Nonlinear optics: the first 50 years

Quang phi tuyến: 50 năm đầu

In the summer of 1961, a landmark experiment was performed at the University of Michigan in which optical second harmonic generation was observed for the first time. This event 50 years ago marked the birth of modern nonlinear optics, and this article celebrates the first half century of what is now a vast and vibrant field at the cutting edge of laser technology. The focus is mainly on nonlinear optics in the 1960s partly because it is appropriate in this anniversary year to remember the genesis of the field, but also because such remarkable progress was made in the first few years. However, a brief review of where things stand at present is included, and one aspect of the field today (high harmonic generation) is taken as a representative example of an area of nonlinear optics that lies at the current frontier of knowledge.

Vào mùa hè năm 1961, tại đại học Michigan, các nhà nghiên cứu đã thực hiện thành công một thí nghiệm mang tính đột phá-thí nghiệm phát sóng hài bậc hai. Sự kiện này đánh dấu sự ra đời của một ngành khoa học mới, đó chính là quang phi tuyến, và bài báo này tổng kết lại những thành tựu đã đạt được của ngành khoa học quang phi tuyến rộng lớn và phát triển không ngừng cùng với ngành công nghệ laser. Ở đây chúng ta chủ yếu tập trung vào quang phi tuyến ở những năm 1960 một phần vì nó phù hợp trong bối cảnh này để mọi người nhớ đến nguồn gốc của lĩnh vực, và cũng vì những bước phát triển vượt bậc đều xảy ra trong những năm đầu tiên này. Tuy nhiên, chúng tôi cũng điểm lại ngắn gọn tầm quan trọng cũng như vai trò của từng hiện tượng ở thời điểm hiện tại và chọn hiện tượng phát sóng hài bậc cao làm ví dụ tiêu biểu cho những hiện tượng quang phi tuyến đang được nghiên cứu bước đầu.

1.Giới thiệu

> This year (2011) is the 50th anniversary of a landmark 1961 experiment in which optical second harmonic generation was observed for the first time. This marked the birth of modern nonlinear optics. Theodore Maiman had demonstrated the first ruby laser the year before; as we shall see, it was the intense optical frequency electric fields delivered by the laser that made the harmonic generation experiment possible, and that powered the rapid growth of nonlinear optics in the years that followed.

Năm nay (2011) là kỷ niệm lần thứ 50 ngày thực hiện thành công thí nghiệm quang phi tuyến đầu tiên-thí nghiệm phát sóng hài bậc hai. Thí nghiệm mở đầu cho sự hình thành quang phi tuyến hiện đại. Trước thời điểm này một năm, Theodore Maiman đã chế tạo thành công laser ruby đầu tiên, và chúng ta sẽ thấy, đây chính là các trường điện tần số quang học cường độ mạnh làm cho hiệu ứng phát sóng hài bậc hai có thể xảy ra và đã góp phần thúc đẩy sự phát triển của quang phi tuyến trong những năm sau đó.

Strong DC electric fields were of course available to scientists in the nineteenth century, and this enabled two early experiments in nonlinear optics to be performed. In 1875, the Rev. John Kerr of the Free Church Training College in Glasgow UK showed that the refractive indices of various liquids and dielectric materials were slightly altered in the presence of a high DC field [1]. The index change varied as the square of the field, and the process is now known as the DC Kerr effect. Then in 1893, Friedrich Pockels at the University of Göttingen published a paper on what we now know as the Pockels effect,[1] in which the refractive index of a non-centrosymmetric crystal changes in direct proportion to the strength of an applied DC field [2].

Trong thế kỉ XIX, các nhà khoa học đã tìm được những điện trường một chiều mạnh cho phép tiến hành hai thí nghiệm quan trọng đầu tiên của quang học phi tuyến. Năm 1875, mục sư John Kerrở trường huấn luyện giáo hội tự do Glasgow, Vương quốc Anh nhận thấy chiết suất của các chất lỏng và vật liệu điện môi thay đổi nhỏ khi đặt vào điện trường một chiều mạnh[1]. Hiện tượng chiết suất thay đổi theo bình phương cường độ điện trường một chiều được gọi là hiệu ứng Kerr một chiều. Sau đó, vào năm 1893, Friedrich Pockels tại Đại học Gottingen xuất bản một bài báo về hiện tượng mà sau này thường gọi là hiệu ứng Pockels, trong đó chiết suất của một tinh thể bất đối xứng tâm tỉ lệ thuận với cường độ điện trường một chiều[2].

The origin of these (and many other) nonlinear optical phenomena can be identified if the optical frequency polarisation P is expanded as a power series in the electric field namely[2]

[redacted]

Writing $E = E_{dc} + E_{\omega}cos\omega t$, and substituting the expression into Equation (1) yields (among other terms)

[redacted]

It is well known that the linear susceptibility $\chi^{(1)}$ is linked to refractive index through the equation $n = (1+\chi^{(1)})^{1/2}$, and it therefore becomes clear from Equation (2) that the terms in $\chi^{(2)}$ and $\chi^{(3)}$ represent *modifications* to n that are respectively proportional to E_{dc} and E_{dc}^2; these respectively represent the Pockels effect and the DC Kerr effect. For nonlinear optics to progress further, a source of intense optical frequency radiation was needed, and this is what the laser provided.

[redacted]

2. 1950–1960 and the invention of the laser

[redacted]

The 1950s had been a difficult decade. The world was still in the shadow of the Second World War, and economic conditions were still marked by austerity. In these circumstances, the dawn of the new decade in 1960 held a peculiar promise: 1945 now lay in the decade *before last*, and suddenly it was time to look to the future rather than to the past.

In physics, the 1950s had seen the development of the maser, the microwave device based on stimulated emission that predated the laser. By the end of that decade, competition to extend the maser principle into the visible part of the spectrum was intense. When

Theodore Maiman observed laser action for the first time on Monday 16 May 1960, he discovered the power source that supports most of today's optical technology.

[redacted]

Maiman's initial experiments on the ruby laser provided only indirect evidence of laser action; what he actually observed was a change in the relative populations of two energy levels in his ruby sample, which he took as evidence that a laser was feeding off one of them. But by the time of the press conference announcing the achievement on 7 July, he had seen the characteristic pencil beam at 694.3 nm, at the far red end of the visible spectrum. His first paper on the laser appeared in Nature [3] on 6 August 1960.[3] By the end of the year, Ali Javan at Bell Telephone Laboratories had also observed laser action, this time using a helium–neon mixture pumped by an electrical discharge; this was the world's first gas laser and the first continuous (CW) laser too [4]. From then on, new lasers came thick and fast. Laser development proceeded amazingly rapidly, and one of the first applications was in nonlinear optics, where the intense and highly-directional nature of laser radiation was precisely what was needed to get the new field off the ground.

4. Second harmonic generation

[redacted]

We now return to the main storyline and to the situation in early 1961 in particular. Remarkably, less than a year after Maiman's pioneering work, a small US company called Trion was already selling ruby lasers, and the strong optical frequency field in the ruby laser beam was precisely what nonlinear optics was waiting for. The link between laser intensity I (= power per unit area) and electric field amplitude E_ω is

[redacted]

By modern standards, the optical power available from the Trion laser was very modest but, if the beam was tightly focused, the field strength E^ω was sufficient for a team led by Peter Franken at the University of Michigan at Ann Arbor to observe optical second harmonic generation (SHG) for the first time in the summer of 1961 [6]. In the key experiment, a schematic diagram of which is shown in Figure 1, a ruby laser at 694.3 nm was focused into a thin quartz crystal, and the output beam analysed in a prism spectrograph for evidence of a second harmonic component at 347.15 nm. Despite a minuscule conversion efficiency of around 1 in 10^8, a weak component at the harmonic frequency was detected on a photographic plate.[6]

[redacted]

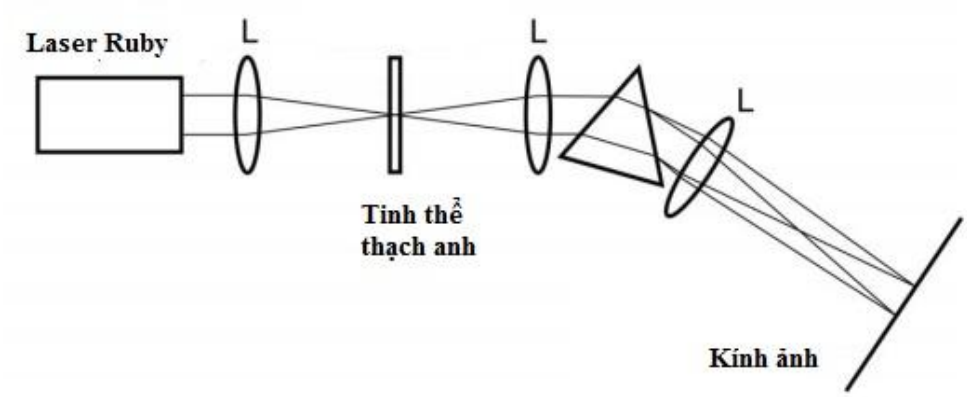

5. Phase matching

Although the observation of second harmonic generation was genuinely ground-breaking, the low harmonic conversion efficiency meant that it was a curiosity rather than a serious way of generating coherent ultraviolet light. Why was the efficiency so poor? The reason was that, due to dispersion, the fields at the fundamental and harmonic frequencies travelled at different phase velocities in the nonlinear crystal, and so quickly got out of step. This point can be appreciated by including spatial dependence in the expression for the ruby laser field by writing $E = E_\omega \cos\{\omega t - k_1 z\}$, where $k_1 = n_1\omega/c$ and n_1 is the refractive index at ω. This leads through Equation (1) to a second-order term in the polarisation of the form

But the space–time dependence of the second harmonic field is $\cos\{2\omega t - k_2 z\}$, where $k_2 = n_2 2\omega/c$, and this does not quite match the final term in Equation (4). In fact, it is easy to see that the two waves are π radians out of step when $(k_2 - 2k_1)z = \pi$, which provides a natural definition of coherence length[7]

In typical optical materials, L_{coh} is $\sim$10–20 μm, so only a small fraction of the quartz crystal in the University of Michigan experiment was participating usefully in the SHG process. This was why the conversion efficiency was so tiny.

The problem of limited coherence length was quickly solved. Independent papers from Joe Giordmaine at Bell Telephone Labs and a group under Robert Terhune[8] at the Ford Motor Company's research laboratories in Dearborn Michigan appeared together in the New Year's Day 1962 issue of *Physical Review Letters* [8]. The trick in both cases was to exploit the birefringence of anisotropic crystals, by making the fundamental beam an ordinary wave, and the harmonic beam an extraordinary wave. And because the refractive index of an extraordinary wave is dependent on the direction of propagation in the crystal, the angle was adjusted to ensure that $n_2^{ext}(\theta) = n_1^{ord}$.

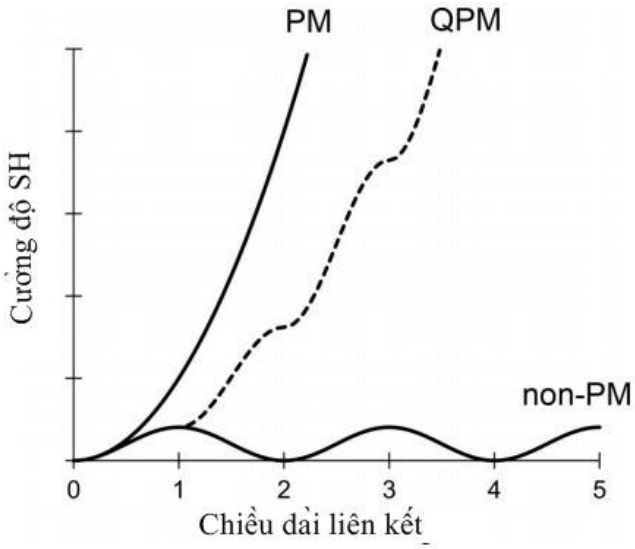

Figure 3 illustrates the dramatic improvement in SHG efficiency that phase matching delivers. The lower solid line (non-PM) that barely manages to lift itself off the bottom axis follows the second harmonic intensity when phase matching is not achieved. By contrast, the upper solid line (PM) shows the harmonic intensity increasing as the square of the distance under phase-matched conditions. So, if the coherence length is increased from (say) 10 μm to 1 mm, the intensity rises by $\sim 10^4$, and this was what changed nonlinear harmonic generation from a curiosity into a practical proposition.

On 10 April 1962, Armstrong, Bloembergen, Ducuing and Pershan (ABDP) of Harvard University submitted a paper to *Physical Review* that is truly remarkable for the breadth and depth of its coverage of nonlinear optics at such an early date [9]. One of the hidden gems in this seminal article is a suggestion for an alternative method of phase matching, one that would take around a quarter of a century to come to full realisation. The essence of the idea is reproduced in Figure 4, and the technique, now known as 'quasi-phase matching', is regarded by many researchers as the preferred method of phase matching. They may indeed see it as 'the best thing since sliced bread', a particularly appropriate aphorism given that Figure 4 looks very like a sliced loaf. However, the slices consist of nonlinear crystal, and are far thinner than anything on the breakfast table! The idea at least is simple. At the end of the first coherence length of the harmonic generation process, the harmonic field has reached a maximum, its phase relationship with the harmonic polarisation has slipped by 180°, and the interaction is just about to go into reverse gear. So, to keep things on track, we reverse the sign of the nonlinear polarisation by turning the next slice of crystal upside down. This enables the harmonic to continue to grow for another

coherence length, after which we turn the crystal right way up again. And so the process goes on. The trouble is that the coherence length is typically only 10–20 μm, so crystal wafers as thin as one hundredth of a millimetre are needed, and that is why it took so long for quasi-phase matching to be reliably implemented. New crystal growth techniques had to be devised and, in particular, a technique known as periodic poling in which an electric field is used to force the growth of crystal structure to reverse direction. Only then did the process envisaged in ABDP[9] become a reality.

[redacted]

The dotted line (QPM) in Figure 3 compares quasi-phase matching to phase-matching based on birefringence, and shows the SHG intensity rising in steps. The step height increases with distance, but this is purely because optical intensity goes as the square of the electric field; see Equation (3). If field rather than intensity were plotted, the steps would be of uniform height.

[redacted]

6. The early growth of nonlinear optics

The Harvard paper set the scene for the gold-rush period in nonlinear optics that occurred in the mid-1960s. Many new nonlinear phenomena were demonstrated between 1962 and 1965, including optical rectification [10], sum and difference frequency generation [11], third harmonic generation [12], optical parametric amplification [9,13], the optical Kerr effect [14,15], and stimulated Raman and Brillouin scattering [16,17]. Laser technology was of course progressing very rapidly at the same time. In particular, 1962 saw the development of laser Q-switching [18], which enabled laser pulses in the 20–50 ns range with peak powers of megawatts to be generated. The attendant increase in peak power had an immediate impact on the efficiency of nonlinear interactions, and the cross-fertilisation between laser physics and nonlinear optics that began at that time continues to this day.

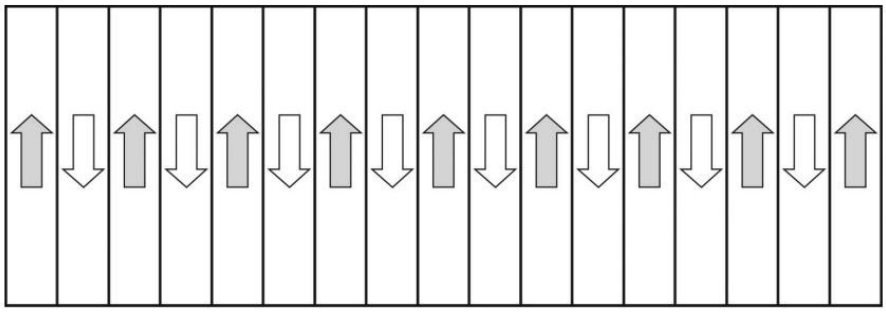

6.1. *Stimulated Raman scattering*

It was in fact during experiments on Q-switching using a nitrobenzene Kerr cell that Woodbury and Ng chanced on the first observation of stimulated Raman scattering [16]. They found that laser action was occurring at 765.8 nm (391.5 THz) as well as at 694.3 nm (431.8 THz), and spotted that the 40.3 THz frequency difference corresponded exactly to a vibrational resonance of nitrobenzene. They called the device a Raman laser, and the principle is now in widespread use, especially in fibre Raman systems.

[redacted]

Stimulated Raman scattering is the stimulated counterpart of spontaneous Raman scattering, a process that was first observed in the 1920s using conventional light sources. In those earlier experiments, the scattered frequency component was called the Stokes wave, and this terminology has been retained for the stimulated process.[10] In practice, the process is frequently based on vibrational resonances, although rotational or electronic states may also be involved. A further possibility is that the laser and Stokes waves interact via an acoustic wave; a process that is called stimulated Brillouin scattering [17].

[redacted]

6.2. *Optical rectification*

[redacted]

Another process first demonstrated in 1962 was optical rectification [10]. The possibility of rectification is evidenced by the presence in Equation (4) of the DC term

where the final step follows from Equation (3). The observation of optical rectification (OR) involves what must surely be the simplest experiment in nonlinear optics. All one has to do is to place a suitable nonlinear crystal between a pair of capacitor plates (see Figure 5), and a voltage appears between the plates in proportion to the laser intensity in the medium, as shown in Figure 6 [19]. Readers with a feeling for symmetry principles will immediately ask what determines the sign of the voltage or, to put it another way, what determines the difference between up and down in this experiment? The answer is that there must exist an inherent 'one-wayness' in the nonlinear crystal or, to put it more formally, the crystal must 'lack inversion symmetry', or 'be non-centrosymmetric'. Indeed, 21 of the 32 crystal symmetry classes possess this property, so it cannot be considered unusual.

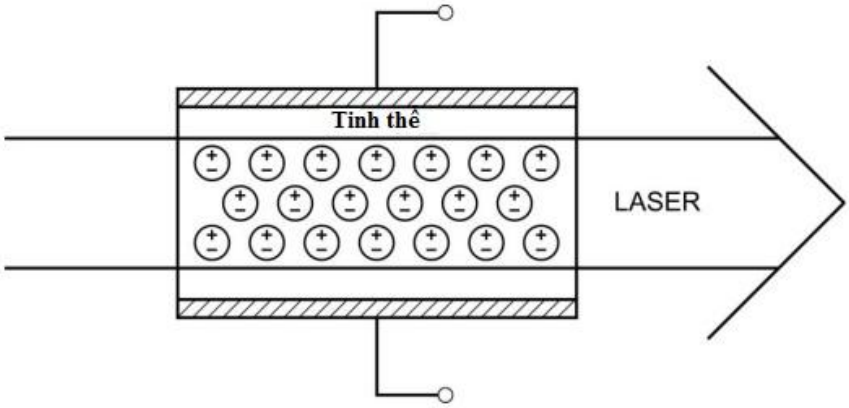

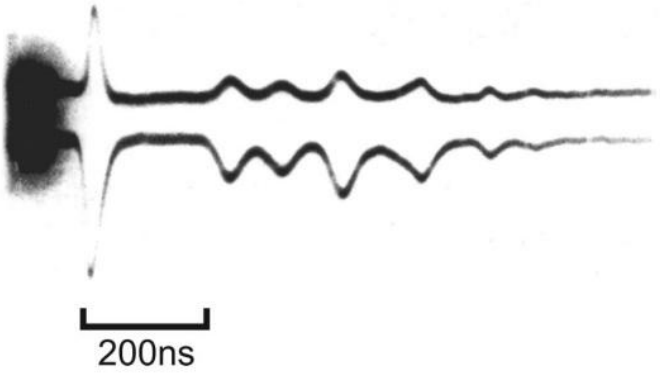

[redacted]

6.3. *Symmetry considerations*

[redacted]

This symmetry principle can be put on a sound mathematical footing by considering the term $P = {}_0\chi^{(2)}E^2$ in the expansion of Equation (1). In a medium possessing inversion symmetry, the sign of P must clearly reverse if the sign of E is reversed. But E squared is positive irrespective of the sign of E, and the conclusion is that all second-order nonlinear effects are forbidden under these circumstances. A non-centrosymmetric medium is therefore essential for $\chi^{(2)}$ to be non-zero. On the other hand, no such restriction applies for $n = 3$, so third-order effects based on $\chi^{(3)}$ exist in all optical materials, irrespective of their symmetry properties.

[redacted]

6.4. *Sum and difference frequency generation*

[redacted]

Sum-frequency generation (SFG) [11] refers to the process $\omega_1 + \omega_2 = \omega_3$, of which SHG is the special case where $\omega_1 = \omega_2 = \omega$ and $\omega_3 = 2\omega$. As an aid to understanding, it is helpful to multiply all terms in the SFG formula by $\hbar$, which then reads $\hbar\omega_1 + \hbar\omega_2 = \hbar\omega_3$. This highlights the fact that, for each photon gained at ω_3, one is lost at the two lower frequencies. But there are other possibilities too. If only the waves at ω_3 and ω_2 were present initially, what about $\hbar\omega_3 - \hbar\omega_2 = \hbar\omega_1$ or, if we started with ω_3 and ω_1, we could have $\hbar\omega_3 - \hbar\omega_1 = \hbar\omega_2$. In fact, in a comprehensive analysis of sum frequency generation, the two difference-frequency generation (DFG) processes need to be considered as well, and all three processes work together to ensure that the energy in the radiation field is conserved. It is the phase relationship between the three waves that determines whether the direction of energy flow is according to $\hbar\omega_1 + \hbar\omega_2 \Rightarrow \hbar\omega_3$ or to $\hbar\omega_3 \Rightarrow \hbar\omega_1 + \hbar\omega_2$.

[redacted]

6.5. Optical parametric amplification and the optical parametric oscillator

[redacted]

A fascinating possibility now presents itself. If the input to the nonlinear medium were a single wave at ω_3, could this be enough to initiate the process $\hbar\omega_3 \Rightarrow \hbar\omega_1 + \hbar\omega_2$? One could for instance argue that energy at ω_1 and ω_2 will certainly be present in the background noise spectrum. But the puzzling question remains: what determines how the ω_3 photon divides into two if there is no specific input at either ω_1 or ω_2 to get the interaction off the ground? After all, there is an infinity number of ways of cutting a cake into two so, if the process we are discussing is going to work, what determines the split?

[redacted]

The answer to this conundrum lies in a process of natural selection governed by the phase-matching conditions. Think of the waves exploring all possible division of $\hbar\omega_3$ into $\hbar\omega_1$ and $\hbar\omega_2$, and it is the one that is phase matched that will win out. Indeed, although not mentioned at the time, it is the critical importance of phase matching in nonlinear optics that allowed us to focus on the SFG combination $\omega_1 + \omega_2 = \omega_3$, and to ignore all the other possibilities (like $2\omega_1$ and $\omega_1 - \omega_2$). The process $\hbar\omega_3 \Rightarrow \hbar\omega_1 + \hbar\omega_2$ is called optical parametric amplification (OPA) [9]. The highest frequency ω_3 is called the pump, and the other two are called the signal and the idler.[11] In practice, a weak initial signal beam is normally used as a seed, to get the process started.

[redacted]

Lots of interesting phenomena occur in OPA, and one can have lots of fun simulating the process. A typical result is shown in Figure 7 where transverse intensity profiles of pump (top), signal (centre) and idler (bottom) are shown at the end of the interaction in a 7 mm crystal of lithium triborate (LBO) [20, 21]. The three beams have been separated purely for display purposes; in reality, they lie on top of each other, and the cross-hairs define the common centre line. LBO is a biaxial nonlinear crystal, which means that it exhibits the more complicated of the two types of birefringence. In this case, however, the geometry is as simple as possible; the signal and idler beams are in effect ordinary waves, while the pump is an extra-ordinary wave. Extraordinary waves exhibit strange

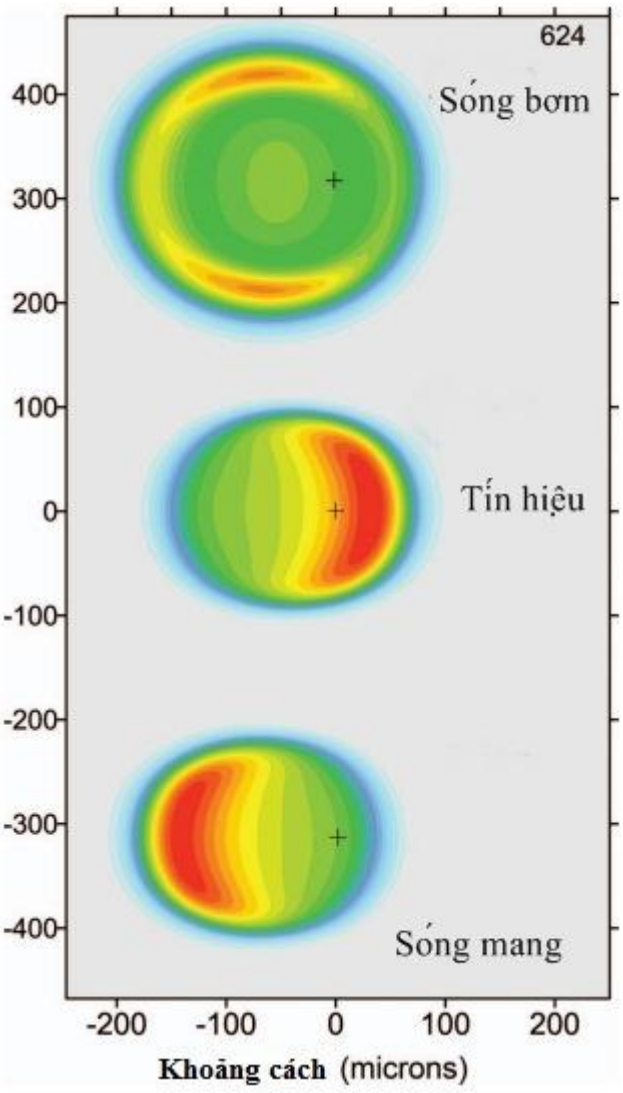
624
Sóng bơm
Tín hiệu
Sóng mang
400
300
200
100
0
-100
-200
-300
-400
-200
-100
0
100
200
Khoảng cách (microns)

[redacted]

propagation properties; for example, the direction of energy flow is not perpendicular to the wavefront. This feature is evidenced in the figure by the fact that the pump has 'walked-off' roughly 60 μm to the left. On the other hand, to optimise the bandwidth, the initial signal seed has been deliberately angled 1.3° to the right, which causes the idler to be angled 2.1° to the left to satisfy the phase-matching condition. Interestingly, these angles would imply respective sideways shifts of around 150 and 250 μm for signal and idler over the 7 mm length of the crystal, far larger than what is seen in the figure. But this ignores the fact that the three fields need to overlap for the interaction to proceed, so signal and idler have to cling together (and to the pump) in order to survive and grow. The pump profile has clearly been depleted fairly uniformly, which suggests that the pump energy has been used quite efficiently in this geometry.

In a particularly exciting extension of optical parametric amplification, the nonlinear crystal is located between high reflectivity mirrors at the signal frequency to create an *optical parametric oscillator* (OPO). Since the signal frequency is determined by the phase-matching conditions, and can therefore be controlled, an OPO solves one of the most vexing questions of laser technology: how to generate coherent radiation at an arbitrary frequency. While some lasers have broad bandwidths that offer limited tunability, laser sources are still to a large extent restricted to the energy levels that nature has given us. But OPOs offer freedom from this fundamental constraint.

Like quasi-phase matching, OPOs went through a long gestation period before becoming standard components in the well-found laser laboratory, as they are today. As for QPM, the problems were in materials technology. The first OPO was demonstrated in 1965 [13], but it was not until the 1980s that nonlinear crystals of the exceptional optical quality required for an efficient and reliable device became available.

[redacted]

6.6. *Third-order nonlinear effects*

[redacted]

Nonlinear optics in the 1960s was not restricted to processes based on the second-order term of Equation (1); indeed, several third-order processes have already been mentioned in this paper [12,14–17]. To see the possibilities, we extend Equations (2) and (4) by inserting $E = E_{dc} + E_{\omega} \cos\{\omega t - k_1 z\}$ into Equation (1). Including selected terms up to the third-order term yields

[redacted]

[redacted]

The term in the bottom line is third harmonic generation which is analogous to second harmonic generation except that it can occur in a centrosymmetric medium [12]. The first term in the third line is of course second harmonic generation, but a term involving $\chi^{(3)}$ now appears as well. SHG normally

requires a non-centrosymmetric medium, but the presence of a DC field provides a preferred direction in a centrosymmetric environment.

[redacted]

6.7. *Intensity-dependent refractive index*

[redacted]

The first three terms in the first line of Equation (7) appeared in Equation (2), but the fourth represents a new effect: the dependence of the refractive index on the square of the optical frequency field or (through Equation (3)) on the optical intensity. This is the basis of *intensity-dependent refractive index* (IDRI), an enormously important process with a wide range of practical applications [14,15]. The associated coefficient is normally positive, which means that refractive index almost always increases with intensity. This has the unfortunate and potentially catastrophic consequence that the attendant wavelength reduction along the axis of an intense optical beam leads it to collapse upon itself [14], causing irreversible damage to expensive optical components. Drastic steps have to be taken in large laser systems to avoid the calamitous consequences. On the other hand, the effect of IDRI on an optical pulse is to cause the peak central region to travel slower than the leading and trailing wings. This affects the carrier wave structure, causing what is known as self-phase modulation (SPM) where the local

[redacted]

frequency is lowered ahead of the peak and raised behind it. The result is that the carrier frequency rises through the pulse, a condition known as an 'up-chirp'. If IDRI becomes strong, the pulse envelope is distorted as well, causing the trailing edge gradient to rise as the peak suffers increasing delay, with the potential formation of a rear-end optical shock. This self-steepening effect was analysed in detail in 1967 [22] although, by a strange quirk of history, the analogous steepening of the trailing edges of the optical carrier waves had been considered two years earlier by Rosen [23]. Like so many things published in the 1960s, Rosen's 1965 paper was rediscovered in the 1990s and, today, 'carrier wave steepening' even has potential applications [24].

Self-phase modulation (SPM) has numerous important uses. Its most significant characteristic is the consequent increase in the spectral bandwidth [25] and, whenever bandwidth goes up, the laser physicist immediately thinks 'and that means that the potential pulse duration goes down, if the chirp can be removed'. The idea of imposing strong SPM and following it with negative group velocity dispersion (GVD), goes back as far as a 1969 paper by Fisher, Kelley, and Gustafson [26]; the operation is now routinely used for optical pulse compression.

A more esoteric outcome occurs if SPM and negative GVD occur simultaneously in an optical fibre, which leads potentially to the formation of *optical solitons*, unique self-contained solutions of the nonlinear wave equation. In 1965, Zabusky and Kruskal [27] performed numerical simulations of soliton pulse propagation, but this was in the pre-fibre era. It was not until 1973 that Hasegawa and Tappert [28] suggested that optical fibres were ideal media for soliton propagation, and it was not until the 1980s, that the concept was actually demonstrated.

In principle, solitons prevent the dispersive spreading that ultimately increases the bit error rate in optical fibre communications. Unfortunately, long-distance soliton transmission is difficult to implement, and soliton-based fibre systems are a rarity, having been overtaken by advances in conventional fibre technology.

6.8. Theoretical foundations

The approach to nonlinear optics adopted in this short article is very simplistic, and would have been regarded as such from the outset. All we have done is to plug expressions such as $E = E_{dc} + E_{\omega}\cos\{\omega t - k_1 z\}$ for the electric field into the hypothetical power series expansion of Equation (1). The procedure has certainly revealed a number of important nonlinear processes, but it has provided no explanation for how the nonlinearity originates and it has led to some misleading conclusions too.

[redacted]

A rigorous quantum mechanical treatment based on time-dependent perturbation theory (TDPT) provides broad justification for Equation (1), with successive terms in the expansion corresponding to different orders of perturbation. The complicated TDPT expressions for the nonlinear coefficients that are generated turn out to be dependent on all the frequencies participating in a given nonlinear interaction. Hence, the second harmonic generation coefficient is (for example) not identical to the coefficient governing optical rectification as Equation (4) suggests, and the coefficients for the DC and optical frequency Kerr effect are not directly related as Equation (7) seems to imply.

[redacted]

As explained in Section 6.3 above, non-centrosymmetric crystals are needed to observe second-order nonlinear phenomena. Many of these are also optically anisotropic, and the associated birefringence is of course exploited in birefringent phase matching. In a complete analysis of nonlinear interactions, the fields and the polarisation are vector quantities, and the coefficients $\chi^{(n)}$ will be $(n + 1)$th rank tensors. These features add extra layers of complexity to the theory of nonlinear optics, when one delves into the subject for real.

7. Was nonlinear optics all done in the 1960s?

[redacted]

Several further observations can be made about the state of nonlinear optics in 2011. Firstly, laser physics and nonlinear optics have always been closely connected, and the relationship is now closer than ever. Not only are lasers inherently nonlinear devices, but virtually all laser systems these days exploit nonlinear optics in one way or another, either as their basis of operation (e.g. Kerr-lens mode-locking [29]), or in the ancillary systems they drive, or at least in a number of key components. Could one perhaps argue that nonlinear optics is the wider field, and see laser physics as a compartment within it?

[redacted]

Secondly, while new principles may be few in number, a glance at the session headings of a leading international conference on lasers and photonics today (e.g. CLEO, or CLEO Europe) is revealing. Typical topics under nonlinear optics are likely to include nanostructures, photonic crystal fibres, metamaterials, high-harmonic generation, attosecond science, spatial solitons, electromagnetically-induced transparency (EMIT), slow light, remote sensing, among others. Some of these terms would have been meaningful 40 years ago but others would not. EMIT for example represents a dramatic extensions of ideas that were well-known by 1970 since it is related to *self*-induced transparency, which was studied in detail by McCall

[redacted]

High harmonic generation (HHG) grew from work on harmonic generation in gases and metal vapours in the 1960s and 1970s. Equation (7) includes a third harmonic term arising from the E^3 term in Equation (1), and third harmonic generation itself was studied in detail in the 1960s (see e.g. [12,33]). Experiments on third and fifth harmonic generation continued in the 1970s and early 1980s,[13] but they remained within the perturbative regime, where the series expansion of Equation (1) remains valid, and the conversion efficiency for successive harmonic orders drops off sharply. By the late 1980s, however, much higher laser intensities were available, and some remarkable experimental results were achieved, from which it was evident that a new strong-field regime was being entered [34]. The conversion efficiency still fell away for the lower harmonics but, after about the seventh or ninth, the efficiency remained essentially constant across a broad plateau that ended in a fairly abrupt high-frequency cut-off. The cut-off could be extended to higher frequencies by increasing the laser intensity, up to a saturation intensity beyond which no further extension of the plateau was possible. By the early 1990s, harmonic orders well into the 100s were generated in neon [35].

[redacted]

The simplest picture of this high harmonic generation (HHG) process is based on the so-called three-step

model [36]. In step 1, an atom is ionised in the optical field, creating a free electron. In step 2, the electron released moves almost freely in the field and, if the conditions are right, it is accelerated first away, and then back towards the parent ion. In the 'recollision' that ensues (step 3), the electron returns to the parent ion, and the accumulated kinetic energy combined with the energy of ionisation is released in a harmonic photon.

[redacted]

Sophisticated quantum mechanical techniques are needed to do justice to steps 1 and 3, but a simple one-dimensional classical model serves well for step 2, which turns out to play a crucial role in the overall process. Figure 8, which requires nothing more than first-year physics to produce, shows the kind of results that one obtains. The time variation of the driving field is shown at the top, while the curves below are (in grey) a set of electron trajectories tracing the electron displacement and (in black) the corresponding electron velocity. The twelve different curves in each set correspond to different times of ionisation within the first quarter-cycle of the field. All curves end at the point of recollision ($X = 0$), and the length of the verticals to the abscissa from the black curves represent the recollision velocity, from which the recollision kinetic energy U_K can be calculated.

[redacted]

[redacted]

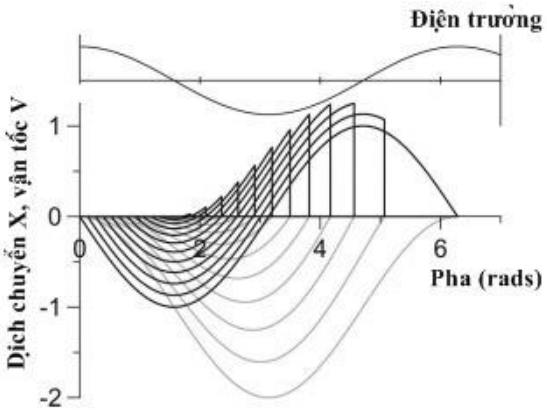

[redacted]

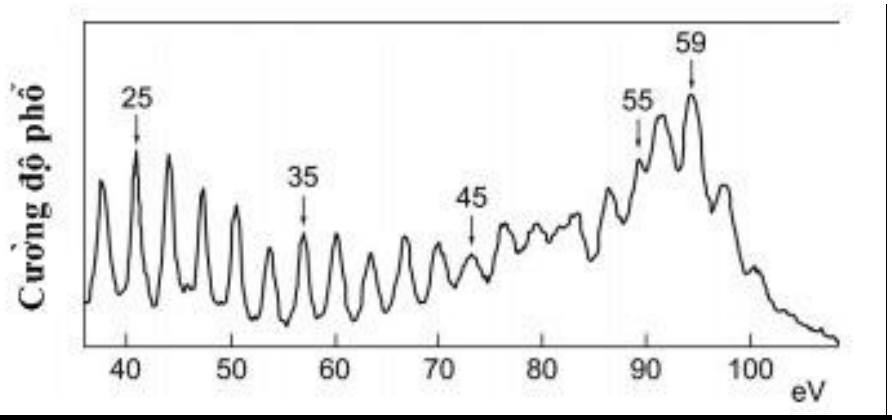

[redacted]

[redacted]