Talk 0: introduction and logistics

Additive number theory seminar

1 On giving (and attending) talks

1.1 On attending talks

In this class, you'll be asked to fill out a feedback form on each presentation, which will also give you some guidelines for what to look for in a talk—what is working well? what isn't? what takeaways can you get? Beyond that, here are some general guidelines for following talks (not just in this class); none are required, but may be useful to you.

- The ideal is to follow the talk all the way through. This is not always possible, partially because many talks are bad and partially because for many you won't have the background. If you lose the thread, it's easy to just stop paying attention; try instead to find it again.
- It can be helpful or unhelpful to try to take notes during talks; see what works for you.
- Keep in mind some basic questions, and try to answer them for yourself by the end of the talk: what question is the speaker trying to answer? Why is it interesting or important? What is the answer, at least very roughly speaking? Is there a simple example where you can concretely understand the question and answer?
- Try to learn at least one (fairly concrete thing) from each talk (e.g. a definition, a theorem of interest, a proof idea). This is not a lofty goal, but it'll add up.
- Try to formulate at least one interesting question (perhaps write it down to remember it). If it isn't answered by the end of the talk, ask it.

For feedback, be precise, critical, and constructive. For example, "your talk sucked" is not useful feedback, but neither is "all in all pretty good!" If your classmates do something you liked in their presentations, it's worth pointing it out so they remember to keep doing it; conversely if they do something that makes it harder for you to follow, to help them improve it's important to tell them.

At the end of each talk, I'll collect your feedback forms. These **must** have your names written on them; rather than being directly passed to the speakers, I'll collect them and pass on your feedback **anonymously**.

1.2 Giving presentations: general principles

This is fairly general advice, some of which will be more or less applicable to this class in particular.

- Structure your talks, similar to a paper: introduce your topic, outline the structure; at the end give some concluding remarks. (Like a paper, think of a talk as telling a story, and never stray far from the plot.) Here it can be reasonable for your conclusion to be more of a summary, as your audience may have forgotten and can't scroll back like in

a paper. In general it can be worth referring to the bigger picture throughout the talk, to remind your audience why we're doing everything. In this class, your talk is part of a larger sequence of talks, so it's useful to relate your talk to whatever came before (and, if possible, what's coming next, to make it easier for your audience to keep track.

- However, avoid spending too much of your talk describing what you're going to do rather than doing it.
- Keep in mind your audience: lay? peers? experts? (In this case peers, but you likely have more background on your specific topic.)
- Often your audience will be mixed. Try to have something accessible and of interest to each (major) segment, e.g. a sketch for laypeople, a somewhat more technical discussion for non-expert peers, and deeper technical stuff for experts. (Less applicable in this class.)
- Prepare in advance—in this class you'll give a practice talk to me, and I'll ask you to have either slides or notes prepared. Your slides or notes will be posted on the class website after your talk as a reference for your classmates, so make sure they're readable. (Notes from this class will likewise be posted, after being edited to include the tentative schedule.)
- Engage the audience: gestures, motion, tone... Look at your audience as much as possible.
- Interact with the audience, if suitable: solicit questions, ask questions of the audience, etc. (Use your judgment here: in formal contexts this is sometimes frowned-upon.)
- Props: notes can be useful, but you should know your talk well enough that you only need to consult them occasionally to find your place, rather than doing the talk from them. (I haven't practiced this talk, so you'll probably notice I'm not doing this very well.)
- Slides are often good, especially if you need to show something—a graph, a figure, a dataset, a diagram. They are *not* good for a lot of text. Never read off your slides; talk about them!
- If you have slides, make sure they're readable. (You can write slides with LaTeX using Beamer.)
- If you're handwriting (e.g. on a blackboard), make sure your handwriting is readable (in particular large enough).
- Think about the layout of your boards, if there are several. Ideally practice in the same room you'll give your talk in, or use a similar board layout.
- Dually to how to attend a talk, think about what you want your audience to get out of it: ideally they'd follow the whole thing, but if not what's one takeaway they can get? What problem are you trying to solve, why should your audience care, what's your method, what's your answer, what are some examples in simple illustrative cases?
- For a one-off talk, it's usually best to avoid technical details as possible, at least for most of the talk. (You can have a more technical section if you want, but be aware

much of your audience may tune out.) The important thing is the big picture: you won't be able to give every detail during a talk. In this class, since your talk is part of a program, you can go into more detail: we don't need every last technical detail of every proof, but you should at least provide the ideas, and be able to go through the details if asked. Important results should be fully proved in most cases.

- For all talks, including in this class: finishing *way* too early isn't great, since if you have the time you should use it, but finishing a little early can be good since it gives extra time for questions, and finishing late is very, very bad. (Starting on time is also important, though not quite as critical.)
- Omit extraneous detail, animations, distracting graphics, writing out overly long formulas; respect your audience's time (and your own).
- Don't talk too fast. (This is difficult when nervous.) Pacing is especially hard if you have slides: since everything is pre-written, it's easy to go through it way too fast. Slow down. It is much better to omit some details or not cover all of your material than to go too fast and leave your audience confused about everything.
- For questions: if you don't know the answer, that's okay and you can admit it; don't imply that you're stupid for not knowing. Conversely never imply that a question is stupid; you can however say something to the effect that it isn't closely related (if true). Partial answers or referrals to sources are also very reasonable.
- A formal academic talk should include at least your main sources (admittedly it can be more or less clear what this means), though not necessarily a full bibliography. A more informal talk can omit them, but if your work is very close to someone else's (or if you're presenting on someone else's work) you should certainly say so. In this class, we know at least your basic sources; if you heavily used another source that you found particularly useful, it's certainly not bad to mention.

A little more about slides vs. chalk talks: the main advantage and disadvantage of slides is that you can present more material. For things like showing off long formulas, graphs, charts, or other visuals, this can be very useful: drawing out your whole figure by hand might take a long time and be difficult to do accurately, but it can still be very helpful for your audience to follow (if it's a good graphic). Slides also act as built-in notes, helping you keep your pace, and they can help make your presentation look more polished. On the other hand, slides also enable you to go too fast, while having to write things by hand keeps you at a slower pace, which is good for audience comprehension. Using a blackboard is also a little more flexible: you can amend things to clarify, answer audience questions more easily, etc., while it's hard to incorporate new stuff into your slides. Slides also take some time and technical background to prepare, though you shouldn't necessarily be preparing any less for chalk talks.

For a 45-50 minute talk, I tend to think chalk talks are usually better, if for no other reason than that you'd need a lot of slides. For shorter talks where you're more crammed for space, the ability of slides to condense your material is much more useful, e.g. I'd probably recommend them for 10-minute talks. In either case though it's a judgment call and you're free to do either.

2 Lagrange's theorem

2.1 Number theory review

Let's quickly sketch some background on prime numbers, as needed. Prime numbers are positive integers with exactly two divisors, 1 and themselves (so we do not consider 1 to be prime): the first few are 2, 3, 5, 7, ... There are infinitely many prime numbers. In fact, we can give a much more precise statement:

Theorem (Prime number theorem, Hadamard–de Vallée-Poussin)**.** *The number $\pi(x)$ of primes less than or equal to x is approximately $\frac{x}{\log x}$, in the sense that*

$$\lim_{x\to\infty} \frac{\pi(x)}{x/\log x} = 1.$$

For example, if $x = 1000000$, the prime number theorem predicts that there should be about $\frac{1000000}{\log(1000000)} \approx 72382.4$ prime numbers less than or equal to 1000000. In fact, one can compute that the true number is 78498, so we're off by about 8%. One of the major goals of analytic number theory is putting bounds on the error in the above approximation (or a slightly more precise version). Since $\frac{x}{\log(x)} \to \infty$ as $x \to \infty$, the infinitude of the primes follows.

Though this is the foundational result on the asymptotic distribution of the primes, their most important property is much older, going back to Al-Farisi in the 13th century or, implicitly, to Euclid in the 3-4th century BCE.

Theorem (Fundamental theorem of arithmetic)**.** *Every positive integer has a unique factorization as a product of prime numbers.*

Here "unique" means unique up to order—e.g. $12 = 2 \cdot 2 \cdot 3 = 2 \cdot 3 \cdot 2$, but we don't count these as different factorizations. Note also that repeating primes is allowed.

This theorem means that we can view the prime numbers as the multiplicative building blocks of the integers. A common use for this theorem is as follows: say we want to prove some property of integers $P(n)$ for every positive integer n. In some cases, it turns out that P is multiplicative, i.e. if $P(m)$ and $P(n)$ both hold then $P(mn)$ also holds. In this case, it suffices to prove $P(p)$ for every prime number p, where we can often apply techniques that wouldn't work more generally; and then since every n can be written as a product of primes, $P(n) = P(p_1 \cdot p_2 \cdots p_r)$ is true since each $P(p_i)$ is true.

Strictly speaking, this omits the case of $P(1)$, since 1 has no prime factors; it is typically trivial, but we can also check 1 separately to be sure.

This will be our broad strategy to prove Lagrange's theorem, though we'll need to do some more work to apply it.

2.2 Linear algebra review

I don't want to get too deep into reviewing linear algebra as it could take all semester. Let's say only this: we can think of vectors (badly) as lists of numbers $(x_1, x_2, \ldots, x_n)$, and

matrices (or linear transforms) as $m \times n$ arrays

$$\begin{pmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{pmatrix}.$$

We can multiply an $m \times n$ matrix with an $n \times k$ matrix to get an $m \times k$ matrix, in a somewhat complicated way which we won't get into now but which makes sense in terms of vector spaces. Note that this multiplication is noncommutative: AB is not in general the same thing as BA, and if A is $m \times n$, B is $n \times k$, then although AB is defined (and an $m \times k$ matrix), BA is only defined if $k = m$, in which case it will be $n \times n$.

A special case of the above is the action of matrices on vectors: if A is an $m \times n$ matrix and $v = (x_1, x_2, \ldots, x_n)$ is a vector, we can think of v as a $n \times 1$ matrix

$$\begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix}$$

so that the multiplication Av makes sense, and is an $m \times 1$ matrix, i.e. a vector with m components.

Given a matrix A, we can consider its transpose A^T given by exchanging the rows and columns, i.e. sending a_{ij} to a_{ji}; for example, if

$$A = \begin{pmatrix} 1 & 0 & 3 \\ 2 & -1 & 1 \end{pmatrix}$$

then

$$A^\mathsf{T} = \begin{pmatrix} 1 & 2 \\ 0 & -1 \\ 3 & 1 \end{pmatrix}.$$

If $A = A^\mathsf{T}$ (possible only for square matrices, i.e. $n \times n$ for some n) then we say that A is symmetric.

The entries of matrices are often taken to be real or complex numbers, but in fact we can define matrices over any (commutative) ring R, such as $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$, $\mathbb{C}$, $\mathbb{Q}[x]$, etc. Let's write $M_{m\times n}(R)$ for the space of $m \times n$ matrices with entries in R, or $M_n(R)$ for short for the space of $n \times n$ matrices (which are the ones we'll see the most).

Any square matrix $A \in M_n(R)$ has a determinant $\det(A) \in R$, which is a certain polynomial in the entries of R. In the case $n = 1$, so $A = (a)$ is really just an element of R, $\det A = a$; in the case $n = 2$, if

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

then $\det A = ad - bc$; for higher n, there exist formulas but we won't bother with them. The key property of the determinant is that as a map $M_n(R) \to R$, it is multiplicative:

$\det(AB) = \det(A) \cdot \det(B)$. (Since multiplication in R is commutative, it follows that $\det(BA) = \det(B) \cdot \det(A) = \det(A) \cdot \det(B) = \det(AB)$.) This lets us restrict to for example matrices A whose determinant is invertible. In this case, the matrix itself is invertible, i.e. there exists another matrix A^{-1} with $AA^{-1} = A^{-1}A = I_n$, where

$$I_n = \begin{pmatrix} 1 & 0 & \cdots & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ \vdots & 0 & 1 & \ddots & \vdots \\ \vdots & \vdots & \ddots & \vdots & 0 \\ 0 & 0 & \cdots & 0 & 1 \end{pmatrix}$$

is the $n \times n$ matrix with 1's on the diagonal and all other entries 0, which has the property that $I_n A = A I_n = A$ for every $n \times n$ matrix A.

Write $\mathrm{GL}_n(R) \subset M_n(R)$ for the set of matrices with invertible determinant, or equivalently invertible matrices over R; this has the structure of a group. Inside it we have the subgroup $\mathrm{SL}_n(R) \subset \mathrm{GL}_n(R) \subset M_n(R)$ consisting of $n \times n$ matrices with determinant 1, which is always invertible.

For example, if $R = \mathbb{Z}$, the only invertible elements in $\mathbb{Z}$ are ± 1, so $\mathrm{GL}_n(\mathbb{Z})$ consists of $n \times n$ matrices with integer entries whose determinants are either 1 or -1; and $\mathrm{SL}_n(\mathbb{Z})$ is the subgroup whose determinants are 1. For most rings, there are other invertible elements and so in general $\mathrm{GL}_n(R)$ is much larger than $\mathrm{SL}_n(R)$, but they are still fairly comparable for many purposes.

2.3 Lagrange's theorem

Our starting point is Lagrange's theorem:

Theorem (Lagrange)**.** *Every positive integer can be written as the sum of four squares.*

We allow 0 and repetition, so for example $6 = 2^2 + 1 + 1 + 0$. Note that we don't require that these representations are unique, e.g. $9 = 3^2 + 0^2 + 0^2 + 0^2 = 2^2 + 2^2 + 1^2 + 0^2$.

To prove this theorem, we'll follow the strategy mentioned when we discussed the fundamental theorem of arithmetic: for a positive integer n, say $P(n)$ is the property of having a representation as a sum of four squares. We ultimately want to show that $P(n)$ is true for every positive integer. So our strategy has two steps: first, show that P is multiplicative, i.e. if $P(m)$ and $P(n)$ holds, then $P(mn)$ also holds; and second, show that $P(p)$ holds for every prime number p. Then by the fundamental theorem of arithmetic (and checking that $P(1)$ is true since $1 = 1^2 + 0^2 + 0^2 + 0^2$) we can conclude that $P(n)$ is always true.

The first step is easier. Suppose that $P(m)$ and $P(n)$ hold, so we can write $m = x_1^2 + x_2^2 + x_3^2 + x_4^2$ and $n = y_1^2 + y_2^2 + y_3^2 + y_4^2$. Then by multiplying out a lot of terms one can see that

$$\begin{aligned} mn &= (x_1^2 + x_2^2 + x_3^2 + x_4^2)(y_1^2 + y_2^2 + y_3^2 + y_4^2) \\ &= (x_1y_1 + x_2y_2 + x_3y_3 + x_4y_4)^2 + (x_1y_2 - x_2y_1 - x_3y_4 + x_4y_3)^2 \\ &\qquad + (x_1y_3 - x_3y_1 + x_2y_4 - x_4y_2)^2 + (x_1y_4 - x_4y_1 - x_2y_3 + x_3y_2)^2 \end{aligned}$$

is also a sum of four squares, i.e. $P(mn)$ also holds. (This polynomial identity is called Euler's four-square identity.) So it remains just to show that $P(p)$ is true whenever p is prime, i.e. every prime number can be written as the sum of four squares.

All primes except for 2 are odd, which we'll want to use, so let's check $p = 2$ separately: $2 = 1^2 + 1^2 + 0^2 + 0^2$, so $P(2)$ is true. Therefore we just need to show that $P(p)$ holds for *odd* primes p.

To do so, we'll again split this into two steps. First, we'll show that there is some integer n such that $P(np)$ holds. Once this is done, we can ask what the smallest such integer n is; and we'll show that in fact the answer is $n = 1$, so that $P(p)$ holds.

For step 1, we need some modular arithmetic. In particular, since p is prime $\mathbb{Z}/p\mathbb{Z}$ is a field: the only element which does not have a multiplicative inverse is 0. (This is not true for $\mathbb{Z}/n\mathbb{Z}$ in general; it's a special property of primes, which is why our broad strategy is a good one.) Therefore things behave well: for example, the squaring function $x \mapsto x^2$ is exactly 2-to-1 except on 0, i.e. for every $x \in \mathbb{Z}/p\mathbb{Z}$, the only elements $y \in \mathbb{Z}/p\mathbb{Z}$ with $y^2 = x^2$ are $y = x$ and $y = -x$ (since $(-x)^2 = x^2$). (If $x = 0$, these are the same.) (If p was not an odd prime, this would not necessarily be true: for example in $\mathbb{Z}/6\mathbb{Z}$ there is only one element satisfying $x^2 \equiv 3 \pmod 6$, namely $x = 3$.)

This means that of the p elements of $\mathbb{Z}/p\mathbb{Z}$, one is zero; and exactly half of the remaining $p - 1$ are squares (remember p is odd). For example, if $p = 5$ then we have 0; $1 = 1^2$ and $4 = 2^2$ are squares; and so 2 and 3 cannot be squares, since $0^2 = 0$, $1^2 = 1$, $2^2 = 4$, $3^2 = 9 \equiv 4$, and $4^2 = 16 \equiv 1$ are all accounted for. The $\frac{p-1}{2}$ elements which are squares are called quadratic residues, and the $\frac{p-1}{2}$ non-squares are nonresidues.

Now, suppose we take just the elements $0, 1, \ldots, \frac{p-1}{2}$ and square them. The resulting residues are all distinct, since if $x^2 = y^2$ with $x \neq y$ we must have $y = -x \equiv p - x$ and so for $0 \leq x \leq \frac{p-1}{2}$, $p - x$ is not in this range; so we've identified exactly $1 + \frac{p-1}{2} = \frac{p+1}{2}$ elements of the form x^2.

Similarly, for $0 \leq y \leq \frac{p-1}{2}$ the elements $-y^2 - 1$ are all distinct, so we have another $\frac{p+1}{2}$ elements of the form $-y^2 - 1$. Combining these with the elements above, if they were all distinct there would be $\frac{p+1}{2} + \frac{p+1}{2} = p + 1$, but there are only p elements of $\mathbb{Z}/p\mathbb{Z}$, so there must be some overlap: in other words, there must be some element c such that $c = x^2$ and $c = -y^2 - 1$, i.e. $x^2 = -y^2 - 1$, or equivalently $x^2 + y^2 + 1 \equiv 0 \pmod p$, with $0 \leq x, y \leq \frac{p-1}{2}$.

This is just the statement that there exist $0 \leq x \leq \frac{p-1}{2}$, $0 \leq y \leq \frac{p-1}{2}$ and an integer n such that $x^2 + y^2 + 1 = x^2 + y^2 + 1^2 + 0^2 = np$. Thus in particular $P(np)$ is true for some n; so we've accomplished step 1. We can also derive a bound on n: since $x, y \leq \frac{p-1}{2}$,

$$np = x^2 + y^2 + 1 \leq 2\left(\frac{p-1}{2}\right)^2 + 1 = \frac{1}{2}p^2 - p + \frac{3}{2} < p^2$$

(as $p > 1$) and so $n < p$.

For step 2, we can consider the set S of all positive integers n such that $P(np)$ is true; we've shown that S is nonempty, and that it contains an element smaller than p. Now let n be the smallest element of S, so we want to show that $n = 1$. We'll prove this by contradiction, so we assume that $n \geq 2$. The idea is to produce another $m \in S$ which is smaller than n, thus contradicting the minimality of n.

Since $n \in S$, there exist integers x_1, x_2, x_3, x_4 such that $np = x_1^2 + x_2^2 + x_3^2 + x_4^2$. Let y_i be representatives of x_i modulo n with $-\frac{n}{2} < y_i \leq \frac{n}{2}$, so we have

$$y_1^2 + y_2^2 + y_3^2 + y_4^2 \equiv x_1^2 + x_2^2 + x_3^2 + x_4^2 = np \equiv 0 \pmod{n},$$

so there exists some integer m such that

$$y_1^2 + y_2^2 + y_3^2 + y_4^2 = mn.$$

At this point we split into cases. Since the left-hand side is nonnegative, so is the right-hand side, and since $n \geq 1$ it follows that $m \geq 0$. We want to show that in fact $m \geq 1$, $m < n$, and $P(mp)$ is true, i.e. $m \in S$ is smaller than n, contradicting its minimality.

If $m = 0$, then $y_1^2 + y_2^2 + y_3^2 + y_4^2 = 0$, so the only possibility is $y_1 = y_2 = y_3 = y_4 = 0$. In this case, $x_i \equiv y_i = 0 \pmod{n}$, so each x_i is divisible by n, so x_i^2 is divisible by n^2, so $np = x_1^2 + x_2^2 + x_3^2 + x_4^2$ is also divisible by n^2; but this would imply that p is divisible by n, which is impossible since p is prime and $2 \leq n < p$. Therefore we can assume $m \geq 1$.

On the other hand, if $m \geq n$ then

$$y_1^2 + y_2^2 + y_3^2 + y_4^2 = mn \geq n^2.$$

By assumption $|y_i| \leq \frac{n}{2}$, so the left-hand side is at most $4 \cdot \left(\frac{n}{2}\right)^2 = n^2$, with equality possible only if each y_i is actually equal to $\frac{n}{2}$ (which requires n to be even). In that case, since $x_i \equiv y_i = \frac{n}{2} \pmod{n}$, $x_i^2 \equiv y_i^2 = \frac{n^2}{4} \pmod{n^2}$, so

$$np = x_1^2 + x_2^2 + x_3^2 + x_4^2 \equiv 4 \cdot \frac{n^2}{4} = n^2 \equiv 0 \pmod{n^2},$$

i.e. np is divisible by n^2, which again would imply that p is divisible by n, which is impossible. Thus we can assume $1 \leq m < n$.

We are now ready to conclude: we have $np = x_1^2 + x_2^2 + x_3^2 + x_4^2$ and $mn = y_1^2 + y_2^2 + y_3^2 + y_4^2$, so by Euler's four-square identity again

$$\begin{aligned} np \cdot mn &= mn^2 p \\ &= (x_1^2 + x_2^2 + x_3^2 + x_4^2)(y_1^2 + y_2^2 + y_3^2 + y_4^2) \\ &= (x_1 y_1 + x_2 y_2 + x_3 y_3 + x_4 y_4)^2 + (x_1 y_2 - x_2 y_1 - x_3 y_4 + x_4 y_3)^2 \\ &\qquad + (x_1 y_3 - x_3 y_1 + x_2 y_4 - x_4 y_2)^2 + (x_1 y_4 - x_4 y_1 - x_2 y_3 + x_3 y_2)^2. \end{aligned}$$

Since $x_i \equiv y_i \pmod{n}$, modulo n the first term (before squaring) is $x_1^2 + x_2^2 + x_3^2 + x_4^2 = (np)^2 \equiv 0 \pmod{n}$ while the other terms all vanish more straightforwardly, so every term (pre-squaring) is divisible by n. This means that there exist integers z_1, z_2, z_3, z_4 such that

$$mn^2 p = (nz_1)^2 + (nz_2)^2 + (nz_3)^2 + (nz_4)^2 = n^2(z_1^2 + z_2^2 + z_3^2 + z_4^2),$$

so dividing by n^2 gives

$$mp = z_1^2 + z_2^2 + z_3^2 + z_4^2.$$

In other words, $P(mp)$ is true. Since we showed above that $m \geq 1$, it follows that $m \in S$; and since we showed $m < n$, this contradicts the minimality of n. Thus our initial assumption that $n > 1$ must have been wrong, and we conclude $n = 1$ is the smallest element of S; in particular, $P(p)$ holds for every odd prime p, and by the above it follows that $P(n)$ holds for every positive integer n.