Appl. Phys. B 74 [Suppl.], S17–S21 (2002)
DOI: 10.1007/s00340-002-0894-8

Applied Physics B
Lasers and Optics

H.G. MULLER

Reconstruction of attosecond harmonic beating by interference of two-photon transitions

FOM-Institute for Atomic and Molecular Physics, Kruislaan 407, 1098 SJ Amsterdam, The Netherlands

Received: 19 September 2001/
Revised version: 7 November 2001
Published online: 5 July 2002 •

ABSTRACT A method is proposed for detailed determination of the temporal structure of XUV pulses. The method is especially suited for diagnostics on attosecond pulses and pulse trains that originate from temporal beating of various harmonics of an ultrashort laser pulse. A recent experiment already showed the feasibility of this method when applied to long attosecond pulse trains, where it measured the average pulse characteristics. Here we argue that the same method is also suitable for determining differences between the individual attosecond pulses in a short train, or the properties of a single attosecond pulse.

PACS 32.80.Rn; 42.30.Rx

1 Measuring spectral phases

Spectral interferometry is one of the best methods currently available for fully characterizing ultra-short optical pulses. If it is known beforehand that an entire pulse does not span more than a time interval T, that pulse is uniquely determined from the knowledge of the amplitude and phase of a set of frequency components spaced by $\Omega = 2\pi/T$. The amplitude of these (or any other) frequency components can be immediately obtained from the spectrum, and the difficulty usually resides in obtaining the phases.

The optical technique known as SPIDER (spectral phase interferometry for direct electric-field reconstruction) [1, 2] determines these phases by measuring the interference between all pairs of frequencies ω, $\omega + \Omega$. To this end two pulses are sent into a spectrometer, the frequency of one of them shifted ('sheared') with respect to that of the other by Ω in a way that conserves the original phase. The spectrally resolved interference between these two pulses then reveals their relative phases.

Short pulses in the optical domain usually have a bandwidth that is small compared to their central frequency, and the independent frequency components have to be spaced by an interval that is only a small fraction of that bandwidth. This makes it inconvenient to up-convert one of the pulses and leave the other pulse unchanged, since the light required for up-conversion by sum-frequency mixing would map in the far or mid infrared. It is much preferable to up-convert both pulses by mixing with an optical photon of slightly different, precisely defined frequency. A convenient way to obtain the narrow-band photons required for the up-conversion process is by selecting different portions of a strongly chirped version of the pulse under study. This pulse is guaranteed to have enough bandwidth to generate a suitable spectral shear Ω.

The major limitation of SPIDER is that to measure pulses that are far from their bandwidth limit, the spectral sampling has to be rather dense (i.e. small Ω). To make a frequency that is sufficiently constant over the (long) duration of the original pulse, one has to chirp out the up-converting pulse so far that it might not have enough intensity left. In cases like this, however, only a small spectral shear is desired, so that up-conversion can be done with a stronger pulse from a stronger, narrower-band source, such as the pump laser of a fiber-compressed pulse (cross-SPIDER or X-SPIDER) [3].

2 Frequency mixing through multi-photon ionization

Implementing SPIDER in the vacuum ultraviolet or soft-X-ray region (XUV) encounters several difficulties. The only proven non-linear conversion involving XUV photons is ionization of a dilute gas of atoms [4, 5]. In contrast to conventional techniques in the optical region, this process does not result in photons, but in electrons. Despite this practical difference, the principle is otherwise exactly the same: the frequency spectrum of the ionizing photons is transferred to the energy spectrum of the photo-electrons. Instead of an optical spectrometer, one needs an electron spectrometer to extract the information.

The ionization process has the advantage that phase matching is not an issue. Unlike photons, electrons originating from different atoms do not interfere, since the quantum states they produce are in principle distinguishable by which atom was left behind in an ionized state. As a consequence, the process is also destructive in the sense that it irreversibly alters the conversion medium, but gas is cheap and simply replaced before the next shot. Another peculiarity of photo-ionization as a non-linear process is that it adds an energy-dependent phase to the phase of the XUV radiation in the conversion process [6, 7]. This phase is rather small, and can be calculated

✉ E-mail: muller@amolf.nl

from atomic theory with very good accuracy [8]. It can thus be easily taken into account when deducing the original XUV phase difference from the photo-electron interference.

The non-linearity of atomic photo-ionization is extremely weak in the XUV, and two-photon processes are out of the question for XUV intensities that are obtainable from harmonic generation, except for the lowest harmonics [9]. Mixed-color non-linear processes involving XUV and infrared (IR) are however much easier to perform. The rate of free–free transitions scales as ω^{-4} [6], favoring the infrared over e.g. its 11th harmonic by more than four orders of magnitude. In addition, the IR can be obtained directly from the laser, which can have six orders of magnitude more pulse energy than the harmonic pulses we want to characterize (mJ versus nJ).

The obvious choice for the frequency conversion is thus mixed-color two-photon ionization with one photon coming from the XUV pulse under study, the other photon coming from the infrared laser used to drive the harmonic-generation process. In other words, X-SPIDER seems the only viable solution. As a fortunate consequence this removes the necessity to find ways for chirping out XUV pulses. Usually the entire period of harmonic generation lasts only a fraction of the duration of the IR pump pulse [10], so that the frequency and intensity of the latter can be considered constant over this duration. Just as in optical SPIDER, the latter condition is required for undistorted transfer of the XUV spectral properties to the two-photon peaks in the electron spectrum. Also, here the IR pulse can be intentionally chirped out to satisfy these requirements if it does not originally do so.

3 The role of pulse delay in SPIDER

An aspect of SPIDER that is comparatively trivial to realize in the optical region, but very hard in the XUV, is the creation of two (delayed) replicas of the pulse under study. In SPIDER the delay τ between the replicas serves two purposes. In the first place, it makes it possible to up-convert the two pulse replicas by a different amount, by overlapping each of them in the non-linear medium with a different frequency of narrow-band light. In the second place, the time delay introduces an extra phase $\omega\tau$ in the frequency spectrum. In SPIDER τ is chosen so large that this extra phase causes several (≈ 2) fringes in the spectrum per independent frequency component (i.e. $\tau > 4\pi/\Omega$). Since the phase change from one independent frequency component to the next is at most $\pm\pi$, these extra fringes make an unambiguous reading of the phase from the spectral interferogram trivial. The recovery procedure only has to Fourier transform the spectrum, and the phase information automatically will appear at a distance τ from the transform of the spectral envelope, completely separated from it if τ is much larger than the duration of the pulse under study [2]. A smaller τ would make phase extraction much more difficult, if possible at all.

Encoding the phase information in the spectrogram in this way is a really elegant trick in the optical case, if we are dealing with pulses that are close to bandwidth-limited. For instance, a pulse with a guaranteed duration of less than 10 times its bandwidth limit is completely characterized by only 10 independent frequency components within this bandwidth. Even the most rudimentary optical spectrometer can resolve many more than 10 frequencies within the bandwidth of a 100-fs pulse, and using the excess resolution to resolve the superimposed fringe pattern due to the delay is thus completely free.

In the XUV we will be working with electron spectrometers, where it is much more difficult to get good energy resolution. Depending on the exact values of central frequency and bandwidth, electron-energy resolution very easily becomes the all-determining bottleneck, and this makes it a really bad idea to encode extra information in this energy. An XUV analog of SPIDER is much more easily realized by making τ just large enough to separate the pulses in time for independent up-conversion. In the absence of sufficient delay, reading the phase information from the spectral interferogram becomes ambiguous, since any variation there can be due to spectral amplitude differences just as well as to phase variations in the contributing pulses. This phase ambiguity can be resolved by repeating the measurement for two or three different relative phases of the up-converting fields. Rather than having all these different phases present in a single shot (as a function of ω) due to the large delay, we probe the various phases by independent measurements taken on separate shots, but in the end we get the same information.

In practice the phase could be changed by very slightly altering the delay of the up-converting (chirped) pulse, just enough to change the relative phase between the places where the two pulse replicas are sitting, but negligibly changing the frequency (which changes in the same way for both pulses anyway if the chirp was linear). In principle the up-converting pulse could be specially tailored to have two regions of constant frequency close to each other by liquid-crystal pulse-shaping techniques [11]. Although this might sound like a rather fancy solution, it is likely to be easier than doubling the resolution of the electron spectrometer, or building Michelson interferometers or etalons in the XUV.

It is interesting to consider what would happen in the case where the two pulse replicas overlapped in time. One obvious consequence is that independent up-conversion is no longer possible. Yet, for undistorted up-conversion of the spectrum of each pulse, it is essential that it is overlapped with the up-converting frequency for its entire duration. The only solution to this dilemma is to have the two frequencies for up-conversion present during the entire duration of both pulses. Each pulse would be up-converted by both frequencies. In the end we would have four up-converted (electron) pulses, which would all interfere with each other. The interference between two pulses up-converted by the same field would be independent of the relative phase of the two up-converting fields. Recording the spectral interferogram as a function of this phase and removing the constant background (not dependent on the phase difference between the two up-converting narrow-band pulses) recovers the same information as with non-overlapping pulses!

In fact the presence of two (delayed) replicas of the pulse in the overlapping case is only detrimental. The pair (overlapping or not) before up-conversion could be considered as a peculiarly shaped single pulse. Too short a delay creates fringes in their combined spectrum that are wide compared to the spacing of the independent frequency components, mak-

ing some of these components so weak that they would be likely to drown in the noise. If the pulses had zero delay, in other words, if there was only a single pulse, there would be no unnecessary fringes in the spectrum and every frequency component would be optimally visible. We will refer to the case where $\tau = 0$ as degenerate SPIDER, which thus overlaps a single pulse under study with two simultaneously present up-converting fields (e.g. obtained by reflecting the chirped up-converting pulse from an etalon).

4 The problem of spectral gaps

The case of measuring short XUV pulses from harmonic generation offers a special challenge. Even in the degenerate-SPIDER case with the single XUV pulse, where we do not create any unnecessary dark fringes in the spectrum at all, the spectrum of such a pulse is full of gaps: all spectral intensity resides at the harmonic peaks spaced by 2ω (Fig. 1). Conventional SPIDER runs into trouble due to such gaps. On the one hand, to characterize a pulse with an envelope of a duration of N optical cycles requires a spectral shear of ω/N, and with such a shear it is impossible to relate the phases in one harmonic peak to those in the next. The shear is not capable of bridging the spectral gap between the peaks in one step, and the vanishing frequencies in the gap have no defined phase that could be used as a 'stepping stone'.

If we are interested only in the average shape of the XUV field during an optical half-cycle, it suffices to sample the spectrum in 2ω intervals. Since the spectral power is exactly spaced by this interval, this solves the gap problem. (This is of course a consequence of the fact that the XUV emission was very similar from one half-cycle to the other, making this coherent average a meaningful quantity.) The practical problem is that the spectral shear of 2ω is rather large, making it unattractive to up-convert with two optical frequencies of that difference (which certainly exceeds the bandwidth available in the optical pulse).

Fortunately nature offers an ideal solution to this problem. Sending in a single optical frequency ω together with the XUV will automatically cause both sum and difference frequency-mixing peaks to appear in the electron spectra, and they are separated by 2ω. In the case of a several-cycle-long XUV pulse both these sidebands appear in a spectral region that did not contain any spectral power originally, and their interference can be studied background-free (Fig. 1b). This method has been successfully used to characterize the average temporal shape of the sub-cycle temporal beating of the harmonics [12], and we will refer to it as RABBITT (resolution of attosecond beating by interference of two-photon transitions). As argued above, RABBITT can be considered as a degenerate version of SPIDER, specifically tailored to the case of harmonic generation where the spectrum is empty except for some equidistant peaks.

5 Partial reconstruction of a pulse train

In [12] a fairly long XUV pulse was analyzed, and no attempt was made to extract information about the pulse envelope. Thus only a single-frequency shear 2ω was used, sampling the spectrum at the harmonic peak centers. To prove the principle of RABBITT a large oversampling was used with respect to the phase scan of the IR. The magnitude of the sideband peak at energy E thus reveals the corresponding part of the spectral interferogram, and its behavior as a function of IR phase (Fig. 2a) reveals the relative phase between the $E + \omega$ and $E - \omega$ electron peaks. The data in Fig. 2a clearly shows the interference fringes, with a contrast better than 2.

The periodicity of this signal really reveals that it is the sought interference signal. Modulation of the sideband intensity can occur through interference between the up-converting IR beam (for which the phase was controlled) and stray IR

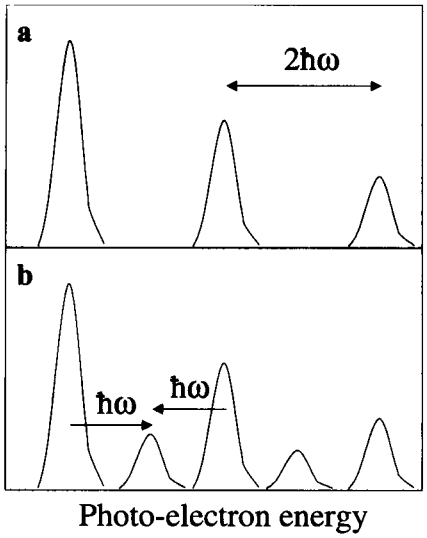

FIGURE 1 Photo-electron spectrum of an atom subjected to a superposition of odd harmonics, in the absence of the fundamental (**a**) and in the presence of it (**b**). The peaks that appear at the even-harmonic positions are in fact two-photon signals, due to both up-conversion of the neighboring lower harmonic, and down-conversion of the next higher harmonic by one fundamental photon. As a consequence, the (energy-resolved) interference in this peak contains information on the relative phases of these harmonics

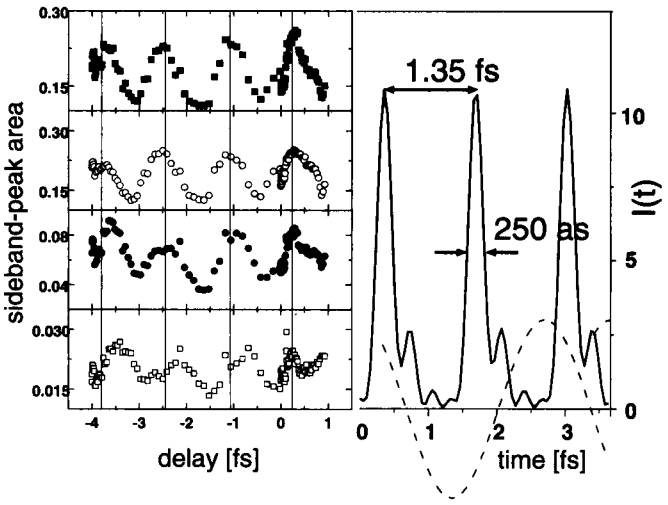

FIGURE 2 Example of a RABBITT measurement adapted from [12]). On the *left* (**a**) the sideband amplitude is shown as a function of the delay of the up/down-converting fundamental beam, clearly showing the interference between the up- and down-conversion processes. On the *right* (**b**) is shown the reconstructed XUV pulse shape collapsed into a single half-cycle

generated by four-wave mixing in the harmonic-generation gas jet (with a fixed phase). This could cause a large modulation of the IR intensity even if the stray IR is fairly weak. The IR intensity modulation would also show up in the sideband magnitude, due to the intensity dependence of the sidebands rather than their phase dependence. This artifact would have a period of 2.7 fs (the 800-nm optical period), though, making it clearly distinguishable from the 1.35-fs interference modulation. Similarly, unwanted contributions of higher sidebands, which occur if the IR used is too strong, would reveal themselves by modulations at higher frequencies.

Assuming that the phase of the XUV is constant within a single harmonic peak is equivalent to calculating the electric-field average of all half-cycles (negating alternate cycles to account for the parity of the atomic emission). The results of [12], repeated in Fig. 2b, show the result of such a construction. Independent measurements, however, have shown that the duration of the harmonic emission in this situation is only about 11 fs [5]. A sampling on a grid with 200-meV spacing (1/8 of a photon energy) would be sufficient to uniquely reconstruct a pulse shorter than eight cycles (i.e. 22 fs). In [12] the peak of the 13th harmonic has a width of about 10% of the 2ω spacing, and after allowance for the instrumental broadening this leaves an energy width of about 200 meV. It thus seems that there is only a single sample point within each peak, leaves little room for amplitude or phase structure in the envelope of this harmonic: at 200 meV from the peak center, where the nearest independent frequency is located, the amplitude has pretty much decayed to zero. The assumption that there is only one independent frequency component per harmonic peak is furthermore supported by the experimental observation that the peak shape of the sidebands does not seem to change with the IR phase. The multi-cycle average of the attosecond beating is likely to give a pretty fair impression of what actually goes on in every half-cycle.

To make any more precise conclusions, we must know the details of the harmonics envelope. RABBITT can in principle also be used to resolve such envelope structure. As mentioned above, the spectrum should be sampled with a step smaller than the width of the harmonic peaks, i.e. much smaller than the optical photon. The electron spectrum would have to be able to resolve the energies on this sampling grid, which becomes the main bottleneck for measuring very long trains of attosecond beats. As argued before, making the shear this small would not allow bridging the spectral gap from one harmonic to the next, but could perfectly resolve the amplitude and phase structure of each of the harmonics separately. For characterizing a single harmonic in the absence of any others, using a small shear would be the only option. Degenerate SPIDER would be the easiest way to obtain it, upconverting the XUV with a superposition of two chirped out and delayed versions of the driving pulse and measuring the spectral fringes in one of the sidebands as a function of the delay.

6 Envelope recovery of a short pulse train

In the presence of multiple harmonics, degenerate SPIDER with the (chirped, replicated, and delayed) fundamental frequency does not work: the interference between the two upper sidebands of one harmonic would be messed up by the lower sidebands of the next harmonic. One could get around this by choosing another frequency for the up/down-conversion, preventing the overlap of such sidebands (e.g. $\omega/2$). Such a SPIDER measurement would then reveal the structure within each harmonic bandwidth, while a complementary RABBITT measurement would relate the phases between different peaks. It is much easier, though, to use the overlap of the two sidebands to our advantage, and simply perform independent RABBITT measurements with two slightly different values of the optical frequency, ω and $\omega+\Omega$. The two sets of phase differences together allow us to step through the entire spectrum (both ω and $\omega+\Omega$ being able to bridge the spectral gaps), by zigzagging between one peak and the next until the entire width of the peaks is covered. There would be some redundancy in the information, which could be used to improve noise rejection or used as a consistency check.

RABBITT seems thus able to handle easily the complete characterization of a several-(fundamental-)cycle harmonics pulse, the limit being set by the resolution of the electron spectrometer. For instance, a (not too taxing) 100-meV resolution would suffice to completely characterize a pulse that is contained entirely within a 16-cycle time window (i.e. 42 fs), i.e. both the envelope structure and the temporal profile of each individual attosecond beat!

7 Characterizing a single attosecond pulse

On the short-duration side RABBITT seems to get into trouble when the individual harmonic peaks get broadened so much that they start to overlap their own sidebands. This will start to happen when the envelope is only a few half-cycles wide. Interference will then start to occur between both sidebands and the original peaks simultaneously. The details depend on how the electron spectrum is measured. In an angle-integrated measurement, the continuum states representing the sidebands and the original peaks are of opposite parity, and thus orthogonal. This means that the interference between the one- and two-photon processes cancels. Unfortunately, the magnitude of the one-photon signals has its own modulation, which does show up in the angle-integrated electron signal, and it would still be impossible to disentangle the sum of the two modulations. This precludes the use of the angle-integrated electron spectrum in the sub-cycle case.

If the complete angular distribution of the photo-electrons is recorded, it is rather easy to deduce the contributions of the continua of different parities separately, and thus determine the contributions from the sideband and the main peak. In this case also the interference between the sideband and the main peak can be unambiguously extracted, reducing the spectral sampling interval from 2ω to ω. This might not always be helpful (it only is if we can be sure that the total pulse duration is less than a single cycle, while we are not sure that it is less than a half-cycle). So we might as well use a poor-man's solution and measure the electron spectrum in a direction for which the angular distribution in the harmonic peaks has a nodal plane. For instance, with helium gas as the non-linear medium the single-photon signals will all

be p-waves, which have a nodal plane perpendicular to the polarization of the XUV. All electrons detected in this direction are thus purely due to the sidebands.

This solution seems to have no lower limit for the envelope duration. If the envelope becomes 'sub-half-cycle' the XUV spectrum becomes completely continuous (i.e. no longer interrupted by gaps and even without modulation within the 'harmonic spacing'). In that case sampling on a 2ω grid is enough to completely characterize the pulse. A second measurement with a slightly different sampling interval for obtaining phases of intermediate frequency components is then not even needed. In fact, the shorter the pulse, the wider the spectrum and the lower the required resolution of the electron spectrum, making the measurement more easy.

The limit to the time resolution with which sub-structure in a short pulse can be measured is determined by the total width over which the phases can be accurately measured. This width is determined by the energy range over which the electron spectrometer can resolve the independent frequencies of the sampling grid. The resolution ΔE of such a spectrometer is usually given as $\Delta E/E$, and is usually better than 1:100. This would make sub-structure 100 times faster than the total pulse duration resolvable! To achieve such fantastic resolution the individual phase differences should be determined precisely enough such that their accumulation over their full width of the spectrum should not produce an intolerable overall error. Uncorrelated phase-difference errors of 1% of a full cycle would grow on average to an error of 10% after 100 steps. This puts severe, but not ridiculous, requirements on the signal levels. The requirements could be relaxed somewhat by doing an additional measurement with a larger frequency shear (e.g. using the third harmonic instead of the fundamental), so that the spectrum can be covered in fewer steps. In addition, the atomic ionization cross section should be large enough over the entire bandwidth of the XUV pulse, unless we are prepared to repeat the measurement on a few different gases that together cover the photon energy range.

In conclusion, we can say that RABBITT seems a comparatively simple way for complete characterization of XUV pulses from harmonic generation, down to the very limit of single attosecond pulses. It does not require any complicated technique for handling the XUV pulse to be characterized, and makes optimal use of the available energy resolution of the electron spectrometer. Its main drawback is that it can not obtain the phase information unambiguously from data recorded in a single shot, but needs a limited amount of scanning (two or three measurements at different optical phasing) to resolve the ambiguity. Such scanning is only meaningful if a guarantee exists that all the measurements of the scan were using the same XUV pulse, which requires stabilization, or at least determination of the carrier-envelope phase [13, 14] of the pulse generating the harmonics.

ACKNOWLEDGEMENTS This work is part of the research program of FOM (Fundamental Research on Matter), which is subsidized by NWO (Netherlands Organization for the Advancement of Research). I acknowledge fruitful discussions with Ian Walmsley, Ferenc Kraus, and Paul Corkum during the Ultra Fast Optics meeting, where we shared our ideas on XUV implementations of SPIDER in a private session.

REFERENCES

1 C. Iaconis, I.A. Walmsley: Opt. Lett. **23**, 729 (1998)
2 C. Iaconis, I.A. Walmsley: IEEE J. Quantum Electron. **QE-35**, 51 (1999)
3 M. Hirasawa, N. Nakagawa, K. Yamamoto, R. Morita, H. Shigekawa, M. Yamashita: Appl. Phys. B. **74** (2002), DOI: 10.1007/s00340-002-0891-y
4 J.M. Schins, P. Breger, P. Agostini, R.C. Constantinescu, H.G. Muller, G. Grillon, A. Antonetti, A. Mysyrowicz: J. Opt. Soc. Am. B **13**, 197 (1996)
5 E.S. Toma, H.G. Muller, P.M. Breger, M. Cheret, P. Agostini, C. Le Blanc, G. Mullot, G. Cheriaux: Phys. Rev. A **62**, 061801(R) (2000)
6 S. Klarsfeld, A. Maquet: J. Phys. B **12**, L553 (1979)
7 M. Aymar, M. Crance: J. Phys. B **13**, L287 (1980)
8 H.G. Muller: 'Photoionisation of atoms in strong radiation fields'. Thesis, Vrije Universiteit Amsterdam, 1985, p. 117
9 T. Sekikawa, T. Ohno, T. Yamazaki, Y. Nabekawa, S. Watanabe: Phys. Rev. Lett. **83**, 2564 (1999)
10 A. Bouhal, P. Salieres, P. Breger, P. Agostini, G. Hamoniaux, A. Mysyrowicz, A. Antonetti, R.C. Constantinescu, H.G. Muller: Phys. Rev. A **58**, 398 (1998)
11 A.M. Weiner: Rev. Sci. Instr. **71**, 1929 (2000)
12 P.M. Paul, E.S. Toma, P. Breger, G. Mullot, F. Augé, Ph. Balcou, H.G. Muller, P. Agostini: Science **292**, 1689 (2001)
13 D. Jones, S.A. Diddams, J.K. Ranka, A. Stentz, R.S. Windeler, J.L. Hall, S.T. Cundiff: Science **288**, 635 (2000)
14 A. Apolonski, A. Poppe, G. Tempea, Ch. Spielmann, Th. Udem, R. Holzwarth, T.W. Hansch, F. Krausz: Phys. Rev. Lett. **85**, 740 (2000)