

Mechanisation of AKS Algorithm: Part 1 – the Main Theorem

Hing-Lun Chan¹ and Michael Norrish²

¹ joseph.chan@anu.edu.au
Australian National University

² Michael.Norrish@nicta.com.au
Canberra Research Lab., NICTA*;
also, Australian National University

Abstract. The AKS algorithm (by Agrawal, Kayal and Saxena) is a significant theoretical result proving “PRIMES in P”, as well as a brilliant application of ideas from finite fields. This paper describes the first step towards the goal of a full mechanisation of this result: a mechanisation of the AKS Main Theorem, which justifies the correctness (but not the complexity) of the AKS algorithm.

1 Introduction

The AKS algorithm is a decision procedure for primality testing. That is, given a number n , it returns “true” if n is prime and “false” otherwise. As *per* the title of AKS paper [3], “PRIMES is in P”, the significance of their work is that the number of steps for the verification is bounded by some polynomial function of the size of n , measured by $\log_2 n$.

There have been several attempts to formalize the AKS Main Theorem (see Section 6), but so far none is complete. In this paper, we describe the first complete mechanization of this result. In subsequent work, we aim to demonstrate that the algorithm built on top of this result does indeed compute its answer in polynomial time.

1.1 Overview

A number n is a perfect-power of another number m if there exists an exponent e such that $n = m^e$, and n is power-free if it is a trivial perfect power, *i.e.*, if $n = m^e$ then $e = 1$ and $m = n$. Given a number n , the smallest positive exponent j such that $n^j \equiv 1 \pmod{k}$ is denoted by $\text{order}_k(n)$. Computation in $(\text{mod } n, X^k - 1)$ means that all numerical as well as polynomial computational results are reduced to remainders after divisions by n and by $X^k - 1$. The constant polynomial arising from constant c is denoted by boldface c . More notation will be covered in Section 1.2. Here is a peek at our HOL4 result.

Theorem 1. *The AKS Main Theorem.*

$$\begin{aligned} \vdash \text{prime } n &\iff \\ &1 < n \wedge \text{power_free } n \wedge \\ &\exists k. \\ &\text{prime } k \wedge (2(\log n + 1))^2 \leq \text{order}_k(n) \wedge \\ &(\forall j. 0 < j \wedge j \leq k \wedge j < n \Rightarrow \text{gcd}(n, j) = 1) \wedge \\ &(k < n \Rightarrow \\ &\quad \forall c. \\ &\quad 0 < c \wedge c \leq 2\sqrt{k}(\log n + 1) \Rightarrow \\ &\quad (X + c)^n \equiv (X^n + c) \pmod{n, X^k - 1}) \end{aligned}$$

* NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

This theorem then justifies the following algorithm¹ for primality testing:

Input: integer $n > 1$.

1. If $(n = b^m$ for some base b with $m > 1$), return COMPOSITE.
2. Search for a prime k satisfying $\text{order}_k(n) \geq (2(\log n + 1))^2$.
3. For each $(j = 1$ to $k)$ if $(j = n)$ break, else if $(\text{gcd}(j, n) \neq 1)$, return COMPOSITE.
4. If $(k \geq n)$, return PRIME.
5. For each $(c = 1$ to $2\sqrt{k}(\log n + 1)$) if $(X + c)^n \not\equiv (X^n + c) \pmod{n, X^k - 1}$, return COMPOSITE.
6. return PRIME.

Given a number n , this version of the AKS algorithm requires a search for another prime k in Step 2. Step 4 suggests that it is not always true that $k < n$. Nevertheless, the theorem can still be viewed as a well-founded recursive definition because it turns out that k is roughly bounded by $(\log n)^5$ [3]. So, for sufficiently large values of n , there will always be a $k < n$. For smaller n (effectively the base cases of recursion), a look-up table might be used.

The rest of this paper is devoted to explaining the mechanised proof of this result. Section 2 covers some necessary background. Sections 3 and 4 describe the proof of the AKS Main Theorem. Section 5 discusses our mechanisation experience. Section 6 compares our work with others. Finally, we conclude in Section 7.

1.2 Notation

All statements starting with a turnstile (\vdash) are HOL4 theorems, automatically pretty-printed to \LaTeX from the relevant theory in the HOL4 development. Generally, our notation allows an appealing combination of quantifiers (\forall, \exists), set notation (\in, \cup and comprehensions such as $\{x \mid x < 6\}$), and functional programming (λ for function abstraction, and juxtaposition for function application).

The cardinality of a set s is written $|s|$; the image of set s under the mapping f is written $f(s)$; we write $f : s \hookrightarrow t$ to mean that function f is injective from set s to set t .

Number-theoretic Notation With n a natural number, \sqrt{n} is its square-root, and $\log n$ its logarithm to base 2. Both logarithm and square-root are integer functions, being the floor value of the corresponding real functions. We use $\varphi(n)$ to denote the Euler φ -function of n , the count of coprime numbers less than n . We write $n \mid m$ when n divides m .

For the AKS algorithm, we shall use n for the input number, p for its prime factor, k for the prime (existentially quantified) “parameter” of the Main Theorem, and $\ell = 2\sqrt{k}(\log n + 1)$ for a computed parameter (the limit for a range of constants). Note that $\text{order}_k(n)$ is nonzero whenever $\text{gcd}(k, n) = 1$.

Ring, Field, and Polynomial Notation A ring \mathcal{R} has carrier set R , with $\mathbf{1}$ and $\mathbf{0}$ its one and zero. The characteristic of a ring \mathcal{R} is written as $\chi(\mathcal{R})$, often abbreviated to χ for a generic ring. A ring homomorphism from a ring \mathcal{R} to another ring \mathcal{S} via a map f is denoted by $f : \mathcal{R} \mapsto_r \mathcal{S}$.

We write $\mathcal{R}[X]$ to denote the ring of polynomials with coefficients drawn from the underlying ring \mathcal{R} . Similarly, the ring $\mathcal{F}[X]$ has polynomials with coefficients from a field \mathcal{F} . Polynomials from those rings are written with the sans-serif font, e.g., p, q, h . The constant polynomial c (in bold) is derived from adding $\mathbf{1}$ repeatedly c times. The degree of p is written $\text{deg } p$, its leading coefficient is $\text{lead } p$, and

¹ The constants involved in this algorithm are based on [10, Algorithm 8.2.1]. They are slightly different from those in the AKS papers [2, 3], but such variations do not affect the conclusion of “PRIMES is in P”.

monic p means its leading coefficient is 1. The polynomial X is the monomial (monic of degree 1) with zero constant. The polynomial field quotiented by modulus with an irreducible polynomial h is $\mathcal{F}_h[X]$, with multiplicative group $\mathcal{F}_h^*[X]$.

Arithmetic (addition, subtraction, multiplication, remainders) on polynomials is written with usual symbols ($+$, $-$, $*$, mod etc.), e.g., $X^k - 1$ is the unity polynomial of degree k . Here we can see HOL4's overloading facilities at work: constant polynomials one and zero are $\mathbf{1}$ and $\mathbf{0}$, the same as those for a ring. More aggressively, we use overloading to conceal "implicit" parameters such as the underlying ring \mathcal{R} in terms such as $p * q$ (polynomial product).

We write $p[q]$ to denote the substitution of q for every X in p . We use $\text{roots } p$ for the set of p 's roots, and $\text{ipoly } p$ to mean that p is an irreducible polynomial, both with respect to its underlying ring \mathcal{R} . The quotient ring formed by $\mathcal{R}[X]$ and irreducible polynomial h is denoted by $\mathcal{R}_h[X]$, which can be shown to be a field. Its multiplicative group is $\mathcal{R}_h^*[X]$. The order of an element in this group, e.g., X , is denoted by $\text{order}_h(X)$.

HOLA Sources All our proof scripts can be found at <http://bitbucket.org/jhlchan/hola/src/aks/theories>.

2 Background

A glance at the algorithm in Section 1.1 shows its most prominent feature: polynomial identity tests in modulo $X^k - 1$. To understand this we need to get a feel for the motivation behind the AKS algorithm.

2.1 Finite Fields

The AKS Main Theorem has a setting in finite fields, since the characteristic of a finite field is always prime. A field is also a ring, and a ring with prime characteristic enjoys some wonderful properties.

Theorem 2. *The Freshman-Fermat Theorem*

$$\vdash \text{Ring } \mathcal{R} \wedge \text{prime } \chi \Rightarrow \forall c. (X + c)^\chi = X^\chi + c$$

Proof. This follows directly from two theorems, (a) Freshman's Theorem:

$$\vdash \text{Ring } \mathcal{R} \wedge \text{prime } \chi \Rightarrow \forall p \ q. \text{poly } p \wedge \text{poly } q \Rightarrow (p + q)^\chi = p^\chi + q^\chi$$

and (b) Fermat's Little Theorem for polynomials:

$$\vdash \text{Ring } \mathcal{R} \wedge \text{prime } \chi \Rightarrow \forall c. c^\chi = c$$

Both theorems (a) and (b) have been mechanized in a previous paper by the same authors [6]. \square

The converse, suitably formulated, is also true:

Theorem 3. *A ring has its characteristic prime iff a Freshman-Fermat identity holds for any constant coprime with the characteristic.*

$$\vdash \text{Ring } \mathcal{R} \Rightarrow \forall c. \text{gcd}(c, \chi) = 1 \Rightarrow (\text{prime } \chi \iff 1 < \chi \wedge (X + c)^\chi = X^\chi + c)$$

Given a number $n > 1$, we can identify \mathcal{R} with \mathbb{Z}_n , and $\chi(\mathbb{Z}_n) = n$. Since $\text{gcd}(1, n) = 1$ always, pick $c = 1$, then this theorem applies, and whether n is prime is just one check of a Freshman-Fermat polynomial identity in \mathbb{Z}_n , i.e., $(\text{mod } n)$.

Therefore, this theorem amounts to a deterministic primality test. But there is a catch: the left-side, upon expansion, contains $n + 1$ terms. Thus this is an impractical primality test for large values of n .

The AKS idea begins with checking such Freshman-Fermat identities, with two twists:

- Instead of just checking in $(\text{mod } n)$, perform the computations in $(\text{mod } n, X^k - 1)$ for a suitably chosen parameter k . Since results are always the remainder after division by $X^k - 1$, the degree of intermediate polynomials (which determines the number of terms) never exceed k —presumably k is much smaller than n .
- Instead of checking just one coprime value, check for a range of coprime values c , for $0 < c \leq \ell$, up to some maximum limit ℓ —presumably ℓ is related to k , and small compared to n .

Of course, the big question is whether after such modifications there is still a primality test. The AKS team answered this in the affirmative—there exist parameters k and ℓ , bounded by some polynomial function of the size of input number n , *i.e.*, $\log n$, giving a polynomial-time deterministic primality test.

2.2 Introspective Relation

Recall from Section 1.1 that AKS computations are done in $(\text{mod } n, X^k - 1)$. This double modulo notation is clumsy. Let us work in a generic ring \mathcal{R} , later to be identified with instances such as \mathbb{Z}_n . The first $(\text{mod } n)$ equivalence becomes equality in the ring \mathcal{R} (*i.e.*, $x \equiv y \pmod{n}$ means $x = y$ in \mathbb{Z}_n); leaving the symbol (\equiv) to indicate the polynomial modulo equivalence in $\mathcal{R}[X]$.

In this context, of a general ring \mathcal{R} , the polynomial identity checks in Theorem 1 take the form:

$$(X + c)^n \equiv X^n + c \pmod{X^k - 1}$$

They look like Freshman-Fermat identities of Theorem 2, only now under modulo by a polynomial. Rewriting with a polynomial substitution, the left and right sides are strikingly similar:

$$(X + c)^n \llbracket X \rrbracket \equiv (X + c) \llbracket X^n \rrbracket \pmod{X^k - 1}$$

The rewrites are trivial, since for any polynomial p , we have $p \llbracket X \rrbracket = p$ and $(X + c) \llbracket p \rrbracket = p + c$. Superficially, the left-hand side is transformed into the right-hand side simply by shifting of the exponent n . Following the terminology of AKS team, we say n is *introspective* to polynomial p , denoted by $n \bowtie p$, when:

$$\vdash n \bowtie p \iff \text{poly } p \wedge 0 < k \wedge p^n \equiv p \llbracket X^n \rrbracket \pmod{X^k - 1}$$

Note that the symbol for introspective relation (\bowtie) hides the polynomial modulus $X^k - 1$, and the underlying ring \mathcal{R} . We shall include a subscript when the underlying ring is of significance, *e.g.*, $\bowtie_{\mathbb{Z}_n}$.

Therefore, the AKS algorithm verifies, for the input number n , the identities $n \bowtie X + c$ in \mathbb{Z}_n for $0 < c \leq \ell$ up to some maximum ℓ . Moreover, Freshman-Fermat (Theorem 2) can be restated as:

Theorem 4. *For a ring with prime characteristic, its characteristic is introspective to any monomial.*

$$\vdash \text{Ring } \mathcal{R} \wedge 1 \neq 0 \wedge \text{prime } \chi \Rightarrow \forall k. 0 < k \Rightarrow \forall c. \chi \bowtie X + c$$

Proof. By introspective definition, this is to show: $(X + c)^X \equiv (X + c) \llbracket X^X \rrbracket \pmod{X^k - 1}$. Transforming the right side by substitution, $(X + c) \llbracket X^X \rrbracket = X^X + c$. Then both sides are equal by Freshman-Fermat (Theorem 2), hence they are equivalent under modulo by $X^k - 1$. \square

The fundamental properties of introspective relation are:

Theorem 5. *Introspective relation is multiplicative for exponents.*

$$\vdash \text{Ring } \mathcal{R} \wedge 1 \neq 0 \Rightarrow \forall k \ p \ n \ m. n \bowtie p \wedge m \bowtie p \Rightarrow n m \bowtie p$$

Proof. Working in $(\text{mod } X^k - 1)$, we have $p^n \equiv p[X^n]$ by $n \bowtie p$, and $p^m \equiv p[X^m]$ by $m \bowtie p$. The latter means $p[X]^m - p[X^m]$ is divisible by $X^k - 1$. Substitute every X of the previous statement by X^n , and noting $X^k - 1 \mid (X^n)^k - 1$ by divisibility of unity polynomials, $p[X^n]^m \equiv p[(X^n)^m]$. Therefore, $p^{n \cdot m} = (p^n)^m \equiv p[X^n]^m \equiv p[(X^n)^m] = p[X^{n \cdot m}]$, or $n \cdot m \bowtie p$. \square

Theorem 6. *Introspective relation is multiplicative for polynomials.*

$$\vdash \text{Ring } \mathcal{R} \wedge 1 \neq 0 \Rightarrow \forall k \ p \ q \ n. \ n \bowtie p \wedge n \bowtie q \Rightarrow n \bowtie p * q$$

Proof. Working in $(\text{mod } X^k - 1)$, we have $p^n \equiv p[X^n]$ by $n \bowtie p$, and $q^n \equiv q[X^n]$ by $n \bowtie q$. Therefore, $(p * q)^n = p^n * q^n \equiv p[X^n] * q[X^n] = (p * q)[X^n]$, or $n \bowtie p * q$. \square

3 Main Theorem

We can now restate the AKS Main Theorem (Theorem 1) in terms of the introspective relation.

Theorem 7. *A number is prime iff it satisfies all the AKS checks.*

$$\begin{aligned} \vdash \text{prime } n &\iff \\ &1 < n \wedge \text{power_free } n \wedge \\ &\exists k. \\ &\text{prime } k \wedge (2 (\log n + 1))^2 \leq \text{order}_k(n) \wedge \\ &(\forall j. 0 < j \wedge j \leq k \wedge j < n \Rightarrow \text{gcd}(n, j) = 1) \wedge \\ &(k < n \Rightarrow \\ &\quad \forall c. 0 < c \wedge c \leq 2\sqrt{k} (\log n + 1) \Rightarrow n \bowtie_{\mathbb{Z}_n} X + c) \end{aligned}$$

Note how the symbol $\bowtie_{\mathbb{Z}_n}$ encapsulates the introspective relation (*i.e.*, $\text{mod } X^k - 1$) within \mathbb{Z}_n (*i.e.*, $\text{mod } n$), the double modulo in the AKS computations. We prove this logical equivalence in two parts.

3.1 Easy Part (\Rightarrow)

Theorem 8. *The if-part of AKS Main Theorem (Theorem 7).*

$$\begin{aligned} \vdash \text{prime } n &\Rightarrow \\ &1 < n \wedge \text{power_free } n \wedge \\ &\exists k. \\ &\text{prime } k \wedge (2 (\log n + 1))^2 \leq \text{order}_k(n) \wedge \\ &(\forall j. 0 < j \wedge j \leq k \wedge j < n \Rightarrow \text{gcd}(n, j) = 1) \wedge \\ &(k < n \Rightarrow \forall c. 0 < c \wedge c \leq 2\sqrt{k} (\log n + 1) \Rightarrow n \bowtie_{\mathbb{Z}_n} X + c) \end{aligned}$$

Proof. The first two goals, $1 < n$ and $\text{power_free } n$, are trivial for a prime n . For the third goal, let $m = (2 (\log n + 1))^2$, then parameter k exists by Theorem 25 in Section 4.6:

$$\vdash 1 < n \wedge 0 < m \Rightarrow \exists k. \text{prime } k \wedge \text{gcd}(k, n) = 1 \wedge m \leq \text{order}_k(n)$$

If $k \geq n$, the coprime checks are subsumed by $\forall j. 0 < j \wedge j < n \Rightarrow \text{gcd}(n, j) = 1$. Otherwise $k < n$, and the coprime checks are subsumed by $\forall j. 0 < j \wedge j \leq k \Rightarrow \text{gcd}(n, j) = 1$. Either way this is true since a prime n is coprime with all values less than itself. When $k < n$, the last check is established by Theorem 4, since a prime n gives a field \mathbb{Z}_n , with $\chi(\mathbb{Z}_n) = n$. \square

A close equivalent of this Theorem 8 was mechanised by de Moura and Tadeu [9] in Coq, and by Campos *et al* [5] in ACL2.

3.2 Hard Part (\Leftarrow)

Theorem 9. *The only-if part of AKS Main Theorem (Theorem 7).*

$$\begin{aligned} & \vdash 1 < n \wedge \text{power_free } n \wedge \\ & (\exists k. \\ & \quad \text{prime } k \wedge (2(\log n + 1))^2 \leq \text{order}_k(n) \wedge \\ & \quad (\forall j. 0 < j \wedge j \leq k \wedge j < n \Rightarrow \text{gcd}(n, j) = 1) \wedge \\ & \quad (k < n \Rightarrow \forall c. 0 < c \wedge c \leq 2\sqrt{k}(\log n + 1) \Rightarrow n \not\equiv_{\mathbb{Z}_n} X + c) \Rightarrow \\ & \quad \text{prime } n \end{aligned}$$

Proof. Based on the given parameter k , let $\ell = 2\sqrt{k}(\log n + 1)$. If $k \geq n$ the coprime checks will verify $\forall j. 0 < j \wedge j < n \Rightarrow \text{gcd}(n, j) = 1$, thus n will be prime since it has no proper factor. Otherwise $k < n$, the coprime checks are $\forall j. 0 < j \wedge j \leq k \Rightarrow \text{gcd}(n, j) = 1$. In Section 3.3 we shall establish:

Theorem 10. *The AKS Main Theorem in \mathbb{Z}_n .*

$$\begin{aligned} & \vdash 1 < n \Rightarrow \\ & \quad \forall k \ell. \\ & \quad \text{prime } k \wedge (2(\log n + 1))^2 \leq \text{order}_k(n) \wedge \ell = 2\sqrt{k}(\log n + 1) \wedge \\ & \quad (\forall j. 0 < j \wedge j \leq k \Rightarrow \text{gcd}(n, j) = 1) \wedge \\ & \quad (\forall c. 0 < c \wedge c \leq \ell \Rightarrow n \not\equiv_{\mathbb{Z}_n} X + c) \Rightarrow \\ & \quad \exists p. \text{prime } p \wedge \text{perfect_power } n \ p \end{aligned}$$

Applying this theorem, $n = p^e$ for some prime p and exponent e by definition of a perfect power. But n is assumed power-free, so $e = 1$ and $n = p$, making n a prime. \square

3.3 Shifting Playgrounds

The AKS verifications take polynomials with coefficients from \mathbb{Z}_n , a ring for general n . Polynomials with coefficients from a ring can have more roots than their degree, due to the possible existence of zero divisors in a ring.² A field has no zero divisors, and polynomials with coefficients from a field have this nice property:

$$\vdash \text{Field } \mathcal{F} \Rightarrow \forall p. \text{poly } p \wedge p \neq 0 \Rightarrow |\text{roots } p| \leq \text{deg } p$$

As we shall see (Sections 4.3 and 4.4), there will be two important injective maps involved in the AKS proof. To establish the injective property, this restriction on the number of polynomial roots by its degree is of utmost importance.

But where to find a field \mathcal{F} to work with, given that we start in the ring \mathbb{Z}_n ?

When the number n is not 1, it must have a prime factor p . This leads to the field \mathbb{Z}_p . If relationships between monomials $X + c$ are carried over unaffected from $\mathbb{Z}_n[X]$ to $\mathbb{Z}_p[X]$, we are in a better place to investigate the nature of n . This shifting of playgrounds is essential in the proof of Theorem 10:

Proof (of Theorem 10).

If n is prime, take $p = n$. Otherwise, n has a proper prime factor p such that $p < n$ and $p \mid n$. Introduce two rings, \mathbb{Z}_n and \mathbb{Z}_p . The latter ring \mathbb{Z}_p is also a field, in fact a finite field. This is because all nonzero elements are coprime to the prime modulus p , hence they have inverses.

There is a homomorphism between these two rings due to that fact that p divides n :

² For example, in \mathbb{Z}_6 , $2 \times 3 = 0$, hence $(X - 2)(X - 3) = X^2 - 5X = X(X - 5)$, which shows a polynomial of degree 2 can have more than 2 roots.

$$\vdash 0 < n \wedge 0 < p \wedge p \mid n \Rightarrow (\lambda x. x \bmod p) : \mathbb{Z}_n \mapsto_r \mathbb{Z}_p$$

This ring homomorphism will preserve monomials $X + c$ if a condition on limit ℓ is met:

$$\vdash 0 < n \wedge 1 < p \wedge \ell < p \Rightarrow \\ \forall c. 0 < c \wedge c \leq \ell \Rightarrow \forall f. f : \mathbb{Z}_n \mapsto_r \mathbb{Z}_p \Rightarrow f(X + c) = X + c$$

Here $f(p)$ denotes applying the ring homomorphism f to each coefficient of a polynomial p . We shall show in Section 4.6 that $\ell \leq k$ (Theorem 27). To meet the condition $\ell < p$, we need only to show $k < p$. Note that the given coprime checks on k are (from the statement of Theorem 10):

$$\forall j. 0 < j \wedge j \leq k \Rightarrow \text{gcd}(n, j) = 1$$

Taking $j = k$, we conclude $\text{gcd}(n, k) = 1$. This will be useful later. Apply the following theorems:

$$\vdash 1 < n \wedge \text{prime } p \wedge p \mid n \Rightarrow \forall j. \text{gcd}(n, j) = 1 \Rightarrow \text{gcd}(p, j) = 1 \\ \vdash 1 < p \Rightarrow \forall k. (\forall j. 0 < j \wedge j \leq k \Rightarrow \text{gcd}(p, j) = 1) \Rightarrow k < p$$

Tracing the transformation of gcd's gives $k < p$, hence $\ell < p$.

Therefore the monomials are preserved by homomorphism, together with the introspective relation:

$$\vdash 0 < n \wedge 1 < p \wedge p \mid n \wedge 0 < k \wedge \ell < p \Rightarrow \\ \forall m c. 0 < c \wedge c \leq \ell \wedge m \bowtie_{\mathbb{Z}_n} X + c \Rightarrow m \bowtie_{\mathbb{Z}_p} X + c$$

Thus the AKS checks in \mathbb{Z}_n are equivalent to checks in \mathbb{Z}_p , a finite field, where p is a prime factor of n . Generalising to arbitrary finite fields, in Section 4.5 we will prove:

Theorem 11. *AKS Main Theorem in finite fields.*

$$\vdash \text{FiniteField } \mathcal{F} \wedge \text{prime } k \wedge k < \chi \Rightarrow \\ \forall n. \\ 1 < n \wedge \chi \mid n \wedge \text{gcd}(n, k) = 1 \wedge (2(\log n + 1))^2 \leq \text{order}_k(n) \wedge \\ \ell = 2\sqrt{k}(\log n + 1) \wedge (\forall c. 0 < c \wedge c \leq \ell \Rightarrow n \bowtie X + c) \Rightarrow \\ \text{perfect_power } n \chi$$

We then identify \mathcal{F} with \mathbb{Z}_p , noting $\chi(\mathbb{Z}_p) = p$, with $k < p$. Knowing $\text{gcd}(n, k) = 1$ from the gcd checks above, we conclude that n must be a perfect power of its prime factor p , as required. \square

4 Introspective Game

There are two useful facts when working in the context of a finite field \mathcal{F} , where χ is necessarily prime:

- We get, for free in $\mathcal{F}[X]$, the result: $\chi \bowtie X + c$, by Theorem 4, since a field is a non-trivial ring.
- The modulus polynomial $X^k - 1$, now in $\mathcal{F}[X]$, will have a monic irreducible factor $h \neq X - 1$.

Both will play significant roles in the proof of Theorem 11. Here are the highlights:

- The finite field \mathcal{F} will enrich the introspective relation, through the interplay between prime χ and n .
- This will give rise to some interesting sets, among them are two finite sets $\widehat{\mathcal{N}}$ and \mathcal{M} (Section 4.2).
- The conditions on parameters k and ℓ will establish an injective map from $\widehat{\mathcal{N}}$ to \mathcal{M} .
- If n were not a perfect power of χ , then we would have $|\widehat{\mathcal{N}}| > |\mathcal{M}|$, contradicting the Pigeonhole Principle.

Summary of the AKS proof (Theorem 11)

Our strategy for the AKS proof can be described as a game between two players (see Figure 1). The introspective relations of n and p , a prime factor of n , give rise to two sets \mathcal{N} and \mathcal{P} (Section 4.1). Taking modulo by k (an AKS parameter) and by h (an irreducible factor of $X^k - 1$), the sets \mathcal{N} and \mathcal{P} map (straight arrows), respectively, to two finite sets \mathcal{M} and \mathcal{Q}_h (Section 4.2). Two finite subsets of \mathcal{N} and \mathcal{P} , shown as $\hat{\mathcal{N}}$ and $\hat{\mathcal{P}}$, can be crafted in such a way that injective maps (curve arrows) between the finite sets can be constructed, if k and ℓ (another AKS parameter) are suitably chosen to satisfy the “if” conditions (Section 4.3 and Section 4.4). The construction of injective maps involves interactions (dashed arrows) between the two players, based on properties of the introspective relation and polynomials in $\mathcal{F}_h[X]$. Once these are all in place, if n were not a perfect power of p , the grey set $\hat{\mathcal{N}}$ will have more than $|\mathcal{M}|$ elements, where \mathcal{M} is the target of the left injective map. This contradicts the Pigeonhole Principle (Section 4.5). Hence n must be a perfect power of its prime factor p .

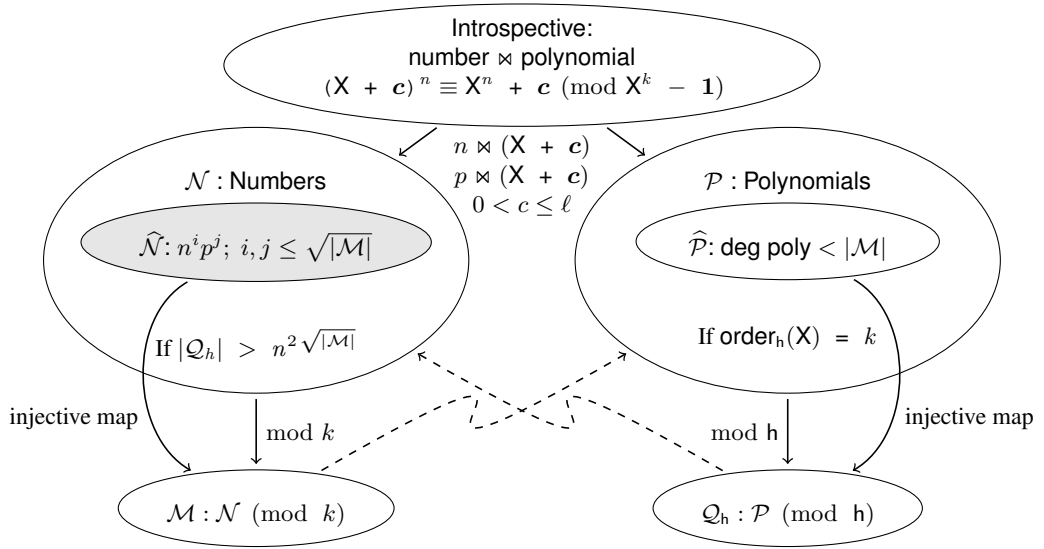


Fig. 1. The AKS proof as a game between numbers and polynomials *via* introspective relation. Refer to summary above for an explanation.

4.1 Introspective Sets

As noted above, after shifting to a finite field \mathcal{F} where $p = \chi$ is prime, for the constants $0 < c \leq \ell$, besides the given $n \bowtie X + c$, we also have $p \bowtie X + c$ by Theorem 4.

In view of this, we define the following two sets:

$$\begin{aligned} \vdash \mathcal{N} &= \{m \mid \text{gcd}(m, k) = 1 \wedge \forall c. 0 < c \wedge c \leq \ell \Rightarrow m \bowtie X + c\} \\ \vdash \mathcal{P} &= \{p \mid \text{poly } p \wedge \forall m. m \in \mathcal{N} \Rightarrow m \bowtie p\} \end{aligned}$$

The set \mathcal{N} captures the introspective exponents. Observe that $n \in \mathcal{N}$, $p \in \mathcal{N}$, and trivially, $1 \in \mathcal{N}$. They are all coprime to k , since the coprime checks in Section 3.3 give $\text{gcd}(n, k) = 1$ and $k < p$. For a prime p , $k < p$ gives $\text{gcd}(p, k) = 1$.

The set \mathcal{P} captures the introspective polynomials, those with introspective exponents in \mathcal{N} . Certainly $\forall c. 0 < c \wedge c \leq \ell \Rightarrow X + c \in \mathcal{P}$, and trivially, $1 \in \mathcal{P}$.

Recall the fundamental properties of introspective relation: there will be multiplicative exponents for \mathcal{N} (Theorem 5) and multiplicative polynomials for \mathcal{P} (Theorem 6). Together they imply that the sets \mathcal{N} and \mathcal{P} will be infinitely large. Our contradiction from the Pigeonhole Principle comes when we have derived some related, and finite sets.

4.2 Modulo Sets

One way to get a finite counterpart from an infinite set is by looking at remainders after division, or image of the set under some modulus. For the exponents set \mathcal{N} , the parameter k provides a modulus:

$$\vdash \mathcal{M} = (\lambda m. m \bmod k) (\mathcal{N})$$

It is easy to estimate the cardinality of \mathcal{M} :

Theorem 12. *The cardinality of set \mathcal{M} is bounded by k and $\text{order}_k(n)$.*

$$\vdash \text{Ring } \mathcal{R} \wedge 1 \neq 0 \wedge 1 < k \Rightarrow \forall n. n \in \mathcal{N} \Rightarrow \text{order}_k(n) \leq |\mathcal{M}| \wedge |\mathcal{M}| < k$$

Proof. Since there are k remainders under modulo k , $|\mathcal{M}| \leq k$. But multiples of k (those n with $n \bmod k = 0$) are not in \mathcal{N} , as all elements of \mathcal{N} are coprime to k and $k \neq 1$. Therefore $0 \notin \mathcal{M}$, making $|\mathcal{M}| < k$. Given $n \in \mathcal{N}$, so are all its powers: $\forall j. n^j \in \mathcal{N}$ by Theorem 5. Hence all the remainders $n^j \bmod k$ are in \mathcal{M} . Since $\text{order}_k(n)$ is the minimal exponent j before such remainders wrap around to 1, there are at least $\text{order}_k(n)$ distinct remainders. Thus $\text{order}_k(n) \leq |\mathcal{M}|$. \square

For the polynomials set \mathcal{P} , the irreducible factor h of $X^k - 1$ provides a modulus:

$$\vdash \mathcal{Q}_h = (\lambda p. p \bmod h) (\mathcal{P})$$

For the cardinality of \mathcal{Q}_h , estimation requires more work, due to the change of modulus to h . Let $z = X^k - 1$, then monic z and $\deg z = k$. Note that $z \equiv 0 \pmod{h}$, since h divides z by being a factor. These facts ensure that polynomial equivalences in $(\bmod z)$ are preserved to $(\bmod h)$:

Theorem 13. *Polynomial modulo equivalence holds for modulus factor.*

$$\vdash \text{Ring } \mathcal{R} \wedge \text{monic } z \wedge 0 < \deg z \wedge \text{monic } h \wedge 0 < \deg h \wedge z \equiv 0 \pmod{h} \Rightarrow \\ \forall p q. \text{poly } p \wedge \text{poly } q \wedge p \equiv q \pmod{z} \Rightarrow p \equiv q \pmod{h}$$

Proof. When $(p - q)$ is divisible by z (due to $p \equiv q \pmod{z}$), and z is divisible by h (due to $z \equiv 0 \pmod{h}$), the difference $(p - q)$ is also divisible by h due to transitivity of division. \square

An irreducible polynomial h gives a polynomial modulo field $\mathcal{F}_h[X]$, and nonzero elements of a field form a multiplicative group. Since $X \neq 0$, it has a nonzero $\text{order}_h(X)$, with the following features.

Theorem 14. *When X is a root of unity, order of X is maximal when unity exponent is prime.*

$$\vdash \text{FiniteField } \mathcal{F} \wedge \text{monic } h \wedge \text{ipoly } h \wedge h \neq X - 1 \Rightarrow \\ \forall k. \text{prime } k \wedge X^k \equiv 1 \pmod{h} \Rightarrow \text{order}_h(X) = k$$

Proof. Let $t = \text{order}_h(X)$. By definition of order, $X^t \equiv 1 \pmod{h}$, and given $X^k \equiv 1 \pmod{h}$. Since t is minimal, t divides k . Given prime k , $t = 1$ or $t = k$. Only 1 has order 1, but $X \not\equiv 1 \pmod{h}$ by assumption. Therefore $\text{order}_h(X) = t = k$. \square

Theorem 15. *In the polynomial field $\mathcal{F}_h[X]$, powers of X are distinct for exponents less than $\text{order}_h(X)$.*

$$\begin{aligned} \vdash \text{FiniteField } \mathcal{F} \wedge \text{monic } h \wedge \text{ipoly } h \wedge h \neq X \Rightarrow \\ \forall m \ n. \ m < \text{order}_h(X) \wedge n < \text{order}_h(X) \Rightarrow (X^m \equiv X^n \pmod{h}) \iff m = n \end{aligned}$$

Proof. Since $\mathcal{F}_h[X]$ is a finite field, its multiplicative group is a finite group. By the given assumption, $X \not\equiv 0 \pmod{h}$, thus X is an element in this group. Its order is the minimal exponent for the powers of X to wrap around to 1 . Given the exponents are less than its order, such powers of X are distinct. \square

We shall see how the distinct powers of X lead to a lower bound for \mathcal{Q}_h . This simple result is helpful:

Theorem 16. *Powers of X are equivalent in $X^k - 1$ if exponents are equivalent in \mathbb{Z}_k .*

$$\vdash \text{Ring } \mathcal{R} \wedge 1 \neq 0 \Rightarrow \forall k. \ 0 < k \Rightarrow \forall m. \ X^m \equiv X^{m \bmod k} \pmod{X^k - 1}$$

Proof. Since $m = (m \text{ div } k)k + m \bmod k$ and $X^k \equiv 1 \pmod{X^k - 1}$, the result follows. \square

4.3 Reduced Polynomials

Referring to Figure 1, we shall see eventually that the right injective map is essential to give a lower bound for \mathcal{Q}_h , and this lower bound is essential to provide the left injective map. These two injective maps are critical in the AKS proof.

To obtain a lower bound for \mathcal{Q}_h , we need another way to get something finite from the infinite set \mathcal{P} , by taking a reduced subset of \mathcal{P} :

$$\vdash \hat{\mathcal{P}} = \{p \mid p \in \mathcal{P} \wedge \text{deg } p < |\mathcal{M}|\}$$

This is a finite subset of \mathcal{P} due to the polynomial degree cut-off. We shall prove that there is an injective map from $\hat{\mathcal{P}}$ to \mathcal{Q}_h , hence a lower bound on $|\hat{\mathcal{P}}|$ will also be a lower bound for $|\mathcal{Q}_h|$.

First, note an interesting interaction from \mathcal{M} to \mathcal{P} , which is relevant to $\hat{\mathcal{P}}$ since $\hat{\mathcal{P}} \subseteq \mathcal{P}$. We know that \mathcal{P} has a lot of elements (Section 4.1), so pick two polynomials $p \in \mathcal{P}$ and $q \in \mathcal{P}$. While it is almost impossible to identify any roots for p or q in $\mathcal{F}_h[X]$, it turns out that introspective relation helps to identify some interesting roots of their difference $(p - q)$, from the elements of \mathcal{M} .

Theorem 17. *Each element in \mathcal{M} gives a root for any difference polynomial from \mathcal{P} in $\mathcal{F}_h[X]$.*

$$\begin{aligned} \vdash \text{Field } \mathcal{F} \wedge \text{monic } h \wedge 0 < \text{deg } h \wedge X^k - 1 \equiv 0 \pmod{h} \Rightarrow \\ \forall p \ q. \\ p \in \mathcal{P} \wedge q \in \mathcal{P} \wedge p \equiv q \pmod{h} \Rightarrow \\ \forall n. \ n \in \mathcal{M} \Rightarrow (p - q)[X^n] \equiv 0 \pmod{h} \end{aligned}$$

Proof. Given $n \in \mathcal{M}$, there is $m \in \mathcal{N}$ such that $n = m \bmod k$. Therefore $m \times p$ and $m \times q$ by definition of \mathcal{P} . Let $z = X^k - 1$. Note that $z \equiv 0 \pmod{h}$ by assumption. We can proceed:

	$p^m \equiv p[X^m] \pmod{z}$	by $m \times p$
and	$p[X^m] \equiv p[X^n] \pmod{z}$	by Theorem 16
so	$p^m \equiv p[X^n] \pmod{z}$	by transitivity
or for p , by $z \equiv 0 \pmod{h}$	$p^m \equiv p[X^n] \pmod{h}$	by Theorem 13—[1]
Repeat the same steps for q	$q^m \equiv q[X^n] \pmod{h}$	by $m \times q$ etc.—[2]
Since	$p^m \equiv q^m \pmod{h}$	by $p \equiv q \pmod{h}$ given
so	$p[X^n] \equiv q[X^n] \pmod{h}$	by [1] and [2] above
or	$(p - q)[X^n] \equiv 0 \pmod{h}$	as claimed.

\square

Due to this, an injective map between the two finite sets derived from \mathcal{P} is possible:

Theorem 18. *There is an injective map from reduced set of \mathcal{P} to modulo set of \mathcal{P} .*

$$\begin{aligned} \vdash \text{FiniteField } \mathcal{F} \wedge 0 < k \wedge \text{monic } h \wedge \text{ipoly } h \wedge X^k - 1 \equiv \mathbf{0} \pmod{h} \wedge \\ k = \text{order}_h(X) \Rightarrow \\ (\lambda p. p \pmod{h}) : \widehat{\mathcal{P}} \hookrightarrow \mathcal{Q}_h \end{aligned}$$

Proof. Let $p, q \in \widehat{\mathcal{P}}$, with $p \equiv q \pmod{h}$ in \mathcal{Q}_h . For our map to be injective, we need to show $p = q$. Since $\widehat{\mathcal{P}} \subseteq \mathcal{P}$, $p, q \in \mathcal{P}$. Theorem 17 applies: each $n \in \mathcal{M}$ gives a root X^n for $(p - q)$. Now $h \neq X$ because, by assumption, $h \mid X^k - 1$, but $X \nmid X^k - 1$, and $n < k$ since $n \in \mathcal{M}$ means n is a remainder in $(\text{mod } k)$. By assumption, $k = \text{order}_h(X)$, hence these roots are distinct by Theorem 15. Thus there are at least $|\mathcal{M}|$ distinct roots for $(p - q)$.

But $\deg p < |\mathcal{M}|$ and $\deg q < |\mathcal{M}|$ since $p, q \in \widehat{\mathcal{P}}$, hence $\deg (p - q) < |\mathcal{M}|$. There are more roots than its degree for the difference $(p - q)$ with coefficients from a finite field \mathcal{F} . This is possible only when the difference is $\mathbf{0}$, i.e., $p = q$. \square

This injective map leads to a lower bound for the cardinality of \mathcal{Q}_h .

Theorem 19. *The modulo set of \mathcal{P} has a nice lower bound.*

$$\begin{aligned} \vdash \text{FiniteField } \mathcal{F} \wedge 1 < k \wedge k = \text{order}_h(X) \wedge \ell < \chi \wedge \text{monic } h \wedge \text{ipoly } h \wedge \\ X^k - 1 \equiv \mathbf{0} \pmod{h} \Rightarrow \\ 2^{\min(\ell, |\mathcal{M}|)} \leq |\mathcal{Q}_h| \end{aligned}$$

Proof. Applying Theorem 18, there is an injective map from $\widehat{\mathcal{P}}$ to \mathcal{Q}_h . As both sets are finite, $|\widehat{\mathcal{P}}| \leq |\mathcal{Q}_h|$. We shall estimate $|\widehat{\mathcal{P}}|$, by counting how many polynomials $p \in \mathcal{P}$ have $\deg p < |\mathcal{M}|$.

Note that $1 < |\mathcal{M}|$, since $\text{order}_k(n) \leq |\mathcal{M}|$ by Theorem 12, and $1 < \text{order}_k(n)$ since $n \neq 1$. A simple estimate for $|\widehat{\mathcal{P}}|$ proceeds as follows:

- For $0 < c \leq \ell$, $X + c \in \widehat{\mathcal{P}}$, since each monomial is in \mathcal{P} , and each has a degree equal to 1.
- Given $\ell < \chi$, these monomials are distinct, as χ is the least additive wrap-around of $\mathbf{1}$ in field \mathcal{F} .³
- By Theorem 6, any product of these monomials will be in $\widehat{\mathcal{P}}$, if the product has a degree less than $|\mathcal{M}|$.
- If $\ell < |\mathcal{M}|$, there are less than $|\mathcal{M}|$ such monomials. Therefore any product drawn from a subset of $\{X + c \mid 0 < c \leq \ell\}$ will have a degree less than $|\mathcal{M}|$. There are 2^ℓ such products.
- If $|\mathcal{M}| \leq \ell$, reduce the constants range to $0 < c \leq |\mathcal{M}|$. Any product drawn from a subset of $\{X + c \mid 0 < c \leq |\mathcal{M}|\}$ will have a degree less than $|\mathcal{M}|$, almost—the product of all such monomials must be excluded. However, $\mathbf{1} \in \widehat{\mathcal{P}}$, but $\mathbf{1}$ is not a monomial product. There are still $2^{|\mathcal{M}|}$ products.

Considering both cases, we conclude that $2^{\min(\ell, |\mathcal{M}|)} \leq |\mathcal{Q}_h|$. \square

4.4 Reduced Exponents

It turns out that an injective map to \mathcal{M} is possible based on the following set of reduced exponents:

$$\vdash \widehat{\mathcal{N}} \ p \ n \ m = \{p^i n^j \mid i \leq m \wedge j \leq m\}$$

This is generated by the two known elements $n, p \in \mathcal{N}$ (Section 4.1), with cut-off m in their exponents. By multiplicative closure of introspective exponents (Theorem 5), $\widehat{\mathcal{N}} \subseteq \mathcal{N}$. Observe the following property:

Theorem 20. *Upper bound of an element in $\widehat{\mathcal{N}} \ p \ n \ m$.*

$$\vdash 1 < p \wedge p \leq n \Rightarrow \forall e \ m. e \in \widehat{\mathcal{N}} \ p \ n \ m \Rightarrow e \leq n^{2m}$$

³ The characteristic χ of a ring \mathcal{R} is defined as the order of $\mathbf{1}$ in the additive group of \mathcal{R} , i.e., $\chi \mathbf{1} = \mathbf{0}$.

Proof. Each $e \in \widehat{\mathcal{N}} \setminus \mathcal{P}$ has the form $p^i n^j$, where $i, j \leq m$. Given $p \leq n$, we can deduce $e = p^i n^j \leq n^i n^j \leq n^m n^m = n^{2m}$. \square

Note another interesting interaction from \mathcal{Q}_h to \mathcal{N} , which is relevant to $\widehat{\mathcal{N}}$ since $\widehat{\mathcal{N}} \setminus \mathcal{P} \subseteq \mathcal{N}$. Pick two exponents $n \in \mathcal{N}$ and $m \in \mathcal{N}$. Consider the special polynomial $X^n - X^m$. It turns out that the introspective relation helps to identify some interesting roots of this special polynomial, from the elements of \mathcal{Q}_h .

Theorem 21. *Each element in \mathcal{Q}_h gives a root for a special polynomial from \mathcal{N} in $\mathcal{F}_h[X]$.*

$$\begin{aligned} \vdash \text{Field } \mathcal{F} \wedge \text{monic } h \wedge \text{ipoly } h \wedge X^k - 1 \equiv \mathbf{0} \pmod{h} &\Rightarrow \\ \forall n, m. & \\ n \in \mathcal{N} \wedge m \in \mathcal{N} \wedge n \equiv m \pmod{k} &\Rightarrow \\ \forall p. p \in \mathcal{Q}_h \Rightarrow (X^n - X^m) \llbracket p \rrbracket &\equiv \mathbf{0} \pmod{h} \end{aligned}$$

Proof. Given $p \in \mathcal{Q}_h$, there is $q \in \mathcal{P}$ such that $p = q \pmod{h}$. Therefore $n \asymp q$ and $m \asymp q$ by definition of \mathcal{P} . Let $z = X^k - 1$. Note that $z \equiv \mathbf{0} \pmod{h}$ by given. We can proceed:

$$\begin{array}{ll} & q^n \equiv q \llbracket X^n \rrbracket \pmod{z} & \text{by } n \asymp q \text{ — [1]} \\ & q^m \equiv q \llbracket X^m \rrbracket \pmod{z} & \text{by } m \asymp q \text{ — [2]} \\ \text{and} & q \llbracket X^m \rrbracket \equiv q \llbracket X^n \rrbracket \pmod{z} & \text{by Theorem 16} \\ \text{so} & q^n \equiv q^m \pmod{z} & \text{by [1], [2], transitivity} \\ \text{Therefore} & q^n - q^m \equiv \mathbf{0} \pmod{z} & \text{by subtraction} \\ \text{by } z \equiv \mathbf{0} \pmod{h} & q^n - q^m \equiv \mathbf{0} \pmod{h} & \text{by Theorem 13—[3]} \\ \text{Since} & (X^n - X^m) \llbracket p \rrbracket \equiv (X^n - X^m) \llbracket q \rrbracket \pmod{h} & \text{by } p = q \pmod{h} \text{—[4]} \\ \text{and the right-side} & (X^n - X^m) \llbracket q \rrbracket = q^n - q^m & \text{by substitution of } q \text{—[5]} \\ \text{Combine [4],[5],[3]} & (X^n - X^m) \llbracket p \rrbracket \equiv \mathbf{0} \pmod{h} & \text{as claimed.} \end{array}$$

\square

Due to this, an injective map between the two finite sets derived from \mathcal{N} is possible:

Theorem 22. *There is an injective map from reduced set of \mathcal{N} to modulo set of \mathcal{N} .*

$$\begin{aligned} \vdash \text{FiniteField } \mathcal{F} \wedge \text{monic } h \wedge \text{ipoly } h \wedge X^k - 1 \equiv \mathbf{0} \pmod{h} &\Rightarrow \\ \forall n, p. & \\ 1 < p \wedge p < n \wedge n \in \mathcal{N} \wedge p \in \mathcal{N} \wedge n^2 \sqrt{|\mathcal{M}|} < |\mathcal{Q}_h| &\Rightarrow \\ (\lambda m. m \pmod{k}) : \widehat{\mathcal{N}} \setminus \mathcal{P} \setminus \sqrt{|\mathcal{M}|} &\hookrightarrow \mathcal{M} \end{aligned}$$

Proof. Let $i, j \in \widehat{\mathcal{N}} \setminus \mathcal{P} \setminus \sqrt{|\mathcal{M}|}$, with $i \equiv j \pmod{k}$ in \mathcal{M} . If the map is to be injective, we need $i = j$. Since $\widehat{\mathcal{N}} \setminus \mathcal{P} \setminus \sqrt{|\mathcal{M}|} \subseteq \mathcal{N}$, both $i, j \in \mathcal{N}$. Theorem 21 applies: every $p \in \mathcal{Q}_h$ is a root of $X^i - X^j$. Hence there are at least $|\mathcal{Q}_h|$ roots.

By Theorem 20, both i, j are bounded by $n^2 \sqrt{|\mathcal{M}|}$, hence $\deg(X^i - X^j) \leq n^2 \sqrt{|\mathcal{M}|}$. Given $n^2 \sqrt{|\mathcal{M}|} < |\mathcal{Q}_h|$, there are more roots than its degree for the polynomial $(X^i - X^j)$ with coefficients from a finite field \mathcal{F} . This is not possible, unless it is $\mathbf{0}$, which means $i = j$. \square

4.5 Punch Line

Given a prime p that divides n , if $n^x = p^y$ for some exponents x, y with $x > 0$, what can we conclude?

Theorem 23. *A condition that implies a number is a perfect power of prime.*

$$\vdash 0 < n \wedge \text{prime } p \wedge p \mid n \wedge (\exists x, y. 0 < x \wedge p^x = n^y) \Rightarrow \text{perfect_power } n, p$$

Proof. Since $p \mid n$, divide n by p as many times as possible, and express $n = p^m q$ where m is the maximum possible, and $p \nmid q$. The equation $p^x = n^y$ becomes $p^x = (p^m q)^y = p^{my} q^y$. By unique factorisation, with prime p and $p \nmid q$ and $x \neq 0$, it must be that $y \neq 0$, and $q^y = 1$, i.e., $q = 1$. \square

When its generators have a special property, the cardinality of $\widehat{\mathcal{N}}_p(n, m)$ is simple to express:

Theorem 24. *Cardinality of $\widehat{\mathcal{N}}$ when generators n and prime divisor p are not related by perfect power.*

$$\begin{aligned} &\vdash \text{Ring } \mathcal{R} \wedge \mathbf{1} \neq \mathbf{0} \wedge \mathbf{1} < k \Rightarrow \\ &\quad \forall n, p, m. \\ &\quad n \in \mathcal{N} \wedge p \in \mathcal{N} \wedge \text{prime } p \wedge p \mid n \wedge \neg \text{perfect_power } n, p \Rightarrow \\ &\quad |\widehat{\mathcal{N}}_p(n, m)| = (m+1)^2 \end{aligned}$$

Proof. Let $f = (\lambda(i, j) \cdot p^i n^j)$, $t = \{j \mid j \leq m\}$. From its definition, it is simple to verify that $|\widehat{\mathcal{N}}_p(n, m)| = \{p^i n^j \mid i \leq m \wedge j \leq m\} = f(t \times t)$. More interesting is that the conditions will imply $f : t \times t \hookrightarrow \widehat{\mathcal{N}}_p(n, m)$. Once this is proved, being the image of an injective map gives $|\widehat{\mathcal{N}}_p(n, m)| = |t \times t| = |t|^2 = (m+1)^2$.

To show that the map is injective, assume $p^i n^j = p^u n^v$ for some i, j and u, v . We need to show $i = u$ and $j = v$. This comes down to analysis by cases.

If $i < u$, only the case $j > v$ is interesting, with $n^{j-v} = p^{u-i}$. As $j - v \neq 0$, Theorem 23 applies, giving $\text{perfect_power } n, p$, which contradicts the assumption. By the symmetric roles of i, j and u, v , the case $i > u$ leads to the same contradiction. The only possible case is $i = u$, giving $j = v$. \square

This property is crucial in order to complete the proof of AKS Main Theorem (Theorem 11).

Proof (of Theorem 11). AKS Main Theorem in finite fields

$$\begin{aligned} &\vdash \text{FiniteField } \mathcal{F} \wedge \text{prime } k \wedge k < \chi \Rightarrow \\ &\quad \forall n. \\ &\quad 1 < n \wedge \chi \mid n \wedge \text{gcd}(n, k) = 1 \wedge (2(\log n + 1))^2 \leq \text{order}_k(n) \wedge \\ &\quad \ell = 2\sqrt{k}(\log n + 1) \wedge (\forall c. 0 < c \wedge c \leq \ell \Rightarrow n \not\equiv \mathbf{X} + c) \Rightarrow \\ &\quad \text{perfect_power } n, \chi \end{aligned}$$

Let $p = \chi$. By assumption, $p \mid n$, so $p \leq n$. The case $p = n$ is trivial, so we shall assume $p < n$.

The finite field \mathcal{F} gives prime p , so $p \not\equiv \mathbf{X} + c$ (Theorem 4). We have $k < p$, so $\text{gcd}(p, k) = 1$. Assuming $\text{gcd}(n, k) = 1$ and $n \not\equiv \mathbf{X} + c$, we have the ingredients for the introspective sets \mathcal{N} and \mathcal{P} (Section 4.1). Their finite counterparts, the modulo sets \mathcal{M} and \mathcal{Q}_h (Section 4.2), and reduced sets $\widehat{\mathcal{N}}$ and $\widehat{\mathcal{P}}$ (Section 4.3 and Section 4.4) can be set up accordingly.

Recall that the introspective relation is based on modulus $\mathbf{X}^k - \mathbf{1}$. By the second useful fact in Section 4, in a finite field \mathcal{F} it has a monic irreducible factor $h \neq \mathbf{X} - \mathbf{1}$, i.e., $\mathbf{X}^k - \mathbf{1} \equiv \mathbf{0} \pmod{h}$. With prime k , we have $\text{order}_h(\mathbf{X}) = k$ (Theorem 14), giving the injective map from $\widehat{\mathcal{P}}$ to \mathcal{Q}_h (Theorem 18), which is essential for the lower bound estimate of \mathcal{Q}_h .

In Section 4.6, we shall investigate the parameters k and ℓ . We shall show that $\ell < k$ (Theorem 27). By assumption, $k < p$, so $\ell < p$. Therefore $2^{\min(\ell, |\mathcal{M}|)} \leq |\mathcal{Q}_h|$ (Theorem 19, which invokes Theorem 18). We shall also show that $n^{2\sqrt{|\mathcal{M}|}} < 2^{\min(\ell, |\mathcal{M}|)}$ (Theorem 26). Hence $n^{2\sqrt{|\mathcal{M}|}} < |\mathcal{Q}_h|$. With $p < n$, these inequalities establish the injective map from $\widehat{\mathcal{N}}_p(n, \sqrt{|\mathcal{M}|})$ to $\widehat{\mathcal{N}}$ (Theorem 22).

Now, given prime p and $p \mid n$, if n were not a perfect power of p , Theorem 24 applies, so that:

$$|\widehat{\mathcal{N}}_p(n, \sqrt{|\mathcal{M}|})| = (\sqrt{|\mathcal{M}|} + 1)^2 = |\mathcal{M}| + (2\sqrt{|\mathcal{M}|}) + 1 > |\mathcal{M}|$$

This means the injective map from $\widehat{\mathcal{N}}_p(n, \sqrt{|\mathcal{M}|})$ to \mathcal{M} , both finite sets, would violate the Pigeonhole Principle. Therefore, n must be a perfect power of $p = \chi$. \square

4.6 Parameters

The AKS Main Theorem contains a parameter k with the property: $\text{order}_k(n) \geq (2(\log n + 1))^2$, from which a related parameter $\ell = 2\sqrt{k}(\log n + 1)$ is computed.

In the original AKS paper [2], parameter k is a prime (for a different set of conditions) while in the revised version [3] this prime requirement on k is dropped. Only the bound on k affects the conclusion “PRIMES is in P”, a general k needs more advanced theory to establish. Our mechanisation effort is based on a prime k , following Dietzfelbinger [10]. We use a prime k to show $k = \text{order}_h(X)$ in Theorem 14.

The existence of such a prime k can be established by generalizing the problem: given a number n , and a maximum m , find a prime modulus k such that $\text{order}_k(n) \geq m$. This is applied in Theorem 8:

Theorem 25. *There is always a modulus k giving big enough order for n in \mathbb{Z}_k .*

$$\vdash 1 < n \wedge 0 < m \Rightarrow \exists k. \text{prime } k \wedge \text{gcd}(k, n) = 1 \wedge m \leq \text{order}_k(n)$$

Proof. First, we define a set of candidates:

$$\vdash \text{candidates } n \ m = \{k \mid \text{prime } k \wedge k \nmid n \wedge \forall j. 0 < j \wedge j < m \Rightarrow k \nmid n^j - 1\}$$

Pick a large prime $z > n^m$, then z cannot divide n or any of the factors $n^j - 1$ for $0 < j < m$, hence $z \in \text{candidates } n \ m$.

Thus $\text{candidates } n \ m \neq \emptyset$, and we can pick a candidate k , say the least value, from the set. Being an element, $\text{prime } k \wedge k \nmid n$. Since a prime is coprime to its non-multiples, $\text{gcd}(k, n) = 1$. Thus n has nonzero order in \mathbb{Z}_k . Let $j = \text{order}_k(n)$, then $0 < j$ with $n^j \equiv 1 \pmod{k}$, or $k \mid n^j - 1$. If $j < m$, by the candidates definition $k \nmid n^j - 1$, a contradiction. Hence $\text{order}_k(n) = j \geq m$. \square

The parameters k and ℓ provide a crucial inequality involving $|\mathcal{M}|$, used in Theorem 11:

Theorem 26. *The AKS parameters meet the inequality condition.*

$$\vdash \text{FiniteField } \mathcal{F} \wedge 1 < k \wedge 1 < n \wedge n \in \mathcal{N} \wedge (2(\log n + 1))^2 \leq \text{order}_k(n) \wedge \\ \ell = 2\sqrt{k}(\log n + 1) \Rightarrow \\ n^2 \sqrt{|\mathcal{M}|} < 2^{\min(\ell, |\mathcal{M}|)}$$

Proof. Let $j = \text{order}_k(n)$, and $m = \log n + 1$, then $2^m > n$ for integer logarithm. By Theorem 12, $j \leq |\mathcal{M}|$ and $|\mathcal{M}| < k$. By the given assumption, $(2m)^2 \leq j$. Taking integer square roots, we have $\sqrt{|\mathcal{M}|} \geq \sqrt{j}$, $\sqrt{k} \geq \sqrt{|\mathcal{M}|}$ and $\sqrt{j} \geq 2m$. Note also $|\mathcal{M}| \geq \sqrt{|\mathcal{M}|} \sqrt{|\mathcal{M}|}$ by integer square root. Therefore:

$$\begin{aligned} - \ell &= 2\sqrt{k}m \geq m(2\sqrt{|\mathcal{M}|}) \\ - |\mathcal{M}| &\geq \sqrt{j} \sqrt{|\mathcal{M}|} \geq m(2\sqrt{|\mathcal{M}|}) \end{aligned}$$

Thus $\min(\ell, |\mathcal{M}|) \geq m(2\sqrt{|\mathcal{M}|})$, and

$$2^{\min(\ell, |\mathcal{M}|)} \geq 2^{m(2\sqrt{|\mathcal{M}|})} = 2^{m^2 \sqrt{|\mathcal{M}|}} > n^2 \sqrt{|\mathcal{M}|}.$$

\square

Incidentally, the choice of k and ℓ ensures that $\ell \leq k$, used in Theorem 10 and Theorem 11:

Theorem 27. *The AKS computed parameter does not exceed the modulus parameter.*

$$\vdash 1 < n \wedge 1 < k \wedge \text{gcd}(k, n) = 1 \wedge (2(\log n + 1))^2 \leq \text{order}_k(n) \Rightarrow \\ 2\sqrt{k}(\log n + 1) \leq k$$

Proof. Since $\text{order}_k(n) \mid \varphi(k)$, and $\varphi(k) < k$ when $k > 1$, we have $\text{order}_k(n) < k$. Taking integer square-roots, with the given $\text{order}_k(n)$, deduce

$$k \geq \sqrt{k} \sqrt{k} \geq \sqrt{k} \sqrt{\text{order}_k(n)} \geq 2\sqrt{k}(\log n + 1).$$

\square

5 Mechanisation and Its Traps

The updated AKS proof [3] is contained within four pages. Mechanisation of such a proof is the process of unwinding the dense mathematics within those pages. It took us about a year to build up the basic libraries, another year to forge the advanced libraries, then about six months to adapt the libraries for the proof of AKS Main Theorem, during which time the missing pieces in the developed libraries were steadily being filled in.

There are always pitfalls during the mechanisation process. One example is the symbol X in various expositions of the AKS proof, *e.g.*, [10, 7, 8, 4]. Usually X starts as an indeterminate or a degree one zero constant monomial, then switches to a root of unity, even to a primitive root of unity. While this is common practice, such changes mean that we needed to prove the switchings are valid.

The substitution by X is fundamental in the introspective relation (Section 2.2). These subtle changes in the role of X presented some difficulties in our initial effort to formalize the AKS proof. Indeed, we first used an inappropriate definition and got carried along until we found that shifting playgrounds (Section 3.3) is impossible with that definition.

Shifting of playgrounds in the AKS proof is pivotal. Most expositions just point this out without further elaboration.⁴ After this shifting, where the playground is now \mathbb{Z}_p , the introspective relation is defined in $\mathbb{Z}_p[X]$, side-stepping the issue. It was in the process of mechanisation that we realized a proper formulation should start by defining the introspective relation in a ring \mathcal{R} (Section 2.2), and then prove that shifting is valid through ring homomorphisms from \mathbb{Z}_n to \mathbb{Z}_p (Section 3.3).

Lessons Learnt Rather than attempting a direct transcription of the AKS proof, we came to understand the proof in the context of finite fields, identifying the key concepts involved in the proof, even comparing various expositions. By reformulations of polynomial theorems in number theory into their counterparts in rings and fields, a clear picture of the proof’s logic emerged, resulting in this succinct presentation.

HOLA and Abstract Algebra This work demonstrates that HOL4’s simple type theory, together with its proof machinery, are sufficient to allow the statement and proof of moderately complicated theorems in abstract algebra. Without dependent types (as in Coq) or locales (as in Isabelle), theorems are slightly more awkward to state, but our experience is that *ad hoc* overloading gets one a long way. Over-annotation of terms so that the parser chooses the “right” meaning of a symbol like $+$ is only necessary occasionally. Exploiting overloading in this way requires a careful understanding of just what the parser is and is not capable of, and one is often on exactly that boundary. Nonetheless, the result gives terms that are not far removed from those that have been pretty-printed in this paper. (Pretty-printing to \LaTeX adds niceties such as superscripts and juxtaposition for multiplication; these could not be handled by the parser.)

Nor should we forget that Campos *et al* [5] proved half of the Main Theorem in ACL2, where the underlying logic is even simpler, and provides no static type-checking.

6 Related Work

Other Pen-and-Paper Proofs The revised proof (2004) of the AKS team [3] takes this approach: use the injective map on \mathcal{Q}_h to establish a lower bound for $|\mathcal{Q}_h|$; assuming that n is not a power of p , use the Pigeonhole Principle to show that a special nonzero polynomial has at least $|\mathcal{Q}_h|$ roots, thus giving an upper bound for $|\mathcal{Q}_h|$; manipulate inequalities to show that the chosen parameters will lead to the lower bound exceeding the upper bound, hence a contradiction.

Other expositions of the AKS Main Theorem [1, 12, 11, 13] take similar approaches, working mainly in $\mathbb{Z}_p[X]$. Our method is equivalent, but is clean in that we: (i) emphasize the important role of shifting

⁴ For example, [3] first stated the computational identity in \mathbb{Z}_n , then “this implies” the corresponding identity in \mathbb{Z}_p . Only [10] proved the shifting from \mathbb{Z}_n to \mathbb{Z}_p as a lemma.

from \mathbb{Z}_n to \mathbb{Z}_p (Section 3.3); (ii) reformulate the AKS Main Theorem in the context of finite fields (Theorem 11); (iii) clarify that the choice of parameters gives injective maps between reduced sets and modulo sets (Theorem 18 and Theorem 22); (iv) bring in the assumption that n is not a power of prime p as late as possible; and (v) use the Pigeonhole Principle as a punch line to force n to be a power of prime p (Section 4.5).

Other Mechanisations We believe that we are the first to mechanise both directions of the central theorem of AKS algorithm. As noted earlier, two other teams (Campos *et al* [5] in ACL2, and de Moura and Tadeu [9] in Coq) have mechanised the fact that if the number being tested is prime, then the AKS algorithm will indeed report “yes”.

We are also aware of preliminary work started by John Harrison, and carried out in HOL Light.⁵

7 Conclusion

It is well-known that the cardinality of a finite field must be a prime power, and it is elementary to check whether a number is power-free. In essence, the AKS team showed that primality testing can be reduced to finite field cardinality testing, and demonstrated that the latter can be done in polynomial time.

Through our mechanisation effort, especially in presenting the AKS proof as an introspective game (Section 4), we hope that this elementary proof of the AKS Main Theorem provides further appreciation of the AKS team’s brilliant ideas.

Future Work While the existence of parameter k in the AKS Main Theorem is assured, to show that it is bounded by a polynomial function of $\log n$ is harder. In future work, we intend to perform the necessary complexity analysis of the AKS algorithm to complete the mechanised proof that PRIMES is indeed in P.

References

1. Manindra Agrawal. Primality tests based on Fermat’s Little Theorem, December 2006. Available from <http://www.cse.iitk.ac.in/users/manindra/presentations/FLTBasedTests.pdf>.
2. Manindra Agrawal, Neeraj Kayal, and Nitin Saxena. PRIMES is in P, August 2002. Original paper.
3. Manindra Agrawal, Neeraj Kayal, and Nitin Saxena. PRIMES is in P. *Annals of Mathematics*, 160(2):781–793, 2004.
4. Alberto Bedodi. Primality tests in polynomial time. Master’s thesis, Università Degli Studi Roma TRE, February 2010.
5. Cynthia Campos, Francois Modave, and Steve Roach. Towards the verification of the AKS primality test in ACL2, November 2004. Fifth International Conference on Intelligent Technologies.
6. Hing-Lun Chan and Michael Norrish. A string of pearls: Proofs of Fermat’s Little Theorem. *Journal of Formalized Reasoning*, 6(1):63–87, December 2013.
7. Richard Crandall and Carl Pomerance. *Prime Numbers: A Computational Perspective*. Springer, 2005.
8. Gabriel Daleson. Deterministic primality testing in polynomial time. Master’s thesis, Portland State University, December 2006.
9. Flávio L. C. de Moura and Ricardo Tadeu. The correctness of the AKS primality test in Coq, July 2008.
10. Martin Dietzfelbinger. *Primality Testing in Polynomial Time: From Randomized Algorithms to ‘PRIMES is in P’*. Lecture Notes in Computer Science. Springer, 2004.
11. Riaal Domingues. A polynomial time algorithm for prime recognition. Master’s thesis, University of Pretoria, January 2006.
12. Benjamin Linowitz. An exposition of the AKS polynomial time primality testing. Master’s thesis, University of Pennsylvania, March 2006.
13. Carl Pomerance. Primality testing, variations on a theme of Lucas, 2008. Available from <http://cm.bell-labs.com/who/carlp/PS/primalitytalk5.ps>.

⁵ John was kind enough to share his approach with us *via* private communication.