Proceedings of the 20th Conference on Computer Science and Intelligence Systems (FedCSIS) pp. 795–800
DOI: 10.15439/2025F0998
ISSN 2300-5963 ACSIS, Vol. 43

A statistical hypothesis test for primality based on random divisor sampling: Principles, properties, adaptive design, and algorithmic analysis

Lubomír Štěpánek[1, 2, 3]
ORCiD: 0000-0002-8308-4304
[1]Department of Statistics and Probability
[2]Department of Mathematics
Faculty of Informatics and Statistics
Prague University of Economics and Business
W. Churchill's square 4, 13067 Prague, Czech Republic
lubomir.stepanek@vse.cz
&
[3]Institute of Biophysics and Informatics
First Faculty of Medicine
Charles University
Salmovská 1, 12000 Prague, Czech Republic
lubomir.stepanek@lf1.cuni.cz

***Abstract*—We propose a statistical hypothesis test for determining whether a given integer $n \geq 2$ is prime. Under the null hypothesis H_0, we assume that n is not a prime (i.e., composite). The test operates by randomly sampling integers from the candidate divisor set $\mathcal{D} = \{2, 3, \ldots, \lfloor\sqrt{n}\rfloor\}$ and checking whether any of them divide n. If a proper divisor is found, H_0 is not rejected and n is declared composite. If no divisor is found among an initial set of k_0 samples, additional k samples are drawn, and a p-value is computed based on the probability of missing all actual divisors under H_0. This probability is calculated exactly via a hypergeometric distribution and approximated using an exponential bound. We derive a closed-form upper bound for the minimal number of trials k required to reject H_0 at a given significance level α under the conservative assumption of only one true divisor ($m = 1$). The algorithm has worst-case time complexity $\mathcal{O}(\sqrt{n})$, matching that of classical trial division, but its expected runtime is substantially lower when n has multiple divisors. The proposed test is simple, statistically interpretable, and well-suited both as an educational tool and as a lightweight probabilistic pre-check in layered primality testing pipelines.**

I. Introduction

THE problem of determining whether a given integer n is a prime number is a fundamental question in number theory and theoretical computer science. Primality testing has widespread applications in modern cryptography (such as RSA encryption [1]), secure key generation, error-correcting codes, hashing schemes, and symbolic computation. In many such contexts, fast and reliable primality tests are crucial for both correctness and efficiency [2].

A straightforward approach to primality testing is to check whether any integer in the range $\{2, 3, \ldots, n-1\}$ divides n, or, more effectively, the range of possible divisors is firstly reduced to $\{2, 3, \ldots, \lfloor\sqrt{n}\rfloor\}$ and each one is tested whether divides n. If no such divisor is found, then n is declared prime. While conceptually simple, this method becomes computationally expensive for large n, as it requires up to $\mathcal{O}(\sqrt{n})$ divisibility tests in the worst case. This naive method does not scale well for applications involving large integers or repeated primality checks.

Several more advanced primality tests have been developed, many of which rely on algebraic or number-theoretic properties. Notable examples include the following.

- The *Fermat primality test*, based on Fermat's little theorem, assumes $H_0 : n$ is prime, and checks whether $a^{n-1} \equiv 1 \pmod{n}$ for a random base a [3]. Failure to satisfy this congruence implies n is composite, but passing it does not guarantee primality due to pseudoprimes and Carmichael numbers [4].
- The *Miller–Rabin test*, a widely used probabilistic test, strengthens the Fermat test by testing additional congruences derived from factorization of $n - 1$. It is efficient and has a controllable error probability [5].
- The *AKS primality test*, a deterministic polynomial-time algorithm, decides primality unconditionally. However, it is relatively complex and not competitive in practice for large inputs [6].

While these methods are powerful, most of them assume the null hypothesis that n is prime and look for algebraic contradictions [7]. In this paper, we propose a novel *statistical hypothesis test* for primality that takes a fundamentally different approach: we assume the null hypothesis H_0 that n is

not prime, and assess how "surprising" the observed outcomes of random divisor tests are under this assumption.

Specifically, we randomly sample integers from the set $\{2, 3, \ldots, \lfloor\sqrt{n}\rfloor\}$ and test whether they divide n. If a proper divisor is found, we do not reject H_0 and declare that n is composite. If no divisor is found, we compute a p-value based on the probability of missing all actual divisors under H_0, using both exact combinatorial formulas and exponential approximations. We also derive closed-form bounds for the minimum number of integers that must be tested as potential divisors to achieve a given significance level α, and we analyze the algorithmic complexity of this test in comparison to the classical exhaustive approach.

The proposed test is simple to implement, statistically interpretable, and often requires significantly fewer checks than exhaustive search – especially when n has many proper divisors. It can serve as a practical and educational tool, or as a component in probabilistic or layered primality testing frameworks.

II. Preliminaries

This section establishes basic mathematical facts used in the formulation and justification of the proposed test.

We begin by showing that it is sufficient to restrict our attention to potential divisors of n that are less than or equal to $\lfloor\sqrt{n}\rfloor$.

Lemma 1 (Sufficiency of checking divisors up to $\sqrt{n}$)**.** *Let $n \geq 2$ be an integer. Then n is composite if and only if there exists a proper divisor d such that $2 \leq d \leq \lfloor\sqrt{n}\rfloor$.*

Proof. Let $n \geq 2$ be arbitrary. Suppose first that n is composite. Then there exist integers d, q such that $n = d \cdot q$ with $2 \leq d < n$ and $2 \leq q < n$. Without loss of generality, assume $d \leq q$. Then it is $d \cdot d \leq d \cdot q$, but $d \cdot q = n$, so $d^2 \leq d \cdot q = n$, or simply $d^2 \leq n$. Consequently, it is $d \leq \sqrt{n}$, and since $d \in \mathbb{N}$, it is also $d \leq \lfloor\sqrt{n}\rfloor$. So, d is a proper divisor of n and satisfies $2 \leq d \leq \lfloor\sqrt{n}\rfloor$, the claim follows.

Conversely, suppose there exists a proper divisor d of n such that $2 \leq d \leq \lfloor\sqrt{n}\rfloor$. Then $d \mid n$, so n is not a prime. Hence, n must be composite. □

The following lemma formalizes the intuitive fact that, when sampling from a finite population with few successes, drawing without replacement decreases the chance of obtaining no successes at all compared to drawing with replacement.

Lemma 2 (Probability of no success is lower or equal without replacement)**.** *Let $N \geq 1$ and $1 \leq M < N$ be integers representing a population of N items, M of which are marked as successes. Consider drawing k items from this population with $1 \leq k \leq N - M$. Then the probability of drawing zero successes without replacement is less than or equal to the probability of drawing zero successes with replacement.*

Proof. The probability of no success in k draws *with replacement* is given by

$$P_{\text{with}}(0) = \left(1 - \frac{M}{N}\right)^k. \tag{1}$$

The probability of no success in k draws *without replacement* is

$$P_{\text{without}}(0) = \frac{\binom{N-M}{k}}{\binom{N}{k}} = \prod_{i=0}^{k-1} \left(1 - \frac{M}{N-i}\right). \tag{2}$$

Since for each i in $0 \leq i \leq k-1$ we have $N - i \leq N$, it follows that $\frac{M}{N-i} \geq \frac{M}{N}$, and thus

$$\left(1 - \frac{M}{N-i}\right) \leq \left(1 - \frac{M}{N}\right).$$

Consequently, the product of k such terms satisfies

$$\prod_{i=0}^{k-1} \left(1 - \frac{M}{N-i}\right) \leq \prod_{i=0}^{k-1} \left(1 - \frac{M}{N}\right)$$

$$\prod_{i=0}^{k-1} \left(1 - \frac{M}{N-i}\right) \leq \left(1 - \frac{M}{N}\right)^k$$

$$P_{\text{without}}(0) \overset{(1,2)}{\leq} P_{\text{with}}(0).$$

Hence, the probability of no success is lower than or equal to when drawing without replacement. □

III. Proposed methodology

We now introduce a statistical hypothesis testing procedure for primality, based on the idea of randomly sampling potential divisors of a given integer n. The method is built upon a probabilistic interpretation of missing actual divisors under the assumption that n is composite. In addition to defining the logic, hypotheses, and operational steps of the proposed test, we derive exact and approximate expressions for the resulting p-value and analyze the algorithm's time complexity in comparison to classical exhaustive methods.

A. Test logic and procedure

Let $\mathcal{D} = \{2, 3, \ldots, \lfloor\sqrt{n}\rfloor\}$ be the set of candidate proper divisors of $n \in \mathbb{N}$ with $n \geq 2$ large. Using Lemma 1, we know that investigating potential divisors not lower than or equal to $\lfloor\sqrt{n}\rfloor$ is sufficient to determine primality. Denote its cardinality by $|\mathcal{D}|$, where obviously

$$|\mathcal{D}| = |\{2, 3, \ldots, \lfloor\sqrt{n}\rfloor\}| = \lfloor\sqrt{n}\rfloor - 1. \tag{3}$$

If n is composite, then it must have at least one divisor in $\mathcal{D}$. This observation motivates a statistical hypothesis test based on sampling elements from $\mathcal{D}$ and checking whether they divide n.

We define the null and alternative hypotheses as follows.

H_0 : n is not a prime, i.e., there exists at least one $d \in \mathcal{D}$ such that $d \mid n$.

H_1 : n is a prime, i.e., no element of $\mathcal{D}$ divides n, or also $\forall d \in \mathcal{D} : d \nmid n$.

The test consists of the following steps.

(i) We randomly select $k_0 \geq 0$ integers from the set $\mathcal{D}$ without replacement, and test each to see whether it divides n.
- If any sampled value divides n, we do not reject H_0 and declare that n is composite.
- If no divisor is found, we proceed to the next step.

(ii) We randomly select an additional $k \geq 1$ integers from the remaining elements of $\mathcal{D}$, again without replacement, and test them for divisibility.
- If any sampled value divides n, we do not reject H_0 and declare that n is composite.
- If no divisor is found, we compute the p-value under H_0.

(iii) We define the p-value as the probability that none of the k additional randomly selected integers divide n, under the assumption that n is composite and has $m \geq 1$ actual divisors in $\mathcal{D}$. Although the exact value of m is unknown in practice – if it were known, we could straightforwardly decide whether n is prime – we can still compute valid upper bounds on this probability. The smaller the p-value, the more extreme the observed outcome (i.e., absence of divisors) appears under H_0, and the stronger the evidence against the assumption that n is composite. If the p-value is less than or equal to a given significance level α, we reject H_0 and declare that n is likely a prime.

This method enables a sequential, interpretable, and probabilistically grounded test for primality. The exact and approximate computation of the p-value, as well as guidance for selecting k given a desired confidence level, are discussed in the following subsections.

B. Exact derivation of the p-value

We now derive the exact p-value used in the proposed hypothesis test. This value represents the probability of observing no divisors among a randomly selected subset of potential divisors, under the assumption that n is composite.

Let $\mathcal{D} = \{2, 3, \ldots, \lfloor\sqrt{n}\rfloor\}$ be the set of candidate proper divisors of n. Suppose – under H_0 – that n is composite and has exactly $m \geq 1$ proper divisors in $\mathcal{D}$. We assume m is fixed but unknown.

After testing an initial set of $k_0 \ll |\mathcal{D}|$ integers from $\mathcal{D}$ with no observed divisor, we draw an additional k values uniformly at random without replacement from the remaining $|\mathcal{D}| - k_0$ elements. We are interested in the probability that none of these k values divide n, conditioned on the assumption that n is composite and has m proper divisors in $\mathcal{D}$.

Let X denote the number of divisors observed in the additional k draws. Under H_0 and fixed m, the distribution of X follows the hypergeometric distribution,

$$X \sim \text{Hypergeometric}(|\mathcal{D}| - k_0, m, k).$$

Then, the p-value is defined as the probability of observing $X = 0$ under H_0,

$$p\text{-value} = \mathbb{P}(X = 0 \mid H_0) = \frac{\binom{|\mathcal{D}|-m-k_0}{k}}{\binom{|\mathcal{D}|-k_0}{k}}, \tag{4}$$

or also

$$p\text{-value} = \frac{\binom{|\mathcal{D}|-m-k_0}{k}}{\binom{|\mathcal{D}|-k_0}{k}} = \prod_{i=0}^{m-1} \frac{|\mathcal{D}| - k_0 - k - i}{|\mathcal{D}| - k_0 - i}. \tag{5}$$

This expression quantifies the likelihood of missing all true divisors in the k additional samples, under the null hypothesis that n is not a prime. Inspecting formulas (4) and particularly (5), we can see that as k increases or m increases, the p-value decreases, providing stronger evidence against H_0. Thus, the exact p-value decreases as the number of actual divisors m increases. That is, the more true divisors exist in the candidate set $\mathcal{D}$, the less likely it is that a random sample of size k misses all of them. Formally,

$$p\text{-value} = \frac{\binom{|\mathcal{D}|-m-k_0}{k}}{\binom{|\mathcal{D}|-k_0}{k}} \leq \frac{\binom{|\mathcal{D}|-1-k_0}{k}}{\binom{|\mathcal{D}|-k_0}{k}},$$

or also

$$p\text{-value} \leq \frac{\binom{|\mathcal{D}|-1-k_0}{k}}{\binom{|\mathcal{D}|-k_0}{k}} \overset{(5)}{=} 1 - \frac{k}{|\mathcal{D}| - k_0} \overset{(3)}{=} 1 - \frac{k}{\lfloor\sqrt{n}\rfloor - 1 - k_0}, \tag{6}$$

for all $m \geq 1$, where $m \leq \lfloor\sqrt{n}\rfloor - 1 - k_0 - k$. Thus, the case $m = 1$ corresponds to the worst-case scenario (i.e., highest possible p-value under H_0), and can be used as a conservative upper bound. This allows the test to remain valid even when the number of actual divisors is unknown. Also, controlling the p-value under $m = 1$ yields a conservative test.

C. Exponential approximation and upper bound of the p-value

While the exact p-value expression in (4) is combinatorially precise, it can be computationally demanding to evaluate, especially when one wishes to determine the minimal value of k for which the p-value falls below a given significance level α, particularly under the worst-case assumption $m = 1$.

To enable more practical reasoning and simplify the estimation of the required sample size k, we now derive an upper bound on the p-value using an exponential approximation.

Let's assume $k \ll |\mathcal{D}| - m - k_0$, which is natural since k is typically chosen much smaller than $|\mathcal{D}| = |\{2, 3, \ldots, \lfloor\sqrt{n}\rfloor\}| = \lfloor\sqrt{n}\rfloor - 1$. Under this assumption, sampling without replacement is numerically close to sampling with replacement. In particular, using Lemma 2, the probability of missing all m true divisors when drawing without replacement is slightly lower than when sampling with replacement, and hence,

$$p\text{-value} = \frac{\binom{|\mathcal{D}|-m-k_0}{k}}{\binom{|\mathcal{D}|-k_0}{k}} \leq \left(1 - \frac{m}{|\mathcal{D}| - k_0}\right)^k. \tag{7}$$

This bound follows from the fact that each draw without replacement strictly reduces the remaining pool size, increasing the probability of hitting a true divisor. On the other hand, when drawing with replacement, the chance of selecting a true divisor remains constant across draws. As a result, sampling with replacement underestimates the chance of detection, making the bound conservative.

D. Minimum number k of potential divisors to satisfy a given significance level

Having derived an exponential upper bound on the p-value in (7), we now invert this inequality to estimate the minimum number k of integers that need to be checked as potential divisors of n to ensure that the p-value falls below a desired significance level α, where $0 < \alpha < 1$.

Combining formulas (4) and (7) and comparing the result to the significance level α, we obtain

$$p\text{-value} = \frac{\binom{|\mathcal{D}|-m-k_0}{k}}{\binom{|\mathcal{D}|-k_0}{k}} \leq \left(1 - \frac{m}{|\mathcal{D}| - k_0}\right)^k \leq \alpha,$$

and focus on the rightmost inequality to solve for k. Taking logarithms on both sides yields

$$k \cdot \log\left(1 - \frac{m}{|\mathcal{D}| - k_0}\right) \leq \log(\alpha). \tag{8}$$

Since $\log(1-x) < 0$ for $x \in (0,1)$, inequality (8) leads to the following lower bound

$$k \geq \frac{\log(\alpha)}{\log\left(1 - \frac{m}{|\mathcal{D}|-k_0}\right)},$$

or also

$$k \overset{(3)}{\geq} \frac{\log(\alpha)}{\log\left(1 - \frac{m}{\lfloor\sqrt{n}\rfloor-1-k_0}\right)}. \tag{9}$$

This expression provides a closed-form estimate of the number of random checks needed after the initial k_0 samples in order to reject H_0 at significance level α. The most conservative bound arises under the worst-case scenario $m = 1$ (i.e., n has exactly one proper divisor in $\mathcal{D}$). Substituting $m = 1$ into (9) gives

$$k \geq \frac{\log(\alpha)}{\log\left(1 - \frac{1}{|\mathcal{D}|-k_0}\right)},$$

or also

$$k \overset{(3)}{\geq} \frac{\log(\alpha)}{\log\left(1 - \frac{1}{\lfloor\sqrt{n}\rfloor-1-k_0}\right)}. \tag{10}$$

This value of k guarantees that the p-value under H_0 does not exceed α for any possible number of actual divisors $m \geq 1$, thus ensuring a valid and conservative decision threshold even when m is unknown.

IV. Practical usage of the proposed test

The proposed statistical hypothesis test for primality can be applied in two main modes: a classical version with a precomputed sample size k, and an adaptive version that incrementally samples potential divisors until a statistically justified conclusion can be reached. We describe both approaches in this section.

A. Fixed-sample usage with conservative decision threshold

A straightforward way to use the test is to choose a fixed sample size k of candidate divisors in advance, based on a desired significance level α. This guarantees that the computed p-value will fall below α unless a divisor is found, as long as k satisfies the inequality derived in formula (10) under the worst-case assumption that n has only one proper divisor in $\mathcal{D}$ (i.e., when $m = 1$). The steps for such usage are as follows.

(i) The user first selects $k_0 \geq 0$ initial values to test.
(ii) Then, they choose k based on the exponential bound from formula (10) to ensure that the p-value will be sufficiently low if no proper divisor of n is found.
(iii) If any of the tested values divide n, the number is declared composite, since H_0 cannot be rejected.
(iv) If no divisor is found and the computed p-value is below α, then H_0 is rejected, and n is declared probably prime (at the given statistical significance level α).
(v) Otherwise (i.e., if k was chosen too small), the test may be inconclusive: no divisors of n are found in $\mathcal{D}$ (so n may be a prime), but p-value $> \alpha$, and thus H_0 (claiming that n is composite) cannot be rejected.

This fixed-sample approach is simple and deterministic, but may be inefficient if k is much larger than necessary, or inconclusive if k is too small.

B. Adaptive usage with incremental k sampling

An alternative approach is to use the test adaptively: incrementally draw one candidate divisor at a time, and terminate as soon as either a divisor is found (in which case H_0 cannot be rejected), or the p-value's upper bound, calculated using formula (6), drops below α (in which case H_0 is rejected). This version avoids the need to specify k in advance and can potentially stop earlier, depending on the observed evidence.

We present this version formally in Algorithm 1, which assumes a conservative worst-case bound with $m = 1$ and computes the p-value after each additional draw.

This variant can be seen as a probabilistic filter with automatic stopping conditions and is particularly useful when computational efficiency is critical or when n is large and sparsely divisible.

V. Algorithmic Complexity Analysis

We now analyze the computational complexity of the proposed statistical primality test. The analysis is based on the number of basic divisibility checks required to reach a statistically justified conclusion, assuming that a single test of the form $d \mid n$ takes constant time.

Algorithm 1: Adaptive statistical primality test via incremental k sampling

Input : Integer $n \geq 2$; significance level α; initial batch size $k_0 \geq 0$
Output: Decision: "n is `composite`" or "n is `probably prime` (at statistical significance α)"

Let $\mathcal{D} \leftarrow \{2, 3, \ldots, \lfloor\sqrt{n}\rfloor\}$
Randomly sample k_0 integers from $\mathcal{D}$ without replacement
foreach d *in the* k_0*-sample* **do**
  if $d \mid n$ **then**
    return do not reject $H_0 \implies n$ is `composite`
Let $\mathcal{D}' \leftarrow \mathcal{D} \setminus \{\text{already tested values}\}$
Set $k \leftarrow 0$ `// counter of post-initial checks`
while $\mathcal{D}'$ *is not empty* **do**
  Randomly select one $d \in \mathcal{D}'$ and remove it from $\mathcal{D}'$
  if $d \mid n$ **then**
    return do not reject $H_0 \implies n$ is `composite`
  $k \leftarrow k + 1$
  Compute p-value's upper bound:
  $$p\text{-value}_{\text{upper}} \overset{(6)}{\leftarrow} 1 - \frac{k}{\lfloor\sqrt{n}\rfloor - 1 - k_0}$$
  if $p\text{-value}_{upper} \leq \alpha$ **then**
    return reject $H_0 \implies n$ is `probably prime` (at statistical significance α)

A. Search space size

Let $\mathcal{D} = \{2, 3, \ldots, \lfloor\sqrt{n}\rfloor\}$ be the set of candidate proper divisors. Its cardinality satisfies

$$|\mathcal{D}| = \lfloor\sqrt{n}\rfloor - 1 = \Theta(\sqrt{n}). \tag{11}$$

A classical exhaustive trial division test checks all values in $\mathcal{D}$, resulting in a worst-case complexity of $\mathcal{O}(\sqrt{n})$.

B. Worst-case bound for the proposed test

Unlike exhaustive search, the proposed statistical test samples values randomly from $\mathcal{D}$ and stops either when a divisor is found or when the computed p-value drops below the significance threshold α.

From formula (9), the minimum number of values k to sample (after initial k_0 values) that guarantees p-value $\leq \alpha$ is

$$k_{\min} = \left\lceil \frac{\log(\alpha)}{\log\left(1 - \frac{m}{|\mathcal{D}|-k_0}\right)} \right\rceil, \tag{12}$$

where m is the number of proper divisors in $\mathcal{D}$. In the worst-case, $m = 1$, yielding

$$k_{\min} = \left\lceil \frac{\log(\alpha)}{\log\left(1 - \frac{1}{|\mathcal{D}|-k_0}\right)} \right\rceil. \tag{13}$$

Using the standard approximation $\log(1-x) \approx -x$ for small $x > 0$ coming from Taylor series, we derive the asymptotic behavior as

$$\begin{aligned}
\Theta(k_{\min}) &= \Theta\left(\left\lceil \frac{\log(\alpha)}{\log\left(1 - \frac{1}{|\mathcal{D}|-k_0}\right)} \right\rceil\right) \overset{(\dagger)}{\approx} \\
&\overset{(\dagger)}{\approx} \Theta\left(\left\lceil \frac{\log(\alpha)}{-\frac{1}{|\mathcal{D}|-k_0}} \right\rceil\right) \approx \\
&\approx \Theta\left(\lceil -\log(\alpha)(|\mathcal{D}| - k_0)\rceil\right) \approx \\
&\approx \Theta\left(\left\lceil \log\left(\frac{1}{\alpha}\right)(|\mathcal{D}| - k_0) \right\rceil\right) \overset{(11)}{\approx} \\
&\overset{(11)}{\approx} \Theta\left(\left\lceil \log\left(\frac{1}{\alpha}\right)\sqrt{n} \right\rceil\right) \approx \\
&\approx \Theta\left(\log\left(\frac{1}{\alpha}\right)\sqrt{n}\right) \approx \\
&\approx \Theta\left(\sqrt{n}\right).
\end{aligned} \tag{14}$$

Thus, the number of divisor checks required to make a valid decision at level α is bounded by $\mathcal{O}\left(\sqrt{n}\right)$.

C. Expected number of checks

In practice, the number of checks is usually smaller due to the following two stopping mechanisms,

(i) a divisor is found early, particularly when m is large;
(ii) no divisor is found, but the p-value falls below α.

Assuming uniform sampling without replacement from $\mathcal{D}$ and a uniform distribution of the m divisors within $\mathcal{D}$, there is, on average, one divisor in every $\frac{1}{m}$-fraction of the set. Therefore, we expect to check approximately $\frac{|\mathcal{D}|}{m} = \frac{\sqrt{n}-1}{m}$ potential divisors. Thus, the expected number of trials until the first divisor is observed is approximately

$$\mathbb{E}[T_{\text{hit}}] = \frac{|\mathcal{D}|}{m} = \frac{\sqrt{n}-1}{m} \approx \frac{\sqrt{n}}{m}. \tag{15}$$

Alternatively, if no divisor is found, the test terminates when $k = k_{\min}$ as given in (12). Therefore, the number of checks is bounded by

$$T(n, m) \overset{(11,12,15)}{\approx} \mathcal{O}\left(\min\left(\frac{\sqrt{n}}{m}, \frac{\log(\alpha)}{\log\left(1 - \frac{m}{\sqrt{n}}\right)}\right)\right).$$

D. Summary

The test adapts to the actual number of divisors: it runs faster when n has many proper divisors. In the extreme case $m = \Theta(\sqrt{n})$, the expected number of checks becomes constant. The worst-case complexity of $\mathcal{O}(\log(1/\alpha) \cdot \sqrt{n})$ arises only when $m = 1$. This adaptivity makes the proposed method significantly more efficient in practice compared to exhaustive search, which always requires $\Theta(\sqrt{n})$ steps.

VI. Illustration of p-value behavior for varying number of divisors

To better understand the behavior of the p-value in the proposed test, we analyze how it changes as a function of the number of random trials k for different values of m, the number of actual proper divisors in the candidate set $\mathcal{D} = 2, 3, \ldots, \lfloor\sqrt{n}\rfloor$. Fig. 1 illustrates this relationship for $n = 10007$, where $|\mathcal{D}| = \lfloor\sqrt{10007}\rfloor - 1 = 99$.

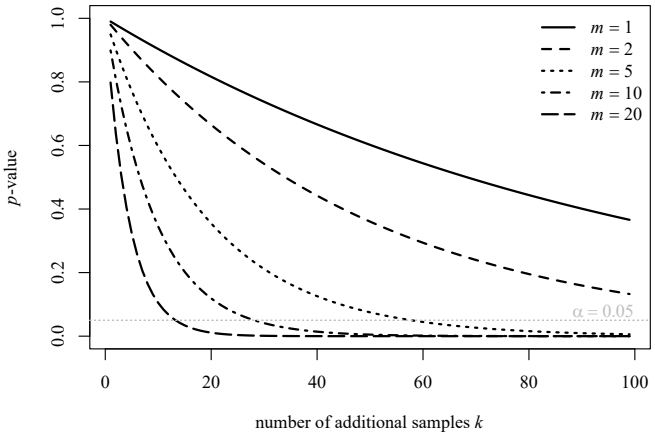

Fig. 1. Decay of the p-value with increasing number k of additionally sampled candidate divisors from $\mathcal{D}$ for various numbers m of true divisors under H_0. The horizontal dotted line marks the significance level $\alpha = 0.05$. Each curve shows the upper bound $\left(1 - \frac{m}{|\mathcal{D}|-k_0}\right)^k$ on the p-value, highlighting how the test becomes more decisive as either m or k increases.

Table I shows corresponding values of the minimum number $k_{\min}$ of samples required to reduce the p-value below the significance threshold $\alpha = 0.05$, based on the exponential bound $\left(1 - \frac{m}{|\mathcal{D}|-k_0}\right)^k$ from formula (7). This bound is conservative, especially for small values of m, and in some cases leads to estimates $k_{\min}$ exceeding the size of the candidate set, i.e., $k_{\min} > \lfloor\sqrt{n}\rfloor - 1$. This inefficiency illustrates that the bound is useful for providing safe, but not necessarily tight, thresholds.

TABLE I
Minimum number of random divisor checks $k_{\min}$ to reject H_0 at significance level $\alpha = 0.05$, for various numbers of actual divisors m, using the exponential bound with $n = 10007$, where $|\mathcal{D}| = \lfloor\sqrt{10007}\rfloor - 1 = 99$.

m	$k_{\min}$ such that p-value ≤ 0.05	is $k_{\min} \leq \lfloor\sqrt{n}\rfloor - 1$?
1	297	no
2	148	no
5	59	yes
10	30	yes
20	15	yes
40	8	yes
75	4	yes

Obviously, the p-value decreases exponentially with k, and this decay is faster when the number of divisors m increases. This demonstrates the adaptivity of the proposed method: when a composite number has more divisors, fewer random checks are typically needed to either discover one or to obtain a statistically significant p-value.

VII. Conclusion

We have proposed and rigorously analyzed a statistical hypothesis test for assessing the primality of a given integer n. The test is based on randomly sampling potential divisors from the set $\mathcal{D} = \{2, 3, \ldots, \lfloor\sqrt{n}\rfloor\}$ and computing a p-value under the null hypothesis H_0 that n is composite. If no divisor is found and the p-value falls below a given significance level α, the test rejects H_0, suggesting that n is likely a prime.

We derived both exact and exponential approximations of the p-value and provided a closed-form expression for the minimum number of additional checks k needed to guarantee a statistically valid conclusion. The exponential upper bound for the p-value decreases as either the number of actual divisors m or the number of random trials k increases. The test allows for an optional initial sample of $k_0 \geq 0$ values, which can be used for early termination or warming up the process, and is fully accounted for in the p-value computations and bounds. Moreover, we analyzed the worst-case and expected algorithmic complexity of the test, showing that the method is adaptive: it performs significantly better when the tested number has more divisors.

Compared to classical exhaustive methods, the proposed approach offers some practical advantages: it is statistically interpretable, simple to implement, and can terminate early in favorable cases. While the worst-case complexity remains $\mathcal{O}(\sqrt{n})$, the average-case performance is often much better, especially for composite numbers.

This paper provides a theoretical introduction to the proposed statistical primality test. Practical evaluations, methodological improvements, simulations, and comparisons with existing methods are part of ongoing work and will be addressed in future research.

VIII. Acknowledgment

This paper is supported by the grant IG410035 with no. F4/51/2025, which has been provided by the Internal Grant Agency of the Prague University of Economics and Business.

References

[1] R. L. Rivest, A. Shamir, and L. Adleman, "A method for obtaining digital signatures and public-key cryptosystems," *Communications of the ACM*, vol. 21, no. 2, pp. 120–126, 1978. doi: 10.1145/359340.359342

[2] R. Crandall and C. Pomerance, *Prime Numbers: A Computational Perspective*, 2nd ed. Springer, 2005. ISBN 9780387252820

[3] G. L. Miller, "Riemann's hypothesis and tests for primality," *Journal of Computer and System Sciences*, vol. 13, no. 3, pp. 300–317, 1976. doi: 10.1016/S0022-0000(76)80043-8

[4] K. H. Rosen, *Elementary Number Theory and Its Applications*, 6th ed. Addison-Wesley, 2011. ISBN 9780321500311

[5] M. O. Rabin, "Probabilistic algorithm for testing primality," *Journal of Number Theory*, vol. 12, no. 1, pp. 128–138, 1980. doi: 10.1016/0022-314X(80)90084-0

[6] M. Agrawal, N. Kayal, and N. Saxena, "Primes is in p," *Annals of Mathematics*, vol. 160, no. 2, pp. 781–793, 2004. doi: 10.4007/annals.2004.160.781

[7] G. Casella and R. L. Berger, *Statistical Inference*, 2nd ed. Duxbury Press, 2002. ISBN 9780534243128