Audio file

119010_T2_W5_Science&Society-21StCenturyCommunicationsByJohnMitchell.mp3

Transcript

Welcome and thank you for coming along. It's all a privilege to do the last one of these. And this is the sort of lecture that I would, we've used this I think with undergraduates, probably in the second year, a little bit of an introduction in the comms course. We've used this as a sort of first lecture for MScs to talk a bit of a background about communications programmes. So it's fairly light. There's not many equations. I think there is one towards the end. But it's more to give an idea about some of the background and some of the history and the reasoning behind the communications technologies that we use today. So communicating is a very human act. They often say it's one of the things that really differentiates us from animals. And it really, the social impact that communications have has been phenomenal over roughly the last 200 years or so. really, communication started with the telegraph. 1844, the first telegraph started linking countries together. This is by 1891. And you'll see the old world order, I guess, borne out by the way, certainly the British Empire connecting together outposts. empires from Europe connecting together. And I think what it's hard for us today to understand is the sort of impact that this changing technology would have had. So if you go back to the early 18th century, certainly transatlantic or from South Africa or from India, correspondence would have been by ship, would have generally taken eight weeks, 12 weeks perhaps, from Australia to return to the UK, correspondence. It was a long, slow process. And if you read any of the letters that were sent back, A, they were very expensive, they cost about a week's wages for the average person in Australia to send a letter back to the family in the UK. But you're also paying by the word. So they were very, very direct. So and so has died, such and such has had a baby. Very direct and not the sort of conversations we would have. The telegraph, while still expensive, crossing the Atlantic, turned that many weeks into many minutes, maybe hours. It wasn't instantaneous. You still had relay stations. Someone would write it down and then put it in a pile and someone else would send it on and it would go stage by stage. But you're now talking about... has gone from weeks to minutes for those that could afford it. Although equally, if anyone ever says about teenagers today using text speak and taking vowels out of words, you weren't the first. The Victorians were doing that because they were paying by the letter for the telegraph. So you didn't put extra letters in you. There was a lot of text speak really came from the telegraph. Because you're paying by the letter, you're truncating every sort of word down, otherwise you're paying for it. The next big shift in communications came in the late 18th century. Alexander Graham Bell, although there's some debate about whether there was a German pair who invented the telephone, first developed the telephone. He studied at UCL. He didn't study telecommunications. He hadn't invented it yet, so I guess that's only fair. He studied vocal anatomy. His mother was profoundly deaf, and he was very interested about the way we enunciate words. And so, which I think he moved to Canada and developed the telephone. Early telephones looked like this. No dials, no buttons to press. you'd pick it up and you'd be connected to someone sitting at something like this where they would ask who you wanted to speak to. Now this was quite interesting because actually in this day and age there was quite a lot of intelligence in the telecommunication system, mainly because there was a person sitting here and you'd pick up the phone and ask for, you know, I need to see the doctor and maybe the person sitting here knows, well, Dr. Brown usually plays cards with so-and-so on a Wednesday night. We'll connect you. I know nowhere to find him. It wasn't really for another 100 years that sort of intelligence of you around actually occurred. So it was all a human operator. Up until about the 1930s, when a man called Strowger invented the automatic switch. And if you've ever seen the old telephones with dials on where you can dial in a number, that took the human out of the loop. And the story of Strowger is quite interesting, that he was an undertaker. One of two undertakers in a village in Germany. And he was convinced that his competitor was getting all the better jobs because his competitor's wife worked in the telephone exchange. And so when someone rang up and said, so-and-so's died, we need someone to take care, they would farm them off to the other, to her husband, not to Strowger. And so Strowger looked to invent a system to replace the person connecting the phone. And that was pretty much it for a long time. We didn't have much change for almost 100 years in telecommunications. More and more people got the telephone. It got cheaper and cheaper. It wasn't actually until about 80 years after this we started had reliable connections across the Atlantic. So the first telegraph cable was 1856. The first telephone cable across the Atlantic was not until 1956. or certainly high bandwidth one. And then in the late 90s, we saw mobile telephony just explode. And now we're seeing most of the world's population with some form of mobile device. It's reached about 80% saturation. So we're seeing huge amounts of the population with these devices. and most of them using them for social media. To bring back to the department here and a little bit about that, the department here was founded in 1885, a man called Sir John Ambose Fleming, and he invented the diode valve, the first electronic device. And at the time he was working for a company called the Marconi Company. Gugliomo Marconi was an inventor, an entrepreneur trying to sell wireless. And at the time, this device to receive the radio waves to detect them was very, very basic. And Fleming discovered a way of making those, of detecting those radio signals with a much more efficient device. He's also, if you remember from your school, physics, left and right hand rules of forces. He sort of claimed at least that these mnemonics, these ideas about you've got force, thrust, and current in a cable, these sort of, the left and right hand rules are associated with Fleming. One of the main things he worked on was also power. And again, we're talking just 125 years in December. This year will be the anniversary of the first time we sent radio waves across the Atlantic. So 12th of December 1901, the first letter S in Morse code was sent across the Atlantic by these huge power station and great big aerials that they built down on the coast of England in a place called Poldu, across the Atlantic to Newfoundland in Canada. And the interesting thing at the time was that people believed you could not transmit radio waves that far. Basically, for the reason, actually, In theory, Canada can't see the UK. The curvature of the Earth means that actually you're over the horizon. And the theory at the time was that radio waves travelled in a straight line. Therefore, if I try and send over the horizon, they'll just go off into space. And there's an interesting thing about this about a lot of scientific discoveries and something about engineering, I think, in this. We've sort of come to this modern conception that engineering uses science, that scientists discover things and then engineers build useful things from it. That's actually a fairly new conception of it. And if you go back to this day and age of the Victorians, it was very common that things would be invented that used science that had not been discovered. So they managed to prove that actually the radio wave would go, would follow the curvature of the Earth and arrive in Canada. It wouldn't go off into space. And it was about another 25 years before they discovered what at the time was called the Heaviside layers, now called the ionosphere, basically that there's a charged layer in the atmosphere that radio waves will bounce off, keeping them following the curvature of the Earth, and that they don't just disappear off. into space if you try and transmit like that. And we didn't know that. We didn't know why it worked at the time. There was a long time before we discovered that. So, yeah. This is the famous picture of Marconi sitting in a little room in what was now called Signal Hill in Newfoundland, listening on his wireless set for the letter S to be transmitted. So that's a bit of the old school of communications. Of course, now we use mobile phones, everything's digital, everything's sending data. So there's a little bit about what's really driven that. And the big shift that enabled a lot of modern communications and what differs from modern communications only in about the last 20, 25 years or so, is the move to digital. So what is digital? Most of the signals we start with are what's known as analog. They're waveforms that change all the time. They can change whenever they like. They can take various different values. And that's what many of our pictures start life as, our voice starts life as, music starts life as, these analog samples. analog signals. And historically, we tried to transmit analog. The problem with it is that we don't know much about this. It depends on what I'm saying and what I'm doing. So it's actually very hard to predict. And if you get any sort of noise, any sort of corruption, it's very hard to remove that from the signal. Digital signals, and technically digital just means it's set levels and set times, We have something we can predict in our signal. Now, we often use binary and digital interchangeably. Technically, binary just means it has two levels. Digital, in theory, could have any number of levels. It's fingers, really. But most of our data these days at least starts life as binary. And the thing here is when we know it's either this level or this level, we know when we receive a signal, it can only be one of two things. And so if any noise has changed it slightly, we can make a prediction about, actually, I think that should be this level or not this level. Actually, if we've tried to predict what this is, we know nothing about it. could be anywhere. So that's partly where digital has allowed us to send much more information and improve our communications. So one of the things you need to know to understand this is how we get from an analog signal to a digital signal. So the first thing, if you talk into your phone, that will happen is there was a microphone that takes it from a movement of air and turns it into an electrical signal, a movement of electrons. And then once you've got that, we convert it to digital. And the way we do that is by a two-stage process. First, we sample it. And so what sampling does is say, what value does it have at each point? And so we turn it into what we call a sampled signal. Now, this is getting close to being digital in that it only changes at set points in time, but it can still take any level. So the next step we do is known as quantization. And that basically says, I'm going to create a band. In this case, I've got five bands here. And I'm going to, for each sample I've got, I'm going to say, which band is it in? And so if it's in the first band, I will send something that tells me it's in the first band. If it's in the second, I'll send that information, and so on and so forth. And this, as you can probably see, loses some of the information. Whereas actually, if you sample properly, you don't lose any information, quantization does. But now I've got something I can send a code for. And so we convert each of these levels into a code. And this is what I send. This is the digital message I send. Now, what code I use is somewhat arbitrary. I've just used 0000001, 010. But in theory, as long as you understand at both ends, it's like a language. As long as you both speak the same language, you can communicate. It doesn't matter what language it is necessarily. But you've agree on what the words mean. In this case, you've got to agree what these codes mean, that 000 means I'm down here, 100 means I'm up the top there. And there's things you can do to make that better, but you won't go into. OK? So that takes our analog signal and turns it into a code. So a reasonable question is, well, how big a set do we need? How many codes do we need? How many levels? Probably unsurprisingly, this depends. This gives you a value of how good the signal quality is. So let's see how this works. So this is a famous person who you may well recognize eventually, saying something. And I've basically cut it down into a certain number of levels. So see if you can work out who it is and what they're saying. Any guesses? Churchill? Sounds a bit Churchillian, yes, but no? Very much not. We'll come back. So this has just turned that audio signal into just one, two, three levels. So I could send this with 00000110. I've just got 3 levels. OK. People don't usually get it at that. People don't usually get it at this one either. We'll go to five different levels then. Any guesses? It is Trump. He's got so many to choose from. Yeah. So people usually start getting me around seven levels. So you can see we've got seven levels of audio. By the time we get to 15, it starts a bit more like speech.

I went to an Ivy League school. I'm very highly educated. I know words. I have the best words. I went to an Ivy League school. I'm very highly educated. I know words. I have the best words.

So that's 31 levels. And you can see that that sounds sort of like speech. The typical speech coding is 8 bits, which is this next one. This is what a typical file often is.

I went to an Ivy League store. I'm very highly educated. I know words. I have the best words.

OK. So this is-- your mobile phone actually starts with a lot more, but then compresses it down. which I'll say something about. But this is a sort of standard OK level for audio, 255 levels. And I need eight zeros and ones to send all those letters. So you can see the samples that are involved here. Each one of those samples, we sample 8,000 times a second for this particular clip, which is standard for audio, will pick up a level. You can't really see it, but there's 255 levels here for that code. OK? So that's how we get to digital. But one of the issues with that is I've now turned each level into 8 bits. I've got a lot of information now to send. So there is a... That gives us a problem, but digital also gives us a possible solution. Now, CD players, MP3s used a lot higher quality. Spotify and things like that now are compressing. So actually, you're doing this sort of compressing and ending up with a lot, lot lower bit rates, which is why sometimes your music doesn't sound so great on Spotify. And there's now lossless compression systems and all sorts of things. But this is something we've been doing for a while as well, about once we've got some information, if we know something about it, we can find ways of having to transmit less information. So this is a slide from Fleming. So this is a thing called Morse code. So when the telegraph was going, you'd tap. You'd tap out your letters. Someone would hear it and write it down and then tap it out for the next one. Does anyone notice anything a bit strange about these letters, and about how they're arranged? Yes, at the back there? They do look like zeros and ones, except when we said we're going to sample everything in that code, we gauge each one at 8 bits. Have a look at these. Are they all the same? Are they all the same length? They're different lengths. So E is just a dot. T is just a dash. Y, however, is dash, dot, dash, dot. Why do you think it's chosen that way? Yeah? Yes, exactly. It's very English-centric, I'm afraid. It's very centric on the English language. French letters, French and Spanish letters with accents on have very long. have very long numbers. But E is the most common letter in written English. So the E is the shortest thing to send. If it's the one you're using most and you make it the shortest, you save some time. So this uses the frequency of letters to decide how long the thing you need to transmit is. And the shortest are given to the more common letters, longer to the less common letters. And you can do the same with speech. So we have things called codecs, which code speech. And this is to try and compress the voice down to make it more easier to send. So this is a mobile phone clip. They're test clips. So they say slightly peculiar things. But the idea is that they test out all the different sounds that are made in spoken English. So this is a standard high rate mobile phone clip.

It's a dense crowd in two distinct ways. The fruit of a fig tree is apple shaped.

So you can tell there was a man speaking, there was a woman speaking. If you knew them, you might be able to recognize their voice and you could understand what they're saying. And that's the requirements. So this when some of the early mobile phones, this was one that reduced that by about a factor of four. And what this does is it sends partly information about the form that we saw, but it also interprets that and looks at something that uses a slight model of what the voice box is. So it can sort of tell as you start making a sound how that sound might play out. That's this long-term prediction part.

It's a dense crowd in two distinct ways. The fruit of a fig tree is apple-shaped.

So it still sounds all right. You can tell it's a man and a woman. You can hear what they say. This last one was one that was used in the very early versions of Skype, if people remember Skype. It was a thing before Zoom and Teams and all that sort of stuff, back when voice over internet was a bit of a novelty. And because bandwidth was low, they had some very low bandwidth. And what this one does is tries to create a model of the voice box and tries to just send the different sounds you might make, and then effectively recreate that at the other end. It doesn't send your voice. And it's got that classic sort of teams running out of bandwidth sort of sound to it, where it goes slightly robotic. And Teams does still use versions of this for very low bandwidth. But actually, they get that low bandwidth by effectively knowing the sort of sounds that it's looking for and just saying, now make an oo sound, now make an r sound, and building up words out of that. Which is all quite clever, so it does get the numbers down. So we've taken our analog signal, we've turned it to digital, and now we've compressed it. And by compressing it, we basically get rid of anything we don't need. Now the danger with that is if you get rid of everything you don't need, and anything goes wrong, and you lose some bits in your communications, everything sort of falls apart quite a lot, and you get some really strange things happening. So the other big piece of technology that really enabled mobile phones to work is what's called error correction. And so this is a coding. So we send our data in a certain sort of way, code it up, that allows the system to recognize if some of the data bits have dropped, have disappeared, or that there's gaps in it. This was first really heavily used in CD players. Now, when I first did the version of this lecture 20 years ago, I'd tell people to go and do something with a CD. I doubt any of you own any CDs, let alone have a CD player. My house has two CDs, Taylor Swift and incoming Harry Styles, but that was purely to get access to tickets to the concerts, nothing for anything else. We don't actually own a CD player anymore. But CDs really started this. And the idea was that if you've got a disk and it gets scratched or it gets dirt on it, it would still play. And you can actually put quite a thick strip of black tape on a

CD, and it will still sound fine because it's got this long-term error correction coding on it. And it's mathematical. It's very computer-driven. And the amount of processing that goes on in your mobile phone to enable it to fill in the gaps of what's being said or what's being transmitted is really quite remarkable. So what are we doing when we have error correction coding? Effectively, we're adding a little bit of redundancy. We've got rid of all that redundancy, but now we're adding some planned extra data in that gives us information where if some things go wrong, we can get it back. Now, actually, our minds are very, very good at this. We process an awful lot. And there's examples of this. Has anyone come across the phonetic alphabet before? If you watch any sort of cop shows, they'll talk about whiskey, tango, November. Again, these are words specially chosen not to sound like each other, because N&M can sound similar if spoken quickly. Various letters sound similar, whereas each of these words starting with that letter is very distinct. And that's why they were chosen that way, to make it very clear. And so this is an example of that. You're adding redundancy. It takes longer to say Golf, Oscar, Papa than it does to say GOP. Not a lot, but it gives you some redundancy. Now, the human mind is very good at this. And it uses this information. It uses its knowledge of language. If you're in a noisy environment, it can also use tells and cues by the way people speak and what they're saying. if you're in a noisy bar or a noisy nightclub or something and you're trying to talk to someone, you're not hearing everything, but you're filling in the gaps. This is a form of error correction. This is quite a nice example of it. Anything over 3 letters is spelt incorrectly on here. But this you interpret because A, most of the words here, but also what the brain is actually looking for when you try and speed read something, it doesn't look at all the letters one by one. It actually looks at the first letter, it looks at the shape of the word, you're using your interpretation of the word. You are not actually going letter by letter to listen to spell the word. So basically what it shows, as long as the first and last letters are usually correct, you can read a jumble of the letters within it. So that was how we moved voice from tapping things out to spoken voice in analog on the telephone to digital voice on the mobile. And then, of course, the Internet came along. So the Internet was formed. It was a US defense project. And this was one of those inventions that come out of military spend. So at the time, telephone networks were felt to be very vulnerable because you had an exchange building that served that served a region. So just close to here under BT Tower, there would have been a telephone exchange. And if much of the telecoms in the country came to Tower and then went off international, if you took that building out, you could have wiped out, certainly in the 1960s, you could have wiped out pretty much the entire telephone network of the UK by doing so. Certainly cut the UK off from anywhere else. The idea of the internet was that it's multiply connected, and that if one route disappears, it will automatically find another route to get to the same place. It started with a number of universities. UCL was one of the first locations outside of the US to be attached to it. Satellite link There's a story that a piece of what's now the Northeast Wing was designated as US soil because it was owned by the US. It had to have a bit like embassies are. I don't know whether that's actually true, but it's quite a good story. And by 1980, we were on the internet here. Not that there was much to send. There was no web, there was no internet as such there. you were sending packets between places. But when we saw that the mobile phone took off, the internet took off even quicker. And 100 million within three years of the World Wide Web appearing, the browser we're all used to searching the internet with, more than 100 million people were using it. When you look at how long it took television or radio to reach that same number of people, it was remarkable. And now almost 5 billion people on the planet are online. Most people are accessing it by their mobile phones, so the two have sort of converged. Which is slightly ironic, because actually it's only really the last couple of generations of mobile phones that have been designed for data. Even 4G, it's really only 5G that was really built for data. 4G had good data capabilities, but was still at its heart. a voice device. Whereas I can guarantee probably you're the same, but certainly if I tot up the number of hours I scroll through things and read emails on here, compared to the amount of time I talk on it, talking on it's virtually nothing by comparison, I would imagine, these days. Now, One of the things that has happened with that, though, is I think we now often people associate the internet with being wireless, because virtually all the devices we have do not connect with a wire to anything. Your mobile phone, probably your laptop, your iPads, most things have gone wireless. Although the irony of that, of course, is it's only really the very last bit that is wireless. so my mobile phone will be communicating somewhere. Where is there an access point? Is there one at the back? One in the middle there? Yeah. So the wireless bit will go from here to there. That's the only bit that's really wireless, that white dome thing up on the top there. There's probably a copper cable that runs along that bit, a trunking that will go to a network cabinet somewhere within 50 to 100 metres of here. And then after that, It will be on a technology called optical fibre. And it will stay on optical fibre for pretty much the rest of its journey. It will go from our machine room here to the main UCL machine room in Catherine Lonsdale, probably on an optical fibre, 10 gigabit link, I would guess. And then from that cabinet in the Catherine Lonsdale building, it will go to one of the big internet exchanges, probably ILX, which is in the London, LIX, the London Internet Exchange, which is down in Docklands. And from then it will go off to whatever website I'm asking for, or mail server, it will go off somewhere else in the world via a range of those. Probably transatlantic, it will go under the sea, and it will be carried on a very, very thin strand of glass. So let me tell you how that works. I said there was one equation, Snell's law. So this actually brings us back to what we found out about wireless, that wireless bounced off reflected off the ionosphere and stayed. So if you remember Snell's law from physics, it basically tells you that if you shine a beam of light on an interface between two materials with different refractive indices, the light beam gets bent. This is the basic rule if you want to test this out. Get a glass of water in a clear glass, put a pencil in it. When you look at it from the side, you'll see it looks like the pencil goes like that. It doesn't look like the pencil's straight. It looks like when the pencil goes into the water, it kinks slightly. The reason for that is you've got two refractive indices, air and water. The light travels at different speeds within them, so you see it as if the pencil was bent. That's what Snell's law gives us. What it also gives us is if you've aim your light at this interface with a steep enough angle, the light doesn't come out the other side. It gets reflected. It bounces back. So you can effectively turn this boundary between 2 materials into a mirror. And that's what an optical fiber does. Now, like a lot of things, we knew about this a long time before we could use it. A chap called John Tyndall did this experiment. I've tried to recreate it like this. It's really hard. With a bottle of water and a laser pointer, it works quite well, if you want to try that. Basically, you make a hole in the-- get a big bottle of water, shine a laser pointer through it from one side to a spot, and then drill a hole where that spot is, and the water will come out. And what you would expect if you shine a laser in here is that the laser would come out the other side. What you actually see is that the laser will be trapped in the water. And you can put your hand in the water down here, and you'll see the laser beam. The only tricky bit, really, is to get what's called laminar flow of the water. If the water is all sort of splashing around, the laser gets scattered out of it. If you get a nice, steady pipe of laminar flow of water out of it, the laser will be nicely trapped inside it. And it will do exactly what an optical fibre does, and it guides the light through it. And so we make glass to do that. We make very, very thin strands of glass with two different types of glass, with two different refractive indices, so we can trap the energy of the light in the middle bit. Now, these Pipes of glass are about 120 micrometres, 120 microns thick across. To give you an idea of scale, a human hair is about 500 microns. So the centre part of an optical fibre, the bit that's actually guiding the light, is thinner than a human hair. We then put some plastic on it and some things to give it some strength and some rigidity. But actually, despite being glass, it's actually quite strong. You could go like that with it. If you tie it in a knot, it breaks. But if you pull it, it's actually quite strong in that direction. And we'll band many of those together and turn that into these networks that go around the world. And you'll see, actually, it follows a lot of the same sort of routes that we had with the telegraph cables. So you've got all the transatlantic routes. carrying data across the Atlantic. I always look at this every year and go, I need to update that because no one's got an 8 megabit internet connection anymore. But assuming you did, that was what typical connections were. A DVD is about 5 gigabytes, so 37, so that's a high quality film. would take about 5 gigabytes of data. The current record for an optical fiber is quite a lot more than that. It runs at about 170 terabits per second. That's down one singular thin span of glass. So yeah, roughly the entire Netflix in a second or two, you could transfer over that. Now that's in labs. There are certainly one terabit links across the Atlantic now. I know they were talking about an upgrade to a 10 terabit. I don't know. I can't remember if it's in service. But certainly we are now getting to the stage of transmitting terabits down these strands of glass. And we do that by, rather than turning one laser on and off a lot of times, we do that by using the fact it's light and sending multiple different, well, I'm going to use the word colour. Now, it's not colours that you can see. The light we transmit down an optical fibre is outside of our visible range. It's ultraviolet, infrared light we transmit. But it's the same principle. It's different colours. And so we send some information on a green, on a red, on a blue. We might have up to 192 of these different coloured channels, and then we split them all out at the other end. So that's how we add all of these together. And we are no longer the record. We had the record of 156 terabits. So if you go up, Professor Bavell's lab up on the 6th floor of the engineering building has some of the fastest communication links in the world. This world record changes about every six months or so. Someone goes slightly faster. So back in 2000, we held the record for six months. I think we were 156 terabits per second at the time. But that slowly crept up, and I think we're now at 170. So that's down in optical fiber. Actually, also in the same lab, there's a group that is trying to break the terabit on a wireless link. So at the moment, it's about 10 meters. But they showed last year 936 gigabits per second. So they're about 8% off breaking the terabit barrier. And that still stands as a world record for the data over a wireless link. So it's quite impressive. These are the sorts of numbers we're now getting. It will be a little while before you get that sort of data to your mobile phone. Although, to be honest, if you've got a terabit, you probably wouldn't know what to do with it on a mobile phone anyway. The latest thing they're talking about is actually AR, VR immersive type stuff. And 6G is talking about certainly getting you close to a gigabit to your device to do almost sort of holographic imaging, this very ultra-high resolution type imaging you might want to give you full immersion. Okay, so hopefully that's given you a bit of a tour from early communications through to where we are today. And if you've got any questions, I'd be happy to take them otherwise. Thank you.