = .271\dots$.

5.4. Variances. Recurrences for variances may be derived in the same fashion as those for the expected values. Let $v(n) = \text{var}(F(n))$. The following elementary property of variance (easiest to find in texts on statistics, e.g. [5]) will be handy: if X is any random variable and \mathcal{A} a σ -algebra, then

$$(16) \quad \text{var}(X) = \text{Evar}_{\mathcal{A}}(X) + \text{var}(\mathbf{E}_{\mathcal{A}}(X)),$$

where $\text{var}_{\mathcal{A}}(X) = \mathbf{E}(X - \mathbf{E}(X|\mathcal{A})^2|\mathcal{A})$ and $\mathbf{E}_{\mathcal{A}}(X) = \mathbf{E}(X|\mathcal{A})$ denote the conditional variance, and the conditional expectation given \mathcal{A} , respectively.

We will use the above with \mathcal{A} being a σ -algebra generated by compositions. That is to say that conditioning on \mathcal{A} means fixing a particular composition of n and if $\kappa = (n_1, \dots, n_t)$ was fixed, the conditional distribution of, say $F(n)$, is that of

$$\sum_{i=1}^t F(n_i) + T(n),$$

where we think of n_i 's, t , and $T(n)$ as deterministic (the last statement is certainly true in all three cases of our interest, and it is a reasonable assumption in general). Thus, the randomness is only in $F(n_1), \dots, F(n_t)$ and once n_i 's are fixed these are independent random variables. This, plus the fact that the variance is invariant under translations by a constant, and that, conditionally on \mathcal{A} , t is nonrandom, yields

$$(17) \quad \text{var}_{\mathcal{A}}(F(n)) = \text{var}_{\mathcal{A}}\left(\sum_{i=1}^t F(n_i) + T(n)\right) = \text{var}_{\mathcal{A}}\left(\sum_{i=1}^t F(n_i)\right)$$

$$(18) \quad = \sum_{i=1}^t \text{var}_{\mathcal{A}}(F(n_i)) = \sum_{j=1}^{n-1} m_j^{(n)} v(j),$$

where $m_j^{(n)}$ is the multiplicity of part size j . By the same reasoning, for the conditional expectation we get

$$(19) \quad \mathbf{E}_{\mathcal{A}}(F(n)) = \mathbf{E}_{\mathcal{A}}\left(\sum_{i=1}^t F(n_i) + T(n)\right) = \sum_{i=1}^t \mathbf{E}F(n_i) + \mathbf{E}_{\mathcal{A}}T(n) = \sum_{j=1}^{n-1} m_j^{(n)} f(j) + \mathbf{E}_{\mathcal{A}}T(n).$$

It follows from (16), (18), and (19) that

$$(20) \quad v(n) = \text{Evar}_{\mathcal{A}}(F(n)) + \text{var}(\mathbf{E}_{\mathcal{A}}(F(n))) = \sum_{j=1}^{n-1} \mu_j^{(n)} v(j) + \text{var}\left(\sum_{j=1}^{n-1} m_j^{(n)} f(j) + \mathbf{E}_{\mathcal{A}}T(n)\right),$$

which is exactly the same recurrence as (15) with toll function equal to $\text{var}(\sum_{j=1}^{n-1} m_j^{(n)} f(j) + \mathbf{E}_{\mathcal{A}}T(n))$. In each of the three cases this term can be quite precisely computed, but

does not appear to have a workable closed formula. For example, for $T_1(n)$, observing that the number of parts in a composition is the sum of all multiplicities, the toll $\text{var}(\sum_{j=1}^{n-1} m_j^{(n)} f(j) + t/2^n)$ can be written as

$$\text{var}\left(\sum_{j=1}^{n-1} m_j^{(n)} (f_1(j) + 1/2^n)\right).$$

However, solving the recurrences is more problematic since now we know the tolls only asymptotically and the recurrences are quite sensitive, since small values of n contribute significantly. Nonetheless, it is readily seen, that the tolls are linear functions of n (that is because the $m_j^{(n)}$ is asymptotically distributed like $\text{Bin}(\lceil n/2 \rceil, 1/2^j)$ and have asymptotically enough independence to show that covariances are negligible). Linearity of tolls implies that the solutions of the recurrences are also linear. Since in the next section we will provide another argument showing asymptotic linearity of the variances we will omit further details and will just state the result.

Proposition 4. *There exist absolute constants ω_0 , ω_1 , and $\omega_2 =$ such that for $i = 0, 1, 2$ we have*

$$v_i(n) \sim (v_i(1) + \omega_i)n.$$

Linearity of the variances is enough to establish convergence in distribution of $F(n)$ to normal random variable.

5.5. Limiting distribution. We will show in this section that the random variables $F(n)$, normalized to have mean zero and the variance 1, converge in distribution to a standard normal random variable. That is,

Theorem 5. *For $k = 0, 1, 2$ we have*

$$\frac{F_k(n) - f_k(n)}{\sqrt{v_k(n)}} \implies N(0, 1),$$

where $N(0, 1)$ denotes the normal random variable with mean zero and variance 1. That is, for all $a < b$ we have

$$\Pr(a \leq \frac{F_k(n) - f_k(n)}{\sqrt{v_k(n)}} \leq b) \longrightarrow \frac{1}{\sqrt{2\pi}} \int_a^b e^{-\frac{x^2}{2}} dx,$$

as $n \rightarrow \infty$.

Proof: This is a consequence of basic properties of random compositions and a central limit theorem for martingales (we refer the reader to ([3]; for all necessary background from probability theory that will be used throughout the remainder of this section). We recall [10, 11] that if all 2^{n-1} compositions are considered, then a randomly chosen composition is equidistributed with

$$(\Gamma_1, \Gamma_2, \dots, \Gamma_{\tau-1}, n - \sum_{j=1}^{\tau-1} \Gamma_j),$$

where $\Gamma_1, \Gamma_2 \dots$ are i.i.d. geometric random variables with parameter $1/2$, $\text{GEOM}(1/2)$. That is

$$\Pr(\Gamma_1 = j) = \frac{1}{2^j}, \quad j = 1, 2, \dots,$$

and τ is a stopping time defined by

$$\tau = \inf\{k \geq 1 : \Gamma_1 + \Gamma_2 + \cdots + \Gamma_k \geq n\}.$$

Restricting attention to compositions with at least two parts amounts to considering $\Gamma_k I(\Gamma_k < n)$'s rather than Γ_k 's and since the probability that these two are different is exponentially small and inconsequential from the point of view of the limiting law, from now on we will consider *all* compositions of n . In that case, τ is distributed like a $1 + \text{Bin}(n-1, 1/2)$ random variable, and thus is tightly concentrated around its expected value which is $(n+1)/2$. In particular, as is well known

$$(21) \quad \frac{\tau}{\mathbf{E}\tau} = \frac{2\tau}{n+1} \longrightarrow 1,$$

in probability as $n \rightarrow \infty$. Again, let us consider a generic quantity

$$\frac{F(n) - f(n)}{\sqrt{v(n)}} = \frac{\sum_{i=1}^t F(n_i) - \mathbf{E} \sum_{i=1}^t F(n_i)}{\sqrt{v(n)}} + \frac{T(n) - t(n)}{\sqrt{v(n)}}$$

Since $v(n)$ is of order n and all three tolls are bounded, $(T(n) - t(n))/\sqrt{v(n)}$ goes to zero and since the limiting distribution is continuous, this term can be neglected. The quantity whose limiting distribution we want to study in the present set-up, is

$$\tilde{S}_n = \sum_{k=1}^{\tau-1} F(\Gamma_k) + F(n - \sum_{j=1}^{\tau-1} \Gamma_j),$$

where for an integer valued, positive random variable W , $F(W)$ is a random variable, whose conditional distribution given $W = k$ is the distribution of $F(k)$, and the random variables $F(\Gamma_j)$, $j \geq 1$ are independent. Since, as follows from computations carried out in [11], $n - \sum_{j=1}^{\tau-1} \Gamma_j$ does not differ much from Γ_τ , (the difference is bounded in expectation), and $F(\Gamma)$ grows linearly with Γ , we may replace \tilde{S}_n by $S_{\tau \wedge n} = \sum_{j=1}^{\tau \wedge n} F(\Gamma_j)$; more specifically, we have

$$\frac{\tilde{S}_n - S_{\tau \wedge n}}{\sqrt{v(n)}} \longrightarrow 0,$$

in probability as $n \rightarrow \infty$. Thus, it suffices to consider a sequence $\{S_{\tau \wedge n} : n \geq 1\}$ and we want to show that

$$\frac{S_{\tau \wedge n} - \mathbf{E}S_{\tau \wedge n}}{\sqrt{v(n)}} = \frac{\sum_{k=1}^n I(\tau \geq k)F(\Gamma_k) - \sum_{k=1}^n \mathbf{E}I(\tau \geq k)F(\Gamma_k)}{\sqrt{v(n)}},$$

converges in distribution to a standard normal random variable. The plan is to apply the martingale central limit theorem, [3, Theorem 35.12]. Let \mathcal{F}_n , $n \geq 0$ be an increasing sequence of σ -fields with $\mathcal{F}_0 = \{\emptyset, \Omega\}$. There is a canonical way of turning a random variable $X - \mathbf{E}X$ into a martingale, by taking successive conditional expectations, i.e.

$$X - \mathbf{E}X = \sum_{k \geq 1} (\mathbf{E}(X|\mathcal{F}_k) - \mathbf{E}(X|\mathcal{F}_{k-1})).$$

Setting

$$\mathcal{F}_k = \sigma(\Gamma_1, \dots, \Gamma_k, F(\Gamma_1), \dots, F(\Gamma_k)),$$

and then for $n \geq 1$, $\mathcal{F}_{n,k} = \mathcal{F}_k$ and turning each for the variables

$$S_{\tau \wedge n} - \mathbf{E}S_{\tau \wedge n},$$

into a martingale $(X_{n,k})$, we obtain a triangular array of martingale difference sequences, just as required for an application of [3, Theorem 35.13]. (Strictly speaking we should have used $(S_{\tau \wedge n} - \mathbf{E}S_{\tau \wedge n})/\sqrt{v(n)}$, so that σ in that theorem is 1, but this is just a matter of normalization, and for the sake of notational convenience we will denote by $X_{n,k}$ and $Y_{n,k}$ the quantities before normalization. In this notation, the conditions (35.35) and (35.36) of that theorem become, respectively,

$$(22) \quad \frac{\sum_{k=1}^n \mathbf{E}(Y_{n,k}^2 | \mathcal{F}_{k-1})}{v(n)} \longrightarrow 1,$$

in probability, and

$$(23) \quad \frac{\sum_{k=1}^n \mathbf{E}Y_{n,k}^2 I(|Y_{n,k}| \geq \epsilon v(n))}{v(n)} \longrightarrow 0,$$

for each $\epsilon > 0$. Writing $\mathbf{E}_m(\cdot)$ for $\mathbf{E}(\cdot | \mathcal{F}_m)$ we have

$$\begin{aligned} X_{n,k} &= \mathbf{E}_k \left(\sum_{j=1}^n I(\tau \geq j) F(\Gamma_j) \right) \\ &= \sum_{j=1}^k I(\tau \geq j) F(\Gamma_j) + \mathbf{E}_k \left(\sum_{j=k+1}^n I(\tau \geq j) F(\Gamma_j) \right) \\ &= \sum_{j=1}^k I(\tau \geq j) F(\Gamma_j) + \mathbf{E}_k \left(\sum_{j=k+1}^n I(\tau \geq j) \mathbf{E}_{j-1} F(\Gamma_j) \right) \\ &= \sum_{j=1}^k I(\tau \geq j) F(\Gamma_j) + \mathbf{E}_k \left(\sum_{j=k+1}^n I(\tau \geq j) \mathbf{E} f(\Gamma_j) \right) \\ &= \sum_{j=1}^k I(\tau \geq j) F(\Gamma_j) + \mathbf{E} f(\Gamma) \cdot \mathbf{E}_k \left(\sum_{j=k+1}^n I(\tau \geq j) \right), \end{aligned}$$

where we have used the fact that both Γ_j and $F(\Gamma_j)$ are independent of \mathcal{F}_{j-1} and thus the conditional expectation $\mathbf{E}_{j-1} F(\Gamma_j)$ is equal to

$$\mathbf{E} F(\Gamma_j) = \mathbf{E} \mathbf{E}(F(\Gamma_j) | \Gamma_j) = \mathbf{E} f(\Gamma_j) = \mathbf{E} f(\Gamma),$$

where Γ is a random variable equidistributed with Γ_j . Hence, the differences $Y_{n,k}$ are

$$Y_{n,k} = X_{n,k} - X_{n,k-1} = I(\tau \geq k) F(\Gamma_k) + \mathbf{E} f(\Gamma) \left(\mathbf{E}_k \sum_{j=k+1}^n I(\tau \geq j) - \mathbf{E}_{k-1} \sum_{j=k}^n I(\tau \geq j) \right)$$

Now, the conditional distribution of $\sum_{j=k+1}^n I(\tau \geq j)$ given \mathcal{F}_k is the number of parts following the first k parts $\Gamma_1, \dots, \Gamma_k$. This is the number of parts in a randomly chosen composition of $n - S_k$ where $S_k = \Gamma_1 + \dots + \Gamma_k$. Thus

$$\mathcal{L} \left(\sum_{j=k+1}^n I(\tau \geq j) | \mathcal{F}_k \right) = 1 + \text{Bin} \left(n - 1 - S_k, \frac{1}{2} \right),$$

provided $S_k \leq n-1$, i.e. $\tau \geq k$. In particular, for every $k \geq 0$,

$$\mathbf{E}_k \sum_{j=k+1}^n I(\tau \geq j) = \frac{n+1-S_k}{2} I(\tau \geq k+1).$$

Using that and then $I(\tau \geq k+1) = I(\tau \geq k) - I(\tau = k)$ we obtain

$$\begin{aligned} Y_{n,k} &= I(\tau \geq k)F(\Gamma_k) + \mathbf{E}f(\Gamma) \left(\frac{n+1-S_k}{2} I(\tau \geq k+1) - \frac{n+1-S_{k-1}}{2} I(\tau \geq k) \right) \\ &= I(\tau \geq k)F(\Gamma_k) + \mathbf{E}f(\Gamma) \left(\frac{-\Gamma_k}{2} I(\tau \geq k) - \frac{n+1-S_k}{2} I(\tau = k) \right) \\ &= I(\tau \geq k) \left(F(\Gamma_k) - \frac{\Gamma_k}{2} \mathbf{E}f(\Gamma) \right) + \frac{\mathbf{E}f(\Gamma)}{2} (S_k - n - 1) I(\tau = k) \\ &:= d_{n,k} + e_{n,k}. \end{aligned}$$

Each of the two terms is a martingale difference, and we will show the total contribution to the sum coming from $e_{n,k}$'s is negligible. Writing \Pr_{k-1} for the conditional probability given \mathcal{F}_{k-1} we have

$$\begin{aligned} \Pr_{k-1}(S_k - n - 1 = m, \tau = k) &= \Pr_{k-1}(\Gamma_k + S_{k-1} - n - 1 = m, S_{k-1} \leq n, S_k \geq n) \\ &= I(S_{k-1} < n) \Pr_{k-1}(\Gamma_k + S_{k-1} - n - 1 = m, \Gamma_k \geq n - S_{k-1}) \\ &= I(\tau \geq k) \Pr_{k-1}(\Gamma_k + 1 - (n - S_{k-1}) - 2 = m, \Gamma_k \geq n - S_{k-1}) \\ &= I(\tau \geq k) \Pr(\Gamma + 1 - (n - S_{k-1}) - 2 = m, \Gamma \geq n - S_{k-1}), \end{aligned}$$

where Γ is a GEOM(1/2) random variable and by independence of Γ_k and \mathcal{F}_{k-1} , in the last line S_{k-1} is considered fixed, and \Pr applies to Γ only. Furthermore, by the memoryless property of Γ (see [5, Section 3]), conditionally on $\Gamma \geq \ell \geq 1$, $\Gamma + 1 - \ell$ is equidistributed with Γ . Hence, the last probability above is equal to

$$\Pr(\Gamma \geq n - S_{k-1}) \Pr(\Gamma + 1 - (n - S_{k-1}) - 2 = m | \Gamma \geq n - S_{k-1}) = 2^{S_{k-1} + 1 - n} \Pr(\Gamma - 2 = m).$$

The first three moments of $\Gamma - 2$ are 0 and 2, and 6 which translates into $\mathbf{E}_{k-1} e_{n,k} = 0$ (confirming that $e_{n,k}$ is a martingale difference),

$$(24) \quad \mathbf{E}_{k-1} e_{n,k}^2 = \frac{\mathbf{E}^2 f(\Gamma)}{4} I(\tau \geq k) 2^{S_{k-1} + 1 - n} \cdot 2 = I(\tau \geq k) 2^{S_{k-1} - n} \mathbf{E}^2 f(\Gamma),$$

and

$$(25) \quad \mathbf{E}_{k-1} |e_{n,k}|^3 = 3I(\tau \geq k) 2^{S_{k-1} - 1 - n} \mathbf{E}^3 f(\Gamma),$$

Hence, we immediately obtain

$$(26) \quad \sum_{k \geq 1} \mathbf{E}_{k-1} e_{n,k}^2 = O(1) \cdot \sum_{k=1}^{\tau} 2^{-k} \leq O(1) \cdot \sum_{k=1}^{\infty} 2^{-k} = O(1).$$

This, in turn implies that

$$\begin{aligned} \left| \mathbf{E}_{k-1} Y_{n,k}^2 - \mathbf{E}_{k-1} d_{n,k}^2 \right| &\leq \mathbf{E}_{k-1} e_{n,k}^2 + 2\mathbf{E}_{k-1} |d_{n,k} e_{n,k}| \\ &\leq \mathbf{E}_{k-1} e_{n,k}^2 + 2(\mathbf{E}_{k-1} d_{n,k}^2)^{1/2} (\mathbf{E}_{k-1} e_{n,k}^2)^{1/2}, \end{aligned}$$

where in the last step we used the conditional version of Cauchy–Schwartz inequality. Summing up yields

$$\left| \sum_{k=1}^n \mathbf{E}_{k-1} Y_{n,k}^2 - \sum_{k=1}^n \mathbf{E}_{k-1} d_{n,k}^2 \right| \leq \sum_{k=1}^n \mathbf{E}_{k-1} e_{n,k}^2 + 2 \left(\max_{1 \leq j \leq n} \mathbf{E}_{j-1} d_{n,j}^2 \right)^{1/2} \sum_{k=1}^n (\mathbf{E}_{k-1} e_{n,k}^2)^{1/2} = O(1),$$

Since $\mathbf{E}_{j-1}d_{n,j}^2 = O(1)$ (uniformly in j) and, by the same argument as for (26),

$$\sum_{k=1}^n (\mathbf{E}_{k-1}e_{n,k}^2)^{1/2} = O(1),$$

we infer that

$$(27) \quad \left| \sum_{k=1}^n \mathbf{E}_{k-1}Y_{n,k}^2 - \sum_{k=1}^n \mathbf{E}_{k-1}d_{n,k}^2 \right| = O(1).$$

Now

$$(28)$$

$$\mathbf{E}_{k-1}d_{n,k}^2 = f(\Gamma)^2 = I(\tau \geq k)\mathbf{E}_{k-1}(F(\Gamma_k) - \frac{\Gamma_k}{2}\mathbf{E}f(\Gamma))^2 = I(\tau \geq k)\mathbf{E}(F(\Gamma) - \frac{\Gamma}{2}\mathbf{E}f(\Gamma))^2,$$

and

$$(29) \quad \sum_{k=1}^n \mathbf{E}Y_{n,k}^2 = v(n).$$

Hence we get

$$\frac{\sum_{k=1}^n \mathbf{E}_{k-1}Y_{n,k}^2}{v(n)} = \frac{\sum_{k=1}^n \mathbf{E}_{k-1}d_{n,k}^2 + O(1)}{\sum_{k=1}^n \mathbf{E}d_{n,k}^2 + O(1)} = \frac{\tau}{\mathbf{E}\tau} + O(1/n)$$

which, in view of (21) implies (22). To prove (23) we just write

$$\begin{aligned} \mathbf{E}Y_{n,k}^2 I(|Y_{n,k}| \geq \epsilon\sqrt{v(n)}) &\leq \mathbf{E} \frac{|Y_{n,k}|^3}{\epsilon\sqrt{v(n)}} I(|Y_{n,k}| \geq \epsilon\sqrt{v(n)}) \\ &\leq \frac{c}{\sqrt{v(n)}} \mathbf{E}(|d_{n,k}|^3 + |e_{n,k}|^3) \\ &= O(1/\sqrt{v(n)}), \end{aligned}$$

since both $d_{n,k}$ and $e_{n,k}$ have uniformly bounded third moments (for $d_{n,k}$'s this is clear, and for $e_{n,k}$'s follows from (25)). Hence

$$\sum_{k=1}^n \mathbf{E}Y_{n,k}^2 I(|Y_{n,k}| \geq \epsilon\sqrt{v(n)}) = O(n/\sqrt{v(n)}) = O(\sqrt{n}),$$

which implies (23) and completes the proof.

Remark: Note that (27), (28), and (29) imply that $v(n) \sim wn$, where $w = \mathbf{E}(F(\Gamma) - \frac{\Gamma}{2}\mathbf{E}f(\Gamma))^2/2$. While we did use the linearity of the variance at the beginning of the proof of Theorem 5 we only needed a superlinearity of $v(n)$, which is evident from the recurrence (15).

6. COMPARISON WITH BINARY SPLITS

It may be of some interest to compare the situation with one in which, at every stage, only random binary splits are allowed. This leads to a quicksort type of recurrence. Such recurrences have been thoroughly analyzed in a series of papers, culminating in [13], which gives the most complete picture to date. Since our toll functions fall into ‘‘small toll function’’ category, the limiting distribution is normal, so one only has to find asymptotic mean and the variance. But this can be

readily done. For example, following the usual steps, we obtain that for the binary splits the expected value $f_0(n)$ satisfies

$$\frac{f_0(n)}{n} = \frac{f_0(1)}{1} + \sum_{j=2}^n \frac{1}{j2^j} - \sum_{j=2}^n \frac{j-2}{(j-1)j2^{j-1}} = f_0(1) + \frac{1}{n2^n} + 2 \sum_{j=2}^{n-1} \frac{1}{j(j+1)2^j},$$

which, writing $1/(j(j+1))$ as $1/j - 1/(j+1)$ and using the fact that

$$\sum_{j=1}^{\infty} \frac{1}{j2^j} = \ln 2$$

yields

$$f_0(n) = n \left(f_0(1) + \frac{3}{2} - 2 \ln 2 \right) + O(1/2^n).$$

Also, for binary splits we have

$$t_1(n) = \frac{1}{2^{n-1}} \quad \text{and} \quad t_2(n) = \frac{1}{n-1} + \frac{1}{2^n} \left(1 - \frac{2}{n-1} \right),$$

which gives

$$f_1(n) = n \left(f_1(1) + 3 - 4 \ln 2 \right) + O(1/2^n)$$

and

$$f_2(n) = n \left(f_2(1) + \frac{5}{2} - 3 \ln 2 \right) + O(1/2^n).$$

Numerically, the coefficients in front of linear terms are: $f_0(1) + .113705\dots$, $f_1(1) + .227411\dots$, and $f_2(1) + .4205584\dots$. To compare with the corresponding values for all compositions with at least two parts, see Theorem 3.

REFERENCES

- [1] V. Alexandrov, J. Dongarra, B. Juliano, R. Renner, and C. Tan, editors. *Computational Science - ICCS 2001*, volume 2073 of *Lecture Notes in Computer Science*. Springer, 2001. Session on Architecture-Specific Automatic Performance Tuning.
- [2] K.G. Beauchamp. *Applications of Walsh and related functions*. Academic Press, 1984.
- [3] P. Billingsley. *Probability and Measure*. Wiley, 3rd edition, 1995.
- [4] J. Bilmes, K. Asanović, C.W. Chin, and J. Demmel. Optimizing matrix multiply using PHiPAC: a portable, high-performance, ANSI C coding methodology. In *Proc. Supercomputing*. ACM SIGARC, 1997. <http://www.icsi.berkeley.edu/~bilmes/hipac>.
- [5] G. Casella and R. L. Berger. *Statistical Inference*. Wadsworth & Brooks/Cole, 1990.
- [6] D. F. Elliott and K. R. Rao. *Fast Transforms: Algorithms, Analyses, Applications*. Academic Press, 1982.
- [7] J.A. Fill and S. Janson. Smoothness and decay properties of the limiting quicksort density function. In D. Gardy and A. Mokkadem, editors, *Mathematics and Computer Science: Algorithms, Trees, Combinatorics and Probabilities*, Trends in Mathematics, pages 53–64. Birkhäuser Verlag, 2000.
- [8] P. Flajolet, G. Gonnet, C. Puech, and J. M. Robson. Analytic variations on quad trees. *Algorithmica*, 10:473–500, 1993.
- [9] Matteo Frigo and Steven G. Johnson. FFTW: An adaptive software architecture for the FFT. In *ICASSP '98*, volume 3, pages 1381–1384, 1998. <http://www.fftw.org>.
- [10] P. Hitczenko and G. Louchard. Distinctness of compositions of an integer: a probabilistic analysis. *Random Struct. Alg.*, 19:407–437, 2001.
- [11] P. Hitczenko and C. D. Savage. On the multiplicity of parts in a random composition of a large integer. preprint.
- [12] Hung-Jen Huang. Performance analysis of an adaptive algorithm for the Walsh-Hadamard transform. Master's thesis, Drexel University, 2002.

- [13] H-K. Hwang and R. Neininger. Phase change of limit laws in the quicksort recurrence under varying toll functions. *SIAM J. Comput.*, 31:1687–1722, 2002.
- [14] H-K. Hwang and T-S. Tsai. An asymptotic theory for recurrence relations based on minimalization and maximization. *Theoret. Comput. Sci.*, 290:1475–1501, 2003.
- [15] J. Johnson and M. Püschel. In Search for the Optimal Walsh-Hadamard Transform. In *Proceedings ICASSP*, volume IV, pages 3347–3350, 2000.
- [16] J. R. Johnson, R. W. Johnson, D. Rodriguez, and R. Tolimieri. A methodology for designing, modifying, and implementing Fourier transform algorithms on various architectures. *Circuits, Systems, and Signal Processing*, 9(4):449–500, 1990.
- [17] B. Kheyfets. Expected multiplicity of parts in a random composition. personal communication.
- [18] F.J. MacWilliams and N.J. Sloane. *The theory of error-correcting codes*. North-Holland Publ.Comp., 1992.
- [19] H. Mahmoud. *Sorting. A Distribution Theory*. Wiley, 2000.
- [20] D. Mirković and S. L. Johnsson. Automatic Performance Tuning in the UHFFT Library. In *Proc. ICCS*, LNCS 2073, pages 71–80. Springer, 2001.
- [21] J. M. F. Moura, J. Johnson, R. W. Johnson, D. Padua, V. Prasanna, M. Püschel, and M. M. Veloso. SPIRAL: Portable Library of Optimized Signal Processing Algorithms, 1998. <http://www.ece.cmu.edu/~spiral>.
- [22] M. Régnier. A limiting distribution for quicksort. *RAIRO: Theoretical Informatics and Its Applications*, 23:335–343, 1989.
- [23] U. Röslér. A limit theorem for quicksort. *RAIRO: Theoretical Informatics and Its Applications*, 25:85–100, 1991.
- [24] S. Roura. An improved master theorem for divide and conquer recurrences. *J. of the A.C.M.*, 48:170–205, 2001.
- [25] C. Van Loan. *Computational Frameworks for the Fast Fourier Transform*, volume 10 of *Frontiers in Applied Mathematics*. Society for Industrial and Applied Mathematics, Philadelphia, 1992.
- [26] R. C. Whaley and J. Dongarra. Automatically Tuned Linear Algebra Software (ATLAS). In *Proc. Supercomputing*, 1998. <http://math-atlas.sourceforge.net/>.