In lecture 4 we learned how to do arithmetic modulo a prime number. We can use modular arithmetic to evaluate expressions like $x^2 + 3x + 1 \mod 5$ for different values of x. For example, when $x = 2$,

$$2^2 + 3 \cdot 2 + 1 \mod 5 = 11 \mod 5 = 1.$$

Such expressions are called *polynomials*.

Throughout this lecture, q will be a prime number and we will use the symbol $\mathbb{F}_q$ to denote the set $\{0, 1, \ldots, q-1\}$ of residues modulo q. For any two numbers x and y in $\mathbb{F}_q$ we will use $x + y$, $-x$ $x \cdot y$, and x^{-1} to denote the corresponding operations modulo q.[1] For example, for $3, 4 \in \mathbb{F}_5$, $3 + 4 = 2$ and $3 \cdot 4 = 2$.

A *polynomial* of degree d over $\mathbb{F}_q$ is a function p from $\mathbb{F}_q$ to $\mathbb{F}_q$ of the form

$$p(x) = a_d x^d + a_{d-1} x^{d-1} + \cdots + a_1 x + a_0$$

where $a_0, a_1, \ldots, a_{d-1}, a_d$ are numbers in $\mathbb{F}_q$ called *coefficients*. We require $a_d \neq 0$ unless $d = 0$.

Polynomials can be added and multiplied. For example, if we work over $\mathbb{F}_7$:

$$(x^2 + 3x + 6) + (x + 5) = x^2 + 4x + 4$$
$$(x^2 + 3x + 6) \cdot (x + 5) = (x^3 + 3x^2 + 6x) + (5x^2 + x + 2) = x^3 + x^2 + 2.$$

If p has degree d and p' has degree d', then their sum $p + p'$ has degree at most the larger of d and d' and their product $p \cdot p'$ has degree $d + d'$.

1 Hashing and collisions

In computer science you sometimes need to index moderately large objects by shorter identifiers. Say you want to maintain a history of previously visited URLs in your web browser. A URL can be hundred of symbols long so it may be convenient to index it by a much shorter string in order to save up on storage and speed up search. A function h that maps a long object into a shorter description is usually called a *hash function*, and the value $h(x)$ is called the hash of object x.

How should one go about choosing a hash function? One possibility is to hash an object by, say, the first 30 bits of its description as a bit string. In the case of URLs this is not a good idea because virtually all URLs share the prefix `http://www` and they would all be indexed by the same string. This is clearly undesirable.

We consider the following mathematical model of hashing. The objects of interest are represented by the integers from 0 to $q-1$ and their hashes will be represented by integers from 0 to $r-1$. Under this convention, a hash function h is a function from the set $\{0, \ldots, q-1\}$ to the set $\{0, \ldots, r-1\}$. Two distinct inputs x and x' are said to *collide* under h if $h(x) = h(x')$.

We will assume that the number r of possible hashes is (much) smaller than the number q of possible objects. The pigeonhole principle then guarantees the existence of a collision for some pair

[1] This structure is called a *finite field*: You can add, subtract, multiply, and divide except by zero, and the usual rules of arithmetic apply.

of objects. In what sense can collisions then be avoided? One strategy is to inject some randomness in the choice of hash function, so that even if collisions always exist they are unlikely to occur.

Specifically, we will consider the functions $h_{a,b}\colon \{0,\dots,q-1\} \to \{0,\dots,r-1\}$ given by the formula

$$h_{a,b}(x) = (ax + b \bmod q) \bmod r$$

where q is a prime number and (a,b) is a random pair of numbers in $\mathbb{F}_q$, chosen equally likely among all q^2 such pairs.

Let x and x' be any pair of distinct objects in $\mathbb{F}_q$. We will argue that x and x' are unlikely to collide under $h_{a,b}$ provided that the number r of possible hashes is sufficiently large.

Lemma 1. *Let x and x' be any two distinct elements of $\mathbb{F}_q$. The number of pairs (a,b) such that x and x' collide under $h_{a,b}$ is at most $(q+r)^2/r$.*

According to this lemma, the probability that x and x' collide for a random choice of a and b is

$$\frac{\text{number of pairs } (a,b) \text{ such that } h_{a,b}(x) = h_{a,b}(x')}{\text{total number of pairs } (a,b)} \le \frac{(q+r)^2/r}{q^2} = \frac{1}{r} \cdot \left(1 + \frac{r}{q}\right)^2.$$

The proof of Lemma 1 relies on the following important property of the degree-1 polynomials.

Lemma 2. *For every pair of distinct elements x and x' of $\mathbb{F}_q$, the function $E\colon \mathbb{F}_q^2 \to \mathbb{F}_q^2$ given by $E(a,b) = (ax+b, ax'+b)$ is bijective.*

Proof. Since the domain and the range both have size q^2, it is enough to prove that E is surjective. To do this, we verify that for every pair (y,y') of (not necessarily distinct) elements of F_q^2,

$$E((y'-y)/(x'-x), (yx'-y'x)/(x'-x)) = (y,y').$$

Indeed,

$$\frac{y'-y}{x'-x} \cdot x + \frac{yx'-y'x}{x'-x} = \frac{y'x - yx + yx' - y'x}{x'-x} = \frac{y(x'-x)}{x'-x} = y.$$

The other equality

$$\frac{y'-y}{x'-x} \cdot x + \frac{yx'-y'x}{x'-x} = y'$$

is verified similarly. We omit the details. □

The meaning of this mysterious formula will become apparent when we discuss interpolation. Let us now use Lemma 2 to prove Lemma 1. Since E is bijective, the pairs $(a,b) \in \mathbb{F}_q^2$ such that $h_{a,b}(x) = h_{a,b}(x')$ are in 1-1 correspondence with the pairs (y,y') such that

$$y \bmod r = y' \bmod r, \qquad \text{where } 0 \le y, y' \le q-1. \tag{1}$$

To prove Lemma 1, it remains to upper bound the number of pairs (y,y') that satisfy (1). To gain some intuition about this number, let us work out $q = 7, r = 3$ as an example. Then the set of pairs (y,y') that satisfy (1) is

$$\begin{aligned}\{\ &(0,0),(0,3),(0,6),(3,0),(3,3),(3,6),(6,0),(6,3),(6,6)\\ &(1,1),(1,4),(4,1),(4,4),\\ &(2,2),(2,5),(5,2),(5,5)\ \}.\end{aligned}$$

This set is the disjoint union of the three product sets $\{0,3,6\}^2$, $\{1,4\}^2$, and $\{2,5\}^2$, each of which is described by the common residue of y and y' modulo r. So the number of solutions to (1) in this example is $3^2+2^2+2^2=13$.

More crudely, we can upped bound the size of each product set by $(7/3+1)^2$ and get the bound of $3\cdot(7/3+1)^2=(7+3)^2/3$ for the number of solutions to (1). The following lemma generalizes this reasoning.

Lemma 3. *The number of pairs (y,y') that satisfy (1) is at most $(q+r)^2/r$.*

Proof. Let A_z be the set of all y between 0 and $q-1$ such that $y \bmod r = z$. Then the set of solutions to (1) is the disjoint union of the sets A_0^2, A_1^2, up to A_{r-1}^2. The function $g(y)=y-z$ is an injective function from A_z to A_0, so A_0 is at least as large as the other sets A_z. By the sum and product rules of counting, it follows that the number of solutions to (1) is at most

$$|A_0|^2+|A_1|^2+\cdots+|A_{r-1}|^2 \le r\cdot|A_0|^2.$$

The set A_0 consists of all non-negative multiples of r that are less than q. There can be at most $q/r+1$ such multiples. It follows that the desired number is at most $r\cdot(q/r+1)^2 \le (q+r)^2/r$. □

Lemma 1 now follows directly from Lemma 2 and Lemma 3.

More generally, we have some larger set X of objects that need to be hashed and we want to upper bound the probability that the hash function $h_{a,b}$ avoids collisions among elements of X. Lemma 1 addresses the special case when X consists of exactly two elements. The next theorem provides the desired bound.

Theorem 4. *Let X be any subset of $\mathbb{F}_q$ of size k. The number of pairs (a,b) such that some distinct elements x and x' of X collide under $h_{a,b}$ is at most $\binom{k}{2}(q+r)^2/r$.*

Therefore the probability that a collision within X exists, assuming a and b are chosen equally likely among all possibilities, is at most

$$\frac{\binom{k}{2}\cdot(q+r)^2/r}{q^2} = \binom{k}{2}\cdot\frac{1}{r}\cdot\left(1+\frac{r}{q}\right)^2.$$

For example, assuming we are hashing 100 objects, each of which is 500-bits long, into a 30-bit hash. Then $k=100$, $q\approx 2^{500}$, and $r=2^{30}$. For these parameters Theorem 4 guarantees that the probability some pair in X collides under $h_{a,b}$ is at most 2^{-18}. The ratio r/q is on the order of $2^{30}/2^{500}=2^{-470}$, which is negligible in the calculation.

The proof of Theorem 4 uses the inequality

$$|S_1\cup\cdots\cup S_N| \le |S_1|+\cdots+|S_N| \tag{2}$$

which is valid for any collection of sets $S_1,\cdots,S_N$. It can be proved by induction on N.

Proof of Theorem 4. Let $S_{\{x,x'\}}$ be the set of pairs (a,b) such that x and x' collide under $h_{a,b}$. The set of pairs (a,b) of interest is the union of the sets $S_{\{x,x'\}}$ over all unordered disjoint pairs $\{x,x'\}$ of elements of X. By Lemma 1, the size of each $S_{\{x,x'\}}$ is at most $(q+r)^2/r$. As there are $\binom{k}{2}$ ordered pairs $\{x,x'\}$, by inequality (2) the size of the union of the sets $S_{\{x,x'\}}$ can be at most $\binom{k}{2}(q+r)^2/r$. □

2 Zeros

A value c in $\mathbb{F}_q$ such that $p(c) = 0$ is called a *zero* of the polynomial p. For example, 1 is a zero of the polynomial $p(x) = x^3 + x^2 + x + 2$ over $\mathbb{F}_5$ because $p(1) = 0$. We can then write

$$\begin{aligned} p(x) &= p(x) - p(1) \\ &= (x^3 + x^2 + x + 2) - (1^3 + 1^2 + 1 + 2) \\ &= (x^3 - 1^3) + (x^2 - 1^2) + (x - 1) \\ &= (x-1)(x^2 + x + 1) + (x-1)(x+1) + (x-1) \\ &= (x-1)\big((x^2 + x + 1) + (x+1) + 1\big) \\ &= (x-1)(x^2 + 2x + 3) \end{aligned}$$

so we get the *factorization* of $p(x)$ as $x - 1$ times a polynomial of degree one lower than p:

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

We can apply a similar procedure to factor a polynomial p with a zero c as $p(x) = (x - c)p'(x)$, where p' is a polynomial of degree one lower than p. We now prove that such a factorization always exists.

Lemma 5. *If p is a polynomial of degree d over $\mathbb{F}_q$ and c is a zero of p, then there exists a polynomial p' of degree $d - 1$ such that $p(x) = (x - c)p'(x)$.*

Proof. First we do the case $c = 0$. Assume $p(x) = a_d x^d + \cdots + a_1 x + a_0$. Then $p(0) = a_0$, so $a_0 = 0$ and we can factor p as

$$p(x) = a_d x^d + \cdots + a_1 x = x \cdot p'(x) \quad \text{where } p'(x) = a_d x^{d-1} + \cdots + a_1.$$

If c is nonzero, look at the polynomial $p_1(x) = p(x + c)$, which also has degree d. Then 0 is a zero of p_1 so there exists a factorization $p_1(x) = x \cdot p_1'(x)$ for some p_1' of degree $d - 1$. Then

$$p(x) = p_1(x - c) = (x - c) \cdot p_1'(x - c) = (x - c)p'(x)$$

where $p'(x) = p_1'(x) - c$ is a polynomial of degree $d - 1$ as desired. □

The *zero polynomial* is the polynomial all of whose coefficients are zero.

Theorem 6. *Every nonzero polynomial over $\mathbb{F}_q$ of degree d has at most d zeros.*

Proof. The proof is by induction of d.

Base case $d = 0$: A nonzero polynomial p of degree 0 has the form $p(x) = a_0$ for $a_0 \neq 0$, so it doesn't have any zeros.

Inductive step: Assume every nonzero polynomial of degree $d - 1 \geq 0$ has at most $d - 1$ zeros. Let p be a polynomial of degree d. If p has no zeros, the theorem is true. Otherwise, p has at least one zero c. By Lemma 5, $p(x) = (x - c)p'(x)$ for some polynomial p' of degree $d - 1$. Every other zero of p is also a zero of p'. By our inductive assumption, p' has at most $d - 1$ zeros, so p has at most d zeros. □

A nonzero polynomial may evaluate to zero everywhere: For example, $x(x - 1)(x - 2) = x^3 - x$ over $\mathbb{F}_3$ evaluates to zero everywhere. However, this is impossible if the degree is smaller than q.

Corollary 7. *If $d < q$, p has degree d, and $p(c) = 0$ for all $c \in \mathbb{F}_q$, then p is the zero polynomial.*

Proof. Under the assumptions, p has $q > d$ zeros. By Theorem 6, p is the zero polynomial. □

3 Interpolation

One nice thing about polynomials is that we can recover the polynomial from its degree and enough of its input-output pairs. Let's see a couple of examples.

Example 1. Say you have a degree one polynomial p over $\mathbb{F}_5$ and you are told that $p(2) = 3$ and $p(3) = 0$. What is this polynomial p? We know p has the form $p(x) = a_1x + a_0$ for some a_0, a_1 in $\mathbb{F}_5$. We are told that

$$\begin{aligned} a_1 \cdot 2 + a_0 &= 3 \quad \text{and} \\ a_1 \cdot 3 + a_0 &= 0. \end{aligned}$$

To figure out p, we need to solve these two equations modulo 5. If we subtract the first equation from the second one we get $a_1 = -3 = 2$. Plugging this in the first equation we get $a_0 = 3-2\cdot 2 = -1 = 4$. So the desired polynomial is $p(x) = 2x + 4$.

Example 2. How about finding a polynomial p of degree 2 over $\mathbb{F}_7$ such that $p(0) = 3$, $p(1) = 0$, and $p(4) = 3$? Again, we know p has the form $p(x) = a_2x^2 + a_1x + a_0$ for some a_0, a_1, a_2 in $\mathbb{F}_7$ and we know that

$$\begin{aligned} a_0 &= 3 \\ a_2 + a_1 + a_0 &= 0 \\ a_2 \cdot 4^2 + a_1 \cdot 4 + a_0 &= 3. \end{aligned}$$

I solved this system of equations on the computer to get $a_2 = 1$, $a_1 = 3$, $a_0 = 3$, so $p(x) = x^2+3x+3$.

It looks like if we are given $d + 1$ values of the polynomial, we can figure out what the polynomial is. It turns out this is always possible to do. In fact, there is a formula for this polynomial called the *Lagrange interpolation formula*. For degree one, if we are given the "data" $p(x_0) = y_0$ and $p(x_1) = y_1$, this formula tells us that p is the polynomial

$$p(x) = y_0 \cdot \frac{x - x_1}{x_0 - x_1} + y_1 \cdot \frac{x - x_0}{x_1 - x_0}.$$

If you plug in $x = x_0$, the second term vanishes and the first term becomes y_0. If you plug in $x = x_1$, the first term vanishes and the second one becomes y_1 as desired. Let's re-solve Example 1 using this formula: We are told that $p(2) = 3$ and $p(3) = 0$ (over $\mathbb{F}_5$), so

$$p(x) = 3 \cdot \frac{x - 3}{2 - 3} + 0 \cdot \frac{x - 2}{3 - 2} = 2(x - 3) = 2x + 4.$$

For a degree 2 polynomial such that $p(x_0) = y_0$, $p(x_1) = y_1$, and $p(x_2) = y_2$, the Lagrange interpolation formula says that p equals

$$p(x) = y_0 \cdot \frac{(x - x_1)(x - x_2)}{(x_0 - x_1)(x_0 - x_2)} + y_1 \cdot \frac{(x - x_0)(x - x_2)}{(x_1 - x_0)(x_1 - x_2)} + y_2 \cdot \frac{(x - x_0)(x - x_1)}{(x_2 - x_0)(x_2 - x_1)}.$$

If we work out Example 2 using this formula, we get

$$\begin{aligned} p(x) &= 3 \cdot \frac{(x - 1)(x - 4)}{(0 - 1)(0 - 4)} + 0 \cdot \frac{(x - 0)(x - 4)}{(1 - 0)(1 - 4)} + 3 \cdot \frac{(x - 0)(x - 1)}{(4 - 0)(4 - 1)} \\ &= 6(x - 1)(x - 4) + 2x(x - 1) \\ &= (6x^2 - 30x + 24) + (2x^2 - 2x) \\ &= 8x^2 - 32x + 24 \\ &= x^2 + 3x + 3. \end{aligned}$$

As you can see, the formula can be cumbersome to use in practice, but it will help us prove that interpolation works.

Theorem 8. *Assume $d < q$. For every set of $d+1$ pairs of numbers $(x_0, y_0), \ldots, (x_d, y_d)$, all in $\mathbb{F}_q$, where $x_0, x_1, \ldots, x_d$ are all distinct, there exists exactly one polynomial p of degree at most d such that $p(x_i) = y_i$ for every i between 0 and d.*

We need to prove two things: That there is at least one p (existence) and there is at most one p (uniqueness).

Proof of existence. Let p be the polynomial

$$p(x) = \sum_{i=0}^{d} y_i \cdot \prod_{j=0, j\neq i}^{d} \frac{x - x_j}{x_i - x_j} \tag{3}$$

Every term in the summation is a product of d degree 1 polynomials, so the polynomial has degree at most d. When $x = x_i$, all but the i-th term in the summation contain the factor $(x - x_i)$ so they vanish. Only the i-th term survives and it evaluates to

$$p(x_i) = y_i \cdot \prod_{\substack{j=0 \\ j\neq i}}^{d} \frac{x_i - x_j}{x_i - x_j} = y_i,$$

so $p(x_i) = y_i$ for every i between 0 and d. □

Proof of uniqueness. Assume that p and p' are two polynomials of degree d such that $p(x_i) = y_i$ for all i between 0 and d. We will show that p and p' must in fact be the same polynomial. Let r be the polynomial $r(x) = p(x) - p'(x)$. The polynomial r has degree at most d and

$$r(x_i) = p(x_i) - p'(x_i) = y_i - y_i = 0$$

for all i between 0 and d. So r is a polynomial of degree d that has at least $d+1$ zeros. By Theorem 6, r is the zero polynomial, so p and p' must be the same polynomial. □

4 Secret sharing

Alice, Bob and Charlie find a stash of cash in the trash. They put it in a box for keeping until they find out who it belongs to. Each of them is afraid they will be conned by the other two, who may just decide to take the money and run. They come up with a safekeeping system: Each one will put their own lock on the outside of the box and keep the key for it. For the box to open, all three will need to agree and present their keys.

Secret sharing is the digital variant of this scenario. It involves n parties who want to share a secret piece of information, like the password to a shared bitcoin account. Each party i gets a *share* $p(i)$, which is also some piece of information. The sharing should be done in such a way that if any $d+1$ out of the n parties reveal their shares then they can recover the secret, but even if one out of these $d+1$ refuses to cooperate the other d cannot obtain any meaningful information about the secret.

It takes a bit of effort to explain all this in precise mathematical language. Instead of doing that, let me tell you somewhat informally how polynomials can be used to achieve this task. (There are also other ways to do it.)

First, we choose a large enough prime q and represent the secret s as a number in the set $\mathbb{F}_q = \{0, 1, \ldots, q-1\}$. For example $q = 5915587277$ is a 10 digit prime number. This is large enough to represent all possible secret numbers up to 9 digits long. The shares will also be numbers in $\mathbb{F}_q$.

To share the secret s, choose a sequence of d numbers $(a_1, a_2, \ldots, a_d)$ in $\mathbb{F}_q$, randomly so that the sequence is equally likely among all q^d possible sequences. Set $a_0 = s$ and evaluate the polynomial

$$p(x) = a_d x^d + \cdots + a_1 x + a_0$$

At $x = 1$, 2, up to n. Party i receives the share $p(i)$.

For example, if Alice, Bob, and Charlie want to share the secret password $s = 123456789$ (and work over $\mathbb{F}_{5915587277}$), I use the computer to choose the random sequence

$$a_1 = 3769551523, a_2 = 775093894$$

which gives

$$p(x) = 775093894 \cdot x^2 + 3769551523x + 123456789.$$

I give Alice, Bob, and Charlie the shares $p(1) = 4668102206$, $p(2) = 4847348134$, and $p(3) = 661194573$, respectively.

We have set things up so that the secret s equals the value $p(0) = a_0$. When a subset S of $d+1$ parties gets together to recover the secret, they use the Lagrange interpolation formula (3) at $x = 0$ to calculate

$$p(0) = \sum_{i \in S} p(i) \cdot \prod_{j \in S \setminus \{i\}} \frac{-j}{i-j}.$$

For Alice, Bob, and Charlie, this formula gives

$$\begin{aligned} p(0) &= 4668102206 \cdot \frac{(-2)(-3)}{(1-2)(1-3)} + 4847348134 \cdot \frac{(-1)(-3)}{(2-1)(2-3)} + 661194573 \cdot \frac{(-1)(-2)}{(3-1)(3-2)} \\ &= 4668102206 \cdot 3 + 4847348134 \cdot (-3) + 661194573 \cdot 1 \\ &= 123456789. \end{aligned}$$

Security Now we want to show that any d parties cannot discover any information about the secret s. To not discover any information does not merely mean that d parties cannot *recover* the secret. They must also be unable to answer questions like "Is it an even number?" or "Does its decimal representation contain a zero"? Knowing such information can compromise the security of the password.

To reason about the power of such questions it helps to carry out a mental experiment. Imagine two scenarios: In scenario 1, the secret is $s = 123456789$. In scenario 2, the secret is $s = 100000000$. If Alice and Bob can get together and answer the question "Is the secret an even number?", they must be able to *distinguish* between these two scenarios.

How can they distinguish? They have to use the information at hand, namely the values of their shares $p(1)$ and $p(2)$. Recall that these shares were obtained by setting a_0 to s, choosing (a_1, a_2) at random, and evaluating the polynomial $p(x) = a_2 x^2 + a_1 x + a_0$ at $x = 1$ and $x = 2$.

The ordered pair of values $(p(1), p(2))$ is random; $p(1)$ and $p(2)$ is determined only after we fix the choice of a_1 and a_2. In Theorem 9 we will show that regardless of the value of the secret, any pair of values $(p(1), p(2))$ in $\mathbb{F}_q \times \mathbb{F}_q$ is as likely as any other. For instance, Alice and Bob are equally likely to observe the pair of shares $(3056292003, 1111111111)$ as they are to observe the pair $(987654321, 123123123)$. Therefore the pair of values $(p(1), p(2))$ does not reveal any information to distinguish whether the secret was 123456789 or 100000000.

Theorem 9. *Assume the sequence $(a_1, \ldots, a_d)$ in $\mathbb{F}_q^d$ was chosen randomly so that every possible sequence of q^n values is equally likely. Then for every $a_0 \in \mathbb{F}_q$ and every d distinct inputs $x_1, \ldots, x_d$, the sequence of values $(p(x_1), \ldots, p(x_d))$ takes all q^d values equally likely.*

It follows that for every possible secret $s = a_0$, if Charlie refuses to cooperate, the sequence $(p(1), p(2))$ consisting of Alice's and Bob's shares is equally likely to take on any of the q^2 possible values. The same is true if Alice or Bob refuse to cooperate. So observing the shares of any two parties does not reveal any information about the secret.

Proof of Theorem 9. For every choice of a_0, let $f\colon \mathbb{F}_q^n \to \mathbb{F}_q^n$ be the function that maps $(a_1, \ldots, a_d)$ to $(p(x_1), \ldots, p(x_d))$, where $p(x) = a_d x^d + \cdots + a_1 x + a_0$. We will show that f is a bijective function. Since all inputs of f are equally likely, its outputs must also then be equally likely.

f is surjective by the existence part of Theorem 8: For any set of values $y_1, \ldots, y_d$, there exists a polynomial p of degree d such that $p(x_i) = y_i$ for $i = 1$ up to d. f maps the coefficients of this polynomial to $(y_1, \ldots, y_d)$.

f is injective by the uniqueness part of Theorem 8: For every sequence of values $(y_1, \ldots, y_d)$, the polynomial p whose coefficients f maps to this sequence is unique. □

I wrote a computer program that implements this secret sharing scheme. Feel free to play with it.

References

The secret sharing scheme described here was first described by Adi Shamir in the article *How to share a secret?* Secret sharing is usually studied in the context of cryptography, a vast subject concerned with the study of secure communication and computation in insecure environments.