

The Femtosecond Frequency Comb: From Seconds to Attoseconds

R. Jason Jones

From the earliest observations of light-matter interaction by Ibn Sahl to Sir Isaac Newton, light has been used to provide exquisite detail concerning the physical world. In the field of spectroscopy, light has been used to probe the inner structure of atoms and molecules, providing some of the earliest and most stringent tests of the developing quantum theories that governed their configurations. The light emitted from lasers has since enabled experiments based on the *coherent* interaction between light and matter. Using advanced laser sources, light-matter interactions have been studied on very disparate time scales, from highly monochromatic lasers whose pure sinusoidal field can induce and follow atomic oscillations unperturbed for up to several seconds without losing track of its relative phase (i.e., greater than 500,000 billion atomic oscillations), to short bursts of laser light lasting only a few femtoseconds (10^{-15} sec, or a millionth of a billionth of a second!) where it is now possible to observe and even control the motion and rearrangement of atoms and molecules during chemical reactions. As state-of-the-art lasers used for these studies are highly specialized, research groups were often known for their work in either precision spectroscopy or ultrafast science, with each area utilizing different tools and often having little technological or scientific overlap prior to the late 1990s.

A paradigm shift has since occurred with the development of the femtosecond (fs) frequency comb. This powerful tool is nothing more than a highly coherent laser source emitting a periodic train of short pulses. Its development has, however, intimately connected the fields of precision spectroscopy and ultrafast science and coherently linked timescales from seconds to attoseconds (10^{-18} sec). Its impact has ushered in a new era of increasingly accurate time and frequency standards based on optical transitions in atoms and ions (optical atomic clocks), and unveiled the world of attosecond dynamics, where the motion, correlation, and rearrangement of electrons in atoms and molecules can be tracked and even controlled in real time. Such breakthroughs promise not only a deeper understanding of the physical world, such as the implications of our fundamental constants, but also has the potential to significantly advance or even replace technology as significant as our global-positioning system (GPS), the use of laser based radar (LIDAR), the temporal and spatial accuracy by which we can remotely guide unmanned vehicles (e.g., spacecraft, drones, robotic automation), and the precise calibration of astronomical telescopes needed for the detection of extra-solar planets, to name a few.

The frequency comb...simplified. The ideas driving the development of the fs frequency comb originated from the desire to both advance high-resolution spectroscopy by enabling the comparison of different atomic and molecular transitions spanning the ultraviolet, visible, and infrared spectrum, as well as to measure their *absolute* frequencies to develop better atomic clocks based on optical transitions. Although modern electronics can easily count the oscillations of microwave signals, measuring optical frequencies requires establishing a phase coherent link between ~ 500 THz atomic transitions and the 9,192,631,770 GHz frequency of the cesium atomic clock.¹ The basic idea behind the frequency comb and how it can link optical frequencies to microwave or radio frequencies is in fact quite simple, and is the result of the Fourier transform relationship, which relates a signal measured in time to its corresponding frequency distribution. For example, consider the pure sinusoidal oscillation of the electric field from a

¹ The current definition of the second is based on the time it takes for the cesium atom to oscillate 9,192,631,770 times between its hyperfine ground states.

linearly-polarized monochromatic laser as plotted in the left hand side of Fig.1a. For light in the visible portion of the spectrum, a single oscillation of the electric field lasts about 2.1 fs. This sinusoidal waveform corresponds to a single frequency of 474×10^{12} Hz, or a “red” wavelength of 633 nm (right column of Fig.1a). Now consider a single light pulse as shown in Fig.1b. This pulse can be viewed as having the same “carrier” frequency as in Fig.1a, but is strongly amplitude-modulated by a pulse envelope such that the waveform only lasts for $\Delta t \sim 8$ fs, in this example. The corresponding frequency spectrum for this pulse is still centered at 474×10^{12} Hz, but due to the short pulse envelope, the spectrum is very broad. This is also a result of the mathematical properties of the Fourier transform, which tell us that the width of the spectrum is inversely proportional to the pulse duration. Now consider the pair of pulses shown in Fig.1c. In this example the pulses have a time delay T between them that is exactly 13 periods of the laser carrier frequency.² Note that the two pulses are identical in this example. In this case we can see




Figure 1. Time (left column) and frequency (right column) relationship of the laser electric field (in blue). Black lines are for reference only. Coherent light emitted from (a) a single frequency cw laser, (b) a single laser pulse, (c) a pair of laser pulses, and (d) 4 laser pulses separated with equal time delays. Note that $1 \text{ THz} = 10^{12} \text{ Hz}$.

that they are “phase coherent” with the continuous-wave (cw) field of Fig.1a, in the sense that the field nodes and antinodes of the second pulse are still aligned with the original cw carrier frequency. Now consider the spectrum of this two-pulse waveform. Each pulse individually has the same spectral amplitude (as shown in Fig.1b). However, the mathematical property of the Fourier transform tells us that, because one pulse is delayed in time relative to the other, there will be a linear phase shift across the spectrum of the second pulse relative to the first, given by $\Phi(f) = 2\pi T \times f$, where f is frequency. The spectral components from each pulse will constructively interfere only when this relative phase shift is an integer multiple of 2π . Spectral interference peaks therefore occur at $2\pi T \times f = 2\pi n$, or when $f = n/T$, where n is an integer.

² This is not a typical example of the pulse period, which is usually on the order of nanoseconds (10^{-9} sec), but is simply used for illustrative purposes.

Now, if instead of a pulse pair we have four pulses, as shown in Fig.1d, or better yet, an *infinite* train of equally spaced pulses, the constructive interference peaks will simply become sharper and sharper until each resembles a narrow frequency similar to that of a cw laser, as in Fig.1a.

The above equation, which determines the location of these peaks, is often written in the following form:

$$f_n = n \times f_{rep}, \quad (1)$$

where f_{rep} is the repetition rate of the laser pulse train and f_n corresponds to one of the lasing frequencies. Modern laser systems are now capable of generating trains of pulses so short that the resulting “comb” of precisely spaced frequencies can span the entire visible spectrum.

The simple expression of Eq.(1) is quite powerful, as it directly relates the individual optical frequencies of the femtosecond comb ($f_n \sim 10^{15}$) to the easily measured microwave/radio frequency of the pulse repetition rate ($f_{rep} \sim 10^7 - 10^9$) by multiples of some large integer n .³ The absolute frequency of a highly monochromatic cw laser that is used, for example, to excite an atomic or molecular transition, can now be directly measured. One must first compare⁴ the frequency of the cw laser with that of a single “tooth” of the femtosecond frequency comb to measure the difference in their frequencies. Then, the frequency of that comb tooth can be determined by simply recording the repetition rate of the pulse train with a standard electronic counter (which can easily be referenced to an atomic clock such as cesium by using, for example, a GPS receiver). In this way one can think of the frequency comb as a set of large gears as found in a bike or clock, which converts the fast oscillations of an optical signal down to a slowly oscillating microwave signal that can be directly counted and/or compared with other transitions. As long as the teeth of these gears do not slip, the measured microwave signal will remain phase coherent with the optical signal.

The phenomenon of spectral interference, the constructive and destructive interference peaks in the frequency spectrum, due to the spacing of the pulses in the time domain, is completely analogous to what is observed in the diffraction of light from a single slit versus a large array of slits (such as a grating). When light illuminates a single slit, the resulting diffraction pattern far away as seen on a screen (known as the “far-field”) is related to the slit dimensions by the same Fourier transformation that relates the time and frequency domains. For a single narrow slit (relative to the wavelength of the light), a wide diffraction pattern will appear on a screen in the far-field. For two slits, this diffraction pattern will have a sinusoidal modulation due to the constructive and destructive interference of the light emerging from each individual slit (similar to the spectrum in Fig.1c). If one now has an infinite array of slits such as a grating, the interference peaks in the far-field will become very sharp. The different diffraction peaks correspond to the different diffraction orders of the grating. This is why the *resolving power* of a grating is related to the number of illuminated slits. Since the spatial width of a single diffraction order for a single color of light depends on the number of slits illuminated, the ability to separate or resolve different colors of light that hit the same grating also depends on the number of grating grooves illuminated. The different colors will be diffracted at slightly different angles, but will only be resolved if their separation is greater than the width of their individual diffraction patterns.

³ Although determining the integer n requires secondary techniques, there is no uncertainty in this quantity. The value of n will be obvious if the frequency to be counted is already known to a precision better than f_{rep} .

⁴ This comparison is usually done by mixing the cw laser frequency with the nearest component of the frequency comb on a photodiode and directly measuring their “beat frequency”, which is just the small difference of their optical frequencies. This is completely analogous to the tuning of musical instruments by comparing a pure and known tone with the corresponding key of the instrument and audibly detecting any difference from the beating between the two tones.

A closer look at the frequency comb. The concept of the frequency comb may seem simple, but the full realization of this powerful technique took nearly thirty years to develop. Why did this take so long? Many of the basic concepts were explored in early experiments conducted in the 1970s by Theodor Hänsch and colleagues.⁵ They were trying to utilize the equally spaced frequencies of a mode-locked dye laser for high-resolution measurements of energy level splitting in sodium atoms (known as fine-structure splitting). This early frequency comb served as a kind of ruler to precisely measure a large frequency interval. However, the laser technology at that time was not sufficiently advanced to answer many lingering questions about the feasibility of expanding the approach, or whether it was possible to measure the *absolute* frequency of optical transitions. In addition, many ultrafast groups focusing on the development of short-pulse laser systems were unaware of the subtle questions raised in this early work due to the minimal overlap of these different research communities. The unresolved question came down to this: “Just what do the pulses coming from the laser *actually* look like?” Laser systems developed to generate a train of ultra-short pulses are usually based on the concept of “mode-locking.” The pulse trains from mode-locked lasers in fact do *not* normally look like the pulse train shown in Fig.1d. Another issue was the intrinsic noise of the laser, which could affect each pulse differently, raising the question: “Is each pulse completely independent from the next?” Is it possible that there can exist some stable phase relationship between successive pulses, or between pulses from a mode-locked laser and a cw laser with the same carrier frequency (as illustrated in Fig.1)? Until the late 1990s, ultra-short pulses from mode-locked lasers were generally treated by most experimentalists and theorists as something like that shown in Fig.2a, where the underlying carrier frequency and/or phase is not well defined or even significant. The majority of research considered mainly the short duration of the pulse envelope. With improving laser technology through the 1980s and 90s (e.g. shorter pulse laser systems, reduced technical noise, lower noise pump lasers), researchers once again began to address this question.




Figure 2. A short pulse with (a) a poorly-defined carrier frequency and phase, and (b) a well-defined carrier-envelope phase.

⁵ See, for example, (i) A. I. Ferguson, J. N. Eckstein, and T. W. Hänsch, *J. Appl. Phys.* **49**, 5389 (1978), (ii) A. I. Ferguson, J. N. Eckstein, and T. W. Hänsch, *Appl. Phys.* **18**, 257 (1979), (iii) J. N. Eckstein, *Ph.D. Thesis*, Stanford University, 1978.

The average over multiple pulses from a mode-locked laser may indeed look something like that in Fig.2a. However, it certainly is possible to define, for a *single* short pulse, a phase like that shown in Fig.2b, where the “carrier-envelope offset” (CEO) phase is just the offset of the nearest crest of the electric field to the peak of the pulse envelope. When this pulse is circulating inside a laser cavity, the pulse envelope will travel at its group velocity through the laser gain medium and various optical elements. The carrier frequency, however, travels at its phase velocity. As it turns out, there is no requirement in the general theory of mode-locking that these two velocities must be the same. Therefore, after each round trip, the pulse that is coupled out of the laser through a mirror may in general have a different CEO phase. A more realistic picture of the pulse train is then shown in Fig.3, where the pulses are no longer identical due to this pulse-to-pulse CEO phase shift.




Figure 3. Relationship between the pulse train in the time domain (top) and the corresponding structure of the frequency comb (bottom).

The change of the CEO phases from pulse-to-pulse modifies the simple expression given in Eq.(1). In that expression, we said that the spectrum from successive pulses will only constructively interfere at frequencies where their relative phase shift is an integer multiple of 2π . We must now include in this analysis not only the relative phase shift due to the pulse delay, but also that due to the change in the CEO phase. The relative phase shift of the spectrum between each pulse is therefore given by $\Phi(f) = 2\pi T \times f + \Delta\phi_{CEO}$, where $\Delta\phi_{CEO}$ is the change in the CEO phase from pulse to pulse. Setting this equal to an integer multiple of 2π results in a new expression for the position of the frequency comb teeth, that is,

$$f_n = n \times f_{rep} + f_{CEO}, \quad (2)$$

where $f_{CEO} = \Delta\phi_{CEO}/(2\pi T)$ is known as the carrier-envelope offset frequency. This extra frequency shift of the entire comb (see Fig.3) is just a result of the time changing phase from pulse to pulse.

The final piece to the puzzle. In the 1990s researchers could determine that newly developed and lower noise mode-locked lasers could produce very narrow frequency comb teeth that were

rigorously spaced at exactly the pulse repetition rate and whose spectrum spanned over 100 nm.⁶ The individual frequency components of the mode-locked laser could be observed by simply overlapping the beam from a single frequency laser with the mode-locked laser. When this combined beam is sent onto a fast photodiode, the resulting photocurrent will only be sufficiently fast to follow the difference frequency (or beat frequency) between the cw laser and a few of the comb's teeth. These broadband frequency combs were immediately used to measure large frequency intervals between atomic and molecular transitions of interest. However, to determine the absolute optical frequencies of the comb components, one still had to find a way to measure f_{CEO} . Many researchers knew *how* to do it, but the technology still seemed too far away for it to be feasible. The technique would require a laser spectrum that spanned an *optical octave* (i.e., a frequency comb bandwidth so broad that, if the frequency of the lowest comb components were doubled, they would overlap with those at the higher frequencies). If this could be achieved, then when the lower comb frequencies were doubled (this could be done using standard nonlinear crystals), their offset frequency would double as well. When these doubled frequency components are mixed with the higher frequency components from the same comb, their difference frequency (beat frequency) would be exactly equal to f_{CEO} !

Although very short laser pulses could be generated with spectral bandwidths of ~ 100 nm, this was not yet the needed optical octave. As it turned out, a fortunate development (almost unbelievably so!) in the field of nonlinear optics had just been reported. Researchers at both Bell Laboratories and the University of Bath had developed specialty fibers with unique internal structure capable of broadening the spectrum of laser pulses to over an optical octave, producing a bright and a stunning rainbow of colors. Figure 4 shows one example produced with a Ti:sapphire laser in our lab after sending the light off a prism to be dispersed. The only remaining question was whether this nonlinear generation process preserved the comb-like structure of the original laser. It turned out that under appropriate conditions this “super-continuum” of colors does indeed preserve the original laser coherence. Within weeks of receiving this new “nonlinear” microstructure fiber, groups were able to directly detect f_{CEO} . Now that both f_{CEO} and f_{rep} could be detected, researchers were able to explore techniques for actively stabilizing these frequencies to stable, low-noise microwave references, thus eliminating any significant noise or drift in the optical frequency comb. Due, in part, to their efforts in developing the optical frequency comb, Theodor Hänsch (Munich, Germany) and John Hall (Boulder, Colorado) each received one quarter of the 2005 Nobel Prize in Physics.



Figure 4. Super-continuum of laser light generated from short pulses sent through a nonlinear micro-structured fiber. A zoomed-in look at this spectrum would reveal a discrete frequency comb structure rather than a continuous distribution of frequencies.

Connecting timescales from seconds to attoseconds. With the ability to directly convert optical frequencies to microwave frequencies based on the relationship in Eq.(2), the long envisioned goal of utilizing narrow optical transitions in atoms or ions as atomic clocks now became a reality. The race was soon on to develop superior atomic clocks with this emerging technology.

⁶ For comparison, the visible spectrum spans about 300 nm.

Why does this have such potential to improve the stability of atomic clocks? To accurately determine 1.0 second currently requires counting ~ 9 billion (9×10^9) oscillations of the cesium atom. There will always be some phase error in counting the final cycle of these oscillations, and that will ultimately determine the uncertainty in the measurement of a time interval (or a frequency). For simplicity let us say that we are off by an entire cycle in our count. This would mean our time measurement is off by $1/(9 \times 10^9)$ seconds, or $\sim 10^{-10}$ seconds. However, if one instead uses an *optical* transition frequency to measure (or define) 1.0 second, one will be counting $\sim 10^{15}$ cycles; the one-second interval will then be divided up into much smaller time periods. Assuming the same error in counting the last cycle, the total error will be only 10^{-15} sec (a femtosecond). This potential improvement in precision by five orders of magnitude assumes a number of technical details, such as the stability of the transition itself. However, all else being equal, developing atomic clocks based on significantly higher optical transition frequencies promised to dramatically increase the precision by which one could measure time and frequency.

This has in fact already turned out to be true. Spectacular improvements in atomic clock performance have been demonstrated in laboratories around the world since the development of the frequency comb. The gains are so significant that the cesium atomic clock will almost assuredly be replaced one day by an optical clock as the standard of time. The improvement in stability can impact many areas in science and technology. Most visible may be the precision by which time and navigation can be remotely determined. For example, GPS satellites all carry microwave-based atomic clocks on board as the position of a GPS receiver is based on the precision timing of signals sent from various locations. Improved timing signals can therefore directly impact navigation precision. In addition to benefits in technology, the frequency comb has become an invaluable tool in fundamental science. Atomic clocks are used not only for determining time and improving navigation; they can also be used as precision sensors of anything that might change the atomic resonance frequency. Examples include the study of collisions between ultra-cold atoms, which can shift their oscillation phase during collisions. Atomic transitions sensitive to magnetic fields can be utilized to make sensitive magnetometers. Optically-based clocks are currently capable of mapping out gravitational fields and/or altitude changes due to the effect known as “gravitational red-shifting,” which is a prediction of general relativity. If you raise a clock a meter higher, it will “tick” at a slightly different rate due to the difference in the gravitational field! Finally, researchers around the world are comparing measurements of various optical atomic clocks in search of any potential variation of Nature’s fundamental constants. Such a possibility is predicted by some physical theories. Since these constants determine the internal structure of atoms and molecules, any variation of certain fundamental constants would result in a change of the atomic/molecular transition frequency and therefore a change in the clock frequency.

We have mentioned how the frequency comb can produce pulse trains that remain phase coherent for time scales up to several seconds. The pulses themselves each consist of optical cycles lasting only a few femtoseconds. Researchers from the precision spectroscopy community who had developed ultra-stable cw lasers for decades quickly worked together with colleagues from the ultrafast community to perfect the frequency comb techniques. This technology, however, not only impacted the field of precision spectroscopy, but also paved the way for the field of attosecond science. In conjunction with Fig.3, it was pointed out that, by producing a well-behaved pulse train, one could generate a useful frequency comb. The reverse is also true. By carefully controlling the optical frequencies, one can precisely control the electric field waveform. As illustrated in Fig.2, without knowledge or control of ϕ_{CEO} , the pulse envelope will

determine the shortest timescales by which one can measure any change in an experiment. By stabilizing and setting f_{CEO} to zero (for example), researchers could now produce trains of *identical* pulses. This opens the door to utilizing a single oscillation of the carrier frequency to measure temporal dynamics.

For example, with extremely short pulses, setting ϕ_{CEO} to zero or π means that the central peak of the electric field will be either positive or negative and can be significantly stronger than subsequent optical cycles. Experiments have since demonstrated that the direction in which an electron is pulled off during ionization of an atom can be controlled by changing ϕ_{CEO} between zero and π . Other experiments use a single cycle of the pulse to produce a strong linear variation of the electric field in time, and use this to determine the exact moment at which an electron is ionized (due to a second pulse, for example). The approach is completely analogous to the use of a streak camera, which also uses a scanning electric field to map out the timing of an event. For example, if an electron is instantly ionized from an atom in the presence of the “streaking field” provided by the short pulse, it would experience a force proportional to the strength of the electric field, which in turn depends on the precise moment it was ionized. This technique enables one to measure when an atom is ionized to within a small fraction of an optical cycle. Since a single optical cycle lasts only a few femtoseconds, the timing resolution in these experiments is on the scale of attoseconds.

A final example is the ability to now generate single pulses of light lasting less than 100 attoseconds through a process known as high harmonic generation. Depending on when the electron is ionized, it is possible that the change in the sign of the electric field after half a cycle can bring the electron *back* to the atom it had just ionized, resulting in a collision. The energy the electron gains while being accelerated by the electric field can be given off as a high-energy photon during the recollision. This recollision happens during a fraction of the optical cycle, resulting in a burst of light on the time scale of attoseconds. By controlling ϕ_{CEO} , researchers can now control this process to generate attosecond pulses of light, which can in turn be used to probe ever-shorter time scales.

The control of ϕ_{CEO} enabled by the frequency comb has enabled new experiments that probe dynamics on attosecond time-scales. It has simultaneously enabled the measurement of atomic transition frequencies by counting their oscillations on time scales of several seconds. From next generation atomic clocks and precision spectroscopy, to attosecond science, the frequency comb has proven to be an invaluable tool that has dramatically impacted and brought together two seemingly unrelated research communities, to the benefit of all.

R. Jason Jones is an associate professor at the College of Optical Sciences at the University of Arizona. He received his PhD from the University of New Mexico in 2001. He continued on as a National Research Council Postdoctoral Award recipient and then senior research associate at JILA (University of Colorado) with Jun Ye and John Hall until July 2006. He joined the College of Optical Sciences as an Assistant Professor in August 2006. His research interests include optical physics, high-resolution and ultrasensitive laser spectroscopy, laser cooling and trapping, optical frequency metrology utilizing femtosecond frequency combs, nonlinear optics, and ultrafast optics. His current research involves the extension of the femtosecond frequency comb into the extreme-ultraviolet using intracavity high harmonic generation. He has over 80 combined journal and conference publications and two patents. He is the recipient of a National Science Foundation CAREER award (2007) and the DARPA Young Faculty Award (2009), as well as a Kavli Fellow of the National Academy of Sciences (Frontiers of Science 2010). He is a member of the American Physical Society and the Optical Society of America.

