DISCRETE LOGARITHMS AND THE DIFFIE-HELLMAN EXCHANGE

MATH 195

[For more information, see §6.4 (and Exercise 6.19), §7.7, and §6.5 in the text.]

Notation

The following are examples of what we will talk about: Multiplicative groups, $(\mathbb{Z}/n\mathbb{Z})^*$ and $\mathbb{F}_q^*$ (an example of both is $(\mathbb{Z}/p\mathbb{Z})^*$), the group of points on an elliptic curve E.

Let G be a finite group. Write G multiplicatively. We may as well assume that G is abelian. Let $g \in G$. Look at the powers of g:

$$g^0 = 1,\ g^1 = g,\ g^2 = g \cdot g,\ \ldots,\ g^n = g^{n-1} \cdot g.$$

For example, taking $G = \mathbb{F}_{17}^*$, $g = 2$, we have the sequence $1, 2, 4, 8, 16, 15, 13, 9, 1$, and then it repeats.

Suppose m is the smallest positive integer such that $g^m = g^i$ for some i, $0 \leq i < m$. (In fact, one has $i = 0$ since otherwise one would have $g^{m-1} = g^{i-1}$, $0 \leq i - 1 < m - 1$.) The number m is called the *order* of g, notation: $m = \mathrm{ord}(g)$.

We also have $g^{-1} = g^{m-1}$, and in general, if $i, j \in \mathbb{Z}$ then $g^i = g^j$ if and only if $i \equiv j \pmod{m}$. For example, $2^{311} = 2^7 = 9$ in $\mathbb{F}_{17} = \mathbb{Z}/17\mathbb{Z}$, since $311 \equiv 7 \pmod{8}$.

Notation: $\langle g \rangle = \{g^n : n \in \mathbb{Z}\}$ is a subset of G of precisely m different elements. It is an abelian group, with $g^i g^j = g^{i+j} = g^j g^i$ and $(g^i)^{-1} = g^{-i}$. Such a group is called *cyclic*.

If $h \in \langle g \rangle$, then the *discrete logarithm* of h to the base g is the unique integer $i \pmod{m}$ such that $h = g^i$. Notation: $i = \log_g h = \mathrm{ind}_g h$. If $h \in G$ but $h \notin \langle g \rangle$ then the discrete logarithm is not defined.

In general, we have that $\mathrm{ord}(g) \mid \#G$. This follows from the theorem of Lagrange: the cyclic group generated by g is a subgroup of G, so its order (m) must divide the order of the group. Or, as above, we have that $g^{\#G} = 1 = g^0$, so $\#G \equiv 0 \pmod{m}$, i.e. $m = \mathrm{ord}(g) \mid \#G$.

In particular, we see that the groups

$$\langle g \rangle \simeq \mathbb{Z}/m\mathbb{Z}$$
$$h \mapsto \log_g h$$

are isomorphic by the logarithm map. This map is a homomorphism because

$$\log_g(h_1 h_2) = \log_g(h_1) + \log_g(h_2),$$

which is the statement that $h_1 h_2 = g^i g^j = g^{i+j}$. The inverse map is given by $i \mapsto g^i$. The logarithm map also has the property that $\log_g g = 1$, $\log_g 1 = 0$,

---

This is some of the material covered April 11–16, in Math 195: Cryptography, taught by Hendrik Lenstra, prepared by John Voight `jvoight@math.berkeley.edu`.

$\log_g(h^{-1}) = -\log_g h$, and $\log_g(h^n) = n \log_g h$, just as with the usual logarithm defined on real numbers. (One also has $\log_g a = (\log_h a) \log_g h$—this multiplication is in $\mathbb{Z}/m\mathbb{Z}$—but we usually consider a fixed base.)

We now define the discrete logarithm problem for the group G. Given $g, h \in G$, decide whether $\log_g h$ is defined and if so, compute it.

The point is: computing the logarithm map is difficult for certain choices of G. Or put another way, for discrete logarithm based systems, we need G and g (the base) to be such that $\log_g$ is difficult to compute.

Finite Fields

First, we begin with the example of a finite field.

Theorem. *If k is any finite field (e.g. $k = \mathbb{F}_p$, p prime), then there exists $g \in k^*$ such that $\langle g \rangle = k^*$.*

Equivalently, the order of g is $m = (\#k) - 1$, or $\log_g$ is defined on all of k^*. Such an element is called a *primitive root* of k (or modulo p if $k = \mathbb{F}_p$).

Indeed, if we take 3 in $\mathbb{F}_{17}^*$, we compute:

$$1, 3, 9, 10, 13, 5, 15, 11, 16, 14, 8, 7, 4, 12, 2, 6, 1$$

Notice that $3^8 = 16 = -1$, so the last 8 elements are the negatives of the first 8. Notice that this is a crazy ordering of these elements, and so to compute $\log_3 7 = 11$ we can do no better than locate this element in the table.

Question: Given a (big) prime number p, how does one find a primitive root modulo p? This question is very difficult and gives rise to a number of unsolved problems in number theory. One may ask instead: given p a prime number and $g \in \mathbb{F}_p^*$, how does one quickly decide whether or not g is a primitive root? This is still too hard:

Theorem. *Suppose k is a finite field, $\#k = q$, and let $g \in k^*$. Then: g is a primitive root for k if and only if for each prime number $\ell \mid q-1$, one has $g^{(q-1)/\ell} \neq 1$.*

For $k = \mathbb{F}_{17}$, $q = 17$, $q - 1 = 16$, so we need only check $\ell = 2$, so g is a primitive root if and only if $g^8 \neq 1$. So checking $3^8 = -1 \pmod{17}$, we see that this is a primitive root without doing any other work.

If p is a prime number, for which the complete prime factorization of $p - 1$ is known, then there is a fast algorithm for deciding if an element is a primitive root modulo p. More generally, if p is such a prime number, then one can quickly compute the order of a given element of $\mathbb{F}_p^*$.

There are several variations on this theme:

- We also may take $\mathbb{F}_p$ replaced by any finite field k (and conditions ensuring that the discrete logarithm is hard in k^*).
- One can also take $\langle g \rangle = k^*$ replaced by the condition that $m = \mathrm{ord}(g)$ is a *large* divisor of $\#k^* = q - 1$, where $q = \#k$, i.e. $(q - 1)/m$ should be small.)
- Without explaining what this means, one can also take $G = E(\mathbb{F}_p)$ (or $E(k)$), the group of *rational points* on an *elliptic curve* E defined over $\mathbb{F}_p$ (or k).

Diffie-Hellman Key Exchange

Alice and Bob wish to agree on a common secret, a key, which for example can be used as the basis for exchanging messages using a cryptosystem requiring such a key. To set up a common key, A and B proceed as follows. They agree on a group G of order m and an element $g \in G$ to be used (e.g. $G = \mathbb{F}_p^*$, p a big prime, g a primitive root). Note: G and g are public knowledge (the eavesdropper Eve knows them).

Now A picks a *secret* number $a \pmod m$ (unknown to everyone other than A) and similarly B picks a secret number $b \pmod m$. A computes g^a (so that g^a is public knowledge); but assuming that the discrete logarithm problem is difficult, it is impossible to find a from g^a. Similarly, B computes g^b and sends it to A.

Then A computes $(g^b)^a = h$. Alice is the only person which knows a, so she is the only one who can do this computation. Similarly, B computes $(g^a)^b$, also equal to h. (Note $h = g^{ab}$.) Now h is the common key.

Eve knows G, g, g^a, g^b and would like to compute $h = g^{ab}$. Can Eve do this? Yes, if she can do the discrete logarithm (given g^a, then she could compute a). What if she cannot solve the discrete log? This is an unsolved problem.