The Generation of Random Numbers that are Probably Prime

Pierre BEAUCHEMIN, Gilles BRASSARD, Claude CREPEAU, Claude GOUTIER *and* Carl POMERANCE

ABSTRACT

In this paper we make two observations on Rabin's probabilistic primality test. The first is a provocative reason why Rabin's test is so good. It turns out that a single iteration has a non-negligible probability of failing only *on composite numbers that can actually be split in expected polynomial time. Therefore, factoring would be easy if Rabin's test systematically failed with a 25% probability on each composite integer (which, of course, it does not). The second observation is more fundamental because is it* not *restricted to primality testing: it has consequences for the entire field of probabilistic algorithms. The failure probability when using a probabilistic algorithm for the purpose of testing some property is compared to that when using it for the purpose of obtaining a random element hopefully having this property. More specifically, we investigate the question of how reliable Rabin's test is when used to* generate *a random integer that is probably prime, rather than to* test *a specific integer for primality.*

Key words: factorization, false witnesses, primality testing, probabilistic algorithms,
Rabin's test.

1. A BRIEF SURVEY OF PRIMALITY TESTING

How difficult is it to distinguish prime numbers from composite numbers? This is perhaps the single most important problem in computational number theory. We do not attempt here an exhaustive review of its long history. Let us only mention some of the most outstanding modern steps. It has been known for several years that the problem of recognizing prime numbers belongs to **P** under the Extended Riemann's Hypothesis [15] and that it belongs to Co-**RP** [21, 23] and **NP** [19] without any assumptions. It can also be solved in *almost* polynomial time by a deterministic algorithm that runs for a number of steps in $O(m^{O(\log\log m)})$, where m is the size of the number to be tested [2, 17, 7]. More recently, it was found to lie in **RP** [11, 1], and therefore in **ZPP** [10] as well. In other words, this problem can be solved in probabilistic polynomial time by a Las Vegas [3] algorithm: whenever an answer is obtained, that answer is correct.

From a theoretical point of view, the problem of primality testing is therefore essentially solved: the only remaining question is to figure out whether or not it belongs to **P** without assumptions. However, the polynomial that gives the running time of [11] is of the twelfth degree and [1] does not improve on this, which makes these algorithms of little practical use. For very large numbers (several hundreds of decimal digits), this leaves us with Rabin's probabilistic test [21] as the best approach. Let prob[$Rabin(n) = verdict$] denote the probability that one iteration of this algorithm on input n returns *verdict*, where *verdict* can either be ''*prime*'' or ''*composite*''. The basic theorem about Rabin's test is that

$$\text{prob}[\, Rabin(n) = \text{``}prime\text{''} \mid n \text{ is indeed prime} \,] = 1$$

whereas

$$\text{prob}[\, Rabin(n) = \text{``}prime\text{''} \mid n \text{ is in fact composite} \,] \le \tfrac{1}{4} \ .$$

One is therefore certain that n is composite whenever any single run of $Rabin(n)$ returns ''*composite*''. On the other hand, one can never be sure that n is prime no matter how many runs of $Rabin(n)$ have returned ''*prime*''. This test is usually run in a loop as follows:

```
function RepeatRabin(n, k)
  { n is an odd integer to be tested for primality;
    k is a safety parameter discussed below }
  var i : integer; done : Boolean
  i ← 0
  repeat
    i ← i + 1
    done ← (Rabin(n) = ''composite'')
  until done  or  i ≥ k
  if done  then return ''composite'' { for sure }
           else return ''prime'' { probably (?) } .
```

There is a trade-off in the choice of the parameter k above: the bigger it is, the more confident we are in the advent of a ''*prime*'' answer but the more time it takes to build up this confidence. This paper addresses two aspects of the question: *just how confident in a number's primality can we be after running this test?*

2. Why is Rabin's Test so Good?

Let n be an odd integer, $n > 1$, and let v and u be integers such that $n-1 = 2^v u$, where u is odd. Define

$$\mathbb{Z}_n^* = \{ a \mid 1 \le a < n \text{ and } \gcd(a, n) = 1 \} \quad \text{and}$$
$$\mathrm{R}_n = \{ a \in \mathbb{Z}_n^* \mid a^u \equiv 1 \pmod{n} \ \text{ or } \ (\exists j)[0 \le j < v \text{ and } a^{2^j u} \equiv -1 \pmod{n}] \} \ .$$

The basic theorem about Rabin's test states that $\mathrm{R}_n = \mathbb{Z}_n^*$ whenever n is a prime, whereas $\#\mathrm{R}_n \leq \#\mathbb{Z}_n^*/4$ otherwise, where $\#X$ denotes the number of elements in set X. Notice that both 1 and $n-1$ always belong to R_n. This theorem is normally used as follows:

function *Rabin*(n)
  $\{n$ is an odd integer, $n>1\}$
  $a \leftarrow$ integer randomly and uniformly selected between 2 and $n-2$
  if $a \in \mathrm{R}_n$ **then return** ''*prime*''
    else return ''*composite*'' .

Whenever n is composite, the error probability of this procedure is clearly given by $(\#\mathrm{R}_n - 2)/(n-3)$, so that elements of R_n others than 1 and $n-1$ are known as *Rabin false witnesses* for n. From the basic theorem, we know that this error probability is always smaller than 25%. However, it is well known to be often *much* smaller. Monier gives an exact formula for this probability [16]; see also [14]. As a corollary of Monier's formula, the error probability never exceeds $(\phi(n)/2^{r-1}-2)/(n-3)$, where r is the number of distinct prime factors of n and $\phi(n) = \#\mathbb{Z}_n^*$ denotes Euler's function [12]. Despite this tightening of the bound on the error probability (at least when n has more than three distinct prime factors), it turns out that the latter is usually *still* much smaller. In other words, Rabin's test performs in practice much better than one might naively expect.

For instance, 42,799 $(=127\times337)$ admits only 880 Rabin false witnesses, compared to $\phi(42{,}799)/4 = 10{,}584$. Even better, Rabin's test *never* fails on integers of the form $3\times5\times7\times11\times\cdots$ such as 15,015: these admit no false witnesses at all. More impressively, it is enough to test deterministically for each $a \in \{2,5,7,13\}$ in order to decide primality without *any* failures up to 25×10^9 (using $\{2,3,5,7\}$ still leaves *one* error in this range [18]). Although ''high risk'' numbers exist, such as $n = 79{,}003$ $(=199\times397)$ or 3,215,031,751 $(=151\times751\times28{,}351)$ with $\#\mathrm{R}_n = \#\mathbb{Z}_n^*/4$ and $(\#\mathrm{R}_n-2)/(n-3) \approx 24.8\%$, these are not the rule (one can nonetheless prove, using Monier's formula, that every composite number of the form $(12m+7)(24m+13)$ is such a high risk number, provided both $12m+7$ and $24m+13$ are prime — but the existence of infinitely many such numbers remains an open question). We address here the following question: why is Rabin's test so good? Said otherwise: what would happen should Rabin's test perform less well?

Consider the set of *Fermat* false witnesses for n:

$$\mathrm{F}_n = \{a \in \mathbb{Z}_n^* \mid a^{n-1} \equiv 1 \pmod{n}\} .$$

Obviously, $\mathrm{R}_n \subseteq \mathrm{F}_n$ for each odd integer $n > 1$. Define $\mathrm{H}_n = \mathrm{F}_n \setminus \mathrm{R}_n$, the set of Fermat false witnesses that are however not Rabin false witnesses. Assume for the moment that n is not a prime power (i.e. not of the form p^m for some prime p and some integer $m \geq 1$). Theorem 1 states that each element of H_n is a *handle* that allows easy splitting of n (i.e. finding at least one non-trivial factor of n) and that there are at least as many such handles as there are Rabin false witnesses. Our provocative interpretation states that it is only possible for even a single iteration of Rabin's test to fail (i.e. declare n prime) with a non-negligible probability if it happens that n is easy to split (and hence obviously composite)! This result extends to every composite n since even numbers and prime powers are easy to split. In other words, there exists a simple probabilistic splitting algorithm whose running time is small on every composite number on which Rabin's test is not extremely effective. More precisely, for any polynomial p, our splitting algorithm succeeds at finding a non-trivial factor in expected polynomial time on all those composite integers n such that prob[$Rabin(n)$ = ''*prime*''] $\geq 1/p(|n|)$, where $|n|$ denotes the size of n (in bits or in decimal digits).

Another consequence of this result is that whenever Fermat's test fails to recognize a composite integer as such whereas Rabin's test would not have failed (with the same random choices), this means that Fermat's test has just missed a golden opportunity to split the given integer! (This phenomenon has already been observed by Baillie and Wagstaff [4, page 1402].)

Theorem 1.

Let n be any odd composite integer that is not a prime power, and let v and u be integers such that $n-1 = 2^v u$, where u is odd.

i) $(\forall a \in \mathrm{H}_n)(\exists j < v)[\gcd(n, a^{2^j u} + 1)$ is a non-trivial factor of n];

ii) $\#\mathrm{H}_n \geq \#\mathrm{R}_n$.

Proof.

i) Consider any $a \in \mathrm{H}_n$. Let i be the smallest integer such that $a^{2^i u} \equiv 1 \pmod{n}$. We have $0 \leq i \leq v$ because $a \in \mathrm{F}_n$. However, $i = 0$ is not possible since $a \notin \mathrm{R}_n$. Let $j = i-1$ and $x = a^{2^j u}$. Clearly, $x \not\equiv 1 \pmod{n}$ because of i's minimality. Moreover, $x \not\equiv -1 \pmod{n}$ since otherwise a would belong to R_n. Hence, $x \not\equiv \pm 1 \pmod{n}$ but $x^2 \equiv 1 \pmod{n}$. Therefore, $\gcd(n, x+1)$ is a non-trivial factor of n [20].

ii) let $n = p_1^{m_1} p_2^{m_2} \cdots p_r^{m_r}$ be the decomposition of n into prime factors. Let $p_i - 1 = 2^{v_i} u_i$ for each i, where u_i is odd. Let $\tau = \min\{ v_i \mid 1 \le i \le r \}$. Monier has proved the following formulae [16]:

$$\#\mathrm{F}_n = \prod_{i=1}^{r} \gcd(n-1, p_i - 1) \ , \quad \text{and}$$

$$\#\mathrm{R}_n = \left[1 + \frac{2^{r\tau} - 1}{2^r - 1} \right] \prod_{i=1}^{r} \gcd(u, u_i) \ .$$

But

$\gcd(n-1, p_i - 1) = \gcd(2^v u, 2^{v_i} u_i) = 2^{\min(v, v_i)} \gcd(u, u_i) \ge 2^{\tau} \gcd(u, u_i)$ because $\tau \le v$ [16]. Hence,

$$\#\mathrm{F}_n \ \ge \ 2^{r\tau} \prod_{i=1}^{r} \gcd(u, u_i) \ .$$

Therefore,

$$\frac{\#\mathrm{F}_n}{\#\mathrm{R}_n} \ \ge \ \frac{2^{r\tau}}{1 + \frac{2^{r\tau} - 1}{2^r - 1}} \ = \ \frac{2^r - 1}{2^{1-r\tau}(2^{r-1} - 1) + 1}$$

$$\ge \ 2^{r-1} \qquad \text{(because } \tau \ge 1) \ .$$

This proves that $\#\mathrm{F}_n \ge 2 \times \#\mathrm{R}_n$ (because $r \ge 2$ by assumption that n is not a prime power), hence $\#\mathrm{H}_n = \#\mathrm{F}_n - \#\mathrm{R}_n \ge \#\mathrm{R}_n$.

A more exact formula is not harder to establish:

$$2^{r+x-1} \le \#\mathrm{F}_n / \#\mathrm{R}_n \ \le 2^x \ (2^r - 1) < 2^{r+x} \ , \quad \text{where}$$

$$x \ = \ \sum_{i=1}^{r} [\, \min(v, v_i) - \tau \,] \ \ge \ 0 \ .$$

In particular, $\mathrm{F}_n = \mathrm{R}_n$ when n is a prime power ($r = 1$). □

There are several ways in which the notion of a *splitting handle* can be generalized. It is clear from part (i) of Theorem 1 that any Fermat false witness that is however not a Rabin false witness can be used to split n efficiently, but other numbers can be used as well. Let us define an *n-splitter* to be an integer x such that $1 < \gcd(x, n) < n$. Obviously, knowledge of any n-splitter can be used to efficiently compute a non-trivial factor of n (using Euclid's algorithm), and there are exactly $n - \phi(n) - 1$ of them between 1 and n. What else could be used as handle?

As usual, let n be an odd composite integer, and let v and u be integers such that $n-1=2^v u$, where u is odd. The following definitions for handles are natural in the sense that they mimic the Rabin and Fermat false witnesses, respectively:

i) a is a *Rabin handle* ($a \in \mathrm{RH}_n$) if $1 \le a < n$ and

- $a^u - 1$ is an n-splitter, or
- $a^{2^j u} + 1$ is an n-splitter for some $0 \le j < v$;

ii) a is a *Fermat handle* ($a \in \mathrm{FH}_n$) if $1 \le a < n$ and

- $a^{n-1} - 1$ is an n-splitter.

We know from Theorem 1 that $\mathrm{H}_n \subseteq \mathrm{RH}_n$, where H_n is the set of handles previously discussed (the Fermat false witnesses that are not Rabin false witnesses), but how many additional handles for splitting n did we miss by concentrating only on H_n? The following theorem tells us that we missed precisely the Fermat handles, and that there are at least $n - \phi(n) - 1$ of them.

Theorem 2.

Let n be an odd integer, $n > 1$, and let $p_1^{m_1} p_2^{m_2} \cdots p_r^{m_r}$ be its decomposition into prime factors.

i) $\mathrm{RH}_n = \mathrm{H}_n \cup \mathrm{FH}_n$;

ii) $\mathrm{H}_n \cap \mathrm{FH}_n = \varnothing$;

iii) $\#\mathrm{FH}_n = n - \#\mathrm{F}_n - n \prod_{i=1}^{r} \left(1 - \frac{\gcd(p_i - 1, n-1)}{p_i}\right) \ge n - \phi(n) - 1$.

(This third part applies as well if n is even.)

iv) $\#\mathrm{RH}_n = n - \#\mathrm{R}_n - n \prod_{i=1}^{r} \left(1 - \frac{\gcd(p_i - 1, n-1)}{p_i}\right)$.

In order to simplify the proof, let us first state some useful facts whose justifications are left to the reader.

a) In general, for any integers $x \ge 2$ and $k \ge 1$, $x^k - 1$ is divisible by $x - 1$; it is also divisible by $x + 1$ provided k is even. In particular, for any integers $a \ge 2$, $i \ge j \ge 0$ and $u \ge 1$, $a^{2^i u} - 1$ is divisible by $a^{2^j u} - 1$; it is also divisible by $a^{2^j u} + 1$ provided $j < i$.

b) It is obvious that $\mathrm{R}_n \subseteq \mathrm{F}_n$, $\mathrm{R}_n \cap \mathrm{H}_n = \varnothing$, $\mathrm{F}_n \cap \mathrm{FH}_n = \varnothing$ and $\mathrm{R}_n \cap \mathrm{FH}_n = \varnothing$; it is somewhat less obvious that $\mathrm{R}_n \cap \mathrm{RH}_n = \varnothing$ (hint: facts (a) above are relevant).

c) Let p be any prime, then
$\#\{a \mid 0 \le a < p \text{ and } a^{n-1} \equiv 1 \pmod{p}\} = \gcd(p-1, n-1)$.

d) Consider any integer $c \ge 1$. It follows from elementary calculus that $\max\{x + c/x \mid 1 \le x \le c\} = 1 + c$.

Proof of Theorem 2.

As always, let v and u be integers such that $n - 1 = 2^v u$, where u is odd.

i) Consider any $a \in \mathrm{RH}_n$. By fact (a), $a^{n-1} - 1$ is divisible by $a^u - 1$ and by $a^{2^j u} + 1$ for each $j < v$. Thus, by definition of RH_n, $\gcd(a^{n-1} - 1, n) > 1$. If this gcd is equal to n, then $a \in \mathrm{F}_n$ by definition. But $a \notin \mathrm{R}_n$ by fact (b). Therefore, $a \in \mathrm{F}_n \grave{e} \mathrm{R}_n = \mathrm{H}_n$. On the other hand, if this gcd is strictly between 1 and n then $a \in \mathrm{FH}_n$ by definition. This shows that $\mathrm{RH}_n \subseteq \mathrm{H}_n \cup \mathrm{FH}_n$.

Now, consider any $a \in \mathrm{H}_n \cup \mathrm{FH}_n$. By fact (b), $a \notin \mathrm{R}_n$. By definition, $\gcd(a^{n-1} - 1, n) > 1$. Let p be a prime factor of this gcd. Clearly, $a^{n-1} \equiv 1 \pmod{p}$, which implies (because p is a prime) that either $a^u \equiv 1 \pmod{p}$ or $a^{2^j u} \equiv -1 \pmod{p}$ for some $j < v$. Thus either $\gcd(a^u - 1, n) > 1$ or $\gcd(a^{2^j u} + 1, n) > 1$ for some $j < v$. But none of these gcd's can be n since $a \notin \mathrm{R}_n$. Thus $a \in \mathrm{RH}_n$, which completes the proof that $\mathrm{H}_n \cup \mathrm{FH}_n \subseteq \mathrm{RH}_n$.

ii) Immediate from the facts that $\mathrm{F}_n \cap \mathrm{FH}_n = \varnothing$ and $\mathrm{H}_n \subseteq \mathrm{F}_n$.

iii) Consider any $0 \le a < n$. It is clear that exactly one of the following three possibilities holds: $a \in \mathrm{F}_n$, $a \in \mathrm{FH}_n$ or $\gcd(a^{n-1} - 1, n) = 1$. Therefore, the equality in part (iii) will follow immediately after we show that

$$\#\{a \mid 0 \le a < n \text{ and } \gcd(a^{n-1} - 1, n) = 1\}$$
$$= n \prod_{i=1}^{r} \left(1 - \frac{\gcd(p_i - 1, n-1)}{p_i}\right).$$

To see this, consider any i, $1 \le i \le r$. By fact (c),

$$\#\{a \mid 0 \le a < p_i \text{ and } a^{n-1} \equiv 1 \pmod{p_i}\} = \gcd(p_i - 1, n - 1).$$

Therefore,

$$\#\{a \mid 0 \le a < p_i^{m_i} \text{ and } a^{n-1} \equiv 1 \pmod{p_i}\} = p_i^{m_i - 1} \gcd(p_i - 1, n - 1).$$

We thus conclude that

$$\#\{a \mid 0 \le a < p_i^{m_i} \text{ and } a^{n-1} \not\equiv 1 \pmod{p_i}\}$$
$$= p_i^{m_i} - p_i^{m_i-1}\gcd(p_i - 1, n-1)\ .$$

From the Chinese Remainder Theorem, it follows that the number of a such that $0 \le a < n$ and $a^{n-1} \not\equiv 1 \pmod{p_i}$ for each $1 \le i \le r$ is given by

$$\prod_{i=1}^{r}(p_i^{m_i} - p_i^{m_i-1}\gcd(p_i - 1, n-1)) \;=\; n\prod_{i=1}^{r}\left(1 - \frac{\gcd(p_i - 1, n-1)}{p_i}\right)\ .$$

Thus the equality in part (iii) holds since $a^{n-1} \not\equiv 1 \pmod{p_i}$ for each $1 \le i \le r$ precisely when $\gcd(a^{n-1} - 1, n) = 1$.

In order to show the inequality in part (iii) of the theorem, let x_i denote $\gcd(p_i - 1, n-1)$, let $\alpha(n)$ denote $\prod_{i=1}^{r} p_i$, the largest square-free divisor of n, and let $\beta(n) = \phi(\alpha(n))$ denote $\prod_{i=1}^{r}(p_i - 1)$. By Monier's formula [16] (also found independently by Baillie and Wagstaff [4]), $\#\mathrm{F}_n = \prod_{i=1}^{r} x_i$. Therefore, the equality we just proved implies that

$$\#\mathrm{FH}_n \;=\; n - \prod_{i=1}^{r} x_i - n\prod_{i=1}^{r}(1 - x_i/p_i)$$
$$= n - \prod_{i=1}^{r} x_i - \frac{n}{\alpha(n)}\prod_{i=1}^{r}(p_i - x_i)\ .$$

Consider now the function Ξ_n in r variables defined by

$$\Xi_n(y_1, y_2, \ldots, y_r) \;=\; \prod_{i=1}^{r} y_i + \frac{n}{\alpha(n)}\prod_{i=1}^{r}(p_i - y_i)\ ,$$

which is defined precisely so that $\#\mathrm{FH}_n = n - \Xi_n(x_1, x_2, \ldots, x_r)$. Because $1 \le x_i \le p_i - 1$ for each $1 \le i \le r$, it is clear that

$$\#\mathrm{FH}_n \ge n - \max\{\Xi_n(y_1, y_2, \ldots, y_r) \mid 1 \le y_i \le p_i - 1 \text{ for each } 1 \le i \le r\}.$$

In order to maximize the function Ξ_n, first notice that it is linear in each variable. Therefore, its maximum value occurs at a point where each y_i is either 1 or $p_i - 1$. Let S be an arbitrary subset of $\{1, 2, \ldots, r\}$ and let x denote $\prod_{i \in S}(p_i - 1)$. Notice that $\prod_{i \notin S}(p_i - 1) = \beta(n)/x$ and $n\,\beta(n)/\alpha(n) = \phi(n)$. Setting $y_i = p_i - 1$ for $i \in S$ and $y_i = 1$ for $i \notin S$, we thus have $\Xi_n(y_1, y_2, \ldots, y_r) = x + \phi(n)/x$. Using fact (d), since x can only take values between 1 and $\beta(n) \le \phi(n)$, we conclude that the value of Ξ_n in the range of interest never exceeds $1 + \phi(n)$. We have thus

established that $\# \mathrm{FH}_n \geq n - \phi(n) - 1$, as claimed.

iv) Immediate from the first three parts of this theorem and from the obvious fact that $\# \mathrm{H}_n = \# \mathrm{F}_n - \# \mathrm{R}_n$. □

3. How Good is Rabin's Test?

We must first ask the following question: *what* is Rabin's test good for? At least two answers come to mind: to *test* the primality of a given integer and to *generate* one or several random integers that are probably prime. We consider these two settings in turns, starting with the second.

3.1. How to Generate Random Numbers that are Probably Prime

The generation of large primes drawn with a uniform distribution from the set of all primes of a given size is of crucial importance in cryptography [22]. Although it is possible to generate such primes with certainty using the algorithms of [2, 1], their running time is currently too high to be used in practice. It is also possible to efficiently generate large certified primes by a variation on Pratt's non-deterministic algorithm [19] (generate the **NP** certificate and the resulting prime hand in hand) or by more sophisticated techniques [8], but the resulting distribution would not be uniform. Again, the most attractive solution in practice is to use Rabin's test as follows:

```
function GenPrime(l, k)
  { l is the size of the prime to be produced;
    k is a safety parameter discussed below }
  repeat
    n ← randomly selected l-digit odd integer
  until RepeatRabin(n, k) = ''prime''
  return n .
```

The resulting output is a *probabilistic prime* in the sense that we can never be assured that it is indeed prime. We can nonetheless increase our confidence in the number's primality by increasing the safety parameter k. (What a shame that Rabin's algorithm can certify those cryptographically useless composite numbers whereas it can only give probabilistic information on the useful primes! — which is precisely why the algorithms in [11, 1] are of such (as yet theoretical) interest.)

In order to use *GenPrime* for cryptographic purposes, it is important that its probability of returning a composite integer be estimated. The popular belief is that

$$\mathrm{prob}[\, \mathit{GenPrime}(l, k) \text{ is composite} \,] \leq 4^{-k}$$

because each of the k rounds of *RepeatRabin* has a probability smaller than ¼ of failing on any given composite number. If we repeatedly use *GenPrime* to produce m distinct ''primes'', we therefore expect on the average that less than $m \times 4^{-k}$ of them will turn out to be composite. For instance, Knuth writes [13, page 379]:

> ''If we certified a billion different primes with such a procedure[1], the expected number of mistakes would be less than $\frac{1}{1000000}$''

This assertion *is* true, but its proof is not so obvious. In particular, the simple minded reasoning given above is fallacious. Indeed, it is *only* true because prob[*Rabin*(n) = ''*prime*''] is much smaller than ¼ on most composite numbers. Should the error probability of Rabin's test be *exactly* ¼ on each and every composite odd integer, the number of expected errors in Knuth's quote would be *significantly larger* than $10^9 \times 4^{-25} \approx 10^{-6}$.

Let X stand for ''n is composite'' and Y_k for ''*RepeatRabin*(n, k) returned ''*prime*'''''. The basic theorem about Rabin's test obviously implies that $\text{prob}[Y_k \mid X] \le 4^{-k}$. However, this is *not* the probability relevant to the procedure *GenPrime* : rather, we are interested in $\text{prob}[X \mid Y_k]$. This latter probability cannot be estimated (or even bounded away from 1) without specification of the probability distribution over n. For the moment, let S be any set of odd integers and let n be chosen randomly and uniformly among S. Let p stand for the probability that n be prime (i.e.: $p = \#\{n \in S \mid n \text{ is prime}\}/\#S$), and let us assume that $0 < p < 1$. Elementary probability theory yields:

$$
\begin{aligned}
\text{prob}[X \mid Y_k] &= \frac{\text{prob}[X] \times \text{prob}[Y_k \mid X]}{\text{prob}[Y_k]} \\
&= \frac{\text{prob}[X] \times \text{prob}[Y_k \mid X]}{\text{prob}[X] \times \text{prob}[Y_k \mid X] + \text{prob}[\textbf{not } X] \times \text{prob}[Y_k \mid \textbf{not } X]} \\
&= \frac{1}{1 + \dfrac{\text{prob}[\textbf{not } X] \times \text{prob}[Y_k \mid \textbf{not } X]}{\text{prob}[X] \times \text{prob}[Y_k \mid X]}} \\
&= \frac{1}{1 + \dfrac{p}{1-p} \dfrac{1}{\text{prob}[Y_k \mid X]}}
\end{aligned}
$$

1. Knuth does not *explicitly* say how he would use Rabin's test to certify those billion primes, except that he would run it ''25-times-in-a-row'' on each of them. It is our interpretation that he meant something along the lines of *RepeatRabin*(•, 25). Of course, Knuth's assertion is vacuously true if taken literally: if the integers thus certified are indeed ''one billion primes'', no mistakes are possible at all!

$$\leq \; p^{-1} \times \text{prob}[Y_k \mid X]$$

In particular, if $\text{prob}[Y_k \mid X] \ll p \ll 1$, as would be the case for the generation of a large probabilistic prime with $k = 25$, we get that $\text{prob}[X \mid Y_k]$ is almost equal to $p^{-1} \times \text{prob}[Y_k \mid X]$. It is thus clear that $\text{prob}[X \mid Y_k] \leq 4^{-k}$ cannot be a direct consequence of the mere fact that $\text{prob}[Y_k \mid X] \leq 4^{-k}$. Worse still, $\text{prob}[X \mid Y_k] \leq 4^{-k}$ would be *false* should p be small and should $\text{prob}[Y_k \mid X]$ be insufficiently smaller than 4^{-k}. This line of thought is carried out to the extreme in [5] and unduly pessimistic consequences are drawn. Nonetheless, the more detailed analysis below shows that the precautions advocated in [5] are not necessary. In particular, the probability of $\mathit{GenPrime}(l, k)$ returning a composite integer actually *decreases* as the value of l increases (for any fixed value of k).

In order to bound $\text{prob}[X \mid Y_k]$ more precisely, a tighter bound on $\text{prob}[Y_k \mid X]$ is needed. Thus, we must study the *average* behaviour of Rabin's test, rather than its much simpler worst case. This question was raised and solved to a large degree by Erdos and Pomerance [9], but only for the case of a *single iteration* of Rabin's test, corresponding to $k = 1$: the average number of false witnesses for integers up to n does not exceed $n / L(n)^{\xi(n)}$, where $L(n)$ is defined as $n^{\ln\ln\ln n / \ln\ln n}$ and $\xi : \mathbb{N} \to \mathbb{R}$ is a function such that $\lim_{n \to \infty} \xi(n) = 1$. (This holds even for the *Fermat* congruence.)

However, successive runs of Rabin's test on a given randomly chosen odd integer n are not independent in the sense that if a first run finds that $a_1 \in \mathrm{R}_n$, this increases very significantly the likelihood that the next run will also find that $a_2 \in \mathrm{R}_n$, where a_1 and a_2 are randomly, uniformly and independently chosen between 2 and $n-2$. (Let us disregard here the fact that the second run would not even take place should the first run find that $a_1 \notin \mathrm{R}_n$.) Consult [18] for numerical evidence about this. This lack of independence prevents us from asserting that $\text{prob}[Y_k \mid X] = (\text{prob}[Y_1 \mid X])^k$. Nonetheless, the basic fact that $\text{prob}[a_2 \in \mathrm{R}_n] \leq ¼$ remains true as long as n is composite, and this applies equally well to each and every run, independently of what happens in the other runs. Therefore,

$$\text{prob}[Y_k \mid X] \;\leq\; \frac{\text{prob}[Y_1 \mid X]}{4^{k-1}} \;.$$

Let us now be more specific about the set S of odd integers among which n is randomly and uniformly drawn, so that our study corresponds to a call on

$GenPrime(l,k)$ for positive integers l and k. Let S be the set of l-digit odd integers. The prime number theorem [12] tells us that

$$p = \text{prob}[n \text{ is prime} \mid n \in S] \approx \beta/l,$$

where $\beta = \log_{10} e^2 \approx 0.87$ (see Lemma 1 in [5]). Moreover, the study in [9] tells us that

$$\text{prob}[Y_1 \mid X] \le 1/P(10^{l-1}),$$

where $P(n)$ stands for $L(n)^{\xi(n)}$. Hence,

$$\begin{aligned}\text{prob}[GenPrime(l,k) \text{ is composite}] &= \text{prob}[X \mid Y_k] \\ &\le p^{-1} \times \text{prob}[Y_k \mid X] \le \frac{4}{\beta} \frac{l}{P(10^{l-1})} 4^{-k}.\end{aligned}$$

But clearly $\lim_{l\to\infty} l/P(10^{l-1}) = 0$, hence $\text{prob}[GenPrime(l,k) \text{ is composite}] \le 4^{-k}$ for all sufficiently large l, thus proving Knuth's previously quoted claim.

Let us stress again that, although one's confidence in the primality of the output of $GenPrime(l,k)$ increases with the value of k, $l/P(10^{l-1})$ is already so small for large values of l that an arbitrarily low error probability can be obtained with a mere call on $GenPrime(l,1)$ if large enough primes are sought. In fact, this remains true even if the simpler Fermat congruence is used instead of Rabin's test, but one should resist the temptation of doing this because Fermat's congruence is only simpler *conceptually* but never faster (and often slower) to compute.

3.2. How to Decide on the Primality of a Given Integer

Suppose some odd integer n is given to you. You are to decide whether you think it is prime or not. You therefore run Rabin's test for some number k of rounds, and it never finds n to be composite. What can you tell from this?

One obviously wrong answer is: ''this number is prime with probability $1-4^{-k}$''. This makes no sense because any *given* integer is either prime or not.

The classic answer is: ''I believe this number to be prime, and my error probability is at most 4^{-k} (in the sense that I expect to be wrong at most once every 4^k such statements if you quiz me long enough)''. This is wrong as well because *no* estimate on the error probability of ''I believe this number to be prime'' can be made without an *a priori* estimate on the probability that the number is prime. If you know that n was chosen randomly and uniformly among the odd integers of some given size, section 3.1 is relevant. However, if you do *not* know where the number comes from, you are at a *complete* loss.

Still, there *is* one thing you can say: ''I believe this number to be prime, and if I am wrong I have observed a natural phenomenon whose probability of occurrence was bounded by 4^{-k} ''. This statement is certainly weak, but we cannot think of anything stronger one can infer in general from running Rabin's test any number of times.

As mentioned in the abstract, this observation is not restricted to primality testing. Whenever one runs *any* probabilistic algorithm that is not Las Vegas, care must be taken as to how to interpret the outcome. This general issue is discussed in [6].

4. Open Problems

It would be interesting to analyse the average and normal behaviour of $\#\mathrm{FH}_n$ and $\#\mathrm{RH}_n$, much like the analyses of $\#\mathrm{F}_n$ and $\#\mathrm{R}_n$ found in [9]. It would also be interesting to analyse more tightly the probability that $\mathit{RepeatRabin}(n,k) = \text{''}\mathit{prime}\text{''}$ when n is a composite integer randomly and uniformly chosen among all l-digit composite integers; in particular, where does it lie between p^k and $p/4^{k-1}$, where p stands for the probability corresponding to $k=1$? (We conjecture that it is much closer to $p/4^{k-1}$ than to p^k, i.e. that our analysis was not too pessimistic.) Finally, a more precise analysis of how and how fast the function $\xi(n)$ converges to 1 (a question left unresolved in [9]) would be crucial in establishing from which values of l it is true that, for every integer $k \geq 1$, prob[$\mathit{GenPrime}(l,k)$ is composite] $\leq 4^{-k}$.

Acknowledgements

The first observation sprung from an idea attributed to Silvio Micali by Charles H. Bennett (private communication) to the effect that any algorithm for extracting square roots modulo a prime can be turned into a probabilistic algorithm to decide primality. However, this origin is no longer perceptible in this paper (but see [5]). The proof of Theorem 1(ii) given here was suggested to one of the authors by Evangelis Kranakis; it is much nicer but far less elementary than our original proof. A very simple proof of Theorem 1 for the case when $n \equiv 3 \pmod 4$ can be found in [5].

We wish to thank all those who showed interest for this work at the Crypto 86 meeting, in particular: Yvo Desmedt, Oded Goldreich, Gus Simmons and Moti Yung. Finally, we are very grateful to Ernest F. Brickell, the *Journal*'s editor-in-chief, for his efficiency in handling our paper.

REFERENCES

[1] **Adleman, L. and M.-D. Huang, ''Recognizing primes in random polynomial time'',** *Proceedings of the 19th Annual ACM Symposium on the Theory of Computing*, **pp. 462-469, 1987.**

[2] **Adleman, L., C. Pomerance and R. Rumeley, ''On distinguishing prime numbers from composite numbers'',** *Annals of Mathematics*, **vol. 117, pp. 173-206, 1983.**

[3] **Babai, L., ''Monte Carlo algorithms in graph isomorphism testing'',** *Rapport de Recherches du Département de Mathématiques et de Statistiques*, **Université de Montréal, D.M.S. #79-10, 1979.**

[4] **Baillie, R. and S. S. Wagstaff, Jr., ''Lucas pseudoprimes''** *Mathematics of Computation*, **vol. 35, no. 152, pp. 1392-1417, 1980.**

[5] **Beauchemin, P., G. Brassard, C. Crépeau and C. Goutier, ''Two observations on probabilistic primality testing'',** *Advances in Cryptology — CRYPTO 86 Proceedings*, **Springer-Verlag, pp. 443-450, 1987.**

[6] **Brassard, G. and P. Bratley,** *Algorithmics : Theory and Practice*, **Prentice-Hall, Englewood Cliffs, New Jersey, 1988.**

[7] **Cohen, H. and A. K. Lenstra, ''Implementation of a new primality test'',** *Mathematics of Computation*, **vol. 48, no. 177, pp. 103-121.**

[8] **Couvreur, C. and J.-J. Quisquater, ''An introduction to fast generation of large prime numbers'',** *Philips Journal of Research*, **vol. 37, nos. 5/6, pp. 231-264, 1982.**

[9] **Erdos, P. and C. Pomerance, ''On the number of false witnesses for a composite number'',** *Mathematics of Computation*, **vol. 46, no. 173, pp. 259-279, 1986.**

[10] **Gill, J., ''Computational complexity of probabilistic Turing machines'',** *SIAM Journal on Computing*, **vol. 6, no. 4, pp. 675-695, 1977.**

[11] **Goldwasser, S. and J. Kilian, ''Almost all primes can be quickly certified'',** *Proceedings of the 18th Annual ACM Symposium on the Theory of Computing*, **pp. 316-329, 1986.**

[12] **Hardy, G. H. and E. M. Wright,** *An Introduction to the Theory of Numbers*, **Fifth edition, Oxford Science Publications, 1979.**

[13] **Knuth, D. E.,** *The Art of Computer Programming*, **volume 2 :** *Seminumerical Algorithms*, **Second edition, Addison-Wesley, Reading, Massachusetts, 1981.**

[14] **Kranakis, E.,** *Primality and Cryptography*, **Wiley-Teubner Series in Computer Science, 1986.**

[15] **Miller, G. L., ''Riemann's hypothesis and tests for primality'',** *Journal of Computer and System Sciences*, **vol. 13, pp. 300-317, 1976.**

[16] **Monier, L., ''Evaluation and comparison of two efficient probabilistic primality testing algorithms'',** *Theoretical Computer Science*, **vol. 11, pp. 97-108, 1980.**

[17] **Pomerance, C., ''The search for prime numbers'',** *Scientific American*, **vol. 247, no. 6, pp. 136-147, 1982.**

[18] **Pomerance, C., J. L. Selfridge and S. S. Wagstaff, Jr., ''The pseudoprimes to** 25.10^9 **'',** *Mathematics of Computation*, **vol. 35, no. 151, pp. 1003-1026, 1980.**

[19] **Pratt, V., ''Every prime has a succinct certificate'',** *SIAM Journal on Computing*, **pp. 214-220, 1975.**

[20] **Rabin, M. O., ''Digitalized signatures and public-key functions as intractable as factorization'',** *MIT/LCS/TR*-**212, 1979.**

[21] **Rabin, M. O., ''Probabilistic algorithm for testing primality'',** *Journal of Number Theory*, **vol. 12, pp. 128-138, 1980.**

[22] **Rivest, R. L., 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.**

[23] **Solovay, R. and V. Strassen, ''A fast Monte Carlo test for primality'',** *SIAM Journal on Computing*, **vol. 6, pp. 84-85, 1977.**