PROFESSOR: Another strange observation about mirrors that I've never really understood-- the mirror plane is reflecting me left to right, so it looks as though a mirror has a grain to it. It knows what line to reflect me back and forth. But if I kept the mirror in fixed orientation and I lay down, the thing should reflect me side to side, but it doesn't. It still reflects me from top to bottom.

So how can that be? Why does a mirror plane, if I hold it in one orientation, appear to have a direction across which I'm reflected but it doesn't follow me if I move? Know what I mean? You have any explanation of that? Hmm?

AUDIENCE: Rotate your eyes, too.

PROFESSOR: Rotate my eyes. I can roll them around. I can't rotate in any other fashion.

That's strange. I mean, you look at yourself every morning-- several times, perhaps, and you're reflected always left to right. And if you turn the mirror, it doesn't reflect you top to bottom. Or conversely, you can leave the mirror alone and you can-- why does a mirror plane just reflect you left to right? Ah. I'll let you stew about that one for a while.

And I can tell you that I know of three papers in scientific journals that tried to explain this. I can give you literature citations, but you think about it until our next meeting. Why does it know which way I'm oriented? Ah, I can't think about it.

Alright, back to more straightforward questions. We are now about to embark on a grand process of synthesis which will take us the better part of half a semester. And we've identified our four basic operations-- four basic one step operations-- namely translation, reflection, rotation, and for the time being I'm going to look just at two-dimensional symmetries, so I'll leave inversion out of the picture. It's only defined in three dimensions, and the logic which I will follow will be to first build up two-dimensional symmetries and then we'll turn them into three-dimensional symmetries by picking another translation that's not coplanar with the first two.

So we're going to look for the time being at just two-dimensional symmetries. The nice thing about doing this is that the number of two-dimensional symmetries is relatively small, so we can derive them rigorously and exhaustively. To do so in three dimensions is a much more time consuming and elaborate exercise. It's no different in principle, there's just more work. So if we do two-dimensional symmetries first, it's an easy case. We can do it rigorously and completely, and then what we'll do is just look at a few examples in three dimensions and look at how the results are designated and tabulated.

So the first combination I will choose to make is one that I set up last time and should not have started when there was no time to finish it. So let me take an initially pristine space and say that to this I'm going to the operation of translation. That immediately implies a string of translations and a string of lattice points, but I'm just going to focus my attention on the first one.

Then what I said I'll do is add to the space a rotation operation A alpha. I'm going to for convenience put it right at the lattice point, but as there is no unique origin to the translation, I can start and stop the translation anywhere I like.

So here's my translation. The rotation operation is going to take that translation and repeat it at angular intervals alpha so that I get a radiating porcupine-like sheaf of translations all coming out of a common point. Alpha, we observed, has to be a submultiple of 2 pi, so I will have a cluster of translations separated by the equal interval alpha-- angular interval alpha-- and I'm going to choose to focus my attention on just the first one. So this, since it's repeated by rotation, has the same length T.

And then I'm going to look at the other end of the translation and say that similarly I must have a set of equally spaced angular-wise translations all separated by alpha, and I'm going to choose to focus my attention just on this one. So this angle is alpha, and this angle here is alpha.

And now there's big trouble in River City. This is a translation, here's a lattice point, here's a lattice point. This is a translation-- a lattice point sits up here. This is a translation-- a lattice points sits up here. Now I've got two lattice points eyeball to eyeball, and they jolly well better be separated by either the interval T or some multiple P times T, where P is some integer. It would be quite all right if there were two translation separation or five translation separation, but it must be some multiple of translation or I have violated my initial premise that everything in this space is periodic at a translational interval T.

So let me take this geometrical constraint and very quickly convert it into analytical form so that we can proceed to systematically find out what the possibilities are. Let me drop a perpendicular down to the original translation. So this distance in here is PT, this distance will be T times the cosine of alpha, and this distance in here will be T times the cosine of alpha. So in analytic form, then, I can say that my original translation T is equal to T cosine of alpha plus T cosine of alpha. That's 2T cosine of alpha plus PT.

And the first thing we can see is that the magnitude of T drops out because this construction and the constraint it embodies in no way depends on the magnitude of the translation. So this says, then, that 1 is equal to 2 cosine of alpha plus the integer P. And if I solve for the value of cosine of alpha, cosine of alpha will be 1 minus an integer P divided by 2. So there is the constraint that must be followed if my construction is to be self-consistent.

So we've got this now in a plug and chug situation. So what I'm going to put down is the value of P, taking all possible integers for which a value of cosine of alpha exists. I'm then going to evaluate cosine of alpha, which is 1 minus P over 2, and then I'm going to identify the n-fold axis that corresponds to that particular value of alpha.

Let's put in P equals 3, and this is as far as we got last time. Well, if we put 4, then cosine of alpha is minus 3/2-- it's not defined. If we drop P down to 3, then cosine of alpha is minus 1, 1 minus 3 over 2. And the angle whose cosine is minus 1 is a-- let me put down the value of alpha and then the n-fold axis. The angle whose cosine is minus 1 is 180 degrees, and that would describe quite nicely the rotational throw of a twofold axis.

If I let P drop down to the value 2, then I have minus 1/2 of the cosine of alpha. The angle whose cosine is minus 1/2 is 120 degrees. And guess what? That's the angular throw of a threefold axis. Let P drop down to 1, and then cosine of alpha is 0. The angle whose cosine is 0 is 90 degrees, and that would be a fourfold axis.

Looks like we're done, except P could be equal to 0. In that case, cosine of alpha is plus 1/2. The angle whose cosine is plus 1/2 is equal to 60 degrees, and that's a sixfold access. What about negative integers? Minus one, will that work? This says that cosine of alpha is 1, and the angle whose cosine is one is 0 degrees or 360 degrees. And that would be no rotational symmetry all. It's rather amusing that the trivial case of no symmetry at all also falls out of this construction.

So this is a momentous result. We've shown that if you're going to have a pattern that has a repetition by translation in it, the number of rotational symmetries that can be added are either no symmetry at all, a twofold rotational symmetry, threefold, fourfold, or sixfold. In other words, the axis-- no symmetry at all, twofold, threefold, fourfold, or sixfold, nothing else.

This tells you a lot about the shapes that you can see macroscopically on crystals. You could have a crystal that had the shape of the trigonal prism, and that would be perfectly fine. You could have a crystal that had the shape of an orthogonal brick that would have twofold axes coming out of the faces. You could have a crystal in the shape of a hexagonal present, or you could have a crystal in the shape of the square prism. But something thing like a crystal with a pentagonal cross section that would be a fivefold access-- that is strictly forbidden, because the external shape of the crystal has to reflect the internal symmetry among the arrangement of atoms.

There are lots of things in nature that have crystallographic symmetry. There is a little cactus that looks like this with some little spines coming out like this on the top, and its proper name is astrophytum myriostigma, also known as the ornamented bishops cap. Beautiful example of fivefold symmetry, but the cells inside of that cactus can not have the same size and shape and be repeated by translation.

There are flowers that have very common examples of fivefold and even sevenfold symmetry. There's one little purple flower that looks like this that comes out in the spring, and that's called periwinkle. Fivefold symmetry-- fine for a plant, but if you got down inside the stem of this flower, the cells cannot be repeated by translation.

There are astronomically high symmetries. These big giant cacti in the Southwest, the saguaro, they have rotational symmetries that run from 18-fold, 19-fold, 23-fold, so that if you picked up one of these guys very carefully, because they're covered with spines, and if you could lift it-- which would be very hard because they weigh a couple of tons-- take that guy and rotate them through 1/27 of a circle, and if he had 27-fold symmetry, you could plop them down and you couldn't tell that it had been moved.

There is no crystallographer who can resist cacti. All sorts of symmetry, even symmetries that violate crystal graphic symmetries-- wonderful textures and colors. And they have another really remarkable property, which commends them as house plants. If they die, this sheath of spines stays intact until someday when you're watering it, you brush against it and you poke a hole right through the spines and there's nothing inside. So a house plant that dies and you can't tell for two years or so is a very good plant to have as a companion.

So cacti have all sorts of strange symmetries. Fine, but you can say something about the internal structure of that flower or that cactus.

But this little almost trivial proof has told us something else about crystals, because the presence of translation imposes a constraint on the rotational symmetry that can be present, and the rotational symmetry tells us something about the nature of the two-dimensional lattice which can accommodate that symmetry. So let's go through this list once more, and let's pay attention to the value of P.

For a twofold rotational symmetry, what we would do would be to take this original translation, we put the twofold axis here, and that takes the translation and rotates it around. So here's a lattice point here, one here, one here, and the twofold axis here, takes the translation and we rotate in a counterclockwise sense. Here's another lattice point here. P was equal to three translations, and I'll be darned if that isn't exactly what we have.

Three translations from this lattice point A-- let me put some labels up here. This was lattice point A, this was lattice point B, and this was our translation PT. Three translations, just as advertised.

What constraint does this put on a lattice? None whatsoever, because all this says is that if you have a translation that translation must be repeated into an extended one-dimensional row. So you can put this an any 2D lattice whatsoever. In other words, the magnitudes of the two translations-- let's call them T1 and T2 are under no constraint to be related one to another-- and this angle between them, alpha, can be anything that it likes. You could have a twofold axis, but that requires simply a lattice row parallel to T1 and a lattice row parallel to T2.

The next integer we hit was minus 1/2. That was cosine of minus 1/2, this was P equals 3, and that corresponded to something that could accept a threefold symmetry. So let me put a guide to the I in here and make an equilateral triangle. If we translate and rotate up by 120 degrees using a threefold axis and then rotate minus 120 degrees about the other lattice point, that puts another translation up here, and now-- I'm sorry, this was P equals 2 And as required, this is PT, and that's equal to two translations.

This has put a constraint on the sort of lattice which can exist in this space because we have two translations, T1 and T2, which are equal in magnitude, identical in magnitude, because they're related by symmetry. And consequently, we have defined a space lattice, a two-dimensional space lattice. We'll call this T1 and call this T2. This lattice has the constraint that magnitude of T1 must be identical to the magnitude of T2, and the angle between T1 and T2 is again identically 120 degrees. Not 119.9, but exactly 120 degrees because there is a threefold axis in here that demands that that be so.

So this is a very specialized kind of lattice, restricted to have two translations identical in magnitude. And if there's a threefold axis there, you should be able to find these two. And if there's a threefold axis there, then the angle between these two specialized translations is 120 degrees.

Our next magic integer was 1, and that corresponded to a fourfold axis. And if we do what we claimed we did in that construction, we'd put a fold axis here. That takes T1 and rotates it exactly 90 degrees to a translation T2. Once again, two noncollinear translations, so we have defined a two-dimensional net. If we complete the cell, this is one translation. PT is equal to one translation in here, and this angle is 90 degrees because it's produced by a fourfold axis.

So we have a very special lattice. Again, the magnitudes of two translations are identical-- not approximately the same, they're identical-- and the angle between them is identically 90 degrees.

Only a couple to go. P could be equal to 0, and that was the case for a 60 degree rotation, a sixfold axis. So again, let's draw what came out of this particular special case. Here's T1, here is T2. This angle is 60 degrees exactly because there's a six-fold rotation axis here, a sixfold rotation axis here, and the rotation of 60 in the opposite sense gives us another translation here. These two lattice points coincide and there is PT equal to 0T, and these two points coincide.

Now if I complete a standard unit cell with T1 T2 as I've done in other cases-- this was T1, this was T2, this was a translation which I'm not going to use. So this is the shape of the lattice and these now are lattice points with a sixfold axis on them. The dimensional specialization is exactly the same as I found for a threefold axis, T1 identical to T2. If I pick this cell, T1 to T2 can be described as an angle of 120 degrees. Exactly the same lattice that we found for a threefold axis.

So with this simple minded little construction we've found two profound things-- that there are five kinds of rotation axes, including the onefold no rotational symmetry at all. And it turns out that there are one, a general lattice, a hexagonal lattice, and the square lattice. There are three kinds of two-dimensional lattices that are required by these symmetry elements.

So these guys require that there be three lattice of different specializations that are able to accommodate them. So let me call this a parallelogram net, and that has T1 not equal to T2, the angle between T1 and T2 general. And this is a lattice that can accommodate either no symmetry at all or a twofold rotation access. Then there was a net that I'll call the hexagonal net, and this had T1 identical to T2 in magnitude, and it had the angle between them, the angle between T1 and T2 as identically 120 degrees. And this could accommodate either a threefold or a sixfold axis.

And then finally, the general net, which I call a parallelogram net, and that has magnitudes of the two translations not equal to one another. They can have any values they like, and the angle between T1 and T2 is completely general. And that's exactly what I had up here for the-- oh, we did that once already. The one that I'm missing, the third one, is the square net.

Square net has T1 identical to T2. In magnitude, the angle between them is exactly 90 degrees, and that is required by a fourfold axis.

Let me pause here to see if there are any questions. Yes, sir.

AUDIENCE: When you write on the last one with the sixfold, it's only 120. You could have also written 60, correct?

PROFESSOR: Yes, I could have. And this lattice is actually the same as what I found for a threefold axis. I could pick either this or this as the cell, but the two translations in those two cells are equal. And again at various stages along the way, we'll need a convention. And if I have a net that looks like this, a parallelogram, whether a specialized parallelogram or not, I have a choice of two angles that I could use.

We call that alpha. This is going to be 180 degrees minus alpha. Which do I pick? You need a rule, and the rule is that for labeling the cell, pick alpha so that it's greater or equal to 90 degrees. That's pure convention, but you want to have a rule just like a language so people use the same words to describe the same thing. Here we want to use the same geometry to describe one and the same thing.

The other convention is that there are no unique translations that define a net. We could take linear combinations of these vectors and define the same lattice, so we need another rule for selecting the standard translations-- again, so that two people can do an x-ray diffraction experiment and report the results in terms of the same lattice, and this is a fairly reasonable thing. This is pick the two shortest translations, and that clearly makes sense. There's absolutely nothing at all to commend a cell that has this as T2 and this as T1 so you get a long, skinny oblique things. Your natural inclination would pick the two shortest translations in the net.

OK, so these are conventions. This has nothing to do with the nature of the symmetry or what makes it unique, but just so that people have one defined way of labeling things.

AUDIENCE: Right. [INAUDIBLE] my question was actually that--

PROFESSOR: It was a good answer, even if was not to the question that you asked.

AUDIENCE: I said the sixfold and the threefold are exactly the same, but then I realized they are because there is no crystal that can have threefold symmetry without sixfold.

PROFESSOR: You were doing fine. You should have quit just before that last statement. How can you have a hexagonal lattice that sometimes has sixfold symmetry and sometimes has threefold symmetry?

Well, let me give you an example for that, and it makes a very useful point. Let me draw two hexagonal nets, and in this one I'll put a threefold axis. So I'll have one motif here, I'll go 120 degrees away. Here is a motif here, and I'll go 120 degrees away, and here is a third motif. So these three guys form a triangle about this lattice point. And we would have about the other lattice points at the corners of the cell exactly the same triangle of motifs, and the same thing over here.

So there is a pattern. It has a hexagonal lattice, and it has a triangle of objects related by a threefold axis. And now let me take exactly the same lattice, and now I'll put in instead a sixfold axis. And that means I'm going to have a hexagon of objects, and it'll do something like this.

And I don't want to push my luck and try to draw that twice, but there would be a hexagon here and a hexagon of motifs here, another one at this lattice point, and another one at this lattice point here. Same lattice, same shape, same dimensions-- although that's not critical-- but one of them has only a threefold axis. One of them has only a sixfold axis in it. Why? Because I decided to put a threefold axis in this pattern, and I decided to put a sixfold axis in this pattern, but they both end up being contented and happy with a lattice with the same degree of specialization.

Now we will come as we progressed a little bit further that we have to go in two dimensions to the reverse situation. We have a particular symmetry and it's happy with two different sorts of lattices with different shapes and different specializations, and that's going to come up directly. Any other questions? Yeah.

AUDIENCE: In this example, the sixfold [INAUDIBLE] axis [INAUDIBLE]. It's just that [INAUDIBLE].

PROFESSOR: Yes. You're saying that here hanging at this lattice point is something that has sixfold symmetry. Here is something has only threefold symmetry. The nature of the lattice and the symmetry that is in that lattice are two inseparable aspects of the pattern. But often, as we've seen here, there's more than one possibility for a given lattice. And as we'll see very shortly, for some other symmetries, for one given symmetry, there are two kinds of lattices.

But nevertheless, the thing to keep fixed in your mind is that we call this a hexagonal lattice. Why? Because this translation is equal to this one, and they are exactly 120 degrees apart. That lattice can have that specialness only if there's either a threefold or a sixfold axis in it. So the specialization of a lattice is inseparable from the symmetry that is in the lattice that demands that specialization of the lattice.

Conversely, a lattice can have the specialization only if you place in a symmetry which demands precisely that specialization. So if you measure a lattice, and this turns out to be 119.99 degrees and these turn out to be 3.21 angstroms and 3.21 angstroms, that is not a hexagonal lattice, because there's no symmetry in there that demands that this angle be 120 degrees.

And that may seem to be an academic fine point, but we'll see that in due course the properties of a crystal depend on the symmetry of that crystal. If the crystal has symmetry, the property also has to have that symmetry. And it is the atoms inside the cell which determine the symmetry of the property, and the properties is one aspect of the symmetry that goes along with lattice dimensions and lattice angles. Is that clear?

So let me say it again, because this is an important point. The specialness of a lattice is inseparable from the symmetry that is existing in that lattice that demands that specialness. So if you have a crystal with three orthogonal translations that is equal in length as you might care to measure, that crystal is not cubic unless there's symmetry in that lattice that demands that the edges of the cell conform to the geometry of a cube. Any other questions?

Let's then in the time that's remaining look at the other symmetry element that could be present in a two-dimensional crystal, and that's the mirror plane. So here's a mirror plane. That's the one remaining symmetry element in two dimensions, and let's ask how we might combine in this space along with the mirror plane a translation.

If I just pop in a translation, and call this T1, and for convenience, I'll take the lattice point on the mirror plane. Here's another lattice point that sits here. The mirror plane acts on everything. It's going to take this lattice point and flip it over to here, and it's going to take the translation that goes from the origin lattice point to this lattice point and give me a T1 prime that sits here.

And now I have two non collinear translations, so these have defined for me a cell that looks like this. This is T1, this is T2, and this is some angle alpha between them. That's a special lattice. This is a lattice which has, just as in the hexagonal or square net, two translations that are identical in magnitude.

Why? Because they're repeated by a mirror plane, and the angle between them-- the angle between T1 and T2-- is completely general. It can be anything it likes. It can close up to a very narrow angle or it can open up to an almost flat angle. That's a new kind of lattice. None of the preceding lattices could make that claim.

Now let me point out that this is for the first time a case where it would be much to our advantage to not choose a cell that contains one lattice point. Let me put some dotted lines in here. And let me submit-- these will also be lattice points-- that I could pick a larger cell that would have this as T1, it would have this as T2, and the angle between them now would be identically 90 degrees because of the geometry that gives us a rhombus here. This would have two translations that are not equal in magnitude, and it would have an angle between these two translations that is identically 90 degrees, but it is no longer a cell that contains one lattice point. It now catches a second lattice point that's in the middle.

So this is a double cell, and it has a rectangular shape. It's a centered rectangular lattice. Being a double cell that's redundant has twice the area that is unique in the pattern, and anything that's hanging up here is going to be hanging down here, so it has a twofold redundancy.

But the thing that you get in return for paying the price of that redundancy is a cell that has a right angle in it. And as we'll see, we're going to use the edges of the unit cell as the basis of a coordinate system for describing what goes on at positions xy within the cell. And the advantages of an orthogonal coordinate system, whenever you can take advantage of it, far outweighs the price of pain-- twice the area to describe the same pattern.

So this is what is generally taken as the standard cell. It's a double cell. But I think you're used to that sort of compromise, because you've all heard of face centered cubic lattices. The primitive cell in a face centered cubic lattice is a rhombohedron. But, oh, that Cartesian coordinate system is so great to use rather than something that has an oblique coordinate system. So is this is a definition of convenience.

Notice the curious duality-- either special relation between the translations angle general or special relation between the translations, an angle special. General translation, special angle, relation between the translation, general angle-- so, it's a curious sort of duality. You can have one but not the other or vice versa.

So this is a fourth sort of lattice. Number one, number two, number three, and now we have number four, which is a centered rectangular lattice. And this particular lattice has T1 not equal to T2 in magnitude, but it has the angle between T1 and T2 exactly 90 degrees, and it's a double cell. It's centered.

Are we done? The reason I asked that silly rhetorical question is that obviously I suspect that we're not done, and what else might we do?

We got our centered rectangular net by starting with a translation-- starting with an mirror plane, really-- and then we added a general translation, and that reflected it across and gave us a diamond shaped net which we could define as a rectangular double shell. Does that always happen? Do I always get that oblique diamond shaped cell? No.

Suppose I put in my first translation deliberately in a fashion such that it was at exactly right angles to the mirror plane. That mirror plane will then reflect the translation and change its direction, and now I have generated a one-dimensional lattice row with translational periodicity T1. And I've got a translation that's exactly perpendicular to the mirror plane.

How do I now make a space lattice, a two-dimensional space lattice? And the answer is very carefully. Suppose I throw in a second translation T2 and the mirror plane reflects it across to here. This interval between lattice points up here is totally in commensurate with the first translation, and that won't work. It's violated my initial choice of the translational periodicity unless I do one of two things, and let's try them both and dispose of this quickly-- this is straight away.

Here's my one-dimensional lattice row. I do not violate this periodicity T1 if I do two things. I could pick T2 so that it fell exactly along the mirror line, and that's going to generate for me a new type of lattice in which this is 90 degrees and these two translations are unequal in length. So this is a lattice that has a rectangular shape. It's a primitive rectangular cell, and it has T1 not equal to T2 in magnitude, and it has T1 and T2 exactly 90 degrees apart.

The second choice that would not result in any contradiction would be to have this as T1, and then pick T2 very carefully so that it spanned the mirror line with one half of T1 exactly up here and one half of T1 exactly here. And this would be compatible with the separation T1 down at the start of this translation. I think you can see that what this is going to give me is the centered rectangular net right back again, so this is the centered rectangular net-- nothing new.

But we did pick up one additional two-dimensional lattice of distinct character. Number 5 is a primitive rectangular network, and it has the characteristics T1 is not equal to T2 in magnitude. Just as in the centered rectangular net, T1 is at exactly 90 degrees to T2, but this is a primitive cell. And that's it.

We really set up the ground rules for the geometry of a periodic two-dimensional space. There are five kinds of rotation axes-- one, two, three, four, and six. Each one requires one or more of the specialized two-dimensional lattices. We have a case where interestingly two different symmetries are compatible with a lattice of the same specialization. In the case of the [? hexagon ?] on that, either three or sixfold symmetry could require that. In the case of the mirror plane we have one symmetry element, M, that can fit into two different kinds of lattices.

So in one case, the same lattice can take two different symmetries. In this case, two different lattices can accommodate the same symmetry.

So that's the story for two dimensions, and we have just one final thing to do. That is to add to the lattices that we have found, and there are five of them, and add to the lattice point the symmetries that we have found require them. And there are a limited number of these-- one, two, three, four, or sixfold rotational symmetry in a mirror plane. And when we have finished these additions, we will end up with a combination of lattice and symmetry, which is something that is called a plane group-- group because the operations that are present in the space follow the requirements of the mathematical entity called a group, plane because this is a distribution of symmetry elements throughout a plane and not just fixed at one particular point.

Let me finish with some general observation, and we will obtain some of these rules later on. Suppose I take a pristine space-- and this blackboard is no longer pristine-- and I put in the space a first operation, and there's a motif in there. This first operation moves around this first motif and gives me a second one, number two.

Then I say let me put in another operation. I'd like to combine these things and see how many different combinations I have, so I put in operation number two. Operation number two will take the second object and repeat it to a third object. Number two is identical to number one, so this gives me a third object reproduced from the second by operation number two.

Now I have a space, and sitting in it are two different operations and three different motifs. Motif number one and motif number three are the same darn thing because they've been repeated by symmetry steps, so there must automatically be some third transformation that is equal to the combined effect of going from one to two and then from two to three-- going from one to three directly.

So this is another truth about these symmetries. Whenever you take two operations and combine them in a space, the net effect of those two operations is equal to a third operation. So a question we're going to ask all along the way until we are at the end of the month-- if you take a translation and combine it with a mirror plane, what new operation has to arise? If you take two rotation operations and put them together, what third net operation has to arrive? If we have some general rules, then we can automatically say, OK, I'm going to take a square lattice and I'm going to put in a fourfold axis. What else is going to pop up elsewhere within the cell? We'll be able to do this systematically but fairly automatically.

So that's where we're going from here. When we're done, we will have derived systematically and rigorously the sorts of symmetries that can combine symmetries such as rotation and reflection with a lattice, and we will know completely the different sorts of patterns that exist around us in two dimensions-- in floor tiles, brick work, wrapping paper, and plaid shirts.

We'll pick up at this point on Thursday. Let me caution you we are going into territory that is not covered in Berger's book. I've just passed it out to you. What we've said in the early parts of the term are in there, but we're going to do things in a slightly different way and then return to his text later on.