Swamp Optics, LLC.
6300 Powers Ferry Rd.
Suite 600-345
Atlanta, GA 30339
+1.404.547.9267
www.swamoptics.com

Swamp Optics Tutorial

Pulse Compression

Recall that different colors propagate at different velocities so that, after passing through the material, different colors experience different delays, a ubiquitous effect called *group-delay dispersion (GDD)*. For wavelengths less than those used in telecommunications (about 1.5 μm), red colors always propagate faster than blue colors and so precede the blue wavelengths in the pulse, and one says that the pulse is *positively chirped*. Chirped pulses are nearly always undesirable in applications, and their increased length and resulting lower intensity make them even less desirable. Because ultrashort pulses have broad spectra, GDD is a large effect for them and generally results in significant pulse lengthening after propagation through even a few mm of material. Worse, the effects of GDD accumulate, and the pulse lengthens more and more and becomes more and more positively chirped as it passes through additional material.

It is therefore important to minimize the amount of material through which a pulse propagates. Indeed, when the pulse is in the UV spectral region, even propagation through *air* can lengthen it significantly. The problem is that, in essentially all applications, the pulse is focused and so must pass through a significant amount of glass *in the lens or microscope objective*, and this glass cannot be avoided. Simply passing through a microscope objective significantly lengthens the pulse, significantly decreasing the intensity of the pulse at the focus. This reduces the sensitivity in all multi-photon microscopy techniques, and it can also hurt the resolution, too. And it is quite deleterious for most other applications for the same reason. Indeed, one of the world's largest laser companies intentionally sells longer-pulse lasers to the imaging community with less bandwidth to the effect of GDD. Of course, because the pulses are longer to begin with, images made with such pulses have less sensitivity. Unfortunately, while laser companies could sell lasers with negative chirp, they don't in practice because each user has a different amount of GDD in his microscope objective and in the beam path from the laser to the sample.

Fortunately, this problem can be solved. Ultrafast labs have for years used devices called *pulse compressors* (see Fig. 1),[1-3] which consist of sequences of four prisms that yield a path for which red colors travel a longer optical path than blue colors, and so have *negative GDD*. Specifically, the trick is that the redder colors pass through the thicker regions of the second and third prisms, allowing the bluer colors to catch up with them.

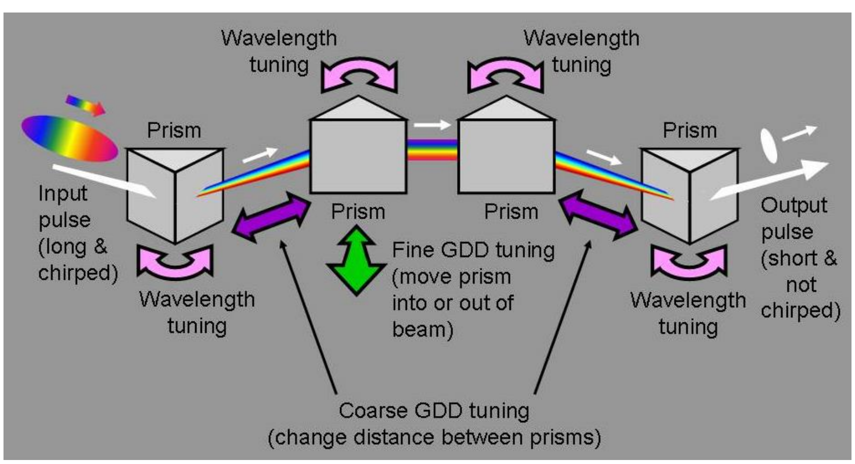

Fig. 1. Four-prism pulse compressor. Longer wavelengths (red) traverse more glass than shorter ones (blue), passing through the thicker parts of the second and third prisms. A problem, however, is that, when the average wavelength of the pulse changes, all four prisms must be rotated by precisely the same amount (thick pink arrows). Also, coarse and fine tuning of the GDD are separate. And coarse-tuning the GDD requires changing the separations between the first two and last two prisms, maintaining them precisely equal (thick purple arrows).

Unfortunately, the pulse compressor is very unwieldy. To vary the GDD slightly, a prism is moved into or out of the beam (the green arrow). But to vary the GDD over a larger range, the separations between the first and second prisms and the third and fourth prisms must be varied (the purple arrows)—and maintained precisely equal to each other—which involves several alignment parameters (knobs to turn). Also, if the input wavelength changes, all four prisms must be rotated by precisely the same amount (the pink arrows), maintaining equal incidence angles for all the prisms. Otherwise the device will involve magnification or demagnification of the beam in one direction (yielding an undesirable elliptical output beam) as well as angular dispersion and the spatio-temporal distortions, spatial chirp and pulse-front tilt. Worse, to obtain the desired amounts of negative GDD, the prism separations must be quite large, and the device can become quite large (a meter or more—larger than the laser itself!). Finally, the device is tricky to align to begin with, and small misalignments can cause the above-mentioned deleterious spatio-temporal distortions in the pulse, which can be difficult to diagnose.

Fortunately, it was realized quickly that the device could be simplified to two prisms using a mirror (or better, a roof mirror) after the second one to reflect the beam back through the first two prisms. Two-prism designs (see Fig. 2) simplify the device significantly, but it remains complex, and it requires a very complex mechanical apparatus to maintain the required equal incidence angles of the two prisms, and it is difficult to do so sufficiently to avoid the various distortions mentioned above. Indeed, Swamp Optics' GRENOUILLEs and FROGs measure these distortions (in addition to the pulse intensity and phase), and we usually find significant amounts of these distortions in pulses that have emerged from pulse compressors.

Fig. 2. Two-prism pulse compressor. This device avoids the need for maintaining equal prism separations (note only one thick purple arrow). Unfortunately, when the average wavelength of the input pulse changes, both prisms must still be rotated by precisely the same amount (thick curved pink arrows). And this design still requires a complex arrangement for coarse and fine-tuning the GDD.

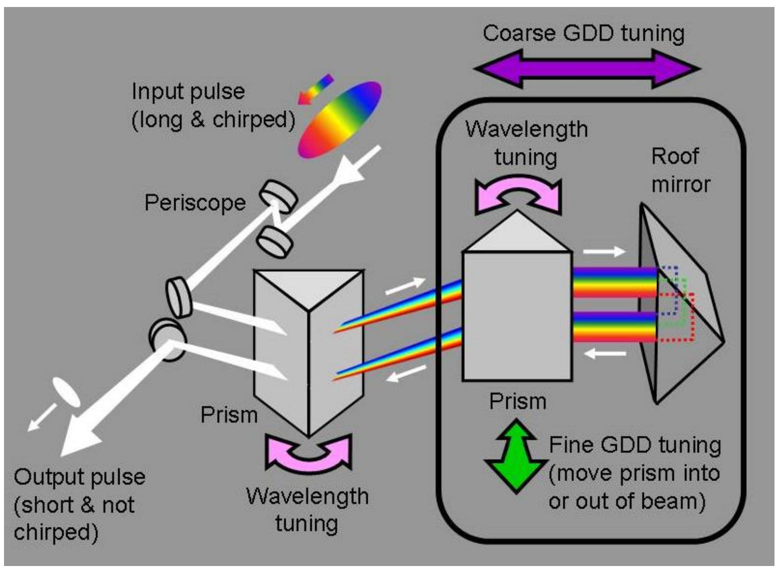

To solve these problems, Swamp Optics has introduced a (patent-pending) *single-prism pulse compressor* (see Fig. 3), which is much simpler, and, even better, it also solves all of the problems described above.[4] It uses a *corner cube* and inverts the beam—not the prism—and so only requires one prism. Indeed, the corner cube inverts the beam between the first and second passes through the prism and also between the third and fourth passes through the prism, precisely as required by a pulse compressor!

It only requires varying the prism angle to tune in wavelength (no complex mechanical apparatus needed). Also, all of the above-mentioned distortions automatically cancel out—even as it is tuned in wavelength. And because it involves double-passing the corner cube, all beam magnifications and dispersions cancel out. In addition, even if the device is bumped, the output beam does not misalign! And because each path between the prism and corner cube is double-passed, the device is only *half* as big as two prism designs and one fourth as big as four-prism designs. In all other respects, it is identical to conventional pulse compressors and so shares their advantages; for example, all passes through the prism occur near the prism tip to avoid introducing unnecessary additional positive GDD from the prism itself.

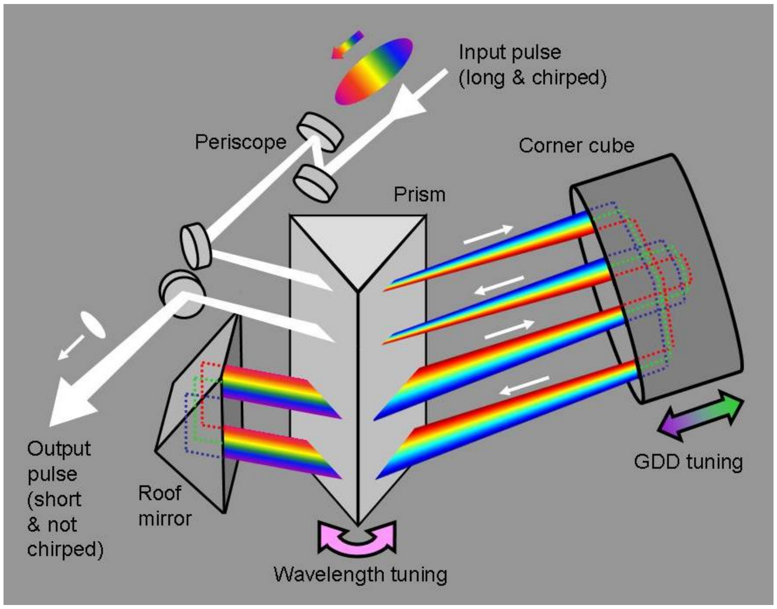

Fig. 3. Single-prism pulse compressor. The corner cube inverts the beam so that the second and third passes through the prism occur as if the prism is inverted, as required for a pulse compressor. Also, corner cubes are manufactured extremely precisely, so the beam emerging from it is precisely parallel to that entering it. This and the use of a single prism are the keys to avoiding all of the distortions that two- and four-prism designs introduce into the pulse. Also note that only the one prism angle need be rotated when the input wavelength changes, and only one distance need be changed to tune the GDD. This simplicity yields a much more accurate, much less distorting, and much less expensive device.

Also, it yields no deviation in the beam and so can be inserted into the beam and removed without misaligning the apparatus, and, unlike previous compressors, this lack of beam deviation will be maintained despite angle-tuning the prism for other wavelengths. Finally, it achieves significant amounts of negative GDD, enough to compensate for even a many-element microscope objective. These features make it ideal for multi-photon imaging systems and essentially all ultrafast-optical applications. Indeed, it requires no alignment, and it has only two knobs, one for the GDD and another for the pulse center wavelength.

Here's more detail on these characteristics:

Magnification: As the wavelength is tuned by rotation of a prism, each pass through the prism necessarily involves unequal incidence and exit angles. While the one-dimensional magnification through a prism is unity when its incidence and exit angles are equal (i.e., Brewster's angle for an ideal pulse compressor), this is not generally the case. When a prism's incidence angle exceeds its exit angle, some one-dimensional magnification occurs, and when the opposite occurs, demagnification occurs. Thus, away from the ideal equal-incidence-and-exit-angle case, each pass through a prism will always yield some one-dimensional magnification or demagnification in the beam. This problem exists in all prism and grating pulse compressors, and, in conventional compressors, it must be addressed in the alignment procedure and unity magnification must be maintained as well as possible by a complex mechanical apparatus when it is tuned.

In our single-prism design, as the wavelength is tuned, the first pass through the prism will introduce, say, some magnification. However, after the beam returns from the corner cube, the beam will retrace its path and experience demagnification by exactly the same amount because the input/output angles are reversed as compared with the first pass. As a result, if M_i is the magnification after the ith pass through the prism, we will have:

$$M_1 = \frac{1}{M_2} = M_3 = \frac{1}{M_4}$$

Multiplying all four magnifications together yields the overall magnification of the device, which will always be unity. As a result, our single-prism pulse compressor automatically retains unity magnification as the wavelength is tuned! We measured the magnification of the device by measuring the input and output spatial profiles of the beam on a CCD camera. Our measured magnification value was 1.01±.02.

Angular dispersion: It is even more important that the output angular dispersion remain zero as the GDD and/or the wavelength is tuned. Because translating the corner cube to vary the GDD will not affect any device angles, corner-cube translation (GDD tuning) need not be considered—if the dispersion is zero for one value of GDD, it will be zero for them all. But the variation of the prism angle with wavelength-tuning could, in principle, cause the output angular dispersion to be nonzero, as it easily can in two- and four-prism pulse compressors. The angular dispersion introduced by a prism depends on which direction the beam propagates through it. It is easy to show that, if the ith prism has dispersion D_i and magnification M in the forward direction (and $1/M$ in the backward direction), it has dispersion MD in the reverse direction. Thus, the angular dispersion added on each pass through the prism is:

$$D_1 = -M_2 D_2 = -D_3 = M_4 D_4$$

where the minus signs take into account the beam flips.

The total dispersion of a four-prism sequence with non-Brewster incidence angles is described by a complex formula in terms of the prism apex and incidence angles, and so it is difficult to see immediately what the total dispersion of the device would be. However, a little-known simple result [5] gives the total dispersion of an arbitrary sequence of prisms in terms of only the prisms' dispersions and magnifications: the total dispersion is the sum of the individual dispersions, each weighted by the reciprocal of the total magnification that follows it. So the angular dispersion for an arbitrary sequence of four prisms and hence the dispersion at the output of the single-prism pulse compressor will be:

$$D_{tot} = \frac{D_1}{M_2 M_3 M_4} + \frac{D_2}{M_3 M_4} + \frac{D_3}{M_4} + D_4$$

Substituting the values for the dispersion and magnification given earlier, it's easy to see that the total device dispersion is necessarily identically zero, independent of the prism angle!

Spatial chirp and pulse-front tilt: The symmetry of the device and lack of angular dispersion after two and four prism passes imply that the spatial chirp is also identically zero. This requires proper alignment of the roof mirror pair, but this is a relatively easy component to align, and it need only be aligned once. Also, when the spatial chirp and angular dispersion are zero, the pulse-front tilt must also be zero. We measured the spatial chirp and pulse-front tilt at the output of the compressor using GRENOUILLE[6, 7], and we found the pulse-front tilt angle to be less than 16 μrad. We can compare this value to that of a two- or four-prism design, in which, typically, a ~ 1 degree misalignment in a prism (which is likely in practice) causes a pulse-front tilt angle of >2000 μrad [8].

Throughput: As with any pulse compressor, we must also consider the throughput as the wavelength is tuned. And in this case, our pulse compressor's properties are identical to those of two- and four-prism compressors (except for a few extra reflections), so there's no need to repeat them here.

Alignment: It is well known that the conventional pulse compressor is an alignment night mare that "burns out" grad students. But it is often claimed (usually by those who don't have to do the alignment themselves!) that a pulse compressor used for only one wavelength need only be aligned once and so can be set up without too much grief. However, keep in mind that, in conventional two- and four-prism designs, each prism can only be aligned by using the minimum-deviation condition: that the beam deviation will be minimal when the incidence and output angles are equal. However, this is a very poor alignment technique. This is because, at the minimum-deviation condition, the beam deviation angle depends on the prism incidence angle only to second order (it's a minimum, after all!) and so is *not* a sensitive indicator of the correct incidence angle. On the other hand, the important quantity, the angular dispersion has a dependence on input angle that remains first order—much stronger. Thus, minimum deviation is a poor alignment technique. Unfortunately, there is no other method for aligning a conventional pulse compressor. Worse, even if it is somehow perfectly aligned on one day, the next day, when the input beam has walked a bit in angle, the pulse compressor will no longer be aligned and so must then be re-aligned. So it's no surprise that conventional pulse compressors burn out grad students. At Swamp optics, we were once grad students, some of us very recently, and we are designing products that make this experience easier. We think our pulse compressor, which requires no alignment of its prism and which need not be re-aligned when the input beam moves, will help greatly, both in eliminating the alignment time and also yielding a distortion-free compressed pulse.

In conclusion, the single-prism pulse compressor has all of the advantages of conventional pulse compressors, and none of their disadvantages. It solves essentially all of their problems, yielding a compact, alignment-free, easy-to-work-with device that is also inexpensive—about half the price of other commercially available pulse compressors.

1. Bor, Z. and B. Racz, *Group Velocity Dispersion in Prisms and Its Application to Pulse Compression and Traveling-Wave Excitation.* OC, 1985. **54**(3): p. 165-170.
2. Rudolph, W. and B. Wilhelmi, *Light Pulse Compression*. 1989, Chur: Harwood Academic Publishers.
3. Johnson, A.M. and C.V. Shank, *Pulse Compression in Single-Mode Fibers—Picoseconds to Femtoseconds*, in *The Supercontinuum Laser Source*, R.R. Alfano, Editor. 1989, Springer-Verlag: Berlin. p. 399-450.
4. Akturk, S., et al., *Extremely simple single-prism ultrashort-pulse compressor.* Opt. Expr., 2006. **14**(21): p. 10101-8.
5. Trebino, R., *Achromatic N-prism beam expanders: optimal configurations.* Appl. Opt., 1985. **24**(8): p. 1130-1138.
6. Akturk, S., et al., *Measuring pulse-front tilt in ultrashort pulses using GRENOUILLE.* Optics Express, 2003. **11**: p. 491-501.
7. Akturk, S., et al., *Measuring spatial chirp in ultrashort pulses using single-shot Frequency-Resolved Optical Gating.* Optics Express, 2003. **11**: p. 68-78.
8. Akturk, S., et al., *The general theory of first-order spatio-temporal distortions of Gaussian pulses and beams.* Optics Express, 2005. **13**(21): p. 8642 - 8661.

About Swamp Optics

Founded in 2001, Swamp Optics, LLC offers cost-effective quality devices to measure ultrashort laser pulses. It specializes in frequency-resolved optical gating (FROG) and GRENOUILLE (an experimentally simple version of FROG), the gold standards for measuring the time-dependent (or, equivalently, frequency-dependent) intensity and phase of an ultrashort pulse.

Swamp Optics also sells an innovative pulse compressor.

For more information, visit us on the Web at www.swampoptics.com.