Benefits of coherent demodulation for eavesdropping on HDMI emissions

Dimitrije Erdeljan, Markus G. Kuhn

Department of Computer Science and Technology
University of Cambridge

A typical TEMPEST attack

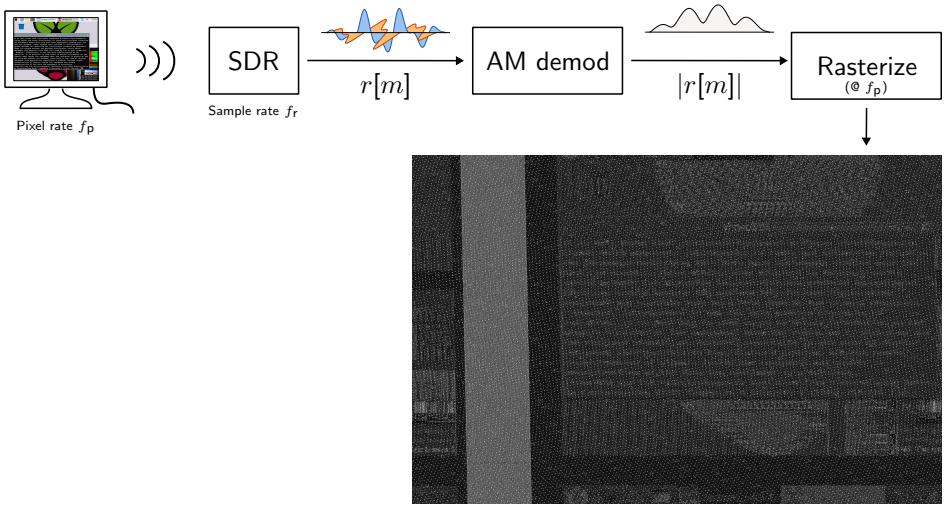

A typical TEMPEST attack

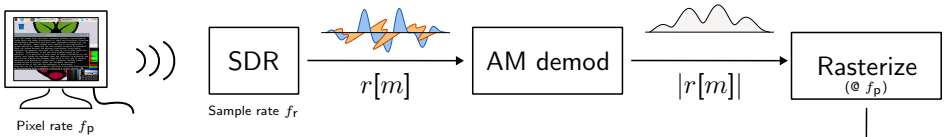

- Estimate pixel rate f_p, e.g. from the autocorrelation $R_{r,r}[d] = \sum_m r[m] \cdot r[m+d]^*$:

$$\hat{d} = \underset{f_{v,\perp} \le f_r/d \le f_{v,\top}}{\operatorname{argmax}} |R_{r,r}[d]|^2$$

$$f_p \approx f_r \cdot \frac{w_t h_t}{\hat{d}}$$

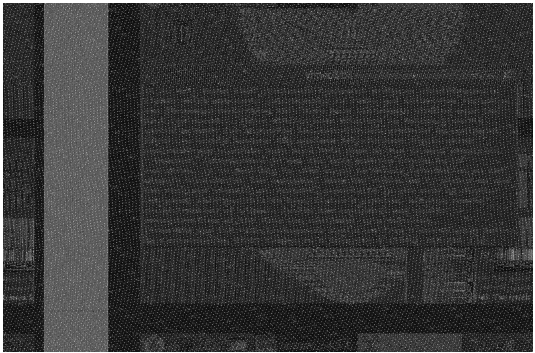

A typical TEMPEST attack

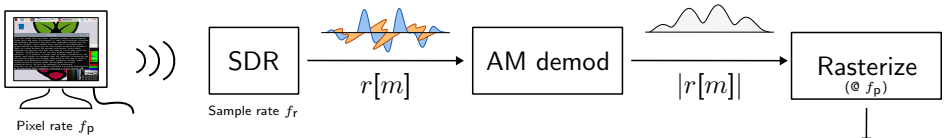

- Estimate pixel rate f_p, e.g. from the autocorrelation $R_{r,r}[d] = \sum_m r[m] \cdot r[m+d]^*$:

$$\hat{d} = \underset{f_{v,\perp} \le f_r/d \le f_{v,\top}}{\operatorname{argmax}} |R_{r,r}[d]|^2$$

$$f_p \approx f_r \cdot \frac{w_t h_t}{\hat{d}}$$

- Resample to $f_s = k \cdot f_p$ for $k \in \mathbb{N}$

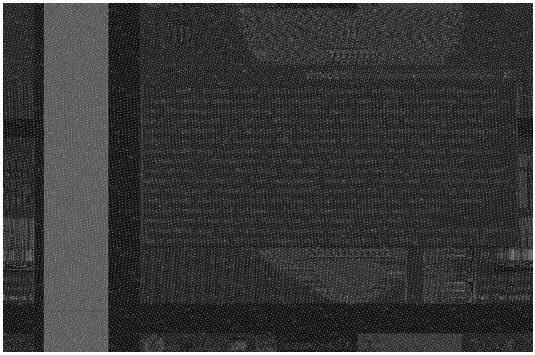

A typical TEMPEST attack

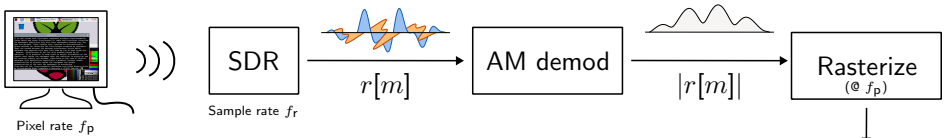

- Estimate pixel rate f_p, e.g. from the autocorrelation $R_{r,r}[d] = \sum_m r[m] \cdot r[m+d]^*$:

$$\hat{d} = \underset{f_{v,\perp} \le f_r/d \le f_{v,\top}}{\operatorname{argmax}} |R_{r,r}[d]|^2$$

$$f_p \approx f_r \cdot \frac{w_t h_t}{\hat{d}}$$

- Resample to $f_s = k \cdot f_p$ for $k \in \mathbb{N}$
- Average several rasterized frames

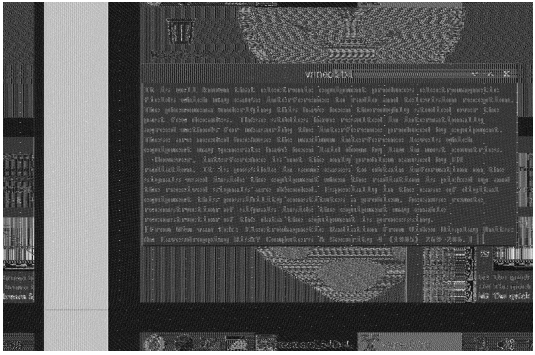

A typical TEMPEST attack

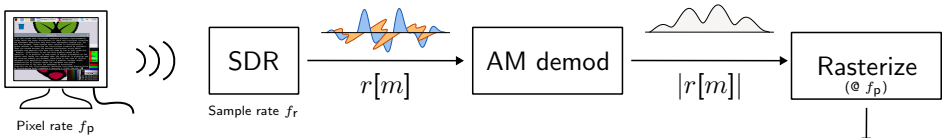

- Estimate pixel rate f_p, e.g. from the autocorrelation $R_{r,r}[d] = \sum_m r[m] \cdot r[m+d]^*$:

$$\hat{d} = \underset{f_{v,\perp} \le f_r/d \le f_{v,\top}}{\operatorname{argmax}} |R_{r,r}[d]|^2$$

$$f_p \approx f_r \cdot \frac{w_t h_t}{\hat{d}}$$

- Resample to $f_s = k \cdot f_p$ for $k \in \mathbb{N}$
- Average several rasterized frames
- Align image

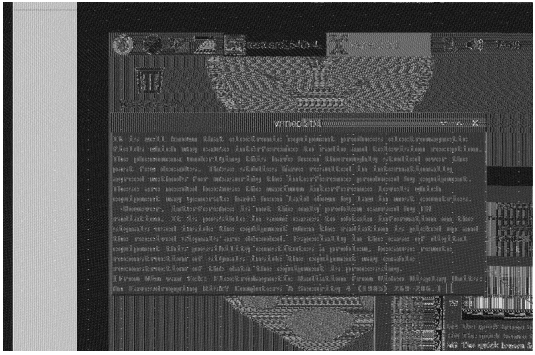

Software-defined radio receiver

Antenna waveform (shown as Fourier spectrum) $s_0(t)$:

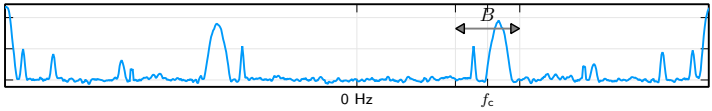

Downconvert: $s_\mathsf{d}(t) = s_0(t) \cdot \mathrm{e}^{-2\pi \mathrm{j} f_c t}$

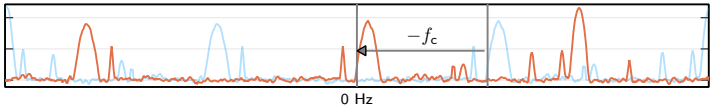

Lowpass filter: $s_\mathsf{f}(t) = \int s_\mathsf{d}(t-\tau) g(\tau) \mathrm{d}\tau$

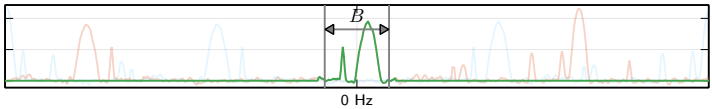

Finally, output sampled $r[m] = s_\mathsf{f}(m/f_\mathsf{r})$.

Rasterizing complex-valued signals: amplitude demodulation

Most eavesdropping demonstrations amplitude demodulate samples $M_{i,j} \in \mathbb{C}$ and visualise them as grayscale pixels.

For example, mapping 1% and 99% quintiles to black and white:

$$\text{Gray}\left(\frac{|M_{i,j}| - q_{1\%}}{q_{99\%} - q_{1\%}}\right)$$

This discards phase information $\angle M_{i,j}$.

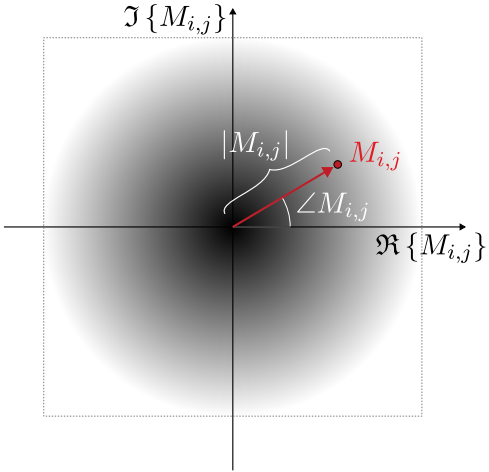

Rasterizing complex-valued signals: HSV visualisation

Using the HSV (hue, saturation, value) colour space allows us to also show the phase:

$$\text{HSV}\left(\angle M_{i,j},\ S,\ \frac{|M_{i,j}| - q_{1\%}}{q_{99\%} - q_{1\%}}\right)$$

(We leave the saturation coordinate S as a user preference.)

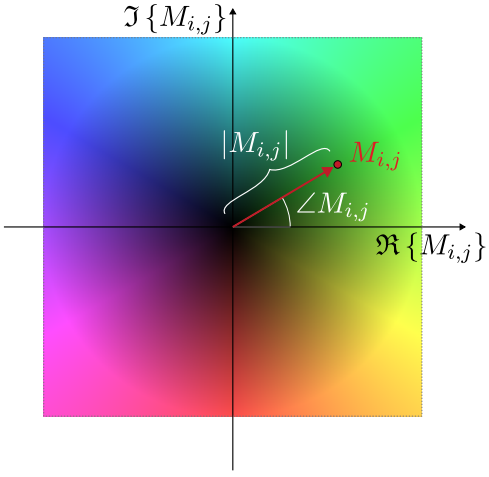

First rasterization attempt

Directly rasterizing an SDR-received signal produces a "rainbow-banding" image:

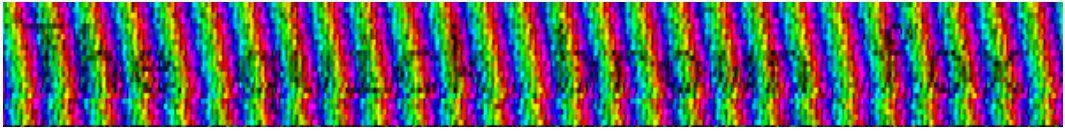

This is due to SDR downconversion from the antenna waveform $s_0(t)$ to $\mathrm{e}^{-2\pi \mathrm{j} f_c t} \cdot s_0(t)$.

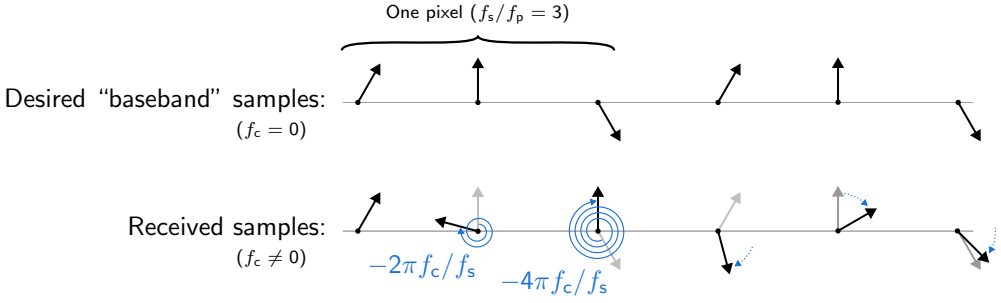

Obtaining consistent phase angles

Shift the centre frequency to a harmonic $h \cdot f_\mathsf{p}$ of the pixel frequency:

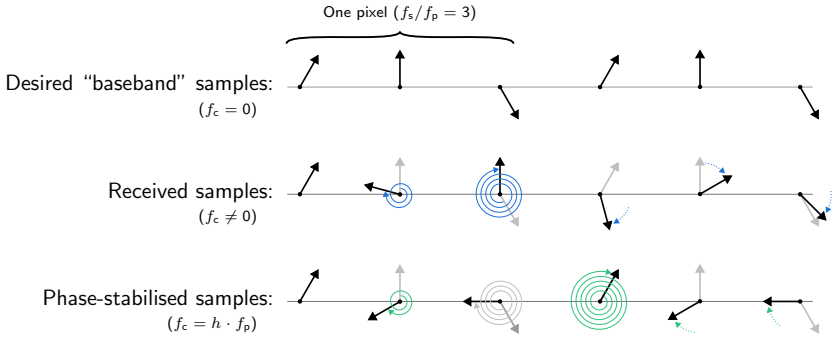

We combine frequency shifting $f_\mathsf{c} \to h \cdot f_\mathsf{p}$ with resampling to $f_\mathsf{s} = k \cdot f_\mathsf{p}$:

$$s[n] \approx s_\mathsf{f}\left(\frac{n+\lambda}{f_\mathsf{s}}\right) \cdot \mathrm{e}^{2\pi \mathrm{j}(f_\mathsf{c} - h f_\mathsf{p})n/f_\mathsf{s}}$$

Obtaining consistent phase angles

Some drift still remains over longer intervals.

Coherent averaging requires consistent phase across many frames, i.e. a more accurate f_p estimate.

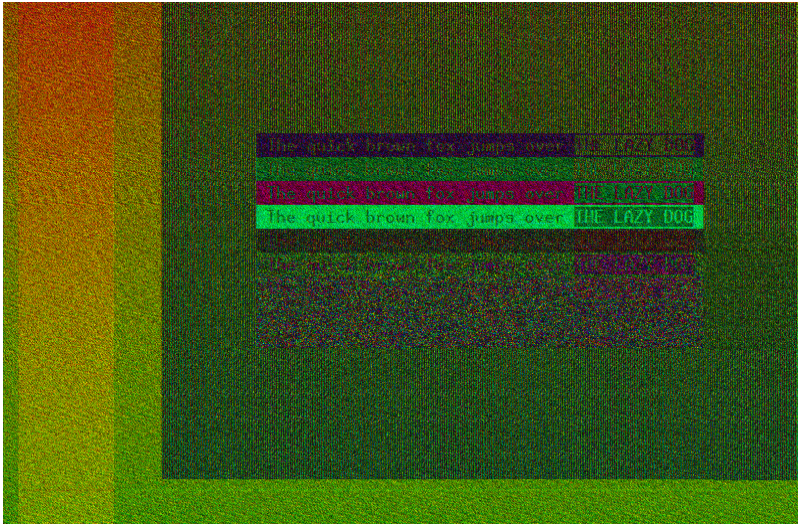

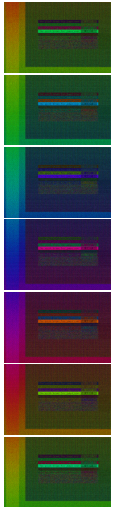

Accurate f_p estimation

We improve the f_p estimate several times until convergence, by iterating over three steps:

1. Resampling and frequency-shifting $f_\mathsf{c} \to h \cdot f_\mathsf{p}$:

$$s[n] \approx s_\mathsf{f}\left(\frac{n+\lambda}{f_\mathsf{s}}\right) \cdot \mathrm{e}^{2\pi \mathrm{j}(f_\mathsf{c}-hf_\mathsf{p})n/f_\mathsf{s}}$$

2. Computing the autocorrelation:

$$R_{s,s}[d] = \sum_n s[n] \cdot s[n+d]^*$$

3. Updating the f_p estimate, with a fine-tuning term which measures phase drift between frames:

$$f_\mathsf{p} := f_\mathsf{p} \cdot \left(\frac{kw_\mathsf{t}h_\mathsf{t}}{\hat{d}} + \frac{k\angle R_{s,s}[\hat{d}]}{2\pi h\hat{d}}\right)$$

(This is a corrected version of Eqn. 11 in the proceedings)

Iterations	f_p
0	25.200000000 MHz
1	25.200096064 MHz
2	25.200096764 MHz
3	25.200096794 MHz
4	25.200096793 MHz
5	25.200096788 MHz
6	25.200096788 MHz

$$\hat{d} = \underset{f_{\mathsf{v},\perp} \le f_\mathsf{r}/d \le f_{\mathsf{v},\top}}{\operatorname{argmax}} |R_{s,s}[d]|^2$$

In later iterations, we can also search for the correlation peak at larger multiples of the frame period.

Obtaining consistent phase angles (with accurate f_p)

With the more accurate f_p, the phase now stays consistent across frames.

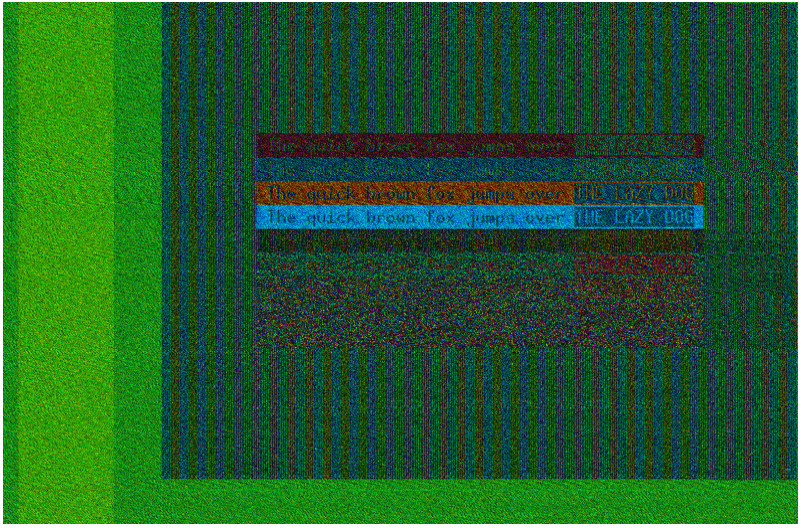

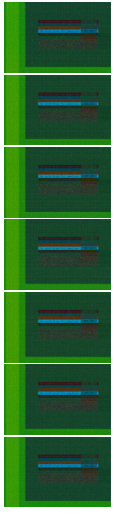

Consistent phase angles enable coherent averaging

We can now periodically average unrotated $M_{i,j} \in \mathbb{C}$ to reduce noise.

This image was rasterized from 30 averaged frames ($\approx$ 0.5 s long).

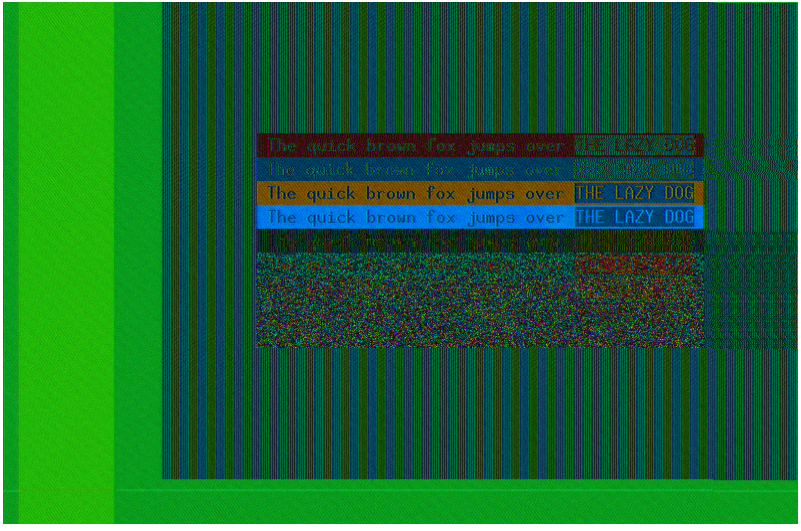

Does the phase provide new information?

To demonstrate how phase information can help with distinguishing colours, we used a test image with two grayscale colours: #101010 and #eeeeee.

These are TMDS-encoded in HDMI as complementary bit sequences:

- #101010 → 0111110000
- #eeeeee → 1000001111

The resulting emissions therefore differ only in their sign.

Displayed:

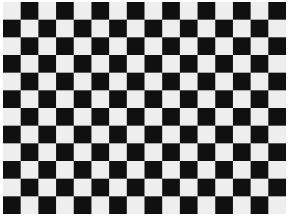

Coherently demodulated:

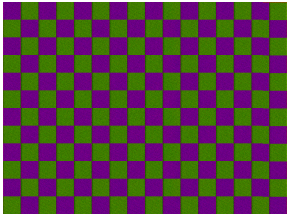

Amplitude demodulated:

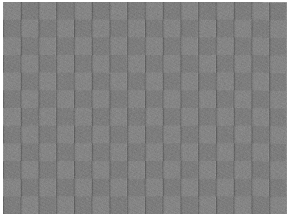

Does the phase provide new information?

The distributions of amplitudes $|M_{i,j}|$ for these two colours overlap significantly, while those for the full $M_{i,j} \in \mathbb{C}$ do not:

Amplitude distribution ($|M_{i,j}|$):

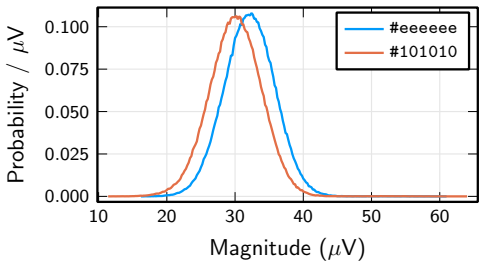

Complex distribution ($M_{i,j}$):

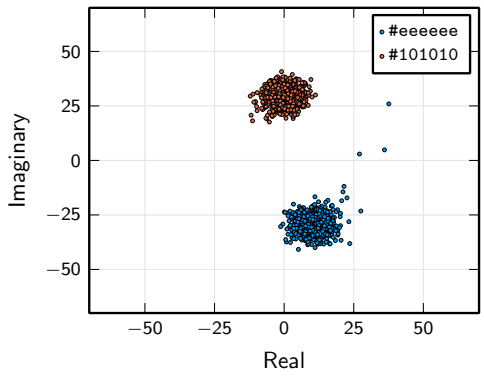

Comparing distributions using dimensionality reduction

We can better distinguish colours by combining information from all $k = 3$ samples for a pixel with dimensionality reduction, e.g. using Principal Component Analysis (PCA).

Amplitude demodulation (3-dimensional PCA on $|M_{i,3j}|, |M_{i,3j+1}|, |M_{i,3j+2}|$):

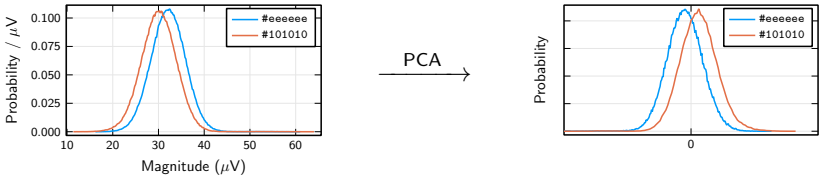

Coherent demodulation (6-dimensional PCA on $\Re\{M_{i,3j}\}, \Im\{M_{i,3j}\}, \ldots, \Im\{M_{i,3j+2}\}$):

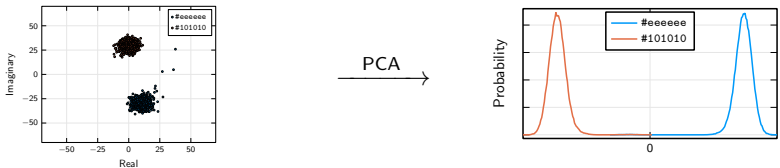

A more complex encoding

For another randomly chosen pair of colours, their more varied TMDS encodings result in different distributions for each sample position in a pixel:

Amplitude distribution ($|M_{i,j}|$):

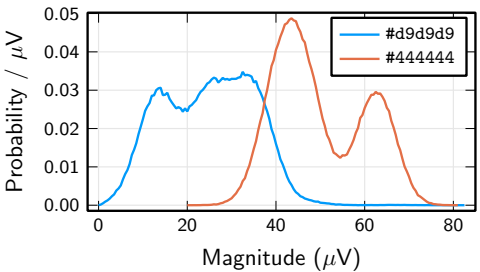

Complex distribution ($M_{i,j}$):

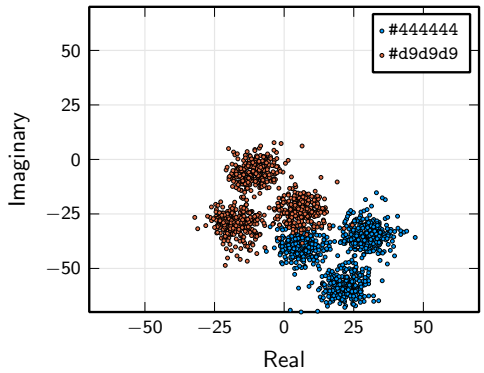

PCA-based rasterization

Rasterizing the projection on the largest-eigenvalue PCA vector can produce grayscale images with better contrast than amplitude demodulation.

1. Black and white:

2. Maximum bit transition contrast:

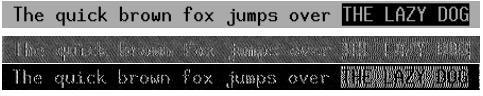

3. Complementary TMDS encoding:

4. Maximum contrast if AM demodulated:

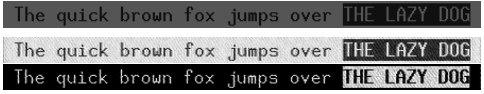

5. Minimum contrast if AM demodulated:

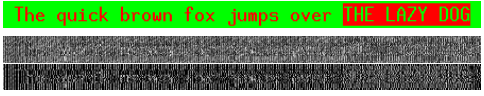

6. Low nibble random:

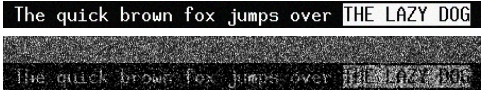

(Top to bottom: displayed, amplitude demodulated, PCA)

Summary

- Phase information can be included in the rasterized image as hue in the HSV colour space
- Shifting the centre frequency precisely to a harmonic of the pixel clock stabilises phase and allows periodic averaging in the complex domain
- We can precisely estimate the target's pixel-clock frequency using both position and phase of the peak of the complex-valued autocorrelation sequence
- Preserving phase information helps better discriminate between colours
- With multiple samples per pixel, dimensionality reduction can further improve contrast