https://doi.org/10.2352/ISSN.2470-1173.2020.16.AVM-038

Describing and Sampling the LED Flicker Signal

Robert C Sumner; Imatest LLC; Boulder, Colorado

Abstract

High-frequency flickering light sources such as pulse-width modulated LEDs can cause image sensors to record incorrect levels. We describe a model with a loose set of assumptions (encompassing multi-exposure HDR schemes) which can be used to define the Flicker Signal, a continuous function of time based on the phase relationship between the light source and exposure window. Analysis of the shape of this signal yields a characterization of the camera's response to a flickering light source–typically seen as an undesirable susceptibility–under a given set of parameters. Flicker Signal calculations are made on discrete samplings measured from image data. Sampling the signal is difficult, however, because it is a function of many parameters, including properties of the light source (frequency, duty cycle, intensity) and properties of the imaging system (exposure scheme, frame rate, row readout time). Moreover, there are degenerate scenarios where sufficient sampling is difficult to obtain. We present a computational approach for determining the evidence (region of interest, duration of test video) necessary to get coverage of this signal sufficient for characterization from a practical test lab setup.

Introduction

Pulse-Width Modulated (PWM) lights change their apparent brightness by turning off for some portion of the time at a frequency too high for humans to perceive. Though they may appear constant-brightness to humans, to cameras which may operate on roughly the same time scale as the PWM frequency the light source fluctuations may become painfully obvious. This results in a "flickering" effect in videos and "banding" effect in still images. Though this effect may be evoked by any temporally-modulated light source, we will simply call this LED flicker–or just flicker–for the remainder due to the increasing prevalence of PWD LED light sources.

The automotive imaging industry has become especially aware of this effect [1] due to a confluence of the wide-ranging integration times used (required by the various very-high- to low-light scenarios automobiles encounter daily) and many non-standardized LED sources "in the wild". These often include head-and tail-lights on the vehicles themselves and information-bearing street signage.

A challenge in determining the flicker susceptibility of a camera is knowing the relevant parameters to measure over and their ranges. Moreover, once the light source's and camera's properties for a test have been determined, there is also a question of what to from image or video data such that the measurements are meaningful and complete.

In this paper we show that this problem is based around sampling of a continuous function on a finite domain- the "flicker signal". We describe how samples of this function can be measured from video data of a standard test target and how to ensure the sampling is sufficient to capture the shape of the function. Using such a testing regimen, the flicker signal can be captured and various metrics can be derived.

Previous work

The phenomenon of illumination-induced image-banding is a well-known one in the world of consumer CMOS sensors due to mains-powered lighting which operates at 50 or 60 Hz. This effect can largely be eliminated [2] due to the predictability of those frequencies, a boon which is not extended to the automotive field where there is no consistency in PWM frequency.

Automotive sensor manufacturers have begun producing sensors with "LED Flicker Mitigation" (LFM) schemes built in [3], but the problem is far from solved for all cases and standards for testing still need to be established.

Imaging Model

The imaging model assumed in this paper is based upon the exposure of a camera pixel to the light source and its subsequent digital number. No spatial effects such as spatial frequency response or stray light in the imaging system are accounted for, only the mapping from light received at the sensor to its output value. We have minimal assumptions about the light source and camera, as follows.

LED assumptions: The PWM LED output pattern, subsequently called the light signal, is assumed to be regular. No assumptions are made about the duty cycle or shape of the light-on pulse, only that it repeats with some minimum regular interval, t_{LED}.

Exposure Scheme assumptions: Like the LED signal, we assume the frame capture timing (*i.e.*, frame rate) of the camera is regular, with no jitter or varying rate. The inverse of the frame rate is called t_{Frame}.

We call the combination of the light-gathering by the sensor during the exposure window and the subsequent digital tone mapping an exposure scheme. Our assumption is that this scheme is the same across all frames, *i.e.*, temporally invariant. For example, this means that auto-exposure is not actively adjusting operation over the video.

Importantly, this model does not assume any particular exposure scheme. There is no assumption about either the way in which the exposure window is used to gather light or about how the sensor output is mapped to the final DN. For example, multiple-exposure schemes used to increase dynamic range [4] may split the exposure window however they want between the sub-exposures, as illustrated in Figure 1.

Likewise, any tone mapping mechanism may be used, under the loose assumption of being spatially non-varying, *e.g.*, not adaptive tone mapping based on local structure.

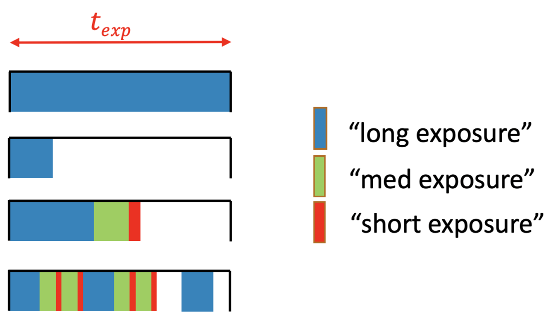

Figure 1. *The exposure window may be used in any way by sub-exposures as long as the scheme is consistent across all frames.*

Timing Diagram

Figure 2 details the timing diagram of the light signal and the camera exposure windows. A "row-readout" style sensor is illustrated where each row of the sensor integrates light over a slightly different window of time. Typically, this is due to the electronic readout architecture being limited in size to hold only a single row's worth of data (due to cost or silicon space constraints). The exposure window of each row is thus offset from the previous row by some fixed amount, t_{row}, to accommodate the amount of time it takes the electronics to transmit this buffer to some other location.

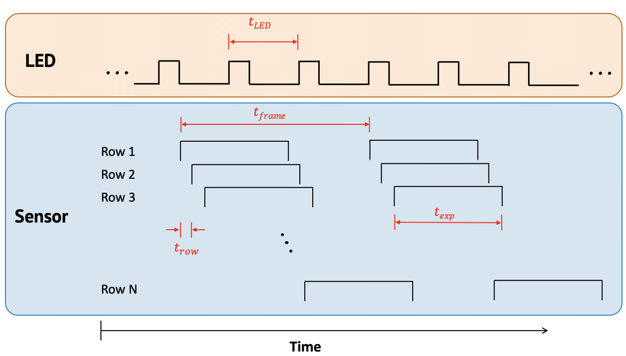

Figure 2. *Timing diagram of a PWM LED light signal and "row-readout" sensor with video frame rate* $\frac{1}{t_{frame}}$.

In general, this definition of a row-exposure camera is not required to define the flicker signal. The camera may be global shutter (equivalent to only one unique row in Figure 2) or not meet the assumption of regular offsets in row-wise exposure windows.

Note that the exposure is shown as a contiguous block of length t_{exp} (which is by definition not longer than t_{Frame}), but this is only for illustrative purposes. This time can be divided into non-contiguous sub-exposures of any design, as previously stated. t_{Frame} is the maximum this "exposure window" can be, though we will use that term interchangeably with whatever subset of this period is actually used to capture light.

Canonical Flicker Period

The single most important observation leading to the subsequent analysis is that the pixel response to the light signal is periodic. This comes directly from the simple model constraints of temporal invariance and periodicity of the camera and light components, as previously stated. As illustrated in Figure 3, the temporal pattern of light the pixel sensor is exposed to repeats after t_{LED} seconds due to the light signal's periodicity. If the exposure window started an integer multiple of t_{LED} seconds later it would be exposed to the same pattern and result in the same pixel value.

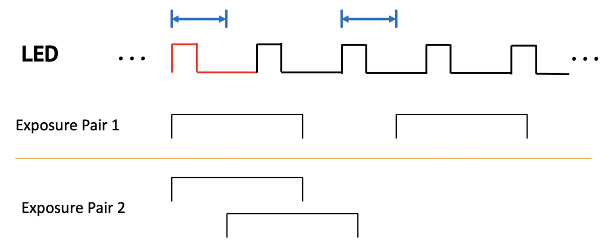

Figure 3. *The offset of the exposure window relative to a given period of the LED light signal defines the flicker signal. The two exposures of the Pair 1 capture the same two sampling points of the flicker signal as Pair 2.*

Since this function is periodic, we can fully characterize it by defining it over a single period. Thus, the only quantity relevant to unique pixel responses is the offset of the exposure window relative to the light signal period. We call this period the **canonical flicker period**.

This period is constrained to being the lesser of t_{Frame} or t_{LED}. In practice, since many sensors' frame rates are either 30 or 60 frames per second (FPS) and the human threshold for perceiving flickering lights as constant-brightness is around 60Hz, t_{LED} is the limiting factor here. As sensor data throughput becomes faster and frame rates increase to 120 FPS and beyond, t_{Frame} may become the limiting factor, but the subsequent analysis will remain the same. For simplicity, we will use t_{LED} as the length of this period in discussions here.

Note that a special case of this is a standard rolling-shutter linear sensor where light is simply integrated by pixels in the row when they are in their window. This linearity assumption combined with the assumed time-invariance reduces the output flicker signal to a convolution of light signal and a box function in time representing the integration. In general, however, the exposure is not a simple integration and thus non-linear, and the output will be periodic but not a convolution.

Flicker Signal

We define the flicker signal simply as a pixel's digital response to the light signal over the canonical flicker period. Due to the nature of the domain on which it is defined, this signal is continuous and only over a finite interval.

While this signal is easily seen as a function of time, it actually has many dimensions since it depends on the light source and camera properties.

$$F(t, \phi_{LED}, \phi_{cam}) \tag{1}$$

ϕ_{LED} is a vector of relevant properties of the LED light signal, including its frequency, power, duty cycle, and pulse shape, and spectral distribution. ϕ_{cam} is a vector of relevant camera parameters, including frame rate, exposure scheme, and tone mapping scheme.

Of the above parameters, LED frequency and camera frame rate hold special positions- these determine the canonical flicker period domain itself while the other parameter define the response over this domain. They do not have any impact on the signal other than determination of this period. Since they do not have any other effect and the common use case is to describe this signal for a given LED rate and a given frame rate, it is often useful to remove these from the sets of variables above (labeling the remaining sets in primed form), and instead identify the function as parameterized by this period.

$$F_{t_{LED}}(t, \phi'_{LED}, \phi'_{cam}) \tag{2}$$

Of course, this function of many variables will typically also need to be mapped over many different LED frequencies which may be encountered in the wild. Also, for the sake of simplicity in the remainder, we will often go the other way and discuss "the flicker signal" as only a function of time, assuming a given set of the other parameters involved.

Note that we do not define this function on any particular channel of image data, such as the linear digital number straight from the sensor or any selection of color channel value at that location in a demosaicked, colorspace-encoded image. The flicker signal may be any one of these, and more, defined over the canonical flicker period. In essence, the flicker signal is more of a way of defining a function of interest rather than the specific function itself.

Though the canonical flicker period on which this signal is defined is fundamentally an interval of time, it is sometimes useful to refer to it defined over "phase angle", often in degrees.

$$\theta = 360 \cdot \frac{t}{t_{LED}} \tag{3}$$

This allows for comparison of relative proportions of the signal defined with one canonical period (*i.e.*, LED period) with that defined on a different canonical period.

Manifestations

A primary benefit of defining the flicker signal as we have is that it unifies commonly observed spatial and temporal components under a single explanation.

Banding within a frame

Spatial effects manifest as bands in otherwise-uniform areas of an image, as seen in Figure 4. This is effectively a sampling of the flicker signal every t_{row} seconds. While this can be an effective means of densely sampling the flicker signal, it is not always possible to capture such tall regions exhibiting this effect in test lab set ups.

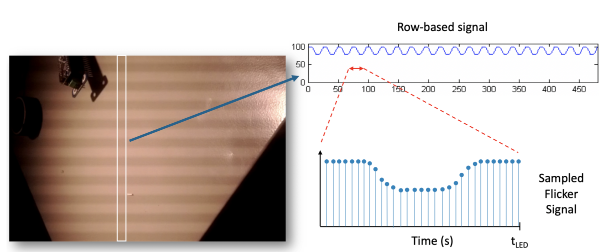

Figure 4. *Plotting data down a column of the image of a row-readout sensor and mapping one period of it to the flicker signal.*

Flickering across frames

It is less straightforward to understand temporal flicker as being generated in the same way as the spatial bands. The process may seem less predictable because the sampling point on the flicker signal may jump around in what seems to an observer to be unpredictable ways, as illustrated in Figure 5.

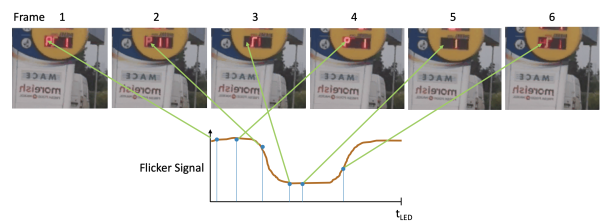

Figure 5. *Plotting one sample location across subsequent frames and mapping it to the flicker signal.*

Key Performance Indicators of the flicker signal

If the flicker signal is known, it is natural to attempt to reduce this continuous function to a few key performance indicators (KPIs) of interest. KPIs are typically chosen to act as metrics of the level of "goodness" of a camera, which in this scenario often means minimal presence of flicker in some sense. Depending on the use case, the goal may be, *e.g.*, to eliminate flicker at a specific LED frequency road signs are known to operate at, or to avoid highly-objectionable sharp transitions in banding effects in areas illuminated by a car's own tail lights in a back-up camera.

Figure 6 illustrates a number of potential metrics which could be derived from the flicker signal for a given LED frequency. We repeat that these could be derived from any form of image data appropriate for the application and stage of the image processing pipeline–*e.g.*, linearly encoded or not, any color channel, etc.

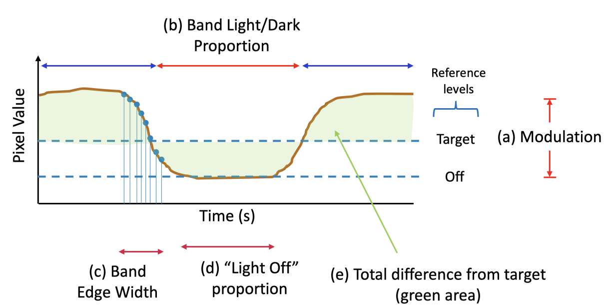

Figure 6. *Illustration of metrics which may be derived from the flicker signal. Pixel value represents and arbitrary channel or encoding of the pixel data to measure.*

Different subsets of these metrics are available if certain things are known about the system. If only the pixel value data is seen, we can measure:

(a) Modulation or Contrast between min and max value. A number of definitions of contrast could be used: simple difference or contrast ratio, Weber, Michelson, etc.

If the row-read time of the sensor is known, we can further determine the following metrics by determining the span of signal features in number of multiples of t_{row}:

(b) Proportion of light and dark band width.

(c) Band transition width, in number of rows.

Furthermore, if some reference pixel values are known for true off and true on (or "target") states of the light, we can define metrics such as the following:

(d) Proportion of time the light is seen as off.

(e) Total area of difference from true level.

Reference-off may be determined from pixel level when the target light is fully off (note that this doesn't equate to zero pixel level because of other reflections off the light source itself). Reference-on value may be determined by measurement of a constant-current light source of the equivalent "brightness" as the PWM light. The definition of equivalent brightness is not fully apparent, but two possible ways of defining this are psychometric study with a human indicating what level is equivalent, or using a slow response luminance meter to determine equivalent light output of the two.

Note that the metric of "rolling band speed" over a video sequence–how fast the bands march up or down the frame–is actually determined by the camera frame rate and light signal frequency. No measurements need to be made.

Any KPI measured from a flicker signal is only relevant for the given LED and camera operating point parameterizing that test. A full report of a flicker-indicating KPI will need to include such measurements for many LED frequencies, LED brightnesses, and camera operating modes (*e.g.*, exposure times). A family of color-coded heat maps indicating the KPI performance over these parameters may prove to be appropriate for conveying such information and indicating trouble areas. Worst-case analysis over a set of relevant use-case parameters will also simplify this space.

Sampling the flicker signal

Knowing that the flicker signal, if known, can describe the spatiotemporal effects which are often seen as undesirable, the logical next question is how to measure it. This turns into a sampling problem, as we only observe discrete samples of this continuous function with each row or frame sampling it.

Patch-ROI from Video measurements

We describe here a method of sampling the flicker signal from video data of a standard test lab setup. The scheme is similar in principle to that in [5], but rather than having a single flat-field flickering target, we recommend using a transmissive test chart with patches of many different densities. Using a target such as the Imatest UHDR 36-patch target shown in Figure 7, the brightness dimension of the space can be sampled with every single capture, effectively reducing the number captures needed to map out the full space.

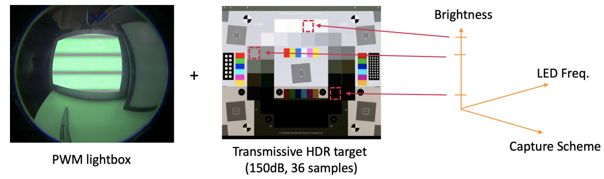

Figure 7. *Flickering lightbox test setup with transmissive target yields patch-based regions, sampling the light source brightness dimension for a given LED frequency with each video capture.*

Total Set of Phase Samplings

Each patch on this target yields a region of interest (ROI) in the image. We assume that the target and lighting is uniform over each ROI, and that the exposure scheme and tone-mapping is the same for all pixels in a given ROI. Finally, we assume in the following analysis that any row-wise architecture means that all pixels in a row are exposed during the same interval and that these exposures are offset by the fixed interval t_{row}.

We consider an ROI which subtends M rows and N columns of the image, extending over F frames of a test video. Having more than one column in an ROI does not add new sampling points to the test, but it does provide redundant data which can be used to reduce noise in the measurement.

The full set of phase samplings over the canonical flicker period for these observed data is given by Eq. 4, where $\delta(t)$ is the discrete delta function: $\delta(t) = 1$ when $t = 0$, and 0 otherwise.

$$s(t) = \sum_{f=0}^{F-1} \sum_{r=0}^{R-1} \delta\left(t - \left(f \cdot t_{frame} + r \cdot t_{row}\right) \bmod t_{LED}\right) \qquad (4)$$

An arbitrary offset term can be added to the inner-most parentheses, representing some unknown offset between the first exposure window and light signal period, but it does not substantially change the analysis.

$s(t)$ represents the sampling points on the canonical period. It is reasonable to want to know the coverage of this set on this period, and how much of the flicker signal may be missed by gaps in this set of samplings. Fortunately, while analysis of this equation as written can be very difficult it is easily simulated.

Maximum Phase Gap analysis

Figure 8 shows a phase gap analysis plot generated by simulation. The comb-like clusters are indicative of the 20 rows of the ROI used, each giving a fine sampling of flicker signal phases. Each cluster comes from a separate frame of video.

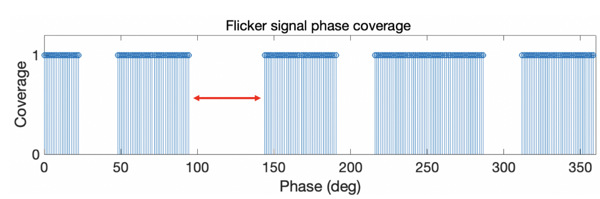

Figure 8. *Phase gap analysis for $t_{LED} = 1/82s$, $t_{Frame} = 1/30s$, $t_{row} = 40\mu s$, $M = 20$ rows, and $F = 10$ frames. Maximum phase gap: 50 deg.*

The most relevant feature of this analysis to observe is the *maximum gap* between any two samples, indicated in with the red arrow. This is the largest blind spot where the flicker signal may or may not exhibit some interesting behavior which we need to measure. No measure of smoothness has yet been guaranteed for flicker signals, so the best way to ensure the validity of metrics defined on it is to keep this worst case gap as small as possible.

Using this sort of target and measurement scheme, the obvious way of minimizing this maximum gap is by increasing the duration of the video or the number of rows in the ROI. Unfortunately, there are some degenerate cases where you cannot get new information by increasing the number of frames- when the flicker rate is an integer multiple of the frame rate. There is also the case of a global shutter camera which does not benefit from multiple rows being sampled per patch.

Figure 9 shows the maximum phase gap derived from a set of these simulations over a range of LED flicker frequencies, again assuming the rest of the parameters of the simulation are the same. Such plots show massive spikes at frequencies which are integer multiples of the frame rate due to the banding effect being temporally stationary in these cases (bands move up and down the frame at the beat frequency of the light and camera frequencies).

Such structure is fairly typical of these maximum-phase-gap-over-frequency plots, with both the width and height of the high-gap "dead zones" increasing as the number of rows and frames sampled go down. Note that while it becomes virtually impossible with this measurement scheme to guarantee that you can test with satisfactory coverage at every LED frequency, it implies that you can typically find a frequency close to any one of interest at which you will be able to gather enough data at to make the maximum gap sufficiently small.

For example measuring the flicker signal at 151Hz instead of 150Hz exactly brings the maximum phase gap in the scenario above from 318 degrees to 2.2 degrees. In practice, we often find that with only 20 rows in an ROI and 2-10 seconds of video it is possible to get a comprehensive set of test frequencies to have maximum phase gap below 10 degrees.

Sampling without understanding the flicker signal

Note that the above sounds very similar to the naïve approach of averaging all pixels in a uniform-patch ROI in each frame to get a time-series measurement from a single value per frame. However, simply constructing a time-series over the video duration from mean spatial data to describe video flicker can allow spatial effects can leak in if you use more than one row. "Soft edges" are often reported in such cases where the ROI is too tall and averages pixel data which straddles multiple parts of the flicker signal. This prompts selection of a smaller number of rows. However, by discarding these rows you lose sampling coverage which could help you uncover the true shape of the flicker by closing the maximum phase gap, possibly requiring significantly longer video data or entirely missing relevant structure.

Our approach is set apart by taking into account where to position each measurement in the canonical flicker period. As seen in Figure 9, adding information *correctly* from multiple rows in an ROI has a significant effect on sampling coverage. Constructing the flicker signal this way and then deriving KPIs from it allows us to make sense of spatial effects and temporal effects simultaneously. This ability comes simply from knowing some key facts about the test setup and system- the frame rate, LED frequency, and row readout time (for a row-wise sensor).

Conclusions

In this paper we put forth a model which unifies the spatial and temporal effects of a high-frequency modulated light source on image and video data. The model is based upon simple assumptions of temporal consistency and periodicity. This model uses knowledge of the relevant timing parameters to align measurements both from within a single frame and across multiple frames to construct this flicker signal, from which meaningful KPI can be derived. This approach is superior to previous work for measuring flickering video data which only considered down-the-image or across-video-frame data directly as the signal of interest because it gives insight into the information necessary to map out the true shape of this function.

Testing with this model in mind helps reduce the likelihood of missing an unexpected effect which more naïve pixel value observation may not produce sufficient test coverage to catch. A method was presented for determining what is a sufficient amount of test data for a standard test lab setup to further guarantee that there is no such "unexpected behavior" which was not observed in testing, and for finding LED frequencies which allow testing with a reasonable amount of video data.

Acknowledgments

The author thanks Dr. Paul Romanczyk, as well as Brian Deegan and the rest of the IEEE P2020 Subgroup 1 members for their discussion and contributions to this work.

References

[1] Members of the IEEE P2020 Working Group, "IEEE P2020 automotive imaging white paper," Tech. Rep., 2018.

[2] D. Poplin, "An automatic flicker detection method for embedded camera systems," *IEEE Transactions on Consumer Electronics*, vol. 52, no. 2, pp. 308–311, May 2006.

[3] On Semiconductor, "LED flicker mitigation for automotive applications featuring the AR0231LFM," 2017, [Online; accessed 10–Feb–2020]. Available: https://www.onsemi.com/blog/automotive/led-flicker-mitigation-for-auto-applications

[4] S. Mann and R. Picard, "On being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures," *Proceedings of IS&T 48th Annual Conference Society for Imaging Science and Technology Annual Conference*, no. 323, pp. 422–428, 1995.

[5] B. Deegan, "LED flicker: Root cause, impact and measurement for automotive imaging applications," *Electronic Imaging*, vol. 2018, no. 17, pp. 146–1–146–6, 2018.

Author Biography

Robert C Sumner is an imaging science manager at Imatest, LLC. in Boulder, Colorado. He holds a BS in Cinema and Photography from Ithaca College and an MS in Electrical Engineering from Boston University. He is a member of the IEEE P2020 working group.

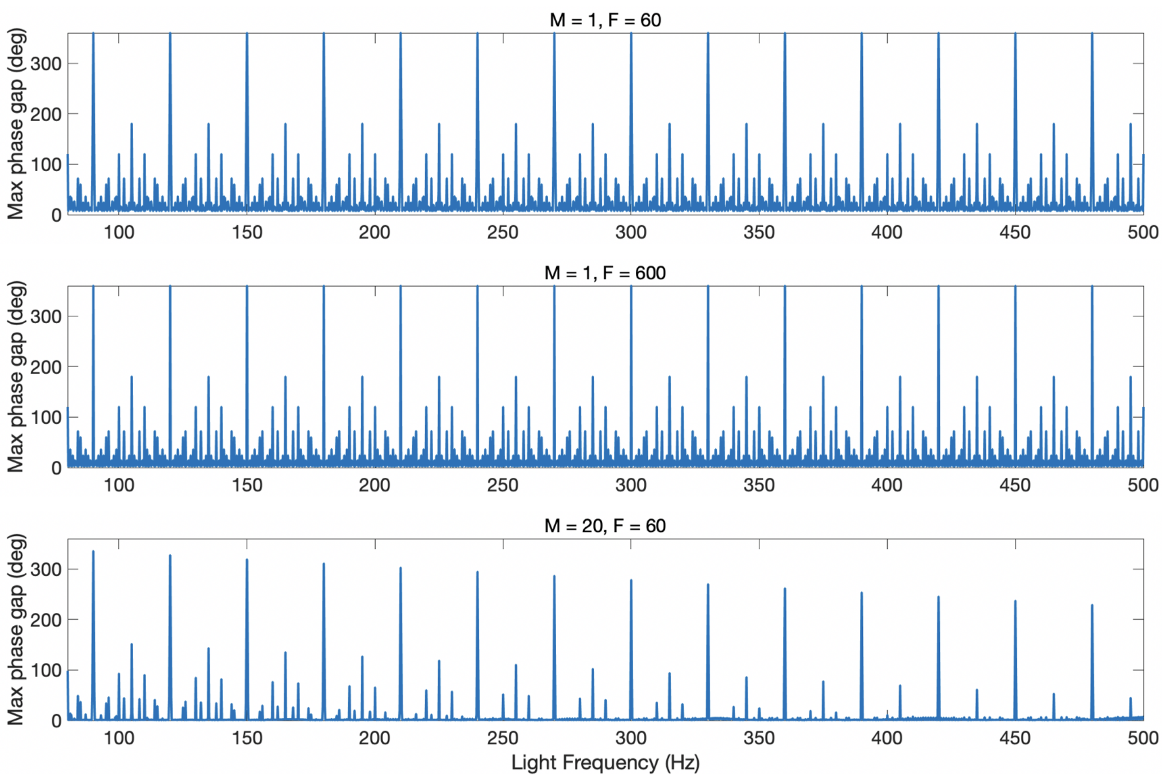

Figure 9. *Maximum phase gap analysis over LED frequencies from 80 to 500Hz, $t_{Frame} = 1/30$s, $t_{row} = 40\mu$s. The top plot shows a single row of data measured for 60 frames, while the middle and bottom plots add more frames and more rows to to the observation, respectively.*