



HAL
open science

Good predictions are worth a few comparisons

Nicolas Auger, Cyril Nicaud, Carine Pivoteau

► **To cite this version:**

Nicolas Auger, Cyril Nicaud, Carine Pivoteau. Good predictions are worth a few comparisons. STACS 2016, Feb 2016, Orléans, France. pp.12:1-12:14, 10.4230/LIPIcs.STACS.2016.12 . hal-01212840

HAL Id: hal-01212840

<https://hal.science/hal-01212840v1>

Submitted on 7 Oct 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Good predictions are worth a few comparisons

Nicolas Auger, Cyril Nicaud, and Carine Pivoteau

Université Paris-Est, LIGM (UMR 8049), F77454 Marne-la-Vallée, France

Abstract

Most modern processors are heavily parallelized and use predictors to guess the outcome of conditional branches, in order to avoid costly stalls in their pipelines. We propose predictor-friendly versions of two classical algorithms: exponentiation by squaring and binary search in a sorted array. These variants result in less mispredictions on average, at the cost of an increased number of operations. These theoretical results are supported by experimentations that show that our algorithms perform significantly better than the standard ones, for primitive data types.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases branch misses, binary search, exponentiation by squaring, Markov chains

1 Introduction

As an introductory example, consider the simple problem of computing both the minimum and the maximum of an array of size n . The naive approach is to compare each entry to the current minimum and maximum, which uses $2n$ comparisons. A better solution, in terms of number of comparisons, is to look at the elements of the array two by two, and to compare the smallest to the current minimum and the greatest to the current maximum. This uses only $3n/2$ comparisons, which is optimal.¹

In order to observe the benefit of this optimization, we implemented both versions (see Figure 3) and measured their execution time² for large arrays of uniform random `float` in $[0, 1]$. The results are given in Figure 1 and are very far from what was expected, since the naive implementation is almost twice as fast as the optimized one. Clearly, counting comparisons can not explain these counterintuitive performances. An obvious explanation could be a difference in the number of cache misses. However, both implementations make the same memory accesses, in the same order. Instead, we turn our attention to the comparisons themselves. Most modern processors are heavily parallelized and use predictors to guess the outcome of conditional branches in order to avoid costly stalls in their pipelines. Every time a conditional is used in a program, there is a mechanism that tries to predict whether the corresponding conditional jump will be taken or not. The cost of a misprediction can be quite large compared to a basic instruction, and should be taken into account in order to explain accurately the behavior of algorithms that use a fair amount of comparisons.




Figure 1 Execution time of simultaneous minimum and maximum searching.

¹ More precisely, an adversary argument can be used to establish a lower bound of $\lfloor \frac{3n}{2} \rfloor - 2$ comparisons, in the “decision tree with comparisons” model of computation.

² We used a Linux machine with a 3.40 GHz Intel Core i7-2600 CPU.



In this matter, our example is quite revealing since the trick used to lower the number of comparisons relies on a conditional branch that is unpredictable (for an input taken uniformly at random) and will cause a substantial increase in the number of mispredictions. As we will see in the sequel, the expected number of mispredictions caused by the naive algorithm is $\Theta(\log n)$, whereas it is $\Theta(n)$ for the “optimal” one.

The influence of branch predictors over comparison based algorithms has already been studied, mostly to acknowledge the over-cost induced by mispredictions. Our approach is quite the opposite as we propose to take advantage of this feature, by proposing predictor-friendly versions of two classical algorithms.


Our contributions. After dealing with our introductory example using combinatorial arguments, we turn our attention towards the classical exponentiation by squaring and give a simple alternative algorithm, which reduces the number of mispredictions without increasing the number of multiplications. The analysis is based on the study of the Markov chains that describe the dynamic local predictors (see next section for a brief description of predictors). Finally, in the same vein, we propose biased versions of the binary search in a sorted array. We analyze the expected number of mispredictions for local predictors and we also give the (first to our knowledge) analysis of a global predictor. For these two different problems, we manage to significantly lower the number of mispredictions by breaking the perfect balance usually favored in the divide and conquer strategy. In practice, the trade-off between comparisons and mispredictions allows a noticeable speed-up in the execution time, when the comparisons involve primitive data types, which supports our theoretical results.

Related work. Over the past decade, several articles began to address the influence of branch predictors, and especially the cost of mispredictions, in comparison based algorithms. For instance, Biggar and his coauthors [1] investigated the behavior of branches for many sorting algorithms, in an extensive experimental study. Brodal, Fagerberg and Moruz reviewed the trade-offs between comparisons and mispredictions for several sorting algorithms [3] and studied how the number of inversions in the data affects statistics such as the number of mispredictions [2]. Moreover, these works introduced the first theoretical analysis of static branch predictors.

Also interested by the influence of mispredictions on the running time of sorting algorithms, Sanders and Winkel considered the possibility to dissociate comparisons from branches in their SAMPLESORT, which allows to avoid most of the misprediction cost [13]. Elmasry, Katajainen and Stenmark then proposed a version of MERGESORT that is not affected by mispredictions [6], by taking advantage of some processor-specific instructions.³ The influence of mispredictions was also studied for QUICKSORT: Kaligosi and Sanders gave an in-depth analysis of simple dynamic branch predictors to explain how mispredictions affect this classical algorithm [9]; however, Martínez, Nebel and Wild pointed out that this is not enough to explain the “better than expected” performances of the dual-pivot version of QUICKSORT [11] implemented in Java’s standard library.

Besides, Brodal and Moruz conducted an experimental study of skewed binary search trees in [4], highlighting that such data structures can outperform well-balanced trees, since branching to the right or left does not necessarily have the same cost, due to branch prediction schemes. Our work follows the same line, as we also want to take advantage of the branch predictions, but we focus on algorithms rather than on data structures.

³ Namely, conditional moves, which are now widely available on computers.



■ **Figure 2** Two different 2-bit predictors (left: saturating counter, right: flip-on-consecutive).

2 Elements of computer architecture

To analyze the complexity of searching or sorting algorithms, the standard model consists in counting the number of comparison operations performed. However most modern processors are pipelined. And to avoid stalling the pipeline when coming across a conditional jump, the processor tries to predict if the jump will occur and proceeds according to its prediction. A correctly predicted jump does not stall the pipeline whereas mispredictions lead this one to be flushed, causing a significant performance loss.⁴ Therefore, the cost of a comparison in an “if” statement actually depends on the quality of the prediction.

For any conditional jump, a branch predictor will “guess” if the corresponding branch will be taken or not. For this purpose, many different strategies have been designed. The simplest one is a *static* branch predictor that does not use information from the code execution. It can, for example, predict that all branches will be taken. To improve its accuracy, a *dynamic* branch predictor uses the outcome of past branches to guess whether a particular branch should be taken or not. We now describe different techniques of dynamic branch prediction.


A *1-bit predictor* is a state buffer which remembers the last outcome of the branch; the guess is that the next outcome will be the same. As an improvement, the *2-bit predictors* try to avoid making two mispredictions when a branch takes an unlikely path. Two slightly different schemes are given by Figure 2. The saturating counter scheme can be further improved by keeping more information (*k-bit predictors* using 2^k states). All these predictors are *local*: there is one for every conditional (up to some limit in practice).

An *history table* has 2^n entries indexed by the sequence of the last n branches (1 for taken, 0 otherwise). The entries themselves are usually k -bit predictors. Such a table is said to be *local* when its entries correspond to the behavior of one specific branch and are used for this one only. On the contrary, in a *global* history table, the outcomes of the most recently executed branches are used to index the table, which is shared by all the conditionals.

To get the best of both worlds, *correlating branch predictors* use local and global information mixed together, and *tournament predictors* use an additional dynamic scheme to decide if they follow the local or the global prediction. These types of predictors are far beyond what we study in this article, but are worth mentioning for further analysis.

Strictly speaking, mispredictions can only be analyzed on a given assembly code, as they occur at conditional jumps. In this article, we use C-style pseudo code. We implicitly work on the non-optimized assembly code, where control structures are translated into conditional jumps in the standard way. For our experimental results, we checked that it was indeed the case. Furthermore, we remarked that our good results still hold when considering fully optimized binaries.

⁴ For instance, on an Intel Core i7, a misprediction causes a penalty of about 15 cycles [8].



■ **Figure 3** Naive and optimized implementations of simultaneous maximum and minimum finding.

3 Simultaneous maximum and minimum finding

In this section, we go back to the example given in the introduction: We consider two algorithms that simultaneously compute the minimum and the maximum of an array of length n . These algorithms are given in Figure 3 (see also [5, Sec. 9.1]). For our analysis, we consider the local predictors presented in Section 2.

In the classical settings for the analysis, the algorithm $\frac{3}{2}$ -MINMAX is optimal (see the footnote¹ in the introduction), the number of comparisons performed being asymptotically equivalent to $\frac{3}{2}n$. Obviously, NAIVEMINMAX needs $2n - 2$ comparisons.

In order to give an explanation of the experimental results presented in Figure 1, where NAIVEMINMAX outperforms $\frac{3}{2}$ -MINMAX, we estimate the expected number of mispredictions for both algorithms. Our probabilistic model is the following: we consider the *uniform random distribution on arrays of size n*, where each element is chosen uniformly and independently in $[0, 1]$. Up to an event of probability 0 (when the elements of the input are not pairwise distinct), this is the same as choosing a uniform random permutation of $\{1, \dots, n\}$, since we only use comparisons on the elements in both algorithms.

Recall that a *min-record* (resp. *max-record*) in an array or a permutation is an element that is strictly smaller (resp. greater) than any element to its left. Obviously, in NAIVEMINMAX, the first conditional at line 3 (resp. the one at line 5) is true for each min-record (resp. max-record), except for the first position. The number of records in a random permutation is a well-known statistics, which we can use to establish the following proposition.

► **Proposition 1.** *The expected number of mispredictions performed by NAIVEMINMAX, for the uniform distribution on arrays of size n, is asymptotically equivalent to $4 \log n$ for the 1-bit predictor and to $2 \log n$ for the two 2-bit predictors and the 3-bit saturating counter. The expected number of mispredictions performed by $\frac{3}{2}$ -MINMAX is asymptotically equivalent to $\frac{n}{4}$ for all the considered predictors.*

In light of these results, we observe that the mispredictions occurring in NAIVEMINMAX are negligible towards the number of comparisons. On the other hand, the additional test used to optimize $\frac{3}{2}$ -MINMAX (line 3) causes the number of mispredictions to be comparable to the number of comparisons performed. We believe this is enough to explain why the naive implementation performs better (Figure 1), since we know that mispredictions can

CLASSICALPOW (x,n)	UNROLLEDPOW (x,n)	GUIDEDPOW (x,n)
<pre> 1 r = 1; 2 while (n > 0) { 3 // n is odd 4 if (n & 1) 5 r = r * x; 6 n /= 2; 7 x = x * x; 8 } </pre> <p>x is a floating-point number, n is an integer and r is the returned value.</p>	<pre> 1 r = 1; 2 while (n > 0) { 3 t = x * x; 4 // n₀ == 1 5 if (n & 1) 6 r = r * x; 7 // n₁ == 1 8 if (n & 2) 9 r = r * t; 10 n /= 4; 11 x = t * t; 12 } </pre>	<pre> 1 r = 1; 2 while (n > 0) { 3 t = x * x; 4 if (n & 3) { 5 if (n & 1) 6 r = r * x; 7 if (n & 2) 8 r = r * t; 9 } 10 n /= 4; 11 x = t * t; 12 } </pre>

■ **Figure 4** Three versions of the exponentiation by squaring, in C. The `&` denotes the bitwise AND in the C language.

cost many CPU cycles and that comparisons are cheap operations. Of course, we are aware that other factors can influence the performances of such simple programs, including cache effects. In our implementation, we took care to fetch each element of the array only once and in the same order, so that the cache behavior should not interfere with our results. We also tried the most commonly used optimization of the gcc compiler (`-O3`) to check that these results withstand strong code optimization.⁵ In this particular case, all the branches but the one at line 3 in $\frac{3}{2}$ -MINMAX are replaced by conditional moves that are not vulnerable to misprediction. Hence, $\frac{3}{2}$ -MINMAX still causes approximatively $\frac{1}{4}n$ mispredictions on average.

4 Exponentiation by squaring

We saw in the previous section that conditional branches with equal probabilities of going one way or another are particularly harmful when using branch prediction. Besides, most divide and conquer algorithms feature such branches, since they tend to split problems into parts of equal size to reach an optimal complexity. In the sequel, we explore two different ways of disrupting this balance, to end up with better performances for two classical algorithms: exponentiation by squaring and binary search.

4.1 Modified algorithms

The classical divide and conquer algorithm to compute x^n consists in rewriting $x^n = (x^2)^{\lfloor n/2 \rfloor} x^{n_0}$, where $n_k \dots n_1 n_0$ is the binary decomposition of n , in order to divide the size n of the problem by two. This is the algorithm CLASSICALPOW of Figure 4. As expected, the conditional branch of line 3 is taken with probability $\frac{1}{2}$, which is what we want to avoid.⁶ In order to introduce some imbalance in the algorithm, we first unroll the loop (UNROLLEDPOW, Figure 4) using the decomposition $x^n = (x^4)^{\lfloor n/4 \rfloor} (x^2)^{n_1} x^{n_0}$. Still, both conditional branches are taken with probability $\frac{1}{2}$, but we can now guide the algorithm by injecting the test that determines whether the last two bits of n are 11 or not. This is the

⁵ Both algorithms are faster, as expected, but the naive version is still almost twice as fast.

⁶ In our model, n is chosen uniformly at random between 0 and $4^k - 1$ for some positive k .

Pow	time (in sec.)	loops $\times 10^9$	mult. $\times 10^9$	branches $\times 10^9$	mispred. $\times 10^9$
classical	7.230	1.250	1.900	1.300	0.674
unrolled	6.316	0.633	1.917	1.317	0.683
guided	5.606	0.633	1.917	1.658	0.554

■ **Figure 5** Some parameters measured during 5.10^7 computations of x^n with the three algorithms of Figure 4, using the PAPI library.⁷ The number of branches is given excluding the ones caused by loops, since these branches do not yield mispredictions.

third algorithm of Figure 4. Note that this conditional branch (line 4) is absolutely unnecessary in the algorithm, as it is redundant with the tests of line 5 and 7. But on the other hand, this branch is taken with probability $\frac{3}{4}$ and the branches of line 5 and 7 are now both taken with probability $\frac{2}{3}$. This is how we aim at using the branch predictions.

To compare their performances experimentally, we computed the floating-point value of x^n using each of the algorithms 5.10^7 times, with n chosen uniformly at random in $\{0, \dots, 2^{26} - 1\}$. We measured the execution time, as well as some other parameters given by the latest version of the PAPI library,⁷ which give access, for instance, to the number of mispredictions occurring during the execution. These results are depicted on Figure 5. The first observation is that GUIDEDPOW is 14% faster than UNROLLEDPOW and 29% faster than CLASSICALPOW and yet, the number of multiplications performed is essentially the same for the three algorithms. The main explanation we have come across for the speed-up between UNROLLEDPOW and CLASSICALPOW is that the number of loops is divided by two. As for GUIDEDPOW, the number of loops is the same as for UNROLLEDPOW and it uses 25% more comparisons, but still the guided version is faster. The main difference between the two is that the test added at line 4 allows to decrease the number of mispredictions by about a quarter. We are in similar settings as for the simultaneous minimum and maximum, where the increased number of comparisons is balanced by less mispredictions. We now proceed with the analysis of this phenomenon.


4.2 Analysis of the average number of mispredictions for GuidedPow

For the analysis, we consider that n is taken uniformly at random in $\{0, \dots, N - 1\}$, for $N = 4^k$ and with $k \geq 1$. This model is exactly the same as choosing each of the $2k$ bits of the binary representation of n uniformly at random and independently. We consider the local predictors presented in Section 2.

Let $L_k(n)$ be the number of loop iterations of GUIDEDPOW. This is a random variable, which is easy to analyze since it is equal to the smallest integer ℓ such that 4^ℓ is greater than n . In particular, we have $\mathbb{E}[L_k] = k - \frac{1}{3} + o(1) \sim k$.

We now recall, using our algorithm as an example, why Markov chains are the key tools for that kind of analysis (as done in [9, 11]). Let us consider the first conditional of line 4. In our model, at each iteration, the condition is true with probability $\frac{3}{4}$, as it is not satisfied when the last two bits are 00. It yields that the behavior of the predictor associated to this conditional is exactly described by the Markov chain obtained when changing the edges labels “taken” by $\frac{3}{4}$ and the labels “not taken” by $\frac{1}{4}$ (see Figure 6). A misprediction occurs whenever an edge labeled by “taken” (resp. “not taken”) is used from a state that predicts

⁷ PAPI 5.4.1.0, see <http://icl.cs.utk.edu/papi>.



■ **Figure 6** The saturating counter and its associated Markov chain for the first conditional of GUIDEDPOW. The bold edges correspond to mispredictions.

“not taken” (resp. “taken”). We also need to know the initial state of the predictor, but it has no influence on our asymptotic results, as we shall see.

Hence, we reduced our problem to counting the number of times some particular edges are taken in a Markov chain, when we perform a random walk of (random) length L_k . We can therefore conclude using the classical Ergodic Theorem [10], which we restated below in order to fit our needs.

► **Theorem 2 (Ergodic Theorem).** *Let (M, π_0) be a primitive and aperiodic Markov Chain on the finite set S . Let π be its stationary distribution. Let E be a set of edges of M , that is, a set of pairs $(i, j) \in S^2$ such that $M(i, j) > 0$.*

For any nonnegative integer n , let L_n be a random variable on nonnegative integers such that $\lim_{n \rightarrow \infty} \mathbb{E}[L_n] = +\infty$. Let X_n be the random variable that counts the number of edges in E that are used during a random walk of length L_n in M (starting from the initial distribution π_0). Then the following asymptotic equivalence holds: $\mathbb{E}[X_n] \sim \mathbb{E}[L_n] \sum_{(i,j) \in E} \pi(i)M(i, j)$.

When considering a given predictor, under the model where the condition is satisfied with probability p , we denote by M_p its transition matrix, by π_p its stationary vector and by $\mu(p)$ its *expected misprediction probability* defined by $\mu(p) = \sum_{(i,j) \in E} \pi_p(i)M_p(i, j)$, where E is the set of edges corresponding to mispredictions. As shown in [11], if we denote by $\mu_1(p)$, $\mu_2(p)$ and $\mu'_2(p)$ the expected misprediction probability of the 1-bit, 2-bit saturating counter and the flip-on-consecutive 2-bit, then we have:

$$\mu_1(p) = 2p(1-p); \quad \mu_2(p) = \frac{p(1-p)}{1-2p(1-p)}; \quad \mu'_2(p) = \frac{2p^2(1-p)^2 + p(1-p)}{1-p(1-p)}. \quad (1)$$

Similarly, the expected misprediction probability $\mu_3(p)$ of the 3-bit saturated counter is

$$\mu_3(p) = \frac{p(1-p)(1-3p(1-p))}{1-2p(1-p)(2-p(1-p))}. \quad (2)$$

Applying these mathematical tools to GUIDEDPOW yields the following results. The theorem is stated for values of N that are not powers of 4, which is more complicated since the bits are not exactly 0's and 1's with probability $\frac{1}{2}$ (and not independent). In Section 5 we show how to deal with the cases where we slightly deviate from the ideal case.

► **Theorem 3.** *Assume that n is taken uniformly at random in $\{0, \dots, N-1\}$. The expected number of conditional tests in CLASSICALPOW and UNROLLEDPOW is asymptotically equivalent to $\log_2 N$, whereas it is asymptotically equivalent to $\frac{5}{4} \log_2 N$ for GUIDEDPOW. The expected number of mispredictions is asymptotically equivalent to $\frac{1}{2} \log_2 N$ for CLASSICALPOW and UNROLLEDPOW, for any kind of predictor. For GUIDEDPOW, it is asymptotically equivalent to $\alpha \log_2 N$, where $\alpha = \frac{1}{2}\mu(3/4) + \frac{3}{4}\mu(2/3)$, where μ is the expected misprediction probability associated to the local predictor.*

<p style="text-align: center; margin: 0;">BIASEDBINARYSEARCH</p> <pre style="margin: 0;"> 1 d = 0; f = n; 2 while (d < f){ 3 m = (3*d+f)/4; 4 if (T[m] < x) 5 d = m+1; 6 else 7 f = m; 8 } 9 return f; </pre>	<p style="text-align: center; margin: 0;">SKEWSEARCH</p> <pre style="margin: 0;"> 1 d = 0; f = n; 2 while (d < f){ 3 m1 = (3*d+f)/4; 4 if (T[m1] > x) 5 f = m1; 6 else { 7 m2 = (d+f)/2; 8 if (T[m2] > x){ 9 f = m2; 10 d = m1+1; 11 } 12 else d = m2+1; 13 } 14 } 15 return f; </pre>
--	--

In both cases, T is an array of floats of size n and x is the number that is searched for. The classical binary search is obtained by replacing line 3 of BIASEDBINARYSEARCH by $m = (d+f)/2$;

■ **Figure 7** Algorithms for the biased binary search and skew search. Both return the position where the element should be inserted.

Using Theorem 3 and Equations (1) and (2), we get that α is equal to $\frac{25}{48} \approx 0.52$, $\frac{9}{20} = 0.45$, $\frac{2045}{4368} \approx 0.47$ and $\frac{1095}{2788} \approx 0.39$ for the 1-bit, 2-bit saturated, flip-on-consecutive 2-bit and 3-bit saturated counter, respectively. These values are to be compared with the $\frac{1}{2}$ of the other two algorithms. In particular, for the 1-bit predictor, the expected number of mispredictions is greater for GUIDEDPOW than for CLASSICALPOW or UNROLLEDPOW. This predictor is not efficient enough to offset the mispredictions caused by the additional conditional. For the 3-bit saturated counter, GUIDEDPOW therefore uses $\approx 0.25 \log_2 n$ more comparisons than UNROLLEDPOW, but $\approx 0.11 \log_2 n$ less mispredictions.


5 Binary search and variants

5.1 Unbalancing the binary search

We first consider the classical binary search which partitions a sorted array of size n into two parts of size $\frac{n}{2}$ and compares the value x that is searched for to the middle of the array in order to determine in which part of the array to continue the search. As before, if we consider arrays of uniform random floating-point numbers, we get a conditional branch that is taken with probability $\frac{1}{2}$. A simple way to change that is to partition another way, for instance with parts of size about $\frac{n}{4}$ and $\frac{3n}{4}$, as in the BIASEDBINARYSEARCH (see Figure 7). Carrying on with the divide and conquer strategy but partitioning the array into three parts of size about $\frac{n}{3}$, gives a ternary search. The main issue with this approach is that, in practice, the division by 3 is costly in terms of hardware. Thus, to limit the cost of partitioning, we choose to slice the array into two parts of size $\frac{n}{4}$ and one part of size $\frac{n}{2}$. This can be done using only divisions by powers of two, which are simple binary shifts, as in the initial binary search (see SKEWSEARCH in Figure 7).

5.2 Experiments

As expected at this point in our work, the BIASEDBINARYSEARCH experimentally performs better than the classical binary search and the SKEWSEARCH performs much better. Unlike our previous examples, the changes we brought in the binary search are quite sensitive to



■ **Figure 8** Execution time of the three searching algorithms of Figure 7 for small-size arrays (that fit in the first-level cache) and medium-size arrays (that fit in the last-level cache).

cache effects, since the way we partition the array influences the location where the memory is accessed. Thus we conducted experiments on arrays that fit in the last-level cache of our machine² in order to mostly measure the effects of branch prediction. The results are given by Figure 8 and we can see that, for medium-size arrays, SKEWSEARCH is up to 23% faster than the binary search (program compiled with gcc without optimization, in order to keep track of what really happens during the execution). Experiments in JAVA using a dedicated micro-benchmarking library⁸ gave roughly the same results (but with a lesser speedup of about 12%), when comparing our skew search to the implementation of the binary search on doubles in the standard library.

5.3 Local predictors analysis

As in Section 4, we aim at using the Ergodic Theorem (page 7) to obtain a good asymptotic estimate of the number of mispredictions. We therefore need to compute the expected number of times each given conditional is performed, in our different algorithms. We consider that each possible output is equally likely (*i.e.* the uniform distribution on $\{0, \dots, n\}$).

A first order estimation of the expected number of times a given conditional is executed can be obtained using the following version of Roura’s Master Theorem [12], which has been simplified for our specific case:⁹


► **Theorem 4** (Master Theorem). *Let $k \geq 1$, and a_1, \dots, a_k and b_1, \dots, b_k be positive real numbers such that $\sum_{i=1}^k a_i = 1$. For every $i \in \{1, \dots, k\}$, let also $\varepsilon_i(n)$ be a real valued sequence such that $b_i n + \varepsilon_i(n)$ is a positive integer and $\varepsilon_i(n) = \mathcal{O}(\frac{1}{n})$. Let $T(n)$ be the real valued sequence that satisfies, for some positive constants c and d ,*

$$T(0) = c \text{ and } T(n) = d + \sum_{i=1}^k a_i T(b_i n + \varepsilon_i(n)) + \mathcal{O}\left(\frac{\log n}{n}\right) \text{ for } n \geq 1.$$

Then $T(n) \sim \frac{d}{h} \log n$, with $h = -\sum_{i=1}^k a_i \log b_i$.

⁸ Benchmark using jmh: <http://openjdk.java.net/projects/code-tools/jmh/>. Our algorithms are compared to `Arrays.binarySearch(double[] a, double key)`.

⁹ For more general statements, we refer the reader to the seminal work of Roura [12].



■ **Figure 9** The decomposition tree of BIASEDBINARYSEARCH for $n = 8$.

Before stating our main result, we describe the main steps of our analysis on the algorithm BIASEDBINARYSEARCH. The expected number of iterations $L(n)$ of BIASEDBINARYSEARCH satisfies the relation

$$L(n) = 1 + \frac{a_n}{n+1}L(a_n) + \frac{b_n}{n+1}L(b_n), \text{ with } a_n = \left\lfloor \frac{n}{4} \right\rfloor + 1, b_n = \left\lceil \frac{3n}{4} \right\rceil \text{ and } L(0) = 0.$$

Thus, Theorem 4 applies and $L(n) \sim \lambda \log n$, with $\lambda = \frac{4}{4 \log 4 - 3 \log 3} \approx 1.78$.


Unfortunately, we cannot directly transform the predictor into a Markov chain as we did in Section 4, because the probabilities $\frac{a_n}{n+1}$ and $\frac{b_n}{n+1}$ are not fixed anymore (they slightly depend on n). However, since $\frac{a_n}{n+1} = \frac{1}{4} + \mathcal{O}(\frac{1}{n})$ and $\frac{b_n}{n+1} = \frac{3}{4} + \mathcal{O}(\frac{1}{n})$, this Markov chain should still yields a good approximation of the number of mispredictions with Theorem 2.

A convenient way to prove this formally is to introduce the *decomposition tree* \mathcal{T} associated with the search algorithms, which is defined as follows. If the input has size n , its root is labeled by the pair $(0, n)$, and each node corresponds to the possible values of d and f during one loop of the algorithm. The leaves are the pairs (i, i) , for $i \in \{0, \dots, n\}$; they are identified with the output of the algorithm in $\{0, \dots, n\}$. There is a direct edge between (d, f) and (d', f') whenever the variables d and f can be changed into d' and f' during the current iteration of the loop. Such an edge is labeled with the probability $\frac{f'-d'+1}{f-d+1}$, which is the probability that this update happens in our model. An example of such a decomposition tree for BIASEDBINARYSEARCH is depicted on Figure 9.

By construction, following a path from the root to a leaf, by choosing between left and right according to the edge probability is exactly the same as choosing an integer uniformly at random in $\{0, \dots, n\}$. Let $u = (u_0, u_1, \dots)$ be a infinite sequence of elements of $[0, 1]$ taken uniformly at random and independently. To u is associated its path $\mathbf{Path}_n(\mathcal{T}, u)$ in \mathcal{T} where, at step i , we go to the left if u_i is smaller than the left child edge probability and to the right otherwise. Let $L_n(\mathcal{T}, u)$ be the length of $\mathbf{Path}_n(\mathcal{T}, u)$. Let also $\mathbf{Path}_n(\mathcal{I}, u)$ be the path following the values in u in the ideal (infinite) tree \mathcal{I} where we go to the left with probability $\frac{1}{4}$ and to the right with probability $\frac{3}{4}$. Then the following result holds.

► **Lemma 5.** *The probability that $\mathbf{Path}_n(\mathcal{T}, u)$ and $\mathbf{Path}_n(\mathcal{I}, u)$ differ at one of the first $L_n(\mathcal{T}, u) - \sqrt{\log n}$ steps is $\mathcal{O}(\frac{1}{\log n})$.*

Hence, the algorithm BIASEDBINARYSEARCH behaves almost like the idealized version, for most of the iterations of its main loop, and we have a sufficiently precise estimation of the error term. This is enough to prove that the idealized version is a correct first order approximation of the number of mispredictions. The same construction can be done for all three algorithms, yielding the following result.




■ **Figure 10** A fully global predictor scheme: The history table of size 2^ℓ keeps track of the outcomes of the last ℓ branches encountered during the execution, the last one corresponding to the rightmost bit. To each sequence of ℓ branches is associated a global 2-bit predictor (shared by all the conditional branches).

► **Theorem 6.** Let C_n and M_n be the number of comparisons and mispredictions performed in our model of randomness. For BINARYSEARCH, $\mathbb{E}[C_n] \sim \frac{1}{\log 2} \log n$ and $\mathbb{E}[M_n] \sim \frac{1}{2 \log 2} \log n$. For BIASEDBINARYSEARCH, $\mathbb{E}[C_n] \sim \frac{4}{4 \log 4 - 3 \log 3} \log n$ and $\mathbb{E}[M_n] \sim \mu(\frac{1}{4})\mathbb{E}[C_n]$. For SKEWSEARCH, $\mathbb{E}[C_n] \sim \frac{7}{6 \log 2} \log n$ and $\mathbb{E}[M_n] \sim (\frac{4}{7}\mu(\frac{1}{4}) + \frac{3}{7}\mu(\frac{1}{3}))\mathbb{E}[C_n]$, where μ is the expected misprediction probability associated with the predictor.

5.4 Analysis of the global predictor for skewSearch

In this section we intend to give hints about the behavior of a global branching predictor, such as the one depicted on Figure 10 (see also Section 2), for the algorithm SKEWSEARCH. Notice in particular that the predictor of each entry is a 2-bit saturated counter. This is not the only possible choice of a global predictor, but it is simple enough without being trivial. We make the analysis in the idealized framework that resemble the real case sufficiently well, by ignoring the rounding effects of dealing with integers. We saw in the previous section why these approximations still give the correct result for the first order asymptotic.

In our idealized model we only consider the sequence of taken / not taken produced by the two conditional tests of SKEWSEARCH. We deliberately do not consider the conditional induced by the test within the “while” loop, which would be always not taken in our settings (except for the very last step). Adding it would complicate the model without adding interesting information to the branch predictor.¹⁰ We encode a taken conditional by a 1 and a not taken conditional by a 0. The trace of an execution of the algorithm is thus a non-empty word on the binary alphabet $B = \{0, 1\}$. Because of the way the two conditional tests are nested within the algorithm, we can keep track of the current “if” by the use of the simple deterministic automaton \mathcal{A}_{if} with two states depicted in Figure 11: **main** stands for the first conditional and **nested** for the second one. In our model, **main** is taken with probability $\frac{1}{4}$ and **nested** with probability $\frac{1}{3}$. As done in Section 4, \mathcal{A}_{if} can be changed into a Markov chain \mathcal{M}_{if} using this transition probabilities. A direct computation shows that its stationary vector π_{if} satisfies $\pi_{\text{if}}(\text{main}) = \frac{4}{7}$ and $\pi_{\text{if}}(\text{nested}) = \frac{3}{7}$.



■ **Figure 11** On the left, the automaton \mathcal{A}_{if} . On the right, the Markovian automaton \mathcal{M}_{if} of transition probabilities $\mathbb{P}(1 \mid \text{main}) = \frac{1}{4}$, $\mathbb{P}(0 \mid \text{main}) = \frac{3}{4}$, $\mathbb{P}(0 \mid \text{nested}) = \frac{2}{3}$ and $\mathbb{P}(1 \mid \text{nested}) = \frac{1}{3}$.

¹⁰ Also, most modern architectures have “loop detectors” that are used to identify such conditionals.

For the same reason as above, in the global table, we only record the history for the two conditionals `main` and `nested`. Let ℓ denote the history length, that is, the number of bits used in the history table of Figure 10. We assume that ℓ is even. An *history* h is thus seen as a binary word of length ℓ . Let 0^ℓ be the history made of 0's only.

When a conditional is tested at time t , the predictor uses the entry at position h_t to make the prediction, where h_t is the current history. To follow the evolution of the algorithm at time $t + 1$, we therefore only have to keep track of (1) the history table T_t , (2) the current history h_t and (3) which of the two conditionals IF_t is under consideration. Knowing IF_t is required in order to compute the probability that the next outcome is 0 or 1. This defines a Markov chain \mathcal{M}_{up} for the updates in the history table. From \mathcal{M}_{up} , one can theoretically estimate the expected number of mispredictions using Theorem 2, as we did for local predictors. The main issue with this approach is that computing π_{up} is typically in $\mathcal{O}(m^3)$, where m is the number of states of \mathcal{M}_{up} . Since the number of states is exponential in ℓ , the computations are completely intractable for reasonable history lengths (such as $\ell \geq 6$), even if we first remove the unreachable states. In the sequel, we therefore use the particular structure of \mathcal{M}_{up} to directly compute the typical number of mispredictions.

Let $h \in B^\ell$ be an history that is not equal to 0^ℓ . There is at least one 1 in h . Since reading a 1 always send to state `main` in \mathcal{A}_{if} , we know for sure the conditional IF_t under consideration when an occurrence of h has just happened at time t . Hence, we know the probability to have a 0 or a 1 at time $t + 1$, given that $h_t = h$. As a consequence, each entry of $h \neq 0^\ell$ in the table T behaves like a fixed-probability local 2-bit saturating predictor, with probability $\frac{1}{4}$ (resp. $\frac{1}{3}$) for histories associated to `main` (resp. to `nested`). Therefore, $h = 0^\ell$ concentrates all the differences between the local and the global predictors.

What happens for the entry 0^ℓ is well described by considering the automaton on pairs (s, i) , where s is a state of the predictor and i is the current conditional. This automaton can be turned into a Markov chain, and the Ergodic Theorem yields a precise estimation of the number of mispredictions. Following this idea yields the following result.

► **Theorem 7.** *For the global predictor, the average number of mispredictions caused during SKEWSEARCH on an input of size n is asymptotically equivalent to $(\frac{12}{35} + \frac{1}{595 \cdot 2^\ell})\mathbb{E}[C_n]$.*

By Theorem 6, if we use a local 2-bit predictor for each conditional, the expected number of mispredictions is asymptotically equivalent to $\frac{12}{35}\mathbb{E}[C_n]$. The difference with the global predictor is therefore extremely small, which is not surprising as there is a difference only when the history is 0^ℓ . However, if there is a competition between a global predictor and a more accurate local predictor (a 3-bit saturated counter for instance), then the local predictor performs better; it is probably slightly disrupted by the global one, as the dynamic selector between both predictors can choose to follow the global predictor from time to time.

6 Conclusion

In this article we propose unbalanced predictor-friendly versions of two very classical algorithms, namely the exponentiation by squaring and the binary search. Using a precise estimation on the expected number of mispredictions, we show that our new algorithms are worth considering when the cost of a comparison is reasonable compared to the cost of a misprediction. This is typically the case for primitive data types.

We believe that these theoretical results, supported by experiments, advocate strongly for considering this particular feature of modern computers in the design and analysis of algorithms: we showed that taking branch prediction into account can yield significant improvements, even on very classical algorithms.

References

- 1 Paul Biggar, Nicholas Nash, Kevin Williams, and David Gregg. An experimental study of sorting and branch prediction. *Journal of Experimental Algorithmics*, 12:1, June 2008.
- 2 Gerth Stølting Brodal, Rolf Fagerberg, and Gabriel Moruz. On the adaptiveness of quicksort. *ACM Journal of Experimental Algorithmics*, 12, 2008.
- 3 Gerth Stølting Brodal and Gabriel Moruz. Tradeoffs Between Branch Mispredictions and Comparisons for Sorting Algorithms. In *Algorithms and Data Structures*, volume 3608, pages 385–395. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.
- 4 Gerth Stølting Brodal and Gabriel Moruz. Skewed Binary Search Trees. In *Algorithms – ESA 2006*, volume 4168, pages 708–719. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.
- 5 Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. MIT Press, Cambridge, MA, third edition, 2009.
- 6 Amr Elmasry, Jyrki Katajainen, and Max Stenmark. Branch Mispredictions Don’t Affect Mergesort. In *Experimental Algorithms*, volume 7276, pages 160–171. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.
- 7 Philippe Flajolet and Robert Sedgewick. *Analytic Combinatorics*. Cambridge University Press, 2009.
- 8 John L. Hennessy and David A. Patterson. *Computer Architecture, Fifth Edition: A Quantitative Approach*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 5th edition, 2011.
- 9 Kanela Kaligosi and Peter Sanders. How Branch Mispredictions Affect Quicksort. In *Algorithms – ESA 2006*, volume 4168, pages 780–791. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.
- 10 David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2008.
- 11 Conrado Martínez, Markus E. Nebel, and Sebastian Wild. Analysis of branch misses in quicksort. In *Proceedings of the Twelfth Workshop on Analytic Algorithmics and Combinatorics, ANALCO 2015, San Diego, CA, USA, January 4, 2015*, pages 114–128. SIAM, 2015.
- 12 Salvador Roura. Improved master theorems for divide-and-conquer recurrences. *Journal of the ACM*, 48(2):170–205, 2001.
- 13 Peter Sanders and Sebastian Winkel. Super scalar sample sort. In *Algorithms – ESA 2004*, volume 3221 of *Lecture Notes in Computer Science*, pages 784–796. Springer Berlin Heidelberg, 2004.

A Appendix

A.1 Programs

The programs we used for our experiments can be found here: <http://www-igm.univ-mlv.fr/~pivotau/benchmarks/>.

A.2 Proofs

A.2.1 Proof of Proposition 1

Proposition 1. *The expected number of mispredictions performed by NAIVEMINMAX for the uniform distribution on arrays of size n is asymptotically equivalent to $4\log n$ for the 1-bit predictor and to $2\log n$ for the two 2-bit predictors and the 3-bit saturating counter.*

The expected number of mispredictions performed by $\frac{3}{2}$ -MINMAX for the uniform distribution on arrays of size n is asymptotically equivalent to $\frac{n}{4}$ for all the considered predictors.

Proof. For a given positive n , let $[n]$ be the set $\{1, \dots, n\}$. It is convenient for this proof to use the correspondence between cycles and records in a permutation. If one sees a permutation σ as a one-to-one mapping from $[n]$ to $[n]$, its underlying graph¹¹ is a set of labeled directed cycles. If $\mathcal{C} = c_1 \rightarrow c_2 \rightarrow \dots \rightarrow c_\ell \rightarrow c_1$ is such a cycle of length ℓ , whose smallest element is c_i , let $f(\mathcal{C})$ be the sequence defined by $f(\mathcal{C}) = c_i c_{i+1} \dots c_\ell c_1 \dots c_{i-1}$. In other words, we read the cycle starting from its smallest element. If the cycles $\mathcal{C}_1, \dots, \mathcal{C}_m$ of σ are ordered by decreasing order of their smallest element, then $f(\sigma)$ is defined by $f(\sigma) = f(\mathcal{C}_1) \cdot f(\mathcal{C}_2) \cdot \dots \cdot f(\mathcal{C}_m)$. Classically, f is a bijection from the set \mathfrak{S}_n of permutations of size n onto itself, such that the number of min-records in $f(\sigma)$ is equal to the number of cycles of σ . Hence, the expected number of records in a uniform random element of \mathfrak{S}_n is asymptotically equivalent to $\log n$ (see [7]).

In the sequel, we will use the fact that f is a bijection by remarking that if ξ is a real valued random variable on uniform random permutations, the expectation of ξ satisfies:

$$\mathbb{E}_n[\xi] := \frac{1}{n!} \sum_{\sigma \in \mathfrak{S}_n} \xi(\sigma) = \frac{1}{n!} \sum_{\sigma \in \mathfrak{S}_n} \xi(f(\sigma)).$$

We first consider the 1-bit predictor. Let σ be a permutation of \mathfrak{S}_n whose cycles, ordered by decreasing order of their smallest element, are $\mathcal{C}_1, \dots, \mathcal{C}_m$. We want to estimate the number of mispredictions $\xi(f(\sigma))$ caused by line 3 of NAIVEMINMAX applied to σ . As $f(\sigma) = f(\mathcal{C}_1) \cdot f(\mathcal{C}_2) \cdot \dots \cdot f(\mathcal{C}_m)$, we count the mispredictions cycle by cycle. Assume that the predictor is on state NT (not taken) just before processing $f(\mathcal{C}_i)$. As the first element of $f(\mathcal{C}_i)$ is a min-record, it causes a misprediction and the predictor switch to state T (taken). If \mathcal{C}_i has length at least 2, the second element of $f(\mathcal{C}_i)$ is greater than the first one: it also causes a misprediction and the predictor goes back to state NT . The remaining elements of \mathcal{C}_i , if any, are all greater than its first element and therefore cause no more misprediction. Then, either \mathcal{C}_i has length 1, and the process of $f(\mathcal{C}_i)$ causes 1 misprediction leaving the predictor on state T , or \mathcal{C}_i has length at least 2 and it causes 2 mispredictions, leaving the predictor on state NT . A similar study can be done when the starting state of the predictor is T , yielding the following table:

¹¹The set of vertices is $[n]$ and there is an edge $i \rightarrow j$ whenever $\sigma(i) = j$.

starting state	cycle of length 1		cycle of length ≥ 2	
	mispred.	ending state	mispred.	ending state
NT	1	T	2	NT
T	0	T	1	NT

For any positive integer i , let $\mathbf{Cyc}_i(\sigma)$ be the number of cycles of length i of σ . Let $\mathbf{Cyc}(\sigma) = \sum_{i \geq 1} \mathbf{Cyc}_i(\sigma)$ be the number of cycles of σ . We see the table above as follows: there is one misprediction if the starting state is NT , plus one misprediction if the cycle has length at least 2. As the starting state is NT when the previous cycle has length at least 2, we get that the number of mispredictions caused by $f(\sigma)$ at line 3 is $\xi(f(\sigma)) = 2 \mathbf{Cyc}(\sigma) - 2 \mathbf{Cyc}_1(\sigma) + \mathcal{O}(1)$, where the $\mathcal{O}(1)$ captures the border effects (initial configuration of the predictor and whether the last cycle of σ has length 1 or not). This concludes the proof for the 1-bit predictor, since $\mathbb{E}_n[\mathbf{Cyc}] \sim \log n$ and $\mathbb{E}_n[\mathbf{Cyc}_1] \rightarrow 1$, as n tends to infinity. There is a factor two since we have to add the number of mispredictions of both conditionals.

For the 2-bit saturating counter predictor, we use the same technique. If, for instance, the starting state is NT when beginning the process of $f(\mathcal{C}_i)$, then there is a misprediction and the predictor switches to T . If \mathcal{C}_i has length at least 2, then the next element also causes a misprediction and the predictor goes back to NT . If it has length at least 3, then the predictor is set to SNT with no further misprediction. All useful cases are depicted in the following table:

starting state	cycle of length 1		cycle of length 2		cycle of length 3		cycle of length ≥ 4	
	mispred.	ending	mispred.	ending	mispred.	ending	mispred.	ending
SNT	1	NT	1	SNT	1	SNT	1	SNT
NT	1	T	2	NT	2	SNT	2	SNT
T	0	ST	1	T	2	NT	2	SNT
ST	0	ST	1	T	2	NT	2	SNT

From this, we readily get that the number $\chi(f(\sigma))$ of mispredictions caused by the first conditional satisfies

$$\mathbf{Cyc}_{\geq 4}(\sigma) \leq \chi(f(\sigma)) \leq \mathbf{Cyc}_{\geq 4}(\sigma) + 3 \mathbf{Cyc}_{\leq 3}(\sigma),$$

since there can be two mispredictions caused by $f(\mathcal{C}_i)$ only if \mathcal{C}_{i-1} has length at most 3. This concludes the proof for this predictor, by summing the contribution of both conditionals, as $\mathbb{E}_n[\mathbf{Cyc}_{\leq 3}] = \mathcal{O}(1)$. The flip-on-consecutive 2-bit predictor and the 3-bit saturating counter are analyzed the same way, yielding the same results.

We now consider the algorithm $\frac{3}{2}$ -MINMAX. Using the model of n random numbers in $[0, 1]$, it is straightforward to see that the first test in the loop of $\frac{3}{2}$ -MINMAX (line 3) causes a misprediction with probability $\frac{1}{2}$, for every predictors. Hence, this first test causes around $\frac{n}{4}$ mispredictions in average, when the algorithm ranges through the whole input. Moreover, the inner tests are true only when a min-record or max-record occurs. Using the same kinds of arguments as for Lemma 1, the expected number of mispredictions caused by these inner tests is in $\mathcal{O}(\log n)$, concluding the proof. \blacktriangleleft

A.2.2 Proof of Theorem 2

Theorem 2. *Let (M, π_0) be a primitive and aperiodic Markov Chain on the finite set S . Let π be its stationary distribution. Let also E be a set of edges of M , that is, a set of*

pairs $(i, j) \in S^2$ such that $M(i, j) > 0$.

For any nonnegative integer n , let L_n be a random variable on nonnegative integers such that $\lim_{n \rightarrow \infty} \mathbb{E}[L_n] = +\infty$. Let X_n be the random variable that counts the number of edges in E that are used during a random walk of length L_n in M (starting from any given distribution π_0). Then the following asymptotic equivalence holds: $\mathbb{E}[X_n] \sim \mathbb{E}[L_n] \sum_{(i,j) \in E} \pi(i)M(i, j)$.

Proof. From the classical Ergodic Theorem [10, p. 58], we get that if F is a subset of S , and if Y_n counts the number of times an element of F is met during a random walk of length n in M , then

$$\frac{1}{n} \mathbb{E}[Y_n] \xrightarrow{n \rightarrow \infty} \sum_{i \in F} \pi(i).$$

We just have to modify it so that it works for edges and for random walks of random lengths.

Let B be the set of all edges of M . We consider the Markov Chain of order two M_2 obtained from M as follows: its set of states is B , and the only positive values of M_2 are $M_2(i \rightarrow j, j \rightarrow k) = M(j, k)$, for every $i \rightarrow j$ and $j \rightarrow k$ in B . It is straightforward to verify that M_2 is irreducible and aperiodic, and that its stationary vector π_2 satisfies $\pi_2(i \rightarrow j) = \pi(i)M(i, j)$. Applying the classical Ergodic Theorem on M_2 for the set of edges E yields, if Z_ℓ counts the number of times an edge of E is used during a random walk of length ℓ in M :

$$\mathbb{E}[Z_\ell] = \ell \left(\sum_{i \rightarrow j \in E} \pi(i)M(i, j) \right) (1 + \epsilon_\ell),$$

with $\epsilon_\ell \rightarrow 0$ as ℓ tends to infinity.

We now prove the result for variable length random walks:

$$\begin{aligned} \mathbb{E}[X_n] &= \sum_{\ell \geq 0} \mathbb{P}(L_n = \ell) \mathbb{E}[X_n | L_n = \ell] \\ &= \sum_{\ell \geq 0} \mathbb{P}(L_n = \ell) \mathbb{E}[Z_\ell] \\ &= \sum_{\ell \geq 0} \mathbb{P}(L_n = \ell) \ell \left(\sum_{i \rightarrow j \in E} \pi(i)M(i, j) \right) (1 + \epsilon_\ell) \\ &= (\mathbb{E}[L_n] + \mathbb{E}[L_n \epsilon_{L_n}]) \left(\sum_{i \rightarrow j \in E} \pi(i)M(i, j) \right). \end{aligned}$$

At this point we only have to prove that $\mathbb{E}[L_n \epsilon_{L_n}] = o(\mathbb{E}[L_n])$ to conclude the proof. For any real $\alpha > 0$, there exists an integer ℓ_0 such that, for every $\ell \geq \ell_0$, $\epsilon_\ell \leq \frac{\alpha}{2}$. Hence, with $\eta_0 = \max \{ \epsilon_\ell : \ell \in \{0, \dots, \ell_0 - 1\} \}$,

$$\begin{aligned} \sum_{\ell \geq 0} \mathbb{P}(L_n = \ell) \ell \epsilon_\ell &= \sum_{\ell=0}^{\ell_0-1} \mathbb{P}(L_n = \ell) \ell \epsilon_\ell + \sum_{\ell \geq \ell_0} \mathbb{P}(L_n = \ell) \ell \epsilon_\ell \\ &\leq \sum_{\ell=0}^{\ell_0-1} \mathbb{P}(L_n = \ell) \ell \epsilon_\ell + \frac{\alpha}{2} \mathbb{E}[L_n] \\ &\leq \eta_0 (\ell_0 - 1) + \frac{\alpha}{2} \mathbb{E}[L_n]. \end{aligned}$$

Since $\mathbb{E}[L_n]$ tends to infinity, for n sufficiently large we have $\frac{\eta_0(\ell_0-1)}{\mathbb{E}[L_n]} \leq \frac{\alpha}{2}$. Therefore, for any $\alpha > 0$, for n sufficiently large we have $\mathbb{E}[L_n \epsilon_{L_n}] \leq \alpha \mathbb{E}[L_n]$. Hence, $\mathbb{E}[L_n \epsilon_{L_n}] = o(\mathbb{E}[L_n])$, concluding the proof. \blacktriangleleft

A.2.3 Proof of Theorem 3

Theorem 3. *Assume that n is taken uniformly at random in $\{0, \dots, N-1\}$. The expected number of conditional tests in CLASSICALPOW and UNROLLEDPOW is asymptotically equivalent to $\log_2 N$, whereas it is asymptotically equivalent to $\frac{5}{4} \log_2 N$ for GUIDEDPOW. The expected number of mispredictions is asymptotically equivalent to $\frac{1}{2} \log_2 N$ for CLASSICALPOW and UNROLLEDPOW, for any kind of predictor. For GUIDEDPOW, it is asymptotically equivalent to $\alpha \log_2 N$, where $\alpha = \frac{1}{2}\mu(3/4) + \frac{3}{4}\mu(2/3)$.*

Proof. The proof is done for $N = 4^k$. Some care is required for other values of N , since the bits are not 0's and 1's with probability exactly $\frac{1}{2}$, and since they are not completely independent. We explain in Section 5 how to deal with this approximations rigorously: it can also be done here in a similar way.

It is straightforward to prove, as we did for UNROLLEDPOW, that the expected number of iterations of CLASSICALPOW is $\log_2 N + \mathcal{O}(1)$. As each iteration performs one conditional test which is mispredicted with probability $\frac{1}{2}$, we get the announced results for this algorithm.

We already saw that the expected number of iterations of UNROLLEDPOW and GUIDEDPOW is asymptotically equivalent to $\log_4 N = \frac{1}{2} \log_2 N$. At each iteration, the first conditional is performed, and the other two are performed with probability $\frac{3}{4}$. This yields that the expected number of conditional tests is asymptotically equivalent to $\frac{5}{4} \log_2 N$.

In UNROLLEDPOW, each conditional is mispredicted with probability $\frac{1}{2}$ yielding the announced result.

For GUIDEDPOW we use the Ergodic Theorem and the fact that $\mathbb{E}[L_k] \sim k \sim \log_4 N$. The first conditional causes an expected number of mispredictions asymptotically equivalent to $\mu(3/4) \log_4 N$ and each of the two nested conditionals causes $\frac{3}{4}\mu(2/3) \log_4 N$ mispredictions, since they are tested with probability $\frac{3}{4}$. Hence, the expected number of mispredictions is asymptotically equivalent to $(\mu(3/4) + \frac{3}{2}\mu(2/3)) \log_4 N$, concluding the proof. \blacktriangleleft

A.2.4 Proof of Lemma 5

Lemma 5. *The probability that $\mathbf{Path}_n(\mathcal{T}, u)$ and $\mathbf{Path}_n(\mathcal{I}, u)$ differ at one of the first $L_n(\mathcal{T}, u) - \sqrt{\log n}$ steps is $\mathcal{O}(\frac{1}{\log n})$.*

Proof. Let (d, f) be a node of \mathcal{T}_n that is not a leaf, and let $t = f - d + 1$. The probability to go to the left child at the next step is $p(d, f) = \frac{m-d+1}{t}$ where $m = \lfloor \frac{d+f}{4} \rfloor$, which is $\frac{1}{4} + \mathcal{O}(\frac{1}{t})$. Hence, there exists a positive α such that $|p(d, f) - \frac{1}{4}| \leq \frac{\alpha}{2t}$. And therefore, the probability that the two paths differ at this step is at most $\frac{\alpha}{t}$.

Let (d', f') be the node reached after $L_n(u) - \lambda_n$ steps. For n sufficiently large, the range of a node is multiplied by a number that lies in $[\frac{1}{5}, \frac{4}{5}]$ (as $\frac{1}{5} < \frac{1}{4}$ and $\frac{4}{5} > \frac{3}{4}$). Hence, we have

$$f' - d' + 1 \geq \left(\frac{5}{4}\right)^{\lambda_n}.$$

As a consequence, all the nodes at distance at most $L_n(u) - \lambda_n$ from the root have a range that is greater than or equal to γ_n , with $\gamma_n = \left(\frac{5}{4}\right)^{\lambda_n}$. Moreover, $L_n(u) \leq \log_{5/4} n$ for n

sufficiently large. Thus, the probability that the $L_n(u) - \lambda_n$ first steps of $\mathbf{Path}_n(\mathcal{T}, u)$ and $\mathbf{Path}_n(\mathcal{I}, u)$ are equal is at least $\left(1 - \frac{\alpha}{\gamma_n}\right)^{\log_{5/4} n - \lambda_n}$. Thus:

$$\begin{aligned} \left(1 - \frac{\alpha}{\gamma_n}\right)^{\log_{5/4} n - \lambda_n} &= \exp\left(\left(\log_{5/4} n - \lambda_n\right) \log\left(1 - \frac{\alpha}{\gamma_n}\right)\right) \\ &= \exp\left(-\frac{\alpha}{\gamma_n} \left(\log_{5/4} n - \lambda_n\right) (1 + o(1))\right) \\ &= 1 - \frac{\alpha \log_{5/4} n}{\gamma_n} + o\left(\frac{\log n}{\gamma_n}\right) = 1 - o\left(\frac{1}{\log n}\right), \end{aligned}$$

concluding the proof. \blacktriangleleft

A.2.5 Proof of Theorem 6

Theorem 6. *Let C_n and M_n be the number of comparisons and mispredictions performed in our model of randomness. For BINARYSEARCH, $\mathbb{E}[C_n] \sim \frac{1}{\log 2} \log n$ and $\mathbb{E}[M_n] \sim \frac{1}{2 \log 2} \log n$. For BIASEDBINARYSEARCH, $\mathbb{E}[C_n] \sim \frac{4}{4 \log 4 - 3 \log 3} \log n$ and $\mathbb{E}[M_n] \sim \mu(\frac{1}{4}) \mathbb{E}[C_n]$. For SKEWSEARCH, $\mathbb{E}[C_n] \sim \frac{7}{6 \log 2} \log n$ and $\mathbb{E}[M_n] \sim \left(\frac{4}{7} \mu(\frac{1}{4}) + \frac{3}{7} \mu(\frac{1}{3})\right) \mathbb{E}[C_n]$.*

Proof. Let $M_n(u)$ be the number of mispredictions produced by the algorithm for u , and let $I_n(u)$ be the number of mispredictions produced by the idealized algorithm when following the $L_n(u)$ first steps in the ideal tree \mathcal{I} . Let \mathcal{G}_n be the set of $u \in [0, 1]^\omega$ such that $\mathbf{Path}_n(\mathcal{T}, u)$ and $\mathbf{Path}_n(\mathcal{I}, u)$ are equal for their first $L_n(u) - \lambda_n$ steps. Observe that if $u \in \mathcal{G}_n$, then we have $|M_n(u) - I_n(u)| \leq \lambda_n$ since they can only differ at the λ_n last steps of the algorithms. If $u \notin \mathcal{G}_n$, we simply use the fact that there are less than $L_n(u)$ mispredictions, which is smaller than $\beta \log n$ for some well-chosen constant β . Using Lemma 5 we have

$$\left|\mathbb{E}[M_n] - \mathbb{E}[I_n]\right| \leq \lambda_n \cdot \left(1 - o\left(\frac{1}{\log n}\right)\right) + \beta \log n \cdot o\left(\frac{1}{\log n}\right) = o(\log n).$$

Hence $\mathbb{E}[M_n] \sim \mathbb{E}[I_n]$, and the first order asymptotic results obtained in the idealized model still hold for the real case.

To obtain the result for SKEWSEARCH, observe that if X_n denotes the number of times the first conditional is executed for an input of length n , then:

$$\mathbb{E}[X_n] = 1 + \frac{a_n}{n+1} \mathbb{E}[X_{a_n}] + \frac{b_n}{n+1} \mathbb{E}[X_{b_n}] + \frac{c_n}{n+1} \mathbb{E}[X_{c_n}], \text{ with } \begin{cases} a_n = \lfloor \frac{n}{4} \rfloor + 1 \\ b_n = \lfloor \frac{n+1}{2} \rfloor - a_n \\ c_n = \lceil \frac{n+1}{2} \rceil \end{cases},$$

thus Theorem 4 applies, yielding the same result that $f(n) = 1 + \frac{1}{2}f(n/4) + \frac{1}{2}f(n/2)$ and

$$\mathbb{E}[X_n] \sim \frac{2}{3 \log 2} \log n.$$

Similarly, if Y_n denotes the number of times the second conditional is executed for an input of length n , then:

$$\mathbb{E}[Y_n] = \frac{a_n}{n+1} \mathbb{E}[X_{a_n}] + \frac{b_n}{n+1} (\mathbb{E}[X_{b_n}] + 1) + \frac{c_n}{n+1} (\mathbb{E}[X_{c_n}] + 1), \text{ with } \begin{cases} a_n = \lfloor \frac{n}{4} \rfloor + 1 \\ b_n = \lfloor \frac{n+1}{2} \rfloor - a_n \\ c_n = \lceil \frac{n+1}{2} \rceil \end{cases},$$

and thus Theorem 4 applies, and yields the same result that $f(n) = 3/4 + \frac{1}{2}f(n/4) + \frac{1}{2}f(n/2)$ and

$$\mathbb{E}[Y_n] \sim \frac{1}{2 \log 2} \log n.$$

The expected total number of comparisons is therefore

$$\mathbb{E}[C_n] = \mathbb{E}[X_n + Y_n] \sim \frac{7}{6 \log 2} \log n,$$

as announced.

By Theorem 2, the expected number of mispredictions is therefore asymptotically equivalent to

$$\mu \left(\frac{3}{4} \right) \mathbb{E}[X_n] + \mu \left(\frac{2}{3} \right) \mathbb{E}[Y_n] = \left(\frac{2}{3} \mu \left(\frac{3}{4} \right) + \frac{1}{2} \mu \left(\frac{2}{3} \right) \right) \frac{\log n}{\log 2} = \left(\frac{4}{7} \mu \left(\frac{3}{4} \right) + \frac{3}{7} \mu \left(\frac{2}{3} \right) \right) \mathbb{E}[C_n],$$

which is the announced result. \blacktriangleleft

A.2.6 Proof of Theorem 7

Theorem 7. *For the global predictor, the average number of mispredictions caused during SKEWSEARCH on an input of size n is asymptotically equivalent to $(\frac{12}{35} + \frac{1}{595 \cdot 2^\ell}) \mathbb{E}[C_n]$.*


Proof. Let $w = w_0 w_1 \dots w_{N-1} \in \{0, 1\}^N$ be the output of the two conditionals, where 1 stands for “taken” and 0 for “not taken”.

Let $h \neq 0^\ell$ be a history. Let $1 \leq \tau_1 < \tau_2 < \dots < \tau_m \leq N - 1$ be the times t such that $h_t = h$ (there is an occurrence of h in w that ends at position t). Of course, the τ_i 's and m depend on N and h , and are random variables for random inputs of the algorithm. By the Ergodic Theorem, there exists a constant $\alpha_h > 0$ such that $\mathbb{E}[m(h)] \sim \alpha_h N$, as N tends to infinity. Indeed, if π_{up} is the stationary vector of the final strongly connected component of \mathcal{M}_{up} , the big Markov chain that captures everything, then α_h is the sum of the $\pi_{\text{up}}(x)$, when x ranges over the states such that the current history is h .

Observe also that the entry h of the global table can only be updated at the times $\tau_i + 1$, that is, $T_t(h)$ is constant for $\tau_i + 1 \leq t \leq \tau_{i+1}$, with the convention that $\tau_{m+1} = N$.

Let $\mathcal{H}_{\text{main}}$ (resp. $\mathcal{H}_{\text{nested}}$) be the set of histories h such that if $h_t = h$ then $\text{IF}_t = \text{main}$ (resp. $\text{IF}_t = \text{nested}$). The sets $\mathcal{H}_{\text{main}}$ and $\mathcal{H}_{\text{nested}}$ are disjoint and contain all the histories of size ℓ but 0^ℓ . Let $h \in \mathcal{H}_{\text{main}}$. From what we have just established, snapshots of $T_t(h)$ at times $\tau_1 + 1, \tau_2 + 1, \dots, \tau_m + 1$ describe a random walk in the predictor of parameter $\frac{1}{4}$. Hence, the expected number of mispredictions caused by this history is asymptotically equivalent to $\mu(1/4)\alpha_h N$, by Theorem 2. Similarly, if $h \in \mathcal{H}_{\text{nested}}$, the expected number of mispredictions caused by this history is asymptotically equivalent to $\mu(1/3)\alpha_h N$.

Let us analyze the behavior of the predictor of the entry 0^ℓ . When the history is 0^ℓ , the current conditional can be either **main** or **nested**. Thus, we have to distinguish the two cases and consider the pairs (s, i) , where s is a state of the 2-bit predictor and i is either **main** or **nested**. Thanks to the specificity of our problem, knowing the starting pair (s, i) and whether the conditional is taken or not is enough to determine the next pair (s', i') reached the next time the history is 0^ℓ : s' is determined by the associated transition in the 2-bit predictor; if the conditional is taken then we have another occurrence of 0^ℓ immediately after and $i' \neq i$, otherwise the next time the history is 0^ℓ is immediately after a 1, then $i' = \text{main}$ as ℓ is even. This yields the automaton (and its associated Markov chain \mathcal{M}_0)



■ **Figure 12** The Markov chain corresponding to the entry 0^ℓ .

depicted on Figure 12. Note that for odd ℓ we obtain a different automaton, which can be studied in the same way; however, ℓ is always even in real computers.

The stationary vector π_0 of \mathcal{M}_0 can easily be computed. In particular, we get that the stationary probability of being in a **main**-state is $p_{\text{main}} = \frac{4}{7}$ and of being on a **nested**-state is $p_{\text{nested}} = \frac{3}{7}$ (this would be different for odd ℓ). Moreover, the expected misprediction probability of \mathcal{M}_0 is $\mu_0 = \frac{41}{119}$. The stationary probability of the history is 0^ℓ is $p_0 = (3/4)^{\ell/2}(1/3)^{\ell/2} = 2^{-\ell}$. Hence, by the Ergodic Theorem, the number G_n of mispredictions for the global predictor has an expectation that is asymptotically equivalent to

$$\mathbb{E}[G_n] \sim \left(\sum_{h \in \mathcal{H}_{\text{main}}} \alpha_h \mu_2(1/4) + \sum_{h \in \mathcal{H}_{\text{nested}}} \alpha_h \mu_2(1/3) + \frac{\mu_0}{2^\ell} \right) \cdot \mathbb{E}[C_n].$$

But, as the stationary probability of being in **main** is $\frac{4}{7}$, we have

$$\sum_{h \in \mathcal{H}_{\text{main}}} \alpha_h = \frac{4}{7} - p_{\text{main}} \cdot p_0 = \frac{4}{7}(1 - 2^{-\ell}).$$

Similarly,

$$\sum_{h \in \mathcal{H}_{\text{nested}}} \alpha_h = \frac{3}{7} - p_{\text{nested}} \cdot p_0 = \frac{3}{7}(1 - 2^{-\ell}).$$

As $\mu_2(1/4) = \frac{3}{10}$ and $\mu_2(1/3) = \frac{2}{5}$, we get the announced result:

$$\mathbb{E}[G_n] \sim \left(\frac{12}{35}(1 - 2^{-\ell}) + \frac{41}{119}2^{-\ell} \right) \mathbb{E}[C_n] = \left(\frac{12}{35} + \frac{1}{595 \cdot 2^\ell} \right) \mathbb{E}[C_n].$$

◀