Occultation Timing Accuracy: Dependence on Frame Rate and SNR

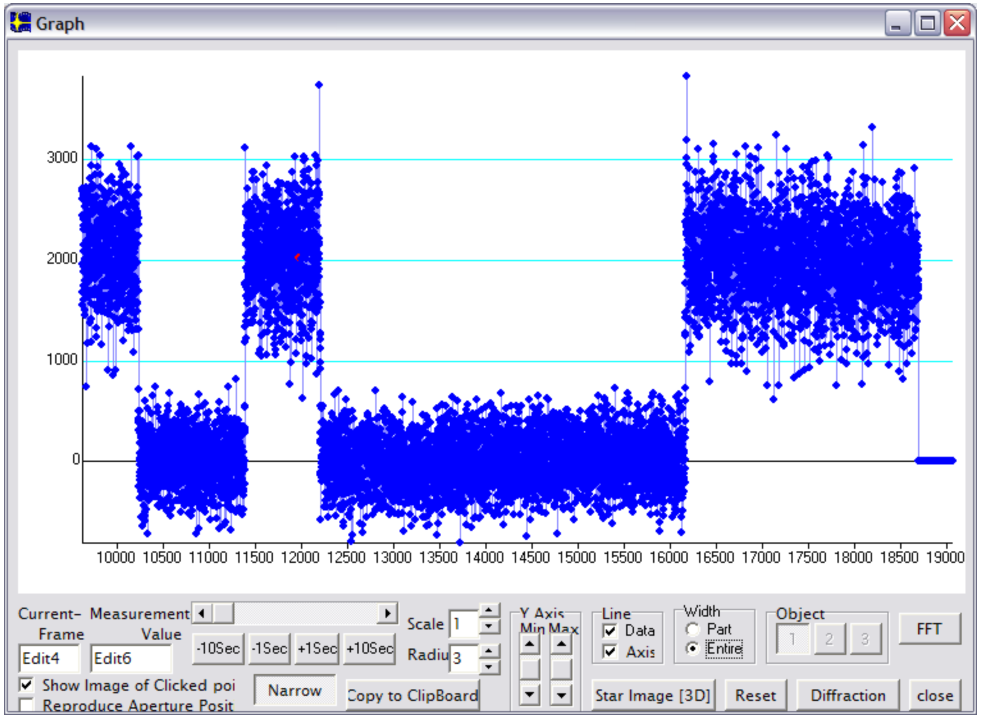

Frank Suits

Outline

- Nature of the problem: timing accuracy in occultation events
- Detection vs. accuracy
- Role of noise and frame rate
- Nature of the fitting and modeling process
- Simulation results
- A real example from a recorded occultation

Nature of the problem

- Occultation timing involves finding the time of an event based on measurements at regular intervals
- Need to know:
  - Best estimate of the event time
  - Accuracy of that estimate
- Key factors limiting accuracy are:
  - Frame rate
  - Noise

Not discussed: Detection

- This talk assumes an event is actually present in a light curve
- Separate problem: determine if the event actually happened
- The “detection” problem is very difficult, particularly with noisy curves and many possible event signatures
- This talk focuses on best estimates of the time under the assumption the event is real

Need to know

- In order to estimate the time of an event you need
  - Model for the noise in the measurement
  - Model for the event itself
  - Method for finding best fit to the data (many choices)
  - Free parameters used in fitting the model
  - Any other prior knowledge of the measurement or event that would change the statistical distribution of outcomes

Rich problem

- This is a general problem of optimal fitting of a model to a time-series, with “event time” as the primary goal
- The unique thing for astronomical events is the trade off between noise and frame rate
- As general as this problem is, I could not find good references specific to the dependency of timing accuracy on this trade off

Perfect occultation:
Reappearance of a star with no measurement noise

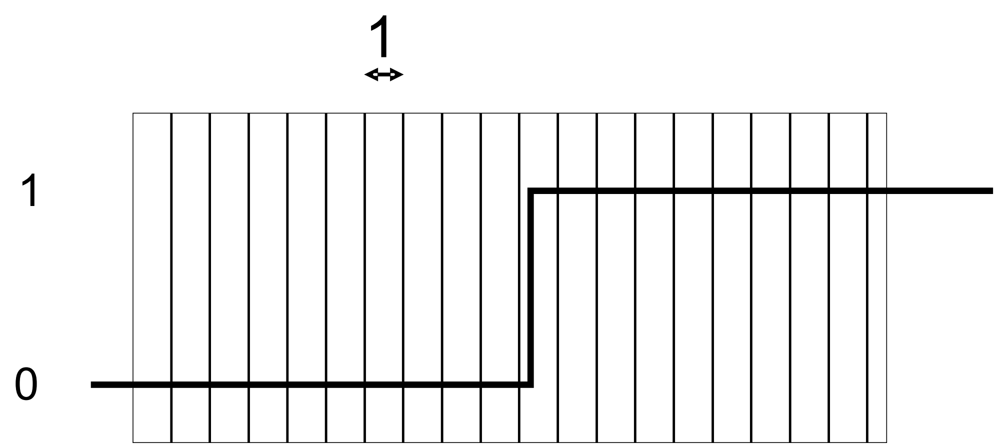

Canonical representation – signal goes from 0 to 1, with unit time steps

With no noise, exact time possible even with slow frame rate

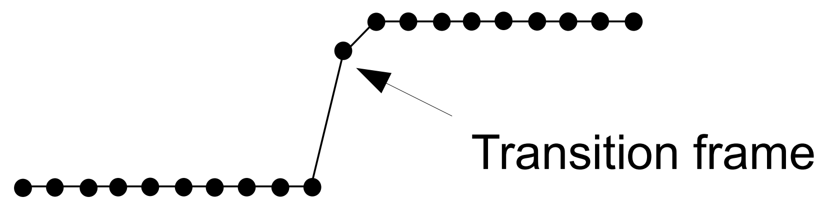

- Light curve shows single transition point
- If the underlying model is an instantaneous transition, the intensity of that one point allows interpolation of the time within that frame
- With no noise and a model with no free parameters, the exact time is known even with a very slow frame rate

- This is always under the assumptions of a model

Diffraction and more

- This talk focuses on the specific case of negligible diffraction effects
- Diffraction and other effects can always be included in the model
- All “prior” knowledge can be folded in as parameters with some expected range and statistical distribution of values

More realistic conditions

Presence of intensity-dependent noise makes event harder to interpret
Noise could have any distribution or intensity dependence
But often a mixture of Gaussian and Poisson

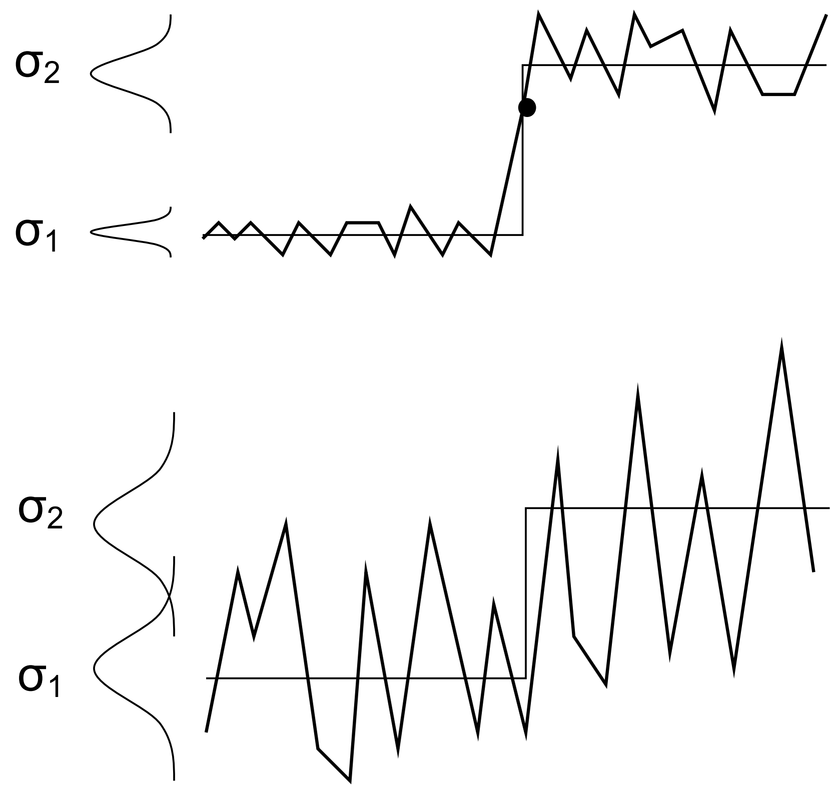

In canonical form, sigmas are scaled to 0->1 mean intensity transition

Example simulation with Gaussian noise

Green line shows a noisy light curve
Blue line shows best fit

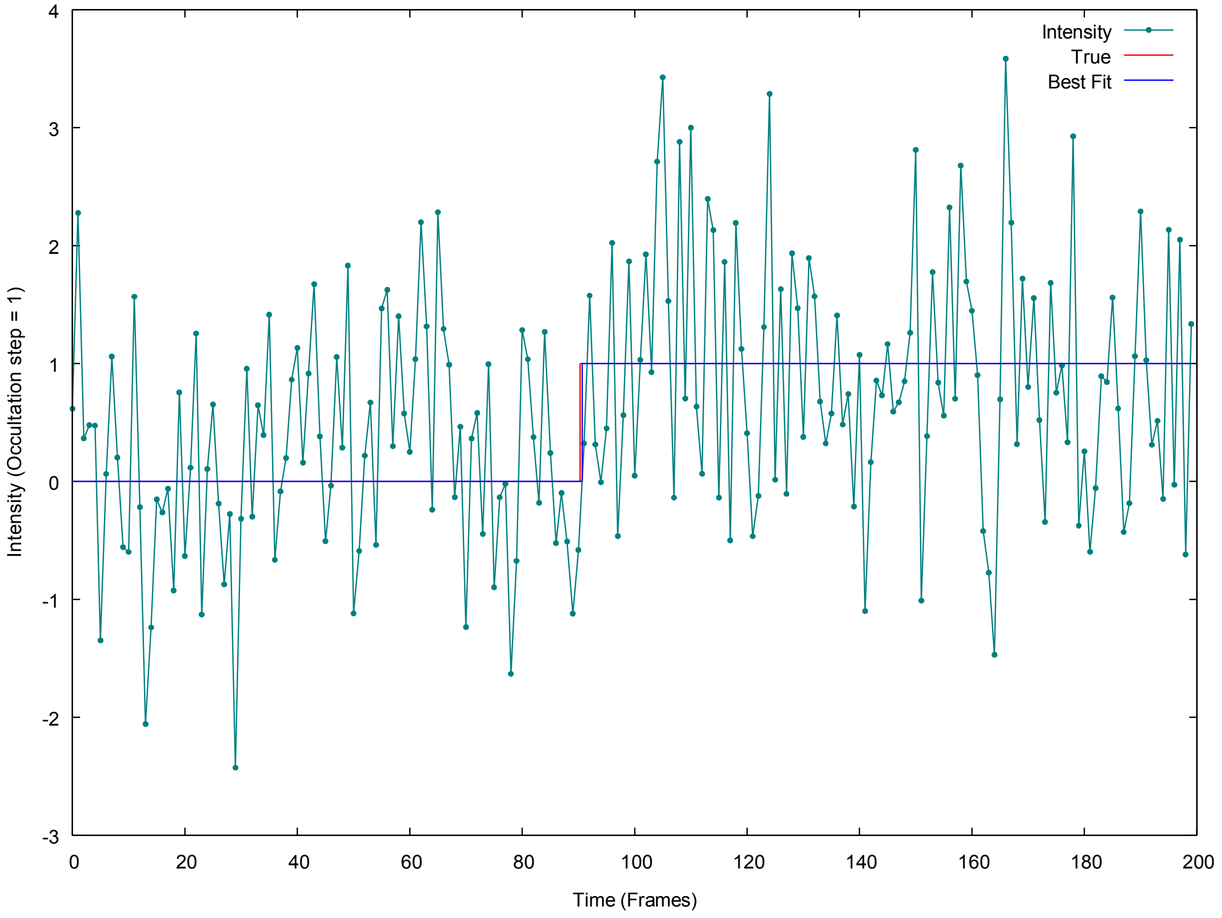

Close up of transition region

Red line shows the true transition (known since this is a simulation)
Difference between red and blue is the exact timing error for this result

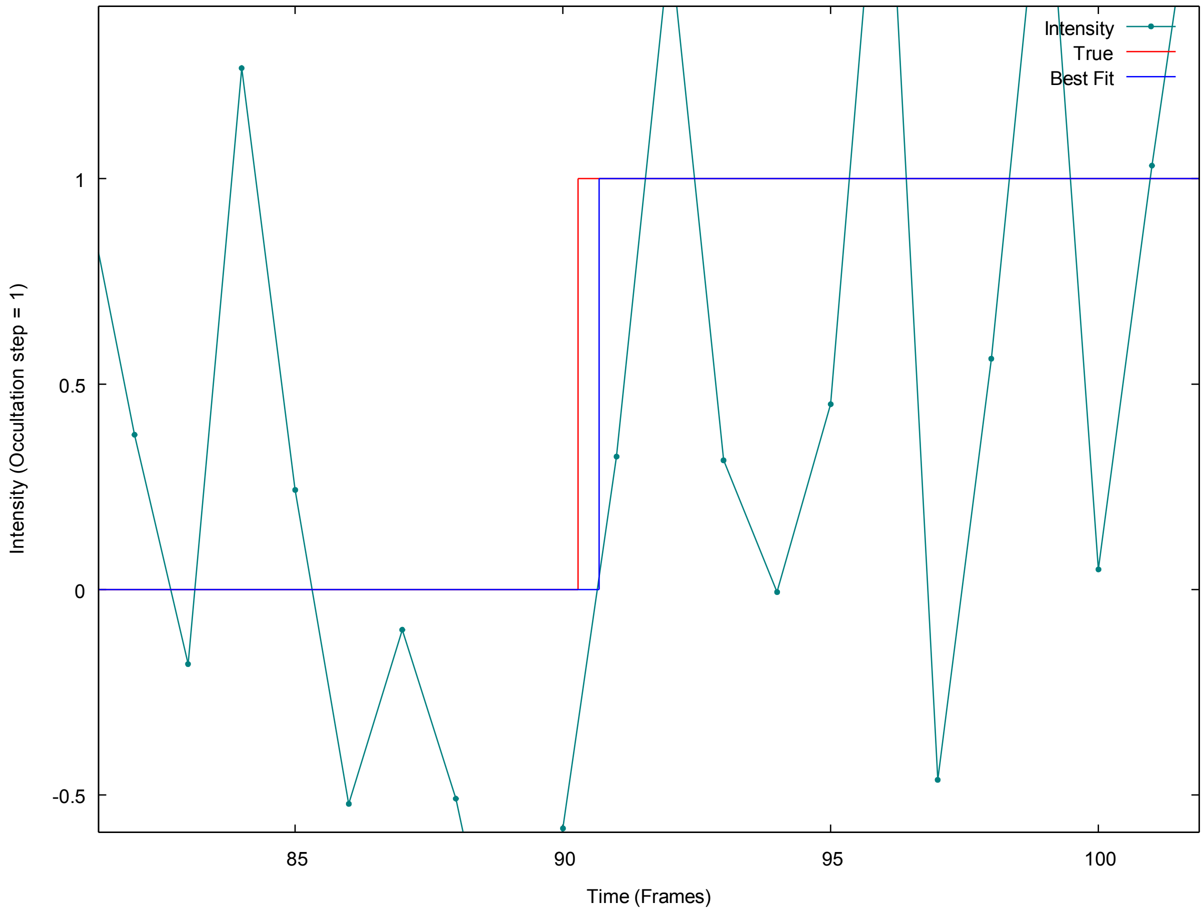

Different noise behavior with intensity

Pins the event away from the region with low noise

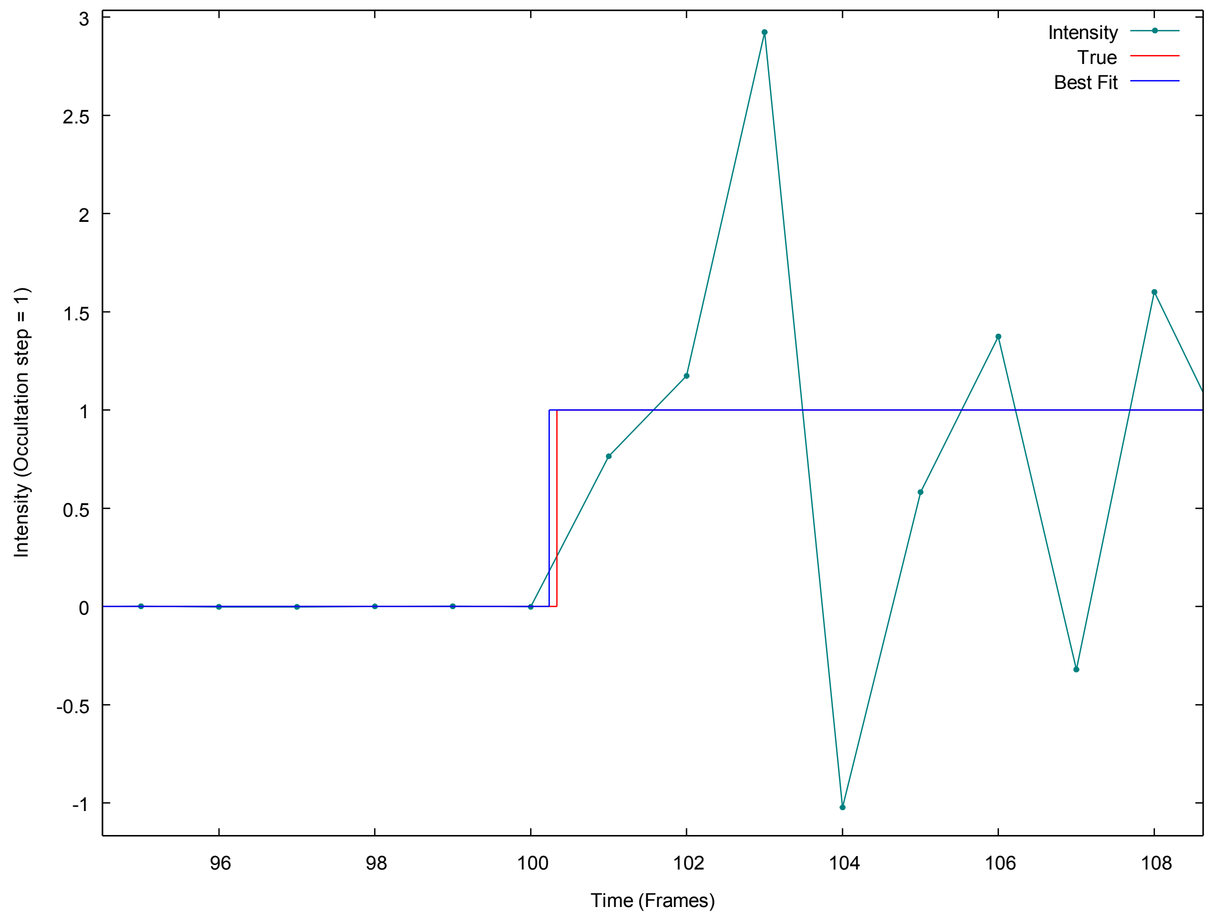

How to find “best” estimate?

- Obvious answer is non-weighted least squares
- Could also weight according to local estimate of noise
- But – since transition point carries much more information – could weight it more

- No clear “best” way to estimate the time
- Each method will have its own statistical distribution of resulting errors

Find error estimate by Monte Carlo

- Given a light curve and models for the noise and event
- Run many simulations to generate a similar light curve
- Apply whatever best-fit procedure you are using to each simulated curve
- You know the exact time for each simulation, so you get an error distribution

- Works for any black box fitting procedure you choose
- Could just use “middle of nearest frame” as method

Timing error as a function of noise

Black curve assumes equal noise levels for high and low intensities
Red curve assumes very little noise on low intensity side
Black is initially linear then roughly quadratic
Red shows threshold caused by “pinning” from low noise on the left

Error bars are from bootstrap on results

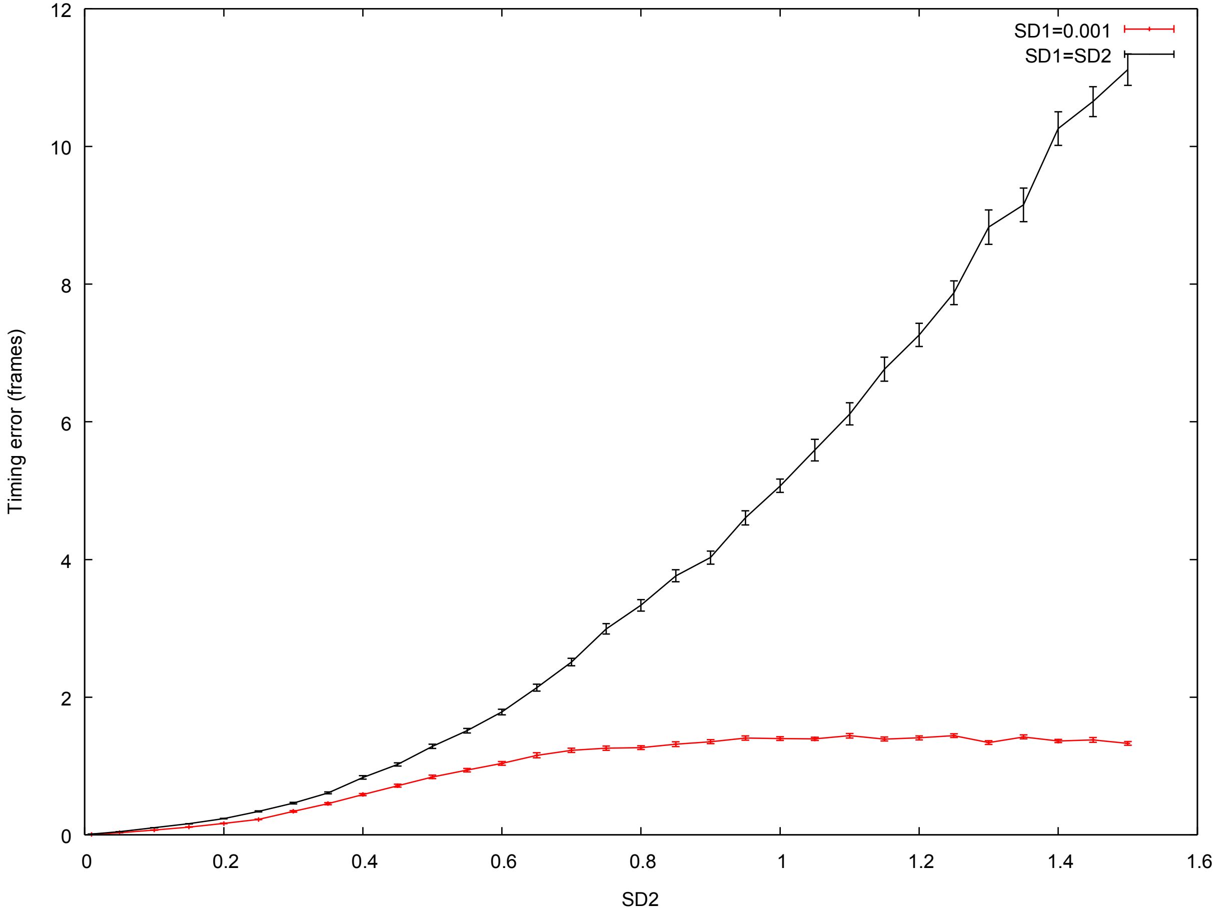

Actual asteroid occultation example

Fairly good SNR reveals clear transition frames

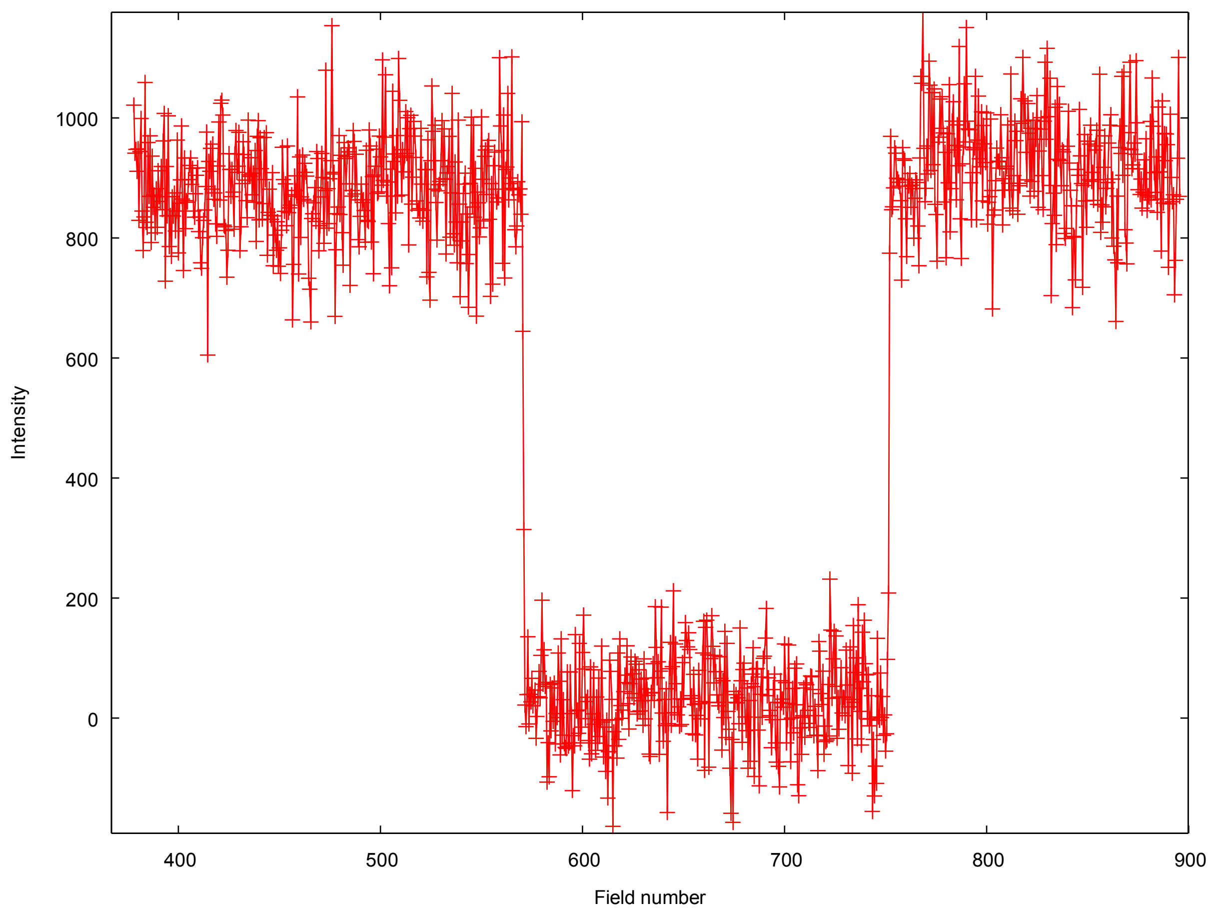

Close up showing transition frame

On a disappearance you can only have one transition frame
Presumably upper one is just noisy, while lower one is a true transition

Statistically it could be the other way around, though

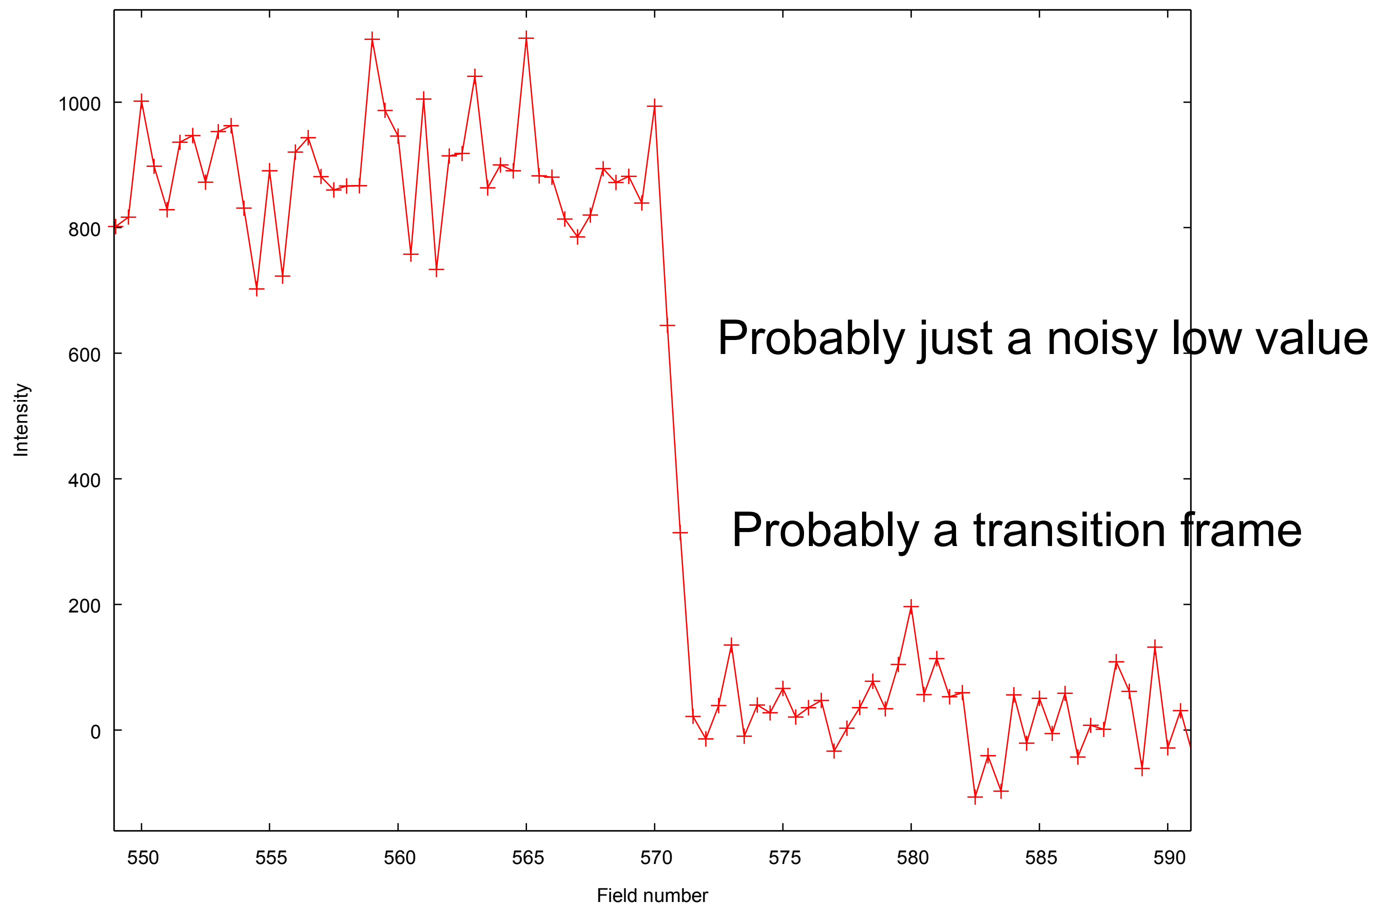

Reappearance

Transition frame is most likely at the beginning of the ascent

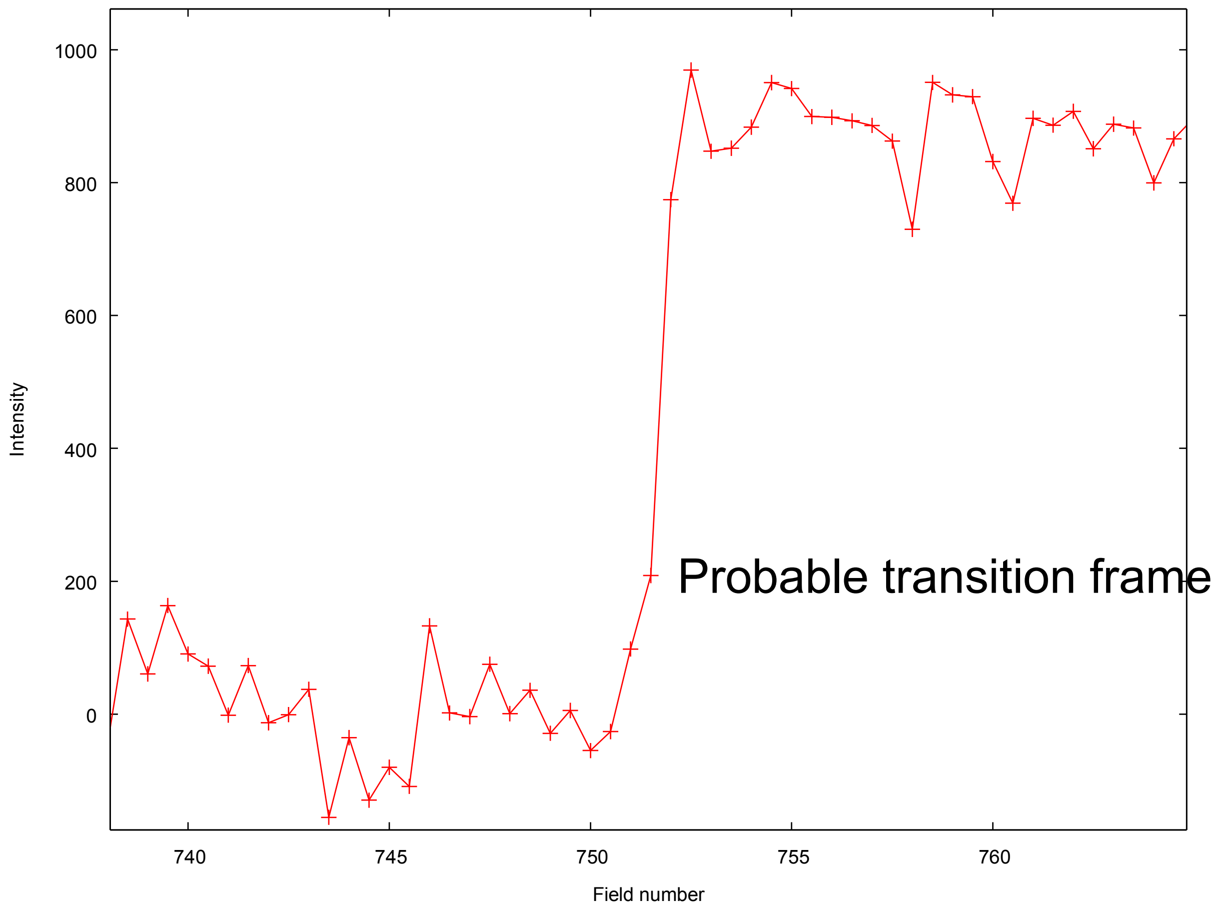

Actual error histogram from Monte Carlo

Although each frame is 17ms, the sigma of the resulting distribution is only 1.46ms

Note that the distribution has wide tails and is not purely Gaussian, so outliers are possible

As always – this result only makes sense under assumptions of the model and noise

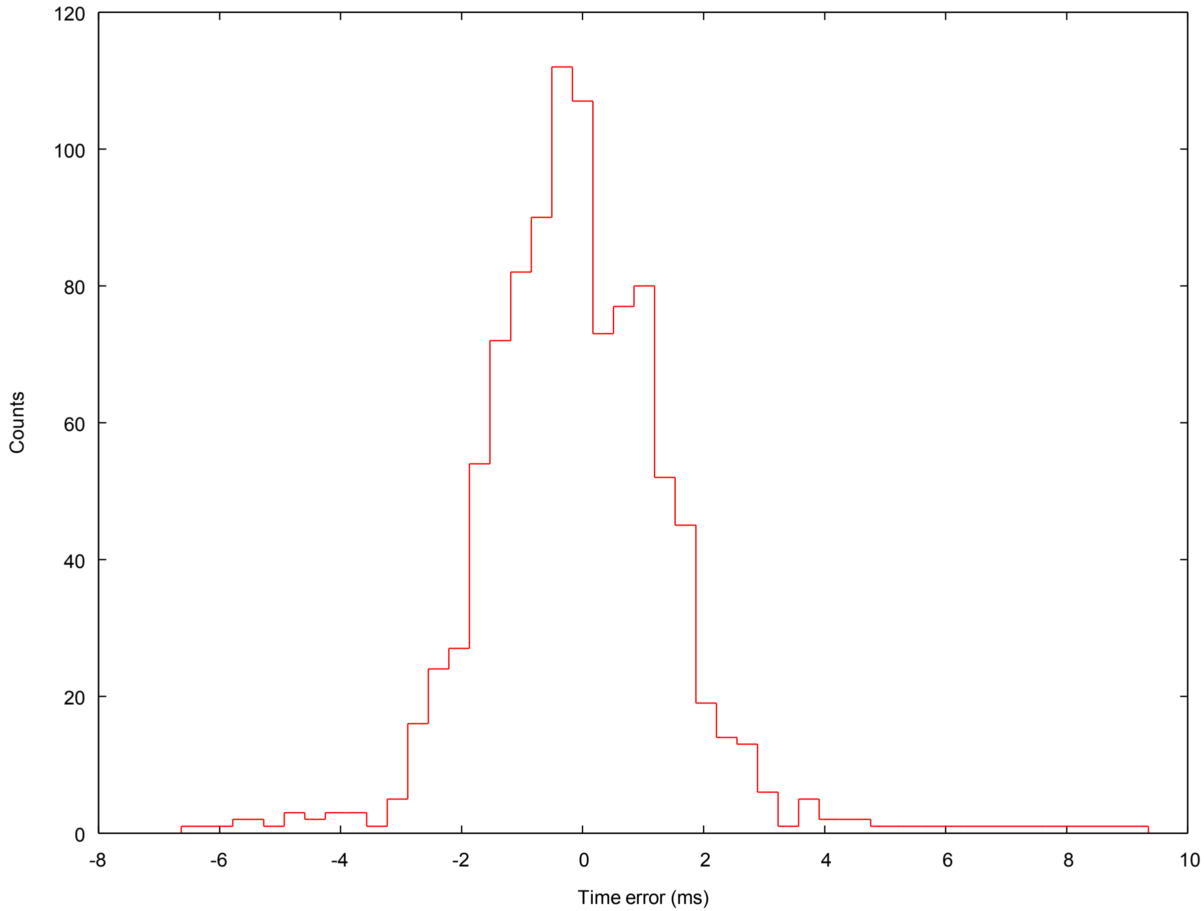

Conclusion

- There is an interesting trade off between noise and frame rate in occultation timing
- You need to model the noise and the phenomenon
- Under the assumptions, resulting accuracy can be much less than the frame time

- Future: Model the noise of a camera carefully to find optimal frame rate for a given event
- With digital video cameras and a given event model, you can dial in the optimal frame rate for the event