[SQUEAKING]

[RUSTLING]

[CLICKING]

JUEJUN HU: All right. So let's get started. So in this lecture, we are going to talk about, essentially, optical waveguides and why glass has such unique properties that make them so useful for waveguide applications.

Now, speaking of waveguides, we'll be talking about basically what we call the integrated photonics. By way of introduction, I always like to draw an analogy with it's a close analog electronic integrated circuits or integrated electronics. So you all know that you have this thing called Moore's law, which basically says that, essentially, the function of electronic circuits would grow exponentially over time.

And this applies not just to CPUs. This also applies to storage units, to other electronic components that similarly leverage the electronic CMOS microfabrication techniques. So just to show you the power of this kind of exponential scaling law, I always like to present this picture here. You can make a guess about what is this big piece of equipment that they are trying to carry onto the aircraft.

AUDIENCE: A GPS?

JUEJUN HU: It's a hard drive with a 5-megabyte capacity. This is back in 1956. So you can see, with exponential scaling, how the capacity of a hard drive has actually evolved towards today. What I'm trying to argue here is that people believe that the same kind or similar kind of paradigm shift is going to take place for optics as well.

What I'm showing you in this picture here is what an optics lab traditionally would be looking like. I don't know if you've ever been to any of these optics labs, but it's going to be a dark room. You have a bunch of discrete elements. You have lenses, prisms, gratings.

One of the big part of the job for a PhD student specializing in optics is to do alignment. You have to carefully align all the pieces together so that you can actually do your experiment. And depending on how complicated the experiment is, the alignment may be the step that actually takes the most of the time.

So you can imagine that this is something that's very complicated, not robust. It's also very expensive because there's a large number of elements involved. So what people envision is that they can essentially leverage, similarly, the microfabrication techniques, try to essentially create microscale counterparts of the discrete elements using microfabrication methods, and therefore put all of them together on a common substrate to create what are called photonic integrated circuits or integrated photonics.

So the value of this kind of technology is pretty clear there. One, of course, as I mentioned, in terms of cost, it costs much less to actually fabricate a chip compared to having all the large assembly of discrete elements. What's also more important is about scalability. As I mentioned, optical alignment is a big challenge when you're dealing with discrete optical systems.

But in photonic integrated circuits, all the elements, essentially, they are precisely aligned together with lithographic precision, because everything is defined basically by lithography. So that way, you can actually create much more scalable optical systems. The small footprint also means that you can actually deploy them for many different applications that previously is not possible.

It can range from, for example, point-of-care diagnosis, if you're using them for sensing, to essentially all type of, for example, sensors. You can deploy onto vehicles, robotics, UAVs, whatever.

Last part, power. Turns out that the power that's consumed by a lot of these devices actually scales with the volume, in lots of cases. So by making them small, you also make them a lot more power efficient at the same time as well. OK.

All right. So this is actually not a new concept. And what you will see is that glasses actually play a very important role from the very beginning of the idea of integrated photonics. So what I'm showing you here, this is actually-- I'm quoting-- this is a paper by Stewart Miller back in 1969, where the concept of integrated optics was actually first proposed.

What he has proposed is that you can actually combine these different elements here, like laser cavities, switch, detectors, on one common glass plate. And then these different components, they are linked together by essentially glass optical waveguides.

So that's what I'm trying to argue here, that glass materials is basically the first material people have actually investigated to create this kind of integrated optics platforms, both as a substrate material and also leveraging its transparency for essentially a low-loss light transmission between the different components.

Now, more than 50 years later, we are obviously looking at much more advanced and complex integrated systems. But despite that, I would want to argue that glass materials still play a pivotal role in state-of-the-art integrated photonics systems. So for a typical waveguide, this is the configuration we are talking about here.

You essentially have a substrate, and then you grow a low-index material that serves as a cladding. You then structure a high-index core on top of substrate here. So this gives rise to many different possible variations of materials that can be used for these kind of waveguide systems.

For substrate, the commonly used substrate platforms, you can use crystalline silicon or III-V semiconductor, but amorphous substrates, such as glass and polymer, are often being used for various applications. They offer the advantage that they tend to be a lower cost compared to their single crystalline counterpart.

Also, if you're using polymers, they also offer additional properties, such as mechanical flexibility as well as biocompatibility. The waveguide core materials, I'm again listing some of the common materials here. Besides crystalline silicon, a lot of amorphous materials are considered viable waveguide materials. These include amorphous silicon, silica glass, silicon nitride, as well as polymers.

So again, compared to a crystalline material, the main advantage people want to use this material for waveguides-- a few things. One is that they are a lot more easier to process. If it's a single crystalline material, you need to grow them either through a bulk single crystal growth process, or you can grow them in thin-film form on a lattice-matched substrate through epitaxy.

This is actually a pretty severe constraint when it comes to integrating devices. You want to be able to have waveguides, for example, in different layers of photonic integrated circuits. But if you use crystalline material, you basically restrict to either you have to use a substrate with a lattice constant matching your material.

Or you can use wafer bonding to bond essentially a single crystalline material onto a different substrate. But again, that's also a relatively complicated process. And you can usually only do that for one time. Amorphous material, on the other hand, you can basically deposit it on practically any kind of substrate without having to worry about lattice-matching constraints.

So that gives rise to a lot of process flexibility. You can essentially place those waveguides at different layers in a stack. You can put it on semiconductor. You can put it on dielectric materials. So it's really making integration a lot more convenient.

So now, before I go ahead and talk about how you actually fabricate waveguides out of glass materials, I do want to highlight some important differences between optical fibers, which we discussed in the last lecture, as on-chip waveguides or planar waveguides. One of the big thing that's different here is essentially loss mechanism and the magnitude of optical loss.

So I mentioned about optical fibers. You can get your loss down to 0.2 dB per kilometer in silicon fiber at the 1,550 nanometer. Very, very low loss. Now, in waveguides on a chip, the loss tends to be a lot higher. The typical loss we are talking about here is of the order of dB per centimeter. So a million times higher compared to optical fiber there.

I mean, there are some recent reports. People are essentially demonstrating very low-loss waveguides with loss of the order of 0.0-something dB per meter. Still, it's higher than fibers. But usually those require very special process.

So bottom line, the losses tend to be a lot higher in on-chip waveguides. This is usually not considered as a showstopper because, for fiber, we have to transmit signals essentially over very long distances, tens of hundreds of kilometers. But for a lot of the on-chip applications, you really only need to send signal across essentially a chip, which is like a centimeter across.

So this level of loss is considered acceptable. I mean, it's not great, but it's actually acceptable for many applications that involve on-chip integrated photonics. Now, the reason why on-chip waveguide has such a high loss, it goes back to essentially the loss mechanisms here.

So there are essentially two kinds of major loss mechanisms. You have loss coming from the material itself, material attenuation, and then also coming from scattering, Now, for scattering, a lot comes from, essentially, any kind of surface roughness. So when you prepare a fiber or when you prepare a waveguide, you're not going to get a perfectly smooth surface. There's always some interface or surface roughness that can then actually divert light propagation, leading to scattering loss.

So a simple equation I'm listing here essentially writes this roughness scattering loss, alpha s. It should be scales with the index difference between the core and the cladding squared and also a square of the RMS roughness. So basically, if you have a waveguide that has a large surface roughness, then the loss increases quadratically. Also, if a waveguide has a core and cladding material with large index differences, then effectively, you have a much larger index perturbation for the same roughness, and therefore, you also have a much higher scattering loss.

So now we can take a look at why optical fiber can have such a low loss. So for optical fibers, normal optical fiber, the core and cladding typically has a very small index difference. Because the core and cladding index difference is generated by essentially doping the glass core or the cladding with some dopant materials, like germanium oxide.

So typically, the index difference we're talking about for fibers is of the order of 0.01 or 0.001. So very, very small index difference. So in that regard, it's inherently less sensitive to any kind of roughness. That's the first thing.

Second thing is if you think about the fiber manufacturing process that we showed in the last lecture, you see that we essentially have a preform. You have to heat it up to basically turn it into a viscous liquid, more or less. And then you apply axial stress to deform it into an optical fiber. So in this whole process, because it's sort of like a viscous liquid, surface tension actually plays a very big role in dictating the surface morphology.

Now, surface tension will always try to smooth the surface, therefore minimizing surface area and lowering the Gibbs free energy of the overall system. So because of presence of surface tension, it means that the roughness, residual roughness, on the fiber surface is minuscule.

So these two factors combine together. One, it has low surface roughness. Two, it also has very low index contrast. And that's why, in a typical optical fiber, this kind of surface scattering loss is usually negligible compared to material loss. And that's why we are getting to essentially material loss limit down to 0.2 dB per kilometer for silica, for example.

Now, however, there's actually one case where the surface roughness loss can become significant for optical fibers-- the microstructured optical fibers. You may remember in the last lecture when we talk about microstructured fibers, these can actually be a hollow core. Essentially, the core is hollow. You have this kind of holes arranged in hexagonal or whatever lattice to create essentially photonic bandgap effect, to trap light in the hollow core.

Now the difference here is that for the hollow core, the index difference is actually a lot higher, effectively, because the core is basically air, index 1. The cladding effectively now is essentially glass, which is 1.5. So the index difference, instead of being like 10 to the minus 3, this is like 0.5.

So in a hollow-core or microstructured optical fiber, the roughness scattering can become the dominant loss mechanism. And in this case, if you think about-- let's think about a simple model of just a glass rod, a glass cylinder. You are going to have some-- despite that you have surface tension, you are still going to have some essentially roughness.

The roughness takes the form of this kind of surface waviness. So it creates basically a periodic modulation of the diameter of the rod. This is called a surface capillary wave. Now, surface capillary wave, there is some energy associated with this surface capillary wave because it's basically associated with some kind of a degree of freedom for the system. OK.

From statistical mechanics, you know that the Hamiltonian, for the Hamiltonian of the system, every quadratic term in the Hamiltonian corresponds to half kT of energy. I don't know if that's how you learn the energy equipartition theorem. The other way to present it is that essentially every degree of freedom you have a 1/2 kT associated with every degree of freedom

I mentioned that this kind of surface modulation of surface capillary wave is associated essentially with the degree of freedom for the system. So it means that whenever temperature is not 0, there's always some finite energy associated with every wavelength of the surface capillary wave. So the wavelength essentially means the axial period for it.

So when we are gauging the residual roughness in a glass fiber, we have to look at-- for example, this has to be somewhere around Tf. Essentially, you have a glass rod. As it cools down, when it goes through glass transition temperature or goes through effective temperature, essentially any kind of roughness due to this kind of energy equipartition is going to get frozen. And those are going to be the loss limit for-- that's going to be a surface roughness limit for a fiber that's prepared using the process.

Now, fortunately, this kind of roughness tends to be pretty small because k, the Boltzmann constant, is very small. 10 to the minus 23. So if you convert that into essentially the surface roughness, it's quite small. But it can still play a significant role when it comes to devices like photonic crystal fibers.

So that's where the roughness comes in for optical fibers. Now, for planar waveguide, on the other hand, the roughness usually comes from, essentially, the manufacturing process So in the next few slides, I will talk about a few different ways how you actually prepare planar optical waveguides out of glass materials. But one common feature is that none of the fabrication process is perfect. They always leave behind some residual roughness, especially on the sidewalls of the waveguide.

A typical fabrication process that I'm going to talk about uses essentially lithography and etching. You lithographically define the waveguide pattern, and then you actually etch it down to actually create the waveguide structure. So during lithographic patterning, when you're trying to pattern the line, it's not going to be a perfect line. There's always a little bit of roughness at the edges, just due to intrinsic imperfection of the photomask creation, due to some kind of variation in the mass fabrication process, as well as variation in the etching process for the material.

So those then get translated into the sidewalls of the waveguides. So the sidewalls of waveguides usually tend to be a lot rougher compared to the case of optical fibers. We are talking about RMS roughness of the order of a nanometer. So this is actually significantly larger compared to optical fibers.

The top surface can be pretty smooth if it's a deposited film. So most of the roughness actually comes from essentially the sidewalls. Sidewall scattering roughness. So in addition to the much larger magnitude of roughness, on-chip waveguides also tend to have a much larger index difference between core and cladding.

So this depends on the material. So if it's, say, a silicon waveguide, then the silicon has an index of 3.45. The core is, say, silicon dioxide. It's 1.45. So index difference is like 2. So much, much higher compared to fibers.

Silicon nitride, index difference is something like 0.5. For other glass materials, depending on fabrication process, the index function varies between somewhere from 0.5 to 0.1. So usually it tends to be a lot larger compared to index difference of optical fibers. And that's why, because of the much larger surface roughness magnitude and the high index difference, this is what leads to the dominance of surface roughness scattering in waveguide loss contributions.

So typically, when you're talking about loss of planar waveguides, a large part of it is basically coming from this kind of a roughness scattering. So this is quite different compared to fibers. Fibers is, in most cases, intrinsic material loss. Waveguides, it's basically coming from this extrinsic scattering-induced attenuation.

All right. So now, in the next few slides, I will just give you some example processes of how people actually can pattern this kind of waveguide devices. So the most standard one, the most commonly adopted process, this is quite generic. This is essentially lithographic patterning.

You start from a substrate. You do a deposition to form this low-index cladding layer. You then, again, do another deposition to form the high-index core layer. You can also do a bonding, wafer bonding, if it's for single crystalline material, and then you get rid of the substrate to get a thin layer of crystalline material.

For glass, it's a lot easier. You usually just do a simple deposition. You then go through a lithographic process. You first pattern a photoresist mask through photolithography. You then do an etching process to essentially transfer the pattern from the resist into the waveguide layer. And in the end, it's customary that people deposit another layer of cladding to encapsulate the waveguide.

People typically do this just to avoid any kind of a surface damage and make the waveguide a lot more robust against any kind of surface contamination. If you want to continue to do the integration, it's also a lot easier because now it creates actually small flat surfaces where you can fabricate other devices on top.

So amorphous materials, again, this is why people still like them for waveguide fabrication. For reasons I've already mentioned, it's a lot easier to deposit these amorphous materials. In addition to the fact that you can deposit essentially any kind of substrate without having to worry about lattice-matching constraints, the deposition rate for amorphous materials also tends to be a lot higher compared to crystalline materials.

So the reason for that is because if you think about it, you want to deposit a single crystalline material. Now what you have to do is that the process goes like-- you have atoms coming from gas space. They essentially adsorb to the surface. They then go through surface diffusion to try to find essentially lattice sites where they essentially can attach to and then actually grow the single crystal.

So in this process, because they rely on this kind of 2D diffusion for the atoms to find the correct lattice sites, you need to have sufficient diffusion length for the surface atoms. Now if you have a very high diffusion flux, then what happens is that these atoms, before they were able to diffuse sufficiently long distances to find the lattice sites, it's going to get disrupted by the incident atoms due to high atomic flux. It's going to get scattered because of very high deposition flux.

And that's why, as you increase the deposition flux, your structure tends to become more defective. And eventually you transition from single crystal to polycrystalline and then to amorphous with an increasing deposition flux. And high deposition rate is actually beneficial from a manufacturing perspective because it means that you can actually make things with high throughput. It takes less time to do deposition. So that's why amorphous materials tend to be a lot easier in this case.

Also, you can actually do the deposition-- you can do the deposition at a lower temperature. Again, for the same reason. High temperature, due to Arrhenius law, it means the surface diffusion is more effective. So the atoms will be a lot more efficient in finding the low-energy sites attached to the lattice.

At low temperature, on the other hand, diffusivity is limited. Then the atom has to more or less settle where they actually land on the substrate, and you tend to get amorphous structure. Again, low temperature is actually beneficial because, number one, it prevents you to disrupt any kind of on-chip structures. Two is that it's also more energy efficient. It's also a faster process. So for all those reasons, that's why amorphous materials is, in lots of cases, considered preferred for waveguide fabrication.

And in fact, if you look at some of the dominant waveguide materials, one of them is actually silicon nitride. You deposit through PECVD, and then you pattern. So that's a very common waveguide material that's amorphous. And also, when I'm talking about the lowest waveguide loss records, they are pretty much all achieved in silicon nitride.

And it also relates to the low loss of the optical materials. So amorphous materials, they tend to have very low intrinsic loss resulting from the fact that despite you have Rayleigh scattering that's present. But Rayleigh scattering comes from atomic-scale inhomogeneities. So because Rayleigh scattering scales wavelengths to the-- scale is essentially structure size to the negative fourth power.

So much smaller length scale of inhomogeneity also corresponds to a much reduced Rayleigh scattering loss. And that's why you can achieve very low-loss waveguides in these amorphous materials.

All right. So basically, for lithography, this is probably the most commonly used technique for fabrication. But there are also some other techniques that are unique only to amorphous materials. So one such technique is called ion exchange. We have talked about ion exchange in the case of strengthening for glass waveguides.

In that case, we are trying to [? indiffuse, ?] say, potassium ion into a sodium-containing glass. Unfortunately, potassium and sodium, all those alkaline ions, they have very similar refractive indices. So when it comes to fabricating waveguides, people usually use different kind of ion combinations.

So one common example is to actually use silver. By [? indiffusing ?] silver into a sodium glass, you can create essentially an optical waveguide in where you have essentially silver substitution. So silver is preferred for a couple of reasons. First thing, we have mentioned that silver has the electron cloud that's very easily deformable.

So a deformable electron cloud optically essentially means that it's very polarizable. So it easily responds to an incident electric field of the light. So it corresponds to essentially a large refractive index. So you can create essentially waveguides with sufficient index difference between the core and the cladding when you substitute sodium with silver. So that's one.

Second thing is that because it's such a deformable electron cloud, it means that silver ions, they tend to have high diffusivity. And that's why, if you recall, when we're talking about the ionic conduction, we said a lot of the superionic conductors or fast ionic conductors, they contain silver or copper ions. Again, because it's very easy to deform the electron cloud, you can easily squeeze these ions in between the solid structure.

And therefore, they tend to have high diffusivity. So they also benefit essentially the throughput of the fabrication. So that's why silver is one of the preferred elements for this fabrication process.

Now specifically, if you look at fabrication scheme here, you start out with essentially just a glass substrate containing sodium. You put a metal mask on top, you do lithography to pattern the metal mask, and then you use the metal mask essentially as a diffusion barrier to allow only ion exchange to take place where you have an opening on the mask.

You then immerse it into a molten salt bath. You can apply electric field to essentially accelerate this kind of a diffusion here. And then, in the end, you can then remove this kind of mask to create a waveguide here. So this process usually, up to this step, essentially creates a waveguide that's actually on the surface, because that's where ions come in.

For many applications, you want the waveguide core to be isolated from the environment to essentially increase its robustness against contamination. In that case, what you can do is you can continue to do a so-called field-assisted burial step that drives the silver ion further into the surface. So then, essentially, you now create this kind of buried or embedded optical waveguide here.

So this is quite common fabrication process for glass waveguides. It's very scalable. It's a simple fabrication process. It also has the advantage that it tends to have a quite low optical loss. Its low optical loss come from two parts here. One is that its index contrast is not as large compared to, for example, silicon dioxide.

Two is that this diffusion process, it creates this kind of a diffused interface. It does not create an abrupt index difference between the core and cladding. This diffuse interface also helps you to alleviate essentially the effective roughness you have and therefore lower the scattering loss. So that's why this is a quite commonly used waveguide fabrication technique. And a lot of the optical waveguide technologies, they rely on this process to create the devices.

All right. Now, next technique that people are also using to create glass waveguides is essentially a nanoimprint process. Some people also call it embossing. But effectively what you would do, for standard nanoimprint process, you basically start out with substrate. You deposit your glass layer. You then deposit a layer of polymer. That's used essentially as a resist for imprint.

You then press a mold with a pattern inverse what you want to get in the end against the surface. You then either apply a UV light, or you apply the heat to introduce reflow of the polymer resist. Once you have the polymer resist defined, you then do an etch pattern transfer to then transfer into the final waveguide material.

This is one way to make waveguides. But for glass materials, there's actually a simpler way to actually create waveguides. Because glass themselves, if you heat them up, they will essentially turn to basically a viscous liquid or viscoelastic solid. So then you can just directly press a mold against glass material.

And heating it up to well above the glass transition temperature, you can then introduce essentially a thermal reflow of the glass materials or viscous flow, which then directly allows you to structure the glasses into a waveguide. So the unique properties of glass enables this much simpler fabrication process using direct nanoimprint without having to use an intermediate resist-based pattern transfer step.

So on the right side, this is just one example. You can see this is actually a waveguide. That's actually created, in this case in arsenic sulfide glass, a pattern using 3D printing technique. Also, one of the advantages here is that the roughness of the waveguide essentially is a direct replica from the mold. So if you can create a high-quality mold with very low roughness, then the resulting waveguide also inherits the property having a smooth surface. And that's beneficial to creating essentially a low-loss fabrication-- for low-loss propagation.

All right. Now, last technique I'm going to talk about here, this is using laser writing. So for what's called direct laser writing or laser direct writing, you can start with a bulk piece of glass materials. You then focus the laser light tightly through basically a lens or microscope objective. And then, at the focal spot, you can induce essentially some refractive index change through a process of so-called two-photon polymerization. Sorry, Two-Photon Absorption, TPA.

So two-photon absorption basically means that if you think about, this is the energy band. This is energy. This is Ec. This is Ev. Normally, in order for a photon to get absorbed, its energy has to be larger than the bandgap.

However, if the light intensity is strong enough, you can introduce so-called two-photon absorption, where instead of having one photon, you use two photons. You combine the energy together to then promote one electron from the valence band to the conduction band. So in this case, you only need to have the photon energy to be larger than half of the bandgap.

Now the unique behavior of this kind of absorption process is that it's nonlinear. So if you think about the absorption coefficient alpha, for typical linear absorption process, which is the regular absorption process that we see in materials, alpha for a given material at a particular wavelength is a material constant. It does not depend on light intensity.

So essentially it means that the amount of photons absorbed scales linearly with the amount of photons you send into a system. That's why it's called the linear absorption process. For nonlinear absorption process for two-photon absorption, the absorption coefficient actually scales with light intensity. So in other words, the amount of light that's absorbed will scale quadratically with the light intensity.

Now, intuitively, that's also easy to understand because this is basically a two-particle process. You need two photons in order to trigger absorption. So that's why its absorption coefficient should scale linearly with density. And the amount of absorption scales quadratically with light intensity.

Now, why I have to mention this kind of TPA process here? So when it comes to waveguide fabrication, if you're trying to write a waveguide in a bulk glass-- so say this is the bulk piece of glass material, and you are trying to focus light down to this very small spot. And the idea is that this very small spot, this absorption induces some material change. You can either increase the index or decrease the index. Then, by translating this focal spot through all the material, you can essentially directly write a waveguide.

Now if you leverage not the two-photon absorption process, by using linear absorption, then it means that, in addition to absorption at the focal spot, the light, as it passes through the material, in this shaded region, you also get absorption. So that means that you cannot spatially confine your material modification to only the focal spot.

On the other hand, if it's two-photon absorption, then because of this kind of quadratic dependence of the absorption with intensity, it means that this absorption is really only significant where you have very high light intensity, which is the focal spot. And essentially, all the absorption in this kind of a material above the focal spot becomes negligible.

So that's why you can actually use this kind of TPA process to create buried waveguides. Also with much higher resolution, again, because you can essentially confine the light interaction volume only to the focal spot very precisely. OK?

OK. So then, depending on the glass material, most glass material exhibits so-called photosensitivity. So if you shine light on it, as long as you put enough energy into it, it's going to absorb light. As it absorbs light, local lattice structure actually gets slightly changed. So due to the local lattice change, it will also change the refractive index slightly.

Some glass materials, their refractive index actually increases upon irradiation. In this case, it's very straightforward to write the waveguide. You just translate it. Essentially, the focal spot trace becomes a waveguide. Some glass materials, the index actually decreases as you expose it to these laser pulses.

In that case, what you do is you go to a so-called depressed cladding design. You essentially write the cladding of the waveguide. You essentially write sort of like a tube. And in the center of the tube where you don't have a light, that's actually where you have the waveguide. You essentially write in this kind of patterns.

So besides fabricating waveguide using this technique here, this is actually a technique now that's getting increasing interest. People like it because it's very versatile. And it's increasingly getting used for many applications, especially to create 3D waveguide structures, which is quite a unique capability that this technique can have.

So besides for waveguide fabrication, I also want to mention that this is also a technique that people actually can use to actually store data in glass materials. So people are essentially also leveraging this photoinduced refractive index change to store digital data.

[VIDEO PLAYBACK]

JUEJUN HU:

So this is a permanent storage. Once you actually write the index change, it's going to stay permanent there in the glass material. They use this kind of femtosecond laser pulse because you can achieve very high intensity in these very short pulses. And that would allow you to actually confine the interaction volume to a very small spot, leveraging a nonlinear absorption.

So now the advantage of using two-photon absorption is that you can write essentially a multi-layer storage, because you don't have any kind of interference coming from light absorption above the focal spot. So you can essentially write multiple layers of this kind of voxels to increase the storage density. It's a 3D storage method rather than 2D storage method.

The motivation there is that people are looking for essentially permanent storage media for the huge amount of data, and this is considered one of the most promising technologies.

So that's about the difference of fabrication techniques. Any questions about that before we move on to talk about applications of fibers and waveguides? All right.

There are lots of applications for this kind of integrated photonics and optics. I will talk about communications, augmented realities. But some other applications that I'm not going to mention in this class in detail. One of them is LiDAR, which refers to a light detection and ranging. It's similar to radar, which radar uses radio wave. LiDAR actually uses laser.

So what it does essentially is send a laser beam or laser pulse, and then the laser pulse is going to get reflected if you see the object. And then, by collecting the reflected pulse, you can then basically calculate what is the distance of the object. So it gives you a way to perform 3D mapping of the surrounding environment with very high accuracy by scanning a laser beam.

So LiDAR has attracted increasing attention in recent years because of the autonomous vehicles. It's considered as a way to actually get very precise three-dimensional data of the surrounding environments, with a lot higher resolution, much higher resolution compared to radar. And also a lot more precise and with much less ambiguity compared to direct imaging or cameras.

So that's why there's a lot of attention to LiDAR there. But they are used not just for autonomous vehicles. They're used in robotics, used for UAVs. They're using for remote sensing. People can also perform chemical analysis using essentially laser spectroscopy techniques, basically based on LiDAR. So it's really attracting lots of attention in recent years.

And the motivation there is that conventional LiDAR essentially uses mechanical components to rotate the laser beam to scan the environment. That's why, if you look at in the early days, when Google has their autonomous vehicles, before Waymo, the older vehicles, on the top of vehicle, they have little rotating things. That's actually LiDAR. It's actually scanning the laser beam by mechanical rotation.

That's very expensive. It cost tens of thousands of US dollars. And that's why this adoption has been severely hampered. But in recent years, people have actually developed ways to replace that whole big mechanical rotating LiDAR with just a small photonic chip, essentially performing the same function. And that's why it's getting lots of interest now for autonomous driving industry.

Sensing, you can use light to actually detect the presence of different kinds of molecules, by doing spectroscopy analysis. So that's also one big area that's actually attracting lots of attention recently. Now what I want to do, I want to focus on communications and augmented reality.

So starting off by talking about optical communication systems. So I'm showing you this Great Wall here because this is what I consider as one of the oldest optical communication systems in the world. Its operation principle is very simple. So if there's no fire, that means that it's peace. Nothing happens.

If there's actually invasion, people will light it up. And essentially you can see it. So now, essentially, it means that something happens. So this is basically a 1-bit or binary optical communication system. All right.

So nowadays, state-of-the-art communication systems still operates on more or less the same principle, except that people do it a lot faster and more efficient. You can still encode digital signals, 1's and 0's, by turning light on and off, on and off. If it's on, it's 1. Off, it's 0. This is called on-off keying. It's the most simple form of doing a digital encoding. On-Off Keying, OOK.

I mean, people have developed more advanced ways to encode information on optical chip, not just only the on-and-off light. You can essentially, for example, divide up light intensity into different levels so that when you send one optical pulse, you can actually encode multiple bits of information, not just like one bit, 1 and 0.

You can also change the phase of the light or change the polarization of the light. All of those parameters of photons can be used to encode digital information. So it can be a lot more complex. But the idea is the same. You essentially turn light or switch light essentially very fast and use that to actually convey information.

So now, if you look into optical communication system, you need to have two basic components. You need to have a transmitter which receives electronic signal, the electronic data, 1's and 0's, and converts it into some kind of optical property change of the photons, some property change of the photons. So it can be on or off, simply.

So this kind of encoding unit is what's called transmitter. It allows you to transfer signal from electronic domain to the optical domain. This signal is then sent down-- this transmitter then sends the signal down through a fiber or waveguide. And then, at the other end you have a receiver, which basically is a high-speed photodetector. It detects a signal and then converts it back into a basically voltage. High-low voltage. Digital signals.

So if you want to do a bidirectional transmission, you want to send signals in both directions, you then need to combine transmitter and receiver in one single unit so you can both send and receive signal. And that's what's called a transceiver. A transceiver is basically a transmitter plus a receiver. Transceiver.

All right. So these transceivers, what I'm showing you here, this is a little picture of one example of optical transceiver. It can all be made on a small chip essentially using this integrated circuit technologies.

Now, why is that important? Why is this kind of optical communication of light to actually send information, why is it important? So if you go back a few decades, people send information pretty much exclusively through copper cables. By essentially directly sending voltage signals across copper cables. What we are seeing today is that these copper cables are getting increasingly replaced by optical fibers and waveguides.

So why is that? I'm showing you this slide basically from Intel. So vertical axis is basically the different hierarchies of the optical links. I mean, for hierarchy, I essentially mean more or less they [? call it a ?] distance. So for this one, this is basically long-distance optical communications. You are sending, for example, signals through a fiber to your home.

So this is a communication distance of the order of hundreds of meters or kilometers, tens of kilometers meters. Now if you go to shorter distances, this could be rack to rack. This could be from one rack of servers to another rack of servers, for example, inside a data center. So this is a shorter distance.

You can go to the extreme. The very extreme is you are trying to send signals from one chip to another chip on the same board. This could be, for example, you're trying to send signal from essentially your CPU to the memory, which are only a few centimeters apart on the board.

So what we also see on the horizontal axis here-- the vertical axis is distance, horizontal axis is basically time as well as the bandwidth. How fast you send your signal. So this is quite an outdated chart, as I mentioned. If you look back 20 years, you are talking about 3 gigabit per second transmission speed. This, of course, increases exponentially over time.

And what we are seeing here is that the optics would gradually penetrate into the domain where, traditionally, people use copper cables to do communication. So rule of thumb is that for longer distances, optics is better. Why? Because you have long distances. In order to send signals through a copper cable, what you have to do, you essentially have to charge up the optical cable. That's how you send voltage signals across the copper cable.

Now, the longer the cable is, you need a higher voltage to properly charge it up. You need to put more charges into it. The capacitance actually increases. So you actually need to put more charges in order to efficiently charge it up.

Also, if the cable gets longer, you have more electrical resistance. So you also have to waste more energy there. So bottom line-- the longer the cable, the less energy efficient it's going to be for electrical communication. So that's why you see that several decades back, people have already replaced copper cables for long-haul telecommunications. You are trying to send signals over kilometers. Decades ago, people are already using light.

But also, if you go to higher and higher bandwidth, essentially faster and faster signal transmission speed, optics also has an advantage. This also comes from the same argument here. When you are trying to do communication through copper cables, you need to charge it up. But if you are trying to go to very high-speed communications, you need to essentially charge and discharge the copper cable at very high speed.

This becomes quite energy efficient because any cables always has an inherent capacitance associated with it. So if you are trying to charge and discharge at a very high speed, it means that you need to spend lots of energy just trying to charge and discharge this cable there. That becomes very, very inefficient when you need to do this repeatedly at very high speed.

Well, for optics, you don't have that limitation here. You can send the pulse through long distances as long as the waveguide or fiber is low loss. Your energy consumption barely depends on the distance you send the signal through. And it also barely depends on how fast you modulate the signal. You can always use the-- in fact, you can use less amount of energy if you go to higher speed in some cases.

So as a result, what you're seeing here, the general trend is that as you go to higher and higher speed communication, optics start to penetrate into a shorter and shorter distance communication links. I mean, despite that this graph is somewhat outdated, this general trend is what I'm trying to show here. Optics is becoming more important as you go to higher high-speed communication.

So now we are not talking about a few gigabit per second. We are talking about, in fact, in this case, terabit per second per communication channel nowadays. And the capacity for optical communication can also be further boosted by using a technique called wavelength-division multiplexing.

So it's not just one single channel the data speed is increasing. You can also essentially throw in tens or hundreds of channels into one single optical fiber. And the reason this wavelength-division multiplexing works is essentially different wavelengths of light or different colors of light, they propagate independently in a fiber.

This goes back to essentially so-called linearity of the system. When the light intensity is not too high, you don't have any kind of nonlinear effects, like two-photon absorption. Then if it's in the linear domain, then all these different wavelengths, there was very little crosstalk in between them. They will essentially propagate independently as they travel down the fiber or waveguide.

So therefore, what you can do is that you can now use a large number of closely packed wavelength where each of these wavelengths you can independently encode one channel of information. So by doing so, you can use the same fiber to achieve much higher capacitance or communication capacity.

So this is one example of how people do essentially wavelength-division multiplexing of optical fibers. You can pack this large number of channels near the 1,550-nanometer wavelength. This is where you have a low-loss fiber channel. So you can essentially squeeze lots of channels around this wavelength range here.

I'm just showing you one example here. This is actually what people have done a few years back, where they were able to demonstrate that one single optical fiber was able to transmit over a petabyte signal per second. A petabyte is 10 to the 15th bits per second-- 10 to 15 bits, which means that it's over a million gigabits. So you are sending over a million gigabits signal through one single optical fiber per second.

So the way they can achieve this remarkably high information capacity is that they actually use WDM techniques. They use a total of 382 wavelength as a separate channel to transmit information. In fact, they also use modes as well. They actually use a multi-mode fiber. They were able to selectively excite each mode here so that each mode can be used as an independent channel to actually send information.

We mentioned that multi-mode fiber usually is not preferred for long-distance communication. That's in the case where all the modes are essentially randomly excited in the fiber. In that case, all the modes essentially travel at different speeds. Then you cannot have a very good signal integrity.

However, in this case, they developed a way to very selectively excite each order mode so that each mode corresponds to one specific information channel. So that's why they were able to achieve this very impressive communication capacity number using essentially one single optical fiber.

All right. So bottom line is that using light, you can achieve a much higher information transmission capacity compared to copper cables. Now, in terms of the examples of this kind of communication, this is one example which has been very widely available. You probably heard about Verizon fiber to the home. This is a technique that, in fact, in my house, we actually use a Verizon fiber to the home.

So they have the so-called Tri Play service. They use one single optical fiber sent to your house, connected to your house. You can have essentially high-definition TV. You can have your phone service. You also have your internet data all through that one single optical fiber.

So to see how that works, I'm going to actually show you a little quick video here. Let me see. Oh. Hold on.

[VIDEO PLAYBACK]

JUEJUN HU:

- You'll get your guaranteed maximum refund.

Now when you look at this video, try to catch what is the wavelength.

JUEJUN HU: Try to catch what is the wavelength that's actually referred to here.

- Actually, it's a variety Fios tech. It brings fiber optics right to your door on three different spectrums of light. You got your 1,310 your 1,490, and the 1,550.

- The light is so clean it's plus 20 dB hot. It's true quantum.

- You should see his truck.

[END PLAYBACK]

JUEJUN HU: OK.

AUDIENCE: What a knowledgeable kid.

JUEJUN HU: Yes. Well, he knows the wavelength, obviously. So how many wavelength number have you heard about? Three. 1,310, 1,490, 1,550. So if you look at how they actually divide up the different wavelength band here, you see that for 1,310-- so if you look at how they divide up the band here, for 1,310, they use it essentially for upstream, essentially uploading of voice and data.

So this includes your phone service as well as internet data upload. And 1,490, this is the waveband used for voice and data download. So this is, again, for your phone and your internet service. At 1,550, they use it for high-definition TV. Now for TV, you only need to download information. You don't need to upload information. That's why there's only one single channel there.

Now if you look at the three wavebands, 1,550 is right at the low-loss band. 1,490 is also pretty close to the low-loss band here. This is 1,550. This is 1,490. 1,310 is used because this is also a low-loss transmission window in the fiber. Turns out, this is also the so-called minimum dispersion window. It means that, essentially, the chromatic dispersion is actually minimum at this particular window here. So that's why 1,310 is also a pretty commonly used wavelength for these kind of communications.

But what you are seeing here is essentially one example where people can use different wavelengths to transmit different channels of data independently. All right. So that's for long-distance fiber-optic communication. As I mentioned, over time, optic communication is increasingly penetrating into shorter-reach communication links.

So one such example is actually data centers. So what I'm showing you here, this is actually a photo of a Google data center. So if you go on to Google, you try to search for something, like search for cat photo, then what's going to happen there is then your query is going to send one to the data centers.

The server is then going to process that, now try to fetch information essentially from either this data center or other data centers around the globe. So it's a pretty complex operation for every search you're actually performing there. There's lots of internet data traffic that's actually involved.

So to efficiently transmit information between the different servers here, that's where actually optical communication comes in. So nowadays, for communicating information between one of these big server racks to another server rack, people used to be using copper cable. But nowadays, exclusively, pretty much, people only use optical fibers to do this kind of communication.

And people are also-- one example for doing this kind of rack-to-rack communication here, people use a technology called active optical cables. So an active optical cable, from the outside, it looks exactly like a copper cable. Its connector would be exactly the same as a copper cable.

So from the outside, it can be used essentially in a technology-agnostic way. You can plug it, use it to connect different server racks without having to worry about whether it's actually a light inside or it's actually electrons running inside. But if you look inside this connector here, what it contains here, this is not a copper cable. It has actually several optical fibers running in the cable here.

And also, in one of these connectors, it actually contains essentially a photonic transceiver. It converts the electrical signal you get from the socket and converts it to optical signal. You then send it through optical fibers. In this case, in this particular example here, there are actually four independent fiber channels. And there are actually four transceivers. So this allows you to essentially increase the communication bandwidth by four-fold because there are four fibers actually involved.

OK. So this is one technology that people commonly use in data centers to connect different racks using essentially what they call active optical cables. This technology has been pretty mature. It has been around for past decade. In more recent years, due to the huge surge of data requirements, especially related to the emergence of AI, people are now trying to push this interconnect to even shorter length scales.

So it's not just rack to rack. People are now talking to essentially communications between the different chips, different GPUs, for example. So in this slide here I'm showing you the architectural tradition for essentially a communication between a large AI computing node. So the key takeaway message here is that AI is actually very communication-intensive.

So for many AI operations, you want to run the task in parallel across many GPUs. But in order to be able to do that, you want to make sure that all the GPUs, they are constantly communicating with each other. So in lots of cases, essentially the efficiency and the speed of computation is not limited by the GPU itself. The bottleneck is actually in the communication-- how well the GPUs can talk to each other.

So conventionally, this is the conventional architecture about how people actually interlink together different GPUs. You essentially have copper cables trying to link all the GPUs into one network switch. The switch essentially is a device that allows you to create different connections and reconfigure the connections between the different units.

And then this switch, then you have optical fibers to use it to connect to essentially other racks. So this is not very efficient because as you can tell, you have this kind of middleman here. In order to communicate from one GPU to another, you always have to go through this kind of a network switch here, which is done electrically.

So as you go to higher data speed, as you go to longer distances, it becomes very inefficient. So what people are trying to push for nowadays is they want all the GPUs to be able to directly talk with each other through light, rather than through copper traces or copper cables.

So that's why now people are now investigating a technology of so-called co-packaged optics. So without going into too much details, what I'm showing you here is essentially people are trying to push optics closer and closer to the electronic chip.

So the earlier generation of optical communication technologies, you essentially will have your ASIC here. Your ASIC is your Application-Specific Integrated Circuit. This could be your GPU, for example. It could be your memory chip. So this GPU or memory chip, they will be mounted on essentially a circuit board, printed circuit board. You have copper traces that link these electrical circuits now to the optical transceiver.

You have essentially a plug here where you can plug in, for example, your active optical cable. And then you use the active optical cable to then communicate to another board. So the downside about this configuration here is that the electrical signal still has to travel through these copper traces on the printed circuit board for pretty long distances, like tens of centimeters, or at least a few centimeters.

So as you go to higher and higher data rate, people are talking about, for example, 100, 400, 800 gigabit. So with higher data rate, this kind of information transmission through the copper trace becomes very energy inefficient. So what people want to do, as I mentioned, for electrical communication, you want to make it shorter so you don't have to charge up very long cables.

So then what people want to do is try to push these optics closer to the electronic chip to essentially minimize the length of the copper trace. And that's why people are now evolving towards what's called the co-packaged optics, where essentially the optical transceiver is packaged essentially in the same package, along with your GPU or your computer chips. They are all inside one single package.

So I just want to show you this video here. This is actually a very recent video. This is something unveiled by NVIDIA just last month. They have demonstrated that they can use this kind of co-packaged optics technologies to enable some of the most powerful AI engines in the world. So let me see. I'm hoping the internet connection works. Come on.

[VIDEO PLAYBACK]

JUEJUN HU: There's no sound. So you see the interface. They have optical fiber bundles going directly into the package, into the ASIC. Going into the ASIC.

Now this is outside. This is outside the server. You see the large number of fiber connections coming out.

OK. So this is what's actually driving today's state-of-the-art AI computing there. Essentially, there's lot of optics communication channels that's involved.

OK. So now next application I'm going to mention here is about augmented reality. I have this little short video here. I don't know why the sound is not working here.

[VIDEO PLAYBACK]

JUEJUN HU: For those of you who are not familiar with AR, or Augmented Reality, I just have a very short video that shows some envisioned applications of AR.

All right. So what AR is basically is you have this kind of see-through glasses where they can be operated just like normal glasses. You can see the things around you. But at the same time, you have essentially this optical gadget that's actually embedded. So they can actually project images into your eye, and it's actually superimposed onto whatever you are seeing in the surroundings.

So if you look closely into how this augmented reality, how it works, essentially you always have a little projector here. These are microdisplay modules with a very high resolution, very small display panels placed on the side. Essentially you have three colors, red, green, and blue. These three colors display and project light into these glass waveguides. You have these sheets of glasses that essentially serve as waveguides.

To couple the incident light into the waveguides, you have essentially some coupling elements here. These are typically diffractive gratings that allow you to bend the propagation direction of light and then allow them to couple into these waveguides here. You have three sheets of waveguides to deliver three different colors.

And then at the output end, you have some outcoupling elements here. Typically these are so-called holographic gratings which, again, redirect the light from being trapped inside the waveguide essentially into free space and then couple into your eye. So that's why you can actually see the projected images here.

At the same time, any kind of a light coming from the surrounding environment, they can also just transmit through this glass deck with minimal amount of attenuation. So that's why you can see superimposed images of this kind of images from this display with the surrounding environment.

Now, about why glass material is important for augmented reality applications. I hope this YouTube video plays.

[VIDEO PLAYBACK]

- Experience augmented reality like never before with SCHOTT RealView. Augmented reality is merging our view of the real world with virtual experiences, and it will be a vital part of our life for daily usage in business and private life. The applications include manufacturing through education to entertainment.

But what role does SCHOTT play in this progress? SCHOTT RealView empowers the optical engine of augmented reality devices. Our glass wafers or substrates for optical waveguides are produced in an innovative fabrication process. Layered over one another, ultra-thin glass slices collect the light waves of the image from the input structure and guide them to a similar output structure, projecting the image to the eye.

Our glass waveguides directly impact the picture quality, simulating reality. The exceptionally high refractive index enlarges the total reflection angle, expanding the field of view. SCHOTT RealView is robust and durable, but also thin and lightweight, suited for everyday use. The excellent glass surface enables high picture quality and contrast, and the added optical coatings create excellent transparency.

The result? The image merges perfectly with the world's real view. Our waveguide guarantees brilliant image quality for full, immersive mixed-reality experiences.

[MUSIC PLAYING]

All these are outstanding qualities are the result of a specialized know-how and innovative fabrication processes of the global SCHOTT network. SCHOTT RealView provides manufacturers with the ability to develop augmented reality devices that will redefine our view of the world.

[END PLAYBACK]

JUEJUN HU: Oh. All right. I think we actually are running out of time here. So I will wrap up about the glass application for augmented reality in the next lecture. Yeah. Sorry for the poor image quality here. That's why we need better optical connections, better interconnects. Yeah. That's why. OK. Anyway, thank you.