Title: Thermoacoustic Engines and Refrigerators

RECEIVED

FEB 14 1997

OSTI

Author(s): G. Swift

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Submitted to: Third Joint Meeting/Acoustical Society of America and
Acoustical Society of Japan, Honolulu, Hawaii
Dec. 2-6, 1996

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

MASTER

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Thermoacoustic engines and refrigerators

Greg Swift

Condensed Matter and Thermal Physics Group
Los Alamos National Laboratory

Lectures given December 1996

Joint ASA/ASJ meeting,
Honolulu

Contents

1	Overview and demonstrations	2
2	Oscillatory pressure and velocity	8
3	Power	21
4	A detailed example	27
5	Efficiency	31
6	Beyond the acoustic approximation	37
7	Practical details	45
8	Summary	55

This transcript was prepared from a "practice" lecture series at Los Alamos in November 1996, before the "real thing" in Honolulu in December. Thanks to Pam Rockage for typing it, and to David Gardner, Mike Hayden, Bob Reid, and Bill Ward for listening to it.

In what follows, "vg" denotes a viewgraph number and "ca" denotes a computer animation. The entire "home-study" course package consists of

- this transcript with bibliography
- the viewgraphs (without photographs)
- a floppy disk with computer animations for dos-style personal computers
- a set of homework problems (most due to Steve Garrett)
- a list of symbols

1. Overview and demonstrations

Let's start with a review of some things that should be familiar to you from undergraduate course work. This will serve as orientation, to show some of the ideas we will combine and elaborate on to build up our understanding of thermoacoustic engines and refrigerators.

(vg 1.1) Remember from thermodynamics class[1, 2] that there two kinds of heat engines, the heat engine or the prime mover which produces work from heat, and the refrigerator or heat pump that uses work to pump heat. Here the circle represents the device, and the device operates between two thermal reservoirs at temperatures T_{hot} and T_{cold} . In the heat engine, heat flows into the device from the reservoir at T_{hot} , produces work, and delivers waste heat into the reservoir at T_{cold} . In the refrigerator, work flows into the device, lifting heat Q_{cold} from reservoir at T_{cold} and rejecting waste heat into the reservoir at T_{hot} .

The laws of thermodynamics put bounds on the efficiency of such devices. The first law of thermodynamics is just energy conservation, which says that, if the device is in steady state, what goes in must come out. So the energy going into the engine has to equal the sum of the energies going out. The second law of thermodynamics is a little more subtle. It says that the entropy of the universe can only increase or stay the same; it can never decrease. The universe here consists of three pieces: the two reservoirs and the device itself. If the device is in steady state then its entropy doesn't change with time, so with respect to entropy we need only consider the two reservoirs. So the entropy increase of the cold reservoir Q_{cold}/T_{cold} has to be greater than or equal to the entropy decrease of the hot reservoir Q_{hot}/T_{hot} . If you combine these two equations by eliminating the uninteresting variable, Q_{cold} , you get an inequality for the efficiency. The efficiency is what you want—the work the engine produces—divided by what you have to spend to get it—the heat that flows out of the hot reservoir. Combining the two equations shows that the efficiency is bounded above by this ratio of temperatures, which is called the Carnot efficiency.

Similarly, for the refrigerator you can go through the same sort of arguments. There, the relevant efficiency is called the coefficient of performance: what you want—heat of refrigeration—divided by what you have to spend to get it—the work that you put in. The coefficient of performance is bounded by this ratio of temperatures.

(vg 1.2) Here is another one of the threads that we'll weave together in this lectures series: sound waves.[3, 4] This should also be familiar to many of you; if not, we will return to it in great detail later. In acoustics, we often express quantities like pressures in complex notation, with the pressure expressed as some mean value plus the real part of a complex function of x times $e^{i\omega t}$. We will use this style of notation throughout the lectures. The simplest acoustics derivation goes like this. Newton's law for a gas looks like this: Force on the gas is the pressure gradient, which equals the mass ρ_m times the acceleration $i\omega u_1$. The continuity equation for a gas says that mass is conserved, so that the time derivative of mass density at any point, $i\omega\rho_1$, arises from the gradient of mass velocity at that point. Many of you know that if you combine those two equations with an equation of the state for the gas that says how pressure and density oscillations are coupled together, you get a simple wave equation. (Helmholz equation in this case since we've turned the time derivatives into ω 's.

Amazingly, it wasn't until about 20 years ago that anyone figured out how to form the correct wave equation for sound waves propagating in a duct in the situation where there is a temperature gradient in the duct along the direction of propagation, and thermal contact

with the side walls of the duct is important. This wave equation[5] (more properly: Helmholtz equation) was published by Nikolaus Rott in 1969. You can see pieces of the original wave equation in it here and here. The other factors in this equation have to do with thermal and viscous contact with the walls of the duct.

(vg 1.3) A third thread to bring into our discussions is the engineering to simplify heat engines and refrigerators by elimination of moving parts. Here is a Stirling cycle machine[6, 7] as an example. In 1816 Stirling figured out how to make these engines. They had lots of moving parts: rotating crankshafts, moving connecting rods, reciprocating pistons. The mechanical parts severely dominate the thermal parts, in volume and in weight, in the early Stirling machines. By the way, this little Stirling engine was the most commercially successful Stirling engine in history.

In recent decades people have started figuring out ways to eliminate moving parts from such devices. In 1969 William Beale was thinking about resonance effects as moving pistons bounced against gas compressibility in Stirling machines, and one day he realized that you could get rid of the crankshaft and the thing could keep on working! He has been largely responsible for the early development of what are called free-piston[8, 6] Stirling engines and refrigerators, in which the moving pistons bounce against gas springs in resonance, and other moving parts such as connecting rods and crankshafts are eliminated.

Later, Peter Ceperley[9, 10] realized that the phasing between pressure and velocity in the thermodynamic elements of Stirling machines is the same as the phasing between pressure and velocity in a traveling acoustic wave. And so he proposed eliminating *all* moving parts, and basically using acoustics to control the gas motion and gas pressure. Here's one of his figures showing ducts with two heat exchangers and a regenerator and no pistons or any other moving parts.

(vg 1.4) So what we want to try to do here is weave all of these threads together into some kind of coherent picture, with both physical and intuitive foundations. A picture of thermodynamic engines and refrigerators, with the kind of mathematics that Rott developed, and with some of the features of Beale and Ceperley, namely the elimination of moving parts whenever possible to achieve simplicity and reliability.

We'll be talking about devices where some kind of acoustic resonator determines the gas motion and pressure, where heat exchangers exchange heat with external heat reservoirs, and where a stack or a regenerator in between the heat exchangers make the device function as a thermoacoustic standing-wave device or a Stirling machine.

The key dimensions to consider are going to be shown in more detail on the next viewgraph, but basically we want to be thinking in terms of the wavelength of sound along the direction of acoustic propagation in these devices and the thermal penetration depth in the perpendicular direction over which heat must be transferred each cycle of the oscillation. (vg 1.5) So here are some of the important length scales to think about in thermoacoustics.

Along the wave-propagation direction, the direction of motion of the gas, the x direction, the wavelength of sound is of course an important dimension or length scale. Typically when the gas itself provides the resonance behavior, as in most standing wave thermoacoustic devices, the whole length of the apparatus might be a half wavelength or a quarter wavelength; but when mechanical components participate meaningfully in the resonance behavior, as in free-piston Stirling systems, the size of the system is typically much smaller than the wavelength. In both cases, the lengths of heat-exchange components are much shorter than

the wavelength.

Another important length scale in the x direction is the gas displacement amplitude $|x_1|$, which is the velocity amplitude $|u_1|$ divided by the angular frequency ω of the wave. This displacement amplitude is a very large fraction of the stack length or regenerator length, and usually larger than the lengths of the heat exchangers. The displacement amplitude is always shorter than the wavelength.

Well those are the dimensions along the direction of motion of the gas. Perpendicular to the direction of motion of the gas, the characteristic lengths are the thermal penetration depth δ_κ and the viscous penetration depth δ_ν . These tell us how far heat and momentum can diffuse laterally during a time interval of the order of the period of the oscillation over π . So for gas at distances much greater than these penetration depths from the nearest solid boundary, the sound wave doesn't feel any thermal contact or viscous contact with the solid boundaries, but in parts of the apparatus whose lateral dimensions are close to the viscous and thermal penetration depths, the gas does feel both thermal and viscous effects from the side walls.

If you form the ratio of these two penetration depths you see that the frequency and the density cancel out so you end up with ratio of viscosity and heat capacity to thermal conductivity. This ratio is called the Prandtl number of the gas, and it's close to one for typical gases, so viscous and thermal and penetration depths are about the same size and you really can't have one without the other; you can't have thermoacoustic heat engines without substantial viscous effects.

In ordinary audio acoustics, like the sound waves we're communicating with, the displacement amplitude of gas is much smaller than the thermal and viscous penetration depths, which in turn are much smaller than the wavelength. In thermoacoustic engines and refrigerators typically the first inequality is switched around, so that the gas displacement amplitudes are typically much larger than the penetration depths, but still much smaller than acoustic wavelengths.

I brought a couple of demonstrations[11] of standing wave thermoacoustic systems. The first one is an engine, a prime mover, so you should have in mind this picture (repeat vg 1.4). So we need a heat source at high temperature—that's the flame we're applying here—and a heat sink at low temperature—that's the cooling water flowing through here. The acoustic resonator is a quarter wavelength long—see, the tube is closed at this end, and open at this end, so there's a velocity node here and a velocity antinode here. In this region, there's a stack between two heat exchangers, just like on the viewgraph. The stack is about 30 stainless-steel sheets, aligned this way, and the heat exchangers are little copper strips, to transfer heat between the resonator body and the ends of the stack. The work output of the engine is the sound we hear, which the engine radiates into the room.

There are a few games we can play with this engine. First, if I use this extension tube to lengthen the resonator, the resonance frequency will go down—just like organ pipes, longer is lower pitch—but the engine still operates ok. Second, what if instead I stick this pencil in the end? Will that make the resonance frequency go up, or down, or stay the same? If you guessed wrong, you have a homework problem to look forward to! Third, notice what happens if we plug the end, and then quickly remove the plug. Listen for the buildup of the sound—it takes a good fraction of a second. When the plug is in place, the oscillations are gone—it's not that they are sealed into the resonator so we can't hear them, no, they really

are gone. Then when we pull away the plug, it takes a while for the oscillations to build up. We'll see later that thermoacoustically generated power is proportional to the (square of the) amplitude of the wave itself, so it has to pull itself up by its bootstraps, slowly.

Now a standing-wave refrigerator demonstration. It is like this figure on the viewgraph. Here is the stack, with heat exchangers on each end. In this demonstration, we don't have any thermal reservoirs, because the heat exchangers aren't connected to anything in the outside world; we just have a differential thermocouple on them, so we can observe the temperature difference between the two heat exchangers. The resonator is of course this tube, and the source of acoustic work is this loudspeaker.

Now I'll turn on the electric power to the loudspeaker, tune the frequency while listening for resonance, and...you see the temperature difference building up, as heat is pumped thermoacoustically along the stack. On a good day, this will generate a 50 degree temperature difference.

Next let's look at a few examples of real thermoacoustics research hardware, to get you better oriented. (Sorry, there will be no paper copies of photograph-viewgraphs.) There's a big difference between the real world of experiment and practical application and the imaginary world of proposal writing, concepts, and theory. This first example indicates that quite clearly. Here is the realm of proposal writing and concepts. The idea, which seemed so simple and elegant, was to have a thermoacoustic engine in the ocean as a useful sound source[12] for the sonar guys who want to have microphones and listen for reflected sound to figure out where the enemy submarines are. Clean, elegant concept: cheap, reliable sound source. Now here's a photo showing what happened when this turned into a real project. The way you can tell a real project is it's really ugly and gross and miserable and complicated. The picture is so poor because it was such a miserable day that no one was really interested in photography. These are some people from Los Alamos and the Navy and the Naval Postgraduate School getting ready to put this thermoacoustic engine in a lake; at a navy sonar test facility in a very cold lake in a northern part of the United States in the winter. The three people in this photo are wearing six hats because it was so cold and miserable; so cold that the photographer wasn't even asking everybody to pose and smile and all that. Anyway this is an example of a rather large thermoacoustic engine. The resonator extends down past the end of these people, so the acoustic wavelength is much longer than in our little demonstrations, which means the frequency was much lower than in the demonstrations. The stack and heat exchangers are hidden under all this stuff.

Next, here is a posed picture—but don't worry, it was a real project.[13] This was a thermoacoustic engine driving a pulse tube refrigerator, assembled and tested at Tektronix in Portland, Oregon. The photograph shows the same hardware as is shown in the line drawing down here. Two thermoacoustic engines here and here taking heat from red hot heat sources here and here and delivering acoustic power into this half wavelength resonator from here to here with acoustic power flowing into this side branch driving a two stage orifice pulse tube refrigerator. This operated at 350 Hertz with 30-atmosphere helium gas. The half wavelength in helium at that frequency is about a meter or so. The plate spacing here in the stack is about a quarter of a millimeter; the penetration depths are less than that, something like a tenth of a millimeter. We'll return to this example in very great detail in lecture 4.

Another example: This is a standing wave thermoacoustic refrigerator build at Ford motor company in Michigan. It has a very large loudspeaker down in this lower can driving basically

a quarter wavelength resonance in this tube with velocity node here and pressure node here. This ran at about 400 Hertz with 10 atmosphere helium gas, or at 260 Hz with a mixture of 80% helium, 20% argon. The cooling power was about 100 Watts at a temperature of about 0 degrees C.

(vg 1.6) Here's a color schematic of another, recent thermoacoustic refrigerator[14] that will give you a better idea of the greatly different length scales involved in these things. This is a half wavelength resonance through this basically U-shaped pipe going through from here to here, using a helium-argon mixture. Two loudspeakers here and here drive the standing wave, and there are two sets of stacks and heat exchangers stacks here. The waste heat was removed by cooling water at these heat exchangers while these heat exchangers at the cold end cooled the cold circulating fluid loop. This piece of hardware is physically about 30 or 40 cm tall, and has roughly the right amount of cooling power for a household refrigerator: about 200 watts or so at zero degrees C. Again notice that the thermal penetration depth, which is a fraction of the width of the gaps in the stack, is a tiny fraction of wavelength, which is about double the length of the resonator.

(vg 1.7) The first really nice standing wave thermoacoustic refrigerator[15, 16] was built by Tom Hofler in the early 1980's and it looked like this. One loudspeaker and essential a quarter wavelength resonance from here to here. One stack; hot heat exchanger; cold heat exchanger. This shows some of the measurements that he made on that system. The cold temperature got as low as seven tenths of the hot temperature with the hot temperature at room temperature— 300 Kelvin—which makes the cold temperature about 200 Kelvin. The coefficient of performance, which you might remember from the first viewgraph is a measure of the efficiency, divided by Carnot's limit on the coefficient of performance is in the range of 0.1 to 0.2 for this refrigerator. The work used for the plots is the acoustic power going into the resonator, not the electric power going into the loudspeaker, so the "real" COP is lower than this because the loudspeaker is not 100% efficient in converting electric power to acoustic power. The points here are measurements taken from that apparatus and the lines of the results of the calculations that just use the geometry of the apparatus and the properties of the gas. This is typical of the agreement you can expect between calculations and real hardware, if you're as smart as Hofler.

(vg 1.8) So with those examples in mind, let's return to the overview picture for just a minute and think about what the advantages and disadvantages of such systems might be relative to the energy conversion devices that are in widespread use in the world, such as the internal combustion engine, the steam turbine, the reverse Rankine or vapor compression refrigeration cycle, etc. All of those systems that are in widespread use in the real world have a lot of moving parts and from that point of view you might think of them as being less reliable and more expensive than thermoacoustic devices; the reliability and low cost of these conventional devices are due to decades of investment in engineering development. Thermoacoustic devices have the immediate potential to be very reliable and low cost because they have no moving parts, no exotic materials, no close tolerances, etc.

The decades of investment in engineering development of conventional devices has also pushed their efficiencies very high. Thermoacoustics is relatively young and immature; it hasn't yet enjoyed such dramatic efficiency improvements, so its efficiency is a little lower than that of conventional devices. And in the case of thermoacoustic systems with gas-based resonators, we're stuck with a large size. So these might be regarded as shortcomings in

technology.

There's another important feature of thermoacoustics that I don't know whether to list as an advantage or a shortcoming. In the workhorse technologies, if you know exactly how the mechanical parts are moving, then you know exactly how the gas is moving and it's really simple to understand what's going on with the gas dynamics and thermodynamics. For the thermoacoustic systems, you don't always have such a direct intellectual grip on the gas motion, which makes it much more challenging to design these things and to build them and to instrument them and make them do what you want. I think of that as an advantage because it's great fun to try to figure it out. But it might be that it's a disadvantage because it's very difficult to diagnose thermoacoustic systems that are behaving in unexpected, undesired ways.

(vg 1.9) Let's say a little more about efficiency and cost, because those are not independent of each other. In any technology, you can choose to spend more when you build the hardware, in order to make the technology more efficient. For instance, qualitatively, the efficiency vs capital cost of a Stirling refrigerator might look something like this, with very high efficiency possible if you're willing to spend a huge amount of money. But in the real world, energy costs are low enough that no one chooses to go with hardware that operates at this point. Everyone chooses some compromise between construction costs and operating costs; the real world might choose to operate here on this curve. So when a new technology comes along, it can sometimes take over in the real world in spite of lower "ultimate" efficiency if it is significantly cheaper. So I think one of our jobs in the thermoacoustics game is to keep moving our technology curve to the left here, to give the real world more choices of what technology to choose to achieve a given purpose. At the moment it looks like the fundamental limitation on efficiency in thermoacoustics—at least, in standing-wave thermoacoustics—may actually be lower than that of previously existing technologies, but if we can make the cost low enough this our new technology will capture some of the market and there will be applications that now sit on this curve right here that will shift to this curve here.

So with that as a very qualitative introduction to thermoacoustics, I'll outline what we'll do in the course of these lectures.. (vg 1.10) We've just done lecture number one: Introduction and demonstrations. The second segment of the course will deal with the oscillatory pressure and velocity. Basically this includes the fundamentals of ordinary acoustics, with viscous damping and with thermal contact to the side walls of the channel. The climax of that lecture will be Rott's wave equation. The third lecture will add concepts of power to the pressure and velocity results, because power is what you really care about in heat-engine and refrigerator applications! In lecture 4, we'll spend a lot of time looking in great detail at the Tektronix thermoacoustically driven pulse-tube refrigerator example—I showed you a photograph of it a little while ago. This example will provide real-world numbers and hardware, to solidify our understanding of pressure, velocity, and power. After that, in lecture 5 we'll look seriously at the efficiency of these technologies, and develop a formal way to account for just exactly what sources of loss are responsible for low efficiency. In lecture 6, we'll spend some time looking at reality beyond the acoustic approximation: things that real thermoacoustic devices do, which are not included in the acoustic-approximation mathematics developed by Nikolaus Rott and others following the same principles. And finally the 7th lecture will cover practical details: Construction details, methods of measurement and diagnostics on thermoacoustics, and calculation methods.

These lectures present thermoacoustics from a particular point of view. Intuition is important because it helps us humans organize our thoughts. Mathematics is unavoidable, because it is the language with which scientists and engineers communicate, and it allows us to interpolate and extend our knowledge quantitatively. But experiment is the source of all real truth. These lectures attempt to weave intuition, mathematics, and experimental results together. The lectures put the most emphasis on the mathematics, because without this common vocabulary we can get nowhere. Intuition gets second-highest emphasis, because oral lectures with questions and answers lend themselves well to intuition-building. Experimental results get the least emphasis in these lectures, because of the short time available. But please remember that the results presented here mathematically really are distillations of many experiments, just as what you learned in undergraduate fluid mechanics or thermodynamics appeared very theoretical in the textbooks but was in fact well established by experiment. In my laboratory, experiments are the most important and time-consuming activity.

These lectures are also strongly biased toward my own personal viewpoint about thermoacoustics, with undue emphasis on Los Alamos methods and results. For example, in the realm of thermoacoustics software, you'll find only DeltaE calculations here, even though I believe Sage and Thermoacoustica are equally valid. I'm sorry, but I can't imagine trying to present a more balanced viewpoint right now, because that would take a great deal more effort, and I've already worn myself out preparing these lectures, as narrow as they are!

References will not appear on viewgraphs; look for them here in the transcript, in the bibliography at the end. Often, a reference will lead you to a series of references; for example, I haven't cited all of the papers in the important series by Rott, but you can find the others listed in the references of the two or three Rott papers cited here.

(General overview references: fluid dynamics,[17, 18] thermodynamics[1, 2, 19], standing-wave thermoacoustics[20, 21, 22, ?], pulse-tube refrigeration[23].)

2. Oscillatory pressure and velocity

This is the second lecture, dealing with the dynamics of gas motion in thermoacoustics: velocity oscillations and pressure oscillations, with thermal conductivity and viscosity. [3, ?, ?] It's going to be a bit mathematical, but I don't see any way around that. We'll try to keep the math to a minimum when possible, but when you try to apply this lecture material to real problems you'll find that mathematics is unavoidable.

We'll start with the simplest sort of thermoacoustic problem as a warm up problem: the ordinary attenuation of sound propagating in a duct[3]. You probably know that sound attenuation has viscous and thermal contributions. We'll start with viscous effects.

(vg 2.1) Consider the following problem: Suppose you have acoustic oscillation along a solid boundary, with gas motion in the x direction. We're concerned what the y dependence of that gas motion is, due to viscous interaction with the wall. If you start with the full momentum or Navier-Stokes equation in all its glory, it looks like this. (I'll refer to this sometimes as the Navier-Stokes equation, and sometimes as momentum equation, throughout the lectures.) We will apply the boundary condition that the gas velocity is zero at the solid surface. In acoustics we always linearize problems, simplifying them by assuming that all the variable such as pressure, density, and velocity are the sum of the average or mean value (subscript m) and an oscillating part (subscript 1) that oscillates at angular frequency

$\omega = 2\pi f$ in time. If you substitute such expressions in here, and keep only terms that have just one subscript 1, you get the acoustic approximation to the momentum equation. It really does look like Newton's law, $F = ma$. There's mass times acceleration, and there's the sum of the forces on the little mass element—both the pressure force and the viscous force. The solution to this little differential equation for $u_1(y)$ with the boundary condition that $u_1 = 0$ at $y = 0$, at the solid boundary, is a simple expression. The first thing to see is that the characteristic dimension in the y dimension, which tells how abruptly the velocity can change in the y direction, is what we call the viscous penetration depth, the square root of twice the viscosity over the angular frequency and density.

I think it's hard to interpret or visualize this sort of a complex function with a complex argument. To get a better feeling of what this function looks like in time and in space, I have a computer animation of it. (ca viscous) Suppose we have a long cylindrical pipe like this, full of gas, with an oscillating piston at one end to drive the lowest-frequency resonance in the gas. Think of the coordinates as the same as is the previous viewgraph, with x measured this way along the system and y measured from the wall that way.

If you're driving this gas resonator at the lowest resonance frequency, then a half wavelength of sound just fits the length from here to here; this is what the pressure and velocity look like as functions of time and space. Basically you see a half of a wave there with the velocity and pressure out of phase both in space and in time. If you look really closely you can see a little bit of nonzero velocity here at the end where the piston is moving, in contrast with the situation over here where the velocity is really zero because the left wall doesn't move.

If we put some markers in the gas to show how individual portions of gas move, then we can appreciate some more features of the standing wave. See that when velocity is highest in the negative direction, the gas moves most quickly to the left. When it's highest in the positive direction, it moves most quickly to the right; and when the velocity passes right through zero the gas is momentarily stationary.

After the time when the gas has moved to the right, the gas density is highest at the right end of the resonator, and lowest at the left end of the resonator. That's when the pressure is highest at the right end and lowest at the left end, because the equation of state of the gas links pressure to density.

Okay, so we're supposed to be understanding viscous effect here, so let's get on with it. For that, we will have to magnify a small region near the wall, a region with dimensions with of the order of the viscous penetration depth, to see closely what's going on there. Because the way this large-scale frame of the animation is shown can't be right, because it shows the gas slipping along right at the wall, and you know that it can't slip right at the wall—the velocity has to be zero right at the wall.

When we magnify the zone within the circle, this is what the motion of the gas looks like. Far from the wall, this matches to the moving lines in the previous frame, where viscosity is irrelevant. Right at the wall, the gas can't move at all, because viscosity locks it to the wall. This strange shape in between, as a function of position and time, is that solution shown in the previous viewgraph, with the exponential of $(1 + i)y/\delta_v$. Each tick mark here represents one viscous penetration depth. By the time that you get to about 4 penetration depths, you're very close to spatially uniform displacement back and forth, with no viscous shear. The really interesting, complicated action takes place right around 1 viscous penetration depth.

(repeat vg 2.1) So that was what this function looks like.

(vg 2.2) Power is of great interest in thermoacoustics, so we should go a little bit beyond just what the velocity looks like in time and space, and think about dissipation of power by viscous effects. Gradients in velocity (with viscosity) dissipate acoustic power, and in this geometry the dissipation per unit volume is given by viscosity times the derivative of the velocity this way (x) with respect to this direction (y). It turns out that the time average of that dissipation per unit volume is largest right next to the surface, where the velocity gradients are, on average, the highest. The time-averaged dissipation falls off exponentially with y .

I want to stop at this point and review the notation [3, 4, 22] that we have used in the last couple of viewgraphs, because we're going to use this notation over and over again in these lectures.

(vg 2.3) Suppose you have a pure traveling wave, so the pressure as a function of position and time could be written as an average pressure p_m plus some amplitude A times cosine of frequency times time, with some arbitrary phase. That would be the most arbitrary wave traveling in the positive x direction. If it's not simply a rightward traveling wave, but a more general wave, you can express it as an average pressure plus an x -dependent amplitude and then a cosine of frequency times time plus some phase which also depends on x . Often in acoustics we choose to write that in the following way. We write this x and t dependent function as the real part of a complex function of x times $e^{i\omega t}$. If you work out the complex arithmetic, you can convince yourself quickly that the magnitude of this complex function $p_1(x)$ is $A(x)$, the amplitude of the wave. And the phase of this complex function $p_1(x)$ keeps track of the phase of the wave. If you've had a formal acoustics class you've seen this notation a lot. All the variables—not just pressure, but also temperature, velocity, density, etc.—will also be written in the same way, breaking each into the sum of a mean value and an oscillating part, with the oscillating part expressed with complex functions to keep track of the magnitude and phase of the oscillations.

Often we get lazy or forgetful, and don't write $\text{Re}[]$ explicitly, and we just say that $p_1(x)$ is the oscillating pressure; but of course we really always mean that $\text{Re}[p_1(x)e^{i\omega t}]$ is the oscillating pressure.

Back to dissipation of sound. The viscous effects that we talked about a minute ago are one component of acoustic dissipation. (vg 2.4) Thermal effects are the other component of acoustic dissipation at a solid boundary. Here's the prototypical picture. Imagine an acoustic oscillation near a wall, toward and away from the wall. The pressure out here is oscillating and the gas is moving toward and away from the wall in response to the oscillating pressure, with the density of the gas going up and down. The equation of heat transfer in all its glory[?] has heat capacity per unit volume times time derivative of temperature, minus pressure derivatives, equaling divergence of thermal conductivity terms and then some velocity terms that don't concern us too much. Generally in thermoacoustics the appropriate boundary condition is that the solid wall has high enough heat capacity to maintain an isothermal boundary condition on the gas.

So we do our usual acoustic-approximation substitution. Substitute for each variable, such as pressure p and temperature T , the usual acoustic approximation and keep only the first order terms. You get a differential equation for the oscillating temperature T_1 that looks just like the differential equation for oscillating velocity u_1 that we had on a minute ago. The

same differential equation, so it has the same solution. Except that now we have the thermal penetration depth instead of the viscous penetration depth as the characteristic length of interest.

(ca thermal) We'll look at this complex solution for T_1 with a computer animation, to appreciate what it really means in time and in space. Here's the standing wave in the cavity again. (I was lazy and didn't draw the oscillating piston this time, but there has to be one to drive this standing wave if there is any dissipation.) Again, the gas moving back and forth, with oscillating density, oscillating pressure, and now also we see oscillating temperature. The temperature goes up when the pressure goes up, because in wide-open spaces the gas oscillates adiabatically. Now we're going to be interested in the details here, so close to the solid wall that the gas cannot oscillate adiabatically because the solid wall's high heat capacity imposes an isothermal boundary condition. We're going to be interested in what happens to the temperature right up against that boundary. The way it's shown in this frame is obviously wrong, because it shows the gas temperature going up and down right up against the wall.

The next frame shows this yellow zone magnified. Out here is the temperature moving up and down at large distance from the wall, where the temperature just experiences adiabatic oscillation in response to the pressure oscillations. Those are in phase: you can see that this line and this line are moving in phase with each other. But right at the wall the heat capacity of the solid maintains an isothermal boundary condition, so the gas temperature can't oscillate. In between, the temperature as a function position and time is quite complicated. Each tick mark here is one thermal penetration depth, so, just like in the viscous case, by the time you get out to about 4 penetration depths you're far enough from the wall that you can't feel the influence of the wall.

In the viscous case a few minutes ago, we ended up talking about the dissipation of acoustic power by viscous shear the wall. There's also dissipation in this thermal picture. It's harder to get a grasp on it intuitively, but we'll try to do so. The picture that I want you to have is the following. Think about whether there's dissipation in this region. You can ask this question: if you put an imaginary piston out here that moves with the gas, will that piston do work on what's in front of it or not? Of course as the piston moves toward the wall it does work on the gas between it and the wall, and as it moves away from the wall it absorbs work from the gas, so what I really mean is this: Averaged over a cycle of the oscillation, does the piston do net work? If the gas in front of it were perfectly springy, then this piston would not do any net work averaged over one cycle. But if the gas in front is in some way a hysteretic spring, then there will be some work done by this piston, and we would say that it is due to dissipation by thermal effects in this thermal boundary layer.

If you want to compute the work done by a piston, you perform a cycle integral of pressure d volume. This little dot is tracing out the intersection between the pressure trace and the volume trace as a function of time. So the area in the little loop is the cycle integral of pressure d volume. So that slim little elliptical area is the work that the piston does on the gas in front of it. It's not zero; it has nonzero area. The reason that it's not zero has to do with the thermal relaxation of the gas to the wall. The gas in here, right at the wall, is experiencing isothermal oscillations, which are reversible. And this gas, very very far from the wall, is experiencing adiabatic oscillations, which are also reversible. But in between, the gas is confused about whether it should be experience adiabatic or isothermal oscillations, so it gets compressed sort of adiabatically and then thermally relaxes to the wall, losing a little

bit of the stiffness in its springiness, and then it gets expanded, again almost adiabatically, and again it thermally relaxes to the wall, again losing some of the stiffness in its spring. The thermal relaxation causes hysteresis in the spring.

You can ask what does the area of this loop look like as a function of position; what would that area be if we put the imaginary piston here instead of where it is? The answer to that question is this curve. Of course if the piston is right up against the wall it's not moving, so there's no work done at all. A penetration depth away, or two, or three, it does more and more work. If you move the piston farther and farther away you get no additional work done, because the gas in between this location and this location is a perfect spring: it's adiabatic. You can see that a little bit more clearly by considering the slope of this curve, which essentially gives the dissipation per unit volume in the gas. That dissipation is the highest here, just a little less than one thermal penetration depth away from the wall, where the slope is the highest. By the time you get out to several penetration depths there just isn't any dissipation in the gas.

It's interesting to notice that right about here, about four penetration depths from the wall, the gas actually *does* work instead of dissipating it. That's just the way the phasing between the motion and the pressure and the temperature and everything else works out.

(vg 2.5) I'll re-emphasize the thermal dissipation. Here is the solution for the oscillating temperatures as function of position. I want you to imagine a little blob of gas of volume V moving back and forth near the wall under the influence of oscillating pressure. One extreme of its motion and size are shown here in red, and the other extreme of its motion and size are shown here in blue, where it is squeezed closer to the wall and smaller. If that little bit of gas that we're focusing our attention on is very very close to the wall, much much closer than the thermal penetration depth, then it's temperature is always the temperature of the wall, which is constant in time. So its pressure oscillates back and forth along a curve following PV equals a constant temperature, like this. If instead this little blob of gas were much much farther from the wall than a thermal penetration depth, then it would be oscillating back and forth along a adiabetic. PV to the gamma equals a constant—that's a steeper trace than the isothermal trace over here. If instead the gas is in between, where its distance from the wall is about a thermal penetration depth, then the gas is confused about whether to be isothermal or adiabetic, and as it chooses somewhere in between those extremes it traces out an elongated loop in pressure-volume like this, with a non-zero area. That area represents the dissipation that this blob of gas is responsible for.

(vg 2.6) For those of you who have had a acoustics class [3] recently enough to remember everything from it: You can tie these ideas that we've just discussed to bulk attenuation of sound rather easily. The textbook expression for the attenuation constant (one over the attenuation length) for sound propagation through free space is this. If you rewrite that in terms of the viscous and thermal penetration depths, you see that you get big viscous damping in bulk sound attenuation if the viscous penetration depth is at all appreciable compared to the wavelength, and big thermal damping if the thermal penetration depth is all appreciable compared to the wavelength. So, as you probably know, for very low frequency, long wavelength sound, bulk attenuation is just not that big. The ratio of penetration depth to wavelength is strongly frequency dependent because the wavelength is proportional to one over frequency while the penetration depths are proportional to one over square root of frequency.

The sound attenuation constant in ducts and the quality factor of a resonator can also be expressed in terms of these penetration depths.

Okay, that's the warm up problem of attenuation of sound. I think you can see the approach that I'm going to take in this segment of the lecture: to alternate between hard-core math and these computer animations that try to provide some intuition.

(vg 2.7) Now we're finally ready to tackle a real thermoacoustics problem. This is the simplest half-way realistic thermoacoustics problem that I can think of. It's not perfectly realistic because we are ignoring viscosity, and (as you should remember) you can't ignore viscosity, because viscous and thermal penetration depths in gases are typically about the same size. Nevertheless, to gain intuition, we going to ignore viscosity for a few minutes. So here's the simplest thermoacoustics problem. A nonzero temperature gradient and a boundary, along which we have oscillating motion, like this. Again we start from the linearized equation of heat transfer, which we saw three or four viewgraphs ago. At that point it did not have this term because there was no temperature gradient along the wall. But now we have this additional term.

So the solution to this differential equation with this new term resembles the one for the previous problem, but with a little more complexity arising from this term. Here's the solution, the product of two factors. The first factor gives the overall magnitude of the oscillating temperature, and the second factor describes the y dependence of that oscillating temperature. The magnitude factor has two terms that are easy to understand. The first term is just the adiabatic temperature oscillation. When the pressure goes up the temperature goes up, and when the pressure goes down, the temperature goes down. The other term is due to motion along x along the temperature gradient. Basically, if you have gas with a temperature gradient, suppose this is temperature and this is position, so this stick represents a graph of temperature vs position, if that gas is moving along x and you're sitting at one point, observing the local temperature, you see the temperature go down, and you see the temperature go up, and you see it go down. That's all this term is.

The y dependent part is the same complex function that we've seen several times before. It's complex enough—in the literal sense of complex with real and imaginary parts—that the combination of this complex y dependent factor with this magnitude factor can give you any sign, phase and size that you want. The magnitude factor is also complex and can have either sign depending on the relative signs of p_1 and u_1 and the relative magnitude of this term relative to this term. So you can get anything that you want out of this simple equation. Much of the richness of thermoacoustics—standing wave behavior, traveling wave behavior, small-pore behavior and well-spaced behavior—arises fundamentally from the complexity of the simple mathematics here on this viewgraph.

(ca fancy) Now we'll return to an animation, to gain intuition about this whole process. Once again, consider a half wavelength resonator, with oscillating velocity, oscillating pressure, oscillating gas position. Again, I've lazily omitted a driving piston. Now we're going to insert a thermoacoustic stack here, and think about what standing wave thermoacoustics has to offer and how it works. We'll zoom in to get a closer look at this small region of the thermoacoustic stack. This is a magnified view of this solid plate here, and we'll focus attention on one small element of gas moving back and forth here, shown magnified here.

First we'll consider the case of zero temperature gradient in the wall, in the stack. In this case, the pressure and volume of the gas element in question does this as a function of time,

following this little white dot. We've seen this before, in the "thermal" animation, with the net work effect of that element of gas being dissipation: That element of gas removes work from its surroundings. The only thing different in this animation that while this element of gas is exchanging heat with the side wall in the y direction, it's moving in the x direction, whereas in the previous animation it wasn't moving parallel to the wall. So, what's going on here is that there's oscillatory heat transfer between gas and the wall, as shown by the red arrows. When the gas comes over to the left, it gets pressurized by the standing wave, so it's heated adiabatically, and its temperature comes up higher than the wall temperature. At that phase of its motion, the gas is hotter than the wall, so heat flows into the wall. That makes the volume of the gas shrink a little, at the high pressure extreme of pressure, as it cools. Similarly, at the low pressure, rightmost extreme of its motion here, it heats up, because it is cooler than the wall. So that oscillating heating and cooling causes oscillating expansion and contraction with the correct phase with respect to the pressure, which in turn causes the dissipation here in the pressure-volume plot.

You have to stare at this by yourself, quietly, for a long time, to really appreciate it, because you have to be able to absorb about five things going on at once.

The overall slope of the pressure-volume plot is due to the standing-wave geometry and the adiabatic temperature changes the gas experiences because of the pressure changes. The fact that the pressure-volume plot is a loop instead of a reciprocating line is due to the thermal relaxation to the wall. It is crucial that the gas we are interested in is approximately a thermal penetration depth from the wall, because that ensures the necessary time-delayed thermal relaxation.

With respect to heat, this element of gas is picking up a little of heat from the solid wall here, carrying along, and depositing it over here. So it moves heat from right to left.

It's not a big step to make a picture like that to turn into a refrigerator. Imagine a slight temperature gradient indicated by this blue line, in a long stack between a cold heat exchanger here and a hot heat exchanger here. You can have a large temperature difference if the stack is long enough. The gas still picks up heat here, deposits heat here, but now it's depositing its heat uphill in temperature just a little. So this is a heat pump or a refrigerator, moving heat from right to left, up the temperature gradient. Overall, every element of gas in this entire stack region, each cycle of the wave, picks up a little heat from solid, moves it a little bit to the left and deposits it again at a slightly higher temperature. So the net effect is to pick up heat from this heat exchanger, and deposit it in the other heat exchanger.

If this temperature gradient is a little steeper, at what we call the critical temperature gradient[20], then nothing interesting happens, because now the oscillating temperature that the gas experiences from the pressure oscillations of the standing wave just exactly matches the local temperature in the solid that the gas sees as it moves because of the standing wave. So there's no longer any temperature difference to drive heat between gas and solid at any time during the oscillation. No heat transfer, no thermal contraction and expansion, no area in the pressure-volume trace.

At an even steeper temperature gradient along the wall, the signs of everything get reversed relative to what they were in the refrigerator. So now even though the gas gets heated adiabatically as it moves left, it's not hot enough to match the local solid temperature when it gets there. So at the left extreme of its motion, at high pressure, it's cooler than the plate, so heat flows into it, which thermally expands it. That thermal expansion gives the volume

a little increase up here when the pressure is high. Similarly, thermal contraction gives the volume a little decrease down here when the pressure is low. The net effect is to make the cycle integral of pressure d volume positive, so that this element of gas no longer absorbs work—it does work on it's surroundings. This is an engine, not a refrigerator. The gas shuttles heat from hot to cold, not cold to hot, and it produces work instead of dissipating it.

A good place for a short break.

A brief review of what we just did: Thermal relaxation processes, across a thermal penetration depth, are key. The stack has gaps of a few thermal penetration depths, and allows gas motion perpendicular to the thermal relaxation. The relaxation plus gas motion can move heat from one end of the stack to another, and can dissipate or produce work. The signs of these powers depend on the steepness of the temperature gradient relative to the critical temperature gradient, which depends only on the gas and the standing wave.

All that was focussed on standing waves. What about traveling waves? Orifice pulse tube refrigerators[23], and other Stirling-like systems[6, 7]? Let's start by reviewing the difference between standing and traveling waves.

(ca travel) Here's a standing wave, like we've seen in the previous animations. I got rid of the endcaps to the resonator, so imagine an infinitely long pipe. Standing waves have velocity and pressure 90 degrees out of phase from each other, both in time and in space.

Here's a traveling wave, traveling to the left. Velocity and pressure are in phase, both in time and in space.

What about real thermoacoustic systems? Typically so far, standing-wave systems have phasing (in time) between velocity and pressure within about 5 degrees of the 90 degree ideal, and traveling wave systems have phasing within about 30 degrees of the traveling-wave ideal. Let's look at a wave that's not ideal. See, it gets complicated. Look in this vicinity: it's not really a velocity node, but there's definitely a minimum there. So it has some standing-wave nature. But you can also see something that sort of travels. You see a peak that sort of moves to the left, across the place that resembles a node; then it sort of dies out. Here's a trough moving to the left and it moves through here and has a diminishing amplitude. This is typical of the actual phasing between pressure and velocity in a Stirling system. But the length of a Stirling system would be far shorter than a wavelength, so you should imagine a short spatial slice here.

(vg 2.8) From the acoustics point of view, the Stirling engine or refrigerator is easy to interpret. I don't have an animation for this, so you will have to use your imagination. This upper picture shows the parts of a typical Stirling refrigerator, and these show the four steps of the thermodynamic cycle. The parts are: a couple of pistons working in cylinders; two heat exchangers, at temperatures T cold and T hot; and in between them is the regenerator, which is kind of like our standing-wave thermoacoustic stack, but with the spacing in the passages much much smaller than the thermal penetration depth, rather than being a few thermal penetration depths in the case of the standing-wave stack.

The four steps of a Stirling cycle refrigerator machine are as follows. The first step is to move this hot piston in, compressing the gas adiabatically, so the temperature of the gas in this cylinder goes up from T hot to something a little bit higher. Then in the second step, both pistons move to the left, so the heat of compression from the first step gets captured in this heat exchanger and rejected out of the system. The regenerator, with its good thermal contact and smooth temperature gradient between hot and cold, then precools the gas to T

cold, so it can accumulate here in the left cylinder at T cold. In the third step, this piston moves out, cooling the gas adiabatically here in the cold cylinder. Finally, in the fourth step, the two pistons move back to the right in unison, so that the cold gas is shoved into the cold heat exchanger, entering it at a temperature a little lower than T cold, so the heat exchanger can absorb heat from the outside world and feed it to the gas, raising the gas temperature from $T_C - \delta T_C$ back up to T_C . That's the refrigeration effect that you're trying to produce. The gas continues through the regenerator, getting preheated back to temperature T hot, so that it accumulates here at T hot ready to start the cycle again.

The regenerator needs heat capacity, just like the stack of standing-wave thermoacoustic systems, in order to store heat for half the cycle. But here, the thermal contact must be really good, because you want to get the heat transfer during the displacement phases of the cycle. In the standing-wave animation we saw a few minutes ago, you want poor thermal contact, because you want the gas to thermally relax *after* the gas motion and pressure change are finished.

The phasing between pressure and motion here is that of a traveling wave. [9, 10] It's traveling to the left here: when the pressure is high, after the compression, the displacement is towards the left; and when the pressure is low, after the expansion, the displacement is towards the right.

With respect to work, the right piston does work on the gas, the left gas piston removes work from the gas. You can see that as follows. The right piston does most of its inward motion under high pressure, after the compression. And it does most of its outward motion at low pressure, after the expansion. So it does net work on the gas during each full cycle. The left piston does just the opposite: It does most of its outward motion at high pressure, and most of its inward motion at low pressure, so it absorbs work from the gas. This is consistent with the picture of the traveling wave going right to left, carrying acoustic power from right to left.

(vg 2.9) In the quest to get rid of moving parts and to achieve higher reliability, the most common simplification of the Stirling machine is the orifice pulse tube refrigerator. It's easy to understand how it works, by comparison to the Stirling refrigerator that we just looked at in the last viewgraph. Averaged over one cycle, a Stirling refrigerator does work on the gas at the hot piston, takes work out of the gas at the cold piston, rejects waste heat at the hot heat exchanger, and absorbs heat at the cold heat exchanger. The orifice pulse tube refrigerator does its thermodynamics essentially the same, but instead of having a mechanical part here at the cold end to remove work, it substitutes a passive, dissipative structure to remove work: namely the orifice here, followed by a reservoir volume or compliance tank. The orifice and reservoir volume are the acoustic analogs of an electrical RC circuit, so the dissipate work just like an RC circuit. Therefore, they are removing work from the cold heat exchanger here. The cold heat exchanger doesn't know whether it's an RC circuit or a piston removing work; it just knows that work is being removed, so it absorbs heat from the load just like in the Stirling machine.

Of course you want the dissipation to take place at ambient temperature, 300 Kelvin or whatever, not at T_C . So you need the pulse tube to act as a temperature buffer between this hot heat exchanger and the cold heat exchanger. The way to think of the gas in the pulse tube is as a moving, insulating piston, which is slightly compressible too.

(ca tadoptr) One thing that we've worked on a lot in Los Alamos is the combination of a

standing wave thermoacoustic engine with a traveling wave refrigerator—a pulse tube refrigerator. [13] Very qualitatively, this is how it looks. Essentially a half wavelength resonator from here to here with a thermoacoustic engine with a stack with spacing of the order of a few thermal penetration depths here, and an orifice pulse tube refrigerator with a regenerator with spacing much smaller than a thermal penetration depth here. The engine accepts heat at the hot heat exchanger at red hot temperatures, rejects heat at the cold heat exchanger at ambient temperature, and thereby produces work that drives the standing wave. The orifice pulse tube refrigerator over here—hot heat exchanger, cold heat exchanger, pulse tube, orifice, and reservoir volume—consumes the work from the engine, and produces cryogenic refrigeration here, rejecting waste heat to ambient temperature here and here. So the complete system is a heat driven cryogenic refrigerator with no moving parts. I showed you a photo of such a system in the very beginning—the posed photo with the three people from Tektronix. We'll return to this system later in very great detail, as a concrete example of calculation methods, real-world numbers, and construction methods.

Another good place for a short break, for questions.

We've been discussing the gas motion and pressure oscillations in thermoacoustics, with gas motion along the x direction, and heat transfer in the y direction causing all the important thermoacoustic effects. I want to reinforce all this one more time, by going over the equations that describe the thermoacoustic process again.

(vg 2.10) Forty-five minutes ago we started out with the Navier-Stokes equation, making the acoustic approximation substitution where each variable consists of an average value plus an oscillating part. We ended up with this differential equation. Boundary condition: zero velocity at solid boundary. At that time, we had one solid, flat boundary, but in general we have channels in a stack with some other geometry, like circular pores or honeycomb. So in general we'd include the z dependence of velocity in the x direction as well as the y dependence. Whatever boundary conditions you have on the side walls on the channel determine the solution for this function $u_1(y, z)$.

Next you can integrate with respect to y and z to get the spatial average of u_1 , which we denote by $\langle u_1 \rangle$, spatially averaged in the y, z plane. That's fine, or you can rearrange the solution to solve for dp_1/dx instead of $\langle u_1 \rangle$ like this. Either way, the details of the solution of the differential equation with the boundary conditions, and the spatial integration, all get buried in this complex constant f_ν . We're not going to spend any time in these lectures deriving f_ν , because other people have done it for various geometries; [5, 24, 25] we'll just look it up and use it.

(vg 2.11) That function f_ν ends up looking like this for various geometries of channel. The upper graph is the real part of f_ν , and the lower part is its imaginary part, as a function of the hydraulic radius of the channel divided by the viscous penetration depth. The hydraulic radius is area divided by perimeter; it ends up being half the gap in an infinitely wide rectangular channel, and half the true radius in a circular channel. For today, just think of it as of the order of the smallest y or z dimension in the problem. You can see that it doesn't matter a whole lot whether you have parallel plate channels or circular pores or rectangular pores. I think that I picked a six to one aspect ratio for these rectangular pores. As long as you're out beyond between two and three penetration depths, you're close to the boundary layer limit, which is the complex exponential solution we looked at earlier, the exponential of $-(1+i)y/\delta_\nu$. The spatial average of that ends up giving $f_\nu = \delta_\nu/(1+i)r_h$.

(vg 2.12) For intuitive purposes, we want to interpret that momentum equation like Ohm's law. Here's the result of the simplification of the momentum equation we had a minute ago. If you rewrite this for a short duct, a duct of length of dx much shorter than wavelength, and area A , and if you use the volumetric velocity U_1 as the velocity variable instead of the spatial average of the velocity $\langle u_1 \rangle$, you can rewrite the equation like this: Pressure difference across this short length of duct is equal to this complex factor times the volumetric velocity through it. That looks like to me like Ohm's law, like voltage drop across an impedance equals the impedance times the current through the impedance. It's really helpful to have this picture in mind.

The details of whether this is a circular pore or a rectangular pore or parallel plates or whatever are all wrapped up in this complex function f_ν , which I've drawn again down here, copying from the previous viewgraph. The behavior of the duct depends a little on what kind of passages you have, but mostly it depends on how large the passages are compared to the viscous penetration depth. If you have very large diameter pipes, like open ducts and resonators or the pulse tubes in pulse tube refrigerators, then you're out at the extreme right edge or off the right edge of this graph. So f_ν —both its real part and its imaginary part—is small, call it zero, so the denominator of the impedance turns into 1 and this impedance Z is purely imaginary and the inertia of the gas in the little duct is all that matters. That's the inertial case, with velocity and pressure difference being 90 degrees out of phase. In really tight passages such as regenerators in Stirling machines, with a very small hydraulic radius, you're down in here with $\text{Re}[f_\nu]$ near one and $\text{Im}[f_\nu]$ near zero, so you must have almost $1 - 1$ here in the denominator of Z ; the small imaginary part turns this ratio into a purely real impedance Z . In that case Z is positive real. If you're in between in r_h , as in standing-wave thermoacoustic stacks, you're in between inertial and resistive. In that case, Z is a complex impedance. If you plot these impedances on a complex plane, with the real part of Z horizontal and the imaginary part vertical, you find wide open ducts here on your imaginary axis and regenerators way out here to the right somewhere with an almost negligible imaginary part compared to a large real part, and stacks somewhere in between with the phase of the impedance Z being somewhere around 45 degrees, give or take 30 degrees.

Well, we've been staring at the acoustic approximation to the Navier-Stokes equation for several minutes now. Let's move on to the continuity equation (vg 2.13), which looks like this. In the acoustic approximation, the continuity equation turns into this: minus the time derivative of the lateral spatial average of the oscillating density equals the x derivative of the spatially averaged mass flux along x . All that says is that, if the density is going up and down at a given point, mass has to be flowing in and out of that point. If you substitute in the equation of state for $\langle \rho_1 \rangle$, so that you express the oscillating density in terms of oscillating pressure and temperature, and solve the equation of heat transfer for the oscillating temperature (which we did earlier today), you can rewrite this continuity equation in terms of p_1 and u_1 only, without T_1 or ρ_1 . This is the result, rewritten with $d \langle u_1 \rangle / dx$ by itself on the left hand side.

The rate of change of x velocity with x has one term proportional to oscillating pressure and another term proportional to oscillating velocity, with the proportionality constants complex and full of f_κ s. But f_κ is nothing but f_ν with δ_κ substituted in for δ_ν , and we've stared at this function already—remember (repeat vg 2.11)?

(vg 2.14) Now we'll look at the continuity equation for a short duct, like we did with the momentum equation. Top line is the thermoacoustic approximation to the continuity

equation, copied from the previous viewgraph. Again we have a short duct, length Δx , area A . A couple of cases are easy to interpret.

Suppose that the duct is very large, so that $r_h/\delta_\kappa \gg 1$. Then both f_κ and f_ν are approximately zero, and this gets really simple. The only thing that survives is "one" in the first term. So, the expression turns into this simple thing: the difference in velocity between the inlet and the outlet of that duct is proportional to the volume of gas in the duct and the oscillating pressure. This is a really simple thing to remember and understand, because all that this is saying is that the gas in there is experiencing adiabatic compression and expansion. The denominator here is the adiabatic compressibility of an ideal gas, which relates the oscillating pressure to the oscillating density, which in turn must be provided by volumetric velocity from either or both ends.

Another case that's really easy to interpret is the opposite extreme, where the hydraulic radius is very near zero, so f_κ is around one. For simplicity, let's pretend that viscosity is zero for now, so we can set f_ν to be zero. So this is like a perfect regenerator for a Stirling machine: Good thermal contact but no viscous dissipation, which would be ideal for a regenerator. Now there are survivors from both terms in the original equation. There's the first term, the p_1 term, which looks the same as it did before but now without the gamma; now we have isothermal compressibility instead of adiabatic compressibility. The second term reduces to this simple expression, which is also easy to interpret. If there is a temperature difference between one end of the duct and the other, and you push mass through, the mass comes out at lower density at the exit where the temperature is higher. So the mass has to be moving at higher velocity where the temperature is higher, by mass conservation. So U_1 has to increase by this amount.

So that should give you flavor of what these terms in the continuity equation are all about. This is the compressibility term. You have adiabatic compressibility for wide open channels, isothermal compressibility for very tight channels, or something in-between for in-between channels. And the details of the in-between come in here in the factor $1 + (\gamma - 1)f_\kappa$. This second term has to do with the thermal expansion of the gas as it moves along the temperature gradient. It's difficult to get an intuitive grasp on the viscous effects here, but the essentials of the thermal effect is clear.

(vg 2.15) It's convenient to think about the continuity and momentum equations in terms of an electrical analog. [3] This viewgraph summarizes my favorite way to make an analogy between acoustics situations and ac electric circuits. You can find this in any standard acoustics textbook. If you make the analogs that the oscillating pressure is like the oscillating voltage in a ac electric circuit, and that the oscillating volumetric velocity is like the ac current, then in both acoustic and electric cases you talk about the complex impedance Z as the ratio of those two quantities. In both cases, you get the time-averaged power by taking the appropriate time-averaged product of the two variables: In the electrical case, the power is half the voltage times the current times the cosine in the phase angle. If you look at the way complex arithmetic works out, you can express that in this form: The real part of the product of these two complex numbers. The tilde above a variable means complex conjugate. In the acoustics case, you have exactly the same kind of expressions, with power being half the product of pressure and volumetric velocity and the cosine of phase angle between them.

The momentum equation for a short duct length Δx and area A , which we looked at a few minutes ago, looks like a series RL impedance. In the momentum equation, we call the

real part of the impedance the resistance, and we call the imaginary part inertance, which is a word chosen to sound a little like the electrical word inductance and a little like inertia, because its origin is inertia in the moving gas. So think of a little duct as having resistance and inertance, which depend on the details of f_ν . It might have a lot of inertance and not much resistance if it's a very fat duct, or a lot of resistance and not much inertance if it's a very skinny duct.

The continuity equation ends up looking like a capacitance to ground in the electrical analog. If you have a different volume velocity going in and coming out of a short duct that experiences oscillating pressure, that's like an electrical circuit in which you have different currents going in and going out, with oscillating voltage in between, like this circuit with a capacitor. The analog to electrical capacitance is acoustic compliance, which is the volume of the acoustic element divided by the appropriate compressibility (which would be the adiabatic compressibility for a big duct, the isothermal compressibility for a tiny duct, and somewhere in between adiabatic and isothermal, depending on f_κ , for ducts that are in between.

(vg 2.16) So what does all that mean for thermoacoustics? Do you want something with inertance, do you want something with resistance, do you want something with a lot of compliance? What do you want? Well, you have some choices.

If you want a traveling-wave thermoacoustics system, a Stirling-like machine, typically you'll find that the phasing between pressure and volumetric velocity is within 30 degrees of that of a traveling wave. And you want good thermal contact in your regenerator so r_h has to be much smaller than δ_κ . If you remember what that graph of f_κ looks like, near its left extreme where r_h/δ_κ is small, then f_κ is very near one. It looks like one minus a small imaginary part. Then r_h/δ_ν is also small, so f_ν is also very near one, so the momentum equation, which gives you the series impedance of the regenerator, looks very resistive. The continuity equation's compressibility term then looks nearly isothermal, and its temperature-gradient term looks much like conservation of mass flux: Instantaneous mass flux into a little piece of regenerator equals mass flux out.

If, instead, you want a standing-wave thermoacoustic engine, typically you'll find that the phasing between pressure and volumetric velocity is within about 5 degrees of pure standing-wave phasing. And you want some thermal contact in your stack, but not perfect thermal contact, so you want r_h/δ_κ of order 1 so that f_κ is somewhere around a half times $1 - i$. The same is true for f_ν , because δ_κ and δ_ν are always about the same size. In that case, the momentum equation, which gives you the series impedance of the little channels in the stack, has comparable resistance and inertance. The continuity equation is somewhere between isothermal and adiabatic compressibility and somewhere between conservation of mass flux and conservation of volumetric velocity through the channel in the stack. The details of "somewhere between" are all buried in the details of f_κ and f_ν .

Finally, there's not much thermo in the thermoacoustics if you're in very large diameter channels. In that case, the f 's are nearly zero, the momentum equation is mostly inertial, and the continuity equation shows nearly adiabatic compressibility and constant volumetric velocity.

(vg 2.17) There's one more little detail I want to add to all of this. The graphs of f versus r_h/δ_κ that I've been showing so far have listed several different geometries of interest mostly to standing wave thermoacoustics people. You can follow through the same kind of analysis for the stacked-screen regenerators that are usually used in Stirling systems. Here's an outline

of the derivation[26]. If you assume that you have a sinusoidal volumetric velocity through your stacked-screen regenerator, and you assume that your handbook data[27] for stacked-screen viscous pressure drop and heat transfer are valid at each instant of time during this oscillation, you can take a Fourier transform of the continuity and momentum equations with those very complicated pressure drop and heat transfer built in, and rearrange the results in a way that give you f_κ and f_ν as a function of r_h/δ_κ (or r_h/δ_ν), the porosity of the screen bed, and the peak Reynolds number of the flow through the screen bed. The results are only valid for small r_h/δ_κ , but that's ok because that's all you care about for Stirling machines anyway. The results look something like this, with slightly different results for different porosities in the screen bed. There's another detail that you can find in the reference: the f_κ that shows up in the compressibility term of continuity equations is actually a little different than the f_κ that shows up in the other term.

(vg 2.18) One more thing that we can do with the continuity and momentum equations is to combine them to produce Rott's 1969 wave equation.[5] Here's the momentum and continuity equation again; if you combine them in such a way that you eliminate u_1 , then you get this beautiful wave equation. (Strictly, I should call it a Helmholtz equation, because the time derivatives have been replaced by $i\omega$'s.) It has pieces of the old familiar wave equation $1 \times p_1$ and two derivatives of p_1 with respect to x . This is one of the principal results of thermoacoustics, and it's equally valid for traveling and standing wave systems. I think that there's more intuition to be gained by studying at the two starting equations separately, rather than the wave equation itself, but that may just be my own personal bias.

3. Power

In the last lecture we looked at the momentum and continuity equations, to understand how the pressure oscillations and the velocity oscillations occur and interact with each other. We spent a lot of time thinking about microscopic details, within the viscous and thermal penetration depths. That's all good, and necessary, but what do you think your boss is really going to care about when you're trying to develop a thermoacoustics engine or refrigerator? Will your boss care that you understand the microscopic details within the penetration depths? No way; we have to go well beyond that. Your boss will want to know how much *power* you'll get, and with what efficiency. So that's what comes next.

(vg 3.1) The first power concept is ordinary acoustic power, which should be familiar to everyone with an acoustics background. [3, 4] Imagine a duct of area A with a sound wave propagating along x , with pressure and velocity waves in the duct like this. If you want to calculate the acoustic power flowing down the duct, you need to make the product of that area and the time average of the product of the oscillating pressure and the spatially-averaged oscillating velocity. (You could think of this as the time average of the spatial average of the product, but we know that p_1 is spatially uniform, so it can be pulled outside the spatial average.)

If you want to consider the details of effects near the walls, viscous and thermal effects within a viscous or thermal penetration of the walls of the duct, you can gain insight by asking how much power is absorbed (or in the thermoacoustics case sometimes is *produced*) in the length dx in this duct. In other words, what is the derivative of the acoustic power with respect to x ? Let's break this product into two pieces with the product rule for derivatives.

[20] Now remember from the previous lecture that our momentum equation give us a nice compact expression for dp_1/dx in terms of u_1 and f_ν . Remember that? It looked like Ohm's law. And our continuity equation gave us a nice expression for $d\langle u_1 \rangle/dx$ in terms of p_1 and $\langle u_1 \rangle$ and f_κ and f_ν , which we interpreted as telling about compressibility and mass flow along a temperature gradient. Now, if you make those substitutions here, you get this big mess. There are three terms. The first one came from the momentum equation, and the other two came from the continuity equation. The f functions are all over the place. Remember that this is what f_κ looks like as a function of hydraulic radius over δ_κ .

The first term, the one from the momentum equation, doesn't have any thermal stuff in it, it only has f_ν , and it represents viscous dissipation of power at the surface of the duct. Basically, the kinetic energy density in the sound wave, a half ρu squared, is getting turned into heat by viscous shear at a rate ω near the wall. The details of how much, and how near the wall, are all tied up in the imaginary part of minus f_ν . over $|1 - f_\nu|^2$. (We derived a similar expression in lecture 2, in the boundary-layer approximation, and talked about it at length while watching the "viscous" computer animation.)

The next term, which comes from the compressibility part of the continuity equation, has $(\gamma - 1)\text{Im}[f_\kappa]$ in it, and no viscous stuff. This is the thermal relaxation term. Basically, the compressive energy density in the sound wave, a half p_1 squared over $\rho_m a^2$, is getting turned into heat by thermal relaxation at a rate $(\gamma - 1)\omega$ near the wall. (The adiabatic compressibility is $\gamma p_m = \rho_m a^2$ in an ideal gas.) The details of how much, and how near the wall, are all tied up in the imaginary part of minus f_κ . (We derived a similar expression in lecture 2, in the boundary-layer approximation, and talked about it at length while watching the "thermal" computer animation.)

So the imaginary parts of the f 's are really important in the first two terms; that's why I highlighted it down here. You can see that you get a large contribution when the size of the pore r_h is comparable to the thermal penetration depth or comparable to the viscous depth.

Notice that the first two terms are always positive, because they have squares of magnitudes of complex numbers, and positive constants, and the imaginary part of minus f . Always positive, which means always dissipative, for the sign conventions we picked here.

The final term in dW_2/dx is the one that comes from the second part continuity equation, which is the part of the continuity equation that's proportional to velocity times dT_m/dx . This term is more complicated than the others. In this one, you can get any sign you want, depending on the phase between $\langle u_1 \rangle$ and p_1 , so it can give power production instead of dissipation. This is the term that gives us all the interesting behavior in standing-wave thermoacoustics. For $dT_m/dx = 0$, this term is obviously zero. As the critical temperature gradient is approached from below, this term causes there to be less and less thermal relaxation dissipation, because this term then has the opposite sign to the second term, the thermal relaxation dissipation term, and so tends to cancel it. For very large dT_m/dx , the third term is larger than the sum of the first two, so the net effect of all three is production of acoustic power—that's a thermoacoustic engine.

It's sometimes convenient to think about acoustic power as a function of position x in a piece of thermoacoustic hardware. [20] In this simple refrigerator, there a lot of acoustic power coming from the left at the driver here, and a little gets dissipated in the duct wall here where there's a little surface area, so the slope here is negative but not too large. When you get into the stack, the power absorption per unit length is greater because of the greater

surface area, so the slope gets steeper. Finally, you get out into the rest of the resonator, with less surface area, and the dissipation per unit length gets smaller again. Outside the stack, only the first and second terms in the equation contribute, because $dT_m/dx = 0$ so the third term is zero. But the first and second terms are fairly small, because r_h/δ is huge so $\text{Im}[-f]$ is small. In the stack, you get r_h/δ of order one, so much larger $\text{Im}[-f]$; in spite of a little cancellation from the third term, you end up with much greater dissipation per unit length.

(vg 3.2) Acoustic power is not the only power that's important in thermoacoustics. We have to think about total energy. So we're going to spend a few viewgraphs deciding exactly what we mean by energy in thermoacoustics. Here are some pictures that should be more familiar to the engineers in the audience than to the physicists, because engineers encounter "control volumes" in their thermodynamics classes. A control volume is basically an imaginary boundary to which you apply conservation of energy.

Here's a first example to motivate us. Suppose you have a thermoacoustic refrigerator and you've wrapped insulation all around it except where the heat exchangers are—the only places that you can exchange heat with the outside world are at the cold heat exchanger and the hot heat exchanger. If you imagine a control volume as shown by this dotted line here, you know that in steady state, time-averaged over a full acoustic cycle, you must have energy conservation: you must have energy flowing into that control volume equal the energy flowing out. What flows in is clearly the acoustic work, the pV work that the oscillating piston does on the gas. So that has to equal the sum of the two outflowing of energies. The heat energy is easy to think about—it's simply heat. It's easiest for me to imagine the heat flowing via conduction through a solid, so I might like to imagine the control-volume boundary being within the metal case that separates the thermoacoustic gas from the outside world. But the energy flowing down the stack is much more subtle; it's not acoustic power, and it's not heat. But we'd better understand it, because it's what comes into the conservation of energy picture here.

Another control volume that's very important to think about sometimes is one like this, that intersects a stack or regenerator in two places, with insulation around the side walls. Here the only energy flows are in and out of these two end surfaces of the control volume. So energy conservation (again, steady state; time-averaged) implies that energy in equals energy out. So energy can't change with x within a stack or regenerator; it has to be constant, independent of x . Even in a situation like we saw at the bottom of the last viewgraph, where the work flow— $A p_1 u_1$ time averaged—may be changing quite dramatically with x , the energy flow must be independent of x . It's tricky; we have to think very carefully about what we mean by energy here.

Another example, an historic example, was the 1975 experiment of Merkli and Thomann. [28] They had an acoustic resonator with variable length—they could slide this thing in and out—and it had an oscillating piston over on this end to excite a resonance. This is a half wavelength shown here. Here's the pressure amplitude as a function of position and the velocity amplitude as a function of position in that standing wave. The dissipation per unit length in this system is highest in the center, where the viscous losses are the highest. (The work plot looks funny, because I adopted the convention that work flow to the right is positive, so this leftward work flow is negative.) See, the work is largest right at the piston and drops to zero at the other end. In between, the slope is steepest right in the center, because the viscous losses are a little more important than the thermal losses.

Here's what surprised those guys at that time (although it doesn't really surprise us much anymore). They had sensitive thermometers all along the tube. When they started this thing oscillating, the central thermometers cooled while the rest of them warmed. That seemed surprising at the time, to see cooling where the dissipation of acoustic power was highest. That won't seem so surprising as soon as we understand the important difference between energy flow and work flow.

(vg 3.4) So, what do we mean by energy flow? In gases (and liquids), the enthalpy is the right energy to consider. [?] This is because the energy flux density in fluids is velocity times the sum of kinetic energy density and enthalpy density. The enthalpy itself is the sum of internal energy plus pressure over density. Look at this expression and imagine these three pieces. Kinetic energy density, internal energy density, and some kind of pressure energy density. The idea is this: as a lump of fluid passes through some cross section fixed in space, the energy that passes through that cross section has those three parts. The kinetic and internal energies are literally convected along with the fluid, but the other energy that you have to include is the work that the pressure does on the fluid ahead of it as it pushes that fluid of its way in order to pass through the cross section. Any fluid mechanics text will explain this more fully.

This is the second-order expression for time-averaged energy flux density, which is all that we usually have to think about in thermoacoustics. The kinetic-energy term drops out, because it's third order—velocity times velocity squared. So it turns out that only the mean density times the time-averaged product of the oscillating enthalpy density and oscillating velocity matters. [29]

There are at least a couple of good ways to look at this simple expression. Professor Thomann of the previous viewgraph once wrote that the energy equation is like a kaleidoscope: Every time you shake it and look at it, it looks completely different. We're going to shake it twice here, to get two different ways of looking at it. First, if you express enthalpy in terms of entropy and pressure, and substitute that in here for that little oscillating enthalpy, then the energy flux density turns into these two pieces. One of them looks like heat, because it has to do with entropy transport, and the other is clearly the old, familiar work term. If instead you express enthalpy not in terms of entropy and pressure but in terms of temperature and pressure, it breaks up in a different way, with the pressure term equal to zero for an ideal gas. In that case, the enthalpy flux density is simply this one term, with oscillating temperature. Both of these are correct. Sometimes it's useful to think in these terms, and sometimes to think in these terms.

For instance, in a perfect regenerator in a Stirling system, there is no oscillating temperature because the heat capacity of the regenerator and the good thermal contact in the pores there maintain locally isothermal conditions all of the time. So the energy flux density for the perfect regenerator is zero. In a real regenerator, it should still be small. That would imply that it might be useful to think of the two terms in this other expression for energy flux density as being essentially the same magnitude but opposite sign so they cancel.

Rott worked out an energy equation[29] for thermoacoustics by taking this $T_1 u_1$ expression and substituting in the expressions that we've already looked at with all the f_k 's and f_v 's, and working out the math. (vg 3.5) If you do that, you get this expression for the energy flux in thermoacoustics. It's a little too complicated to be very intuitively clear. There are a few pieces that we can easily recognize. The "1" here in the first line is just regular, familiar

acoustic power flow. The third line is easy; it's just ordinary dumb conduction of heat in the x direction. The rest of all this stuff holds the details that make thermoacoustics so rich. Clearly, the fact that there's a term with temperature gradient and a velocity squared looks a lot like something that came from one of the two terms in the continuity equation.

It's important to realize that you can regard this equation as a differential equation for $T_m(x)$ in the stack or regenerator, because of the control-volume picture, with thermally insulated side walls, which tells you that the energy flux H_2 dot has to be constant, independent of x , throughout the stack or regenerator. If you know pressure and velocity from the previous lecture, then you can regard this as a differential equation for $T_m(x)$.

(vg 3.6) Energy flux in pulse tube refrigerators provides an interesting example, because they are a little more complicated, so this example will sharpen our intuition about all of this a little more. Remember that we had two ways of looking at the energy equation: one where we broke it up into a heat flux term and a work flux term, and the another with only one term, with oscillating temperature and velocity.

If you consider a control volume like this in the basic pulse tube refrigerator (which nobody builds anymore because they don't work as well as the orifice variety), you find that the cooling power has to be the difference between the energy flux that goes out through the pulse tube and the energy flux that comes in through the regenerator. This is the expression of the first law of thermodynamics for the control volume. At the regenerator end of the control volume, if it's a nearly perfect regenerator, there is almost no oscillating temperature, so the energy flux coming in there is essentially zero. Coming out here in the basic pulse tube refrigerator, you don't have much acoustic power—only whatever is dissipated at the walls of the pulse tube. So the larger part of the energy flux in the pulse tube is the entropy flux. But that's small too, because most of the gas in the pulse tube is far from the walls, so it's adiabatic, so s_1 is zero almost everywhere. (I've cheated here, giving the wrong reason for the right conclusion. Homework assignment: Discover the cheat, and work out the right reason.) So the energy flux in the basic pulse tube is mostly the entropy transport along the wall that we have all the time in standing-wave thermoacoustics. Essentially the cooling power at the cold heat exchanger, which is what you want to achieve, is equal to this standing-wave heat pumping along the wall of the pulse tube in a basic pulse tube refrigerator.

We'll be able to understand why the orifice pulse tube refrigerator works so much better than the old basic pulse tube refrigerator, because it has more cooling power. In the orifice pulse tube refrigerator, you pull substantial acoustic power out of the system into this dissipative structure—the orifice and the compliance volume—making the $p_1 u_1$ term dominate the $s_1 u_1$ term in the pulse tube itself. You still have essentially zero energy flux coming in through the regenerator. So now your cooling power is nearly equal to the acoustic power that you extract with the orifice.

(vg 3.7) Two viewgraphs ago, I described Rott's energy equation as a differential equation for $T_m(x)$. Let's combine that idea with the momentum and continuity equations from lecture two. This results in a standard calculation method in thermoacoustics. The momentum equation can be written as dp_1/dx equals a bunch of stuff. (Remember? It looked like Ohm's law? with f_v and u_1 on the right.) The continuity equation can be written as $\langle u_1 \rangle$ equals a bunch of stuff. (Remember? two terms, proportional to p_1 and $\langle u_1 \rangle$?) Rott's energy equation can also be rewritten, solving for dT_m/dx equals a bunch of stuff. We regard these three equations—which are actually five real equations because the first two are complex

equations—as describing the evolution of p_1 , $\langle u_1 \rangle$, and T_m with x . These equations can be numerically integrated along x to tell you everything you want to know in a thermoacoustic system, within the acoustic approximation. I'll show you an elaborate example of that in the next lecture.

One software package to do this is documented and freely available from us: it's called DeltaE, [30] and in the bibliography at the end of the lecture transcript you'll find a reference giving Bill Ward's email address if you want a copy. I think the physics of our DeltaE, and Tominaga's Thermoacoustica, and David Gedeon's Sage, [31] and Ray Radebaugh's REGEN3 for regenerators [32] is essentially the same, so it doesn't really matter which you use. If you want to, you can write your own code. You'll learn a lot if you write your own code, but you'll get to the answer faster and get on with your projects faster if you use someone else's code.

Many approximate methods exist, and some are easy enough to do with spreadsheets and pocket calculators. [33, 20] Typical approximation that can be made, relative to the starting point shown on this viewgraph, are to use boundary-layer expressions for the f 's, to let $T_m(x)$ be simply linear in x , or to pretend that the gas properties are independent of temperature so that you evaluate everything on the right-hand side of these equations only at the x -averaged temperature, the temperature in the center stack, instead of integrating everything as a function of x .

(vg 3.8) The crudest approximate methods, summarized on this viewgraph, are crude enough to do in your head while you're driving to work. For an order-of-magnitude estimate, assume that a regenerator is perfect, so the energy flux through it is zero, and that a pulse tube experiences no convection or conduction, so the energy flux through it is the pV work through it. That's really all you need to know to estimate the cooling power of a pulse tube refrigerator at the very crudest level. In thermoacoustic stacks, if you made all of approximations that I discussed a minute ago, at the bottom of the last viewgraph—boundary layer approximation, gas properties evaluated at the average temperature, exactly standing-wave phasing, linear temperature distribution as a function of x , you can get these simplified expressions of Rott's energy equation and the work equation. [20] These include a temperature-gradient factor, with the actual temperature gradient in the hardware divided by a critical temperature gradient, which takes us back to the computer animation that showed a thermoacoustic refrigerator with a shallow temperature gradient, a thermoacoustically uninteresting situation where nothing was happening with an intermediate temperature gradient, and a thermoacoustic engine with a steeper temperature gradient. That intermediate temperature gradient was this critical temperature gradient, where the adiabatic temperature change that the oscillating gas experiences as it moves in the standing wave just match the temperature gradient in the stack. That gradient separates the refrigerator and engine regimes; this factor changes sign.

Notice that all these powers are second order in acoustic amplitude—they're all products of two p_1 's or u_1 's. That simple fact can help you estimate roughly how powers scale with amplitude. You can also use these simple expressions as starting points to estimate how powers scale with frequency, or mean pressure, or sound speed. To do that, you have to think carefully about what you want to imagine holding constant. Often, you can gain insight by holding the Mach number constant, so you would want to rewrite these in terms of p_1/p_m and u_1/a . For example, suppose your thermoacoustic engine is driving a load with a fixed acoustic impedance. Then the power consumed by the load is proportional to p_1^2 . But the

power consumed by the load has to equal the power produced by the engine, which is given by this approximation, also proportional to p_1^2 . Set them equal; the p_1^2 's cancel, and you find that the engine's temperature gradient should be approximately independent of amplitude. Then look at the H equation: for fixed temperature gradient, H and hence the required heater power should be approximately proportional to p_1^2 .

4. A detailed example

(vg 4.1) Let's briefly summarize what we've done in the last two lectures. The starting physics in thermoacoustics is: the equation of the state of the gas, other properties of gas (viscosity and thermal conductivity), the momentum equation, the continuity equation, and the heat transfer equation (which is essentially interchangeable with the energy equation). We understand very well how to apply the acoustic approximation to these, integrating perpendicular to x analytically to get the functions f_κ and f_ν which describe the viscous and thermal details in stacks and regenerators. We didn't spend any time looking at the details of deriving f 's in various geometries. We rewrote the acoustic versions of the momentum, continuity, and energy equations in forms with dp_1/dx , $d\langle u_1 \rangle/dx$, and dT_m/dx on the left-hand side, and we said that you can numerically integrate those equations with respect to x to get detailed information about the x dependences of those variables.

Now I'll show you a very detailed example of doing that, and of the hardware that the calculations resulted in, and how well the hardware actually worked.

The example is the project that we did with Tektronix several years ago, to make a thermoacoustic engine driving an orifice pulse tube refrigerator. [13] Here's that photo again, showing one of the final hardware setups. Basically the half-wavelength resonator extends along here, there are two standing-wave thermoacoustic engines, and there's a side branch close to one of the engines, that goes to a two-stage orifice pulse tube refrigerator.

(vg 4.2) When we started out on this project, Tektronix came to us and told us what they wanted for the thermoacoustic driver; they told Ray Radebaugh at NIST what they wanted for the orifice pulse tube refrigerator. Many of these specs were a compromise between the needs of the standing-wave engine and the refrigerator; many of the specs were driven by Tektronix's goal of a particular useful device to work toward. They wanted us to give them a thermoacoustic engine using 30 bar helium gas, running at 400 cycles per second, which would deliver a kilowatt of acoustic power to the pulse tube refrigerator at a pressure amplitude of 3 atmospheres. They said that the pulse tube refrigerator was mostly a real acoustic impedance with some compliance, with its impedance having a phase angle of about thirty degrees.

And of course like all real-world sponsors that we all have to satisfy, these sponsors wanted this system to be as small as possible and as efficient as possible. At first, they didn't understand that you have to compromise between those two goals. We went back and talked with them a little more, and we all realized that what they really wanted was a system no larger than a certain size; and given that size constraint, they wanted it as efficient as possible.

We made some of the initial design choices by seat of the pants. For instance, we chose this double-engine configuration with the pulse tube refrigerator on a side branch. We also fixed the stack diameter, to keep the size down to their goal. I don't know if these were the best choices, but you can spend huge amounts of time, whole master's theses worth of effort, justifying each such decision—or you can trust luck and instinct, and get on with other

challenges.

Anyway, once having made these decisions, we then trusted our computer calculations [30] to determine the details of the design, to pick things like stack length and spacing and so on, to make the system as efficient as possible. Now I'll show you some of the results of these calculations, to reinforce what has come before, and to show you that you really can calculate some of these quantities and think about them. (vg 4.3) The top of the figure here shows the outline of the apparatus, with the side branch to the refrigerator here, one of the thermoacoustic engines here, the mid line of the standing wave here, and the other thermoacoustic engine over on the right. (vg 4.4) Our design calculation gave us this for the magnitude of the pressure as a function of position in the hardware; and this is the magnitude of velocity as a function of position. Basically this is what you know a half-wavelength standing wave looks like: cosine kx in pressure and sine kx in velocity, with little discontinuities here at the area discontinuities. At those locations, you actually have continuity of volumetric velocity, so the velocity itself changes abruptly. That's the standing wave picture.

Now we'll add more detail to that. In standing wave engines, the wave is very nearly a standing wave, so the magnitude of the pressure is very nearly equal to the real part of the pressure (with the arbitrary choice for the zero of phase I made here) and the magnitude of the velocity is very nearly equal to the imaginary part of the velocity. What I've added in green here are the out-of-phase components of pressure and velocity. These parts give power; pure standing waves carry no power. I've had to multiply both of these out-of-phase components by 20 to get them to show up clearly on the same vertical scale as the others. Here, in the region of the stack, is where things are changing most dramatically with x , as the stack produces acoustic power. This big step here is where acoustic power goes out into the side branch where the pulse tube refrigerator is.

(vg 4.5) You can also learn what the temperature as a function of position in the apparatus should be. It drops off through the stack from the high temperature of about 760 Kelvin down to room temperature. Notice that the temperature profile is not even close to linear: It drops off much more steeply at the hot end than at the cold end. That's why I like to use a "complete" calculation like this, instead of one of the simplified approximations such as I mentioned at the end of the last lecture, when working with a engine spanning a large temperature difference. Next, if you look at the total energy flux and the pV work flux in the system as a function of position, you get curves like this. As we've seen in the previous lecture, the total energy has to be constant through the stack because of the sidewall insulation. On the other hand, work can build up through the stack, as a function of x . A lot of work is consumed by the side branch. The negative work on the right is that which is provided from the other engine; it's flowing from left to right so it's negative here.

This is just an example of the kind of information that you can get from these straightforward acoustic-approximation calculations with DeltaE or whatever other Rott-based computer program you choose to use.

(vg 4.6) At the design operating point, this is the overall calculated performance of the system. In order to deliver one thousand watts to the pulse tube refrigerator, the calculations say we have to provide 5,380 watts—total—to the heaters in the two thermoacoustic engines. This includes all the acoustic effects we just discussed, with a couple of other easy effects added in: heat leak from the hot ends to the room through the insulation, and heat leak from the hot ends to the cold heat exchangers along the pressure vessel walls around the stacks.

If you take the ratio of these two numbers, you get an efficiency of 18%. The Carnot efficiency would be 60% at the temperatures we chose, so this system is predicted to have an efficiency that is 30% of the Carnot efficiency. That's quite respectable, compared to the workhorses of our society: The internal-combustion engine and the electric-power generating plants with steam turbines and so on are around 50% of Carnot. But I'm greedy. Why is this so inefficient? Why isn't it a greater fraction of Carnot? I've listed the sources of inefficiency in rough order of importance. The most important source is the thermal relaxation loss in the stack, which is what makes standing-wave thermoacoustic engines work. We are stuck with that. Viscous loss in the stack, and viscous and thermal losses in the heat exchangers, are also important. In this system, heat leak to the room was very important, because we were so constrained in size that we couldn't put a lot of insulation around the hot ends. Resonator-wall acoustic losses, and a couple other conduction losses and imperfections in the heat exchanger accounted for the rest. We'll come back to this list in more quantitative detail later.

(vg 4.7) First, let's look at some of the hardware we built. This is a scale drawing of one of the thermoacoustic engines. The only thing not to scale in this drawing is the stack spacing. Had I drawn the stack spacing to scale these lines would have had to be three or four times more dense, and it would be too cramped to see. We'll see pictures of all these components later. This is the boundary of the pressure vessel containing the acoustic standing wave. This is the hot heat exchanger, stack, and cold heat exchanger. This space was filled with ceramic-fiber thermal insulation to insulate the red hot hot end from room temperature out here. In order to keep this pressure-vessel wall as thin as possible, to minimize heat leak down this stack-housing wall, we maintained this insulation space at the same average pressure as the gas in here. This was 30 atmosphere argon, and this was 30 atmosphere helium, so this wall only had to support the oscillating pressure, not the entire average pressure. This outermost wall is very thick, because it supported the average pressure—but it's easy out there, because it's at room temperature, where metals are strong.

Next we'll look at the various components, starting with the hottest and moving toward the coldest. Here's that inner pressure vessel, for the stack and the heat exchangers and the standing wave. This shows us wrapping the ceramic-fiber insulation around the outside of it. Then the outer pressure vessel went on: The outer vessel looked like this, and when it was welded to the other part of the case, the assembly ended up looking like this. The hot heat exchanger, the source of heat in the system, was this Nichrome ribbon wound on a ceramic frame. Here you see the ceramic frame and you can just barely just see the Nichrome ribbon in there which goes back and forth in a zig zag pattern like this, turning its corners around little Nichrome rods. This view from the back side shows the electrical connections a little bit more clearly. Nichrome pins, which plugged into nickel sockets in the resonator hot end.

We use two different stacks at different times. We started out with commercially available Hastelloy honeycomb. Hastelloy is an alloy with more nickel than stainless steel has, so it retains strength at a little higher temperature than stainless steel. Some of the data I'll show you later were taken with these three Hastelloy honeycombs stacked in series in there. We also built parallel-plate stacks. This is about 6 cm diam, to give you a sense of scale. This stack was spirally wound, with a Hastelloy strip wound around a central mandrel, with copper as a sacrificial spacer material to maintain the gaps, which were about a quarter of a millimeter. When the spiral was wound up nice and tight, we cut little slots across the ends—here and

back on the rear end—, dropped in little radial ribs, and welded each intersection of a rib with a turn of the Hastelloy spiral. Then we could etch out the sacrificial copper, to leave a free-standing spiral of Hastelloy supported only by these ribs on the ends. It's a very beautiful thing to look at: the geometry looks nice and uniform. If you pick a place where you can see clearly in this photograph, you'll see how uniform the spacing is.

The cold heat exchanger was basically a finned tube heat exchanger, with water going through the CuNi tubes, and copper fins making thermal contact to the helium. Here you see four little water tubes, and here are four other little water tubes. The copper fins go across. The assembly was simply soft soldered together by hand. Water also flowed in these passages around the perimeter, so that a place on a fin like this was cooled by conduction through copper to water both in this direction to the tubes and also in this direction to the perimetral case. Finally, this is just a photograph of that same heat exchanger stuck into the central portion of resonator, with the water inlet line here and water outlet line over there. That's what all the parts look like.

(vg 4.8) We built that at Los Alamos, and tested it before shipping it to Tektronix. We tested it without a pulse tube refrigerator load, with the two thermoacoustic engines driving only the resonator. At the time of this particular set of measurements, we only had one hot heat exchanger and one stack (and both water-cooled heat exchangers!) rather than stacks on both ends of the resonator. But we had both a honeycomb stack and a spiral stack that we could interchange. This shows you how well the hardware agrees with the calculations based on the acoustic approximation, with all the calculation methods we've been talking about in the previous lectures. To make the measurements, we set the voltage to the heaters at a chosen value, and waited an hour or two until the hot temperature and pressure amplitude settled down to a steady state. Look at the upper left graph first. For the spiral stack, the green line shows the calculations using the DeltaE computer program, and the green points are measurements in the real hardware, showing pressure amplitude as a function of heater power to the hot heat exchanger. The red symbols are the same for the honeycomb stack. Hot temperature as a function of heater power is shown down here. The right graphs are the same quantities, plotted in a dimensionless way that we like a lot at Los Alamos; we'll come back to the reason for this choice of dimensionless groups in lecture 6. We want to think in terms of pressure amplitude normalized by the mean pressure, and squared, as the independent variable for both graphs. The dependent variables are heater power, normalized by mean pressure, sound speed, and cross-sectional area of the stack; and hot temperature, normalized by cold temperature (300 Kelvin). So plotting it this way, we are asking how much heat, and how much hot temperature, are needed to get the thing to oscillate at a given pressure amplitude. Anyway, the data are the same here and here, just plotted in a different way.

In the limit of low amplitude, the agreement between calculations and measurements is good, much better than it is out at 10% amplitude. That is the sort of behavior you expect from an acoustic-approximation theory, which you only expect to work best in the limit of low amplitude.

The temperature crosses over here. I don't really know what that means, but our thermometers were external to the case of the thermoacoustic engine, so as not to threaten the reliability of the hardware we were shipping to Tektronix; internal thermometers would've been more accurate. I think that we could easily have had 100 degree temperature differences

between inside and outside. Instrumentation compromises like this are part of the reason why these data never got published.

Anyway, overall good agreement between prediction and fact. In our experience at Los Alamos, you can expect this level of agreement. The customer was satisfied.

Those data were unloaded: with no pulse tube refrigerator on the system. We shipped the hardware to Tektronix, and they started using it to drive pulse tube refrigerators. They were never very interested in checking exactly whether the thermoacoustic engine worked the way we said it would under load, so quantitative information under load is not great. (vg 4.9) This viewgraph summarizes the most careful data point taken with that thermoacoustic engine assembly when it was at Tektronix under load, with a pulse tube refrigerator drawing 490 watts of acoustic power from the resonator at 5% amplitude. Both stacks were spiral. The total measured heater power to the two heaters on the two thermoacoustic engines was within 10% or so of the calculated heater power under those circumstances. The hot temperature was also within about 10%. As far as we know, the system met our goals. Tektronix continued to use the hardware to drive various pulse tube refrigerator modifications, until their management canceled the program.

(vg 4.10) That was the hardware from the point of view of the thermoacoustic engine. I'll say only a few words about the pulse tube refrigerator, because we weren't involved in that directly at Los Alamos so I don't know as much about it. There was a one stage pulse tube refrigerator, and later a two stage refrigerator; each with several modifications. I'll just talk about the one stage 'fridge. It cooled to 147 Kelvin with no heat load under the circumstance that I described in the last viewgraph. Here are the parts of the pulse tube refrigerator: a first heat exchanger that we came to call the aftercooler just because you need a different name for some heat exchangers since there are so many heat exchangers in these systems. The stacked-screen regenerator was a few centimeters long and a few centimeters in diameter, made of stainless-steel screen, 325 wires per inch—a very fine screen. The cold heat exchanger didn't have much heat transfer metal in it, but it had a thermometer for measuring the cold temperature, and a heater. The pulse tube itself was less than a centimeter in diameter and about 20 centimeters long. These show the locations of various sensors that the Tektronix people put in to help diagnose what was happening.

This is a photograph of that one-stage pulse tube refrigerator, much later, when we started using it at Los Alamos for other work. The regenerator is between the level of this surface and the level of this surface, and there's some superinsulation stuffed into the space here. You can't really see the side wall of the regenerator pressure vessel, because it's hidden by these massive flanges and that superinsulation. The cold heat exchanger is in here somewhere, and here's the pulse tube going back up to a water-cooled heat exchanger. The orifice was a needle valve right here with the shaft going out of the vacuum can so you could turn the valve from outside the vacuum can when the system was operating. And here's the reservoir volume. What you see down here is the O ring seal to the vacuum can that went around this whole thing to provide thermal insulation.

5. Efficiency

In the last segment we looked in great detail at one example, the heat-driven refrigerator assembled at Tektronix. Remember, at one point we talked about the efficiency of that

system: the design calculations predicted that the thermoacoustic engine would have 30% of the Carnot efficiency, and I listed sources of loss, such as the inherent thermal relaxation losses in the stack, viscous losses in the stack, etc. Now I want to spend some time talking about more sophisticated ways to look at efficiency; then we'll return to the Tektronix example at the end and make a more quantitative list of the source of loss in that system.

(vg 5.1) Percent of Carnot serves as useful measure of efficiency for simple systems, such as the simple heat engine or the simple refrigerator. [1] This goes back to viewgraph one, lecture one. The efficiency was what you want, divided by what have to spend to get it, and it was bounded above by the Carnot ratio of temperatures, which we derived from the first and second laws of the thermodynamics. The same kind of arguments applied to the refrigerator; it's efficiency is called coefficient of performance: what you want—cooling power—divided by what you have to spend to get it—work—bounded by this ratio of temperatures. So, in both cases, you can write down a "second-law" efficiency, that's the percent of Carnot efficiency: the ratio of the actual efficiency to the maximum efficiency allowed by the laws of thermodynamics.

That's fine for such simple systems. What are you going to do for a more complex system? For instance, suppose you have a machine that takes air at atmospheric temperature and pressure, and separates and cools it to produce liquid oxygen, liquid nitrogen, liquid argon, solid dry ice, and whatever. You drive the machine with some work, and it dumps waste heat to ambient temperature. The question is: How should we talk about the efficiency of a complicated system like this? How should we understand the limit on its efficiency that is imposed by the first and second laws of thermodynamics?

One way to do so is based on the following principles. [19, 34] First of all, you measure all losses in the system in term of lost work, not lost heat or any other measure of loss. Second, you specify one temperature T_0 as a very special temperature, the environment temperature where you dump waste heat, such as I indicated here. I'm going to elaborate on these two principles for several viewgraphs.

(vg 5.2) Suppose we have a simple a heat engine as shown here. We've seen this before; the only difference now is that I've labeled the cold temperature with the subscript 0 to specify it as the special temperature where we dump waste heat. We can write the work as the hot heat times the Carnot efficiency minus something I'll call W_{lost} . The first term is the work that would be produced if everything were ideal so this operated as a Carnot engine. If it's not ideal, the actual work is less, and we'll call the difference the lost work. This example shows what I meant by (1) identifying one temperature as the special dump of uninteresting waste heat, and (2) measuring losses in terms of lost work.

For the simple refrigerator, you write it like this. The work that you put in is the work you would have had to put in if the refrigerator were ideal, with Carnot COP, plus the extra work that you have to put in to account for the fact that it's not an ideally efficient system.

A third example, a very simple example of lost work: If you have friction in machinery at ambient temperature, whatever work is dissipated by friction in the machinery is clearly lost work.

Now here's a more subtle example, that really points out how tricky and how powerful this way of looking at efficiency can be. Suppose you have a refrigerator, and within the refrigerator there's some frictional dissipation at a temperature less than ambient temperature. For example, a piston in the cold part of the refrigerator rubbing as it moves. How much lost work

must be assigned to that frictional dissipation? Here's the idea. Suppose that this friction is the only loss in the system. Then you can imagine the system like this: an ideal refrigerator with Carnot efficiency, driven by a drive mechanism with friction, with the drive mechanism thermally anchored to the cold temperature. You drive the system with work W , you dissipate some work into heat by friction, and that heat goes into the low-temperature reservoir. The remaining work drives the ideal refrigerator, which now must pump the frictional dissipation plus the external cooling load Q_C out of the low-temperature reservoir. In the figure, I've written things around the frictional element using energy conversation.

Over to the left, I've written down the heat-pump equation for this ideal refrigerator, using $W - W_{fric}$ as the work into the ideal refrigerator, and $Q_C + W_{fric}$ as the heat being lifted from the low-temperature reservoir. That "inner" refrigerator is ideal, so its heat and work are related by the Carnot COP ratio of temperatures. Now we do a little algebra, to shift W_{fric} entirely to the right side of the equation. See how this cancels this? So when you get the expression to this state, you can identify the work that you're putting into the whole system as the sum of two pieces: the ideal work that would be needed to pump the desired heat load Q_C with an ideal Carnot device, plus another term. Principle two says that we want to identify this second term as the lost work. So the lost work is the frictional dissipation times this ratio of temperatures.

This temperature ratio makes sense, because intuitively you would expect that dissipating power at low temperature should be worse than if dissipating at room temperature. If you're dissipating work at low temperature, you have to pump that heat up to room temperature, and that costs you even more work. Further, if you were dissipating the friction at a high temperature, higher than T_0 , such as in an engine instead of a refrigerator, this temperature ratio would be smaller than one. That makes sense too, because if you dissipate some work into heat through friction at high temperature, you haven't lost all of the work, because you can use that high-temperature heat to run an engine and recover some work, in principle. That's the idea here.

(vg 5.3) This idea of lost work is so important because you can express the lost work as the product of the special temperature T_0 (the dump temperature) times the sum of all the entropy generating processes in the system. Let's go back to some of our examples in the previous viewgraph to make sure that this theorem makes sense. If we have friction at a temperature T , then entropy generation must be the heat generated—which is the frictional dissipation—divided by T . We divide by T because that's the definition of entropy in terms of heat and temperature. So if you multiply by T_0 according to the theorem, you get this expression: and that's what we had on the last viewgraph. There's the temperature ratio, T_0/T .

Here's another example: imperfect heat transfer. Suppose we have two thermal reservoirs, at $T + \delta T$ and T . Think of one of these as the real thermal reservoir at one end of an engine, and the other is the "inside temperature" of the working gas in your engine, so the little difference δT represents the imperfect nature of your heat exchanger. How much entropy is generated in this process? The entropy generated is the difference between the entropy gained here and the entropy lost here. The entropy gained here is Q/T and the entropy lost here is $Q/(T + \delta T)$. Do the mathematics, and you get this. The theorem says multiply that by T_0 to get lost work: Lost work is the heat that you're transferring times this ratio of temperatures.

This method and theorem are important because entropy is an extensive quantity. You

can do this sum in any of a number of ways. You can sum it component by component throughout a piece of experimental hardware: heat exchangers, stack, resonator and so on. Or you can add it up by location within a component: the hot end of the stack, the middle of the stack, the cold end of the stack. You can sum it up by process: viscous processes, thermal processes, frictional processes and so on. You can break it down even finer, asking how much of the thermal entropy generation is due to processes in the x direction, and how much due to processes perpendicular to x . Or you can sum it up by time inside a thermodynamic cycle, if you want to know whether the entropy generation due to thermal relaxation in a regenerator is coming mostly from the compression and expansion phases of the cycle, or the gas-motion phases of the cycle, or what.

If you choose to sum it by both process and location, you can think of it like this. The rate of entropy generation—the overdot means per unit time, like we've been doing with heats and works—is an integral over the whole volume of your apparatus of the sum of per-unit-volume effects, each coming from a different process. This term, with viscosity, comes from the viscous dissipation: this term, from thermal conduction, this from electrical dissipation, friction, chemistry, more processes if you can think of them. Having this in mind will let you think about where and how the inefficiencies arise in a system.

(vg 5.4) I want to show you one more textbook example, because it reinforces some things from this lecture and from lecture 3. Suppose we have this system, processing a fluid stream from one temperature and pressure to another temperature and pressure. Maybe you're liquefying nitrogen or methane or something like that. You drive the system with work, and you dump your waste heat to T_0 . The first law of thermodynamics says that what goes out, the heat going out plus the enthalpy going out, equals what goes in, the sum of the work going in and the enthalpy going in. The second law of thermal dynamics says that the entropy going out of the system, which is Q_0/T_0 plus the entropy carried out by mass flow, has to equal the sum of the entropy carried in by mass flow plus any entropy generation in the system. So inside the system you can be imagining viscous and frictional and other dissipative processes generating entropy like we saw in the last viewgraph. Okay, combine these two equations by eliminating the uninteresting quantity, the waste heat. This is what we get. The work that you need to accomplish this process is this. What does it tell us? First lesson, a review. Enthalpy is the correct energy for the first law of thermodynamics; that's what we used up here. Second lesson: it looks like $h - T_0s$ is another important, special kind of energy. That is the energy whose difference gives the minimum work, or the ideal work, required to process the fluid from one state to another. What do I mean by minimum or ideal work? If there's no entropy-generating process in the system, so this term is 0, then the system is ideal, Carnot-like. Then the amount of work that it takes to process the fluid is given by the difference between this energy going in and going out. Any real-world process must use at least that much work, according to the first and second laws of thermodynamics. Look closely: this is not the Gibbs energy. It sort of looks like Gibbs energy, but if it were the Gibbs energy we would have T in here instead of T_0 , because the Gibbs energy out is $h_{out} - T_{out}s_{out}$, not $h_{out} - T_0s_{out}$. In America, I think this energy is usually called availability, and in Europe I think exergy is a more common term. I don't know what it's called in Japan. This illustrates again that you have to think carefully about energy in engines and refrigerators, because there are a lot of quantities that have the units of energy, and which quantity you want depends on what you're trying to answer!

(vg 5.5) Here's a table that I want to give you for future reference, so if you choose to think about entropy generation in terms of different dissipative processes, or dissipative processes per unit volume, you can use these expressions. The left column is for a macroscopic point of view, and the right column is for a differential, per-unit-volume point of view. I won't talk about everything in this table, just a few examples. Suppose you have dissipation of electrical power by resistance. Macroscopically, the dissipated work is $I^2 R$, so the lost work is T_0/T times that. Microscopically, the entropy generation is given by this, per unit volume. If you integrate that over the whole volume of a resistor, and multiply by T_0 to convert from entropy generation to lost work, you get back to the macroscopic expression. Same idea for friction and viscous flow, although the microscopic viscous expression is really complicated, including bulk viscosity. If you only want to think about viscous flow down a pipe, and not worry about microscopic details like whether it's turbulent or laminar, you can compute the lost work macroscopically with this expression. If instead you want to think about it in a differential sense, perhaps to understand whether the fluid closest to the wall in the channels of a stack is more or less dissipative than the fluid in mid channel, then use the microscopic expression.

You can bet that imperfect heat transfer is an important loss mechanism in heat engines and refrigerators! Think for a moment about this $k(\nabla T)^2/T^2$ term. If $k = 0$ —perfect insulation—there is no loss. It turns out that in the other extreme, k infinite, there is also no loss, because in that case ∇T is zero. The inbetween case is lossy. So all real gases and materials are lossy! And recall that in standing-wave systems, we rely on being in the inbetween case to get any useful thermodynamic action at all, with heat being transferred across a thermal penetration depth throughout the stack, every cycle of the wave. That's why we sometimes call standing-wave systems "intrinsically irreversible". [35]

Here are a couple of examples that we haven't seen yet. If you have two fluid streams of different temperatures that mix, you generate entropy. If you have a Joule-Thompson expansion, you generate entropy. The list goes on and on; it's scary realizing how many different ways the entropy of the universe is increasing all around us!

A homework exercise: Show that, within the acoustic approximation, in stacks and regenerators, you can compute the entropy generation using these expressions, without bothering break it down into thermal and viscous components. Interesting exercise.

(vg 5.6) Let's go back to the Tektronix example to finish this lecture. In that system, at the one well-measured operating point we looked at, we were putting 3,000 watts of electric power into the system. And we were putting 3,000 watts of waste heat out into our water stream in all the room-temperature heat exchangers. That's really all that heat-driven refrigerator was doing. The cold end was sitting at 147 Kelvin, but it wasn't removing any heat from an external load at that temperature. In principle, you could make a cold thing sit at 147 Kelvin without an external load simply by insulating it perfectly, without providing any work at all. So all 3,000 watts is lost work in this system!

(vg 5.7) So we can think about how we want to break that full 3000 watts into different entropy-generating processes, to learn quantitatively what are the worst components, or the worst processes, or whatever. Think about it: Do we want to break it up by location, by process, or both? I made an arbitrary choice of how to break it up, and here's the list that I generated. Mostly I chose to break it down by location, but some of it is broken down further by process. The total of all the numbers in this column is 3,000 watts. Basically 2,500 watts

is lost in the thermoacoustic engine plus the resonator, and 500 watts is lost in the orifice pulse tube refrigerator.

The second-biggest loss is I^2R in the NiChrome heaters in the two thermoacoustic engines. Now what does that mean? What does it mean, that 730 Watts is lost to I^2R , when 3000 watts is converted from electricity to heat in the heaters? We're turning 3000 watts of electrical power into 3000 watts of heat at a temperature that's not infinite. So the formula says that the lost work in that process is I^2R , times T_0/T_H . The first and second laws allow you to do better, in principle. If you turn that 3,000 watts into heat at infinite temperature, you haven't lost anything, because you can in principle run a Carnot engine off that heat and turn it all back into useful work. The fact that we're at a finite temperature instead of an infinite temperature means that some of the 3,000 watts can't be turned back into useful work, even with a Carnot engine. That's how much is lost by our choice of operating temperature.

(vg 5.8) Here is another way of looking at it. In order to deliver 3000 watts of heat to the engine at 1230 Kelvin, you don't have to spend 3000 watts of electric power. You could, in principle, spend only 2270 watts of electric power, driving a heat pump with Carnot's COP, which would deliver 3000 watts of heat. The difference, $3000 - 2270$, is how much power is lost.

(repeat vg 5.7) You have to think through everyone of these that carefully to really appreciate this lost-work concept. This is an honest way to account quantitatively for the different losses in the system, to decide what components are responsible for losses, and maybe to decide where it would be most effective to spend your research budget or your construction budget to improve the overall system efficiency. You could use this type of analysis to decide what to try to improve first. You could argue with the sponsor about sacrificing some of the size constraint to get more efficiency, or you could do more research to go after the unknown losses, or you could do some engineering to improve the heat exchanger temperature defects, or whatever. An honest way to account for losses in the system. I've just started thinking in these terms this year, and I'm finding it fun, and useful so far. It keeps you focussed on the possibilities for improvement, rather than on your successes of the past! I think in the case of the combustion-driven natural-gas liquefaction system, it's going to be extremely useful, because that's such a complicated system, especially when you include multiple staging in the refrigeration and cryogenic separation of impurities and stuff like that.

In the Tektronix system, the stacks accounted for the biggest loss: 950 watts of lost work. I didn't break it down further in the viewgraph, but I remember that most of that loss is the thermal relaxation, not the viscous losses. That loss is large in this system because of the size constraint that Tektronix put on us. If we had the freedom to make the system bigger, we could have operated closer to the critical temperature gradient and therefore had a more efficient thermoacoustic engine.

(vg 5.9) This viewgraph summarizes the types of processes that are responsible for irreversibilities in thermoacoustics. First, you're stuck with some thermal relaxation losses within the thermal penetration depths within the stack in any standing-wave system, because the phase lags between gas motion and thermal relaxation are what make the standing wave systems work. So you're stuck with those, and they are inherently dissipative. In traveling-wave systems, those losses are in principle smaller and in practice smaller, but still there. You're also stuck with viscous penetration depths and viscous losses. Whenever you have a velocity gradient, whenever you have flow past stationary solid boundaries, you're going to have

viscous dissipation, and it's an important one in both standing- and traveling-wave systems. In standing-wave systems, it's easy to visualize why you're stuck with viscous loss, because for anything except liquid metals [36] the Prandtl number is of order one so the viscous penetration depth is of the same order as the thermal penetration depth. Ordinary conduction of heat in the x direction from hot to cold is a conceptually very simple loss that you're generally stuck with. All of those processes occur not only in stack and regenerators but also on resonator walls and in heat exchangers. You can have very serious thermal bottlenecks in heat exchangers, either in the solid metal parts of the heat exchangers or in the flowing fluids to carry heat to or from the system. Finally, there are some things that we don't yet understand.

A third of Carnot is typical of today's designs of thermoacoustic systems. Experimental hardware is lagging a little bit behind that, but catching up fairly rapidly.

6. Beyond the acoustic approximation

(vg 6.1) Last time, we looked at the elaborate Tektronix example in detail, and at the end we had some plots comparing measurements and calculations for the engine, showing that as you apply more and more heat, you get higher pressure amplitude and higher hot temperature. But there was some difference between the calculations, based on the acoustic approximation, and measurements. In this lecture, we'll focus on the differences, to try to understand where and how to begin to study those. It's not clear to anyone what the causes of these differences between measurement and experiment are. I don't think they're just experimental blunders, although they might be. I think that they're actually inadequacies in our understanding of the behavior of real thermoacoustic systems.

(vg 6.2) The first place to look for physics beyond the acoustic approximation is back at the starting equations. I think we believe that we believe that all the truth about the thermodynamic behavior and fluid-mechanic behavior of gases in thermoacoustics is summarized in the equations in the top half of this viewgraph. You need to know the gas properties: specific heat, and the equation of state. I guess I forgot to list the viscosity and thermal conductivity; they should be here too. Next, the continuity and Navier-Stokes equations. Finally, the heat-transfer equation. (The energy-flux equation can be derived from these last three, so we don't have to list it separately. [?]) These equations, plus boundary conditions, should include everything we need to know about the gas behavior, right? But now think about how much we've thrown away when we've made the acoustic approximation by substituting in expressions like this. For example, in the Navier-Stokes or momentum equation, we threw away this term because it ends up being second order, the product of two velocities, each a quantity that is first order. So that disappears. You might know that that essentially means we're doing a low-Reynolds-number approximation here, so we're throwing away all possibility of including effects of turbulence. That's one possible candidate piece of physics that's left out of our calculations so far. How important might that be? How many more effects have we excluded in the acoustic approximation?

(vg 6.3) One general approach to keeping track of physics beyond the acoustic approximation is the use of similitude. This is the science of picking the right dimensionless groups of variables, such as Reynolds number or Mach number, to describe a problem. [?, ?] One approach to similitude is illustrated here for thermoacoustics. [37]

Imagine any acoustic engine or refrigerator. I've drawn an engine here, with a stack here, some heater power coming in to a hot heat exchanger here, and acoustic power going out here. Imagine that you build it, fill it with a gas to an average pressure you choose, anchor this location at a temperature you choose, and put in a chosen heat here. The engine then springs to life, oscillating at some frequency, with some amplitude, and with some temperature distribution in the stack. All the dimensions of that apparatus are included in a long list called x_j . There are a few dimensions that we'll pick out for special notice and give special names. One is the hydraulic radius h_{ref} in the stack. (Oops—sorry. I've switched notation on you. This should be called $r_{h,ref}$ to be consistent with notation from earlier lectures.) Another is the cross sectional area A of the stack. The subscript ref means at this reference location here. A third specially-named dimension is the total length L of the system. One approach to similitude is to write down all the variables that you think matter in the problem, and then form dimensionless groups of variables and see what you get. (There are formal procedures [?] to ensure that you don't miss anything important, but we don't have time to discuss those details here.) Here's the list of all the variables that might matter in a thermoacoustics problem. All the dimensions of the hardware. The properties of the gas: the ratio of specific heats, the sound speed at whatever temperature exists at this reference location, the thermal conductivity and viscosity at whatever temperature exists at the reference location, and these two constants which give the temperature dependence of the viscosity and thermal conductivity. You have to think about it for a while to assure yourself that this is a long enough list to give *all* relevant gas properties: For example, where's the specific heat? the density? You can calculate them from the items in this list, plus the pressure and temperature listed down here. Solid properties that you might think really matter in thermoacoustics include the thermal conductivity, density, and specific heat, but usually the conductivity and specific heat can be regarded as essentially infinite so I've omitted them here. Some miscellaneous variables here that don't logically fit in the lines above: the heater power that you put into this engine; the temperature at the reference location, the average pressure in the gas

Those are the variables we'll think of right now as independent variables, under our independent control. Next are all the dependent variables. These are the variables whose values nature chooses. Basically these can be thought of as experimental results. The frequency that this engine chooses to operate at, the temperature everywhere (as a function of time too), the velocity everywhere, the pressure everywhere, and for example the acoustic power flowing down the duct. The power is an example of an "extra" dependent variable, because it is directly calculated from other variables.

There's more than one way to form these variables into dimensionless groups, but this is a way that we think makes a lot of sense. The dependent variables are on the left. Frequency is normalized by the length of the resonator and the sound speed of the gas at the reference location. Temperatures are normalized by the reference temperature. Pressures by the mean pressure. Velocities by the sound speed at the reference location. Powers—both acoustic and heater power, and any other power you're interested in—normalized by the product of mean pressure, sound speed, and cross sectional area. Over here on the right, here's another normalized power; and lot's of ratios of lengths. Notice this length ratio: δ_κ/h_{ref} . The Prandtl number σ is a ratio of some of the gas properties.

It might look like this is just a bunch of weird bookkeeping; what good is this going to do for us? Let me show you some examples that I hope will help convince you that this is a lot

more interesting than bookkeeping.

(vg 6.4) We had an experimental engine running for a few years, a simple electrically heated thermoacoustic engine driving a resonator, with no other acoustic load on it. One of the things that we looked at was the harmonic content of the oscillating pressure at the end point here. At this location, we looked at the $2f$ component of the wave, which is plotted here normalized by p_m . [37] It's displayed as a function of the square of the fundamental component of the wave divided by mean pressure. That's a simple example, almost trivial, of displaying data using dimensionless groups of variables, as chosen at the bottom of the previous viewgraph. Look: it turned out that it didn't matter what gas you put in this resonator—helium, neon, argon, nitrogen or a helium-argon mixture—when plotted in this normalized fashion, the data all fell on this same curve. Furthermore: The different symbols lined up in columns here represent different mean pressures for each gas. At different mean pressures, the data still fall on the same curve. If we hadn't plotted it this way, with mean pressure divided out, we wouldn't have noticed this universal behavior.

Does this universal behavior tell us anything about the physics responsible for the presence of this $2f$ component? Remember, we haven't see anything yet in these lectures, in the acoustic approximation, that would give us any idea that there should be any harmonic content in the wave, or how big it should be. Well it looks like to me like the fact that it doesn't matter if you have a monatomic gas or a diatomic gas means that the ratio of specific heats γ doesn't matter in generating this $2f$ component. And the fact that it doesn't matter whether you have a helium-argon mixture or pure helium or argon separately means that the Prandtl number doesn't matter. So if you're looking for explanations for the presence of the $2f$ component you should probably start with physics that doesn't include these. I'm not sure, but I think that means that boundary-layer processes aren't important, because they would depend on Prandtl number. It might also mean that compressibility is not important, because γ is the ratio between isothermal and adiabatic compressibilities, and γ doesn't seem to matter. The fact that it doesn't matter when you switch mean pressure corroborates the suspicion that boundary-layer processes aren't important, because the thermal and viscous penetration depths change with pressure. We've eliminated a lot of the physics. I don't have the answers, but I notice that one thing we haven't eliminated is the $u du/dx$ term in the momentum equation, so that might be a good place to start. I hope this gives you the idea that thinking in terms of dimensionless groups and similarity can be very productive.

(vg 6.5) Here are some more measurements with that same engine. This shows how hot the engine had to operate to produce a given oscillatory pressure, and, under those circumstances, how much heater power was being consumed. The three different sets of symbols here are for three different gases: helium, neon, and argon. Each with a different mean pressure, chosen so as to make the ratio of thermal penetration depth to hydraulic radius in the stack equal to this particular number 0.89. This ratio is one of our similitude normalized variables from two viewgraphs ago. When you choose the mean pressures that way and plot things in this normalized way, you find that all the experimental data lie along the same curve. And that curve is quite a bit different from where the acoustic approximation calculations would predict, at least for hot temperature. (vg 6.6) If you then switch to a different set of mean pressures for the three gases, each different by about a factor of two from before, so that the ratio of thermal penetration depth to hydraulic radius is different by a factor of about the square root of two, you get a different set of measured points and a different set of calculated

lines. The points still cluster together, but under those conditions it's the heater power that's far off from calculation while the hot temperature is closer to calculation.

What can you learn about the difference between acoustic-approximation calculations and reality from these data? The fact that plotting things this normalized way causes the data sets to cluster, independent of gas type, means that the physics that included in our variable list is complete. For instance, the excess heater power here can't be due to blackbody radiation carrying heat from hot to cold. Heat carried by blackbody radiation would scale with area, but not with gas pressure or sound speed! This normalization factor changes by almost a factor of 10 from helium to argon (with neon inbetween), so if all this difference between experiment and calculation was due to blackbody radiation, the experimental data would fan out here, with the difference covering a factor of 10. Similarly, flapping, buzzing, and rattling of the stack plates, which might not be perfectly rigid, probably wouldn't scale according to pressure times sound speed, so that's not to blame either. It has to be something within the hydrodynamics and thermodynamics of the gas itself, something included in our similitude analysis, that's responsible for this difference.

But it can't be $2f$ content in the wave, because the previous viewgraph showed that harmonic content was independent of gas and of $\delta\kappa/h_{ref}$.

Here's a cute thing that we found with that engine. [37] I don't know what to make of this, but I'll put it up to show you that something really odd can also obey the similitude rules. Under certain circumstances, that engine would not operate stably. We would put a steady heat into the electric resistance heaters, but the engine would shut off, turn on, shut off, turn on; cycling every 10 or 20 minutes. This map shows the boundary between the unstable region to the left and the stable region to the right. The map is plotted with dimensionless axes: thermal penetration depth over hydraulic radius horizontally, and pressure amplitude over mean pressure vertically. When plotted that way, changing gases makes no difference to the boundary of the instability region. But the frequency of shutting off, turning on did *not* scale with sound speed when we switched gases! So that suggests that the instability mechanism arises from variables included in the list that generated our similitude analysis (mostly gas properties), but the details of the instability, once it has arisen, involve something not included in the list, such as the heat capacity of the solid material in the hot heat exchanger.

(vg 6.8) Let's think about some possible candidates for physics that might be important to thermoacoustics, beyond the acoustic approximation. I'd like to break that into two categories. The first category is physics that's included in the equations that describe ideal-gas thermodynamics and hydrodynamics. These would be effects that are included in the similitude analysis: harmonic content in the sound wave, turbulence in the gas, shock-wave formation in the gas, acoustic streaming in the gas, joining conditions (you'll see what I mean by that a little later), entrance effects (meaning the fact that as gas enters the end of a stack or heat exchanger there must be some entry region where the velocity and temperature profiles as functions of y and z won't be quite the same as deep in the heart of the stack. And then there are some things that would not be included in the gas dynamics, like blackbody radiation, solid parts vibrating, thermal bottlenecks (not in the gas, but in the heat exchanger metal or the cooling water stream), things like that. Looking at your data with dimensionless variables, and switching gases, and that kind of thing can clearly differentiate between these two categories of effects.

Next I'd like to discuss some of these effects individually, in part to prompt some of you

to consider what kind of research you want to do for the next decade or two.

(vg 6.9) This is to continue our discussion on harmonic content a little. This is the same engine I was talking about before, in which we measured the $2f$ component of the wave. [38] We put in heat with electric resistance heaters to make the engine oscillate, and measured the acoustic pressure amplitude here at the end, and the hot temperature. Originally the resonator was a simple right circular cylinder. The $2f$ component, as a function of the square of the fundamental component, went like this, which are a subset of the data I showed about four viewgraphs ago. The corresponding heater and hot-temperature data are these and these. With the resonator a simple cylinder, the higher resonance frequencies of the resonator are integral multiple of fundamental, because you're just fitting half wavelengths between here and here in the resonator. So, if there are any important quadratic non-linear terms in the equations, like $u du/dx$, that are squaring the fundamental solutions to the equations and generating secondary oscillations at double the fundamental frequency, those processes hit resonances and you might expect that they would generate a lot of $2f$ pressure amplitude. Next, we put in some little tapered inserts in the resonator to ensure that the higher resonances were no longer integral multiples of the fundamental. With those inserts in place, the harmonic content of the wave dropped dramatically. What effect did that have on the other quantities we could measure? Here are heater power and hot temperature (minus cold temperature) as a function of pressure amplitude squared. In every case, reducing the harmonic did what I wanted it to do, it gave us a high pressure amplitude with less hot temperature and less heater power.

Clearly these harmonic effects can have an important effect on thermoacoustics. We should strive for a quantitative understanding. This work was very incomplete: I have no idea whether bigger inserts, shifting the second resonance further from $2f$, would've reduced T_H and heater power even more. I don't know which of the many possible quadratic nonlinearities in the starting equations are responsible for generating the $2f$ pressure in the first place.

(vg 6.10) This viewgraph may seem to contradict some of the previous viewgraph. That's an indication of the primitive level of our understanding of harmonics. This argument implies that harmonics should have only a small effect on power. Suppose you have a moving piston, sweeping out a volume as a function of time written as a Fourier series like this. The sum of a fundamental and higher harmonics. Suppose the pressure that the piston is pushing against as it moves can be written in the same way. If you write down the mechanical power that the piston does on the gas, which is the time average of pressure times the volumetric velocity, and you work through what the time average of the product of these sums works out to be, you end up with an expression like this in which there are no cross terms. It's not a double sum. There aren't any p_3 's times V_5 's or anything like that, only $p_3 V_3$ and $p_5 V_5$ and so on, because of the way integrals of products of trig functions are zero. That would imply that if a piston moves sinusoidally, so there are no higher components to the piston velocity, then the pressure harmonics have no effect on power. If instead the velocity harmonics are small (but not zero), then the higher terms in power are small squared, so again they shouldn't have very much effect. For enthalpy flow, through a regenerator or a stack or anything else, you similarly find an expression like this with no cross terms, and again you would think that the harmonics don't have much effect. But the measurements that I showed you in the last viewgraph showed quite a large effect of harmonics on power carried through a stack.

(vg 6.11) Let's leave harmonics behind, and consider turbulence. Those of you who have

had an engineering fluid-mechanics course recognize this graph. [?] It shows the friction factor for steady flow in a pipe. Vertical axis is the ratio of pressure drop in the pipe to the square of the velocity through the pipe times some other stuff, and horizontal axis is the Reynolds number of the flow through the pipe. Way over here at low Reynolds number you have laminar flow, which is the only flow that physicists study as undergraduates. And out here in the real world at high Reynolds number you have turbulence, with pressure drop far far higher than you would predict by extrapolating the laminar-flow friction factor. The actual pressure drop depends strongly on the roughness of the wall of the pipe. .

It would be great if we had data like this for oscillatory flow in pipes in the regime of interest in thermoacoustics. But we don't. In fact, what we use is the acoustic approximation, which is probably as naive as just extra extrapolating this laminar flow curve out as far as you want and ignoring all this other reality.

(vg 6.12) Here is some of what is known about turbulence in oscillatory flow. [39, 40, 41] This is much more complicated than turbulence in steady flow. In oscillatory flow, there are many regimes of turbulent and non-turbulent flow. The map of those regimes is shown here, with the Reynolds number on the horizontal axis and the pipe diameter over viscous penetration depth on the vertical axis. (This whole map is for smooth-walled pipes so in reality there's some third axis, pipe wall roughness, that we don't have a clue about.) For smooth-walled pipes there are seven regimes of flow. I've marked a few special ones. At low Reynolds number and small pipe diameter is laminar flow. If you put a hot wire anemometer in the gas and look at velocity as function of time at different radii in the pipe, you get what looks like laminar flow here. (These data look like rectified sine waves because hot-wire anemometers measure the absolute value of the velocity.) Let's look at two other regimes. Weakly turbulent, up here at modest Reynolds number but larger pipe diameter, still behaves basically like the laminar flow, but with little wiggles and jiggles in the velocity. Conditionally turbulent, out here, looks smooth while velocity is rising, but just at the peak of velocity changes abruptly to wildly turbulent behavior. These regimes have been mapped out qualitatively like this, but nobody really knows what the pressure drop is in a duct that experiences flow like this, much less what the pressure drop is in flow like this with rough walls. Maybe I'm exaggerating a little—there is a little information scattered throughout the literature. But I don't think there's a textbook summary or tabulation comparable to that single figure that exists for steady flow.

(vg 6.13) Does this matter for thermoacoustics? Yes. Regenerators operate down in here, so they're OK.. But the standing-wave thermoacoustic stacks are in the ugliest region you can imagine, right at the transition zone between several of these regimes. Pulse tubes in pulse tube refrigerators have large diameter, and not very large Reynolds numbers, so they're up here in the weakly turbulent regime. Acoustic resonators in standing-waves systems are often somewhere up in here. So thermoacoustics situations lie all over the map, and all we really understand about power dissipation for thermoacoustic processes, in detail, is the laminar regime down in here.

(vg 6.14) Let's continue down our list of gas phenomena beyond the acoustic approximation. We've looked at harmonic content and turbulence. Streaming is another gas phenomenon. [42] We have found one instance where streaming is very important in these things: in the pulse tube of a pulse tube refrigerator. In the pulse tube itself it turns out that you can get big convective loops driven by boundary layer streaming at the wall. One of the physical

origins of that boundary layer driven streaming is illustrated here. We've magnified the picture so that the viscous penetration depth looks big; it would look like a tiny fraction of the diameter over here. Because of oscillating pressure, and the axial temperature gradient and oscillating velocity, the gas is likely to have one temperature when it's moving upward and a different temperature when it's moving downward. The viscosity of the gas is temperature dependent. Hence the gas is more viscously locked to the wall during one direction of its motion than during the other direction of its motion. That leads to a net drift in gas motion after each cycle of the wave. In other words, in the simple acoustic approximation we think of the gas as just oscillating back and forth, but it's actually oscillating back and forth and ratcheting its way up just a little as a function of time. That ratcheting or slipping along the wall drags the nearby central gas along, driving a big convective loop in the pulse tube, carrying heat from hot to cold. We've made some measurements[43] that convince us that this really is what happens in pulse tube refrigerators, and that by picking the correct taper angle in the pulse tube you can suppress that streaming. I bet that there are other situations in thermoacoustic where streaming has an important effect, but I don't know what they are.

(vg 6.15) Here's another thing that must be in the gas dynamics that we don't understand: joining conditions. Remember in these acoustic-approximation calculations that we numerically integrate the momentum equation in this form, the continuity equation in this form, and the energy equation in this form. We discussed this earlier in terms of stacks, but the same is true in ducts, heat exchangers, or any other 1-d acoustic element. What conditions do we use to join numerical solutions like this to each other at junctions between segments, such as the junction between the end of a stack and the beginning of the adjacent the heat exchanger, or the junction between the end of a heat exchanger and the beginning of an open duct? In the acoustic approximation, we simply match pressure amplitude (complex), volumetric velocity amplitude (complex), and mean temperature. But that can't be right beyond the acoustic approximation. For instance, it's easy to see that in a standing wave, at the junction between a heat exchanger and an open duct, you can't have continuity of mean temperature, because as the gas moves out into open space it gets adiabatically heated, then it moves back and gets adiabatically cooled. So, on average, the gas out here has to be hotter than the gas in here. So mean temperature doesn't match at the junction. As a second example, suppose you have two abutting heat exchangers, each with small pore size so the gas temperature is the same as metal temperature in each heat exchanger. Suppose the two have different temperatures. As gas flows from one to the next, you can't have continuity of U_1 ; you must have continuity of $\rho_m U_1$ so that mass doesn't accumulate at the junction. Rott proposed a general alternative to continuity of U_1 ; [5] but it assumes you know how large the discontinuity in T_m is, and we don't always know that ahead of time. So this is another category of physical phenomenon that's beyond our present understanding of thermoacoustics.

(repeat vg 6.8) So we've talked about these. We've skipped shock waves because I don't know of any thermoacoustic situations so far that have had shock waves. Entrance effects means, for example, that deep within the heart of a parallel-plate stack, the velocity u_1 is given by a hyperbolic cosine of y . But that can't be true right at the entrance to the stack, where the gas remembers that it just left behind a different geometry with different boundary conditions. Another situation we don't understand well.

(vg 6.16) Let me remind you again that one way to categorize these things and to begin sorting them out is to think in terms of dimensionless groups and similitude. There's a second

important use of similitude that I'd like to mention next. To introduce this, we need to realize that for two different thermoacoustic systems to be similar, in the strictest sense of the word, they need to have all ratios of lengths identical. (They also need to have the ratio of specific heats, and the Prandtl number, and some things like that identical.) In order to get all possible ratios of lengths identical, you certainly need all geometrical length ratios to be identical, so the two systems must truly be geometrical scale models of each other. But you also need any ratios involving gas-property lengths to be identical. And the only gas-property lengths that fall out of our ideal-gas similitude analysis are the wavelength and the two penetration depths.

This tells you how to make a scale model of a thermoacoustic system. For instance, suppose you are interested in a full size thermoacoustic engine using helium at 450 psi = 3.1 MPa pressure. It turns out if you build a half size model and fill it with argon at 355 psi, the gas-property lengths are also cut in half. The wavelength of sound in the argon system is automatically half the original helium wavelength, because I'm thinking of a resonant system so it picks its frequency automatically to make the wave size and shape fit the geometry. The thermal and viscous penetration depths will also be half what they are here, because of the special choice we made of mean pressure in the argon. So all the length ratios here are the same as they were here. And they're both monatomic gases, so the ratio of specific heats is also the same. Luckily, it turns out that the Prandtl numbers are also nearly the same, and the exponents that give the temperature dependences of viscosity and thermal conductivity are also nearly the same! Under those circumstances the similarity arguments say that all the temperatures should be the same in the half size model as they were in the full size model. All of the pressure ratios should be the same. And all powers end up being divided by a factor very close to 16 in the half-size model—that's the ratio of $p_m a_{ref} A_{ref}$ for the two systems.

This might be a very useful thing to know and use if you're interested in a thermoacoustic system that's big like this and won't fit in your laboratory and requires a lot of heater power to run, and you'd like to study it conveniently in your laboratory. The resonator diameter in this photo is half of what it was in the previous photo. The stack length, the stack pore size, the heat exchanger lengths... every dimension in this system is half what it was in the previous photo. We run this one with 355 psi argon instead of 450 psi helium.

By the way, there are some things that don't similitude scale, like the requirement for thermal insulation here. That insulation doesn't look like a half scale model of the other system's insulation, and it's not, because, of course, the thermal conductivity of the insulation doesn't change properly when you change the gas in the resonator. In fact the insulation requirements here are harder than they are in the full size system because the powers are all down by a factor of about 16 here, and the hot temperature is the same, so you have to insulate more carefully here to lose the same percentage of your heat to heat leak as in the full-size system.

We put this half scale model to use to test things such as modifications to the resonator that our sponsor is interested in. These tests are a lot cheaper and faster to test in the laboratory, with small hardware than they would be to test outdoors with large hardware. And if we believe that the only physics that matters in bending a resonator like this is included in the hydrodynamics and thermodynamics of an ideal gas, then we have to believe that our half-scale results will be directly applicable to the full-scale hardware.

7. Practical details

(vg 7.1) The previous lectures have heavily emphasized theory, mathematics, and intuition. This lecture is going to cover the practical arts of construction, measurement, and practical calculation in thermoacoustics.

(vg 7.2) Let's start with stack and regenerator construction. These are some techniques people have used to build stacks and regenerators so far. I'll go through the list showing you examples of each one. To date, about half the stacks in the experimental literature have been parallel-plate or rectangular with a large aspect ratio; according to calculations of f_κ , this is the second-most efficient stack geometry. Most people have built parallel stacks by making spirals. We'll start with that.

(vg 7.3) One way of making a spiral is to wrap up a sheet around a central axle with tiny rods acting as spacers every so often, so that the channels for the thermoacoustics are actually like little rectangular cross section channels between the spacer rods. [15] These look like they have about a 10 to 1 aspect ratio. Looks like the spacers here had too much glue; or maybe the camera wasn't lined up just right. In the case of the plastic stacks that are popular at the Naval Postgraduate School, the stack is kapton and the spacers are nylon fish line segments glued onto the sheet before the sheet is rolled up. That's one way of making parallel plate stacks.

I showed you another way back in lecture 4, the engine stacks for the Tektronix system. [13] There, the spacers were only at the very ends of the stack, so most of the length of the stack was about as close to ideal parallel-plate geometry as we're likely to ever achieve.

(vg 7.4) Here's another way of making parallel plate stacks. We used stainless steel wires as spacers, and stainless steel sheets as plates, all stacked up. The wires had been copper plated so that when we heated all of this up in a hydrogen atmosphere to the melting point of copper, the copper acted as a braze alloy to bond the whole thing together. Unfortunately, when it came out of the furnace things were pretty wrinkled and warped and disgusting. This stack actually didn't perform anywhere near as well as the calculation for parallel plate geometry would suggest, so I conclude that any stack that looks this irregular is not good enough. We never studied whether you could avoid this mess by using annealed stainless or heating and cooling slowly or something; we just abandoned this method. But we've had good luck building up glued plastic stacks this way.

(vg 7.5) You can buy metal honeycomb. If you look very closely at this photograph, [38] you'll see the details of how this honeycomb is made. Sheets are bent into a zig zag pattern like this, and then they are welded together wherever one layer's zigs touches the other layer's zags. You see how there is a double thickness of sheet metal on each of the vertically aligned segments, and only a single layer on the others? The double thickness is where the little welds are. This stuff is fairly cheap, and they'll make it for you very fast. But they only make it slightly smaller pore size than this. There are probably a dozen manufacturers in the US—just look in the Thomas Register under "metal honeycomb". We model this as circular pores believing that the gas can't tell the difference between hexagons and circles; that seems to do a decent job of modeling it. The performance of these honeycombs is not as good as parallel plate stack performance, based on f_κ calculations and on real hardware, but they are very easy and cheap.

(vg 7.6) One newcomers to the menagerie of stacks in thermoacoustics is the pin array.

[25, 44] The idea is to somehow hold a triangular array of thin wires parallel to the acoustic oscillation. Calculations show that this is the most efficient stack yet, and the Naval Postgraduate school guys have made some encouraging preliminary measurements on these. It is very challenging to make such a thing, holding the pins aligned without blocking the gas motion with the holders. Remember that the transverse length scale is of the order of δ_κ , which is typically much smaller than a millimeter.

(vg 7.7) The Naval Postgraduate school people have also been trying some porous stacks. [45] One that they've been trying is wire mesh stacks that look like this, with rather open screen—much more open than the screens used in regenerators in Stirling machines. They've also been working with a porous carbon material, reticulated vitreous carbon, that looks a lot like the side surface of this pile of screens looks. I think that the results there are also still preliminary.

In the Stirling machine business, the most commonly used regenerator is a stacked screen like this, typically with very fine wires and very close spacing. People will also sometimes use a felt, which looks to me like very fine, crushed stainless steel wool. [46]

(vg 7.8) There's some new work on microscopically engineered regenerators for Stirling cryocoolers. [47] The idea is to make a regenerator by rolling up a photochemically milled sheet. To make the sheet, they use a different mask on the back side and the front side of the sheet, so that you end up with a geometry like this. These are tiny dimensions, small fractions of a millimeter, probably less than a tenth of a millimeter typically, so I think that the real stuff doesn't look nearly as beautiful as this drawing. The authors claim that by doing serious numerical calculations you can design the spacing between these ribs, the thickness of the ribs, and so on, to get thermal contact comparable to that in stacked screen regenerators, but with less viscous pressure drop, so it should work much better than the stack screen regenerators do. They published some experimental results which were very encouraging, but I think no one has yet built another one and reproduced the results. So this is something that still needs work, but it seems very promising.

By the way, if you calculate the performance of a simple parallel plate regenerator in a Stirling cryocooler, you get fantastically better performance than in the stack screen regenerators in common use. But I believe no one has yet succeeded in building a parallel-plate regenerator that works anywhere near as well as the stack screen regenerators. That's another puzzle that needs more research.

(vg 7.9) Moving on from stacks and regenerators to heat exchangers: As in other heat engines and refrigerators, in these thermoacoustic systems you generally have to move a lot of heat across a solid boundary, from the thermoacoustic working gas to circulating water or something like that. When you look in the standard heat exchanger textbooks, [48] you'll find figures like this that are a good summary of the things that people have been doing in thermoacoustics. You can have finned tubes, you can have other finned geometries, or you can just have tube arrays. You have the choice of putting the working gas in the space between the fins or putting the working gas inside the tubes. I think that all of these things have been tried in thermoacoustics, and it's not clear to anyone what's best or cheapest.

There's an important difference between the heat exchangers that you find in the standard handbooks and the heat exchangers that we use in thermoacoustics. In thermoacoustics, we have oscillatory flow in the gas, but steady flow in the water. All of the handbook stuff is for steady flow in both the gas and the water. It's not obvious exactly how best to apply

the steady correlations that you find in the handbooks to the case that we're interested in. I think there's a lot of room for research in this area, good fundamental engineering research to map out the details of heat transfer in oscillatory flow.

Here are some pictures of a couple of heat exchangers. If you have a really small refrigerator or engine, you don't really need to put water channels across the diameter of the apparatus. You can get by with copper fins spanning the whole diameter of the resonator. But as soon as you get a little bit larger in diameter, you find that the thermal conductivity of copper isn't enough to carry heat from the center of the heat exchanger all the way out to the perimeter. So you have to put water lines through like shown here.

(vg 7.10) Besides the choices of overall heat-exchanger geometry, such as tube and shell, or finned tube or whatever, you also have the choice of the regime of flow for the helium or the other gas in a thermoacoustic system. In the early days of standing-wave systems, we built heat exchangers with the length of the heat exchanger in the direction of the acoustic oscillation comparable to the oscillation amplitude of the gas, and with the gaps in the heat exchanger comparable to the thermal penetration depth in the gas, and with laminar flow, because these are the situations that are relatively easy to understand for physicists. More recently some people [?] have proposed and studied heat exchangers that are much shorter in the direction of oscillation, and much tighter in the perpendicular direction; they claim that you get a better ratio of heat transfer to viscous dissipation that way. Meanwhile at Los Alamos we've been moving in another direction, keeping the length of the heat exchanger comparable to the displacement amplitude, but going to much larger gaps in the gas, and deliberately using turbulence to enhance the heat transfer. All three of these work, but we don't know which works best. I think the "lost-work" ideas we discussed in lecture 5 are good ones to consider to decide what you mean by "the best" in a given circumstance.

(vg 7.11) Another important part of standing wave thermoacoustic systems is resonators. Here are some of the issues that come up in resonator design and construction. We'll talk in turn about harmonic suppression by shaping resonators, dissipation reduction by shaping resonators, concerns about size and weight, and finally I'm going to stick in a viewgraph about pressure vessel safety.

(vg 7.12) Harmonic suppression first. If you have a uniform diameter cylinder of length L , with a half wavelength frequency of sound speed over twice the length, the higher resonances in that resonator are all integral multiples of that fundamental frequency. Under these circumstances, non-linear terms in the hydrodynamic equations, which generate $2f$ and higher pressure oscillations and velocity oscillations from the fundamental pressure and velocity oscillations, these processes will find resonances and you end up with oscillations with a lot of harmonic content. We discussed this in lecture 6. It doesn't take much reshaping of the geometry of the resonator to disturb this integer-multiple situation. [38] For an example, I worked out the resonance frequencies for this resonator, in which the length of the end sections is one fourth of the length of the center section, and the area of the end sections is twice the area of the center section. If you look at these resonance frequencies closely, you'll see that they are no longer equally spaced. This is something that you probably want to do in thermoacoustic resonators, to ensure that the $2f$ oscillations that are generated by the fundamental don't find a resonance frequency to get magnified in amplitude. Quantitative details? I don't think we know exactly how the $2f$ oscillations are generated, or how far you have to shift the resonances to be safely away from $2f$.

(vg 7.13) A bonus of such a change in shape is that it can actually reduce the dissipation in the resonator. Tom Hofer figured this out. [15] Here are three resonator designs, all of which do the same thing for the stack. But this is the one with the lowest dissipation in the resonator. The first one is a half wavelength resonator from here to here. Clearly you can replace that with a quarter wavelength resonator plus an infinite volume, because the center of this resonator has a pressure node and velocity maximum, and that can match quite naturally to an infinite volume. It turns out that a large volume is just about as good as an infinite volume. So far, so good. Then Tom showed that if you reduce the area and length of this portion of the resonator, in just such a way as to maintain the resonance frequency the same, the total dissipation in the resonator decreases at first, until at some particular area-to-length ratio you find a minimum of dissipation in that segment.

(vg 7.14) This is the ratio of the diameters of the large and small parts of the resonator. This is some normalized dissipation in the resonator. It turns out that as you make the diameter of this portion smaller than this one the viscous losses and thermal losses both go down, because the surface area goes down. But at a diameter ratio of about 0.6 you arrive at a minimum in the dissipation. Beyond that, as the surface area continues to go down, the thermal losses continue going down, but the viscous losses go up because the velocity goes up, and the harmful effect of higher velocity overcomes the beneficial effect of lower surface area. It buys you the difference between this much dissipation and this much dissipation. In small thermoacoustic systems the resonator dissipation can be a large fraction total, and it's worth going after this improvement.

Often, another aspect of resonator design is trying to fold or bend or coil the resonator so that it doesn't take up so much room. This one would have been a ridiculously long resonator if it hadn't been coiled up into three turns here, so it would fit into this room. In the references you'll find a little work on the extra dissipation that you get from coiling a tube in a system such as this, but we really don't know enough about it quantitatively. [49]

(vg 7.15) One more issue about resonators: If you calculate how much stored energy there is in the compressed gas, in even a small thermoacoustic refrigerator or engine, it is deadly. Suppose you have a liter of helium, or other monatomic gas, at 20 atmospheres pressure. The amount of energy stored in that is the same as the energy stored when you lift a 100 kilogram weight 1 meter. So if you don't think these things can hurt you, I guess you are willing to fall on your head from a height of one meter. The stored energies are dangerous, even in small thermoacoustic systems. So be careful. What careful means in our laboratory is that we design our systems according to the US pressure vessel code, the ASME code. In the big systems we make sure to get certified welding. The ASME code tells you under what circumstances the welds have to be x-ray inspected. If you don't get your welds x-rayed you have to have extra safety factors. When all that is done we pressure test the system with nobody in the vicinity, according to the code rules, to a pressure higher than we will ever use in the lab. When we can get the system wet inside, we hydrostatically test it, to 50% higher pressure than we will ever use again. If we can't get it wet, we do a pneumatic test; I think the rules call for a 30% overpressure in that case. Please take this seriously.

(vg 7.16) Next: Transducers in thermoacoustics. There are many kinds of acoustic to electric transducers, [3] such as piezoelectric, magnetostrictive,... but so far the one that is in almost universal use in thermoacoustics is the electrodynamic transducer, so that's the only one that I'm going to discuss here. The physicist's prototype of this is the loudspeaker, with

a magnet structure making a radial magnetic field in an air gap, and a coil of wire in that air gap. The basic physics of that is, of course, that the force on a current-carrying wire is proportional to the magnetic field and the length of the wire. Also, the voltage induced by a moving wire at velocity v , is also proportional to magnetic field and the length of the wire. If you add a little bit more physics to that picture, you can write all this down in the acoustic approximation. You end up with equations giving the pressure difference across the moving speaker equal to the sum of the terms proportional to current and volumetric velocity, and giving the oscillating voltage induced in the coil also equal to the sum of two terms proportional to current and volumetric velocity. These are complex equations; we use them to keep track of both magnitude and phase of the oscillating variables, just like in the acoustics in the gas. The proportionality constants include the area of the diaphragm, the spring constant of the suspension, the moving mass, any kind of mechanical hysteresis or friction that damps the system, and the electrical resistance and inductance of the coil. You can use this either as a speaker to generate sound from electricity, or as a microphone to generate electricity from sound.

If you had an acoustics class you've seen this, but the engineering realities of doing this for thermoacoustic engines and refrigerators are much different from the engineering realities in ordinary audio acoustics. Even the vocabulary is different: in thermoacoustics we should be thinking about these things as linear motors and linear alternators, [50] which is what they are called in the Stirling world, not speakers and microphones. (vg 7.17) This table shows many of the differences between audio loudspeakers and the kind of transducers, which are based on the same physics, that we should use in thermoacoustics. Audio loudspeakers have to operate over a broad frequency range, but in thermoacoustics we can design for a single frequency and stick with it. Audio loudspeakers generally have low efficiency because of this requirement, whereas in thermoacoustics we're trying for the highest efficiency possible. Audio speakers usually have a very low moving mass, because of the broad frequency range requirement, and a weak spring, and a large damping. Because in thermoacoustics we can aim for a single frequency, and we want high efficiency, we tend to make linear motors and linear alternators mechanically resonant, with a large mass and a strong spring, resonating at the single frequency that we're interested in. In thermoacoustics we try to keep the mechanical damping as small as possible, so that the Q of the mechanical resonance is very high.

Hi-fi speakers usually have a large area and a short stroke, because of the nature of the flexure at the suspension. In trying to achieve high efficiency in thermoacoustics, we find that we want a smaller area and a larger stroke, which makes the engineering of the flexure much more difficult—if indeed a flexure is used at all. Finally, all the audio speakers that I know of are have moving coils. In contrast, in linear motors you can have any of the key components of the transducer moving.

(vg 7.18) For example, here's the loudspeaker diagram from two viewgraphs ago, with a moving coil. Below it, this is one of the configurations that's used in the Stirling machine community quite frequently, in which the permanent magnet and some of the iron core move, and the coil of wire and the rest of the iron core are stationary. It is a puzzle to figure out how the physics of this configuration ends up being equivalent of physics of our familiar moving-coil loudspeaker, but you can see some it. If the moving portion is up about that far (relative to where it's shown in the drawing), then the magnetic path is north to south around this way, so it encloses the coil here with flux in this sense. If the moving parts move down here,

then the flux path is this way, which is the opposite sense. Since the flux changes sign, the motion of this induces an emf in the coil. In terms of force: if you put a current in here, the current goes one way, and turns these surfaces into north poles and these into south poles, so it pulls the moving parts up. Switch current direction, and you switch field polarity, pulling it down. This type of linear alternator is used widely in the Stirling community. Avoiding a moving coil lets them avoid figuring out how to get the current-carrying leads to flex.

(vg 7.19) To give you an idea of how advanced all of this can be: this is one version [51] of the flux-switching configuration, in which the nice cylindrical symmetry of the previous picture is gone. Here, they have eight pairs of permanent magnets. These are rectangular blocks with magnetic field this way. The speckled part is the iron yoke. The moving parts here are the magnets and this central, circular frame for them, which is an extension of the piston. Notice that none of iron yoke moves, so this folded design helps keep moving mass to a minimum. These are coils of wire that go around this way around the iron yolk. I believe the physics here is the same as in this cylindrically-symmetrical picture, but the packaging is quite different.

I suspect that such designs are new to the standing-wave thermoacoustic community, because we have been trained looking at loudspeakers, not looking at these high-power, high-efficiency packaging options. They routinely get over 90% efficiency with these linear alternators and linear motors in the Stirling literature; but to morph a loudspeaker design to that level of efficiency would be a technical and costly challenge. Another way to compare the Stirling community's devices with audio loudspeaker-style devices is to compare Bl/R , the ratio of Bl product with the electrical resistance of the coil. You want high Bl , because that provides the transduction mechanism; you want R as low as possible, because I^2R is loss. Typically this ratio is 30 times higher in the Stirling hardware than in audio hardware.

(vg 7.20) Another feature of thermoacoustics is the fact that gas mixtures have a lower Prandtl number than pure gases. [52] I'm going to stick this in the lectures here, because I couldn't think of where else it fits. This plot shows Prandtl number as a function of helium concentration for mixtures of helium with other inert gases. Out here, on the right edge, is pure helium, with a Prandtl number of about two thirds. If you mix in just a little argon or xenon, maybe 10% or 20% argon or xenon, and 80% or 90% helium, you get a dramatic reduction of Prandtl number. This drops the viscous losses in thermoacoustics systems relative to the thermal effects that are beneficial. As far as I know, gas mixtures have not yet been used in pulse tube refrigerators, because in cryocoolers you're generally working below the liquefaction temperature of these heavy gases. But it might be worth trying a helium-neon mixture sometime, to get just a little reduction in Prandtl number.

That concludes the hardware portion of lecture 7. Now let's talk about measurement.

(vg 7.21) This is the way I see it. It's very easy to measure the average temperature as a function of position in your apparatus. And it's very easy to measure the average pressure, and the magnitude and phase of the oscillating pressure. We usually use thermocouples to measure temperatures, an ordinary Bourdon tube pressure gauge to measure mean pressure, and piezoresistive pressure sensors for measuring oscillating pressure. We get ours from Endevco. We use a lockin amplifier to get the magnitude and phase of the oscillations.

Of course it's very easy to measure heat if you're putting in heat electrically. If you're taking heat out with a water stream that's reasonably easy to measure; we'll get back to that in a second. It's also not too difficult to measure the acoustic power flow in some circumstances,

which we'll talk about in a second. But I think it's very difficult to make direct measurements of the acoustic velocity or of the oscillating temperature. The easiest tools we have available for measuring velocity, hot wire anemometers, are extremely tricky in the presence of pressure oscillations and temperature oscillations going on simultaneously. The measurement of the oscillating temperature is also very difficult at our typical high frequencies, and the length scales over which oscillating temperature changes are very tiny. So we tend to emphasize measurements of the easier variables, and try to infer the velocity from the easy measurements when we need information about it.

(vg 7.22) I'd like to show a few details of some of the things you can do to measuring these powers. If you're trying to measure how much heat power goes out into a water stream, such as the hot heat exchanger of a thermoacoustic refrigerator, the straightforward method that you might think of first is to measure the water temperature before and after the heat exchanger as the water flows through, measure the volumetric flow rate with a flow meter, look up the heat capacity of water, and multiply it all together to get the heater power. That works fine, if you have a flowmeter that's well calibrated.

There's a second technique that I learned from Ray Radebaugh. You don't have to know the volumetric flow rate in Ray's technique, and that can be the least accurate of the quantities in this equation. What Ray likes to do is to put in another sort of heat exchanger down stream of the real heat exchanger; the new one simply has an electric resistance heater in it. You make three temperature measurements here, and you adjust the current in the heater until this middle temperature is the average of these other two. In that circumstance, you know that the power that this new heat exchanger is putting in is equal to the power that's going in in this heat exchanger. And it's easy to measure the second heat exchanger's power very accurately, if you power the resistor dc. A further trick is to wire this up as a differential thermocouple mess, so you just have two wires coming out that give you zero volts when the center temperature is the average of the other two. Let the physics do the subtracting, instead of subtracting measured temperatures with a calculator.

In either of these methods, take care to ensure that the water is thermally well mixed where you measure its temperature. Sometimes if you stick a thermocouple into the water stream too close to the heat exchanger outlet, before the water turns through a couple of elbows or something, you won't measure a true, flow-averaged temperature.

(vg 7.23) Here's another important type of power measurement in thermoacoustics: in pulse tube refrigerators. It's easy to figure out the acoustic power flowing out of the hot end of the pulse tube and into the orifice-compliance network, if you measure the oscillating pressures here in the pulse tube and here in the compliance volume. That's easy because, if you know the volume of the compliance V , then knowing the oscillating pressure in it gives you the volumetric velocity into it. The orifice has negligible compliance of its own, so the volumetric velocity at the orifice entrance equals that into the compliance. And the product of volumetric velocity and the pressure at the hot end of the pulse tube is the power flowing out of the hot end of the pulse tube. So that's something that's very easy to determine experimentally.

Next, if you insulate all of this and draw a control volume around it, the only place energy can go in and out of the control volume are at the hot heat exchanger, where the cooling water can extract heat, and up the pulse tube, where the acoustic processes are carrying energy into the control volume. So the energy flux up the pulse tube has to equal the heat

removed at that hot heat exchanger. A perfect pulse tube would have no convection carrying heat from hot to cold, and no stupid conduction of heat from hot to cold down the metal walls. In a perfect pulse tube, this would really be simply a slug of gas, experiencing adiabatic pressure oscillations, transmitting work from here to here. Thus a perfect pulse tube would have the energy flowing up the pulse tube exactly equal to the work flow, the acoustic pV power flow up the pulse tube. So, if you determine the acoustic power by measuring the pressure amplitudes, and you measure the heat that's carried out in this heat exchanger, you have a measure of how perfect your pulse tube is. Or how much heat is being convected from hot to cold by streaming or by whatever other lossy process is there. That's an important, simple measurement to make in pulse tube refrigerators.

Now consider the big control volume shown here. It cuts through the regenerator and goes around the whole system. Across the boundary of that big control volume, the energy flows are the enthalpy flow $\dot{H}$ up the regenerator, the heat $\dot{Q}_C$ that the refrigeration load puts in at the cold heat exchanger, and the heat $\dot{Q}_{H2}$ that the water stream here carries away. Conservation of energy says that it has to add up like this. A perfect regenerator carries no energy through it, because the temperature oscillations are zero in a perfect regenerator and one of expressions we saw in lecture 3 for $\dot{H}$ in thermoacoustics was the product of oscillating temperature and oscillating velocity, with no other term. So the measurements of these two quantities tells you how perfect your regenerator is; it tells you something about regenerator losses.

People should make measurements like this more often. Very careful measurements of the heat that's rejected in this heat exchanger, because that tells you something about whether the performance of your pulse tube refrigerator can be improved by working on the regenerator or by working on the pulse tube. Generally people don't bother to make these measurements and interpret them in this way; they just present data, show how the overall system performed, without trying to understand the system at this level of detail, separating regenerator losses from pulse tube losses.

(vg 7.24) Another way to measure acoustic power: This figure comes from of Tom Hoffer's Ph.D. thesis. [15] You can measure acoustic power at a moving transducer without too much difficulty. If you put an accelerometer on the transducer, to measure its motion, and you put a pressure sensor nearby, you can determine the acoustic power that this transducer launches into this tube as a time-averaged product of the pressure and the volumetric velocity. You can get the volumetric velocity from the acceleration of the moving part.

(vg 7.25) Another way to measure to acoustic power: In a long duct, put two pressure sensors, A & B, along the duct and use them in the following way. [53] The acoustic power flowing down the duct is, of course, the area of the duct times the time average of the product of pressure and velocity. At this point, the midpoint of the two sensors, the oscillating pressure is roughly the average of the oscillating pressures at the two sensors. And the oscillating velocity at the midpoint can be roughly determined from the measurements, by using the acoustic approximation to the momentum equation: mass time acceleration is pressure gradient. So the difference in pressures between these two sensors tells us the velocity of the gas, and the sum tells us the pressure. Put that all together and you get the acoustic power flowing past that point given by this expression: with magnitude of the pressures at the two sensors and the sine of phase angle between those two pressures. The proper derivation for all this, which you'll find in the bibliography, also includes the viscous and thermal effects at the walls.

They're very important, there just isn't enough time to include them in this discussion. It makes the equation a little more complicated.

This method only works well if the sine of the phase angle between the two pressures is not too close to zero. In practice, in standing wave systems that means you want to have this location at a pressure node of the standing wave. That really limits the usefulness of this technique to situations where you have a source of acoustic power at one end of the resonator and an acoustic load at the other end.

So far in lecture 7, we've covered construction methods and measurement techniques. The third part of lecture 7 will cover one of the standard calculation techniques in the acoustic approximation, using the computer code DeltaE. [30]

You should recall from lecture 3 that the hard-core way to do it is to numerically integrate the y, z -averaged momentum, continuity, and energy equations in the x direction. Those equations were rewritten in terms of the x derivatives of oscillating pressure, oscillating volumetric velocity, and mean temperature. These are really five variables, because two of them are complex functions, so they each count double.

You just numerically integrate from one end of the apparatus to the other, integrating those through whatever geometry you encounter. If you come to a conical duct, you might use Webster's horn equation, [4] including losses. If you come to a straight duct, you might use the analytic results for damped propagation in a duct. In a stack, you use Rott's equations.

(vg 7.26) Here is an example I want to do in some detail for you. This refrigerator design is mostly due to a guy named Merrick Lockwood in Bangladesh, with me doing the calculations while he provided the good sense. He wants to make a heat-driven standing wave thermoacoustic refrigerator. It will burn kerosene, and its purpose is to keep vaccines cold in rural medical clinics where there's no electricity. Merrick got started building this thing, but has postponed it for a while, because a Stirling-engine project is keeping him very busy right now.

This is what his refrigerator looks like. Overall it is a half-wavelength resonator, with thermoacoustic parts symmetrical on both ends. There will be a hot heat exchanger here. This is going to have a kerosene burner out here, with flames licking through these fins. Inside the working gas will be air at pressure of 5 or 6 atmospheres, something you can get with a tire pump. The engine stack will be between here and here—it is not shown, because that would clutter up the drawing too much. There will be a room-temperature heat exchanger here, with cooling water freely convecting through this to cool it. Another stack, the refrigerator stack, here; and then the cold heat exchanger and the long duct of the resonator will be inside the refrigerated box. I think the resonator will be curved, to bring the burner on the far end around, adjacent to the one on this end.

(vg 7.27-28) To do a quantitative design calculation for this system, I used DeltaE. DeltaE accepts an input file that has mostly the geometry of the apparatus you have in mind, and some other stuff as well. It integrates through the parts of the apparatus in sequence. On the previous viewgraph it is from left to right. The rules this program follows in reading an input file like this are that the only thing that matters is what it encounters first on each line—either one of these segment names that tells you what kind of geometry to expect, or a number that tells you a value. The order of entry is crucial. All the rest of the stuff out here to the right is just comments for us humans to keep track of everything.

So we start here. 6 atmospheres mean pressure; 60 Hz; I guessed that we might start out

at 600 Kelvin on the hot end; 5×10^4 pascals pressure amplitude; a bunch of zeros; air; and solid type is ideal. That is a minor nuisance: You always have to have a solid type, even for segments that do not need them.

You remember 10 cm was the typical dimension on the drawing I just showed you. So after an endcap, we have a duct, 6-inch diameter. It has some perimeter. The hot heat exchanger comes next. I have guessed that we might have to put 400 W of heat in from the burner, at a temperature of about 675 Kelvin.

Next, the stack. Let's see, we have 5-mm gaps in this stack. That is going to be easy to build. It is only 3 inches long. We will try 0.002 inch sheet metal to make it.

Next is the heat exchanger in the middle, which takes the waste heat out to room temperature, whatever temperature is room temperature in Dhaka.

Another stack: the refrigerator stack. It is a little longer and with smaller gaps.

Where did all these numbers come from? Lots of iterating between Merrick and me. You are seeing what we arrived at near the end of a long process.

There is the cone that adapts from the end of the cold heat exchanger to the final duct; and that is the end of the DeltaE file, at the mirror-image symmetry plane at the center of this system.

This computer file has all the geometry in it. If we want to run DeltaE on that, there are a bunch of different things we can do, like plot variables as functions of other variables, or ask it what the thermal conductivity of the air in the middle of the stack is. I'll just show you a very small fraction of what DeltaE can do with this.

What we are going to do now is just run this file. DeltaE is going to integrate those equations of Rott from one end to the other using that sequence of geometry in the input file. That is all it is going to do.

(vg 7.29) Now we can display that it has done. Here is this sequence, again, of all the segments: heat exchanger, stack, heat exchanger, stack, etc. The data that are in columns here are the real and imaginary parts of the complex pressure amplitude, as functions of position all the way through the apparatus, and the real and imaginary parts of the volumetric velocity; the temperature; the energy flux; and the work flux i.e. the acoustic power. This is the sort of computer results I used to make the plots I showed you in lecture 4 of the standing wave and everything for the Tektronix 350 Hz engine.

This is the kind of thing that DeltaE will compute for you. It took that input file and it computed along and—oops, look, it did some things that I am not very happy with. I marked them in red here. When it integrated along, starting at the beginning, it gave us a temperature here in the middle, where it's supposed to match up with room temperature in Dhaka, of only 290 K. I think that is too cold for rooms in Bangladesh. So, as computed here, this thing is not going to be able to dump its waste heat, unless we change something. We will have to try to make this temperature come out warmer.

Another thing DeltaE did when it integrated: look at this: it has 21 watts of work going out the end of the file. That's the mirror-image plane in the center of the resonator. You cannot have 21 W of work going across a mirror-image plane—that's not symmetrical. What is going on here?

What is going on is that often, in solving differential equations, you have mixed-up boundary conditions. You do not know all the boundary conditions at the beginning in this problem. You know some of them at the other end. You know the impedance at the far end: It is a

mirror-image plane. The pressure amplitude is supposed to be zero at that mirror-image plane in a half-wavelength resonator. You also know the temperature that the room-temp heat exchanger will be anchored at.

So you have to use a shooting method, just like people use to integrate most other differential equations. You have to fiddle with something up in the beginning here to try to get these things later in the calculation to come out right.

(vg 7.30) DeltaE will do that fiddling for us, automatically. It lets you select some target boundary conditions near the end of the apparatus, and some fiddle variables (we call them guess variables) near the beginning of the apparatus. I am going to introduce two targets and the two fiddle variables right now. I am going to use ω , the frequency, as one of the guess variables, to try to get to the end of the apparatus in a quarter of a wavelength instead of some other fraction of wavelength. I am going to use the heat that is put in from the burner, and the heat that is taken out at the middle heat exchanger, as two more guess variables. Then, as targets at the far end, I am going to use the impedance at the mirror-image plane—both real and imaginary parts—and the temperature at the room-temp heat exchanger. DeltaE will display the table of these guess and target variables; here it is. These values of the guess variables: 60 Hz, 400 watts, and -400 watts, were my personal guesses, based on experience; but you can see they weren't perfect, because the targets didn't come out exactly right.

Now, if you run DeltaE on this file again, it will fiddle around with the guess variables until it gets all the targets correct. If we display the guess/target table again, you can see that it hit the targets. It didn't have to change the guesses too far: from my guess of 60 Hz to 61 Hz; from my guess of 400 watts to 332 watts; etc.

(vg 7.31) Now let's look at the output printout again. Remember these were the two things we were upset with. Now it has gotten the impedance at the end and the room-temp temperature correct. Now we have 320 K for our dump temperature, and we have zero pressure amplitude, both real and imaginary parts, at the mirror-image plane. This should give you an idea of what this computer program can do.

Now we can tell the guy in Bangladesh that he should expect to have to put in 300-and-something watts in his burner at a temperature like about 660 Kelvin. If he does that, he can expect to get 35 W of cooling power in his refrigerator at -10 degrees C. That ought to be cold enough to keep medicine cold.

8. Summary

Here is what we know how to do well. Start with the basic fluid-mechanics equations. Make the acoustic approximation. Integrate analytically in the plane perpendicular to the sound-wave propagation, to get Rott's equations. Then integrate numerically along x to get everything else you want.

I think this is mostly understood in the acoustic approximation, although I think we do not really understand a lot of details at the heat exchangers yet.

Other things for the future, besides the heat exchanger details, include high-amplitude corrections to the acoustic approximation, such as turbulence, streaming, harmonic generation. High amplitude is important for practical systems, because that gives us more power per unit volume of apparatus, and applications are always going to be sensitive to power-per-unit volume or unit weight.

The future will also bring challenges in the real applications, both in picking applications that make sense from an economic point of view, and also in inventing whatever will be required to make them real, to make them work economically.

That is it. Thanks.

References

- [1] H. B. Callen. *Thermodynamics and an introduction to thermostatistics*. Wiley, 1985.
- [2] P. M. Morse. *Thermal physics*. Benjamin, 1969.
- [3] L. E. Kinsler, A. R. Frey, A. B. Coppens, and J. V. Sanders. *Fundamentals of acoustics*. John Wiley and Sons, 1982.
- [4] A. D. Pierce. *Acoustics: An introduction to its physical principles and applications*. Acoustical Society of America, Woodbury, New York, 1989.
- [5] N. Rott. Damped and thermally driven acoustic oscillations in wide and narrow tubes. *Z. Angew. Math. Phys.*, 20:230, 1969.
- [6] G. Walker. *Stirling engines*. Clarendon, Oxford, 1960.
- [7] G. Walker. *Cryocoolers*. Plenum, New York, 1983.
- [8] I. Urieli and D. M. Berchowitz. *Stirling cycle engine analysis*. Adam Hilger, Bristol UK, 1984.
- [9] P. H. Ceperley. A pistonless stirling engine—the traveling wave heat engine. *J. Acoust. Soc. Am.*, 66:1508–1513, 1979.
- [10] P. H. Ceperley. Gain and efficiency of a short traveling wave heat engine. *J. Acoust. Soc. Am.*, 77:1239–1244, 1985.
- [11] J. C. Wheatley, T. Hofler, G. W. Swift, and A. Migliori. Understanding some simple phenomena in thermoacoustics with applications to acoustical heat engines. *Am. J. Phys.*, 53:147, 1985.
- [12] T. B. Gabrielson. Radiation from a submerged thermoacoustic source. *J. Acoust. Soc. Am.*, 90:2628, 1990.
- [13] K. M. Godshalk, C. Jin, Y. K. Kwong, E. L. Hershberg, G. W. Swift, and R. Radebaugh. Characterization of 350 Hz thermoacoustic driven orifice pulse tube refrigerator with measurements of the phase of the mass flow and pressure. *Advances in Cryogenic Engineering*, 41:(to be published), 1996.
- [14] S. L. Garrett, D. K. Perkins, and A. Gopinath. Thermoacoustic refrigerator heat exchangers: Design, analysis, and fabrication. In G. F. Hewitt, editor, *Heat Transfer 1994: Proceedings of the 10th International Heat Transfer Conference. Vol. 4*, pages 375–380, Rugby, England, 1994. Institution of Chemical Engineers.

- [15] T. J. Hoffer. *Thermoacoustic refrigerator design and performance*. PhD thesis, Physics department, University of California, San Diego, 1986.
- [16] T. J. Hoffer. Concepts for thermoacoustic refrigeration and a practical device. In Chairman Paul Lindquist, editor, *Proceedings of the 5th International Cryocoolers Conference*, pages 93–101, Wright-Patterson AFB, OH, August 1988. Wright-Patterson Air Force Base. Monterey, California.
- [17] L. D. Landau and E. M. Lifshitz. *Fluid Mechanics*. Pergamon, 1982.
- [18] R. W. Fox and A. T. McDonald. *Introduction to fluid mechanics*. Wiley, 1985.
- [19] A. Bejan. *Entropy generation minimization: The method of thermodynamic optimization of finite-size systems and finite-time processes*. CRC press, 1995.
- [20] G. W. Swift. Thermoacoustic engines. *J. Acoust. Soc. Am.*, 84:1145–1180, 1988.
- [21] G. W. Swift. Thermoacoustic engines and refrigerators. *Physics Today*, pages 22–28, July 1995.
- [22] G. W. Swift. *Encyclopedia of Applied Physics*, chapter Thermoacoustic engines and refrigerators, page ?? American Institute of Physics, to be published in 1997.
- [23] R. Radebaugh. A review of pulse tube refrigeration. *Adv. Cryogenic Eng.*, 35:1191–1205, 1990.
- [24] W. P. Arnott, H. E. Bass, and R. Raspet. General formulation of thermoacoustics for stacks having arbitrarily shaped pore cross sections. *J. Acoust. Soc. Am.*, 90:3228–3237, 1991.
- [25] G. W. Swift and R. M. Keolian. Thermoacoustics in pin-array stacks. *J. Acoust. Soc. Am.*, 94:941–943, 1993.
- [26] G. W. Swift and W. C. Ward. Simple harmonic analysis of regenerators. *Journal of Thermophysics and Heat Transfer*, 10:652–662, 1996.
- [27] W. M. Kays and A. L. London. *Compact heat exchangers*. McGraw-Hill, New York, 1964.
- [28] P. Merkli and H. Thomann. Thermoacoustic effects in a resonance tube. *J. Fluid Mech.*, 70:161–177, 1975.
- [29] N. Rott. Thermally driven acoustic oscillations, part iii: Second-order heat flux. *Z. Angew. Math. Phys.*, 26:43, 1975.
- [30] W. C. Ward and G. W. Swift. Design environment for low amplitude thermoacoustic engines. *J. Acoust. Soc. Am.*, 95:3671–3672, 1994. Fully tested software and users guide available from Energy Science and Technology Software Center, US Department of Energy, Oak Ridge, Tennessee. For a beta-test version, contact ww@lanl.gov (Bill Ward) via Internet.

- [31] D. Gedeon. A globally implicit Stirling cycle simulation. In *Proceedings of the 21st Intersociety Energy Conversion Engineering Conference*, page 550. American Chemical Society, 1986.
- [32] J. Gary, A. O'Gallagher, and R. Radebaugh. A numerical model for regenerator performance. Technical report, NIST-Boulder, 1994. Also known as REGEN 3.1 Users Guide.
- [33] A. A. Atchley. Standing wave analysis of a thermoacoustic prime mover below onset of self-oscillation. *J. Acoust. Soc. Am.*, 92:2907-2914, 1992.
- [34] R. C. Tolman and P. C. Fine. On the irreversible production of entropy. *Rev. Mod. Phys.*, 20:51-77, 1948.
- [35] J. C. Wheatley, T. Hofer, G. W. Swift, and A. Migliori. An intrinsically irreversible thermoacoustic heat engine. *J. Acoust. Soc. Am.*, 74:153-170, 1983.
- [36] A. Migliori and G. W. Swift. Liquid sodium thermoacoustic engine. *Appl. Phys. Lett.*, 53:355-357, 1988.
- [37] J. R. Olson and G. W. Swift. Similitude in thermoacoustics. *J. Acoust. Soc. Am.*, 95:1405-1412, 1994.
- [38] G. W. Swift. Analysis and performance of a large thermoacoustic engine. *J. Acoust. Soc. Am.*, 92:1563, 1992.
- [39] S. M. Hino, M. Sawamoto, and S. Takasu. Experiments on transition to turbulence in an oscillatory pipe flow. *J. Fluid Mech.*, 75:193-207, 1976.
- [40] M. Ohmi and M. Iguchi. Critical Reynolds number in an oscillating pipe flow. *Bulletin of the JSME*, 25:165, 1982.
- [41] M. Ohmi, M. Iguchi, and I. Urahata. Flow patterns and frictional losses in an oscillating pipe flow. *Bull. JSME*, 25:536-543, 1982.
- [42] W. L. M. Nyborg. Acoustic streaming. In W. P. Mason, editor, *Physical Acoustics, Volume IIB*, page 265. Academic Press, 1965.
- [43] J. R. Olson and G. W. Swift. Acoustic streaming in pulse tube refrigerators: Tapered pulse tubes. *Cryogenics*, ???, 1997.
- [44] R. J. Gibson. An experimental comparison of a pin stack to a conventional stack in a thermoacoustic prime mover. Master's thesis, US Naval Postgraduate School, 1996.
- [45] M. S. Reed. Measurements with wire mesh stacks in thermoacoustic prime movers. Master's thesis, US Naval Postgraduate School, 1996.
- [46] D. Gedeon and J. G. Wood. Oscillating-flow regenerator test rig: Hardware and theory with derived correlations for screens and felts. Technical Report 198442, NASA Lewis Research Center, 1996.

- [47] M. P. Mitchell, R. Yaron, S. Shokralla, J. Yuan, P. E. Bradley, and R. Radebaugh. Etched foil regenerator. *Adv. Cryogenic Eng.*, 41, 1996. (to be published).
- [48] F. P. Incropera and D. P. DeWitt. *Introduction to heat transfer*. Wiley, 2nd edition, 1990.
- [49] J. R. Olson and G. W. Swift. Energy dissipation in oscillating flow through straight and coiled pipes. *J. Acoust. Soc. Am.*, 100:2123–2131, 1996.
- [50] Linear motor and linear alternator papers from the Stirling community can be found in the proceedings of several conference series: The Intersociety Energy Conversion Engineering Conferences, the International Stirling Engine Conferences, the Cryocooler Conferences, and the Cryogenic Engineering Conferences.
- [51] J. A. Corey and G. A. Yarr. HOTS to WATTS: The LPSE linear alternator system re-invented. In *Proceedings of the 27th Intersociety Energy Conversion Engineering Conference*, pages 5.289–5.294, Warrendale PA, 1992. Society of Automotive Engineers. No editor; general chairman Buryl McFadden.
- [52] F. W. Giacobbe. Estimation of Prandtl numbers in binary mixtures of helium and other noble gases. *J. Acoust. Soc. Am.*, 96:3568–3580, 1994.
- [53] A. M. Fusco, W. C. Ward, and G. W. Swift. Two-sensor power measurements in lossy ducts. *J. Acoust. Soc. Am.*, 91:2229–2235, 1992.