PHY 208 – Atoms and Lasers

Lecture 3: Is Laser Light Really Special?

Introduction

In this course, so far, we have looked into how a laser device produces a stream of photons, that we also refer to as a laser beam. We have taken it as a given that this is a useful and interesting process, but this lecture will actually address the question: is laser light really special? Is it unique? Can we only obtain light with the properties of that emitted by a laser by having a laser device? Or could we obtain a reasonable facsimile any other way?

To do so, we need to examine in more details some properties of laser light, and then see if these properties can be achieved in another fashion.

These lecture notes rely heavily throughout on the two following textbooks:

[S&T] Saleh & Teich, Fundamentals of Photonics, Chapter 10, Statistical Optics

[Brooker] Brooker, Modern Classical Optics, Chapter 11, Optical Practicalities: Etendue, interferometry, fringe localization

I. Coherence

Temporal Coherence

When discussing light, one of the most important properties differentiating laser light from that emitted by a thermal source is coherence. Coherence can be conceived as the predictable similarity of some property at a given point in time or space to another point in time or space. In our case, we could say,"if I perfectly know the amplitude of the electric field at some point, then how perfectly do I know it at every other point in time and space"?

So far, we have treated the idealized laser beam as a well-defined, single frequency (monochromatic) complex wavefunction, as described by equation 1.

$$U(\boldsymbol{r},t) = U(\boldsymbol{r})e^{j\omega t} \tag{1}$$

With this nomenclature[1], borrowed throughout this lecture from [S&T], the intensity of the beam is easily calculated.

$$\mathrm{I}(\boldsymbol{r}) = |U(\boldsymbol{r})|^2 \tag{2}$$

[1] NB: I am using the notation of S&T, where U is the complex wavefunction. It can be replaced by E for a propagating TEM wave, but then we need a factor of 2η to get optical intensity.

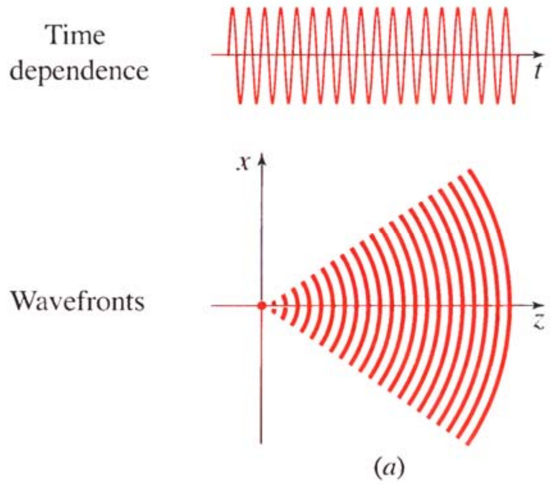

Figure 1. Time dependence and wavefront map for an ideal, coherent, monochromatic light source [S&T].

Taking and plotting the real part of eqn. 1 at a given point in space, as shown in Figure 1 (top), we obtain a sinusoid, typical of a propagating wave. If we consider the spatial distribution of the wavefronts coming from a point source (Fig 1 bottom), we see that they are perfectly spaced by the wavelength of the light. Knowing the value of the wavefunction at a particular spot in time and space, I know it perfectly at all other spots.

This idealized situation does not exist in reality, even for laser light. In reality, the exact instantaneous value of the wavefunction, U, will fluctuate in both amplitude and phase (around some central value), and has to be described statistically.

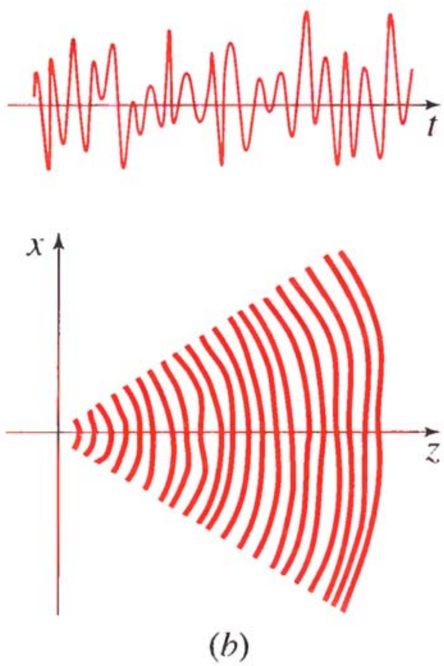

Figure 2. Time dependence and wavefront map for a partially coherent light source [S&T].

If we take a cartoon wavefunction where these fluctuations are great enough to significantly change the shape of the wavefunction (but where traces of the central frequency can still be visually observed), we could get a graph such as Figure 2 (top).

Tracing the wavefronts coming from a point source (Figure 2 bottom), we see that they are not perfectly and regularly spaced, but they do roughly maintain an average spacing.

We can say that the wavefunction representing this light is partially coherent. Intuitively, we can see the traces of the base frequency that reminds us of a sinusoid, but it is clearly not a sinusoid. These are helpful cartoon images, but we will need a more mathematical way to describe the difference with the ideal case.

Mathematical Description of Partially Coherent Light

As we consider such partially coherent light, we need to establish some further ground rules to describe and limit the type of wavefunctions that we will consider. To begin, we are going to consider **random light**, but light that is **statistically stationary**. This means that the statistical descriptors of the wave function (like mean, skewness, etc) do not change over sufficiently long time scales. Another way of describing this condition is to say that for some quantities, we will consider an ensemble average to be the defining quantity. That is to say, if we could repeat a given experiment numerous times, and the measured value would take on a constant value, then we will consider that value representative of the wavefunction.

Doing so for an arbitrarily varying wavefunction $U(t)$, the average intensity of the light becomes:

$$\mathrm{I}(t) = \langle |U(t)|^2 \rangle \tag{3}$$

Where the $\langle - \rangle$ indicates an ensemble average over many realizations of the random function under the same conditions. The requirement that such ensemble averages are constant in time puts a good limit on the type of wavefunctions that we are considering. For example, figure 3 shows two wavefunctions, both varying rapidly. Upon ensemble averaging, the stationary wavefunction (a) gives a constant value of intensity, while the nonstationary one (b) gives a pulse in intensity. This is a small constraint to put on the wavefunctions we consider, but a useful one.[2]

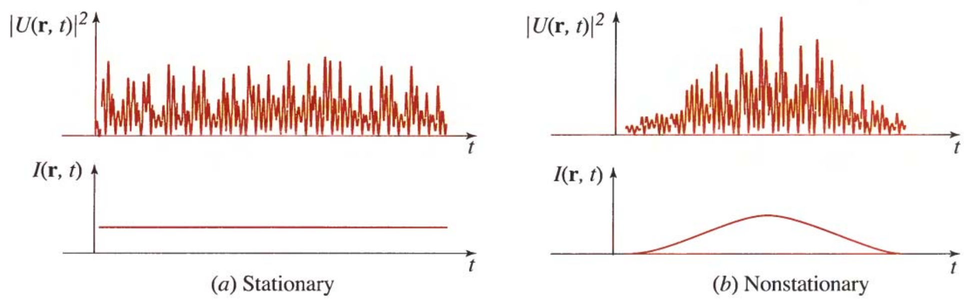

Figure 3. Demonstration of calculation of optical intensity for (a) stationary and (b) nonstationary wavefunction [S&T].

Now that we have narrowed the family of wavefunctions we are looking at, we must now look at this family in more detail. Although we have said that the "random fluctuations" will average out for quantities like intensity, these random fluctuations are very important. We therefore need a way to quantify the "randomness" of the light for a wavefunction *U(t)*.

For this we will use the **temporal coherence function,** or **autocorrelation function,** which we define as follows.

$$\mathrm{G}(\tau) = \langle U^*(t+\tau)U(t) \rangle \tag{4}$$

One can clearly see that this function compares a wavefunction to itself, but shifted by a given amount of time, τ. This function can be normalized (using the intensity) to get $g(\tau)$, which will vary between 0 and 1. We easily get $g(\tau = 0) = 1$. If the wavefunction looks nothing like its prior self after a given time, τ , then $g(\tau)$ will drop to zero.

---

[2] Note again that these are not quantum wavefunctions we are considering. These are functions describing the amplitude of the electric field of the propagating light beam, within a corrective constant.

We can now take a few examples of partially coherent waves and consider their autocorrelation function, $g(\tau)$. The figure below shows two such examples. For Figure 4(a), one can still recognize the underlying fundamental frequency of the wavefunction. One can also see that random fluctuations in amplitude and phase are occurring. Translated into an autocorrelation function, we would obtain a peaked function with a characteristic width, τ_C. I can now compare this to a slightly different wavefunction, also showing a central frequency, but with a much slower variation in amplitude and phase, Figure 4(b). Displaying the autocorrelation function for this wavefunction, we should expect a great coherence time, as the fluctuations are smaller. The autocorrelation function for (b) confirms this.

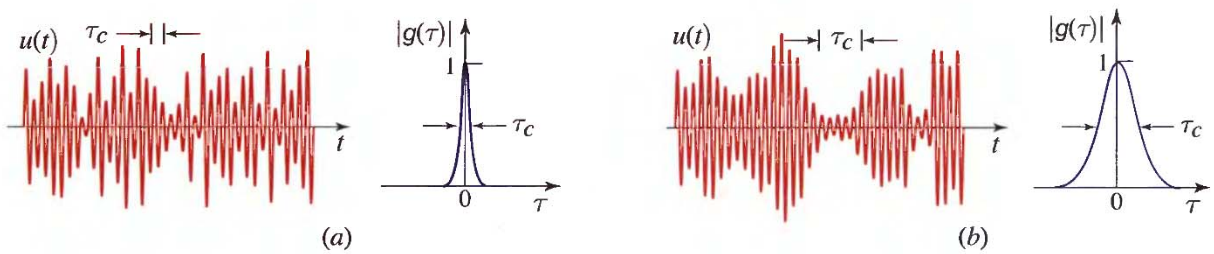

Figure 4. Wavefunction examples and concurrent autocorrelation functions [S&T].

To quantify the information in these graphs, we define the time scale over which $g(\tau)$ falls to 0.5 as the **coherence time, $\boldsymbol{\tau_c}$**.

The coherence time is a difficult value to picture intuitively. An easier way to visualize it is though the propagation of light with such a coherence time. In vacuum, in a coherence time $\boldsymbol{\tau_c}$, an electromagnetic wave will travel a distance $\mathbf{c}\boldsymbol{\tau_c}$ (**c** is the speed of light). The distance it travels is called the **coherence length, l_c**. A good way to picture the coherence length is as the **maximum path length difference** (after splitting a beam) allowed after which one could still observe **interference** effects. More accurately, for path differences much less than l_c, perfect interference fringes will form. The two beams are coherent, and the electric fields will add, and so can cause constructive or destructive interference. For path differences much greater than l_c, the two beams are incoherent, so only the intensities will add up. No fringes will form.

Spectral Width and Wiener-Khinchin Theorem

Another descriptor of light that we are used to seeing is the **optical power spectral density** $S(\nu)$, given in $[\mathrm{W/Hz}]$, where the total average power, $P = \int_{-\infty}^{\infty} S(\nu) d\nu$. This quantity may also be expressed as optical intensity spectral density, in $[\mathrm{W/Hz\ m^2}]$, giving the intensity of the beam when integrated, and may also be expressed in terms of wavelength rather than frequency.

This quantity is often plotted to graphically show the spectral density of an optical source. For example, shown in Figure 5 is the power spectral density of a green light emitting diode, sold by Thorlabs. Notice that in this case, the information is presented in terms of wavelength (as this quantity is typically more familiar to people). To transform from one quantity to the other, we must use the relationship:

$$\Delta\lambda = \frac{\lambda^2 \Delta\nu}{c} \quad (5)$$

Furthermore, if the data being plotted was not normalized, the units would be $\left[{}^{\mathrm{W}}/_{\mathrm{nm\ m^2}}\right]$, as it is an intensity spectral density being plotted.

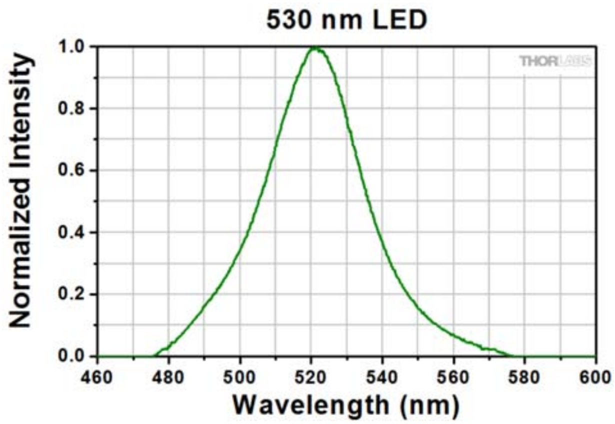

Figure 5. Power spectral density for a green LED, (www.thorlabs.com).

The Wiener-Khinchin theorem states that the **temporal coherence function** (or autocorrelation function) $G(\tau)$ and the **power spectral density** $S(\nu)$, form a **Fourier transform pair**:

$$\mathrm{S}(\nu) = \int_{-\infty}^{\infty} G(\tau)\exp(-j2\pi\nu\tau)\,d\tau \tag{6}$$

This means that the spectral width (**linewidth**) of a source and its **coherence time** are inversely related. The exact relation will depend on the definition of linewidth that is used.

If we use the definition $\Delta\nu_c = \frac{\left(\int_0^\infty S(\nu)d\nu\right)^2}{\int_0^\infty S^2(\nu)d\nu}$, then we obtain the simple relationship $\Delta\nu_c = \frac{1}{\tau_c}$. Again, we note that this relationship holds when the spectral density is given as a function of frequency.

[S&T] provides a table, reproduced below; with typical linewidths (given in terms of frequency), coherence times, and coherence lengths for a few light sources with which we are familiar.

Table 1. Spectral widths of a number of light sources together with their coherence times and coherence lengths in free space [S&T].

Source	$\Delta\nu_c$ (Hz)	$\tau_c = 1/\Delta\nu_c$	$l_c = c\tau_c$
Filtered sunlight ($\lambda_o = 0.4$–$0.8\ \mu$m)	3.74×10^{14}	2.67 fs	800 nm
Light-emitting diode ($\lambda_o = 1\ \mu$m, $\Delta\lambda_o = 50$ nm)	1.5×10^{13}	67 fs	20 μm
Low-pressure sodium lamp	5×10^{11}	2 ps	600 μm
Multimode He–Ne laser ($\lambda_o = 633$ nm)	1.5×10^{9}	0.67 ns	20 cm
Single-mode He–Ne laser ($\lambda_o = 633$ nm)	1×10^{6}	1 μs	300 m

Notice in Table 1 that even a single mode laser (which we have treated so far as a perfectly monochromatic source) has a finite bandwidth. Luckily, from the previous lectures and with the Wiener Khinchin theorem, we now have the tools to understand why.

Laser Linewidth and Spontaneous Emission

To understand the origin of the finite spectral width of a laser beam, we will now use a wavetrain model to look at the beam resulting from a set of atoms. We consider an ensemble of two level atoms and will specifically look at the excited state population. Spontaneously, one of the atoms can decay and release radiation at a frequency, ω_0 .

This radiation can then cause the stimulated emission of the surrounding atoms and result in a coherent wave.

$$u_0\, cos(\omega_0 + \varphi) \tag{7}$$

We will also assume that the gain is fully saturated, so the amplitude of the wave stays constant.

With each new spontaneous emission event, the phase of the radiation will change randomly. The total radiation emitted by the ensemble of atoms will therefore take the form:

$$u(t) = U(t) + U^*(t) \tag{8}$$
$$U(t) = U_0 e^{i(\omega_0 t + \varphi(t))} \tag{9}$$

Where $\varphi(t)$is a random variable that translates spontaneous emission into random jumps in phase.

Probability Law for Spontaneous Emission

Let's first derive the probability *p(t)* that the phase has not jumped within a time *t*.

As spontaneous de-excitation is a random, memoryless process, we can define γdt as the probability that the phase jumps between time *t* and *t+dt*. It can be shown (but we won't do it here) that $\gamma = \Gamma/2$, where Γ is the excited state lifetime. The probability p(t) must therefore satisfy

$$p(t + dt) = p(t) \times (1 - \gamma dt) \tag{10}$$

And if we assign p(0) = 1, we obtain

$$p(t) = e^{-\gamma t} \tag{11}$$

If we plot this function, as in figure 6, it makes sense intuitively. Initially, before any time has elapsed, the phase has definitely not jumped. However, as time progresses, it is more and more likely that the phase has jumped.

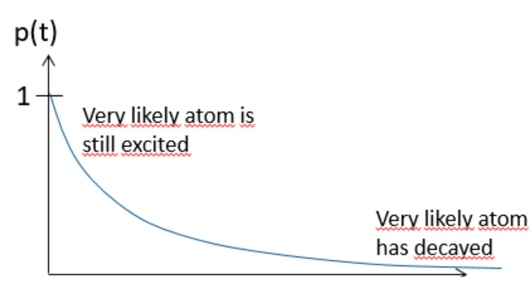

Figure 6. Probability vs time of an excited atom being in its excited state.

We can now use this probability to weight the autocorrelation functions for each of the two scenarios.

Autocorrelation function for Spontaneous Emission

Let us now calculate the autocorrelation function $G(\tau)$ of the emitted radiation.

(Rather than use *U* and U^* to get a real value, we will leave it as imaginary until the end).

We can distinguish between two contributions to the autocorrelation function:

$$G(\tau) = \underbrace{\langle U^*(t)U^*(t+\tau)\rangle}_{\text{with phase jump}} + \underbrace{\langle U^*(t)U^*(t+\tau)\rangle}_{\text{without phase jump}}$$

If the phase doesn't jump during the time $|\tau|$:

$$\langle U^*(t)U^*(t+\tau)\rangle = U_0^2 e^{i\omega_0\tau} \tag{12}$$

This case has a probability $p(|\tau|)$ of occurring.

If the phase does jump during the time $|\tau|$:

$$\begin{aligned}\langle U^*(t)U^*(t+\tau)\rangle &= U_0^2 e^{i\omega_0\tau}\left\langle e^{-i\varphi(t)}\right\rangle \\ &= 0\end{aligned} \tag{13}$$

Because the random variable $\varphi(t)$ means the wave is no longer coherent with itself at the later time, τ. This case has a probability $1 - p(|\tau|)$ of occurring.

The full autocorrelation function $G(\tau)$ can now be calculated, using the probability of each event.

$$\begin{aligned}G(\tau) &= \langle U^*(t)U^*(t+\tau)\rangle \\ &= p(|\tau|) \times U_0^2 e^{i\omega_0\tau} + (1 - p(|\tau|)) \times 0 \\ &= e^{-\gamma|\tau|} \times U_0^2 e^{i\omega_0\tau} \\ &= U_0^2 e^{(i\omega_0)\tau - \gamma|\tau|}\end{aligned}$$

And now we can use the Wiener-Khinchin theory to transform this into a spectrum

$$\mathrm{S}(\omega) = \int_{-\infty}^{\infty} G(\tau)\exp(-i\omega\tau)\, d\tau \tag{14}$$

$$\mathrm{S}(\omega) = \frac{I_0}{(\omega - \omega_0)^2 + \gamma^2} \tag{15}$$

This expression provides an important source of insight. Barring all other potential sources of broadening (and there are others), the **randomization of phase** caused by **spontaneous emission** is the fundamental, unavoidable source of broadening for lasers.

Is Laser Light Special? Coherence Edition

The Wiener-Khinchin theory tells us that saying a laser is **temporally coherent** and saying that it has **a narrow bandwidth** are two ways of saying the same thing. In this sense, the nature of laser light being coherent (more precisely, coherent with itself over long time scales) is equivalent to it simply having a much narrower spectral width than other light sources. We could therefore take a sufficiently intense, thermal light source and use a great number of optical filters to decrease the bandwidth, and therefore increase the coherence length of the beam. We would be throwing away a huge amount of optical power, but could recuperate at least the approximate optical spectral width of a laser beam.

In-class Example: In the lecture, we will take the example of a halogen headlight lamp, and see if optically filtering it down, we can achieve the same optical power as a small laboratory laser (HeNe).

However, the coherence/linewidth of a laser beam is only one of two properties that make laser beams useful. The second property is the ability to focus all that optical energy into a very small area, making possible the ablation of material or tissue, the reading or writing of tiny features, etc. In the next section, we will examine the physics surrounding the limits of focusing or concentrating light to see if it is this feature that makes laser light unique.

II. Focusing Light

The second useful and special property of a beam of laser light is the ability to focus it down to a very small spot, and concentrate a huge amount of optical power into a very small area. We will now look at the physics that limits our ability to cram photons into a very small area. To do so, we need to remind ourselves of a few descriptors of optical systems that help us define the limits of focusing light.

Acceptance Angle and Numerical Aperture

The first descriptor we will use is Numerical Aperture. The NA, or acceptance angle, of a lens or optical fibre describes the cone of light (in 2D) that can be collimated by that optical component.

$$NA = n\, sin\theta_A \tag{16}$$

Where *n* is the refractive index of the medium in which light is propagating, and θ_A is the acceptance angle, as defined in Figure 7. The NA accounts for both the **focal length** and **size** of lens, and note that it is dimensionless, always less than *n.*

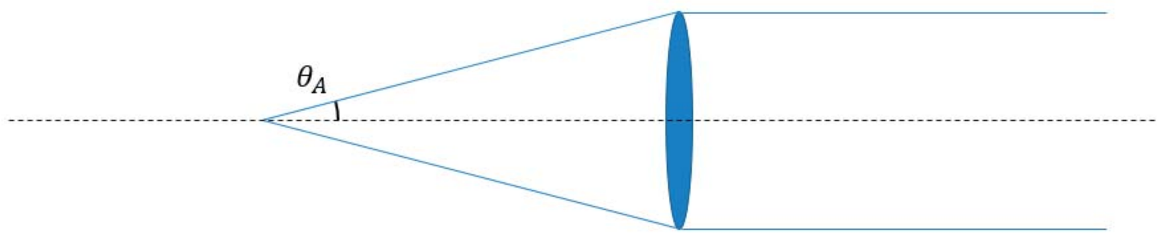

Figure 7. Depiction of acceptance angle for an optical lens

In this case, NA is described using the ray optics model, which can treat light as coming from a point source (and therefore assumes one can focus light to a point). Although the quantity NA will be useful to us later on, here it seems to be telling us that we could focus light down to an arbitrarily small area, making it seem like infinite concentration is possible. However, we know that the ray optics model fails at small sizes, and that we need to take the electromagnetic nature of the travelling wave into account. So let us now look at NA using the more accurate model of a Gaussian beam profile.

NA for Gaussian Beams

For propagating Gaussian modes (such as a single mode laser beam), the entire beam profile is fully defined by only the beam waist and the wavelength. That is to say, if we know those two quantities, we can describe the width of the beam at all other points in space (until it encounters some optical components). For such a Gaussian beam, the beam diameter, w, is defined as a function of z by

$$w(z) = w_0\sqrt{1+\left(\frac{z}{z_R}\right)^2} \tag{17}$$

Where z_R is the Rayleigh length, and is given by

$$z_R = \frac{\pi w_0^2}{\lambda} \tag{18}$$

Knowing this, and looking at Figure 8, we can define the numerical aperture for a focused Gaussian beam, in analogy with the ray optics definition, as $NA \approx \frac{\lambda_0}{\pi w_0}$.

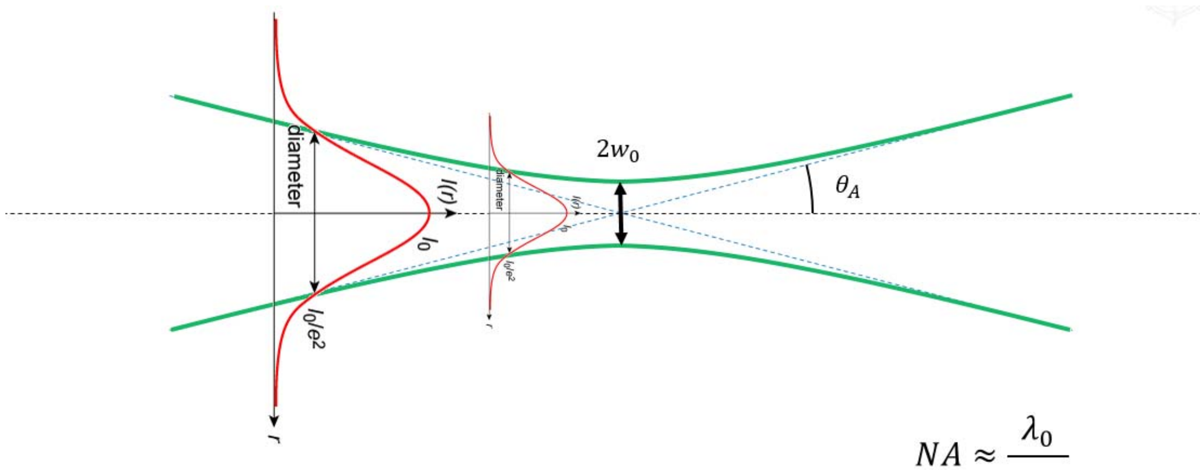

$$NA \approx \frac{\lambda_0}{\pi w_0}$$

Figure 8. Beam width vs position for a propagating Gaussian beam.

This can tell us the spot size for a lens (or mirror) of a given NA. This seems to put a limit on concentration, at least, for a coherent Gaussian beam. Whereas the ray optics model led us to believe infinite concentration would occur for a collimated beam and any lens, this Gaussian beam model directly links the concentration to the parameters of the lens focusing the beam.

However, it does imply that if with a sufficiently large lens/mirror with a very short focal length, we could achieve an arbitrarily small spot size for our filtered, monochromatic beam[3]. We have not yet determined a limit to focusing, so we have not yet found the factor that makes laser light special. To determine the fundamental limit to focusing, we need to introduce one more property, which is the optical étendue.

Optical Etendue

The **optical étendue** of a system is most easily pictured by imagining the collection of optical radiance of an object (or surface) with finite physical dimensions, which I will call the emitter, by another surface, which I will dub the collector. This is depicted in Figure 9, with the radiating surface in yellow, and the collecting surface in blue.

[3] This is not strictly true – the equations used here to describe a laser beam are only true in the paraxial approximation where numerical apertures are small. For very short focal lengths, a fuller EM description is necessary. Nevertheless, this is a useful path to follow for now.

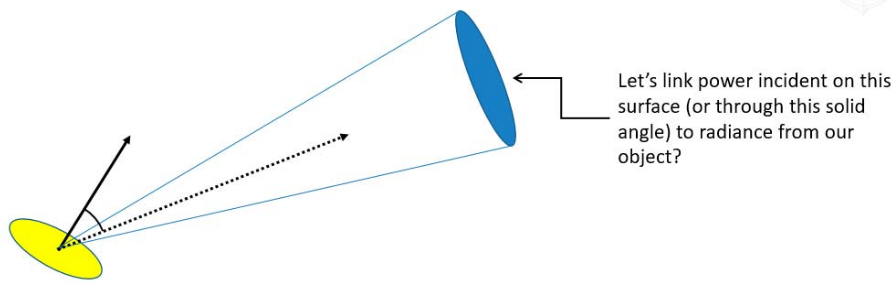

Figure 9. Model for étendue derivation: Two surfaces, an emitter in yellow and a collector in blue.

The system has two important quantities to keep in mind – the angle between the normal of the radiating surface and the collecting object (θ), and the solid angle that the collecting surface represents ($\delta\Omega$), as seen from the radiating surface. These are denoted in Figure 10, taken from [Brooker].

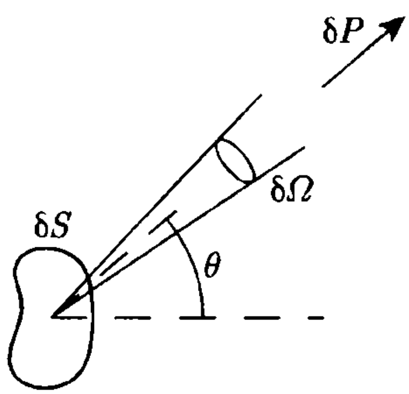

Figure 10. Model for étendue derivation including mathematical quantities [Brooker]

The infinitesimal area element δS radiates optical power with some radiance, B $\left[\frac{W}{str\ m^2}\right]$. (The importance of "infinitesimal" here is that all points on the emitter surface δS can be considered to have the same values of θ and $\delta\Omega$ for the same surface on the collector surface.)

The power δP that we can collect in some solid angle $\delta\Omega$ is proportional to $\delta\Omega$ and to $\delta S\ cos\theta$ (the area of the radiating object projected along the line of sight). The power δP collected is now given by:

$$\delta P = B \times n^2\ \delta S\ \delta\Omega \cos\theta \tag{19}$$

We define a quantity **étendue** (in a medium with refractive index n), as

$$\delta(\text{étendue}) \equiv\ n^2 \delta S\ \delta\Omega \cos\theta \tag{20}$$

So far, I have described the system with an "emitter" and a "receiver", but notice that we could treat either surface as the emitter and get the same étendue, so it is describing the **system**, not really the propagation of light in one direction.

Let us take the simpler case when the normal of the surface (with a total area of ΔS) and the central axis of the "collection cone" are collinear. If we gather light radiated symmetrically into a cone of semi-angle θ_{max} , then we have:

$$\text{étendue} =\ n^2 \Delta S \int \cos\theta\ d\Omega =\ \Delta S\ \pi(n \sin\theta_{max})^2 \tag{21}$$

Here we can recognize the numerical aperture (NA) from classical optics, $NA = n \sin\theta_{max}$, and so étendue can also be defined (in this limited case) as

$$\text{étendue} = \Delta S\, \pi (NA)^2 \tag{22}$$

If we make small-angle approximations as well, we get a useful, even simpler version:

$$\text{étendue} \approx n^2 \Delta S \Delta\Omega \tag{23}$$

This simplest version reminds us why other nomenclature for étendue includes "the $A\Omega$ product".

Etendue and focusing light

All this brings us to a very important fact: **étendue is conserved in a lossless optical system**. This is true for all perfect optical components (mirrors, lenses, changes in refractive index,…).

A simple embodiment of this can be seen in Figure 11, where a "cone" of light passes from a medium of refractive index n_1 to another of refractive index n_2. As the étendue is conserved (but the surface is the same), the change in refractive index entails a change in the solid angle of the cone.

As it is conserved and it contains information about the physical size of the beam, étendue is therefore a very important quantity when discussing the limits of focusing light.

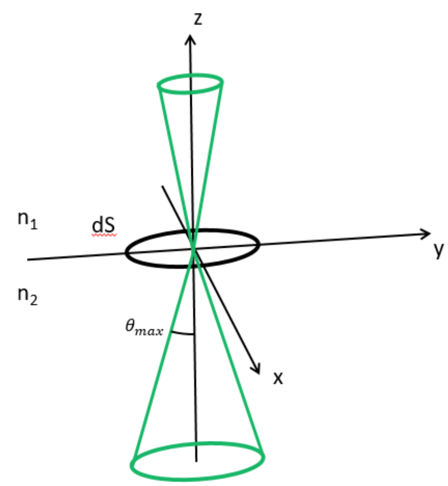

Figure 11. Demonstration of conservation of étendue for change in *n*

Essentially, my thermal source (which could be a filament, the sun, a flame) will have a physical size and emits light into the solid angle that I collect. That gives it a finite optical étendue that I cannot decrease. More accurately, I can decrease it, but only by adding losses; in other words, by throwing away optical power. I cannot arbitrarily focus the light down, as I will always be constrained by the optical étendue.

In-class example: In the lecture we will take the halogen headlight lamp example from before, calculate its optical étendue, and calculate to what size we could theoretically focus that light (after optical filtering, of course)

Optical étendue and Gaussian beams

So far, we have only needed to refer to ray optics to appreciate the limits that étendue puts on focusing. But what significance does optical étendue have for Gaussian beams? We can now use the previously obtained results to get this answer.

For the “on-axis” case, we previously linked the étendue to the numerical aperture

$$\text{étendue} = \Delta S\, \pi NA^2 \tag{24}$$

We also had the relationship between NA and the beam waist for a propagating single mode Gaussian beam

$$NA \approx \frac{\lambda_0}{\pi w_0} \tag{25}$$

If we express this rather as the minimum area at the beam waist, A_0

$$NA^2 \approx \frac{\lambda_0 2}{\pi A_0} \tag{26}$$

We get an expression for étendue for a propagating Gaussian beam

$$\text{é}tendue = \left(\frac{\Delta S}{A_0}\right)\lambda_0^2 \tag{27}$$

Equation 27 tells us that for a given propagating monochromatic beam, **étendue** informs us how many times larger it is than the minimum spot size, relative to the ideal case of a single mode. **But why would it be larger than the minimum?**

The answer lies in the fact that so far, we have only been considering the simplest, Gaussian mode. This is just one solution for a propagating EM wave. Many other solutions exist, and they form a full basis set with which we can build up any monochromatic light beam. A part of the full basis set of Hermite-Gaussian modes is presented in Figure 12, and it important to note that **photons may also occupy these modes**, not just the first transverse mode (TEM_{00}), which is the one we are discussing when we describe a Gaussian beam.

The more modes that are occupied, the larger the smallest possible spot size. The first order Gaussian mode has the smallest possible spot size.

Expressed in practice as the dimensionless Beam Parameter Product (BPP), which describes the **actual** spot size vs **minimum theoretical** size (ie for TEM_{00} alone).

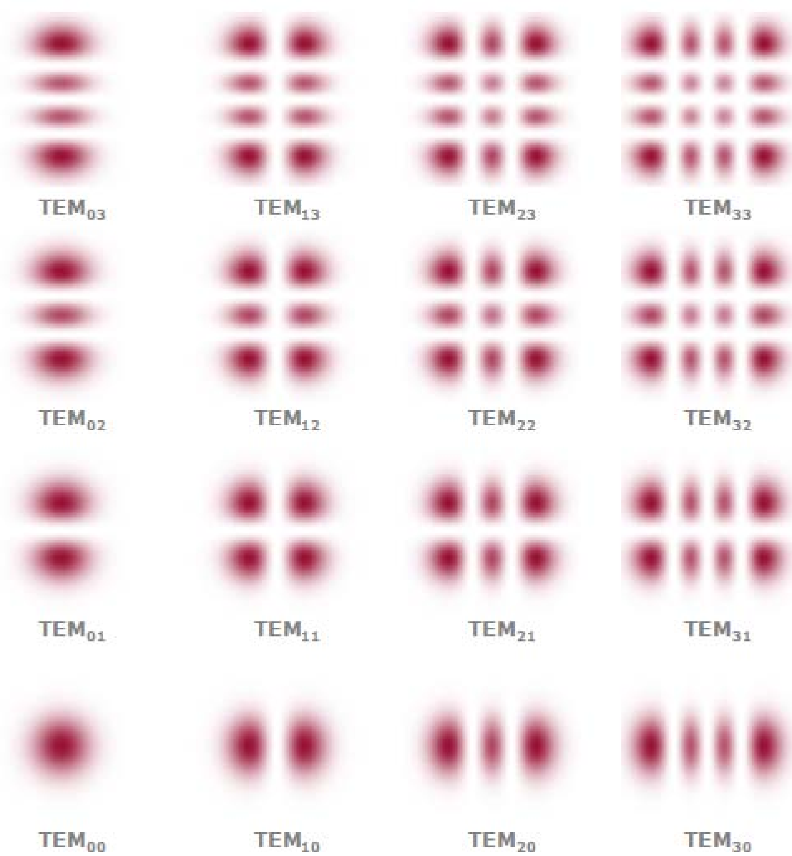

Figure 12. Intensity distribution patterns of Hermite-Gaussian solutions (modes) for a propagating EM wave [S&T].

For the Hermite-Gaussian beam profiles, the optical étendue is another way of expressing the occupation by photons of transverse propagating modes. It can be shown that (and it feels obvious):

$$\#\ \text{occupied transverse modes} = \left(\frac{\text{étendue}}{\lambda_0^2}\right) \tag{28}$$

And so…

The minimum achievable spot size (focus area) will scale with the number of occupied transverse modes, thus also with étendue.

Is Laser Light Special? Focussing Edition

We now come to the key concept that we needed to understand étendue for a laser beam: it is describing the occupation of the various optical modes. We cannot « corral » photons into a mode, this would decrease entropy. Random scattering will distribute them, increasing entropy. The best we can do is filter them out, such as by focusing a multimode laser beam onto a pinhole to remove all but the lowest order mode.

III. Final Discussion and Take-Away Messages

So finally, is laser light actually unique? No, but it is special…

Unlike electrons (which are Fermions), photons (which are Bosons) can happily all occupy the same state (in this case, a transverse mode). Lasers are special because the « stimulated » process can **selectively** put many, many photons into **the same transverse mode** (TEM_{00}) with the same phase, so they can then be perfectly focused.

Although with enough filtering (first optical to leave only a narrow bandwidth, then spatial to remove all higher order modes), we could mimic a laser beam, the losses would be spectacular and we would be left with a tiny amount of optical power.

Take-away messages:

Laser light is almost perfectly coherent:

- Stimulated emission processes leads to photons with same frequency and phase
- Coherence time, length, and spectral width are all linked
- Spontaneous emission is an unavoidable source of incoherence

Laser light can be perfectly focused:

- Minimum spot size depends on étendue, which represents the number of occupied transverse modes, and the wavelength
- Stimulated emission selectively puts photons into one transverse mode (or a few at most)