Number Theory in Function Fields - Notes (by Evan Dummit, 2025, v. 0.30)

Contents

0 Number Theory in Function Fields

These are lecture notes for the graduate course Math 7360: Number Theory in Function Fields, taught at Northeastern in Fall 2025.

0.1 (Sep 3) Overview + Fermat's Last Theorem for Polynomials

- The goal of this course is to elucidate some of the many analogies between number theory in number fields and number theory in function fields.
  - ○ Some things from classical number theory: primes, factorizations, congruences and modular arithmetic, Fermat's and Euler's theorems, the prime number theorem, quadratic reciprocity (and higher reciprocity), Dirichlet's theorem on primes in arithmetic progressions, zeta functions.
  - ○ Some things from the more modern take on algebraic and analytic number theory: algebraic number fields and their rings of integers, Galois theory and its interplay with number fields, discriminants, class groups, Dirichlet's unit theorem, cyclotomic fields, ramification, L-functions, the Riemann hypothesis.
  - ○ Our goal is to do as much of these things as possible in the context of function fields, where many of the results are more approachable, because the function-field setting has a major kit of additional tools (namely, algebraic geometry).
  - ○ Though do note: number theory in function fields is a beautiful subject in its own right, and not just because it has so many similarities to algebraic number theory.
  - ○ We will illustrate how things can become simpler by proving Fermat's Last Theorem, which is quite notoriously difficult over $\mathbb{Z}$, for polynomials using only elementary techniques.
- To start, let $q = p^f$ be a prime power, and let $\mathbb{F}_q$ be the finite field with q elements. The story begins with the polynomial ring $A = \mathbb{F}_q[t]$.
  - ○ We have the degree map on A: explicitly, for coefficients $a_i \in \mathbb{F}_q$ and an element $f = a_0 + a_1 t + \cdots + a_n t^n$ with $a_n \neq 0$, we define $\deg(f) = n$ and $\operatorname{sgn}(f) = a_n$. (We also set $\deg(0) = -\infty$ and $\operatorname{sgn}(0) = 0$.)

  - Exercises (trivial): $\deg(fg) = \deg(f) + \deg(g)$, $\mathrm{sgn}(fg) = \mathrm{sgn}(f)\mathrm{sgn}(g)$, and $\deg(f+g) \leq \max(\deg f, \deg g)$ with equality whenever $\deg f \neq \deg g$.
  - The polynomials with sign 1 (i.e., monic polynomials) behave analogously to the integers with positive sign (i.e., the positive integers).
  - We also note that the degree properties easily give a characterization of the units of A: they are the nonzero constant polynomials.

- Our first basic result is the standard division-with-remainder algorithm for polynomials, which we record over arbitrary fields for no extra cost:

- Exercise (Polynomial Division): If F is any field, then for any $f, g \in F[t]$ with $g \neq 0$, there exist unique $q, r \in F[t]$ such that $f = qg + r$ and $\deg r < \deg g$.
  - The idea is simply to prove that the usual long-division algorithm works by induction on the degree of g.
  - As a consequence, $F[t]$ is a Euclidean domain, meaning that it is also a principal ideal domain (all ideals are principal) and a unique factorization domain (every element can be factored uniquely into a product of irreducibles up to reordering and unit factors).

- As it turns out, unique factorization is essentially all we need to prove Fermat's Last Theorem for polynomials.
  - We would like to show that the equation $f^n + g^n = h^n$ has no nontrivial solutions in polynomials f, g, h. Aside from the case $n = 4$, it is enough to treat the situation where n is a prime.
  - But we do need to be a little bit careful to write down exactly what the trivial solutions look like, beyond the obvious ones where one of f, g, h is zero.
  - For example, if f, g, h are all constants, we can certainly have lots of solutions to $f^n + g^n = h^n$, depending on the field and on n (e.g., $1^5 + 1^5 = 2^5$ inside $\mathbb{F}_3$).
  - We need to avoid the situation where n is divisible by $p = \mathrm{char}(\mathbb{F}_q)$, since $f^p + g^p = (f+g)^p$ for any polynomials $f, g \in \mathbb{F}_q[t]$.
  - Also, since the equation is homogeneous, we can scale solutions to get new solutions.
  - To avoid all of these situations, we can consider only the case where f, g, h are relatively prime (since if they are not, then any common divisor of two of them also divides the third, so we could cancel it) and where the exponent n is not divisible by the characteristic p.

- Theorem (FLT for Polynomials): Suppose that $f, g, h \in F[t]$ are pairwise relatively prime and that $p \geq 3$ is prime with $p \neq \mathrm{char}(F)$. Then the only solutions to $f^p + g^p = h^p$ are when f, g, h are all constants.
  - We will remark that $p \geq 3$ is needed, since the usual parametrization of Pythagorean triples also works for polynomials: if we take $f = a^2 - b^2$, $g = 2ab$, $h = a^2 + b^2$ for any polynomials $a, b \in F[t]$, then $f^2 + g^2 = h^2$.

- We will give two different proofs: the first uses a classical-style infinite descent argument, while the second uses a more function-field type of argument.
  - Proof 1: Without loss of generality, we may assume that F is algebraically closed, since any solution to $f^p + g^p = h^p$ over F is still a solution over the algebraic closure $\overline{F}$.
  - We show the result by inducting on $d = \deg f + \deg g$. The base case $d = 0$ is trivial, since there is nothing to prove. So now suppose we have a solution with $d > 0$.
  - By the assumption that $p \neq \mathrm{char}(F)$, there are p distinct pth roots of unity in F: say, $1, \zeta_p, \zeta_p^2, \ldots, \zeta_p^{p-1}$, and we can factor $f^p + g^p = (f+g)(f+\zeta_p g)(f+\zeta_p^2 g)\cdots(f+\zeta_p^{p-1}g)$.
  - Next, note that all of the terms $f + \zeta_p^i g$ are relatively prime: if e divides both $f + \zeta_p^i g$ and $f + \zeta_p^j g$, then e also divides the difference $(\zeta_p^i - \zeta_p^j)g$ hence divides g, hence also divides $(f + \zeta_p^i g) - \zeta_p^i g = f$, but f and g are relatively prime by assumption.
  - Then by unique factorization inside $F[t]$, since all of the terms in the product $(f+g)(f+\zeta_p g)(f+\zeta_p^2 g)\cdots(f+\zeta_p^{p-1}g)$ are relatively prime and their product is a pth power (namely, h^p), each term must be a pth power up to a unit factor. But since F is algebraically closed, everything in F has a pth root in F, so the unit factor is also a pth power.

  - Thus, in particular, we see that $f+g=a^p$, $f+\zeta_p g=b^p$, and $f+\zeta_p^2 g=c^p$ are all pth powers.
  - Using basic linear algebra to eliminate f and g yields the relation $-\zeta_p a^p+(1+\zeta_p)b^p=c^p$, so if we set $a'=(-\zeta_p)^{1/p}a$, $b'=(1+\zeta_p)^{1/p}b$, and $c'=c$, then we have $(a')^p+(b')^p=(c')^p$.
  - Note that a', b' cannot both be constant, since then f, g would have been constant. But we also have $\deg(a')+\deg(b')=\deg(f+g)/p+\deg(f+\zeta_p g)/p \le 2\max(\deg f, \deg g)/p < d$, so we have constructed a solution with smaller positive degree, but this contradicts the induction hypothesis. Therefore, there are no nonconstant solutions.
  - Exercise: For any field F of characteristic p, we have exhibited nontrivial polynomial solutions to $f^p+g^p=h^p$ in $F[t]$. Where and why in the proof of FLT above does the argument break down when $\mathrm{char}(F)=p$?

- Before giving the second proof, we need a few preliminary results.
  - First, if f has prime factorization $f=\prod_i p_i^{a_i}$, define $\mathrm{rad}(f)=\prod_i p_i$, the product of the monic irreducible polynomials dividing f.

- Lemma: We have $\deg\gcd(f,f') \ge \deg f-\deg\mathrm{rad} f$, where f' is the derivative of f.
  - Proof: Suppose $f=p^a q$ where p is irreducible and doesn't divide q. Then $f'=ap^{a-1}p'q+p^a q'=p^{a-1}(ap'q+pq')$ is divisible by p^{a-1}. Therefore, $\gcd(f,f')$ is divisible by p^{a-1}.
  - Taking the product over all primes dividing f shows that $\prod_i p_i^{a_i-1}$ divides $\gcd(f,f')$, so $\gcd(f,f')\cdot\mathrm{rad}(f)$ is divisible by $\prod_i p_i^{a_i-1}\prod_i p_i=\prod_i p_i^{a_i}=f$, so taking degrees yields the inequality.
  - Exercise: Determine when equality holds, namely when $\deg\gcd(f,f')=\deg f-\deg\mathrm{rad} f$.

- Next, we show a result due independently to Mason and Stothers:

- Proposition (Mason-Stothers): Suppose that $f,g,h\in F[t]$ are nonconstant, relatively prime, that $f+g=h$, and that not all of f', g', h' are zero. Then $\max(\deg f, \deg g, \deg h) \le \deg\mathrm{rad}(fgh)-1$.
  - Proof: If $f+g=h$ then $f'+g'=h'$, and then $fg'-f'g=(f+g)g'-(f'+g')g=hg'-h'g$.
  - Note also that $fg'-f'g$ is nonzero: if $fg'=f'g$ then f must divide $f'g$ hence that f must divide f' since f,g are relatively prime.
  - Exercise: Suppose $f\in F[t]$. Show that f divides its derivative f' if and only if $f'=0$.
  - By the exercise we see then that $f'=0$. But now by the same argument we would also have $g'=0$ and $h'=0$, contradicting the assumption that not all of f', g', h' are zero.
  - Now let $d_f=\gcd(f,f')$, $d_g=\gcd(g,g')$, $d_h=\gcd(h,h')$. Then d_f, d_g, d_h all divide $fg'-f'g=hg'-h'g$, and they are all relatively prime since they are divisors of the relatively prime polynomials f,g,h.
  - This means $d_f d_g d_h$ divides $fg'-f'g$, so taking degrees yields $\deg(d_f d_g d_h)\le\deg(fg'-f'g)\le\deg(f)+\deg(g)-1$.
  - By the Lemma, we have $\deg(d_f)\ge\deg(f)-\deg\mathrm{rad} f$, $\deg(d_g)\ge\deg(g)-\deg\mathrm{rad} g$, $\deg(d_h)\ge\deg(h)-\deg\mathrm{rad} h$, so summing yields $\deg(f)+\deg(h)+\deg(h)-\deg\mathrm{rad}(fgh)\le\deg(d_f d_g d_h)\le\deg(f)+\deg(g)-1$, and therefore $\deg(h)\le\deg\mathrm{rad}(fgh)-1$.
  - By rearranging we obtain the same bounds on $\deg(f)$ and $\deg(g)$, and so we are done.

- At last, we can finish the second proof of Fermat's Last Theorem for polynomials:
  - Proof 2: Suppose $f^p+g^p=h^p$. By the assumption on the characteristic, we have $(f^p)'$, $(g^p)'$, $(h^p)'$ are not all zero.
  - Then by Mason-Stothers, we see $\max(\deg f^p, \deg g^p, \deg h^p)\le\deg\mathrm{rad}(f^p g^p h^p)-1$, which is equivalent to $p\cdot\max(\deg f, \deg g, \deg h)\le\deg\mathrm{rad}(fgh)-1\le\deg f+\deg g+\deg h-1$ since the radical ignores powers.
  - Now apply the simple observation that $\max(a,b,c)\ge(a+b+c)/3$ and set $d=\deg f+\deg g+\deg h$ to see that $p\cdot d/3\le d-1$, which is impossible, since $d\le p\cdot d/3$ by the hypothesis that $p\ge 3$.

0.2 (Sep 8) Quotients of $\mathbb{F}_q[t]$

- We now return to study the structure of quotient rings of $A = \mathbb{F}_q[t]$, which (re-posed) is simply studying modular arithmetic in this ring.
  - In particular, we will recover almost identical versions of Fermat's little theorem, Euler's theorem, and Wilson's theorem.
  - We will also take some time to look at the structure of the unit group of A/gA, which turns out to be a bit more complicated to write down than the unit group of $\mathbb{Z}/m\mathbb{Z}$.
- As noted last lecture, A is a Euclidean domain, so it is a PID and also a UFD. Since every ideal is principal, if we want to understand the structure of the quotient rings of A, we only have the quotients of the form A/gA to consider.
  - We can also assume g is monic by replacing it with its unique monic associate, which does not change the quotient ring A/gA.
- Using the division algorithm, we can write down the residue classes in A/gA, and in particular compute its cardinality, quite easily:
- Proposition: Let $g \in \mathbb{F}_q[t] = A$ be nonzero. Then the residue classes in A/gA are uniquely represented by the polynomials of degree less than $\deg(g)$. In particular, $\#(A/gA) = q^{\deg g}$.
  - Proof: If $f \in \mathbb{F}_q[t]$ is any polynomial, then by the division algorithm we can write $f = qg + r$, and so inside A/gA we see $\overline{f} = \overline{r}$. So the possible remainders give a complete set of residue class representatives – but by the uniqueness of the quotient and remainder, no two remainders are equivalent mod g, so in fact they give all of the residue classes exactly once.
  - For the counting, if $\deg(g) = n$, then the remainders are of the form $c_0 + c_1 t + \cdots + c_{n-1} t^{n-1}$ with $c_i \in \mathbb{F}_q$. Since there are n coefficients each of which has q possible values, there are $q^n = q^{\deg g}$ possible ways to select a remainder.
- The size of the quotient ring gives a convenient way of measuring the "size" of a polynomial that behaves pleasantly under multiplication:
- Definition: For $g \in \mathbb{F}_q[t]$, we define $|g|$, the norm of g, to be $q^{\deg g}$. By the calculation above, $|g| = \#(A/gA)$ when $g \neq 0$.
  - Exercise: Show $|fg| = |f| \cdot |g|$ and $|f + g| \leq \max(|f|, |g|)$ with equality whenever $|f| \neq |g|$.
- Our next goal is to understand the units of A/gA, since this is the context in which to pose Fermat's and Euler's theorems.
  - Regardless of the polynomial g, the units of A/gA will contain an isomorphic copy of the constant polynomials (i.e., the units of A), which is the multiplicative group $\mathbb{F}_q^*$.
  - As is well-known, the multiplicative group of a finite field is cyclic. We record a few proofs of this fact, for completeness:
- Proposition (Multiplicative Group of $\mathbb{F}_q$): If G is a finite multiplicative subgroup of a field F, then G is cyclic.
  - All known proofs of this fact are essentially nonconstructive, to varying degrees: there does not seem to be a nice algorithm for writing down a multiplicative generator of a finite field that is appreciably better than a brute-force search.
  - Proof 1: Let G be a finite multiplicative subgroup of F. By the fundamental theorem of finite(ly generated) abelian groups, G is isomorphic to a direct product of cyclic groups.
  - Let m be the lcm of the orders of these cyclic groups: then $x^m = 1$ for all $x \in G$. Since $F[t]$ has unique factorization, the polynomial $t^m - 1 \in F[t]$ has at most m roots in F, so $\#G \leq m$. On the other hand, by Lagrange's theorem, the order of every element in G divides $\#G$, so m divides $\#G$. We must therefore have $m = \#G$.

- But since $\#G$ is equal to the product of the orders of the cyclic groups, we see that the product of these orders equals their lcm, so the orders are all relatively prime. This means G is cyclic, as claimed.
- Proof 2: Let M be the maximal order among all elements in G: we claim that the order of every element in G divides M. To see this, suppose g has order M, and let h be any other element of order k. If k does not divide M, then there is some prime q which occurs to a higher power q^f in the factorization of k than the corresponding power q^e dividing M.
- By properties of orders, the element g^{q^f} has order M/q^f, and the element h^{k/q^e} has order q^e. Since these two orders are relatively prime and $gh = hg$ (since these are elements in a field), we see that the element $g^{q^f} \cdot h^{k/q^e}$ has order $M \cdot q^{f-e}$. This is a contradiction because this element's order is larger than M. Thus, k divides M as claimed.
- For the second claim, any element of order M generates a subgroup of G having M elements, so $M \le \#G$.
- Furthermore, by the above, we know that all elements in G have order dividing M, so the polynomial $t^M - 1$ has $\#G$ roots in $F[t]$. By unique factorization, this requires $M \ge \#G$, and so we have $M = \#G$. Now select any element of order M: it generates G.
- Proof 3: Observe by Lagrange's theorem that $t^{\#G} - 1$ factors as the product $\prod_{d|\#G} \Phi_d(t)$, where $\Phi_d(t) = \prod_{\mathrm{order(g)=d}} (t - g)$ is the dth cyclotomic polynomial.
- By an inductive argument, or by observing invariance under the Galois action, all of the polynomials $\Phi_d(t)$ have coefficients in $F[t]$.
- By induction on d using the fact that $t^d - 1$ has at most (hence exactly) d roots in F and in G, one has that $\deg(\Phi_d) = \varphi(d)$. In particular, $\deg(\Phi_{\#G}) = \varphi(\#G) > 0$, so there is an element of order $\#G$ in G.

- Now we tackle the question of the units of A/gA.
  - We can simplify the problem first: if we factor $g = p_1^{a_1} \cdots p_d^{a_d}$ where the p_i are distinct monic irreducible polynomials, then all of the ideals $(p_i^{a_i})$ are pairwise comaximal, so by the Chinese remainder theorem, we see $A/gA \cong (A/p_1^{a_1}A) \times (A/p_2^{a_2}A) \times \cdots \times (A/p_d^{a_d}A)$.
  - Taking units on both sides then gives $(A/gA)^* \cong (A/p_1^{a_1}A)^* \times (A/p_2^{a_2}A)^* \times \cdots \times (A/p_d^{a_d}A)^*$. So it is enough to study the structure of the ring A/p^aA where p is irreducible.

- Proposition (Structure of A/p^aA): For $A = \mathbb{F}_q[t]$ where $\mathrm{char}(\mathbb{F}_q) = \tilde{p}$, and $p \in A$ is a monic irreducible polynomial, we have the following:
  1. The cardinality of $(A/p^aA)^*$ is $\#(A/p^aA)^* = |p|^{a-1}(|p| - 1) = |p^a|(1 - 1/|p|)$.
     - Exercise: Show that a commutative ring R with 1 has a unique maximal ideal M if and only if the set of nonunits in R forms an ideal, which is then a unique maximal ideal M. A ring with this property is called a local ring.
     - Proof: The ring A/p^aA has a unique maximal ideal, namely pA/p^aA, and is therefore a local ring, because the quotient $(A/p^aA)/(pA/p^aA) \cong A/pA$ is a field by the third isomorphism theorem.
     - By the exercise above, evvery element not in the maximal ideal is a unit, and the cardinality of the maximal ideal is $1/|p|$ times the cardinality of the entire ring (since the elements in the ideal are just the multiples of p). The formula follows.
  2. $(A/p^aA)^* \cong$ [cyclic group of order $|p| - 1$] $\times$ [an abelian $\tilde{p}$-group].
     - Proof: The reduction-mod-p map is a surjective group homomorphism from $(A/p^aA)^* \to (A/pA)^*$, and the latter is the multiplicative group of the field A/pA hence is cyclic of order $|p| - 1$.
     - Pulling back a generator yields that $(A/p^aA)^*$ contains a cyclic subgroup of order $|p| - 1$. By the cardinality calculation in (1), the remaining piece has order $|p|^{a-1}$ and is therefore a $\tilde{p}$-group (and it is clearly abelian).
     - Remark: The direct product decomposition writes each element modulo p^a as [its residue modulo p] times [an element congruent to 1 modulo p].
  3. The $\tilde{p}$-part of $(A/p^aA)^*$ has exponent at most $\tilde{p}^s$ where $\tilde{p}^s \ge a$.
     - Proof: By the above, the elements in the $\tilde{p}$-part are of the form $1 + bp$ for some $b \in \mathbb{F}_q[t]$.

    - Since we are in characteristic $\tilde{p}$, we then have $(1+bp)^{\tilde{p}^s} = 1+(bp)^{\tilde{p}^s}$, and since $p^{\tilde{p}^s}$ is divisible by p^a by assumption, we see $(1+bp)^{\tilde{p}^s} \equiv 1 \pmod{p^a}$, which is to say, the element $1+bp$ modulo p^a has order dividing $\tilde{p}^s$ (as required).

4. As $a \to \infty$, the number of cyclic factors in the $\tilde{p}$-part of $(A/p^aA)^*$ goes to infinity.
    - The point here is that we get a different kind of behavior than over $\mathbb{Z}$: over $\mathbb{Z}$, we see that $(\mathbb{Z}/p^a\mathbb{Z}) \cong \begin{cases} \mathbb{Z}/(p^a - p^{a-1})\mathbb{Z} & \text{for odd primes } p \\ (\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/2^{a-3}\mathbb{Z}) & \text{for } p=2 \end{cases}$, and so even for large prime powers, the quotient is either cyclic or basically cyclic.
    - For polynomials, we end up getting a large number of cyclic factors when we take a large power, regardless of the prime.
    - Proof: Since the exponent of the $\tilde{p}$-part is at most $\tilde{p}^s$, if we have a total of j cyclic factors then the order of the group is at most $\tilde{p}^{sj}$. So we need $\tilde{p}^{sj} \geq |p|^{a-1} = q^{\deg(p)\cdot(a-1)} = \tilde{p}^{f\cdot\deg(p)\cdot(a-1)}$ and so $j \geq f \cdot \deg(p) \cdot (a-1)/s$.
    - Since $s \sim \log_p a$, we see that for a fixed field $\mathbb{F}_q$ (i.e., fixed f) and fixed prime p (i.e., fixed $\deg p$), we have $j \sim C(a-1)/\log_p a \to \infty$ as $a \to \infty$.

- Now that we have established some basic things about the unit group of A/p^aA, we can establish the analogues of Fermat's little theorem, Euler's theorem, and Wilson's theorem.
    - First, we need the analogue of the Euler phi-function. We define $\Phi(f) = \#(A/fA)^*$ to be the number of polynomials of degree less than $\deg f$ that are relatively prime to f.
    - By our calculations with the unit group earlier, we have the usual formula $\Phi(f) = |f| \prod_{p|f\text{prime}}(1-1/|p|)$, which is the analogue of $\varphi(n) = n\prod_{p|n \text{ prime}}(1-1/p)$ for the phi-function over $\mathbb{Z}$.

- Proposition ("Euler"): If $f \in \mathbb{F}_q[t]$ is nonzero and g is relatively prime to f, then $g^{\Phi(f)} \equiv 1 \pmod{f}$.
    - Proof 1: Apply Lagrange's theorem to $\overline{g}$ in $(A/fA)^*$.
    - Proof 2: Multiplication by $\overline{g}$ is a bijection on the cosets in $(A/fA)^*$. Thus, $\prod_{u\in(A/fA)^*} u = \prod_{u\in(A/fA)^*}(ug) = g^{\Phi(f)}\prod_{u\in(A/fA)^*} u$ inside $(A/fA)^*$, and cancelling the unit factor $\prod_{u\in(A/fA)^*} u$ yields $g^{\Phi(f)} = 1$ inside $(A/fA)^*$.

- Proposition ("Fermat"): If $p \in \mathbb{F}_q[t]$ is irreducible, then $a^{|p|} \equiv a \pmod{p}$ for any $a \in \mathbb{F}_q[t]$.
    - Proof: If $p|a$ the result is trivial. Otherwise, a is a unit modulo p and the result follows from Euler above.

- We can use the analogue of Fermat's theorem to prove an analogue of Wilson's theorem:

- Proposition (Factoring, 1): If $p \in \mathbb{F}_q[t]$ is irreducible of degree d, then $x^{|p|} - x \equiv \prod_{\deg f<d}(x-f) \bmod p$.
    - Proof: As we have noted, in A/p the polynomials of degree $< d$ represent all of the residue classes modulo p.
    - By Fermat, each of these polynomials is a root of $x^{|p|} - x$. But by unique factorization, this polynomial has at most $|p|$ distinct roots, and we have just exhibited $|p|$ roots, so these are all of the roots, and the factorization follows.

- Corollary ("Wilson"): If $p \in \mathbb{F}_q[t]$ is irreducible of degree d, then $\prod_{\deg f<d, f\neq 0} f \equiv -1 \pmod{p}$.
    - Proof 1: Dividing the result above by x yields $x^{|p|-1} - 1 \equiv \prod_{\deg f<d, f\neq 0}(x-f) \bmod p$.
    - Now set $x=0$: if the characteristic is odd, then the number of minus signs on the RHS is even and the result follows, while if the characteristic is even, then $1=-1$ so the result still follows.
    - Proof 2: If $\overline{f}$ does not have order 2 in A/pA, then $\overline{f} \neq \overline{f}^{-1}$ and so we can pair up and discard $(\overline{f}, \overline{f}^{-1})$ without affecting the product.
    - When we have done this for all possible pairs, the only elements left are the elements of order dividing 2 (i.e., the solutions to $x^2 = 1$), which are $x = \pm 1$. In characteristic not 2, the product is -1, while in characteristic 2, the product is $1 = -1$.

- Exercise: Generalize proof 2 of Wilson's theorem to show that if G is a finite abelian group, then the product of all elements in g is the unique element in G of order 2 (if there is one), or is otherwise the identity.

- We also record a useful result about roots of unity:

- Proposition (Roots of Unity): If $p \in \mathbb{F}_q[t] = A$ is irreducible and d divides $|p| - 1$, then there are d dth roots of unity in A/pA; equivalently, $x^d \equiv 1 \pmod{p}$ has exactly d solutions.
  - Exercise: For positive integers a, b, show $\gcd(x^a - 1, x^b - 1) = x^{\gcd(a,b)} - 1$ in $F[x]$.
  - Proof: As shown above, $x^{|p|-1} - 1$ splits completely mod p . By the exercise, $x^d - 1$ divides $x^{|p|-1} - 1$ when d divides $|p| - 1$, and so $x^d - 1$ also splits completely, which is to say, it has d roots mod p.
  - Exercise: Prove the converse: if there are d dth roots of unity in A/pA, then d divides $|p| - 1$.

0.3 (Sep 10) Prime-Counting and The Zeta Function

- Now that we have established most of the classical results for modular arithmetic, we move to our next item: counting primes.
  - We will do things in a more ad hoc manner first, and then give a more general approach using zeta functions that will allow us to go further.

- Our first step is to write down a generalization of the fact we used to establish Wilson's theorem above:

- Theorem (Factoring, II): For a positive integer m, the polynomial $t^{q^m} - t$ factors in $\mathbb{F}_q[x]$ as the product of all monic irreducible polynomials of degree dividing m.
  - Proof 1 ("Elementary"): We will show that $t^{q^m} - t$ has no repeated factors, that each of the claimed polynomials does divide it, and that no other polynomials divide it.
  - Exercise: A polynomial in $F[t]$ is separable (i.e., has no repeated factors) if and only if it is relatively prime to its derivative.
  - Since $(t^{q^m} - t)' = q^m t^{q^m - 1} - 1 = -1$ in characteristic p, the polynomial is relatively prime to its derivative, so it has no repeated factors by the exercise.
  - Exercise: For positive integers q, a, b, show that $\gcd(q^a - 1, q^b - 1) = q^{\gcd(a,b)} - 1$ in $\mathbb{Z}$. (This is almost identical to the polynomial version mentioned earlier.)
  - Next, suppose p is irreducible of degree dividing m. If $p = t$ the result is trivial, and otherwise, in A/pA we have $t^{q^m - 1} \equiv 1$ mod p because $q^m - 1$ is a multiple of $|p| - 1 = q^{\deg p} - 1$ by the exercise above along with Euler's theorem. This means $t^{q^m - 1} - 1$ is divisible by p as required.
  - Finally, suppose p is irreducible of degree not dividing m. Then in A/pA we have $t^{q^m - 1} \equiv t^{q^{\gcd(m, \deg p)}} \neq 1$ mod p by the exercise above along with Euler's theorem and the fact that $q^{\gcd(m, \deg p)} < q^{\deg p}$. This means $t^{q^m - 1} - 1$ is not divisible by p as required.
  - We have shown that $t^{q^m} - t$ has no repeated factors, that each of the claimed polynomials does divide it, and that no other polynomials divide it. Since the polynomial is monic, its factorization must therefore be as claimed.
  - Proof 2 ("Galois"): By basic Galois theory, $\mathrm{Gal}(\mathbb{F}_{q^m}/\mathbb{F}_q)$ is a cyclic group of order m generated by the Frobenius map $x \mapsto x^q$[1].
  - By the Galois correspondence, the intermediate fields of $\mathbb{F}_{q^m}/\mathbb{F}_q$ are $\mathbb{F}_{q^d}$ for $d|m$. Therefore, p is irreducible of degree dividing $d \iff \mathbb{F}_q[t]/(p)$ is (isomorphic to) an intermediate field of $\mathbb{F}_{q^m}/\mathbb{F}_q \iff p$ divides $x^{q^m} - x$.
  - Since $x^{q^m} - x$ is separable, its factorization must therefore be as claimed.

[1] This follows by noting that $\mathbb{F}_{q^m}$ is the splitting field of $x^{q^m} - x$ over $\mathbb{F}_q$ and since this polynomial is separable as noted in proof 1, the order of the Galois group is m. The Frobenius map is an injective field map from $\mathbb{F}_{q^m}$ to itself, hence an automorphism by finiteness, and its order is clearly at least m (since $x^{q^d} - x$ has at most q^d solutions) and at most m (by Lagrange).

- Corollary: If a_d is the number of irreducible monic polynomials in $A = \mathbb{F}_q[t]$ of degree d, then $\sum_{d|n} da_d = q^n$.
  - Proof: Count degrees in the theorem above.
- We can use this recurrence to write down an exact formula for a_d using Mobius inversion.
- Definition: The Mobius μ-function is defined as $\mu(n) = \begin{cases} 0 & \text{if } n \text{is not squarefree} \\ (-1)^r & \text{if } n \text{is the product of } r \text{distinct primes} \end{cases}$. Note $\mu(1) = 1$.
  - Exercise: Show that $\sum_{d|n} \mu(d) = \begin{cases} 1 & \text{for } n = 1 \\ 0 & \text{for } n > 1 \end{cases}$.
- Proposition (Mobius Inversion): If f, n are integer functions such that $g(n) = \sum_{d|n} f(d)$, then $f(n) = \sum_{d|n} \mu(d) g(n/d)$.
  - Proof: Induct on n. The base case $n = 1$ is trivial.
  - For the inductive step, we have $\sum_{d|n} \mu(d) g(n/d) = \sum_{d|n} \mu(d) \cdot \sum_{d'|n/d} f(d') = \sum_{dd'|n} \mu(d) f(d') = \sum_{d'|n} f(d') \sum_{d|(n/d')} \mu(d) = f(n)$ because the last inner sum is zero except for when $n/d' = 1$.
- By using Mobius inversion on the sequence $\{da_d\}$, we can write down formulas for the number of monic irreducible polynomials of degree d.
- Proposition (Prime Counting): If a_n is the number of monic irreducible polynomials in $\mathbb{F}_q[t]$ of degree n, then $a_n = \frac{1}{n} \sum_{d|n} \mu(d) q^{n/d}$.
  - The first few values are $a_1 = q$, $a_2 = \frac{1}{2}(q^2 - q)$, $a_3 = \frac{1}{3}(q^3 - q)$, $a_4 = \frac{1}{4}(q^4 - q^2)$, $a_5 = \frac{1}{5}(q^5 - q)$, $a_6 = \frac{1}{6}(q^6 - q^3 - q^2 + q)$, ...
  - Proof: Immediate from applying Mobius inversion to the sequence $\{na_n\}$.
- We can also do some basic asymptotic analysis using the formula above.
  - The main term is $\frac{1}{n} q^n$, and then the next biggest possible term is $\frac{1}{n} q^{n/2}$, so we see that $a_n = \frac{1}{n} q^n + O(q^{n/2}/n)$.
  - If we write $X = q^n$ (which is the total number of monic polynomials of degree n), we see that the number of "primes" in A of "size" $\sim X$ is $a_n = \frac{X}{\log_q X} + O(\frac{\sqrt{X}}{\log_q X})$.
  - This is quite in the spirit of the prime number theorem over $\mathbb{Z}$, which says that the number of primes $\leq X$ is $\Pi(X) = \frac{X}{\log X} + O(\frac{X}{(\log X)^2})$. If we replace $X/\log X$ with the logarithmic integral $\text{li}(x) = \int_2^x \frac{dt}{\log t}$, then as shown by von Koch, the Riemann hypothesis is equivalent to the error estimate $\Pi(X) = \text{li}(x) + O(\sqrt{X} \log x)$.
  - Qualitatively, then, we have already obtained a prime-counting result that is closely analogous to the best possible one predicted by the Riemann hypothesis.
- Up until this point, our approach has been purely algebraic. However, by introducing analytic methods, we can give even easier solutions to these (and other) counting problems. The necessary object of study is the zeta function, which we now define:
- Definition: For $A = \mathbb{F}_q[t]$, the zeta function of A is $\zeta_A(s) = \sum_{f \in A \text{ monic}} \frac{1}{|f|^s}$ for $s \in \mathbb{C}$.
  - Compare to the definition of the Riemann zeta function $\zeta(s) = \sum_{n>0} \frac{1}{n^s}$ for $s \in \mathbb{C}$.

  - Unlike the Riemann zeta function, however, we can actually just evaluate the zeta function for A: since there are q^d monic polynomials of degree d, we see that $\sum_{\deg(f)\le d \text{ monic}} \frac{1}{|f|^s} = 1 + \frac{q}{q^s} + \frac{q^2}{q^{2s}} + \cdots + \frac{q^d}{q^{ds}} = \frac{1-q^{(d+1)(1-s)}}{1-q^{1-s}}$, and so taking $d \to \infty$ we see that $\zeta_A(s) = \frac{1}{1-q^{1-s}}$ whenever $\mathrm{Re}(s) > 1$ (to ensure convergence).
  - We have an obvious meromorphic continuation for $\zeta_A(s)$ to the complex plane (i.e., via the formula above), and it is clear that ζ is analytic everywhere except for a simple pole at $s = 1$.
  - Exercise: Show that the residue of $\zeta_A(s)$ at $s = 1$ (which is to say, the value of $\lim_{s\to 1}(s-1)\zeta_A(s)$) is $1/\log q$.
  - We also have a functional equation for $\zeta_A(s)$: if we set $\xi_A(s) = q^{-s}(1-q^{-s})^{-1}\zeta_A(s)$, then $\xi_A(s) = \xi_A(1-s)$.
  - Exercise: Do the algebra to establish the functional equation.

- We can also represent $\zeta_A(s)$ as an Euler product, just as with the Riemann zeta function.
  - Explicitly, by the uniqueness of prime factorization, we can formally write $\zeta_A(s) = \sum_{f\in A \text{ monic}} \frac{1}{|f|^s} = \prod_{p\text{monic irred}}(1 + \frac{1}{|p|^s} + \frac{1}{|p|^{2s}} + \cdots) = \prod_{p\text{monic irred}}(1 - 1/|p|^s)^{-1}$, and both sides are absolutely convergent for $\mathrm{Re}(s) > 1$.
  - To prove this equality rigorously, we need to do some estimations on tails of the respective series, but since everything converges absolutely, this is not so difficult; we leave the precise details as an exercise.

- We can use the Euler product for the zeta function to obtain the same prime counts that we got earlier.

- Proposition (Prime Counting, Again): If a_d is the number of irreducible monic polynomials in $A = \mathbb{F}_q[t]$ of degree d, then $\sum_{d|n} da_d = q^n$, and so by Mobius inversion as before, we see $a_n = \frac{1}{n}\sum_{d|n}\mu(d)q^{n/d}$.
  - Proof: Group the terms in the Euler product together by degree: if $\deg p = d$ then $|p|^s = q^{ds}$.
  - Thus, since there are a_d monic irreducibles of degree d by definition, we see that $\zeta_A(s) = \prod_{p\text{monic irred}}(1 - 1/|p|^s)^{-1} = \prod_{d=1}^{\infty}(1-q^{-ds})^{-a_d}$.
  - Noting from earlier that $\zeta_A(s) = \frac{1}{1-q^{1-s}}$, if we substitute $u = q^{-s}$, we obtain the equality $\frac{1}{1-qu} = \prod_{d=1}^{\infty}(1-u^d)^{-a_d}$.
  - Taking the log-derivative of both sides yields $\frac{q}{1-qu} = \sum_{d=1}^{\infty}\frac{da_d u^{d-1}}{1-u^d}$. These expressions are equal as power series in u, and thus corresponding coefficients must also be equal.
  - The LHS is $\frac{q}{1-qu} = q\sum_{k=0}^{\infty}(qu)^k$ while the RHS is $\sum_{d=1}^{\infty} da_d u^{d-1}\sum_{l=0}^{\infty} u^{dl} = \sum_{d=1}^{\infty}\sum_{l=0}^{\infty} da_d u^{d(l+1)-1}$. So the coefficient of u^{n-1} on the LHS is $q \cdot q^{n-1} = q^n$, while the coefficient of u^{n-1} on the RHS is $\sum_{d(l+1)=n} da_d = \sum_{d|n} da_d$.
  - Thus, $q^n = \sum_{d|n} da_d$ as claimed.

- Of course, we have already proven this result by counting irreducible polynomials algebraically. However, this approach using the zeta function also extends to solve other counting problems quite conveniently.

- Proposition (Squarefree Counting): The number of monic squarefree polynomials of degree n over $\mathbb{F}_q[t]$ is equal to $b_n := q^n - q^{n-1}$. Equivalently, a randomly-chosen degree-n polynomial is squarefree with probability $1 - 1/q = 1/\zeta_A(2)$.
  - Compare this result to the corresponding fact about integers (which is a little harder to pose because we have to phrase it over a range): if α_n is the probability that a randomly-chosen integer in $[1, n]$ is squarefree, then $\lim_{n\to\infty}\alpha_n = 6/\pi^2 = 1/\zeta(2)$.

- ◦ Proof: Consider the product $\pi = \prod_{p \text{ monic irred}} (1 + |p|^{-s})$.
- ◦ By multiplying out the terms, we see that for $\text{Re}(s) > 1$, we have $\pi = \sum_{f \text{ monic}} \dfrac{\delta(f)}{|f|^s}$ where $\delta(f) = \begin{cases} 1 & \text{if } f \text{is squarefree} \\ 0 & \text{if } f \text{is not squarefree} \end{cases}$, since the denominators in the Euler product only include prime factors of exponents 0 and 1.
- ◦ Now, since $1 + |p|^{-s} = \dfrac{1 - |p|^{-2s}}{1 - |p|^{-s}}$, taking the product over monic irreducibles and using the fact that the resulting numerator and denominator products converge absolutely allows us to write $\pi = \prod_{p \text{ monic irred}} \dfrac{1 - |p|^{-2s}}{1 - |p|^{-s}} = \dfrac{\prod_{p \text{ monic irred}} 1 - |p|^{-2s}}{\prod_{p \text{ monic irred}} 1 - |p|^{-s}} = \dfrac{\zeta_A(2s)}{\zeta_A(s)}$.
- ◦ Setting $u = q^{-s}$ yields $\dfrac{1 - qu^2}{1 - qu} = \dfrac{\zeta_A(2s)}{\zeta_A(s)} = \pi = \sum_{f \text{ monic}} \dfrac{\delta(f)}{|f|^s} = \sum_{n=0}^{\infty} b_n u^n$.
- ◦ But as a power series in u, we have $\dfrac{1 - qu^2}{1 - qu} = (1 - qu^2)(1 + qu + q^2u^2 + \cdots)$, and so comparing coefficients yields $b_n = q^n - q^{n-1}$ as claimed.

- In a similar way, we can use the zeta function to write down formulas for the number of monic kth-powerfree polynomials of a given degree over $\mathbb{F}_q[t]$.
  - ◦ Specifically, these values are packaged as the coefficients in the Euler product $\prod_{p \text{ monic irred}} (1 + |p|^{-s} + |p|^{-2s} + \cdots + |p|^{-(k-1)s}) = \dfrac{\zeta_A(ks)}{\zeta_A(s)}$, and then by doing a calculation like the one above, one can write down an explicit formula.
  - ◦ Exercise: Finish this calculation and give the actual formula for the number of cubefree polynomials of degree n.
  - ◦ It is also worthwhile interpreting this Euler product calculation heuristically in terms of probabilities.
  - ◦ Explicitly, we would expect (under suitable probability assumptions) that the probability of a given polynomial not being divisible by f is $(1 - 1/|f|)$.
  - ◦ So, assuming independence (which can be made rigorous by appealing to the Chinese remainder theorem), the probability that a given polynomial is not divisible by any prime power p^k for all monic irreducible p is $\prod_{p \text{ monic irred}} (1 - 1/|p|^k) = 1/\zeta_A(k)$: this is why the 1/zeta factor shows up in the answer.

0.4 (Sep 15) Dirichlet Series and Multiplicative Functions

- Another classical object of study in elementary number theory over $\mathbb{Z}$ are arithmetic functions related to divisors, such as the Euler φ-function, the divisor-counting function, and the sum-of-divisors function.
  - ◦ All of these are examples of multiplicative functions, which have the property that $f(ab) = f(a)f(b)$ whenever a, b are relatively prime. (Note the infelicitous terminology: if $f(ab) = f(a)f(b)$ for all a, b, f is instead called completely multiplicative.)
  - ◦ In particular, if n has prime factorization $n = \prod_i p_i^{a_i}$ and f is multiplicative, then $f(n) = \prod_i f(p_i^{a_i})$.
  - ◦ We will briefly review some results about multiplicative functions in the classical setting, and then redo them in the function-field setting.

- It is a standard combinatorial principle that if we want to understand a function with domain $\mathbb{N}$, we should look at its generating function.
  - ◦ A natural first guess would be to use the standard power series $\sum_{n=0}^{\infty} f(n)x^n$.
  - ◦ However, this type of generating function is useful primarily for functions that behave additively. For number-theoretic functions, we instead want to use a Dirichlet series.

- Definition: If $h : \mathbb{N} \to \mathbb{C}$ is a complex-valued function defined on positive integers, then its associated Dirichlet series is $D_h(s) = \sum_{n=1}^{\infty} \dfrac{h(n)}{n^s}$.
  - Example: If $h(n) = 1$ for all n, then $D_h(s) = \zeta(s)$, the Riemann zeta function.
  - In order for this series to converge, we need h not to grow too fast. One may check that if $h(n) = O(n^\alpha)$ then $D_h(s)$ is absolutely convergent for $\mathrm{Re}(s) > 1 + \alpha$. (We will mostly ignore issues of convergence, since our functions will grow polynomially at worst, and so we may manipulate the series as if they were formal power series.)
  - If h is multiplicative, then it is a straightforward calculation to see that $D_h(s)$ has an Euler product expansion: $D_h(s) = \prod_{p \text{ prime}}(1 + \dfrac{h(p)}{p} + \dfrac{h(p^2)}{p^2} + \cdots)$, on the appropriate domain of convergence.

- The key property of Dirichlet series is that they reproduce desired behaviors under multiplication:

- Proposition (Dirichlet Multiplication): If $f, g : \mathbb{N} \to \mathbb{C}$ are functions, then $D_f(s) \cdot D_g(s) = D_{f\star g}(s)$ where $f * g$ is the Dirichlet convolution defined via $(f * g)(n) = \sum_{d|n} f(d)g(n/d)$.
  - Proof: $D_f(s)D_g(s) = \sum_{a=1}^{\infty} \sum_{b=1}^{\infty} \dfrac{f(a)g(b)}{(ab)^s} = \sum_{n=1}^{\infty} \dfrac{1}{n^s} \sum_{ab=n} f(a)g(b) = \sum_{n=1}^{\infty} \dfrac{(f*g)(n)}{n^s} = D_{f*g}(s)$.

- The Dirichlet convolution, owing to the fact that it is merely multiplication of the underlying Dirichlet series, has various nice properties.
  - Exercise: Show that Dirichlet convolution is commutative and associative, and has an identity element given by $I(n) = \begin{cases} 1 & \text{for } n = 1 \\ 0 & \text{for } n > 1 \end{cases}$.
  - Exercise: Show that f has an inverse under Dirichlet convolution if and only if $f(1) \neq 0$.
  - Exercise: If $f(1) \neq 0$ and f is multiplicative, then its Dirichlet inverse f^{-1} is also multiplicative.
  - Exercise: Show that if two of f, g, and $f * g$ are multiplicative, then the third is also.

- By exploiting Dirichlet convolution, we can find the Dirichlet series for many basic multiplicative functions in terms of the Riemann zeta function.
  - Recall $I(n) = \begin{cases} 1 & \text{for } n = 1 \\ 0 & \text{for } n > 1 \end{cases}$ and the Mobius function $\mu(n) = \begin{cases} 0 & \text{if } n\text{is not squarefree} \\ (-1)^r & \text{if } n\text{is the product of } r\text{distinct primes} \end{cases}$.
  - Also define $N(n) = n$ and $1(n) = 1$ (for all n).
  - Exercise: Show that $D_I(s) = 1$, $D_1(s) = \zeta(s)$, and $D_N(s) = \zeta(s - 1)$.
  - First, we note that $\mu * 1 = I$, since $(\mu * 1)(n) = \sum_{d|n} \mu(d)1(n/d) = \sum_{d|n} \mu(d) = \begin{cases} 1 & \text{for } n = 1 \\ 0 & \text{for } n > 1 \end{cases}$ as noted in an exercise previously. Therefore, by multiplicativity of the Dirichlet series, we see that $D_\mu(s)D_1(s) = D_I(s)$, so that $D_\mu(s) = \dfrac{1}{\zeta(s)}$.
  - Exercise: Use $\mu * 1 = I$ to establish Mobius inversion: if $g(n) = \sum_{d|n} f(n)$ then $f(n) = \sum_{d|n} \mu(d)g(n/d)$.
  - Exercise: For the Euler φ-function, show that $\sum_{d|n} \varphi(d) = n$.
  - The previous exercise says that $\varphi * 1 = N$, and so by composing with μ and using associativity, we see that $\varphi = \mu * N$. Then we have $D_\varphi(s) = D_\mu(s)D_N(s) = \dfrac{\zeta(s-1)}{\zeta(s)}$.
  - In principle, we could have established this fornula for $D_\varphi(s)$ by manipulating the zeta function directly, but this method is both more difficult and requires knowing the actual (non-obvious) formula for the answer ahead of time.
  - We can also find the Dirichlet series for the divisor-counting function $d(n) = \#\{d \in \mathbb{N} : d|n\}$ quite easily by noting that $d(n) = \sum_{d|n} 1(d)1(d/n)$: this means $d = 1 * 1$, so $D_d(s) = D_1(s)^2 = \zeta(s)^2$.

- Exercise: If σ is the sum-of-divisors function $\sigma(n) = \sum_{d|n} d$, show that $D_\sigma(s) = \zeta(s)\zeta(s-1)$.
- Exercise: If σ_k is the sum-of-kth-powers-of-divisors function $\sigma_k(n) = \sum_{d|n} d^k$, find and prove a formula for $D_{\sigma_k}(s)$ in terms of the Riemann zeta function.

- One of the main applications of computing the Dirichlet series for these various arithmetic functions is that we can extract information about average growth rates from them.
  - In the classical case, obtaining average-growth results is moderately delicate, so we will instead just focus on the function-field case.

- Here are the function-field analogues of these classical multiplicative functions, which are now complex-valued functions on monic polynomials rather than positive integers:
  - The identity: $I(f) = \begin{cases} 1 & \text{for } f = 1 \\ 0 & \text{for } f \neq 1 \end{cases}$.
  - The norm: $N(f) = |f|$.
  - The Mobius μ-function: $\mu(f) = \begin{cases} 0 & \text{if } f \text{is not squarefree} \\ (-1)^r & \text{if } f \text{is the product of } r \text{distinct primes} \end{cases}$.
  - The Euler Φ-function: $\Phi(f) = \#(A/fA)^* = |f| \prod_{p|f \text{ prime}} (1 - 1/|p|)$.
  - The divisor-counting function: $d(f) = \#\{\text{monic } d|f\}$.
  - The sum-of-divisors function: $\sigma(f) = \sum_{d|f \text{ monic}} |d|$, or more generally the sum-of-kth-powers-of-divisors function $\sigma_k(f) = \sum_{d|f \text{ monic}} |d|^k$. (Note here that we take the norm of the divisors, since we want a $\mathbb{C}$-valued function.)
  - It is easy to check that all of these functions are multiplicative, and to write down formulas for all of them in terms of the prime factorization of $f = p_1^{a_1} \cdots p_k^{a_k}$.
  - Exercise: Verify that $d(f) = (a_1+1)\cdots(a_k+1)$ and $\sigma(f) = \dfrac{|p_1|^{a_1+1} - 1}{|p_1| - 1} \cdots \dfrac{|p_k|^{a_k+1} - 1}{|p_k| - 1}$.

- We have essentially the same definition for the Dirichlet series in the function-field case:

- Definition: If $h : \{\text{monics}\} \to \mathbb{C}$ is a complex-valued function defined on monic polynomials in $\mathbb{F}_q[t]$, then its associated Dirichlet series is $D_h(s) = \sum_{f \text{ monic}} \dfrac{h(f)}{|f|^s}$.
  - As before, we will mostly ignore issues of convergence, but just as in the classical case, one may check that if $h(f) = O(|f|^\alpha)$ then $D_h(s)$ converges absolutely for $\mathrm{Re}(s) > 1 + \alpha$.
  - We also have the same Dirichlet convolution operator: if $g, h : \{\text{monics}\} \to \mathbb{C}$ are functions, then $D_g(s) \cdot D_h(s) = D_{g\star h}(s)$ where $(g * h)(f) = \sum_{d|f \text{ monic}} g(d)h(f/d)$.
  - Dirichlet convolution is commutative, associative, and has the identity element $I(f) = \begin{cases} 1 & \text{for } f = 1 \\ 0 & \text{for } f \neq 1 \end{cases}$.
  - All of the same formulas for our arithmetic functions in terms of the zeta function follow through just as before. Here, however, we can actually write out the expressions explicitly, since we have a formula $\zeta_A(s) = \dfrac{1}{1 - q^{1-s}}$.

- Proposition (Some Dirichlet Series): For $u = q^{-s}$, we have the following formulas: $D_I(s) = 1$, $D_N(s) = \zeta_A(s-1) = \dfrac{1}{1-u}$, $D_1(s) = \zeta_A(s) = \dfrac{1}{1-qu}$, $D_\mu(s) = \dfrac{1}{\zeta_A(s)} = 1 - qu$, $D_\Phi(s) = \dfrac{\zeta_A(s-1)}{\zeta_A(s)} = \dfrac{1-qu}{1-q^2u}$, $D_d(s) = \zeta_A(s)^2 = \dfrac{1}{(1-qu)^2}$, and $D_\sigma(s) = \zeta_A(s)\zeta_A(s-1) = \dfrac{1}{(1-qu)(1-q^2u)}$.
  - Proof: Exercise.

- Using these formulas we can recover average-value results quite easily.

- Definition: If $h : \{\text{monics}\} \to \mathbb{C}$ is a function, the average value of h on degree-n polynomials is $\text{Avg}_n(h) = \dfrac{1}{q^n} \sum_{\deg(f)=n \text{ monic}} h(f)$. If the limit $\lim_{n\to\infty} \text{Avg}_n(h)$ exists, we call it the "average value" of h.

  - We can also easily average h on polynomials of degree $\leq n$: the desired sum is instead $\dfrac{1}{1+q+\cdots+q^n} \sum_{\deg(f)\leq n} h(f)$.
  - Exercise: Show that if $\lim_{n\to\infty} \text{Avg}_n(h) = \alpha$, then $\lim_{n\to\infty} \dfrac{1}{1+q+\cdots+q^n} \sum_{\deg(f)\leq n} h(f) = \alpha$ as well, so it is irrelevant whether we average over degree exactly n or $\leq n$.
  - The nice result here is that we can read off the value of $\text{Avg}_n(h)$ from the coefficients of the Dirichlet series for h: explicitly, we have $D_h(s) = \sum_{n=1}^{\infty} \dfrac{\sum_{\deg(f)=n} h(f)}{q^{ns}} = \sum_{n=1}^{\infty} \dfrac{q^n \text{Avg}_n(h)}{q^{ns}} = \sum_{n=1}^{\infty} q^n \text{Avg}_n(h) u^n$ for $u = q^{-s}$.
  - So we can calculate these averages by simply expanding out the Dirichlet series calculated above as power series in $u = q^{-s}$ and then dividing by q^n.
  - For example, $D_\mu(s) = 1 - qu$, so the average value of μ is 1 on degree-0 polynomials, -1 on degree-1 polynomials, and 0 on higher-degree polynomials.
  - Similarly, $D_d(s) = \dfrac{1}{(1-qu)^2} = (1 + qu + q^2u^2 + \cdots)^2 = 1 + 2qu^2 + 3q^2u^3 + \cdots$, so the average value of d on degree-n polynomials is $n+1$.
  - Likewise, $D_\Phi(s) = \dfrac{1-qu}{1-q^2u} = (1-qu)(1+q^2u+q^4u^2+q^6u^3+\cdots) = 1 + (q^2-q)u + (q^4-q^3)u^2 + \cdots$, so the average value of Φ on degree-n polynomials is $(q^{2n} - q^{2n-1})/q^n = q^n - q^{n-1}$.
  - Exercise: Show that the average value of σ on degree-n polynomials is $(q^{n+1}-1)/(q-1)$.

0.5 (Sep 17) Primes in Arithmetic Progressions, Part 1

- Our next task is to prove the function-field analogue of Dirichlet's theorem on primes in arithmetic progressions.

  - Over $\mathbb{Q}$, Dirichlet's theorem says that for any positive integer m and any a relatively prime to m, there exist infinitely many primes in the arithmetic progression $\{a, a+m, a+2m, a+3m, \dots\}$: in other words, congruent to a modulo m.
  - Exercise (easy): Show that if a is not relatively prime to m, then there are only finitely many primes congruent to a modulo m.

- There are $\varphi(m)$ residue classes modulo m that contain infinitely many primes, so one can ask more precisely about how the primes are distributed among these residue classes.

  - In fact, the primes are asymptotically uniformly distributed among these residue classes: the proportion of primes congruent to a modulo m approaches $1/\varphi(m)$ upon taking an appropriate limit.
  - Explicitly, define the natural density of a set S of primes to be $\lim\limits_{n\to\infty} \dfrac{S \cap \{1,2,\dots,n\}}{\{\text{primes}\} \cap \{1,2,\dots,n\}}$, provided the limit exists.
  - Then, as first proven by de la Vallée Poussin, the natural density of the primes congruent to a modulo m is $1/\varphi(m)$ when a is relatively prime to m.

- However, the natural density is somewhat difficult to handle with analytic methods. From the standpoint of zeta functions, a more natural choice is the Dirichlet density:

- Definition: If S is a set of primes, the Dirichlet density of S is the value $\delta_S = \lim\limits_{s\to 1+} \dfrac{\sum_{\text{primes } p\in S} p^{-s}}{\sum_{\text{primes } p} p^{-s}}$, assuming the limit exists.

- Note that the sum in the numerator is always finite for $\mathrm{Re}(s) > 1$ by comparison to the sum for the zeta function.
- Exercise: If S is finite, show that its Dirichlet density is 0.
- One may prove that if a set has natural density δ, then its Dirichlet density is also δ. The converse is not true, however: a simple counterexample due to Serre is the set S of primes whose leading digit is 1 in base 10.
- Exercise (hard): Show that the set of primes whose leading digit is 1 in base 10 has undefined natural density, but has Dirichlet density $\log_{10} 2$. (The answer works out the same if you use integers with leading digit 1.)

- The corresponding definition for function fields is as follows:
- Definition: If T is a set of monic irreducibles in $\mathbb{F}_q[t]$, its Dirichlet density is $\delta_T = \lim_{s \to 1+} \dfrac{\sum_{p \in T} |p|^{-s}}{\sum_p |p|^{-s}}$, assuming the limit exists.
  - We note that both the numerator and denominator sums converge for $\mathrm{Re}(s) > 1$.
- Our main result is the following:
- Theorem (Analogue of Dirichlet's Theorem): Let $m \in \mathbb{F}_q[t]$ have positive degree and let a be relatively prime to m. Then the Dirichlet density of the set of primes congruent to a (mod m) exists and is $1/\Phi(m)$. In particular, there are infinitely many such primes.
  - The fundamentally hard part of proving this theorem is to establish the nonvanishing of the L-functions for nontrivial characters at $s = 1$.
  - In order to explain what this means (and then do it), we will begin with a brisk discussion of Dirichlet characters and their properties.
- Definition: Let G be a finite abelian group. A group character χ of G is a homomorphism $\chi : G \to \mathbb{C}^\times$.
  - Note that $\chi(1) = 1$ for every character, and also if $g \in G$ has order d, then $1 = \chi(1) = \chi(g^d) = \chi(g)^d$, so $\chi(g)$ is a dth root of unity. Thus in general, χ is a map from G to the group of complex $|G|$th roots of unity.
  - Example: For any G, the trivial character χ_{triv} has $\chi_{\text{triv}}(g) = 1$ for all $g \in G$.
  - Example: If $G = (\mathbb{Z}/p\mathbb{Z})^\times$, the quadratic residue symbol $\chi(a) = \left(\dfrac{a}{p}\right)$ is a group character.
  - Example: If $G = (A/pA)^\times$ for $A = \mathbb{F}_q[t]$ and d divides $q-1$, the dth-power residue symbol $\chi(a) = \left(\dfrac{a}{p}\right)_d$ gives a group character, provided we identify the dth roots of unity in $\mathbb{F}_q$ with the dth roots of unity in $\mathbb{C}$ (simply choose any fixed isomorphism).
- We will be interested in the case where G is the group of units $(\mathbb{Z}/m\mathbb{Z})^\times$ or $(A/fA)^\times$, in which case we call χ a Dirichlet character.
  - In some situations it is slightly more convenient to work with extended Dirichlet characters, which we extend to have domain $\mathbb{Z}/m\mathbb{Z}$ or A/fA by setting $\chi(a) = 0$ whenever a is not relatively prime to the modulus.
  - Exercise: Extended Dirichlet characters modulo m are the same as functions $\chi : \mathbb{Z} \to \mathbb{C}$ (or $A \to \mathbb{C}$) such that (i) $\chi(a + bm) = \chi(a)$ for all a, b, (ii) $\chi(ab) = \chi(a)\chi(b)$ for all a, b, and (iii) $\chi(a) \neq 0$ iff a is relatively prime to m.
- We can multiply two group characters on G pointwise, and this operation makes them into a group:
- Proposition (Dual Group of G): The set of group characters on G forms a group under pointwise multiplication. The identity is the trivial character and the inverse of χ is its complex conjugate $\overline{\chi}$. This group is called the dual group of G and is denoted $\hat{G}$.

- Proof: These properties can be checked directly (exercise), or one may simply note that $\hat{G} = \mathrm{Hom}(G, \mathbb{C}^\times)$.

- The dual group $\hat{G}$ is also an abelian group, so it is natural to wonder how its structure relates to G. In fact, it is isomorphic to G:

- Proposition (Dual Group, II): If G is a finite abelian group, its dual group $\hat{G}$ is isomorphic to G.
  - Proof: First consider the special case where G is a cyclic group of order n generated by g. Then $\chi(g^d) = \chi(g)^d$ for all d, so any group character χ is uniquely determined by the value of $\chi(g)$, which must be some nth root of unity.
  - Conversely, any such selection $e^{2\pi i a/n}$ for $\chi(g)$ yields a valid group character χ_a, namely with $\chi_a(g^d) = e^{2\pi i a d/n}$. Since $\chi_a \chi_b = \chi_{a+b}$ and χ_1^n is the trivial character, we see that the dual group $\hat{G}$ is cyclic of order n (the map $a \mapsto \chi_a$ is an isomorphism of $\hat{G}$ with $\mathbb{Z}/n\mathbb{Z}$).
  - Now suppose $G = H \times K$ is a direct product. If $\chi : H \times K \to \mathbb{C}^\times$ is a homomorphism, let $\chi_H : H \to \mathbb{C}^\times$ and $\chi_K : K \to \mathbb{C}^\times$ be the projections $\chi_H(h) = \chi(h, 1)$ and $\chi_K(k) = \chi(1, k)$. Then χ_H is a group character of H, χ_K is a group character of K, and $\chi = \chi_H \chi_K$.
  - Conversely, any pair $(\chi_H, \chi_K) \in (\hat{H}, \hat{K})$ yields a character $\chi = \chi_H \chi_K \in \hat{G}$, so we see $\hat{G} \cong \hat{H} \times \hat{K}$.
  - Since every finite abelian group is a direct product of cyclic groups, and the result holds for cyclic groups and direct products, we are done.

- Exercise: If H is a subgroup of the finite abelian group G, define $H^\perp = \{\chi \in \hat{G} : \chi(H) = 1\}$. Show that $H^\perp \cong \widehat{G/H}$ and that $\hat{G}/H^\perp \cong \hat{H}$. Use these results along with $\hat{G} \cong G$ to conclude that the subgroup lattice of G is the same when turned upside down.

- The isomorphism between $\hat{G}$ and G above is non-canonical (i.e., it is not "coordinate-free" in the sense that we must pick specific generators for G and $\hat{G}$ to obtain the isomorphism).
  - However, there is a canonical isomorphism between $\hat{\hat{G}}$ (the double dual) and G given by the "evaluation map" φ, which maps an element $g \in G$ to the "evaluation-at-g" map e_g on characters $\chi \in \hat{G}$, defined by $e_g(\chi) = \chi(g)$.
  - Exercise: Verify that the evaluation map $\varphi : G \to \hat{\hat{G}}$ with $\varphi(g) = \{\chi \mapsto \chi(g)\}$ is an isomorphism from $\hat{\hat{G}}$ to G.
  - This result is a special case of Pontryagin duality, and has an analogous statement for duals of finite-dimensional vector spaces.
  - In fact, it is really the algebraic analogue of Fourier inversion (the reason being that Fourier analysis on finite abelian groups involves sums over group characters in lieu of integrals). For a brief taste of the analogy, the main idea is to note that the map $e^{inx} : \mathbb{R} \to \mathbb{C}^\times$ is a group homomorphism, and thus is an "$\mathbb{R}$"-character.

- We can also put the structure of an inner product on group characters. To establish this we first show some simple orthogonality relations:

- Proposition (Orthogonality Relations): If G is a finite abelian group and χ is a group character, the following hold:

  1. The sum $\sum_{g \in G} \chi(g) = \begin{cases} |G| & \text{if } \chi \text{is trivial} \\ 0 & \text{otherwise} \end{cases}$.
     - Proof: If χ is trivial the sum is clearly $|G|$. If χ is not trivial, say with $\chi(h) \neq 1$, then $\sum_{g \in G} \chi(g) = \sum_{g \in G} \chi(gh) = \chi(h) \sum_{g \in G} \chi(g)$ by reindexing (since $G = Gh$), and so $\sum_{g \in G} \chi(g) = 0$.
  2. The sum $\sum_{\chi \in \hat{G}} \chi(g) = \begin{cases} |G| & \text{if } g = 1 \\ 0 & \text{otherwise} \end{cases}$.
     - Proof: Apply Pontryagin duality to (1).

3. (Orthogonality 1) For any characters χ_1 and χ_2, $\sum_{g\in G} \chi_1(g)\overline{\chi_2(g)} = \begin{cases} |G| & \text{if } \chi_1 = \chi_2 \\ 0 & \text{otherwise} \end{cases}$.
   - Proof: Apply (1) to $\chi = \chi_1\overline{\chi_2}$.
4. (Orthogonality 2) For any elements g_1 and g_2, $\sum_{\chi\in\hat{G}} \chi(g_1)\overline{\chi(g_2)} = \begin{cases} |G| & \text{if } g_1 = g_2 \\ 0 & \text{otherwise} \end{cases}$.
   - Proof: Apply (2) to $g = g_1 g_2^{-1}$, or apply Pontryagin duality to (3).
5. The pairing $\langle f_1, f_2\rangle_G = \frac{1}{|G|}\sum_{g\in G} f_1(g)\overline{f_2(g)}$ is a complex inner product on functions $f : G \to \mathbb{C}$, and the elements of the dual group $\hat{G}$ are an orthonormal basis with respect to this inner product.
   - Proof: The inner product axioms are straightforward, and the fact that $\hat{G}$ yields an orthonormal basis follows from (3).
6. The pairing $\left\langle \hat{f}_1, \hat{f}_2\right\rangle_{\hat{G}} = \frac{1}{|G|}\sum_{\chi\in\hat{G}} \hat{f}_1(\chi)\overline{\hat{f}_2(\chi)}$ is a complex inner product on functions $\hat{f} : \hat{G} \to \mathbb{C}$, and the elements of G are an orthonormal basis with respect to this inner product.
   - Proof: The inner product axioms are straightforward, and the fact that $G \cong \hat{\hat{G}}$ yields an orthonormal basis follows from (4), or apply Pontryagin duality to (5).
7. (Fourier Inversion) For any function $f : G \to \mathbb{C}$, with the Fourier transform $\hat{f} : \hat{G} \to \mathbb{C}$ defined by $\hat{f}(\chi) = \langle f, \chi\rangle_G = \frac{1}{|G|}\sum_{g\in G} f(g)\overline{\chi(g)}$, we have $f(g) = \sum_{\chi\in\hat{G}} \hat{f}(\chi)\chi(g)$ for all $g \in G$.
   - Proof: This follows immediately from (5), since the elements of $\hat{G}$ are an orthonormal basis.

- Exercise: Prove Plancherel's theorem $\langle f_1, f_2\rangle_G = \frac{1}{|G|}\left\langle \hat{f}_1, \hat{f}_2\right\rangle_{\hat{G}}$ and deduce Parseval's theorem $\sum_{g\in G} |f(g)|^2 = \frac{1}{|G|}\sum_{\chi\in\hat{G}} \left|\hat{f}(\chi)\right|^2$.
- With the fundamentals taken care of, we can now focus on Dirichlet characters.
  - Studying primes congruent to a modulo m naturally leads to a question about Dirichlet characters via Fourier inversion, since we may decompose the characteristic function of [primes congruent to a modulo m] as a sum over Dirichlet characters for the group $G = (A/mA)^*$.
  - Explicitly, if $\delta_a(p)$ is 1 when $p \equiv a \pmod{m}$ and 0 otherwise, then $\hat{\delta_a}(\chi) = \frac{1}{\Phi(m)}\sum_{g\in G} \delta_a(g)\overline{\chi(g)} = \frac{1}{\Phi(m)}\overline{\chi(a)}$, since the only nonzero value of $\delta_a(g)$ occurs when $g \equiv a \pmod{m}$.
  - Then by Fourier inversion we have $\delta_a(p) = \sum_{\chi\in\hat{G}} \hat{\delta_a}(\chi)\chi(p) = \sum_{\chi\in\hat{G}} \frac{1}{\Phi(m)}\overline{\chi(a)}\chi(p)$. So the numerator for the Dirichlet density is $\sum_{p\equiv a \pmod{m}} |p|^{-s} = \sum_p \delta_a(p)\,|p|^{-s} = \frac{1}{\Phi(m)}\sum_{\chi\in\hat{G}} \left[\overline{\chi(a)}\sum_p \chi(p)\,|p|^{-s}\right]$.
  - This is a bit complicated, but the point is that we have a sum over the Dirichlet characters of constants (namely $\overline{\chi(a)}$) times $\sum_p \frac{\chi(p)}{|p|^s}$, which is quite close to the Dirichlet series for the character χ (the only difference is that we are only summing over primes, rather than all monic polynomials).
  - As we will see, we will be able to extract this sum over primes from the full Dirichlet series, which we now examine more closely.
  - The main reason we go to this effort to use Fourier inversion is that the Dirichlet series for Dirichlet characters behave very nicely (far more nicely than the original series over primes congruent to a modulo m) because Dirichlet characters are multiplicative.
- Definition: If χ is a Dirichlet character modulo m, we define its associated Dirichlet L-series $L(s,\chi) = \sum_{f \text{ monic}} \frac{\chi(f)}{|f|^s}$.

  - Note that this is just the Dirichlet series for $\chi(f)$, as we defined it previously. It is traditional to denote these series with the letter L (which was the letter Dirichlet used for such functions).
  - As usual, the series converges absolutely for $\mathrm{Re}(s) > 1$, since $|\chi(f)| \le 1$ for all f.
  - Furthermore, because Dirichlet characters are completely multiplicative, the L-series has a very simple Euler product: explicitly, $L(s,\chi) = \prod_{p \text{ irred}} \left[1 - \dfrac{\chi(p)}{|p|^s}\right]^{-1}$, for $\mathrm{Re}(s) > 1$.
  - The Euler product is the key to calculating the Dirichlet density we wanted earlier: taking the logarithm of the Euler product gives $\log L(s,\chi) = -\sum_{p \text{ irred}} \log(1 - \chi(p)/|p|^s) \approx \sum_{p \text{ irred}} \dfrac{\chi(p)}{|p|^s}$ using the Taylor approximation $-\log(1-x) \approx x$ which is accurate for small $|x|$.
  - So our main task is to determine what happens to $\log L(s,\chi)$ as $s \to 1$, since this is the required input for calculating the Dirichlet density of the primes congruent to a modulo m.

0.6 (Sep 22) Primes in Arithmetic Progressions, Part 2

- Our main task is to determine what happens to $\log L(s,\chi)$ as $s \to 1$, since this is the required input for calculating the Dirichlet density of the primes congruent to a modulo m.
- Example: For the trivial character χ_{triv}, we have $L(s,\chi_{\text{triv}}) = \prod_{p|m \text{ irred}}(1 - |p|^{-s}) \cdot \zeta_A(s)$, since the terms with $p|m$ are missing from the Euler product for $L(s,\chi)$.
  - In particular, we see that $L(s,\chi_{\text{triv}})$ has an analytic continuation (since $\zeta_A(s)$ does) and a single simple pole at $s = 1$.
- For other characters, the L-series is essentially finite.
- Proposition (L-Series for Nontrivial Characters): Let m be a monic polynomial of positive degree and χ be a nontrivial Dirichlet character modulo m. Then $L(s,\chi)$ is a polynomial in q^{-s} of degree at most $\deg m - 1$, and in particular has an analytic continuation.
  - Proof: Let $A(n,\chi) = \sum_{\deg f = n} \chi(f)$ and note, as we have previously done in working out average-value results, that $L(s,\chi) = \sum_{n=0}^{\infty} A(n,\chi) q^{-ns}$. The claimed result is then equivalent to saying $A(n,\chi) = 0$ for $n \ge \deg m$.
  - For this, suppose $\deg f = n \ge m$ and write $f = hm + r$ with $\deg r < \deg m$, where $\deg h = \deg f - \deg m$ and $\mathrm{sgn}(h) = 1/\mathrm{sgn}(m)$. Conversely, given such an h and r, we get a unique $f = hm + r$. Note that $\chi(f) = \chi(r)$, and also that there are $q^{n-\deg m}$ possible h.
  - Then $A(n,\chi) = \sum_{\deg f = n} \chi(f) = \sum_{\deg f = n} \chi(r) = q^{n - \deg m} \sum_{\deg r < \deg m} \chi(r) = 0$ where the last sum is zero by the orthogonality relation (1).
  - The observation about the analytic continuation is immediate (simply take the analytic continuation as the given polynomial in q^{-s}).
- Exercise: Choose a modulus $m \in \mathbb{F}_q[t]$ and a nontrivial Dirichlet character χ, and verify explicitly that $L(s,\chi)$ is a polynomial in q^{-s}.
- As a consequence, we see that $L(s,\chi)$ has no pole at $s = 1$ when $\chi \ne \chi_{\text{triv}}$. Our next major goal is to prove that $L(1,\chi) \ne 0$ for $\chi \ne \chi_{\text{triv}}$.
- Lemma: Let χ be any Dirichlet character modulo m. Then for each monic irreducible p not dividing m, there exist $f_p, g_p > 0$ with $f_p g_p = \Phi(m)$ such that $\prod_{\chi \in \hat{G}} L(s,\chi) = \prod_{p \nmid m}(1 - |p|^{-f_p s})^{-g_p}$.
  - Proof: For a fixed monic irreducible $p \nmid m$, as we have previously noted the evaluation-at-p map $\chi \mapsto \chi(p)$ is a homomorphism from $\hat{G}$ to $\mathbb{C}^{\times}$.
  - Let the image be a cyclic group of order f_p and the kernel have size g_p: then $f_p g_p = \#\hat{G} = \#G = \Phi(m)$ by the first isomorphism theorem.

  - For this p, by grouping the fibers of the evaluation-at-p map together, for $\zeta = e^{2\pi i/f_p}$ we have $\prod_{\chi\in\hat{G}}(1-\chi(p)/|p|^s)^{-1} = \prod_{j=0}^{f_p-1}(1-\zeta^j/|p|^s)^{-g_p}$, and this last product equals $(1-|p|^{-f_p s})^{-g_p}$ since it is the evaluation of the polynomial $(1-t)(1-\zeta t)\cdots(1-\zeta^{f_p-1}t) = 1-t^{f_p}$ at $t = |p|^{-s}$.
  - Thus, taking the product over all monic irreducibles $p \nmid m$ yields the claimed $\prod_{\chi\in\hat{G}} L(s,\chi) = \prod_{\chi\in\hat{G}}\prod_{p\nmid m}(1-\chi(p)/|p|^s)^{-1} = \prod_{p\nmid m}(1-|p|^{-f_p s})^{-g_p}$ after reversing the order of the products.

- We next show that $L(1,\chi) \neq 0$ for nonreal Dirichlet characters χ:

- Lemma (Nonvanishing, I): Let χ be any Dirichlet character modulo m such that $\chi \neq \overline{\chi}$. Then $L(1,\chi) \neq 0$.
  - Proof: If we expand the product $\prod_{\chi\in\hat{G}} L(s,\chi) = \prod_{p\nmid m}(1-|p|^{-f_p s})^{-g_p}$ from the Lemma above, it yields a Dirichlet series with nonnegative coefficients and constant term 1.
  - Thus, if s is real and greater than 1 (so that the product converges), the value of the product is real and greater than 1.
  - If $\chi \neq \overline{\chi}$, then $\prod_{\chi\in\hat{G}} L(s,\chi) = L(s,\chi_{\rm triv})L(s,\chi)L(s,\overline{\chi}) \cdot$ [other terms].
  - Now suppose $L(1,\chi) = 0$: then we would have $L(1,\overline{\chi}) = 0$ also. But this would mean the product $\prod_{\chi\in\hat{G}} L(s,\chi)$ vanishes at $s = 1$, because the only term that has a pole at $s = 1$ is $L(s,\chi_{\rm triv})$ and that pole has order 1, but we have two zeroes at $s = 1$ arising from $L(s,\chi)$ and $L(s,\overline{\chi})$.
  - But this is impossible because the value of the product is real and greater than 1 for $s > 1$. Thus, $L(1,\chi) \neq 0$.

- The case where $\chi = \overline{\chi}$ and $\chi \neq \chi_{\rm triv}$ (i.e., when χ has order 2 in $\hat{G}$) is quite a bit trickier, since we cannot get away with such a simple order-of-vanishing argument.

- Lemma (Nonvanishing, II): Let χ be any Dirichlet character of order 2 modulo m (i.e., such that $\chi = \overline{\chi}$ but $\chi \neq \chi_{\rm triv}$). Then $L(1,\chi) \neq 0$.
  - Proof: Suppose that $\chi = \overline{\chi}$ but $\chi \neq \chi_{\rm triv}$, so that $\chi(p) \in \{\pm 1\}$ for $p \nmid m$, and define the function $G(s) = \dfrac{L(s,\chi_{\rm triv})L(s,\chi)}{L(2s,\chi_{\rm triv})} = \prod_{p\nmid m}\dfrac{(1-|p|^{-s})^{-1}(1-\chi(p)\,|p|^{-s})^{-1}}{(1-|p|^{-2s})^{-1}} = \prod_{p\nmid m}\dfrac{1+|p|^{-s}}{1-\chi(p)\,|p|^{-s}} = \prod_{p\nmid m,\chi(p)=1}\dfrac{1+|p|^{-s}}{1-|p|^{-s}} = \prod_{p\nmid m,\chi(p)=1}[1+\sum_{k=1}^{\infty}|p|^{-ks}]$.
  - By expanding this last expression for G, we can see that its Dirichlet series has all coefficients nonnegative.
  - We also have $\dfrac{L(s,\chi_{\rm triv})}{L(2s,\chi_{\rm triv})} = \dfrac{\zeta_A(s)}{\zeta_A(2s)} \cdot \prod_{p|m}\dfrac{1-|p|^{-s}}{1-|p|^{-2s}} = \dfrac{1-q^{1-2s}}{1-q^{1-s}}\prod_{p|m}(1+|p|^{-s})^{-1}$. Substituting this into the expression for G yields that $\dfrac{1-q^{1-2s}}{1-q^{1-s}}L(s,\chi) = \dfrac{L(s,\chi_{\rm triv})L(s,\chi)}{L(2s,\chi_{\rm triv})}\prod_{p|m}(1+|p|^{-s})^{-1} = G(s)\prod_{p|m}(1+|p|^{-s})^{-1}$ is a Dirichlet series with all coefficients nonnegative.
  - Suppose $G(s)\prod_{p|m}(1+|p|^{-s})^{-1} = \sum_{f \text{ monic}} \dfrac{h(f)}{|f|^s}$.
  - Rewriting in terms of $u = q^{-s}$, and noting that $L^*(u,\chi) = L(s,\chi)$ is a polynomial in u as we proved earlier, we obtain the equality $\dfrac{1-qu^2}{1-qu}L^*(u,\chi) = \sum_{d=0}^{\infty}[\sum_{\deg(f)=d} h(f)]u^d$.
  - Now suppose that $L(1,\chi) = L^*(q^{-1},\chi)$ is equal to zero. Then $1-qu$ would divide $L^*(u,\chi)$, which would mean that $\dfrac{1-qu^2}{1-qu}L^*(u,\chi)$ is a polynomial in u. But then the right-hand side would also be a polynomial in u. All of its coefficients are nonnegative (as noted above), which means it cannot have a positive root for u.
  - But, finally, notice that $\dfrac{1-qu^2}{1-qu}L^*(u,\chi)$ is zero when $u = 1/\sqrt{q}$. This is a contradiction, and so $L^*(q^{-1},\chi) = L(1,\chi)$ must be nonzero.

- Now that we know $L(1,\chi)$ vanishes for nontrivial characters χ, we can prove Dirichlet's theorem:

- Theorem (Analogue of Dirichlet's Theorem): Let $m \in \mathbb{F}_q[t]$ have positive degree and let a be relatively prime to m. Then the Dirichlet density of the set of primes congruent to a (mod m) exists and is $1/\Phi(m)$. In particular, there are infinitely many such primes.

    - We have already obtained all of the necessary ingredients, so the proof is mostly a matter of putting them all together.
    - Proof: Recall the power series $-\log(1-x) = \sum_{k=1}^{\infty} x^k/k$, valid for $|x| < 1$.
    - Then for any Dirichlet character χ, we have $\log L(s,\chi) = \sum_p -\log(1-\dfrac{\chi(p)}{|p|^s}) = \sum_p \left[\sum_{k=1}^{\infty} \dfrac{\chi(p)^k}{k} |p|^{-ks}\right] = \sum_p \dfrac{\chi(p)}{|p|^s} + \sum_p \sum_{k=2}^{\infty} \dfrac{\chi(p)^k}{k} |p|^{-ks}$. The absolute value of the second term is bounded by $\sum_p \sum_{k=2}^{\infty} \dfrac{1}{k} |p|^{-ks} \le \sum_{k=2}^{\infty} \sum_{d=1}^{\infty} q^d q^{-kds} \le \sum_{n=1}^{\infty} (n+1) q^{-ns}$, which is bounded as $s \to 1+$.
    - Therefore, as $s \to 1+$, we have $\log L(s,\chi) = \sum_p \dfrac{\chi(p)}{|p|^s} + O(1)$. In particular, we see that $\sum_p |p|^{-s} = \log(s-1) + O(1)$ as $s \to 1+$, since $L(s,\chi_{\text{triv}})$ has a simple pole at $s = 1$.
    - Now, by Fourier inversion (as we previously worked out) we have $\sum_{p \equiv a \pmod{m}} |p|^{-s} = \sum_p \delta_a(p) |p|^{-s} = \dfrac{1}{\Phi(m)} \sum_{\chi \in \hat{G}} \left[\overline{\chi(a)} \sum_p \dfrac{\chi(p)}{|p|^s}\right]$.
    - So, the quotient for the Dirichlet density is $\dfrac{\sum_{p \equiv a \pmod{m}} |p|^{-s}}{\sum_p |p|^{-s}} = \dfrac{\dfrac{1}{\Phi(m)} \sum_{\chi \in \hat{G}} \left[\overline{\chi(a)} \sum_p \dfrac{\chi(p)}{|p|^s}\right]}{\sum_p |p|^{-s}} = \dfrac{1}{\Phi(m)} \left[\dfrac{\sum_{p \nmid m} |p|^{-s}}{\sum_p |p|^{-s}} + \dfrac{\sum_{\chi \neq \chi_{\text{triv}}} \overline{\chi(a)} \sum_p \dfrac{\chi(p)}{|p|^s}}{\sum_p |p|^{-s}}\right] = \dfrac{1}{\Phi(m)} \left[1 - \dfrac{\sum_{p \mid m} |p|^{-s}}{\log(s-1) + O(1)} + \dfrac{\sum_{\chi \neq \chi_{\text{triv}}} \log L(s,\chi) + O(1)}{\log(s-1) + O(1)}\right]$.
    - Now, taking the limit as $s \to 1+$ makes the second term go to zero (since the numerator is finite) and the third term go to zero (since $L(1,\chi) \neq 0$ for $\chi \neq \chi_{\text{triv}}$), and so the value of the limit is just $1/\Phi(m)$, as claimed.

- We can, in fact, improve this argument to show that the natural density of the primes congruent to a modulo m is equal to $1/\Phi(m)$, not just the Dirichlet density.

    - To do this requires showing that $L(s,\chi)$ is zero-free on a larger region: specifically, we need it to be zero-free for $\text{Re}(s) = 1$, rather than just $s = 1$.
    - The L-function is in fact zero-free on a much larger region: as we will eventually prove, the only zeroes of $L(s,\chi)$ are on the line $\text{Re}(s) = 1/2$; this is the Riemann hypothesis for function fields.
    - Taking this zero-free result for granted, we again need to manipulate the series expressions for the $L(s,\chi)$. This time, we will use in a more substantial way the fact that the $L(s,\chi)$ for $\chi \neq \chi_{\text{triv}}$ are polynomials in $u = q^{-s}$ and compare the Euler products with their factorizations.

- Theorem (Strengthened Dirichlet Analogue): Let $m \in \mathbb{F}_q[t]$ have positive degree and let a be relatively prime to m. Then the number of primes congruent to a (mod m) having degree N is equal to $\dfrac{1}{\Phi(m)} \dfrac{q^N}{N} + O(\dfrac{q^{N/2}}{N})$, where the implied constant is independent of q and N.

    - If we only know that the L-function is zero free for $\text{Re}(s) > \theta$ for some $\theta \in (1/2, 1)$, we instead get an error term of $O(\dfrac{q^{\theta N}}{N})$, which is still good enough to establish that the natural density of primes congruent to a (mod m) equals $1/\Phi(m)$.
    - Proof: For convenience, we first note the identity (*) $u \frac{\partial}{\partial u} \log(1 - \alpha u^d)^{-1} = \sum_{N=1}^{\infty} d\alpha^k u^{dN}$.
    - As we showed previously, if $\chi \neq \chi_{\text{triv}}$ then $L^*(u,\chi) = L(q^{-s},\chi)$ is a polynomial in $u = q^{-s}$ of degree at most $m - 1$. Since its constant term is 1, we obtain a factorization of the form $L^*(u,\chi) = \prod_{i=1}^{m-1} (1 - \alpha_i(\chi) u)$ for some constants $\alpha_i(\chi) \in \mathbb{C}$.

  - From the Euler product, we also have $L^*(u,\chi) = \prod_{p\nmid m}(1-\chi(p)u^{\deg p})^{-1} = \prod_{d=1}^{\infty}\prod_{p\nmid m,\deg p=d}(1-\chi(p)u^{d})^{-1}$.
  - Now apply the operator $u\frac{\partial}{\partial u}\log$ to the equality $\prod_{i=1}^{m-1}(1-\alpha_i(\chi)u) = \prod_{d=1}^{\infty}\prod_{p\nmid m,\deg p=d}(1-\chi(p)u^{d})^{-1}$ and compare coefficients of u on both sides.
  - For the LHS, using the identity (*) with $d=1$ yields $u\frac{\partial}{\partial u}\log L^*(u,\chi) = -\sum_{i=1}^{m-1}\sum_{N=1}^{\infty}\alpha_i(\chi)^N u^N = -\sum_{N=1}^{\infty}\left[\sum_{i=1}^{m-1}\alpha_i(\chi)^N\right]u^N$.
  - Letting $c_N(\chi) = -\sum_{i=1}^{m-1}\alpha_i(\chi)^N$ yields the expansion $u\frac{\partial}{\partial u}\log\prod_{i=1}^{m-1}(1-\alpha_i(\chi)u) = \sum_{N=1}^{\infty}c_N(\chi)u^N$. For $\chi=\chi_{\text{triv}}$, we have $c_N(\chi) = q^N + O(1)$, while for $\chi\neq\chi_{\text{triv}}$, by the Riemann hypothesis we have $|\alpha_i(\chi)|\in\{q^0,q^{1/2}\}$ for each i, and so $c_N(\chi) = O(q^{N/2})$.
  - For the RHS, we have

$$\begin{aligned} u\frac{\partial}{\partial u}\log L^*(u,\chi) &= \sum_{d=1}^{\infty}\sum_{p\nmid m,\deg p=d} u\frac{\partial}{\partial u}\log(1-\chi(p)u^d)^{-1} \\ &= \sum_{d=1}^{\infty}\sum_{p\nmid m,\deg p=d}\sum_{k=1}^{\infty} d\,\chi(p)^k u^{kd} \\ &= \sum_{N=1}^{\infty}\left[\sum_{d\mid N}\sum_{\deg p=N/d} d\,\chi(p)^{N/d}\right]u^N \end{aligned}$$

    by applying the identity (*) and then grouping together all of the terms of the same degree. This means $c_N(\chi) = \sum_{d\mid N}\sum_{\deg p=d} d\,\chi(p)^d$.
  - Now, by separating out the terms with $d=1$ from the others, we see $c_N(\chi) = \sum_{d\mid N}\sum_{\deg p=N/d} d\,\chi(p)^{N/d} = N\sum_{\deg p=N}\chi(p) + \sum_{d\mid N,d\geq 2}\sum_{\deg p=N/d} d\,\chi(p)^d$. The absolute value of the second term is at most $\sum_{d\mid N,d\geq 2}\sum_{\deg p=N/d} d \leq \sum_{d\mid N,d\geq 2}\frac{q^{N/d}}{N/d} = O(q^{N/2})$.
  - Therefore, we see $c_N(\chi) = N\sum_{\deg p=N}\chi(p) + O(q^{N/2})$.
  - Now we use our Fourier decomposition from earlier: we have $\frac{1}{\Phi(m)}\sum_{\chi\in\hat{G}}\overline{\chi(a)}c_N(\chi) = N\cdot\#\{\text{primes } p\equiv a \pmod{m}\} + O(q^{N/2})$ using the expression we just computed.
  - Also, we have $\sum_{\chi\in\hat{G}}\overline{\chi(a)}c_N(\chi) = q^N + O(q^{N/2})$ by directly summing over characters: $\chi=\chi_{\text{triv}}$ contributes the q^N term and the other characters each contribute $O(q^{N/2})$.
  - Setting these two equal to one another yields $\#\{\text{primes } p\equiv a \pmod{m}\} = \frac{1}{\Phi(m)}\cdot\frac{q^N}{N} + O(\frac{q^{N/2}}{N})$, as claimed.

- Exercise: For $a,m\in\mathbb{F}_q[t]$ with a relatively prime to m, show that the proportion of primes of degree N congruent to a (mod m) is $\frac{1}{\Phi(m)} + O(q^{-N/2})$, where the implied constant is independent of q and N.

0.7 (Sep 24) dth Powers and dth-Power Residue Symbols

- Our next task is to discuss the analogue of another famous result from elementary number theory: Gauss's celebrated law of quadratic reciprocity, along with its higher-order generalizations. A brief recap of the story over $\mathbb{Z}$:
  - If $a\in(\mathbb{Z}/p\mathbb{Z})^*$, we say a is a quadratic residue if $a\equiv b^2 \pmod{p}$ for some b, and otherwise we say a is a quadratic nonresidue.
  - Since the quadratic residues are simply the image of the squaring map on $(\mathbb{Z}/p\mathbb{Z})^*$, by the first isomorphism theorem there are $(p-1)/2$ of them. (One may also simply enumerate them as $1^2, 2^2, \ldots, [(p-1)/2]^2$.)

- The Legendre symbol $\left(\frac{a}{p}\right)$ is defined to be $+1$ on quadratic residues and -1 on quadratic nonresidues. By writing a as a power of the generator of $(\mathbb{Z}/p\mathbb{Z})^*$, one then obtains Euler's criterion: $a^{(p-1)/2} \equiv \left(\frac{a}{p}\right)$ (mod p), from which one sees that the Legendre symbol is multiplicative. Equivalently, it is a group homomorphism from $(\mathbb{Z}/p\mathbb{Z})^*$ to $\{\pm 1\}$.
- Exercise: Another group homomorphism from $(\mathbb{Z}/p\mathbb{Z})^*$ to $\{\pm 1\}$ is obtained by calculating the signature of the permutation associated to multiplication by a, as an element of the symmetric group S_{p-1}. Prove Zolotarev's lemma: this homomorphism is the same as the Legendre symbol.

- The law of quadratic reciprocity gives an unexpected relation between the Legendre symbols $\left(\frac{p}{q}\right)$ and $\left(\frac{q}{p}\right)$ for distinct odd primes p and q.
  - Explicitly, as first proven by Gauss, we have $\left(\frac{p}{q}\right)\left(\frac{q}{p}\right) = (-1)^{(p-1)(q-1)/4}$. Equivalently, $\left(\frac{p}{q}\right) = \left(\frac{q}{p}\right)$ if p or q is 1 mod 4, and otherwise $\left(\frac{p}{q}\right) = -\left(\frac{q}{p}\right)$ if both p, q are 3 mod 4.
  - A priori, it would seem that there is no reason for the values of $\left(\frac{p}{q}\right)$ and $\left(\frac{q}{p}\right)$ to be related to one another, since they are discussing seemingly independent questions (whether p is a square mod q and whether q is a square mod p).
  - But in fact, these questions are related: for $p^* = (-1)^{(p-1)/2}$, the value of $\left(\frac{p}{q}\right)$ determines whether the ideal (p) splits in the ring of integers $\mathcal{O}_{\sqrt{q^*}}$ of the quadratic extension $\mathbb{Q}(\sqrt{q^*})$ while the value of $\left(\frac{q}{p}\right)$ determines whether the ideal (q) splits in the ring of integers of the quadratic extension $\mathbb{Q}(\sqrt{p^*})$.
  - These two questions are related because there are several ways to understand the splitting of (q) in $\mathcal{O}_{\sqrt{p^*}}$.
  - First, from basic algebraic number theory, to determine whether (q) splits in $\mathcal{O}_{\sqrt{p^*}}$, one can study the splitting of the minimal polynomial $x^2 - x + \frac{1-p^*}{2}$ modulo q, which splits precisely when its discriminant p^* is a square: in other words, when $\left(\frac{p^*}{q}\right) = 1$.
  - Alternatively, one may look at the action of the local qth-power Frobenius map inside the Galois group of the cyclotomic field $\mathbb{Q}(\zeta_p)$, whose unique quadratic subfield is $\mathbb{Q}(\sqrt{p^*})$. Since the Galois group is cyclic, the Frobenius element Frob_q fixes $\mathbb{Q}(\sqrt{p^*})$ if and only if $q \in (\mathbb{Z}/p\mathbb{Z})^\times$ lies in $\mathrm{Gal}(\mathbb{Q}(\zeta_p)/K)$. But this group is the unique index-2 subgroup of $(\mathbb{Z}/p\mathbb{Z})^*$, which is simply the quadratic residues, so this means (q) splits precisely when $\left(\frac{q}{p}\right) = 1$.
  - Comparing these two statements yields that $\left(\frac{p^*}{q}\right) = 1$ if and only if $\left(\frac{q}{p}\right) = 1$, and this can be shown to be equivalent to the usual version of quadratic reciprocity.
  - Exercise: For distinct odd primes p, q, show that $\left(\frac{p^*}{q}\right) = \left(\frac{q}{p}\right)$ is equivalent to $\left(\frac{p}{q}\right)\left(\frac{q}{p}\right) = (-1)^{(p-1)(q-1)/4}$, where $p^* = (-1)^{(p-1)/2}$.
  - There are very many other proofs of quadratic reciprocity, many of which involve lengthy formal manipulations of various sums and (generally) yield little to no intuition about why the result is actually true. There is a fairly nice proof using Gauss sums that, suitably interpreted, is really the same as the one given above.

- We would like to generalize the reciprocity law to handle general dth powers in $\mathbb{F}_q[t]$. We begin by describing the dth powers:

- Definition: If $f \in \mathbb{F}_q[t]$ is nonconstant and a is relatively prime to f, we say that a is a dth-power residue modulo f when $x^d \equiv a$ (mod f) has a solution for x. (In other words, when a is the dth power of something mod f.)

- Example: Over $\mathbb{F}_2[t]$, we see $t+1$ is a quadratic residue modulo t^3+t+1 since $t+1 \equiv (t^2+t+1)^2 \pmod{t^3+t+1}$.
- Example: Over $\mathbb{F}_5[t]$, we see $3t^2+3t+4$ is a cubic residue modulo t^3+t+1 since $3t^2+3t+4 \equiv (t^2+2t)^3 \pmod{t^3+t+1}$.
- By the Chinese remainder theorem, $x^d \equiv a \pmod{f}$ has a solution if and only if $x^d \equiv a \pmod{p^d}$ has a solution for each prime power p^d in the factorization of f.
- Thus, we need only consider the case where the modulus is a prime power, and we can handle this case fairly easily using our earlier analysis of the structure of $(A/p^dA)^*$.

- We can start by looking at the prime-modulus case, since it is the simplest.
  - As we have mentioned previously, $(A/pA)^*$ is the multiplicative group of the finite field A/pA, so this group has order $q^{\deg p} - 1 = \tilde{p}^{f \deg p} - 1$.
  - If d does not divide $|p|-1$, then the dth power map on $(A/pA)^*$ is injective by Lagrange's theorem, so it is a bijection, and so everything in $(A/pA)^*$ is a dth power.
  - This means we can ignore divisors of d that aren't factors of $|p|-1$, and so essentially we are reduced to the situation where d divides $|p|-1$.
  - By analogy with Euler's criterion in $\mathbb{Z}$, we would expect that the value of $a^{(|p|-1)/d}$ will identify whether or not a is a dth power. This is indeed the case:

- Proposition (dth Roots Mod p): If $p \in \mathbb{F}_q[t]$ is irreducible, a is not divisible by p, and d is a divisor of $|p|-1$, then $x^d \equiv a \pmod{p}$ is solvable if and only if $a^{(|p|-1)/d} \equiv 1 \pmod{p}$.
  - Proof 1: First, if $x^d \equiv a \pmod{p}$ then $a^{(|p|-1)/d} \equiv x^{|p|-1} \equiv 1 \pmod{p}$ by Euler.
  - For the converse, recall that we showed previously that $x^d \equiv 1 \pmod{p}$ has d solutions mod p whenever d divides $|p|-1$.
  - Therefore, the kernel of the dth-power map on $(A/pA)^*$ has size d, so by the first isomorphism theorem, the image, which is precisely the set of dth powers, has size $(|p|-1)/d$.
  - But by the same observation, there are exactly $(|p|-1)/d$ solutions to the equation $x^{(|p|-1)/d} \equiv 1 \pmod{p}$, so by the above, these must be exactly the dth powers.
  - Proof 2: As shown previously, $(A/pA)^*$ is cyclic of order $|p|-1$. Let u be a generator.
  - Since every element in $(A/pA)^*$ is a power of u, it is easy to see that for any d dividing $|p|-1$, the dth powers in $(A/pA)^*$ are precisely $\{u^d, u^{2d}, u^{3d}, \dots, u^{d(|p|-1)d} = 1\}$. All of these elements clearly satisfy $x^{(|p|-1)/d} \equiv 1 \pmod{p}$.
  - Conversely, if $a = u^k$ has $a^{(|p|-1)/d} \equiv 1 \pmod{p}$, then $u^{k(|p|-1)/d} \equiv 1 \pmod{p}$ so since u has order $|p|-1$, d must divide k.

- Now that we have analyzed the prime case, the prime-power case follows by "lifting" the solutions from the prime case.
  - This is a consequence of a much more general result known as Hensel's lemma, which we might as well do in general.

- Proposition (Hensel's Lemma): If $p \in \mathbb{F}_q[t]$ is irreducible, $a \in \mathbb{F}_q[t]$, and $r(x)$ is any polynomial such that $r(a) \equiv 0 \pmod{p^d}$ and $r'(a) \not\equiv 0 \pmod{p}$, then there is a unique k modulo p such that $r(a+kp^d) \equiv 0 \pmod{p^{d+1}}$. Explicitly, if $u = f'(a)^{-1} \pmod{p}$, then $k = -\dfrac{uf(a)}{p^d}$.
  - By repeatedly applying Hensel's lemma, we can lift a solution of $r(a) \equiv 0 \pmod{p}$ to a solution modulo p^2, and then lift that to a solution modulo p^3, and so on and so forth, until we have a solution to the equation modulo any power of p.
  - This iteration process yields a sequence of solutions $x \equiv a_j \pmod{p^j}$ for each j, where $a_{j+1} = a_j - \dfrac{1}{r'(a)} r(a_j)$, which one may recognize as the iteration procedure from Newton's root-finding method. In fact, if we instead think of solving the polynomial $r(x) = 0$ p-adically (which amounts to taking the inverse limit $\varprojlim(A/p^dA)$), this lifting procedure is precisely Newton's method with starting point $x = a$.

- ○ Proof: First, by the binomial theorem we have $(a+p^d k)^n = a^n + na^{n-1}p^d k + [\text{terms divisible by } p^{2d}] \equiv a^n + na^{n-1}p^d k \pmod{p^{d+1}}$.
  - ○ Then if $r(t) = \sum c_n t^n$ we see that $r(a+p^d k) \equiv \sum c_n (a^n + na^{n-1}p^d k) \equiv \sum c_n a^n + p^d k \sum nc_n a^{n-1} \equiv r(a) + p^d k \cdot r'(a) \pmod{p^{d+1}}$.
  - ○ By hypothesis, $r(a)+p^d k\cdot r'(a)$ is divisible by p^d. So dividing the congruence $r(a+kp^d) \equiv 0 \pmod{p^{d+1}}$ by p^d yields $\dfrac{r(a)}{p^d} + kr'(a) \equiv 0 \pmod{p}$, which has the unique solution $k \equiv -\dfrac{uf(a)}{p^d} \pmod{p}$, as claimed.

- This version of Hensel's lemma is quite a bit more than we really need here, but it will be helpful to have it available later.

- Corollary (dth Roots Mod p^e): If $p \in \mathbb{F}_q[t]$ is irreducible, d divides $|p|-1$, and p does not divide a, then $x^d \equiv a \pmod{p}$ has a root if and only if $x^d \equiv a \pmod{p^e}$ has a root for every $e \geq 1$.
  - ○ Proof: If there is a solution to $x^d \equiv a \pmod{p^e}$ then clearly there is a solution mod p.
  - ○ Conversely, if there is a solution mod p, then we claim we may lift the solution mod p^e using Hensel's lemma.
  - ○ We just need to check that the derivative is not zero: for $r(x) = x^d$ we have $r'(a) = da^{d-1}$. Then $d \neq 0$ mod p because d divides $|p|-1 = \tilde{p}^{f \deg p} - 1$ and so d cannot be divisible by the characteristic $\tilde{p}$, and also $a \neq 0$ mod p because p does not divide a. Thus, Hensel's lemma applies, and we are done.

- Corollary (Counting dth Powers): If $p \in \mathbb{F}_q[t]$ is irreducible and d divides $|p|-1$, then there are $\Phi(p^e)/d$ total dth-power residues modulo p^e.
  - ○ Proof 1: Count residue classes: as shown earlier there are $(|p|-1)/d = \Phi(p)/d$ total dth-power residue classes modulo p. By the corollary above, the dth-power residue classes modulo p^e are precisely those that reduce to a dth power modulo p. So the probability of selecting one is $\Phi(p)/(d\,|p|)$, and thus the total number is $|p|^e \cdot \Phi(p)/(d\,|p|) = \Phi(p^e)/d$.
  - ○ Proof 2 (sketch): The dth-power homomorphism commutes with reduction modulo p. Then just count the sizes of the various kernels and images and use the first isomorphism theorem.
  - ○ Exercise: Show that for any monic polynomial m, there are $\Phi(m)/d^{\lambda(m)}$ total dth powers modulo m, where $\lambda(m)$ is the number of distinct monic irreducible factors of m.

- Returning back to the prime case, in the particular case where d divides $q-1$, then the dth roots of unity in $(A/pA)^*$ actually lie inside $\mathbb{F}_q$, because $x^d = 1$ already has d solutions inside $\mathbb{F}_q$ (since $\mathbb{F}_q^*$ is cyclic of order $q-1$).
  - ○ We have shown above that a is a dth power modulo p if and only if $a^{(|p|-1)/d} \equiv 1 \pmod{p}$.
  - ○ We can use this as the basis for our definition of the dth-power residue symbol, in analogy with Euler's criterion over $\mathbb{Z}$.

- Definition: If $p \in \mathbb{F}_q[t]$ is irreducible and d divides $q-1$, then we define the dth-power residue symbol $\left(\dfrac{a}{p}\right)_d$ to be the unique element of $\mathbb{F}_q$ congruent to $a^{(|p|-1)/d}$ modulo p.
  - ○ Example: For $d=2$ over $\mathbb{F}_3[t]$, we calculate $\left(\dfrac{t}{t^2+t+2}\right)_2 \equiv t^4 \equiv 2 \pmod{t^2+t+2}$.
  - ○ Example: For $d=3$ over $\mathbb{F}_7[t]$, we calculate $\left(\dfrac{t}{t^2+2t+2}\right)_7 \equiv t^{16} \equiv 4 \pmod{t^2+2t+2}$.
  - ○ Example: For $d=3$ over $\mathbb{F}_7[t]$, we calculate $\left(\dfrac{t}{t^2+t+6}\right)_7 \equiv t^{16} \equiv 1 \pmod{t^2+t+6}$, which means t is a cube modulo t^2+t+6.

- Proposition (Properties of Residue Symbols): If $p \in \mathbb{F}_q[t]$ is irreducible and d divides $q-1$, the following hold:

  1. $\left(\dfrac{a}{p}\right)_d = 0$ if and only if p divides a.

2. If $a \equiv b \pmod{p}$ then $\left(\frac{a}{p}\right)_d = \left(\frac{b}{p}\right)_d$.

3. The residue symbol is multiplicative: for any a, b, $\left(\frac{ab}{p}\right)_d = \left(\frac{a}{p}\right)_d \left(\frac{b}{p}\right)_d$.

4. $\left(\frac{a}{p}\right)_d = 1$ if and only if a is a dth-power residue modulo p.

5. If ζ is any dth root of unity in $\mathbb{F}_q$, then there exists $a \in \mathbb{F}_q[t]$ with $\left(\frac{a}{p}\right)_d = \zeta$.

6. The residue symbol is a surjective group homomorphism from $(A/pA)^*$ to μ_d, the group of dth roots of unity in $\mathbb{F}_q$.

7. If $d|d'$ then $\left(\frac{a}{p}\right)_d = \left(\frac{a}{p}\right)_{d'}^{d'/d}$.

8. If $\alpha \in \mathbb{F}_q$ then $\left(\frac{\alpha}{p}\right)_d = \alpha^{(q-1)/d \cdot \deg p}$.

   - Proofs: (1)-(4) are trivial from the definition or results previously shown. (5) follows by the first isomorphism theorem, since the kernel of the $(|p|-1)/d$th-power map has size $(|p|-1)/d$ hence the image has size d. (6) is a rephrasing of (3) and (5).
   - (7) follows by noting $\left(\frac{a}{p}\right)_{d'}^{d'/d} \equiv (a^{(|p|-1)/d'})^{d'/d} = a^{(|p|-1)/d} \equiv \left(\frac{a}{p}\right)_d \pmod{p}$, and then observing that since the residue symbols are both elements of $\mathbb{F}_q$, the congruence mod p forces actual equality.
   - For (8), first note that $\frac{|p|-1}{d} = \frac{q^{\deg p}-1}{d} = (1 + q + q^2 + \cdots + q^{\deg p - 1})(q-1)/d$. Then since $\alpha^q = \alpha$ by Fermat's little theorem in $\mathbb{F}_q$, we have $\left(\frac{\alpha}{p}\right)_d \equiv \alpha^{(|p|-1)/d} = (\alpha \cdot \alpha^q \cdot \alpha^{q^2} \cdot \cdots \cdot \alpha^{q^{\deg p - 1}})^{(q-1)/d} = \alpha^{\deg p \cdot (q-1)/d} \pmod{p}$. Then as in (7), the congruence modulo p forces equality.

- We can now state the dth-power reciprocity law, which we will prove next time:
- Theorem (dth-Power Reciprocity): If d divides $q-1$ and P, Q are monic irreducible polynomials in $\mathbb{F}_q[t]$, then $\left(\frac{Q}{P}\right)_d = (-1)^{(q-1)(\deg P)(\deg Q)/d} \left(\frac{P}{Q}\right)_d$.

0.8 (Sep 29) The dth-Power Reciprocity Law

- To prove the reciprocity law, we first need a reciprocity result about roots of polynomials known as Weil reciprocity:
- Lemma (Weil Reciprocity): If $P(t) = (t - r_1) \cdots (t - r_n)$ and $Q(t) = (t - s_1) \cdots (t - s_n)$ are monic polynomials over a field F, with the $r_i, s_j \in F$, then $\prod_{i=1}^n Q(r_i) = (-1)^{(\deg P)(\deg Q)} \prod_{j=1}^m P(s_j)$.
   - Proof: Note that $Q(r_i) = \prod_{j=1}^m (r_i - s_j)$ so $\prod_{i=1}^n Q(r_i) = \prod_{i=1}^n \prod_{j=1}^m (r_i - s_j)$. In the same way, $\prod_{j=1}^m P(s_j) = \prod_{j=1}^m \prod_{i=1}^n (s_j - r_i)$.
   - These expressions are the same up to switching the order of the products and scaling each of the $mn = (\deg P)(\deg Q)$ terms by -1, so the result follows.
- We can now prove the dth-power reciprocity law:
- Theorem (dth-Power Reciprocity): If d divides $q-1$ and P, Q are monic irreducible polynomials in $\mathbb{F}_q[t]$, then $\left(\frac{Q}{P}\right)_d = (-1)^{(q-1)(\deg P)(\deg Q)/d} \left(\frac{P}{Q}\right)_d$.
   - The main idea of the proof is to exploit properties of the Frobenius map on the roots of P and Q in their splitting field over $\mathbb{F}_q$, and then use Weil reciprocity.

- Proof: From property (7) of the residue symbol, we have $\left(\frac{a}{p}\right)_d = \left(\frac{a}{p}\right)_{d'}^{d'/d}$, so it is enough to prove the reciprocity law when $d = q-1$.
- Now let α be a root of P and β be a root of Q in a splitting field $E/\mathbb{F}_q$ for the polynomial PQ.
- Since $E/\mathbb{F}_q$ is a finite-degree extension of a finite field, its Galois group is cyclic and generated by the qth-power Frobenius map.
- Also, since P and Q are irreducible over $\mathbb{F}_q$, we must have the factorizations

$$P(t) = (t-\alpha)(t-\alpha^q)(t-\alpha^{q^2})\cdots(t-\alpha^{q^{\deg P-1}})$$
$$Q(t) = (t-\beta)(t-\beta^q)(t-\beta^{q^2})\cdots(t-\beta^{q^{\deg Q-1}})$$

since α, α^q, α^{q^2}, ... are all the Galois conjugates of α and P is irreducible (with the same logic applying to β and Q).
- Inside $E[t]$, we have $\left(\frac{Q}{P}\right)_{q-1} \equiv [Q(t)]^{(q^{\deg P}-1)/(q-1)} = [Q(t)]^{1+q+q^2+\cdots+q^{\deg P-1}} = Q(t)Q(t)^qQ(t)^{q^2}\cdots Q(t)^{q^{\deg P-1}} \equiv Q(t)Q(t^q)Q(t^{q^2})\cdots Q(t^{q^{\deg P-1}}) \pmod{P}$ since $Q(t^q) = Q(t)^q$ in characteristic q.
- Reducing both sides modulo the factor $t-\alpha$ of P (equivalently, evaluating both sides at $t=\alpha$) then yields $\left(\frac{Q}{P}\right)_{q-1} \equiv Q(\alpha)Q(\alpha^q)\cdots Q(\alpha^{q^{\deg P}-1}) \pmod{t-\alpha}$. Since the right-hand side of this expression is the product of the values of Q evaluated at the roots of P, it is the same for any other root of P we choose in place of α.
- So by the Chinese remainder theorem, in fact $\left(\frac{Q}{P}\right)_{q-1} \equiv Q(\alpha)Q(\alpha^q)\cdots Q(\alpha^{q^{\deg P}-1}) \pmod{P}$. But the right-hand side is an element of E, and since it is a $(q-1)$st root of unity (or alternatively, since it is Galois-invariant), it must actually be in $\mathbb{F}_q$. So since these quantities are congruent modulo P, they must actually be equal as elements of $\mathbb{F}_q$.
- This means $\left(\frac{Q}{P}\right)_{q-1} = Q(\alpha)Q(\alpha^q)\cdots Q(\alpha^{q^{\deg P}-1})$. In the same way, $\left(\frac{P}{Q}\right)_{q-1} = P(\beta)P(\beta^q)\cdots P(\beta^{q^{\deg Q}-1})$.
- Weil reciprocity then says $Q(\alpha)Q(\alpha^q)\cdots Q(\alpha^{q^{\deg P}-1}) = (-1)^{(\deg P)(\deg Q)}P(\beta)P(\beta^q)\cdots P(\beta^{q^{\deg Q}-1})$, so we see $\left(\frac{Q}{P}\right)_{q-1} = (-1)^{(\deg P)(\deg Q)}\left(\frac{P}{Q}\right)_{q-1}$, which establishes the case $d = q-1$.
- The case where d divides $q-1$ follows immediately and gives the general statement above.

- Just as in the case of $\mathbb{Q}$, to give a convenient method for calculating residue symbols, we can extend the definition to include nonprime moduli (i.e., generalizing the Jacobi symbol):
- Definition: If $b \in \mathbb{F}_q[t]$ has prime factorization $b = uq_1^{b_1}\cdots q_n^{b_n}$ for distinct monic irreducible q_i and $u \in \mathbb{F}_q^\times$, then we define the general residue symbol as $\left(\frac{a}{b}\right)_d = \prod_{j=1}^n \left(\frac{a}{q_i}\right)^{b_i}$.
- Proposition (Properties of Residue Symbols, II): If $b \in \mathbb{F}_q[t]$ is nonzero and d divides $q-1$, the following hold:

1. $\left(\frac{a}{b}\right)_d$ is either 0 or a dth root of unity, and $\left(\frac{a}{b}\right)_d \neq 0$ if and only if a, b are relatively prime.
2. If $a_1 \equiv a_2 \pmod{b}$ then $\left(\frac{a_1}{b}\right)_d = \left(\frac{a_2}{b}\right)_d$.
3. The residue symbol is multiplicative on the top: $\left(\frac{a_1a_2}{b}\right)_d = \left(\frac{a_1}{b}\right)_d\left(\frac{a_2}{b}\right)_d$.
4. The residue symbol is multiplicative on the bottom: $\left(\frac{a}{b_1b_2}\right)_d = \left(\frac{a}{b_1}\right)_d\left(\frac{a}{b_2}\right)_d$.
5. If $\gcd(a,b) = 1$ and a is a dth-power residue modulo b, then $\left(\frac{a}{b}\right)_d = 1$. (The converse need not hold.)

6. If $d|d'$ then $\left(\frac{a}{b}\right)_d = \left(\frac{a}{b}\right)_{d'}^{d'/d}$.
7. If $\alpha \in \mathbb{F}_q$ then $\left(\frac{\alpha}{b}\right)_d = \alpha^{(q-1)/d \cdot \deg b}$.

  - Proofs: (1)-(4) follow straightforwardly from the definition, while (6) and (7) follow the same way as for the residue symbol with prime modulus. For (5), if $a \equiv c^d \pmod p$ then $\left(\frac{a}{b}\right)_d = \left(\frac{c^d}{b}\right)_d = \left(\frac{c}{b}\right)_d = 1$ since $\left(\frac{c}{b}\right)$ is a dth root of unity (since it is not zero since a, b are relatively prime).
  - We will remark that the residue symbol $\left(\frac{\star}{b}\right)_d : (A/bA)^* \to \mu_d$ is still a group homomorphism since it is multiplicative by (3), but it is not necessarily surjective when b is not prime. For example, if $b = p^d$ is a dth power, then by (4) we see that $\left(\frac{a}{b}\right)_d = \left(\frac{a}{p}\right)_d^d = 1$ for all $a \in (A/bA)^*$. (This also shows that the converse of (5) is false, as noted above.)

- We can write down the reciprocity law for general dth-power residue symbols:
- Theorem (General Reciprocity Law): If d divides $q-1$ and a, b are any nonzero polynomials in $\mathbb{F}_q[t]$, then $\left(\frac{a}{b}\right)_d = (-1)^{(q-1)(\deg a)(\deg b)/d}[\text{sgn}a]^{(q-1)/d \cdot \deg b}[\text{sgn}b]^{-(q-1)/d \cdot \deg a}\left(\frac{b}{a}\right)_d$.
  - Proof (sketch): As in the prime case, reduce to the case $d = q-1$. Then pull out the leading coefficients of a, b (these are where the sgn$a$ and sgn$b$ terms come from) and then apply the definition of the general residue symbol to write $\left(\frac{a}{b}\right)_{q-1}$ and $\left(\frac{b}{a}\right)_{q-1}$ as products of residue symbols with prime moduli, apply the prime-modulus reciprocity law, and tally up the results. The full details are left as an exercise.
- A standard application of quadratic reciprocity over $\mathbb{Z}$ is to characterize all of the prime moduli for which a given integer m is a quadratic residue.
  - Typical examples of such statements: -1 is a quadratic residue mod p when $p \equiv 1 \pmod 4$, 3 is a quadratic residue mod p when $p \equiv 1, 11 \pmod{12}$, 5 is a quadratic residue mod p when $p \equiv 1, 4 \pmod 5$, and so forth.
  - Aside from the special cases of -1 and 2, one may answer this question simply by factoring m as a product of primes $m = q_1 \cdots q_k$, so that $\left(\frac{m}{p}\right) = \left(\frac{q_1}{p}\right) \cdots \left(\frac{q_k}{p}\right)$, and then applying quadratic reciprocity to flip each of the quadratic residue symbols. The end result is that the statement $\left(\frac{m}{p}\right) = +1$ is equivalent to a congruence condition for p modulo $4m$, which one may calculate explicitly if desired.
- We can use this same type of argument to solve the analogous problem in function fields:
- Theorem (Criterion for dth-Power Residues): Let $m \in \mathbb{F}_q[t]$ be monic and $d|(q-1)$, and let $\{a_1, \dots, a_k\}$ be coset representatives for the residue classes in $(A/mA)^*$ with $\left(\frac{a}{m}\right)_d = +1$ and $\{b_1, \dots, b_k\}$ be coset representatives for the residue classes in $(A/mA)^*$ with $\left(\frac{b}{m}\right)_d = -1$ (if there are any). Then the following hold:

  1. If $\deg(m)$, $(q-1)/d$, or $\text{char}(\mathbb{F}_q)$ is even, then m is a dth power modulo an irreducible monic polynomial p if and only if $p \equiv a_i \pmod m$ for some i.
  2. If $\deg(m)$, $(q-1)/d$, and $\text{char}(\mathbb{F}_q)$ are all odd, then m is a dth power modulo an irreducible monic polynomial p if and only if either $\deg(p)$ is even and $p \equiv a_i \pmod m$ for some i, or $\deg(p)$ is odd and $p \equiv b_i \pmod m$ for some i.

  - Proof: Note that $p \equiv a_i \pmod m$ is equivalent to saying $\left(\frac{p}{m}\right)_d = 1$, while $p \equiv b_i \pmod m$ is equivalent to saying $\left(\frac{p}{m}\right)_d = -1$.

- Since p and m are monic, by the reciprocity law we see $\left(\frac{m}{p}\right)_d = (-1)^{(q-1)/d \cdot \deg(m)\deg(p)} \left(\frac{p}{m}\right)_d$.
- First, if q is even, then $\mathrm{char}(\mathbb{F}_q) = 2$: then $-1 = 1$ over $\mathbb{F}_q$, so $\left(\frac{m}{p}\right)_d = \left(\frac{p}{m}\right)_d$. Likewise, if $\deg(m)$ or $(q-1)/d$ is even, then the exponent of -1 is even, so again we see $\left(\frac{m}{p}\right)_d = \left(\frac{p}{m}\right)_d$. Together with the observation above, (1) follows.
- For (2), if $\deg(m)$, $(q-1)/d$, and $\mathrm{char}(\mathbb{F}_q)$ are all odd, then $-1 \neq 1$ and $(-1)^{(q-1)/d \cdot \deg(m)\deg(p)} = (-1)^{\deg p}$. So $\left(\frac{m}{p}\right)_d = \left(\frac{p}{m}\right)_d$ if $\deg(p)$ is even while $\left(\frac{m}{p}\right)_d = -\left(\frac{p}{m}\right)_d$ if $\deg(p)$ is odd. This yields (2).

- Example: Identify all monic irreducibles $p \in \mathbb{F}_3[t]$ such that t is a square modulo p.
  - There are two residue classes in $(A/tA)^*$, namely 1 and 2, and we see $\left(\frac{1}{t}\right)_2 = 1$ while $\left(\frac{2}{t}\right)_2 = -1$.
  - Since $\deg(m) = 1$, $(q-1)/d = 1$, and $\mathrm{char}(\mathbb{F}_q) = 3$, we are in case (2). Thus, m is a quadratic residue modulo the monic irreducible polynomial p precisely when $\deg(p)$ is odd and $p \equiv 2 \pmod{t}$, or when $\deg(p)$ is even and $p \equiv 1 \pmod{t}$.
  - For example, we see that t is a square modulo the irreducible polynomial $t^3 + 2t + 2 \in \mathbb{F}_3[t]$, and indeed with some more work, one may calculate $t \equiv (t^2 + t + 2)^2 \pmod{t^3 + 2t + 2}$.
  - Exercise: Extend this example to describe all monic irreducibles $p \in \mathbb{F}_q[t]$ such that t is a square modulo p for arbitrary finite fields $\mathbb{F}_q$.

- Another interesting application of the dth-power reciprocity law is to establish a "Hasse principle"-type result for dth powers.
  - Obviously, if a polynomial with integer coefficients has a solution in $\mathbb{Z}$, then it also has solutions modulo p^k for all prime powers p^k (equivalently, it has a p-adic solution for each p) and it also has a real solution.
  - The Hasse principle asks when the converse of this observation is valid: if a polynomial has a p-adic root and a real root, does it necessarily have a rational root? The general idea is that one may try to piece together information modulo the prime powers for many primes p using the Chinese remainder theorem, but it is not clear when this actually forces the existence of a global solution.
  - As first proven by Minkowski for integer coefficients (and then later extended by Hasse for number-field coefficients), for quadratic polynomials this local-global principle holds: if a quadratic polynomial has a p-adic root and a real root, it necessarily has a rational root.
  - The result is known to be false for cubic forms: Selmer's famous counterexample is the cubic equation $3x^3 + 4y^3 + 5z^3 = 0$, which has no rational solution but does have real solutions and p-adic solutions for all p.
  - Even in the absence of a literal Hasse-principle statement, in many cases one can analyze the precise obstructions to lifting local solutions to global solutions. (An example of this sort of obstruction can be found in the statement of the Grunwald-Wang theorem.)

- Theorem (Hasse Principle for dth Powers): Let $m \in \mathbb{F}_q[t]$ have positive degree and $d|(q-1)$. If $x^d \equiv m \pmod{p}$ is solvable for all but finitely many irreducible polynomials p, then $x^d = m$ has a solution in $\mathbb{F}_q[t]$ (i.e., m is globally a dth power).
  - Proof: Let $m = \beta q_1^{d_1} \cdots q_k^{d_k}$ where the q_i are distinct monic irreducibles and β is a constant. We first show that if any d_i is not divisible by d, then there are infinitely many irreducibles p such that $\left(\frac{m}{p}\right)_d \neq 1$.
  - To show this, suppose without loss of generality that d_1 is not divisible by d. We inductively construct an infinite set of irreducibles $\{r_i\}$ with $\left(\frac{m}{r_i}\right)_d \neq 1$, so suppose we have a set (possibly empty to start) $\{r_1, \ldots, r_s\}$ of monic irreducibles not dividing m with $\left(\frac{m}{r_i}\right)_d \neq 1$ for all i.

- Select any primitive dth root of unity ζ_d: then there exists an element $c \in \mathbb{F}_q[t]$ with $\left(\frac{c}{q_i}\right)_d = \zeta_d$ by our properties of the dth-power residue symbol.
- By the Chinese remainder theorem, there exist solutions a to the system of congruences $a \equiv c \pmod{q_1}$, $a \equiv 1 \pmod{q_2 \cdots q_k}$, $a \equiv 1 \pmod{r_1 \cdots r_s}$. Select any such solution that is monic and has degree divisible by $2d$.
- For this a, we have $\left(\frac{a}{m}\right)_d = \prod_{i=1}^k \left(\frac{a}{q_i}\right)_d^{d_i} = \zeta_d^{d_1} \neq 1$ since d_1 is not divisible by d.
- Then by the reciprocity law, we then have $\left(\frac{m}{a}\right)_d = (-1)^{(q-1)/d \cdot (\deg m)(\deg a)} \left(\frac{a}{m}\right)_d = \left(\frac{a}{m}\right)_d \neq 1$, since the exponent of -1 has a factor of 2 from $\deg a$.
- Since the general dth-power residue symbol is multiplicative on the bottom, there must be some monic irreducible factor r_{s+1} of a such that $\left(\frac{m}{r_{s+1}}\right) \neq 1$ since $\left(\frac{a}{m}\right)_d \neq 1$. This monic irreducible factor is relatively prime to $r_1 \cdots r_s$ since $a \equiv 1 \pmod{r_1 \cdots r_s}$, so we have found another monic irreducible to add to our list.
- By induction, we can construct infinitely many such irreducibles.
- Now, if $x^d \equiv m \pmod{p}$ is solvable for all but finitely many irreducible polynomials p, then by the above, each of the exponents d_i must be divisible by d. This means $m = \beta \cdot \tilde{m}^d$ for some monic polynomial $\tilde{m}$, so all that remains is to show that β is a dth power.
- For any irreducible p not dividing m, we have $\left(\frac{m}{p}\right)_d = \left(\frac{\beta}{p}\right)_d = \beta^{(q-1)/d \cdot \deg p}$ as we have previously shown. Since there are irreducibles of any desired degree in $\mathbb{F}_q[t]$, select p to be one of degree relatively prime to d with $\left(\frac{m}{p}\right)_d = 1$: then $\beta^{(q-1)/d \cdot \deg p} = 1$ implies $\beta^{(q-1)/d} = 1$, which is equivalent to saying that β is a dth power. Then m itself is a dth power, as claimed.

0.9 (Oct 1) Transcendence and Localization

- We now move into the second major part of the course, which deals with algebraic function fields: these are function fields of transcendence degree 1 over a general constant field F.
  - Later, we will specialize to function fields over $\mathbb{F}_q$ (equivalently, these are the finite-degree field extensions of $\mathbb{F}_q(t)$), which along with algebraic number fields (the finite-degree field extensions of $\mathbb{Q}$) constitute the global fields.
  - Global fields (to be considered as parallel to local fields) share a number of common properties that we will elucidate and study.
- We begin by reviewing some basic facts about transcendental extensions.
- Definition: Let K/F be a field extension. We say a subset S of K is algebraically dependent over F if there exists a finite subset $\{s_1, \ldots, s_n\} \in S$ and a nonzero polynomial $p \in F[x_1, \ldots, x_n]$ such that $p(s_1, \ldots, s_n) = 0$. If there exists no such p for any finite subset of S, we say S is algebraically independent.
  - The general idea here is that a set of elements is algebraically dependent if they satisfy some algebraic (i.e., polynomial) relation over F.
  - Example: If $x_1, \ldots, x_n$ are indeterminates inside $F(x_1, \ldots, x_n)$, the function field in n variables, then the set $\{x_1, \ldots, x_n\}$ is algebraically independent over F.
  - Example: Over $\mathbb{Q}$, the set $\{\pi, \pi^2\}$ is algebraically dependent, since $p(x, y) = x^2 - y$ has $p(\pi, \pi^2) = 0$.
  - Example: Over $\mathbb{Q}$, the set $\{\sqrt[3]{2}\}$ is algebraically dependent, since $p(x) = x^3 - 2$ has $p(\sqrt[3]{2}) = 0$.
  - More generally, the set $\{\alpha\}$ is algebraically independent over F if and only if α is transcendental over F.
  - Exercise: Show that the set $\{x + y, x^2 + y^2\}$ is algebraically independent in $F(x, y)$ for any field F of characteristic not 2, but is algebraically dependent if F has characteristic 2.

- Example: In $F(x,y)$, the set $\{x+y, x^2+y^2, x^3+y^3\}$ is algebraically dependent, since $p(a,b,c) = a^3 - 3ab + 2c$ has $p(x+y, x^2+y^2, x^3+y^3) = 0$.

- The notion of algebraic independence generalizes the notion of linear independence, and as such the two concepts are related in various ways.
  - It is easy to see that any subset of an algebraically independent set is algebraically independent, while any set containing an algebraically dependent set is algebraically dependent.
  - Since having a basis of a vector space is very convenient for calculations, we might therefore hope to define an analogous "transcendence basis" to be an algebraically independent set that generates the extension K/F.
  - Unfortunately, such a set need not exist: for example, $\mathbb{Q}(\sqrt{2})/\mathbb{Q}$ has no such set, because there are no transcendental elements at all.
  - The correct analogy is instead to observe that a basis for a vector space is a maximal linearly independent set:

- Definition: Let K/F be a field extension. A transcendence base for K/F is an algebraically independent subset S of K that is maximal in the set of all algebraically independent subsets of K.
  - Remark: The term "transcendence basis" is also used occasionally.
  - By a straightforward Zorn's lemma argument, every extension has a transcendence base. (Exercise: Write down this argument.)
  - Example: The empty set $\emptyset$ is a transcendence base for $\mathbb{Q}(\sqrt{2})/\mathbb{Q}$. More generally, K/F is algebraic if and only if $\emptyset$ is a transcendence base.
  - Example: The set $\{x\}$ is a transcendence base for $F(x)$ over F.

- Here are some of the fundamental properties of transcendence bases, many of which are analogous to properties of vector spaces:

- Proposition (Transcendence Bases): Suppose K/F is a field extension and S is a subset of K.
  1. If S is algebraically independent and $\alpha \in K$, then $S \cup \{\alpha\}$ is algebraically independent over F if and only if α is transcendental over $F(S)$.
     - This is the algebraic analogue of the statement that if S is linearly independent, then $S \cup \{\alpha\}$ is linearly independent if and only if α is not in the span of S.
     - Proof: Suppose $S \cup \{\alpha\}$ is algebraically dependent. Then there exists $s_i \in S$ and $p \in F[x]$ with $p(\alpha, s_1, \dots, s_n) = 0$ and $p \neq 0$. View p as a polynomial in its first variable with coefficients in $F[s_1, \dots, s_n]$: there must be at least one term involving α, as otherwise p would give an algebraic dependence in S. Then α is the root of a nonzero polynomial with coefficients in $F[s_1, \dots, s_n] \subseteq F(s_1, \dots, s_n) \subseteq F(S)$, so it is algebraic over $F(S)$.
     - Conversely, suppose that α is algebraic over $F(S)$. Then α is the root of some nonzero polynomial with coefficients in $F(S)$. Each coefficient of this polynomial is an element of $F(S)$; clearing denominators yields a nonzero polynomial p with coefficients in $F[s_1, \dots, s_n]$ for the elements $s_i \in S$ that appear in these coefficients. This polynomial yields an algebraic dependence in $S \cup \{\alpha\}$.
  2. S is a transcendence base of K/F if and only if K is algebraic over $F(S)$.
     - Proof: This follows from (1) and the maximality of transcendence bases: S is a transcendence base if and only if no elements in K can be adjoined to S while preserving algebraic independence, and by (1) this is equivalent to saying that all elements in K are algebraic over $F(S)$.
  3. If T is a subset of K such that $K/F(T)$ is algebraic, then T contains a transcendence base of K/F.
     - Proof: Apply Zorn's lemma to the collection of all algebraically independent subsets of T, partially ordered by inclusion.
     - A maximal element M in this collection must then be a transcendence base for K/F: if $\beta \in K$ then β must be algebraic over $K/F(M)$ by the maximality of M, and then M is a transcendence base by (2).

4. If T is an algebraically independent subset of K, then T can be extended to a transcendence base of K/F.
   - Proof: This is the analogue of the fact that every linearly independent subset can be extended to a basis, and the proof follows from a similar Zorn's lemma argument.
5. If $S = \{s_1, \dots, s_n\}$ is a transcendence base for K/F and $T = \{t_1, \dots, t_m\}$ is any algebraically independent set, then there is a reordering of S, say $\{a_1, \dots, a_n\}$, such that for each $1 \le k \le m$, the set $\{t_1, t_2, \dots, t_k, a_{k+1}, \dots, a_n\}$ is a transcendence base for K/F.
   - Proof: This is the analogue of the replacement theorem for linearly independent sets, and the proof proceeds inductively in essentially the same way. (We will omit the details.)
6. If S is a (finite) transcendence base for K/F, then any subset T of K having larger cardinality than S must be algebraically dependent.
   - Proof: If $S = \{s_1, \dots, s_n\}$ is finite, apply the replacement theorem (5) to S and T. At the end of the replacement, the result is that $\{t_1, \dots, t_n\}$ is a transcendence base. But then by (2), any additional element of T would be algebraic over $\{t_1, \dots, t_n\}$, contradicting the algebraic independence of T.
7. Any two transcendence bases S and T for K/F have the same cardinality.
   - Proof: If the bases are infinite the result is immediate. If S has finite cardinality n, then the result follows by applying (6), since then T's cardinality m must satisfy $m \le n$ (since T is algebraically independent and S is a transcendence base) and also $n \le m$ (since S is algebraically independent and T is a transcendence base).

- The result of the last part of the proposition shows that any two transcendence bases have the same cardinality, and in analogy with the situation for vector spaces, this cardinality behaves somewhat like an extension degree:

- Definition: Let K/F be a field extension. The transcendence degree of K/F, denoted $\text{trdeg}(K/F)$, is the cardinality of any transcendence base of K/F.

- The key property of transcendence degree is that it is additive in towers:

- Proposition (Transcendence in Towers): If $L/K/F$ is a tower of extensions, then $\text{trdeg}(L/F) = \text{trdeg}(L/K) + \text{trdeg}(K/F)$.
   - The idea here is quite simple: we want to show that the union of transcendence bases for K/F and L/K gives a transcendence base for L/F.
   - Proof: First suppose that both $\text{trdeg}(K/F)$ and $\text{trdeg}(L/K)$ are finite, and let $S = \{s_1, \dots, s_n\}$ and $T = \{t_1, \dots, t_m\}$ be transcendence bases for K/F and L/K. Then $S \cap T = \emptyset$ since each t_i is transcendental over K.
   - Furthermore, K is algebraic over $F(S)$, so $K(T)$ is algebraic over $F(T)(S) = F(S \cup T)$ by our results on algebraic extensions.
   - Then since L is algebraic over $K(T)$, we deduce that L is algebraic over $F(S \cup T)$, also by our results on algebraic extensions.
   - Thus, by property (3) above, $S \cup T$ contains a transcendence base of L/F.
   - Finally, we claim $S \cup T$ is algebraically independent over F, so suppose that $p(s_1, \dots, s_n, t_1, \dots, t_m) = 0$ for some $p \in F[x_1, \dots, x_n, y_1, \dots, y_m]$.
   - Separate monomial terms to write $p(s_1, \dots, s_n, t_1, \dots, t_m) = 0$ as a sum $\sum f_i(s_1, \dots, s_n) g_i(t_1, \dots, t_m) = 0$ with $f_i \in F[x_1, \dots, x_n]$ and $g_i \in F[y_1, \dots, y_m]$.
   - Now, since T is algebraically independent over $F(S) \subseteq K$, all of the $f_i(s_1, \dots, s_n)$ must be zero (as elements of K). But since S is algebraically independent over F, that means all of the polynomials $f_i(x_1, \dots, x_n)$ must be zero (as polynomials).
   - This means p is the zero polynomial, and so $S \cup T$ is algebraically independent.

- Fields that are generated by a transcendence base are particularly convenient:

- Definition: The extension K/F is purely transcendental if $K = F(S)$ for some transcendence base S of K/F.

- Equivalently, K/F is purely transcendental when it is generated (as a field extension) by an algebraically independent set.
- If $S = \{s_1, \dots, s_n\}$, then the purely transcendental extension $K = F(S)$ is ring-isomorphic to the function field $F(x_1, \dots, x_n)$ in n variables: it is not hard to check that the map sending s_i to x_i is an isomorphism.
- If K/F has transcendence degree 1 or 2 and E/F is an intermediate extension, then in fact E is also purely transcendental: the degree-1 case is a theorem of Lüroth that we will prove later, while the degree-2 case is a theorem of Castelnuovo. In higher degrees, there do exist extensions that are not purely transcendental, but it is not easy to verify this fact.

- Now let F be a field and K be an extension of F of transcendence degree 1.
  - By the results above, there exists $x \in K$ such that $K/F(x)$ has transcendence degree 0, which is to say, it is algebraic.
  - Since we do not want to worry for the moment about infinite-degree algebraic extensions, we will make the further assumption that this extension $K/F(x)$ has finite degree.

- Definition: We say K is an (algebraic) function field over F if there exists $x \in K$ such that x is transcendental over F and $K/F(x)$ is finite.
  - Example: $\mathbb{Q}(x)$ is an algebraic function field over $\mathbb{Q}$.
  - Example: $\mathbb{C}(x, \sqrt{x^2 - 1})$ is an algebraic function field over $\mathbb{C}$.
  - Note that the algebraic closure of F inside K has finite degree over F: this follows by noting that if E/F is algebraic inside K, then $[E : F] = [E(x) : F(x)] \le [K : F(x)] < \infty$.
  - So, without loss of generality, we may replace F by its algebraic closure inside K. In this case we call F the constant field of K.
  - If F is the constant field of K, then there are no elements of K that are algebraic over F other than the elements of F themselves. Equivalently, every element of $K\backslash F$ is transcendental over F.
  - Finally, since the transcendence degree of K/F is 1, for any two $a, b \in K\backslash F$, there is some nonzero polynomial $g \in F[x, y]$ such that $F[a, b] = 0$.

- Now that we have some very basic facts about function fields, our next goal is to do number theory.
  - In order to do this, however, we need to know how to define primes in the function field context.
  - Over $\mathbb{Q}$, the primes arise as the prime ideals of the ring of integers $\mathbb{Z}$, which we can define starting from $\mathbb{Q}$ purely in terms of integral closures. For other number fields, we also define their primes using integral closures.
  - However, this approach will not work for function fields, because (as noted above) everything in K not in F is transcendental over F, so there is no sensible way to define a "ring of integers" inside K using integrality.
  - Instead, we need to use a different sort of construction to give a sensible notion of a prime: that of a discrete valuation on K.

- In order to develop all of this properly, we also need to review some facts about localization.

- Proposition (Localization): Let R be a commutative ring with 1 and D be a multiplicatively closed subset of R containing 1. Then there exists a commutative ring $D^{-1}R$, the localization of R at D, and a ring homomorphism $\pi : R \to D^{-1}R$ such that any for any ring homomorphism $\psi : R \to S$ sending 1 to 1 and such that $\psi(d)$ is a unit in S for every $d \in D$, there exists a unique homomorphism $\Psi : D^{-1}R \to S$ such that $\Psi \circ \pi = \psi$.
  - More succinctly, any homomorphism $\psi : R \to S$ such that ψ maps all of the elements of D into units necessarily extends to a homomorphism $\Psi : D^{-1}R \to S$.
  - The main idea is simply to define "fractions" r/d with $r \in R$ and $d \in D$ via an appropriate equivalence relation, and then to write down the usual rules of fraction arithmetic and verify that all of the definitions are well posed.

- Proof (outline): Define an equivalence relation on elements of $R \times D$ by setting $(r, d) \sim (s, e)$ whenever there exists $y \in D$ such that $y(ds - er) = 0$; it is straightforward to check that $\sim$ is an equivalence relation.
- Denote the equivalence class of (r, d) by the symbol r/d and the set of all equivalence classes by $D^{-1}R$, and define the two operations $r/d + s/e = (re + ds)/(de)$ and $r/d \cdot s/e = (rs)/(de)$ on $D^{-1}R$. It is tedious but straightforward to see that these operations make $D^{-1}R$ into a commutative ring with 1.
- Now define $\pi(r) = r/1$ and suppose $\Psi : D^{-1}R \to S$ is a homomorphism with $\Psi \circ \pi = \psi$.
- Then we must have $\Psi(r/1) = (\Psi \circ \pi)(r) = \psi(r)$, and also $1 = \Psi(1/1) = \Psi(1/d)\Psi(d/1)$, meaning that $\Psi(1/d) = \psi(d)^{-1}$. Then $\Psi(r/d) = \Psi(r/1)\Psi(1/d) = \psi(r)\psi(d)^{-1}$.
- But it is easy to see that this choice of Ψ does work, so it is the only such homomorphism.

- The point here is that $D^{-1}R$ is the smallest ring in which all elements of D become units.
  - When D contains no zero divisors (which is automatically the case if R is a domain and D does not contain zero), then R injects into $D^{-1}R$ via $r \mapsto r/1$.
  - A particular useful case of localization is to construct $\mathbb{Q}$ from $\mathbb{Z}$ (we take $D = \mathbb{Z}\backslash\{0\}$ and $R = \mathbb{Z}$) or more generally to construct the field of fractions of an integral domain R (take $D = R\backslash\{0\}$).
- We also note in passing that we can localize any R-module M in the same way: one simply writes down the same construction using pairs (m, d) with $m \in M$ and $d \in D$ in place of pairs (r, d).
  - Alternatively, one can obtain the localization of an R-module using tensor products: $D^{-1}M \cong M \otimes_R D^{-1}R$. (This tensor product just extends scalars from R to $D^{-1}R$, which is exactly what $D^{-1}M$ is.)
  - Exercise: Show that localization commutes with sums, intersections, quotients, finite direct sums, and is exact.
  - Exercise: Show that if I is an ideal of R, then $D = R\backslash I$ is multiplicatively closed if and only if I is prime.
- Our main situation of interest is that of localizing at a prime: this is the case where R is an integral domain and $D = R\backslash P$ is the complement of a prime ideal P of R.
  - Exercise: Show that if P is a prime ideal and $D = R\backslash P$, then $D^{-1}R$ is a local ring with unique maximal ideal $D^{-1}P = \pi(P) = e_P$, the extension of the ideal P to $D^{-1}R$.
  - The utility of localizing at a prime is that it isolates the ring's behavior at that prime.
  - Example: The localization of $\mathbb{Z}$ at the prime ideal (p) is the ring $\mathbb{Z}_{(p)} = \{a/b \in \mathbb{Q} : p \nmid b\}$ of rational numbers whose denominator is not divisible by p. Its unique maximal ideal is $p\mathbb{Z}_{(p)}$, the set of multiples of p. The quotient ring $\mathbb{Z}_{(p)}/p\mathbb{Z}_{(p)}$ is isomorphic to $\mathbb{Z}/p\mathbb{Z}$.
  - Note that $\mathbb{Z}_{(p)}$ is not the ring of p-adic integers $\mathbb{Z}_p$: the p-adic integers are obtained by taking a completion of the localization $\mathbb{Z}_{(p)}$ under the p-adic metric (which we will define later).
  - Example: Let k be a field and take R to be the ring of k-valued functions on a set S. If we let M_a be the set of functions vanishing at a point $a \in S$, then M_a is a maximal ideal of R. The localization $R_{M_a} = \{f/g \in R : g(a) \neq 0\}$ is the ring of k-valued rational functions defined at a. The unique maximal ideal of M_a is the ideal of all k-valued rational functions vanshing at a.
- This second example illustrates the utility of localizing at a prime, because it allows us to study the local behavior of a rational function near the point a.
  - For example, the elements of M_a are precisely those rational functions vanishing at a, while the elements of M_a^2 are the rational functions that vanish to order 2 at a (i.e., have a double root), and so forth.
  - More generally, if we localize a domain at a principal prime ideal, by looking at powers of the maximal ideal, we can measure what power of a prime a given element is divisible by.
- We will formalize all of this using discrete valuations, which provide us a way to identify primes using only the field structure, next time.

Well, you're at the end of my handout. Hope it was helpful.