The analytic class number formula

January 4, 2014

Let F be a number field, say of degree $d = [F : \mathbb{Q}]$. Our goal today is to study the Dedekind zeta function

$$\zeta_F(s) = \sum_{I \neq 0} N(I)^{-s}$$

as $s \to 1^+$. Here I runs over all nonzero ideals of $\mathcal{O}_F$, and $N(I) = \#\mathcal{O}_F/I$.

The result will be that $\zeta_F(s)$ converges for $s > 1$, but that convergence past this region is obstructed by a "simple pole" at $s = 1$. Furthermore, and this is the key, this simple pole has a residue which can be expressed in terms of arithmetic invariants of F. More precisely:

Theorem 0.1. *Let F be a number field, as above. Then $\zeta_F(s)$ converges for real $s > 1$, and we have*

$$\lim_{s \to 1^+} (s-1)\zeta_F(s) = \frac{M \cdot h}{A},$$

where:

1. *h is the* class number *of F. Recall that two nonzero ideals $I, J \subset \mathcal{O}_F$ are called "isomorphic" if there is an $\alpha \in F^\times$ with $I = \alpha J$, and that the class number is the (finite!) number of such isomorphism classes. So $h = 1$ if and only if every nonzero ideal of $\mathcal{O}_F$ is principal, or if and only if $\mathcal{O}_F$ has unique prime factorization.*

2. *A is the volume of the torus $F_\mathbb{R}/\mathcal{O}_F$, where $F_\mathbb{R}$ is the Minkowski space of F with its canonical inner product. Recall that A is the square root of a natural number (the absolute value of the "discriminant", in usual terminology).*

3. *M is the "multiplicative mystery factor". (Related to the "regulator", in usual terminlogy.) We'll describe this M in the course of the proof. For now let me just point out two features: one, it appears to be transcendental in nature, unlike h and A; and two, it's related to the group of units $\mathcal{O}_F^\times$.*

Now let's give the proof. Actually, we will skip a step or two regarding the factor M. In the end, M is not really our main concern: we'll be most interested in situations where, by design, it more or less cancels out.

The first step is a simple one. The sum defining $\zeta_F(s)$ is indexed by the nonzero ideals. We start by partitioning these ideals into their isomorphism classes. We will write the isomorphism class of a nonzero ideal J as $[J]$. In the following expression, J runs over a set of representatives for the isomorphism classes of ideals.

$$\zeta_F(s) = \sum_J \sum_{I \in [J]} N(I)^{-s}.$$

The first sum has h terms. Thus we see that it suffices to prove the analog of the above theorem for the series defined by the inner sum, but where for the residue we instead get M/A, independently of the fixed ideal J.

So let's work on this. The first question to ask is, how do we parametrize the ideals $I \in [J]$? We know that $I \in [J]$ if and only if there exists an $\alpha \in F^\times$ such that $\alpha J = I$. So every ideal in the class $[J]$ is of the form αJ for some $\alpha \in F^\times$. But this is not an exact parametrization, for two reasons:

1. First, not every α actually defines an ideal $I = \alpha J$. We need to require that $\alpha J \subset \mathcal{O}_F$.

2. Second, two different α, α' can give rise to the same ideal $\alpha J = \alpha' J$. This happens if and only if there is a unit $u \in \mathcal{O}_F^\times$ with $\alpha' = u\alpha$. (Proof: by invoking unique prime ideal factorization and clearing denominators in α and α', one can reduce to the fact that $(x) = (y)$ if and only if $x = uy$ for some unit u whenever $x, y \in \mathcal{O}_F$.)

Returning to the first point, we denote the set of $\alpha \in F$ with $\alpha J \subset \mathcal{O}_F$ by J^{-1}. Here is a lemma about this set. It is also tied up with the unique prime factorization, but we will give a proof, just to give a taste of those kinds of arguments.

Lemma 0.2. *Let J be a nonzero ideal of $\mathcal{O}_F$, and define*

$$J^{-1} = \{\alpha \in F \mid \alpha J \subset \mathcal{O}_F\}.$$

Then J^{-1} is a free abelian group of rank $d = [F : \mathbb{Q}]$, and in fact $\mathcal{O}_F \subset J^{-1}$ with index $N(J)$.

Proof. Note that J^{-1} is an abelian group, in fact a subgroup of F under addition. In particular J^{-1} is torsion-free, so by the classification of finitely generated abelian groups, the first assertion follows from the second. So it suffices to see that $\mathcal{O}_F \subset J^{-1}$ has index $N(J)$.

First, note that multiplication by $N(J)$ does kill the quotient $J^{-1}/\mathcal{O}_F$: given any $\alpha \in J^{-1}$, we have $N(J)\alpha \cdot J \subset N(J)\mathcal{O}_F \subset J$. But J is a finite rank abelian group, so this means that the characteristic polynomial of multiplication by $N(J)\alpha$ has integer coefficients, so $N(J)\alpha \in \mathcal{O}_F$, as claimed.

The next observation is that the claim is true if J is principal, say $J = (x)$. That's because multiplication by x gives an isomorphism $(x)^{-1}/\mathcal{O}_F \simeq \mathcal{O}_F/(x)$, so the cardinalities are the same.

Now, not every ideal is principal. But we note that multiplication by $x \in J$ can give an isomorphism $J^{-1}/\mathcal{O}_F \simeq \mathcal{O}_F/J$ even if x doesn't exactly generate J. Indeed, since both of these finite abelian groups are killed by $N(J)$, it suffices to ensure that the index of (x) in J is relatively prime to $N(J)$.

Briefly, this can be arranged follows: if $J = P_1^{n_1} \dots P_k^{n_k}$ is the maximal ideal factorization of J, then choose (using the "Chinese remainder theorem") an $x \in \mathcal{O}_F$ with $x \in P_i^{n_i} - P_i^{n_i+1}$ for all $i = 1, \dots k$. Then we'll have $x \in J$, and the index $[J : (x)]$ will be prime to $N(J)$. Thus multiplication by x gives the desired isomorphism, so the claim holds. ☐

OK, now let's come back to our analysis. We have seen that the set of ideals I in $[J]$ is in bijection with the set

$$(J^{-1} - 0)/\mathcal{O}_F^\times,$$

the quotient of the set $J^{-1} - 0$ by the action (by multiplication) of the group of units $\mathcal{O}_F^\times$. Under this bijection, an $\alpha \in J^{-1} - 0$ corresponds to the ideal $I = \alpha J$. Thus our above inner sum is

$$\sum_{I \in [J]} N(I)^{-s} = \sum_{\alpha \in (J^{-1}-0)/\mathcal{O}_F^\times} N(\alpha J)^{-s}.$$

Now, we have previously seen that the norm of ideals is multiplicative, and that the norm of a principal ideal (x) is the same as the absolute value of the norm of the element $|N(x)|$ (this came from problem 1 of PSet 2). (Recall that for $x \in F$, the number $N(x) \in \mathbb{Q}$ can be defined as the product over all the complex conjugates of x, or as the determinant of multiplication by x acting on $F/\mathbb{Q}$.) It follows that $N(\alpha J) = |N(\alpha)| \cdot N(J)$, and so we can pull out $N(J)^{-s}$ from the sum, getting

$$\sum_{I \in [J]} N(I)^{-s} = N(J)^{-s} \cdot \sum_{\alpha \in (J^{-1}-0)/\mathcal{O}_F^\times} |N(\alpha)|^{-s}.$$

Now, as $s \to 1^+$ the factor $N(J)^{-s}$ tends to the finite number $N(J)^{-1}$. Thus we see that it will suffice to prove the analog of our theorem for the sum $\sum_{\alpha \in (J^{-1}-0)/\mathcal{O}_F^\times} |N(\alpha)|^{-s}$, except for the residue at $s = 1$ we should find the number $N(J) \cdot M/A$. Note, however, that due to the above Lemma, J^{-1} is a full lattice in Minkowski space $F_\mathbb{R}$, and the quantity $A/N(J) = Vol(F_\mathbb{R}/\mathcal{O}_F)/N(J)$ is just the same as the volume $Vol(F_\mathbb{R}/J^{-1})$. So we can rewrite our target number $N(J) \cdot M/A$ as $M/Vol(F_\mathbb{R}/J^{-1})$.

Thus we observe (following the Oslo notes I linked to on the website) that our claim can be viewed intrinsically to Minkowski space. Namely:

1. J^{-1} can be viewed as a full rank lattice in Minkowski space $F_\mathbb{R}$.
2. The action of $\mathcal{O}_F^\times$ on J^{-1} is induced from action of $\mathcal{O}_F^\times$ on $F_\mathbb{R}$: just embed and multiply on each coordinate.
3. The function $\alpha \mapsto |N(\alpha)|$ also extends to Minkowski space: viewing Minkowski space as a subset of $\prod_{F \to \mathbb{C}} \mathbb{C}$, we just take the absolute value of the product of these $\mathbb{C}$-coordinates.

Now, we will need one bit of geometric/arithmetic input from this situation, which we will not fully explain. It is the following:

Theorem 0.3. *There exists a fundamental domain D for the action of $\mathcal{O}_F^\times$ on $F_\mathbb{R}$ with the following two properties:*

1. *D is a cone: for any $t \in \mathbb{R}_{>0}$, we have $tD = D$.*
2. *The subset $D_{\leq 1} = \{x \in D \mid |N(x)| \leq 1\}$ of $F_\mathbb{R}$ is bounded and has finite volume.*

Instead of giving a proof, let me just make some remarks, and then we'll look at two simple examples. The first remark is that both properties can be simultaneously arranged as follows. Consider the subset of Minkowski space consisting of the elements x with $|N(x)| = 1$ (the "norm-1 hypersurface"). Then $\mathcal{O}_F^\times$ still acts on this subset. If we can find a bounded, finite volume fundamental domain $D_{=1}$ for this action, then by taking $D = \mathbb{R}_{\geq 0} \cdot D_{=1}$ we'll fulfill both of the above requirements.

The second remark is that this norm-1-hypersurface is itself a group under multiplication, and except for a finite kernel, $\mathcal{O}_F^\times$ is actually a subgroup of it. Thus the problem of finding this $D_{=1}$ is very much like that of finding a fundamental domain for a lattice in a real vector space, just with multiplication instead of addition. But actually, by taking logarithms, you can exactly move into the familiar additive context.

What we find is a "logarithmic" embedding of $\mathcal{O}_F^\times/\{\text{finite group of roots of unity}\}$ into a real vector space of dimension $\mathbb{R}^{r+s-1}$. Then the main claim (due to Dirichlet, like everything we've been discussing recently) is that this embedding actually gives a full lattice. (This is the "Dirichlet unit theorem"). That's what lets us find the right kind of fundamental domain. It also tells us exactly how big $\mathcal{O}_F^\times$ is: it has a finite torsion subgroup, and the quotient is free of rank $r+s-1$.

The last remark is that we can now define our "multiplicative mystery factor": it is the volume $M = Vol(D_{\leq 1})$, if D is chosen as outlined above.

OK, let's look at two examples, namely the imaginary quadratic fields and the real quadratic fields. They exhibit different behavior, and the general case is sort of a combination of the two of them.

First, suppose F is an imaginary quadratic field, say $F = \mathbb{Q}(\sqrt{-d})$ with d a squarefree natural number. In this case, the Minkowski space is just $\mathbb{C}$. Also, and this is very special, the group $\mathcal{O}_F^\times$ is finite. In fact, we have $\mathcal{O}_F^\times = \{\pm 1\}$ unless $d = 1$ or $d = 3$, in which cases $\mathcal{O}_F^\times = \{1, i, i^2, i^3\}$ and $\mathcal{O}_F^\times = \{1, \zeta_6, \ldots, \zeta_6^5\}$ respectively. This is because, recall, $x \in \mathcal{O}_F$ is a unit if and only if $N(x) = \pm 1$. But $N(a + b\sqrt{-d}) = a^2 + db^2$, and squares are nonnegative. So even in the case where a and b can be half-integers, there are only finitely many possible a and b for which $N(a + b\sqrt{-d})$ can be ± 1, and it's easy to explicitly enumerate these, getting the desired claim.

In all cases, our group $\mathcal{O}_F^\times$ is finite cyclic and acts by rotations on Minkowski space $\mathbb{C}$. What this means is that when $d \neq 1, 3$, we can choose our fundamental cone D to just be any half-plane in $\mathbb{C}$. When $d = 1$, we instead take D to be a quarter-plane, and when $d = 3$ we make it a $1/6$-plane. Thus $D_{\leq 1}$ is the appropriate ratio of a unit circle, meaning we have $M = \pi$ for $d \neq 1, 3$ and $M = \pi/2$ for $d = 1$ and $M = \pi/3$ for $d = 3$. (If you thought a unit circle had area π instead of 2π, that's because you were using the naive Minkowski inner product.)

If F is a real quadratic field instead, say $F = \mathbb{Q}(\sqrt{d})$ with d a squarefree natural number, then Minkowski space is $\mathbb{R} \oplus \mathbb{R}$, and the group of units is no longer finite. Instead, we have $\mathcal{O}_F = \pm 1 \times u^{\mathbb{Z}}$, where $u \in \mathcal{O}_F^\times$ is a *fundamental unit* for F. Here the norm function is $N(a + b\sqrt{d}) = a^2 - db^2$, so the norm-1 hypersurface is not a circle, but a hyperbola, and u acts by translating along this hyperbola (and maybe switching the branches simultaenously). We find that the region $D_{\leq 1}$ can be written as a union of two little wedges which extend from 0 to the norm-1-hypersurface. That's all we'll say about the real quadratic case.

Now, granting the above theorem, we are reduced to the following abstract claim:

Proposition 0.4. *Let V be a finite-dimensional real vector space with a notion of volume (say, given by an inner product), and let L be a full lattice in V. Suppose also given the following:*

1. *A subset $D \subset V$ which is a cone ($tD = D$ for $t > 0$);*
2. *A continuous function $F : V \to \mathbb{R}_{\geq 0}$ which is homogeneous of degree $d = dim(V)$, i.e. $F(tx) = t^d F(x)$ for all $t \geq 0$.*

Also, assume that the set $D_{\leq 1} = \{x \in D \mid F(x) \leq 1\}$ is bounded and has finite volume.

Then the series

$$z(s) = \sum_{x \in L \cap D} F(x)^{-s}$$

converges for $s > 1$, and

$$\lim_{s \to 1^+} (s-1) z(s) = Vol(D_{\leq 1})/Vol(V/L).$$

The proof will be in two stages. Let us use the notation $f(s) \simeq g(s)$ to indicate that f and g have the same first order behavior as $s \to 1^+$, meaning that $f(s) - g(s)$ gives a continuous function for $s > 1$, and $(s-1)(f(s) - g(s))$ extends to a continuous function for $s \geq 1$, with value 0 at $s = 1$. The steps are as follows:

1. First,
$$\sum_{x \in L \cap D} F(x)^{-s} \simeq \frac{1}{Vol(V/L)} \int_{x \in D} F(x)^{-s}.$$

2. Second,
$$\int_{x \in D} F(x)^{-s} \simeq \frac{Vol(D_{\leq 1})}{s-1}.$$

The idea behind the first step is simple: consider the fundamental parallelogram $\mathcal{P}_x$ for the lattice L, centered at a given point $x \in L \cap D$. As x varies over $L \cap D$, these parallelograms neatly cover the region D, except for some error near the boundary ∂D. Now, define a function $\widetilde{F} : V \to \mathbb{R}_{\geq 0}$ by having its value inside $\mathcal{P}_x$ be the constant $F(x)^{-s}$, and having it be 0 outside the union of these parallelograms $\mathcal{P}_x$. Also let $F|_D : V \to \mathbb{R}_{\geq 0}$ denote the function which agrees with F on D and is 0 elsewhere. Then the idea is that $\widetilde{F}$ and $F|_D$ are fairly close to each other. The difference is twofold: first, there are discrepancies near ∂D; and second, the values are still only nearby, not exactly equal, in the interior of D. However, one can see, by using the "mean value theorem" $|x^{-s} - y^{-s}| \leq -s \cdot \max\{x^{-s-1}, y^{-s-1}\}$ and following the kind of argument we'll give in establishing the second step, that these discrepancies don't matter to first order, so that $\int_V \widetilde{F} \simeq \int_V F|_D$, or equivalently

$$Vol(V/L) \cdot \sum_{x \in L \cap D} F(x)^{-s} \simeq \int_{x \in D} F(x)^{-s}.$$

That gives step 1 (modulo us not actually carrying out the details). As for step 2, we can calculate $\int_{x \in D} F(x)^{-s}$ by breaking up the integral over the level sets of F, say $D_t = \{x \in D \mid F(x) = t\}$ for $t \geq 0$. This gives

$$\int_{x \in D} F(x)^{-s} = \int_{t > 0} Vol(D_t) \cdot t^{-s} dt.$$

We can calculate the volume $Vol(D_t)$ as

$$Vol(D_t) = \frac{d}{dt} Vol(D_{\leq t}) = \frac{d}{dt} \left(t \cdot Vol(D_{\leq 1})\right) = Vol(D_{\leq 1}),$$

where the second step is because F is homogeneous of degree d, so that $D_{\leq t} = t^{1/d} \cdot D_{\leq 1}$. Thus we have

$$\int_{x \in D} F(x)^{-s} = Vol(D_{\leq 1}) \int_{t > 0} t^{-s} dt.$$

For $s > 1$, we can explicitly evaluate this integral by an antiderivative of t^{-s}, getting

$$\int_{x \in D} F(x)^{-s} = Vol(D_{\leq 1}) \cdot \frac{s}{s-1}.$$

Letting $s \to 1^+$, we evidently get the desired result.

Thus we've proven the main theorem: we've seen that as $s \to 1^+$,

$$(s-1)\zeta_F(s) \to M \cdot h/A,$$

where h is the class number, A is the additive volume $Vol(F_{\mathbb{R}}/\mathcal{O}_F)$, and M is the multiplicative volume, $Vol(D_{\leq 1})$ where D is a fundamental cone for the action of $\mathcal{O}_F^\times$ on $F_{\mathbb{R}}$.

The next thing we'll want to do is try to apply this in the case $F = \mathbb{Q}(\zeta_p)$. Recall that we previously had the factorization

$$\zeta_{\mathbb{Q}(\zeta_p)}(s) = \zeta(s) \cdot \prod_{\chi \neq 1} L(s, \chi).$$

Now we've analyzed the behavior of each of these pieces as $s \to 1^+$, and we're all set up to plug this in and make the comparison.

But the transcendental factors on both sides (coming from M on the left, and the log's of sin's on the right) make this a little bit delicate. So we'll actually try to move to a context where those factors don't appear. This can be done by cleverly selecting certain subfields $F \subset \mathbb{Q}(\zeta_p)$ for consideration. On the right-hand side, this will amount to throwing out certain of the χ's. Namely:

Proposition 0.5. *Let p be a prime, and $F \subset \mathbb{Q}(\zeta_p)$ a subfield of the p^{th} cyclotomic field. Let $H \subset (\mathbb{Z}/p\mathbb{Z})^\times$ be the subgroup to which F corresponds in the Galois correspondence. Then*

$$\zeta_F(s) = \zeta(s) \cdot \prod_{\chi \neq 1, \chi(H)=1} L(s, \chi).$$

The extreme examples are $F = \mathbb{Q}$ and $F = \mathbb{Q}(\zeta_p)$. The first case corresponds to $H = (\mathbb{Z}/p\mathbb{Z})^\times$, so there are no $\chi \neq 1$ with $\chi(H) = 1$. Thus we just find the Riemann zeta function on the right, as we should. The second case corresponds to $H = 1$, so every χ satisfies $\chi(H) = 1$, and we recover the result we already knew. In general, we can identify the characters $\chi : (\mathbb{Z}/p\mathbb{Z})^\times \to \mathbb{C}^\times$ satisfying $\chi(H) = 1$ with the characters of the quotient $(\mathbb{Z}/p\mathbb{Z})^\times/H$, or in other words the characters of the Galois group $Gal(F/\mathbb{Q})$. If we think in those terms, the above factorization actually exists in a great generality, as noticed by Artin. But we won't get into that story, and actually we'll just take the above proposition on faith.

To get rid of the transcendental factors on the right, we'll want to pick $H = \pm 1$, to isolate the even characters from the odd ones. But for now let's look at a simpler case. For p odd, let H denote the index two subgroup of the cyclic group $(\mathbb{Z}/p\mathbb{Z})^\times$, or in other words the subgroup of nonzero *squares* modulo p. Then there is exactly one character χ of $(\mathbb{Z}/p\mathbb{Z})^\times$ which is trivial on H, namely the Legendre symbol

$$\chi(\cdot) = \left(\frac{\cdot}{p}\right),$$

defined by $\left(\frac{a}{p}\right) = 1$ if a is a square, and -1 if it is not.

This Legendre symbol is an odd character if and only if $\left(\frac{-1}{p}\right) = -1$, i.e. if and only if -1 is not a square (mod p), or if and only if p is 3 (mod 4).

So now assume p is 3 (mod 4). The subgroup H of index 2 corresponds to the unique quadratic subextension F of $\mathbb{Q}(\zeta_p)$. We know from the first homework (though there are easier ways...) that

$$F = \mathbb{Q}(\sqrt{-p}).$$

Thus we deduce the following:

$$\zeta_{\mathbb{Q}(\sqrt{-p})}(s) = \zeta(s) \cdot L(s, \chi),$$

when p is 3 (mod 4) and χ is the Legendre character (mod p). Note here that $\mathbb{Q}(\sqrt{-p})$ being imaginary quadratic exactly matches up with χ being odd: the first implies the lack of transcendental factors on the left as $s \to 1^+$, and the second implies the same for the right.

Now let's let $s \to 1^+$ and see what we get. You should be holding your breath right now, because there's always the chance that we'll end up with a tautology: perhaps our separate analyses of both sides of this equation were really just doing the same thing in different language. But instead we'll get lucky, and find a very simple but mysterious formula for the class number of $\mathbb{Q}(\sqrt{-p})$.

Assume $p > 3$ for simplicity. When we multiply by $s - 1$ and let $s \to 1^+$, on the left we get

$$M \cdot h / Vol(\mathcal{O}_{\mathbb{Q}(\sqrt{-p})}).$$

We already computed M above: it's just π, since $D_{\leq 1}$ is a half-disc in $\mathbb{C}$. As for $Vol(\mathcal{O}_{\mathbb{Q}(\sqrt{-p})})$, it's not hard to calculate: $\mathcal{O}_{\mathbb{Q}(\sqrt{-p})}$ as a $\mathbb{Z}$-basis given by $(1, \frac{1+\sqrt{-p}}{2})$, so the volume is the absolute value of the determinant of the matrix whose first row is that vector and whose second is the conjugate vector $(1, \frac{1-\sqrt{-p}}{2})$. This gives $Vol(\mathcal{O}_{\mathbb{Q}(\sqrt{-p})}) = \sqrt{p}$. Thus on the left we get

$$\frac{\pi}{\sqrt{p}} \cdot h_p.$$

On the right, on the other hand, since $\lim_{s \to 1^+}(s-1)\zeta(s) = 1$, we just get $L(1, \chi)$. But we know a formula for $L(1, x)$; it's

$$L(1, \chi) = \pi \cdot \frac{i \cdot g(\chi)}{p^2} \sum_{r=1}^{p-1} \overline{\chi(r)} \cdot r.$$

We can get rid of the complex conjugation on $\chi(r)$, because $\chi(r) = \pm 1$ is always real. The only mystery, then, is the Gauss sum

$$g(\chi) = \sum_{k=1}^{p-1} \chi(k) \cdot \zeta_p^k.$$

However, it turns out that $g(\chi)$ is just $i \cdot \sqrt{p}$. It doesn't take too long to prove that $g(\chi)^2 = -p$, so the only issue here is whether $g(\chi)$ is $i \cdot \sqrt{p}$ or $-i \cdot \sqrt{p}$. This is the problem of the "sign of the Gauss sum", which kept Gauss entertained on and off for four years. So it's not easy to resolve. But actually we could do it by comparison with the first problem of problem set one. Instead let's just take for granted that $g(\chi)$ is indeed $i \cdot \sqrt{p}$. Then by comparing the two sides, we get a formula for the class number h_{-p} of $\mathbb{Q}(\sqrt{-p})$ when p is a prime congruent to 3 (mod 4):

$$h_{-p} = -\frac{1}{p} \sum_{r=1}^{p-1} \left(\frac{r}{p}\right) r.$$

For example, let's use this formula to show that $\mathbb{Z}[\frac{1+\sqrt{-19}}{2}]$ has unique prime factorization. To calculate $\left(\frac{\cdot}{19}\right)$, we need to figure out what the squares are $\pmod{19}$. We can just keep generating squares until we get the correct number of them, namely $18/2 = 9$. To get from n^2 to $(n+1)^2$ you can just add $2n+1$, so this is pretty easy. We get the list

$$\{1, 4, 9, 16, 6, 17, 11, 7, 5\} = \{1, 4, 5, 6, 7, 11, 16, 17\}$$

of squares mod 19. Thus our sum is

$$-\frac{1}{19}\left(1-2-3+4+5+6+7-8-9-10+11-12-13-14-15+16+17\right) = -\frac{-19}{19} = 1!$$