

Lecture 8: Provolution in the Beam


To accompany https://youtu.be/ktqB1-j7m_s
v1, April 20, 2022

Paul Mirsky

paulmirsky633@gmail.com, symmetryoptics.com

1 Introduction

1.1 Beam in the lens-limited configuration




Hello, and welcome to Lectures on Symmetry Optics. I'm Paul Mirsky.

Lectures 1-7 of the series belong to the introductory course in symmetry optics. This is lecture 8, and there's no longer a course – just individual lectures. The topic of this one is, provolution in the beam.

Most of the computations we did in the introductory course concerned the free configuration. This means, the space was bounded by the flat and the far limit.

Now we'll look at the lens-limited configuration, which consists of two component modes. The first mode is just like the one in the last slide, but the front flat is at the front focal plane of the lens, and the far limit is at the lens plane. The second mode is inverted front-back from the first one. Its flat is at the rear focal plane, and its far limit is also at the lens plane. The region behind the lens is called the flip field.

1.2 Fourier transform

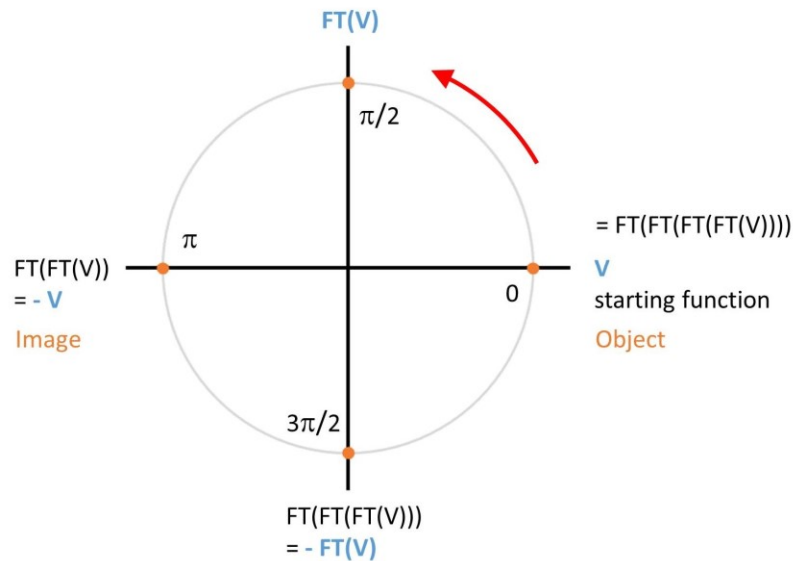


A truly remarkable fact about this arrangement is that the pattern at the rear flat is the *Fourier transform* of the pattern of the front flat. The entire field of Fourier optics is based on this one fact.

Most of the time, calculating the Fourier transform is a discrete mathematical operation, it's either all or nothing. But on the other hand, the pattern of light actually changes gradually from the front flat to the rear flat. So how can we describe this gradual change, and how do we quantify it, and model it? That's the subject of this lecture.

One way, which is conventional and is not a part of symmetry optics, is the idea of the *Fractional Fourier transform*.

1.3 Fourier transform, applied many times



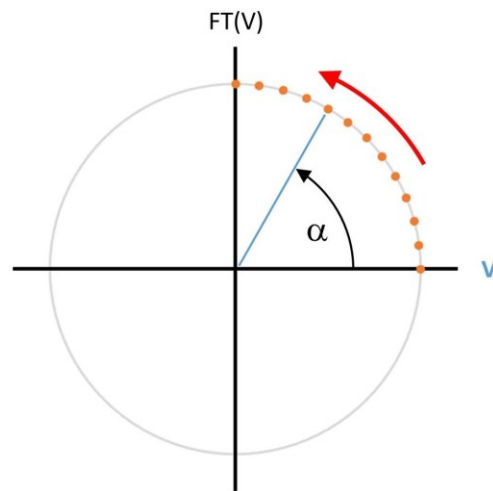
But before we introduce that idea, let's back up and put the Fourier transform into a wider context.

We begin with a pattern of light, or an object. In mathematical terms, it's a complex function of amplitude vs position. We can envision the transform as a $\frac{1}{4}$ rotation around a circle, which is an angle of $\pi/2$ radians. This dot represents the starting pattern at the front flat, and this dot represents the final pattern at the rear flat.

The lens-limited configuration only goes up to this point, but it is possible to go further. If we take the transform of the transform, or a half-rotation, we get the negative of the starting function. In optical terms this represents an *image* of the object. For most optical systems, forming an image is the entire purpose – so, this is very significant.

A third transform yields the negative of the single transform, and a fourth transform returns us to the original starting pattern. But this particular slide is just background information – we're actually interested in just a single application of the transform.

1.4 Fractional Fourier transform



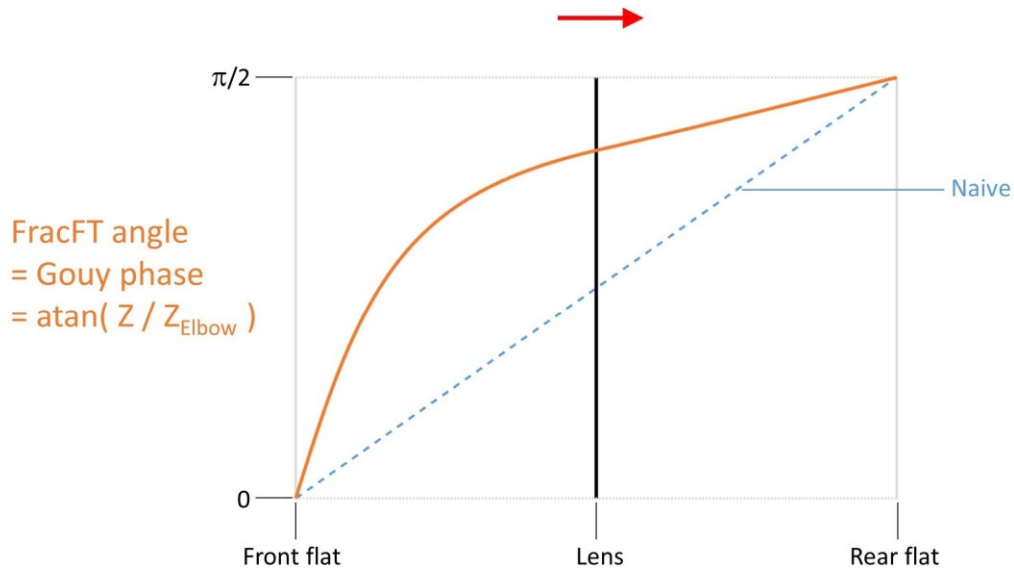
The Fractional Fourier transform extends the basic transform to include all possible angles, rather than just increments of a quarter-rotation. Simply put, a starting pattern can be transformed to any arbitrary degree, which can be quantified by an angle α . This angle is also called the *order* of the transform

Similarly, a quarter-rotation of $\pi/2$ radians can be broken up into discrete steps of any size – for example, this shows 12 equal steps of $\pi/24$.

Many authors have contributed research on this topic, but the ones I've learned the most from are Ozaktas and Mendlovic, and I cite some of their work at the end of this video.

- Ozaktas, H.M. and Mendlovic, David. *Fractional Fourier transform as a tool for analyzing beam propagation and spherical mirror resonators*. Optics Letters, Vol. 19, No. 21, 1994
- Ozaktas, H.M. and Mendlovic, David. *Fractional Fourier Optics*, J. Opt. Soc. Am. A, Vol. 12, No. 4, April 1995
- Ozaktas, H.M., Zalevsky, Zeev, and Kutay, M. Alper. *The Fractional Fourier Transform with Applications in Optics and Signal Processing*. Wiley and Sons, 2001

1.5 Fractional Fourier transform vs space



Now, let's apply this concept to the beam. The horizontal axis represents distance along the optical axis, and light is propagating from left to right. The vertical axis represents the angle of the Fractional Fourier transform.

Naively, you might suppose that the relationship follows a straight line, and that the transform angle simply equates to propagation distance – so that, for example the lens is half of the distance and so it's half of the transform angle, etc. But this is not correct at all.

Actually, it typically looks something like this curve, where most of the transform happens close to the front flat. For the beam, the transform angle is physically manifested as the *Gouy phase*. We won't discuss Gouy phase in detail here, but in brief: it's when the wavefronts are shifted along the Z axis by a small distance.

Up to the lens, the angle follows a simple equation: it's the arctangent of Z divided by Z_{Elbow} ; Z_{Elbow} is also known as the *Rayleigh range*.

1.6 Provolution

Fractional Fourier transform	Provolution
Mathematical description	Mathematical and physical description
Amount of change = $\pi/2$	Amount of change = 2
Same infinitesimal transform everywhere	Two distinct stages, two types of change

But the Fractional Fourier Transform is just one way to think about gradual change. Symmetry optics uses an analogous concept called *provolution*.


This chart shows some of the differences. First, the transform is basically a mathematical description, and while provolution is described in mathematical terms it has a much clearer physical interpretation as well.

Next, the total amount of change for the transform is $\pi/2$, while the amount for provolution is 2. These differ by a correction factor of $4/\pi$ which is about 1.27. Actually, this correction factor appears in many places throughout symmetry optics.

Last and most importantly, the transform is the same at all points. In other words, a change from 0.5 to 0.6 is the same transform as a change from 1.5 to 1.6. But in provolution there are two distinct stages where the light undergoes two distinct types of change.

At symmetryoptics.com, you can find a link to an appendix for this lecture, where I do a quantitative comparison between the Fractional Fourier transform and provolution.

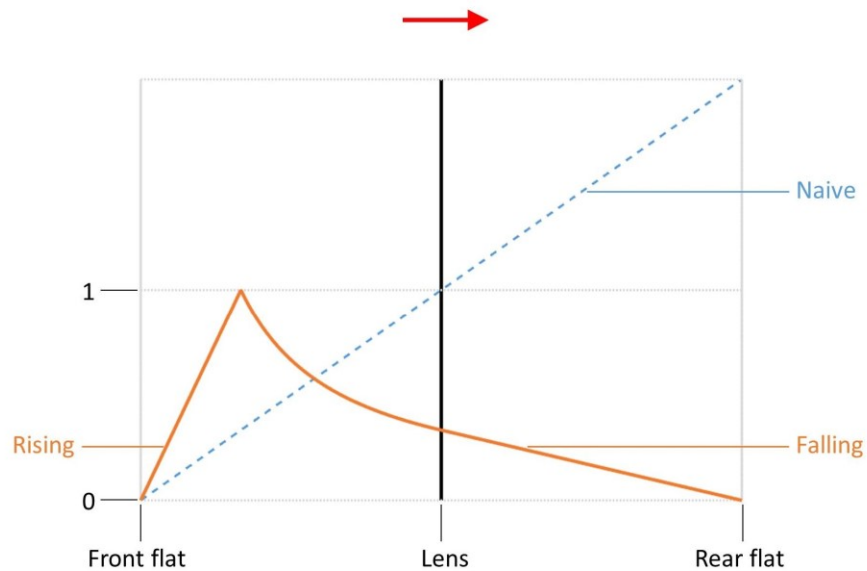
1.7 Provolution, in two stages



A few slides ago, we drew the Fourier transform as a quarter-turn around a circle. In contrast, we can draw provolution as a quarter-turn around a square. The two half-sides are the two different stages of provolution.

In the first stage, the vertical distance increases while the horizontal distance stays constant. Then in the second stage, the horizontal distance decreases while the vertical distance stays constant.

1.8 Provolution vs space



When we draw provolution in spatial terms, we see that it has a rising and a falling stage. It rises from 0 to 1, then falls back from 1 to 0. The total amount of change is 2, but it winds up back where it started.

Most of this lecture will be devoted to explaining the shape of this curve. It's a piecewise function. The first segment is a straight line that is proportional to Z . The second segment is inversely proportional to Z . And the third segment is another straight line.

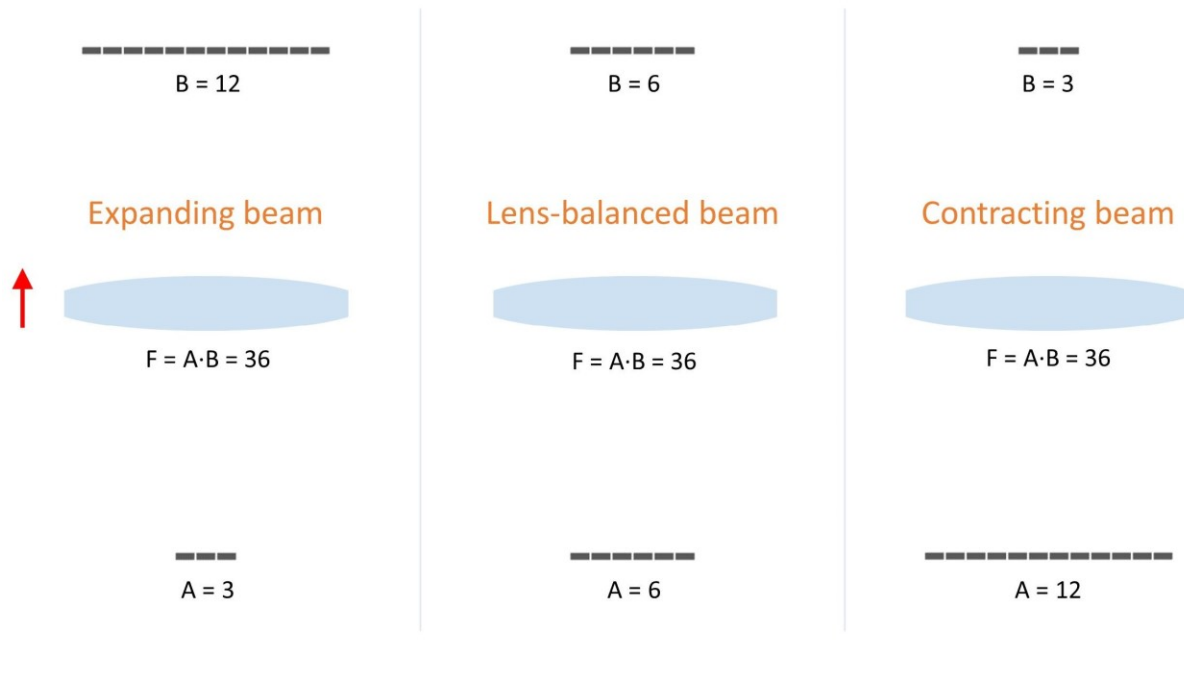
1.9 Preview of model cases

We're going to study this curve in several stages. We'll start with simpler cases, and then build on them.

- Lens-balanced beam
- Free configuration
- Lens and flip field

2 Lens-balanced beam

2.1 Beam cases, lens-limited



We begin with the *lens-balanced beam*.


Up to now we have mostly dealt with *expanding beams* like this, where the beam starts off relatively small, but it expands to be much larger. For instance, in this example the front-flat width is 3 patches, and the rear-flat width is 12 patches. The product of those two is 36, which is the lens focal length.

But coherent optics always works in reverse, meaning that we can turn it around front-to-back. So if we start with a front-flat width of 12, it will *contract* down to a rear-flat width of 3.

Intuitively, you can probably guess that in between these two cases, there must exist a beam that is *balanced* with respect to the lens. In other words, it's the same size in both focal planes. That width is the square root of the focal length. For our lens with focal length 36, the lens-balanced beam is 6 patches wide at both flats.

It turns out that the lens-balanced beam is the simplest case for studying provolution, so this is where we will start.

2.2 Lens-balanced beam

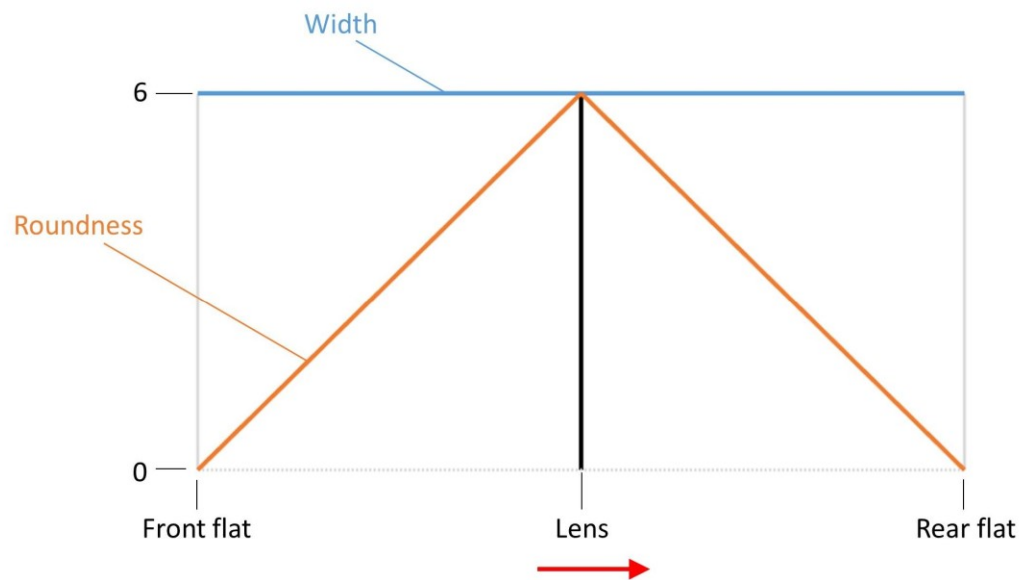


This slide shows the distribution at various planes in the lens-balanced beam. We computed these using the continuous-factor model, so the width stays perfectly constant at 6 the entire time.

The roundness grows from 0 at the front flat up to 6 at the elbow, which occurs in the lens plane.

The flip field has the same trend, but reversed. The roundness falls from 6 down to 0 at the rear flat.

2.3 Width and roundness



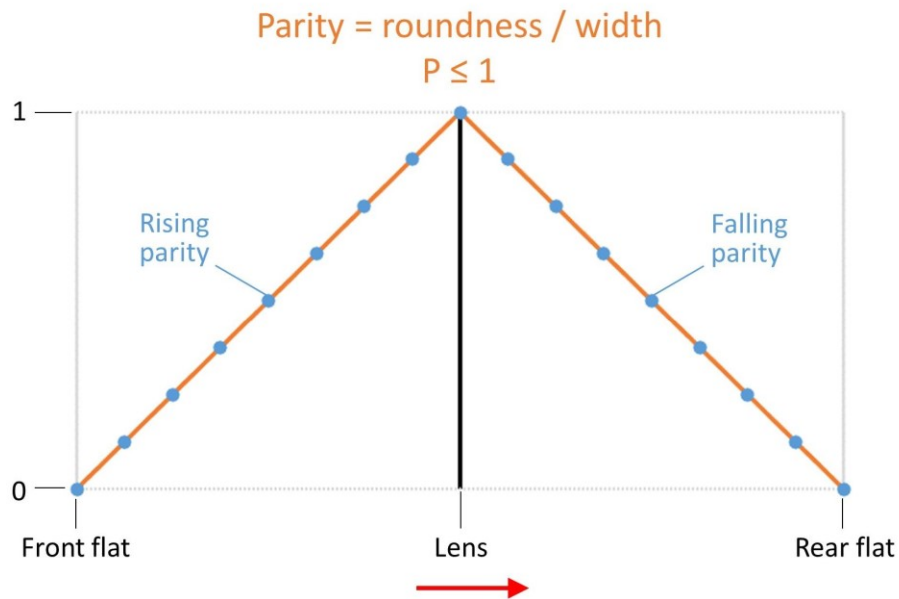
Here are several important points about the trend:

First, the roundness is initially proportional to distance from the front flat, so it's a straight line.

Second, at the elbow the width and roundness are equal.

Third, the roundness in this flip field is also a straight line, and it's proportional to absolute distance from the rear flat.

2.4 Parity



But to quantify provolution, we're actually concerned not with roundness or width per se, but rather with the *parity*. The word 'parity' means different things in different contexts, but here it's defined as the ratio of the roundness divided by the width.

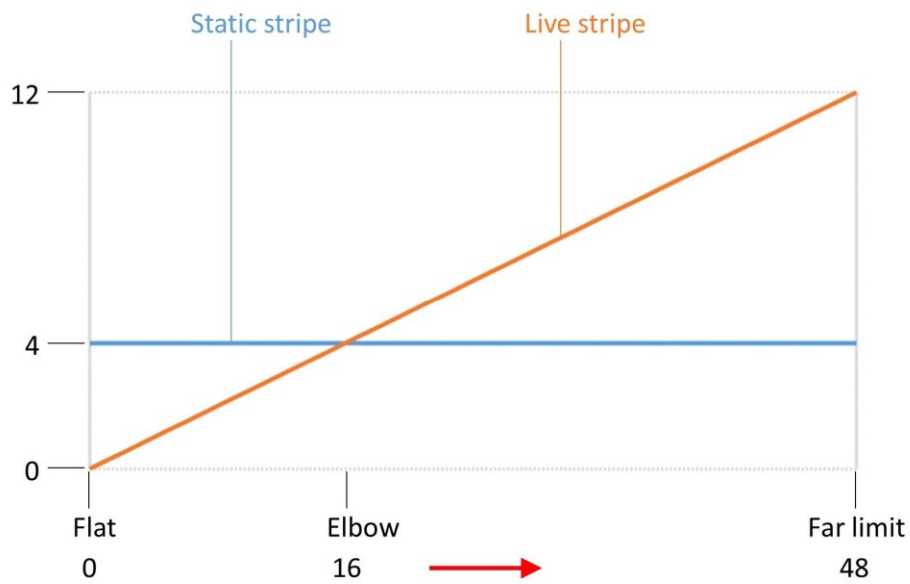
In this case it looks quite similar to roundness, but it's normalized so that its peak is at 1. Roundness is always equal to or smaller than the width, so parity can never be greater than 1.

Fundamentally, provolution occurs in steps of parity. So actually, in the one case of the lens-balanced beam, the provolution actually *does* relate to space in the naive manner. In other words, these even steps of parity are evenly spaced along the Z axis. At $\frac{1}{4}$ of the way to the rear flat, the provolution is $\frac{1}{4}$ done, etc.

This is the simplest case.

3 Free configuration

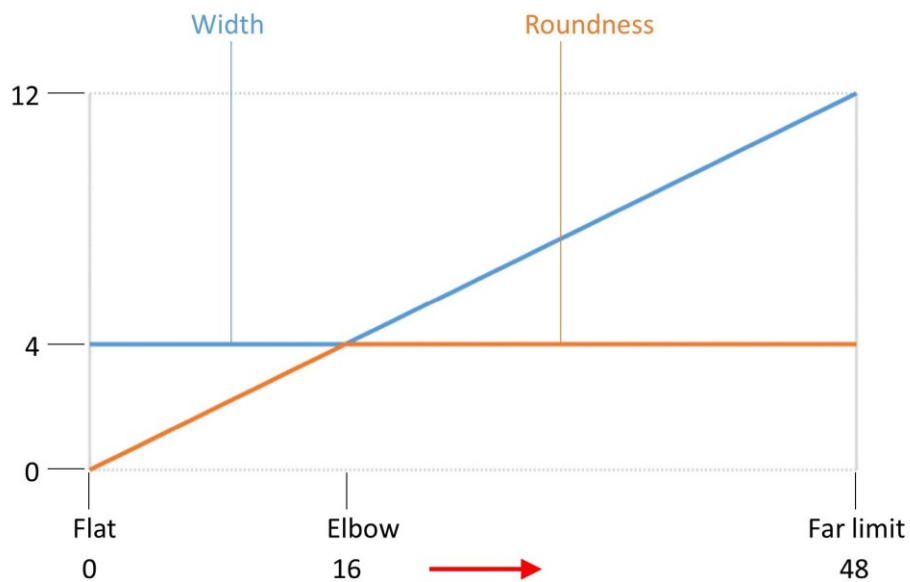
3.1 Static and live in the free configuration



Next, we'll study the *free configuration*.

First, we'll review the beam in terms of static and live. The static stripe remains constant in the free configuration. The live stripe starts at zero and grows linearly with distance from the flat. The two are equal at the elbow.

3.2 Width and roundness in the free configuration




Now let's review the same beam in terms of width and roundness.

From the flat to the elbow, the free configuration looks just like the lens-balanced beam: the width is constant, while the roundness increases linearly. The roundness here is equal to the live stripe.

But, after the elbow, the width starts to expand, while the roundness remains constant. The roundness here is equal to the static stripe.

In all planes, the roundness is smaller than or equal to the width.


3.3 Stripe ratios in the free configuration



To understand the parity, we first need to look at two different functions, which are two different ratios between the stripe widths.

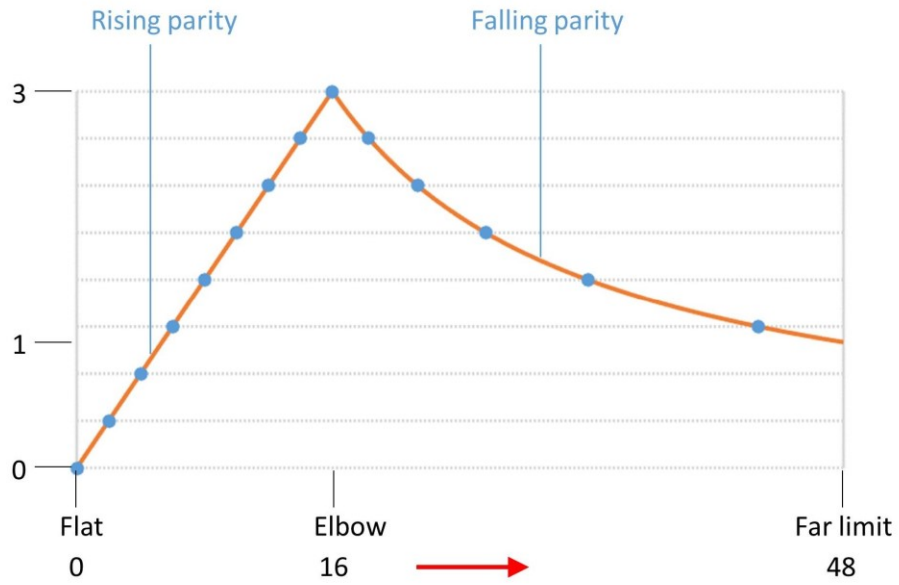
The first function is the live divided by the static, and it looks very similar to the live stripe itself except that it's normalized to 1 at the elbow. The second function is the static divided by the live. Apart from a scale factor, it's the function $1/z$. These two curves are reciprocals of one another, so they always multiply to 1 in any plane.

3.4 Parity in the free configuration




The parity is a piecewise function formed from these two. The rising parity is the live divided by the static, and the falling parity is the static divided by the live.

3.5 Steps of parity are not evenly spaced in Z



Up to the elbow, the steps are just like in the lens-balanced beam: even steps of provolution are also evenly-spaced in Z. But the shape of the falling parity curve is different, so even steps of provolution are *not* evenly-spaced in Z.

3.6 Rising vs falling steps have a different logic



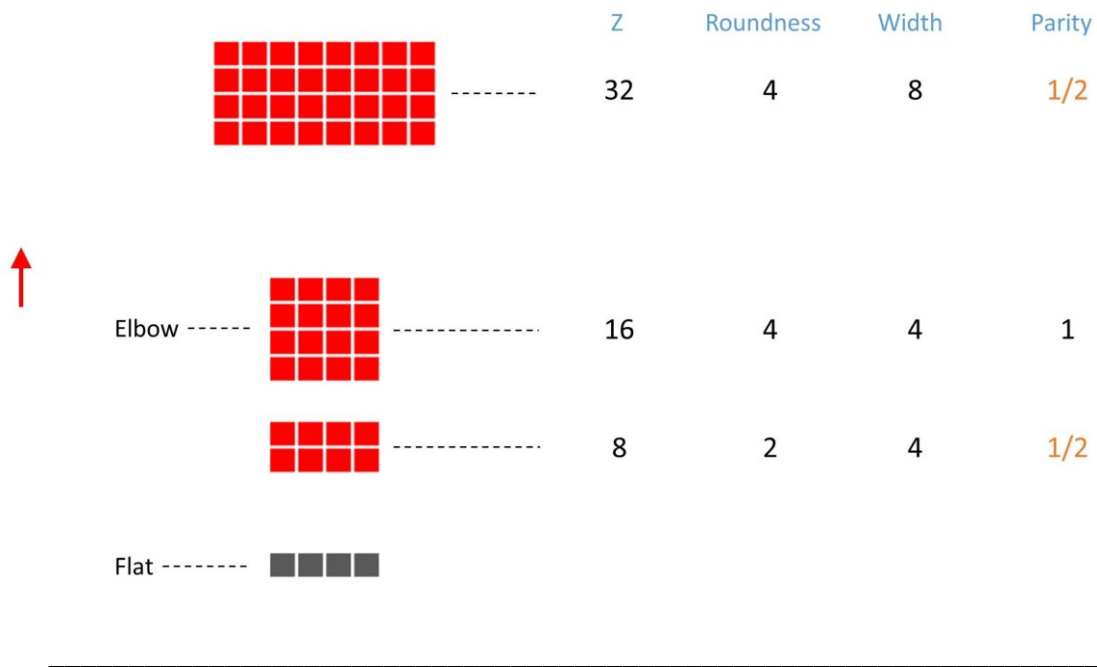
The logic of rising steps is different from the logic of falling steps.

When the parity is rising, the step distances are simply the elbow distance times the parity. So for instance, in this example the rising steps are at $1/4$, $2/4$, $3/4$, and $4/4$ of the elbow distance.

But when parity is falling, the step distances are the elbow distance *divided* by the parity, which is the same as multiplying by the inverse of the parity. So the steps are at $4/4$, $4/3$, $4/2$ and $4/1$ of the elbow distance.

Another way to think about it is that at 2 times the elbow distance, the parity is $1/2$. At 4 times the elbow distance, the parity is $1/4$, and so on for any multiple.

3.7 Same parity, different planes



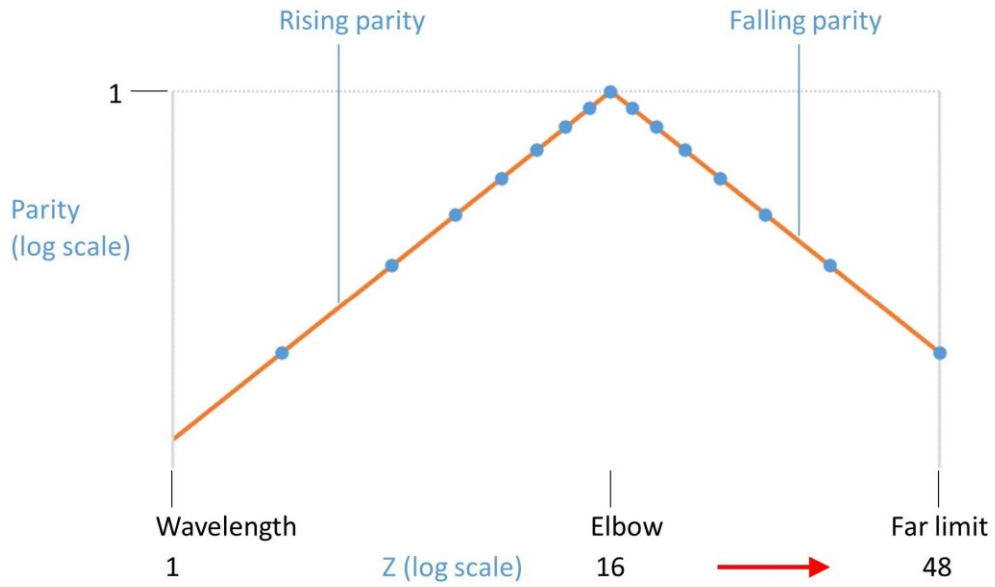
This shows how different distributions can have equal parity.

The elbow occurs at $Z = 16$, where the roundness and the width are both equal to 4, so parity is equal to 1.

This plane is at $Z = 8$, which is $\frac{1}{2}$ of the elbow distance. The roundness divided by the width is $\frac{2}{4}$, so the parity is $\frac{1}{2}$.

This plane is at $Z = 32$, which is 2x the elbow distance. The roundness divided by the width is $\frac{4}{8}$, so the parity is also $\frac{1}{2}$.

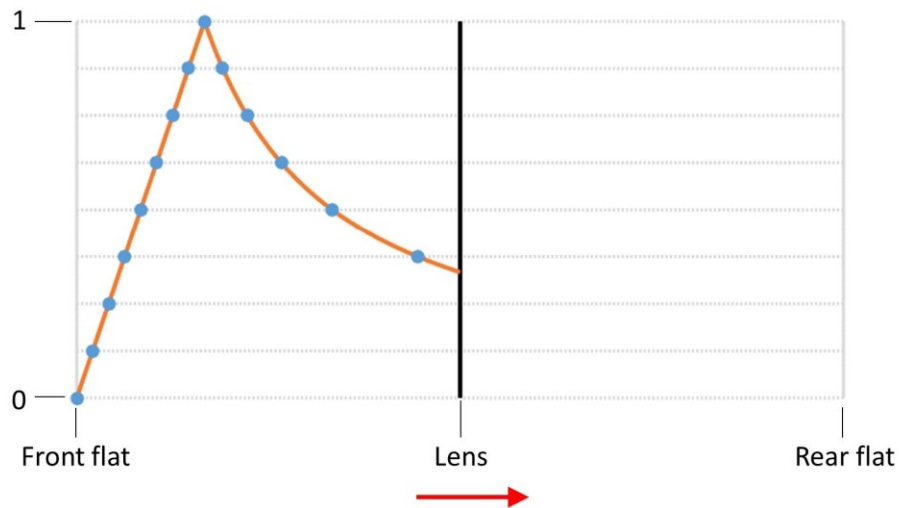
3.8 Parity on a logarithmic scale



One more thing: the falling parity may seem to have an odd shape, but it looks more natural when the axes are logarithmic. Then, the falling parity appears as a straight line. Note, this is without a lens. Sometimes this can be useful.

4 Lens and Flip Field

4.1 Parity cannot fall to zero without a lens

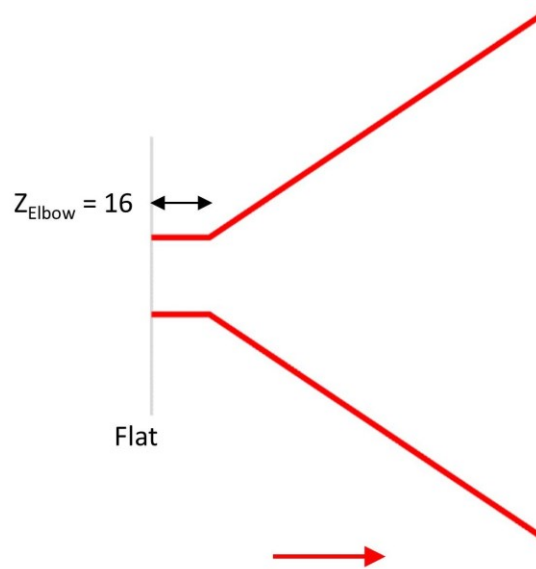


The final portion of our analysis concerns the lens and the flip field. We're now considering the lens-limited configuration, but without the restriction that the beam is balanced to the lens.

The first section, before the lens, is exactly like a free configuration. If the focal length is much longer than the elbow distance, the parity may fall close to zero. But it never gets all the way to zero, because that would require an infinite distance. So, it's not possible to completely finish the Fourier transform without a lens.


Now we'll discuss the effect of the lens

4.2 First beam mode



First, consider this beam propagating to the right. It's expanding as it propagates, and in the absence of any lens it would expand indefinitely. In this example, the elbow distance is 16.

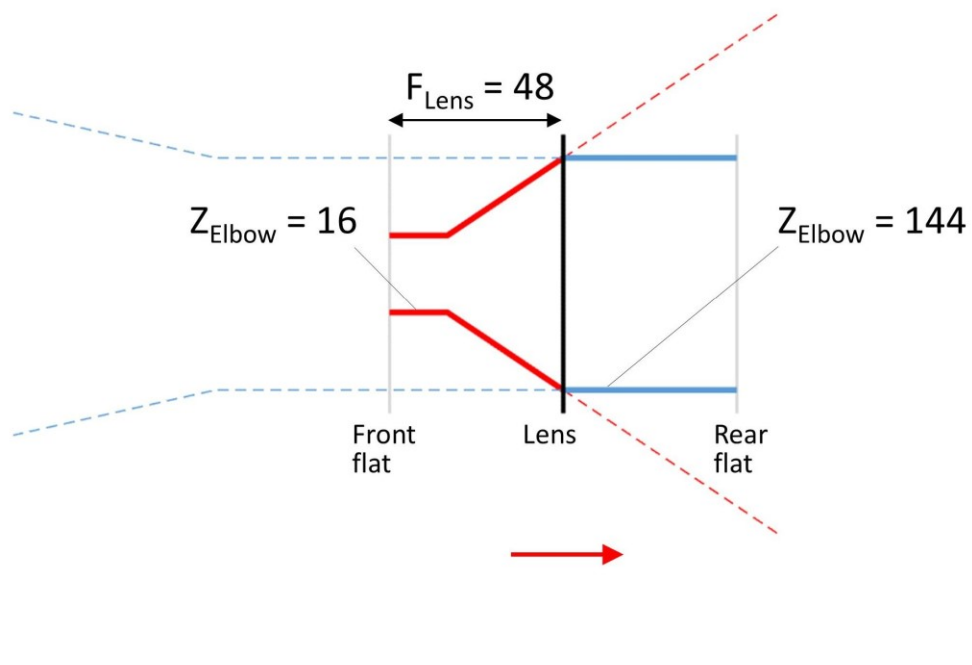
4.3 Second beam mode



Now consider a different beam, whose flat is in a different plane. This beam is wider, with an elbow distance of 144.

It's also oriented backwards. You can interpret this in either of two ways: You can think of it as a beam propagating leftward. Or, even better, you can think of it as a focused or converging beam, propagating rightward and forming a beam waist. Everything in coherent optics can be turned around this way.


4.4 Lens changes one beam mode into another



The lens acts to change the first beam into the second beam. The first beam is in its far field and the second beam is in its near field.


Also, the focal length in this case is 48. So, it's a factor of 3 *larger* than the elbow distance of the first beam, and it's a factor of 3 *smaller* than the elbow distance of the second beam.

4.5 Flip field



In the flip field, the width stays constant but the roundness decreases, and it reaches zero at the rear flat.

4.6 Parity in free vs lens-limited configurations




For every plane in the free configuration, there exists a corresponding plane with the same parity in the lens-limited configuration.

Before the lens plane, the two beams are the same, so all planes correspond exactly.

But after the lens plane, the correspondence follows this simple rule: first, take the ratio of the distance to the focal length. For instance, this plane is 4 focal lengths from the front flat. Then, take the reciprocal of that ratio. In this case, the corresponding plane is located $\frac{1}{4}$ of a focal length from the rear flat. Both planes have the same parity, which in this case is $1/12$.

You can also calculate the corresponding plane using the standard imaging equation.

4.7 Parity falls to zero at rear flat

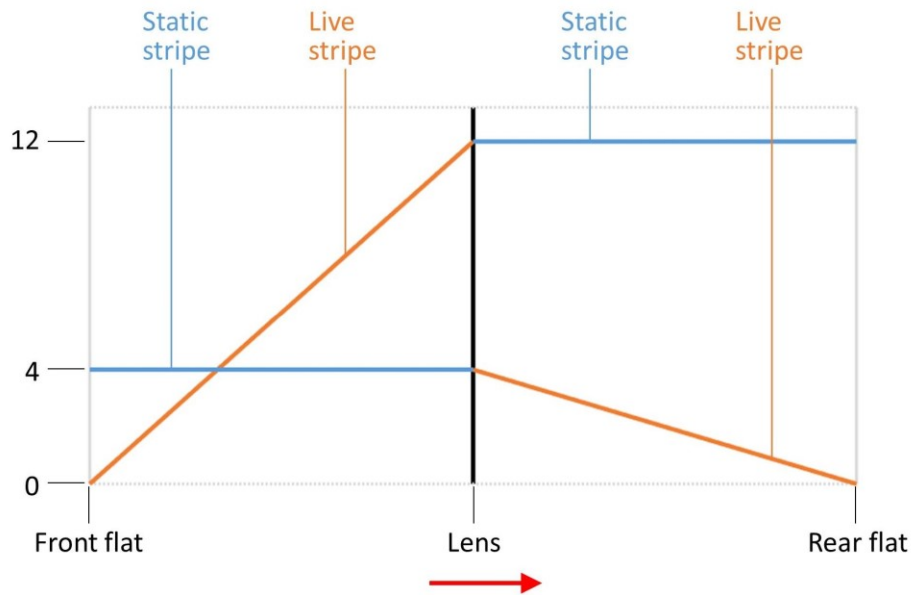


Either way, the result is a straight line in the flip field. It's actually very similar to this region, because they are both near fields, but they belong to two different beam modes. The slope is shallower in the flip field because that mode is much wider. The parity reaches zero at the rear flat, and the Fourier transform is always completed precisely in that plane.

Note that merely passing through the lens does not change the parity, so parity is a continuous function. Moreover, the *slope* of the parity is also continuous at the lens.

Also, note that we've been assuming that the input beam is narrower than balanced. If the input beam is wider than balanced, then the curve is the same basic shape, but simply turned around.

4.8 Lens exchanges static and live

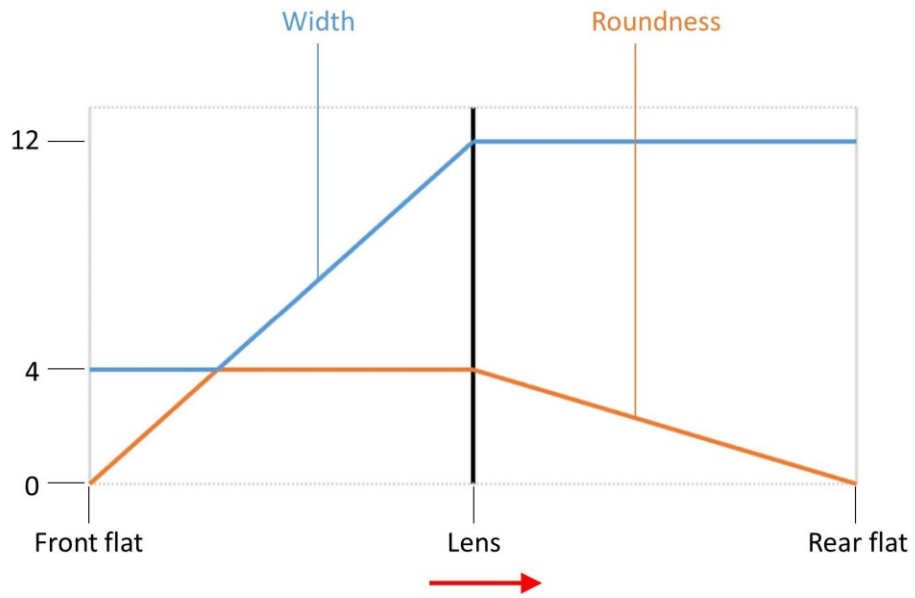


The principle of the lens is that it changes static into live, and vice-versa.

In this example, just before the lens, the static stripe is 4, and the live stripe is 12. The lens instantaneously exchanges these values, so that the static stripe is 12 and the live stripe is 4.

The other effect of the lens is to change the beam from diverging to converging.

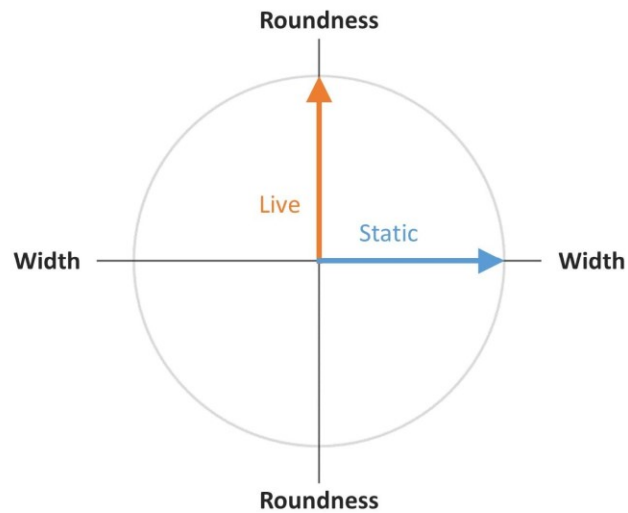
4.9 Lens does not affect width or roundness



Note that while static and live are exchanged at the lens, width and roundness both stay constant.

5 Inclination

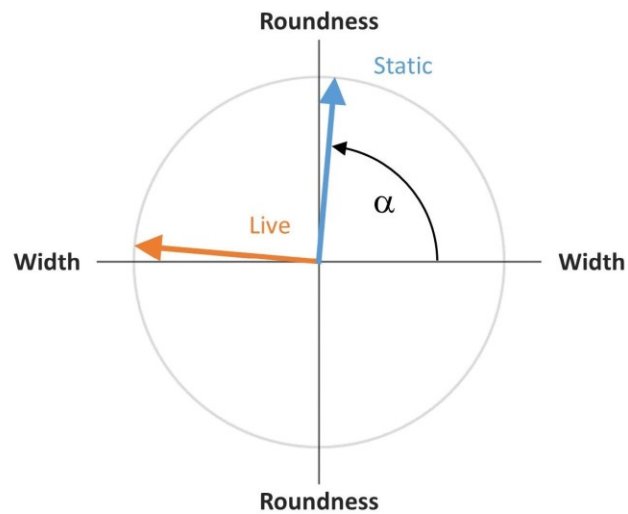
5.1 Provolution and inclination, front flat



This leads us to the subject of *inclination*. The inclination of a factor (*incorrectly read as 'feature' in the video*) is whether it corresponds to the width of a pattern, or to the roundness, or somewhere in between.

This diagram represents the beam at the front flat. The axes are width and roundness, and the static and live stripes are vectors. The static stripe inclines to width, because it appears as a linear distance. The live stripe is actually zero at the flat, but as soon as it does appear, it corresponds to the roundness.


5.2 Provolution and inclination, far field



After some provolution, neither stripe corresponds directly to width or to roundness. Instead both incline somewhere in between, which we can roughly quantify by the 'angle' of provolution, but it's not a clear physical angle that you could measure with a protractor.

In the far field of the free configuration, the live stripe inclines almost exactly to width, and the static stripe inclines almost exactly to roundness.

5.3 Provolution and inclination, flip field

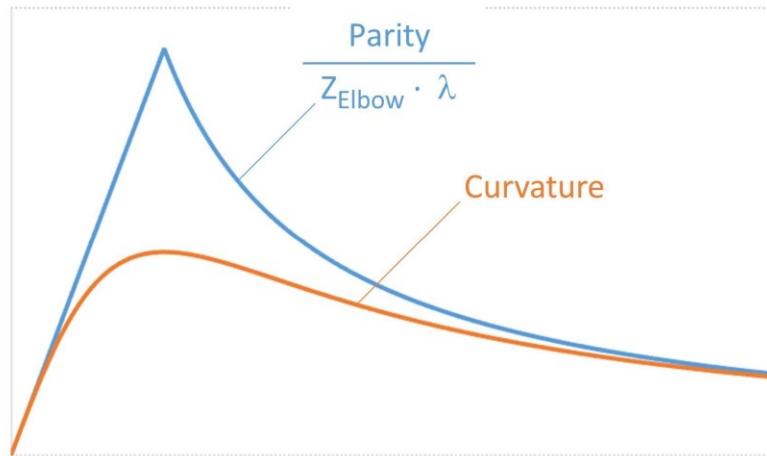


But for those inclinations to become exact, the beam needs to pass through the lens, which exchanges the inclinations of the two stripes. Then the static stripe inclines almost to width, and the live stripe almost to roundness.

And at the rear flat, the provolution is complete. The static inclines to width, just like at the front flat.

6 Curvature

6.1 Curvature vs scaled parity, free

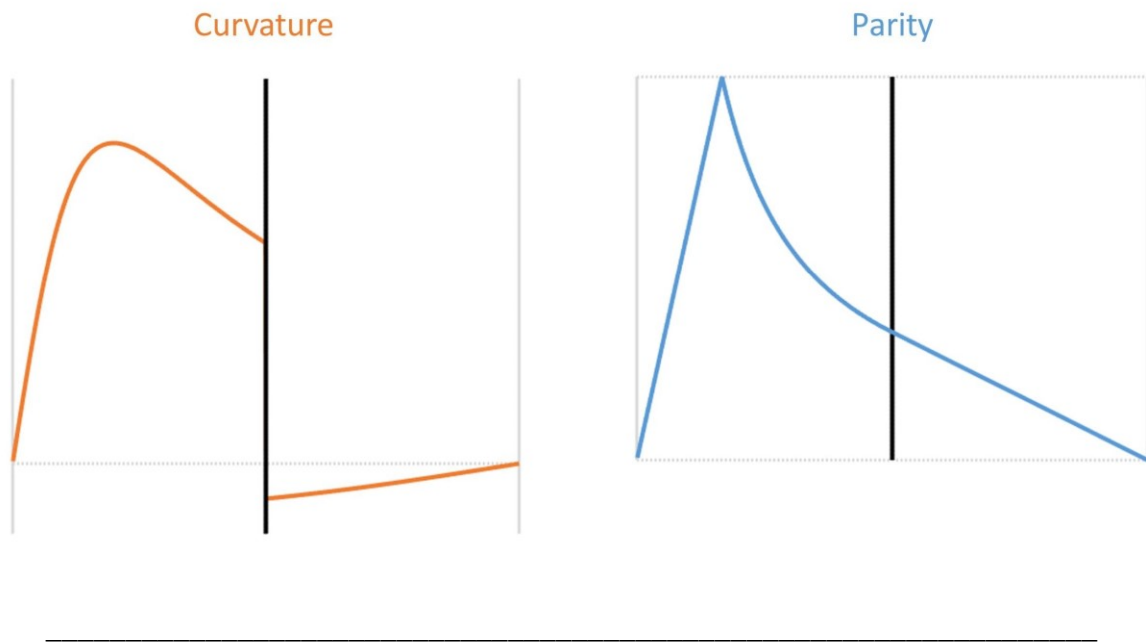


The final topic in this lecture is *curvature*. In conventional optics, a beam wavefront is shaped like a small portion of a sphere. The wavefront curvature is the inverse of the sphere's radius.

In symmetry optics, parity is a more fundamental idea than curvature, but the two are somewhat related. This diagram shows two functions: the orange is the curvature of the beam in the free configuration. The blue is the parity, scaled by a factor of the elbow distance and the wavelength.

Both functions have a peak at the elbow, but the scaled parity is exactly twice as large. Near the flat, and deep in the far field, the two functions are asymptotically close.

6.2 Curvature vs scaled parity, lens-limited



In the lens-limited configuration, even more differences are apparent. First of all, the curvature is negative in the flip field, but there is no such thing as negative parity.

Second of all, because the lens changes a narrow beam mode into a wide beam mode, the scaling factor doesn't stay constant. Therefore, the magnitude of curvature also decreases by passing through the lens.

More details can be found in the appendix.

7 Conclusion

7.1 Reviewing key points

That's all for this lecture, so let's review the key points:

- Provolution is the continuous change of light from one pattern to its Fourier transform
- Parity in a given plane is equal to roundness divided by width.
- The progress of provolution is quantified by parity rising from 0 to 1, then falling from 1 back to 0.
- The lens changes static into live, and vice-versa.

7.2 Thanks for listening

I hope you've found this class informative and interesting. To learn more about symmetry optics, please check out www.symmetryoptics.com.

If you have specific questions about this or other lectures, please post them on reddit, at www.reddit.com/r/symmetryOptics/, and I'll try to answer them.

This is a new field, and there's a lot of opportunity to discover new science and develop new applications. I hope you'll take advantage.

I'm Paul Mirsky, thanks for listening.