CDM

(Semi-)Rings

Klaus Sutner

Carnegie Mellon University

Fall 2024

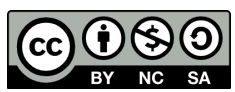

1 **Semirings and Rings**

2 **Polynomials: The Idea**

3 **Polynomials: Formal Definition**

4 **Roots**

Comment

We are interested in algebraic structures that support both addition and multiplication.

Typical examples in classical algebra are integers, rationals, reals, complexes, and are referred to as rings and fields.

In the computational universe one often encounters similar but weaker structures. Hence, it is a good idea to start at a class of structures that are somewhat more general, and then home in on rings and fields later.

A semiring is a structure $\langle X; \oplus, \otimes, 0, 1 \rangle$ that satisfies the following conditions:

1. $\langle X; \oplus, 0 \rangle$ and $\langle X; \otimes, 1 \rangle$ are monoids, the former is commutative.
2. Operation $\otimes$ distributes over $\oplus$ on the left and right:
   $x \otimes (y \oplus z) = (x \otimes y) \oplus (x \otimes z)$ and
   $(y \oplus z) \otimes x = (y \otimes x) \oplus (z \otimes x)$.
3. 0 is a null (or annihilator) with respect to $\otimes$: $x \otimes 0 = 0 \otimes x = 0$.

A semiring is commutative if $x \otimes y = y \otimes x$ and idempotent if $x \oplus x = x$.

All the standard examples (integers, rationals, reals, complexes) are semirings.

The Boolean semiring has the form $\mathbb{B} = \langle \{0, 1\}; \vee, \wedge, 0, 1 \rangle$ where the operations are logical 'or' and 'and'.

The relation semiring has as carrier set all binary relations over some set A, set union as the additive operation, and relational composition as the multiplicative operation. 0 is the empty relation, and 1 is the identity relation.

$$\mathcal{R}_A = \langle \mathsf{Rel}_A; \cup, \circ, \emptyset, I_A \rangle$$

The language semiring over some alphabet Σ has the form

$$\mathcal{L}(\Sigma) = \langle \mathfrak{P}(\Sigma^\star); \cup, \cdot, \emptyset, \varepsilon \rangle$$

It is easy to check that all the equations for a semiring are satisfied by $\mathcal{L}(\Sigma)$. Similarly one can restrict the languages to be, say, regular or context-free and so on.

One can even introduce a metric on $\mathcal{L}(\Sigma)$ by setting $\text{dist}(L, K) := 2^{-n}$ where n is minimal such that $L \cap \Sigma^n \neq K \cap \Sigma^n$ for $L \neq K$ and $\text{dist}(L, L) = 0$.

It is not hard to see that $\mathcal{L}(\Sigma)$ is a complete metric space with respect to this distance function.

The tropical semiring is defined by

$$\mathsf{TS} = \langle \mathbb{N}_\infty; \min, +, \infty, 0 \rangle$$

Here ∞ is an "infinitely large" element that is adjoined to $\mathbb{N}$ and that behaves properly with respect to $\min$ and $+$. E.g., $\min(x, \infty) = x$ and $x + \infty = \infty$.

Dijkstra's famous algorithm for minimal cost paths operates in this semiring.

Suppose we have a semiring $\mathcal{S} = \langle S; \oplus, \otimes, 0, 1 \rangle$.

The matrix semiring over $\mathcal{S}$ has the form

$$\mathcal{S}^{n,n} = \langle S^{n,n}; \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$$

where $\oplus$ and $\otimes$ are the matrix operations that are directly obtained from $\mathcal{S}$; $\mathbf{0}$ and $\mathbf{1}$ are the appropriate null and identity matrices over S.

One can show that $\mathcal{S}^{n,n}$ is again a semiring.

Consider again the language semiring $\mathcal{L}(\Sigma)$ of all languages over Σ. The multiplicative operation here is concatenation of languages.

There is another critical operation on languages, Kleene star $x^\star$, a sort of infinite sum of products of arbitrary length:

$$x^\star = \sum_{i \geq 0} x^i = x^0 + x^1 + x^2 + \ldots + x^n + \ldots$$

This also makes sens e.g. in the relation semiring: in this case, reflexive transitive closure plays the role of Kleene star.

The star operation is useful since it makes it possible to solve linear equations of the form $x = a \cdot x + b$: the solution is $a^\star b$.

Suppose $\mathcal{S}$ is an idempotent semiring. To define Kleene star, we need an infinitary operation $\sum_{i \in I} a_i$ where $(a_i \mid i \in I)$ is any family of elements in S, I an arbitrary index set.

$\mathcal{S}$ is a closed semiring iff:

$$\sum_{i \in [n]} a_i = a_1 + \ldots + a_n$$

$$\Big(\sum_{i \in I} a_i\Big)\Big(\sum_{j \in J} b_j\Big) = \sum_{(i,j) \in I \times J} a_i b_j$$

$$\sum_{i \in I} a_i = \sum_{j \in J} \Big(\sum_{i \in I_j} a_i\Big)$$

In the last equation I is the disjoint union of the sets I_j. These conditions are really all sanity checks.

The star operation in a closed semiring is then defined as above. We can then verify somehat strange equations such as

- $(x+y)^* = (x^*y)^*x^*$
- $(xy)^* = 1 + x(yx)^*y$

Overall, trying to axiomatize these structures is quite difficult, much harder than the classical examples from algebra.

At any rate, one can check that $\mathcal{L}(\Sigma)$ is in fact a closed semiring, as is $\mathcal{R}$.

In these examples we also have the super idempotency property: $\sum a_i = a$ if for all $i \in I$: $a_i = a$.

It follows that $(x^*)^* = x^*$ and $1^* = 1$.

Definition

A ring is a semiring

$$\mathcal{R} = \langle R; +, \cdot, 0, 1 \rangle$$

where

- $\langle R; +, 0 \rangle$ is a commutative group (additive group),
- $\langle R; \cdot, 1 \rangle$ is a monoid (not necessarily commutative),
- multiplication distributes over addition:

$$x \cdot (y + z) = x \cdot y + x \cdot z$$

$$(y + z) \cdot x = y \cdot x + z \cdot x$$

Note that we need two distributive laws since multiplication is not assumed to be commutative. If multiplication is commutative the ring itself is called commutative.

One can relax the conditions a bit and deal with rings without a 1: for example, $2\mathbb{Z}$ is a ring without 1[†]. Rings with 1 are then called unital rings.

For our purposes, all rings are unital and we want $0 \neq 1$.

[†]Some authors call these things *rngs*. No comment.

Example (Integers)

The integers $\langle \mathbb{Z}, +, *, 0, 1 \rangle$ with the usual addition and multiplication for a ring.

Example (Modular Numbers)

The integers modulo n, $\langle \mathbb{Z}_n, +, *, 0, 1 \rangle$ with the usual addition and multiplication form a ring. If n is prime, this ring is actually a field. In particular there is a two-element field consisting just of 0 and 1. Note that these fields are finite.

Example (Standard Fields)

The rationals $\mathbb{Q}$, the reals $\mathbb{R}$, the complex numbers $\mathbb{C}$.

Example (Univariate Polynomials)

Given a ring R we can construct a new ring by considering all polynomials with coefficients in R, written $R[x]$ where x indicates the "unknown" or "variable".
For example, $\mathbb{Z}[x]$ is the ring of all polynomials with integer coefficients.

Example (Matrix Rings)

Another important way to construct rings is to consider square matrices with coefficients in a ground ring R.
For example, $\mathbb{R}^{n,n}$ denotes the ring of all n by n matrices with real coefficients.
Note that this ring is not commutative unless $n = 1$.

Definition

An inverse u' of a ring element u is any element such that $uu' = u'u = 1$.

A ring element u is called a unit if it has an inverse u'.

Proposition

0 is an annihilator in any ring and cannot be a unit.

Proof. Note that $a0 = a(0+0) = a0 + a0$, whence $a0 = 0$. If $0'$ is an inverse of 0 we have $1 = 0\,0' = (0+0)0' = 1+1$, so $0 = 1$.

□

The multiplicative identity in a ring is uniquely determined: $1 = 1 \cdot 1' = 1'$.

Proposition

If u is a unit, then its inverse is uniquely determined.

Proof.

Suppose $uu' = u'u = 1$ and $uu'' = u''u = 1$. Then

$$u' = u'1 = u'uu'' = 1u'' = u''.$$

□

As usual, lots of equational reasoning. At any rate, by uniqueness it makes sense to write the inverse in functional notation as u^{-1}.

Notation

$$R^\star = R - \{0\}$$

$$R^\times = \text{units of } \ R$$

Clearly, $R^\times \subseteq R^\star$ but can be much smaller: $\mathbb{Z}^\times = \{\pm 1\}$.

On the other hand, $\mathbb{Q}^\times = \mathbb{Q}^\star$.

One first step towards organizing rings into some kind of classification is to consider sums of 1s (after all, 1 is the only element other than 0 we know to exist).

$$\mathbf{1}_n = \sum_{i=1}^{n} 1 = \underbrace{1 + \ldots + 1}_{n}$$

There are two possibilities: all the $\mathbf{1}_n$ are distinct, as in $\mathbb{Z}$ or $\mathbb{Q}$.

Otherwise, there must be a repetition, say, $\mathbf{1}_n = \mathbf{1}_{n+k}$ for some $k > 0$. But then $\mathbf{1}_k = 0$.

Definition

The characteristic of a ring R is defined by

$$\mathrm{chr}(R) = \begin{cases} \min\big(k > 0 \mid \mathbf{1}_k = 0 \big) & \text{if } k \text{ exists,} \\ 0 & \text{otherwise.} \end{cases}$$

In calculus, characteristic 0 is the standard case: $\mathbb{Q} \subseteq \mathbb{R} \subseteq \mathbb{C}$ all have characteristic 0.

For us, rings of positive characteristic are even more important. Typical example: $\mathbb{Z}_m$ or $\mathbb{Z}_m[x]$.

We are interested in rings that have lots of units. One obstruction to having a multiplicative inverse is described in the next definition.

Definition

A ring element $a \neq 0$ is a zero divisor if there exist $b, c \neq 0$ such that $ab = ca = 0$.

Recall the old multiplicative map $\widehat{a} : R \to R$, $x \mapsto ax$.

Then $\widehat{a}$ is injective iff a fails to be a zero divisor[†].

[†]Strictly speaking, just a left zero divisor, but we won't get into the weeds

Definition

A commutative ring is an integral domain if it has no zero-divisors.

In other words, in an integral domain, $\langle R^\star; \cdot, 1 \rangle$ is a monoid.

Proposition (Multiplicative Cancellation)

In an integral domain we have $ab = ac$ where $a \neq 0$ implies $b = c$.

Proof. $ab = ac$ iff $a(b - c) = 0$, done. □

Example (Standard Integral Domains)

The integers $\mathbb{Z}$, the rationals $\mathbb{Q}$, the reals $\mathbb{R}$, the complex numbers $\mathbb{C}$ are all integral domains.

Example (Modular Numbers)

The ring of modular numbers $\mathbb{Z}_m$ is an integral domain iff m is prime.

Example (Non-ID)

The ring of 2×2 real matrices is not an integral domain:

$$\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$$

Definition

A field $\mathbb{F}$ is a ring in which the multiplicative monoid $\langle F^*; \cdot, 1 \rangle$ forms a commutative group.

In other words, every non-zero element is already a unit. As a consequence, in a field we can always solve linear equations

$$a \cdot x + b = 0$$

provided that $a \neq 0$: the solution is $x_0 = -a^{-1}b$. In fact, we can solve systems of linear equations using the standard machinery from linear algebra.

As we will see, this additional condition makes fields much more constrained than arbitrary rings. By the same token, they are also much more manageable.

Example

In calculus one always deals with the classical fields: the rationals $\mathbb{Q}$, the reals $\mathbb{R}$, the complex numbers $\mathbb{C}$.

Example

The modular numbers $\mathbb{Z}_m$ form a field for m is prime.

We can use the Extended Euclidean algorithm to compute multiplicative inverses: obtain two cofactors x and y such that $xa + ym = 1$. Then x is the multiplicative inverse of a modulo m.

Note that we can actually compute quite well in this type of finite field: the elements are trivial to implement and there is a reasonably efficient way to realize the field operations.

Note that one can axiomatize monoids and groups in a purely equational fashion, we do not need complicated formulae to describe these structures.

For rings we want to be able to say $0 \neq 1$, so we need one inequality.

Alas, for fields things get more complicated: the inverse operation is partial and we need to guard against argument 0:

$$x \neq 0 \;\Rightarrow\; x * x^{-1} = 1$$

One can try to pretend that inverse is total and explore the corresponding axiomatization; this yields a structure called a "meadow" which does not quite have the right properties.

Products Fail

One standard method in algebra that produces more complicated structures from simpler one is to form a product (operations are performed component-wise).

This works fine for structures with an equational axiomatization: semigroups, monoids, groups, and even rings.

Unfortunately, for fields products to not work. For let

$$F = F_1 \times F_2$$

where F_1 and F_2 are two fields. Then F is a ring, but never a field: the element $(0, 1) \in F$ is not $(0, 0)$, and so would have to have an inverse (a, b).

But $(0, 1)(a, b) = (0, b) \neq (1, 1)$, so $F_1 \times F_2$ is not a field.

1 Semirings and Rings

2 **Polynomials: The Idea**

3 Polynomials: Formal Definition

4 Roots

Informally, a (univariate) polynomial is an expression of the form

$$x^3 - 2x^2 + 3x - 1$$

First off, there is a mysterious unknown or variable x.

The summands $a\,x^k$ are monomials[†] and the coefficient a is supplied by the ground ring.

Lastly, the whole polynomial is a finite sum of such monomials.

Division is not allowed, all we have is addition and multiplication. In the example, the ground ring is presumably $\mathbb{Z}$.

[†] Some authors refer only to x^k as a monomial.

Hence we can represent a polynomial by its coefficient list:

$$\boldsymbol{a} = (a_0, a_1, \ldots, a_{n-1})$$

represents

$$p(x) = a_0 + a_1 x + a_2 x^2 + \ldots + a_{n-1} x^{n-1}$$

If $a_{n-1} \neq 0$ then $n-1$ is the degree of p.

The coefficient list is the explicit form of the polynomial, a term in a special normal form: sum-of-products.

Suppose we have a univariate polynomial $p(x)$ over ring R. If we think of $p(x)$ as an expression, it is quite natural to substitute a ring element a for x, written simply $p(a)$ or perhaps $p[a/x]$ or $p[x \mapsto a]$.

Once the substitution has been made we can evaluate to obtain another element in the ring:

$$p(x) = x^3 - 2x^2 + 3x - 1$$

Upon substitution $x \mapsto 2$ produces

$$p(2) = 2^3 - 2 \cdot 2^2 + 3 \cdot 2 - 1 = 5$$

Hence each polynomial p is associated with a polynomial function

$$\widehat{p} : R \to R \qquad a \mapsto p(a)$$

This may seem like splitting hairs, but it is often important to keep the two notions apart.

The polynomial and the associated polynomial function really are two different objects. Consider the ground ring $\mathbb{Z}_2$. The polynomial

$$p(x) = x + x^2$$

has the associated function

$$\widehat{p}(a) = 0$$

for all $a \in \mathbb{Z}_2$. In fact, any polynomial $p(x) = \sum_{i \in I} x^i$ produces the identically 0 map as long as $I \subseteq \mathbb{N}_+$ has even cardinality.

Exercise

Describe all polynomial functions over ground rings $\mathbb{Z}_2$ and $\mathbb{Z}_3$.

If the polynomial is given in coefficient form

$$\boldsymbol{a} = (a_0, a_1, \ldots, a_d)$$

we can efficiently evaluate it by rewriting it in a nested product form:

$$f(x) = \Big((\ldots(a_d x + a_{d-1})x + a_{d-2})x + \ldots + a_1\Big)x + a_0$$

Proposition

A polynomial of degree d can be evaluated in at most d ring multiplications and at most d ring additions.

Suppose we wish to construct a polynomial f that evaluates to given target values at certain points. Say we want $f(a_i) = b_i$ for $i = 0, \ldots, n$, where all the a_i are required to be distinct (often called support points). Define the Lagrange interpolant

$$L_i^n(x) = \prod_{j \neq i} \frac{x - a_j}{a_i - a_j}$$

Proposition

$L_i^n(a_i) = 1$ and $L_i^n(a_j) = 0$ for $i \neq j$.

Hence we can choose

$$f(x) = \sum_{i \leq n} b_i L_i^n(x)$$

Note that f has degree bound n.

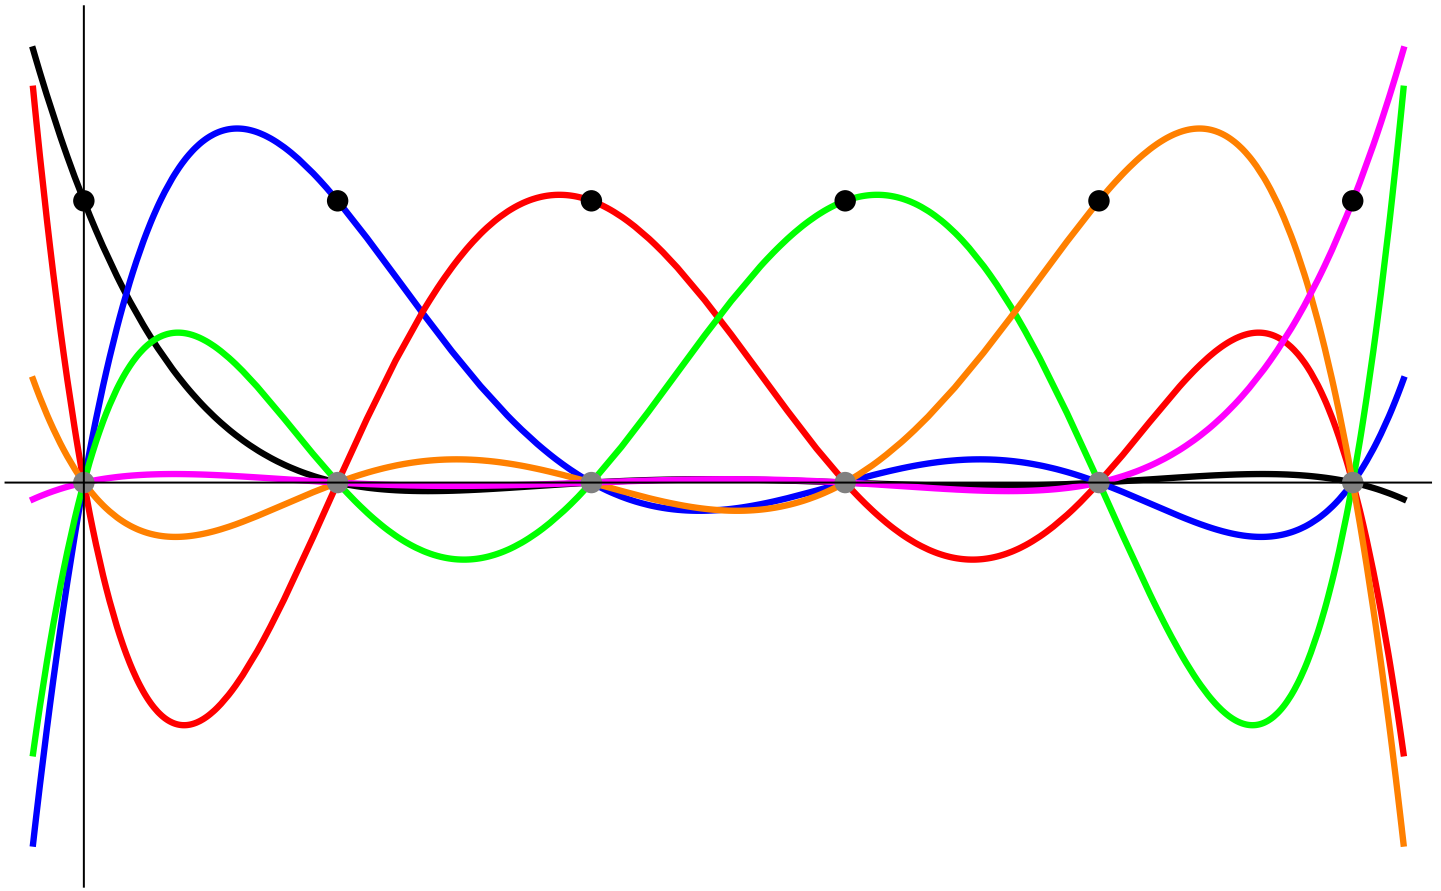

Example

Suppose we want $f(i) =$ the ith prime for $i = 0, \ldots, 5$.

The Lagrange interpolation looks like

$$f(x) = 2L_0^6 + 3L_1^6 + 5L_2^6 + 7L_3^6 + 11L_4^6 + 13L_5^6$$

which, after expansion and simplification, produces

$$\frac{1}{120}(240 - 286\,x + 735\,x^2 - 425\,x^3 + 105\,x^4 - 9\,x^5)$$

Suppose you have a "secret" a, a natural number, that you want to distribute over n people in such a way that no proper subgroup of the n persons can access the secret but the whole group can.

We may safely assume that a is a m-bit number. Generate $n-1$ m-bit random numbers a_i and give number a_i to person i, $i = 1, \ldots, n-1$. Lastly, person n receives

$$a_n = a \oplus a_1 \oplus a_2 \oplus \ldots \oplus a_{n-1}$$

where $\oplus$ is bit-wise xor.

Clearly all n secret sharers can compute a, but if one is missing they are stuck with a random number. This is very similar to one-time pads in cryptography.

A better organized approach is built on the following idea: pick a prime $p > a, n$. We will use the ground ring $\mathbb{Z}_p$ (which is actually a field).

- Generate random numbers $0 < a_i < p$ for $i = 1, \ldots, n-1$.
- Define the polynomial

$$f(x) = a + a_1 x + \ldots + a_{n-1} x^{n-1}$$

f is completely determined by the n point-value pairs (i, b_i), $i = 1, \ldots, n$.

By interpolation we can retrieve f from the point-value pairs, hence we can determine $a = a_0$.

On the other hand, $n-1$ persons can obtain no information about the zero coefficient; every coefficient is equally likely.

1 Semirings and Rings

2 Polynomials: The Idea

3 **Polynomials: Formal Definition**

4 Roots

Recall our original "definition" of a polynomial.

> Informally, a (univariate) polynomial is an expression of the form
>
> $$x^3 - 2x^2 + 3x - 1$$

Wisdom: If you come across a definition of the kind

> A foobag is an expression of the form blah-blah-blah.

run for the hills.

Most of the time, all you get is one example of the expressions in question, containing at least one ellipsis. No explanation of the underlying language, no definition of what exactly an expression is, certainly no formal grammar that defines everythin, zip.

Here is the "definition" of a real number from an otherwise great introductory analysis text.

> Now we begin formally. What is a real number? It will be an expression of the form
>
> $$\pm a_1 a_2 a_3 \ldots a_m . b_1 b_2 b_3 \ldots .$$
>
> Here the $\pm$ represents a choice between plus and minus. The digit a_1 is an integer between 0 and 9 inclusive (unless m is different from 1 in which case it is restricted to being between 1 and 9, since it is the leading digit.) All other digits are integers between 0 and 9 inclusive.

Oh dear, there is nothing formal about this alleged definition. According to the author, a real number literally is an *expression*, an infinite string in this case. On that understanding, the real number π is just the string

$$+3.14159265358979323846264338327950288419 7 \ldots$$

where we have omitted a few digits at the end to save paper.

Sorry, this is just malpractice. The reals are not a bunch of "expressions" like the ones in the definition. No way, never, ever.

The only justification for this line of attack is that the author does not want to get involved with a real definition, based on Cauchy sequences or Dedekind cuts. That's fine, the text is about analysis after all, but it should be clearly stated, with links to an actual definition.

And, for crying out loud, don't call it "formal," it's anything but.

And Polynomials?

In the case of polynomials, what on earth is

- the magic unknown x?
- a monomial ax^k?
- a sum of monomials?

All these objects are supposed to live in some algebraic structure, the ring of polynomials, but we haven't constructed that object yet.

Since the whole purpose of the definition is to pin down this ring, the StringWorld definition is hopelessly circular and just a naked appeal to intuition.

Just to be clear, intuition is the power that drives everything. And, the StringWorld approach can be helpful to get one's intuition going. In fact, it is often the right place to start. To paraphrase Knuth:

> Premature formalization is the root of all evil.

But, but, but ...

One absolutely, totally cannot stop there. Wishy-washy land is where algorithms go to die–to compute, we need to build data types, and those need to refer to an actual definition, not just some vague appeal to intuition and analogy.

We start with a definition that tries to home in on the critical algebraic properties of a polynomial.

Let R be a commutative ring with 1 throughout.

Definition

Given a ring R, the $\mathbb{N}$-coproduct of R is defined by

$$\coprod_{\mathbb{N}} R = \{\, (a_n) \in R^{\mathbb{N}} \mid \text{only finitely many } a_n \neq 0 \,\}$$

An element of the coproduct is a sequence $(a_i)_{i\geq 0}$ of elements of R, subject to the condition that $a_n = 0$ for all $n \geq m$, for some threshold m.

So, in a way, we are still dealing with finite sequences.

Since almost all the terms a_n are 0, it makes sense to write

$$a_0 + a_1 x + \ldots + a_n x^n$$

instead of $(a_0, a_1, \ldots, a_n, 0, 0, 0, \ldots)$.

In fact, one often insists tha a_n is the last non-zero element in the sequence (or $n = 0$ if they all are 0).

Note that the "unknown" x is nothing but syntactic sugar, all we really have is a sequence with finite support. We might as well use X, y, z, `fred`, whatever.

Operations on Coproducts

Again, there is no "unknown" in the definition of the coproduct. That makes it easier to give clean definitions of the algebraic structure. Addition is easy:

$$(a_n) + (b_n) = (a_n + b_n)$$

The sum $(a_n) + (b_n)$ is again an element of the coproduct and it is not too hard to check that this operation is associative and commutative.

But multiplication is somewhat more complicated (Cauchy product):

$$(a_n) \cdot (b_n) = \left(\sum_{i+j=n} a_i \cdot b_j \right)$$

Proposition

The product $(a_n) \cdot (b_n)$ is an element of the coproduct.

Write $\mathbf{0}$ and $\mathbf{1}$ for the sequences $(0, 0, 0, \ldots)$ and $(1, 0, 0, \ldots)$, respectively.

We have $\boldsymbol{a} + \mathbf{0} = \boldsymbol{a}$ so that

$$\left\langle \coprod R; +, \mathbf{0} \right\rangle$$

is a commutative monoid and even a group.

Likewise $\mathbf{1} \cdot \boldsymbol{a} = \boldsymbol{a} \cdot \mathbf{1} = \boldsymbol{a}$ and

$$\left\langle \coprod R; \cdot, \mathbf{1} \right\rangle$$

is also a commutative monoid (assuming that R is commutative).

Here is a much more interesting element: let

$$x = (0, 1, 0, 0, 0, \ldots)$$

Then $x^2 = (0, 0, 1, 0, 0, \ldots)$, $x^3 = (0, 0, 0, 1, 0, \ldots)$ and so forth.

This justifies the all the syntactic sugar in the notation

$$a_0 + a_1 x + \ldots + a_n x^n$$

instead of the actual coproduct element

$$(a_0, a_1, \ldots, a_n, 0, 0, \ldots)$$

We have quietly used the fact that we can embed all of R in the coproduct:

$$a \mapsto (a, 0, 0, \ldots)$$

Moreover, this map is (trivially) a ring monomorphism.

Lemma

$\langle \coprod R; +, \cdot, \mathbf{0}, \mathbf{1} \rangle$ *is a ring. This ring is commutative whenever* R *is.*

This is unsurprising, but note that a proof requires a bit of work: we have to verify e.g. that multiplication as defined above really is associative.

We ignore the details.

Definition

The ring $\langle \coprod R; +, \cdot, \mathbf{0}, \mathbf{1} \rangle$ is the polynomial ring with coefficients in R and is usually written $R[x]$.

In calculus one studies $\mathbb{R}[x]$.

For our purposes, $\mathbb{Q}[x]$, $\mathbb{Z}[x]$, $\mathbb{Z}_m[x]$ or $\mathbb{F}[x]$ where $\mathbb{F}$ is a finite field will be more important.

Definition

Let R and S two rings. A ring homomorphism is a map $f : R \to S$ such that

$$f(a+b) = f(a) \oplus f(b)$$
$$f(a \cdot b) = f(a) * f(b)$$
$$f(1_R) = 1_S$$

A homomorphism is an epimorphism if it is surjective, an monomorphism if it is injective, and an isomorphism if it is bijective. An endomorphism is a homomorphism $R \to R$, and a automorphism is an isomorphism $R \to R$.
The kernel of a homomorphism f is $\{ x \in R \mid f(x) = 0 \}$.

One can show that $f(0_R) = 0_S$ and $f(-a) = -f(a)$. For any unit $u \in R$, $f(a)$ is a unit in S and $f(a)^{-1} = f(a^{-1})$.

Definition

Given two polynomials f and g, g divides f if for some polynomial q: $q \cdot g = f$.

For the integers, the most important algorithm associated with the notion of divisibility is the Division Algorithm: we can compute quotient q and remainder r such that $a = qb + r$, $0 \leq r < b$. The situation for polynomials is very similar.

Theorem (Division Algorithm)

Assume that F is a field. Let f and g be two univariate polynomials over F, $g \neq 0$. Then there exist polynomials q and r such that

$$f = q \cdot g + r \qquad \textit{where } \deg(r) < \deg(g).$$

Moreover, q and r are uniquely determined.

Proof Sketch

For existence consider the set of possible remainders

$$S = \{\, f - q \cdot g \mid q \in F[x] \,\}.$$

If $\mathbf{0} \in S$ we are done, so suppose otherwise.

Trick: let $r \in S$ be any element of minimal degree, say $r = f - qg$.

Write $m = \deg(r)$ and $n = \deg(g)$, so we need $m < n$.

Assume $m \geq n$ and define

$$r' = r - \frac{a_m}{b_n}\, x^{m-n}\, g$$

where a_m and b_n are the leading coefficients of r and g, respectively.

But then $\deg(r') < \deg(r)$ and $r' \in S$, contradicting minimality.

Uniqueness is left as an exercise.

An important application of the Division Algorithm for integers is the Euclidean algorithm for the GCD.

Likewise we can obtain a polynomial GCD algorithm from the Division Algorithm for polynomials.

In fact, essentially the same algorithm works, just replace $\mathbb{Z}$ by $\mathbb{Z}[x]$.

For example, we can obtain cofactors s and t such that

$$\gcd(f, g) = sf + tg.$$

1 Semirings and Rings

2 Polynomials: The Idea

3 Polynomials: Formal Definition

4 **Roots**

Definition

A ring element $a \in R$ is a root of $p(x) \in R[x]$ if $p(a) = 0$.

In other words, a root is any solution of the equation $p(x) = 0$.

Finding roots of polynomial equations is often very difficult, in particular when several variables are involved. For univariate polynomials over the reals good numerical methods exist, but over other rings things are problematic.

For example, computing square roots, i.e. solving $x^2 - a = 0$, over $\mathbb{Z}_m$ is surprisingly difficult. Of course there is a brute-force algorithm, but think of modulus m having thousands of digits.

And for $\mathbb{Z}[x_1, x_2, \ldots, x_n]$ it is even undecidable whether a root exists.

Lemma

Let a be a root of $f \in F[x]$. Then $(x - a)$ divides $f(x)$.

Proof. Write

$$f = q(x - a) + r$$

where $\deg(r) < 1$. But then r must be 0, done. □

Lemma

Any non-zero polynomial $f \in F[x]$ has at most $\deg(f)$ many roots.

Proof. Use the last lemma and induction on the degree. □

So if $\deg(f) = n$ and f has n roots we decompose f completely into linear terms:

$$f = c(x - a_1)(x - a_2)\dots(x - a_n)$$

Of course, there may be fewer roots, even over a rich field such as $\mathbb{R}$:
$f = x^2 + 2$ has no roots.

This problem can be fixed by enlarging $\mathbb{R}$ to the field of complex numbers $\mathbb{C}$ (the so-called algebraic completion of $\mathbb{R}$).

Note that over arbitrary rings more roots may well exist.

For example over $R = \mathbb{Z}_{15}$ the equation $x^2 - 4 = 0$ has four roots: $\{2, 7, 8, 13\}$.

But over the integers this fails:

$$(x-2)(x-7)(x-8)(x-13) = 1 + 7x^2 + x^4 \not\equiv x^2 - 4$$

Exercise

Using the Chinese Remainder theorem explain why there are four roots in the example above. Can you generalize?

The fact that a non-zero polynomial of degree n can have at most n roots can be used to show that the interpolating polynomial

$$f(x) = \sum_i b_i \prod_{j \neq i} \frac{x - a_j}{a_i - a_j}$$

is unique: suppose g is another interpolating polynomial so that $g(a_i) = b_i$.

Then $f - g$ has $n + 1$ roots and so is identically zero.

Hence we have an alternative representation for polynomials: we can give a list of point-value pairs rather than a list of coefficients.

To the naked eye this proposal may seem absurd: why bother with a representation that is clearly more complicated? As we well see, there are occasions when point-value is computationally superior to coefficient list.

Suppose we have two univariate polynomials f and g of degree bound n.

Using the brute force algorithm (i.e., literally implementing the definition of multiplication in $\coprod R$) we can compute the product fg in $\Theta(n^2)$ ring operations.

Now suppose we are dealing with real polynomials. There is a bizarre way to speed up multiplication:

- Convert f and g into point-value representation where the support points are carefully chosen.
- Multiply the values pointwise to get h.
- Convert h back to coefficient representation.

It may seem absurd to spend all the effort to convert between coefficient representation and point-value representation. Surprisingly, it turns out that the conversions can be handled in $\Theta(n \log n)$ steps using a technique called Fast Fourier Transform.

But the pointwise multiplication is linear in n, so the whole algorithm is just $\Theta(n \log n)$.

Theorem

Two real polynomials of degree bound n can be multiplied in $\Theta(n \log n)$ steps.

Take a look at an algorithm book for details.

Here is another look at conversions between coefficient and point-value representation, i.e., between evaluation and interpolating.

Definition

Define the n by n Vandermonde matrix by

$$\mathsf{VM}(x_0, x_1, \ldots, x_{n-1}) = \begin{pmatrix} 1 & x_0 & x_0^2 & \ldots & x_0^{n-1} \\ 1 & x_1 & x_1^2 & \ldots & x_1^{n-1} \\ 1 & x_2 & x_2^2 & \ldots & x_2^{n-1} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & x_{n-1} & x_{n-1}^2 & \ldots & x_{n-1}^{n-1} \end{pmatrix}$$

Lemma

$$|\mathsf{VM}(\boldsymbol{x})| = \prod_{i<j} x_j - x_i$$

It follows that the Vandermonde matrix is invertible iff all the x_i are distinct. Now consider a polynomial

$$f(x) = c_0 + c_1 x + \ldots + c_{n-1} x^{n-1}$$

To evaluate f at points $\boldsymbol{a} = (a_0, \ldots, a_{n-1})$ we can use matrix-by-vector multiplication:

$$\boldsymbol{b} = \mathsf{VM}(\boldsymbol{a}) \cdot \boldsymbol{c}$$

But given the values $\boldsymbol{b}$ we can obtain the coefficient vector by

$$\boldsymbol{c} = \mathsf{VM}(\boldsymbol{a})^{-1} \cdot \boldsymbol{b}$$

None of the implicit descriptions of a polynomial match the actual definition in terms of a coproduct.

But, we can recover the explicit polynomial (i.e., the coefficient list) from these explicit representations. E.g., the implicit polynomial

$$(x_1 - x_2)\,(x_3 - x_4)\,(x_5 - x_6)$$

expands to

$$x_1\,x_3\,x_5 - x_2\,x_3\,x_5 - x_1\,x_4\,x_5 + x_2\,x_4\,x_5 - x_1\,x_3\,x_6 + x_2\,x_3\,x_6 + x_1\,x_4\,x_6 - x_2\,x_4\,x_6.$$

We just have to expand (multiply out) to get the "classical form".

What exactly is meant by "expanding" a polynomial?

Expanding Polynomials

We want to bring a multivariate polynomial $f(x_1, x_2, \ldots, x_n)$ into coproduct form. First we apply rewrite rules to push multiplication to the bottom of the tree until we have a sum of products:

- $\alpha(\beta + \gamma) \mapsto \alpha\beta + \alpha\gamma$
- $(\beta + \gamma)\alpha \mapsto \beta\alpha + \gamma\alpha$

Then we collect terms with the same monomial and adjust the coefficient.

$$\ldots + c\boldsymbol{x}^{\boldsymbol{e}} + \ldots + d\boldsymbol{x}^{\boldsymbol{e}} \ldots \rightsquigarrow \ldots + (c + d)\boldsymbol{x}^{\boldsymbol{e}} + \ldots$$

Some terms may cancel—we don't keep monomials with coefficient 0.

The problem is that it may take exponential time to perform the expansion: there may be exponentially many terms in the actual polynomial.

Dire Warning

There are $\mathbb{NP}$-complete problems like Graph-3-Coloring that could be solved in polynomial time if we could somehow get polynomial expansion under control and perform it in polynomial time.

Take this statement with pounds of salt, obviously the expansion cannot be handled in polynomial time when the coefficient form has exponential size.

More technically, it turns out that a graph is 3-colorable iff a certain polynomial does not vanish. Checking that a polynomial in coefficient form vanishes is trivial, but in implicit form it is hard.