880 Second Harmonic Generation Lab

The goal of this lab is to give you a chance to observe Second Harmonic Generation (SHG) and to use it to measure a short laser pulse.

You may not take this lab unless you have read the laser safety manual for my laboratory. A link to this manual is on the class web page.

Introduction.

This lab uses a KD^*P crystal to double the frequency of the light from a short pulse laser under varying conditions. The short pulse laser system is based on Ti:Sapphire. See the appendix at the end of this write-up for a short description of this system. The output of the laser system is as follows: pulse width 60fs, pulse energy 0.3 to 1 mJ, repetition rate 1 kHz, center wavelength 800 nm, polarization horizontal. The pulse energy cannot be varied for this lab and the intensity is more than high enough to cause blindness. The light cannot start a fire or burn unless focused, however.

For this experiment, you will use silver coated mirrors that have a dielectric coating over the silver. These mirrors reflect well over a large bandwidth. In general, all optics used with short pulse lasers must operate over a large bandwidth and must not introduce significant phase errors or the pulse will broaden.

As always, do as much of this laboratory as you like. A number of options are suggested.

SHG.

Your first goal is to double the pump light using a KD^*P crystal. The laser is sufficiently intense that you do not need to focus it into the doubling crystal. In fact, you will damage the crystal if you do. (An interesting detail: Short pulses may not remain short if focused by a single element lens. Critical applications use achromatic lenses or reflective optics.)

A thin and a thick KD^*P crystal oriented for type I doubling, a zero-order half wave plate, a Glan-Thompson polarizer, and several silver mirrors are provided. You will also have a filter that blocks the fundamental, but transmits the second harmonic. The SHG process is dramatic enough using the short pulse laser that you can see the second harmonic even when the fundamental is left in the beam. However, to get started you should place the filter after the crystal to block the fundamental and place a white card downstream from the filter. Turning the room lights off also helps to see a weak signal. The ordinary axis is marked on the crystal housing by a long straight line across the entrance and exit faces. Start with the thin crystal and optimize the SHG until it can easily be seen by eye.

Some things to try:

- Optimize the SHG by using the compressor to adjust the pulse chirp.

- Send the light into a monochometer. The monochrometers in my lab use line cameras as detectors and are configured to display the spectrum on an oscilloscope. A 90 nm wavelength range centered on the wavelength selected by the micrometer adjust on the monochrometer is displayed. Explain the spectrum's dependence on the incident angle of the beam into the doubling crystal.

- Replace the thin crystal with the thick one, look with the monochometer, and explain the differences.

Obtain A Measure Of The Pulse Width.

The Autocorrelation Function

You will do this by constructing a single shot autocorrelator. For weak pulses, an autocorrelation must be built up using many pulses. Our system is bright enough to obtain a measurement on every laser shot. Let's begin by seeing how to optically generate the temporal autocorrelation function of the laser pulse and determining how this relates to the pulse width.

Suppose we split our short pulse into two identical pulses and mix them in a doubling crystal. What we get will depend on the temporal overlap between them. The nonlinear polarization will look something like (ignoring spatial overlap, polarization, d_{ijk}, and various constants):

$$P^{2\omega}(t,\tau) = E^{\omega}(t)E^{\omega}(t-\tau)$$

where one of the pulses is delayed in time by τ with respect to the other. P(t,τ) and E(t) are slowly varying envelope functions. "Slow" means relative to the optical period. The envelope time variation is at least an order of magnitude faster than available detector electronics. The factors representing the carrier frequency, eg. $e^{i\omega t}$, have cancelled out as have the phase matching factors. For the short pulse laser system, E(t) is closely approximated by a Gaussian. $P^{2\omega}$ will give rise to a field $E^{2\omega}$, and so squaring both sides, we can write:

$$I^{2\omega}(t,\tau) = I^{\omega}(t)I^{\omega}(t-\tau)$$

All of the intensities vary on a ~100 fs time scale, but if we measure the output using a power meter it will simply integrate the energy over time, leaving us with a function only of delay:

$$S(\tau) \equiv \int_{-\infty}^{\infty} I^{2\omega}(t,\tau)\,dt = \int_{-\infty}^{\infty} I^{\omega}(t)I^{\omega}(t-\tau)\,dt$$

S(τ) is the autocorrelation function of the pulse intensity. Although I(t) varies rapidly, S(τ) varies only as rapidly as we choose to vary the delay. In this way the fast optical time scale is mapped onto a slow, easily measured, laboratory time scale. If τ is large and negative then I(t-τ) is early and S = 0 due to no pulse overlap and likewise if τ is large and positive. Only when τ is close to zero will there be signal. Although I(t) cannot be determined given S(τ), limits can be placed on the form of I(t). If we assume a Gaussian pulse shape for I(t), and this has been verified using a more sophisticated measurement techniue call Frequency Resolved Optical Gating (FROG), then the pulse width can be determined from S(τ). It turns out that S(τ) is also a Gaussian function in this case and the pulse width is related to the autocorrelation width as:

$$\tau_{I,FWHM} = \tau_{S,FWHM} / \sqrt{2}$$

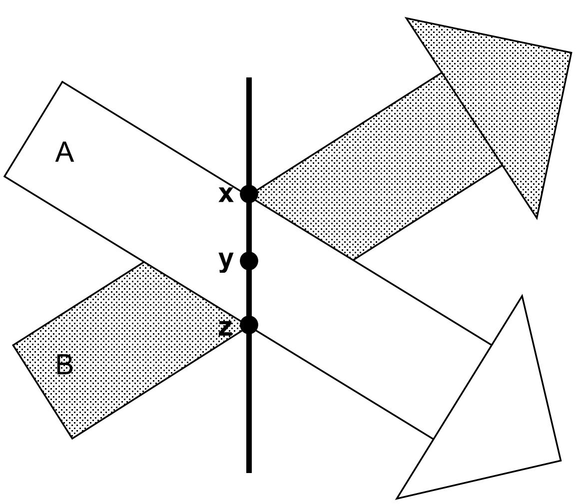

Originally, S(τ) was measured by slowing varying the delay between the two pulse and measuring over many laser shots. There is a way to do it in a single shot using a much simpler apparatus than that used for multi-shot measurements. Suppose the two beams, A and B, cross at some angle as shown in the figure. Consider the three points x,y,z labeled on the vertical line drawn through the region where the beams overlap. Beam A will reach point x before B. It will reach point y at the same time as B and it will reach point z after B. In principle, then, all delays needed can be present at the same time. An autocorrelation is obtained by placing a doubling crystal in the same location as the vertical line in the figure and measuring the *spatial* variation in the SHG signal. Since the delay varies with position in space, now the rapid time variation of the pulse has been mapped onto slowly varing function of position.

Build The Autocorrelator

Use the zero-order half wave plate to make the short pulse laser vertically polarized. We will employ Type I phase matching. Now set up the following:

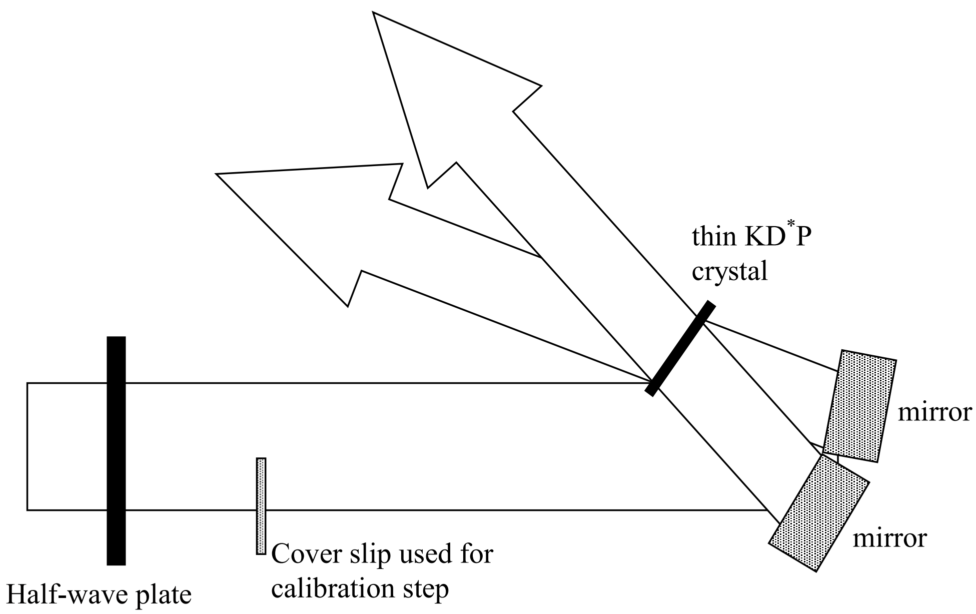

The beam is split into two beams by the two mirrors which then overlap. You may wish to enlarge the beam. Ideally this would be done with a telescope, but for purposes of this lab a long focal length lens placed upstream of the apparatus will suffice. Try f = -1000 mm or -2000 mm.

I have not given any indication of where the second harmonic appears or how phase matching works in this geometry. I will be happy to help you with all aspects of this set-up, but try and see if you can think this through first. First, assume you have phase matching and use momentum conservation (or symmetry for that matter) to predict where the SHG from both beams will appear. Now find the phase matching condition of each beam acting alone. Finally, find the phase matching condition that will combine one photon from each beam to generate the autocorrelation.

Calibrate The Autocorrelator

Place a thin microscope cover slip through half of the beam as indicated in the figure. You will see the autocorrelation image move to one side. The delay introduced by the microscope slide is (n – 1)t/c where n and t are the index of refraction and thickness of the slide respectively. If the distance moved by the image is d, then the delay time per unit distance along the autocorrelation image is (n-1)t/cd. Use n = 1.6 and measure the thickness of the microscope slip with a micrometer. Use this calibration and the spatial width of the autocorrelation image to estimate the pulse width.

- If you like, have someone change the pulse chirp by adjusting the compressor while you observe the autocorrelation image.
- This set-up can also be made to work using horizontal polarization with appropriate changes to the KD*P crystal. The operation is somewhat peculiar, however. If you wish to try this, *see me first*, since there is an eye hazard from reflected beams off the doubling crystal in this case.

Appendix: The Short Pulse Laser System

The short pulse laser system is based on the Ti:Sapphire gain medium, which consists of a sapphire crystal (Al_2O_3) doped with titanium. (Ruby is sapphire doped with chromium and Sapphire jewelry is sapphire doped with a little titanium [blue] or iron [green, yellow] or some combination.) In the following, refer to the figure below.

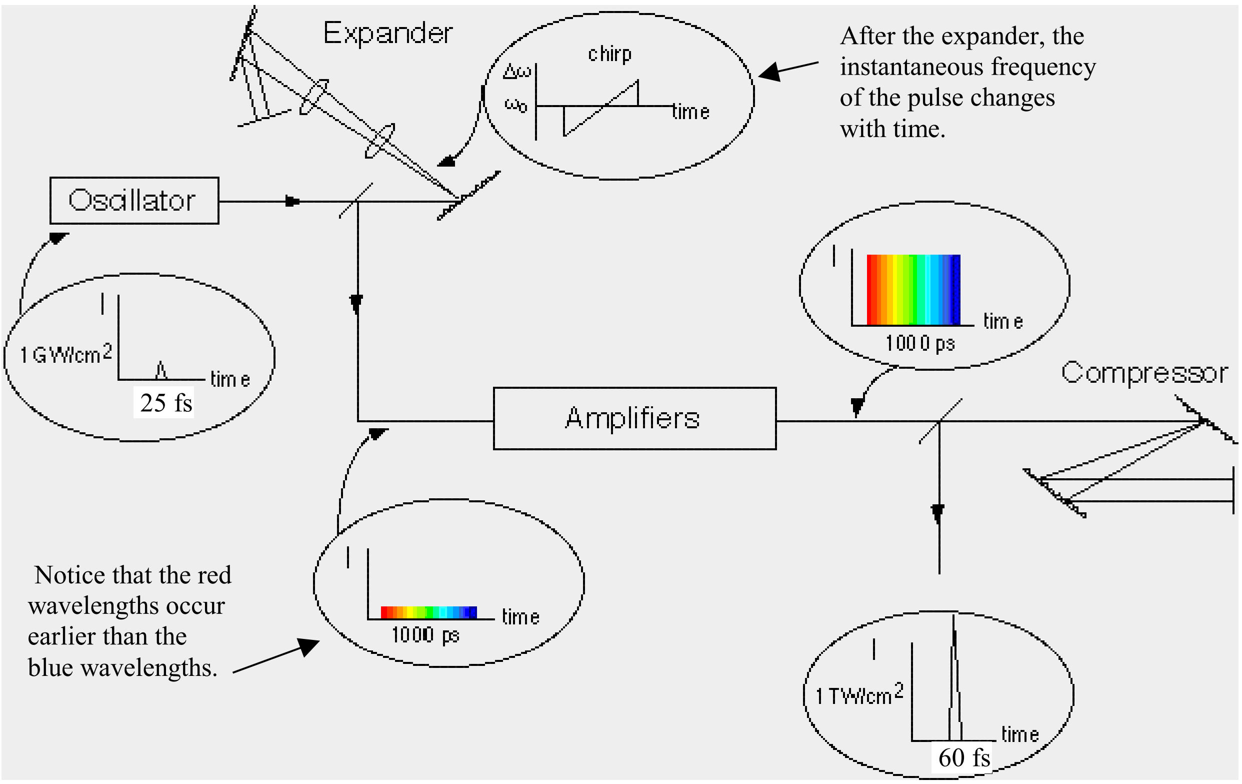

Schematic figure of Ti:Sapphire laser system (adopted from http://www.physics.utoronto.ca/~marj/CPAgen.html).

The system begins with a Kerr Lens Modelocked laser ("oscillator" in the figure). As I discussed in the first lecture, this laser modelocks via non-linear self-focusing. The non-linear process involved is called the Kerr effect which I hope to cover before class ends. The gain medium is Ti:Sapphire and the oscillator is pumped cw using 4.6 W of 532 nm light from an intracavity doubled Nd:YVO_4 laser. The output specifications are: 800 nm center wavelength, 25 fs pulse width, >1 nJ/pulse.

This light cannot be directly amplified. Well before the target pulse energy of 1 mJ was reached, the intensity would be sufficient to drive various non-linear processes that would destroy the light or even damage the amplifier (see the web sites listed under the lab write-up on the class web page). This problem is solved using Chirped Pulse Amplification (CPA). In CPA, the pulse is broadened so that the peak power is kept to a safe value, amplified, then compressed back to a short pulse.

The pulse is first broadened to roughly 50 ps by a pair of so-called anti-parallel gratings ("expander" in the figure). In the expander, the light is dispersed by the lower grating and relayed to the upper grating by a one-to-one telescope. After bouncing off the second grating, the different wavelengths in the beam are traveling parallel to each other, but spatially separated. A mirror retro-reflects the light and the operation is reversed so that after bouncing off the lower grating again, for a total of four grating bounces, the light has been recombined. A slight misalignment of the retro-reflector separates the beam leaving the expander from the incoming beam allowing the output beam to be picked off by another mirror (shown in the figure as a diagonal line). Although this is not obvious from the foregoing discussion, the path taken by red wavelengths through this system is shorter than that taken by blue wavelengths. Effectively, redder lights travels through the system more quickly than bluer light. This is the same situation that occurs when two different wavelengths propagate through a transparent optical medium. The redder wavelengths have a larger phase (and group) velocity than the bluer wavelengths. Put another way, the expander effectively introduces positive group velocity dispersion (GVD). However, unlike an optical medium, the dispersion it introduces is readily adjustable and can be made very large using only a small amount of table space. After passing through the expander the pulse still has sufficient bandwidth to form a short pulse, but the phases have been altered so that the pulse is 1000 times longer than it was. The intensity inside the amplifier will now be 1000 times less and amplification can safely occur. Finally, the instantaneous frequency of the pulse now varies almost linearly with time. It is low (red) at first and then increases. This is called a "chirp."

The chirped pulse is amplified in a regenerative amplifier in the manner discussed last quarter. A Pockels cell is used to selected a single pulse from the 90 MHz pulse train coming from the oscillator and injects it into a stable laser cavity. The cavity contains another Ti:Sapphire crystal which has been pumped by a pulsed, intra-cavity doubled Nd:YLF laser. This is the same Nd:YLF laser you used in the previous labs. Left to itself, the cavity would lase, but the same Pockels cell that injects the seed pulse also acts as a Q-switch preventing the cavity from lasing. Instead, when the pulse is injected the Q is simultaneously raised to a large level and the pulse is amplified by a factor of roughly 10^6 in 100 to 200 ns. A second Pockels cell is used to switch the pulse out of the cavity.

The pulse is then compressed back to a short pulse by a pair of parallel gratings ("compressor" in the figure). Unlike the expander, the compressor is straightforward to understand. A little thought should convince you that redder wavelengths travel a *longer* path through the compressor than bluer wavelengths, so the compressor effectively introduces *negative* group velocity dispersion. The distance between the gratings is tuned to compensate the positive GVD introduced by the expander as well as the GVD accumulated from passing through the various optical media in the laser system. Because of gain narrowing in the regenerative amplifier the pulse width after pulse compression is 60 fs instead of the 25 fs we started with. This can be corrected, again as I discussed last quarter, but shorter pulses are actually not desired in my laboratory currently.

A detail: Although schematically the above description is correct, the short pulse laser system differs in some important ways. For example, the expander actually uses reflective optics instead of lenses. Also, the expander only uses a single grating. An extra mirror is added to get the requisite four bounces.