slide1:

Today, we're beginning a very unique chapter in this course, focusing on ultrasound imaging. This is exciting because ultrasound is rather cost-effective, and becomes increasingly more important due to hardware improvements and AI techniques. In this lecture, we'll go through the fundamental principles behind ultrasound: what it is, how it works, and how it behaves in biological tissues.

Then, in the next lecture, we'll move on to ultrasound imaging modes — how the system collects, processes, and displays data. So, today is all about understanding the principles, the physics, and the key ideas that make ultrasound imaging possible.

slide2:

Here's our course schedule to remind you where we are. After these two lectures on ultrasound, we'll finish the green book, which means we'll have covered everything from X-ray, CT, MRI, PET, to ultrasound imaging.

After that, we'll move beyond the textbook to discuss optical imaging. This topic is not included in the green book, but they're very important to know. Optical imaging is fascinating, and like US imaging it becomes increasingly relevant with advancement of AI methods in this area.

For now, your main focus should be on the ultrasound chapter in the green book. It's about twenty pages long, so it's quite manageable. I recommend reading it carefully — it's actually quite an enjoyable read, almost like a story.

So today, let me walk you through the first part of that chapter — the foundation that connects the physics of sound with medical imaging.

slide3:

In this lecture, our main focus is on the physical principles of ultrasound. We'll start by understanding what an ultrasound wave really is — how it's generated, how it travels through biological tissue, and how it interacts with different types of materials inside the body. These are the physical foundations of ultrasound imaging.

Then, we'll shift to the engineering aspects — how we actually generate and detect these waves using an ultrasound transducer. You'll hear me mention the term "transducer" a lot today. The word transduction means converting one form of energy into another. In our case, the transducer converts electrical energy into mechanical vibration to generate the ultrasound wave. When the reflected wave comes back, the same device converts the mechanical vibration back into electrical signals. So it's a two-way process — we transmit and receive using the same component.

We'll look at both single-unit and array-based transducers, and we'll talk about different types of resolution specific to ultrasound — axial resolution, lateral resolution, and spatial resolution. Each of these determines how finely we can distinguish structures in an image.

We'll also talk about imaging contrast, which is very important for ultrasound. You'll learn about natural contrast, as well as how we can enhance contrast using microbubbles. Microbubbles are a very clever and modern idea — tiny gas bubbles injected into the bloodstream to improve visibility in ultrasound. They're not discussed much in your textbook, but I'll expand on them to give you a sense of current developments in ultrasound imaging.

So this is our outline for today. And as always, for every imaging modality we study, pay close attention to the outline — it tells you exactly what to focus on. Try to understand each of these key ideas clearly — the wave, the transducer, the resolution, and the contrast — because together, they form the foundation for everything that follows in ultrasound imaging.

slide4:

Okay, now let's talk about medical imaging again, focusing on the five major modalities we've been studying.We've already gone through X-ray imaging, nuclear imaging, and MRI — the three big ones. The next one is ultrasound, and after that, we'll finish with optical imaging, which, as you'll see later, is also very unique and quite powerful.

Ultrasound imaging has several clear advantages. First, it involves no ionizing radiation. That's very important because, unlike X-rays or gamma rays, ultrasound uses mechanical sound waves, not electromagnetic radiation. So there's no risk of DNA damage and no concern about radiation-induced cancer.

Another major strength is speed. Ultrasound provides images in real time — you can literally move the probe and watch the image update instantly. The transducer, or ultrasound probe, can be gently moved across the body, for example, over the abdomen to monitor a baby during pregnancy, or over the chest to examine cardiac function. It's a quick, flexible, and low-cost imaging modality.

Of course, like every imaging method, ultrasound has its limitations. One limitation is the penetration depth. It can't go very deep into the body, especially through hard tissue or bone. X-ray imaging, in contrast, is excellent for bone-tissue or tissue-air interfaces — like lungs or bones — because those materials have very different densities and attenuate X-rays differently. Ultrasound, on the other hand, depends on acoustic impedance, which determines how sound travels through a material. When there's a large mismatch in impedance, such as between air and bone, most of the ultrasound energy reflects back instead of penetrating.

That's why it's difficult to image structures behind the lungs or skull. To get a clear image of the heart, for example, you have to find an acoustic window between the ribs — an area where sound can travel through soft tissue effectively.

Ultrasound imaging mainly works on the principle of reflection and scattering. The system sends out sound waves; those waves hit structures inside the body — tissue boundaries, organs, or small particles — and the echoes are detected to form an image. There are also transmission-based ultrasound methods, but those are less common in clinical imaging.

Beyond diagnosis, ultrasound is used for image-guided procedures, such as biopsies and minimally invasive surgeries, where live ultrasound images can be registered with prior CT or MRI scans to help guide the doctor precisely.

And in recent years, ultrasound has even found a role in neuromodulation — using focused ultrasound to stimulate neural tissue, potentially helping to treat conditions like depression or Parkinson's disease. This is a very exciting and emerging field, showing that ultrasound is far more than just an imaging tool.

slide5:

Now, I know DARPA has funded huge research programs in this new frontier of ultrasound technology. It's important to remember that ultrasound is not an electromagnetic wave. This makes it fundamentally different from X-ray, CT, nuclear imaging, or MRI.

MRI, for example, uses radio-frequency (RF) signals, which are simply alternating magnetic fields — still part of the electromagnetic spectrum. X-ray and nuclear imaging use high-energy photons — X-rays and gamma rays — which can behave like both waves and particles. All of those imaging types relate to Maxwell's equations if we describe them as waves, or to Boltzmann's equations if we describe them as particles.

But ultrasound is different — it's a mechanical wave, not electromagnetic. It physically vibrates particles in a medium — air, water, or biological tissue. That's why it needs a medium to travel; it can't move through a vacuum.

As you can see in this frequency chart, the audible range for humans is roughly from 20 hertz to 20 kilohertz — that's what you and I can hear, like my speaking voice or birds singing. Ultrasound goes far beyond that, starting from around 20 kilohertz and reaching into the megahertz range.

In medical imaging, typical frequencies are between 2 MHz and 20 MHz, depending on the application. Lower frequencies penetrate deeper but give lower resolution; higher frequencies provide sharper images but can't go as deep. Industrial and destructive-testing ultrasounds can reach up to hundreds of megahertz.

So, unlike the sound waves we can hear, these ultrasound waves are way too high-frequency for our ears. And with very high-power, focused ultrasound, we can actually go beyond imaging — we can cut or heat tissue, like a kind of ultrasound knife. It's used in certain treatments to destroy tumors or cauterize tissue.

But in this lecture, our focus remains on diagnostic ultrasound imaging — understanding how these mechanical waves are formed, transmitted, and received to create medical images.

slide6:

Now let's look at this picture, which illustrates a classic experiment that helps us understand what sound really is. This is Robert Boyle's experiment, one of the earliest demonstrations of how sound needs a medium to travel.

So imagine this setup: we have a ringing alarm clock placed inside a large glass chamber with a valve on top. At first, the clock is ringing, and you can clearly hear the sound through the air inside the container. Then, we connect a vacuum pump and start to remove the air from the chamber.

As the air is pumped out, the sound of the clock becomes weaker and weaker until eventually, you can hardly hear anything at all. The clock is still ringing, but with almost no air inside the chamber, there's nothing to carry the sound waves to your ears.

Then, when you open the valve and let air rush back into the chamber, the sound suddenly becomes loud again.

This simple experiment beautifully demonstrates that sound waves require a medium — such as air, water, or tissue — to propagate. Unlike light or X-rays, which can travel through a vacuum, mechanical waves cannot.

So, in ultrasound imaging, the same principle applies: the sound waves travel through biological tissue, reflecting and scattering as they go, allowing us to measure and form images. Without a proper medium — say, if there's air between the transducer and the skin — no ultrasound signal would be transmitted effectively. That's why, as you'll see later, we always use ultrasound gel to eliminate air gaps and ensure good acoustic coupling between the probe and the body.

slide7:

Now, let's take a closer look at what an ultrasound wave actually is.Here, the medium — whether air, water, or biological tissue — is made up of countless tiny molecules. When a transducer or some vibrating source starts pushing on those molecules, it moves them back and forth along one direction — say, the z-axis.

When the transducer pushes the molecules forward, they become compressed, and when it pulls back, they spread apart. This alternating process of compression and rarefaction creates a longitudinal wave that travels through the medium.

If you look carefully at the top figure, the denser areas represent compressions, where molecules are packed tightly together and pressure is high. The lighter areas show rarefactions, where the molecules are spread apart and the pressure is lower.

Now, something very interesting happens here: the wave itself travels forward, but the individual molecules don't actually move along with it. Each molecule simply vibrates back and forth around its equilibrium position — it oscillates, but it doesn't go anywhere overall.

So, the sound wave is really a transfer of energy through the medium, not a transfer of matter. The wave moves, but the particles stay roughly where they are, just vibrating locally. That's an essential idea to keep in mind when we talk about how ultrasound propagates through tissue — the energy travels forward, not the molecules themselves.

slide8:

Now that we know what a wave looks like, let's think about how a wave is formed. This is a really fundamental question — what makes a wave a wave?

To understand this, we need to bring together three important physical principles.

First, the conservation of mass, which tells us that mass cannot be created or destroyed — it must remain constant within the system.Second, there's the relationship between pressure and volume — when you squeeze something, the pressure inside changes in a predictable way.And third, we have Newton's second law, which relates force to mass and acceleration.

If you imagine a perfectly rigid body, and you push on one end, the entire object moves together as one piece — there's no wave because nothing inside deforms or compresses. But in real materials, there's some elasticity — things can compress slightly and bounce back. That elasticity allows small disturbances to propagate as waves.

So, when you push on a small region, it compresses and transmits that force to its neighbors, setting off a chain reaction that moves through the medium. The result is a wave.

Now, if we were talking about electromagnetic waves, such as light or X-rays, we'd describe them using Maxwell's equations. For mechanical waves, like ultrasound, we instead rely on these three mechanical principles: conservation of mass, the pressure–volume relationship, and Newton's second law.

In more advanced, graduate-level courses, we can actually derive the full wave equation from these laws step by step. But for now, the key point is understanding the intuition — waves exist because a local disturbance in pressure or motion creates a self-sustaining pattern that propagates through the medium.

And for me personally, this is the fun part — not just using formulas, but really understanding why these waves behave the way they do. Once you grasp these core principles, you can always work out the details later. The details may change from one problem to another, but the underlying physics stays the same. That's what gives you real insight — and the ability to think beyond just using a device as a "black box."

So now, with this foundation, let's look specifically at how we form an ultrasound wave.

slide9:

Alright, so how exactly do we form an ultrasound wave? Let's break this down step by step.

Think about how ordinary sound travels through air. You can hear my voice because the sound waves — or pressure waves — travel from me to you through the air. Now imagine replacing the air with water molecules, or even biological tissue. The same principle applies.

Let's picture a small volume element, or what we might call a voxel, inside that material. Under normal conditions, that voxel is in equilibrium — it has a certain pressure, volume, and position. Nothing's changing.

But now, suppose we push on one side of that element toward the right. This causes a displacement, denoted by the lowercase letter w. When we push it, the volume of that little element decreases because the particles are squeezed closer together.

According to the conservation of mass, the total amount of material inside that voxel stays the same — but the pressure inside increases. That higher pressure creates a net force on neighboring elements, pushing them forward.

By Newton's second law, the net force generates acceleration, which we call a. The acceleration then changes the velocity, denoted by u, and as the velocity builds up, the displacement w also changes further.

So, you can see that these quantities — displacement (w), velocity (u), acceleration (a), pressure (p), and volume (v) — are all linked together. They interact continuously as the disturbance moves through the medium.

This is not like pushing a rigid stick, where moving one end instantly moves the other. In a rigid body, everything moves together. But in an elastic medium, local motion depends on these dynamic relationships — one region moves, affects the next, and that disturbance travels as a wave.

So, an ultrasound wave forms because of this ongoing interplay between pressure, volume, and motion. Push a little, the volume shrinks, pressure rises, acceleration increases, and that disturbance keeps propagating through the material — and that's your ultrasound wave.

slide10:

Now let's take a closer look at how these relationships actually work.When we push on one side of the medium — for example, this piston here — we create a velocity, which we call u. The displacement, denoted as w, shows how far the boundary moves. In a simple case, displacement is proportional to velocity times time. Mathematically, we write this as w equals u times t.

That's a linear relationship, meaning if you double the velocity, the displacement doubles in the same time period.

But, on the other side of the same element, the motion doesn't happen immediately. The molecules there need time to respond, and that brings in acceleration, which we represent as a. The displacement caused by acceleration is a second-order relationship, expressed as w equals one-half times a times t squared.

So the first one — w equals u times t — is linear, and the second one — w equals one-half a t squared — is quadratic in time.

This means when you push on one side of a small tissue element, that side moves first, while the far side lags behind slightly. Because of this delay, the material between them becomes compressed, and pressure increases.

This process isn't perfectly synchronized — one side moves first, then the other side catches up. That's why we get a pressure increment. The delay in motion creates compression, and that's how the ultrasound wave starts to form.

So, putting it all together: the combination of displacement, velocity, and acceleration, together with changes in pressure and volume, gives us the basic mechanism of wave formation in tissue.

slide11:

Once we start driving one side of the material — continuously pushing and pulling — the wave begins to propagate through the medium.

Each small element of the tissue responds a little later than the one before it. So, if you look along the depth of the tissue, every point is moving up and down in the same pattern, but with a small phase delay.

That delay is what makes the wave travel forward. You still see the same waveform shape, but each point along the medium is slightly out of sync — slightly shifted in phase.

That's what we call wave propagation — the motion is passed from one particle to the next, not instantaneously, but through a sequence of tiny, delayed reactions.

If the material were completely rigid, like a solid stick, every part would move together at once, and there would be no wave. But because biological tissue is elastic, parts of it move at slightly different times — and that's exactly what allows the wave to form and travel.

Mathematically, if we focus on one physical quantity — say displacement w, or pressure p — and express all the other variables, like acceleration and velocity, in terms of that single variable, we can combine them into a single second-order differential equation.

That equation is called the wave equation, and it beautifully describes how the disturbance moves through the medium as a function of both space and time.

slide12:

Alright, now we've arrived at one of the most elegant equations in all of physics — the wave equation. In one dimension, it looks like this: partial squared W over partial Z squared equals one over C squared times partial squared W over partial T squared. So what does that mean? It means that the second derivative of displacement, W, with respect to space, Z, is equal to one over C squared times the second derivative of W with respect to time, T. In plain words, the way the wave changes in space is directly linked to how it changes in time. Space and time are tied together through this relationship. You can even think of it conceptually as "space equals time" — not literally, but in a second-derivative sense. This symmetry is what gives rise to those beautiful, smooth, repeating sinusoidal waves we see in sound or vibration.

Now, let's talk about the speed of sound, denoted by C. C depends on two key material properties — density, which we call rho, and compressibility, which we call kappa. The relationship is given by this equation: C equals one divided by the square root of kappa times rho. Here, rho — written as the Greek letter ρ — represents the density of the tissue, and kappa — the Greek letter κ — represents compressibility, which is actually the inverse of the bulk modulus, or in other words, how stiff the material is.

Now here's the interesting part. The stiffer the tissue, the smaller the compressibility. And when compressibility is small, the wave travels faster. So, in biological tissues, sound typically moves at about fifteen hundred meters per second — that's roughly the speed in soft tissue or water. But in bone, which is much stiffer, the speed can reach around four thousand meters per second. And clinically, this becomes quite important. For example, a malignant tumor tends to be harder than normal tissue. That means sound travels faster through it. So in ultrasound elastography, we can actually measure that difference in speed to assess tissue stiffness — and that helps identify abnormalities like tumors.

So this simple equation — the wave equation — doesn't just tell us how sound moves; it also connects directly to the mechanical properties of the body — its density, its stiffness, and its elasticity. That's the physical foundation of how ultrasound works.

slide13:

Now I want to mention two important types of mechanical waves — the so-called P-wave and S-wave, also known as longitudinal and transverse waves. Remember what we discussed earlier — when you push a volume of air in one direction and then release it, the molecules move back and forth along that same line. This back-and-forth motion produces a longitudinal wave, where regions of compression and rarefaction travel in the same direction as the wave. That's what we call a P-wave, or pressure wave. In ultrasound, this is the type of wave we use — the longitudinal wave traveling through soft tissue.

Now, if we compare that to a transverse wave, the vibration direction is completely different. Instead of moving along the same line, the oscillation happens perpendicular to the wave's direction. So if the wave moves horizontally, the vibration is up and down. This is exactly what happens with electromagnetic waves, like light or radio waves — their electric and magnetic fields vibrate at right angles to the direction the wave is moving.

So, to summarize: in a longitudinal wave, the vibration and propagation directions are the same, while in a transverse wave, they are orthogonal. Both are just different ways energy can propagate through a medium.

slide14:

Here again, we see how a longitudinal wave behaves in ultrasound. The sound originates from a source, like a vibrating tuning fork. In fact, in the old days, doctors would use a tuning fork to test hearing — by striking it and moving it near the ear, they could tell if the patient could sense sound through air or through bone conduction. That tuning fork generates a longitudinal sound wave, where the compressions and rarefactions travel along the direction of vibration.

Now, both longitudinal and transverse waves have some key characteristics. They each have a wavelength, which is the distance between two corresponding points — like from one peak to the next, or from one compression to the next. That's the spatial period of the wave. They also have amplitude, which represents how strong the wave is — the height of a peak or the depth of a trough.

If you look at the shape of these waves, you'll notice they form a sinusoidal pattern — this is not a coincidence. Sinusoidal waves naturally arise when solving the wave equation, whether it's for mechanical or electromagnetic waves. Mathematically, they are fundamental solutions. And in Fourier analysis, we use sinusoidal components — sines and cosines — as building blocks to represent all kinds of signals. So in a way, every complex sound or image signal can be broken down into a sum of simple sine waves.

Finally, remember that the speed of sound is directly connected to the wave equation. The constant you see there — 1 over c squared — represents how the wave travels through a given medium.

slide15:

So when we solve the wave equation, we find that the constant c represents the speed of sound in the medium. And this speed depends entirely on the material properties — how dense and how stiff the material is. For example, sound waves travel very differently in air, in steel, or in biological tissue.

If we look at this table, we can see that the speed of sound in air is around 330 meters per second, in lung tissue about 600 meters per second, in fat about 1,450 meters per second, and in soft tissue around 1,540 meters per second. Water is very similar to soft tissue, around 1,480 meters per second, which is why we often use water-based gels in ultrasound coupling. In muscle, the speed is a bit higher, about 1,600 meters per second, and in bone, it reaches as high as 4,000 meters per second.

So, the denser and stiffer the material, the faster the sound travels through it. But within the same medium, like soft tissue, the speed stays constant — it's a property of the material itself. What we can vary is the frequency of vibration. When we increase frequency, the wavelength becomes shorter. So, higher frequency means shorter wavelength, and lower frequency means longer wavelength.

This relationship between frequency and wavelength is fundamental in ultrasound physics — it helps determine both resolution and penetration in imaging.

slide16:

Now, it's nice for you to know this — as we look deeper into how a sound wave propagates, we can see that when the wave travels into a field, energy is actually being carried through the medium. Earlier, we talked about the wave equation in terms of displacement, and if you take the first derivative of displacement with respect to time, you get velocity. That's what's shown here in this equation: u z equals dW over dT.

A typical value for u z, the particle velocity, is around one to ten centimeters per second, which is much smaller than the speed of sound, c. So that's something to remember — the medium's particles move slowly, even though the sound wave itself travels fast.

Now, because the sound wave is propagating, it also introduces pressure into the field. That pressure at a given point can be computed as shown here: p equals rho times c times u z. It's a very straightforward relationship — the pressure is proportional to the density of the medium, the sound speed, and the particle velocity.

Since the wave equation can be solved, and as I explained before, both the pressure and the velocity take on sinusoidal waveforms, we can represent each of them using an oscillating form — exponential or sinusoidal — like this: p of t equals p naught e to the j omega t, and u z of t equals u naught e to the j omega t. So now, you have both quantities — pressure and velocity — in a sinusoidal oscillation form.

From here, we can define intensity, represented by I. The average intensity over a period, capital T, is defined as the product of p and u, averaged over time. Now, let's think about this physically. Pressure, p, has the unit of force per unit area, so you can think of it as force distributed across a surface.

If you apply a force to move something over a certain distance, that's work. So, this is essentially the work you did. But notice that in this case, we're not talking about displacement or distance directly — u is the rate of change of displacement, or in other words, velocity. That means we're really talking about the rate at which work is done — and the rate of doing work is power.

So, if you multiply pressure (which is force per area) by velocity (which is displacement over time), what you get is power per unit area, and that's exactly what intensity represents. It's how much power, or energy per unit time, is transmitted through a certain area in the medium.

So, putting it all together:Pressure, p, corresponds to force per area.Velocity, u, is displacement per time.Multiply them together — p times u — and you get work per time per area, or power density, which we call intensity.

That's how these three physical quantities — velocity, pressure, and intensity — are all linked. And even though these can be derived from the wave equation, here we're just focusing on their basic definitions and physical meanings.

slide17:

So now you understand how we measure intensity, and how energy propagates through a medium. As the ultrasound wave travels through tissue, it doesn't stay constant — it gets attenuated. That means the amplitude decreases gradually, and the intensity drops as the wave penetrates deeper. Later, we'll see that the amplitude or intensity decays exponentially with depth, very similar to Beer's law in optical physics.

Because of this, we often use relative intensity and relative power to describe how intensity changes with distance. And since these values can vary across a huge range — sometimes millions of times between the incident pulse and the reflected echo — it's much easier to express them on a logarithmic scale.

This is where the decibel, or dB, scale comes in. It compresses large ratios into smaller, more manageable numbers. So, instead of dealing with enormous values, we talk about relative changes in decibels.

For ultrasound, pressure intensity is expressed in decibels. The unit of pressure is the Pascal, which is one Newton per square meter. The atmospheric pressure is about 0.1 megaPascal, while diagnostic ultrasound can have a peak pressure of around 1 megaPascal.

Now, we describe relative intensity using the formula 10 log of I2 over I1. This represents the change in intensity between two points. If you're dealing with power, however, we use 20 log of P2 over P1, because power is proportional to the square of amplitude.

Let me explain that part carefully. In circuit theory, we know that power equals resistance times current squared — P equals R times I squared. The same idea applies here. Since power is proportional to amplitude squared, we need a factor of two in the logarithmic expression. That's why, for power, we multiply by 20 instead of 10.

So, in summary:– When we talk about intensity, we use 10 log because intensity is directly proportional to energy flow.– When we talk about power, we use 20 log because it depends on amplitude squared.

A 10 dB change corresponds to a factor of 10 change in intensity.A 20 dB change corresponds to a factor of 100 change in intensity.

This logarithmic relationship is used throughout ultrasound physics because it helps us describe very large variations in intensity and power in a compact, practical way.

slide18:

So now, let's do a quick review and apply what we just learned. We'll look at how to calculate remaining intensity and half-value thickness for ultrasound.

For the first example, we start with a 100 milliwatt ultrasound pulse that loses 30 decibels as it travels through tissue. Using the formula, negative 30 dB equals 10 log of I2 over 100 milliwatts, we can solve for I2. Rearranging, we get I2 equals 10 to the power of negative three times 100 milliwatts, which equals 0.1 milliwatt. So, after losing 30 dB, the intensity is reduced by a factor of one thousand.

Now, let's look at the second example — the half-value thickness. This represents the depth at which the ultrasound intensity drops by 50 percent. That means I2 over I1 equals 0.5. Using the decibel formula again, we get 10 log of 0.5 equals negative 3 dB. So, when the intensity is reduced by half, it corresponds to a 3 decibel change.

This concept is very similar to what we use in X-ray or gamma-ray imaging, where we also talk about a "half-value layer." It tells us how deep the wave can travel before its intensity is halved.

In ultrasound, this is just a convenient way to express relative changes using logarithms — there's no physical constant behind it; it's simply a mathematical convenience. The decibel scale helps us describe very large or very small intensity ratios in a practical and readable way.

slide19:

Okay, now we'll talk about acoustic impedance. This concept is actually not dramatically different from what we've already learned in electrical circuits. Remember when we studied circuit analysis? We defined resistance as the ratio between voltage and current. Later, we introduced impedance, which included the effects of capacitance and inductance, and again it was defined as voltage over current. The only difference was that, in general, impedance can have a phase shift, so it's expressed as a complex number — meaning that both voltage and current can be complex, or we can simply describe them in terms of their phase relationship.

Now, think about this analogy in the context of ultrasound. In an electrical circuit, voltage is the driving force — the potential difference that pushes electrons through the conductor. In ultrasound, the driving force is pressure, P. The mechanical wave you push into the medium — that's the pressure wave. So, just as voltage drives electron motion, pressure drives the particle motion in the tissue.

In an electrical circuit, current is the flow of electrons. But in ultrasound, no electrons are flowing; instead, the medium's molecules — air molecules or soft tissue molecules — move back and forth along one direction. At any given location, they have an instantaneous velocity, which we call u z. So, at each point in space, the particles have an instantaneous velocity that's changing with time.

Both P and u keep changing with time, but when we solve the wave equation, we find something very interesting. It's beyond the scope of this lecture to show the full derivation, but the result is that the ratio between pressure and velocity — P over u — is constant at any given location. Even though both pressure and velocity oscillate as the wave propagates, their ratio remains constant in amplitude. Because this ratio stays constant, we can define it as a characteristic property of the medium.

This constant ratio is what we call acoustic impedance, denoted by Z. If the ratio kept changing, it would depend on position or time, and we couldn't define it as an intrinsic property of the material. But because it's stable and characteristic, it gives us something meaningful — a property that tells us how the medium resists the passage of sound.

Acoustic impedance is extremely important in ultrasound imaging. When we use an ultrasound transducer, we have to make sure that the acoustic impedance of the transducer matches the acoustic impedance of the tissue we're imaging — for example, the abdomen. If there's a big mismatch between them, the energy cannot enter the body effectively, and you can't collect the echo signals efficiently. That's why we use ultrasound gel — it acts as a coupling layer that bridges the impedance gap between the probe and the skin.

So, to define it heuristically: acoustic impedance is like voltage over current in an electrical system. Here it's pressure over particle velocity. We can show this mathematically, though I won't go through the full derivation now. By solving the mechanical wave equation, we can prove that Z is constant and depends only on the material's intrinsic properties — specifically, density (rho) and sound speed (c).

Mathematically, we can write this as Z equals rho times c. And we know that the sound speed, c, can be expressed as one over the square root of rho times kappa, where kappa is the compressibility. So, substituting that into the equation, we get Z equals rho over the square root of rho times kappa, which simplifies to Z equals the square root of rho over kappa.

This shows that acoustic impedance depends on both density and compressibility, and it is a fundamental property of the material.

slide20:

Now, once we have this equation, Z equals the square root of rho over kappa, we can calculate or measure the acoustic impedance for different types of biological tissues. Equation 3.8 and the table shown here give you typical values.

If you look at the numbers, you can get a feeling for how impedance and sound speed vary across materials. In air, the sound speed is only about 330 meters per second, because air is very light and compressible. In biological tissues like the kidney or liver, the sound speed is much higher — around 1,560 to 1,570 meters per second — because these materials are denser and less compressible.

Now, if you ask where sound travels the fastest inside the human body, the answer is bone, because bone is the densest and stiffest material we have. The table shows a speed of about 3,500 meters per second in bone, which is several times faster than in soft tissue.

These are just typical examples of tissue acoustic properties, but they are extremely important for understanding how ultrasound interacts with different parts of the body. The speed of sound, the density, and the impedance all influence how much of the sound wave is transmitted, reflected, or absorbed at tissue boundaries.

slide21:

So next, let’s talk about ultrasound–tissue interaction. Earlier, we mentioned how different imaging modalities interact with biological tissues. For example, X-ray and CT involve attenuation and scattering, while gamma-ray imaging behaves similarly. In MRI, we talk about T1 and T2 relaxation parameters to describe energy interactions within tissues.

For ultrasound, the same basic idea applies — it’s all about energy propagation and how that energy interacts with biological structures. The first major effect we talk about is reflection. When a sound wave hits a boundary where there is a difference in acoustic impedance — for example, between air and tissue, or between bone and tissue — some of the energy is reflected. That’s what creates the echoes we use to form ultrasound images.

Reflection happens because of impedance mismatch. If the impedance difference is large, reflection is strong. If the impedances are well matched, reflection is weak, and more energy passes through.

Now, with X-rays and gamma rays, reflection is not really an issue, because their wavelengths are extremely short and the photons are so energetic that they mostly pass straight through. In ultrasound, however, the wavelength is much larger, so reflection becomes significant.

So, remember this simple rule: reflection at a boundary means a change in direction.

The second effect is scattering. When ultrasound waves travel through tissue, they encounter many small structures — cells, organelles, connective fibers — that scatter energy in different directions. So, instead of traveling as a narrow beam, the energy spreads out or diffuses. You can think of this like Compton scattering in X-rays — energy is diffused, and the propagation direction changes randomly.

The third effect is absorption, which converts mechanical energy into heat. In the case of X-rays, this corresponds to the photoelectric effect, where photon energy is absorbed completely. For ultrasound, as the wave passes through tissue, the molecules oscillate back and forth, and because of internal friction, some of the mechanical energy is lost as heat.

This leads to attenuation, which is the overall reduction in intensity. Attenuation happens because energy is both absorbed and scattered away. The wave's amplitude becomes smaller, its pressure decreases, and the particle displacements become weaker as it moves deeper.

If we look at Figure 3.3, it illustrates how the beam behaves when it meets a boundary between two materials with impedances Z1 and Z2. If the wave hits perpendicularly, part of it is reflected back and part is transmitted forward. In more general cases, where the wave hits the surface at an angle, we define the incident angle (theta i), the reflected angle (theta r), and the transmitted angle (theta t). The incident angle equals the reflected angle, but the transmission angle can differ depending on the material properties.

The energy divides between reflection and transmission, and the total energy always equals the incident energy. The ratios of reflected to transmitted energy can be expressed in terms of pressure amplitude, and when you square those amplitudes, you get the intensity ratios.

These ratios depend on the acoustic impedance values of the two materials, Z1 and Z2, as well as the angles of incidence and transmission.

If Z1 equals Z2 — meaning the materials are identical — there will be no reflection at all, and all the energy will pass straight through. But if the impedances are very different — for example, if one of them is close to zero — then the reflection coefficient becomes one, meaning all the energy is reflected back.

That's why impedance matching is so important in ultrasound imaging — it ensures that most of the sound energy enters the body instead of being reflected away at the surface.

So, in summary, ultrasound–tissue interaction involves five main effects: reflection, refraction, scattering, absorption, and attenuation. Together, these determine how the ultrasound wave propagates, how it loses energy, and how we detect it for image formation.

slide22:

Alright, as I mentioned before, the intensity of the acoustic wave — and likewise the pressure amplitude — decreases as you go deeper and deeper into the medium. And the reason both intensity and pressure amplitude become smaller and smaller is due to attenuation.

Now, why do we have attenuation? Because we have two main effects: absorption and scattering. Some of the energy is absorbed and turned into heat, while some is scattered — redirected into other directions. Along the main beam direction, the intensity keeps getting smaller, and the pressure amplitude also decreases.

This is not a strange concept to us — we've seen exponential decay many times before. In nuclear imaging in X-ray imaging, exponential decay appears everywhere. It's a very common phenomenon. And every time we see exponential decay, we can use it to understand how strongly our target interacts with the probing beam. Based on the attenuation rate, we can tell how dense or how light certain tissues are.

That's really the imaging mechanism — attenuation gives us contrast.

Now, the same logic applies here as with X-rays. You don't want something that absorbs energy too quickly. If all the X-rays — or all the ultrasound — are absorbed by the body, then the detector sees nothing. Every photon, or every bit of energy, gets eaten up, and no information is carried out.

But on the other hand, you also don't want the ultrasound beam to pass through the body without any attenuation. If everything goes through unchanged, again, you get no information.

The ideal situation is somewhere in between — you want about half the energy absorbed, and half transmitted. Not exactly fifty–fifty in every case, but roughly half. That balance gives you the maximum contrast and allows you to reconstruct a meaningful tomographic image.

Now, I often make this analogy with exams. When I give you an examination, I try to set it up so that the class average is around fifty percent. Why? Because that gives me the best spread — I can see which questions are tricky and which ones you can solve easily. It allows me to judge relative performance much more clearly. So, don't worry if your average score isn't high. I purposely don't aim for a high average, because it gives me better information about the class as a whole. And in the end, your final letter grades — A, B, or C — are assigned consistently with the overall performance distribution. So nobody is unfairly graded.

Coming back to imaging, the same principle applies. The attenuation curve gives us the contrast we need — not too much, not too little. And in ultrasound, we often use two equivalent forms to describe it: μ, the attenuation coefficient for intensity, and α, the attenuation coefficient for pressure. These are just two different ways to express the same exponential decay behavior. On a decibel scale, μ and α become interchangeable.

slide23:

Now, let's talk about how attenuation depends on frequency.

Here, the coefficient α is proportional to a constant, which we call α-naught, and it's also proportional to frequency. Mathematically, we often write this as alpha of f equals alpha-naught times f to the power of b.

For diagnostic ultrasound, the value of b is about one, which means the relationship is almost linear. In other words, as the frequency increases, attenuation increases linearly.

So, higher frequency means stronger attenuation. And that makes perfect sense physically. Think about it: when you vibrate molecules more frequently, they move back and forth faster. Faster motion means more friction inside the tissue, and that friction converts energy into heat. So, higher frequency means you lose energy faster — the wave decays in intensity and amplitude more quickly.

In simple terms: higher frequency equals higher energy loss equals less penetration.

Now, for different types of materials, the constant α-naught is different. If you look at the table on the slide, you'll see that fat has a smaller α-naught value, while liver is a little higher. Cardiac muscle is higher still, and bone has a very large value — meaning sound doesn't penetrate bone very well.

You can use these numbers to get a sense of how attenuation works in different tissues. For example, let's take fat at 5 megahertz. You can calculate attenuation as five megahertz times zero point six three decibels per centimeter per megahertz, which equals about three point one five decibels per centimeter. After traveling four centimeters, that's twelve point six decibels of attenuation. The relative amplitude of the wave becomes ten to the power of negative twelve point six divided by twenty, which is about zero point two three four.

So, as frequency increases, attenuation increases. That's why in medical ultrasound we always have to make a trade-off: high frequency gives you better resolution but less depth, while low frequency penetrates deeper but gives you lower resolution.

The takeaway is simple and important: higher frequency → larger attenuation coefficient → less penetration depth.

slide24:

Now let's talk about some practical implications of reflection coefficients — this is really an extension of what I already mentioned earlier.

You cannot penetrate ribs or air bubbles easily, because their acoustic impedances are very different from those of soft tissue. When there's a large mismatch in impedance, most of the ultrasound energy is reflected, and very little gets through.

Also, if you try to do brain imaging in vivo, the skull will reflect almost all of the ultrasound. That's why we don't use ultrasound to image the brain the same way we use CT or MRI. The bone–tissue boundary is such a strong reflector that almost no energy gets inside.

Even inside the brain, gray matter and white matter are acoustically very uniform. In a uniform medium, the ultrasound wave passes straight through without much change — kind of like a transparent window. But that also means you don't get contrast. It's like when everyone in the class scores a perfect 100 — I can't tell who is better, because everyone looks the same.

On the other extreme, if you have a strong mismatch — for example, at a tissue–bone interface — the ultrasound wave cannot penetrate at all. Everything is reflected. That's like an exam where everyone scores zero — again, I can't tell who's better.

So, what we really need is a good balance — a useful imaging window. In practice, this means finding acoustic windows that allow sound to enter effectively. For instance, when imaging the heart, we position the probe around the ribs to find a small gap where the sound can pass through. We also make sure to use ultrasound gel between the transducer and the skin, because that gel helps match the impedance between the probe and the tissue. Without it, the air gap would cause almost total reflection, and no signal would enter the body.

So the key idea is this: reflection depends on impedance mismatch. If there's too little mismatch, there's no contrast — everything looks uniform. But if the mismatch is too large, nothing penetrates — you can't see anything. The useful imaging window lies in between.

This is something I've mentioned before, and you can read through the bullet points on the slide to reinforce the idea. The goal is always the same: find the right balance between reflection and transmission to get good imaging quality.

slide25:

Alright, now let's talk about something quite interesting — the biological effects of ultrasound. When we introduce ultrasound energy into the body, it can lead to two kinds of effects: thermal effects and non-thermal effects.

I've already mentioned the thermal effects earlier, but let's quickly recall them. When ultrasound travels through tissue, part of its mechanical energy gets converted into heat. This heating can cause several physiological changes — for instance, it can increase collagen tissue extensibility, which makes soft tissues more flexible. It can also increase blood flow, which improves circulation and oxygen delivery. In addition, it can increase nerve conduction velocity, meaning that nerve impulses travel faster. This effect can also raise the pain threshold, allowing patients to tolerate more discomfort during therapy.

Moreover, heating can enhance enzymatic activity — many enzymes in the body function more efficiently at slightly elevated temperatures. And finally, ultrasound heating can help decrease muscle spasm, relaxing the tissue and improving mobility.

Now, in addition to these thermal effects, ultrasound also causes non-thermal effects, and these are especially fascinating. Non-thermal effects are not caused by heating; instead, they result from mechanical interaction with the tissue — the actual pressure waves and oscillations in the medium.

For example, ultrasound can increase cell membrane permeability, allowing molecules to move in and out of cells more easily. It can also increase vascular permeability, which means that blood vessels become slightly more "leaky," allowing nutrients or drugs to diffuse through more effectively. Both of these effects are very useful in therapeutic and drug delivery applications.

Ultrasound can also increase local blood flow without significant heating, helping tissue repair and regeneration. It can assist in the reduction of edema, which is the accumulation of fluid in tissues, and that makes it very useful for treating inflammation or swelling.

One of the most remarkable non-thermal effects is something called cavitation. Cavitation refers to the formation and oscillation of tiny bubbles — microscopic gas pockets — in the tissue fluid. At certain frequencies and intensities, these bubbles can oscillate and even collapse, releasing bursts of mechanical energy that can affect nearby cells or molecules.

As shown in the figure on this slide, you can see different frequency ranges where ultrasound is applied — from very low frequencies used in physiotherapy to higher frequencies used in imaging, Doppler studies, and even lithotripsy for breaking kidney stones. The likelihood of cavitation increases with lower frequency and higher acoustic pressure. However, for standard medical diagnostic ultrasound, which operates between about 3 and 20 megahertz, cavitation is highly unlikely, so it's very safe.

Also notice the lower figure — it shows how ultrasound can enhance drug delivery. The red curve shows a larger uptake of drug when ultrasound is applied, compared to the blue curve, which shows little uptake without ultrasound. This demonstrates how ultrasound can help drugs penetrate biological barriers more effectively.

So, in summary, ultrasound has both thermal and non-thermal effects, and both can be used for different therapeutic and diagnostic purposes. The non-thermal effects, like increased permeability and cavitation, open up some exciting possibilities — such as ultrasound-mediated drug delivery and neurological stimulation. These are emerging areas of research and clinical application.

slide26:

Now, let's shift gears a little bit. Up to this point, we've been focusing mainly on the physics of ultrasound — the wave equation, acoustic impedance, and sound–tissue interactions. We've covered how ultrasound propagates, how energy is absorbed or reflected, and how contrast is generated in imaging.

But now, we're moving into a different section — a more engineering-oriented discussion. These next few slides focus on the ultrasound scanner, or sometimes we call it the ultrasound probe or transducer. All these terms basically refer to the same device.

So, in this next part, we'll first talk about the single ultrasound unit, and then how multiple such units are combined to form arrays — either linear arrays for one-dimensional scanning, or two-dimensional arrays for volumetric imaging. This is actually quite similar to how a CT scanner is organized — you have multiple detectors arranged in a specific pattern to collect information from different angles.

We'll also talk about image resolution, contrast, and microbubbles, which are tiny gas-filled spheres used as ultrasound contrast agents. These help us visualize blood flow and enhance tissue contrast in medical imaging.

So again, just to summarize — the first part of this lecture covered the physical foundations, while this next part is all about the engineering principles of the ultrasound scanner itself: how it's built, how it operates, and how it detects signals.

slide27:

So, how do we actually generate ultrasound waves?

Earlier, I mentioned that if you push and pull repeatedly — like stretching and compressing a spring, or vibrating a rope — you can generate waves. Those are good heuristic examples to understand the idea. But for medical ultrasound, we need something that can vibrate millions of times per second — far beyond what we can do by hand. So, we use a physical mechanism called the piezoelectric effect.

The piezoelectric effect is truly fundamental to ultrasound imaging — it's the enabling technology behind both the generation and detection of ultrasound waves. In fact, this phenomenon was so important that it earned recognition at the Nobel Prize level.

Now, the piezoelectric effect works in two directions, which we call the direct effect and the reverse effect.

Let's start with the direct effect. When sound vibrations — that is, pressure waves — hit a piezoelectric crystal, the alternating pressure causes the crystal to deform mechanically. As the pressure oscillates from positive to negative, the crystal expands and contracts slightly. This mechanical deformation generates an electrical potential across its surfaces. In other words, mechanical energy — the sound vibration — is converted into electrical energy. This is how the ultrasound echo signal is detected by the transducer and turned into an electrical signal for processing.

Now, the reverse effect works the other way around. When we apply a voltage across the piezoelectric crystal, it causes the crystal to deform — to change shape. When the voltage alternates — positive, then negative, then positive again — the crystal vibrates back and forth. If that electrical signal has a frequency in the ultrasound range, those vibrations produce ultrasound waves that propagate into the surrounding medium.

So you can see that the same piezoelectric element can do both jobs — it can generate ultrasound waves and also detect the returning echoes.

Here's how it all fits together: when you send a sinusoidal electrical signal to the crystal, it vibrates and emits an ultrasound wave that travels into the body. The wave reflects from tissue boundaries and returns to the transducer. Those returning pressure waves deform the crystal again, and the deformation creates an electrical signal that can be measured.

That's why we call it a two-way transduction process — it goes from electrical energy to mechanical energy when we transmit, and back from mechanical to electrical energy when we receive.

This process — the piezoelectric effect — is the heart of every ultrasound transducer. It allows the probe to both send and receive sound waves, enabling us to form detailed medical images.

So, when you see the term "piezoelectric ceramics" or "piezoelectric crystals," just remember: that's what's inside every ultrasound probe, converting electrical energy to sound and sound back to electrical energy, millions of times every second.

And that's the essence of how ultrasound imaging becomes possible.

slide28:

You might wonder, why does the piezoelectric effect actually happen? What's going on inside the material that allows it to generate electricity when we apply pressure?This little cartoon, or infographic, explains it beautifully.

Imagine you have a crystal structure, such as quartz, that is in perfect static balance — both mechanically and electrically. In its no-stress state, the positive and negative charges are arranged symmetrically. The positive charges — usually from the silicon atoms — and the negative charges — from the oxygen atoms — balance one another out, so the overall crystal is electrically neutral. There is no net polarization, no voltage across the surfaces.

Now, let's apply a horizontal tension — we're stretching the crystal, pulling its sides apart. The crystal is stable but elastic, so it deforms slightly. When that happens, the charge centers shift — the positive charges move upward, and the negative charges move downward. That means the top surface becomes slightly positive, and the bottom surface becomes slightly negative. In other words, the mechanical strain has produced an electrical potential difference between the two surfaces.

Now imagine the opposite — a compression instead of tension. When we compress the crystal, we push those atoms closer together, but in the opposite direction. The negative charges are forced upward, and the positive charges are pushed downward. Now the top surface becomes negative, and the bottom surface becomes positive.

This alternation between tension and compression causes repeated charge separation, which is what generates an alternating voltage when the material is vibrated continuously.

So what you see in this diagram — the shift of positive and negative charge centers — is the fundamental reason behind the piezoelectric effect. Mechanical deformation creates electrical polarization, and electrical excitation can, in turn, create mechanical motion.

This is the microscopic picture that explains how piezoelectric crystals act as transducers, converting mechanical vibration into electrical energy and back again. It's this static-electric imbalance, produced by strain, that forms the very foundation of ultrasound imaging.

slide29:

Now, let's take a look at the single-crystal transducer, which is one of the simplest yet most important components in ultrasound systems.

If you look at the diagram from your textbook, you can see all the key parts clearly labeled. At the heart of the transducer is the piezoelectric crystal, the active element that generates and receives ultrasound waves. On one side, there is a matching layer, which I mentioned earlier — it's used to match the acoustic impedance of the crystal to that of the tissue. Without this layer, much of the energy would be reflected at the interface, and very little sound would enter the body.

Behind the crystal, you'll see a damping material. This is essential because it absorbs unwanted echoes and stops the crystal from ringing after each pulse. The damping material ensures that the emitted pulse is short and well-defined, giving us better axial resolution in the image.

Surrounding everything, there's an acoustic insulator or backing, which directs the ultrasound wave forward into the tissue rather than allowing it to leak backward into the housing. There's also a conducting wire, or contact lead, which connects to the electrical circuitry, allowing us to apply voltage across the crystal and receive electrical signals from it.

All of these components — the crystal, the matching layer, the damping material, and the insulator — are enclosed in a metal or plastic housing, often with an ergonomic handle. Together they form what we call a single-element transducer.

Mathematically, this crystal has a natural resonant frequency, denoted as f-zero, given by the formula f-zero equals the speed of sound in the crystal divided by twice its thickness.In other words,F-Zero equals c-crystal over 2 d.

So, the frequency of the ultrasound wave is determined by how fast sound travels in the crystal and how thick the crystal is. A thinner crystal produces higher frequency sound, while a thicker one produces lower frequency sound.

This simple relationship between frequency, sound speed, and crystal thickness is fundamental to designing ultrasound probes for different applications. High-frequency probes, which give better resolution, use very thin crystals, while low-frequency probes, which penetrate deeper, use thicker ones.

So, this diagram gives you the complete picture — from physical design to mathematical principle — of a single-crystal transducer.

slide30:

Now, once you put that single-crystal transducer inside a handheld housing, you get what we commonly call a probe — or a transducer head.

You simply place this probe on top of the area you want to examine — it could be the chest, the abdomen, the breast, an arm, or any other part of the body — depending on what you are imaging. These probes are used for a variety of diagnostic and therapeutic purposes.

On this slide, you can see several examples of commercial ultrasound probes. Each one is designed for a different purpose — some are curved for abdominal imaging, some are linear for vascular or musculoskeletal scans, and others are very small for intracavitary or ophthalmic use.

Now, I'd like to share a more forward-looking idea — something I've been thinking about for quite a while. I believe that in the future, our everyday devices — maybe even the iPhone — could incorporate this kind of piezoelectric technology directly into the screen.

Just imagine that the phone's surface could act as a two-dimensional acoustic transducer array. Each tiny pixel on the screen could both generate and receive ultrasound signals. That means you could, in principle, perform a quick body scan right at home — maybe before going to bed.

You could place the phone on your skin, and the device would automatically send out ultrasound waves, collect the echoes, and the AI software inside would analyze the data for you. It might tell you, "Everything looks fine," or "You may want to check this area further." It could even show you which parts of the body were already scanned and which ones weren't.

By combining piezoelectric technology, machine learning, and digital signal processing, we could one day make personal, professional-level imaging accessible to everyone — right from their own phone.

This might sound futuristic, but remember — all great innovations start as ideas. And this one could truly revolutionize preventive healthcare by bringing ultrasound diagnostics into the hands of ordinary people.

slide31:

So now let's imagine what the damping effect really means — and more importantly, why we need it in ultrasound imaging.

Think about how an ultrasound pulse is sent into the body. Most of the time, ultrasound imaging doesn't use a continuous wave; instead, it operates in what we call a pulsed mode. You send out a pulse, it travels into the body, reflects from different tissue layers, and then comes back to the transducer as an echo. This pulsed approach is extremely useful because it allows us to measure both depth and velocity — that is, how deep a structure is and how fast it's moving.

Now, let's say you apply a short electrical signal to the piezoelectric crystal — it vibrates for a moment and produces a burst of ultrasound waves. Then, you stop the input signal. Ideally, the crystal should also stop vibrating immediately, just like striking a bell and having it stop ringing right away. But in reality, if there's no damping, the crystal keeps vibrating for quite a while after the excitation stops, almost like a bouncing ball that never settles down.

If there's no damping — no internal friction — the ball keeps bouncing up and down forever. That's exactly what would happen to your transducer: it would continue ringing, creating a long, drawn-out pulse. But we don't want that, because a long pulse means poor resolution — the returning echoes overlap and you lose clarity.

Now, if the transducer is placed in a medium with some friction — like air or fluid — those vibrations die out faster. The amplitude drops quickly, and you get a much shorter pulse. This reduction in vibration amplitude is called damping.

So, when we apply a damping material behind the crystal, it acts like friction for the vibrating transducer. It absorbs the leftover energy, preventing it from ringing too long. In the time domain, that means the vibration decays faster, giving us a short, clean pulse. In the frequency domain, the pulse now covers a wider range of frequencies — in other words, a broader bandwidth with a lower amplitude peak.

Without damping, you have a narrow and tall frequency spectrum — that's a very sharp resonance around the central frequency. With damping, the spectrum is broader and flatter — the amplitude is smaller, but the bandwidth is much wider, which is exactly what we want for imaging.

We define this relationship using something called the Q factor, or the quality factor. Mathematically, Q equals two pi times the central frequency divided by the bandwidth. So, Q equals $2\pi f_0$ divided by BW.

A high-Q system means low damping — long vibration, narrow frequency range. A low-Q system means strong damping — short vibration, wide frequency range.

For diagnostic ultrasound, we actually prefer low-Q systems because we want short pulses and broad bandwidths. Typically, the Q factor is around 1 or 2, which provides an ideal balance between good axial resolution and manageable signal strength.

So, damping is what makes our ultrasound pulses short, our echoes sharp, and our images clear.

slide32:

Now that we understand damping, let's move on to another important component — the matching layer. You can see from this slide that the transducer structure includes this special layer right in front of the piezoelectric crystal, and this layer is absolutely essential for getting ultrasound energy efficiently into the body. Let me start with a quick personal story. I once mentioned that I had a kidney stone, and when you go for an ultrasound scan like that, you'll notice the doctor always applies a gel on the skin. The gel is completely harmless and painless, though yes, it can feel a little messy. But that gel is doing something very important — it's not just for comfort or for better contact. It actually helps solve what we call a matching problem.

Here's why. The piezoelectric material inside the transducer, usually something like PZT — that stands for lead zirconate titanate — has a very high acoustic impedance, which we call Z P or Z transducer. The human skin, on the other hand, has a much lower acoustic impedance, which we call Z skin. And between the two surfaces there are often tiny air bubbles, because the skin isn't perfectly smooth. Now, whenever two materials with very different acoustic impedances are in contact, most of the ultrasound energy reflects back instead of transmitting forward. So the ultrasound wave generated by the crystal would mostly bounce back at the surface — kind of like shouting into a canyon and hearing your echo. That's bad for imaging because very little energy actually enters the body.

So what do we do? We introduce a matching layer — a thin material layer whose acoustic impedance lies between those two values. This layer acts as a bridge, transferring acoustic energy more efficiently from the transducer into the skin. Here's how it works. The vibration from the crystal first transfers its energy into the matching layer, and then from the matching layer into the skin. So instead of one big impedance jump, we now have two smaller steps — from crystal to matching layer, and from matching layer to skin.

Mathematically, the amount of energy that gets transmitted from one medium into another depends on their respective acoustic impedances. The transmitted intensity between two media is proportional to the product of their impedances divided by the square of their sum — that's shown in the equation on the slide. To optimize this transmission, we choose the impedance of the matching layer carefully. The ideal acoustic impedance of the matching layer, which we call Z M, is the geometric mean of the two adjacent impedances. So we say, Z M equals the square root of Z transducer multiplied by Z skin. Again, that's Z M equals the square root of Z transducer times Z skin.

This relationship minimizes reflection and maximizes transmission, allowing most of the ultrasound energy to enter the body and come back to the transducer as echoes. That's why every ultrasound probe has a matching layer built in, and that's also why we always use gel on the skin — the gel acts as an acoustic coupling medium that functions as part of the matching layer system. And by the way, as you can see on the slide, this topic might appear in your homework, so make sure you remember it.

To summarize: without the matching layer, most of the ultrasound energy would just reflect at the skin's surface. With the matching layer — and with the gel acting as a coupling medium — the energy is transmitted efficiently, first from the crystal to the gel, then from the gel to the body. This ensures stronger echoes, clearer signals, and higher-quality ultrasound images.

slide33:

Now, let's take a closer look at something very familiar in every ultrasound exam — the ultrasound gel. You've probably seen this many times during a scan, whether it's for the abdomen, the heart, or during pregnancy imaging. The sonographer or the doctor squeezes out some gel and spreads it on the patient's skin before placing the transducer on top. It may seem simple, but it actually plays a crucial physical role in ultrasound imaging.

Remember, we just discussed the matching layer — that thin layer designed to help ultrasound energy travel efficiently from the transducer into the body. Well, the ultrasound gel is part of that same idea. It acts as an acoustic coupling medium, helping to eliminate the air gap between the transducer and the skin. Without the gel, there would be a very thin layer of air between the probe and the body surface, and that's a serious problem.

Why? Because air has an extremely low acoustic impedance — much lower than either the piezoelectric crystal or human tissue. When an ultrasound wave tries to go from the crystal into air, almost all of the energy reflects. In fact, more than ninety-nine percent of the signal would be lost. So the ultrasound beam would never enter the body; you'd see nothing but noise on the screen.

The gel solves this problem beautifully. It has an acoustic impedance value that lies roughly between the impedance of the transducer surface and that of the skin, just like a matching layer. When the gel fills in the microscopic irregularities on the skin and removes any trapped air bubbles, it creates a smooth, continuous path for the sound waves to travel through. That means more ultrasound energy gets transmitted into the body, and stronger echoes come back to the transducer.

So, in simple terms, the gel makes sure there's no air, only sound transmission. It improves image quality dramatically, ensuring clear boundaries and accurate echo measurements. It also helps the probe glide smoothly across the skin, which is important for maintaining good contact during a scan.

And practically speaking, the gel is made of a water-based polymer that's safe, non-toxic, and easily cleaned off after the exam. Some gels are even warmed up before use to make the patient more comfortable. But whether it's warm or cold, the gel is essential — without it, even the best transducer couldn't produce a good image.

So the next time you see that ultrasound gel being applied, remember it's not just for comfort — it's a critical part of the acoustic system. It's what allows the ultrasound energy to move efficiently from the transducer, through the skin, and into the body, giving us clear, high-quality medical images.

slide34:

So this is your homework, and here we start to look at what modern ultrasound transducers actually look like inside. In most modern ultrasound systems, the transducers are not just a single piezoelectric element like we saw earlier — they're made up of arrays of elements, either one-dimensional or two-dimensional. The image you see here is a schematic of a one-dimensional linear array transducer, which is very common in medical ultrasound imaging.

You can see that each piezo element is one tiny rectangular piece of piezoelectric material. Beneath these elements, we have the electrodes, which are used to apply oscillating sinusoidal electrical signals. These electrical signals cause the piezo elements to vibrate and generate ultrasound waves — and when echoes return from the body, those same elements work in reverse, converting mechanical vibration back into electrical signals.

Now, above these piezoelectric elements, you can see matching layers — sometimes there are two, a first and a second matching layer. These are just like what we discussed earlier: they help transfer ultrasound energy more efficiently into the body by gradually matching the acoustic impedance between the crystal and the skin. The entire system is connected to a custom-designed flexible circuit that allows precise timing and control for each element, enabling electronic focusing and beam steering.

Behind the piezo elements is the backing layer, which provides damping — it absorbs unwanted vibrations so that each pulse is short and clean. This improves axial resolution by preventing long "ringing" of the crystal. So, when everything is working together — the electrode, the matching layers, the gel, and the backing layer — the ultrasound energy is efficiently transmitted into the body, and the echoes return through the same path.

In essence, you have a system that is symmetrically optimized for both transmission and reception. The signal path forward is carefully matched, and the reverse path follows the same principle, leading to a higher signal-to-noise ratio and better overall image quality.

slide35:

Now, let's talk about the trade-off between resolution and penetration in ultrasound imaging. This is a very important concept that determines how clearly you can see structures at different depths inside the body.

If you want to get higher resolution, generally you need to increase the frequency of your ultrasound wave. Higher frequency means a shorter wavelength, and that allows you to distinguish smaller structures — that's why high-frequency ultrasound is used for imaging shallow organs like the thyroid or the eye.

However, there's a price you pay for this. Higher frequency sound waves are attenuated much faster as they travel through tissue. Remember that the attenuation coefficient, alpha, is proportional to frequency — we can say alpha is proportional to f. So, as frequency increases, attenuation also increases. That means the ultrasound wave loses energy faster, and it can't penetrate as deeply into the body.

In contrast, a lower frequency transducer can send energy much deeper, but it comes with lower resolution. For example, a 12 megahertz transducer provides excellent resolution but can't reach very deep — maybe just a few centimeters. On the other hand, a 3 megahertz transducer can reach much deeper into the abdomen, but the image will be coarser.

So, in practice, we always balance these two factors. You choose your transducer frequency based on the target you want to image — high frequency for fine detail near the surface, and low frequency for deeper penetration. This is one of the most fundamental trade-offs in ultrasound imaging.

slide36:

Now, let's move on to beam geometry, which gives us a deeper understanding of how ultrasound beams behave as they propagate through space. This figure from your textbook shows the beam shape produced by a single crystal transducer.

When the transducer vibrates, it sends energy downward in a conical shape. The region closest to the transducer is called the near field, and beyond a certain point — called the near field boundary — the beam starts to diverge, entering what we call the far field.

In the near field, the beam diameter is roughly the same as the transducer's aperture — so if the transducer has a diameter "a," the beam width is about the same. But after the near field boundary, the beam starts to spread out or diverge. This divergence angle, which we call theta, is determined by the wavelength lambda and the transducer aperture "a." It's given approximately by the equation:

theta equals arc sine of zero point six one times lambda divided by a.

This relationship shows that smaller apertures or larger wavelengths give wider beams — meaning poorer lateral resolution.

Now, within any cross-section of the beam, the intensity isn't uniform. Along the central axis, the signal is strongest, and as you move away to the sides, the intensity decreases. The intensity profile across the beam can often be approximated by a Gaussian distribution, meaning it has a peak at the center and falls off symmetrically.

To describe this beam width, we use a quantity called full width at half maximum, or F-W-H-M, which literally means the width of the beam where the intensity drops to half of its peak value. For a Gaussian profile, the FWHM is approximately two point three six times sigma, where sigma is the standard deviation of the Gaussian. Mathematically, we write it as:

FWHM equals two times the square root of two times the natural log of two, multiplied by sigma — approximately equal to two point three six sigma.

This measurement gives us a good estimate of the lateral resolution of the ultrasound beam. The narrower the beam, the better the lateral resolution — meaning you can distinguish two nearby structures more clearly.

So, beam geometry, frequency, aperture size, and attenuation — all these factors together define how sharp and how deep your ultrasound image will be.

slide37:

Now, this slide explains the concept of lateral resolution — sometimes called cross-beam resolution — and it's illustrated again by a figure from your textbook. Remember earlier, we talked about the beam profile and how it has a certain width, which we describe using the full width at half maximum, or F-W-H-M. Ideally, we want this width to be as small as possible, because that determines how well we can distinguish two objects that are side by side.

If two scatterers in the body — say, two small reflecting points — are separated by a distance smaller than the beam width, their echoes will overlap when received by the transducer. When that happens, the signals blend, and we can't tell them apart. That's what you see on the right side of the figure: the two curves merge into one, and the system records them as a single echo.

On the other hand, if the two scatterers are separated by a distance greater than the beam width — specifically, greater than the full width at half maximum — then the two echoes appear as two distinct peaks. That's what we want. So the smaller the F-W-H-M, the better your lateral resolution, because you can separate two adjacent points in space.

So in summary, lateral resolution is determined primarily by the beam width. Narrow beams resolve closely spaced objects; wide beams cause overlap. This concept applies directly to your image sharpness — particularly in the direction perpendicular to the beam path.

slide38:

Now let's move on to ultrasound focusing, which is another crucial concept for improving lateral resolution. You can think of this just like focusing light through a lens in photography. The figure here from your textbook shows three cases: strong focusing, weak focusing, and a reference to how the focal plane and focal distance change.

If we have a single-crystal transducer — often shaped like a disk — it can have a concave surface, meaning it curves inward. The curvature acts like an optical lens, helping to converge the ultrasound beam to a smaller focal spot. When the curvature is stronger, the beam converges faster, giving us a shorter focal distance and a tighter focus. That means, at that focal point, you have very high resolution — the beam is extremely narrow, and you can separate fine structures clearly. However, beyond that focal plane, the beam quickly diverges again, and the resolution deteriorates.

If the transducer surface is only slightly curved, the focusing is weaker. The focal distance becomes longer, and the beam stays more uniform downstream. The resolution isn't as sharp at the focal point, but it's more consistent along the depth of field. So there's always a trade-off — strong focusing gives you very high resolution over a small region, while weak focusing gives you moderate resolution over a wider range.

Now, in both optics and acoustics, we use a quantity called the F-number, or F-ratio, to describe focusing strength. The F-number is defined as the radius of curvature divided by twice the aperture radius. Mathematically, that's:

F equals R divided by two A.

Here, R is the radius of curvature, and A is the aperture radius. So, a smaller R or a larger aperture A gives you a smaller F-number — meaning stronger focusing. Conversely, a large F-number corresponds to weaker focusing.

So what we're really describing here is how lateral resolution changes as a function of focusing strength and imaging depth. When you design an ultrasound system, you must decide whether you want a narrow beam with strong focusing or a broader, more uniform beam with weaker focusing, depending on your imaging application.

slide39:

Alright, so we've talked about lateral resolution — how well we can distinguish two points that are side by side. Now let's move on to the next concept, which is axial resolution, or sometimes called longitudinal resolution — meaning how well we can distinguish two points along the direction of beam propagation.

If you look at this figure, it might seem a little tricky at first glance, but let's walk through it carefully together. To understand axial resolution, remember that in ultrasound imaging we always send out short bursts, or pulses, of sound — not continuous waves. We need these pulses, and as we mentioned before when talking about damping, that damping helps shorten the pulse. A shorter pulse means better resolution.

So imagine you send out one pulse. You already know the speed of sound in the medium, so after a short while, the pulse travels outward, interacts with the tissue, and part of it comes back as an echo. Based on the time interval between when we send the pulse and when we receive the echo, we can calculate the round-trip distance the sound wave has traveled.

Now, if I send another pulse right after the first one, the same process repeats — it goes out, reflects, and comes back. The key here is that axial resolution tells us how well we can resolve separate boundaries along that same line of travel. So the big question is: what's the smallest distance between two reflecting surfaces that we can still tell apart? That minimum separable distance is what we call the axial resolution.

If those two boundaries are too close together, the echoes from them will overlap — just like what you see in this figure — and when that happens, all the signals merge, so you can't tell which one is which. They just blur into one single reflection.

From this physical reasoning, we can define the formula for axial resolution. It's given by one-half times the pulse duration times the speed of sound. Let me say that clearly — axial resolution equals one-half times pulse duration times c, where c represents the speed of sound in the medium.

So, mathematically, that's:Axial resolution equals one-half multiplied by pulse duration multiplied by c.

This comes from a simple idea. For a first-order approximation, we assume the speed of sound is uniform throughout the tissue. Then we can look at which factors matter most. There are two major players here: the pulse duration and the sound speed. The smaller the pulse duration, the better the resolution. And actually, a smaller sound speed also improves resolution, because the wave doesn't travel as far during each time interval.

Now, picture this in your mind. You send out a pulse — it has a certain length, from the front, which we call the leading edge, to the back, which we call the trailing edge. This whole pulse moves through the tissue. As it moves, it hits boundaries — let's call them surface A and surface B. The leading edge of the pulse reaches surface A first and reflects back, while the trailing part continues forward. Some portion of the pulse keeps going and hits surface B, and then that part reflects back too.

So the leading edge reflects first, and a moment later, the trailing edge and the transmitted component reach the second boundary. The reflected waves from surface A and surface B both travel back toward the transducer. Now, if these two boundaries are so close that the time difference between those reflections is shorter than the pulse duration, the two echoes overlap — and we cannot tell them apart. But if the separation is large enough that the round-trip time between them is longer than half the pulse duration, we can distinguish the two.

That's exactly how we define axial resolution — it's the distance between surface A and surface B that produces two distinguishable echoes. In symbols, we can write it as:Delta x equals one-half times pulse duration times c.

Let me read that again in a smooth way:"Delta x equals one-half times pulse duration times c."

The one-half factor is there because the sound wave travels a round trip — out to the boundary and then back — so only half of that total distance corresponds to the actual separation between two reflectors.

So here's the physical picture:If the pulse duration is short, the leading and trailing edges of the pulse are close together, which means the echoes from A and B will be well separated in time — giving you better axial resolution.If the pulse duration is long, or if the boundaries are too close, then the echoes will merge together, and you'll lose that resolution.

So when we design ultrasound systems, we try to make the pulse as short as possible — and that's where the damping material plays a critical role. The damping shortens the pulse, improves temporal separation, and gives us sharper axial resolution.

To sum it up:Axial resolution represents the minimum distance between two boundaries along the beam path that can be resolved separately. It depends mainly on pulse duration and the speed of sound in the tissue. The shorter the pulse and the slower the sound, the finer the detail we can see along that direction.

So just review this figure carefully yourself. Once you visualize how the pulse moves forward, hits surface A, then surface B, and the echoes return, the logic becomes crystal clear. We now have both key ideas of spatial resolution — lateral resolution and axial resolution — and next, we'll move on to discuss how all this affects image contrast, which determines how clearly structures appear on the ultrasound image.

slide40:

Now, let's talk about image contrast in ultrasound imaging. The term "contrast" here simply refers to how different tissues appear relative to one another based on the strength of the returning echoes. In ultrasound, acoustic waves interact with different tissues in different ways — and that's why we see some areas appearing brighter and others darker on the screen.

We use the prefixes "hyper" and "hypo" to describe this difference in the scattering amplitude of echoes. When a tissue produces echoes that have a higher scattering amplitude than the average, we call it "hyperechoic." That means the reflected signal is stronger than normal, and it shows up as a bright region on the ultrasound image. Conversely, when the tissue has a lower scattering amplitude than average, we call it "hypoechoic." In that case, the reflected signal is weaker, and it appears darker on the image.

These variations in brightness — from hyperechoic to hypoechoic — are what give us contrast. They allow us to distinguish one structure from another. For example, on the ultrasound image shown here, you can see how the liver, renal sinus fat, renal pyramids, and fluid-filled bowel loops all appear with different brightness levels. The renal sinus fat appears bright because it is hyperechoic — it reflects ultrasound strongly. The renal pyramids or fluid-filled bowel loops appear darker because they are hypoechoic — they scatter less energy.

Some tissues, like the liver or renal cortex, tend to have uniform, mid-level echogenicity, meaning they reflect ultrasound in a moderate, consistent way. Others, like the anterior surface of the liver or the anterior surface of the kidney, are more reflective, showing distinct bright boundaries on the image.

So, in summary — the way ultrasound waves interact with various tissues determines the level of brightness on the image. Hyperechoic areas are brighter and correspond to stronger scatterers. Hypoechoic areas are darker and correspond to weaker scatterers. This difference in signal amplitude is the foundation of contrast in ultrasound imaging.

However, sometimes the natural contrast between tissues is not sufficient to clearly separate structures or identify pathology. In such cases, we use contrast agents to improve visibility. One of the most effective types of contrast agents in ultrasound imaging is the microbubble. These microbubbles, typically ranging from two to five micrometers in diameter, are introduced into the bloodstream to enhance contrast. They act as additional reflective surfaces that increase backscattering, especially in regions where the natural contrast is weak.

By introducing microbubbles, we can significantly improve the visualization of internal structures, making it easier to detect lesions, tumors, or subtle tissue changes. This approach is particularly useful in clinical cases where precise imaging is needed to make accurate diagnoses.

slide41:

Let's look more closely at these microbubbles and how they work as ultrasound contrast agents.

Microbubbles are very small — typically two to five micrometers in diameter — and are encapsulated by a shell made of either albumin or a lipid layer. This shell makes them stable enough to circulate through the bloodstream without immediately collapsing. Their stability can be further enhanced by filling them with a high molecular weight gas, which diffuses out more slowly than regular air.

These microbubbles are produced outside the human body and then introduced into the bloodstream through an intravenous injection. Once inside the bloodstream, they interact with ultrasound waves. When the ultrasound beam hits them, the microbubbles start to oscillate — they expand and contract in sync with the pressure variations of the sound wave.

Because the interface between the gas inside the bubble and the surrounding liquid is not acoustically transparent, a strong reflection occurs. This makes microbubbles act as powerful scatterers, producing bright, well-defined echoes. As a result, microvasculature and even small blood vessels become visible on ultrasound images, greatly improving the detection of blood flow and perfusion.

Moreover, these microbubbles can be customized. Scientists can modify the shell using polymers or special coatings that allow the surface to be functionalized with antibodies. When certain antibodies are attached, the microbubbles can specifically bind to biological targets — for example, tumor cells or inflammatory markers.

When that happens, the microbubbles accumulate around those target regions. On the ultrasound image, those regions appear as bright, localized signals — making it possible to identify molecular features, not just structural ones. In other words, microbubbles transform ultrasound from a purely anatomical imaging modality into a molecular and cellular imaging tool.

So, when you see very bright, clustered signals after microbubble injection, it often indicates the presence of something biologically distinct — such as a tumor. This development marks a major advancement in medical ultrasound, extending it from traditional imaging into targeted diagnostics.

slide42:

Now let's move one step further — using microbubbles not just for imaging, but also for therapy.

This concept is known as acoustically enabled drug delivery. The idea is both simple and powerful. Suppose we inject a drug together with microbubbles into the bloodstream. These microbubbles circulate throughout the body, including through the small vessels, or microvasculature, around a tumor site.

When we know the approximate location of a tumor, we can focus ultrasound waves precisely on that region. The ultrasound energy — tuned to the right amplitude and frequency — causes the microbubbles in that area to oscillate vigorously. If we increase the energy slightly, the bubbles rupture or collapse, releasing their drug payload right at the target site.

This process delivers the medication directly to the tumor, minimizing side effects and maximizing treatment effectiveness. It's like having millions of tiny delivery capsules that release the drug only when and where an ultrasound activates them.

This technique opens the door to an entirely new field known as theranostics — combining therapy and diagnostics. With this approach, ultrasound imaging can both visualize and treat disease simultaneously. The precision and non-invasiveness of this method make it an extremely promising direction for the future, particularly for treating cancers and other localized diseases.

So, using ultrasound and microbubbles together, we can not only image internal structures in great detail but also control drug delivery dynamically — by sound waves themselves. This is an exciting frontier in biomedical engineering and medical imaging.

slide43:

Alright, let's wrap up today's lecture. As we conclude, I'd like to give you the homework assignment for Ultrasound I.

You’ll be working on two related problems, both designed to reinforce your understanding of the transducer and matching layer concepts we discussed earlier. Please refer to your green textbook for additional details and supporting equations — it’s not difficult, and I’m confident you can handle it.

The first problem, number three point five, asks you to plot the transmitted frequency spectrum of an ultrasound beam from a transducer operating at a central frequency of one point five megahertz. Assume that the transducer is damped, and then repeat the plot for the beam returning to the transducer after passing through tissue and being backscattered.

The second problem, number three point six, focuses on the concept of the matching layer, which is one of the most important topics we covered today. The question is:“To improve the efficiency of a given transducer, the amount of energy reflected by the skin directly under the transducer must be minimized. A layer of material with an acoustic impedance Z M L is placed between the transducer and the skin. If the acoustic impedance of the skin is denoted by Z S, and that of the transducer crystal is denoted by Z C, show mathematically that the value of Z M L that minimizes the energy of the reflected wave is given by Z M L equals the square root of Z C multiplied by Z S.”

This derivation is straightforward, and you can find the full explanation in your textbook. The due date for this homework is one working week from today, so please make sure to submit it on time.

We’ll continue our discussion next time, where we’ll wrap up our study of ultrasound imaging. Remember — the matching layer is crucial, both conceptually and practically, because it determines how efficiently energy moves from the transducer into the body. Understanding it deeply will also help you as we move toward more advanced imaging systems in the upcoming lectures.