The Particle Melting Pot

Welcome back for our second guided walk into the quantum mechanical wood! Last week, we saw how particles move like waves and hit like particles, and how a single particle takes multiple paths. While surprising, this is a well explored area of quantum mechanics—it is on the paved nature path around the visitor's center.

This week I'd like to get off the paved trail and go a bit deeper into the woods, to talk about how particles meld and combine while in motion. This is a topic that is usually reserved for physics majors, and is rarely discussed in popular articles. But the payoff is understanding how precision lidar works and getting to see one of the great inventions making it out of the lab, the optical comb. So let's go get our hiking boots a little dirty.

Two particles

Let's start with the questions: if particles move like waves, what happens when I overlap the paths of two particles? Or said another way, do particle waves only interact with themselves, or do they mix together?

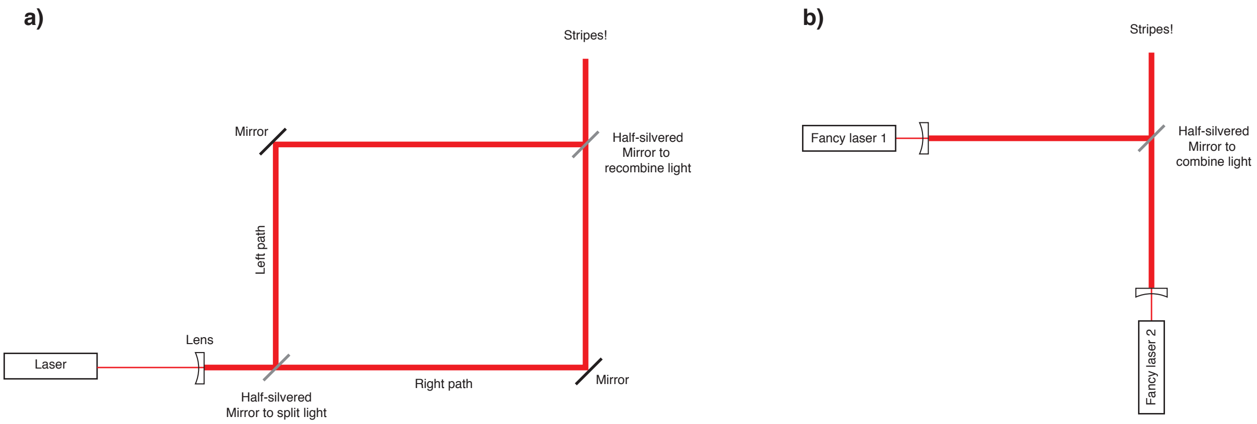

On the left is the interferometer setup from last week, where the light and photons from laser is split into two paths which are then recombined to show the interference stripes. The right shows a simplified setup where the light from two separate (high quality) lasers is combined together.

We can test this in the lab by modifying the setup we used last week. Instead of splitting the light from one laser into two paths, we can use two separate lasers to create the light coming into the final half-silvered mirror.

We need to be careful about the lasers we use, and the quality of your laser pointer is no longer up to the task. If you carefully measure the light from a normal laser the color of the light and the phase of the wave (when the wave peaks occur) wander around. This color wander is not discernible to our eyes—the laser still looks red—but it turns out that the exact shade of red varies. This is a problem money and modern technology can fix,—if we shell out enough cash, we can buy precision mode-locked lasers. Thanks to these, we can have two lasers both emitting photons of the same color with time-aligned wave crests.

When we combine the light from two high quality lasers we see exactly the same stripey pattern that we saw before. The waves of particles produced by two different lasers are interacting!

So what happens if we again go to the single photon limit? We can turn the intensity of the two lasers down so low that we see the photons appear one at a time on the screen, like little paintballs. If the rate is sufficiently low, only one photon will exist between the lasers and the screen at a time. When we perform this experiment we will see the photons arrive at the screen one at a time; but when we look at the accumulated pointillism painting, we will see the same pointillism stripes we saw last week. Once again, we're seeing particle interference.

It turns out that all the experiments we performed before give exactly the same answer. Nature does not care if one particle is interacting with itself or if two particles are interacting with each other—a wave is a wave and particle waves acts just like any other wave.

But now that we have two precision lasers, we have a number of new experiments we can try.

[Sidebar: Mode locking **Q: How do the mode locked lasers emit photons at the same time so that they'll beat together?** Interestingly there is *no* coordination between the lasers about when photons are emitted. Remember that particles move like waves and, in addition to being broad, waves have a length too. So try not to picture this as ball bearings—that mental image is inaccurate. We'll directly tackle what we mean by the 'length' of a particle wave in next week's article. Patience grasshopper.]

Two colors

First, let's try interfering photons of different color. Let's take the color of one of the lasers and make it slightly more blue (shorter wavelength). When we look at the screen we again see stripes, but now the stripes walk slowly sideways. Both the appearance of stripes and their motion are interesting.

First, the fact that we see stripes indicates that particles of different energy still interact.

The second observation is that the striped pattern is now time dependent; the stripes walk to the side. As we make the difference in color between the lasers larger, the speed of stripes increases. The musicians in the audience will already recognize the beating pattern we are seeing, but before we get to the explanation let's improve our experimental setup.

If we are content to use narrow laser beams, we can use a prism to combine the light streams. A prism is usually used to split a single light beam and send each color in a different direction, but we can use it backwards, and with careful alignment use the prism to combine the light from two lasers into a single beam.

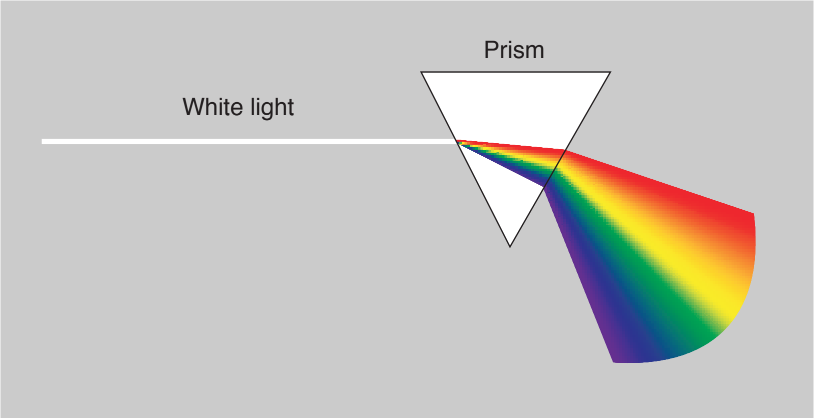

A beam of white light going through a prism to create a rainbow.

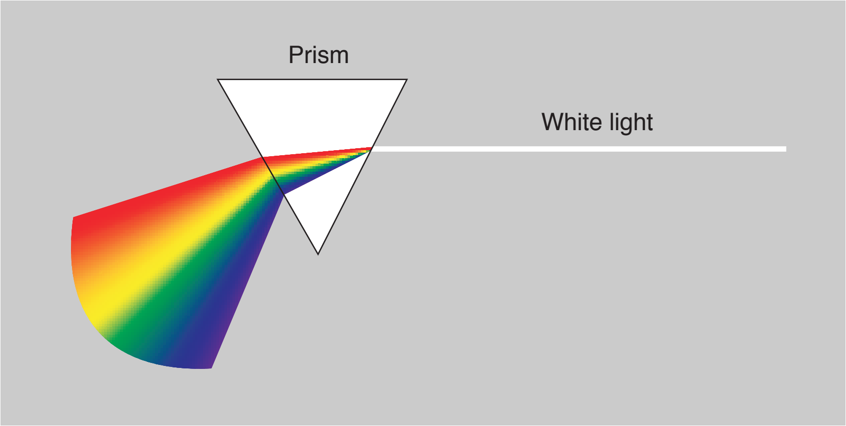

A rainbow of light being sent into a prism (left to right) to create a beam of white light.

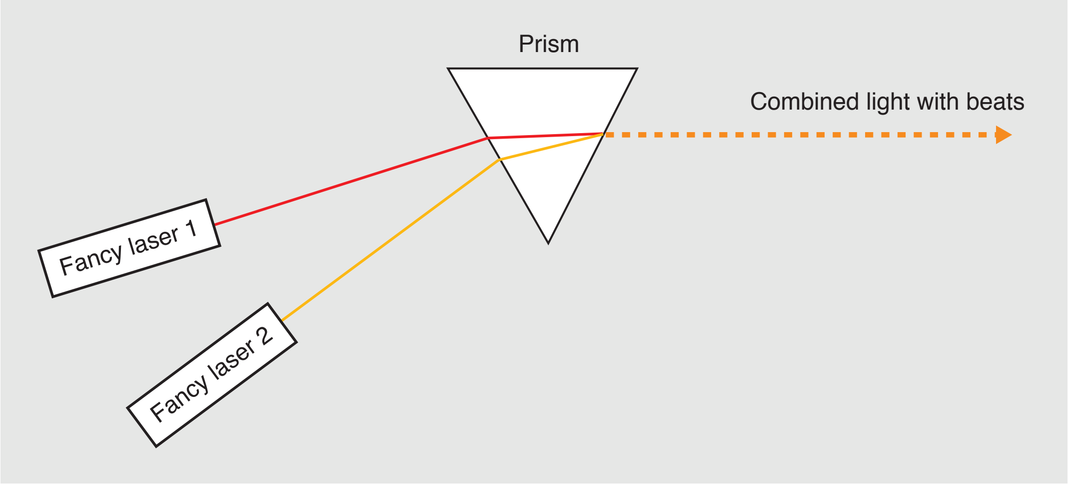

Light from two lasers of different color being combined with a prism. The resulting beam beats from bright to dark just as the sound of two slightly off-tune notes will beat. This can also be seen as stripes in time.

If we look at the intensity of the combined laser beam, we will see the intensity of the light 'beat.' While the light from each laser was steady, when the light from two lasers of slightly different color are combined the resulting beam oscillates from bright to dim. Musicians will recognize this from tuning their instruments. When the sound from a tuning fork is combined with the sound of a slightly out-of-tune string, one can hear the 'beats' as the sound oscillates between loud and soft. The speed of the beats is the difference in the frequencies, and the string is tuned by adjusting the beat speed to zero (zero difference in frequency). Here we are seeing the same thing with light—the beat frequency is the color difference between the lasers.

While this makes sense when thinking about instrument strings, it is rather surprising when thinking of photons. We started with two steady streams of light, but now the light is bunched into times when it is bright and times when it is faint. As the colors of the lasers are made more different (de-tuned), the faster the pulsing becomes.

Paintballs in time

So what happens if we again turn down the lasers really low? Again we see the photons hit our detector one at a time like little paintballs. But if we look carefully at the timing of when the photons arrive, we see that it is not random—they arrive in time with the beats. It does not matter how low we turn the lasers—the photons can be so rare that they only show up one every 100 beats—but they will always arrive in time with the beats.

This pattern is even more interesting if we compare the arrival time of the photons in this experiment with the pattern of stripes we saw with our laser pointer last week.

One way of picturing what is happening in the two-slit experiment is that the wave nature of quantum mechanics is directing where the photons can land side-to-side: the paintballs are allowed to hit in the bright regions and not in the dark regions. We see the same pattern in the paintball arrival in the two-color beam, but now the paintballs are being pushed forward and back in time.

A little over the top

Well, if two lasers was fun, what would happen if we used a lot of fancy lasers? With a prism we can in principle add the light of any number of different colored lasers, and the theory of what we should see is pretty clear. As we add more lasers, each locked in time and with even steps in color, the duration of the light pulses in the combined beam gets smaller and smaller. All of the photons have to show up, so when there is a pulse of laser light, it is very bright. But the dark times between the pulses gets wider and wider as we add more lasers. As we add more lasers the beating becomes more intense, and it starts to look like a strobe light—bright flashes separated by long periods of darkness.

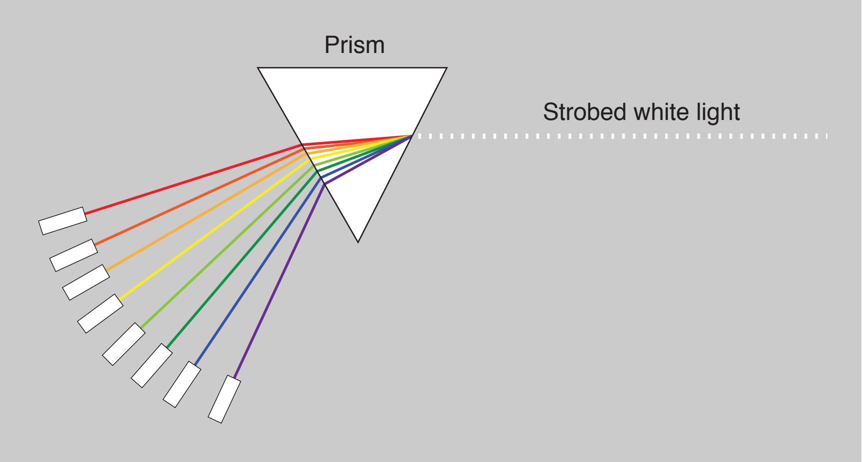

Many lasers of different colors combined with a prism. If the lasers are equally spaced in color and well-timed, the light coming out of of the prism will look like a white strobe light—bright flashes separated by relatively long periods of darkness.

Now, while this kind of laser cascade is entirely possible in theory, in practice they are a pain in the butt to set up. Lasers of this precision are expensive and fickle beasts. They are like Italian sports cars—incredible when they are running, but they spend as much time in the shop as lazing. Chaining them together and keeping them all working in synch requires incredible patience—a set of ten locked lasers is a major technical achievement.

But there is a kind of laser that emits very short pulses of light (creatively called a pulsed laser). By repeatedly firing our laser like a precision strobe light, we can create a stream of light pulses that looks *just* like the light stream after the prism in our hypothetical many-hundred-laser setup. So let's reverse our prism again, and send the strobe of laser pulses through the prism.

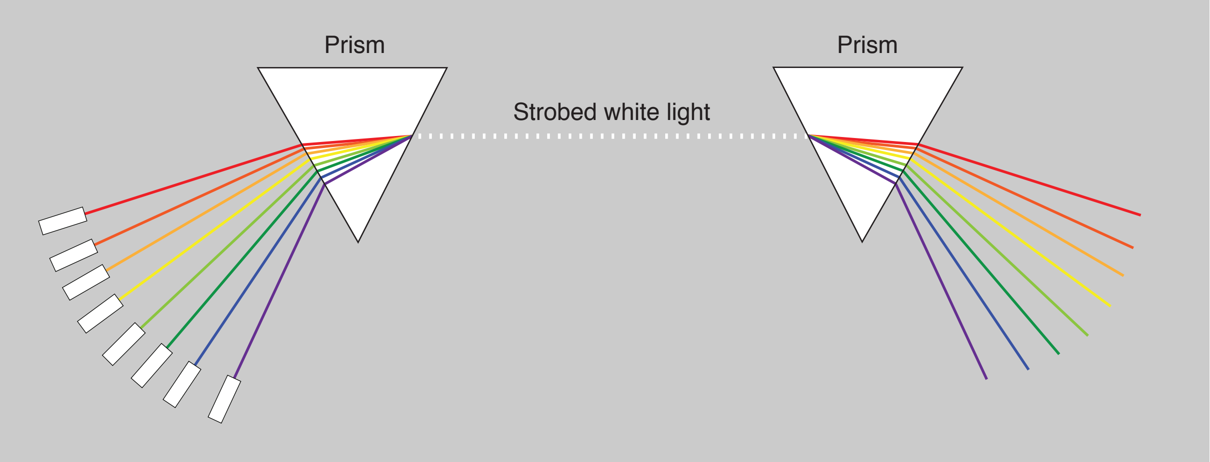

The light from many colored lasers combined with a prism to make a strobed beam. If this strobed beam is sent through another prism the light is again split into its component parts, revealing the many continuous lasers we started with.

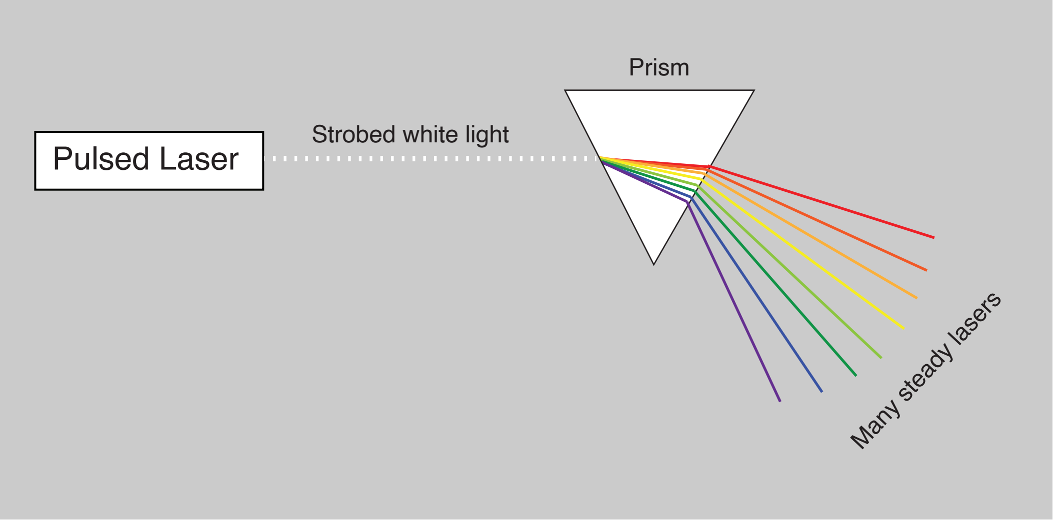

Instead of making the strobed light by combining the light of many lasers, we can make it directly using a pulsed laser with the pulses timed out with an atomic clock. And if we send this beam through a prism we see the same result —a comb of equally spaced continuous lasers.

When we look at the light from the strobed laser after the prism, it looks like a set of steady lasers equally spaced in frequency. In certain ways, this makes sense—if steady lasers can beat in time to make strobed light, the reverse should be true too. In other ways, it makes no sense at all. If I look at the light from one of the colors after the prism and time when the photons arrive, they arrive steadily in time. This means most of the photons arrive at times *between* the original laser pulses. The light in the individual 'lasers' after the prism is perfectly steady and is just as bright between the strobe pulses as during them. This is a purely quantum effect.

This strobed laser is called an Optical Frequency Comb, because the colors look like the teeth of a hair comb. The optical frequency comb is one of the great inventions of our century, and it is hard to overstate the importance it is having on emerging measurements; its development was awarded the 2005 Nobel Prize in Physics. To work properly, an optical frequency comb requires timing the pulses with an atomic clock and exquisite control of the shape of each pulse, but you can now just *buy* one of them. They aren't cheap, at least not yet, but several companies will sell you one complete with a warranty and a service plan.

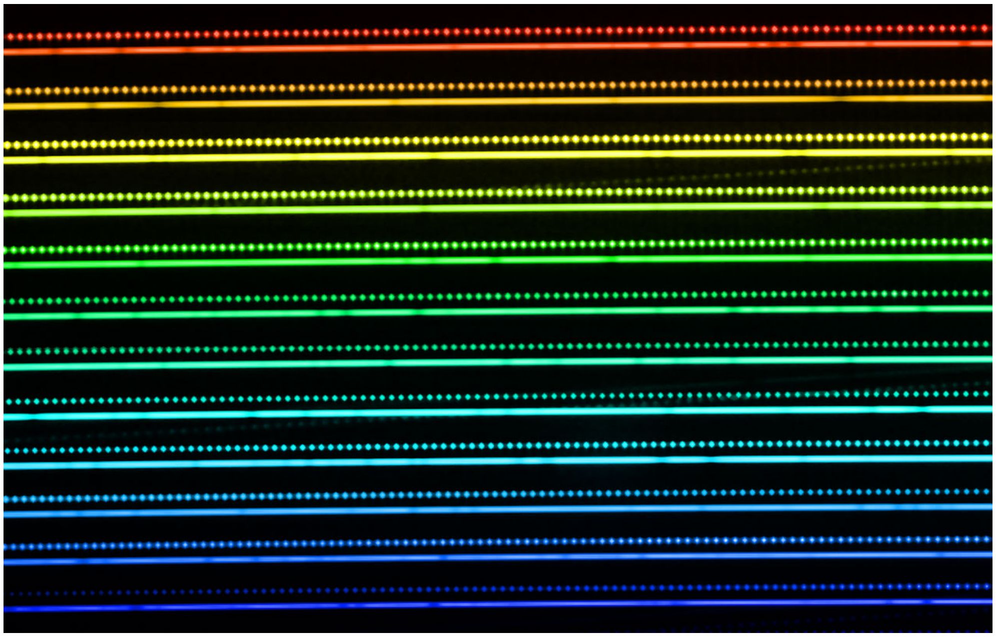

A folded spectrum from the High Accuracy Radial velocity Planet Searcher (HARPS) at the La Silla Observatory in Chile. The bottom line is the spectrum of a star with characteristic absorption lines (dimmer/narrower regions), while the top line is the spectrum of an optical frequency comb from a pulsed laser to provide absolute color reference.

Back at the Visitor's Center

I'm very excited that we got to see temporal interference this week, and how we can interfere different particles together. This is the kind of fun quantum effect that we rarely get to share with non-professional physicists.

We're building on the idea of particles moving as waves by showing that particles from different sources can blend together. Temporal beats can be viewed as a time analog of the stripes we saw in the slits in aluminum foil from the first article. And just like those stripes, temporal beats persist even when the number of particles from the two sources is less than one particle at a time. The mixing of colors to form beats is reversible, and an optical frequency comb uses strobes of light to create steady laser sources at many precise colors.

There are two neat applications I'd like to highlight: coherent lidar and optical clocks.

Lidar is the optical or infrared light analog of radar. Like any really useful technology, there are multiple versions and implementations. Many lidars work by bouncing pulses of light off distant objects. By measuring how long it takes for the light to return, they can determine how far away the objects are. But coherent lidars work on a different principle and are particularly well suited for precise speed measurements such as imaging the air flow near wind turbines.

In coherent lidar, a single color laser is bounced off an object and the doppler induced color shift of the reflected light indicates the relative speed. The trick is that the color shift is small, and it would be very difficult to actually measure the color with the necessary accuracy—the prisms would be prohibitively expensive. So instead these devices combine the reflected light with a copy of the outgoing light. Because the color of the reflected light was doppler shifted, we observe temporal beats when it is combined with the original laser beam. Measuring the speed of the beats measures the doppler shift of the reflected beam, and thus the speed of the object that beam is bouncing off of. Coherent lidars use the temporal quantum mechanical interference we explored at the beginning of this article to measure speed.

Which brings us to optical clocks, one of the new wonders of the world. All clocks work by counting 'swings.' In a grandfather clock, a pendulum slowly swings back and forth, once a second, and the clock works by counting the swings. In a mechanical watch, it is the twisting of a small wheel on a spring (typically 3 swings/second); in a quartz watch, it is the vibrations of a quartz crystal, typically 32,768 vibrations/second.

While many factors such as temperature affect the accuracy of a clock, one of the key contributors is simply how many swings it makes per second. It is easier to make an accurate quartz clock than a grandfather clock because it is oscillating more than 30 thousand times more quickly.

Atomic clocks oscillate a few billion times a second. In an atomic clock what is being counted are the oscillations of microwave light absorbed by an atom (cesium and rubidium are favorite targets). Fundamentally we are still just counting swings like in a grandfather clock, but because we get billions of oscillations per second an atomic clock can be vastly more accurate.

But there are optical atomic transitions that are even faster—hundreds of trillions of oscillations per second. How do you count that fast? Even the fastest computers have no hope of counting a hundred trillion oscillation a second.

The answer is to count beats instead. Because the pulses of light in an optical frequency comb were timed out with an (old fashioned) atomic clock, the colors of the 'lasers' after the prism are at known stable frequencies. So we can take an isolated mercury atom and select the light from particularly stable oscillation of electrons deep inside the atom. The light from this transition can then be combined with the light from the nearest 'laser' of the optical comb, and we can measure the beat frequency.

While we cannot count a 100 trillion oscillations a second, if we know the light from a laser in the comb is exactly a 100 trillion oscillations a second and we see 12 beats a second when we combine it with the light from the mercury atom, then we know the mercury light is oscillating 100 trillion + 12 times a second. We can use the combination of measured beats and a known reference to count *very* fast. And because the mercury atom oscillations are even more stable than our atomic clock, we can reverse the problem and use the beat frequency to correct the 'huge' errors in our atomic clock.

The precision of current optical clocks is *astounding*. You may have heard that time goes more slowly when gravity is strong due to general relativity. Optical clocks are so sensitive they can measure the different flows of time *2 cm apart* in height. If I lay a book on the table, the bottom of the book is slightly closer to the center of the earth than the top, so experiences slightly stronger gravity. This difference is measurable with an optical clock. Optical clocks are so sensitive we can no longer average the time of multiple clocks together—the ground you or a clock are sitting on typically rises and falls by ~5 cm a day due to *land tides*. The seismic motion of the ground currently limits our ability to measure time.

Precise optical clocks are but one application of the optical comb. Optical combs are transforming precision measurement in many areas, from finding planets around distant stars (precision doppler measurements), to potentially measuring the expansion of space itself (time dependence of redshift). Optical frequency combs are one of the next big things working their way out of the laboratory, and rely both on the mixing of particles and the measurement of beats between particles of different color.

Our next expedition

Congratulations on surviving another expedition into the quantum mechanical woods, this time to see effects rarely explored outside of advanced physics classes. In next week's expedition, I'd like to head into a different part of the woods. The first two articles looked at how particles move and mix. One of the natural questions that arises is if a particle can take two paths, how big is a particle? This will be our question for the next two articles. Along the way we'll learn why you buy 'bandwidth' when you want a lot of data, and how all particles can be divided into 'introverts' and 'extroverts'. So keep your boots dry, and we'll see you again next week.

FAQ

But the photons from the pulsed laser were emitted in pulses, it's *inconceivable* the light after the prism is not pulsed! I don't think that word means what you think it means. We do the experiment, this is the result. As an experimentalist it is not my job to tell nature what they can or cannot do. ¯_(ツ)_/¯

But it turns out nature is even more confounding (maniacal physics professor laugh). What happens if we take the pulsed laser beam and turn the brightness way, way down? Say by sending it through an absorber like some welding glass, so that on average only one pulse in ten even has a photon? Interestingly, none of our experimental results change at all. Now the light in each 'laser' color after the prism is very faint and we have to wait quite a while for a photon to show up. But the photons *still* show up at random times (mostly between pulses), and the spectrum *still* has the distinct comb-like color pattern we saw before. To really get you worried, if one photon at a time was coming in, what was its color before the prism? This is a question we'll tackle next week…