Uncertainty can be Better than Certainty:

Some Algorithms for Primality Testing[*]

Richard P. Brent
Australian National University

[*]Presented at ANU, 27 November 2006.
Copyright ©2006, R. P. Brent. AdvCom2t

Abstract

For many years mathematicians and computer scientists have searched for a fast and reliable primality test. This is especially relevant nowadays, because the popular RSA public-key cryptosystem requires very large primes in order to generate secure keys. I will describe some efficient randomised algorithms that are useful, but have the defect of occasionally giving the wrong answer, or taking a very long time to give an answer.

In 2002, Agrawal, Kayal and Saxena (AKS) found a deterministic polynomial-time algorithm for primality testing. I will describe the original AKS algorithm and some improvements by Bernstein and Lenstra. As far as theory is concerned, we now know that "PRIMES is in P", and this appears to be the end of the story. However, I will explain why it is preferable to use randomised algorithms in practice.

First, some notation

As usual, we say that

$$f(n) = O(n^k)$$

if, for some c and n_0, for all $n \geq n_0$,

$$|f(n)| \leq cn^k \ .$$

We say that

$$f(n) = \widetilde{O}(n^k)$$

if, for all $\varepsilon > 0$,

$$f(n) = O(n^{k+\varepsilon}) \ .$$

The "$\widetilde{O}$" notation is useful to avoid terms like $\log n$ and $\log\log n$. For example, when referring to the Schönhage-Strassen algorithm for n-bit integer multiplication, it is easier to write

$$\widetilde{O}(n)$$

than the (more precise)

$$O(n \log n \log\log n) \ .$$

Complexity Classes (informal)

P is the class of decision ("yes"/"no") problems that have a *deterministic polynomial time algorithm*, i.e. an algorithm running in time $O(\lambda^k)$ on inputs of *length* λ, for some fixed k.

RP is the class of decision problems that have a randomised algorithm running in time $O(\lambda^k)$, with (one-sided) error probability at most $1/2$ (if the *correct* answer is "yes").
co-RP is similar, but permitting an error on the other side (if the correct answer is "no").

BPP is similar but allows errors (with probability at most $1/4$) on both sides.

ZPP is the class of decision problems that have an error-free randomised ("Las Vegas") algorithm running in *expected* time $O(\lambda^k)$.

NP is the class of decision problems for which, if the answer is "yes", there is a "certificate" that can be *verified* in time $O(\lambda^k)$.
co-NP is similar (for the answer "no").

Containment Relationships

$$P \subseteq ZPP = RP \cap \text{co-}RP\,,$$

$$RP \cup \text{co-}RP \subseteq BPP\,,$$

$$P \subseteq RP \subseteq NP\,,$$

$$P \subseteq \text{co-}RP \subseteq \text{co-}NP\,.$$

It is not known if any of the inclusions is strict, or whether $BPP \subseteq NP$ or $NP \subseteq BPP$.

Note that $P \neq RP$ or $RP \neq NP$ implies

$$P \neq NP\,.$$

This inequality is one of the seven $1,000,000 *Millenium Prize Problems*, in fact the only one that is directly relevant to Computer Science.

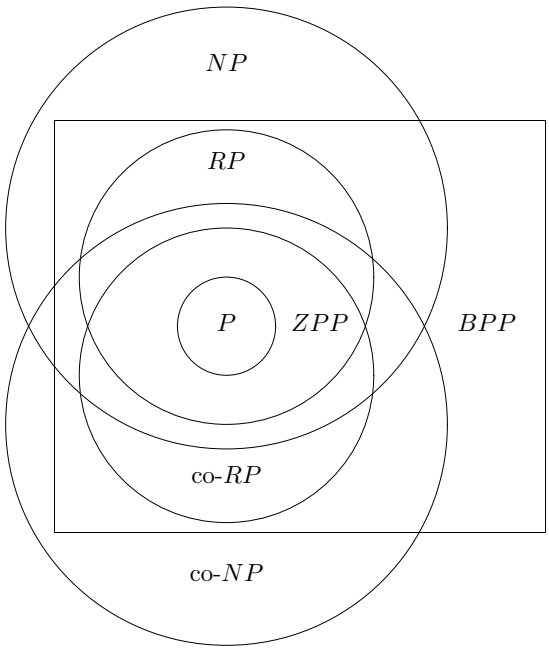

Reducing Probability of Error

Given a problem in RP, co-RP, or BPP, we can reduce the probability of error below any given $\varepsilon > 0$ by iterating $O(\log(1/\varepsilon))$ times with independent random choices at each iteration. In this context an iteration is usually called a "trial".

For example, if an algorithm in RP has (one-sided) error probability at most $1/4$, we can perform 10 independent trials to get an algorithm with error probability at most

$$\frac{1}{4^{10}} < \frac{1}{1000000}\,.$$

The Rabin-Miller primality-testing algorithm is an example of such an algorithm. We'll show later that it is actually an RP algorithm for testing *compositeness*, so if PRIMES is the problem of testing primality, we have:

$$\text{PRIMES} \in \text{co-}RP.$$

Factoring

Every positive integer has a unique factorisation into a product of prime powers. That is the main reason why primes are important: they are "building blocks" for the integers.

However, there does not seem to be any simple connection between algorithms for factoring and algorithms for testing primality. We'll see that there are algorithms that can answer the question

Is n prime?

much faster than any known algorithm for finding the prime factors of n (if it is composite).

The popular RSA cryptosystem needs two large primes p and q whose product $n = pq$ is difficult to factor. Primality testing algorithms are useful for generating p and q, but they are not useful for *cracking* RSA by factoring n (if p and q are kept secret).

Factoring Primes ?

Primality testing and factorisation are often confused, even by such luminaries as Bill Gates.

> *The obvious mathematical breakthrough would be development of an easy way to factor large prime numbers.*
>
> – Bill Gates *et al*, 1995[1]

Presumably Gates did not have in mind factorisations such as

$$13 = (2+3i)(2-3i)\ .$$

[1] Bill Gates, Nathan Myhrvold and Peter M. Rinearson, *The Road Ahead*, Viking Press, 1995, p. 265.

Fermat's Little Theorem

If n is prime and a is any integer, then

$$a^n = a \mod n.$$

Thus, if we find a such that $a^n \neq a \mod n$, we can be sure that n is composite. We say that

"a is a *witness* to the compositeness of n".

Note: we can guarantee that n is composite *without knowing the factors of* n.

The converse of Fermat's little theorem is false: if $a^n = a \mod n$ we can not be sure that n is prime. There are infinitely many examples (called *Carmichael numbers*) of composite n for which a^n is always $a \mod n$. The smallest example is

$$n = 561 = 3 \cdot 11 \cdot 17\ ,$$

$$\phi(n) = \mathrm{LCM}(2, 10, 16) = 80 \mid (n-1)\ .$$

The number of Carmichael numbers up to N is at least of order $N^{2/7}$ (Alford, Granville and Pomerance, 1994) so we can't ignore them!

Certificates

In the definition of NP we mentioned certificates. A *certificate* for a property is some information that enables a proof of the property to be completed in polynomial time.

For example, a "witness" a such that $a^n \neq a \mod n$ provides a certificate of the compositeness of n. A nontrivial factor of n would also provide a certificate.

Pratt (1975) showed that every prime p has a certificate of length $O((\log p)^2)$. The idea is to write

$$p - 1 = p_1^{\alpha_1} \cdots p_\nu^{\alpha_\nu}$$

and give a *primitive root* a of p. This can be verified by checking that

$$a^{p-1} = 1 \mod p$$

and

$$a^{(p-1)/p_\beta} \neq 1 \mod p \text{ for } \beta = 1, \ldots, \nu\ .$$

We (recursively) give certificates for $p_1, \ldots, p_\nu$ unless they are sufficiently small (say < 100).

Consequence of Pratt's Theorem

From Pratt's result,

$$\text{PRIMES} \in NP.$$

Note that we don't claim that Pratt's certificate of primality of n is easy to find. Finding it is as difficult as factoring, since it requires the factorisation of $n-1$. The mere *existence* of the certificate is sufficient.
If n is composite, a nontrivial factor f of n (satisfying $1 < f < n$) provides a certificate of compositeness. Again, we don't claim that this certificate is easy to find! Its existence is enough to show that

$$\text{PRIMES} \in \text{co-}NP.$$

Thus, we know that

$$\text{PRIMES} \in NP \cap \text{co-}NP.$$

$NP \cap \text{co-}NP$ is believed to be smaller than NP, but a proof of this would imply that $P \neq NP$, so is likely to be difficult.

First Extension of Fermat

A slight extension of Fermat's little Theorem is useful for primality testing. The idea is simple. Suppose we are testing the primality of n, and we find x such that

$$x^2 = 1 \mod n .$$

If n is prime, then $x = \pm 1 \mod n$.

However, if n is composite, then it is possible (and quite likely) that $x \neq \pm 1 \mod n$. For example, consider $n = 21$ and $x = 8$.
More generally, if $n = p_1 p_2$, take
$x = -1 \mod p_1$, $x = +1 \mod p_2$.

When applying the Fermat test to n, we compute $a^{n-1} \mod n$ for some choice of a. The Fermat test is passed if $a^{n-1} = 1 \mod n$. If n is odd, then the exponent $n - 1$ is even, and we can take x to be a suitable power of a. This gives the following primality test.

First Extension of Fermat cont.

If $n = 2^k q + 1$ is an odd prime, and $0 < a < n$, then either $a^q = 1 \mod n$, or the sequence

$$S = (a^q, a^{2q}, a^{4q}, \ldots, a^{2^k q}) \mod n$$

ends with 1, and the value just preceding the first appearance of 1 must be $-1 \mod n$.

That is, the sequence S looks like

$$\begin{array}{c} (1, 1, \ldots 1) \quad \text{if } a^q = 1 \mod n, \\ (?, \ldots, ?, -1, 1, \ldots, 1) \quad \text{otherwise.} \end{array}$$

Proof: If $x^2 = 1 \mod n$ then $n|(x-1)(x+1)$. Since n is prime, $n|(x-1)$ or $n|(x+1)$. Thus $x = \pm 1 \mod n$. □

This fact has been known for a long time. Relevant names are Dubois, Selfridge, Artjuhov, and Miller. Rabin was the first to prove its usefulness in a randomised algorithm, usually called the *Rabin-Miller* algorithm.

The Rabin-Miller Algorithm

The extension of Fermat's little Theorem gives a *necessary* (but not sufficient) condition for primality of n. The Rabin-Miller algorithm checks if this condition is satisfied for a random choice of a, and returns "yes" if it is, "no" otherwise.

If n is prime, the answer is always "yes" (correct).

If n is composite, the answer could be "yes" (wrong) or "no" (correct), but it is "no" with probability greater than 3/4, i.e. the probability of error is $< 1/4$ (this is a theorem of Rabin).

Thus we have an RP-algorithm for testing compositeness. We can say that compositeness is in RP, or (equivalently)

$$\text{PRIMES} \in \text{co-}RP.$$

The Rabin-Miller algorithm gives a certificate for compositeness, but not a certificate for primality.

Popular Primality-Testing Algorithms

Popular algorithms include:

- The Rabin-Miller algorithm (1976).
- The *Jacobi Sums* algorithm of Adleman, Pomerance and Rumely (1983).
- The *Elliptic Curve Primality Proving* algorithm ECPP of Atkin and Morain (1993), based on a proposal by Goldwasser and Kilian.

These algorithms all have their good (and bad) points. We'll look at each more closely before discussing

- The AKS[2] algorithm of Agrawal, Kayal and Saxena (2002).

It will be convenient to define

$$\lambda = \log n ,$$

where n is the number being tested for primality.

[2]Not the AKS parallel sorting algorithm of Ajtai, Komlós and Szemerédi!

Jacobi Sums

The *Jacobi Sums* algorithm runs in time

$$\lambda^{O(\log\log\lambda)} .$$

This is *almost* polynomial time[3].

We can be more precise: Odlyzko and Pomerance have shown that, for all large n, the running time is in

$$[\lambda^{A\log\log\lambda},\ \lambda^{B\log\log\lambda}] ,$$

where A, B are positive constants. The lower bound shows that the Jacobi Sums algorithms is definitely not polynomial-time (in theory anyway).

The Jacobi sums algorithm is deterministic and practical: it has been used for numbers of at least 3395 decimal digits (Mihailescu: 6.5 days on a 500 Mhz Dec Alpha).

[3]Recall that $\lambda = \log n$ so $\log\log\lambda = \log\log\log n$. While it has been *proved* that $\log\log\log n \to +\infty$ with n, it has never been *observed* doing so [Pomerance].

ECPP

The *Elliptic Curve Primality Proving* algorithm ECPP runs in *expected* polynomial time under some plausible assumptions, but the time bound has not been proved rigorously. With an improvement suggested by Shallit, the (conjectured) expected time is $\widetilde{O}(\lambda^4)$.

ECPP is a Las Vegas algorithm: the running time is random but the result is error-free.

ECPP is practical and has been used to prove primality of a number $4405^{2638} + 2638^{4405}$ of 15071 decimal digits. The total CPU time was 5.1 Ghz-years (Franke, Kleinjung, Morain, and Wirth, July 2004).

In practice ECPP is comparable to the Jacobi Sums algorithm, but ECPP has the advantage of producing an easily-checked certificate of primality. In fact, ECPP produces a certificate of size $O(\lambda^2)$ that can be checked in deterministic polynomial time $\widetilde{O}(\lambda^3)$.

Adelman-Huang

There is a complicated algorithm[4], due to Adleman and Huang (1992), that gives a certificate of primality in *expected* polynomial time. Thus

$$\text{PRIMES} \in RP.$$

It follows from Adelman-Huang and Rabin-Miller that

$$\text{PRIMES} \in RP \cap \text{co-}RP = ZPP.$$

Recall that ZPP is very close to P. The difference is that for ZPP the *expected* running time is polynomial in λ, but for P the *worst-case* running time is polynomial in λ.

In practice no one uses the Adelman-Huang algorithm because ECPP is much faster. Adelman-Huang is of theoretical interest because we can *prove* that its expected running time is polynomial in λ.

[4]Using abelian varieties, which are generalisations of elliptic curves.

Rabin-Miller

Rabin-Miller is a Monte Carlo algorithm: there is a nonzero probability of error. In practice the probability of error is negligible (less than 10^{-6}) if we take at least ten independent trials.

The algorithm is fast: one trial takes time

$$\widetilde{O}(\lambda^2)$$

(or $O(\lambda^3)$ with classical $O(\lambda^2)$ multiplication).

Rabin-Miller is feasible for numbers of 10^6 decimal digits[5]. It produces a certificate of compositeness, but not a certificate of primality.

Rabin-Miller can be turned into a deterministic $\widetilde{O}(\lambda^4)$ algorithm if we assume the "Extended Riemann Hypothesis" (ERH). However, no one knows how to prove ERH.

[5]The work for one trial is comparable to the work for one Lucas test as routinely performed by the GIMPS Project for numbers of this size. A *Lucas test* is a special primality test for Mersenne numbers.

Combination Algorithm

Recall that ECPP produces a certificate of primality. Thus, using a combination of Rabin-Miller and ECPP, we can get a randomized algorithm that produces a certificate to prove that its result (whether "prime" or "composite") is correct.

All we have to do is run the Rabin-Miller and ECPP algorithms in "parallel" until one of them produces a certificate[6]. The expected running time is believed to be $\widetilde{O}(\lambda^4)$, although we can't *prove* this.

To be guaranteed an expected polynomial runtime, add a parallel thread for the Adelman-Huang algorithm.

[6]That is, run both algorithms simultaneously using time-sharing.

Extension of Fermat to Polynomials

If n is prime and a is fixed, then

$$(x+a)^n = x^n + a \mod n \qquad (1)$$

holds, where each side of the equality is a polynomial in x. Formally, we are working in the ring $(\mathbb{Z}/n\mathbb{Z})[x]$ of polynomials whose coefficients are in the ring $\mathbb{Z}/n\mathbb{Z}$ of integers mod n.

Agrawal and Biswas (1999) noticed that (1) (for $a \neq 0 \bmod n$) is necessary and sufficient for the primality of n.

In general, we can not compute $(x+a)^n \mod n$ in time polynomial in λ.

Nevertheless, Agrawal and Biswas were able to give an RP-compositeness test based on (1). The idea (similar to hashing) is to compute each side of (1) modulo some polynomial of low degree. The resulting test is slower than the Rabin-Miller test, so it did not attract much attention at the time.

The AKS Algorithm

In August 2002, Agrawal, Kayal and Saxena announced a *deterministic* polynomial-time primality test based on (1). Thus,

$$\text{PRIMES} \in P.$$

The idea is to compute $(x+a)^n \mod (x^r - 1, n)$ for $1 \leq a \leq s$ and sufficiently large r, s. The not-so-obvious fact is that it is sufficient to choose

$$r = O(\lambda^6), \quad s = O(\lambda^4)\ .$$

Thus, we can do everything in time $\widetilde{O}(\lambda^k)$.

The precise value of the exponent k depends on details of the implementation. In the original AKS paper, $k = 19$ if classical algorithms are used for multiplication and division; $k = 12$ if faster algorithms are used. A revision of the AKS paper has $k = 7.5$. It is conjectured that k can be reduced further. Lenstra and Pomerance have announced $k = 6$ with a variant of AKS based on Gaussian periods.

Example

Take $n = 1729 = 7 \times 13 \times 19$. This is a Carmichael number, so

$$a^n = a \mod n$$

for all integers a. However,

$$(x+1)^n \neq x^n + 1 \mod n\ .$$

In fact, working mod n (i.e. in $(\mathbb{Z}/n\mathbb{Z})[x]$),

$$(x+1)^n = x^{1729} + 247x^{1722} + \cdots + 247x^7 + 1\ .$$

We can more easily verify that $(x+1)^n \neq x^n + 1 \mod n$ by working mod $(x^5 - 1)$ as well as mod n: we find that

$$(x+1)^n = 134x^4 + 1330x^3 + 532x^2 + 1330x + 134$$

in $(\mathbb{Z}/n\mathbb{Z})[x]/(x^5 - 1)$.

Here $x^5 - 1$ acts rather like a hash function: it lets us sum every fifth term in the binomial expansion of $(x+1)^n$, thus reducing $n+1$ terms to five. The computation involves polynomials of degree at most eight.

The Key Theorem

Theorem (AKS-Bernstein-Morain, 2002)

Suppose that $n, r, s > 0$, where r is prime and q is the largest prime factor of $r-1$. Suppose that

$$\binom{q+s-1}{s} > n^{2\lfloor\sqrt{r}\rfloor}\ , \tag{2}$$

that n has no prime factor $\leq s$, and that $n^{(r-1)/q} \mod r \notin \{0,1\}$. Finally, suppose that

$$(x-a)^n = x^n - a \tag{3}$$

in $(Z/nZ)[x]/(x^r-1)$ for $1 \leq a \leq s$. Then n is a prime power.

Remarks

This formulation is given by Morain. In a primality test, after selecting r and s satisfying (2), it takes time $\widetilde{O}(rs\lambda^2)$ to check (3), so we want to minimise rs.

The original AKS algorithm has $r = O(\lambda^6)$, $q \geq 2s$, $s = O(\lambda\sqrt{r}) = O(\lambda^4)$, time $\widetilde{O}(\lambda^{12})$. The proof is "elementary", but too long to give in this lecture.

Reducing the Exponent

An improvement by AKS has $r = O(\lambda^3)$, giving time $\widetilde{O}(\lambda^{7.5})$. However, the proof is no longer entirely elementary, and the constant implied by the $\widetilde{O}$ is ineffective (it exists, but we don't know how to compute it). The best result with elementary techniques and effective constants is $\widetilde{O}(\lambda^{10.5})$.

A *Sophie-Germain prime* is a prime q such that $r = 2q+1$ is also prime, e.g. $q = 11$, $r = 23$. It is conjectured that there are infinitely many Sophie-Germain primes, and that the number up to N is asymptotic to $2C_2N/(\ln N)^2$, where $C_2 \approx 0.66016$ is the Hardy-Littlewood twin-prime constant. Proving this conjecture is probably as difficult as proving the same conjecture for twin primes.

If the Sophie-Germain conjecture is true, then $r = O(\lambda^2)$, $s = O(\sqrt{r}\lambda) = O(\lambda^2)$, and the time bound is reduced to $\widetilde{O}(rs\lambda^2) = \widetilde{O}(\lambda^6)$. This is what Lenstra and Pomerance obtain with their Gaussian-period variant of AKS (independent of Sophie-Germain).

Reducing the Constant Factor

Lenstra & Bernstein (2002) reduced the implicit constant in the time bound of the original AKS paper by a factor of 60,000 (from 1024 to 0.01764).

Constant factors of this order do matter! A factor of 60,000 turns days into seconds[7] (or vice versa) and is much larger than $\log\log\lambda$ ($\log\log 10^8 < 3$).

Bernstein (Jan 2003) claims to have reduced the constant to 0.0005027 (a factor of 2.037×10^6 smaller than the original). No *proof* of this has appeared; nor is there an implementation that can achieve (close to) this speedup over the original AKS algorithm. Thus, the comparisons below disregard this claim.

[7] Approximately, as there are $86,400$ seconds in a day.

Experimental Results

The following table gives some times for a Magma implementation of the AKS algorithm (with Lenstra & Bernstein's improvements of August 2002) on a 1 Ghz Pentium.

Times marked "(est)" are estimated from the time taken for one of the s iterations.

Times marked "(?)" are estimated by extrapolation, assuming the exponent $k = 6$.

n	r	s	time
$10^9 + 7$	43	315	1.0 sec
$10^{19} + 51$	67	5427	750 sec
$10^{49} + 9$	491	28801	32 hours
$10^{100} + 267$	3541	58820	1 year (est)
$2^{511} + 111$			13 years (?)
$2^{1023} + 1155$			840 years (?)

Table 1: The (improved) AKS algorithm

Comparison with Other Algorithms

The following table gives times for Magma implementations[8] of the Rabin-Miller, ECPP and AKS algorithms on a 1 Ghz machine. In all cases the number tested was $10^{100} + 267$.

Algorithm	trials	time
Rabin-Miller	1	0.003 sec
Rabin-Miller	10	0.03 sec
Rabin-Miller	100	0.3 sec
ECPP		2.0 sec
AKS		37 weeks (est)

Table 2: Various algorithms

[8] Our Magma programs may be inefficient, but they are not biased towards any one algorithm. Thanks to Paul Zimmermann for his assistance with the Magma implementation of the AKS algorithm.

Bernstein's Randomised Algorithm

Bernstein (2004) has announced an AKS-like algorithm that finds a certificate of primality in *expected* time $\widetilde{O}(\lambda^2)$. The certificate can be verified in (deterministic) time $\widetilde{O}(\lambda^4)$. However, finding the certificate involves a randomised algorithm, which is contrary to the AKS philosophy of using only deterministic algorithms!

Bernstein has proved $2^{1024} + 643$ prime by his improved (but no longer deterministic) AKS-like algorithm in 13 hours on an 800 Mhz PC. This is *much* faster than the original AKS algorithm. However, it is still slow.

Rabin-Miller takes less than one second, and ECPP takes about 50 seconds (Morain, on a 450 Mhz PC). Thus, ECPP is about 1660 times faster than Bernstein for numbers of this size (ignoring implementation differences).

Reliability of the Result

ECPP gives a certificate of primality, and the certificate can be checked quickly. It *should* be checked, to guard against hardware and/or programming errors. It is conjectured that the expected time to find a certificate is $\widetilde{O}(\lambda^4)$. The time to check a certificate is $\widetilde{O}(\lambda^3)$.

Rabin-Miller takes time $\widetilde{O}(\lambda^2)$ per trial. T trials give probability of error less than 4^{-T} if n is composite; there is no error if n is prime.

Theoretically, AKS is error-free. However, in a long computation there is a significant probability of a hardware error. Such an error would in most cases make a prime n "appear" composite; a composite n would usually still "appear" to be composite.

To be safe, one should make an independent check of the certificate (if available), or use at least two different algorithms.

Summary – Primality Testing and Complexity Classes

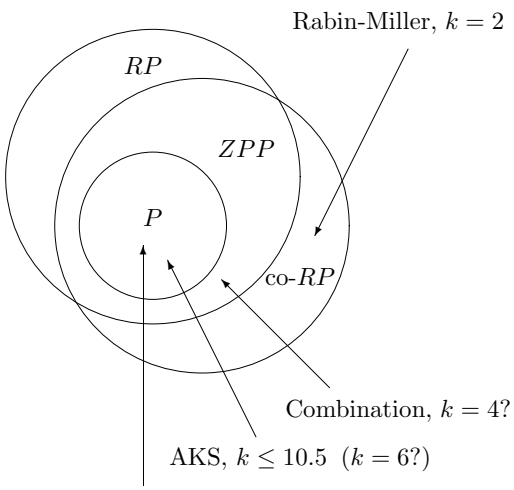

(Conjectured) running time is $\widetilde{O}(\lambda^k)$.

Conclusions

- The AKS algorithm is theoretically significant, since it shows that PRIMES $\in P$.
- AKS is not a practical algorithm. ECPP is *much* faster. Rabin-Miller is even faster, at the price of a *minute* probability of error. A combination of Rabin-Miller and ECPP avoids this possible error, and provides an easily-checked certificate.
- The assumption that problems with polynomial-time algorithms are feasible, and other problems are intractible, is too simplistic.

  In practice, "expected" or "conjectured" or "almost" polynomial algorithms can be better than deterministic polynomial-time algorithms. The exponent k is important.

  For NP-complete problems (not today's topic), polynomial-time *approximation algorithms* might also be relevant.

Last Words

A crude application of theory says that the AKS algorithm is best, but a realistic analysis shows that Rabin-Miller and/or ECPP are much to be preferred. In general, Monte Carlo or Las Vegas algorithms may be better than deterministic polynomial-time algorithms!

I would be more confident in the security of a cryptosystem using large primes certified by Rabin-Miller than (necessarily smaller) primes certified by AKS!

References

[1] L. M. Adleman and M. A. Huang, Primality testing and abelian varieties over finite fields, in *Lecture Notes in Mathematics*, vol. 1512, Springer-Verlag, 1992.

[2] L. M. Adleman, C. Pomerance and R. Rumely, On distinguishing prime numbers from composite numbers, *Ann. of Math.* 117 (1983), 173–206.

[3] M. Agrawal and S. Biswas, Primality and identity testing via chinese remaindering, *Proc. IEEE Symposium on Foundations of Computer Science*, 1999, 202–209.

[4] M. Agrawal, N. Kayal and N. Saxena, PRIMES is in P, preprint, 6 August 2002; also revision of 15 April 2003. Final version in *Annals of Mathematics* 160(2): 781-793, 2004. See `http://www.cse.iitk.ac.in/users/manindra/`

[5] W. Alford, A. Granville and C. Pomerance, There are infinitely many Carmichael numbers, *Ann. of Math.* 139 (1994), 703–722.

[6] A. Atkin and F. Morain, Elliptic curves and primality proving, *Math. Comp.* 61 (1993), 29–68.

[7] D. Bernstein, An exposition of the Agrawal-Kayal-Saxena primality-proving theorem, version of 20 Aug 2002. Superseded by [8].

[8] D. Bernstein, Proving primality after Agrawal-Kayal-Saxena, version of 25 January 2003, available from `http://cr.yp.to/papers.html#aks`

[9] D. Bernstein, Proving primality in essentially quartic expected time, version of 13 February 2004, available from `http://cr.yp.to/papers.html#quartic`

[10] R. Crandall and C. Pomerance, *Prime Numbers: A Computational Perspective*, Springer-Verlag, New York, 2001.

[11] S. Goldwasser and J. Kilian, Primality testing using elliptic curves, *J. ACM* 46 (1999), 450–472. Preliminary version *STOC*, 1986, 316–329.

[12] N. Kayal and N. Saxena, *Towards a Deterministic Polynomial-time Primality Test*, Tech. Report, IIT Kanpur, 2002. Available from `http://www.cse.iitk.ac.in/research/btp-reports.html`

[13] D. E. Knuth, *The Art of Computer Programming*, Vol. 2, 3rd edition, Addison-Wesley, Menlo Park, 1997, §4.5.4.

[14] J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, Elsevier, Amsterdam, 1990.

[15] A. K. Lenstra and H. W. Lenstra, Jr., Algorithms in number theory, in [14], 673–715.

[16] H. W. Lenstra, Jr. and C. Pomerance, Remarks on Agrawal's conjecture, unpublished, 2003. `http://www.aimath.org/WWN/primesinp/articles/html/50a/`

[17] F. Morain, Primalité théorique et primalité pratique ou AKS vs. ECPP, preprint, 4 October 2002. Also other papers and transparencies for séminaire Bourbaki, available from `http://www.lix.polytechnique.fr/Labo/Francois.Morain/`

[18] R. Motwani and P. Raghavan, *Randomized Algorithms*, Cambridge University Press, 1995.

[19] V. Pratt, Every prime has a succinct certificate, *SIAM J. Computing* 4 (1975), 214–220.

[20] M. O. Rabin, "Probabilistic algorithms", in *Algorithms and Complexity* (J. F. Traub, editor), Academic Press, New York, 1976, 21–39.

[21] M. O. Rabin, Complexity of computations (1976 Turing Award Lecture), *Comm. ACM* 20 (1977), 625–633. Corrigendum: *ibid* 21 (1978), 231.

[22] M. O. Rabin, Probabilistic algorithms for testing primality, *J. Number Theory* 12 (1980), 128–138.

[23] R. L. Rivest, A. Shamir and L. Adleman, A method for obtaining digital dignatures and public-key cryptosystems, *Comm. ACM* 21 (1978), 120–126.

[24] R. Solovay and V. Strassen, Fast Monte Carlo test for primality, *SIAM J. on Computing* 6 (1977), 84–85. Erratum: *ibid* 7 (1978), 118.

[25] V. V. Vazirani, *Approximation Algorithms*, Springer-Verlag, Berlin, 2001.