

Glow and Shine in Feature Optics

Paul Mirsky

paulmirsky633@gmail.com

May 20, 2020

Abstract

Optical features exist in two types: *Glow* features remain constant, and describe a pattern of light at a flat plane. *Shine* features expand in proportion to their distance from the flat plane. Glow and shine have inverse feature chains. Glow and shine are convolved to produce the patterns observed in the critical planes of the grating.

1 Introduction

Feature Optics (FO) was recently proposed¹ as a new framework for studying coherent light in the diffraction grating. In reference 2, we showed that FO provides an apt language to describe the trends in the critical planes of the grating. However, that work simply presented ad-hoc rules that roughly agree with experiment – it did not give any underlying theory explaining why those particular trends occur.

In this work we propose a set of physical principles that connect and explain all the trends observed in the critical planes. We define new basic concepts including glow, shine, and depth. We describe an algorithm which takes any grating-like flat pattern as input, and computes the patterns in the critical planes.

This article presumes that the reader is completely familiar with references 1 and 2. We first introduce the fundamental concepts of glow and shine. Then, we apply them to the simple beam. Next, we apply them to the grating, progressing through all of the critical planes.

2 Principles

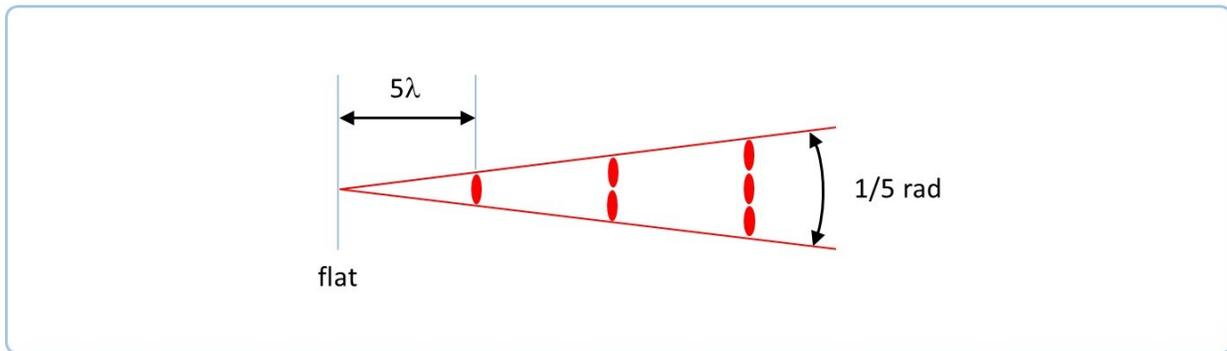
2.1 Glow and shine

In earlier work, we considered only two types of features: bright and dark. We now introduce a new distinction between *glow* and *shine* features. A given feature has one type from each trait; for example, it might be a bright glow feature. Glow features describe a pattern with *flat* wavefronts, as we typically specify the initial conditions. Shine features describe the light that

emanates from the glow. The two types interact to create the patterns at various planes, as we will show.

As light propagates through space, glow remains constant but shine changes. As Figure 2.1 shows, the shine forms a cone of light with its apex at the flat plane. The width of the cone is proportional to the Z-displacement from the flat; the constant of proportionality is the divergence angle (given in radians). Equivalently, the divergence is the inverse of the Z-displacement (given in wavelengths) needed to sprout one patch. Note that negative shine can also exist to the left of the flat, but is beyond the scope of this work.

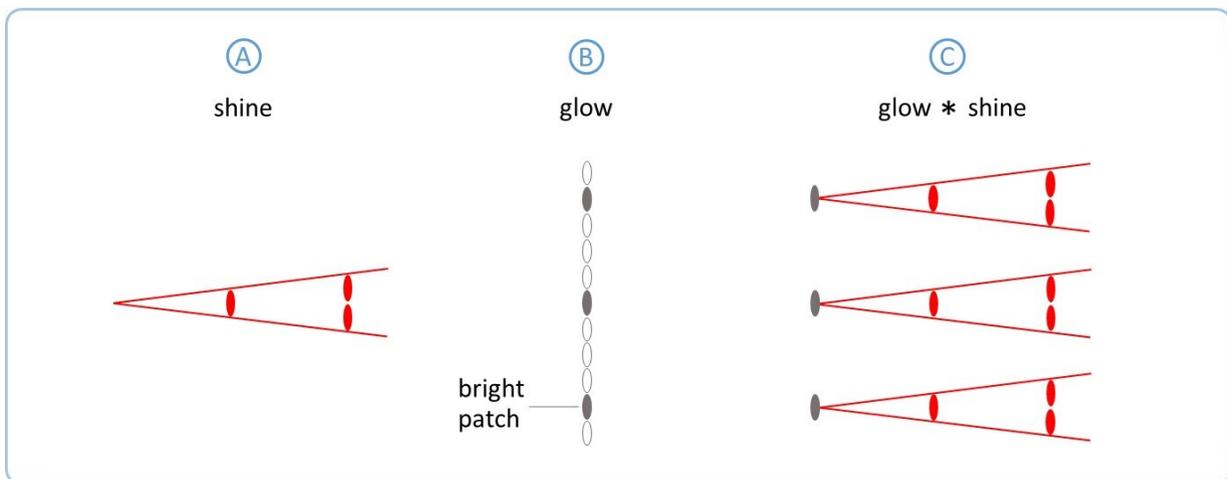
Figure 2.1, Shine



Shine divergence is similar to far-field beam divergence. But the shine diameter decreases to zero at the flat, while the beam has some nonzero waist size.

Metaphorically, the glow patches are like 'seeds' and the shine is like the 'flower' which grows from each seed. The glow sets the starting pattern; each and every bright patch of the glow emanates one *instance* of the shine, as Figure 2.2 shows. Note that shine does *not* emanate from dark glow.

Figure 2.2, Shine and glow

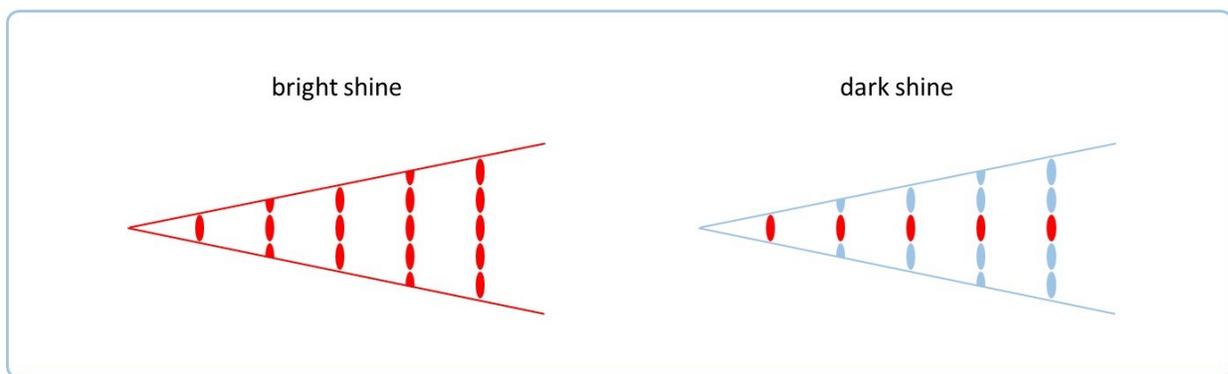


By tabulating all of the shine received at each point, we obtain the pattern for that plane. Equivalently, we convolve the glow and the shine. Further details will be given below.

2.2 Bright vs dark shine

Just as glow exists in both bright and dark types, there also exist bright and dark shine. These two are contrasted in Figure 2.3. Bright shine consists of a bright feature whose size increases with displacement in Z; the examples in the previous section were of this type. Dark shine is analogous; it is a dark feature of increasing size. Dark shine consists mostly of dark, empty space; it contains only a single bright patch.

Figure 2.3, Bright and dark shine



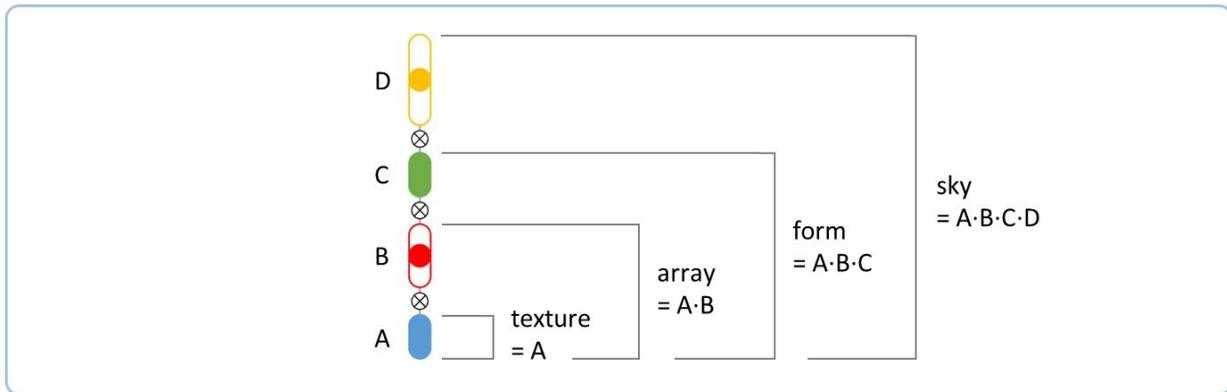
Note that bright and dark features are the same from the apex up to the sprouting of the first patch, because a dark feature always contains exactly one bright patch.

2.3 Shapes

The patterns in all grating planes share the same basic structure: a periodic array of identical bright regions, with dark space in between. FO usually describes such patterns in terms of a ranked chain of four features (from highest to lowest rank, their types are dark, bright, dark, bright). But they can also be described in terms of the *shapes* texture, array, and form¹.

In the most general possible case, a shape may be an arbitrary function of position, such as a star or a letter. But in the elementary case of the grating which this work studies, the shapes are simply rectangular functions of various widths. As Figure 2.4 shows, each shape width is the product of a subset of the feature chain.

Figure 2.4, Texture, array, form, and sky shapes



The fourth shape, *sky*, is more abstract than the others because it represents the border of the space under consideration. While there may exist no physical border, we always consider planes at some finite distance from the flat, which effectively defines a limit. To model a freely-propagating system, we can make the sky arbitrarily large. In a lensed system, the sky is usually the lens focal length.

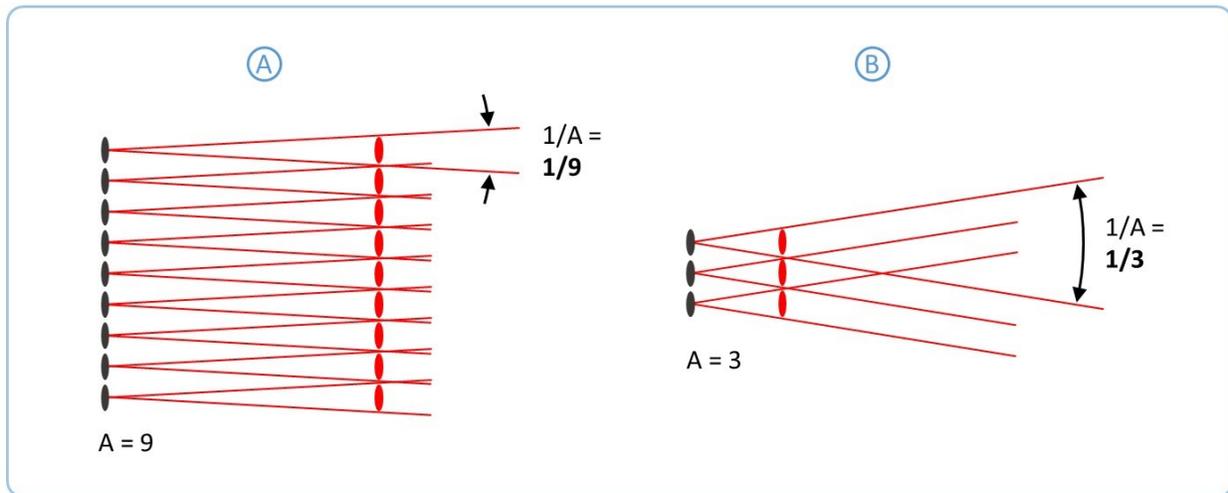
The highest-ranking dark feature (labeled D) is called the *free sky*. It represents the amount of empty, dark space which must be included to make the pattern as large as the sky. Often, feature and spatial diagrams will omit the free sky, because it takes up too much room in the diagram and often plays a very passive role. But, it exists implicitly.

The concept of shape applies in 3 distinct contexts: glow, shine, and pattern. These will arise in later sections.

2.4 Shine divergence and components

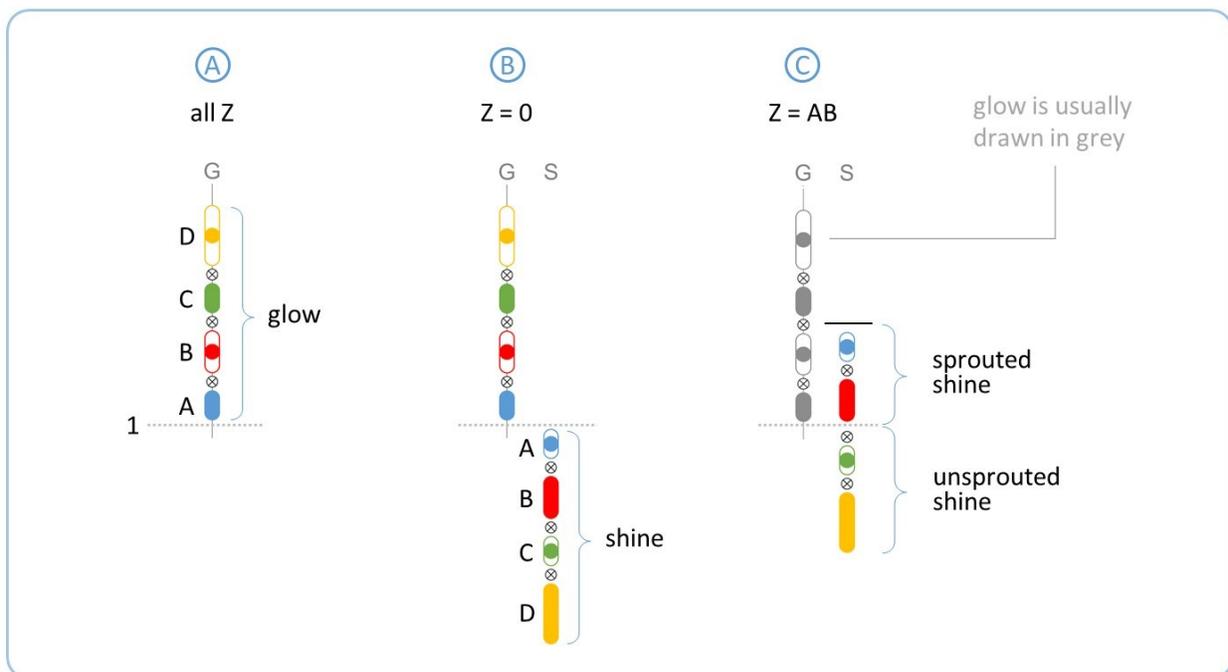
The structure of the shine is determined by the structure of the glow. In the case of the simple beam the rule is very simple, as shown in Figure 2.5: the glow is the beam waist width A (in patches), and the shine divergence is the inverse of the glow (in radians). The larger the glow, the smaller the divergence.

Figure 2.5, Glow width and shine divergence



However, this rule for the simple beam is actually only a special case of the more general rule which governs the grating. Figure 2.6a shows a feature diagram of the glow. It consists of 4 ranked features labeled A, B, C, and D from lowest to highest. A and C are bright features, B and D are dark features. The letter G at the top of the column indicates that it represents the glow.

Figure 2.6, Shine derived from glow



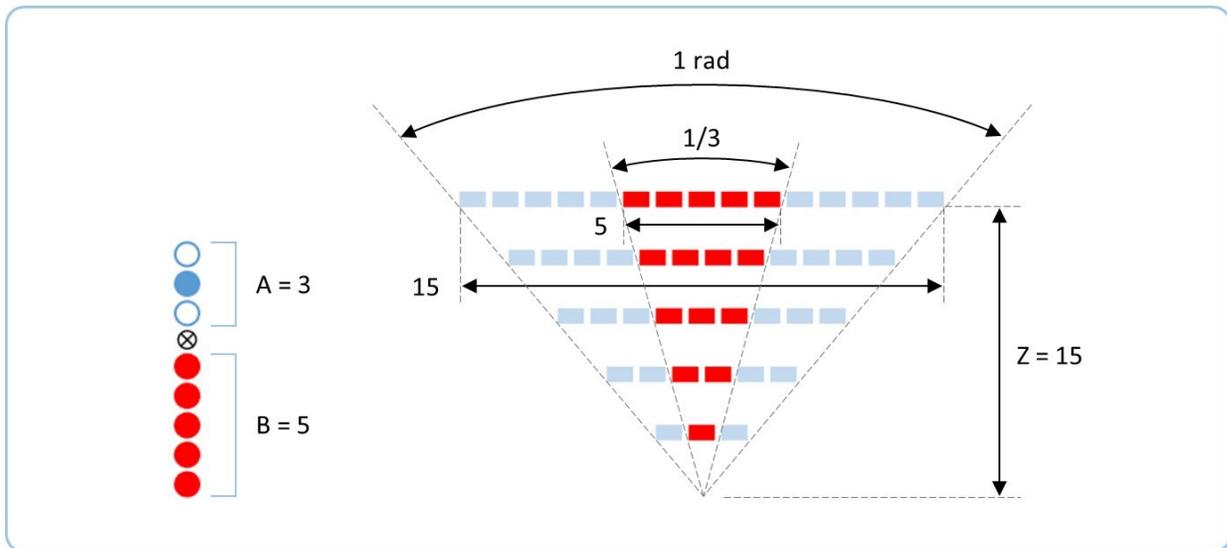
In Figure 2.6b we add a second column, capped with the letter S to indicate that it represents the shine. The shine is the inverse of the glow, meaning two things: first, the shine chain has the same feature sizes as the glow chain, but with the rank order reversed so that A is ranked highest, rather than lowest. Second, each feature is of opposite type from its corresponding

glow feature, so that feature A is dark rather than bright, etc. Note that the adjacent G- and S-columns refer to the *same* plane, even though they are drawn some horizontal distance apart in the diagram. Finally, note that diagrams with G- and S-columns are *not* the only type of 2-column feature diagram; other types exist, but lie outside the scope of this work.

Figure 2.6b represents the flat, where the shine is has not yet sprouted; there, the highest-ranking shine still ranks below the lowest-ranking glow. However, the shine *rises* in rank as the light propagates away from the flat. Figure 2.6c shows the plane $Z = AB$. The shine expands in a cone of 1 radian; so, at distance AB from the flat, the shine sky is also AB patches in width. Note that the glow is now drawn in grey; we will follow this convention going forward.

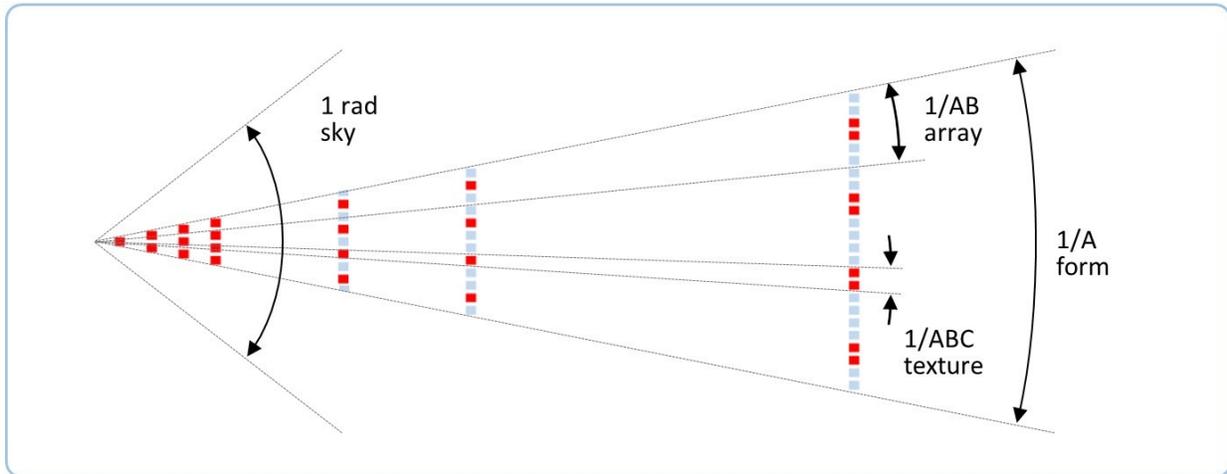
The shine features must fit within this 1-radian cone; this determines the divergences, as Figure 2.7 shows. We compute the plane at $Z = 15$, so the shine sky is also 15. The shine contains a free sky of size 3 ranked above a bright feature of size 5. The bright region is 5 patches in width, filling a divergence angle of $1/3$ radian. (Note that we usually omit the free sky from such diagrams).

Figure 2.7, Divergences from feature diagram



We can extend the above reasoning to calculate divergences for all shine shapes, as shown in Figure 2.8. We use the symbols A, B, and C to refer to the *sizes* of the corresponding features; all angles are given in radians.

Figure 2.8, Shape divergences



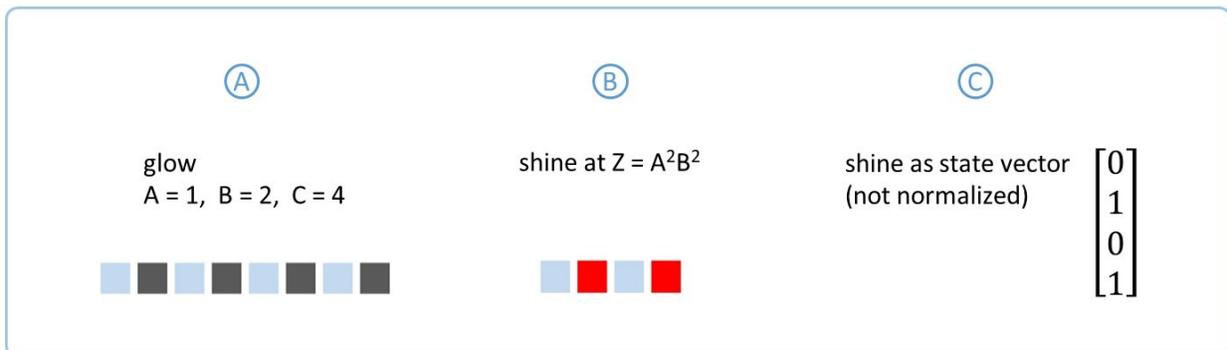
In the shine the shapes have *angular* sizes or divergences which remain fixed for all Z . At any particular Z , the shine shapes also have *linear* sizes. In contrast, glow shapes have linear sizes only.

2.5 Computing tracked-shine diagrams

From any grating pattern specified at the flat, we can use the foregoing principles to compute the pattern in any plane. The computation is conceptually simple and can be performed by hand for very simple cases, although a computer is required beyond the very smallest systems.

In this section we work through an example case, showing the details of the algorithm. We begin by specifying the glow pattern. For this example we choose a grating with very small feature sizes, as shown in Figure 2.9a. By convention, the bright glow patches are drawn in dark grey rather than red, to differentiate them from shine patches. Next, we choose the plane in which to compute the pattern; here, we have chosen the plane at $Z = A^2B^2$.

Figure 2.9, Example glow and shine

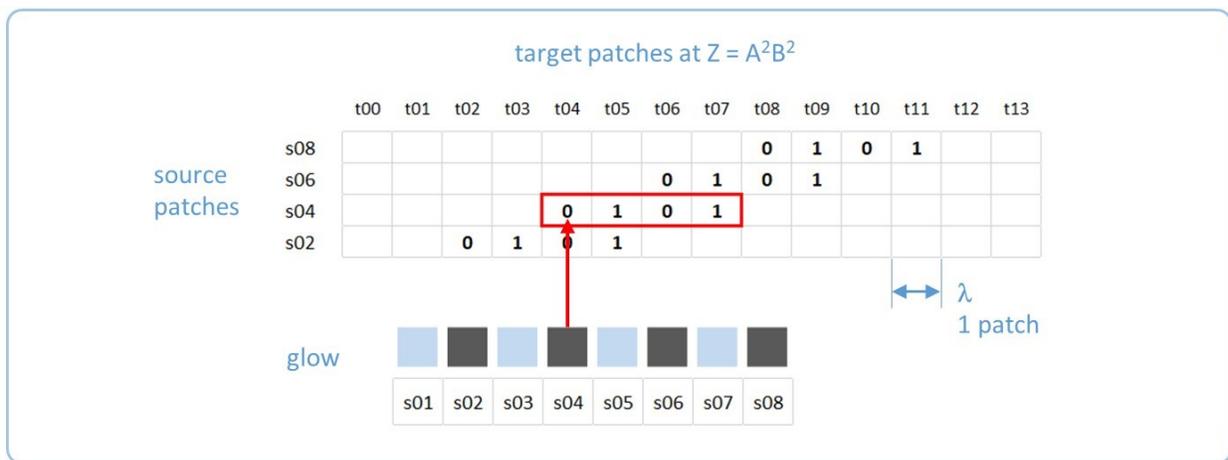


The shine divergences are calculated by inverting the glow sizes, as described above. Those divergences are multiplied by the distance Z to yield the shine pattern, shown in Figure 2.9b. In

this case the pattern is four patches wide; the array has sprouted but the texture has not (see next section). The shine can also be expressed as the state vector shown in Figure 2.9c. The zeros correspond to dark shine, and the ones correspond to bright shine; to minimize clutter, we leave the bright entries at 1 rather than follow the convention of normalizing every state vector to a total length of 1.

For the next step we create a 2-d array of cells, as shown in Figure 2.10. Each column represents one *target patch* – that is, one wavelength-sized region in the plane where we are calculating the pattern. Each row represents one *source patch* – that is, one of the bright glow patches from which the shine instances emanate. Each cell has a row and a column; its contents specify the light that *emanates* from one source point and is *received* at one target patch.

Figure 2.10, Cells by source and target



As stated earlier, one instance of shine emanates from each bright patch of glow. This is represented by writing one copy of the state vector into the cell array for each source patch. In Figure 2.10, one instance of shine is highlighted in red. It is written in the second row, corresponding to its source patch s04. It is written in target patches (columns) t04 through t07; the leftmost target patch t04 is aligned horizontally with the source patch s04.

Note that we draw a distinction between *empty* cells – which represent background space – and cells containing zeros, which represent dark shine. While both of these represent dark patches of space, we track them separately in the computation.

The tracked-shine diagram, shown in Figure 2.11, fundamentally contains the same information as the cell array in Figure 2.10, but it is formatted differently for clarity. These differences are:

- The cells are rendered as colored rectangles rather than as numbers (all of which are either 1 or 0 anyway). Each instance of shine is drawn with a different color for bright patches; all dark shine is drawn in light blue.
- Most tracked-shine diagrams show many planes. The vertical axis of the diagram is used in two ways: over longer distances it represents the Z (propagation axis) position of each plane.

But within a given plane, the different rows of the cell array are drawn vertically staggered, even though they are all actually in one physical plane.

- All patterns are drawn centered in the horizontal direction, so the horizontal alignment of different planes does not necessarily match the calculation. FO neglects this degree of freedom, not making any predictions.
- We do not draw the free skies for glow or shine, which take up a lot of space and are not important for this computation. The current example does not show this clearly because $A = 1$ and thus the shine free sky is trivial anyway.

Figure 2.11, Tracked-shine diagram



The number of bright cells in a given column is called the *depth*. It is the number of different source patches that emanate shine to that target patch. Depth is not observable in an experiment. While we will necessarily compute depth in every tracked-shine diagrams, we will not treat depth in this work; it will be addressed in future work.

If a target patch (column) contains only dark shine, it is observed as dark. If it contains bright shine or a mixture of bright and dark, then it is observed as bright. Thus, this algorithm yields observable predictions which can be compared with experiment.

We have followed the procedure of placing an instance of shine at each bright patch of glow. However, we would obtain the same result if instead we were to place an instance of *glow* at each bright patch of *shine*. In fact, the pattern in any plane is the *convolution* of the shine and the glow, and convolution is commutative. This is similar to convolving a spherical wave with the source pattern, following Huyghens' principle. However, the shine has no wavelike properties, so there is no cancellation or interference between instances; instead, those phenomena are accounted for by dark shine.

Throughout this work we will compute the spatial patterns for various planes, following the principles articulated in this section. Our examples will be artificially simple and small so that the reader can follow the computations and interpret the diagrams in detail.

The computations and diagrams are implemented in the companion code, written in Matlab. The reader may experiment with the code, available at Github³, to become familiar with how the computations work and to explore diverse examples. For systems any larger than the toy models in this work, computing by code is a practical necessity.

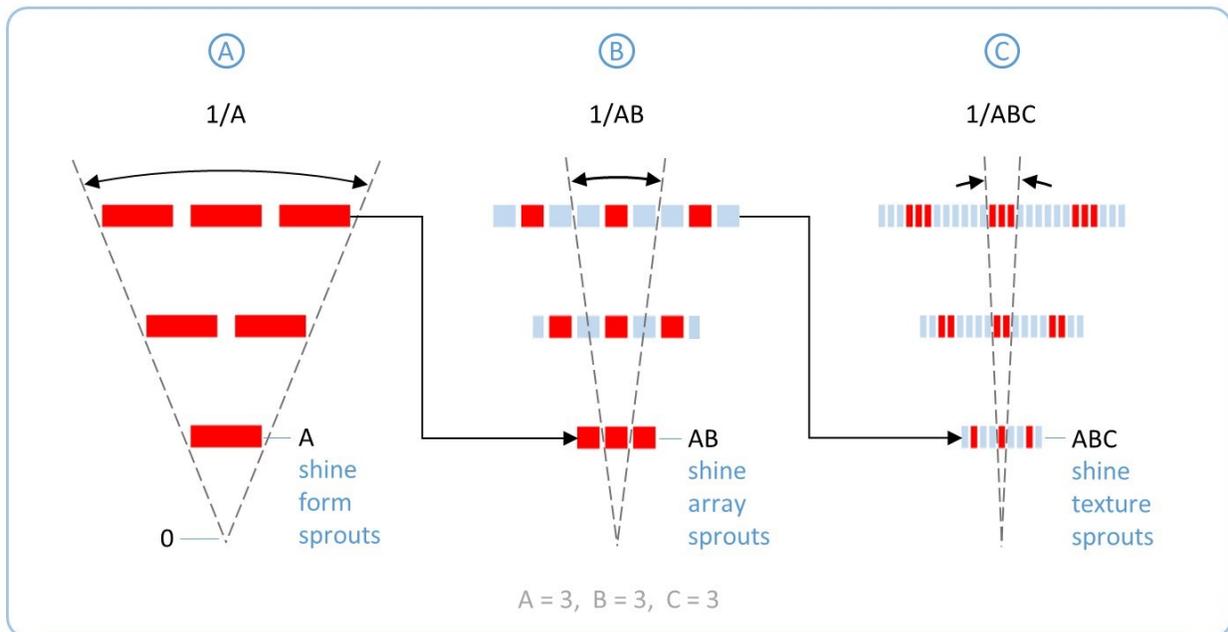
2.6 Sprouting planes

The divergences of the shine shapes are set immediately once the glow is specified at the flat. However, the shapes do not immediately manifest physically because their widths are smaller than a single patch, and light cannot have meaningful features below that width. The shine shapes exist only in potential.

As the shine grows, the hidden features eventually become large enough to manifest physically. A feature *sprouts* when it grows to be one patch in size, which is the dividing line between hidden potential and manifested shine.

Figure 2.12 illustrates an example of one cone of shine at 3 different zooms. Near the flat in Figure 2.12a, all shine is bright. Later in Figure 2.12b, dark shine and bright shine are mixed. The sprouting of the dark shine occurs at plane AB. While the patches in plane AB appear bright, they are actually the first patch of a dark feature; the subsequent patches are truly dark.

Figure 2.12, Sprouting planes



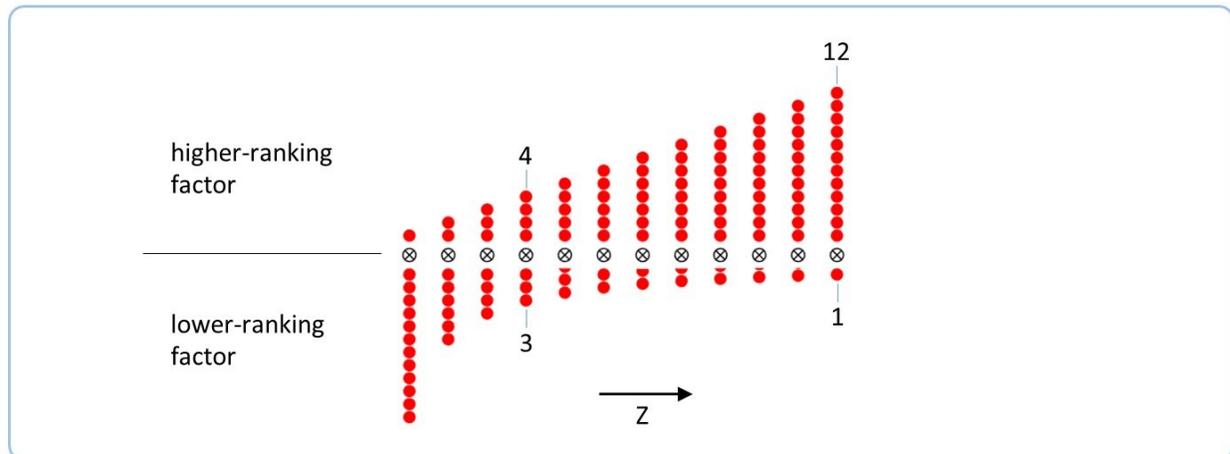
As the dark feature grows, the bright patches grow further and further apart from one another with a divergence of $1/AB$. During this region, the dark factor grows from 1 to C; however, the number of bright patches stays constant at B. These are an embryonic manifestation of the far-field pattern, which contains B orders separated by an angle of $1/AB$.

Like the array, the texture also begins as unsprouted potential, less than a single patch wide in Figure 2.12b. The texture sprouts at plane $Z = ABC$ as the shine texture becomes 1 patch in width. As Figure 2.12c shows, it grows with a divergence of $1/ABC$. Meanwhile, the dark factor has stopped growing and remains constant at C .

2.7 Factor features, rank types

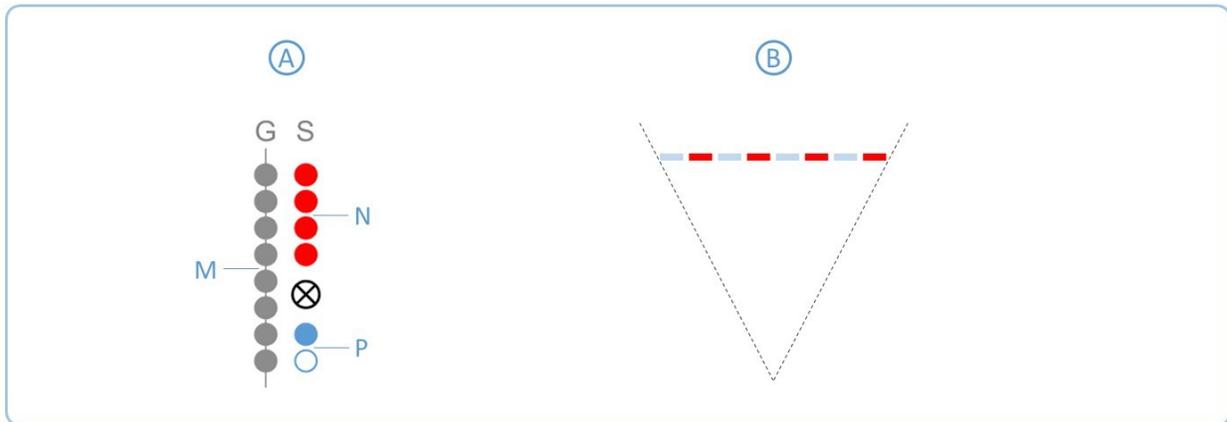
A single feature can also be considered as the product of two or more *factor* features, as shown in Figure 2.13. The sizes of the two factors multiply – they do *not* add – to equal the size of the full feature. As one consequence, trivial factors of size 1 may be added or deleted with no physical meaning. Note that a bright feature has all bright factors, while a dark feature has all dark factors.

Figure 2.13, Factor features



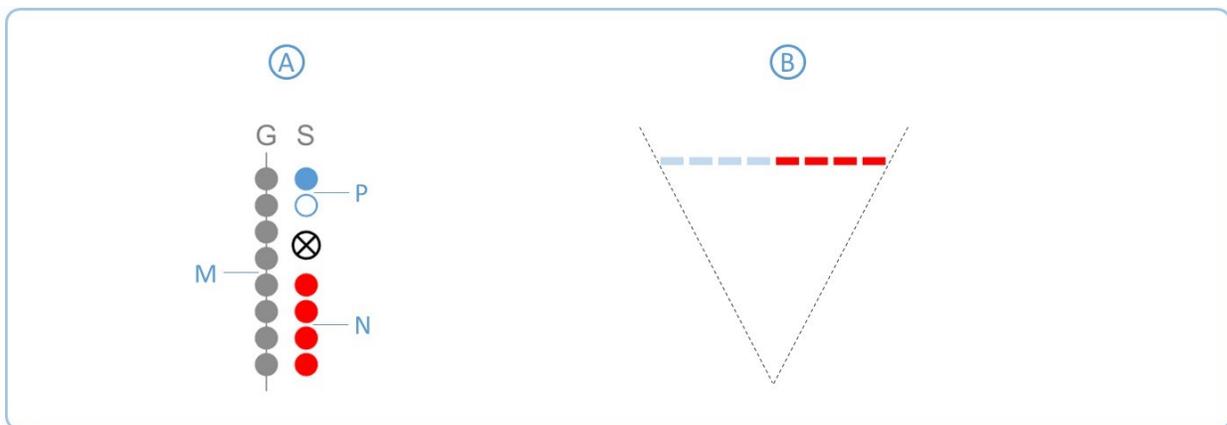
A division into factors occurs when glow and shine features are of different sizes, but partially share a rank. Figure 2.14a shows a glow feature M that is larger than a shine feature N . Here, an additional shine feature P lies *below* N ; the product of the two shines ($2 \cdot 4$) makes up the full size of the glow (8). In such a case we say that N *ranks at the top of* M . Figure 2.14b shows the shine in spatial terms.

Figure 2.14, Ranking at the top



The reverse case is shown in Figure 2.15. Here, feature P ranks *above* N; we say that that N *ranks at the bottom of* M. Figure 2.15b shows the shine in spatial terms.

Figure 2.15, Ranking at the bottom



When M and N are the same size, we say that N *ranks with* M.

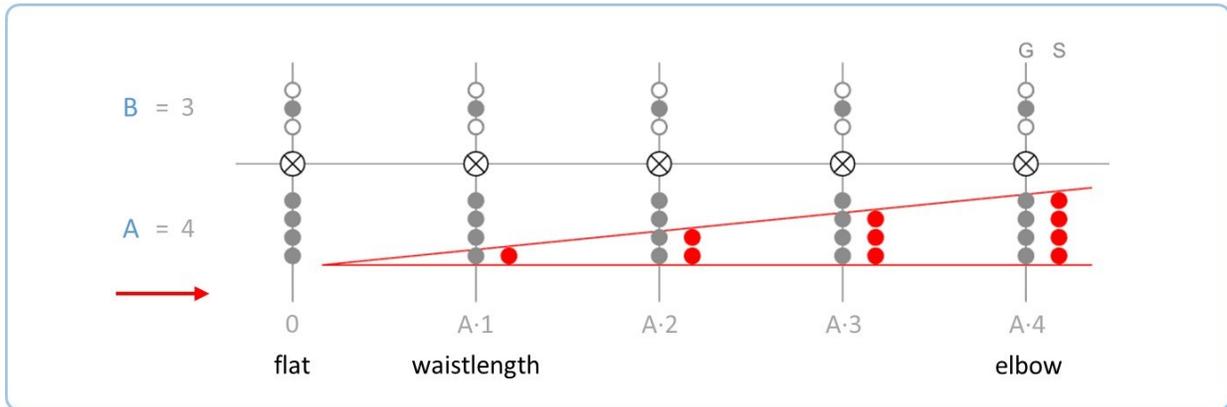
3 The Simple Beam

3.1 Near field and elbow

This chapter analyzes the simple beam in terms of the principles developed in the previous chapter.

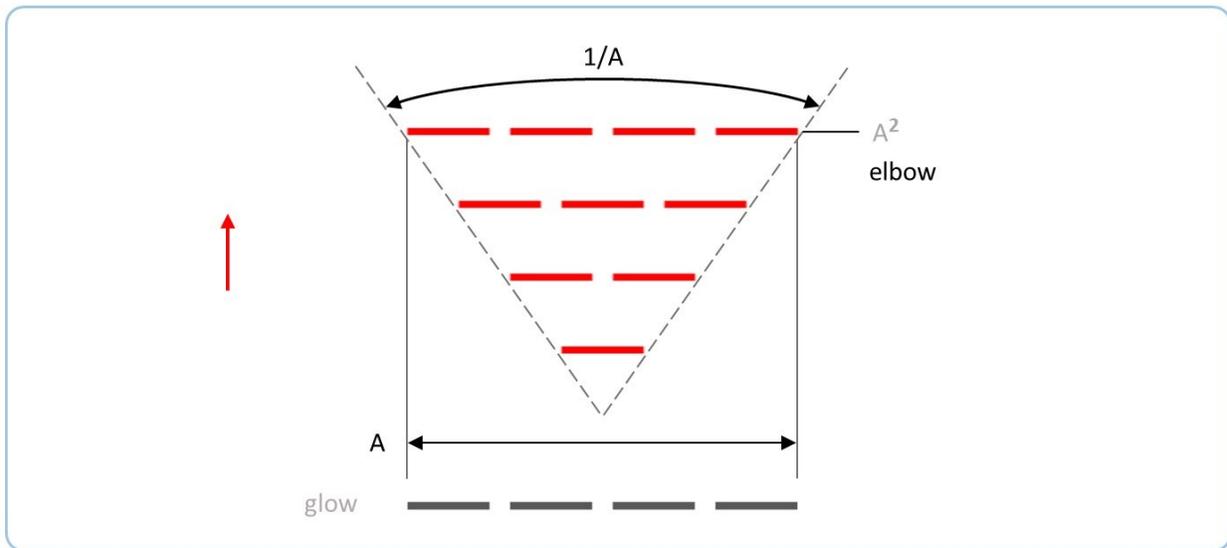
The feature diagram in Figure 3.1 shows the beam *near field*, which extends from the flat to the elbow. The glow consists of a free sky B, and a bright feature A which corresponds to the beam waist width. The shine is a single bright feature which increases linearly from 0 to size A.

Figure 3.1, Feature diagram at the beam near field



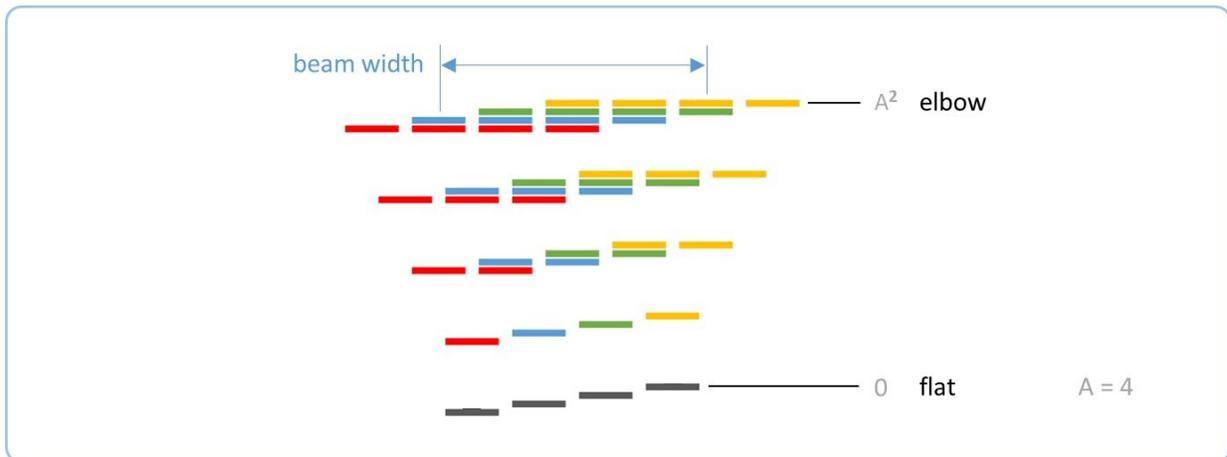
A single instance of shine is drawn in Figure 3.2. The shine has divergence angle $1/A$, and a linear width of A at the elbow. This equals the size of the glow, which is also drawn for comparison's sake.

Figure 3.2, Shine at the beam near field



The tracked-shine diagram is shown in Figure 3.3. While it appears to show the overall pattern growing moderately wider, we measure the width of the beam where the depth is half of its maximum value. Using this approximation, the beam width remains constant in this region. This is one of many approximations used in FO, most of which make FO conceptually simpler at the expense of some fidelity to experiment.

Figure 3.3, Tracked shine in the beam near field



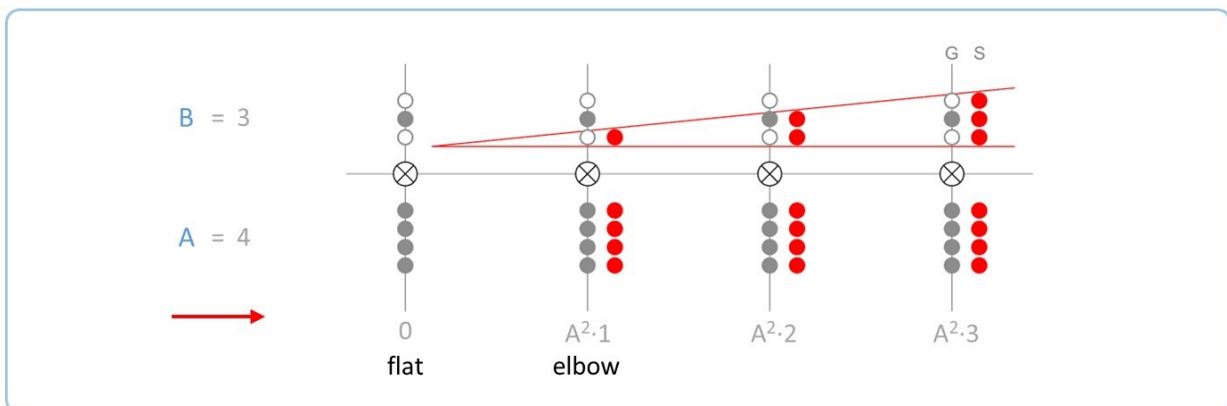
See code in `beam_nearField.m`

In the near field, the cones of shine grow wider while the spacing of their apices remains constant. As a result, the instances of shine overlap and each target patch receives shine from many source patches; the further from the flat, the greater the depth. By computing tracked-shine diagrams we necessarily compute depth; but we will not analyze it within this work. Depth will become important in future work.

3.2 Far field

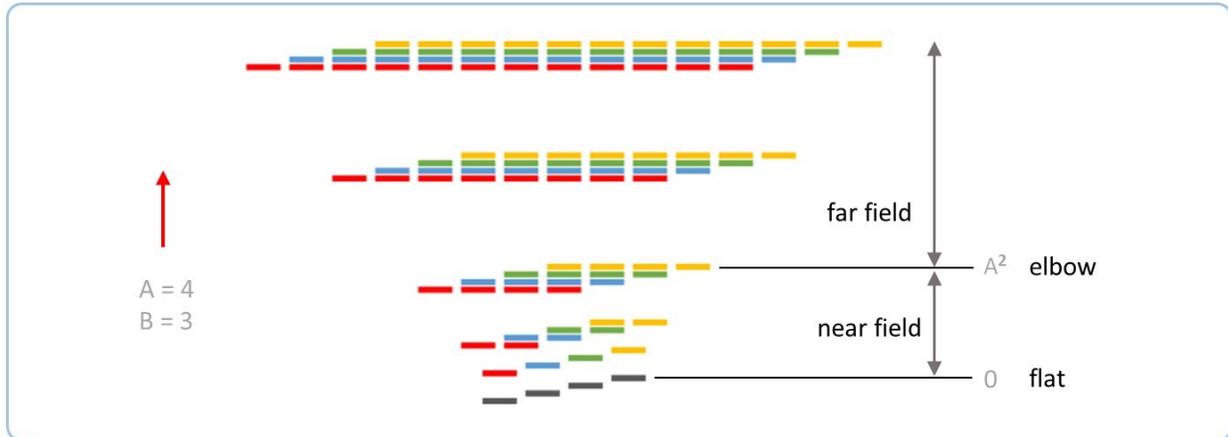
The elbow is the end of the near field, and the start of the *far field* region. In Figure 3.4 we see that at the elbow, the shine has grown as large as the glow, and thus no more expansion is possible at that rank. Instead, the expansion is described by a new factor at a higher rank (while new shine actually grows at the bottom rank, all shine rises and the change appears at the highest rank). Each patch of the new feature represents an entire multiple of width A . At the elbow the new feature consists of just one (trivial) patch, so the shine is A patches wide. When the new feature reaches 2 patches in size, the shine is $2A$ patches wide, etc.

Figure 3.4, Feature diagram at the beam far field



As shown in the tracked-shine diagram in Figure 3.5, the shine continues to widen. Here the overall beam pattern also widens; the width after the elbow is proportional to distance Z . Notably, the depth does not increase further after the elbow.

Figure 3.5, Tracked shine in the beam far field



See code in `beam_farField.m`

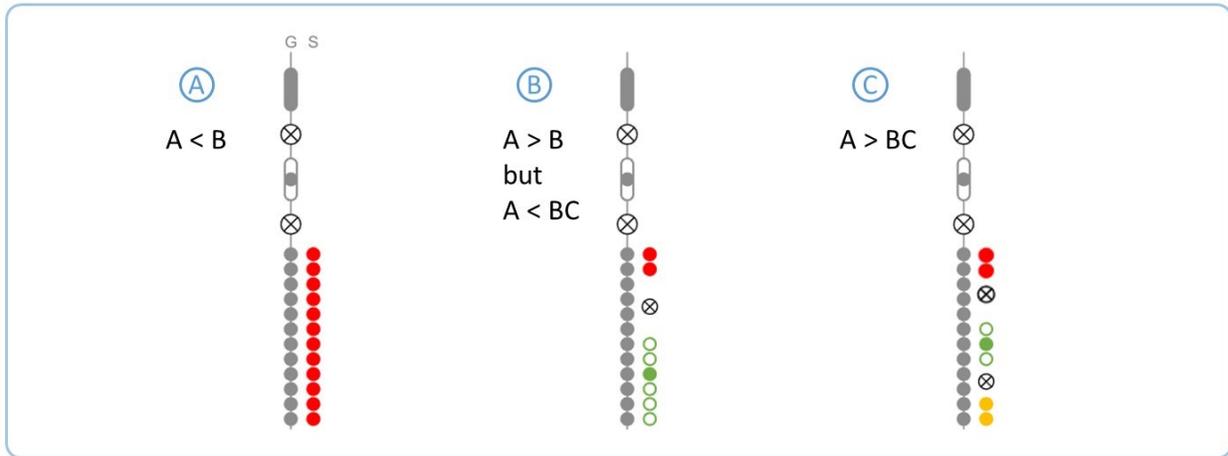
4 The Grating

4.1 Overview and conventions

We now switch to studying the critical planes of the grating, previously analyzed in reference 2. We will study the stages of the grating one at a time, and compute the pattern for each one following the principles postulated above. The results will loosely, though not precisely, match the trends that we simply asserted ad-hoc in reference 2. To avoid duplication, we do not reproduce all of the diagrams from that reference, so the reader is advised to refer there to confirm that these results approximately match those. Future work based on the present work will add additional layers to the algorithm which make the match exact.

In the earlier planes, there exist multiple possibilities for which shapes have sprouted. For example, Figure 4.1 shows three different possibilities for the early elbow which depend on the values of feature sizes A , B , and C . Therefore it is impossible to present a single example that covers all cases, and it is impractical to cover all possible cases within this work.

Figure 4.1, Sprouting cases at the early elbow



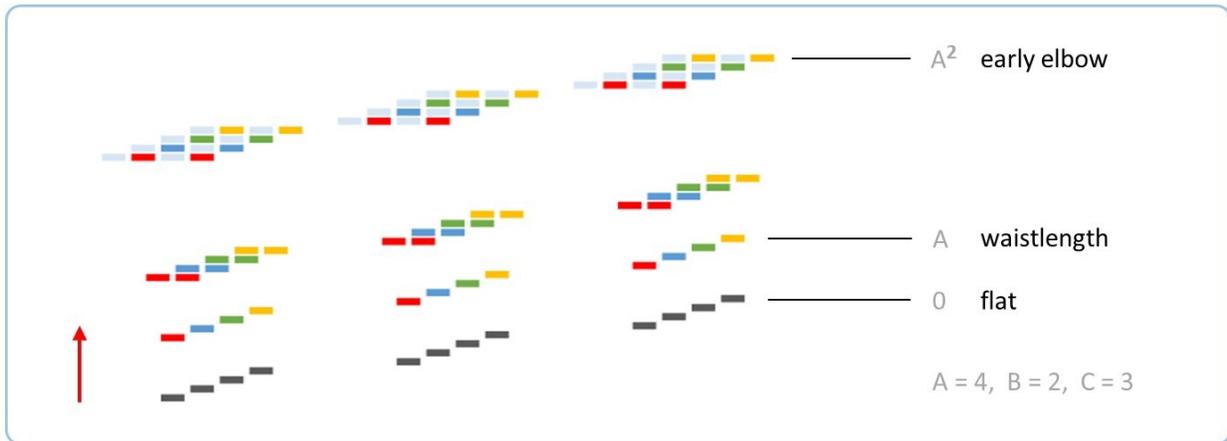
Instead, we proceed by studying the middle case Figure 4.1b, in which the array has sprouted, but the texture has not. This offers the easiest point from which to extrapolate to the others. In fact, the sprouting case does not affect the observable pattern; it only affects the depth, which this work mostly ignores.

Now we establish some conventions to refer to features in the figures. We refer to a full feature by its label, such as B. If we need to refer to both the glow and the shine features with the same label, we will label them with letter plus numbers, such as B1 and B2. For partially-grown features, or for factor features, or any other case where the labels A, B, C, D are not appropriate, we use letter M, N, P, and Q. Note that labels are only used consistently within a single diagram, not between diagrams.

4.2 Early elbow

After the flat, the first critical plane is the early elbow. As Figure 4.2 shows, bright shine feature B ranks at the top of feature A. In this example, we have assumed that $A > B$, and so dark shine feature M has already sprouted.

Figure 4.4, Tracked shine at the early elbow

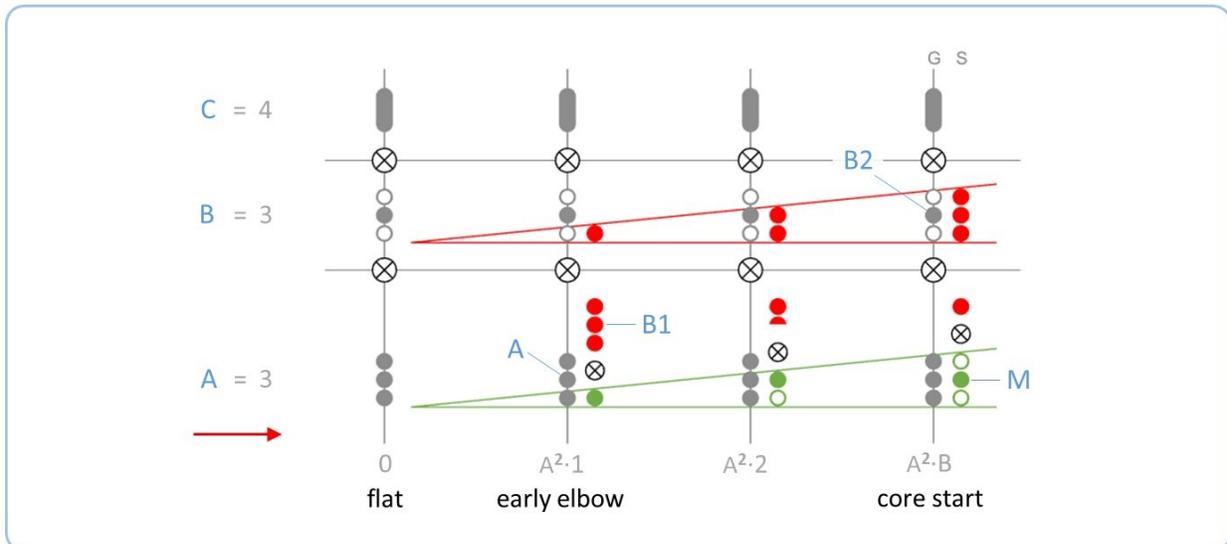


See code in `earlyElbow.m`

4.3 Core start

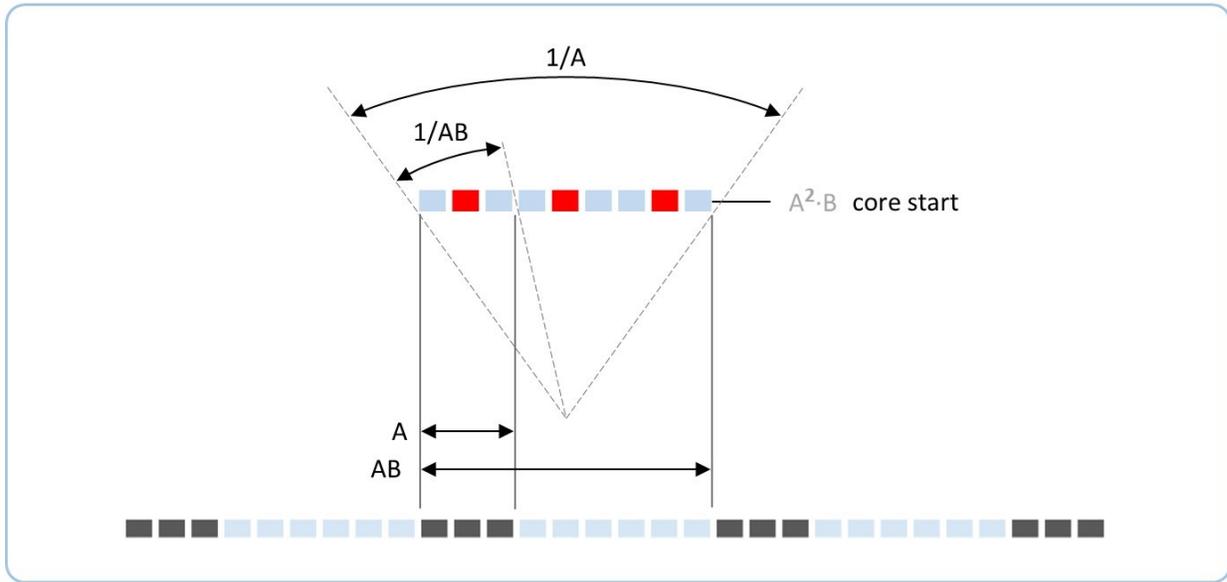
The feature diagram in Figure 4.5 shows the stage approaching the core start. At the early elbow, feature B1 (bright shine) ranks at the top of feature A (bright glow); note that A and B are generally not of equal size as in this example. The shine grows until B1 ranks with B2 (dark glow) at the core start, where dark shine M always ranks with feature A.

Figure 4.5, Feature diagram at the core start



The corresponding shine spatial diagram in Figure 4.6 shows that the shine form has a width of AB at the core start, which equals the width of the glow array. Also, the shine array has width A , which is the same as the glow texture.

Figure 4.6, Shine at the core start



Because the glow texture is a continuous string of A source patches, whenever a target patch receives dark shine from one source patch, it also receives bright shine from an adjacent source patches. Thus, the dark patches are never visible, as Figure 4.7 shows. The pattern appears like C copies of a simple beam, expanding after its elbow to fill the entire interstitial space between the grating slits. The pattern at the core start is a continuous bright width; the periodicity is not observable.

Figure 4.7, Tracked shine at the core start

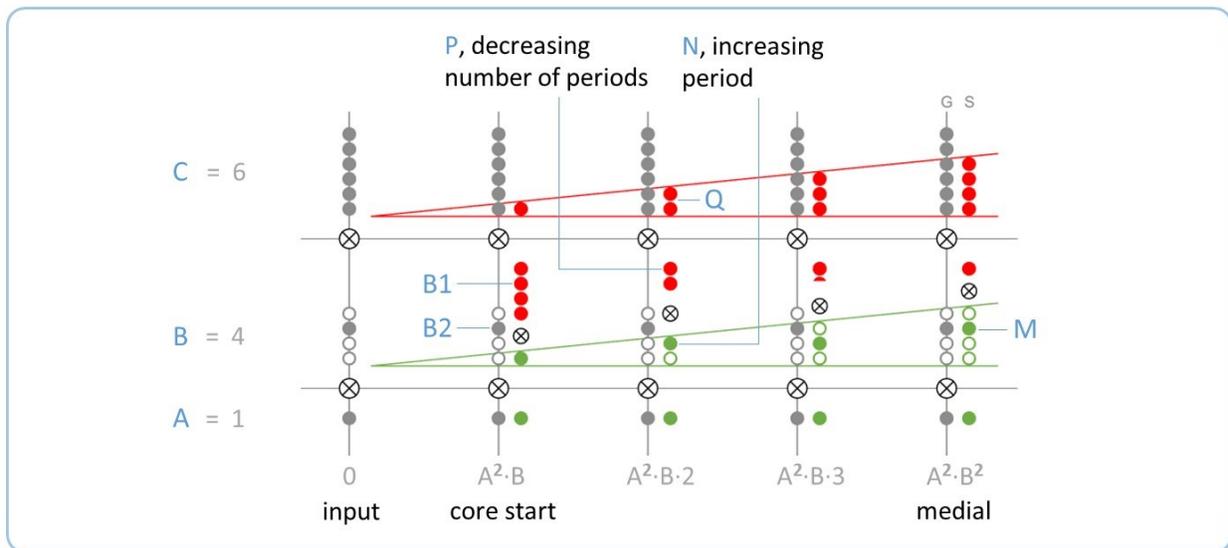


See code in `coreStart.m`

4.4 Early revivals and medial

The feature diagram in Figure 4.8 shows the next stage. Dark feature M rises so that its factor N ranks above feature A (in this example, $A = 1$ and so $M = N$). This dark shine produces the *early revivals*, which appear as a periodic pattern. The bright region of each period has the width A of the starting slit; the value N indicates each revival period as a *multiple* of A, or equivalently the dark factor of the revival (see reference 1). Viewed this way, the solid bright width at the core start is actually a periodic pattern in which the dark factor $N = 1$, i.e. the bright region fills the entire period. The revival periods grow progressively larger until the medial, where the period equals the glow array.

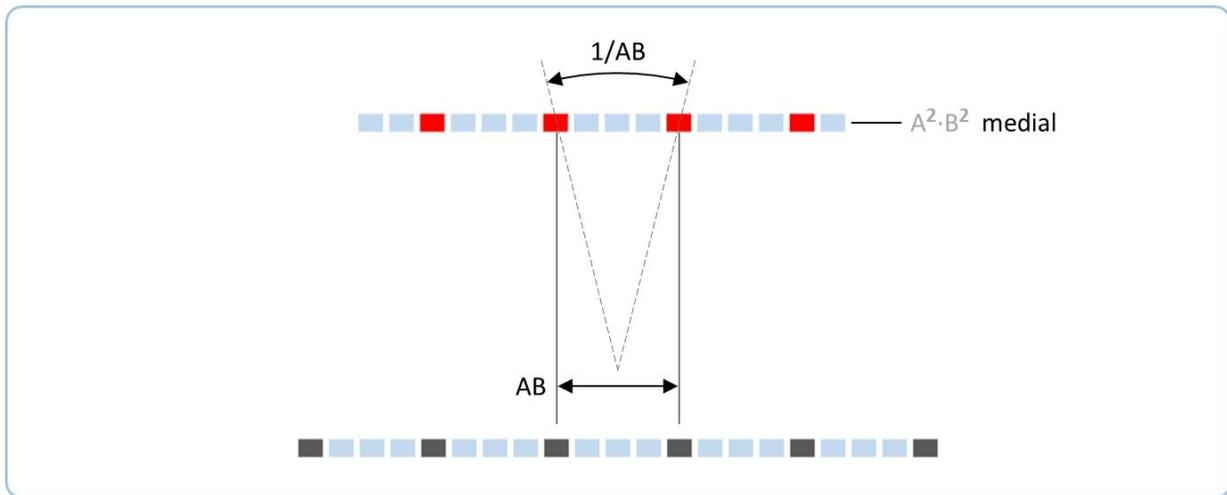
Figure 4.8, Feature diagram at the medial



Simultaneously, as bright shine B1 rises in rank, it divides into two factor features: factor P which remains ranked at the top of glow B2, and the remainder Q which ranks at the bottom of feature C. A revival occurs whenever factor P takes an integer value. Factor P also has a simple physical interpretation: it is the number of revival periods in each glow period. Equivalently, it is the spatial frequency of the early revivals, relative to the glow frequency. It decreases over the course of this stage, falling to 1 at the medial.

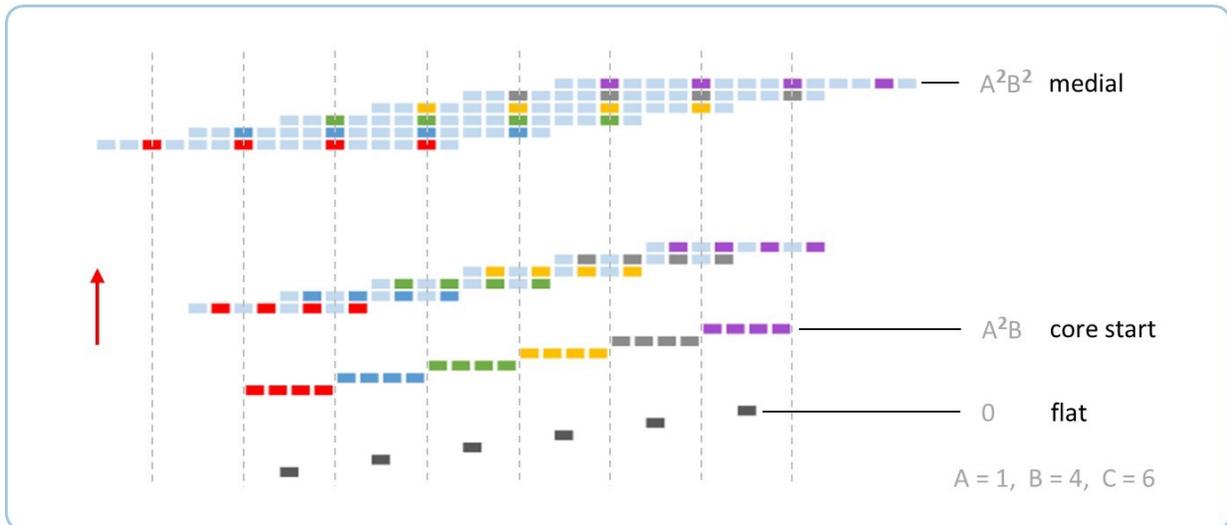
Figure 4.9 shows the shine at the medial. In this example, the bright region of each period consists of a single speck; however, in other cases the shine texture has sprouted at the medial and the bright region consists of multiple patches. In all cases, the shine array is AB, which equals the glow array.

Figure 4.9, Shine at the medial



The consequences of this are visible in the tracked-shine diagram in Figure 4.10. When the glow array is equal to a simple fraction of the shine array, the shine from neighboring slits tends to line up such that dark shine from one slit coincides with dark shine from the other. This renders the dark shine visible, whereas in earlier planes the dark shine was always concealed behind the overlapping bright shine from other sources.

Figure 4.10, Tracked shine at the medial



See code in *medial.m*

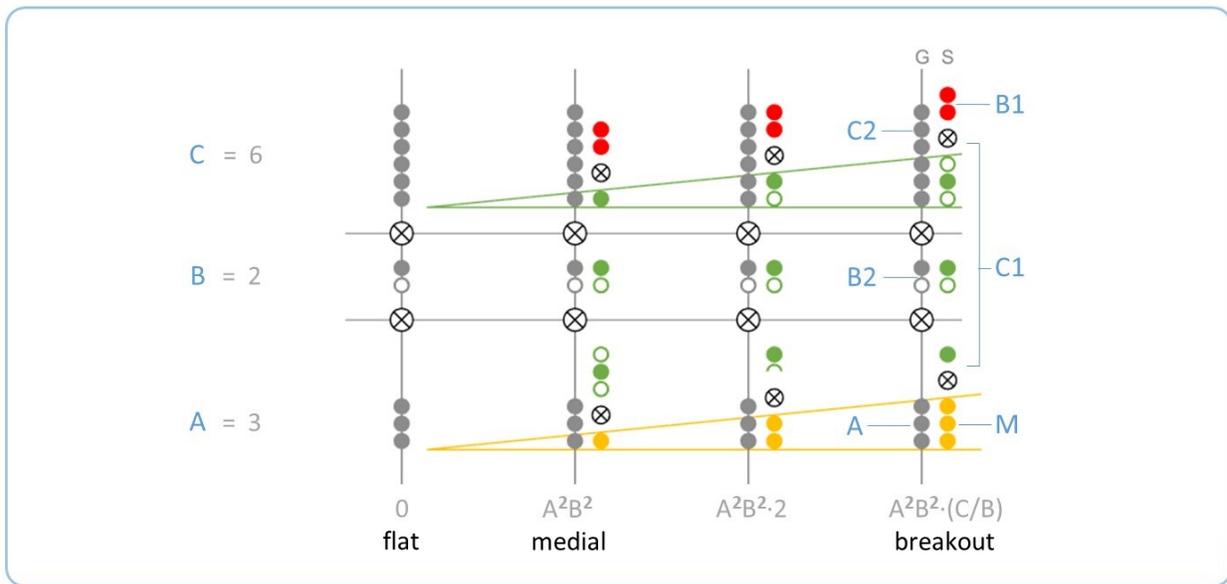
4.5 Breakout

In choosing examples in the previous stages, we were forced to narrow the space of possibilities by making arbitrary assumptions about which shine features had sprouted (see section 4.1). From this point forward such assumptions are not necessary, because all shine features have necessarily sprouted by breakout.

Also note, we have assumed that $C > B$, which specifies a *late-breaking* grating, meaning that breakout occurs *after* the medial. Early-breaking gratings are also possible, but they are beyond the scope of this work.

Figure 4.11 shows the feature diagram of this stage. At the medial, feature B1 ranks at the *bottom* of feature C2. Over the course of this stage, B1 moves upwards until it ranks at the *top* of C2 at breakout. Also at the breakout, M has grown as large as A.

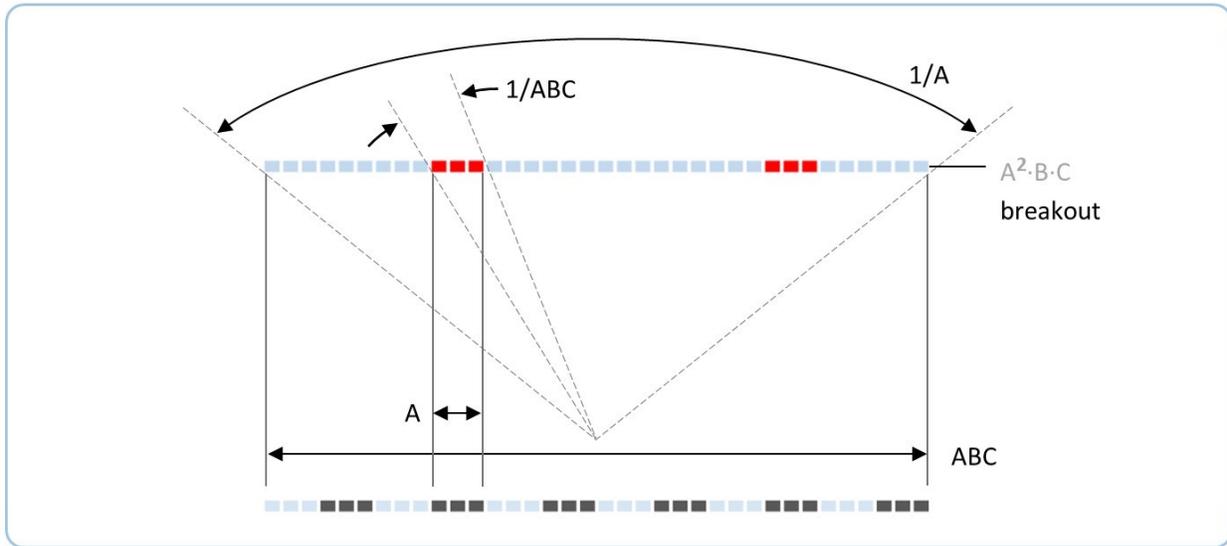
Figure 4.11, Feature diagram at breakout



The *late revivals* (also called *Talbot images* or *self-images*) occur in this stage. One late revival occurs each time a patch of C1 rises above B2 to the rank of C2; there are a total of C/B such revivals in total, which includes the medial and breakout.

The shine drawn in Figure 4.12 illustrates these same points in a slightly different way. Firstly, the shine form reaches width ABC , overtaking the glow form. At this transition, the shine appears to emerge from inside the glow; this is the source of the name 'breakout'.

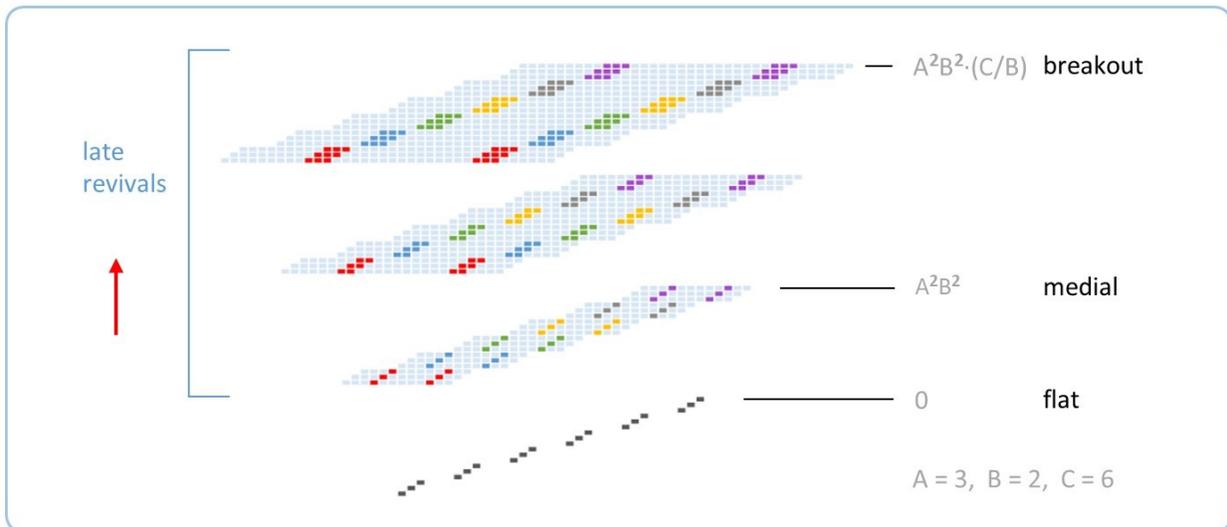
Figure 4.12, Shine at breakout



A similar transition occurs in the texture. At the flat, early revivals, and late revivals, all patterns have had texture of the same width A , i.e. the glow texture, spaced at various periods. At breakout, the shine texture reaches width A , and in all subsequent planes the texture width is $> A$.

As the tracked-shine diagram in Figure 4.13 shows, the changes in this stage occur mainly in the depth and in the distribution of dark shine. Neither of these are observable, and so the pattern appears to recur cyclically with no change – hence, the term ‘revival’.

Figure 4.13, Tracked shine at breakout

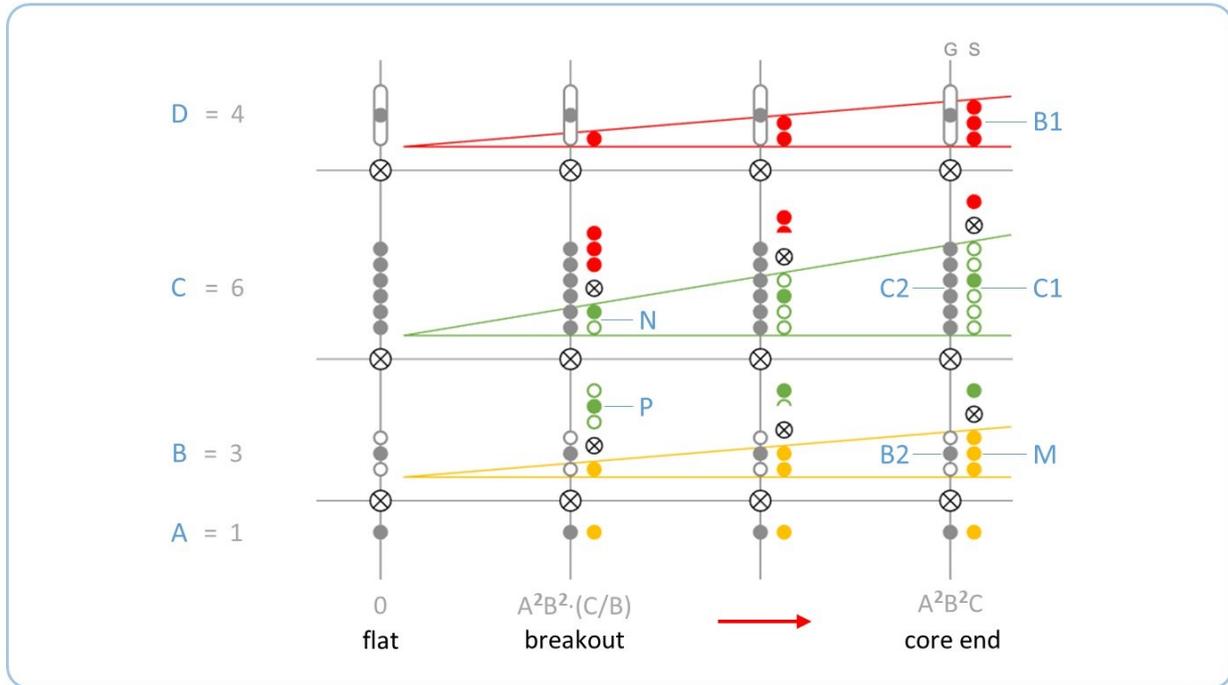


See code in breakout.m

4.6 Core end

The feature diagram in Figure 4.14 shows several trends between breakout and the core end. Firstly, bright feature M is initially ranked at the top of feature A; then, M rises until it is ranked at the top of feature B2 at the core end.

Figure 4.14, Feature diagram at the core end



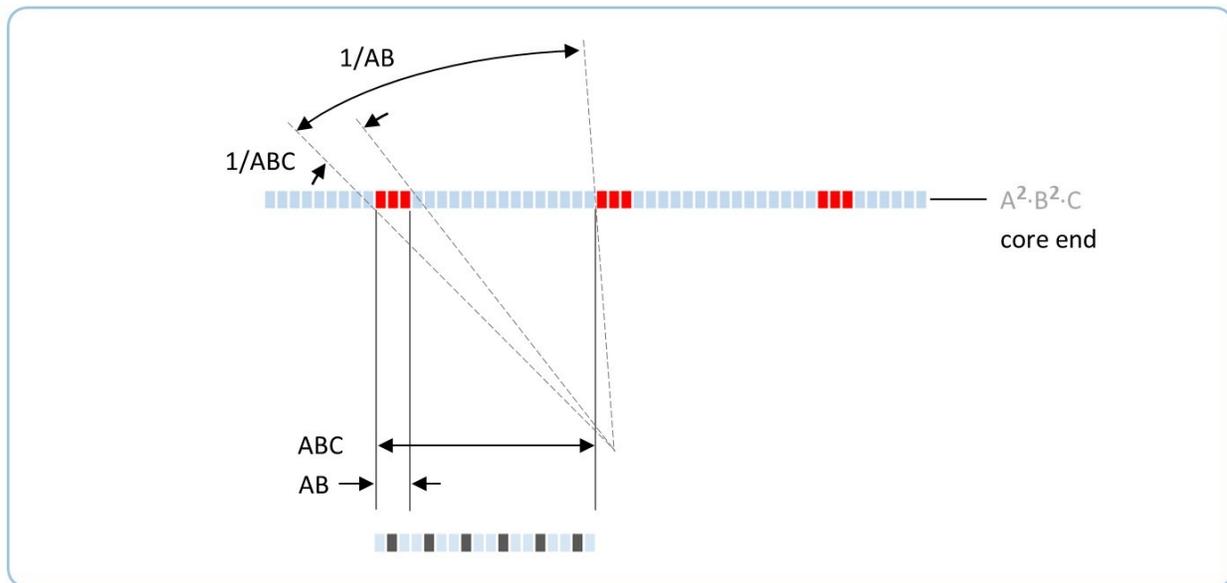
Secondly, dark feature C1 is initially divided between two factors N and P. P is initially of size B and ranks with B2, where it manifests as the dark space in the periodic pattern at breakout. Factor P gradually rises until all of C1 is ranked with C2.

Thirdly, factor B1 is initially ranked at the top of C2. Over the course of the stage, B1 rises until it is ranked immediately *above* C2.

All three of these trends are unified by a common mechanism – the chain of shine features, rising in rank. The chain acts as a single unit whose parts do not change size or arrangement, except for the shine growing at the very bottom of the chain.

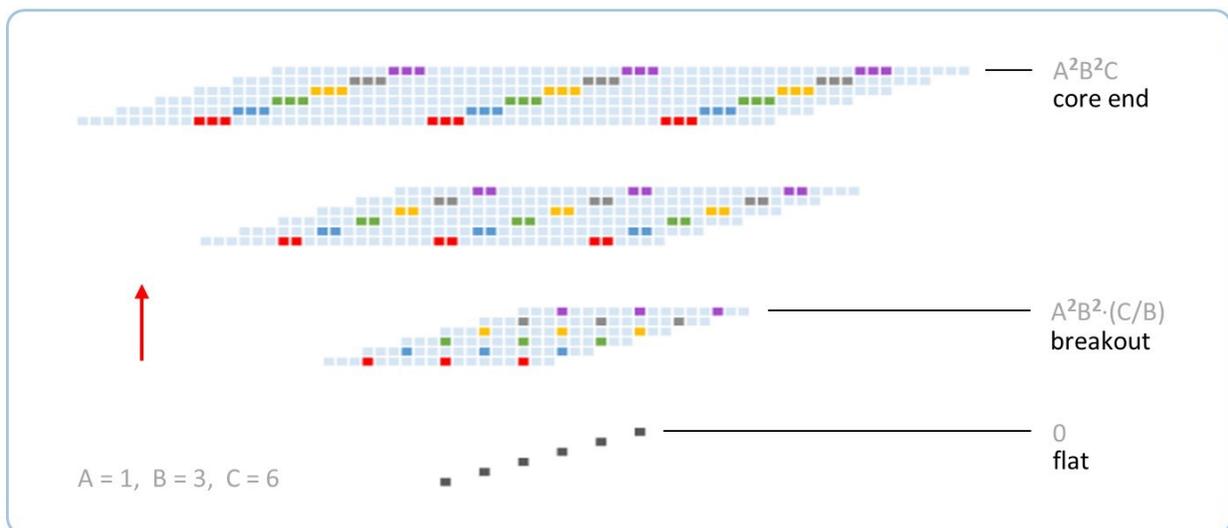
The shine at the core end is shown in Figure 4.15. The shine texture is AB, the width of the glow array. The shine array has width ABC, which matches the glow form; the physical interpretation is that the far-field orders are just barely overlapping. Note that while the shine array is much wider than at breakout, this is *not* due to the growth of the dark shine. Rather, the constant dark shine maintains the array at a fixed dark factor, i.e. a fixed multiple of the shine texture. As the (bright) shine texture gets wider, the shine array grows wider too.

Figure 4.15, Shine at the core end



The tracked-shine diagram in Figure 4.16 shows the same trends as the feature diagram. Several trends occur simultaneously: Firstly, as the individual bright widths expand from width A to AB , they grow to be as wide as the glow array, gradually degrading the revivals. As another aspect of this same trend, the dark shine from each slit is concealed by the bright shine from adjacent slits, so that the dark space is no longer visible. At the core end, they finally form one continuous bright region, similar to the one seen at the core start.

Figure 4.16, Tracked shine at the core end



See code in `coreEnd.m`

Lastly, the overall width of the pattern grows by a factor of B . For the first time, the shine rather than the glow determines this overall width. B distinct far-field orders emerge, each one a single period of the shine, and each one as large as the glow form.

4.7 Late elbow

The final trend in the propagation of the grating is shown in Figure 4.17. At the core end, (shine) feature C1 ranks with (glow) feature C2. In this stage, C1 rises to the rank above C2.

Figure 4.17, Feature diagram at the late elbow

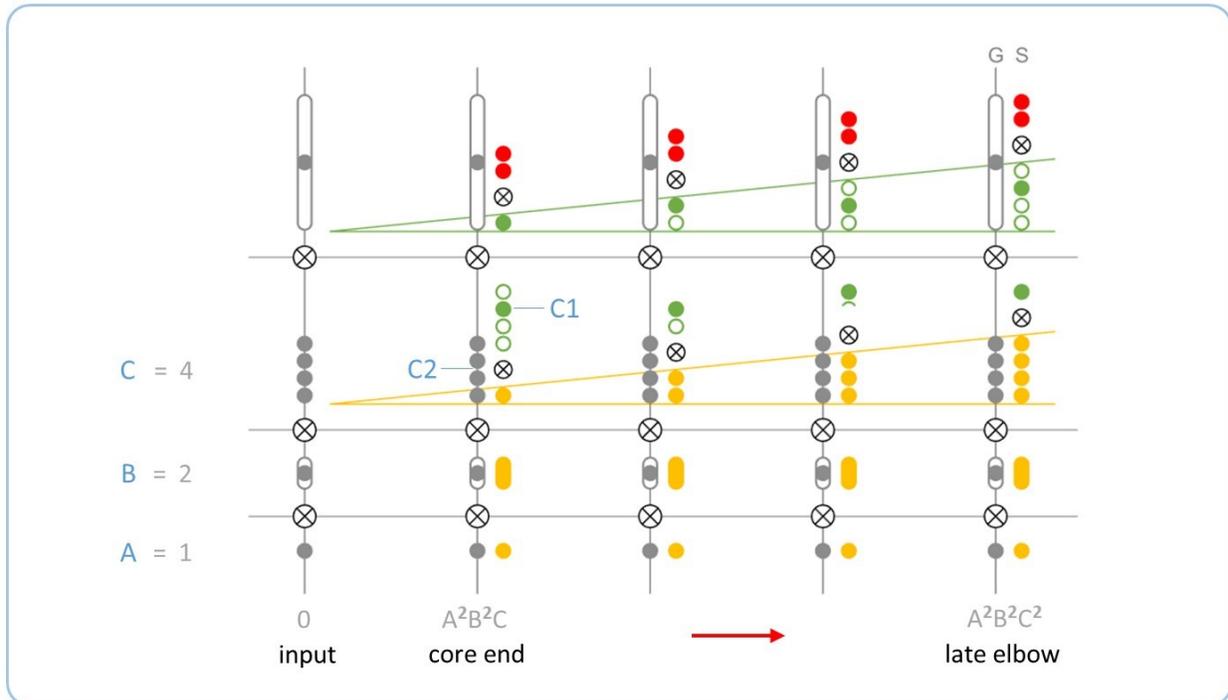
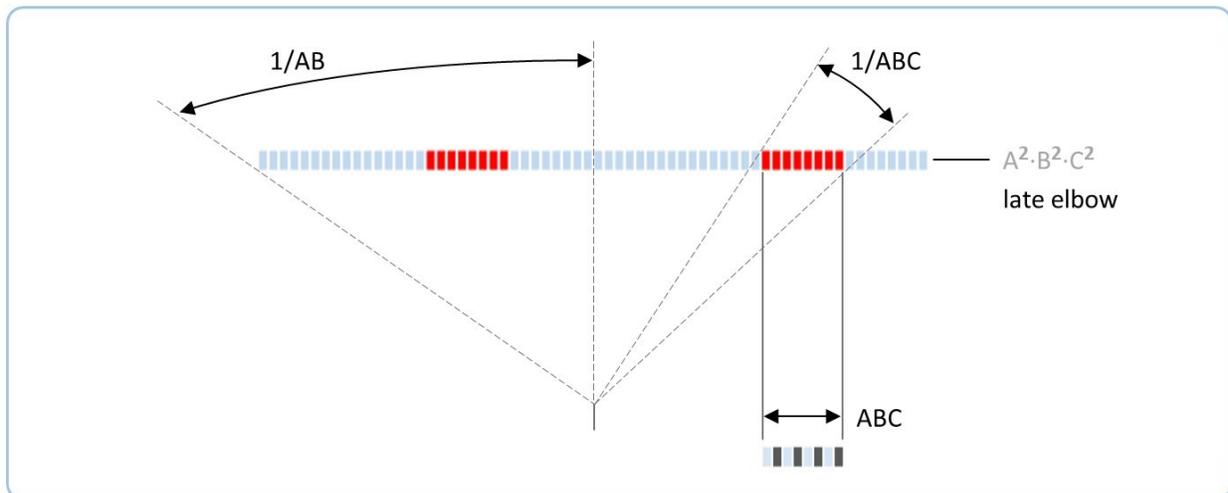


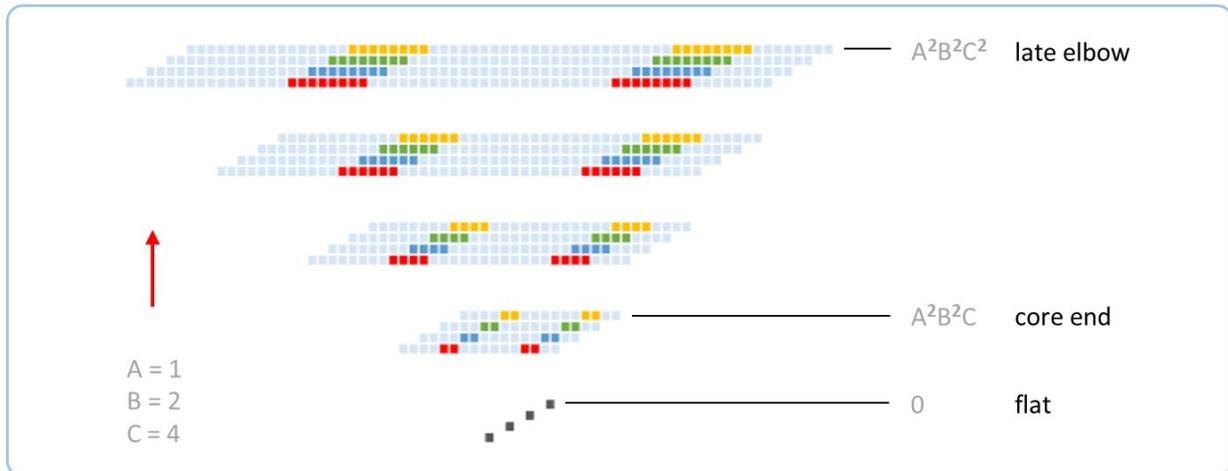
Figure 4.18 shows that the shine texture grows to the width of the glow form.

Figure 4.18, Shine at the late elbow



The tracked-shine diagram in Figure 4.19 shows that the B separate large beams at the core end grow apart from one another, forming B distinct far-field orders. Dark feature C is revealed, producing dark space in between the orders as the shine array grows from ABC to C·ABC.

Figure 4.19, Tracked shine at the late elbow



See code in *lateElbow.m*

In this region, the pattern is dominated by the shine. The structure of the glow is not even observable; it is concealed by the shine texture.

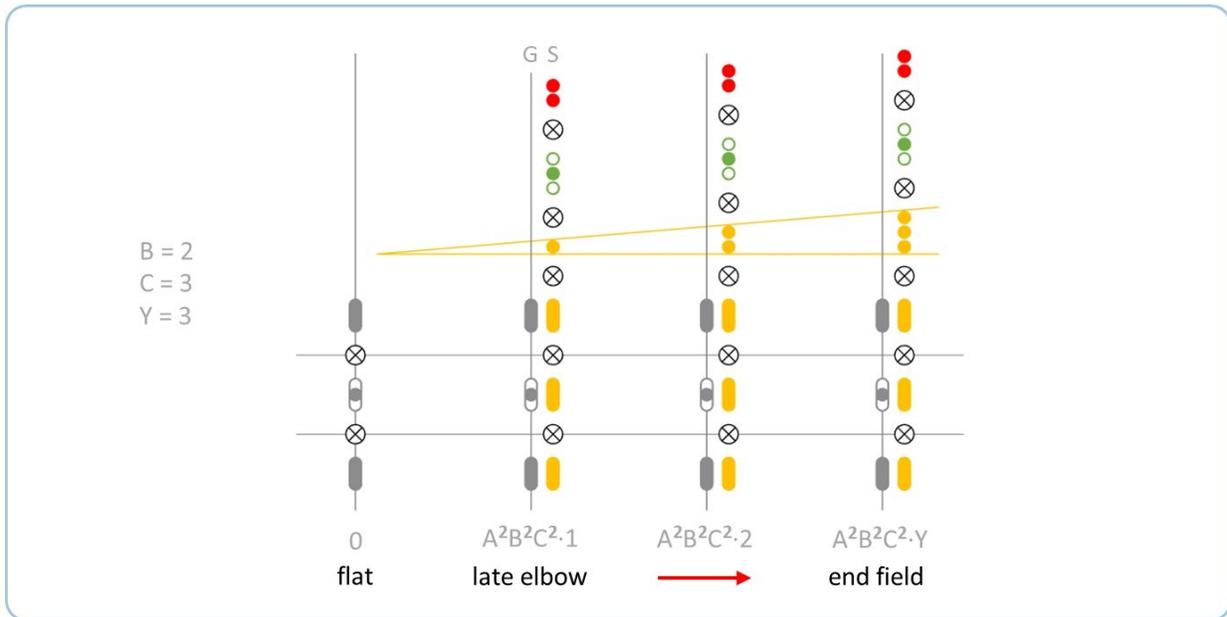
4.8 End field

At the late elbow, the pattern reaches a kind of stasis. The Fourier transform is then nearly complete, apart from a change in curvature. Further propagation brings only a uniform magnification in width, as new bright features push existing ones higher in rank.

We see the essence of this behavior in the Zero-padding theorem⁴, which states that if a discrete input vector to the FT is lengthened by adding additional zeros at the end, the output of the FT is correspondingly lengthened by ideal interpolation, i.e. by keeping the function the same but increasing the sampling density. In our case, the zero-padding is effected by increasing the size of the free sky D , which determines the size of the space under consideration and thus the length of the input vector. This results in a larger bright shine D at the bottom of the shine chain; this is the uniform magnification of the final shape, i.e. the ideal interpolation.

The feature diagram in Figure 4.20 shows the reason for the stasis. At the late elbow, the glow features stand ranked with an unbroken chain of bright shine. As more bright shine sprouts, it merely displaces other equivalent bright shine. We see none of the various effects that we saw earlier, which resulted from changing patterns of shine interacting with the stationary glow.

Figure 4.20, Feature diagram at the end field

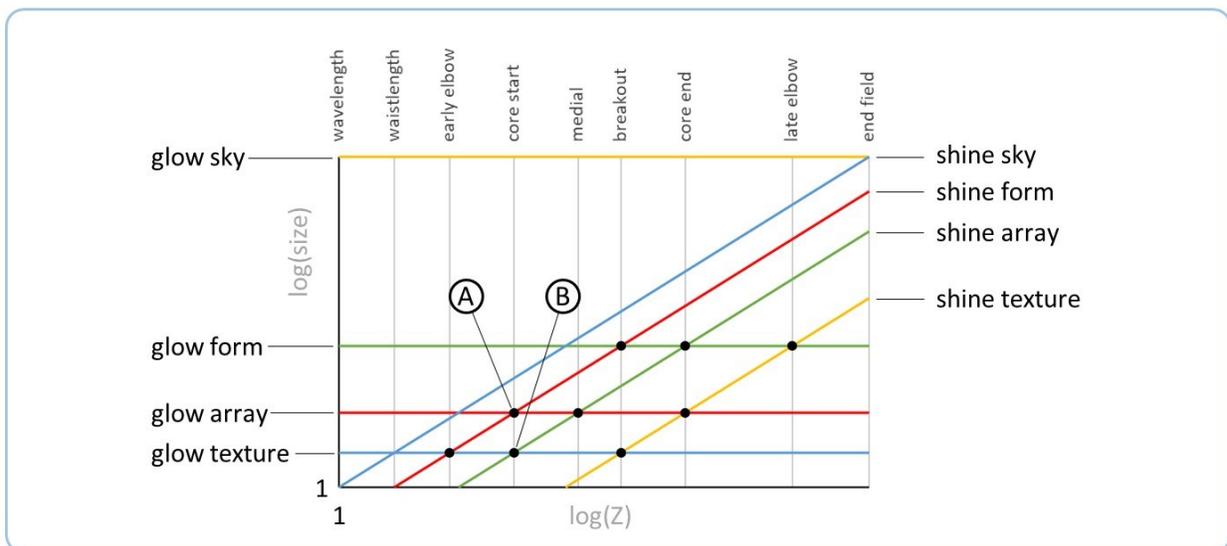


5 Shapes and beams

Earlier we introduced the shapes texture, array, form, and sky. The shine, the glow, and the pattern each have all four of these shapes.

The glow shapes remain constant, while the shine shapes expand with propagation. At many of the critical planes, one or more shine shapes reaches the size of a glow shape, as discussed in the plane-by-plane analysis of the grating and as shown in Figure 5.1.

Figure 5.1, Glow and shine shapes



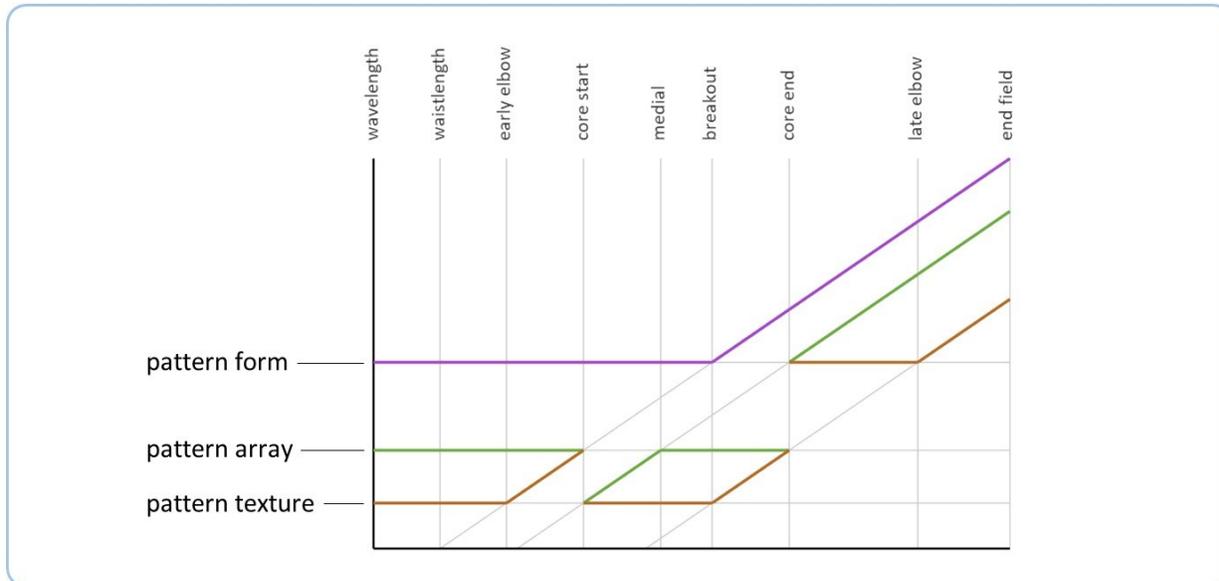
For example, at the core start, the shine form reaches the size of the glow array (label A). At the same plane, the shine array reaches the size of the glow texture (label B).

Note that both axes of the figure use a logarithmic scale. Note also that the colors here partly correspond to the colors in Figure 2.4, and Figure 5.1 looks nearly identical to the corresponding feature diagram. However, there is a crucial distinction: here, each color corresponds to a *shape*, which is a chain of multiple features. In the other figure, each color corresponds to a single *feature* – namely, the highest-ranking feature in the shape, which gives the two a one-to-one correspondence.

The shine sky always rises to the glow sky at the end field. This might seem to be unphysical, since freely-propagating light does necessarily hit a limit. However, while the other 3 shapes are ‘set’ at the flat when the grating is placed in its initial state, the glow sky is not. Rather, it is determined at the end field, i.e. the last plane that we compute the shapes for. It is ‘set’ mathematically at the beginning of the problem, but not physically.

The pattern also has texture, array, and form. However, it follows the much more complex trend seen in Figure 5.2 (sky is not shown). This diagram is a preview only, and the details of these trends will be discussed in detail in future work.

Figure 5.2, Pattern shapes

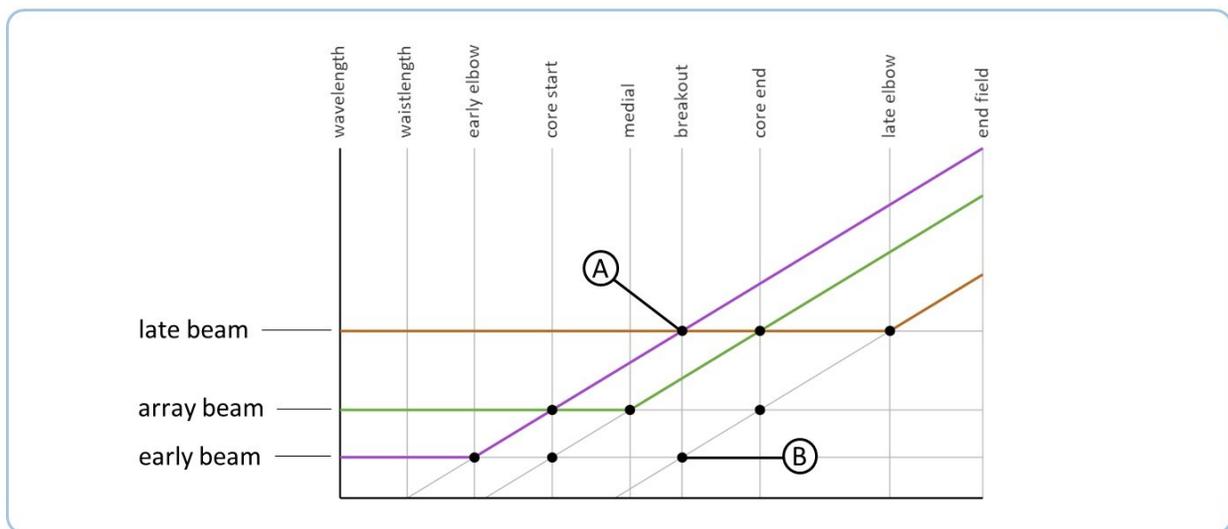


Alternatively, one may analyze the grating in terms of 3 conceptual beams: the *early beam*, the *array beam*, and the *late beam*. The beams are only conceptual or virtual, in that they have some but not all properties of an actual beam. Each beam has the lowest possible space-bandwidth product permitted by the uncertainty principle. The grating, which is in some sense

composed of the three beams, has a space-bandwidth product that is far in excess of that calculated for the individual beams.

The 3 beams are depicted in Figure 5.3. The early beam begins as the (pattern) texture, has its elbow at the early elbow, and finally becomes the (pattern) form. The array beam begins as the array, has its elbow at the medial, and remains the array. The late beam is the exact opposite of the early beam – it begins as the form, has its elbow at the late elbow, and finally becomes the texture. Outside of the core, the pattern is created from the 3 beams, combined by a simple rule – the array convolved by one beam, multiplied by the other beam. But inside the core, the beams cannot be related to the pattern by any such simple rule.

Figure 5.3, Conceptual beams

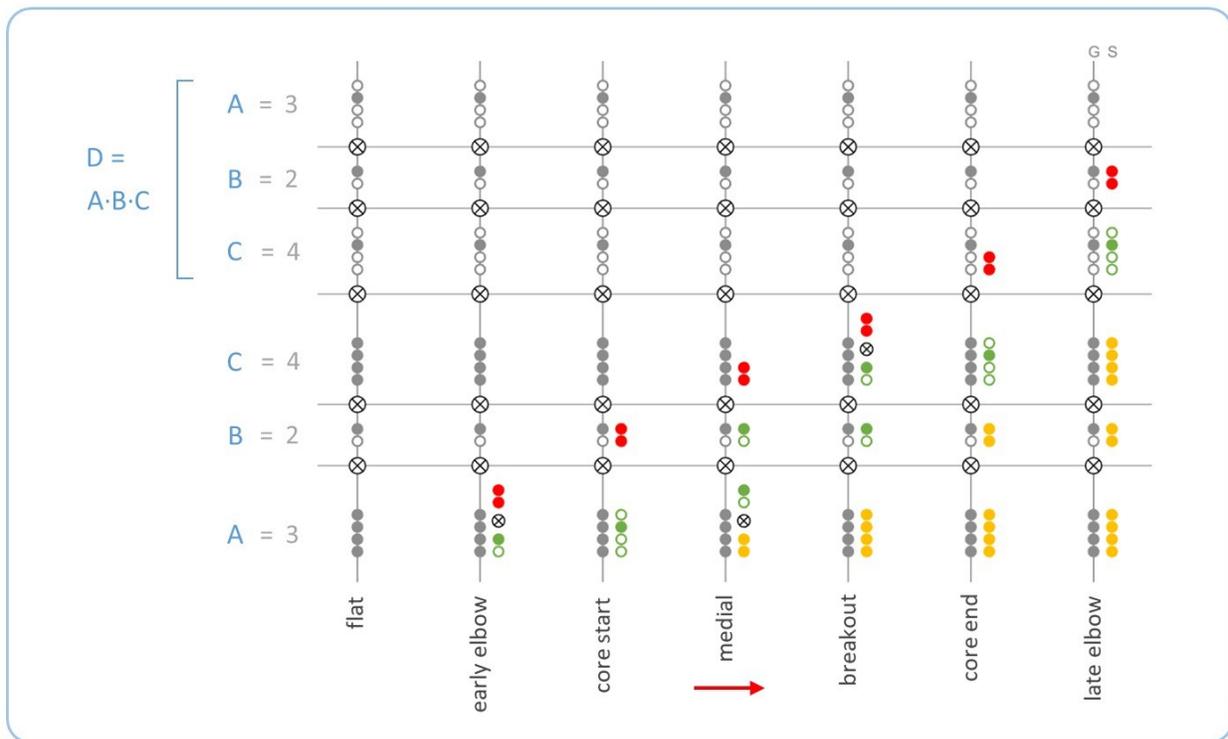


The figure shows a dot at each match between a shine and a glow. In some planes there is only one match; others have two. The 3 planes with only one match each are the elbows of the 3 beams – where its shine matches its glow. At the remaining three planes, one beam's shine matches a *different* beam's glow. The relationship is always mutual – for example, at breakout, the shine of the early beam matches the glow of the late beam (label A). Symmetrically, the shine of the late beam matches the glow of the early beam (label B).

6 Conclusion

Figure 6.1 summarizes the entire progression from the input to the end field. Note that the scale of the horizontal axis is *not* linear, but rather is roughly logarithmic.

Figure 6.1, Summary of critical planes



This work represents the beginning, but not yet the completion, of an effort to place the ad-hoc rules of reference 2 on a theoretical foundation. Here, we have introduced foundational principles and shown how they can be rigorously applied to compute tracked-shine diagrams. These resemble the spatial diagrams in reference 2, but they are not yet exactly the same. The remaining steps will be taken in a subsequent work.

7 References

1. Mirsky, Paul L. A First Look at Feature Optics. <http://vixra.org/abs/1909.0157>, 2019
2. Mirsky, Paul L. Feature Optics and the Critical Planes of the Grating. <http://vixra.org/abs/1909.0185>, 2019
3. Github repository at <https://github.com/paulmirsky/glewAndShine>
4. Smith, Julius O., III. Mathematics of the Discrete Fourier Transform (DFT) with Audio Applications, 2nd edition. See <https://www.dsprelated.com/freebooks/mdft/>