Transcript

Lecture 7: Multimedia

Hour 1

(00:00:00)

DAVID MALAN: Welcome back to Computer Science E-1. My name is David Malan. This is Lecture 7, our lecture in "Multimedia." And what better way to start such a lecture than with a snippet of multimedia itself. I turn our attention to the above.

(Video of the Week "Macs vs. PCs" plays onscreen.)

DAVID MALAN (in video): Hello. My name is David Malan, and I'm the instructor for Computer Science E-1: "Understanding Computers and the Internet," at Harvard University's Extension School. You're watching one of our Videos of the Week. For more such videos or information about this course, visit us on the Web, at computerscience1.org. Enjoy the show.

(music in video plays)

(video ends)

(students applaud)

Well, as you probably realize, this was an excerpt from the teaching fellows volumes of Videos of the Week. So if you haven't been following along from home already, do remember that, per the Syllabus, in addition to these videos of lectures and, in the future, of some sections and workshops, each week the teaching fellows have been doing a remarkable job, along with Chris Thayer, a former E-1 student, in producing, scripting, shooting, editing, and delivering these videos to you, the students, via both the Podcast that we've been distributing as well as via the course's Website.

In fact, if you haven't drawn your attention to such already, surf on over sometime this weekend to the course's Website. You'll notice that the front of the Website has recently changed, which is something we'll come back to shortly. But you now have two means of accessing the course's content essentially: the old-fashioned way of the World Wide Web, on the left here, and then the newfangled way on iTunes U, iTunes University, which again, is a topic we'll come back to.

If we do forge ahead to the course's Website here, immediately at left, you will see that not only in CVS style, similarly is E-1's Website now sporting a Thanksgiving theme. How many of you noticed not only the Halloween motif but some things fluttering around the past few days? All right, so good, a little implicit test of who's been checking out the course's Website.

Next time you do, go on do to Videos of the Week, and what you will see is an increasingly large set of these videos. In fact, the "Macs versus PCs" one that you saw was way back from Volume 2. So if that's the first you've heard or seen that, do realize that you're a bit behind on some of this great content. And by no means are these videos required. They haven't been integrated, per se, to the

Problem Sets. But what they are meant to be are supplements to these lectures, and particularly more focused than these lectures are.

Realize of course tonight is about multimedia, a fairly broad topic. Our first two lectures were about hardware. Well, in contrast to those broad topics, these Videos of the Week focus in on very narrow topics in maybe three to ten minutes maximum. And we sort of think of them as bite-sized pieces of knowledge or instruction, where… I can't imagine when any of you feels up to sitting down to two hours of this on your computer or your iPod, but certainly these, delivered by the teaching fellows, is something you can, you know, sneak a peek at during commercials, perhaps right after dinner, or the like.

So do check those out. You can watch them in two different ways. Not only can you download the QuickTime version—and we will define later tonight what QuickTime is—but you can also watch them right in your browser.

And I think you've seen me do this once or twice, or you may have done this yourself once or twice. But if you simply follow the Flash version of the link, what you will actually get is these videos embedded right in your Web browser. And for the most part…

DAVID MALAN (in video): Hello. My name is David Malan, and I'm the instructor for . . .

DAVID MALAN: Enough of that guy. In fact, I do love the "Macs versus PCs" video. It takes sort of a strange turn, I think, toward the end there, but even I was smiling by the end. So that is the work of your teaching fellows and former E-1 student.

Recall, a couple weeks ago we passed around these surveys, which really just ask you for some candid feedback. It was useful, and we did process and absorb that, and act on some of the suggestions, particularly one related to the course's Website.

A theme that came up two or three times in these surveys was that, even though there's a lot of information on the course's Website, there's apparently a real lot of information on the course's Website, such that a few of you, at least, have felt a little overwhelmed perhaps on first visit to the Website. And even though we try to categorize things nicely and alphabetically via the menu at left, and by design nothing on the Website is meant to be more than two clicks away—choose a menu, choose the content; that's been the design—we admit that the menu, if nothing else, has been somewhat long, increasingly long.

So we actually trimmed that down by a few options, merged some things, and generally just eliminated stuff that probably wasn't very much in demand. If you didn't notice a thing—great. If you haven't visited the Website recently, do check it out, especially if you are among those who felt a bit overwhelmed. Should be a little more simple.

And there'll be more content ripped out in the future as we focus in on certain aspects of the course's Website: those strange Google ads that you seem to see at the bottom of the Website sometimes. More on that in a week or two.

So two comments that I did want to relay verbatim, since they were quite memorable, if nothing else. So one of the questions, recall, in these surveys was, “What do you think of lectures specifically?”

Well, this answer was, “I have yet to nod off.” (Malan chuckles.) So we appreciate that. That is the bar that we set for ourselves here.

One other comment was fun, too: “What do you think of lectures specifically?” “I enjoy free food and candy.” Both of them (Malan chuckles) very much about lectures, apparently.

There is no candy today but we can’t have everything all of the time. But all of that stuff is on sale now, I'm sure, at CVS.

So with that said, let’s forge ahead with our lecture on multimedia. And let’s start perhaps with the easiest of questions, or the most obvious of questions, which is, multimedia, what is it?…without looking at the yellow cheat sheets.

Multimedia: What is it?

STUDENT: (inaudible response)

DAVID MALAN: Stuff that’s not text. Okay, I’ll take that. Let’s see if we can elaborate a bit more. That’s fair, though—stuff that’s not text.

What else is multimedia?

STUDENT: Video, audio…

DAVID MALAN: Ooh, oh, an excellent definition, better than ours perhaps on the yellow sheet there.

So it’s a composition of audio, and video, and maybe text, but certainly text in a more animated way, a more dynamic way.

You saw an example moments ago of multimedia in the form of video, which obviously also contained audio. What else specifically might you described as multimedia? What kinds of files, what kinds of content would you say, “Hey, that’s a piece of multimedia"?

(00:10:01)

Yeah?

STUDENT: (inaudible response)

DAVID MALAN: Like a JPEG. So it…

STUDENT: (inaudible response)

DAVID MALAN: Or a sound itself. We'll home in on each of those topics more narrowly in a bit.

When else in this course have we seen examples of multimedia? What comes to mind? Yeah?

STUDENT: (inaudible response)

DAVID MALAN: Videos. Be more specific. What have we seen?

STUDENT: (inaudible response)

DAVID MALAN: Yeah, good, so PowerPoint slides. We can chalk those up as a form of multimedia, especially if you've seen, among the more stylistically designed PowerPoints that perhaps have less content and more swishes, and zooms, and fades, and page tears. You can do a lot with PowerPoint, including audio and video. We tend to use them just for the static content up here.

Anything else?

STUDENT: Skype and Google Earth?

DAVID MALAN: Skype and Google Earth could be said to be a form of multimedia in that it's sort of interactive, it's graphical. Perhaps in the future you'll be able to do even more than that.

But for the most part there's no steadfast definition that we're looking for here. Multimedia, as the name implies: multiple media—audio, video, animations, text, but in more interesting way than a simple Web page, and so forth.

So let's make this more real, especially since you probably see some form of multimedia every day. In fact, at work this morning or this afternoon. Give me an example of a form of multimedia that you had on display in front of you, anything at all. Yeah?

STUDENT: (inaudible response)

DAVID MALAN: Okay, so different Shockwave files—SWF files, Shockwave files. We'll come to an example of one of those in a bit.

What else? What did you have on your screen? What did you do today at work using your computer that you could say, "Ah, you know what? That was an instance of multimedia."

STUDENT: (inaudible response)

DAVID MALAN: You got an Evite, then? Okay, multimedia, sure, because those things sometimes maybe have a little animation to them perhaps.

Well, what about a simple Web page. How many of you, by a quick… an easy show of hands pulled up a Web page today?

All right, so odds are you were not looking at the most boring of Web pages, but there was at least some graphic in that Web page. So why don't we start here, perhaps, with the most obvious, the most common form of multimedia, the graphics that you interact with, say, on the World Wide Web.

What are these graphics… what is the format in which these graphics come?
Pull up a Web page. You see text, you see images, maybe on CNN's site, maybe on Google's site. In what format, so to speak, are those images?

STUDENT: (inaudible response)

DAVID MALAN: Digital photos, and let's be even more technical. I heard "JPEGs, GIFs."

All right, so let's start with those, and see if, by the end of tonight, you can't tell a bit to your family or friends about each of these formats, which, in and of itself, isn't a particularly useful skill, to be able to say, "The GIF format supports 256 colors, including interlaced frames as well as animation." But rather, as you begin to progress and do things more technical yourself—whether it's in this course, playing around with Photoshop, as you will in due time, or developing graphics, or just generally talking about something that you're seeing on the screen, or describing a problem to someone in a sort of tech-support form—well, let's just see how specific you could get in detail.

So a GIF is perhaps one of the most common file formats on the Web.

What do I mean by file format? Well, it's just the sort of definition of what it means to have an extension of .gif. So if you see a file, like David.gif, well, that appears to be, then, a file in the GIF file format.

Just to contrast this with something more familiar, if you have something like resume.doc, DOC, well, the file format in which your resume is in, is what?

STUDENT: (inaudible response)

DAVID MALAN: Document, or more specifically, Microsoft Word's file format.

So when we talk about file formats, we're pretty much talking about some kind of standard arrangement of zeros and ones inside that file. Who defined the standard? Well, usually it's a person, or a group of people, or a company. For instance, Microsoft decided long ago that the .doc file format will be laid out in a certain way. And it takes Microsoft software, or compatible software, to understand that what, at the end of the day, are just zeros and ones arranged on your disk in the form of a .doc file, well, they have some special meaning and you have to read them in a certain order.

Similarly in the world of the multimedia is a GIF just another file, a file on your hard drive comprising lots of zeros and ones—bits—but they're arranged in patterns that sort of define what it means to be a file format.

So let's make this a little less abstract. So a GIF—you can get these details not only off of your cheat sheets but off of the rapid-fire explanation I gave—are interesting in what sense? Tell me something about a GIF.

STUDENT: (inaudible response)

DAVID MALAN: Ah, okay!

STUDENT: (inaudible response)

DAVID MALAN: Close. So there were some key components that you just said that are correct, but not quite. So a GIF is lossless. Let's come back to that in a moment.

But a GIF is just a file format made up of zeros and ones, and it's a graphical file format, to be clear. Right? This is a file format for storing graphics on disk.

Well, pretty much any image, any graphic these days on a computer is laid out as a rectangle, and a rectangle of pixels.

So a typical image might be, for instance, 100 pixels tall and maybe, let's say, 200 pixels wide. Well, what's a pixel? Who's, uh, who's seen a pixel today?

STUDENT: (inaudible response)

DAVID MALAN: It's a measurement. Specifically it's one of those tiny dots on your screen. Ideally you are not aware that your screen is composed of pixels. But if you put your eyes really close tonight… If you haven't noticed this before, put your eyes really close to your laptop or even your desktop's monitor and you should be able to see that your screen is actually made up of hundreds or thousands of tiny little squares—dots, aka, pixels.

In fact, let's just relate this to our hardware lecture. When we have our old-fashioned computer here, with a monitor on top of it, what were some of the common resolutions, did we say, for a monitor?

What's a common resolution for a monitor? Anything come back to you?

Or as of Exam 1, is that part of the semester over? Give me a resolution, anyone at all, a common resolution. What is resolution? Well, that's just the number of pixels across and the number of pixels down.

So what's a common resolution? Well, something like 1024 across by 768 down; something like 800 x 600; something even like 1600 x 1200. For those of you who have tackled the current problem set, and have gone to a computer store, for instance, and have looked at the specs of your LCD display or your CRT monitor, well, you probably saw some mention of the default, or perhaps the maximum resolution possible on that screen.

So all this is saying is that, if your resolution of your monitor is 1024 x 768, when you look at your screen you are literally seeing a 1,024 pixels across times 768 downward. In fact, if you want to put this into perspective or even play around on your own computer at home, on Windows, if you go to your "Display Control Panel"; go to "Settings"—and I may have pulled this long ago—under "Screen Resolution" you see exactly this. I can pull the slider, and the next tick down changes it to 800 x 600.

(phone rings)

Oh! I wonder who that could be?

Hello? Yes, just one moment, please. Oh, it's the pizza guy!

So, no candy tonight, though!

So, if I pull this slider to the left or to the right, that would effectively change my resolution. Unfortunately, this doesn't make for a great demo in class because these projectors are really designed to only work well with one resolution. But if any of you have ever connected a computer at some meeting at work, or some presentation at school, and things aren't quite working right with your monitor, and maybe when you plug it in the thing goes blank, even though you can see things here, and you've hit the requisite keys to make things go there, well, sometimes, though not always, that's simply because your computer is trying to display a resolution that the projector doesn't support. That's one possible scenario.

I'm going to leave things alone because it will mess up our projection if I change it now. But again, the takeaway is that you have a finite amount of screen real estate when you're looking at your monitor. And the relevance, then, to the world of multimedia is that multimedia—specifically graphics and videos—also have resolutions. And if the resolution of your GIF or your movie is bigger than that of your screen, well, obviously that's going to not be a pleasant viewing experience, since you're trying to look at more pixels than you actually have room for on your display.

Moreover, graphics are typically laid out in rectangles. So this graphic might be 200 pixels—dots—across; 100 pixels up and down. But each of those pixels, meanwhile, can be of some different color. In other words, we have sort of a mapping of bits, if you will.

What ultimately comprises a graphic these days is just a whole bunch of dots. And each of those dots has a color. And the number of colors possible for each of these file formats differs. And that's what ultimately defines GIFs and other file formats.

How many possible colors are there for the GIF file format? In other words, each dot in a GIF: how many different colors can it be?

STUDENT: (inaudible response)

(00:20:00)

DAVID MALAN: All right, good. So that was one of the things I rattled off earlier. GIFs are capable of, let's say, 256 colors. They are also capable of animation. And they are also capable of transparency.

Well, what does this mean? Well, have you ever visited on the Web the most hideous of Websites, where little things are dancing around and things are twirling, and it's clearly made by someone who was a bit overzealous when it came to the development of their site?

Well, you might have seen things like—and I'm just Googling right now for "animated GIFs." And I'm pulling up a site that, if all goes well—let's say "food and drinks" category; let's say "barbecue and picnics." And if for some reason you would like a running puppy on your Website, the only file format that you really have at your choice is the GIF file format for the reason that it allows you to have animation.

Now take a guess: How do animated GIFs work? What do you think is going on that gives it that animation? Yeah?

STUDENT: (inaudible response)

DAVID MALAN: Good, so actually it's not a group of JPEGs going together, it's a group of GIFs going together. Each frame has to meet the same requirements.

But you remember, as a kid perhaps, those little flip-books—the faux comic books that, if you flip through them really quickly, you get the illusion of animation, because each page represented like a frame of video. And if you flash them through your eyes quickly enough, you get the illusion of animation. Well, that's precisely how that dog is working. Someone has drawn maybe four or ten separate images. In each one the puppy's legs are in a slightly different position. So when you play them again and again in a loop, you get the illusion of animation.

What do you I mean by "transparency"? Well, it would be nice, if you wanted to put that puppy on to your Web page, for instance—and your Web page has a blue or a green background, representing graphs, or representing, say, the sky—well, things would start to look pretty ugly if the requirement for having animation on your page is that your animation has to be in the form of a rectangle.

In other words, if you had a blue sky and green grass, and to put this puppy on the grass you had to put a white rectangle with the puppy running inside that rectangle—sort of defeats the purpose of having some nice multimedia in the first place.

So what's nice about GIFs—and in your section on Photoshop, this coming week, you'll get to experiment with this, if you so choose—GIFs also support transparency. Which is to say you can specify that a certain part of the rectangle that your image ultimately is, is actually supposed to be transparent. And whatever else is on, say, the Web page should be allowed to shine through.

And this is in contrast with other file formats. What is perhaps the second or the just as popular file format on the World Wide Web?

STUDENT: JPEG.

DAVID MALAN: JPEG. All right, so a number of you rattled that off earlier. What is JPEG good for?

STUDENT: Photographs.

DAVID MALAN: Photographs. So, how many of you have a digital camera, for instance? So, for the most part, when you have exported those photos to your hard drive, and then maybe uploaded them to Flickr, or to Kodakgallery.com, or other such sites, odds are you are uploading, or at some point you're converting your photos to the JPEG file format, the reason being JPEGs support millions of colors.

Contrast that with 256. It's hard to capture the spirit and the color of an arbitrary photo of people, and landscape, and so forth if you only have a pallet of 256 colors at your disposal.

JPEGs, by contrast, then, allow you to have millions of colors. But they don't give you animation. They don't give you transparency. So it's a sort of trade-off.

Yeah?

STUDENT: (inaudible question)

DAVID MALAN: Ah, a good question. So, does making the background of this dog image transparent, does it perhaps affect the background that's behind the dog? If perhaps the dog itself has some white in him, you don't want blue sky or green grass shining through the dog—is that the idea?

So, the way—long story short—that animated GIFs essentially work is you essentially could fill the background of that dog with a random color, like pink, that simply does not appear in the rest of the image. And then essentially you tell the file format, you know what? Everything that's actually pink—don't show pink; make that transparent. So in that way, anything that's not pink is going to be opaque, and nothing is going to be allowed to shine through from the Web page's background. So that's a good question.

And the nice thing—and you'll get to play again with something called Adobe Photoshop, if you attend this coming week's section—is that software makes these processes fairly easy. Making images these days is not like drawing a dot for every pixel—you know, 200 times in one direction; 100 times in the other direction.

Well, back to JPEGs for a moment, and the relevance to today's cameras. Those of you with digital cameras, how good is your camera? Interpret that question as you will. What are the specs of your camera?

STUDENT: (inaudible response)

DAVID MALAN: Ah, so seven-point-something megapixels. All right, how about the rest of you? Are your cameras also rated in megapixels? Okay, so you've all got megapixels—maybe slightly less than seven of them. What does that mean? Why did you buy all those megapixels?

STUDENT: (inaudible response)

DAVID MALAN: Okay, helps with the resolution—or rather it *is* the resolution, in effect. But that's true. Why, in more technical terms, do you want more resolution to your photos? Why do you want more megapixels?

STUDENT: (inaudible response)

DAVID MALAN: Sharper images, more depth, you can enlarge it—that's exactly right. So think of it this way. If you took a photograph of something like the Grand Canyon, and your image, your camera only supported photos of 200 pixels x 100 pixels, you could take the photo and, from a distance you could probably get a sense that this is a picture, albeit not very compelling, of the Grand Canyon.

But suppose then you wanted to print that photo, and you've only got 200 dots this way, 100 dots that way—it's hard to imagine, even intuitively, stretching those 200 dots and those 100 dots to fill, like, a 5 x 7 photo, or even a 4 x 6 photo.

What would happen, in fact, if you took an image that, you know, in real terms is only this big on film, so to speak, and you try to blow it up to this size? What does it look like?

STUDENT: (inaudible response)

DAVID MALAN: It's very blurry or maybe splotchy, where all those square dots are now really big square dots. And that might look fine if you're looking at the photo like this. But you start to look up close and the thing just looks ridiculous.

And think about the analogue, by the way, to these graphical file formats in your daily newspaper. If you pull up the newspaper on Sunday, from the *Boston Globe*, for instance, and look at—here's your excuse to regress for a bit—pull up the Sunday comics, those comic strips often come in color on Sunday. But what do you notice when you look really close at colorful comic strips?

For the most part, comic strips to this day are colored in, in the form of dots. From a distance, looks fine, because your eye doesn't really distinguish that there are dozens or hundreds of dots there. You look at it up close, it doesn't really look so good. And that's sort of an analogue to zooming in on the image and trying to print it at a larger resolution than it was meant to.

So if you have a camera that only supports 200 pixels x 100 pixels, I apologize. You can't buy cameras of this size. When the first digital cameras became popular a few years ago, you'd get maybe a one-megapixel camera. I in fact still have one that's five years old that's 2.1 megapixels. But what does that mean? If my camera is 2.1 megapixels, what does that mean?

STUDENT: (inaudible response)

DAVID MALAN: It means what I just said. So, yes, it means you can print them bigger. You can zoom in. You have better quality. But now relate it to this—like, 2.1 megapixels. Well, that seems to be unidimensional; one-dimensional.

STUDENT: It's the area.

DAVID MALAN: It's the area. So when the Canon companies of the world, and Sony says that this is a 2.1-megapixel camera, that essentially means that, if you multiply the resolution's width by the resolution's height, what you get is roughly 2.1 million.

If instead we use Eric's camera, which is much newer, clearly, than mine, he has seven-some-odd megapixels. That means that he's presumably got way more horizontal pixels, way more vertical pixels. And the beautiful thing about that is that you just have sort of a bigger canvas on which to draw that photograph of the Grand Canyon. You can crop out portions, and if you have enough of these dots, well, even if you just take a snippet of it, maybe that's enough resolution, enough dots to go to press with them from, like, Kodakgallery.com. He can perhaps zoom in on just part of it, too, and crop away parts he's not interested in.

So, in short, if you're in the market for a digital camera, or perhaps finally realize, "Oh, that's why I paid for the more expensive digital camera," what you get is, one, more quality, but also sort of more flexibility, more discretion.

For me, if I want to go to press with my photographs, I pretty much have to frame it the way I want it because I don't have the luxury of opening that photo later on my computer, opening a program called Photoshop—tweaking things, zooming in, cropping out—because I just don't have enough dots, enough paint to play with and to manipulate.

(00:30:02)

Yeah, question in back?

STUDENT: (inaudible question)

DAVID MALAN: When you have more megapixels or a larger resolution, the image is literally larger. There are more pixels going this way and more pixels going that way.

STUDENT: (inaudible question)

DAVID MALAN: No, you're simply getting… So, for instance, suppose I frame Rei, here, with my camera, and I'm using the real cheap camera with just… (chuckles)—Rei, the only guy eating the pizza, mind you, right now. We'll do that in a minute or so. So if I just frame Rei with my really cheap camera—and essentially I'm shooting Rei and, let's say, Dan, and Eugenia here. Well, 200 pixels across do I get; 100 pixels up and down do I get. Well, Rei, let's say, just fills the bottom corner of the photo, the way I framed it. So think of it as, I only now have… (chuckles)—Rei's trying to get out of the photo. So think of this as meaning that I only have—because of the way I've framed my lens—one pixel to devote to Rei. And the other 199 this way and the other 99 this way are being used to capture Dan, and Eugenia, and part of the wall.

Well, that's if I have really low resolution. So if I then look at that photo after taking it, on my computer, Rei is, in effect, going to be represented as one reddish dot, and that's it.

But now suppose that I go with a more expensive camera—same lens, I stand in the same place, I hold the camera in the same way. Now, because I have higher resolution, I get to now squeeze… They're not shrunk, but conceptually I get to squeeze more megapixels in between my left hand and my right hand both horizontally and vertically. Which means now in this bottom corner, rather than just having one dot, you know, I have maybe a few thousand. They're smaller in this context, because I've taken my image and said, spend them this way. But now I have a few hundred or a few thousand dots so that I can capture more of what's going on here with Rei. I can see the dots on his shirt and the remaining slice of pizza on his plate, and so forth.

But even though I'm framing it this way and shrinking them conceptually with my hands, when you go to press, it's like printing a photo that's this big. You're zooming in, effectively, and spending those pixels in a more effective way.

Yeah?

STUDENT: (inaudible question)

DAVID MALAN: Good question. Is it possible, then, to compress the image in certain places more than others? In theory, yes, right? If I'm not that interested in getting Rei in this photo, but I'm trying to be nice in keeping him in frame—but, you know what? Spend most of the pixels on Dan and Eugenia. Let's just spend that one red dot on Rei. In theory, yes. But this is what a file format does for you. The JPEG file format—beyond just being a sort of standard for how the zeros and ones are laid out—it also specifies how you compress an image, what you do to compress an image.

We'll come back to this in a bit. But the short answer is that, yes, that's possible. You can compress different parts of the image in different ways, but that's not something that you, the photographer, tend to have discretion over. That's a function of your camera and its implementation of, say, the JPEG standard.

I don't think I can hold you off any longer. Why don't we take a five-minute break here. There are six cheese pizzas and four pepperoni pizzas, less two slices, in back, along with some plates. Feel free to bring them to your seats and we'll resume in a few minutes.

Hour 2

(00:33:55)

All right, we are back. So lest you be wondering why we just served pizza. This is a thank you, of sorts, from the Harvard Institute for Learning and Retirement. Recall that, this past Sunday, one of our E-1-sponsored workshops invited a number of you, as well as the staff and me, to work with some of the members of Harvard's Institute for Learning and Retirement. And we had a great time helping some of those folks learn how to master the Internet.

Two students among you, Dawne and Ray, were kind enough to join us, as well Chris, from last year. So this pizza is a thanks, of sorts, not only for their participation, but also just from ILR for their appreciation of how the day went. So thank ILR. Maybe, uh, a few years from now, if some of you choose to join them after E-1.

So, just to put this into perspective, here's an example of a shot, not of the Grand Canyon but, let's say, of a nice mountain and some trees. Well, suppose that this image is in the GIF file format, and it's a few pixels wide by a few pixels tall—maybe, eh, maybe like 100 pixels wide by 300 pixels tall, give or take? Well, if this is in the GIF file format, that means that it is a bitmapped file format, which, as we said before, this image is made up of tiny little dots called, again, what?

STUDENT: (inaudible response)

DAVID MALAN: Called pixels.

The problem, though, if you represent an image, like this photograph here, using a GIF file format, or even the JPEG file format—because the JPEG file format, too, is a bitmapped file format… It's good for photos because each dot, each pixel in a JPEG can be any number of millions of colors, which means you can get much nicer gradients; whereas, GIF only has 256 to work with. But both file formats are bitmapped, which means they're composed of pixels.

The problem, though, to put this into perspective, then, is if you try to zoom in on a bitmapped file format—be it GIF or JPEG—what you get is the effect there. What you see on the slide above you is a zooming in of the top of one of the peaks of that mountain. Because this image is composed of dots, or pixels, and you only have a finite number of them, you can't zoom in forever. Right?

You know how they, in *Law and Order*, and *CSI*, when the security camera has taken photographs of, like, the bad guy fleeing the scene, or of the license plate and, you know, one of the cops will say, "Can you enhance that? Can you zoom in on that?" No! Like, you cannot just zoom in forever just to improve the quality of the license plate that was captured across the street from a camera dozens of meters away! Right? If a camera is, by definition, just storing its photos in a bitmapped file format, which pretty much they do, whether it's GIF or JPEG or something proprietary, for CCTV cameras, at the end of the day, you only have so much information there. And if it's not there to begin with, you can't zoom in and capture what information you would like to see.

So the next time you watch *Law and Order*, *CSI*, and the computer tech in forensics says, "Could you just enhance that for me?" Like, no, that is not possible to do ad infinitum! What you eventually get is a license plate that looks like that. Or there, for instance, is your suspect. That tends to be what happens in reality. Yeah?

STUDENT: (inaudible response)

DAVID MALAN: (chuckles) They claim that, yes! So, yes, it is certainly true that with software and with intelligent software you can interpolate information And you can kind of smooth out an image and maybe get a better sense of what the license plate is saying by sort of smoothing things out so it's not as blotchy as this. But at the end of the day, if you only have a fixed amount of information to start with, you can't come up with new information. And in effect this is what these programs are doing, and it's an interesting question for court. Because if you have computer-enhanced images, computer-enhanced is another way of saying the computer added information, or changed the information that you were given. And while that might help in a practical sense—giving you the first two digits of the license plates, and that might be admissible, for instance—as an I.D. of the suspect, probably couldn't or shouldn't hold up if things are sufficiently blurry in the first place, or a good defense attorney could pick holes in something like that.

But the takeaway for our purposes tonight is not so much why *Law and Order*, one of my favorite shows, and *CSI*, aren't so realistic, but rather to at least relate it to, again, those images you've been seeing on your Web page.

Well, if most of these file formats that you see every day—GIFs and JPEGs—are bitmapped, well, what's the alternative? Are there file formats in which you *can* zoom in ad infinitum and get just a sharper and sharper image, even of just a small piece of that image?

Well, yeah, there exist graphical file formats that are based not on bitmaps—you know, horizontal rows and columns of dots—but rather vectors. Now, this is just a way of saying that you can implement graphics using mathematical formulae. In other words, rather than representing the sun as a circle of dots, literally drawn in a bitmapped file format, like this (draws circle of dots on the board), well, wouldn't it be better to take advantage of some of the algebra or geometry you might remember from high school, and say, "Well, actually I know that the formula for a circle is $x^2 + y^2 = r^2$, the radius."

Well, wouldn't it be nice if, in your graphic, you don't represent circles as dots. You literally represent them with formulas as circles. Well, what does this imply if you desire to zoom in now on that sun? Well, if you're representing a circle with a formula, all you have to do is plug in bigger numbers to that formula and you'll get a bigger but still perfect circle.

By contrast, if you try to zoom in on a sun comprised of dots, what you start to get is what we saw a moment ago—you get blotchiness.

So what you start to get is what we saw a moment ago—you get blotchiness.

So what you have here, and as an example on this second slide is a wire frame of sorts of some blocks of cheese.

(40:00:03)

Well, there's no mathematical formula, to my knowledge, for a block of cheese, but you could certainly break it down into things like circles and triangles, or fragments thereof. So what this wire frame is suggesting, that, rather than representing this block of cheese, if it's in a vector-based file format, with just dots; rather, let's break each of the components of this picture, each of the shapes within the picture down into some polygons—circles, squares, rectangles, straight lines, and so forth—and embed that information in our file format. So if you want to zoom in, all you do is scale things, and you scale them in a way that the quality in effect is preserved.

Let's take a look at that in this example. A gentleman earlier mentioned the Shockwave file format—a file format ending in .swf. This is a file format for animations, essentially, which means it's not just graphical, necessarily. It can also have audio embedded in it. But it's a file format in which an example like this is stored. So here is… it's already a little strange. Here is a bunch of Swedish horses, okay? So maybe you've even received this as an Internet forward. Yes? No? So right now it's very small, but arguably, it's of good quality, right? It's not blotchy and it looks like it's maybe, what, 200 pixels x 100 pixels?

And where am I getting these numbers from? Well, what's my screen resolution now, if you recall from before? It's 1024 x 768, which daresay is the most common resolution today. Though that's debatable as hardware gets better and screens get larger. So if I have 1024 across, well, it looks like it takes up maybe a third of the screen. So maybe it's 300, 400 pixels wide. That's where I'm coming up with these numbers. Well, it looks pretty sharp.

Suppose then I want to zoom in. It would be unfortunate if these cute little horses all of a sudden become blotchy and more squarelike, like our mountaintop does. But it turns out that they don't. In fact, this probably looks even sharper now because we're really spreading out how the dots are being spent on the screen, or how they're being computed. If I zoom in even further, making this full-screen, it's a pretty sharp image. In short, even if I had a bigger monitor, I could keep scaling this again and again and again, and it would still—the lines, the curves, the colors would still look just as sharp because these horses, and everything around them is being represented somehow or other mathematically, not with just dots.

Now this, in and of itself, might be cute. Wait till you see this! How many of you have seen this? One, two?

(first animated horse sings)

Turns out this is an a cappella quartet.

So, to put this into technical context, Shockwave Flash . . . Shockwave is an animation file format. Again, it supports audio and it also supports interaction.

So if you've ever been to a particularly sexy Website, which doesn't just have graphics and text, but rather, if you click something, something swooshes away and maybe something zooms in… If you visit a lot of fashion sites, like Banana Republic or Gap, or if you visit a lot of . . . Furniture stores that tend to have Websites based in Shockwave Flash, or similar technologies. They're interactive.

So I clicked on that horse to get him singing. And I'm going to try to time this correctly.

(second animated horse joins in)

Ah, I got it wrong.

(second animated horse joins in again)

That's not bad. Okay?

(third animated horse joins in)

That's wrong. Try again.

(third animated horse joins in again)

Not bad so far. I think we have to add the soprano, perhaps?

(fourth animated horse joins in)

It's so simple and yet it's funny. That's it, that's all this thing does.

(students laughs as horses continue singing)

So, anyhow. We'll make a link. This will perhaps be one of the most popular links visited after this lecture.

(singing ends)

But that is an example of a file format that doesn't just use bitmaps. It actually uses vectors. And by vectors we mean mathematical formulae. Mathematics, in some sense, such that the visual upsides of that usage are quite clear, or hopefully more clear from an example like that.

Well, where else do mathematics come into play? Well, this isn't a graphics class, and so we won't go into detail as to how the latest and greatest video games are made. But certainly, if you've looked over your kid's shoulder at the latest PlayStation game, or GameCube game, or even PC game these days, these graphics are far beyond the Pole Position, and Asteroids, and Donkey Kong that at least I grew up with, which were very much more bitmapped types of games.

Well, these are three figurines here that perhaps just capture the spirit of how some of today's games are implemented, whereby characters are not just implemented, like Mario and Luigi were in the first Nintendo—as just a bitmap of colors that collectively represented Mario's face, and mustache, and hat, and clothes—but rather they're represented more clearly as a collection of polygons, of mathematical shapes.

The beauty of that is that, if you want to implement a game, where the player can go anywhere on the screen and can do different things with his arms and legs, and you want to have the actual illusion of movement, well, it would not be ideal if you had to, just to get your player to do this (stands with fists held up), store a graphic, using pixels, that precisely depicts the player in this way.

If instead you wanted your player to be able to lift his hand slightly, you don't want to have to recreate an entire image—a GIF, really, or a JPEG—just to capture that slight bit of movement. In short, you don't want to create a sort of flipbook of all possible permutations of characters on the screen, and physical locations because, if nothing else, you would need a DVD or more—multiple DVDs to store all of those darn images, certainly at the resolution that people expect on today's games.

A much more efficient and faster way of storing information, then, tends to be using mathematics of some sort—linear algebra, and other things that you've probably happily forgotten from years ago—but then, at the end of the day, help you model things in a way that's much more scalable. And so if you want your character to move into some direction, that essentially means that your video card or your CPU simply has to perform some mathematics and create on the fly the representation of that character in his new position.

In short, a lot of today's craziest video games are developed in this much more dynamic, rather than static, way. It's fun, perhaps, rather than playing one of today's latest and greatest games, which you yourselves might own, what I thought I would do is draw our attention to the course's Website. There's another underpublicized link there. Perhaps you have stumbled across it late at night when looking for something better to do. But if you click on the course's Website the link entitled "Games," what you have here is a little archive of free games that are actually implemented in, I think, the Shockwave Flash file format, which was developed by Macromedia; is now owned by Adobe. And you have some examples of games here from yesteryear. In fact, we seem to have Asteroids up top; Donkey Kong, if you're familiar; Frogger, Moon Patrol, Pac-Man, Pong—it doesn't get any simpler than Pong, really—Space Invaders, and Tetris, which is a little newer.

From the audience, which was your favorite game from yesteryear?

STUDENT: (inaudible response)

DAVID MALAN: Tetris. Space Invaders. I heard Frogger is the favorite game. Who said Frogger? Can I entice either of you to come down for a moment and show us just how good you are at Frogger?

STUDENT: (inaudible response)

DAVID MALAN: That's fine. Would you like to…

STUDENT: I don't remember…

DAVID MALAN: Remember! Up, down, left, right!

(students chuckle)

There aren't twenty buttons in this game, like today. Can I entice anyone to come down here and show us how good you are at Frogger? Yes? No? Come on! All right, come on down. What's your name?

STUDENT: Heather.

DAVID MALAN: Heather. So Heather… Oh, I guess… Do we get to type? Heath… (tries to type "Heather.") Oh, sorry, yesterday's games only supported six-character names.

(music from Frogger starts playing)

Okay, so you're this frog here. And you can use the arrow keys, I think, to move up, down, left, right. Have you placed this before?

Heather: Yes.

DAVID MALAN: Oh, all right!

(sound effects and music from Frogger continue)

(students applaud)

DAVID MALAN: Keep going.

(sound effects and music from Frogger continue)

(sound effect signaling defeat plays)

(students sigh)

DAVID MALAN: Oh! Try once more! Redeem yourself. Make us proud.

(sound effects and music from Frogger continue)

(sound effect signaling defeat plays)

DAVID MALAN: Once more?

Extra life! Once more.

(sound effects and music from Frogger continue)

(sound effect signaling defeat plays)

DAVID MALAN: All right, once more. Let's end it with a positive note.

(students chuckle)

No, that's fine. You're just getting worse, it seems. (chuckles)

All right.

(sound effects and music from Frogger continues, then ends)

(00:50:05)

DAVID MALAN: All right!

(students applaud)

DAVID MALAN: Very well done. Congrats!

All right, so there's plenty of other games there that you can play with later. But what are some of the takeaways? All right, can we spin this as an academic exercise?

So what was interesting about that game, in the context of all this stuff?

STUDENT: (inaudible response)

DAVID MALAN: All right, it's definitely bitmapped, right? I made the thing so big on my screen that you could see that everything was the result of putting little squares of color together—green, for the case of our frog.

What else did you notice?

STUDENT: (inaudible response)

DAVID MALAN: So very flat, two-dimensional, right? There's animation, but when Heather was hitting the Up key, that whole GIF, or whatever it was back in the day, was just moving entirely up, you know, a few pixels. It wasn't gradually sliding. It's certainly much easier and much more efficient to just have these things jump across the screen. And fortunately the frog is moving relatively short distances so it looks like he's just taking one fluid step; when, in reality, he's hopping from one step to the other, and there's nothing in the middle.

Anything else come to mind?

Well, this is good that we had a note of video there. Let me ask Dan to come up and, as we proceed, to connect his Macintosh, so we have the duality, as we promised in our introductory video earlier.

For a moment, though, there was a question from before about compression, and I said we'll come back to that later.

But the nice thing about these graphical file formats—both GIF and JPEG—is that they are "compressed." Well, what does that mean? Well, what does it mean to compress a file?

STUDENT: (inaudible response)

DAVID MALAN: Make it… yeah, to make it smaller. So what are some of the considerations there, right? Like, I could make a file really small. I could make your resume really small by Control-A and then Delete. Is that compression?

STUDENT: No.

DAVID MALAN: All right, well, I mean, it is, but it's what we would call "lossy compression." Lossy—and this a technical term—lossy in the sense that, yeah, you're making the file smaller, but by throwing away information. So it's sort of a corner case in that, yeah, you can make the file smaller if you get rid of the file's content all together. But that's sort of an extreme case of what is in fact done today.

Well, what, rather, do you want to keep in mind? When you compress a file, what are your concerns, then?

STUDENT: (inaudible response)

DAVID MALAN: Ideally you want to keep all of the information. Unfortunately, that seems like a Catch-22. You want to use less information—fewer bits—to represent the same amount of information. That is to say, if you don't want to lose information, you've got to keep it around.

Well, if you want to keep all of your information, what could you possibly throw away becomes the question.

Well… (aside to Dan) and actually, if you want to sit for a moment, rather than stand there awkwardly, we'll finish this.

Don't worry. The Mac will be back in a moment.

Um, so how can you do it? Well, consider this. Here is a flag of Germany and a flag of France. Suppose that I store both of these images as GIFs, which means, again, they can only have 256 colors, maximally. That's fine. We only need three for each of them, so that's not a constraint.

They can only be bitmapped, therefore. So essentially these flags, though you can't quite see it, are made up of a bunch of little dots: left, right; top, down. And suppose that the file on top is, um… suppose that it's, let's say, 300 pixels x 200 pixels. Well, if it's 300 pixels x 200 pixels, then what do we have?

So 300 x 200. What's the total resolution, or how many pixels are in this flag all together?

All right, so six… one, two, three, four. So 60,000, right? It's kind of impressive that you go from 300 across, 200 down and you get 60,000 pixels, but that's how many little dots are making up that, say, German flag and France flag.

So the question on the table is, can you represent these flags with the same aesthetic effect but using fewer than 60,000 pixels? So to be clear, each of these flags is made up of 300 dots horizontally, 200 dots vertically. Each of those dots in turn has a color.

For instance, in Germany's flag, all of the pixels up top—the top third—have the color black. Okay, all the colors in the middle—red; and then yellow. So you need to remember the size of this thing, the shape of this thing, the colors being used. So what can you throw away? Can you use fewer than 60,000 pixels to represent Germany's flag?

STUDENT: (inaudible response)

DAVID MALAN: True. Yes is the answer. You get the harder one. How?

STUDENT: (inaudible response)

DAVID MALAN: Okay, so a 3 x 3 image—that would distort…

STUDENT: (inaudible response)

DAVID MALAN: 3 x 6 would distort the dimensions, too.

STUDENT: (inaudible question)

DAVID MALAN: So can you make it smaller? You could certainly in this case. What do you mean by smaller?

STUDENT: (inaudible response)

DAVID MALAN: Sure. So you could certainly just make the image smaller so that therefore you need fewer bits. But I would say you're essentially telling me to delete parts of your resume, right? Make the file size smaller by making your flag appear smaller.

But my goal is to put this flag on Germany's home page. Right? www.de is where I want to put this thing. So I want it that size. I want it really big. I can't cheat and save space by just making the flag smaller.

You have a question in back, or a suggestion.

STUDENT: (inaudible response)

DAVID MALAN: Okay, good, so a slightly more mathematical approach. In short—and I'll tweak this just to lead us in the factual direction. So keep around information for some of the pixels—for instance, the left-most column of pixels. So you've got, you know, 200 pixels on the vertical there. So that means you've got, like, 66 black pixels, beneath which are 66 red pixels, beneath which are 66 yellow pixels, and then plus something else, right? That gives us 200.

So what if we just stored that column, and then said, you know what? Multiply this column by 300. Well, would that do the trick?

Well, if you think about it intuitively, there's a lot of redundant information in that flag. What is redundant about the flag?

STUDENT: (inaudible response)

DAVID MALAN: It's just solid bands of color. But in effect, with GIF or JPEG, as we've defined it thus far, you're essentially saying, if you're storing it as a bitmap, "Make this pixel black. Make this pixel black. Make this pixel black," and you're doing this 300 times. Then you're going back here, stepping down this way once and doing it 300 more times, eventually saying, "This one is red. This one is red." Then eventually you get down to the yellow: "This one is yellow." I mean…

STUDENT: (inaudible response)

DAVID MALAN: Good, that's essentially the answer. You retain as much information as you need and then you multiply it by some factor. And this in fact is how GIF works. GIF—in contrast with our Microsoft Word proposal, earlier—is a lossless compressed file format. Lossless, as that word implies, when you compress it, you do make the file size smaller, but the upside is you do not take away from the quality of that image at all.

Well, how is that even possible? Well, in effect, all the GIF file format needs to remember for Germany's flag is that the leftmost column is 66 black pixels, 66 red, and 66 yellow, then one extra pixel left over. Then multiply those by 300.

So now in effect I only have to say this pixel is black 66 times. I only have to say that this pixel is red 66 times, and then the same thing for yellow. And then I, in effect, just shout out "repeat 300 times." And you can say that, just as I could verbally so much more efficiently then you could by just repeating yourself 300 times by 200 times.

Contrast that now with the flag of France, at right. Taking on faith that the GIF file format pretty much works as we just described—defines as many pixels as are necessary vertically and then says "repeat"—does France's flag compress using GIF better or worse than Germany's flag?

STUDENT: (inaudible response)

DAVID MALAN: Okay, you say worse. Why?

STUDENT: (inaudible response)

DAVID MALAN: Takes 300 instead of 200 what?

STUDENT: Pixels. (inaudible response)

DAVID MALAN: Right. Because the colors are not contiguous throughout, you essentially have to say, "This pixel is blue. Repeat 100 times. This pixel is white. Repeat 100 times. This pixel is red. Repeat 100 times."

Contrast that now with the story over here. This top left pixel is black. Repeat 300 times. You're done with that row, saying a third as many sentences. And we can confirm as much in tonight's GIF directory, which is linked on the course's Website.

(01:00:00)

Here are precisely these images. If I go to Details mode, notice how large is the France flag, frLarge?

STUDENT: (inaudible response)

DAVID MALAN: Five kilobytes. It's a little small on screen. How big is Germany's flag?

STUDENT: (inaudible response)

DAVID MALAN: Three kilobytes. So we have a bit of empirical evidence here, too, that the mathematics do in fact corroborate the conceptual explanation we gave a second ago. Yeah?

STUDENT: (inaudible response)

DAVID MALAN: Correct, the GIF file format is designed to read this left-right manner.

Let me pause for a moment. Let's sit still, unless you'd like to grab a piece of pizza and we'll swap out tapes. But let me take a question or two offline now.

Hour 2

DAVID MALAN: "What about PNGs?" is an excellent question.

So we've looked at GIFs, JPEGs. Also on your yellow cheat sheet and also on the World Wide Web is this third file format that's not nearly as popular as these two but, in a sense, is better. And it is called the Portable Network Graphics format. PNG is its file extension.

This too, like JPEG, supports millions of colors, but it is lossless. So think of this sort of an amalgam between what GIF and JPEG offer. Moreover, PNG is not encumbered by various patent issues. The company that invented the GIF file format for a while stomped their feet about people using it without paying royalties to them to output these GIFs. So PNG was partly a response to that.

That particular issue is not so hot anymore. So the short of it is that GIFs and JPEGs remain the most popular file formats today. And if you would like even more technical detail on any of these three, actually a really good resource would be just the computer dictionary linked on the course's Website. Type in "P-N-G," type in "G-I-F," type in "J-P-E-G," and you'll get back some more technical detail from the Webopedia site as well.

Well, how else might you compress information, or what are some of the other considerations?

How about this apple? This apple is on a blue background at left. How much of that information do you need to keep around? Well, just to reiterate what we considered verbally with the flags, you pretty much only have to keep around as much information as is currently in the right-hand version of this image, right? If you're saving this as GIF, you've got to remember a column of blue, and then, in effect—and the details of this are not so important for our purposes tonight—but in effect, you can say, you know what? Anything that's just white: fill it in with blue. And so with a graphical file format can you use tricks like that as well.

Are we losing any information if the right-hand version doesn't represent the visual result of compression, but this is depicting how much information we're physically storing in the file? Can you reconstruct, using that much information at right, the original image, is the question.

And the short answer, using the heuristics we've discussed: Yes, because we've remembered just enough. Especially in this context, if we advance to the level where we can say, you know what? "Don't just repeat horizontally. Go fill in anything that's white with blue," that would be sort of a neat trick to employ as well.

Well, graphics are one thing. We've seen a tad of audio. We'll come back to that. But video is certainly in vogue these days. And as bandwidth speeds increase and people have faster connections to the Internet, and as disk storage gets even cheaper, you see sites like Google Video and YouTube, the last of which we pulled up an example of a couple of weeks ago.

What video did we watch on YouTube?

STUDENT: (inaudible response)

DAVID MALAN: Yeah, it was that hard drive, where some guy had opened up his hard drive but still left it in operable condition, and deleted all of those files. Well, that was an example of YouTube. YouTube uses a file format callled Flash video to store all of their videos. What this means for you, the user, is that the wonderful thing about YouTube and, say, Google Video today, and even E-1's Website—if, like before I did, you clicked the Flash version of our Videos of the Week, or the lectures—you don't need special software. Because where did that Video of the Week start playing?

STUDENT: (inaudible response)

DAVID MALAN: Right on the course's Web page, right? There was no prompt to say, "Do you want to open this or save this?" There was no new window that came up. It was just the course's Website with the video embedded therein.

Now that was a bit of a white lie. You do need special software to play Flash videos. You do need special software to play Frogger, as we did before. But according to a company called Adobe, 97.2 percent of computers on the Internet today are thought, according to their marketing folks, to have what's called the Flash player installed already.

What you've seen happening these days is, when you buy a Dell, or you download software from Yahoo! or Google, a lot of these companies tend to embed other people's software into your computer for you, the result of which is usually kickbacks of some sort or some financial gain for them. But the nice thing about this—because the browsers have done this as well—most of today's browsers, if you install IE or you install Firefox, they too come with support for Flash. That is to say, they come with the "plug-in" that allows you to play those videos inside of them.

So back in the day, when this was not as popular, you would instead, upon visiting the course's Website—if you don't have the Flash plug-in installed, you'd get like a broken icon, or you'd get a prompt, saying, "You need this plug-in. Do you want to install it?"

These days we've reached a point where so much stuff do you get for free, automatically, when you buy your computer or download some other program, that you have the illusion of everything just working seamlessly.

But contrast that with, for instance, QuickTime. QuickTime is the file format in which we played "Dan's Soapbox," at the start of tonight's lecture. And if you recall, we did need a special player for that. A new window came up and that was the QuickTime Player, made by Apple.

QuickTime Player, by contrast to Flash, does not come on every Windows PC these days. Rather, if you have wanted to download the course's Podcasts, you may have in fact had to yourselves go to apple.com and download iTunes, or QuickTime. Funny thing is, there, even Apple does it. You download iTunes, guess what you get for free?

STUDENT: QuickTime.

DAVID MALAN: QuickTime, right? So good and bad things about that. But for practical purposes it tends to be a useful thing that this stuff just seems to work.

So why don't we motivate this foray into video by pulling up a video. I know what's about to happen, so I don't want to be the guy on video, ironically. But I need another volunteer who wouldn't mind being not only on video but also on video, and ergo, video.

(students chuckle)

DAVID MALAN: Come on. Anyone at all. All right, we go through this every week, and then there's always someone that volunteers. Come on!

It's a nonspeaking role.

STUDENT: I will.

DAVID MALAN: All right, come on down, Dawne. All right. Dawne is going to stand in front of Dan's MacBook Pro, here—a nice fancy Macintosh laptop. You will see the effects of what we're doing, but we will see them even bigger up here.

I'm going to have Dan use his Mac savvy to touch the right buttons. This is obviously a Web camera of sorts built into Dan's laptop. You might use this for video conferencing, like we did with—albeit in one direction—with Victor, of the "Typical PC User Podcast," a while back.

Dan is going to use it as an example not only of just multimedia and video, but what kind of software exists today and how fast processors are today such that they can render the images of Dawne you're about to see in real time.

There was a time where all of there morphing examples . . . Does anyone remember the Michael Jackson video from five, ten years ago, where each of the actors in the video morphed—and by that I mean sort of blended into one another? That, as I understood it, took weeks if not months to develop on computers because it took so damn long for the computers to generate all of the intermediate pictures.

What you're about to see is what was now possible in 2006.

I give you Dawne.

(students chuckle)

DAN ARMENDARIZ: So this is good for anyone who didn't happen to go Trick or Treating, I guess. Uh, let's see, here's the fun one. So, like I said, this is good for anyone who didn't get to go Trick or Treating.

(Students chuckles.)

DAN ARMENDARIZ: So just... here, just move around. You can move make some great faces. Okay, right there. Hold it. Move to the left just a touch. Okay, let's see.

(Dawne laughs)

All right, wait. Wait for it. Wait, wait. Oh, there you go. Okay, it takes a photo.

Dawne: Oh, that's lovely!

(students laugh)

DAN ARMENDARIZ: Okay. Let's see what else we've got here. We have... oh, uh, Fisheye.

Dawne: Oh!

DAN ARMENDARIZ: Yeah, if you get really close you can make your nose quite large.

Dawne: Oh, yeah! Yeah, let's do that.

DAN ARMENDARIZ: There's probably a better one. Oh, here's a good one, here's a good one: Bulge.

(David Malan chuckles)

DAN ARMENDARIZ: So one more photo, I guess.

Dawne: Okay.

DAVID MALAN: We'll put these on the course's Website after, perhaps?

Dawne: Yeah, thanks for that! (chuckles)

DAN ARMENDARIZ: We do actually have one of David here that he so kindly took for us.

(students laugh)

I think he was going for the Jay Leno look, but I'm not quite sure.

DAVID MALAN: Uh, this is... We'll get to privacy and security in a couple of weeks. I had no idea this photo was being taken when we were experimenting earlier.

But I guess in fairness we'll put that on the course's Website.

DAN ARMENDARIZ: Yup.

STUDENT: (inaudible question)

DAN ARMENDARIZ: So the question was, if you can change an artwork in the same that you're changing these photos. And basically all there is, it's just a camera right here. So anything that you can hold up in front of the camera will change with these different effects.

So there were color effects in this first... uh, where was it? There was color effects in this first section here. So if we held up, for example, David's notes that we have, we could just hold it up and change the coloring. Or we could go into the Twisting effect, and the various other things, and we can actually see, in real time, the different effects that we have.

(01:10:15)

It's not really the highest quality. Basically what you see for the resolution—this is about 640 x 480. So it's just probably a shade over one megapixel, the ability of the camera.

DAVID MALAN: (chuckling) You know, this is really disturbing. The woman on the overhead is...

DAN ARMENDARIZ: Okay, here, we'll make it a little bit better. Normal, okay.

DAWNE: Okay.

STUDENT: (inaudible response)

DAN ARMENDARIZ: This program comes with new Macs that have the WebCam built in. This includes the laptops and some of the desktops as well. It's in the applications.

STUDENT: (inaudible response)

DAN ARMENDARIZ: No, you cannot layer effects; just one at a time.

DAVID MALAN: Other questions?

Well, a big round of applause for our guinea pig here.

(students applaud)

Thank you.

So on this same note, you'll notice that distributed tonight is Problem Set 5, "Multimedia." This is a particularly fun problem set in that it's very hands-on and graphically . . . Where is this?

DAN ARMENDARIZ: You didn't switch it.

DAVID MALAN: Oh. (chuckles.) And, um... graphically oriented.

Among the tasks ahead of you are the following. On top of the course's Website, as of yesterday, is a big turkey. And you may have noticed for quite a while, if you click on the banner at top left on the course's Web page, it actually changes automatically for you.

Now we have—gobble, gobble—two turkeys. If we click it again, we have a dog dressed as a Native American and a cat dressed as a Pilgrim.

So I wish I could say this was our handiwork. But these are actually images developed by last year's student body in E-1. Because if you've skimmed the problem set being distributed tonight, the first challenge of this problem set is to task you with learning a bit, or applying a bit of knowledge about

Adobe Photoshop, which, though we try to be software- and platform-independent in this course, hands down, the de facto standard for graphics design these days is Adobe Photoshop, which is a very expensive program, normally, but for which there exists a free trial that you can download off of Adobe's Website.

We have a link on our own course's Website, under "Software" and under the "Multimedia" category. And it works perfectly for thirty days, after which they try to coax you to buying it. But with your Harvard I.D. numbers you can actually use Harvard's version of the software, too, so long as you're registered in or affiliated with the university, and so long as you have a network connection to the campus.

This week's section, if you are the sort who does not like dabbling on your own, or feels a little overwhelmed by pulling up a new piece of software, a la Google Earth, and learning yourself how to manipulate images, this week's coming sections will focus in particular on designing GIFs, JPEGs, and PNGs, using Adobe Photoshop. And these will really be hands-on tutorials, hand-holding tutorials on using software like Adobe Photoshop. And this problem set will invite you to apply that newfound knowledge to the design of your own very banner.

And, in fact, what we will do, after the images are all submitted as part of the problem set, is your own work will be exhibited for a number of weeks on the course's Website. And if you get tired of looking at someone else's image, right?, you can always click it and change it to your own.

And if you don't click the image up top, it actually changes automatically, using something called "Cookies," which we'll come to in our Website Development lecture, to actually remember which one you saw yesterday so that you don't show the same one again.

Moreover, as extra credit in this problem set, you'll notice specifications, toward the back, of the shape of a mousepad. So each semester we have a contest of sorts, whereby, for extra credit, any number of students are welcome to submit candidate designs for this semester's fall 2006 mousepad.

So long as the mousepad somehow embodies the theme of survival of Computer Science E-1, per the specification, we will then enter you in a voting process that we'll have in a few weeks. After this problem set is submitted, you, the audience, will vote on your favorite design. And at the very last lecture, you—as well as our distance students, who we'll take good care of via postal mail—will walk home with a souvenir of sorts from E-1.

As if the candy and pizza weren't enough, we'll give you something nonedible to remember the course by. And that will be the mousepad that wins the popular vote. And that will be coming up in the form of this Problem Set 5.

Let me also draw your attention to the blue handout tonight, which is the Final Project, which, per the syllabus, challenges you to develop your very own presence on the World Wide Web. Funny we should hand this out now, given that you have yet to see your Website Development lecture, or content thereof.

That's okay. The first task in the Final Project, this so-called Part 1, tasks you with thinking about what kind of Website you want to design. And fortunately, the only prerequisite for answering this part of the problem is: What sites have you visited before? What do you like? What would you personally like to do?

Literally at the end of this course, you will have your own presence on the Web, something of the form www.davidmalan.com, but with your name inserted. Or you can have it be something completely random, like www.myhomepage.com. I would wager that both of those, actually, are already taken.

And as we mention in this part of the project, perhaps daresay the hardest part of this project, frankly—because the rest is all fun. This part is the frustrating part: finding a domain name that someone or some squatter hasn't already bought. But we don't constrain you to only choosing a dot-com. You can choose a dot-net, dot-org. Heck, you can have, you know, danarmendariz.tv, if you wish. I would wager that that one is still available as of now.

With that said, though, there will be... the cost of buying a domain name for the period of a year is like ten or so dollars, and it's all spelled out there. And we would ask that you take care of that part of it. There are alternatives, if you'd rather not incur the expense. But we, E-1, the course, will host your Website for you up through spring semester.

So we will provide you with an account. We will provide you with an email address of the form, you know, username@davidmalan.com, but substituting your domain name, of course.

So, literally, at course's end and through May, you will have your own presence on the Web. We'll provide you with the disk space. You'll get a gigabyte of storage, which is very useful if you want to have, say, a Website of photographs of your recent trip. Right? Gone are the days where people come over and watch slide shows of their latest trip. You can just put that on the Web and email all of your friends those photos. This would be a wonderful application for the Final Project to do something like that. Really, the sky is the limit. But more on that in technical detail in our next lecture.

The first part of this project, which is due at the end of November, invites you to start thinking about your Website and also to choose your domain name, because we'll need a week or so to set up on our end all of your Website domains so that you have some place to put all of your content when you start writing HTML yourself.

So any questions on that, just let us know.

In the meantime, any questions on this or another thing? Yes? No? Yes? No?

All right, so let's go back to where we began. If we pull up the very first clip from tonight, this was this so-called QuickTime video. QuickTime is a video file format that was developed by Apple. If you have a Macintosh these days, you already have support for QuickTime installed, which is to say, if you visit some Website and they say, "Hey, download this video; it's in QuickTime format," all you have to do is download that video, double-click it, and it will start playing.

If by contrast you're on a PC, odds are you first have to go to apple.com, download the software, install it, and then you can double-click on the .mov file, which is the file extension for QuickTime movies, and then it too will play.

This software here that I just pulled up is the so-called QuickTime Player. It's the software that Apple developed, and if you want to play QuickTime movies, you need this software, or compatible software.

When I click "play," what it starts playing is that Video of the Week that we began tonight doing.

But rather than play this again, I'm going to go up to its Window menu, and choose "Show movie info." And there's a bit of interesting information here. It's a little small for you to see. But is it big enough to tell me what the resolution of this video is?

What's the resolution of this video?

STUDENT: (inaudible response)

DAVID MALAN: Five... nine... current size—that's the current size. That's not its native resolution. What is its built-in resolution?

STUDENT: It's 320.

DAVID MALAN: It's the 320 x 240. So 320 x 240 is the resolution of this thing. Ignore, for tonight's purposes, the 356 x 240. But 320 x 240—that should sound somewhat familiar, at least if you multiply that by two. What do you get if you multiply those dimensions by two?

STUDENT: 640.

DAVID MALAN: 640 x 480, so we're back to a roughly 4 x 3 aspect ratio, if you do the math there, which is to say that these videos that we're distributing are actually pretty small, right? When we blew up those Swedish horses, they filled the screen quite nicely and quite clearly.

By contrast, if I blow up this video, the projector doesn't really give you a sense of just how blotchy things get. But can you already tell that at this resolution the text becomes a little blurry?

(01:20:03)

And if we, say, fast-forward to this, well, this stuff's not bad. It's pretty clear, but it's also a bit . . . We'll pick this one randomly. So tell me about this photo. Critique this photo, ignoring the subject. But talk to me in technical terms. What's interesting about this? Anything?

STUDENT: The outline looks blurry.

DAVID MALAN: So the outlines of Rei—like, of his profile?

STUDENT: His hair.

DAVID MALAN: Yeah, and again, these overhead projectors don't really do it justice. Things look much nicer here. Though Rei looks nicer at any resolution, so... But things are a little... It's not terribly sharp, certainly, right? The Swedish horses, daresay, looked better, at least in terms of their sharpness on the screen.

Rei, by contrast—you know, no slight on his profile—it's just a little blurry. It's a little washed out. Because what file format is probably being used here. It is a video. But for this individual frame of the movie, what's probably being used?

STUDENT: (inaudible response)

DAVID MALAN: A JPEG or something like it.

So within videos, you effectively have frames, literally like you have in a TV show, or in a movie. Right? When you had the old reel projector, a movie is in fact just like a flipbook. But the book is just so long and has so many frames that it really looks fluid.

How many frames does a movie or does TV typically show you every second?

STUDENT: Twenty-nine?

DAVID MALAN: 29.7, 29, 29.7, 30, give or take—usually around there, depending on the encoding scheme. But that is to say, literally, when you're watching TV or a movie, it's as though you're watching a really long flipbook, but one that shows you 29 or 30 images within one second. And once you start to see that many frames per second, your eyes can't even tell that it's 30 different frames because your eyes just can't keep up—your brain can't up with that, so you have the illusion of moving pictures.

But think about moving pictures. What does that mean? Well, they're just pictures. They create the appearance of motion or of movement.

So QuickTime is one file format, developed by Apple, for video. We mentioned another file format for video a little bit ago. What was that?

STUDENT: Flash?

DAVID MALAN: So Flash is another one. So QuickTime, and this ends in .mov. Flash video, which you don't usually see this file extension, because you yourselves don't tend to download these files. Rather, as you saw, they play in the browser, so these details are usually hidden from you. But they're usually .flv.

Yeah?

STUDENT: (inaudible response)

DAVID MALAN: Sorry?

STUDENT: (inaudible response)

DAVID MALAN: Shockwave—we could put it in this category. But for now we'll keep it somewhat distinct because .swf files . . . The Swedish horses was a file called .swf. For now let's just distinguish that ever so slightly, by saying that's more of an animation format, and not put it on this list for now. But you could make an argument there.

What's another video file format you've seen?

STUDENT: RealPlayer?

DAVID MALAN: Yeah! So if you've watched this own course's lectures, via the course's Website, you've watched RealVideo, whose file extensions—there are several—tend to end in .rm, for RealMedia. RealVideos were developed by a company called Real.

How many of you have actually installed the Real software on your computer? So many of you; most of you. So... and presumably all of you who are tuning in from afar, at least with the RealVideo versions. What do you get in this course with the RealVideo versions?

Sorry?

STUDENT: (inaudible response)

DAVID MALAN: So they're bigger, at least if you choose that "bigger" option. What else is interesting about the RealVideo version of this course?

What's going on over on the right side of your screen? So you have that split-screen effect, if you've tuned in that way. You have, like, a list of shortcuts to the ten-minute mark; the twenty-minute mark—little things like that. But the slides are synchronized, right?

We've started to move away from using slides, and I've been moving more toward demonstrations. But for the PowerPoint slides, they are roughly time-synchronized. In fact, perhaps unbeknownst to you, every time we use a PowerPoint presentation here, there's a piece of software running on my laptop that writes to a file the exact time at which we changed slides each time. So after each lecture, I actually email that small file to the video production staff, and they use it to synchronize those slides with the video, which is kind of a cool trick.

But to play these RealVideos, you need special software. In fact, how many of you, now that you have the RealPlayer installed, seem to get little pop-ups in the bottom right of your screen, from Real, telling you what's late and great and newsworthy. Anyone? Yes? No? So . . . yes? No? All right, so those of you who do, you can get rid of all of that stuff.

Real has long taken it upon itself—this is just a company, a multimedia company—to put a whole lot of stuff in software that, at the end of the day, you probably only want to play movies in RealVideo format. So for those of you who have PCs, I would actually recommend going to our course's Website, going to the Software link, and download something called Real Alternative, which, needless to say, is an alternative. And it's sort of like a program—it's the bare bones. All it does is play the videos and the audio files. None of these pop-ups, and news, and weather, and so forth, that come with it.

But the purpose for tonight is that RealVideo is certainly an alternative to these other video formats. And clearly using RealVideo, or the technology made by the Real media company, you can do other things: synchronization, and time stamps, and so forth. But you can do something else with RealVideo. And actually, you can do it with all three of these formats.

When you start playing a RealVideo, those videos tend to be about 200 to 400 megabytes (MB) each week. How quickly do the videos start playing for you when you click on them?

STUDENT: (inaudible response)

DAVID MALAN: Yeah, for the most part, right away—maybe a five-second delay; maybe a twenty-second delay, right? There's that little icon, if you've seen it, that sort of says "buffering." or "loading," or something?

Well, what is going on there? Well, if these videos are 200 to 400 MB, is your computer simply that fast? Is it downloading a 200 or 400 MB file within five or ten seconds, storing it on your hard drive, and then playing it for you to watch?

Probably not. Right? Sort of a leading question. So what's going on?

STUDENT: (inaudible response)

DAVID MALAN: I hear that "str . . . " Streaming. So these three formats, actually, can be "streamed." So these are all streaming file formats, which simply means that you can start playing them before the files are downloaded in their entirety. Which is to say, when a program like Real is buffering, or when a video is loading, even when you go to, like, Google Video, or you go to YouTube, what those Websites do is they start to download just a bit of the video. And then once they have a buffer of, like, five seconds or ten seconds, then you start to get watching the video. And as you're watching the video, the program is ahead of you by five or ten seconds, downloading the next five or ten seconds.

Sometimes, though, as beautiful as this setup is, what happens to you, the user? What do you see?

STUDENT: (inaudible response)

DAVID MALAN: The video—it starts jumping or it just stops, and then you watch some little annoying icon, or whatever. And then eventually it starts to proceed again.

Well, why in the world, if you buffer for ten seconds and then start watching, and you continue to buffer ahead, how could my two hands ever meet?

STUDENT: (inaudible response)

DAVID MALAN: So that's true. So eventually you will download the entire file.

But my question for now is, why in the world, given this seemingly elegant design, would we ever experience what many of you have probably experienced, which is, like, the movie pauses for a second, or freezes for a moment; for several seconds. Why can sometimes real time catch up to where the file is still being downloaded?

STUDENT: (inaudible response)

DAVID MALAN: Memory? Be more specific.

STUDENT: (inaudible response)

DAVID MALAN: Be more detailed.

STUDENT: (inaudible response)

DAVID MALAN: Okay, not bad. It's not that type of memory that's so much at issue. A suggestion back here?

STUDENT: (inaudible response)

DAVID MALAN: So, yeah. I mean, intuitively, that's got to be what's going on, right? The movie is playing faster than it's being downloaded, which is to say ten seconds of buffering just wasn't enough.

Well, at first you only needed to wait ten seconds to download ten seconds of video. But you know as well as I do that sometimes the Internet just slows down, right? Or there's congestion at routers.

We already know that because of IP, routers can drop packets all together. Well, if the packets get dropped, and you're using TCP, for instance, to transport this information, what does TCP guarantee, did we say, a couple of lectures ago? It'll do what? Well, it will retransmit that data. But if you're going to… if that data's going to be retransmitted, well, someone's got to wait for it, and that someone is you. And that's why sometimes these programs would stop.

There's a difference here, though. Notice that all of the videos that we've been discussing, for the most part are prerecorded, right? We record them either right now, and would put it online by tomorrow or Friday. The "Videos of the Week" that the teaching fellows have been shooting—well, those are produced each week, and then they're put online a week later in the entire volume.

So what if, by contrast, you're trying to watch a baseball game live, or a presidential address live, or listen to it live? Well, similarly can you tune in these days via the Internet to live music, or live video broadcasts.

In fact, the groovy music that we tend to play at the start of lectures is the result of my pulling up iTunes, which is a multifaceted multimedia program today.

(01:30:08)

If you've never used it before, it's fairly straightforward. On the left, there's a link these days to radio, and then you see all of these different categories.

And notice, what's the column described as here?

STUDENT: (inaudible response)

DAVID MALAN: A "stream." So what is a stream? A stream is sort of an available feed of bits that you can tune in to, and they're streaming in the literal sense. They'll start coming to your computer. But you're going to tune in wherever that stream currently is. And so in effect what these are is online radio stations.

And if I go to "Ambient," or Ambient here, and choose… uh, which one do we usually use? You usually listen to "Groove Salad" here in E-1.

(music playing)

I think this will be the last night you listen to "Groove Salad," because it's starting to sound like this.

But this is a radio station that is streaming itself over the Internet. Those of you who like Boston radio, there are a number of radio stations even in Boston where, if you don't have access to your car radio or your stereo, but you are sitting at work perhaps, with headphones on, on your laptop, with an Internet connection, some of these local radio stations actually allow you to tune in via the Internet, and they just stream out the information that way.

An interesting thing to note is that there seems to be three different versions of "Groove Salad" here.

(music continues)

What's the distinction among these three lines, if you can see the small print?

STUDENT: (inaudible response)

DAVID MALAN: Yeah, so kilobits per second (kbps). So, good, here, too, is sort of a building block from a few weeks ago. What does this mean if the one we're tuned into, per the little speaker icon there, is apparently… has a bit rate of 128 kbps?

STUDENT: (inaudible response)

DAVID MALAN: Finer sound. So the more bits per second, just intuitively, the more information you're getting, which means the better it sounds. That's true.

And what does it literally mean if this stream is 128 kbps? What is my computer doing each second?

STUDENT: (inaudible response)

DAVID MALAN: It's downloading 128,000 bits every second. And with a decent Internet connection, that's fine.

If instead I were on a slow Internet connection, which is hard to simulate since we do have a pretty fast connection here, well, sometimes bits might get dropped, or maybe the server would get busy, and it just can't keep up with the demand right now. So what might happen?

Well, we could buffer again and again. But when you're talking about live feeds, like a presidential address or a baseball game, or really, a radio station, which, by definition, streams in real time, right? It's not like TiVo, where you download it and play it back later. Radio is sort of by nature real-time.

Podcasts, by contrast, are like offline radio, to make a simple distinction.

Well, what would you rather happen during a baseball game if you're whole purpose is in trying to watch this game live, albeit over the Internet, or your whole purpose is to listen to the game live, over the Internet? Would you like your software to… wait ten seconds. Let me buffer, and then I'll show you what happened ten seconds ago? Right? It's a great way for you to lose bets, right, if this is how you're tuning into the game, betting someone else who's watching live TV.

So, by contrast to a lot of the stuff we do in the course, which really—you do not need to know about computers and the Internet right now! Right? It's okay to wait until tomorrow. It's okay to wait ten more seconds, twenty more seconds. But if you're watching something that, by nature, is meant to be live, like a baseball game or like a radio station, or a presidential address, well, you'd probably rather the software do what in those instances?

STUDENT: (inaudible response)

DAVID MALAN: Stream it, yes. But streaming just means deliver those bits in real time. But what's the problem? Server gets busy. Router starts dropping packets.

STUDENT: (inaudible response)

DAVID MALAN: So that's a clever idea. If there are just… If it's just not possible to get all of those bits to your computer, what could you sacrifice?

Well, you could just turn the thing off, right? That's a wonderful way to compress the information—just stop playing Bush's speech. No comment!

(students chuckle)

But what else could you do? You could just degrade the quality. And in fact we see this here, even though, it's not very dynamic, right? If I tuned in to 128 kbps, that's the stream I'm going to get. And if you play this radio station long enough, you'll see eventually iTunes sometimes pops up a message that literally says, I think, "buffering stream," or something to that effect. And essentially you just have to wait. iTunes does buffer for a few seconds, just to give you that sort of buffer, literally.

So what iTunes does is… Those of you with Discmen—if you even carry around those ridiculously large CD players anymore—well, remember how they used to have, or still have Shock Guard, or something like that, where theoretically you can shake the heck out of the thing for five seconds, ten seconds, whatever it tells you on the box? Well, what are those things doing?

STUDENT: (inaudible response)

DAVID MALAN: Right, those devices, too, are buffering the song, maybe five or ten seconds' worth of the song. So that if you're jogging, or shaking the heck out of the thing, it can withstand that. But—and you can try this at home if you have one of these things, or maybe your roommate has one of these things—after eleven seconds, you start no longer to hear the audio.

So CD players don't just degrade the quality. But radio stations, by contrast, might tune you down to a lower bit rate.

Let's see if we can see the difference—hear the difference.

(music plays)

So this is 128 k…

(Music stops; starts playing again; quality is thinner, less rich)

Good. You can really hear the difference. This is 24 kbps—so an order of magnitude smaller. And even, though, if I had turned this on this way when lecture first began, you might still think the song is bad, but you might not realize that the stream is of low quality.

But contrast this again now with…

(Music plays at higher rate)

Okay, this song sounds bad at any bit stream, really! But the point hopefully is clear.

So to recap: Streaming technology just lets you watch a video in near real time. But if you're watching something that by definition is live, you tend not to just get this buffering effect. Rather the song will degrade in quality. Or what's one other option for a baseball game? Annoying as it might be, if your goal is to continue watching this game live, what's another option the software could do besides just degrading the images, or the quality of the announcer at the game?

Drop bits. And so what a lot of these streaming protocols and programs do is, if they just can't keep up with you, you'll just miss that sentence that Bush just spoke, or you'll miss that save in the game because it's a tradeoff. The designers of the software figure, well, this guy would probably rather we just forge ahead rather than get out of step with real time.

Other video file formats that you might be familiar with?

STUDENT: (inaudible response)

DAVID MALAN: AVI is another one. And I'll jot this on the board. We won't spend as much time on here. You have a whole cheat sheet, if you would like, on your yellow sheet tonight.

AVI is another older file format. Another popular one: DivX is popular. Those of you who deal in online movies, or downloads thereof, might be more familiar with that.

MPEG? In fact, how many of you have ever watched an MPEG-2 movie before? Two of you. Three of you. Anyone else?

Okay, this is always going to be a trick question, right? So, how many of you have watched a DVD before? Bam! All of you have watched an MPEG-2 video.

So we've been talking all tonight about file formats. Well, all a DVD is, is a piece of plastic that has a bunch of bits—pits and lands, if you recall—stored on it. In what format are those bits? Well, DVDs use this format (points to "MPEG-2" on chalkboard), and it's just a standard for video, and it's pretty high quality.

And when you start talking about things like Blu-ray, and HD DVD, and all of these things, for the most part, they're not necessarily just talking about differences in hardware. They're also talking about different schemes for encoding the bits.

And what does that mean? It just means different options. Just as GIF and JPEG gave us different features, so do different formats for video give you different features that, at the end of the day, aren't as important to you, the user, because, for most part, when you, the user, pull up a Website, you don't usually have choice over the format of the video. It's either watch this, or not watch this, though some Websites do let you choose the formats, including our own.

But let's pose this one question about video before ending on a quick chat about audio. You've got a video. Clearly these things are big. These lectures themselves are how large did we say, in terms of videos—in terms of bits?

STUDENT: (inaudible response)

DAVID MALAN: Large. Good E-1 answer. So this five seconds of awkward silence, how many bits did we just waste on those tuning in at home?

So each of these videos, which clocks in around two hours long, less break time, are about 200… 250 MB. That is compressed. That is watching these lectures in 320 x 240 resolution at 32 kbps audio. And it's a funny thing that this music sounds so bad at 24 kb. Well, and maybe I do myself sound bad at 32 kb, but voice does not need as high of a bit rate because I'm just not… I'm just not delivering that much information per second, so to speak.

But how do you get away with using… It's up to you. How do you get away with using fewer bits to store video? Well, what's the analogue to our graphical world?

Well, consider these movie frames? So suppose that, at the top, you have sort of a reel of film, and that film has captured the act of some RV driving past a house and a tree.

(01:40:06)

Well, how in the world, if you want to show users that same reality, that same event, can you do it in fewer bits? That is to say, rather than keeping around a bitmap for every frame, can you throw away any of the bits?

But what does the top reel of film suggest one technique is for compressing video?

STUDENT: (inaudible response)

DAVID MALAN: Good. So just like GIF sort of uses, within one frame, within one image this sort of cheat sheet, and says, "Hey, that was blue, make the rest of this blue," similarly can you do that trick in video, whereby you have what's called "interframe compression." "Inter," like an interstate highway goes within multiple states? Interframe compression goes within multiple frames.

Well, if you already know from the first frame what the house looks like, what the tree looks like, why do you need to keep that information around again? And so what file formats like these tend to do to compress themselves losslessly is just throw away the stuff that we can just repeat from before.

But there's a lossy approach, too, to interframe compression. What could you also do if you want to save space but still give a pretty good approximation of reality?

STUDENT: (inaudible response)

DAVID MALAN: Skip frames, right? The answer's right there. You can skip frames, or drop frames all together. Now, this doesn't mean drop half the frames and then show the video twice as fast, because that wouldn't fly too well.

But think about the Frogger game that we saw before. That in effect had skipped frames, right? At one point in time the frog is here. The next moment, when Heather clicked the up mouse, it then went there. But we didn't see anything in between. So in spirit, that's sort of like interframe compression that's just skipping frames. You don't know what happened to the frog between here and here, but your eyes and your brain sort of decide he clearly went from here to here.

Besides interframe compression, you can also do "intraframe" compression. Inferring, simply from these two reels here (pointing to slide), what is this very much like?

STUDENT: (inaudible response)

DAVID MALAN: Well, it's sort of like the previous example, whereby what's changing on the bottom reel, for instance, from frame one to two, to three, to four?

STUDENT: (inaudible response)

DAVID MALAN: Right, in this case, there's no background, yes. But what that suggests is that we don't need it. We don't need the several thousands, or tens of thousands of bits that might otherwise tell us what the leaves' colors are, and what the petals' colors are. Why? Because the flower's not moving, the bee is moving. So in this way could we similarly compress video. But in this case, are we losing information, or is this lossless?

STUDENT: (inaudible response)

DAVID MALAN: Lossless? Lossy? Lossless?

STUDENT: (inaudible response)

DAVID MALAN: Good question: When you play it, can you tell there's no frame there? It depends. It depends partly on how sharp your eye is. And if you slow the video down, you'll actually see… And you can do this with the course's QuickTime movies. Just pause it, and then drag the little slider. And you won't see fluid motion, usually. You'll see this.

So, yes, you can see it, depending on the precision with which you can navigate the video.

So I hesitated to say that this is lossless, only because, as of this frame, do we know what that flower looks like in its entirety?

STUDENT: (inaudible response)

DAVID MALAN: Well, I mean, as of frame one, if this is the start of the video—just looking at this bottom reel—technically we don't know what is behind that bee. Right? Probably, just if we let our brains fill in the blanks, what is it? It's another pink leaf and maybe some more yellow to the center there. But that's interpolation.

And we said this earlier. What can a computer do to sort of get the "CSI" effect and zoom in ad infinitum to the license plate? Well, it can sort of interpolate and guess, you know, this looks like this and this looks like this. I'm going to guess that this thing in the middle is like the average of those two things. It's a little purple and it's a little yellow. And so a video could use a bit of lossy compression and just fill in those holes for us. It's ultimately a tradeoff, but it's the file formats that decide how best to do this.

A final note. We've looked at audio tonight. We've heard a bit of audio. Consider, though, the following. There are multiple types of file formats for audio. What is the one that most of you… or many of you have probably downloaded before?

So MP3s, right? This is sort of what's on your iPod often, though not always. Apple has its own format for that. MP3s—this is what was popularized by Napster. If you weren't quite tuned in to what was going on, or didn't really know much about computers and the Internet back in the day of Napster, which was—what, like '98, '99-ish, give or take? What was all the rage was the sharing—illegal perhaps—of MP3s.

MP3, actually, as trivia, is MPEG-1, layer 3, though that's technical detail that is not terribly important. But this is only to say that it's related, even to the stuff we've been talking about.

An MP3 is just a file format for audio. But the neat thing is, is that it's compressed, but it's lossy compressed. But turns out, human ears—not so good. Which is to say, if you play a song from a CD, that tends to be, for the most part, uncompressed, though that's a bit of a white lie. But if you play an MP3, that's very much compressed.

Consider this—and you've answered this question before—how big is a typical MP3 file?

STUDENT: (inaudible response)

DAVID MALAN: 2 MB, 6 MB; ten, maximally, at least for reasonably long or short songs. Well, a CD—we talked about CDs in Lectures 1 and 2. How big is a typical CD?

STUDENT: (inaudible response)

DAVID MALAN: So it's like 700 MB, give or take? Well, how many songs are on a typical CD that you buy at the store?

STUDENT: (inaudible response)

DAVID MALAN: Sixteen—I'm going to do easy math. Let's say ten is typical, and that's kind of reasonable; maybe a little low. So if you have 700 MB available to you on a plastic disc, called a CD, and you want to fit ten songs on it, how big does that suggest each song is on a CD?

STUDENT: (inaudible response)

DAVID MALAN: Yeah, roughly 70 MB. That's pretty big. Downloading a 70 MB file is not so fast on a lot of Internet connections. But an MP3 is how big, did you say?

STUDENT: (inaudible response)

DAVID MALAN: Yeah, roughly, what? So, let's keep the math simple. Let's say 7 MB, just as a rough estimate, right? So it's a tenth the size of the original song.

So what does that imply? How do you achieve a compression ratio of 70 MB originally to 7 MB afterward?

STUDENT: (inaudible response)

DAVID MALAN: So it doesn't sound as good because you presumably have to drop information. Intuitively, why is that the case? Well, if you're not dropping information, why was the original song 70 MB in the first place? What was it—just a bunch of zeros? Probably not.

But, again, the human ear—not so finely tuned such that MP3s, for the most part, are difficult for at least non-crazy audiophiles to distinguish from, say, the original CD version. I, for one, can never really tell the difference between a store-bought CD and, say, a set of MP3s of the exact same CD, even though they're a tenth the size.

Consider this example.

(music plays)

So again, we don't necessarily have the latest and greatest audio hardware here, but this is an MP3 of one of Enya's songs. Still sounds pretty good, even in this audio system. But it's a tenth of the size that it originally was when I ripped it from its original CD.

Well, how about this for contrast? Suppose that back in the day, we lived in a world, where, if you wanted to play music on your computer, what you heard was…

(MIDI of Beethoven's Symphony No. 5 plays)

So this is an example of what's called the MIDI (Musical Instrument Digital Interface) file format. Those of you who are musicians and might have digital keyboards at home, and compose music yourself, typically you're storing the songs you compose in MIDI format. MIDI, essentially . . . M-I-D-I, stores the musical notes—A, B, C, D, and so forth; sharps and flats—in some kind of file. And then your computer just plays them back literally.

Well, that's one thing. But if instead you want to actually enjoy the music you're playing, you can't just store all of the notes as notes themselves. Rather, you need, say, the MP3 version or a non-MIDI version that doesn't just represent each sound with an individual note, but sort of stores the information in the aggregate, and you have the same song this time in a much higher-fidelity version.

(higher-fidelity version of Beethoven's Symphony No. 5 plays)

Now we actually have a recording of actual strings, and so forth.

So, your homework for this week, realize, is, one, to familiarize yourself more with a bit of multimedia, particularly in your first banner, and perhaps a mousepad project.

And also, do take a moment, if you haven’t already, to pull up the course’s Website and tune into the Videos of the Week, since those, too, are not meant to just repeat things that I may have said in lecture, but rather are meant to introduce you to topics that we did not cover in lectures.

So with that said, we will see you next week.

(end)