

A shallow arithmetisation of first-order validity

Murdoch J. Gabbay

BCTCS

14 April 2025

Arithmetisation of first-order logic

To **arithmetise** a computational task is to convert it into a task of finding roots of some polynomial. Arithmetisation is the essential first step for cryptographic treatments of logic and computation — as required e.g. for *verifiable* and *zero-knowledge* computation, motivated by modern distributed systems including blockchain ('weak' *verifier* verifies 'strong' *prover*; shielded transactions).

Arithmetisation is now a billion-dollar industry.

In a **recent draft paper** "*Arithmetisation of computation via polynomial semantics for first-order logic*"

(<https://eprint.iacr.org/2024/954>) I arithmetise first-order logic via a compositional shallow translation to assertions about roots of polynomials.

We will need two elementary Lemmas:



Lemma 1 (a standard result)

Write $1..n$ for the set $\{x \in \mathbb{N} \mid 1 \leq x \leq n\}$.

Lemma 1. Suppose $\vec{v} \in \mathbb{Q}^{len}$ is a tuple of length $len \in \mathbb{N}$. Then there exists a (unique) **interpolating polynomial**

$$ip(\vec{v}) \in \mathbb{Q}[X]^{<len}$$

(meaning that $ip(\vec{v})$ is a polynomial over X with rational coefficients and degree strictly less than len) such that

$$\forall x \in 1..len. ip(\vec{v})(x) = \vec{v}[x]$$

where $\vec{v}[x]$ is the x th element of \vec{v} .

Example. $ip(1.5, 2, 2.5, 3) = 1 + 0.5 * X$.

See [diagram on Wikipedia](#).

Lemma 2 (an easy arithmetic fact)

Lemma 2. Suppose $x, y \in \mathbb{Q}_{\geq 0}$. Then:

1. $x + y = 0$ iff $x = 0 \wedge y = 0$.
2. $x * y = 0$ iff $x = 0 \vee y = 0$.
3. $(x - y)^2 = 0$ iff $x = y$.

Proof. Facts of arithmetic.

This Lemma, easy as it is, suggests that we can treat $\mathbb{Q}_{\geq 0}$ as a domain of truth-values, such that zero represents 'true' and non-zero values represent 'false'.

A FOL syntax

Fix some data:

- ▶ An **index variable symbol** X .
- ▶ A set of **polynomial function symbols** $F \in \text{PolyFunc}$, each with **arity** $\text{arity}(F) \in \mathbb{N}_{\geq 1}$.

Write $F : n$ for “ $F \in \text{PolyFunc}$ and $\text{arity}(F) = n$ ”.

Then define **terms** and **predicates** by:

$$t ::= X \mid q \in \mathbb{Q} \mid t + t \mid t * t \mid F_i(t) \mid \text{len}(F) \mid \text{reify}(\phi)$$

$$\phi ::= t = t \mid \phi \wedge \phi \mid \phi \vee \phi \mid \forall_F X. \phi$$

Above, $i \in 1.. \text{arity}(F)$.

Valuations (we'll use them in the next slide)

A **valuation** ζ maps each $F : n$ to an $n \times l$ matrix $\zeta(F) \in \mathbb{Q}^{n \times l}$.
Write $length(\zeta(F))$ for l (number of columns).

Think of $\zeta(F)$ as a finite sample from a function graph, possibly with auxiliary information.

Example: Assume $Sqr : 2$ (length 3) and $Fact$ (length 5) and set:

$$\zeta(Sqr) = \begin{pmatrix} 1 & 2 & -3 \\ 1 & 4 & 9 \end{pmatrix} \quad \zeta(Fact) = \begin{pmatrix} 0 & 1 & 2 & 3 & 4 \\ 1 & 1 & 2 & 6 & 24 \\ 0 & 1 & 2 & 3 & 4 \end{pmatrix}$$

1. Row 2 is the square/factorial of row 1.
2. For factorial, row 3 uses a 'pointer trick': each entry points to a column with the relevant inductive step, if required. E.g. $4! = 4 * 3!$ and $\zeta(Fact)[3, 5] = 4$.
3. By Lemma 1, each matrix is representable as a tuple of two/three interpolating polynomials of degree at most $2/4$.

Let's characterise these matrices in our logic

$$\varsigma(\text{Sqr}) = \begin{pmatrix} 1 & 2 & -3 \\ 1 & 4 & 9 \end{pmatrix} \quad \varsigma(\text{Fact}) = \begin{pmatrix} 0 & 1 & 2 & 3 & 4 \\ 1 & 1 & 2 & 6 & 24 \\ 0 & 1 & 2 & 3 & 4 \end{pmatrix}$$

satisfy the following predicates, for every value of X in $1..\text{len}(\varsigma(\text{Sqr}))$ and $1..\text{len}(\varsigma(\text{Fact}))$ respectively:

$$\text{Sqr}_2(X) = \text{Sqr}_1(X) * \text{Sqr}_1(X)$$

$$(\text{Fact}_1(X) = 0 \wedge \text{Fact}_2(X) = 1) \vee$$

$$(\text{Fact}_2(X) = \text{Fact}_1(X) * \text{Fact}_2(\text{Fact}_3(X)))$$

Denotation $[t]_{\zeta}, [\phi]_{\zeta} \in \mathbb{Q}[X]$

$$[X]_{\zeta} = X$$

$$[q]_{\zeta} = q \quad (q \in \mathbb{Q})$$

$$[t + t']_{\zeta} = [t]_{\zeta} + [t']_{\zeta}$$

$$[t * t']_{\zeta} = [t]_{\zeta} * [t']_{\zeta}$$

$$[F_i(t)]_{\zeta} = ip(\zeta(F_i))([t]_{\zeta})$$

$$[\text{len}(F)]_{\zeta} = \text{length}(\zeta(F_i))$$

$$[\text{reify}(\phi)]_{\zeta} = [\phi]_{\zeta}$$

$$[t=t']_{\zeta} = ([t]_{\zeta} - [t']_{\zeta})^2$$

$$[\phi \wedge \phi']_{\zeta} = [\phi]_{\zeta} + [\phi']_{\zeta}$$

$$[\phi \vee \phi']_{\zeta} = [\phi]_{\zeta} * [\phi']_{\zeta}$$

$$[\forall_F X. \phi]_{\zeta} = \sum_{x \in 1..length(\zeta(F))} [\phi]_{\zeta}$$

$$\zeta \models \phi \quad \text{when} \quad [\phi]_{\zeta} = 0 \quad (\phi \text{ closed})$$

$$t ::= X \mid q \in \mathbb{Q} \mid t + t \mid t * t \mid F_i(t) \mid \text{len}(F) \mid \text{reify}(\phi)$$

$$\phi ::= t=t \mid \phi \wedge \phi \mid \phi \vee \phi \mid \forall_F X. \phi$$

Soundness and completeness

Proposition 3:

1. $\models_{\zeta} t=t'$ iff $[t]_{\zeta} = [t']_{\zeta}$.
2. $\models_{\zeta} \phi \wedge \phi'$ iff $\models_{\zeta} \phi \wedge \models_{\zeta} \phi'$.
3. $\models_{\zeta} \phi \vee \phi'$ iff $\models_{\zeta} \phi \vee \models_{\zeta} \phi'$.
4. $\models_{\zeta} \forall_{\mathbb{F}} X. \phi$ iff $\models_{\zeta} \phi[X:=x]$ for every $x \in 1..length(\zeta(\mathbb{F}))$.

Proof. Just from Lemma 2.

Now we can express that $\zeta(\text{Sqr})$ represents a partial graph for squaring, and $\zeta(\text{Fact})$ for factorial, by asserting:

$$\models_{\zeta} \forall_{\text{Sqr}} X. (\text{Sqr}_2(X) = \text{Sqr}_1(X) * \text{Sqr}_1(X))$$

$$\models_{\zeta} \forall_{\text{Fact}} X. ((\text{Fact}_1(X) = 0 \wedge \text{Fact}_2(X) = 1) \vee \\ (\text{Fact}_2(X) = \text{Fact}_1(X) * \text{Fact}_2(\text{Fact}_3(X))))$$

A flavour of how this can be cryptographically applied

(There's much more to it than what's on this slide.)

Suppose a prover wants to convince a verifier that $\models_{\zeta} \forall_{\mathbb{F}} X. \phi$ for some $\zeta(\mathbb{F})$ it has computed, with 10^9 columns (approx 1 terabyte).

On the face of it we would have to evaluate ϕ on 10^9 claimed roots.

Alternatively, a prover could compute polynomials

$zeroes = \prod_{x \in 1..10^9} (X-x)$ and $h = [\phi]_{\zeta} / zeroes$.

The verifier picks a random $q \in \mathbb{Q}$ and forces the prover to produce a cryptographically assured proof that $h(q) * zeroes(q) = [\phi]_{\zeta}(q)$.

That's *one* polynomial evaluation, instead of 10^9 .

By the Schwartz-Zippel lemma, for low-degree polynomials P and Q , equality $P - Q = 0$ is (to a high degree of certainty) the same as equality of their evaluation at a random point $P(x) - Q(x) = 0$, since for low-degree polynomials, $P(x) - Q(x)$ is zero nearly nowhere in the ambient field.

What's going on? How expressive is the logic?

What's going on is that the matrices pack derivation trees. Each node in the tree gets a column, and the 'pointer trick' allows us to reassemble the tree structure. Thus we can encode arbitrary inductive definitions.

We can specify computationally significant functions like Ackermann's Function, Gödel encodings, and SK-combinator reduction.

With more fiddling, we can express general recursion, negation, and arbitrarily nested function-definitions. Using reify, which reifies predicates as terms, we can express arbitrary Skolem functions.

This logic is *very* expressive; surprisingly so.

Future work: application to concrete schemes

This work is impractical as presented, but the ideas could be operationalised. They'd look different: practical cryptographic schemes use $\mathbb{F}_p[X]$ for a prime p , instead of $\mathbb{Q}[X]$.¹ I'm working on that now.

Many companies offer cryptographic virtual machines. All the ones I know of are essentially deep embeddings of VMs in polynomials: programming on 'virtual chips' built in finite fields instead of silicon.

To my knowledge, nobody's thought of doing a shallow embedding of FOL-style validity. This would permit all the goodness of logic in computer science — inductive definitions, high-level programming, correctness-by-construction, etc — directly in a compilation target, which is logical validity encoded in polynomials.

¹I'm not entirely convinced this is necessary, but it's where the industry is.  

Future work

In this talk, I fixed a univariate logic based on the expressivity required to specify functions of interest, and arithmetised it. Obvious generalisation is the multivariate case: admit X , X' , and Y .

Also, other direction: what logics naturally correspond to validity assertions on roots of polynomials? I used $[\phi]_{\zeta} = 0$ in this talk. An obvious thing is to set $\phi \vDash_{\zeta} \psi$ when $[\phi]_{\zeta} \mid [\psi]_{\zeta}$. This would lead to a resource-conscious logic.

Other flavours and variations are possible.

There's an interesting design space of logics here!

<https://eprint.iacr.org/2024/954>

Much to do. Thank you for listening.

