MITOCW | MITRES_10_S95F20_0201_300k

PROFESSOR: So now, let's start talking about airborne disease transmission, thinking especially of viral diseases where we expect the transmission to occur through aerosol droplets or at least smaller droplets, which may sediment but are going to be suspended in the air for a significant amount of time.

So the first approximation of such a situation is to assume a well-mixed room, meaning that the air in the room is well mixed.

And even if there is, let's say, one person who's infected and that infected person is breathing, talking, singing, respiring, and exhaling infected aerosol droplets, then these droplets start to spread around the room.

And they do so because the room is-- and we assume that the room is well mixed.

So even though, as I've sketched, there's certainly going to be a higher concentration of the droplets near the infected person, that is there are significant airflows in the room, which induce mixing.

That is partly driven by the flow of fresh air through the room.

Typically, we may have some forced ventilation coming through ducts with fans.

There could be open windows.

And even when a room is not open, meaning that the windows are closed, maybe even the door is closed, there's always some leakage and exchange of air with the outside, which typically happens on the order of hours.

So there's always at least some kind of flow rate.

Also, there's a movement of people in the room, and there's also breathing itself, which basically imparts momentum to the fluid and causes swirling motions of the fluid.

So the situation actually is more like this where there's these flows that are generated either by the ventilation, by the movement and respiration of the people, also by thermal flows.

So when you're breathing, the air coming out of your lungs is warmer, and that tends to rise.

But, also, if it's very humid air and very warm air, then actually the weight of all the humidity and droplets that come with air could actually cause some settling as well.

So you have all these different processes going on.

And your first approximation, just assume that those lead to a well-mixed situation.

And so we will proceed to analyze disease transmission from that assumption, and then we'll come back at the end to consider what would happen if there are actually fluctuations and think about the distance from an infected person and how that might play a role in departures from the predictions of a well mixed room.

OK.

So defining some variables, then.

We'll let C of t be the, basically, infectiousness of the air, if you will.

But more specifically, if we think of virus, it'll be the virions per air volume.

So they're contained in droplets.

And later we will consider what size droplets they may be contained in, but for now, let's just average over everything and say there's some concentration in the air.

We have a production rate of these infectious droplets, which can be broken down into many terms.

So this is the production rate, so this is the number of virions per time that are produced.

So that would be Qb is our breathing flow rate.

So that's basically the volume of exhaled air per breath cycle, so that's how much air is being pushed out in your breathing.

And this is typically-- ranges from 0.5 up to around 3 meters cubed per hour.

So 0.5 is a typical resting breathing rate.

And, in fact, even if you're just kind of calmly speaking and sitting in a room breathing through your nose, that'll be your typical breathing rate.

But if you start exercising and start exerting yourself, then that could go up to maybe around 3.

And so that's the typical range of the breathing rate.

The next thing we need is Nd is the droplet concentration, so this is the number of drops per air volume.

So that's sort of the number density of drops.

So if I have various visualization techniques, I can actually see all the droplets, and I can count them.

I can say how many droplets are in a given volume.

The next thing is I need the volume of a drop.

So as I mentioned, of course there are different drop volumes, and we'll come back to assessing the effects of a droplet size distribution.

But for simplicity, why don't we just take these sort of average size drops that we discussed coming from respiration, And that might be a number around 1 micron in size.

And there's a volume that corresponds with that, which, of course, is the 4/3 pi r cubed where we assume for now only one size r for all drops just for the moment.

So now we have the volume here, so now we sort of what is the total amount of drops.

So the next thing we need now is Cv.

Cv will be the number of virions per liquid volume or per drop volume, so there's certain-- so if I could take the pure liquid in the drop-- let's say it's mucus-- and it's been coming out of your pharynx and has sort of fragmented and taken some virions with it, then that would be the viral load, essentially.

Sometimes that's another word for this.

And the viral load varies with time.

So when you first get infected, at first the viral load is very low in the fluids that you're breathing out.

And then that raises up.

And during the period when you're most infectious, which for COVID-19 and SARS-CoV-2 virus, that time ends up being around a week or so when you-- well, within a few days, you reach the peak infectiousness.

And then when you're at the peak viral load, this ends up being about 10 to the ninth virions per milliliter of fluid when you're at your peak infectiousness for SARS-CoV-2.

OK.

So just to give you a sense, but of course, a lot of times it might be less than that.

But if you're very infected individual, that's kind of a worst case.

And we're going to be interested in calculating safety guidelines and probabilities of transmission, so to be conservative, it's good to have an idea of how big this number can actually be.

So what we just calculate here-- so, basically, this Nd, Vd is the amount of drop or liquid volume per air volume, so that's essentially the volume fraction of liquid.

When we go times Cv, we're essentially getting the number of virions in the air per volume of air.

And then Qv is the volume per time, so this is basically virions per time.

And then one other factor that we should also consider is what if we're actually filtering those droplets right at the source?

And that'll be the case if you're wearing a mask.

So this is an important quantity we'll come back to.

This will be the mask penetration probability for a droplet.

So this is-- of course is size dependent, but we'll come back to that.

But for the moment, we're just saying it's one drop size.

And so for the size of drops of interest, we're asking, do they go through the mask?

So 1 minus Pm is also called the filtration efficiency.

Yes.

So a very good mask might be 99% of droplets are filtered.

A very poor cloth covering might be 10% of droplets are filtered, and we'll come back to that in just a moment.

So this here is our production rate capital P.

And already with the variables that we've written down here, we can write down a mass balance for the virions.

So virions are being produced.

They end up in these droplets.

The droplets are being swept out of the room at a flow rate Q, and the room has a volume V. So if I write down just the conservation of mass, making sure that I'm not losing any virus yet-- I'm not allowing them to stick to the walls or do anything else just yet but just looking at the mass balance of one infected person breathing out.

This is, I should say, production rate per infector, or an infected person.

So if there are more infected people, then you'll have this production rate for each person.

They also might be at different stage of the disease, so maybe the viral load will be a little different for each person.

But let's not worry about such details right now.

We want to keep things general.

So for the mass balance, we write down the total number of virions in the room and how that changes in time.

So that'll be the concentration per room air volume, which is well-mixed times the volume of the room per time.

So this is the change in the number of virions per time, and that can change in two ways.

One, we have the production P. So for every infected person, we have production P, but this will be-- I should say virions per air volume per infector.

So if we want to think about having multiple infected people in the room, we can always just basically increase this concentration.

A well-mixed room, it doesn't matter where people are placed.

You're just getting more and more droplets, and it's assumed to be mixed.

So we produce at a rate P, but then the outdoor flow is taking away droplets and, hence, removing virions at a rate Q. So we have a Q times C removal rate.

So this is our equation.

So let's divide through by V.

So we can write this as dC dt is P/V minus, and then Q/V I'll write as lambda a C.

So lambda a, which is Q/V, this is the outdoor air change or exchange rate.

So that is the rate at which the entire volume of the room is replaced with outdoor air.

So the outdoor air is refreshing the air in the room at this rate Q/V. OK.

And so that's what appears here.

And if you compare these two, you can see that this is dC dt.

This is C times lambda, so lambda is units of 1 over time.

It's a t.

This is also sometimes called-- the ACH is air changes per hour if you write it in per hour.

So that's a typical way that this is written.

And, in fact, while we're just talking about, this is a very important concept.

So lambda a is around 0.3 per hour, so roughly every 3 hours for a closed room.

So closed room or what you might call natural ventilation where there's no attempt to deliver air to that room.

Of course, this number depends on the tightness of the construction and whether there is cracks in the windows, whether doors are being opened to the hallway.

So, of course, that's not a perfect number, but that's a rough estimate how quickly air is escaping from typical construction.

But then it can be, also, in a different range.

And this a very important parameter for the theory, so let's pause just to look at some of the numbers.

So it's in the range-- it's typically 3 to 8 per hour for mechanical ventilation.

So this could be-- it could be open windows with fans blowing in and out, which might give you 3 or even 6 on this number.

It could also be a ventilation system, which is delivering fresh air to the space.

And for typical classrooms, offices, and even homes, this is a typical range.

So for example, if it's 3, then that would be every 20 minutes the room gets its air to be fully exchanged.

But of course, if you have situations where you need to have better air quality and you have more risk of, say, transmission of disease or passage of pollutants or contaminants, then we need higher values.

A typical number in the United States for hospitals is 18 air changes per hour.

And then it can be even larger.

So if you have a laboratory, which is dealing with toxic chemicals or even, let's say, virus and pathogens, then you need even higher air changes.

And typical rates, then, can be as high as 20 to 30 for labs that are dealing with toxins of various types because-- and any airborne toxins have to be quickly removed so that if they happen to be leaked into the air from your experiment or from your hood, they need to be quickly sucked out.

Also, even parking lots where you have cars in an enclosed space that are generating carbon monoxide and other fumes, which have to be quickly rushed out, parking lots tend to have this number around 30.

So that's a full air change of the entire room in 2 minutes.

That's a very fast flow rate.

So this is kind of the range of this lambda a.

That's, obviously, a very important parameter.

So let's now solve this equation here.

So, first of all, you can see when dC dt is 0, then the steady state is just P over lambda a V.

But if I write Q equals lambda a V, I can see the steady state is just P/Q.

So I can write the solution like this-- that if we're given time dependence, that the P/Q is the steady state.

And if I started out with my initial condition, was that C of 0 equals 0.

So let's say time equals 0 is when the infected person enters the room and starts breathing.

And then there's a mixing process and there's a build-up of concentration until there's a balance between the production of virus, virions, or infectious air and the removal of infectious air by the ventilation.

And that gives you this ratio P/Q. But the way the relaxation happens, though-- if you balance these two terms here, that's just an exponential decay with a decay rate lambda a, so e to the minus lambda a t.

So this is basically the way the concentration builds up.

If I plot this, then at a certain time here, which is sometimes called Tres-- in chemical engineering, these kinds of models are commonly used to design chemical reactors.

In fact, this kind of model is in chemical engineering called a continuous stirred-tank reactor.

We don't worry about the details, but we have a flow of some reactants and various chemical species going into a tank.

We assume it's well mixed and then it leaves, and this is the mass balance that we use.

And the residence time, Tres, is the inverse of lambda a in this case.

So that is the-- as soon as you know the volume and the ventilation flow rate, there's a typical speed going through here.

And the time that fresh air spends in the room interacting with all the droplets and people and then leaving, carrying some of those droplets, is this.

And in a simple model like this, it's just an exponential relaxation at that time scale approaching the steady state, which is P/Q. So there's always that kind of balance which is reached.

Now let's also ask ourselves, briefly at this point, how reasonable is the well-mixed approximation?

We will come back to this and analyze it much more carefully, taking into account all the different processes I described at the beginning, including breathing, motion of people.

But let's just think about the motion caused by the airflow itself.

So in the case of mechanical ventilation where that flow rate can be sometimes rather high, that can be a significant source of mixing in the system.

And the way we'll think about that is by writing down the typical velocity of the air due to the outdoor air flow.

That is Q. We can write it as either Q divided by the area, so some kind of representative area of the room.

It could be, let's say, the floor area, but depending on the shape of the room, it might be a little bit different value than that.

We can also write this as Q times H over V, where V is the volume of the room and H is some characteristic height, like, for example, a ceiling height of the room.

OK.

So this is the mean airspeed due to ventilation.

And we'll come back to a more deeper investigation of these fluid mechanics of the room.

But as a first example of what we'll be interested in is we'd like to calculate the Reynolds number due to the airflow.

So I'll put a subscript a there.

And that is the typical velocity times a length scale, which could be, let's say, the height of the room or some linear length scale of the room, depending on the direction of the airflow.

And then the kinematic viscosity of air, so I'll write that as new a.

So this here is the Reynolds number.

This is basically telling us how important inertia of the fluid is compared to viscous stresses that slow the fluid down.

So, basically, it tells us how quickly-- how much of a tendency there is for momentum to be carried in the fluid, which then leads to sort of swirling motions and complex flows, as I've sketched here.

And this new a is the viscosity of the air that we've already talked about when looking at Stokes flow but divided by the density of the air.

OK.

So that's the kinematic viscosity.

And for air, the kinematic viscosity is 1.5 times 10 to the minus 5 meters squared per second.

And so if I plug these numbers in and I pick, let's say, a typical ceiling height of maybe a few-- I think this-- if the H is of order 3 meters-- or 2 meters might be a typical scale.

Just to get an approximate sense of the scale here, the Reynolds number will be varying from around 50 or tens, up to 5,000, which would be in the case of very fast ventilation like the 30 ACH.

So this would be if we have 0.3 ACH up to around 30 ACH.

That's just a rough number.

And from fluid mechanics, we know the significance of these large Reynolds numbers is that the flows really do look a bit like I've shown here.

So, basically, when the air is sitting in the room-- you know this from looking at the smoke from a candle or other flows that you can visualize-- it's not just a sort of uniform flow.

But instead there are all these sort of plumes and swirls and vortices.

And when the Reynolds number is on the order of tens, when there is a motion, it tends to lead to a shedding of a vortex and to some kind of swirling flows.

And when the Reynolds number gets up as high as several thousand, then in most geometries that starts to lead to a transition to turbulent flow.

And that's when the flow is getting so complicated there are eddies of different sizes and very rapid mixing.

So I just wanted to show this right at the beginning of the discussion to point out that for typical flows we would expect that there is going to be a decent amount of mixing occurring just because of the ventilation.

And now if you add to the fact that people are breathing, imparting momentum to the fluid, we have people moving and other kinds of activities in a room, that all of those processes lead to giving us a reasonable assumption of a well-mixed room.