MATH 371 – Class notes/outline – September 3, 2013

Integers

The set of integers $\{\ldots, -3, -2, -1, 0, 1, 2, 3, \ldots\}$ is denoted by $\mathbb{Z}$. The natural numbers $\{1, 2, 3, \ldots\}$ is denoted by $\mathbb{N}$. We'll assume the definitions of addition and multiplication in $\mathbb{N}$, also subtraction in $\mathbb{Z}$.

$\mathbb{N}$ and $\mathbb{Z}$ are ordered by the relation $x \geqslant y$ means $x - y \in \mathbb{N}$. They are totally (linearly) ordered in the sense that for each x, y in $\mathbb{N}$ or $\mathbb{Z}$, we have that either $x \geqslant y$ or $y \geqslant x$ (or both, in which case $x = y$) is true.

Additionally, $\mathbb{N}$ is *well-ordered*, which means every non-empty subset of $\mathbb{N}$ has a smallest (first, initial) element. This statement is equivalent to mathematical induction (you prove that the smallest number for which the statement in question is false cannot exist, so the set of numbers for which the statement is false must be empty).

Division with remainder: For every $x \in \mathbb{Z}$ and $d \in \mathbb{N}$ there is a *unique* $r \in \mathbb{Z}$ and a $q \in \mathbb{Z}$ such that $x = qd + r$ with $0 \leqslant r < d$.

(proof of uniqueness: suppose $x = q_1 d + r_1 = q_2 d + r_2$ with $0 \leqslant r_1 < r_2 < d$. Subtract and get $r_2 - r_1 = (q_1 - q_2)d$, but $0 \leqslant r_2 - r_1 < d$, and so $r_2 - r_1 = 0$. Existence: r is the smallest element of $\{x - qd \,|\, q \in \mathbb{Z}\} \cap \mathbb{N}$.)

Divisors: $c \mid a$ means the remainder on division of a by c is zero — say "c divides a". Write $[x]_d$ (or just $[x]$ when d is understood) to be the remainder when x is divided by d, so $[x]_d \in \{0, 1, \ldots, d-1\}$.

Congruences: Write $a \equiv b \pmod{c}$ to mean that $c \mid b - a$. Note that $a \equiv [a]_c \pmod{c}$, and $a \equiv b \pmod{c}$ if and only if $[a]_c = [b]_c$.

Proposition. If $x_1 \equiv x_2 \pmod{c}$ and $y_1 \equiv y_2 \pmod{c}$ then

$$x_1 + x_2 \equiv y_1 + y_2 \pmod{c} \quad \text{and} \quad x_1 y_1 \equiv x_2 y_2 \pmod{c}$$

Therefore, $[xy] = [[x][y]]$, which comes in handy for calculations. For example, repeated squaring to do high powers of (big) numbers:

Example: What is $[12^{11}]_{21}$? Note that $11 = 2^3 + 2 + 1$ so $[12^{11}] = [[12^{2^3}][12^2][12]]$ (all mod 21). By squaring, compute (mod 21):

$$\begin{aligned}
[12^1] &= 12 \\
[12^2] &= [144] = 18 = [-3] \\
[12^4] &= [18^2] = [(-3)^2] = 9 \\
[12^8] &= [9^2] = [81] = [-3] = 18.
\end{aligned}$$

Therefore $[12^{11}] = [18 \cdot 18 \cdot 12] = [9 \cdot 12] = [9 \cdot (-9)] = [-81] = 3$. Note that $12^1 1 = 743,0080,370,688$ so the long division way would have taken a while.

Greatest common divisor: Let $n \in \mathbb{N}$ (or $\mathbb{Z}$). Define $\mathrm{div}(n) = \{d \in \mathbb{N} \mid d \mid n\}$ (set of (positive) divisors of n.

$$\mathrm{div}(18) = \{1, 2, 3, 6, 9, 18\}$$
$$\mathrm{div}(24) = \{1, 2, 3, 4, 6, 8, 12, 24\}$$

so $\mathrm{div}(18) \cap \mathrm{div}(24) = \{1, 2, 3, 6\} = \mathrm{div}(6)$. Not a coincidence: $\gcd(18, 24) = 6$.

Fact: (Euclid) Given $m, n \in \mathbb{Z}$, there is a unique $d \in \mathbb{N}$ such that $\mathrm{div}(m) \cap \mathrm{div}(n) = \mathrm{div}(d)$.

(proof: Assume that $m, n \in \mathbb{N} \cup \{0\}$, get the result for all integers by fiddling with signs. Do induction on $\min(m, n)$.)

Essential fact about gcd's: For any $m, n \in \mathbb{Z}$ there exist $\lambda, \mu \in \mathbb{Z}$ such that

$$\lambda m + \mu n = \gcd(m, n).$$

In fact you can characterize $\gcd(m, n)$ as the smallest positive number in $\{\lambda m + \mu n \mid \lambda, \mu \in \mathbb{Z}\}$.

Additional fact: $\gcd(m, n) = 1$ if and only if there are λ and μ such that $\lambda m + \mu n = 1$. (In this case say m and n are *relatively prime*.

Can find m and n using the Euclidean algorithm (and keeping careful track along the way). The idea is to replace the larger of m and n by the remainder on dividing by the smaller. Keep doing this until you get to zero, then then next-to-last remainder is the gcd. Algorithmically express as follows. Assume $m > n$, and set $r_{-1} = m$, $r_0 = n$. Also set $\lambda_{-1} = 1$, $\lambda_0 = 0$, $\mu_{-1} = 0$ and $\mu_0 = 1$. Then keep increasing i until $r_i = 0$, where $r_i = r_{i-2} - q_i r_{i-1}$ (this is division with remainder), also using the same q_i already defined, set $\lambda_i = \lambda_{i-2} - q_i \lambda_{i-1}$ and $\mu_i = \mu_{i-2} - q_i \mu_{i-1}$. Check that at each stage $\lambda_i m + \mu_i n = r_i$, so at the next-to-last stage you'll have $\lambda_i m + \mu_i n = \gcd(m, n)$.

Here's an example for the computation of $\gcd(312, 81)$:

i	-1	0	1	2	3	4	5
r_i	312	81	69	12	9	3	0
q_i	$-$	$-$	3	1	5	1	3
λ_i	1	0	1	-1	6	-7	$-$
μ_i	0	1	-3	4	-23	27	$-$

The result is the next-to-last entry in the r_i row, namely, $\gcd(312, 81) = 3$. We also get that

$$-7 \cdot 312 + 27 \cdot 81 = 3 = \gcd(312, 27).$$

A few important corollaries of this characterization of the gcd:

1. If $a \mid bc$ and $\gcd(a, b) = 1$ then $a \mid c$.
2. If $\gcd(a, b) = 1$ and $a \mid c$ and $b \mid c$ then $ab \mid c$.
3. If $\gcd(a, b) = 1$ and $\gcd(a, c) = 1$ then $\gcd(a, bc) = 1$.

First, think about why these should be true, and perhaps try and prove them yourself, before looking at the following proofs:

Proof of 1: Since $\gcd(a,b)=1$, we have $\lambda a+\mu b=1$ for some numbers λ,μ, therefore $\lambda ac+\mu bc=c$. Since a divides both terms of the sum on the right, it divides the sum, therefore $a \mid c$.

Proof of 2: We are given $\lambda a+\mu b=1$ and $c=ax=by$ for some numbers λ,μ,x,y. But then $c=\lambda ac+\mu bc=\lambda aby+\mu bax=ab(\lambda y+\mu x)$ so $ab \mid c$.

Proof of 3: We know $\lambda a+\mu b=1$ and $\rho a+\sigma c=1$ for some numbers λ,μ,ρ,σ. Therefore $(\lambda a+\mu b)(\rho a+\sigma c)=1$, i.e., $\lambda\rho a^2+\lambda\sigma ac+\mu\rho ab+\mu\sigma bc=1$, i.e., $(\lambda\rho a+\lambda\sigma c+\mu\rho b)a+\mu\sigma bc=1$. So we've found an integer linear combination of a and bc that sums to 1, therefore $\gcd(a,bc)=1$.

For the moment, define $\mathbb{Z}/n=\{x\in\mathbb{N}\,|\,0\leqslant x<n\}=\{0,1,2,\ldots,n-1\}$ simply as a set. You probably remember giving it the structure of a cyclic group (and in fact it has a ring structure as well), thanks to the proposition on page 1. Now suppose n factors as $n=n_1n_2\cdots n_k$, and define the "remainder map"

$$r\colon \mathbb{Z}/n\to\mathbb{Z}/n_1\times\mathbb{Z}/n_2\times\cdots\times\mathbb{Z}/n_k$$

to be the map that sends $x\in\mathbb{Z}/n$ to the k-tuple of remainders $([x]_{n_1},[x]_{n_2},\ldots,[x]_{n_k})$.

Lemma: If $\gcd(n_i,n_j)=1$ for all $i\neq j$, then the map r defined above is a bijection (i.e., it is one-to-one and onto).

To prove this, we'll repeatedly use corollaries 2 and 3 above. To show the map is injective (one-to-one), suppose $r(x)=r(y)$. Then $n_1 \mid x-y$, $n_2 \mid x-y$, $\ldots n_k \mid x-y$ since x and y leave the same remainders when divided by any of the n_i. From the first two of these, corollary 2 tells us that $n_1n_2 \mid x-y$. But then corollary 3 tells us that $\gcd(n_1n_2,n_3)=1$, so we can use corollary 2 to get that $n_1n_2n_3 \mid x-y$. Proceed inductively and finally arrive at $n_1n_2\cdots n_k \mid x-y$, i.e., $n \mid x-y$. But both x and y are non-negative and less than n, so this can't happen unless $x-y=0$, i.e., $x=y$. So the map is injective. And since it's a map of finite sets of the same size, it's surjective as well.

This lemma leads naturally to:

The Chinese Remainder Theorem: (Sun-Tzu [c. 400 A.D.]) Let $n=n_1n_2\cdots n_k$ with $\gcd(n_i,n_j)=1$ for $i\neq j$. Then, given any integers $a_1,a_2\ldots,a_k$ the system of congruences

$$\begin{aligned}
x &\equiv a_1 \pmod{n_1}\\
x &\equiv a_2 \pmod{n_2}\\
&\vdots \quad\ \vdots \quad\ \vdots\\
x &\equiv a_k \pmod{n_k}
\end{aligned}$$

has solutions x. The system has exactly one solution x in $\mathbb{Z}/n$, and any two solutions are congruent mod n. Also if x is one solution then $x+kn$ is another solution for any $k\in\mathbb{Z}$.

The uniqueness etc statements at the end are obvious from the Lemma, so we'll concentrate on existence of a solution. For each i between 1 and k, set $m_i=n/n_i$. Since m_i is the product of all the n_j's except for n_i, and the n_j's are relatively prime, repeated application of corollary 3 above shows that $\gcd(m_i,n_i)=1$, therefore there are numbers λ_i and μ_i so that $\lambda_in_i+\mu_im_i=1$. Let $\beta_i=\mu_im_i=\mu_in/n_i$. Then we have $\lambda_in_ia_i+\beta_ia_i=a_i$, which shows that $\beta_ia_i\equiv a_i \pmod{n_i}$ and since $m_i\equiv 0 \pmod{n_j}$ for $j\neq i$, we have $\beta_ia_i\equiv 0 \pmod{n_j}$ for $j\neq i$. From these observations it is easy to see that $x=\beta_1a_1+\beta_2a_2+\cdots+\beta_ka_k$ is a solution to the problem.

As an example, let's find the number x less than 30 that leaves remainder 1 when divided by 2, leaves remainder 2 when divided by 3 and leaves remainder 4 when divided by 5. So we need to solve $x \equiv 1 \pmod 2$, $x \equiv 2 \pmod 3$, $x \equiv 4 \pmod 5$ (note that $2 \cdot 3 \cdot 5 = 30$). The Euclidean algorithm tells us that $-7 \cdot 2 + 1 \cdot 15 = 1$, so we take $\beta_1 = 1 \cdot 15 = 15$. Then $-3 \cdot 3 + 1 \cdot 10 = 1$ so we take $\beta_2 = 1 \cdot 10 = 10$. Finally $-1 \cdot 5 + 1 \cdot 6 = 1$ so we take $\beta_3 = 1 \cdot 6 = 6$. So one number that solves the congruences is $15 \cdot 1 + 10 \cdot 2 + 6 \cdot 4 = 59$, which is congruent to 29 mod 30. And 29 is the unique solution of our problem.

Euler's φ-function: Write $(\mathbb{Z}/n)^*$ for the subset of $\mathbb{Z}/n$ consisting of numbers $x \in \{0, 1, \ldots, n-1\}$ for which $\gcd(x, n) = 1$. By corollary 3 above, this set is closed under multiplication (mod n), since $\gcd(1, n) = 1$ we have $1 \in (\mathbb{Z}/n)^*$, and since in the expression $\lambda n + \mu x = 1$ it is possible to choose μ such that $0 \leqslant \mu < n$ (this is homework problem 6(a)), we get that μ also satisfies $\gcd(\mu, n) = 1$ and $\mu x \equiv 1 \pmod n$, so multiplicative inverses exist in $(\mathbb{Z}/n)^*$. Hence $(\mathbb{Z}/n)^*$ is a group with the operation of multiplication.

What is the order of the group $(\mathbb{Z}/n)^*$? This doesn't seem so obvious. It is obviously a function of n, and Euler decided to give this function the name φ, so $\varphi(n)$ is now called Euler's φ-function. Here is a table for some small values of n:

n	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
$\varphi(n)$	0	1	1	2	2	4	2	6	4	6	4	10	4	12	6	8	8	16	6

It looks like $\varphi(n)$ is even for all $n > 2$ — deciding whether this is true will be part of next week's homework. On the other hand, it might be hard to discern any other patterns here. One important property of φ is that it is one of a number of remarkable number theoretic functions called *multiplicative functions*, which is the content of the following:

Proposition: If m and n are relatively prime, then $\varphi(mn) = \varphi(m)\varphi(n)$.

The hypothesis of m and n being relatively prime is essential, for instance note $4 = \varphi(3)\varphi(6) \neq \varphi(18) = 6$. To prove the proposition, we'll use the remainder map r defined above, in this case

$$r \colon \mathbb{Z}/mn \to \mathbb{Z}/m \times \mathbb{Z}/n,$$

and the fact proved in the lemma that r is a bijection.If we can show that the restriction of $r(x) \in (\mathbb{Z}/m)^* \times (\mathbb{Z}/n)^*$ if and only if $x \in (\mathbb{Z}/mn)^*$ we'll be done, since the restriction of a bijection is still a bijection (so these two sets must therefore be the same size). In other words we have to show that $\gcd(x, m) = \gcd(x, n) = 1$ if and only if $\gcd(x, mn) = 1$. The "only if" part is corollary 3 from the bottom of page 2. And if $\gcd(x, mn) = 1$ then there are λ and μ such that $\lambda x + \mu mn = 1$, in other words $\lambda x + (\mu n)m = 1$ (so $\gcd(x, m) = 1$) and $\lambda x + (\mu m)n = 1$ (so $\gcd(x, n) = 1$) and we are done. It's a bit subtle, but note that the place where we used $\gcd(m, n) = 1$ is that r is a bijection.

Euler's theorem: Given $a \in \mathbb{Z}$ and $n \in \mathbb{N}$ such that $\gcd(a, n) = 1$, we have $a^{\varphi(n)} \equiv 1 \pmod n$.

This is simply the fact that the order of an element of a group divides the order of the group (in this case the group is $(\mathbb{Z}/n)^*$.

Prime numbers

A number $p \in \mathbb{N}$ is prime if it is not expressible as the product of natural numbers less than itself. In other words, p is prime if and only if $\mathrm{div}(p) = \{1, p\}$ and $1 \neq p$. Because of this, we have

that if p is prime, then for any natural number x, either $p|x$ or else $\gcd(p, x) = 1$. Also, thanks to corollary 1 on page 2, we have:

Basic fact about prime numbers: If p is prime and $p \mid ab$, then either $p \mid a$ or $p \mid b$ (or both).

Prime numbers are the multiplicative building blocks of $\mathbb{N}$ and $\mathbb{Z}$. In particular

Unique factorization: Every natural number can be expressed as a product of prime numbers in a unique way (up to the order of the factors).

Existence is by induction; show that the smallest natural number that can't be factored must itself be prime. Uniqueness is by cancellation: if x has two factorizations $x = p_1 p_2 \cdots p_k = q_1 q_2 \cdots q_\ell$, then cancel all the common factors from both sides. If there's anything left, use the basic fact above to show that anything that remains on one side must in fact also be on the other, contradicting that all common factors were cancelled.

Euclid's theorem: There are infinitely many prime numbers.

Calcuating $\varphi(n)$: Use prime factorization to calculate $\varphi(n)$. Since we can factor $n = p_1^{e_1} p_2^{e_2} \cdots p_k^{e_k}$, from the proposition above we only need a formula for $\varphi(p^e)$. But the numbers less than but not relatively prime to p^e are simply the multiples of p, namely $p, 2p, 3p, (p^{e-1})p$, so there are p^{e-1} of them. Therefore $\varphi(p^e) = p^e - p^{e-1}$ if p is prime. From this we get:

If $n = p_1^{e_1} p_2^{e_2} \cdots p_k^{e_k}$, then

$$\varphi(n) = n\left(1 - \frac{1}{p_1}\right)\left(1 - \frac{1}{p_2}\right)\cdots\left(1 - \frac{1}{p_k}\right).$$

The RSA cryptosystem

The RSA system is a "public-key" system. This means that the method of *encrypting* messages is made public, but the method of *decrypting* encoded messages is kept secret. This is possible because the system is based on the difficulty of factoring large (on the order of 10^{200}) numbers. The method starts with two large (on the order of 10^{100}) prime numbers p and q (which are kept secret), and their product $N = pq$ (which is made public).

The "message" to be sent is cast in the form of a number X (this can be done in many ways and can be quite simplistic, because the power of the system is in the large primes p and q — so a system like A= 01, B= 02, etc is perfectly fine), where $0 \leqslant X < N$. A very long message can be broken up and represented by a sequence of numbers $X_1, X_2, \ldots$.

To encrypt the message, a special exponent e is chosen (and published, so anyone can encrypt a message), and the encrypted message is $[X^e]_N$ (we know a reasonably efficient way to compute this, using repeated squaring). The key point in the workings of the RSA system is that there is a unique number d with the property that $[(X^e)^d]_N = [X^{ed}]_N = X$, so an encrypted message can be deciphered by raising it to the dth power mod N. And it's important to keep the number d a secret, so that only those who know it can decode encrypted messages. As it turns out, figuring out d, given N and e is an impracticable task, equivalent to factoring N into its constituent primes p and q. So as long as p, q and d are kept secret, only the intended recipient will be able to decode messages, even though anyone can send them.

So the challenge of the RSA system is to come up with two big primes p and q to multiply together (more on this below), and to come up with appropriate exponents e and d, so that $[X^{ed}]_N = X$. To do this, we'll be using Euler's theorem, together with our method for computing the φ-function.

In particular, since $N = pq$, we have that $\varphi(N) = \varphi(p)\varphi(q) = (p-1)(q-1)$.

Proposition: For any $k \in \mathbb{N}$, we have $X^{k(p-1)(q-1)+1} \equiv X \pmod{N}$ for all $X \in \mathbb{Z}$.

In other words, $N \mid (X^{k(p-1)(q-1)+1} - X$. By good old corollary 2, it's enough to show this with N replaced by p and q individually, since they're (relatively) prime. The proof is the same for p and q so we just do it for p. Since p is prime, we either have $p \mid X$ or $\gcd(p, X) = 1$. In the first case, $X \equiv 0 \pmod{p}$ and so any power of X is congruent to zero mod p, so the result is trivial. If $\gcd(p, X) = 1$, note that Euler's theorem tells us that $X^{\varphi(p)} = X^{p-1} \equiv 1 \pmod{p}$. Raise both sides to the $k(q-1)$ power (1 raised to any power is still 1 mod p), and get $X^{k(p-1)(q-1)} \equiv 1 \pmod{p}$. Multiply both sides of this by X to get the result in this case, and we are done.

Now, for the encryption exponent e, choose any number such that $0 \leqslant e < N$ and $\gcd(e, \varphi(N)) = \gcd(e, (p-1)(q-1)) = 1$. There are lots of them! By problem 6(a) of the first homework, we can thus choose λ and μ with $0 < \mu < (p-1)(q-1)$ so that $\lambda(p-1)(q-1) + \mu e = 1$. Claim the decryption exponent d is equal to μ. To see this, put $k = -\lambda$ and see that $de = \mu e = k(p-1)(q-1) + 1$, which is just what is needed to apply the proposition above.

Finding large primes

To implement the RSA system, we need a method for producing those large, 100-digit prime numbers to multiply together. Large prime numbers are rare, but not exceedingly so among numbers that size, so the standard method for finding them involves testing a bunch of large odd numbers (a few hundred, maybe) until you find prime ones.

So the question is, how do you efficiently test whether a large number is prime? In the 1980s, Miller and Rabin came up with a surprising (even revolutionary) approach. They realized that even though it seems quite hard (expensive) to prove whether a number is prime, there is a not-so-expensive way to render determinations of primality that are right almost all of the time (in that their error rate is less than 1 in 10^{50}, say. And their method uses little more theory than what we have already developed.

We'll start with an easy corollary of Euler's theorem:

Fermat's little theorem: If $p \in \mathbb{N}$ is a prime number and $a \in \mathbb{Z}$ is not a multiple of p, then $a^{p-1} \equiv 1 \pmod{p}$.

Unfortunately, Fermat's little theorem is not an if and only if statement; if it were then it would give an easy test for primality. As it is, we can sometimes use it to tell whether a number is composite. For example, if 10 were a prime number then $2^9 = 512$ would be congruent to 1 mod 10, but $2^9 \equiv 2 \pmod{10}$. Therefore 10 is composite. But it is not necessarily the case that p is prime if $2^{p-1} \equiv 1 \pmod{p}$, for example $2^{340} \equiv 1 \pmod{341}$, but $341 = 11 \cdot 31$.

We'll say that n is a *base-a pseudoprime* (or n is a *pseudoprime relative to a*) if $a^{n-1} \equiv 1 \pmod{n}$. So we have that 341 is a base-2 pseudoprime. That's the official lingo. But we'll use more intuitive

language and say that if n is a composite number, then a is a *Fermat witness* to the compositeness of a if $a^{n-1} \not\equiv 1 \pmod{n}$ (since this fact is a proof that n is composite), and that n is a *Fermat liar* if $a^{n-1} \equiv 1 \pmod{n}$ (even though n is not prime).

The fact that 2 is a Fermat liar for some n prevents the following test from being very reliable:

> Given n, determine whether $2^{n-1} \equiv 1 \pmod{n}$. If not, then declare n to be composite (with certainty); if so, then declare n to be prime (but this might not be so).

The Miller-Rabin primality test works somewhat like this, but reduces the doubt in the declaration that n is prime to a microscopic amount. This is achieved by two improvements: First, the method checks several bases (not just 2) to see if they are Fermat witnesses to the compositeness of n. Second, it uses one additional test to reduce the probability that the method will be fooled even further.

The second part of the Miller-Rabin test checks to see whether there is a "fake square root of 1", in other words, a number $x \in \mathbb{Z}/N$ other than 1 or $N-1$ (which is congruent to $-1 \bmod N$) such that $x^2 \equiv 1 \pmod{N}$. Such a number cannot exist if N is prime, since $x^2 - 1 \equiv 0 \pmod{N}$ means that $N \mid (x-1)(x+1)$. But if N is prime then this implies that either $N \mid x+1$ or $N \mid x-1$, i.e., $x \equiv \pm 1 \pmod{N}$.

For example, if $N = 8$ we have four solutions to $x^2 - 1 \equiv 0 \pmod{8}$, namely 1, 3, 5 and 7. This shows that 8 is not a prime number.

So now we'll say that a number b is a *Miller-Rabin witness* to the compositeness of N if b is either a Fermat witness for N or else if the following holds: Express the even number $N-1$ as $2^t u$, where u is odd and $t \geqslant 1$. Set $x_0 = b^u$, and then for $i = 1, 2, \ldots, t$ set $x_i = x_{i-1}^2 \pmod{N}$ (so that at the end of this process, $x_t = b^{N-1}$, which is needed for the Fermat test). If for any i we have $x_i \equiv 1 \pmod{N}$ but $x_{i-1} \not\equiv \pm 1 \pmod{N}$, then we've found a fake square root of 1 mod N (namely $b^{2^{i-1}u}$), and so b is a Miller-Rabin witness.

If N is composite but b is not a Miller-Rabin witness then b is called a *strong pseudoprime* relative to N — but we'll just call b a *Miller-Rabin liar* for N.

The Miller-Rabin algorithm for primality testing works like this: Choose several (maybe 100) values of b from $1 \leqslant b \leqslant N-1$ *randomly*. If any of the chosen values of b is a Miller-Rabin witness then N is definitely composite. If none of the values of b are MR-witnesses, then N is almost surely prime.

How certain is "almost surely" here? To answer this we need:

Theorem: For any odd composite number N, the number of witnesses to the compositeness of N is at least $\frac{1}{2}(N-1)$.

First, we note that any non-witness b must be relatively prime to N, since a non-witness would satisfy $b^{N-1} \equiv 1 \pmod{N}$, so there is a solution to $bx \equiv 1 \pmod{N}$, namely $x = b^{n-2}$. Therefore $N \mid bx - 1$, or $bx - 1 = qN$, i.e., $bx - qN = 1$ which means $\gcd(b, N) = 1$.

Thus, all non-witnesses are contained in $(\mathbb{Z}/N)^*$. We'll show that they are contained in a *proper*

subgroup $B \subset (\mathbb{Z}/N)^*$ — since the order of a subgroup divides the order of the group, B can have at most half the number of elements $(\mathbb{Z}/N)^*$ does, which will show that the number of non-witnesses is at most $\frac{1}{2}(n-1)$.

There are two cases to consider. The first case, which applies most of the time, is pretty easy to handle. It is the case where there is at least one Fermat witness x to the compositeness of N within $(\mathbb{Z}/N)^*$, i.e., $x^{N-1} \not\equiv 1 \pmod{N}$. In this case, we'll take B to be the set

$$B = \{b \in (\mathbb{Z}/N)^* \mid b^{N-1} \equiv 1 \pmod{N}\}.$$

We know that $B \neq \varnothing$ since $\pm 1 \in B$, and B is clearly closed under multiplication, so B is a subgroup of $(\mathbb{Z}/N)^*$. Every non-witness belongs to B by definition, but $x \notin B$, therefore B is a proper subgroup of $(\mathbb{Z}/N)^*$, so $|B| \leqslant \frac{1}{2}|(\mathbb{Z}/N)^*| < \frac{1}{2}(N-1)$ and so the number of witnesses must be greater than $N - \frac{1}{2}(n-1) = \frac{1}{2}(n+1)$.

The second case (rare but harder) is when every $x \in (\mathbb{Z}/N)^*$ is a Fermat non-witness, i.e., for every x such that $1 \leqslant x < N$ and $\gcd(x, N) = 1$ we have $x^{N-1} \equiv 1 \pmod{N}$. When this occurs, N is called a *Carmichael number*. We're going to need some preliminary results to deal with this case. The first one shows that Carmichael numbers are at least somewhat unusual.

Lemma: If N is a Carmichael number, then N has at least three distinct prime factors and N is not divisible by the square of any prime.

We'll deal with the second part first. So suppose that $p^2 \mid N$ for some prime p, so we can write $N = p^k x$ with $k > 1$ and $p \nmid x$. Since clearly $\gcd(p^k, x) = 1$, we can use the Chinese Remainder Theorem to find a number z such that $1 \leqslant z < N$, $z \equiv p + 1 \pmod{p^k}$, and $z \equiv 1 \pmod{x}$. Since $\gcd(z, p^k) = 1$ (which implies $\gcd(z, p) = 1$) and $\gcd(z, x) = 1$, we know that $\gcd(z, N) = 1$ by good old corollary 3.

Now, we'll show that z is a Fermat witness for N. To do this, we'll show that $z^{N-1} \not\equiv 1 \pmod{p^2}$ (and so it cannot be the case that $z^{N-1} \equiv 1 \pmod{N}$, since if $p^2 \nmid z^{N-1} - 1$ then N, which has p^2 as a divisor, can't divide $z^{N-1} - 1$ either).

To show that $z^{N-1} \not\equiv 1 \pmod{p^2}$ we begin with the fact that $(p+1)^p \equiv 1 \pmod{p^2}$ (see this by expanding the left side out, the only tricky term is the linear one), which implies (recall $z \equiv p + 1 \pmod{p^2}$) that

$$z^N \equiv (p+1)^N \equiv (p+1)^{p(p^{k-1}x)} \equiv 1^{p^{k-1}x} \equiv 1 \pmod{p^2}.$$

Thus $z^N \not\equiv z \pmod{p^2}$ which implies that $z^{N-1} \not\equiv 1 \pmod{p^2}$ and so $p^2 \nmid N$ for any prime p.

Now we have to show that a Carmichael number must have at least three distinct prime factors. So we'll show that a product of two distinct primes $N = pq$ with $p < q$ cannot be a Carmichael number (i.e., there must be a Fermat witness to the compositeness of N). We will need to use a fact about polynomials whose proof we are going to postpone until a week or so from now, when we take up a more systematic study of polynomials: A polynomial $p(x)$ of degree d can have at most d distinct roots (solutions of $p(x) \equiv 0 \pmod{p}$ in $\{0, 1, \ldots, p-1\}$) if p is prime (we saw a counterexample to this above is p is not prime, namely $x^2 - 1 \equiv 0 \pmod{8}$). The specific instance of this we are going to use for our Carmichael numbers allows us to assert that there are at most

$p-1$ solutions of $x^{p-1} \equiv 1 \pmod{q}$. Since $p < q$, this means we can choose $x \in \{1, \ldots, q-1\}$ such that $x^{p-1} \not\equiv 1 \pmod{q}$. We claim that such an x is a Fermat witness for N. In particular, we have

$$x^{N-1} \equiv x^{pq-1} \equiv x^{pq-q}x^{p-1} \equiv (x^p)^{q-1}x^{p-1} \equiv x^{p-1} \not\equiv 1 \pmod{q}$$

(we know that $(x^p)^{q-1} \equiv q \pmod{q}$ by Fermat's little theorem because q is prime). And since $x^{N-1} \not\equiv 1 \pmod{q}$, we have $x^{N-1} \not\equiv 1 \pmod{N}$ and we are done with the proof of the lemma.

Now we're in the home stretch:

Proposition: If N is a Carmichael number, then at least $\frac{3}{4}$ of the numbers in $\{0, 1, \ldots, N\}$ are Miller-Rabin witnesses to the compositeness of N.

Correspondingly, we're going to show there are at most $\frac{n-1}{4}$ Miller-Rabin liars. As before, factor $N-1$ as $N-1 = 2^t u$ with $t \geqslant 1$ and u odd. Partition the set $\{1, 2, \ldots n-1\}$ into disjoint subsets $X, Y, Z_1, Z_2, \ldots, Z_t$ as follows:

- $x \in X$ if $\gcd(x, N) > 1$.
- $x \in Y$ if $x^u \equiv 1 \pmod{N}$
- $x \in Z_j$ if $x^{2^j u} \equiv 1 \pmod{N}$ but $x^{2^{j-1}u} \not\equiv 1 \pmod{N}$.

The set X contains no Miller-Rabin liars, since all the elements of X are in fact Fermat witnesses for N. We claim that $|Y| < \frac{1}{8}(N-1)$ and each Z_j contains at most $\frac{1}{7}|Z_j|$ Miller-Rabin liars. If this claim is true, then

$$\text{number of MR-liars} \leqslant |Y| + \tfrac{1}{7}(N-1-|Y|) = \tfrac{6}{7}|Y| + \tfrac{1}{7}(N-1) \leqslant \tfrac{6}{7}(\tfrac{1}{8}(N-1)) + \tfrac{1}{7}(N-1) = \tfrac{1}{4}(N-1)$$

which will complete the proof of the proposition.

So, let $N = p_1 p_2 \cdots p_k$ where $k \geqslant 3$ and the p_i are distinct odd primes. By the Chinese Remainder Theorem, choosing $x \in \{1, 2, \ldots, n-1\} \backslash X$ is equivalent to choosing numbers $x_i \in \{1, 2, \ldots, p_i - 1\}$ for each $i = 1, 2, \ldots, k$ (i.e., it must be the case that $x \not\equiv 0 \pmod{p_i}$ for all i because x and N have no common factors other than 1).

Now, $x \in Y$ means $x^u \equiv 1 \pmod{N}$. By the uniqueness part of the Chinese Remainder Theorem, this means that $x_i^u \equiv 1 \pmod{p_i}$ for all i (since 1 must be the only number in$\{1, \ldots, N-1\}$ with $x_1 = x_2 = \cdots = x_k = 1$). But we can show that at most half of the numbers y in $\{1, \ldots, p_i - 1\}$ satisfy $y^u \equiv 1 \pmod{p_i}$: This is because the set B_i of y satisfying $y^u \equiv 1 \pmod{p_i}$ forms a group under multiplication, but $p_i - 1 \notin B_i$ since u is odd and so $(p_i - 1)^u \equiv (-1)^u \equiv -1 \not\equiv 1 \pmod{p_i}$. Therefore $|B_i| \leqslant \frac{1}{2}(p_i - 1)$.

Therefore, when we specify x by picking $x_i \in \{1, \ldots, p_i - 1\}$ uniformly and at random for each $i = 1, 2, \ldots, k$, the probability that all of the x_i satisfy $x_i^u \equiv 1 \pmod{p_i}$ is at most $(\frac{1}{2})^k \leqslant \frac{1}{8}$ since $k \geqslant 3$. Thus $|Y| \leqslant \frac{1}{8}(N-1)$.

Finally, we turn to the sets Z_j. As above, an element x of Z_j can be represented by the k-tuple $(x_1, x_2, \ldots, x_k)$, where $x_i \in \{1, \ldots, p_i - 1\}$, but now $x \in Z_j$ means that $x_i^{2^j u} \equiv 1 \pmod{p)_i}$ for all i. Because p_i is prime, this implies that $x_i^{2^{j-1}u} \equiv \pm 1 \pmod{p_i}$ for each i. So to each $x \in Z_j$ we can

associated a k-tuple of $+$ and $-$ signs according to the signs of $x_i^{2^{j-1}u} \pmod{p_i}$. Note that the sign sequence is *never* $(+,+,\ldots,+)$ or else we'll have $x \in Z_\ell$ for some $\ell < j$.

Since $-1 \pmod{N}$ is represented by the sign sequence $(-,-,\ldots,-)$, the only way for x to be a Miller-Rabin liar is for its sign sequence to be all minus signs. If Z_j contains a Miller-Rabin liar, then the equation $x^{2^{j-1}u} \equiv -1 \pmod{p_i}$ has a solution which we'll call y_i for each $i = 1,\ldots,k$. Now use the Chinese Remainder Theorem to find a number $w_i \in \{1,\ldots,N-1\}$ satisfying $w_i \equiv y_i \pmod{p_i}$ and $w_i \equiv 1 \pmod{p_j}$ for all $j \neq i$. Then multiplication of $x \in \{1,\ldots,N\}$ by w_i flips the ith sign in the sign sequence of x and leaves the others fixed — and thus gives a $1-1$ correspondence between the elements of j with the original and flipped sign sequences. Since you can get from one sign sequence to any other by a succession of individual sign flips, each of the $2^k - 1$ possible sequences occurs an equal number of times (namely $|Z_j|/(2^k-1)$. And since $k \geqslant 3$, we have that the number of liars (i.e., the size of the set with all negatives in the sign sequence) in Z_j is at most $\frac{1}{7}|Z_j|$.