Annihilation of the 600-rule once and for all.

I guess there are plenty of information about this rule in the Internet but I would like to write my own debunk of this old myth. At least in order to systematize my thoughts on the subject.

Let's start from the fact: it does not work! At least nowadays and with modern cameras it drastically overestimates the exposure: so all stars that should be pinpoint-like — look like star-streaks.

As this rule goes, user can set the exposure inversely proportional to the focal length of the lens used on the camera-body. For example, if one has a ***35mm-film camera*** and the lens with the 50 mm focal length the limiting exposure in seconds could be found as: $600/50 = 12s$. Especially funny when someone recommends to reduce the exposure by the crop factor with absolutely no explanation whatsoever. Or by saying that crop factor somehow changes the focal length of your lens! As an astrophotographer the only emotion I can feel at such moments is the limitless cringe… The only thing that the crop factor changes is the field of view. But if you use lens that was designed for the full-frame cameras (i.e., it covers its sensor entirely) no one can stop you from using it on a cropped-sensor camera. Or you can even switch your full-frame body to a crop-mode. That would be equivalent to cropping your images in post. And it is obvious that switching to a crop-mode on a full-frame camera have zero effect to the limiting exposure. So what factors are in play here?

In order to properly deconstruct this rule, we need to cover the basics: how our cameras capture light and how fast the sky is rotating.

The simplest camera that exists is the pinhole camera, also known as camera obscura. It does not have any lens — only tiny hole. And this hole limits the amount of light hitting the sensor and at the same time it provides an angular resolution to the image. Light can only reach the sensor when it comes from a certain direction, in any other case it is blocked.

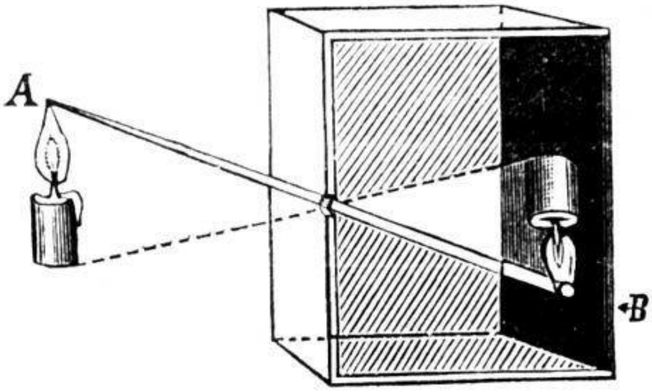

Figure 1 Camera obscura.

One of the main advantages of such a camera scheme is that it preserves straight lines. This is quite useful if one shoots some buildings or any other day-to-day objects. Also, pinhole camera does not require focusing the image: all objects in the shot remain sharp. The main drawback, however, is permanent lack of light. And this is obvious, since we block most of it and let only a tiny amount of it through the hole to hit the sensor. So, lenses are primarily used to compensate for that. They collect more light to produce a brighter picture which is much easier to capture. Most of the lenses do not change the geometry all that much i.e., straight lines remain straight. In order to prove it I took some pictures with lenses at my disposal. For the shooting I chose an A4 letter with the square grid printed on it. This is the easiest way to estimate if the perspective remains linear or not.

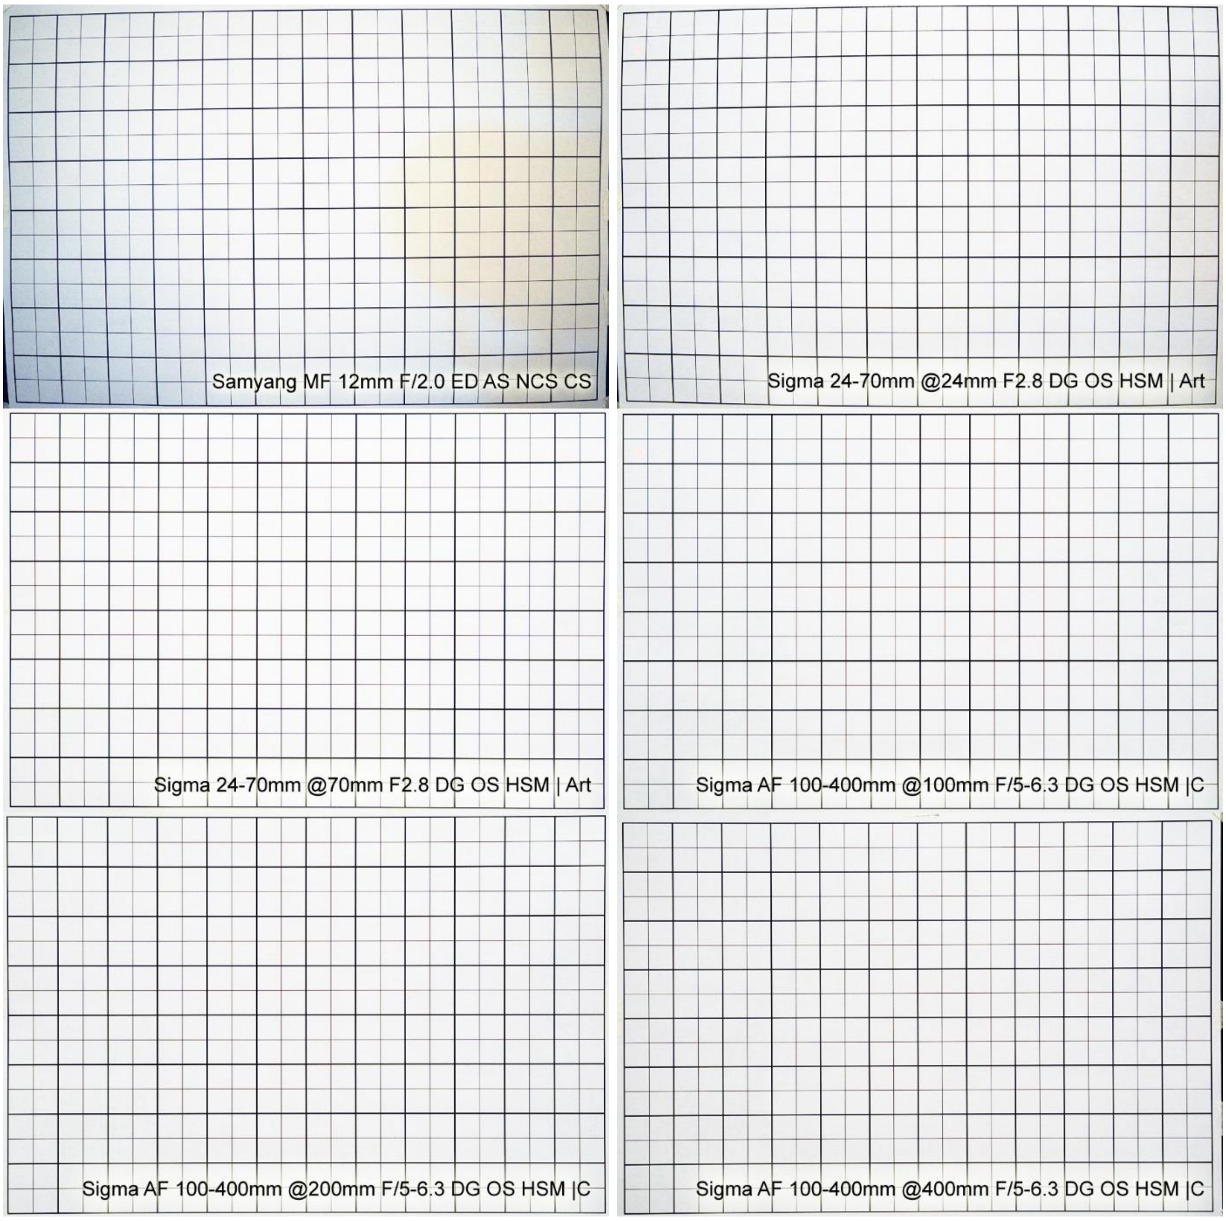

Figure 2 Distortion tests.

From the look of it, despite wide-angle lenses project moderate barrel distortion, overall images seem pretty rectangulary. From this result we can draw two conclusions. First: rectilinear lenses can be well approximated by their pinhole counterparts. Changing the distance between the hole and the sensor is equivalent to changing the focal length of the ordinary lens. And the second, not so obvious one: since the middle squares of the grid have the same size as ones near the corner scale of the image is not constant. It gradually increases from the middle to corner. By the scale here I mean the number of pixels that covers some arbitrary angle. If we take for example 1 degree sector and count the number of pixels it covers, we would find that near the edge of the frame the same angle covers more pixels. It can be illustrated on a simple diagram for the pinhole camera (this is why I started from this example due to its extreme usefulness).

It is clearly visible that the same blueish-purple rectangle occupies very different angles β and γ depending of its placement in the frame.

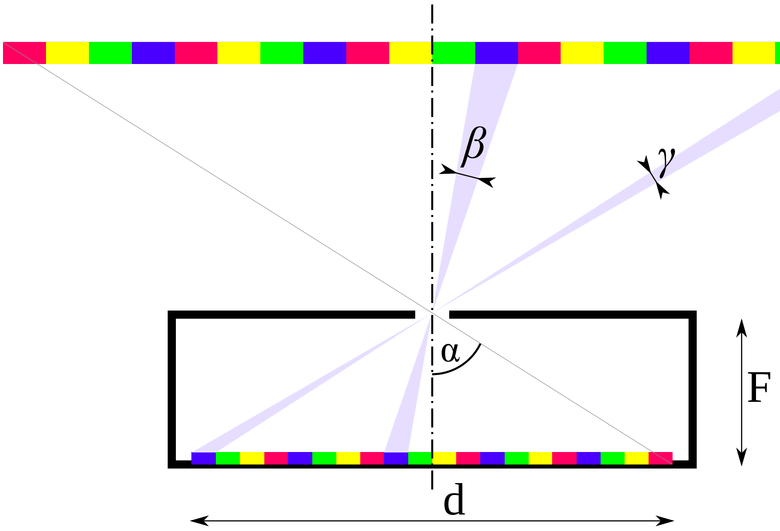

Figure 3 Illustration of the fact that objects projected onto the same number of pixels in a different part of the frame may occupy significantly different solid angles.

In order to estimate how different these angles truly are let's take a closer look:

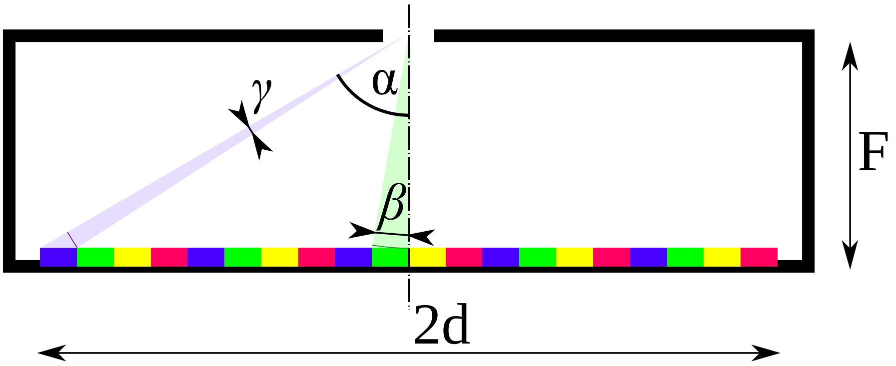

Figure 4 Closeup of Figure 3.

One can clearly see that most left point of our blueish-purple rectangle is defined by the focal length F, and by the sensor size, which here I chose it to be $2d$, so that the d would be equal to the horizontal distance from the optical axis. By the definition of tangent function:

$$\tan\alpha = \frac{d}{F} = \frac{x_{b_{left}}}{F} \tag{1}$$

Right corner of our rectangle of interest will be:

$$\tan(\alpha-\gamma) = \frac{x_{b_{right}}}{F} \tag{2}$$

And the linear distance between these points:

$$x_{b_{right}} - x_{b_{left}} = F(\tan(\alpha-\gamma) - \tan\alpha) \tag{3}$$

From the rectangle one of whose corners is right onto the optical axis we get the following expression:

$$x_{g_{right}} - x_{g_{left}} = F(0 - \tan\beta) \tag{4}$$

And we want these two distances to be equal:

$$x_{g_{right}} - x_{g_{left}} = x_{b_{right}} - x_{b_{left}} = q = -F\tan\beta = F(\tan(\alpha-\gamma) - \tan\alpha) \tag{5}$$

From here we can obtain our required relation:

$$\tan\alpha - \tan\beta = \tan(\alpha-\gamma) \tag{6}$$

After some trigonometric rearrangement which aren't that important here, we get:

$$\tan\gamma = \frac{\tan\beta}{\tan^2\alpha + 1 - \tan\alpha\tan\beta} \tag{7}$$

And we've arrived to a complete nightmare… However, if we choose our β and γ to be so small so we could use an approximation of $\tan x \sim x$, then:

$$\gamma = \frac{\beta}{\tan^2\alpha + 1 - \beta\tan\alpha} \tag{8}$$

A little bit more rearrangements bring us here:

$$\gamma = \frac{\beta}{\frac{1}{\cos^2\alpha} - \frac{\beta\sin\alpha\cos\alpha}{\cos^2\alpha}} = \frac{\beta\cos^2\alpha}{1 - 0.5\cdot 2\beta\sin\alpha\cos\alpha} = \frac{2\beta\cos^2\alpha}{2 - \beta\sin 2\alpha} \tag{9}$$

And since we chose β to be in radiance and it is small then $2 \gg \beta\sin 2\alpha$ and we can neglect that:

$$\gamma = \frac{2\beta\cos^2\alpha}{2 - \beta\sin 2\alpha} \approx \frac{2\beta\cos^2\alpha}{2} = \beta\cos^2\alpha \tag{10}$$

I.e., angle γ must be smaller by the factor of $\cos^2\alpha$ to equalize their projections onto the focal plane.

If it was too difficult to follow (and I wouldn't be surprised) I found an alternative and much easier way to derive the same result.

We know our desired β. And since it is small and it projects a section of the arc which is almost parallel to the focal plane, we can safely assume that:

$$q = x_{g_{right}} - x_{g_{left}} = F(0 - \tan\beta) \approx F\beta \tag{11}$$

Now, we can see that slanted beam drops onto a surface at an angle. And even if we pretend it to be parallel, it would spread over a larger area like so:

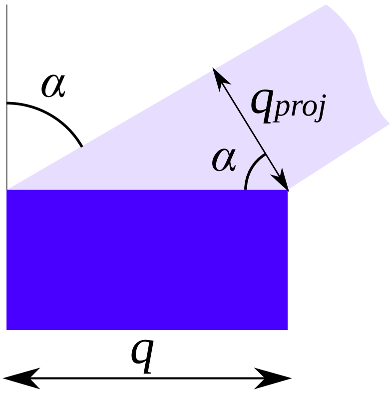

Figure 5 Closeup of the single rectangle near the frame corner.

We know (or rather set) our q. So, it's not hard to see that if we want to find the original width q_{proj} of the light beam, we must multiply q by $\cos\alpha$. But let's not forget that this beam traveled a longer distance as well. By how much? Again, it is greater than the original focal length by the factor of $1/\cos\alpha$ (see Figure 3). Therefore, in order to cover the same length on the sensor, angle γ must be smaller than the angle β by the factor of $\cos^2\alpha$. This fact will turn out to be important later.

For now, let's concentrate on the sky. It is rotating with the constant speed of 1 revolution per day and stellar angular velocity on the celestial equator is:

$$v = \frac{360° \cdot 60' \cdot 60"}{24h \cdot 60m \cdot 60s} = 15\frac{"}{s} \tag{12}$$

Here I use solar day duration instead of the stellar one for simplicity. Indeed, true stellar speed of $15.041\frac{"}{s}$ differs from its rounded value just for 0.27%. It is obvious that for a steady camera exposure can't be limitless. It must be short enough that stars wouldn't be able to leave the smudged tracks. Knowing that the modern cameras have the Bayer filter and that no lens can render point-like stars (diffraction is still a thing in our Universe) it is safe to assume that the star will still be round if it traveled less than $q = 3$ pixels through the sensor. And further we will confirm or disprove this suggestion.

In order to derive some quantitative predictions let's choose a lens with a sufficiently long focal length so that change of scale, I've talked about, is negligible. If we take the cropped sensor with dimensions of $22.2 \times 14.8\,mm$ (more commonly known as APS-C with the crop factor of $36/22.2 = 1.622$) and $200\,mm$ lens then we get total field of view of:

$$\alpha = 2\operatorname{atan}\frac{d}{2F} = 6.353° = 22872'' \tag{13}$$

Angle at which light hits the sensor's edge is only half of that. So, the change of scale relative to the middle is only $1/\cos^2(3.177°) = 1.0031$ i.e., is less than 0.31%. Therefore, we can safely assume that the angular resolution per pixel is the same across the frame and is equal to:

$$R = \frac{\alpha}{x} \tag{14}$$

For convenience I'll express x in terms of megapixels M. Taking into account the aspect-ratio of the sensor $(3{:}\,2)$ we have the following expression:

$$10^6 \cdot M = x \cdot y = x \cdot \frac{2}{3}x \quad => \quad x = 10^3\sqrt{\frac{3M}{2}} \tag{15}$$

Here x — is the number of pixels in the longest side. Linear pixel size is:

$$\mathrm{d}_{pix} = \frac{36\,mm}{x \cdot C_f} \tag{16}$$

In this equation $36\,mm$ is the traditional full-frame sensor size and C_f is the crop-factor of your camera model. Since we chose the limiting travelling distance of $q = 3\,px$ then we have to calculate the respective angle for that case by using the equation (13) but for the sensor size we should take the combination of two factors of $d_{pix} \cdot q$:

$$\alpha = \operatorname{atan}\frac{d_{pix} \cdot q}{F} = \operatorname{atan}\frac{q \cdot 36\,mm}{x \cdot C_f \cdot F} = \operatorname{atan}\frac{q \cdot 12\,mm \cdot \sqrt{6M}}{10^3 \cdot C_f \cdot F \cdot M} \tag{17}$$

Knowing that the stars angular velocity $v = 15\frac{''}{s}$ at the equator, in order to obtain the exposure duration, it is necessary to express the angle α from (17) in arc seconds and divide it by the v:

$$t = \frac{\alpha}{v} = \frac{3600}{15}\operatorname{atan}\frac{q \cdot 12\,mm \cdot \sqrt{6M}}{10^3 \cdot C_f \cdot F \cdot M} = 240\operatorname{atan}\frac{q \cdot 12\,mm \cdot \sqrt{6M}}{10^3 \cdot C_f \cdot F \cdot M} \tag{18}$$

This formula of course works when arctangent is calculated in degrees. For the radians one should use this slightly modified version of it:

$$t = \frac{\alpha}{v} = \frac{180 \cdot 3600}{15\pi}\operatorname{atan}\frac{q \cdot 12\,mm \cdot \sqrt{6M}}{10^3 \cdot C_f \cdot F \cdot M} = 13751\operatorname{atan}\frac{q \cdot 12\,mm \cdot \sqrt{6M}}{10^3 \cdot C_f \cdot F \cdot M} \tag{19}$$

And since the argument under the arctangent function is very small due to the 10^3 in the denominator we can use the approximation $\operatorname{atan} x \sim x$:

$$t = 13751\operatorname{atan}\frac{q \cdot 12\,mm \cdot \sqrt{6M}}{10^3 \cdot C_f \cdot F \cdot M} = 404.2\frac{q \cdot \sqrt{M}}{C_f \cdot F \cdot M} \tag{20}$$

Multiplying by F we get:

$$Ft = 404.2\frac{q\sqrt{M}}{C_f \cdot M} \tag{21}$$

So, by setting our tolerance to the star-tracks q, and by knowing our crop-factor and sensor resolution we can easily determine which "rule" suits us the best. For example, for 24 MP sensor, crop-factor of 1.622 and $q = 3$ we get: $Ft \approx 155$. Nowhere near to 600! And no wonder, this rule was derived into film era, when no one suspected about digital revolution that was about to happened.

At this stage we've done most of the work. There couple of things to be discussed. First one is what happens at the edge of wide-angle lenses, when the angular resolution depends on the position in the frame. But even this case was implicitly covered by us in formula (10). If we want to freeze the stars across the whole frame, we just need to reduce the exposure which is given by formulae (18), (19) or (20). And the amount of that reduction is exactly that $\cos^2(\alpha/2)$ which we obtained earlier. Here α — is the total field of view of the lens from (13) therefore we need only half of that. By using trigonometry, it is possible to obtain the following equality:

$$\cos[\operatorname{atan}(x)] = \sqrt{\frac{1}{(x^2+1)}} \tag{22}$$

So, taking into account, that $x = d/(2F) = \frac{36\,mm}{2 \cdot F \cdot C_f} = \frac{18\,mm}{F \cdot C_f}$ we get:

$$\cos\left[\operatorname{atan}\left(\frac{18\,mm}{F \cdot C_f}\right)\right] = \sqrt{\frac{1}{\left(\left[\frac{18\,mm}{F \cdot C_f}\right]^2 + 1\right)}} = \sqrt{\frac{F^2 C_f^2}{324 + F^2 C_f^2}} \tag{23}$$

And the exposure for the corners is equal to:

$$t = 404.2 \frac{q \cdot F \cdot C_f}{\sqrt{M}\left(324 + F^2 C_f^2\right)} \tag{24}$$

And just for convenience let's put these two formulae next to each other:

$$t_{middle} = 404.2 \frac{q}{C_f \cdot F \cdot \sqrt{M}} \qquad t_{corner} = 404.2 \frac{q \cdot F \cdot C_f}{\sqrt{M}\left(324 + F^2 C_f^2\right)} \tag{25}$$

The only unknown that remained so far is the factor q in our calculation. And instead of theorizing about how star is rendered through the Bayer filter and so on and so forth I suggest to make a simple experiment. Go to the field and shoot some stars. Or, if the weather isn't great: set the smartphone with the glowing pixels as stars and make some pictures of it at various exposures, while camera is standing on a tripod and rotating with the stellar speed. Which is exactly what I did. You can see the results presented in the table below:

Table 1 Snippets of the shots containing artificial star(s).

12 mm		24 mm		70 mm		100 mm	
middle	corner	middle	corner	middle	corner	middle	corner
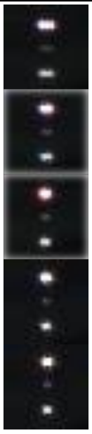	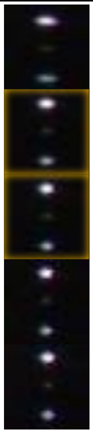	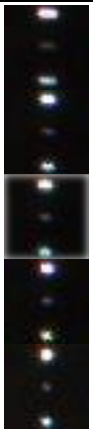	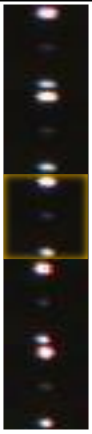	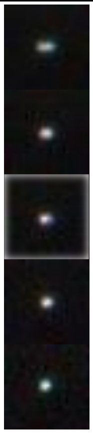	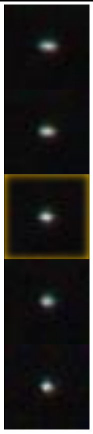	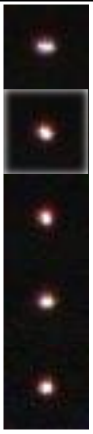	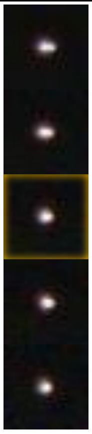

Here I must point out that although exposure in each column decreases from top to bottom it differs not only between various focal lengths but even for the middle and corner images. And it is pretty difficult to distinguish between the real star movement and the lens aberration for wide angle lenses even though I've set the aperture at $f/8$ for all images.

Therefore, these results aren't precision-like but still, we can obtain what we wanted (q value).

Images which were taken with the longest exposure and didn't show significant motion blur (or comparable with the aberration levels) have been highlighted.

Finally, I've plotted theoretical curves for 24-megapixel sensor and value of $q = 3.5$ and added points which correspond to my shots. This is what I got:

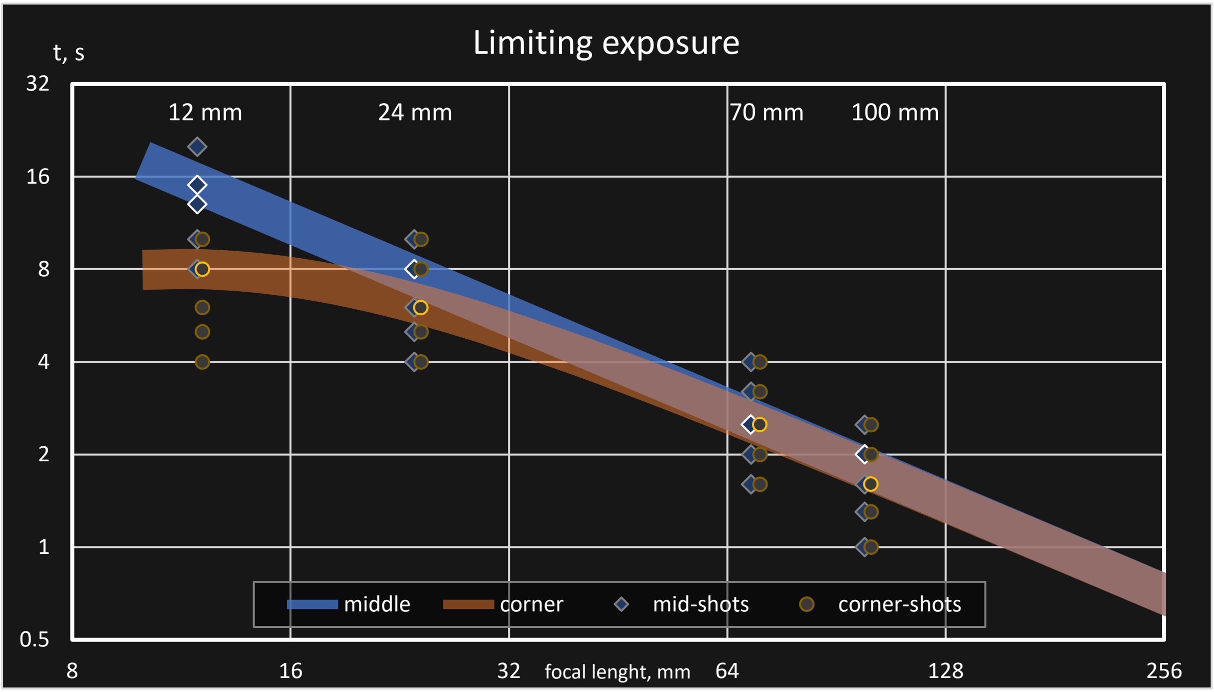

Figure 6 Experimental results of shooting an artificial stellar field on a moving equatorial mount.

One can see that the agreement between experimental points and theoretical curves (for middle of the frame and its corner) is pretty reasonable for the chosen value of $q = 3.5 \pm 0.5$.

So, we can simplify our formulae (25) with this result:

$$t_{middle} = \frac{1400}{C_f \cdot F \cdot \sqrt{M}} \qquad t_{corner} = \frac{1400 \cdot F \cdot C_f}{\sqrt{M}\left(324 + F^2 C_f^2\right)} \tag{26}$$

Since we've already established that 600 rule was valid only for the full-frame cameras equipped with the moderately telephoto lenses let's try to estimate the resolution of the film. Even if we allow the stars to be somewhat elongated i.e., set q to be 4.5 instead of 3.5 our product of $F \cdot t_{middle}$ has to equal 600:

$$t_{middle} = 404.2\frac{q}{C_f \cdot F \cdot \sqrt{M}} \quad => \quad F \cdot t_{middle} = 600 = 404.2\frac{4.5}{1 \cdot \sqrt{M}} \tag{27}$$

From here we can derive effective resolution:

$$600 = \frac{1800}{1 \cdot \sqrt{M}} \quad => \quad M = \left(\frac{1800}{600}\right)^2 = 9 \tag{28}$$

So, the effective resolution is somewhere around 9 megapixels which corresponds to an average graine size of 10 microns. I find this result quite plausible due to my experience of scanning our home archives of film. Most of the scanned images have optimal resolution at 2400 dpi and they didn't benefit from setting it higher than that. And this DPI corresponds exactly to a 10-micron step. One could suggest that all higher DPI are fictitious and obtained via digital interpolation but even visual observations of the raw scans confirmed that for the film with the high ISO DPI could be set even lower because the grain was very prominent and the resolution was limited by the film and not the scanner.

To summarize: the 600 rule is a relic of the film-age. It was based both on the resolution of the film and the old lenses most of which have much more collection value than the optical one. In order to determine which rule is best suited to your setup you must take into account real resolution of your camera. And even when you do that you have to remember that with wide-angle lenses some corrections may be required. At the end of this article, I would like to present a couple of plots which may be helpful to you at night, when the calculations aren't that easy to perform.

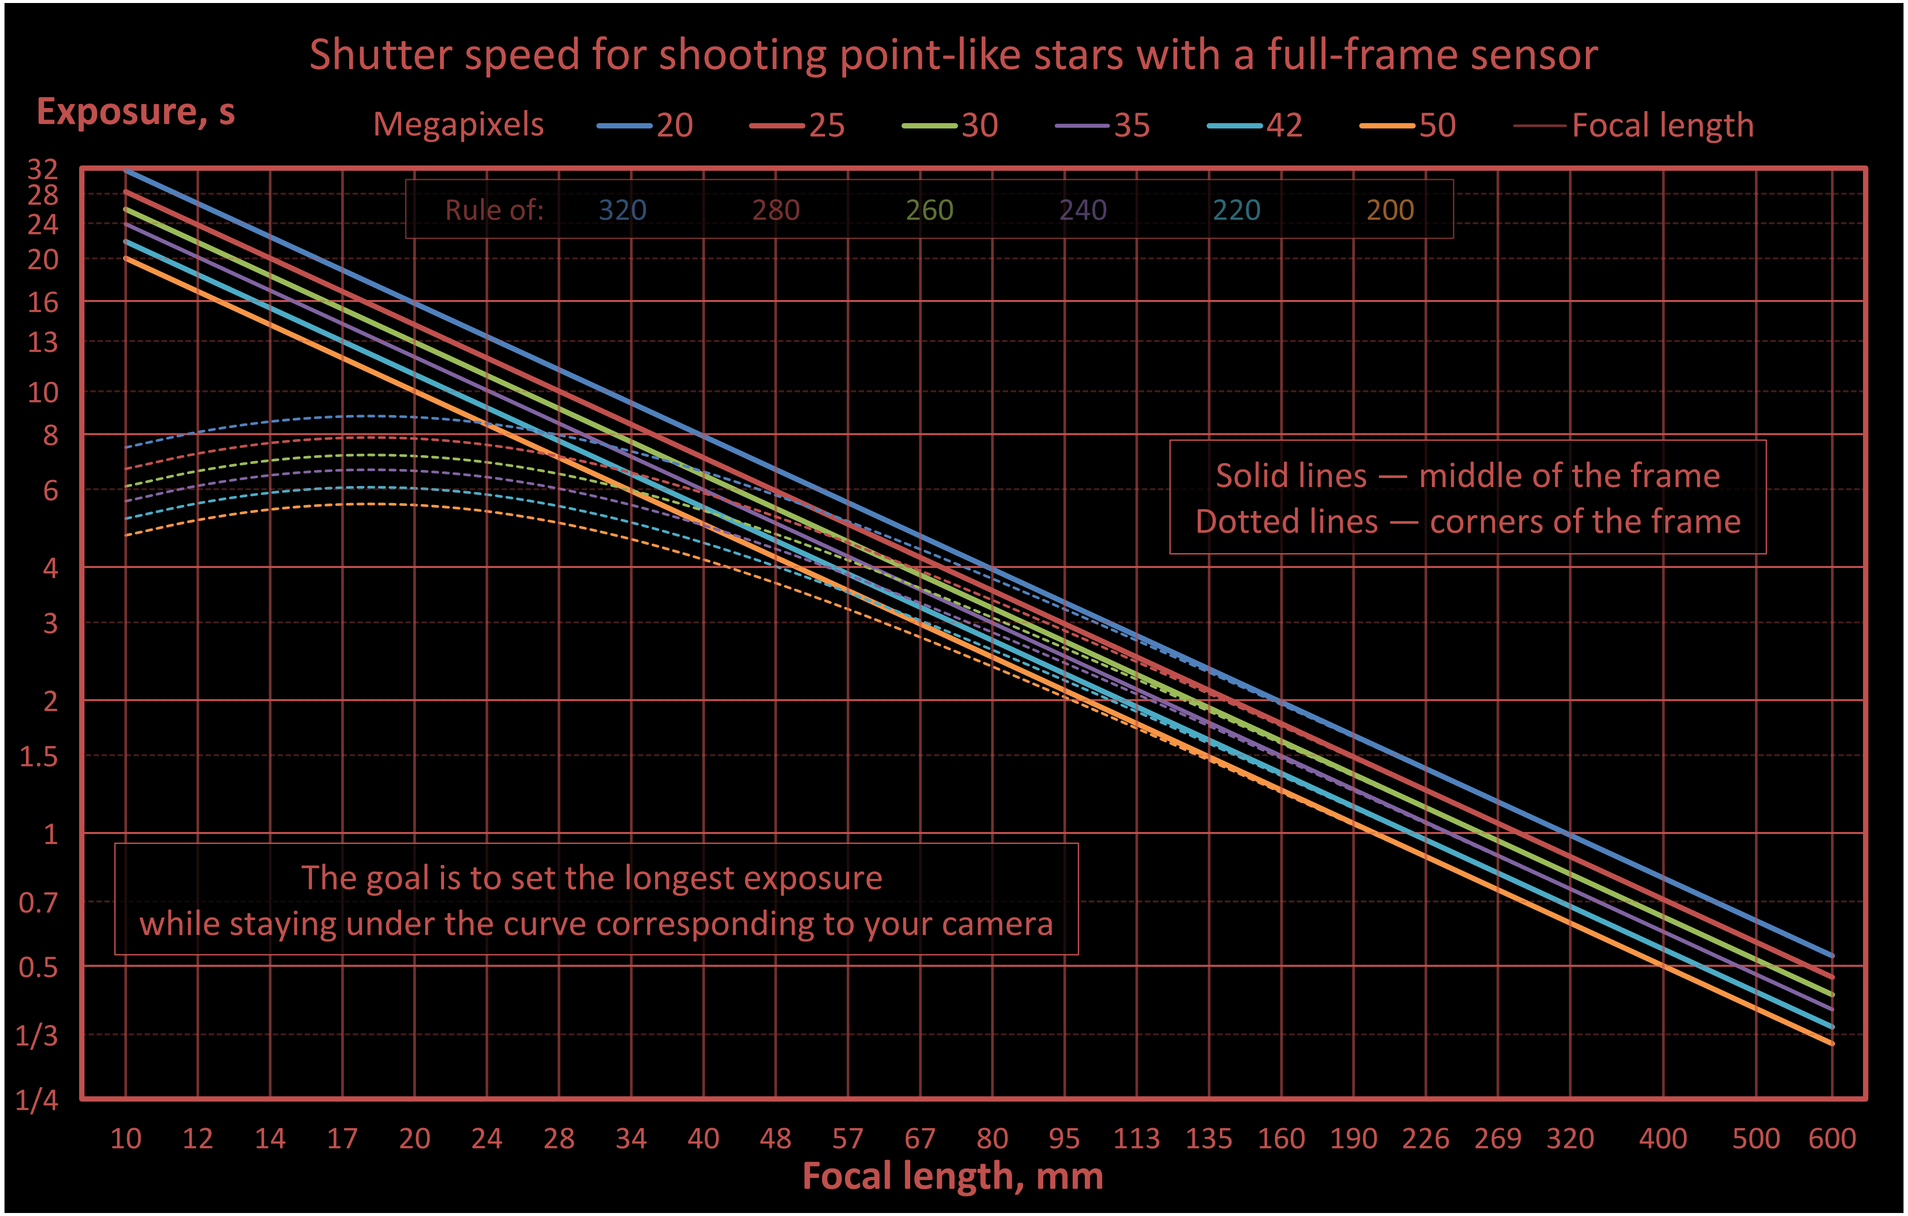

Shutter speed for shooting point-like stars with a full-frame sensor

Solid lines — middle of the frame
Dotted lines — corners of the frame

The goal is to set the longest exposure
while staying under the curve corresponding to your camera

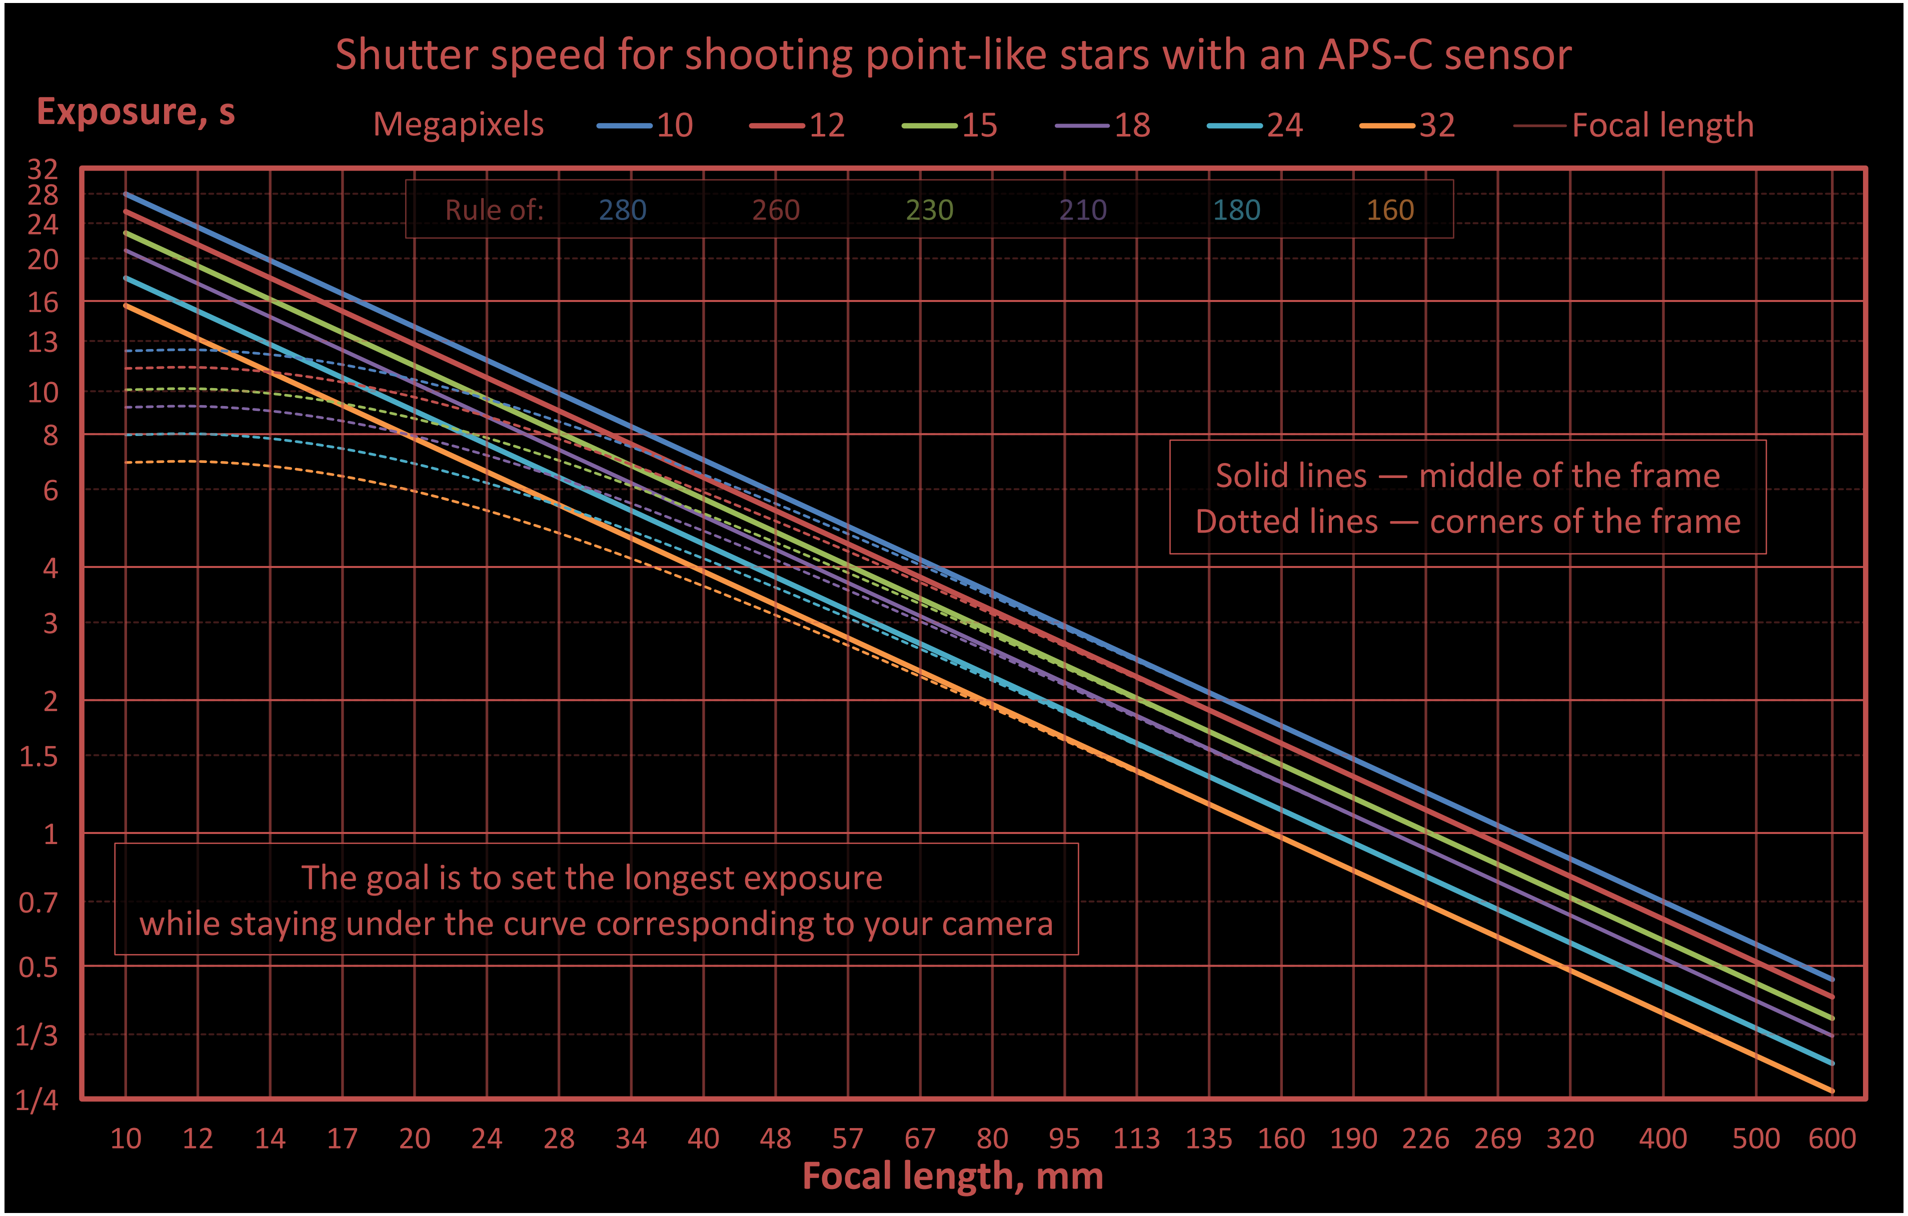

Shutter speed for shooting point-like stars with an APS-C sensor

Solid lines — middle of the frame
Dotted lines — corners of the frame

The goal is to set the longest exposure
while staying under the curve corresponding to your camera