Light waves: diffraction and interference

In Ch. 22 we learned light is a *wave*, and then in Ch. 23 we promptly ignored this fact! It took people a long time to notice it's a wave, because λ is so small.
(Recall, violet light has λ=400 nm, red light is 700 nm)

Here is a "Ch. 23" picture of light coming in from the left, hitting a wall with a hole or slit in it, and forming a shadow on a distant wall.

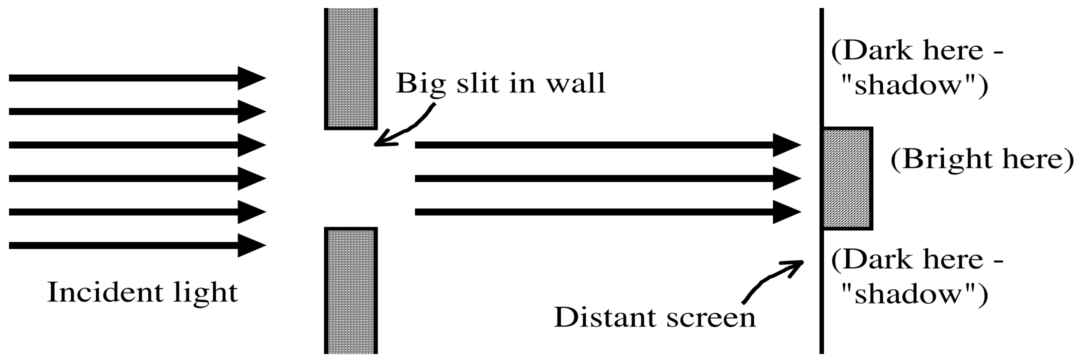

This was basically Newton's idea (and is quite accurate!)

However, if the size of the slight is very small - somewhere on the same scale as the λ of light, then the picture is very different.

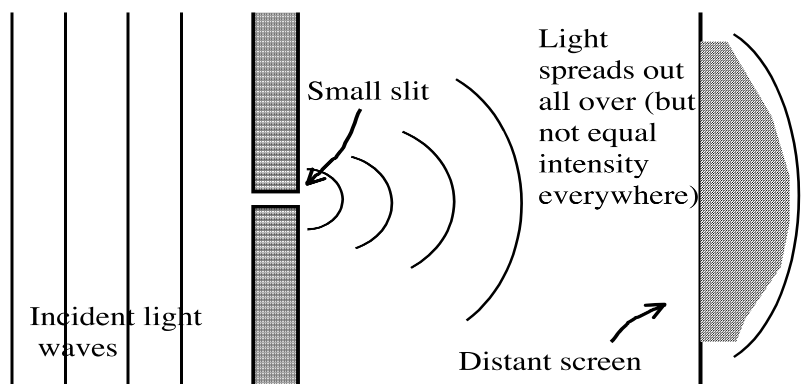

The lines I drew are NOT the same as in the previous picture, they are not "rays". Now we're trying to indicate the wave nature by drawing "wavefronts". You might think of this as connecting points where |E| is cresting. These lines are now *perpendicular* to the direction of travel of the waves. (Think of it as looking down at waves on the ocean, and what I'm drawing is the lines where the waves are highest) In the picture above, the waves move right.

(The image on the screen isn't quite right - we'll make a more accurate picture soon!)

<u>Brief review of interference of waves</u> (Phys 2010 Giancoli Ch. 11.11)

Imagine two traveling waves (same frequency, speed, and direction)

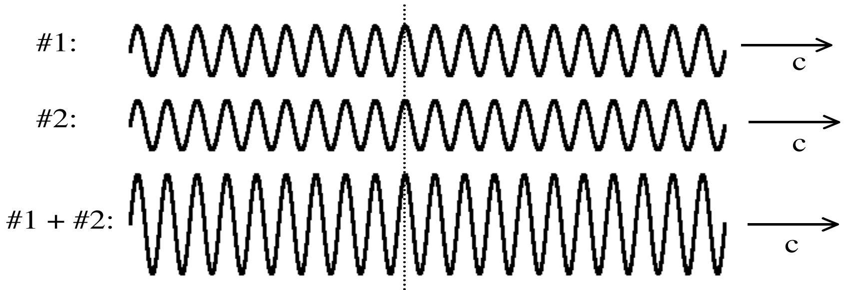

If they are "in phase" (meaning, the peaks are lined up) and if the waves are traveling at the same place and time, they will add up, they **superpose**, as waves always do. The resulting wave, the sum of the two, has the same f, c, and direction, but *twice* the amplitude. The waves have **constructively interfered.**

Now take those *same* two waves, but let them be exactly "out of phase" (meaning, the peak of one is at the same place as the *trough* of another!) Then the result when they superpose is cancellation, the "+" of a peak adds to the "-" of a trough giving zero, no wave at all!

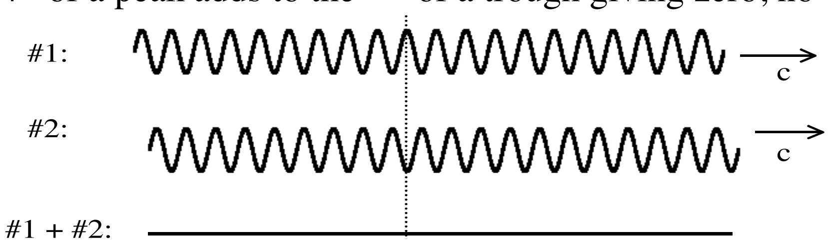

The waves have **destructively interfered.**

You can have two waves heading towards a common point. At that point the waves will interfere. If they are "in phase", they will add up to double the amplitude.

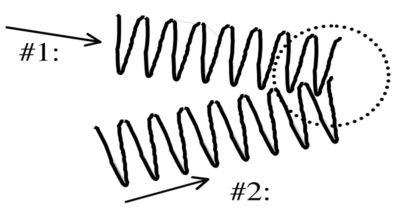

In this figure, both waves peak at the point where they meet - they add up: *constructive interference*.

If you were to shift *one* of the waves a little, so one peaks where the other "troughs" at the meeting point, they would instead be completely "out of phase", they would exactly cancel: destructive interference at that one point.

[If they are just partly out of phase, they will add (interfere, superpose) but not to either extreme. Two waves of amplitude A can add up to a wave of anything between 0 and 2A.]

Huygens (a contemporary of Newton) came up with a neat way to think about traveling waves:

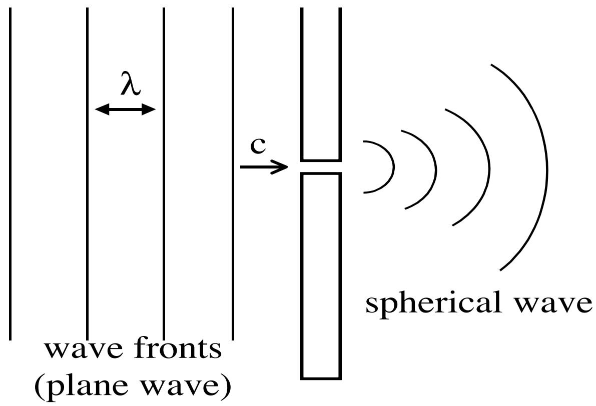

Each point on *any* wavefront can be considered to be the source of a new outgoing spherical wave.

In this figure, only one tiny point (center of the gap) of the incoming wavefront is allowed to propagate to the right.

What you get is an outgoing *spherical wave*, with the same f, λ , and c as the incident *plane wave*. The little gap acts as a single "point source" of wave of frequency f...

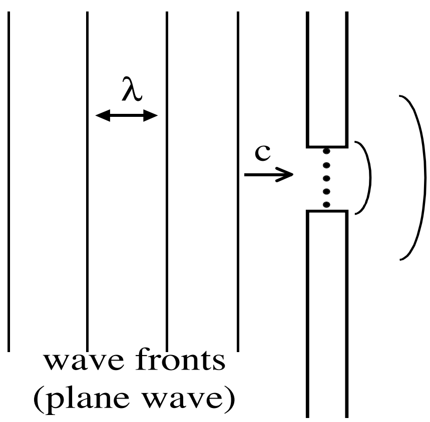

With a wider gap, each point in the gap acts as a point source. Each of those points emits spherical waves!

These waves are "on top of each other", you have to superpose them, *add them up*, to figure out the resulting outgoing wave.

(The outgoing wave here is not spherical any more, nor is it a plane wave. It's something in-between)

We won't do that math in this class (it's a little tricky, though certainly not impossible. Huygens worked it out 350 years ago!) but the result in the end if often very simple, and quite reasonable.

(Giancoli's figures are better than mine, check them out)

As an example, Huygen's principle fully predicts the phenomenon of refraction (!) The details are, as I said, tricky (again, Giancoli has somewhat better pictures), but the result is straightforward:

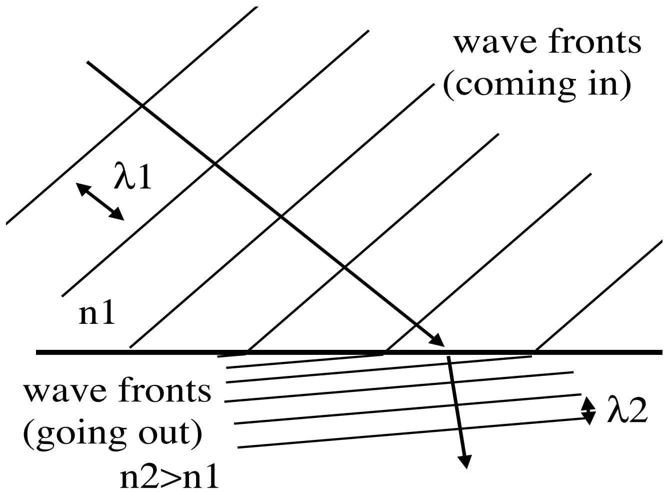

Waves always have λf = speed, so

$\lambda 1\ f1 = c/n1$
$\lambda 2\ f2 = c/n2$.

Now, here's an important consequence of Huygen's ideas: f1 = f2. Why? As the incoming wave reaches the boundary and creates the outgoing wave, it is "jiggling" with frequency f1, and that's the same frequency that the new medium will jiggle with. (Frequency is just counting, it does not depend on what medium you're in. But the resulting λ *does* depend on the medium.)

The equations above tell us that if f1=f2, then
$\lambda 1 / \lambda 2 = n2/n1$ (can you see how I got that?)
That means in the picture above: $\lambda 2$ is *smaller* because n2 is *bigger*.
Wavelength depends on speed. In a "fast medium" (smaller n), the wavefronts are more spread out, λ is bigger.
The picture shows wavefronts, remember that the "rays" these represent are perpendicular to the wavefronts. (I have drawn two sample rays in the figure) The angle at the boundary is bent. Refraction!

Huygens principle predicts that if light hits small slits, you should be more likely to observe the "wave nature" of light. (Look back at the pictures on the previous page - the outgoing wave is more noticeably different from the incoming wave when the slit size is small)
To really see this effect in the lab, it helps (a lot!) to have **monochromatic** light (one pure wavelength, which means one pure color) and also **coherent** light (meaning nice incoming plane waves with all rays "in synch" with each other). (Lasers are a great source of monochromatic, coherent light, but they're not the only one!)

Young's Double Slit Experiment: The first observation of interference of visible light!

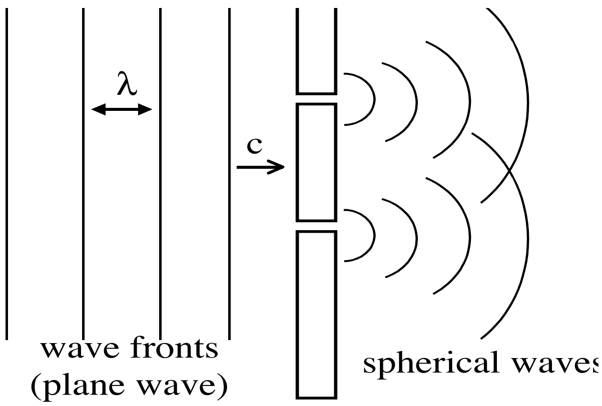

Huygens says each of the two slits will act as a point source for outgoing spherical waves.

On the right, light is emerging (heading off in *all* directions!) from each slit.

The incoming light is monochromatic (one λ), and coherent, which means that if the E field is peaking at the top slit at some instant in time, it is also simultaneously peaking at the bottom slit.

What will we see on a screen far off to the right? The old "particle" or "ray" model predicts this: two bright spots, one directly behind each slit.

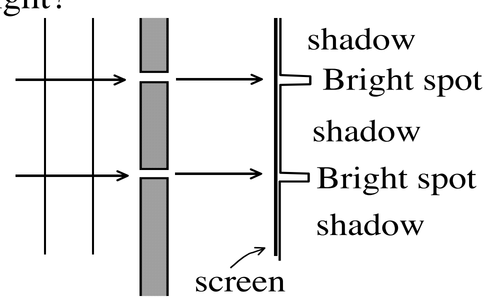

This is experimentally dead wrong!! (At least, if the slits are very small compared to λ.)

Huygens model of interfering waves predicts this - (We'll see why soon): And, this is what you really see. Never mind Giancoli Fig 24-9 and 10, for now! For two SMALL slits, each bright "bump" in the image is equally bright. You see a *bunch* of bright spots over there!! (We'll explain exactly why you get this pattern on the next page) They are *not* necessarily bright right behind the slits, either! And it's brightest RIGHT BEHIND THE "WALL" between the 2 slits, where you'd expect a shadow.

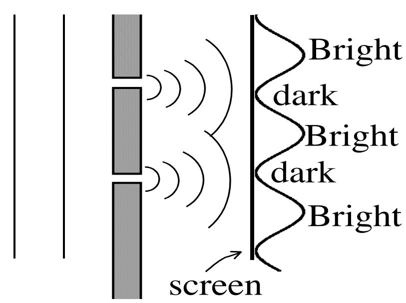

This is weird. "Shadows" are what we're used to, but the wave nature of light says it can effectively bend around corners. This phenomenon is called **diffraction**. It's not so intuitive (for light.)

How can this be? What is going on?
And, how do you figure out where the bright spots will be?

What we have is two slits - two point sources of light, and they are in phase with each other. If light goes the *same* distance from each (e.g. the waves leave the sources and head to the "midpoint" on the screen) then the two waves will arrive in phase (in synch) - they will *constructively interfere,* and you get a bright spot on the screen there.

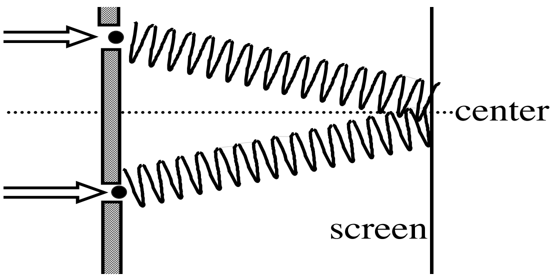

Yes, it's *bright* right smack in the center, just exactly where you'd expect a dark shadow from that central wall!!

Are there any *other* bright spots? What really matters is whether the TWO waves reach the screen *in phase*.

Draw the rays: the light from each slit travels *different* distances.
If the two waves BEGIN in phase, and travel different distances, they might still be in phase, or not, it all depends on the difference in path lengths. (This is tricky to visualize!)

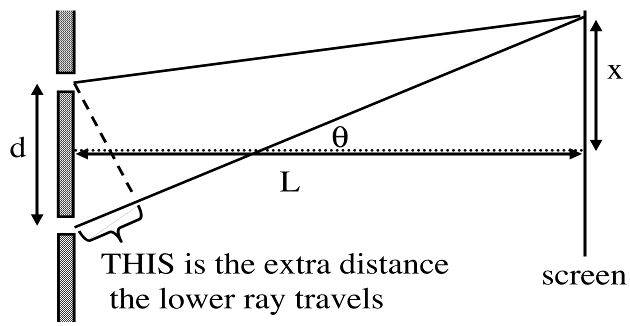

Let's look at a spot on the screen a height "x" above the center line, a distance L (far) away from the sources.
Is it bright or dark there?

It all depends on the *difference* in the path lengths of the two rays...

(I'm making an assumption here, a limit... that L is HUGE compared to d. That means that the figure *really* looks like this:

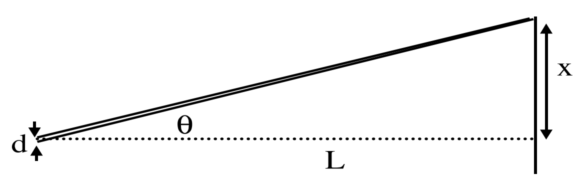

I just exaggerated "d", above, so we could *look* at the geometry more easily.)

The question is: is the extra distance EXACTLY a whole number of full wavelengths? If it is, then the two waves are still in phase!

Think about the condition for waves to be in phase. If a wave travels through exactly ONE full wavelength, it's back in phase. Same if it travels through exactly TWO full wavelengths,...

So if the extra distance traveled by one ray compared to the other exactly equals λ (or 2λ, or 3λ, or $m\lambda$, where "m" is any integer) then the two waves are still (or back) in phase, and they constructively interfere, which means the spot we labeled "x" will be bright.

If, on the other hand, one wave travels through HALF a wavelength extra, it's exactly out of phase! That means the two waves destructively interfere, it would be totally *dark* at point "x". (Same whether the path length difference is $0.5\ \lambda$, or 1.5λ, or 2.5λ, etc...)

We need to look at the picture again, this time a "zoom in" right down next to the slits...

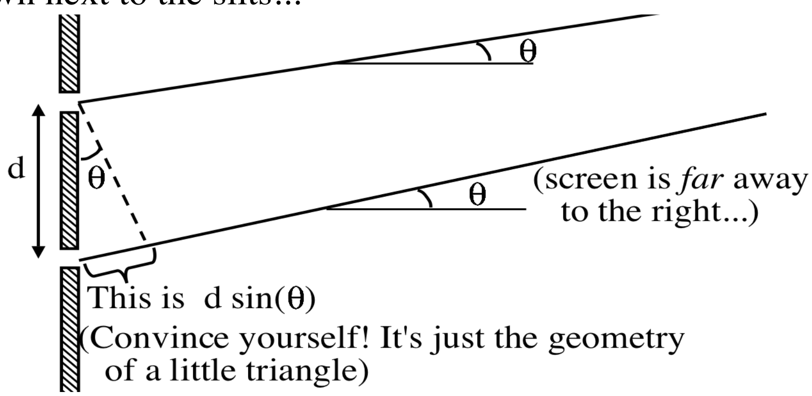

The picture shows that the path length difference is $d \sin(\theta)$.
The geometry is tricky! Think about it, compare to the previous pictures, see if you can really get it for yourself.

What we've just argued is that if that path length difference is an INTEGER number of λ's, we get constructive interference, and thus the spot x will be *bright*:

Constructive: $d \sin(\theta) = m\ \lambda$, (where m=any integer, 0, 1, 2, 3, ...)

Similarly, if the path length difference is a HALF integer number of λ's, we get destructive interference, the spot x will be *dark:*
Destructive: $d \sin(\theta) = (m+1/2)\ \lambda$, (m=any integer, 0, 1, 2, 3,...)

Quick review of radians, and the "small angle" approximation:

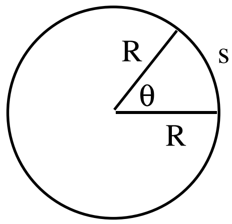

R = radius
s = "arc length" along circle

$\theta = s/R$ defines the angle in radians.

Full circumference is $s = 2\pi R$.

That means θ (all the way around) $= s/R = 2\pi R/R = 2\pi$ rad.

Or, $360 \text{ deg} = 2\pi \text{ rad.}$

Now look at a small wedge, and "square off" the curvy side:

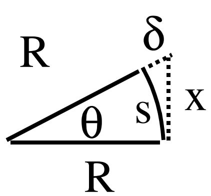

We know $\theta = s/R$
and we know $\tan\theta = x/R$
and we know $\sin\theta = x/(R+\delta)$.
(Those all come from "sohcahtoa")
But, if θ is very small, we see $s \approx x$ (roughly, just look at the picture)
and $\delta \approx 0$, too. Conclusion: if θ is very small, $\theta \approx \tan\theta \approx \sin\theta$!
(That would also mean $x = R\tan\theta \approx R\theta$)
Example: (try this on your calculator!) 1 deg = .017453 rad,
sin(1 deg) = .017452 and tan(1 deg) = .017455. These are very close!

Finally, let's go back to our "2 slit interference" picture. I'm going to simplify the picture by only drawing ONE "average" ray, not both of them. (Remember, d is really tiny compared to L!)

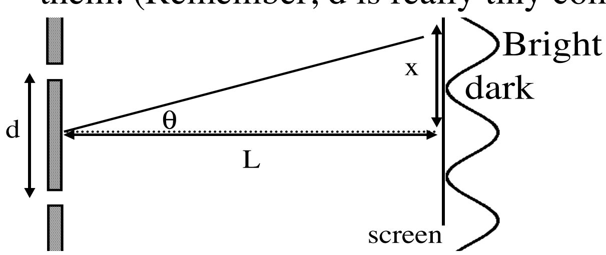

We already argued there are bright spots whenever
$d \sin(\theta) = m\lambda$.

Note: $\theta = \tan^{-1}(x/L)$,

Since L is huge, we frequently have x/L small, and then $\theta \approx x/L$
So, for small x, and big L, $d\sin(\theta) = m\lambda$ becomes $d\, x/L \approx m\lambda$, or

BRIGHT SPOTS: $x \approx m\lambda L/d$ (m = 0, 1, 2, 3, ...)

DARK SPOTS $x \approx (m+1/2)\lambda L/d$ (m = 0, 1, 2, 3, ...)

[If x/L is not small, just use the *correct* formula, $d\sin(\theta) = m\lambda$
Use this to find θ, then solve for $x = R\tan(\theta)$.]

What if we had more than 2 (equally spaced) slits?

Draw the picture - draw ALL the rays coming from all the slots towards a common point at the screen. You should convince yourself that as long as d sin(θ) = m λ, (m=0,1,2,...) then the light from ALL the slots will be in phase, constructive interference, VERY bright! At any other angle, you'll get destructive (or at least partially destructive) interference.

If θ is "slightly" off one of those magic bright values, then, say, the rays from slit #1 and slit #2 won't be very much out of phase (so if there are *only* two slits, it's darker, but only a little, if you shift a small angle away from the brightest point). But with many slits, slit #1 will be way out of phase with some other slit, maybe not #2, but one a few further down the row.

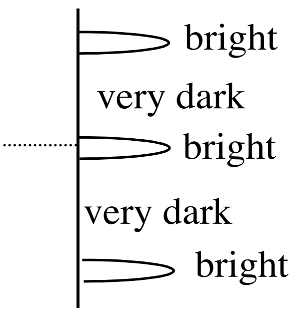

So you end up getting VERY bright spots at the places where the formula for 2 slits says it should be bright, and they're much *sharper*. (It's *very* dark almost everywhere else.)

This is handy for clearly separating the bright spots.

We call a device like this, with many slits, a **diffraction grating**.

They're very useful! Why? Notice the positions $x \approx m \lambda L / d$ for the bright spots depend (linearly) on the wavelength λ. Different colors incident on the same 2 slits (or grating) will make bright spots at *different locations*! The separation of the colored spots (x) is easy to measure, and is proportional to λ! That gives you a very handy and easy way to measure wavelengths of light. (Think about that a bit - how can you measure something so fantastically tiny? λ of light is just a few thousand times larger than an atom.)

Example: Consider 2 slits, a distance d=0.01 mm apart. There is a screen L=2m away. Monochromatic red light (λ=700 nm) shines on the 2 slits. What will the brightness pattern look like over at the screen?

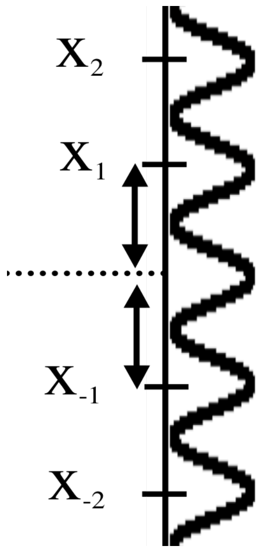

Lingo: x_2 is the "m=2" or "2nd order" line.

x_1 = 1st bright spot above the center. The "first order" line, or "m=1" line. $x_1 = 1^* \lambda L / d$
Here, x_1 = (700E-9 m)*(2m)/(.01E-3 m) = .14 m

"zeroth" order, or m=0 line, in middle. (m=0, x=0)

x_{-1} = 1st bright spot below the center, "m=-1" line.
$x_{-1} = -1^* \lambda L / d = -.14$ m
(Note the symmetry, $|x_{-1}| = x_1$.)
Etc.

Note: $\theta 1 = 1\lambda / d = .07$ rad, $\theta 2 = 2 \lambda / d = .14$ rad, etc.
(If x starts getting big, you should go back to the *correct* formula for the "angles of maxima", namely $d \sin(\theta) = m \lambda$.)

• If we'd used *blue* light (500 nm) instead of red, the pattern is a little different. x1 is slightly smaller, the bright spots are squeezed together more, towards the center (which is unchanged)
• If d was bigger, the pattern is also squeezed, in a similar way.

That's why it's so hard to notice this pattern in normal life. For visible light, λ is SO small, and d is so much bigger, that the bright spots all squeeze right next to each other - you can't even NOTICE that there are dark and bright spots. To make it worse, all the colors are normally present, with their bright spots in slightly different places... so it really just looks bright all over!

What if we used a grating with 1000 lines/cm instead of 2 slits? First, we have to think about the spacing between these lines... 1000 lines/cm means 1cm/1000 lines, or 1/1000 cm per line, or .01 mm per line - that's the same distance between lines as we had in the previous problem! ("lines" in a grating can be though of as analogous to "slits")

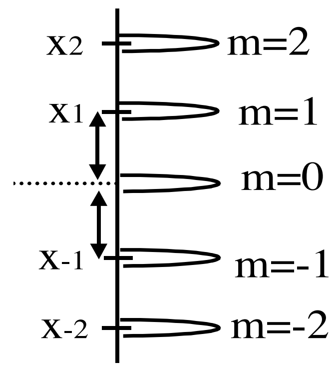

So the pattern is *identical* to that in the previous problem, except that the spots are brighter, and sharper.
The x's (positions) of the bright spots are the same.

If you shine white light (rather than monochromatic) there are many λ's present. Each color has bright spots at different positions x. So, the colors appear to get spread out on the screen. You get a little rainbow, a **spectrum** at each order. (One spectrum for each "peak" in the picture above)

Since $x1 = \lambda L/d$, measuring "x" for different colors can tell you what the wavelength of each color is.

Some light sources have different colors, but *not* all colors. (E.g., Mercury arc-lamps used on highways look yellowish-white, but do *not* contain all wavelengths with anything near equal intensity)

A diffraction grating which spreads out colors allows you to quickly and easily see which colors are present in the source. You will see *lines* of different colors at each order. This is a unique "fingerprint" of the light source. This technique is called ***spectroscopy***, and we use it to determine the chemical composition of materials. In this way, we can even deduce what elements are present in distant stars, by analyzing the light they emit!

Another way to get a **spectrum** is with a **prism**. How does this work? Different colors (λ) have slightly different indices of refraction "n" in normal glass. So, they get bent slightly differently.

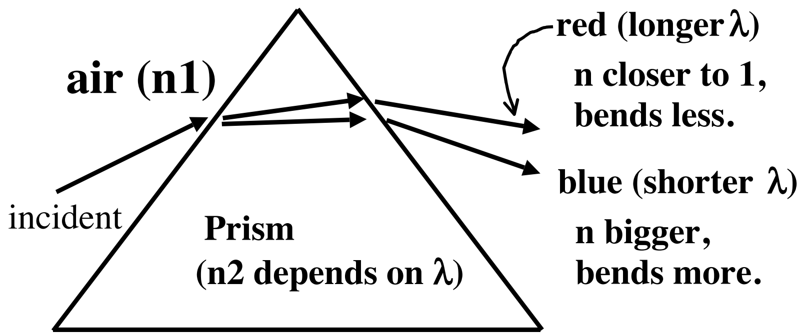

Notice that the spread of colors is *opposite* the way it goes with a grating. (In a grating, red has longer wavelength, and so bends more. In a prism, red has smaller "n", so it bends *less*.) The mechanisms involved are really totally different.

Rainbows are caused by the prism effect. Different colors bend through different angles as they pass from air into water droplets. The result is that white light gets dispersed into a spectrum of colors. The sun produces all the colors of light, so the rainbow has them all. (See Giancoli Fig 24-16)

Take a look at a CD. It's got a lot of closely spaced lines (and is reflective), it acts like a grating! It's a "reflection" grating rather than a "transmission" grating, but the physics is the same. Again, at different angles, you see different colors. Check it out - hold it at an angle from a single light source. Can you see "ROYGBIV"? If you keep tilting, can you see the second order? How many orders does it have? Can you estimate the spacing between the lines?