

# Measurement of the Speed of Light with a Pulsed Laser

## Background

---

Attempts to measure the speed of light date back at least to Galileo during the Late Renaissance. In his experiment (reported in 1638), he flashed a lantern to a friend on a distant hill who responded by flashing his own lantern back as soon as he saw the light. Alas, the experiment failed because the time necessary for light to travel from one hill to the other was much shorter than the time it took for each person to perceive and respond to the flash. But the idea of Galileo's—to measure the time between sending and receiving a flash of light over a known distance—is perfectly valid and indeed is the principle of this experiment.

The apparatus used here replaces the persons used by Galileo with electronic circuits, which have much more rapid response times that can be much more precisely measured. The flash of light is created by a pulsed laser diode, a device very similar to a laser pointer, except this laser is switched on and off—*pulsed*—at a very high rate: 1 million times per second (1 MHz). The pulsed beam can be bounced off of a mirror and the return pulses detected by a very responsive (“fast”) detector. Each pulse of the laser creates two signals, one from the laser itself (the “trigger”) and one from the detector when it picks up the returning pulse. These two signals can be compared to each other with a fast oscilloscope, and the time between them can be measured as a function of the distance between the laser/detector apparatus and the mirror. By varying the total distance from the laser to the detector  $L$  and monitoring the time difference  $\Delta t$  on the scope, one can determine the speed of light  $c$ . If the time difference between the two signals  $\Delta t$  depended only on the time it took for the light to travel that distance, then  $c$  would be given by

$$c = \frac{L}{\Delta t}. \quad (1)$$

Now, if you have studied much physics, you may be bothered by one fact: by common assent, since 1983 the speed of light (in vacuum) is *defined* to be 299,792,458 m/s. Thus, an experimental measurement of it seems out of line; the unit of distance, the meter, is taken to be 1/299,792,458 of the distance that light travels in exactly 1 second. So, how can we say we are honestly *measuring* the speed of light when we use the meter to express the distance traveled by light in this experiment? Indeed, what is the point of the exercise?

These are fair questions, and they raise fundamental issues that are important in every experiment. First, any experimental design must confront the question of *standards*: how does one know that the experiment actually makes a good measurement? One of the most common ways to check is to measure a quantity that is already known, for example, a spectroscopist might check the hydrogen or mercury spectrum with her instrument to see that it gives the expected spectral lines before turning her attention to an unknown sample. Second, every real measurement has some uncertainty: there are limits to how well any measurement can be made. By measuring a known quantity, one can get an understanding of how much of the uncertainty comes from basic statistical variations—*random* uncertainty—(which can range from the impossibility of clicking a stop watch exactly the same way twice to the fundamental randomness inherent in quantum phenomena) and how much comes from reproducible but unaccounted-for physical effects—*systematic* uncertainty—(like wind resistance in projectile motion, or the afore-mentioned spectrometer expanding slightly on a warm afternoon).

The point of this exercise is therefore not to make an “independent” measurement of the speed of light, but, rather, to confront and resolve questions of experimental design, data analysis, and uncertainty analysis. The idea of the experiment is simple: shoot a pulse of light to a mirror and see how long it takes

to return; and the answer is known: it should take  $L/299,792,458$  seconds (with  $L$  in meters). The real questions are, does it? How well? And if it doesn't, why?

## Apparatus

Refer to Figure 1, a photograph of the laser and detector. Each unit is powered by a small plug-in power supply, +5 VDC for the laser and +9VDC for the detector. (Note that the connectors for each type of power supply are different so that it is hard to mix them up.) On each unit there is a BNC connector that will allow you to connect to an oscilloscope with a  $50\Omega$  BNC cable.

To operate the laser and detector, plug in the appropriate power supply and attach a BNC cable to an oscilloscope channel. There are no electronic adjustments to make on these devices.




Figure 1: Pulsed laser and detector units, looking from above. The front edge of the detector box has a small photodiode that can sense the laser beam reflected from a mirror (not shown).

The setup and measurement is conceptually simple:

1. Power up the laser. Attach a BNC cable to CH1 of the scope. “Terminate” the connection with  $50\Omega$  (see the discussion below).
2. Set the mirror a known distance away from the laser. **Do not touch the surface of the mirror.** Set the detector near the laser. Adjust the laser, mirror and detector so that the beam goes from the laser to the mirror and then back to the detector.
3. Power up the detector. Attach a BNC cable to CH2 of the scope. “Terminate” this connection with  $50\Omega$  too.
4. Look at the signals on the scope and optimize the laser/mirror alignment and scope controls to give the strongest, clearest signal.
5. Use the scope to measure the time difference between the laser’s trigger signal (CH1) and the detector’s output (CH2). Record this number  $\Delta t$ . (Hint: use the time, i.e., *vertical*, cursors.)
6. Measure (with a tape measure) and record the distance  $L$  traveled by the beam from the laser to the mirror to the detector.
7. Repeat steps 4 through 6 for different distances  $L$ .

## Complications

In spite of the conceptual simplicity of the setup, there are some tricky parts:

- Special care must be taken when measuring extremely fast, i.e., high frequency, signals. In particular, the cables from the laser and detector must have a  $50\Omega$  impedance, which matches the output impedance of the two devices. In addition, the inputs to the oscilloscope must be “terminated” with a  $50\Omega$  resistor. Fortunately,  $50\Omega$  cables are the most common type, and proper termination can be accomplished by using a BNC “tee” with a  $50\Omega$  terminator resistor attached to the open end.

**Explanation:** The reason for this special precaution comes from the fact that the signals are so fast that the propagation time of the electrical signals down the cables and into the oscilloscope is not that different from the time of the propagation of the light pulse from the laser to the detector. For ordinary “slow” signals, the two ends of a cable always have the same potential; a change of the voltage on one end of the cable is, for practical purposes, immediately felt at the other end. But for extremely rapid changes, the two ends of the cable may not be at the same voltage. In this extreme, one must consider the propagation time for a signal down the cable (and indeed, through the rest of the electronics).

To understand how this works, say the voltage at one end of the cable were to change instantaneously. Charge at that end would be displaced, and the amount of charge needed to change the potential at that same point by a given amount would depend on the cable’s *capacitance* at that point. In addition, the motion of the charge creates a local magnetic field that tends to oppose the charge motion (“Lenz’s Law”), and the size of this effect depends on the cable’s *inductance*. Both the capacitance and the inductance taken together control the speed of propagation of the signal down the cable. In detail, this becomes the physics of a “wave guide”. The capacitance and inductance are given by the dimensions of the cable and the insulation used, and these are picked so that the cable appears to act like a simple resistor of a  $50\Omega$  value if the cable were infinitely long. This insures that the propagation of the signal is “dispersion free”: all frequencies travel along the cable at the same speed (about  $2/3$  of the speed of light in a vacuum), and consequently, a pulse traveling along the cable keeps its shape. Any  $50\Omega$  cable that is not infinitely long can be made to act that way by attaching a  $50\Omega$  resistor to the end of it—the aforementioned  $50\Omega$  terminator.

- The propagation of the signals down the cables and through the rest of the electronics actually does take a significant time. You can confirm this by measuring the trigger and signal pulses when the laser and detector are pointed right at each other with almost no distance between them. So, you need to design your experiment and data analysis so that this part of the total time difference can be extracted from the results. (Hint: since this effect depends on the cables and electronics, and not on the distance that the light travels, it should be a constant. What happens when you add a constant to every measurement of the time  $\Delta t$ ?)
- You will measure the distance from the laser to the mirror back to the detector with a tape measure, and you will need to estimate the uncertainty in this measurement of the total distance  $L$ . Some things to think about: How much does the sag in the extended tape affect your measurement? How reproducible is it? (For example, how different are the measurements made by different people in your group?)
- You will measure the time difference between the trigger pulse and the detector pulse with the oscilloscope, and these pulses will be pretty close together in time. You should estimate what time/division setting will give you a clear picture of the time difference. What time difference do you expect for a pulse that travels, 1 meter, 3 meters, or 5 meters? Work this out before you attempt to set the scope controls.

- The pulses that you see on the scope will be somewhat smeared out and show some funny “wiggles”. You will need to adopt a consistent method for picking the points on the different pulses that you will use to get  $\Delta t$ . You will also need to estimate the reproducibility in these measurements and obtain an uncertainty for them.

**Strongly recommended:** Before you try to record careful data, *play with the apparatus* to see what the signals look like for different mirror distances. Try to optimize the signal by adjusting the alignment of the mirror and laser position. Think about how you would measure the pulse separation, and confirm that you can see a change in this value for different mirror positions. Do NOT just set the apparatus up and expect it to work perfectly; you need to confirm that you are measuring a real effect before you invest time analyzing possibly faulty data.

## Data Analysis

---

At the end of your measurements, you should have a **table of at least five different values** (although more is better) of  $L$  and  $\Delta t$ , along with the uncertainties in each quantity, neatly tabulated in your notebook. You should also have a **clear diagram of the apparatus** along with some **notes about how you determined these numbers** in sufficient detail such that anyone with your current training in physics could reproduce your exact procedure.

### Make a plot

It is always a good idea to plot raw data. You can much more easily spot trends and isolate data-taking mistakes by looking at a picture rather than a list of numbers. In fact, it is excellent technique to get in the habit of plotting your data as you take it; that way you can fix experimental mistakes while you are still working with the apparatus.

### Include uncertainties in your plot

Put error bars for both x- and y-values in your plot. The error bars should be twice as long as the uncertainty value (reflecting the “plus or minus” interpretation of uncertainty). Think carefully about and double-check the units on the uncertainty to insure that the error bar length is correct. If you are using Excel or another computer program to make your plot, you will need to review how error bars are created and manipulated. (For example, Excel allows error bars to be plotted as a fixed value, a percentage of the data point, or according to the value in a specified cell.)

Now study your plot carefully. What are the units for the x- and y-axes? Are they correct (and shown)? Do any data points lie outside of the trend of the others? Is this discrepancy larger than the error bars? If you do see a potentially “bad” point, make sure that it is plotted correctly; it is easy to misplot a point by hand or mistype a value into a spreadsheet. Next, look at the error bars and note whether the horizontal bars (x-uncertainty) are significantly different from the vertical bars (y-uncertainty), especially in comparison to the x- and y-values themselves (in other words, consider the relative sizes of the *percent uncertainty* in each quantity).

### Make a preliminary calculation

The equation for the speed of light  $c$ , Eq. (1), is a simple fraction, and thus you should be able to find  $c$  from a line drawn through your data points. Do this quickly: make a line (by hand or with the line-fitting capability of your computer program, i.e., a “trendline” in Excel). Calculate  $c$  and check to see if your value is close to what it should be. It should, in fact, be pretty close. If it isn't, you have made a mistake. Typical errors are incorrect units (remember, you measured time in microseconds or nanoseconds and, depending on the measuring tape you used, the distance data may be in inches), or incorrect interpretation

of the constants in your line (Do you want the slope or inverse slope? What is the role of the y-intercept?).

### Make a weighted fit, including uncertainty

After you have verified by your preliminary calculation that you are on the right track, and you have resolved any basic mistakes that might have been made, you should revise your analysis to include a calculation of the uncertainty in your result. *Remember: your goal is to get a measurement of the speed of light and its uncertainty, because the purpose of the experiment is to understand the measurement technique.* This is where the data analysis can get tricky, so you need to think the procedure through clearly. Here are questions to consider:

- Which error bars on your plot, the horizontal or vertical ones, are larger? If one type is significantly larger (say 4 or 5 times) than the other, you can safely ignore the smaller error bars when you calculate your final result. If the x and y error bars, on the other hand, are about the same length, then you should consider both in your analysis (see below).
- Are the error bars for each data point the same length, or do different points have different uncertainties? If the uncertainties are all the same (that is, all the y-uncertainties are the same, and all the x-uncertainties are the same), then the analysis can be simplified.

You need to answer these questions before you can make a proper *weighted fit* to your data. A weighted fit includes the uncertainties in the data points into the fit calculation and produces fit parameters with uncertainties that can be used to find the *random* (or statistical) uncertainty in  $c$ . Only after you have a good value for the random uncertainty can you properly assess any *systematic* uncertainty that may be affecting your experiment.

Weighted-fit calculations are made with the assumption that the x-values are the independent variable and that these numbers are precisely known. The uncertainty in each data point should be in the y-value only. Thus, when making your fit, you may need to change the way that the data are plotted and/or modify the uncertainty in each data point so that you have y-errors only. The following cases outline the procedure you should follow to make a correct weighted fit. The theory behind the formulas given comes from the text by Bevington and Robinson. Only some results are stated; consult the book (see the references) for a full discussion.

**Case 1: The uncertainty in each data point is the same.** In this simplest case, your data set is already optimized to obtain the best weighted fit. In this case, the fitting theory shows that the uncertainty in each data point does not matter; the result of the calculation is the same regardless of whether you include the data-point uncertainty or not. The uncertainty in the fit parameters depends only on the amount of scatter of the data points around the fit line. The quantity known as the sample standard deviation for the fit line gives the uncertainty in the fit parameters. Specifically, if one fits a line  $y = mx + b$  to a set of data of  $N$  ordered pairs  $\{x_i, y_i\}$ , one obtains by the method of least squares the values of the constants  $m$  and  $b$  that minimizes the following quantity:

$$s^2 = \frac{1}{N-2} \sum_{i=1}^N [y_i - (mx_i + b)]^2 \quad (2)$$

The quantity  $s$  is the sample standard deviation, and it is a measure of how far the data points are, on average, from the fit line in the vertical direction.

In this case, the values of the fit parameters  $m$  and  $b$  are the same as for the preliminary calculation fit, and the uncertainties in the parameters  $m$  and  $b$ ,  $\sigma_m$  and  $\sigma_b$ , are given by the equations

$$\sigma_m^2 = \frac{Ns^2}{\left[N\sum_i x_i^2 - (\sum_i x_i)^2\right]} \quad \text{and} \quad \sigma_b^2 = \frac{s^2 \left(\sum_i x_i^2\right)}{\left[N\sum_i x_i^2 - (\sum_i x_i)^2\right]}. \quad (3)$$

It is tedious to calculate these quantities by hand. Most computer fitting programs will do this automatically if you can find the right procedures. Look in the help files for “weighted fits” or “uncertainties in fit parameters”. With Excel, you can use the array function LINEST to give fits with uncertainties in the parameters. Information on using LINEST is on the course website under “Notes on making a least-squares fit to a line in Microsoft Excel”.

**Case 2: The uncertainty in each data point is different and most of the uncertainty is in the y direction.** In this case, you should follow the prescription for the fully weighted fit described in Bevington and Robinson, section 6.4. The line fit will minimize the  $\chi_v^2$  (“reduced chi-square”) function:

$$\chi_v^2 = \frac{1}{N-2} \sum_{i=1}^N \left\{ \frac{[y_i - (mx_i + b)]}{\sigma_i} \right\}^2. \quad (4)$$

Equation (2) is a version of this function with all of  $\sigma_i$  being the same. By including each point’s uncertainty, the fit is weighted toward those data points with lower uncertainty. As in Case 1, many computer fitting programs have functions that can deal with fully weighted fitting, and you should investigate whether the one you use can do this kind of fitting. (It can be done in Mathematica, MATLAB, Origin, and Kaleidagraph, for example.) Unfortunately, the LINEST Excel function does not handle this type of problem.

**A note about  $\chi_v^2$ :** If you look carefully at the form of the function, you can see that in each term of the sum the numerator gives the deviation of the fit line from the data point and the denominator gives the uncertainty in the data point. So if the deviation from the line is about the same as the uncertainty itself, then the term will be about 1. If this is true for every point, then the whole sum will be equal to  $N$ , the number of data points, and thus  $\chi_v^2 = N/(N-2) \approx 1$  if  $N$  is large. We can say that if the reduced chi-square function is close to 1, then our fit matches the data set about as well as the quality of the data itself, that is, the fit is “good.” If the function is significantly larger than 1, then the fit is poor: the model applied to the data does not match it within the uncertainty of the data. If the function is significantly less than 1, then it may be that the uncertainty in the data points has been overestimated, or that the model itself is “over determined:” one is trying to extract more information from the model than can be reliably claimed to be present in the data; in other words, nature is acting more randomly than we are assuming with our hypothesis. Thus, it is a good idea to calculate the reduced chi-square function even if all of the data-point uncertainties are the same.

**Case 3: The uncertainty in each data point is different and most of the uncertainty is in the x direction.** In this case, it is easiest to simply flip the  $x$  and  $y$  data, so that the  $x$ -error bars become  $y$ -error bars. Do the fit, and obtain the fit parameters along with their uncertainty following the method of Case 2, and then transform the results to get the parameters and uncertainties that you want. For example, in the equation  $y = mx + b$ ,  $x$  is the independent variable and  $y$  is the dependent variable, that is,  $y$  is a function of  $x$ . When we transform the equation to give  $x$  as a function of  $y$  we get  $x = y/m - b/m$ , and the slope in one form becomes the inverse slope in the other, with the uncertainty in the slope is related to the uncertainty in the inverse slope by

$$\sigma_m = m^2 \sigma_{1/m}. \quad (5)$$

You can derive this relationship from basic error propagation formulas. See Bevington and Robinson, Chapter 4 or “Notes on data analysis and experimental uncertainty” on the course website for more information.

**Case 4: The uncertainty is different in each data point and the x and y uncertainties are of similar magnitude.** As in Case 3, the goal is to transform the data so that all you have are y-error bars, and then to follow the plan of Case 2. However, here you cannot simply choose the other variable to be independent, you need to work on the point uncertainties directly. If you have a reasonably good idea of what functional form the data follow, you can use an approximation to make the uncertainty in the x direction act like an uncertainty in the y direction. Assume that the data do follow a line and that the slope of the line is given approximately by  $M$ .

Since we are assuming that  $y$  is a function of  $x$ , i.e.,  $y(x)$ , any small uncertainty in the data point  $x_i$ ,  $\sigma_{x_i}$  could give rise to an uncertainty in the associated  $y$  value,  $\sigma_{y(x_i)}$ , by an amount

$$\sigma_{y(x_i)} = \left| \frac{dy}{dx} \right| \sigma_{x_i} = M \sigma_{x_i} . \quad (6)$$

Then, under the assumption that the error bars in the x and y directions are independent of each other, we can add this *transformed* uncertainty to the original uncertainty in the y direction by using the sum-of-the-squares rule (also known as “adding in quadrature”). The new y-error bar is related to the old x- and y-error bars by the equation

$$\sigma_{y_i} = \sqrt{\sigma_{y_i}^2 + \sigma_{y(x_i)}^2} = \sqrt{\sigma_{y_i}^2 + M^2 \sigma_{x_i}^2} . \quad (7)$$

After you have recast the x-error bars as effective y-error bars, you can then do your fit. (you could also use this method in Case 3.)

## Present and discuss your final result

You must make an attempt to calculate the uncertainty in your result for  $c$  from by following one of the cases above; failure to do so will result in a significant grade penalty, since it is, in large part, the point of the experiment. After you have calculated your final result for  $c$  with its uncertainty, make sure you express it with correct significant figures (and units). Uncertainties should have at most 2 significant figures and the associated result should be written with the same precision, that is, the same number of figures past the decimal point. See the section on “Significant Digits” in the “Notes on data analysis and experimental uncertainty” on the course website for more information.

After you have your final result, compare it to the defined value of  $2.998 \times 10^8$  m/s. Does it overlap this value within your (random) uncertainty? If it does, your measurement is successful to within the precision of your data. If there is a disagreement of more than three times the uncertainty (“three-sigma”), there is a definite problem to resolve; you have either made a mistake or there is a source of systematic error. (According to statistical theory, there is only about a 0.3% chance that a three-sigma difference could be due to merely random variations in the data set.) You may want to look through your data set and analysis again to see if you can spot any likely bad data points or mistakes in calculation. If the difference is between one and three sigma, then there is less certainty about the quality of your result. You may have a systematic error, or you may have underestimated your random uncertainty. If possible, you should try to double-check some data points by having different people measure them.

Your report should contain a discussion of the result that focuses on how the uncertainty was found and what might be the cause of any disagreement between your result and the defined value. If you are unable to determine the source(s) of systematic error that may be in your result, you should propose hypotheses

about possible causes and what measurements or further investigations you would carry out, were you able to do so.

In addition, your report should conform to the requirements in the document “Lab precautions, recording data and writing reports” on the course website.

## Questions

---

Give complete answers to the following questions in your report.

1. When light propagates through media other than a vacuum, the speed is reduced by the refractive index  $n$  of the medium:  $c' = c/n$ . The refractive index of air is 1.00029 at standard temperature ( $0^\circ$  C) and pressure (1 atmosphere). Is this experiment, in the manner that you have done it, able to measure the effect of air on the speed of light? Explain, using a quantitative argument based on the uncertainty in your result.
2. Assume that you are given a 1 meter rod of clear plastic, whose refractive index is claimed to be about 1.5. You are asked to measure the refractive index of the plastic both to confirm this claim and to get a better value for it. Explain how you could measure the refractive index with your setup without cutting the rod. (You may assume that the ends of the rod are flat and polished.)
3. A fellow student proposes to run the experiment without using a mirror, but to instead point the laser beam directly at the detector and vary its position from the laser to get different values of  $L$ . Explain how to make this approach give good results. What would have to be different in the setup? What affect would the change(s) have on the scope measurements?

## References

---

Bevington, Philip R. and Robinson, D. Keith, *Data Reduction and Error Analysis for the Physical Sciences*, 3<sup>rd</sup> edition, McGraw-Hill Higher Education (2003). ISBN 0-07-247227-8. [www.mhhe.com](http://www.mhhe.com)

Aoki, K, and Mitsui, T., “A tabletop experiment for the direct measurement of the speed of light,” *Am. J. Phys.* **76**, 812-815 (2008). This paper describes the apparatus that we adapted for our setup.

“Notes on data analysis and experimental uncertainty” by David Pengra and L. Thomas Dillman at [http://courses.washington.edu/phys431/uncertainty\\_notes.pdf](http://courses.washington.edu/phys431/uncertainty_notes.pdf)

“Lab precautions, recording data and writing reports” on the physics 331 course website at [http://courses.washington.edu/phys331/lab\\_practice\\_and\\_report.pdf](http://courses.washington.edu/phys331/lab_practice_and_report.pdf)

“Notes on making a least-squares fit to a line in Microsoft Excel” on the Physics 331 course website at <http://courses.washington.edu/phys431/linestExcel.pdf>

New\_SoL\_experiment.doc  
D. B. Pengra - 8 September 2011