

University of Toronto
ADVANCED PHYSICS LABORATORY

LPP

Linear Pulse Propagation and Dispersion

Originally created by R.S. Marjoribanks

Write-up modified by J.H. Thywissen v1.1: 30 Sept 2004

Minor modifications by D. Bailey v1.2, 1.3: January 2016 and October 2018

Rewritten by R.S. Marjoribanks v4: September 2020

Objectives

To observe and understand the physics of dispersion in the propagation of waves.

Section 1. Background

Crowds disperse, fumes disperse -- in either case, they just spread out and dissipate until they are unrecognizable. What does it mean for *waves* to disperse?

Imagine listening to a symphony orchestra or an excellent home sound system. You depend on the sound arriving at your ear in pretty much the same form that it left the sound system. If sound dispersed like fumes, as they spread out around the room, what would be the point of paying a lot of money for tickets, or a sound system with less than 0.02% total harmonic distortion?

For sound waves, things tend to work out pretty well as long as waves of different frequency all travel at exactly the same speed. But imagine waves for which high frequencies traveled faster than low ones. If the orchestra all hit the first note together, you'd hear the piccolo start before the base viol did. But it's worse than that: you know from Fourier analysis that arbitrary sounds can be decomposed into constituent waves, with a spectrum of frequencies. That would mean that a sound leaving your home sound system speakers would decompose as it propagated towards you -- the constituent frequencies would be teased apart, in the sound, with the high ones arriving first, and low ones coming along later. Seats closer to the sound would be much better, because the sounds would virtually fall apart by the time they got to the cheap seats in the back.

Fortunately, this doesn't usually happen with sound. But it can, actually, for sound in a sewer pipe, say, or other sort of waveguide structure -- satisfying boundary conditions in space can change the relationship of the waves to themselves, and make wavefronts of one frequency travel at a different speed than wavefronts of another frequency. In free space, sound waves are *nondispersive*, but in a waveguide, or for sound waves in materials other than air, for transverse waves or for extreme frequencies, waves can be *dispersive*.

In a waveguide there are different ways (modes) in which waves can propagate. Each mode of propagation can have a different relationship between frequency and wavefront speed. The same is true for optical waves. In telecommunications in fiber-optic waveguides (optical fibers), light waves can be dispersive, both because of the glass they travel in (material dispersion) and also because of the waveguide effect (modal dispersion). It may even be possible for the two kinds of dispersion to behave oppositely, and partially cancel out.

This is of tremendous practical importance. Imagine sending digital information in which the *bits* were ultrafast optical pulses. What's the information capacity of

such a system? Well, if the pulses are dispersive, and each optical pulse spreads out, with high frequency components within each pulse running ahead of the low-frequency components, eventually the adjacent bits will spread out until they overlap each other. If this means that you can't tell the bits apart, this will imply a serious limit for how much information can be passed per second. Or, equally, how *far* a fiber optic link can reach before the bits must be detected and generated again afresh.

So these experiments will introduce you to dispersion, beginning with an acoustic waveguide in which you can identify modes, measure modal dispersion relations, and discover the effects of dispersion for wave packets of sound.

[See the web site for sound clips demonstrating many of these concepts]

Experiments on dispersion, following up on LPP – The 100 fs Er-fiber laser experiment (UFF) will let you measure dispersion (material, and modal dispersion) in optical fibers, as well as a number of linear and nonlinear optical effects that make such a laser possible.

Safety Reminders

- Ear protection (either plugs or mugs) must be worn when the experiment is generating loud sound levels; this ear protection must be worn by all who are in the room for more than a short period, e.g. students working on the SONO or UFF experiments.

NOTE: This is not a complete list of every hazard you may encounter. We cannot warn against all possible creative stupidities, e.g. juggling cryostats. Experimenters must use common sense to assess and avoid risks, e.g. never open plugged-in electrical equipment, watch for sharp edges, don't lift too-heavy objects, If you are unsure whether something is safe, ask the supervising professor, the lab technologist, or the lab coordinator. When in doubt, ask! If an accident or incident happens, you must let us know. More safety information is available at <http://www.ehs.utoronto.ca/resources.htm>.

Section 2. Outline of Experiment.

There are six “steps” outlined below. Although they should be followed in order, each of them can be pursued with more or less care (e.g., data, attention to systematic effects, parameters explored, etc.). Steps I to IV cover the basics of dispersive waves, and you should be able to complete them in your first four labs. Step II is especially important. Steps V and VI are for “the experts” who are returning to this lab after the standard 3 weeks.

Step I. Understanding (and playing with) the equipment

- Read the theoretical derivation of the dispersion relation of modes in an acoustic waveguide. These follow directly from the fact that all repeated reflections from the walls must not interfere destructively, and that leaves only resonance conditions.

- Read Appendix I, below, describing the equipment.
- The experiment is run from one LabVIEW Virtual Instrument (VI) on the experiment's computer desktop: **AcousticWaveguideExperiment.vi** (see This VI sends digital waveforms of your design to a digital-to-analog converter on the National Instruments card in a PCI slot in the computer. That converter will drive signals on AO0 and AO1 on the break-out panel next to the computer. You can use BNC-to-banana-plug cables to connect those outputs to the blue patch-box which connects to three speakers at the end of the acoustic waveguide. How you drive these speakers will determine what mode you can excite in the waveguide, depending on frequency and launch-angle.
- Read Appendix II, below, describing the control and acquisition software
- Start by driving all three speakers identically -- patch one DAC output into speaker **A**, and then use jumpers to put the other speakers **B & C** in parallel. Call this mode (+,+,+), where the ordered triple stands for the speakers (A,B,C) and the + signs mean that all have the same phase and amplitude.
- To begin, set the *OutputWaveform* tab to drive a *steady* sine wave oscillation; note the frequency and amplitude controls. Start at 3 kHz frequency and about 200 mV amplitude. Monitor the drive signal (use a BNC Tee) on CH 1 on the digital oscilloscope, using it to trigger the scope and to show a reference that doesn't change.
- There is a travelling microphone on a tiny cart that can move up and down the acoustic waveguide, to sample the oscillations at different places. It is labelled **RoverMike** and it's connected to a tape measure that will let you push and pull it along inside the waveguide while measuring its displacement in the z-direction. Use the microphone bias-voltage box to power that microphone and to connect its output to CH 2 input on the digital oscilloscope. (Turn off the bias voltage switch, when finished)
- With scope CH 1 as a reference, work out how the signal from **RoverMike** will help you find the effective wavelength of sound inside the waveguide – for this mode (+,+,+). As you drag **RoverMike** down the waveguide, you'll see the measured waveform slide in time, as the relative phase changes. When it has slid a whole cycle, and looks like the waveform you began with, then you have moved **RoverMike** through one whole wavelength – *for this mode* (+,+,+).
- Other possible drive configurations are (+,0,-) and (+,-,-,+), where the double minus sign means opposite phase and twice the strength. For the state "0" you need only jumper an input to ground. The "-“ modes are controlled from the *OutputWaveform* tab: drive the *second* analog output channel AO1 in the opposite phase relative to AO0 -- use the multiplier "-1" to do this. The wiring of the output channels to the speakers must also be changed around, to produce the different drive configurations.
- Last part of this Step 1, get a big picture – a bird's-eye view or survey map of what you'll be working with: listen to the signal after the waveguide, at a variety of different frequencies high and low, and configurations (+,+,+), (+,0,-), and (+, - -, +). What differences do you hear? What is systematic about the differences you hear?

Experimentalist's thinking: *trust, but verify*

- How do you know for a fact that the speakers are all the same?
- How important is it, that they each drive the same amplitude? How much could you tolerate them doing different things?
- How will you check for yourself that the speakers all perform the same, or else characterize their differences? You have enough tools, for 10 minutes work you don't have to rely on assumptions.

Step II. Dispersion relations for different modes

- With a (+,+,+) drive, make a series of measurements at decreasing frequencies, down at least to 500 Hz. Plot this as waveguide wave number $k = 2\pi/l$ on the x-axis, and angular frequency $\omega = 2\pi\nu$ on the y-axis. This plot is the *dispersion relation* of the fundamental mode $(n_x, n_y) = (0,0)$ of the waveguide.
- Now drive a different mode: use the *OutputWaveform* tab to drive the second analog output channel AO1 out of phase by π , or 180 degrees, relative to AO0 -- use the multiplier "-1" to do this. Then wire the speakers to drive in this pattern: (+,0,-). Repeat for this drive the experiment you made above for (+,+,+). Can you still get to 500 Hz and hear the signal, and record the amplitude? Add your new data to the dispersion plot you made above, for mode
- Arrange again, to provide a drive which is (+,- -, +), where the double minus sign means opposite phase and twice the strength. Repeat your measurements, again going to as low a frequency as you find feasible. Plot this different drive condition, too, on the dispersion plot you made above).

Analysis and Interpretation

- How do your observations compare to the analytic theory formulae in the Primer on Dispersion? Be quantitative – overlay experiment with theory if possible.
- How have you treated experimental error? How much should be expected of your experimental data? Within this expectation, does it agree with theory? Why, or why not? Is there systematic or random error that you've neglected?
- What is the nature of the different modes driven by (+,+,+), (+,0,-) and (+,- -,+)? Why do these different drives give different modes? Can you construct a dot-product between these ordered triples that helps to explain?

Experimentalist's thinking: *trust, but verify*

- If you give a value, but no limits plus/minus around that value, then it could mean *anything*. If it could mean *anything*, then it means nothing.
- We compare to theory, not because it is right and we want to fit in. We compare to theory because we've done an experiment so competently that we already know the truth – we want to see if the *theory* is right, then it may help explain.

Iterate: Going back to experiment with revised/refined ideas

- If you have new questions, don't hesitate to go back to the experiment to answer them.
- There are microphones set into the aluminum channel that can be moved *transversely* across the waveguide. Can they give information about the differences between different drive-modes?
- The array of 8 microphones [currently downgraded to 3 but we're working on restoring them] on a different travelling cart, **RoverArray**, can be fed into 8 channels of analog-to-digital conversion on the National Instruments multifunction card, to give resolution in the vertical (*y*) direction.
- Do you see the connection to vector-spaces here? Only certain modes are possible. The wave equation is a linear differential equation – any sound excitation you make in the waveguide must decompose over these modes, so we can *analyze* (like Fourier analysis) the sounds we launch into amplitudes of a complete set of basis vectors.

Mode Analysis

The goal of the quantitative analysis for the dispersion relationships of different modes should be to compare your experimental data to the theory for the experiment. Use a fitting program which, if possible, gives you the Chi Squared of the fit, as well as uncertainty on your fitting parameters.

Most importantly, you should compare the dispersion curves you found to the expected shape,

$$\omega = c_s \sqrt{k_z^2 + (m\pi/L_x)^2},$$

where ω is the frequency, c_s is the speed of sound, k_z is the wave number along the waveguide, m is the mode number, and L_x is the transverse size of the waveguide (because L_y is smaller, we look at the lowest few modes and avoid exciting the $n_y > 0$ modes). Use the speed of sound and the size of the waveguide as fitting parameters.

Group speeds should also be compared to theoretical expectations.

Dispersion relations: *everything is here*

- Everything about the pulse propagation in a waveguide is here, in the dispersion curves, by using values, slopes of chords ω/k , and derivatives of different order. Your dispersion curves are the most fundamental part of this experiment
- For linear waves, every kind of pulse can be described as a linear combination of waves of different frequencies – basis states of a vector space of infinite dimension, with pulses decomposed over this basis by Fourier transforms
- For waves in a waveguide, every kind of propagation can be described as a linear combination of normal modes of the waveguide, a countably infinite basis in principle, but since atoms and molecules are finite, there is a real cutoff for frequency and wavelength.

References

Meykens *et al.*, Am. J. Phys. **67**, 5, (1999).

Step III. Phase speed

The primer will have introduced you to the difference between phase speed and group speed, if you haven't already seen it in lectures in physics or engineering. Click [here for an animation](#) [go to website if you are reading a printout of this document] which shows a pulse traveling by at the group speed, with the carrier wave traveling inside the envelope at the phase speed.

How can you measure phase speed, using **RoverMike** or the other microphones? You'll need to track a wave-crest, and follow it; see what time it takes to travel a measured distance.

Measure phase speeds for different modes, and for different drive-frequencies. Use the dispersion relations you found in Section I, above, and determine from them the phase speed which follows, for your mode and frequency. How does your *measured* value compare to the value you *deduce or infer* from your dispersion relation measurements? How does your measured phase speed compare to predictions from analytic theory, starting with the dispersion formulae in the Primer.

Track your error estimates in all cases, and see whether your measurements give comparable answers, within estimated error.

Experimentalist's thinking: *a step ahead for how to do things*

- We think of phase speed as being the speed of a wave crest, but that's a problem, for crests in a pulse — the crests skew one way on the leading edge of a pulse and the other way on the trailing edge, because of the slope of the envelope. But really phase speed is the *speed of a surface of constant phase*. So, actually, it's much better to follow the zero crossings of a wave, if you want to follow the phase speed – zero crossings do not get skewed left or right by pulse envelopes.

Step IV. Group speed

How can you measure group speed, using **Rover** or some other(s) of the microphones? Hint: the *OutputWaveform* tab lets you drive not just a sine wave, but a pulse of a certain carrier frequency. And you have microphones, or can arrange microphones, at different distances. How fast does the pulse travel, for different carrier frequencies? Be careful to use at least 20 cycles in your pulse -- if you go shorter, you'll have larger and larger frequency content in your shorter and shorter pulses, so you will have big error bars in actual frequency, and weird effects which are part of the next section. The bandwidth of your pulses, shown

in the PowerSpectrum display within the *OutputWaveform* tab, will be part of your error-bars on the frequency.

Start at frequencies well above the cutoff, where the dispersion-curve is nearly linear, and then work down toward the cutoff frequency. What happens when you try to propagate a pulse that has a center-frequency *below* the cut-off frequency? How do you explain this? Can you quantify the effect, by measuring the amplitude of your pulse as it propagates down the waveguide, for different frequencies? Can you connect this to results from the simple analytic theory in the Primer?

Compare your measured group speeds, for different centre-frequencies and different modes, to the value which you can determine from your *measured* dispersion relations (how is that found?). Then compare also to values found from analytic theory, for the dimensions of the lab waveguide. As always, track your error estimates, which will serve as a prediction for how closely you might expect experimental results to match theory, or the values expected from other experimental results.

Step V. Pulse dispersion (Group velocity dispersion, GVD)

The **AcousticWaveguideExperiment.vi** is designed mainly to do a very nice job to look in detail at short pulses that have propagated down the waveguide. This Step will allow you to look at the most important phenomenon of this experiment, the phenomenon that dominates all aspects of ultrafast lasers and laser-matter interaction, and which dictates the main methods of telecommunications using optical fibers. In a different world for which the air we breathe absorbed energy in the human acoustic range, we would be unable to listen to musical performances as we do today.

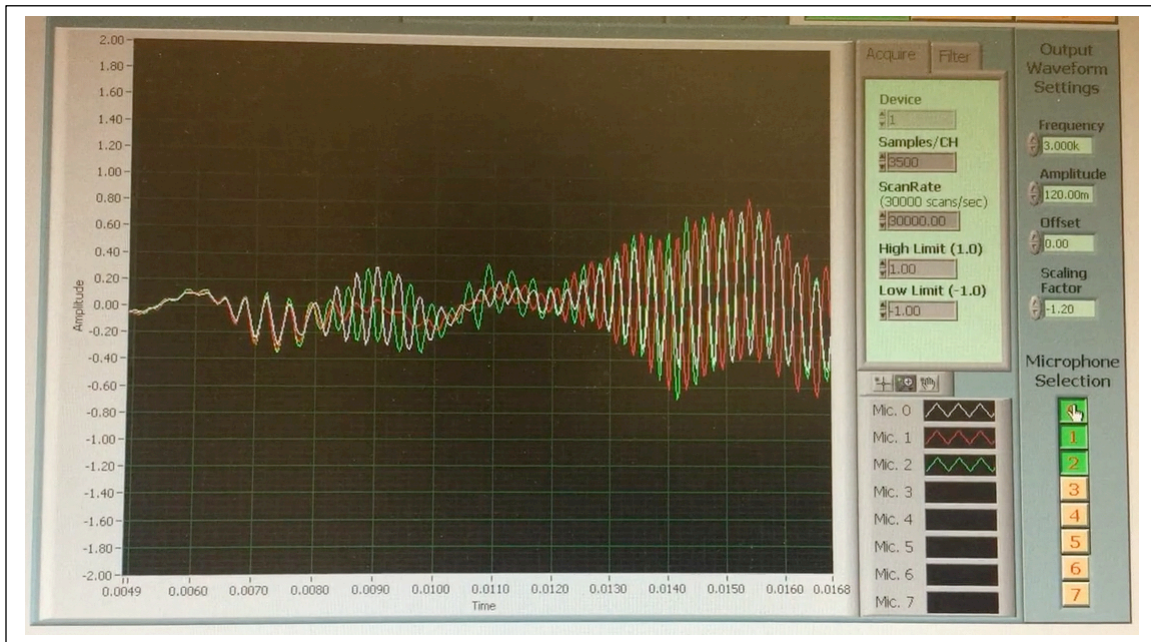
You know, from the dispersion relation, that the energy in different frequencies will travel at different group speeds, according to the slope of the tangent to the dispersion curve at each frequency. If you make a short pulse (only a few cycles, or even less if you wish to make a half-cycle pulse or an impulse), it will all at once contain a substantial range of frequencies, each travelling at its own speed. So the pulse will unravel, or stretch out, as it propagates.

Generate a short pulse from the tab *Output Waveform*, and then capture the waveform from different microphones into the tab *Acquire Waveform*. Listen to the sound you hear in the waveguide, and try different frequencies for the carrier wave of your pulse, until you detect things worth investigating.

Draw a selection box over the time-span of the most interesting part of the waveform you capture, and then go to the tab *Spectrogram* to convert your selection from waveform to joint time-frequency spectrogram. Look very closely also using the tab *Intensity Plot* and the tab *Animated Plot*.

Explain the shapes you see, in terms of the dispersion relations you measured above, or the analytic formulae. What is the lowest group speed you see? You may want to make a series of rough measurements from the spectrogram and

from these reconstruct the dispersion relations for each mode. Then compare these to the dispersion relations you measured directly.



Modal Dispersion: The pulse that was launched here oscillated for just a half-dozen cycles, at the source. The pulse was not launched in a pure mode, but was mostly (+, - -, +), and the carrier frequency was near the dispersion-curve cut-off. Consequently the pulse decomposed into the three different modes, which travelled separately while themselves stretching out.

Step VI. Pulse compression [A 'reach' goal]

What's the next thing that you find interesting in this experiment? Ask your T.A. or supervising professor for suggestions, or for help with your own ideas.

One hint: What can happen forward in time, in a linear physics context, can happen backwards in time equally well. This LabVIEW VI can be fairly easily edited, if you can program in LabVIEW, and the option could be created to take a sound file that you have recorded and time-reverse — thus reversing the *sign* of the frequency chirp. In that case, you could then play it back into the driving speakers in the same mode drive-configuration it was created in. What then happens to the pulse?

APPENDIX I — Guide to parts and components



Acoustic waveguide

The waveguide is an aluminum channel, about 5 cm wide and 15 cm high, approximately 6 m long. Microphone are located at the $1/4$, $1/2$ and $3/4$ marks, which can be slid up and down the height of the waveguide, and pushed in and out to sample almost every point in the cross-section.


The waveguide is mounted elastically, to isolate it from noise transmitted in the structure of the building.

At the far end, in the picture at right, there are three speakers mounted in a row up the height of the waveguide, which can be driven with different amplitudes.



Drive box -- speakers

Three 5 cm speakers with small amplifier-boards are housed in a plexiglass box, with 3.5mm stereo jacks for their connection. The speakers are driven by a National Instruments DAQ card in the computer, which includes a digital-to-analog converter. With this, a LabVIEW program on the computer drives two output lines with arbitrary waveforms. The software provided will let you control the frequency of sound applied, and the duration of sound-pulses, from continuous to pulses of a fraction of a cycle. The interface between the patch-panel for the computer card and the speakers is a small blue box with binder terminal posts which can accept banana plugs or wires.



Patch box -- speaker drive. Black terminals indicate signal ground for the speakers. Connect drive-signals from DAC0OUT and DAC1OUT by black banana plug to black terminal, and likewise red to red.

This box has three pairs of binder terminal posts which connect to the drive-speakers, a pair for each speaker. A BNC-to-banana plug adapter then lets you connect two drive-outputs of the digital-to-analog card to the speakers. Banana-to-banana wires let you connect the third speaker as you wish: one wire to *ground* the input ("0" drive), or two wires to *gang* two speakers together, with the same amplitude and phase, as you could do for (+,+,+)

Pro Tip for Experimentalists: It's not always obvious, but connections on equipment like the patch box and oscilloscope are not only +/- but these may not be merely relative, or *floating*: one terminal of lots of lab apparatus may instead actually be tied to earth-ground, the U-ground of a three-prong plug provides such an earth-ground, the white 'cold' conductor on a polarized 2-prong plug should also be referenced to ground. On an oscilloscope, for instance (unless it is battery-powered), the outside connector of a BNC cable is actually grounded back to earth, via the oscilloscope chassis and the power cord. So be aware: if you swap two connectors like banana plugs, you may not simply be reversing polarity, you may actually be short-circuiting your signal to ground, without noticing.




Break-out panel / Patch panel (BNC 2090) -- from DAC/ADC card

This patch panel facilitates connections to the multifunction acquisition and generation card (NI MIO-16) installed in the computer. It presents BNC connectors for outputs from the digital-to-analog converters (DAC), and accepting inputs to the analog-to-digital converters (ADC), as well as connectors for other timing and digital input/output functions.

We use channels DAC0OUT and DAC1OUT (analog output) to drive the speakers, and channels ACH0 through ACH7 (analog input) to record signals from microphones.

The manual for [BNC-2090](#) is available online.



RoverMike (microphone) [draft photo]

This is a single tiny microphone on wheels which can be inserted into the waveguide. It has a long fine cable attached, which can reach the length of the waveguide. The microphone is positioned using a metal spooled tape-measure, which is rigid enough to push and pull the microphone to different positions while providing at the same time a measure of the position down the waveguide. The tape measure end attaches to strong magnets on the RoverMike (be careful of your credit cards near those magnets).

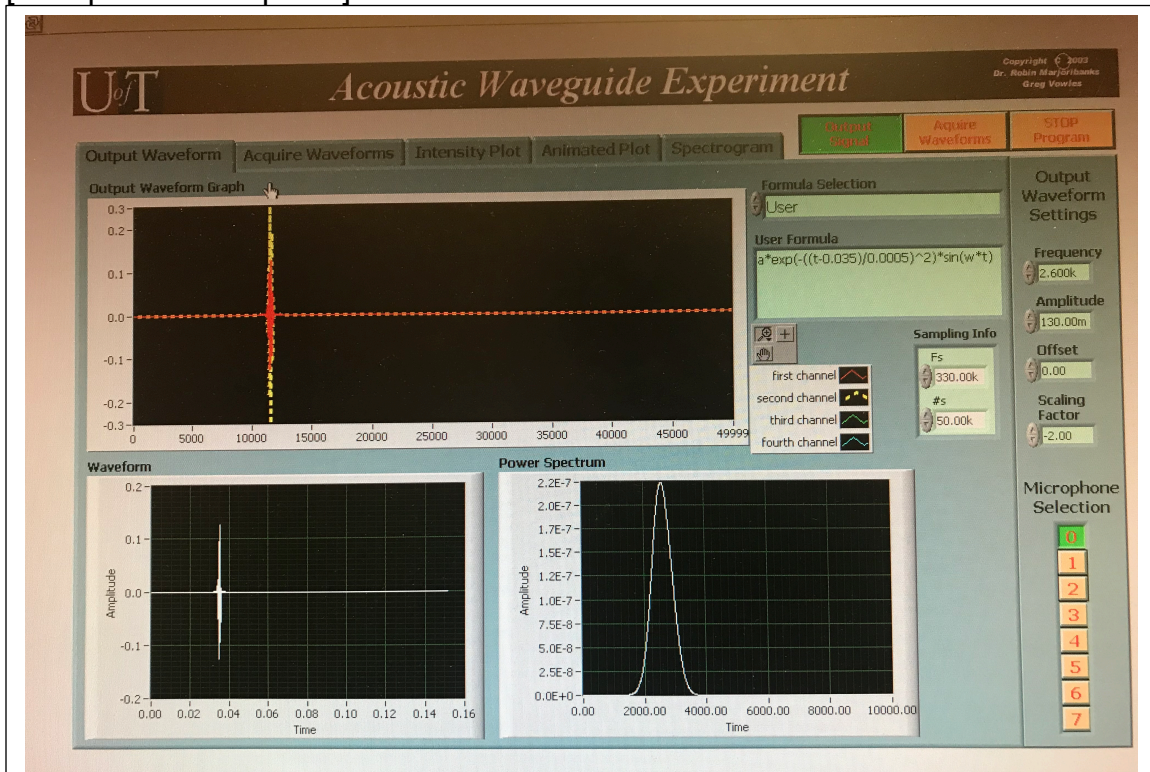
RoverArray (microphone array)

This is a microphone array on wheels which can be inserted into the waveguide. It has a long fine cable attached, which can reach the length of the waveguide. The microphone array is positioned using a metal spooled tape-measure, which is rigid enough to push and pull the microphone to different positions while providing at the same time a measure of the position down the waveguide. The tape measure end attaches to strong magnets on the RoverArray (be careful of your credit cards near those magnets).

The eight microphones of the array feed into eight input channels of the BNC-2090 patch panel; their signals are digitized by the multifunction data acquisition card of the computer, and displayed using special LabVIEW software, which makes a profile of the pressure variations along the array, and then shows the evolution of this profile, of the wave within the waveguide, in time and for different positions.

APPENDIX II — Guide to control and acquisition software

[draft pictures to update]



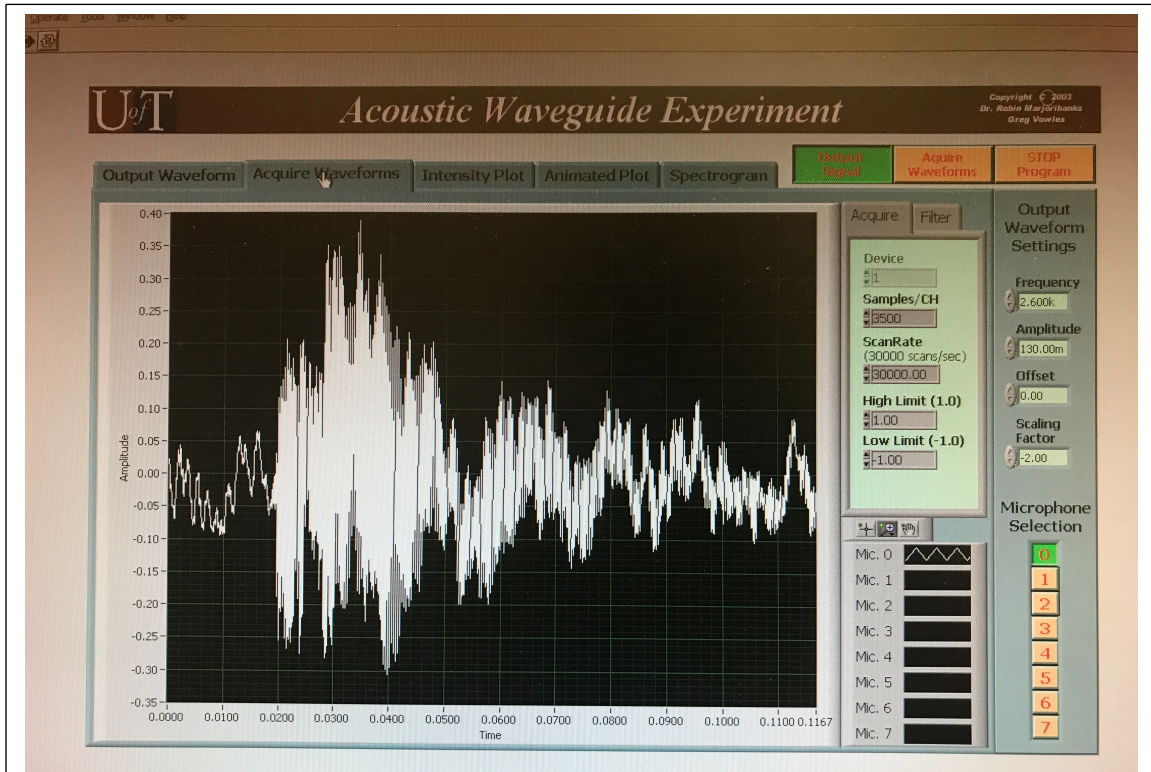
Output Waveform: This tab lets you control the waveform sent to the amplifiers for three speakers, and to visualize those signals in time and frequency. The same general waveform is sent to all three speakers at the far end of the waveguide, but you can control the relative amplitude (including sign) of one speaker, using the control box “scaling factor”.

The VI also lets you configure the acquisition of sound on up to 8 microphones you can configure.

Across all tabs, you can toggle the speaker drive-signals on and off with one button (*Output Signal*); another button lets you toggle on/off the acquisition of microphones you likewise toggle on/off individually; a Master STOP button stops all operations of the VI (to restart, click the RUN arrow at upper left). In the first part of the experiment, you’ll use a single-frequency sine wave, and control the speaker drive in different ways.

In the second part of the experiment, you’ll use the *User Formula* box to specify a pulse of different-duration envelopes, and different carrier-wave frequencies.

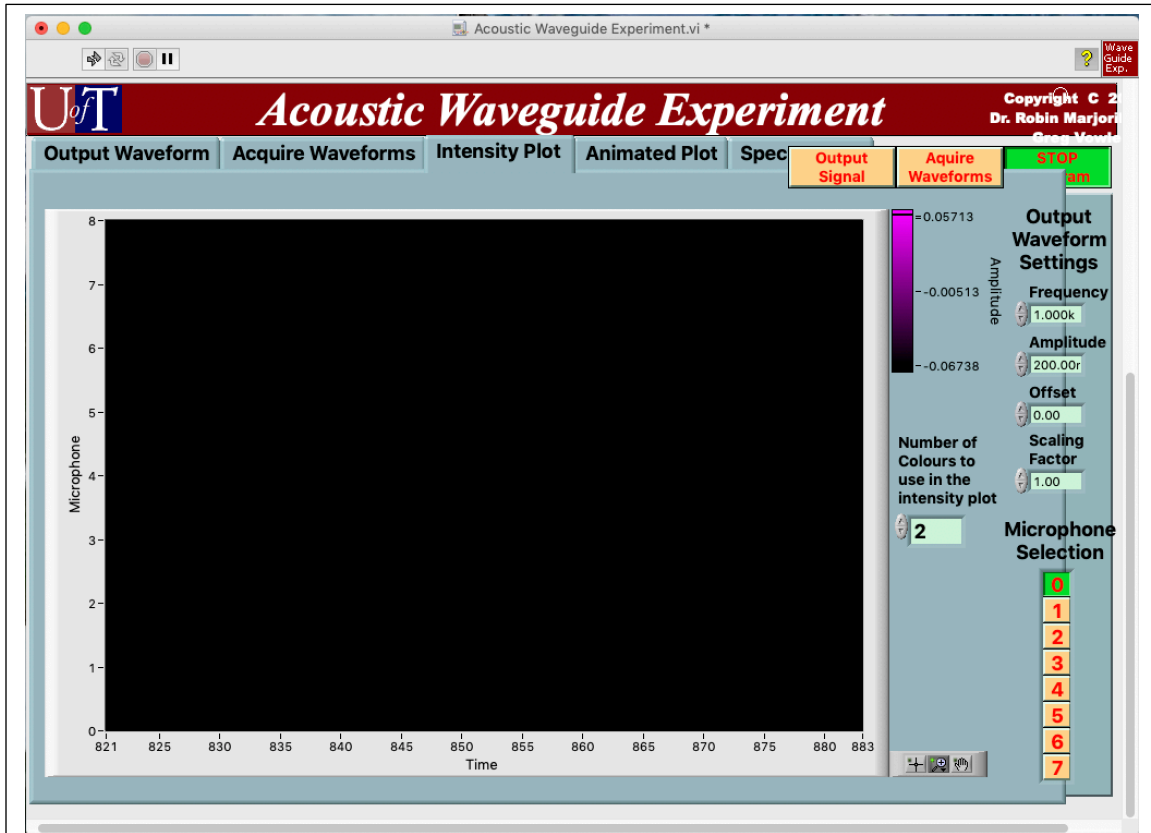
PHY325,326,327 / 425,426,427: Advanced Physics Laboratory
Experiment 42: Linear Pulse Propagation and Dispersion



Acquire Waveform: This tab displays the signals acquired on all microphones. You can toggle the acquisition/display of all microphones by their index, 0–7. Default scan rate and record length are optimal for most purposes. The purpose of multiple microphones really only becomes clear when characterizing pulse propagation.

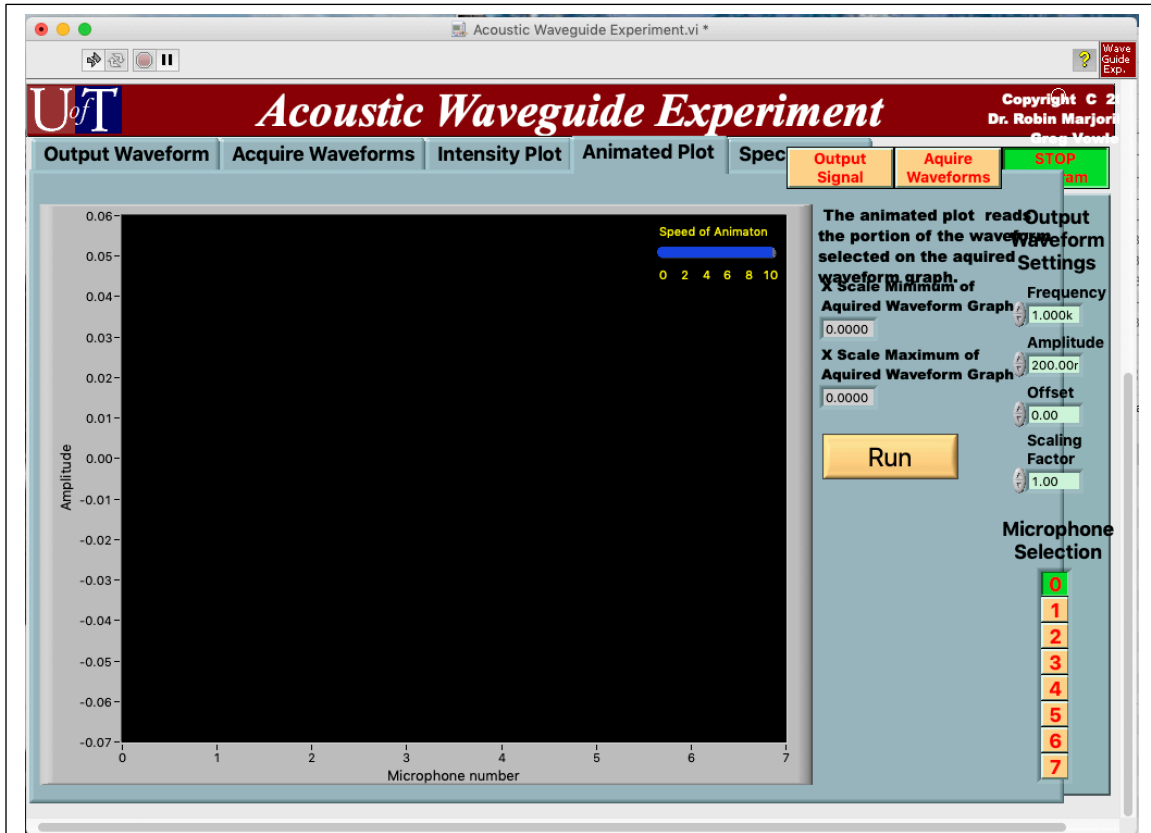
Note: once you have got a nicely framed waveform (you can start/stop *Acquire Waveforms* until you get what you want), you can then use the mouse to select a range within that data, to process further – handing your selected data off to be treated in the next two tabs.

PHY325,326,327 / 425,426,427: Advanced Physics Laboratory
Experiment 42: Linear Pulse Propagation and Dispersion



Intensity Plot: This tab uses the data as acquired and sub-selected in the previous tab, *Acquire Waveforms*.

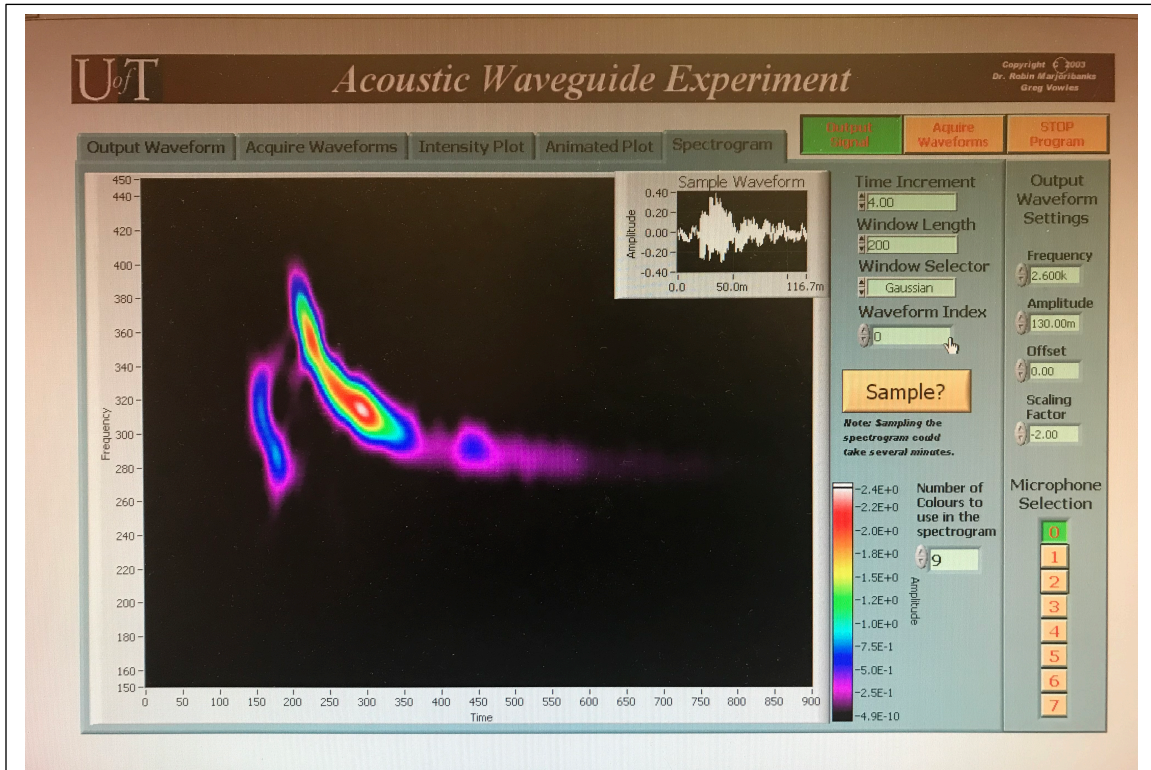
In the window is a 3D false-colour snapshot of the signals on each microphone against time. So any vertical line lets you visualize at any particular time how the air pressure is distributed from top to bottom of the waveguide.



Animated Plot: This tab complements the previous one (which showed the whole history of microphone signal intensities vs. time). Here, across the plot is the instantaneous signal on each microphone, linked by a line, to sketch for you the profile of intensities across the waveguide, top to bottom.

If you launch a pulse that does not excite a pure mode, and if that pulse suffers dispersion in the waveguide, what this plot will animate for you is the one pulse as it arrives at different times, in three different modes. Whatever the decomposition of the pulse, over different mode-components (basis vectors), you'll see the pattern of each transverse mode wiggling on the display, each with its component amplitude.

A blue slider-bar at upper-right of the data-plot controls the speed of playback.



Joint Time-Frequency Spectrogram: As for the previous two tabs, this tab processes your acquired data in just the subset you have formatted in the tab *Acquire Waveform*. This tab displays changing sounds, graphically, the way that birdsong references sometimes display the chirps and whistles of different species of birds.

A temporal waveform can be written as an equivalent Fourier transform – one or the other, temporal domain or frequency domain – but to match our perception of changing frequencies at different times we can also take time intervals, and make separate Fourier transforms for each interval, piecemeal, and then plot the result as an evolving power spectrum, interval-by-interval. This is a joint time-frequency spectrogram.