

MITOCW | mitmas_531f09_lec06_3.mp4

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation or to view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

- RAMESH** All right, so that was light waves. And we're going to switch topics a little bit. We saw, if you think about light
- RASKAR:** fields, you can build really cool, like, CI devices. If you have a camera array, you can do all kinds of crazy things like-- of course you can do refocusing and see-through uploaders. But you can do also other things. Like if I have a light field camera-- let's say this one-- which is interesting. There's actually a camera here--
- [LAUGHS]
- as we're talking about that. Anyway, if this was an array of cameras and you shake it, you have image stabilization problems typically. But if you have a light field camera then-- uh, OK. If you want to stabilize the image, then you have multiple solutions. You can either put big mechanical rigs so that they stabilize your motion to get rid of camera shake, or you can put some optical stabilization techniques to dampen out the effects and so on. But if you have-- sorry?
- AUDIENCE:** For taking many pictures [INAUDIBLE] do all that?
- RAMESH** No, this is a video.
- RASKAR:**
- AUDIENCE:** Oh, video, all right.
- RAMESH** Yeah, this is video.
- RASKAR:**
- AUDIENCE:** Cool.
- RAMESH** So you know, there is Schwarzenegger and he's jumping in his truck and there's a camera moving and it's
- RASKAR:** following. How do you keep the camera stabilized although it's being shot in a very rough terrain? That's the classic video stabilization. Maybe that's what-- I should say video stabilization, not image stabilization.
- But if you have a light field camera, then you can let it shake and create a very smooth video of that. How would you do it? Remember, there's a five by five camera array. And I'm going to shake it but I want to create an illusion as if there was a single camera that's traveling through this space in a very slow manner. Yes?
- AUDIENCE:** [INAUDIBLE] a bunch of [INAUDIBLE] So what you can do is take two frames of video and try to find the last one, the closest previous one.
- RAMESH** Exactly.
- RASKAR:**
- AUDIENCE:** [INAUDIBLE]
- RAMESH** That's it. As simple as that. You have 25 cameras and in the very next frame it was shaken. So the camera from
- RASKAR:** which we are shooting the-- let's say the center of the five by five is your actual camera. But in the next frame, that may be too jerky but some other camera took its place in a straight line. So you can just switch between different cameras and find the camera that looks like it's going in a straight line. Yes?

AUDIENCE: So can you switch and you'll see the blur where it's like--

[MAKING SSH-ING NOISES]

RAMESH Yes, yes. So you'll do a little bit more.

RASKAR:

AUDIENCE: Yeah, yeah.

RAMESH You'll do a little bit more. But as opposed to doing traditional video stabilization, where you have to shift an

RASKAR: image and do some warping or something, here you get very nice clear fits. So there was a-- this was a ICCV paper this year-- 2009-- International Conference on Computer Vision. This simple idea-- take 25 cameras and create a stabilized view.

Because the way things are going, 25 cameras are going to be much cheaper than buying a video stabilization rig. I mean, again, it's a dollar right now-- under a dollar camera. So you can imagine in the future, your phone has 25 cameras on the other side. And it can do anything you want, and from that you can create all these effects. And you can do stabilization. And, as we'll see, we'll be able to do [INAUDIBLE] come back [INAUDIBLE].

Really interesting things are going to happen. And we'll be at the forefront of that. So we're going to talk about cameras-- cameras looking at people, cameras looking at fingers, you know-- when you have camera, how we can think about complexity of emitters and receivers, and motion capture. These are all classic problems in cameras for HCI.

And this is a chart we saw early in the-- we saw early in the class about, what are some HCI projects that were actually boring? Because in the field of HCI and the field of cameras and multimedia, is not that new anymore. And so you are to be very careful when-- even at MediaLab, we had a lot of great things were invented in the '90s. Back then, very few people did this kind of HCI.

So anything you do is likely to be new. You don't have to worry about it. You just start the project and you work on it. And it's pretty much guaranteed that what you're doing is brand new.

But that's not the case anymore. Everybody has access to these technologies and computing power and so on. And if you just start a HCI project, it's very less likely that what you're doing is actually new. So you would expose yourself by going to events and seeing what other people are doing and so on.

Maybe your technology's new but your application is not new. Maybe your application is new but the technology is not new. And you have to just have the right mix of all of that to be really impressive. So always remember this-- you cannot impress all the people all the time.

And I just gave a talk to the first year RA class-- I guess you were there [INAUDIBLE] class. And one of the things I mentioned was that for HCI, it's very tempting to work on HCI projects because you get instant gratification, first of all. And second, especially for a place like MIT MediaLab, you get a lot of press. People look at it and say, wow, cool, it's so interesting, I wish I had it, and so on.

And that is good-- any positive reinforcement is good-- but it's also very detrimental in the long run. Because if what you're doing is not really new or impactful, then you get this positive reinforcement because the press says it's interesting. And people who come and see your demo say it's interesting. You think, wow, this is exactly what I should be doing.

But be careful. Talk to people. Try to publish it. Try to expose yourself and subject yourself to some peer reviews, and then see if it's still worth doing it.

So submit to conferences-- high level conferences, not the ones that are organized by your friends or former colleagues-- high level conferences. And just go to those events and see what's going on. So that's my high level advice for doing HCI.

And what we're going to do is a little bit quick run-through of several technologies. But you know, you're very welcome to explore an HCI demo as your final project it's always fun. I mean, I personally love seeing really cool interaction on the devices. But again, the field is saturated and you have to make sure you're doing the most interesting things.

OK, so by HCI, I'm mainly going to focus on device-oriented HCI. And I will not be talking about traditional cameras and HCI done with that. So one of the most classic examples of camera-based HCI is interaction with touch screens.

And the classic project from [INAUDIBLE], now almost 10, 15 years old, is-- you have a projector that's projecting on the screen like here. You have an IR camera that's looking at the screen. And then you have some lights that are illuminating the scene.

If your hand is very far away from the screen, it does not reflect light back. But if your hand is sufficiently close to the rear projection screen, then it will reflect light back to the camera. And because the projector is working in the visible spectrum and the camera's working in the near IR spectrum, they don't interfere with each other.

And then you can recognize a picture of the output-- recognize the blobs and from that, you can figure out. Very beautiful, very classic, but at this stage, not worth pursuing. Too much has been done in this space. If you want to build a product, maybe. But if you want to do research, don't follow this path too much.

Some other interesting thing-- this was done, I believe, by '99-- yeah, '99. As you know, this mouse is-- the camera in a mouse is one of the largest markets for image sensors. If you remember the chart at the beginning of the class, about 100-- 300-- 300 million-- I think-- I believe-- sensors are sold in mouse, which is amazing.

And it's actually a camera that's very low resolution-- about 20 pixels by 20 pixels-- and it's running at very high frame rate-- close to 1,000 frames per second. And all it's doing is looking at optical flow. If you take an image at one instant versus the next instant, is anything changing?

And that's why you can work it-- you don't have to have it on a rough surface, like you used to have with those spinning balls. You can use it on a smooth surface as long as there is texture. If you put it on a completely white piece of paper that has no texture, this will not work because it's looking at image difference-- frame to frame difference.

And it has everything embedded in it. It's taking the two pictures-- I mean, it's taking a video picture. It has a DSP that's doing this optical flow comparison to basically get this 2D vector. Are you moving? What's the X vector and what's the Y vector? That's it.

And of course, you can buy optical mouse for what? Just tens of dollars now. What's the cheapest one? But really, really cheap. Right?

AUDIENCE: Yeah, for sure.

RAMESH Sorry?

RASKAR:

AUDIENCE: \$10.

RAMESH \$10? OK, yeah. And then this group at Microsoft-- Ken Hinkley and others-- said OK. What if you-- this is only two degrees of freedom. Can you add the other four degrees of freedom? So rotation-- retinal rotation is not covered because the image will not change if you stay in one place-- but also able to lift it and tilt it just from the mouse.

RASKAR:

So they just placed a surface on which they can pivot. And then the grid itself was this array of dots, which you can see in a grid form. And the distortion of that tells you-- if you're head up, then all the disks will look the same shape. If tilted one way or the other, you'll get some [INAUDIBLE] distortion. And in addition, they placed this markers within each disk to get absolute positions. And then between these [INAUDIBLE] shells, they can get the related position.

So it's a pretty clever idea-- slightly difficult to implement because this all has to run in real time. And if you have a printer pattern like that-- you must have a mouse pad of that shape. And people just like to have it completely free. The mouse should not have any constraints on where it is sitting.

All right, then FTIR-- which has become really popular now-- Frustrated Total Internal Reflections. And let me show the demo first and then see how this goes. This is Jeff Han's demo from 2005. And FTIR is used in many other fields. So FTIR itself is not new, but the way it's used for [INAUDIBLE] interaction is really beautiful here.

OK, so how does this work? You have an LED that's emitting light. And if the refractive index of the glass is sufficiently different from air-- so the air is 1, the glass is about 1.5-- and then the light is hitting this interface between glass and air at a sufficiently steep angle, it's going to continuously reflect within the bounds of this glass.

And that's exactly how all these other technologies work as well. So fiber optics, light pipes, and all these things also work on principle. So fiber optics, you have-- in the simplest case, you just have a glass pipe constant [INAUDIBLE] air and glass, and you have sometimes coupling of laser. Light goes in. If it goes at the right angle, it will just reflect back to the photo and will be transmitted over hundreds of kilometers this way [INAUDIBLE] imagination. But if you hit it at the wrong angle, there is what? [INAUDIBLE] refractor. [INAUDIBLE].

And of course the [INAUDIBLE] fibers are a part, single fiber, multi fiber. You have air and 1 and 2. So light goes in, hits the flags. The flags [INAUDIBLE] and so on. So by getting different indexes of refraction, you get [INAUDIBLE] something [INAUDIBLE] as opposed to pure [INAUDIBLE]. And [INAUDIBLE] losses that we assume [INAUDIBLE]. If you just have glass and air, some amount of light will actually get through. And a field like [INAUDIBLE] ratios-- [INAUDIBLE] numbers and ratios [INAUDIBLE] will tell you how much of light you're actually [INAUDIBLE]. But if you use this gradient index, then it turns out everything was stable in [INAUDIBLE].

Light pipes work the same way. Instead of a projector, you might have a bulb, but it actually might just-- [INAUDIBLE] pretty far away. So they use light pipes to guide light in a very similar fashion, and so on. So for low light transmission, the idea of FTIR is very well known.

So that's total internal reflection, or frustrated total internal reflection. So this is just a sequence of total internal reflection. And now what we want to do is frustrate the internal reflection and let the light leak from them. And in this case, at the interface between the glass and air, if you allow the light to escape somehow-- you know, put a finger here, or just put some dust particles here. OK? On the fiber-- if you just put dust particles, your light will try to reflect, but because of the diffusion here, it will also go out.

The same principles that you may have seen in some of the menus that they have at restaurants or in movies-- they have this glass slabs. They'll have a slab like this and the lights at the top. This is just a pipe. And then when it's just clear glass, you just see it.

But you can just take chalk or any simple marker, and you can write on it. And it starts to glow. Again, because light is trapped within the [INAUDIBLE], it just goes away. But when you put some pigment on it, light [INAUDIBLE] internal reflection, at the end it gets diffused and it comes out. Everybody heard of this? Many restaurants will have their menu written [INAUDIBLE] which is the same exact principle of frustrating, at some stage, the total internal reflection.

There's even a company called [INAUDIBLE] Display that built a TV that's based on this principle. And the way that works is-- you can probably look up [INAUDIBLE] Display. Sorry. Right.

AUDIENCE: Thank you.

[INTERPOSING VOICES]

I got it now.

RAMESH OK. Do you know what we just did?

RASKAR:

AUDIENCE: Sorry?

RAMESH Do you know what we-- no, never? OK, it's fun. It's really cool, cool idea. So let's say home projector from which
RASKAR: you want to create a big flat screen TV. If you think of a projector, there a couple of solutions. Either I can do it this way, or I can put some mirrors to create another functioning TV.

But this company in UK came up with a very clever idea, where you have a [INAUDIBLE] glass and then you have a projector. And it projects the light on this side. [INAUDIBLE] still here will, you know, [INAUDIBLE] a steep angle. And they put a film here. It's actually [INAUDIBLE].

And the angle between this wedge, it's chosen in such a way that if you shoot an array at a particular position, it will reflect, reflect, reflect. But eventually, the angle will not be shallow enough, so that the light will actually come on.

So the projector reflection doesn't-- if it was a straight pipe, if it comes in, it will maintain that angle forever. But with a wedge, eventually, it will drop off. And if you shoot from other angle, it'll come out [INAUDIBLE].

So you might have a projector-tied image here to reflect and end up creating this nice, 15 inch wide image. A very beautiful product that just-- I believe it got acquired by Microsoft-- the Microsoft Surface. [INAUDIBLE]. And those people are working with-- yes?

AUDIENCE: Is there [INAUDIBLE]?

RAMESH
RASKAR: You know, it just-- I just saw, actually. I believe they have a diffuser. But it's not [INAUDIBLE]. And now I just realized that because of the diffuser, it would frustrate at every point--

[INTERPOSING VOICES]

AUDIENCE: So is there a gap between the [INAUDIBLE] diffuser? Then--

RAMESH
RASKAR: Oh, that's a good one. All right, perfect, OK. That's right. That's right. If it-- so if there's a gap between the [INAUDIBLE] and the diffuser, extra debate is on me. It's a good one. [INAUDIBLE]

[LAUGHTER]

So if there's a gap here, then you still get internal reflection. But when it eventually comes out, it will diffuse. Then it can be seen.

AUDIENCE: [INAUDIBLE] presumably you could also get the same conditions [INAUDIBLE].

RAMESH
RASKAR: What are you [INAUDIBLE]?

AUDIENCE: So like [INAUDIBLE] and you have-- you have one projector projecting both images--

RAMESH
RASKAR: Right.

AUDIENCE: Because they come out of the [INAUDIBLE]

RAMESH
RASKAR: Excellent. Did everybody get his question? He's saying, OK, if I don't have a diffuser-- so let me zoom in on this part. Light comes in and then it goes out. And it just [INAUDIBLE]. It's not internal reflection anymore. It just straight goes out.

And I think what you're asking is, can I make it in such a way that it goes out this way but also goes on this way? And then if I just control this, maybe I can [INAUDIBLE]. Is that it?

AUDIENCE:

Yes.

RAMESH

It's an interesting challenge. It's a very interesting challenge.

RASKAR:

There's some properties you could exploit. So typically because of the [INAUDIBLE] and so on, [INAUDIBLE] sensor with a dual photography. If a light goes in one part and it does a strange thing and comes out, you could shoot array in the same direction and come out the same.

And this duality can be very quickly used to-- it's like the perpetual machine. If someone starts proposing to you a design that sounds like perpetual motion, you can ignore that. And the duality of light can also be used as a quick check to see if your design is going to work. So here it seems like it could be breaking duality, right? Because when light comes in, and it goes into two different directions, right?

[INAUDIBLE] so that if I shine light this way, both of them will come out in the same way. And usually when the splitting and combining happens, the duality principle says [INAUDIBLE] aren't possible. So that's the only exception that's [INAUDIBLE].

A beam splitter, for example-- then light can go here but light can also go here. So both of them are traveling along the same direction. Light will shine like this way. Light will also go this way and this way. So there are some places where you cannot say, OK, one ray will correspond to one ray out. Or how can the same ray be in two places?

So sometimes it can be explained. This is a situation where maybe you can't explain it. [INAUDIBLE]

AUDIENCE:

[INAUDIBLE] project. I think that's why we suggest first, I mean, actually--

RAMESH

So two different projectors.

RASKAR:

AUDIENCE:

Yes, or two different images somehow are different angles or something.

RAMESH

[INAUDIBLE] also very easy. But they are very easy because I believe there's a sweet spot. If you project from the right angle, you'll get this image. If you project from a different angle, this phenomenon [INAUDIBLE]. There may be a very tight sweet spot that works. But it's interesting. It's interesting.

RASKAR:

So this notion of frustrated total internal reflection is used to make this. Also, some of the best head-mounted displays also use these principles. You have your eye, and you have a single thin sheet of glass. And then you form an image here. And then this pixel travels back and forth here. And then you have [INAUDIBLE] observer [INAUDIBLE] refractor index, so this comes out. [INAUDIBLE]. Very, very quick, the image is projected here and seen [INAUDIBLE].

I think they're one of the best head-mounted displays. [INAUDIBLE]. So how do you use it for HCI? Again, the same principle. We're going to do the simple reflection.

In this case, instead of the diffuser, we're going to use a finger. And the finger is going to prevent it from doing pure internal reflection, and will also [INAUDIBLE] light on the other side. And then in this case, you can just put a scatterer so that this is reflected in [INAUDIBLE] direction. So which way is the projector here?

AUDIENCE: [INAUDIBLE]

RAMESH
RASKAR: So this is only for capture. There's a camera down here that's looking at the scatterer. And there's a projector as well that's projecting light on this diffuser. So in this case, the finger is on the right and the camera and projector are on the left of the screen.

Only 2005, that's when the paper came out. By 2007, 2008, [INAUDIBLE] company, it was a big hit during the elections. In 2008, CNN was using this product. All started with a paper in [INAUDIBLE] conference. Anyways.

So once you have scattered light, put it on the camera, and then only where you have physical contact, you see very bright points. If you remove it, if you hover it or-- a glass-- you can see-- again, the same principle. But if you provide a [INAUDIBLE]. [INAUDIBLE] in sheets of water, there's liquid crystal where you have droplets of water and then you shine it underneath into that. And the light travels through it so you see some-- on this side, it is completely transparent. But if you break a beam by putting your hand or [INAUDIBLE] like that, the light comes out.

OK, so one disadvantage of this particular scheme is that you need to have sufficient separation between the screen and the projector camera. Because it's looking at this-- it needs sufficient area to create this-- you could do, like, a [INAUDIBLE] TV and [INAUDIBLE] and so on. But nevertheless, it's not as convenient if you want to create this on a laptop or on a mobile screen.

So that's why this project started [INAUDIBLE]. You might be able to create that and put it in the same [INAUDIBLE]. [INAUDIBLE] programs-- this group at Microsoft [INAUDIBLE] and others, they are building FTIR mounts, where you just have an acrylic colored piece of glass, and there's a camera here that's looking at the gestures on top of it [INAUDIBLE].

AUDIENCE: Think you got the Best Paper Award.

RAMESH
RASKAR: Sorry?

AUDIENCE: Think you got the Best Paper Award.

RAMESH
RASKAR: It was just presented a couple of weeks ago.

AUDIENCE: Couple weeks ago?

RAMESH
RASKAR: [INAUDIBLE] beautiful. So yeah, this is still happening.

AUDIENCE: Where was that?

RAMESH UIST, I believe. Is that right? Yeah. User interface-- the same place where Jeff Han presented his paper four years ago. And the same paper also talks about a site mouse where-- you can see the two fingers here. Just barely? And two fingers here?

What they're doing is they're creating a sheet of light that's traveling almost parallel to the table. So if I put my finger-- only the bottom part of my fingers will be enough. And it's looking at how the light is being obstructed by the fingers. And there's a camera that's just looking straight at it. But most of the pixels are wasteful, and other [INAUDIBLE] pixels are being used.

AUDIENCE: But this is actually only one of these things.

RAMESH It's only one of these things. That's right. That's right.

RASKAR:

AUDIENCE: Sort of, but I think they can tell how far away your fingers are based on the size of the radius of the--

AUDIENCE: OK.

RAMESH I think they just moved the camera and there's some triangulation. But they want to keep it cheap. And this principle of having a sheet of light and see where you cut it, very obstructive-- it's used in many other projects. So if you remember, there was a canister projector where they wanted to create a keyboard anywhere in space.

So all they had was a canister projector which projects a keyboard pattern out here. And then at the bottom, again they had a sheet of light that goes out the same way. So the sheet of light kind of goes here like that. And there's the camera. So when you put your fingers on it, it gets blocked. But because the projector is also projecting the keyboard here, you know exactly where you are touching. So it's kind of a keyboard out of a pen. So again, here's the projector, camera, and the laser sheet. And because it's a slide projector, it's very cheap. [INAUDIBLE].

And then those big Smart Board-- these pens you may have seen, have also got a similar structure. They have a big board and there's a little bit of a baffle. And the corner of this screen, that are cameras that are looking out.

And this camera actually mostly [INAUDIBLE] looking at [INAUDIBLE] and illuminating the whole board with a [INAUDIBLE] sheet again. So when you put your finger, one or more of these cameras will see where you're touching with your finger or you're pointing with your pen. They match the colors. And the triangulation they can figure out [INAUDIBLE] screen.

So it's huge-- I don't know, \$100 million, [INAUDIBLE] million dollars [INAUDIBLE] Smart Board. [INAUDIBLE]. The board that's in [INAUDIBLE] room, is that a Smart Board? The big one? [INAUDIBLE]. Anyway, this is very powerful. [INAUDIBLE].

The aim is to put this in every classroom of K-12 [INAUDIBLE]. And of course, all this space will have projector and whiteboard. So you have a projector that's projected on the whiteboard, and then you [INAUDIBLE].

All right, and then of course you are familiar with the Wiimote, which was quite revolutionary because it placed a camera in the remote. So the Wiimote actually has a XGA camera-- runs at 100 Hertz. And it can track four plots of light.

So when you play with the Wii, there's a sensor bar on top which is actually not a sensor bar. It's an image bar. See how the bar at the top-- which has, I believe, four LEDs. And then when you're playing with your Wiimote here, there's a camera that's looking at this [INAUDIBLE], and it's just like a four-point sensor. And that's, like, the most [INAUDIBLE] with respect to this kind of [INAUDIBLE].

And because IR-- you may have seen examples where people remove the sensor bar and just put a camera. [INAUDIBLE] anything that's near IR [INAUDIBLE]. But to me it was really amazing because this is completely wireless. So they had to do all the processing onboard of these four points, and they transmit it back over Bluetooth, or whatever other wireless they have, to the Wiimotes in your hand at 100 Hertz. And you can buy the whole system for, again, just tens of dollars.

And the fact that they were really great, sufficiently sophisticated processing was just amazing. And so Johnnie Lee from CMU-- he was working with me on some topics I'll show you later on-- really exploited this emerging platform for automation. And he said, OK, let's use this camera to do [INAUDIBLE]. So [INAUDIBLE]. So it's not transmitting back a whole [INAUDIBLE] image back to the base station. It's just transmitting the coordinates of these photons, and that's what makes it very low bandwidth [INAUDIBLE].

And a similar trick is also used in the Vicon motion capture that we have [INAUDIBLE], where you have a high speed camera that's emitting idle light, and then the [INAUDIBLE] reflector that's [INAUDIBLE] back to that point. Except in this case, you may have an image that's 2,000 megapixels, but I think you have [INAUDIBLE]. And you might [INAUDIBLE] and send only the coordinates from the [INAUDIBLE] that uses the pattern.

So you know, Johnnie Lee, I'm sure, is involved with demos where he created Smart Board, sort of Wiimote, and he changed [INAUDIBLE]. Really interesting things he's done-- even harder, 3D tracking. Because now this Wiimote, which is going over here-- this is a camera. And we know where you are in the field. [INAUDIBLE].

OK, this is looking at this multiflash camera. The one trick I didn't mention was, you can use colored lights. So previously we saw that using multiflash camera, you can get very high quality [INAUDIBLE], which are ideal for visual interactions. But you can actually use three lights at the same time. So shadows are cyan, magenta, and yellow. [INAUDIBLE] is on top. [INAUDIBLE]. Just by displacing the lights with respect to the main camera, we have a very robust image.

All right. Then we have these techniques where we're not looking at the coordinates, the transform between the camera and the point. But we're looking at a transform between a camera and a fixed spot. So this [INAUDIBLE] has a camera, which is very good for the design because it's really close to the fan. And looking at the patterns that are encoded can [INAUDIBLE] ink. So it can be printed on-- you can print on top of a document.

But because usually carbon is transferred through near IR, [INAUDIBLE] dots-- dot patterns on the printer, it can be still seen through with this pattern. And the way they do that is they simply displace these dots from a uniform grid. That gives them basically a [INAUDIBLE] encoding. So a six by six encoding for them apparently is sufficient to create [INAUDIBLE] 36 spots-- 36 dots or 36 [INAUDIBLE].

And then-- so if you print something that occupies about half the use, they can provide unique coordinates-- every point in this 1,600 [INAUDIBLE] 1,600 [INAUDIBLE]. So you can print a lot of paper, and every six by six block on this paper is completely unique. So the camera will take a look at a picture of the pattern, decode what that pixel space really is, and record that. And then if you just move your pen on top of this printed pattern, it will know what trajectory you're following.

It's pretty amazing idea. I believe they have a lot of [INAUDIBLE] in this space. And of course, the pen not just writes where it is, but it can also write. Sometimes we forget [INAUDIBLE].

A pen that can write. And of course, they have other [INAUDIBLE] Bluetooth. So that they can just write with a pen and [INAUDIBLE] all the trajectories have been saved, all your nodes are uploaded [INAUDIBLE] to your [INAUDIBLE]. I think it's a very powerful idea, that you can encode a pattern in a fixed [INAUDIBLE] object, and be able to uniquely identify it with a [INAUDIBLE].

So building a camera like this is almost like building a microscope. And the problem with that is, your depth of field is going to be very narrow if you go out of focus. Like a microscope, if you shift the dial a little bit, of everything goes out of focus. So this camera does have a problem.

So they're engineered well enough so that, most of the time, the pen is looking straight down. But then it's [INAUDIBLE] the distance between the camera and the [INAUDIBLE] is changing a little bit. [INAUDIBLE]. You generate it well enough so that you can [INAUDIBLE].

All these problems of looking at binary factors and be able to deploy them is-- we will see that when we start talking about coding and [INAUDIBLE] and so on. It's really very interesting because [INAUDIBLE], I think, goes on a popular-- iPhone [INAUDIBLE].

AUDIENCE: I don't know about popular, but--

RAMESH Popular amongst a few people.

RASKAR:

AUDIENCE: Yeah.

RAMESH You know, recording software. And everybody always complains that, oh, it has to be our focus. Our focus is

RASKAR: difficult to decode and so on. But mathematically speaking, if you know that the image that you captured is binary, then you can let it blur significantly, and you can still recover the original environment. If you give me an arbitrary photo and look, it's difficult to keep blurring it back and forth. But if it's a binary image, it's very easy to go back and recover [INAUDIBLE].

And so a lot of the techniques are exploring that. I still haven't seen an iPhone app that actually exploits smart image [INAUDIBLE] to recover in the future. So the stakes are [INAUDIBLE]. Somebody needs to [INAUDIBLE].

AUDIENCE: It's like the third most popular data capture app, and all it is a one [INAUDIBLE]. But it works really well.

RAMESH But I'll give you a good example. It does a little bit of-- what name is it-- [INAUDIBLE]. But it still requires

RASKAR: sufficient data to do something here. I don't know. Obviously [INAUDIBLE].

So this is a challenge in a lot of these ad hoc efforts, that they aren't thinking about the whole problem the whole time. They just realize the type of-- they can just study the iPhone blur an iPhone camera produces, and then write new software that's suitable for the camera. They'll get much better outcomes.

So I really feel like starting my own-- writing my own app and [INAUDIBLE] just for myself. Because it's not going to be well engineered for [INAUDIBLE].

AUDIENCE: [INAUDIBLE]

RAMESH RASKAR: So when we come to the coding, you'll see that's how-- all the steps are extremely powerful. So optical mouse, of course-- 20 by 20 pixels, 9,600 frames per second. The camera you have on your mouse is 10,000 Hertz camera. And I took it from the [INAUDIBLE] electronics manual. And the resolution is 800 dots per inch. Even if you shift a little bit, it can detect the variations.

It's pretty amazing. It can deal with motion blur up to 40 inches-- 40 counts per-- oh-- 40 inches per second. It's just amazing, all the specifications of how beautiful this whole thing has been [INAUDIBLE].

Another popular problem in HCI is being able to track the case of [INAUDIBLE] tracking. But [INAUDIBLE] move your [INAUDIBLE] it's quite challenging. So there are several solutions. This is one of the problems that are still not solved as well.

Most of the solutions they have are active. So they eliminate [INAUDIBLE]. They eliminate some [INAUDIBLE] light and look at basically the reflection of that bright light using the camera. And if you move your heads, the reflection will appear a different place. You can pick up your eye of the [INAUDIBLE]. As it moves around, the reflection of that bright spot changes. [INAUDIBLE].

So for big issues-- basically they're looking at [INAUDIBLE] all the time, I think, to be able to do that. And another problem with [INAUDIBLE] tracking is that, because of [INAUDIBLE] and because of things, it's very difficult to predict. [INAUDIBLE] has the camera [INAUDIBLE] first, and also the motions are discontinuous. So you cannot just put something simple-- small [INAUDIBLE] figure out what [INAUDIBLE].

These are very interesting problems. And if you do have really lightweight case tracking solutions, then this will completely transform how we think about interacting with machines. You sit there and [INAUDIBLE] with my eyes, because wherever I'm looking, I want to go. Of course, if there is some interesting distraction, [INAUDIBLE].

[LAUGHTER]

But you can imagine in non-critical scenarios, this would be extremely difficult, because we get tired speaking. We get tired typing. There are fatigue issues. But with convenience, you can do a lot of things. All right, so let me show one last one and then we'll stop.

So thermal IR motion detectors-- I really love them because the principle behind it is so simple and it works so beautifully. Those thermal detectors-- is basically a one-pixel camera. It's able to detect a motion of an object-- just one. Actually, there are two pixels. So something that opens the door for you, or turn the light on and off, and so on. It's just two [INAUDIBLE] pixels.

How does it work? This is a nice cartoon. Let's say each of those two pixels that are represented here have an aperture that's looking at very narrow regions. As some warm body-- because it's looking at body heat, like modern radiation-- moves through the space, first you will trigger this detector, and then you'll trigger the second one.

Now, if you didn't have this differential measurement-- and by the way, the whole point of that is just the difference between the two. So in the beginning, let's say, the first one gets a spike where the second one has no signal. So the difference is positive.

When the one part is in the middle, distance is back to zero. When the one part is here, this one is positive, this one is negative. The difference is as I discussed. This one is a signal, this one's no signal, the difference is negative. OK?

So if something moves in front of this, you get this signal that has a very high radiation [INAUDIBLE]. On the other hand, if just the temperature in the room increases uniformly, this is looking at background radiation. So all the [INAUDIBLE] in the room will also start radiating energy at [INAUDIBLE]. But then both of them will increase by the same factor so the difference between the two will remain constant.

So this is a very simple way of eliminating slow changes and picking out only the first one. So this is a simple cartoon. But in general, you have much more complicated configurations where you have a Fresnel lens, usually in front of the detector. And instead of just two zones, each of them actually look at an overlapping [INAUDIBLE]. So they go from zone A to zone B to zone A to zone B. You're going to get very high variations in this output. Although again, it's filtered and simplified, but you can easily distinguish it from an increase in the DC level of the [INAUDIBLE].

So that's how a thermal detector works. It's basically a one-pixel camera because output is just one screen, although it's been sensed with two pixels. So think about how you can build-- this is basically a coded aperture-- it's a smart aperture that's allowing you to use just a single detector to do measuring of motion-- not [INAUDIBLE] but motion. And many animal eyes also use the same [INAUDIBLE] to detect motion. They have just a few [INAUDIBLE], but by using a clever aperture, [INAUDIBLE].

Any questions here? All right. I'll send out the third assignment later tonight or this weekend. And remember your second assignment is not due today, but is due on Tuesday. And feel free to ask me questions, or Professor [INAUDIBLE] or Professor Oliveira. And also [INAUDIBLE].