Looking Around the Corner using Ultrafast Transient Imaging

IJCV Supplementary

Ahmed Kirmani, Tyler Hutchison, Ramesh Raskar
Media Laboratory
Massachusetts Institute of Technology
Cambridge, MA 02139

James Davis
Computer Science
UC Santa Cruz
CA 95064

Can we infer about the hidden (green) card just using flash photography from this fixed viewpoint?

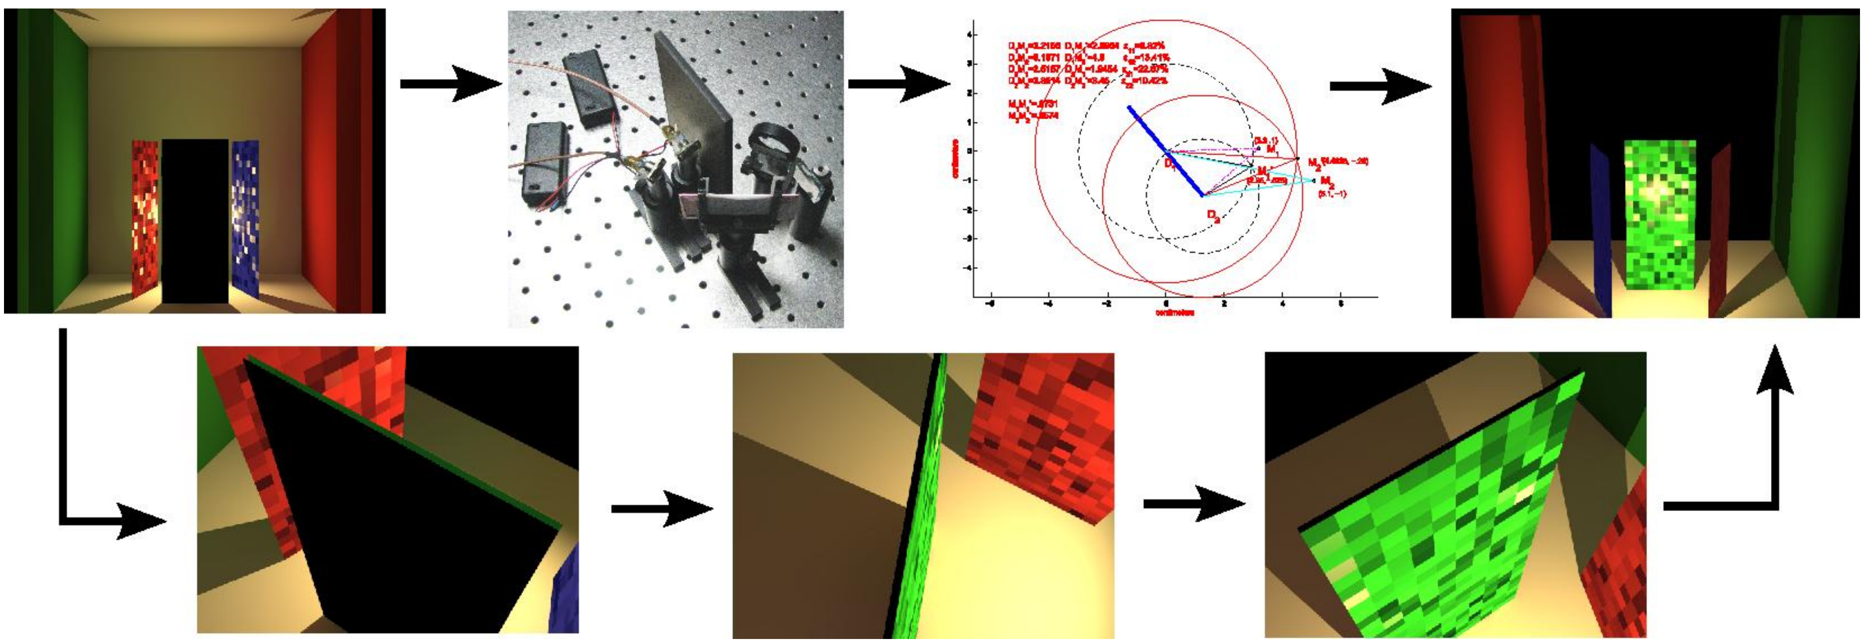

Transient Nature of Light

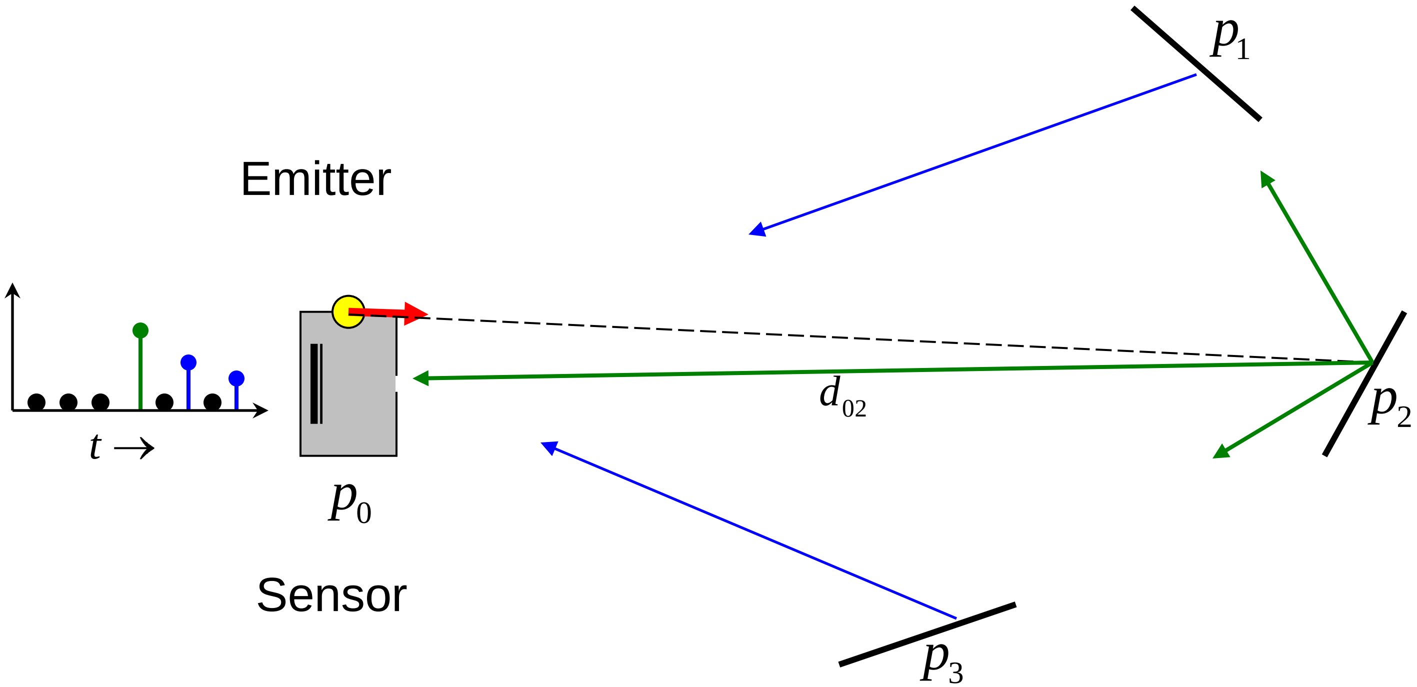

Even in a simple scene with 3 patches, a single light ray towards patch P_2 creates a complex time profile

Transient Light Transport

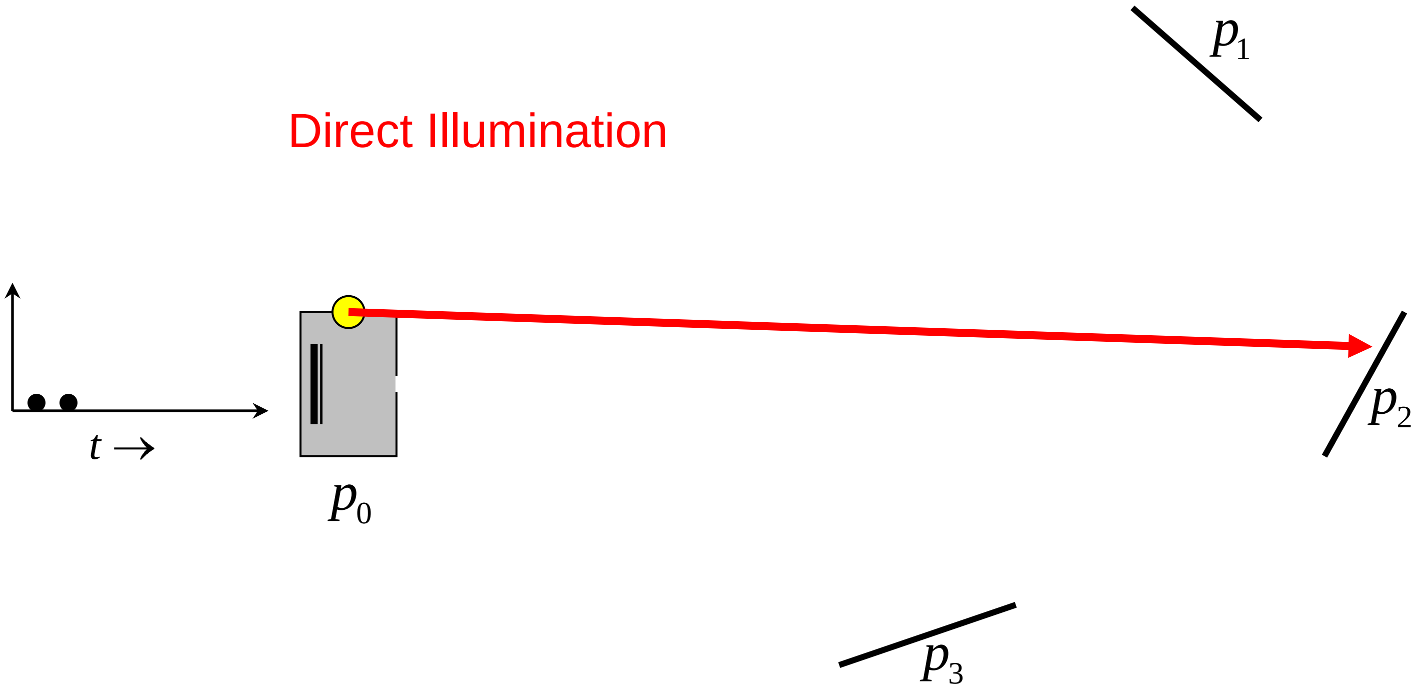

Consider the direct illumination of the second patch

Transient Light Transport

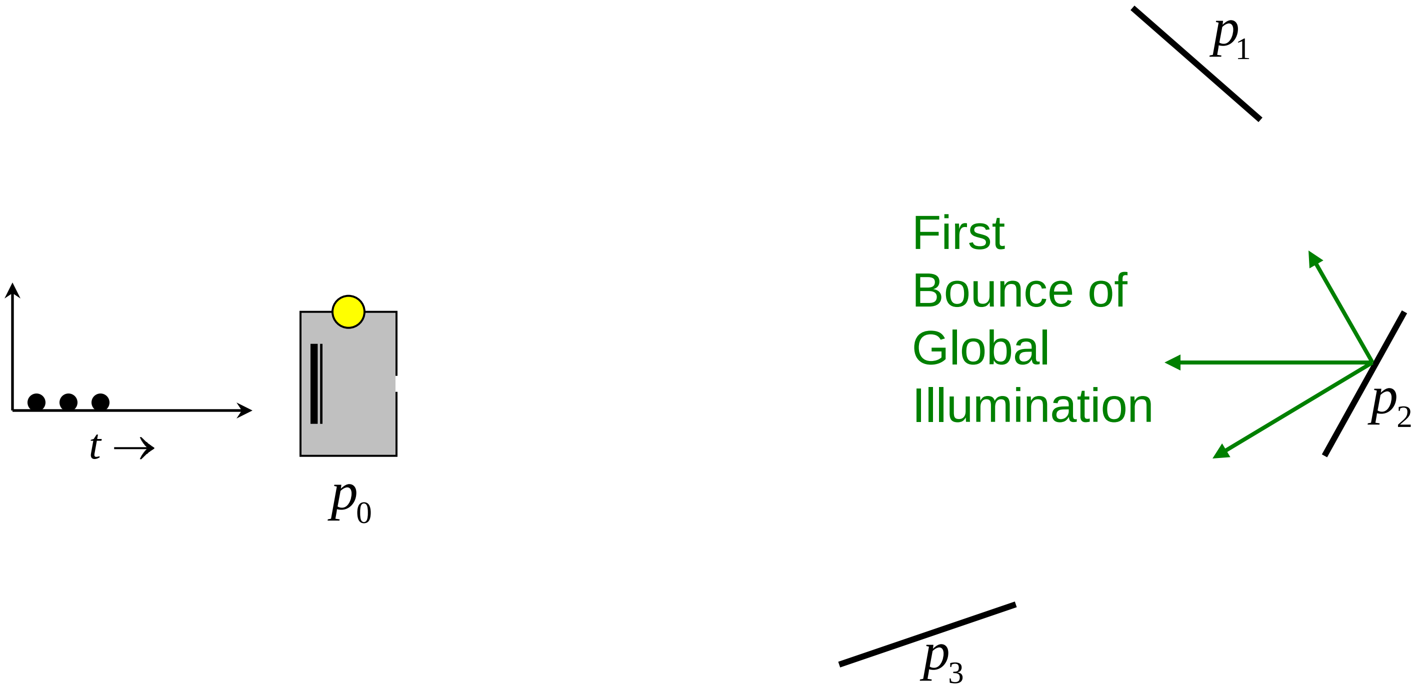

The radiance due to first bounce is distributed in multiple directions

Transient Light Transport

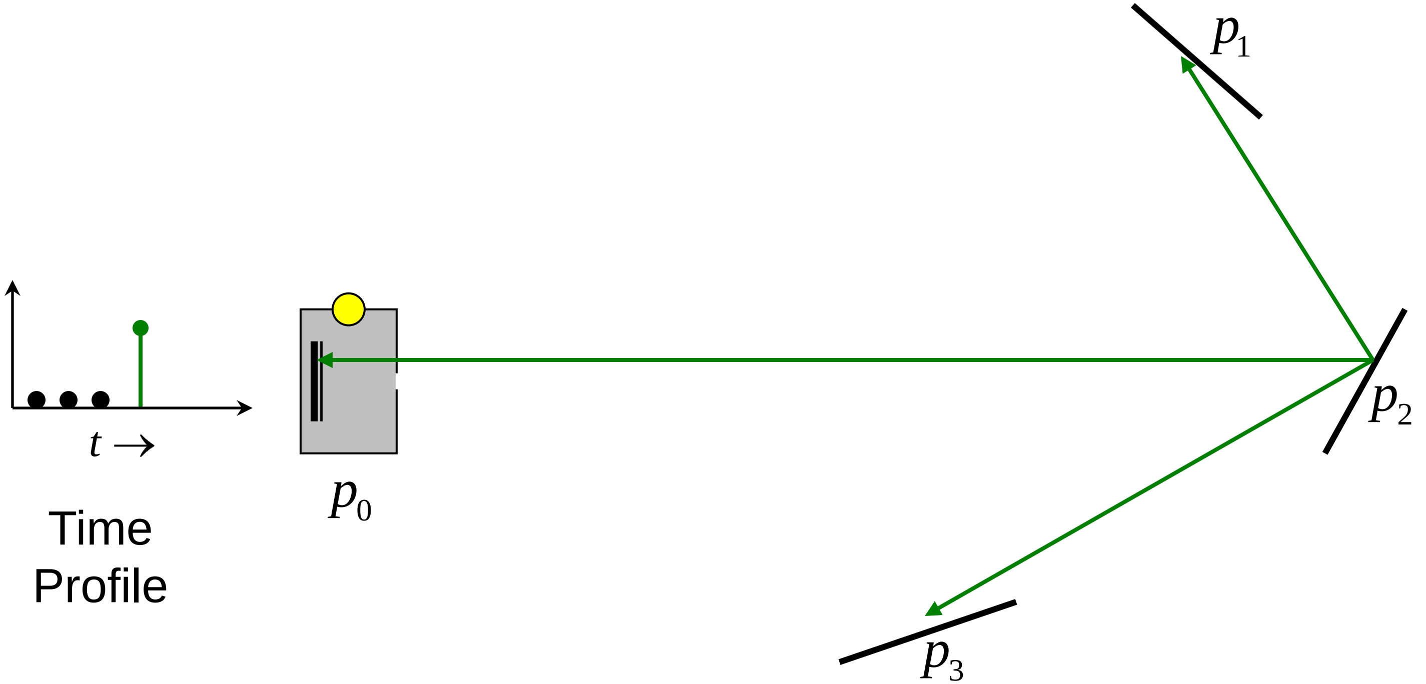

In the time profile, we receive the first bounce after a delay

Transient Light Transport

Second Bounce of Global Illumination

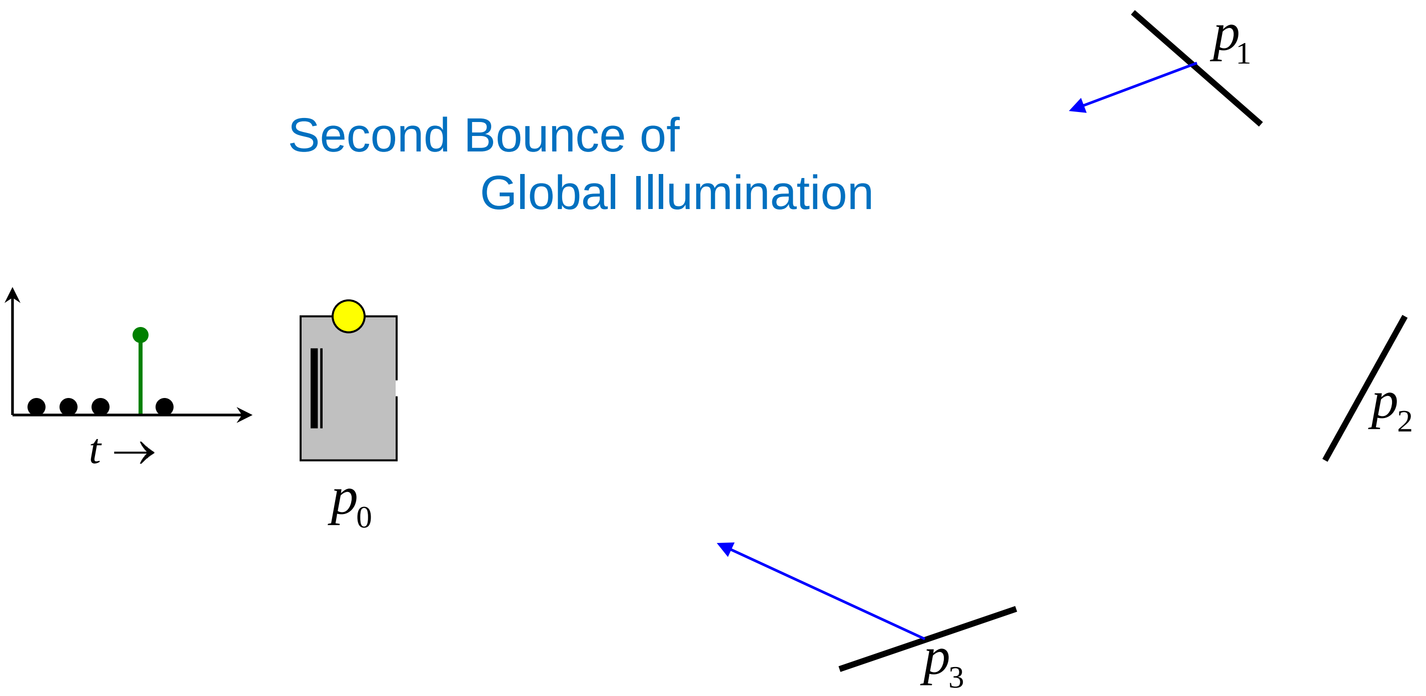

Meanwhile, light bounces a second time and part of the energy is reflected back to the sensor

Transient Light Transport

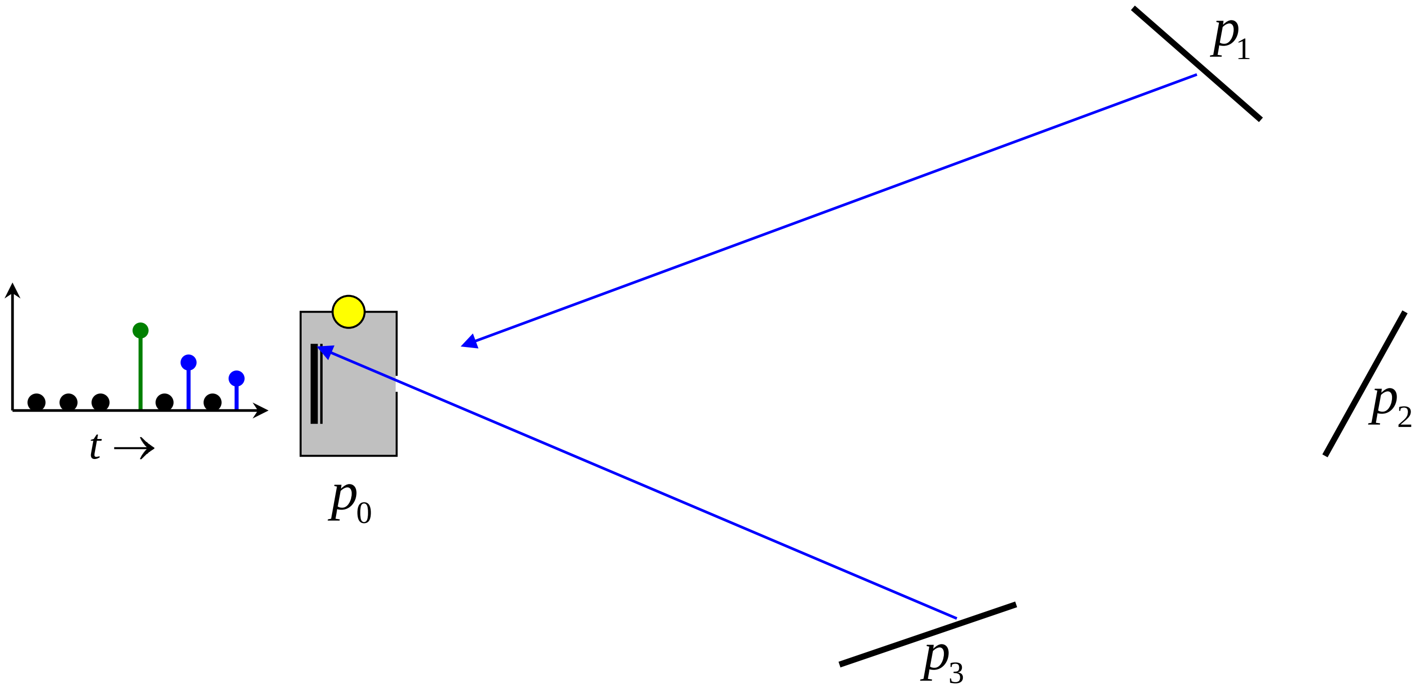

In time profile, we see these weaker second bounces

Transient Imaging Camera

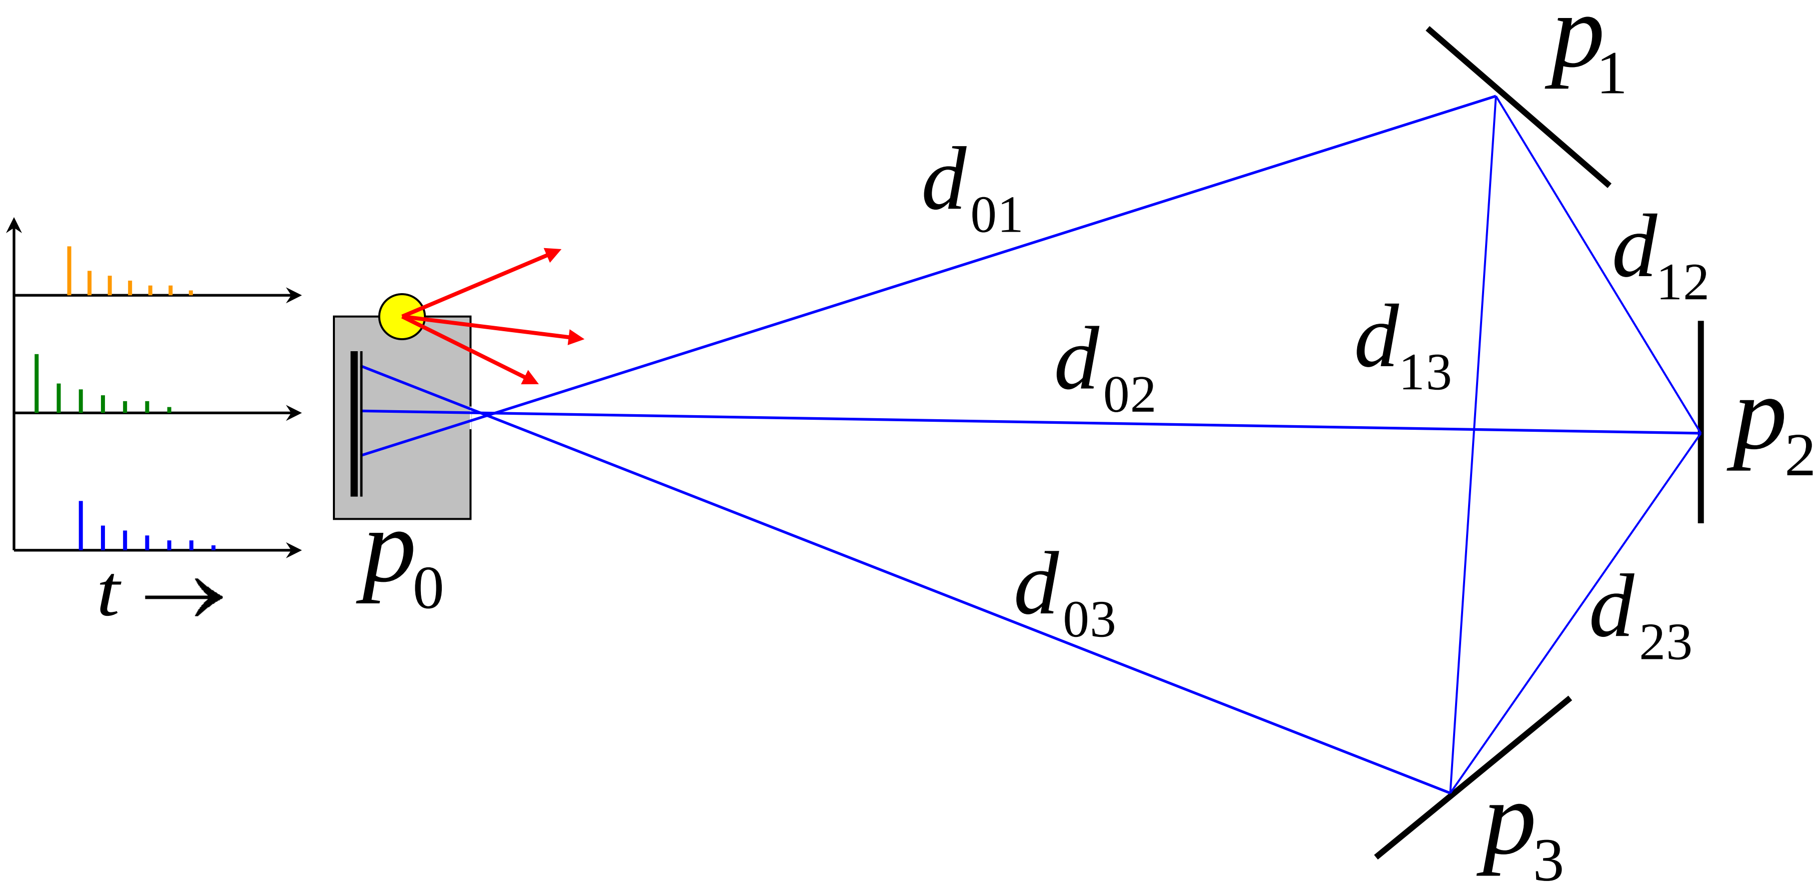

The transient imaging camera create a time-image, with a time profile at each pixel

Space Time Impulse Response

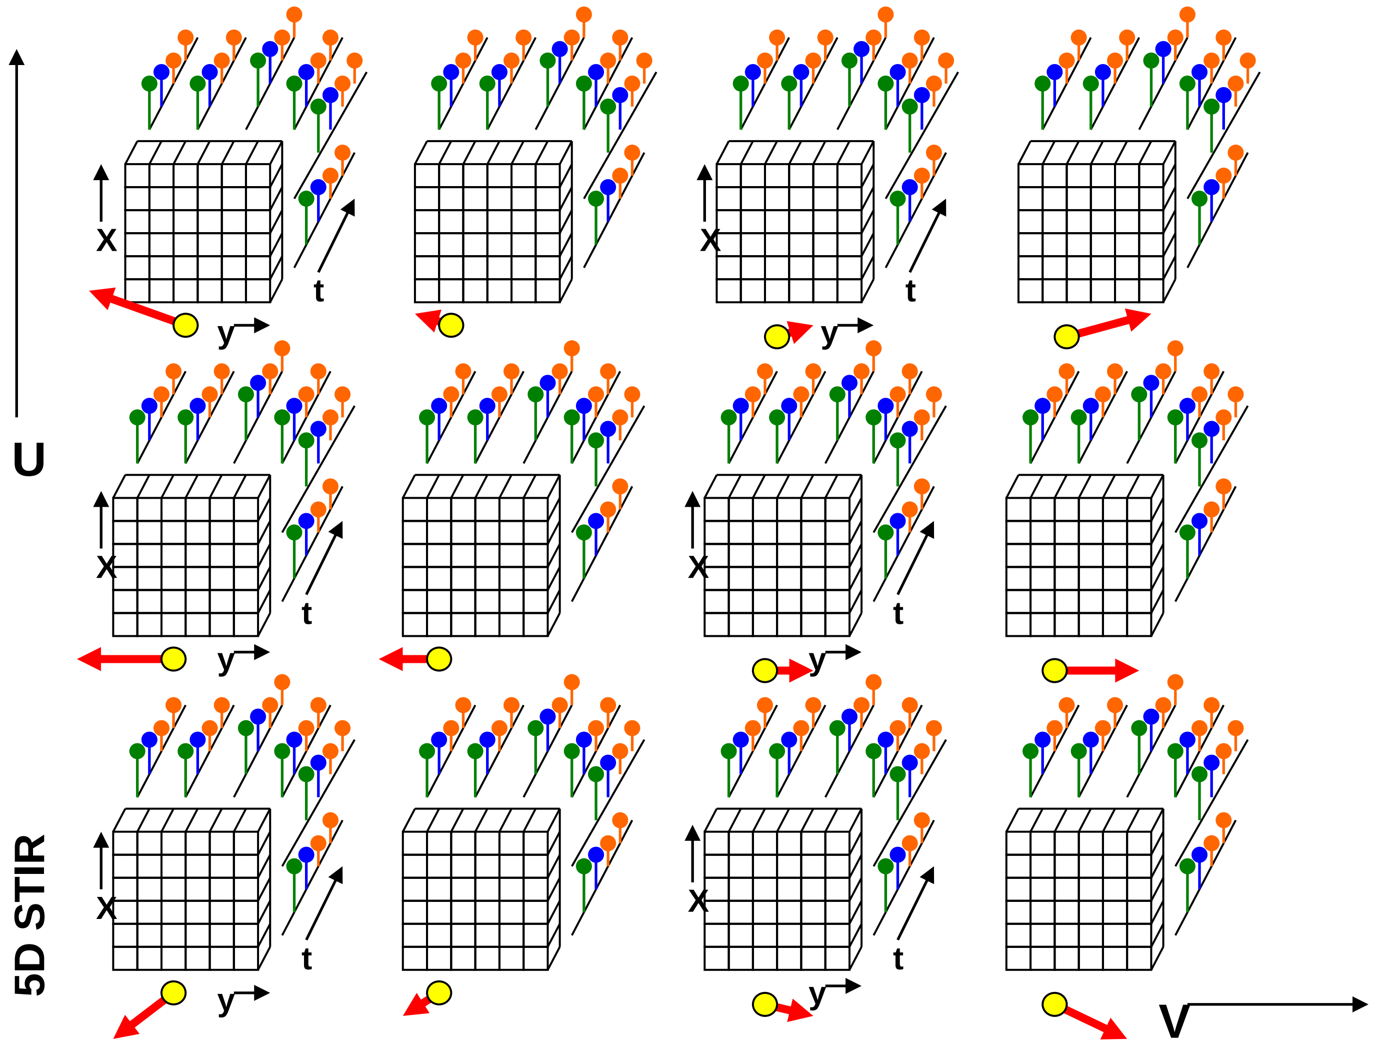

Space time impulse response (STIR) is a 5-dimensional function. It is a collection of time images recorded under varying ray impulse illumination

Recording Space Time Impulse Response

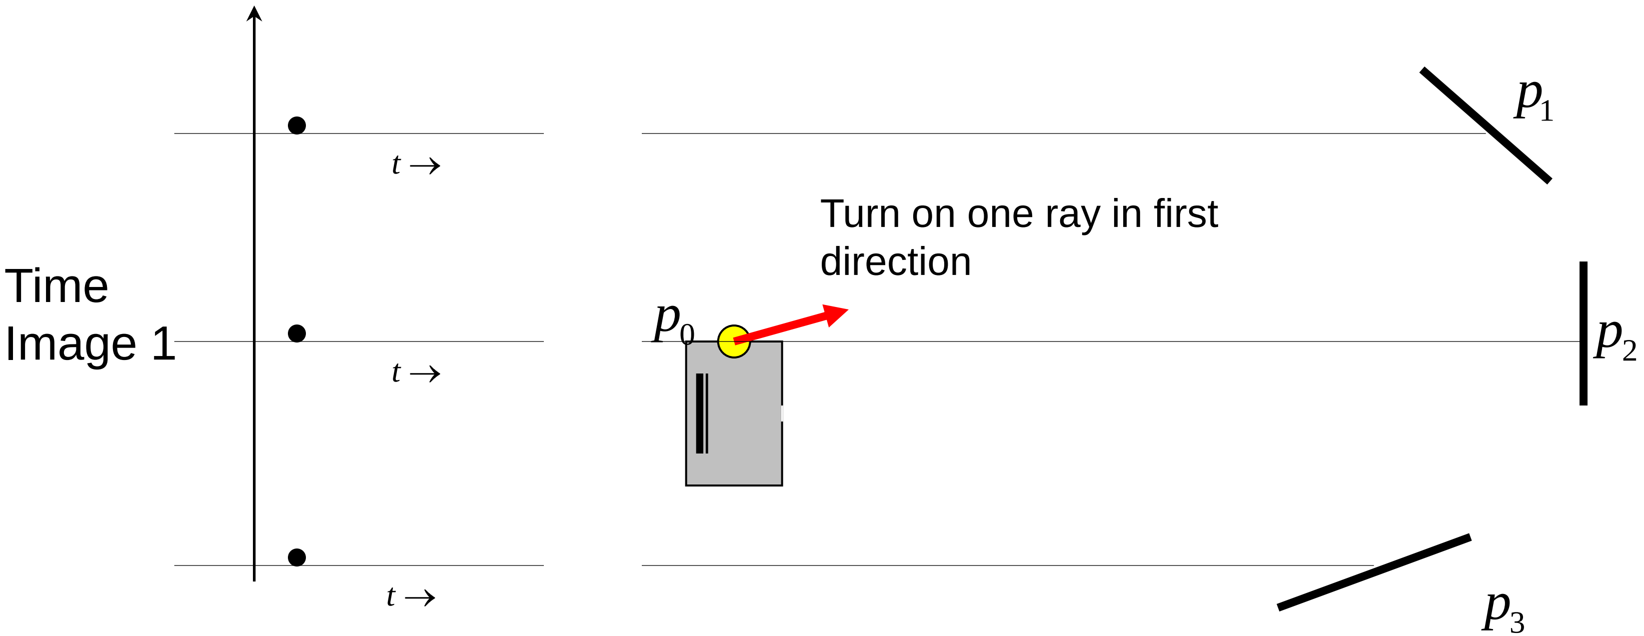

Space time impulse response (STIR) is obtained by imaging the scene one patch at a time and recording a time image. A time image is sequence of images captured at pico-second resolution

Recording Space Time Impulse Response

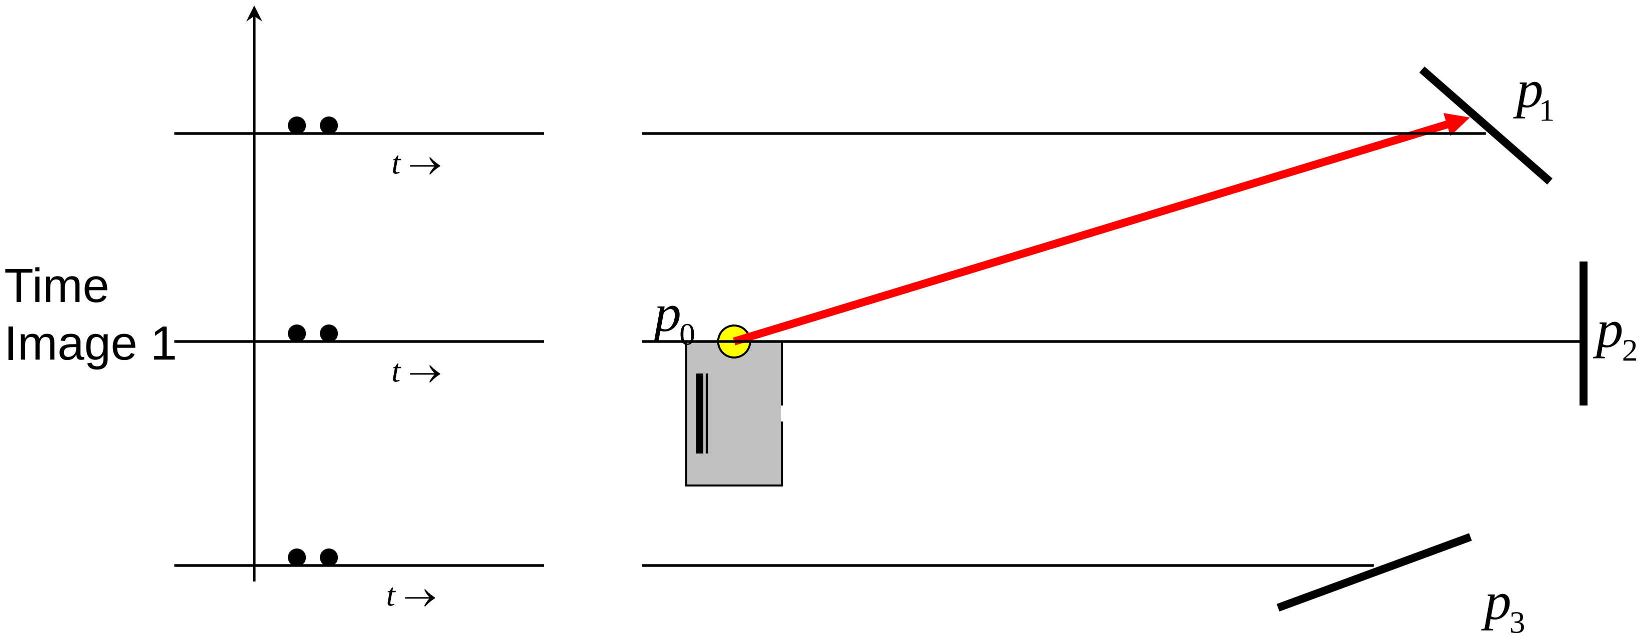

We illuminate one patch with a femto-second ultra-short duration laser

Recording Space Time Impulse Response

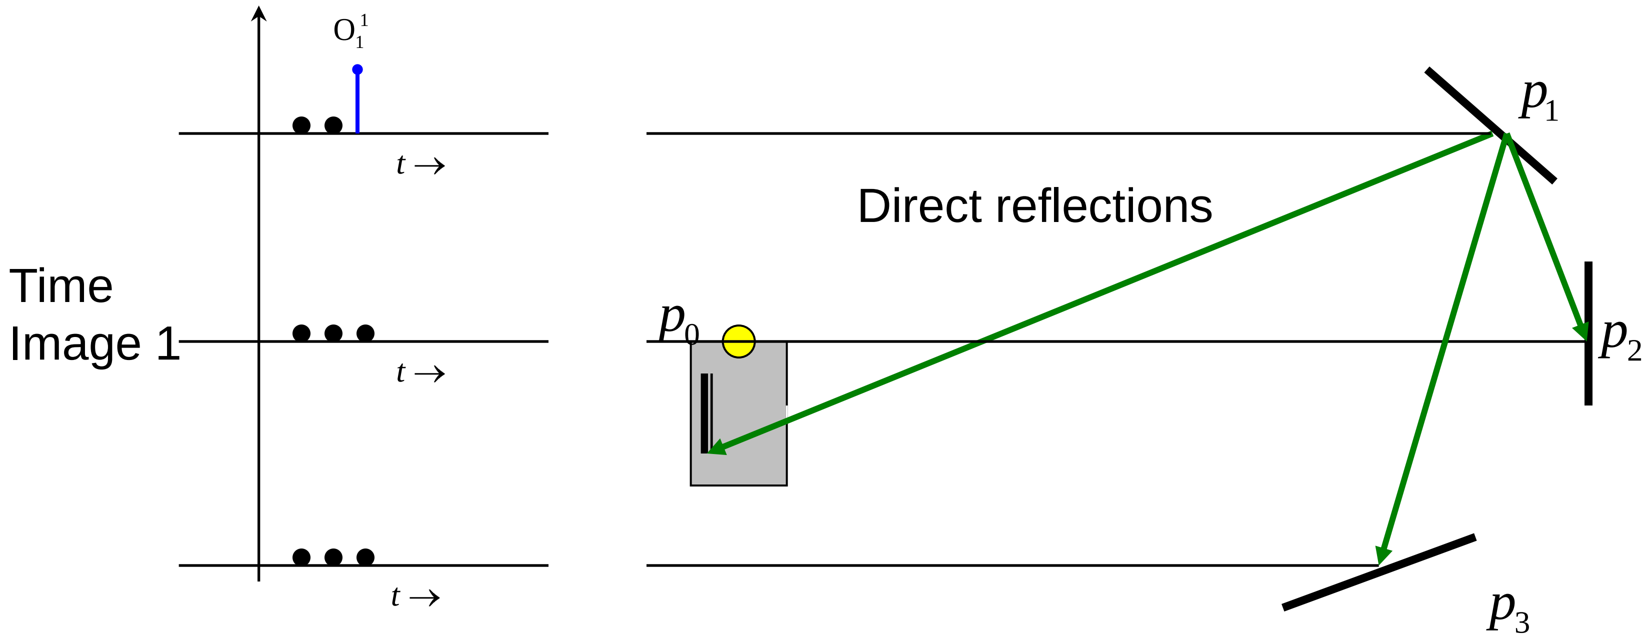

We record what's coming back towards the camera from all other patches

Recording Space Time Impulse Response

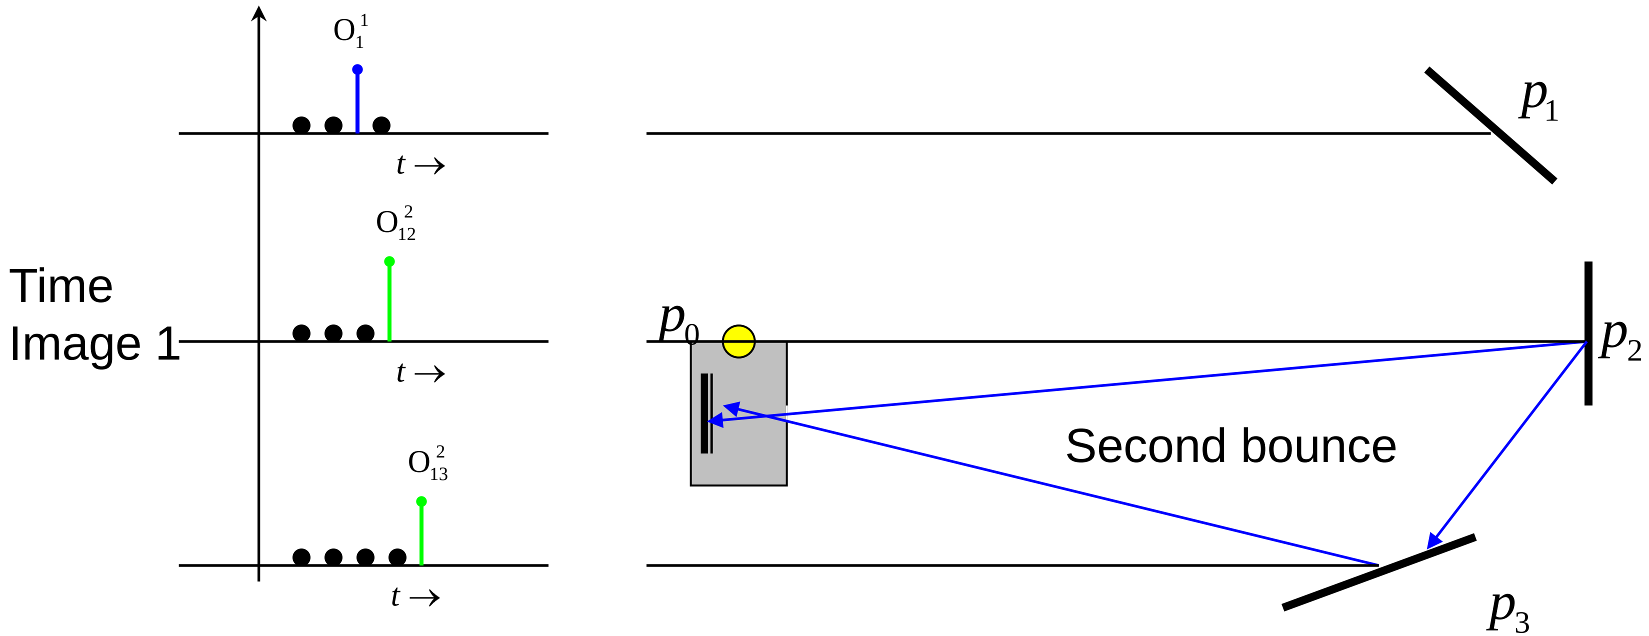

In addition we record second and third bounces shown in green and orange

Recording Space Time Impulse Response

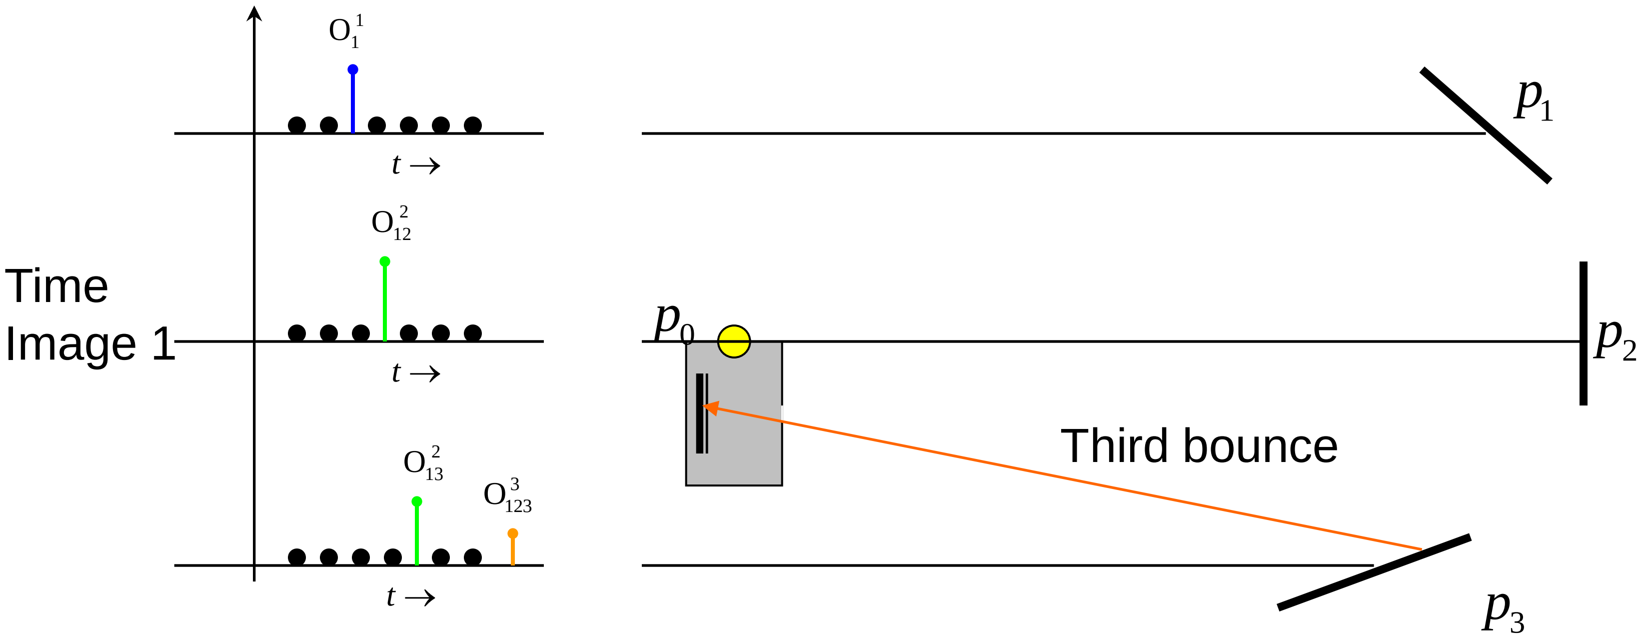

These bounces are recorded on a time profile at each pixel. This is the impulse response for a ray and is a time-image I(x,y,t) in camera

Recording Space Time Impulse Response

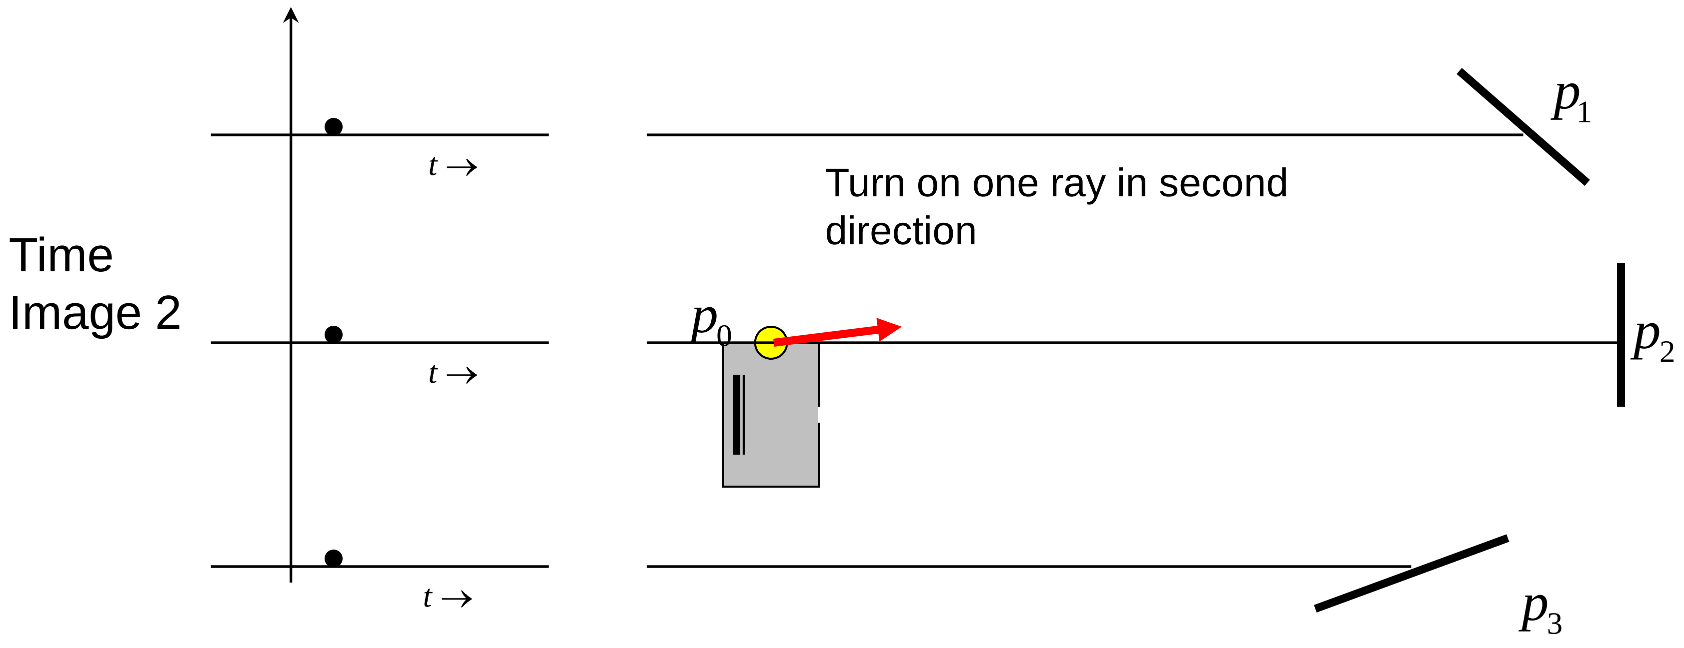

We turn on the second ray and repeat the procedure to capture a second time-image. And so on. The Space Time Impulse Response is union of all such time-images

Forward Problem

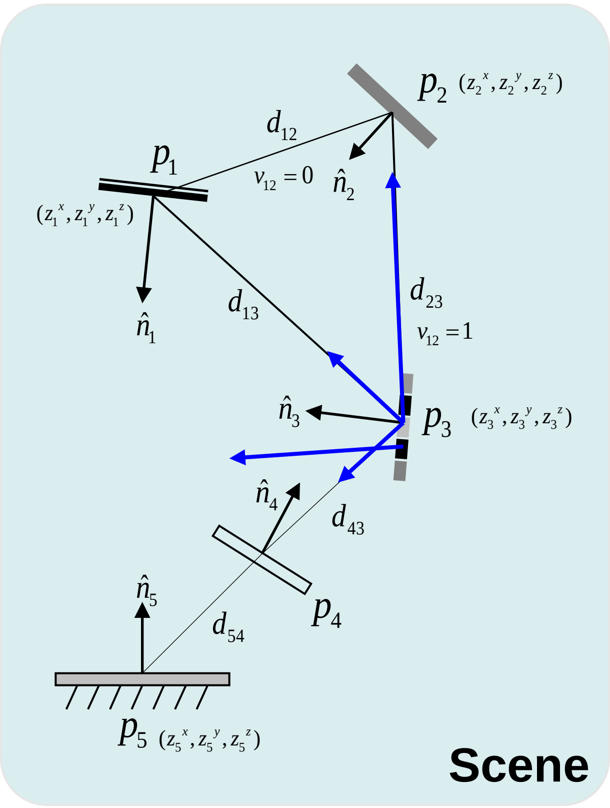

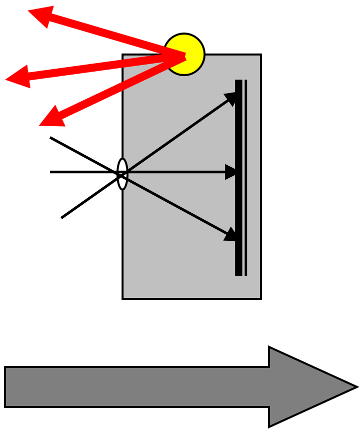

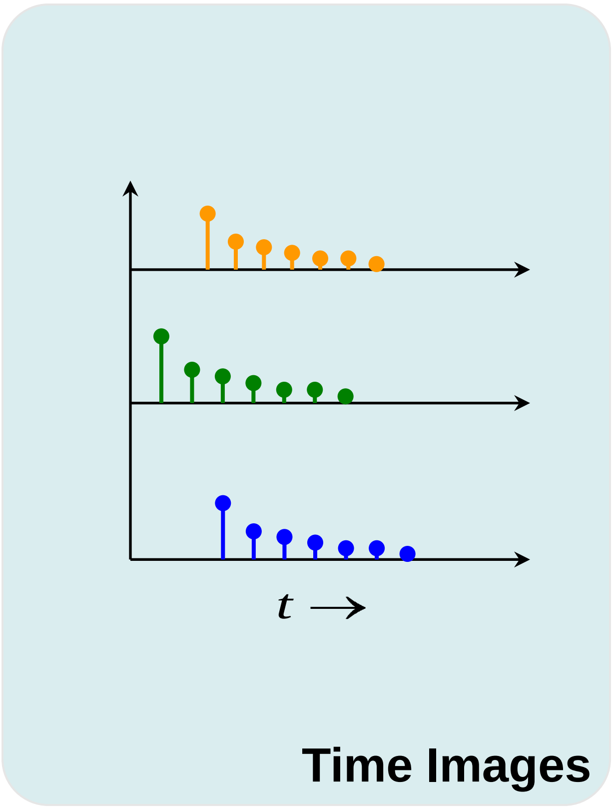

The transient imaging camera maps scene parameters to time-images

The Inverse Geometry Problem

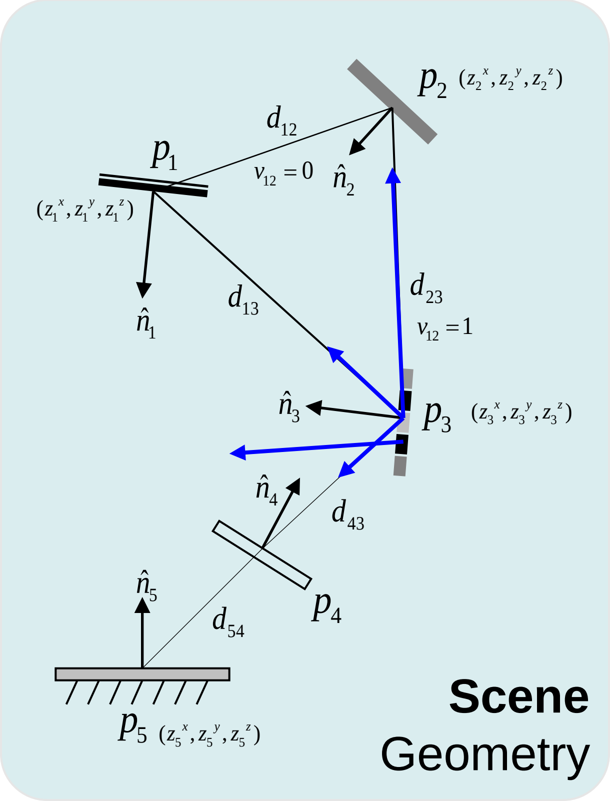

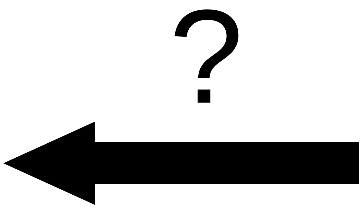

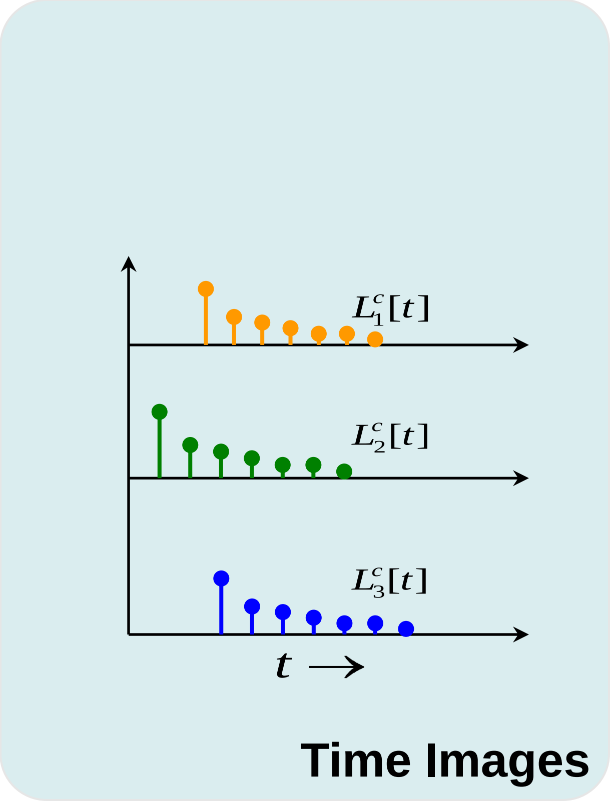

How can we infer scene geometry from its the time image?

Inverse Geometry

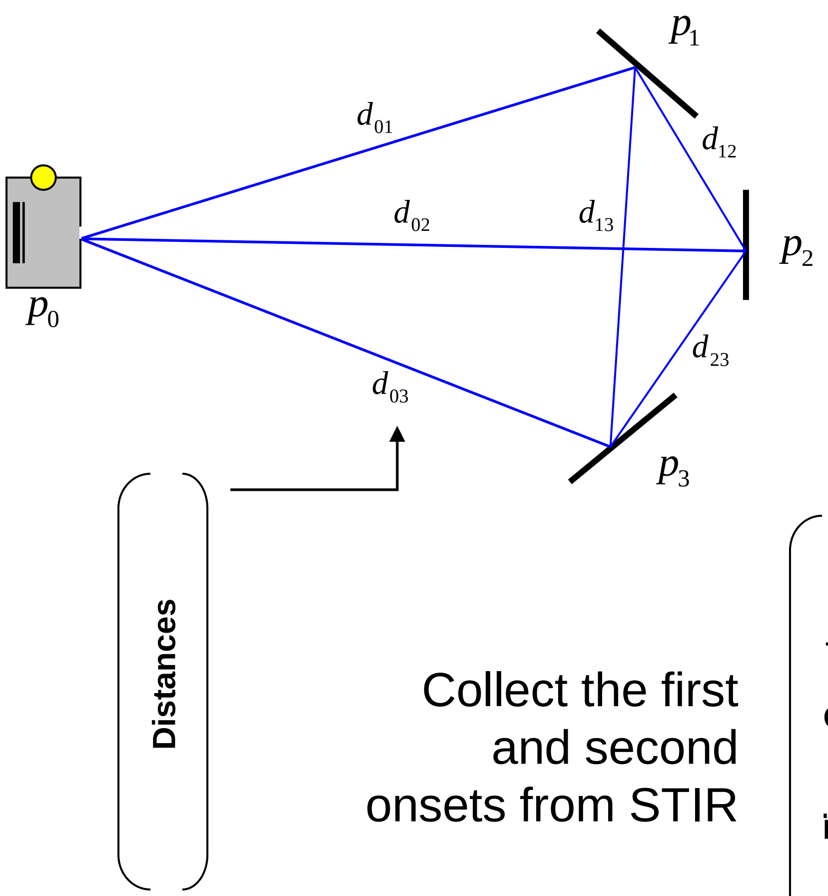

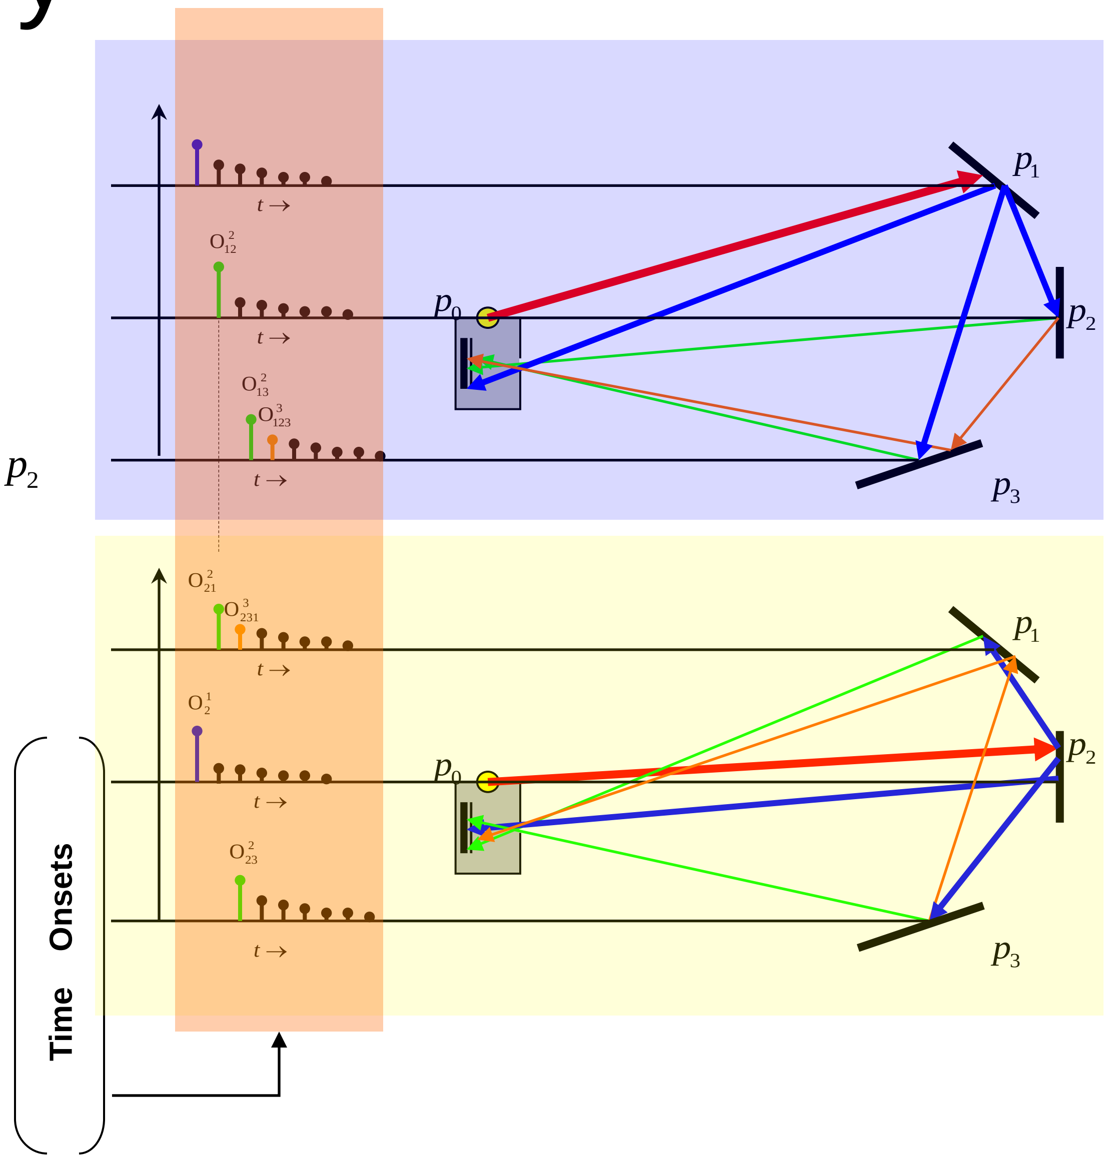

Collect the first and second onsets from STIR

We find camera to patch distance from first bounce. Then we find inter-patch distances from second bounce. It is easy to identify first and second bounces since they are measured in different time images

Relate time onsets to scene distances using the Second Order Distance transform $\mathbf{T_2}$

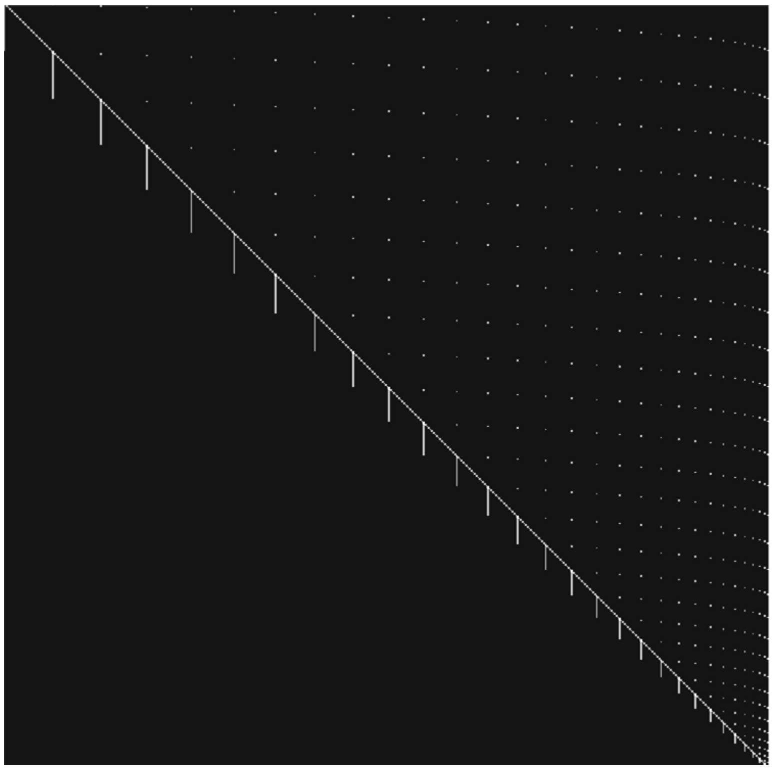

$$\left(\mathbf{T_2}\right)^{-1} \begin{pmatrix} \mathcal{O}^1 \\ \mathcal{O}^2 \end{pmatrix} = \begin{pmatrix} d_{0i} \\ d_{ij} \end{pmatrix} + \left(\epsilon\right)$$

$\mathbf{T_2}$	$\mathbf{O}$	$\mathbf{d}$	ϵ
Distance Transform	**First and Second bounce Time Onsets**	**Direct and inter-patch distances**	**Distance Uncertainty**

We create a distance transform matrix T_2 and invert the linear system to obtain estimates

Geometry Reconstruction

- We recover coordinates using isometric embedding
- We reconstruct surface normals using smoothness assumptions

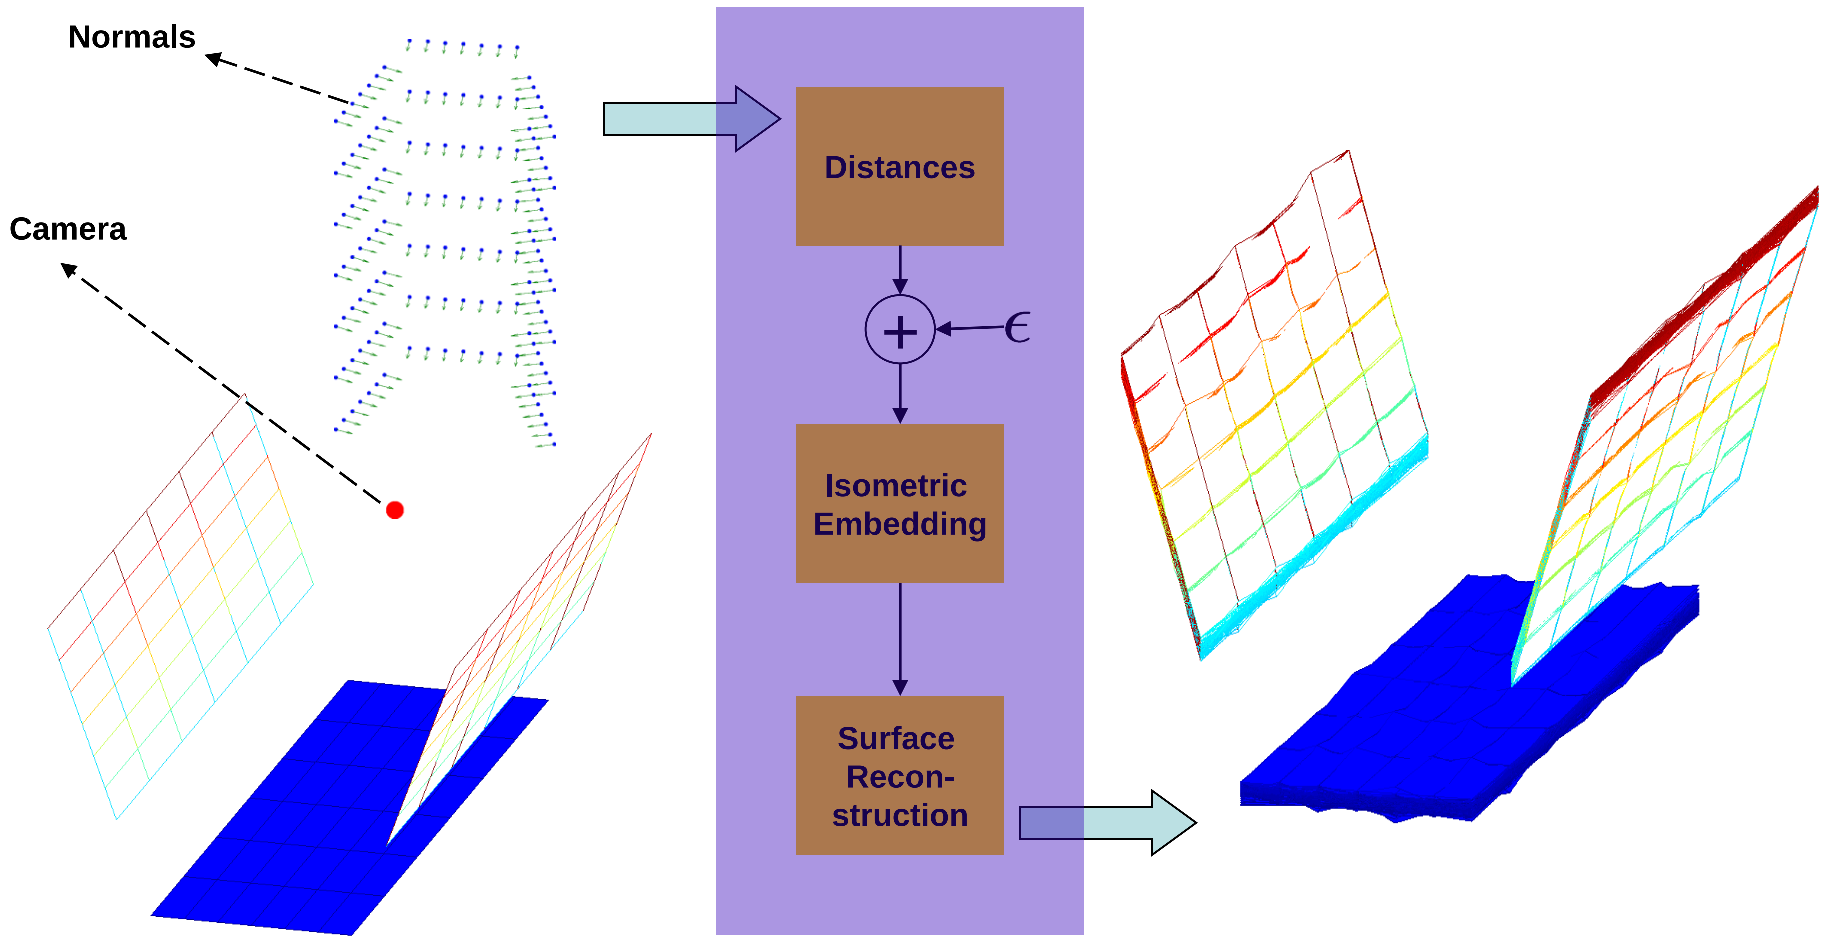

Estimating Hidden Patch Geometry

Patches invisible to both camera and the illumination source

Measuring the STIR of a simple scene with 2 hidden patches

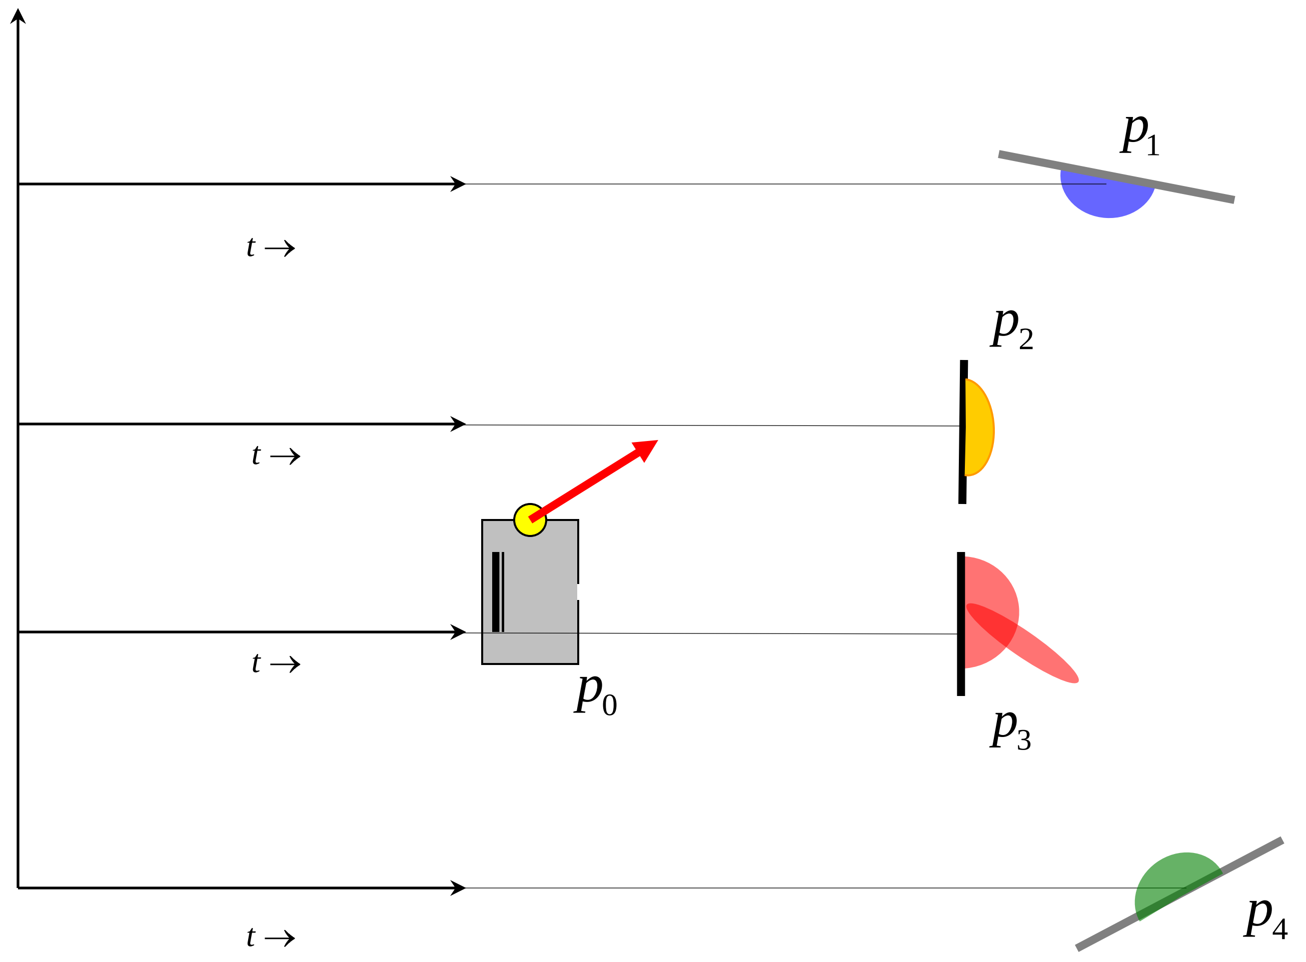

For estimating hidden patch geometry, we use third bounces. Consider two hidden patches, P_2 and P_3

Direct bounce

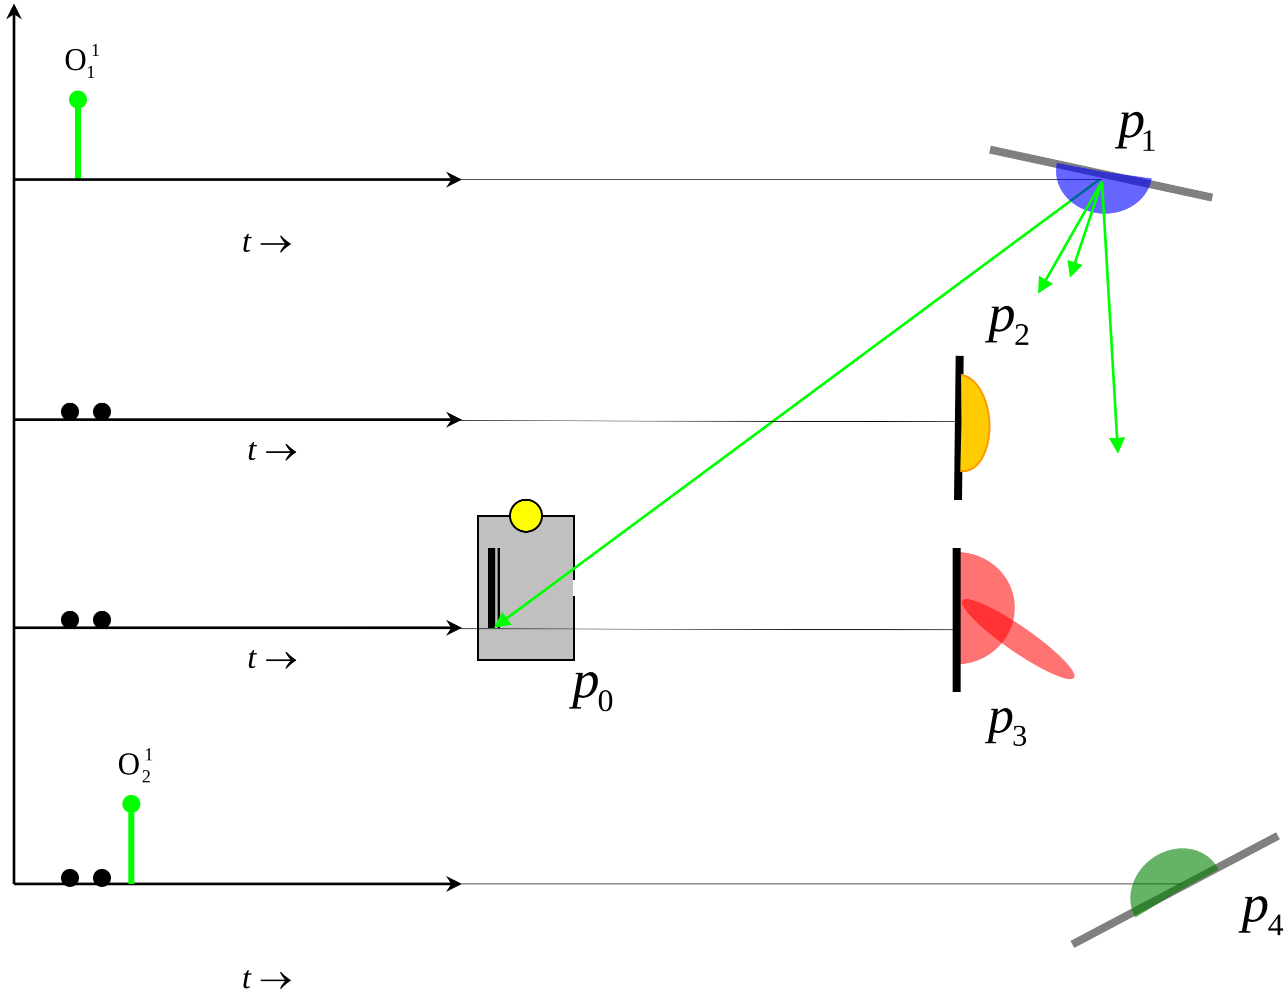

We illuminate P_1 with a ray impulse and we record the first bounce from P_1. The first bounce also creates reflections towards all other patches

Direct bounce

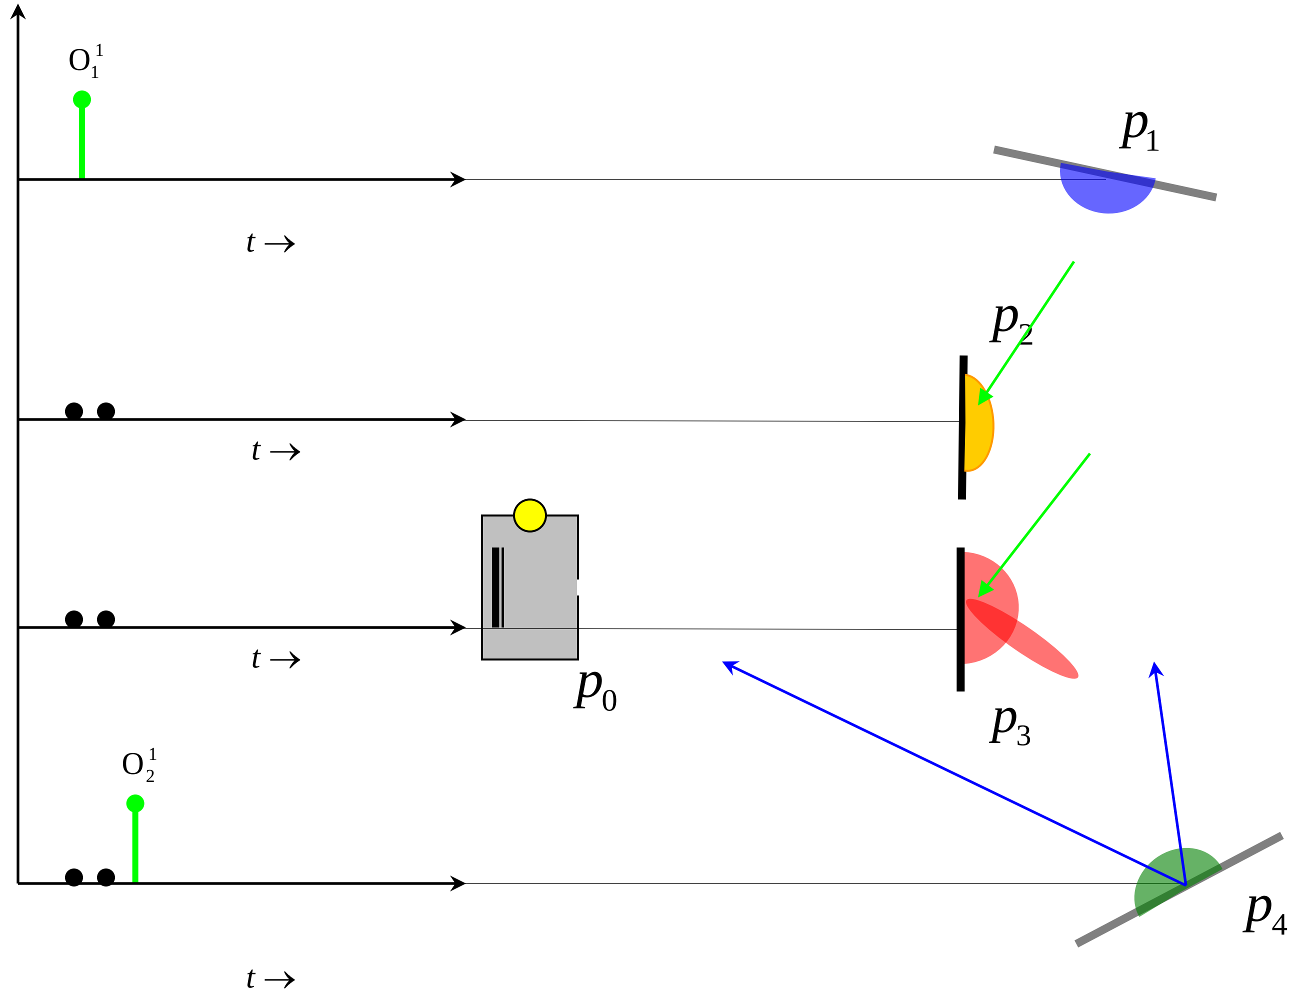

Second bounce reflections are created from both visible and hidden patches.
Only reflections from P4 and P1 reach the camera

Second bounce

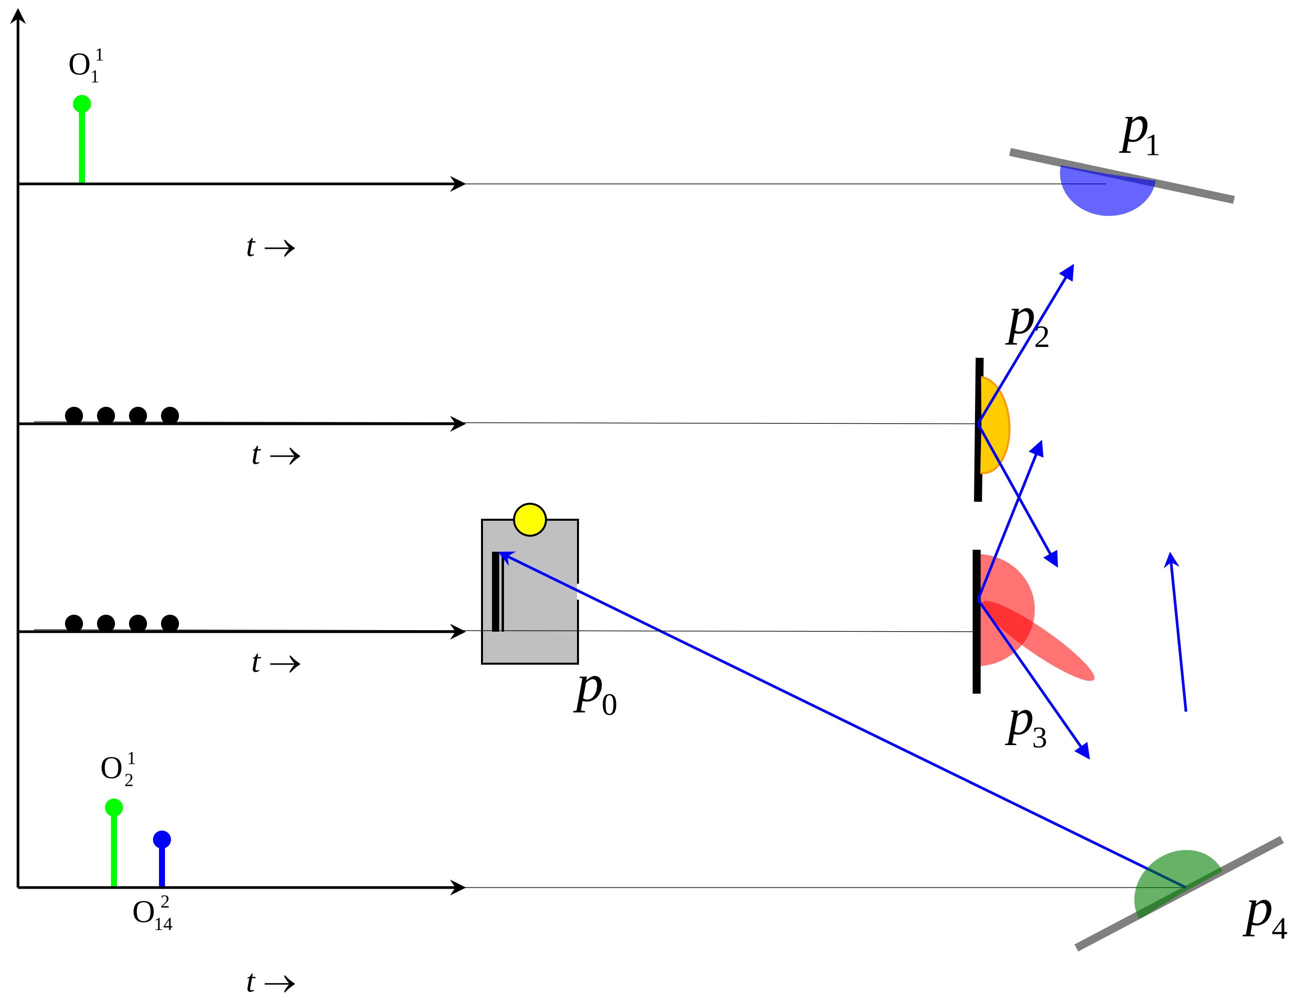

Second bounce reflections from hidden patches reach the visible patches

Second bounce

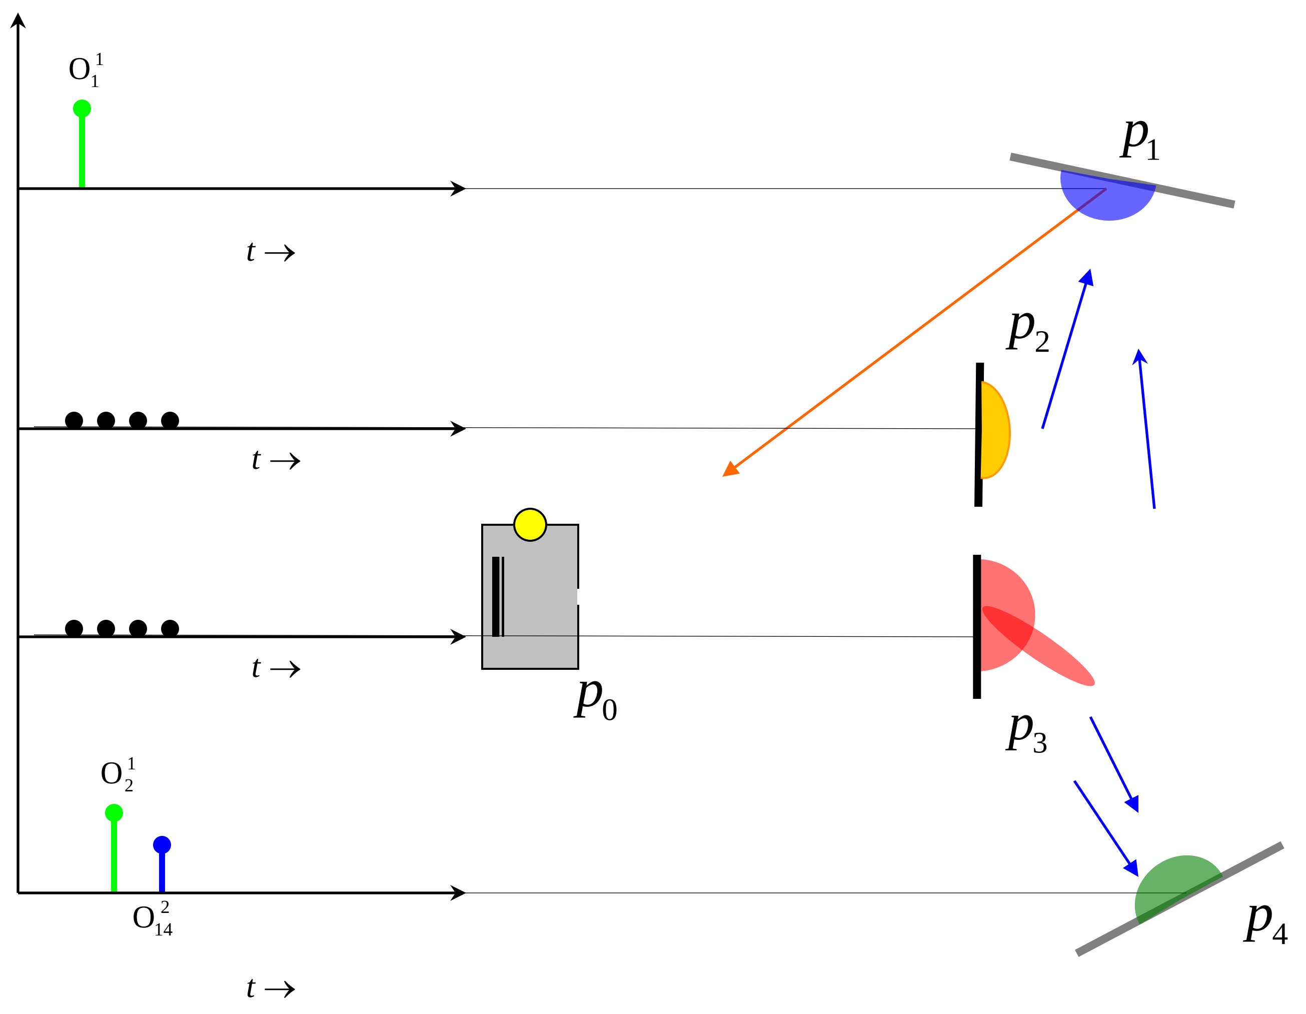

The third bounces contain information about the distances to the hidden patches

Key idea: Use third Bounce

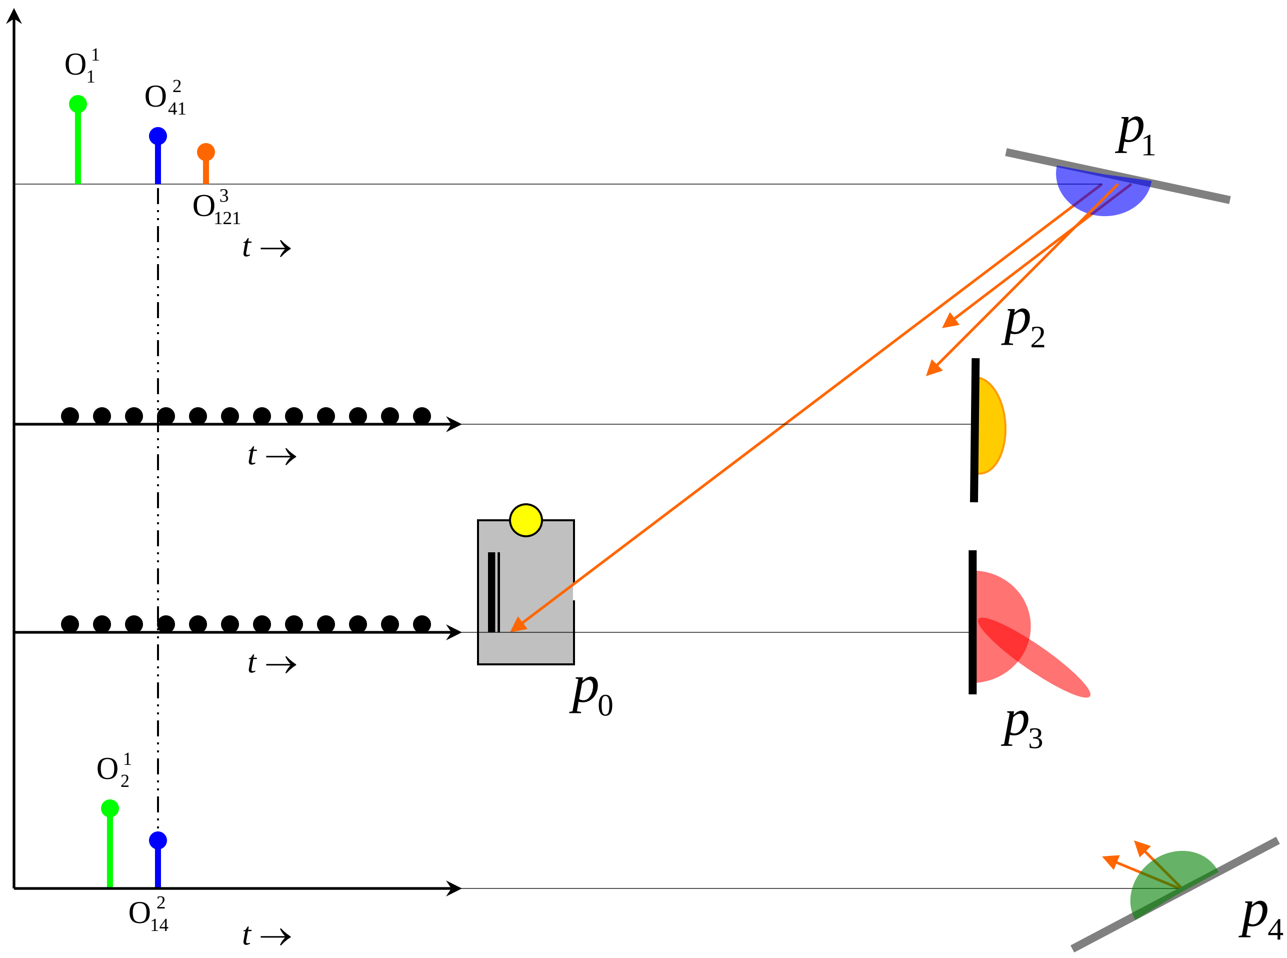

We then record the third bounce corresponding to $P_o \rightarrow P_1 \rightarrow P_2 \rightarrow P_1 \rightarrow P_o$

Key idea: Use third Bounce

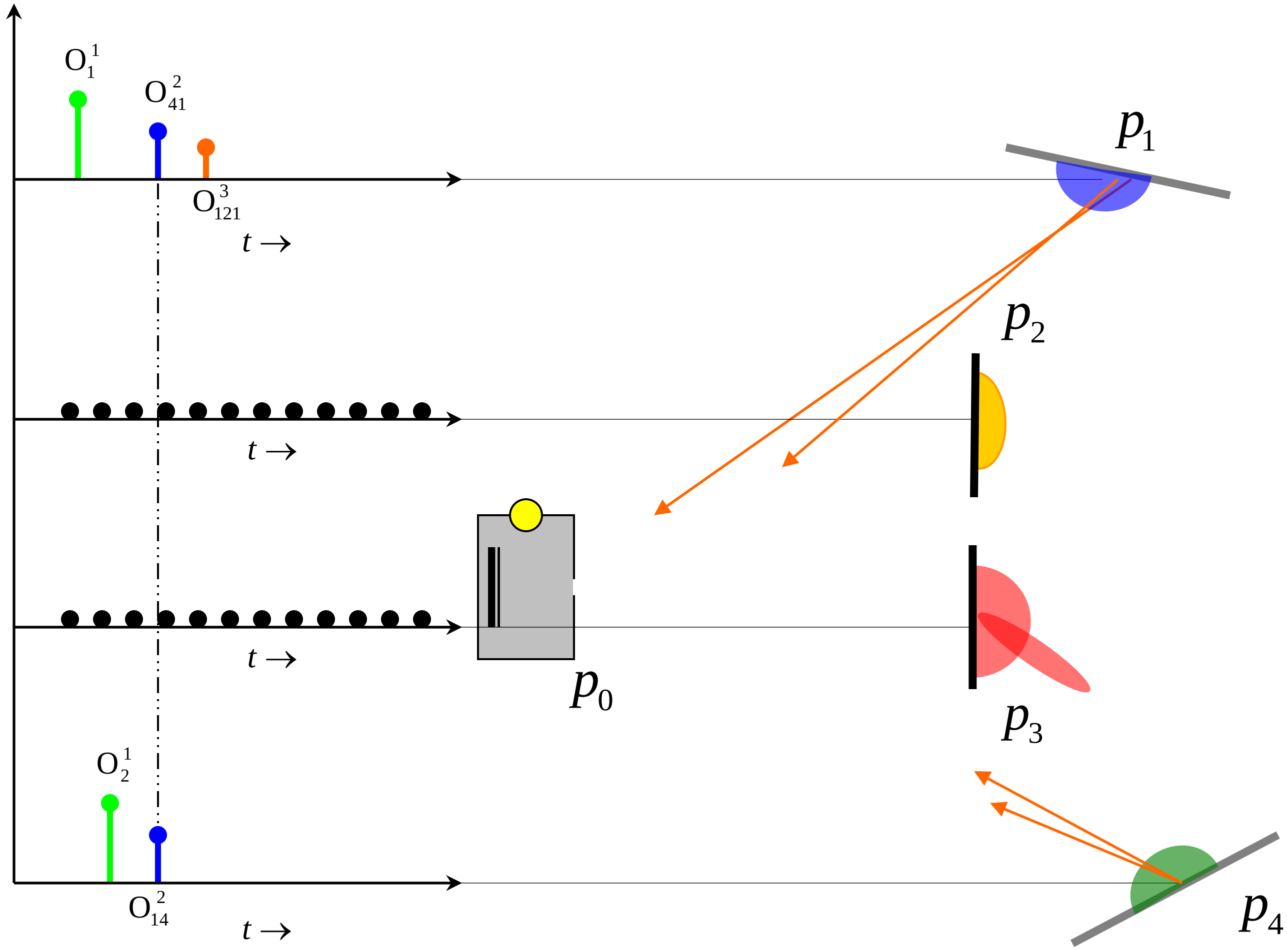

Note that the time of arrival of bounces depends on the Euclidean distance between patches

Key idea: Use third Bounce

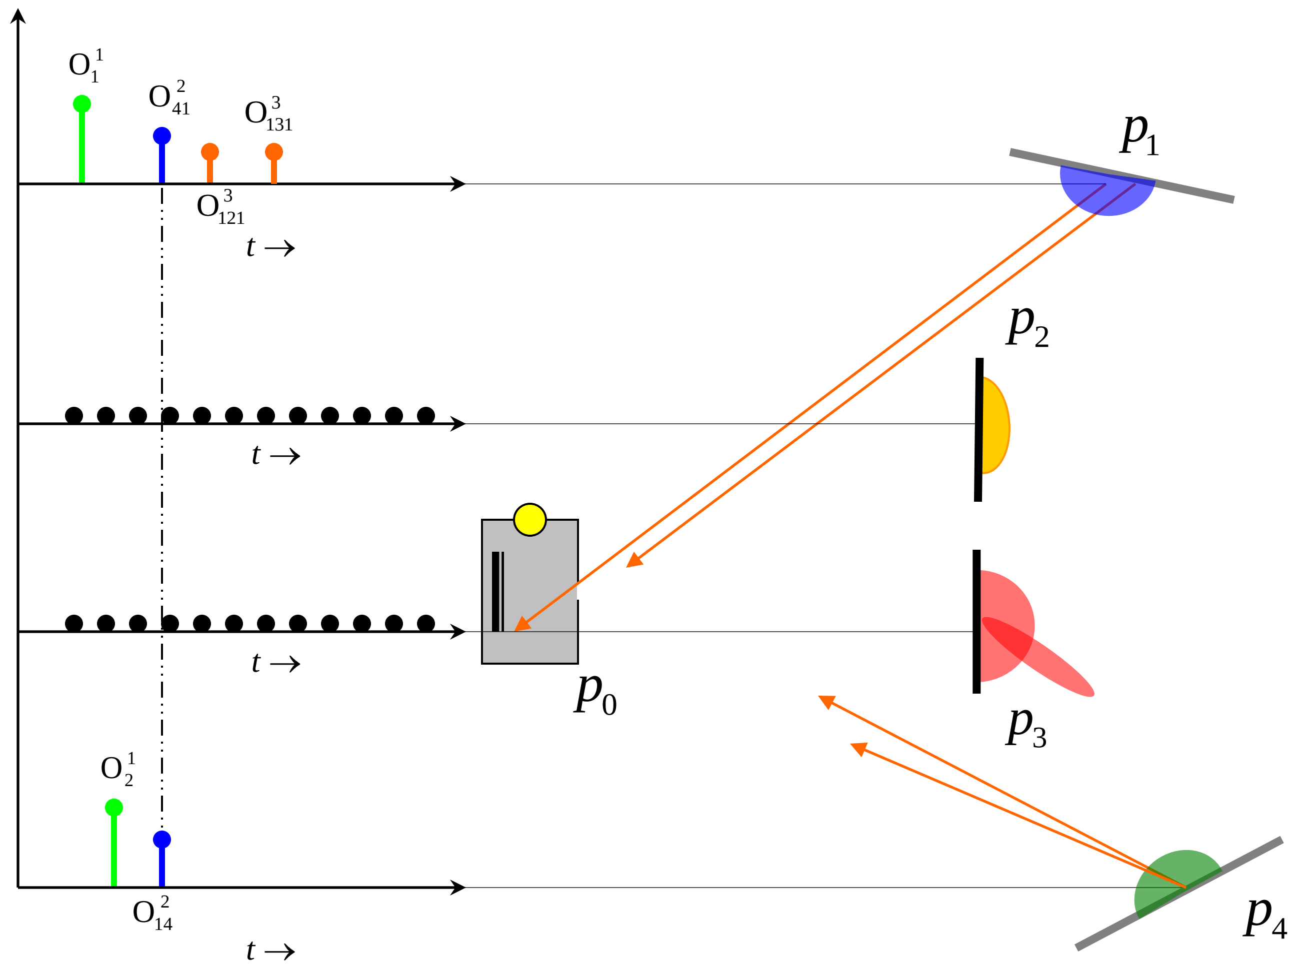

We now record the third bounce corresponding to $P_o \rightarrow P_1 \rightarrow P_3 \rightarrow P_1 \rightarrow P_o$

Key idea: Use third Bounce

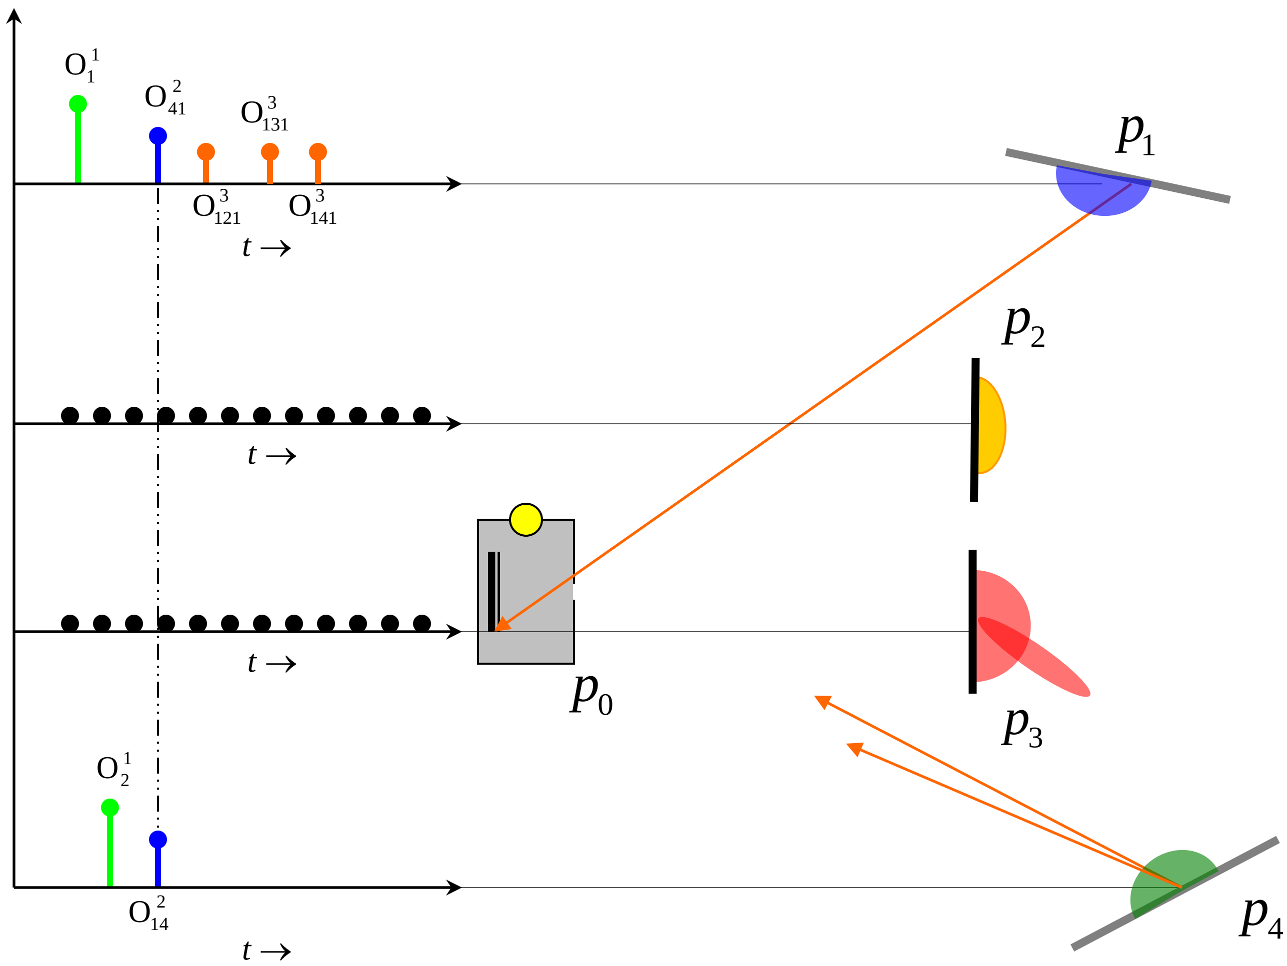

We now record the third bounce corresponding to $P_o \rightarrow P_1 \rightarrow P_4 \rightarrow P_1 \rightarrow P_o$

Key idea: Use third Bounce

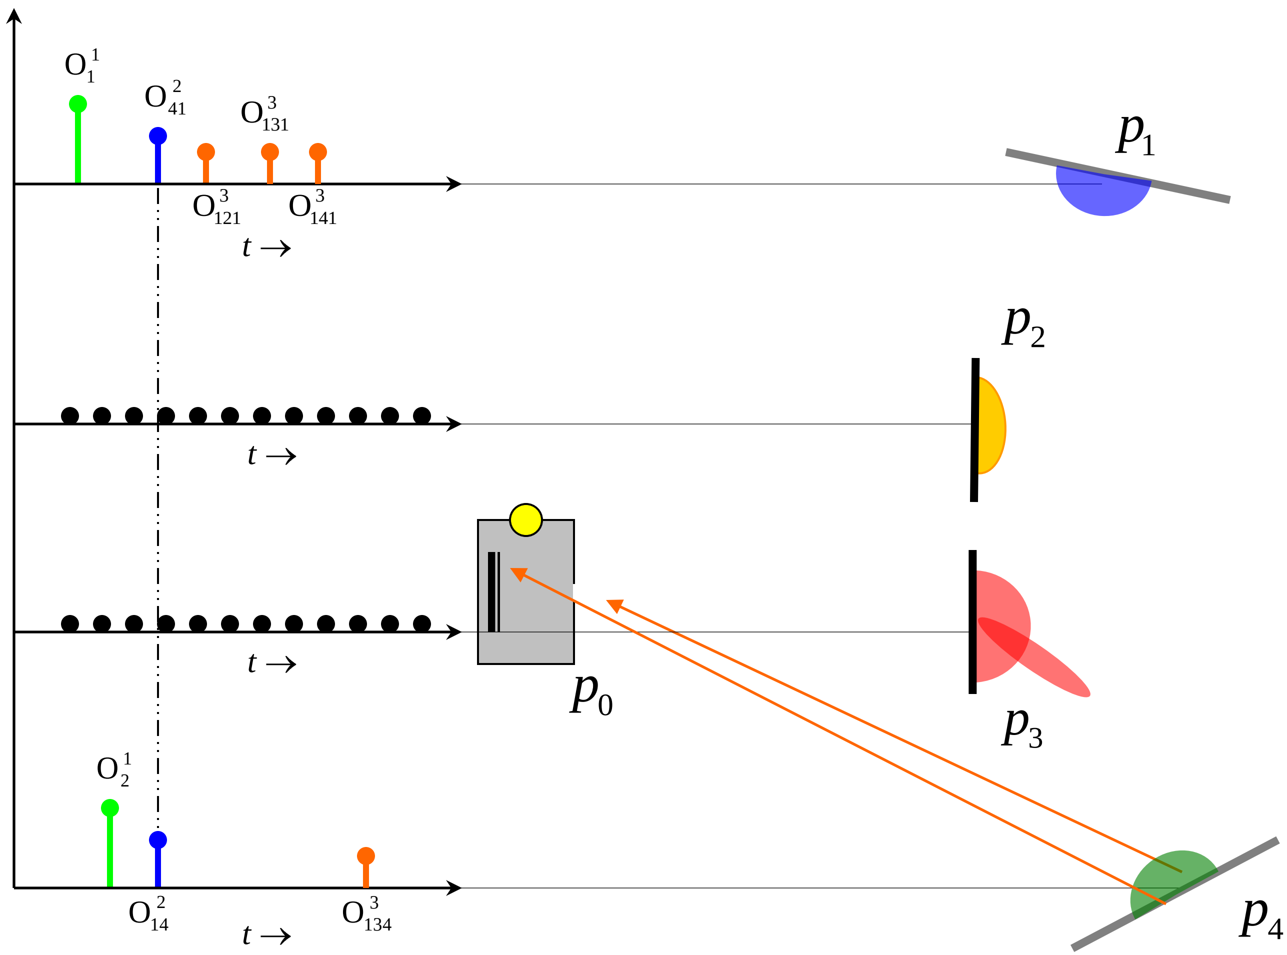

Similarly we record the third bounce corresponding to $P_o \rightarrow P_1 \rightarrow P_3 \rightarrow P_4 \rightarrow P_o$

Key idea: Use third Bounce

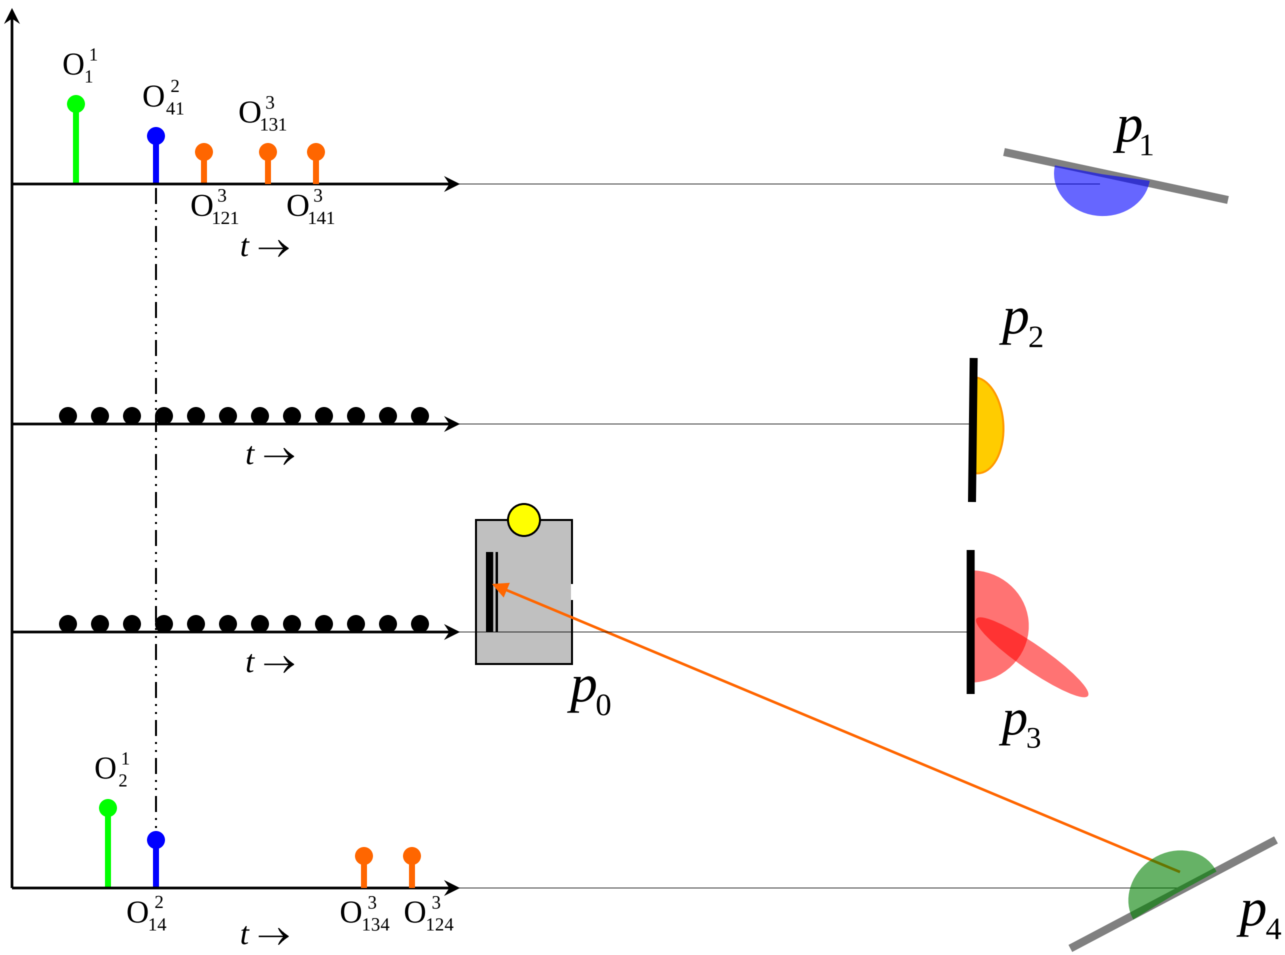

… and $P_o \rightarrow P_1 \rightarrow P_2 \rightarrow P_4 \rightarrow P_o$

Key idea: Use third Bounce

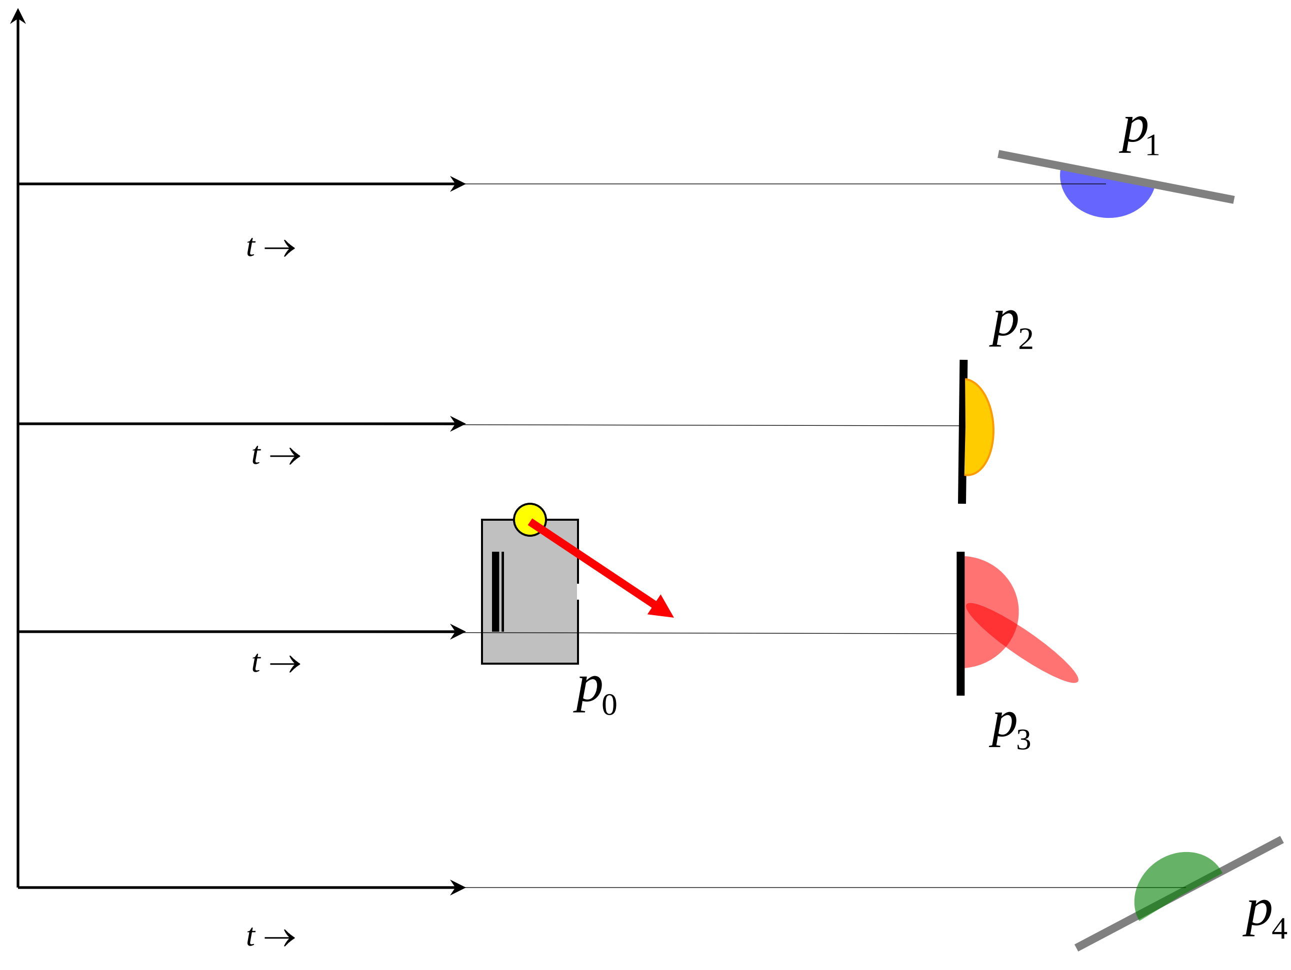

We repeat this process for visible patch P_4

Labeling third bounces

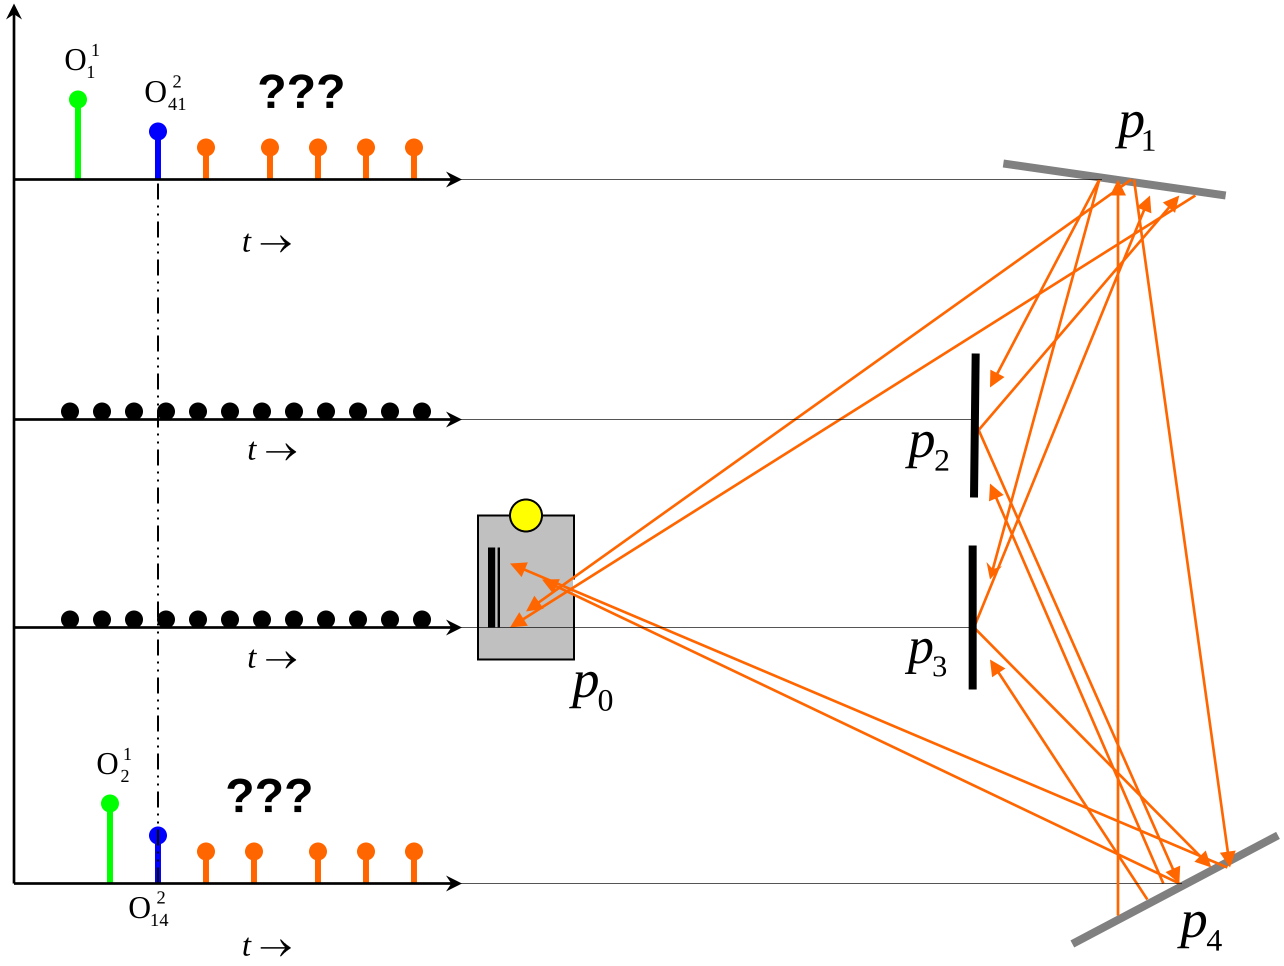

Disambiguating third bounces is non-trivial.

Labeling third bounces

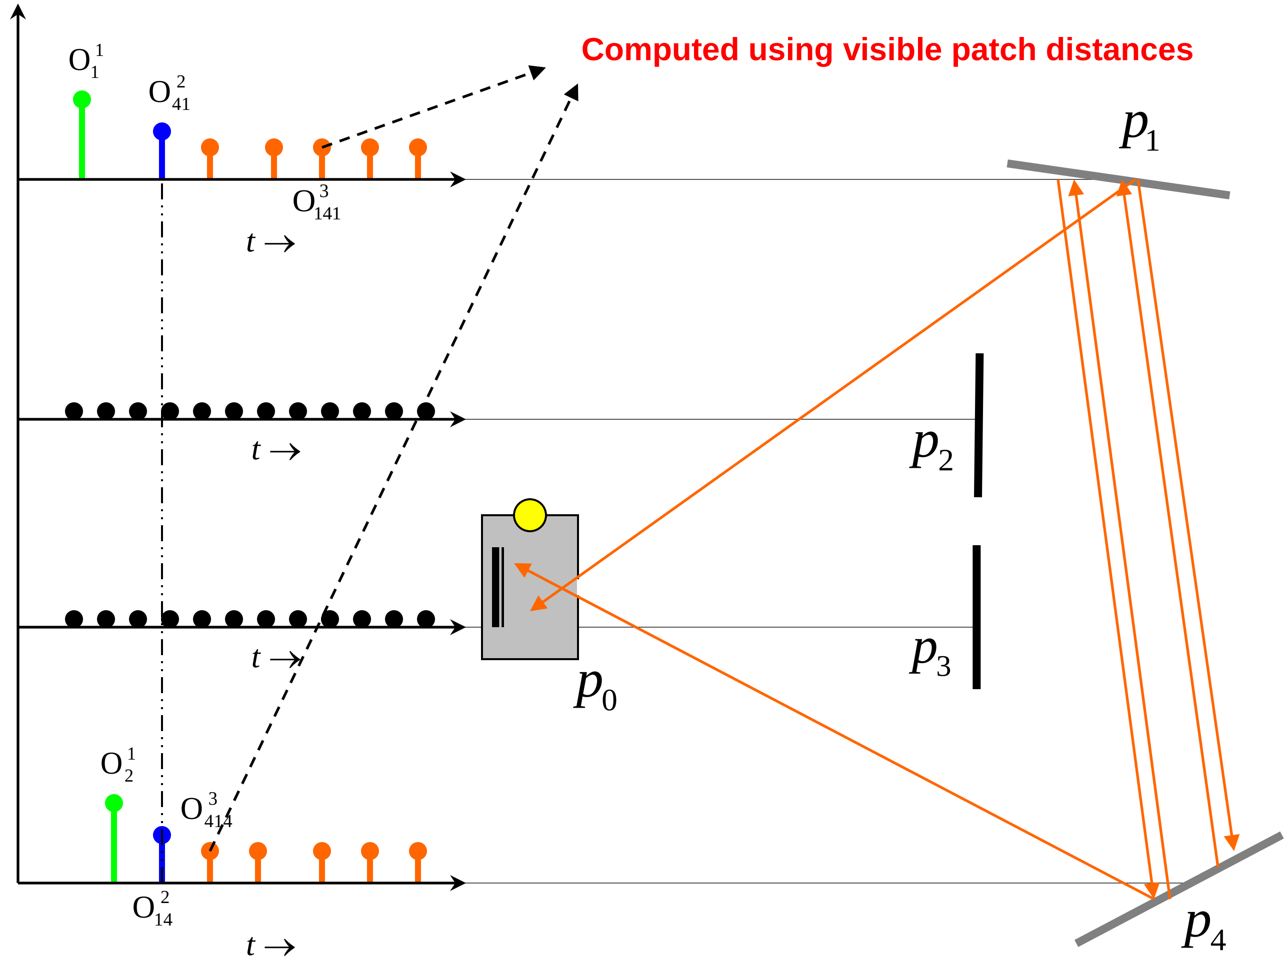

Labeling third bounces

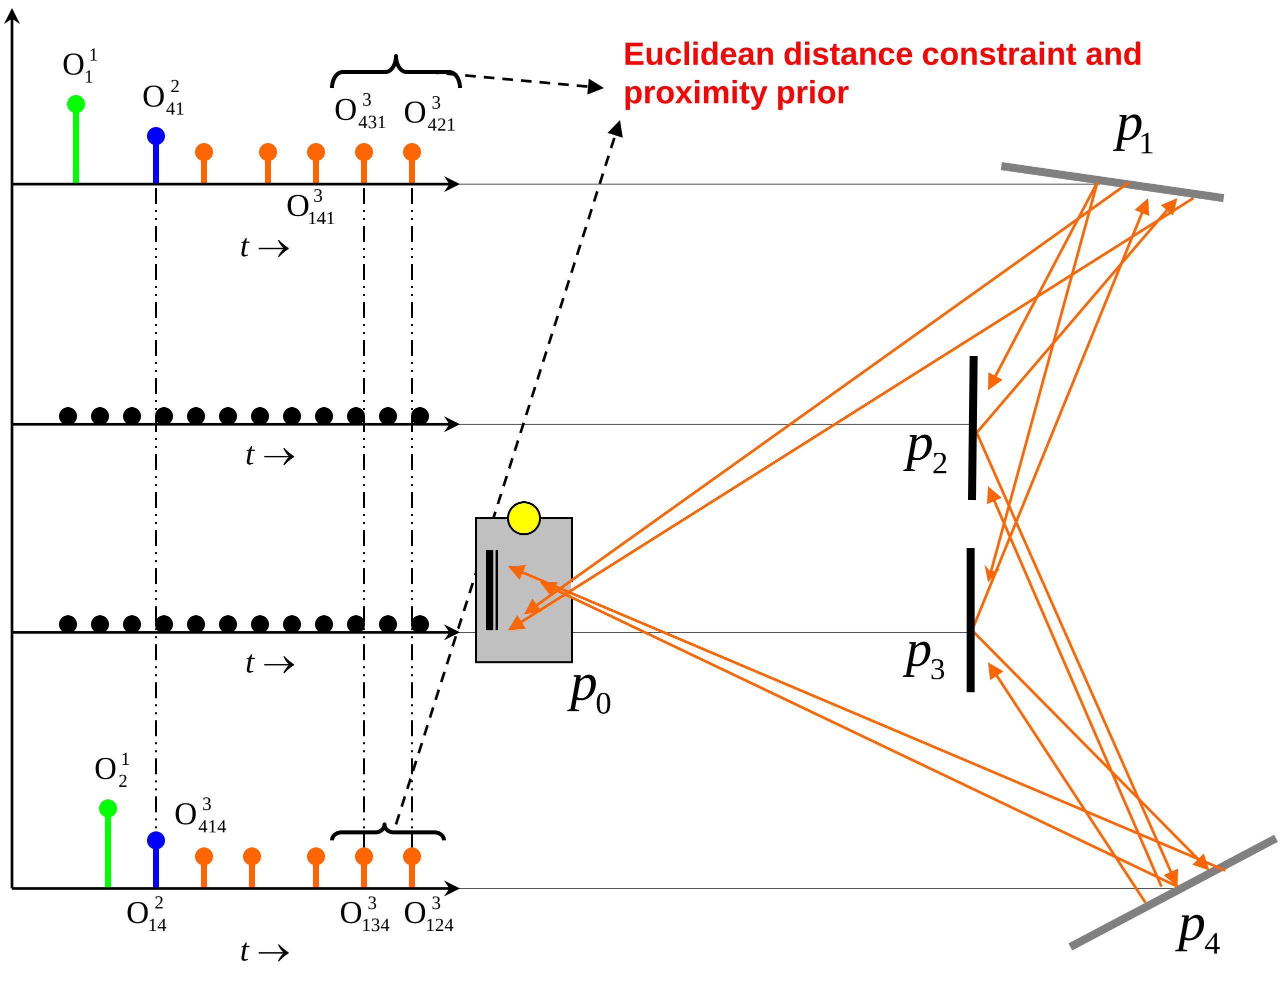

Proximity prior: Hidden patch P2 is closer to P1 than patch P3

Labeling third bounces

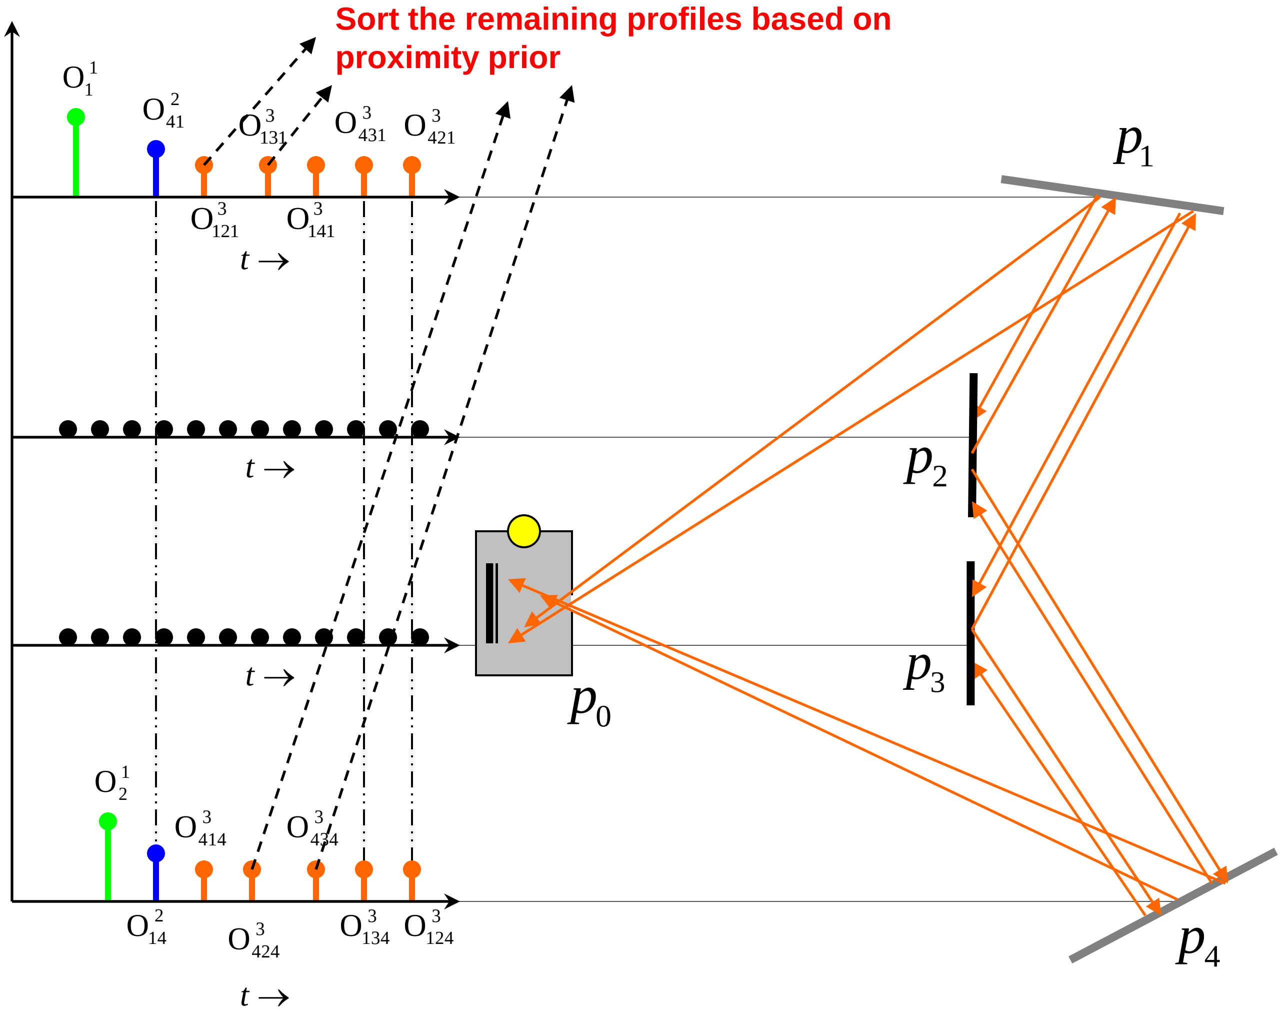

Labeling third bounces

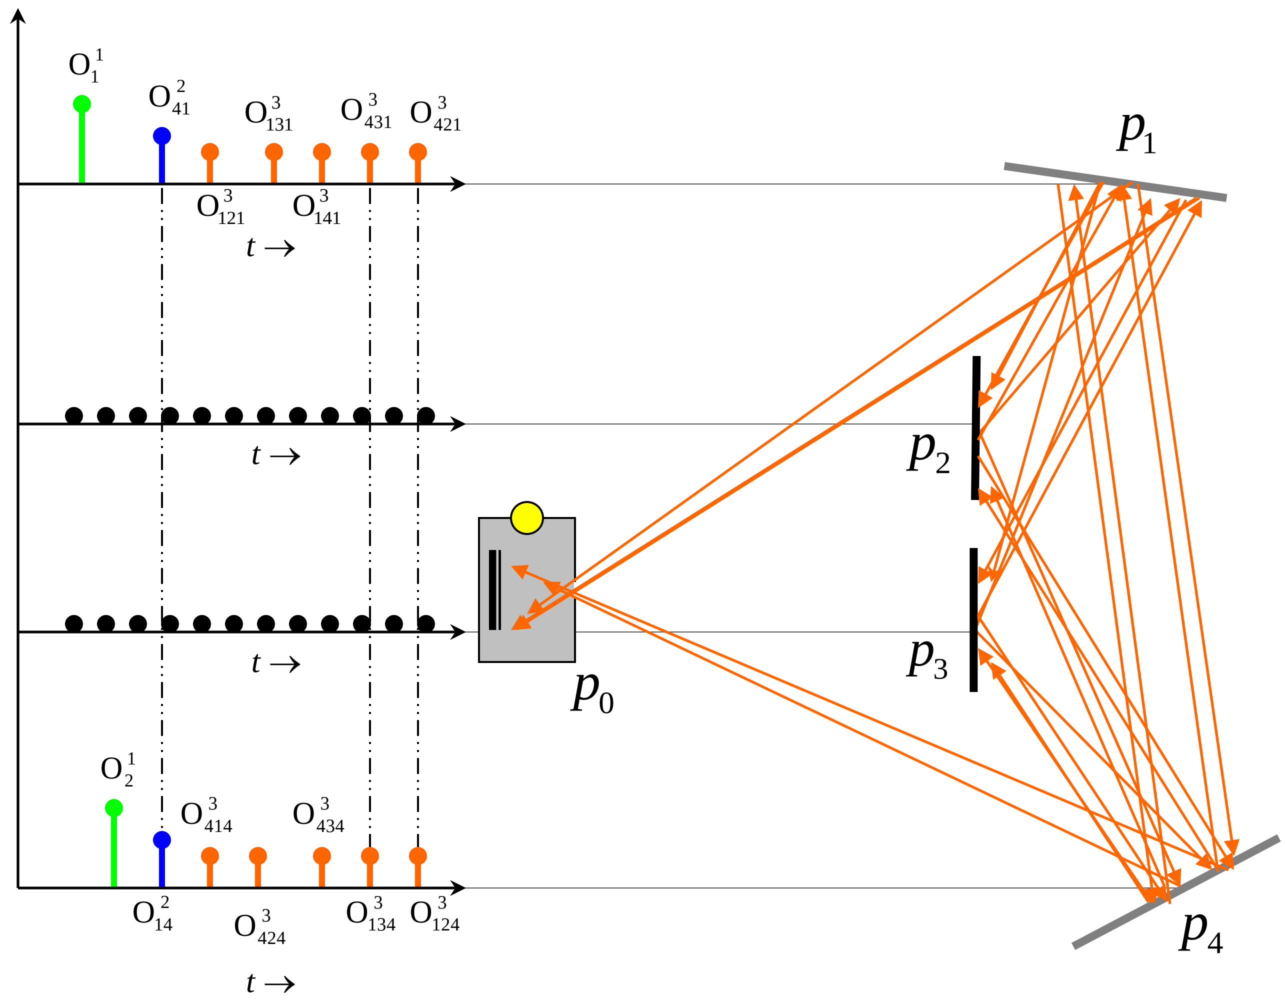

We made use of Proximity prior, Euclidean distance constraints and estimation of visible patch distances to label third bounces

Relate time onsets to scene distances using the Second Order Distance transform $\mathbf{T_2}$

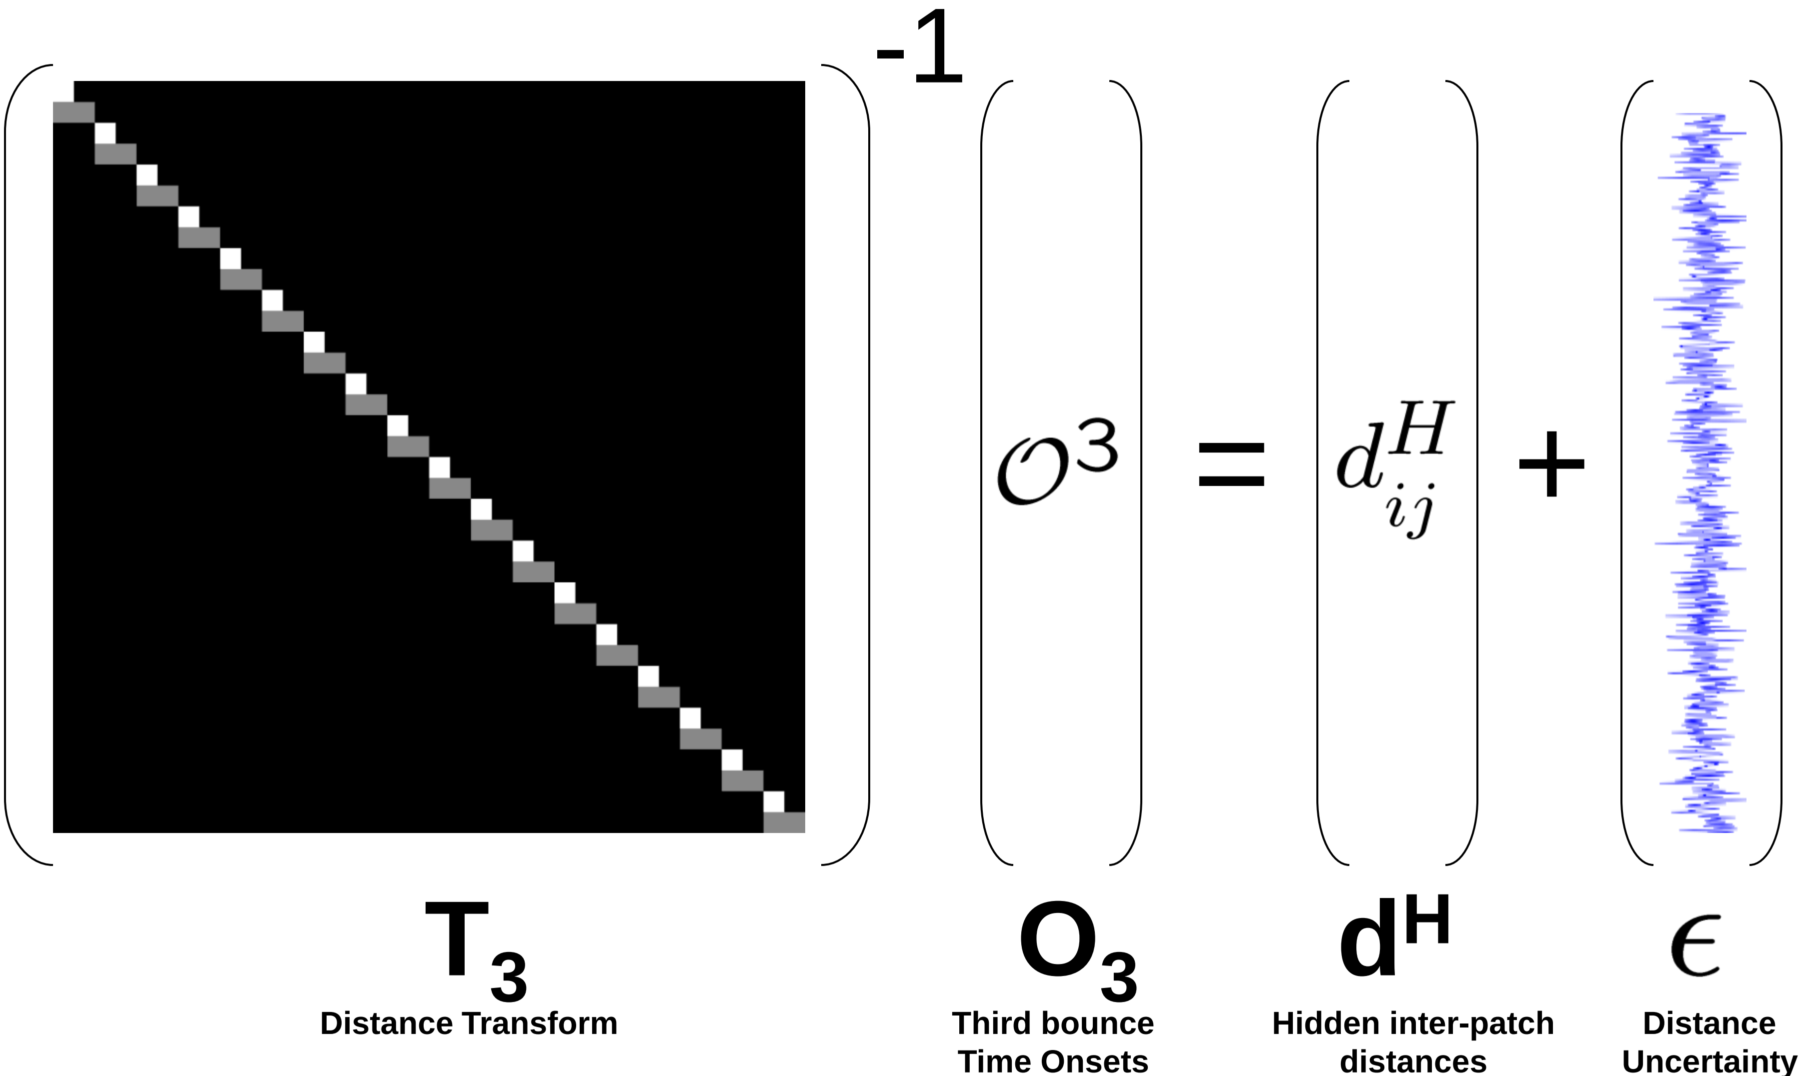

As shown in the paper, we setup a transform matrix, T_3 and solve for hidden patch geometry. Note the geometry computation is possible even when the reflectance is not diffuse and is unknown.

Hidden Surface Geometry Reconstruction

- Isometric embedding and smoothness assumptions

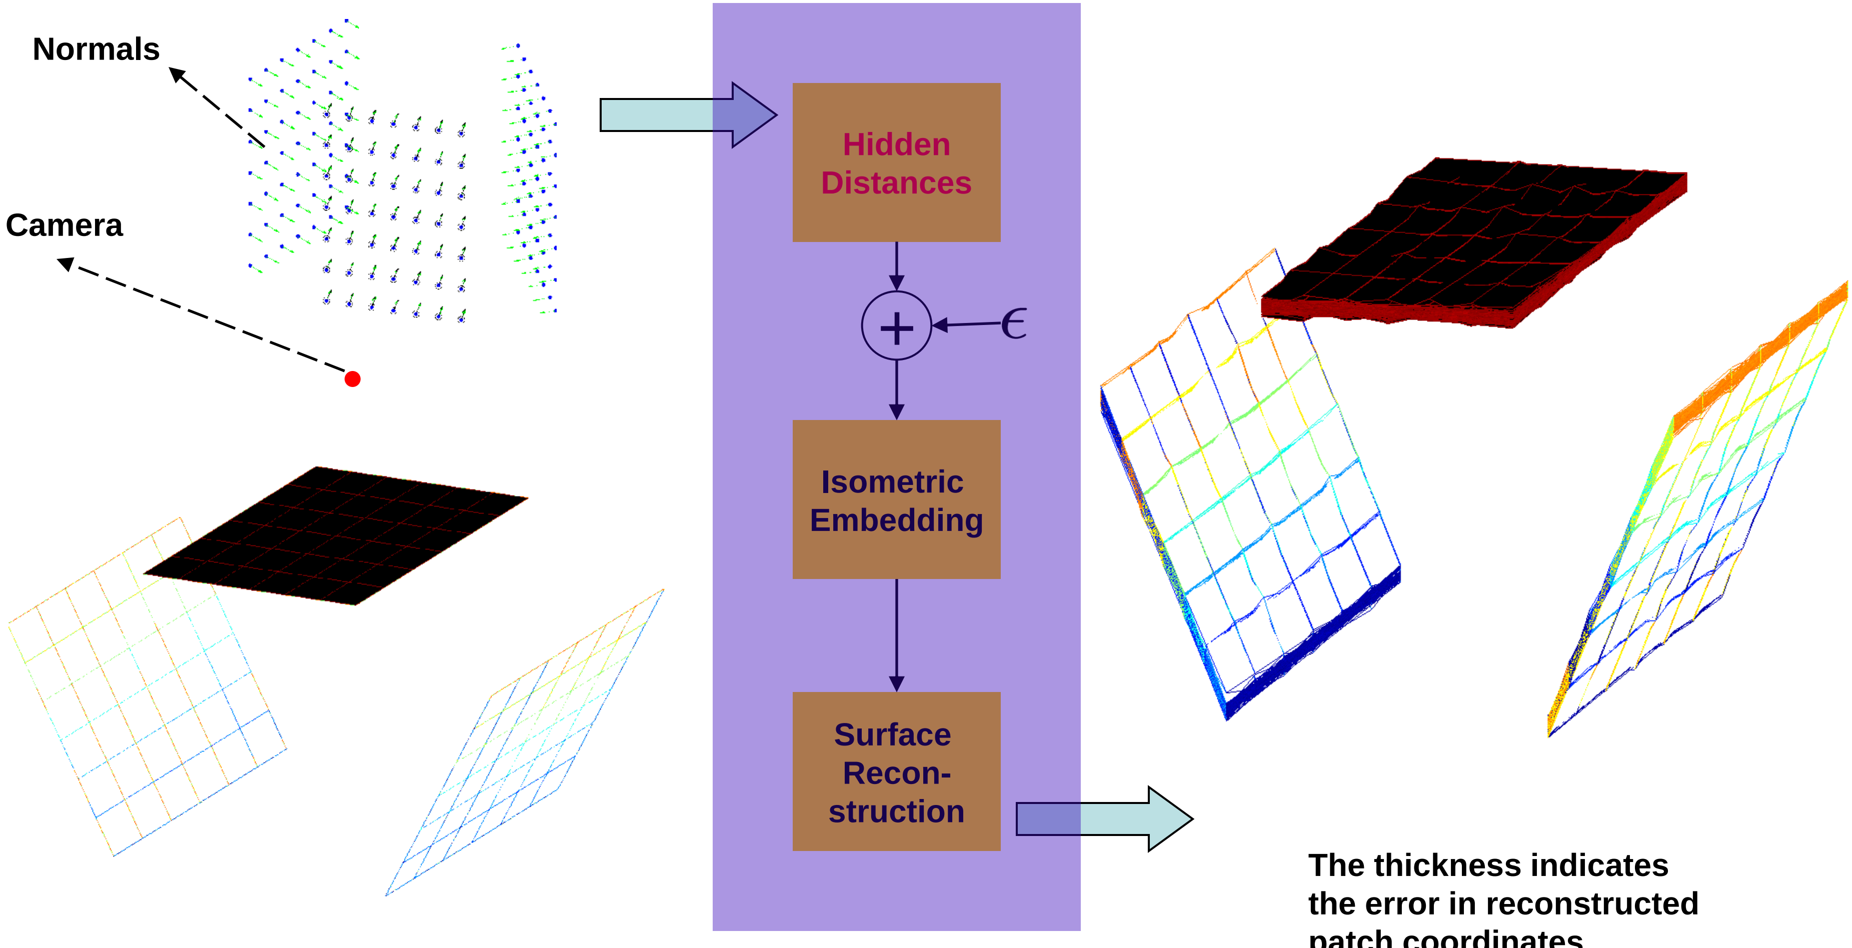

Transient Imaging Camera:
Hardware Experiments

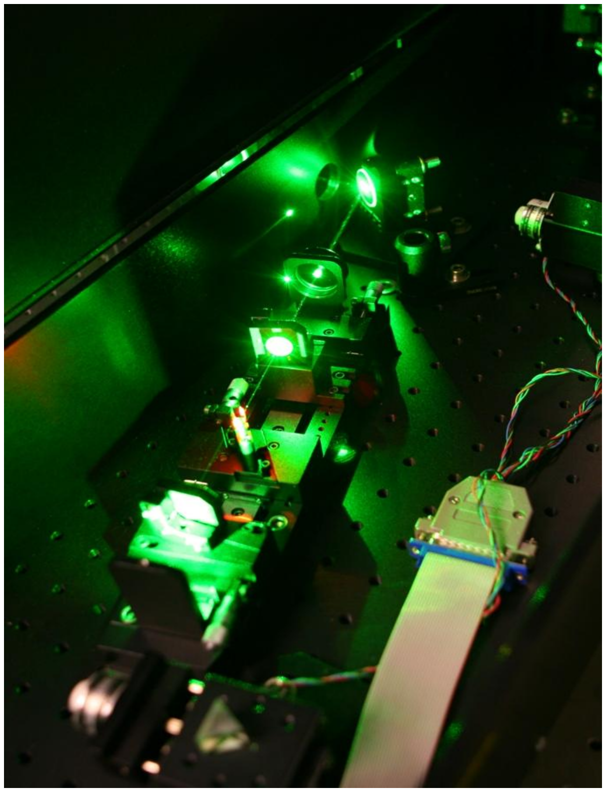

LASER

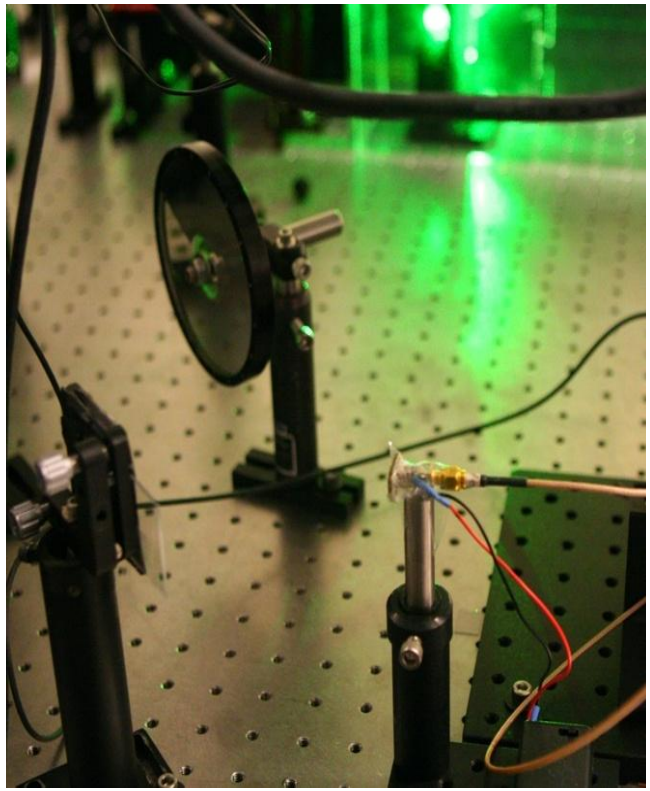

SENSOR

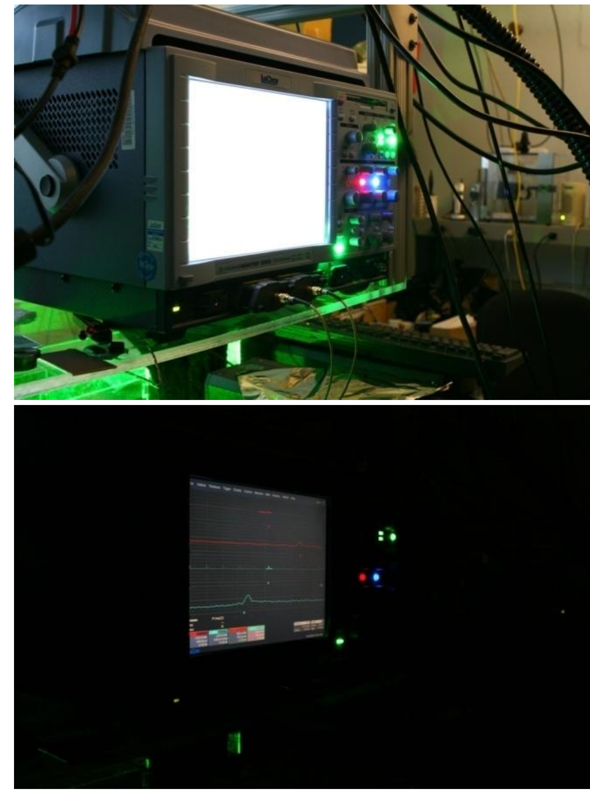

ADC (OSCILLOSCOPE)

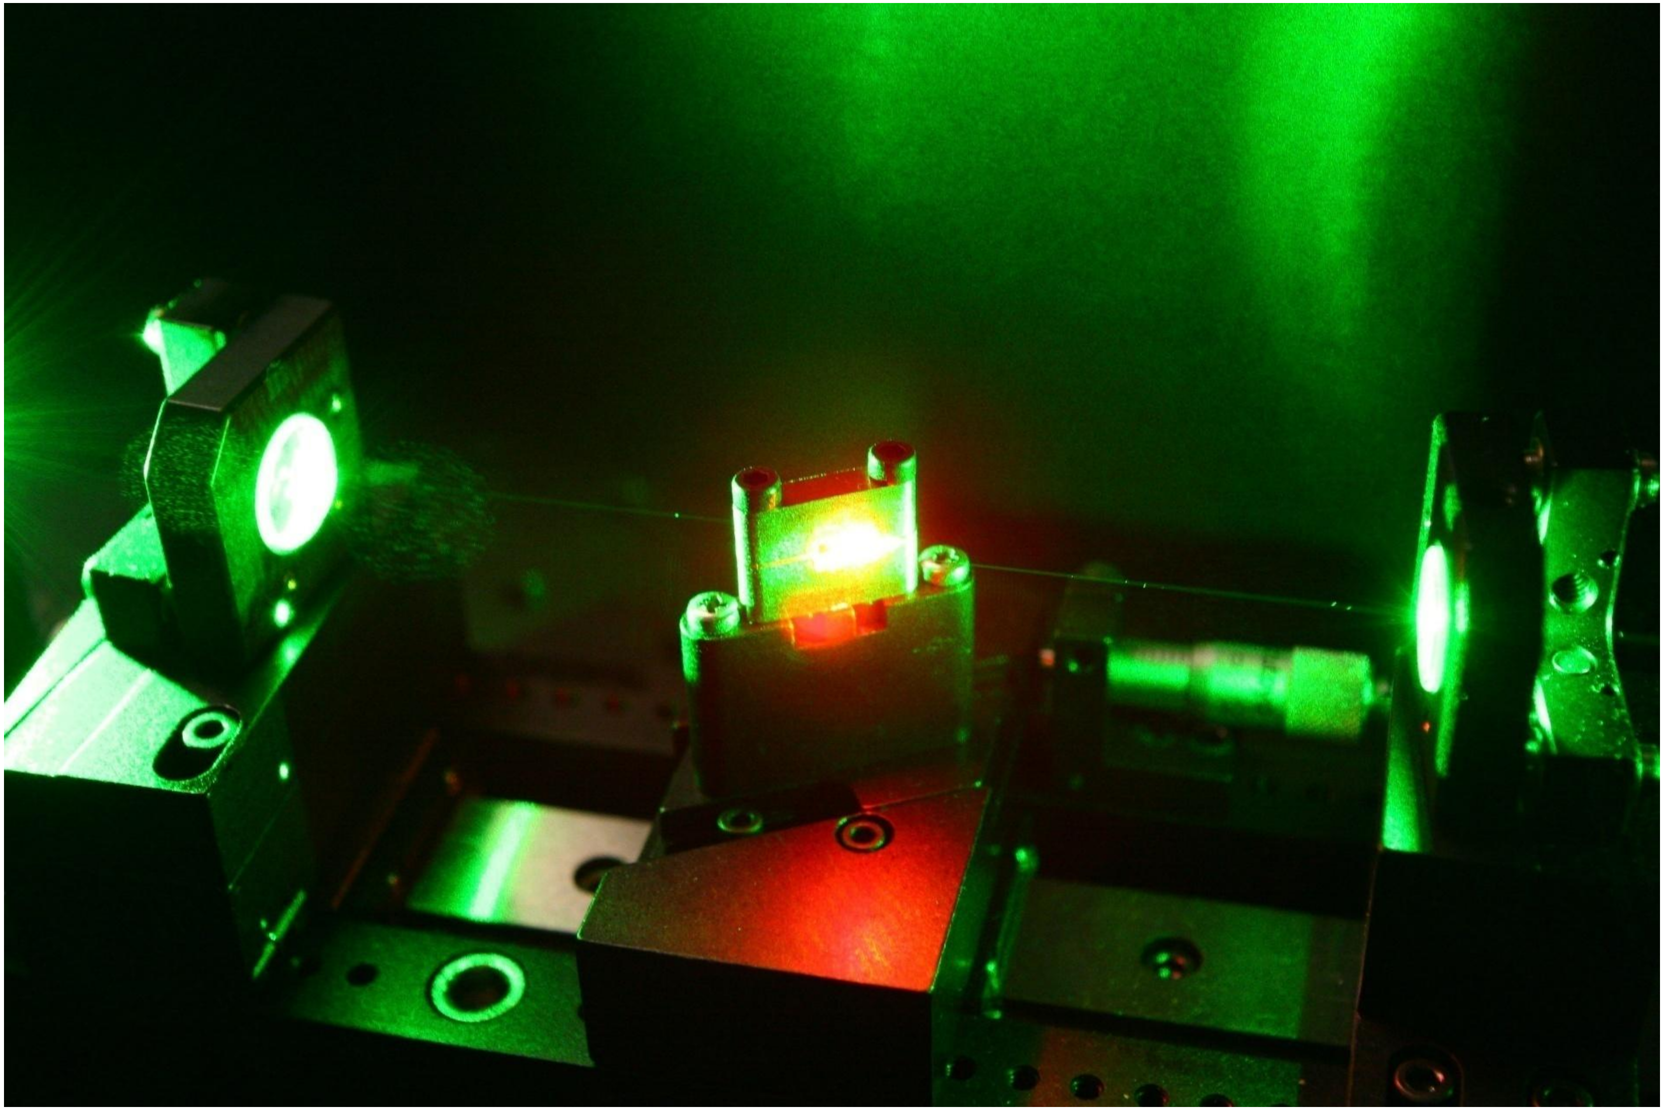

Our source is an ultra short laser with pulse duration of 50 femtoseconds, repeating about every 10 nanoseconds, with average power of 420 milliwatts.

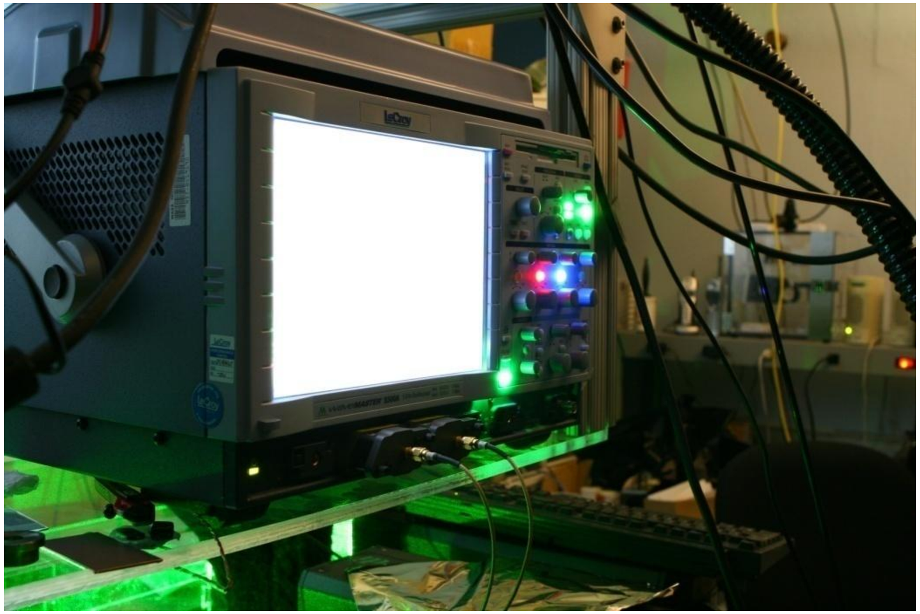

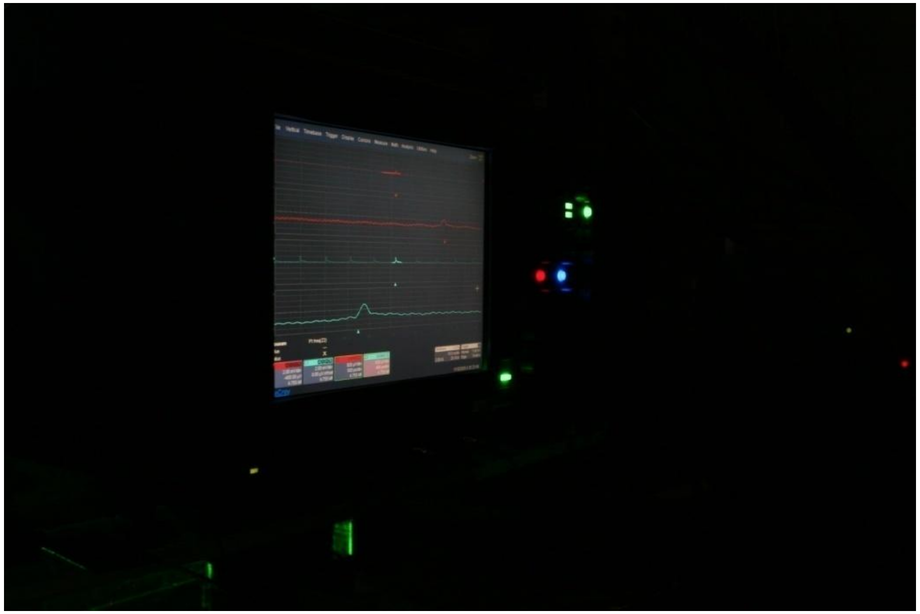

On this 5 GHz scope, we can trigger via a reference sensor and make key measurements.

We used a commercially-available reverse-biased silicon photodiode (Thorlabs FDS02, inexpensive at $72) to capture incoming light.

Design Experiments

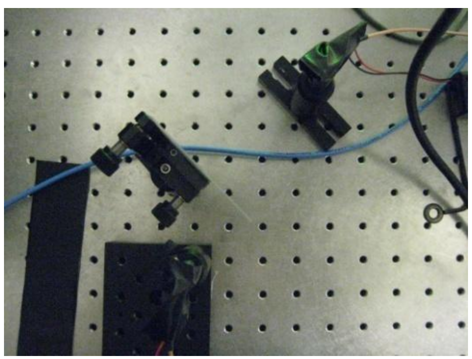

Synchronization

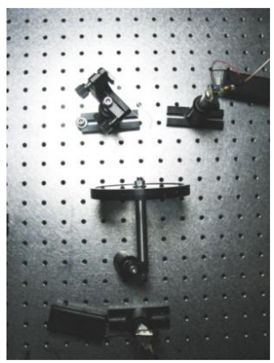

Linearity

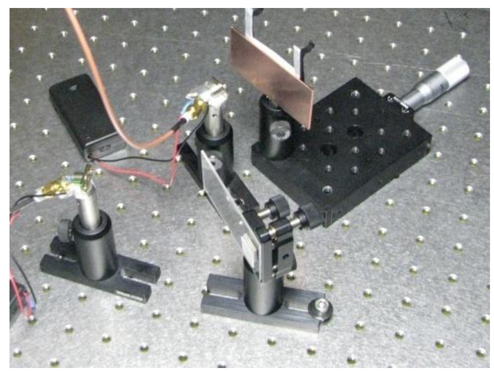

Radial Falloff

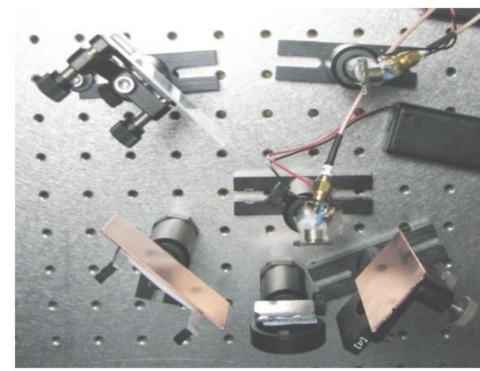

Multi-path bounces

Synchronization

See video for basic experimental verification

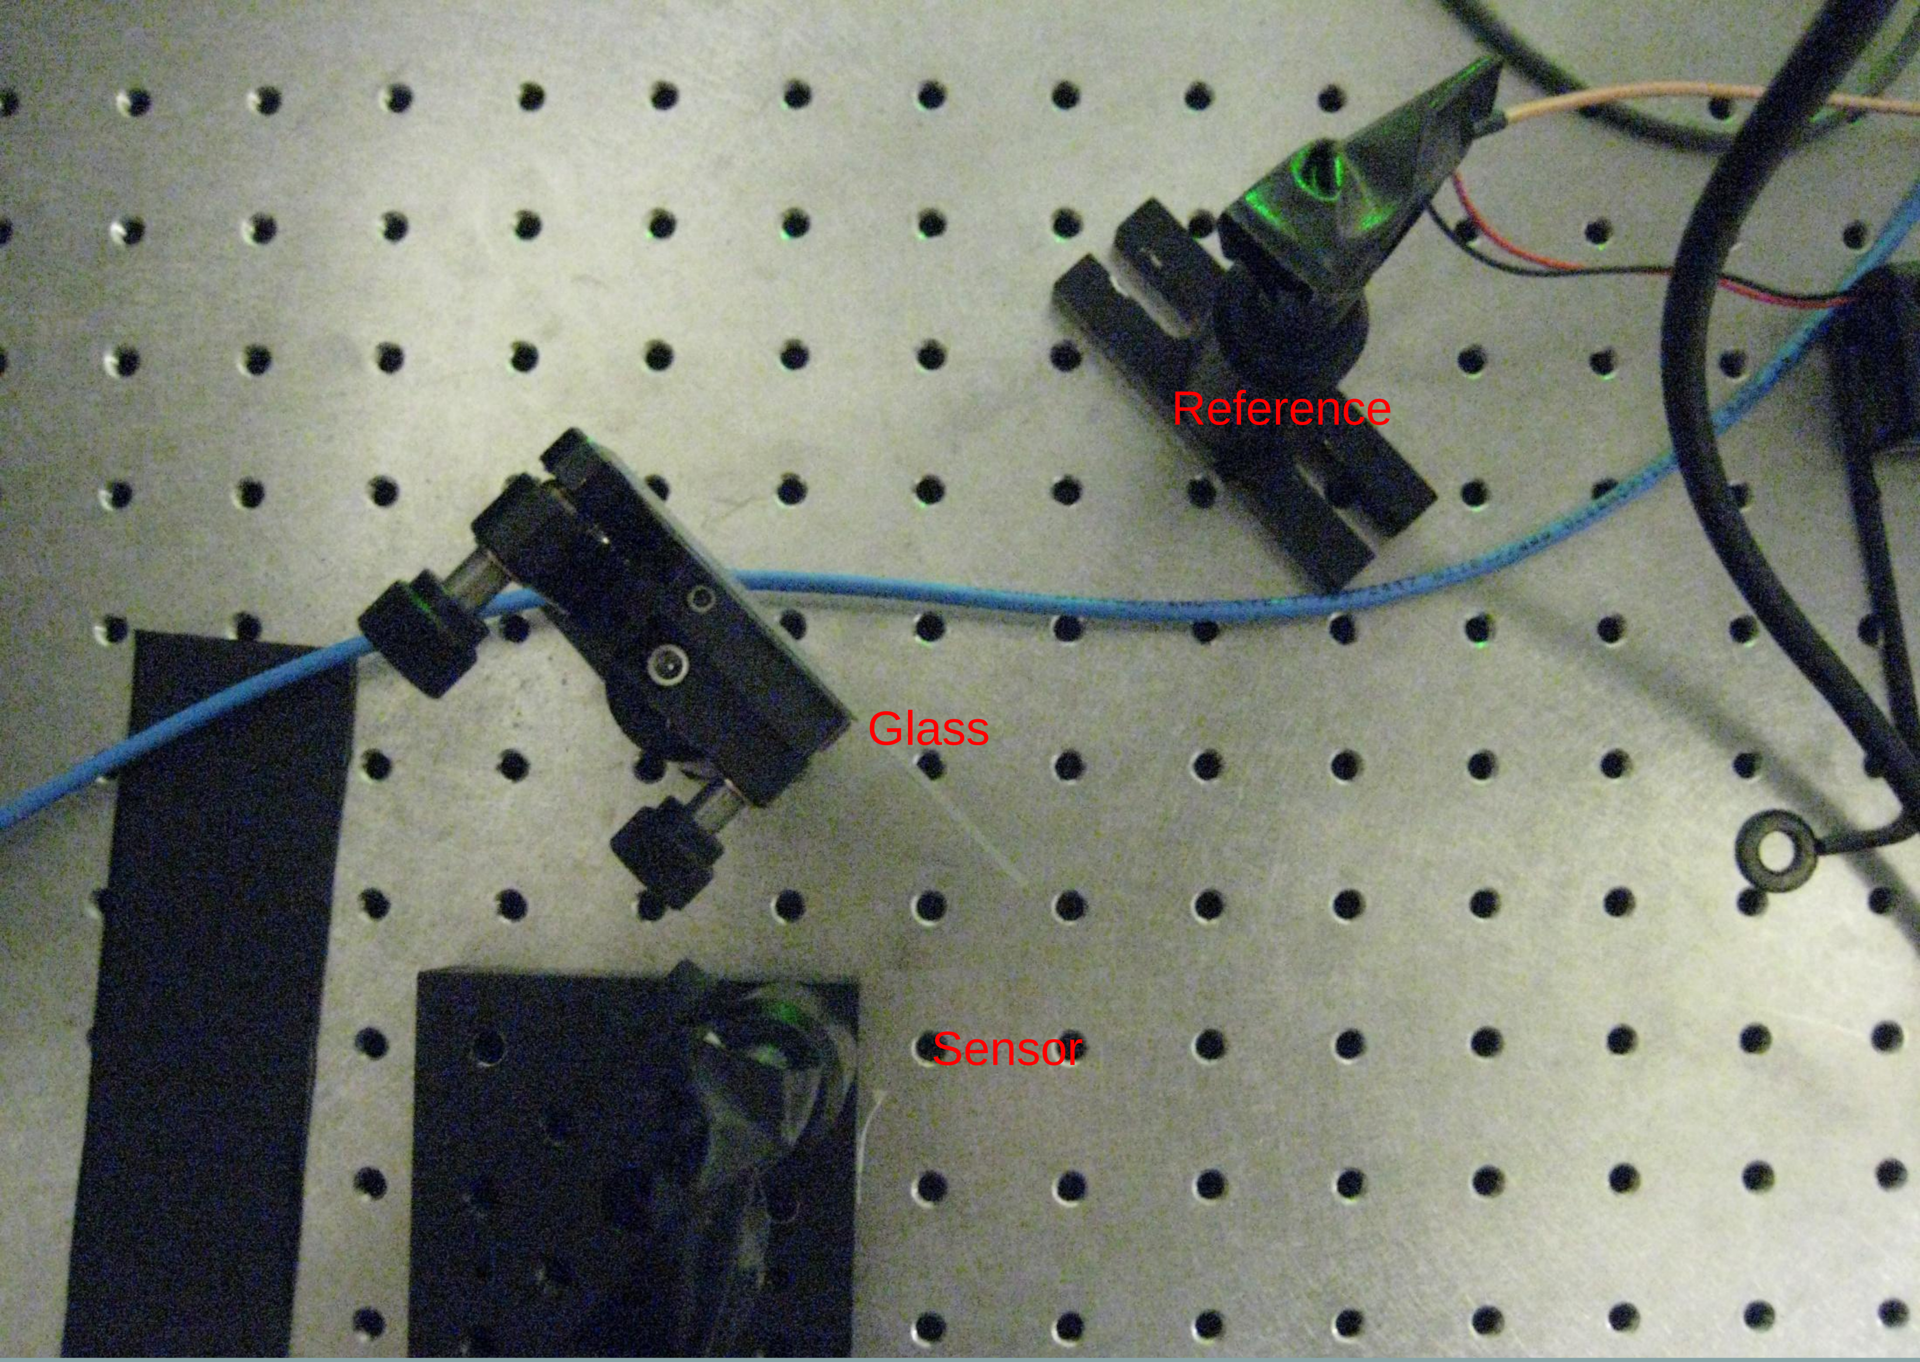

Here we show the setup for triggering the system synchronization using a reference sensor.

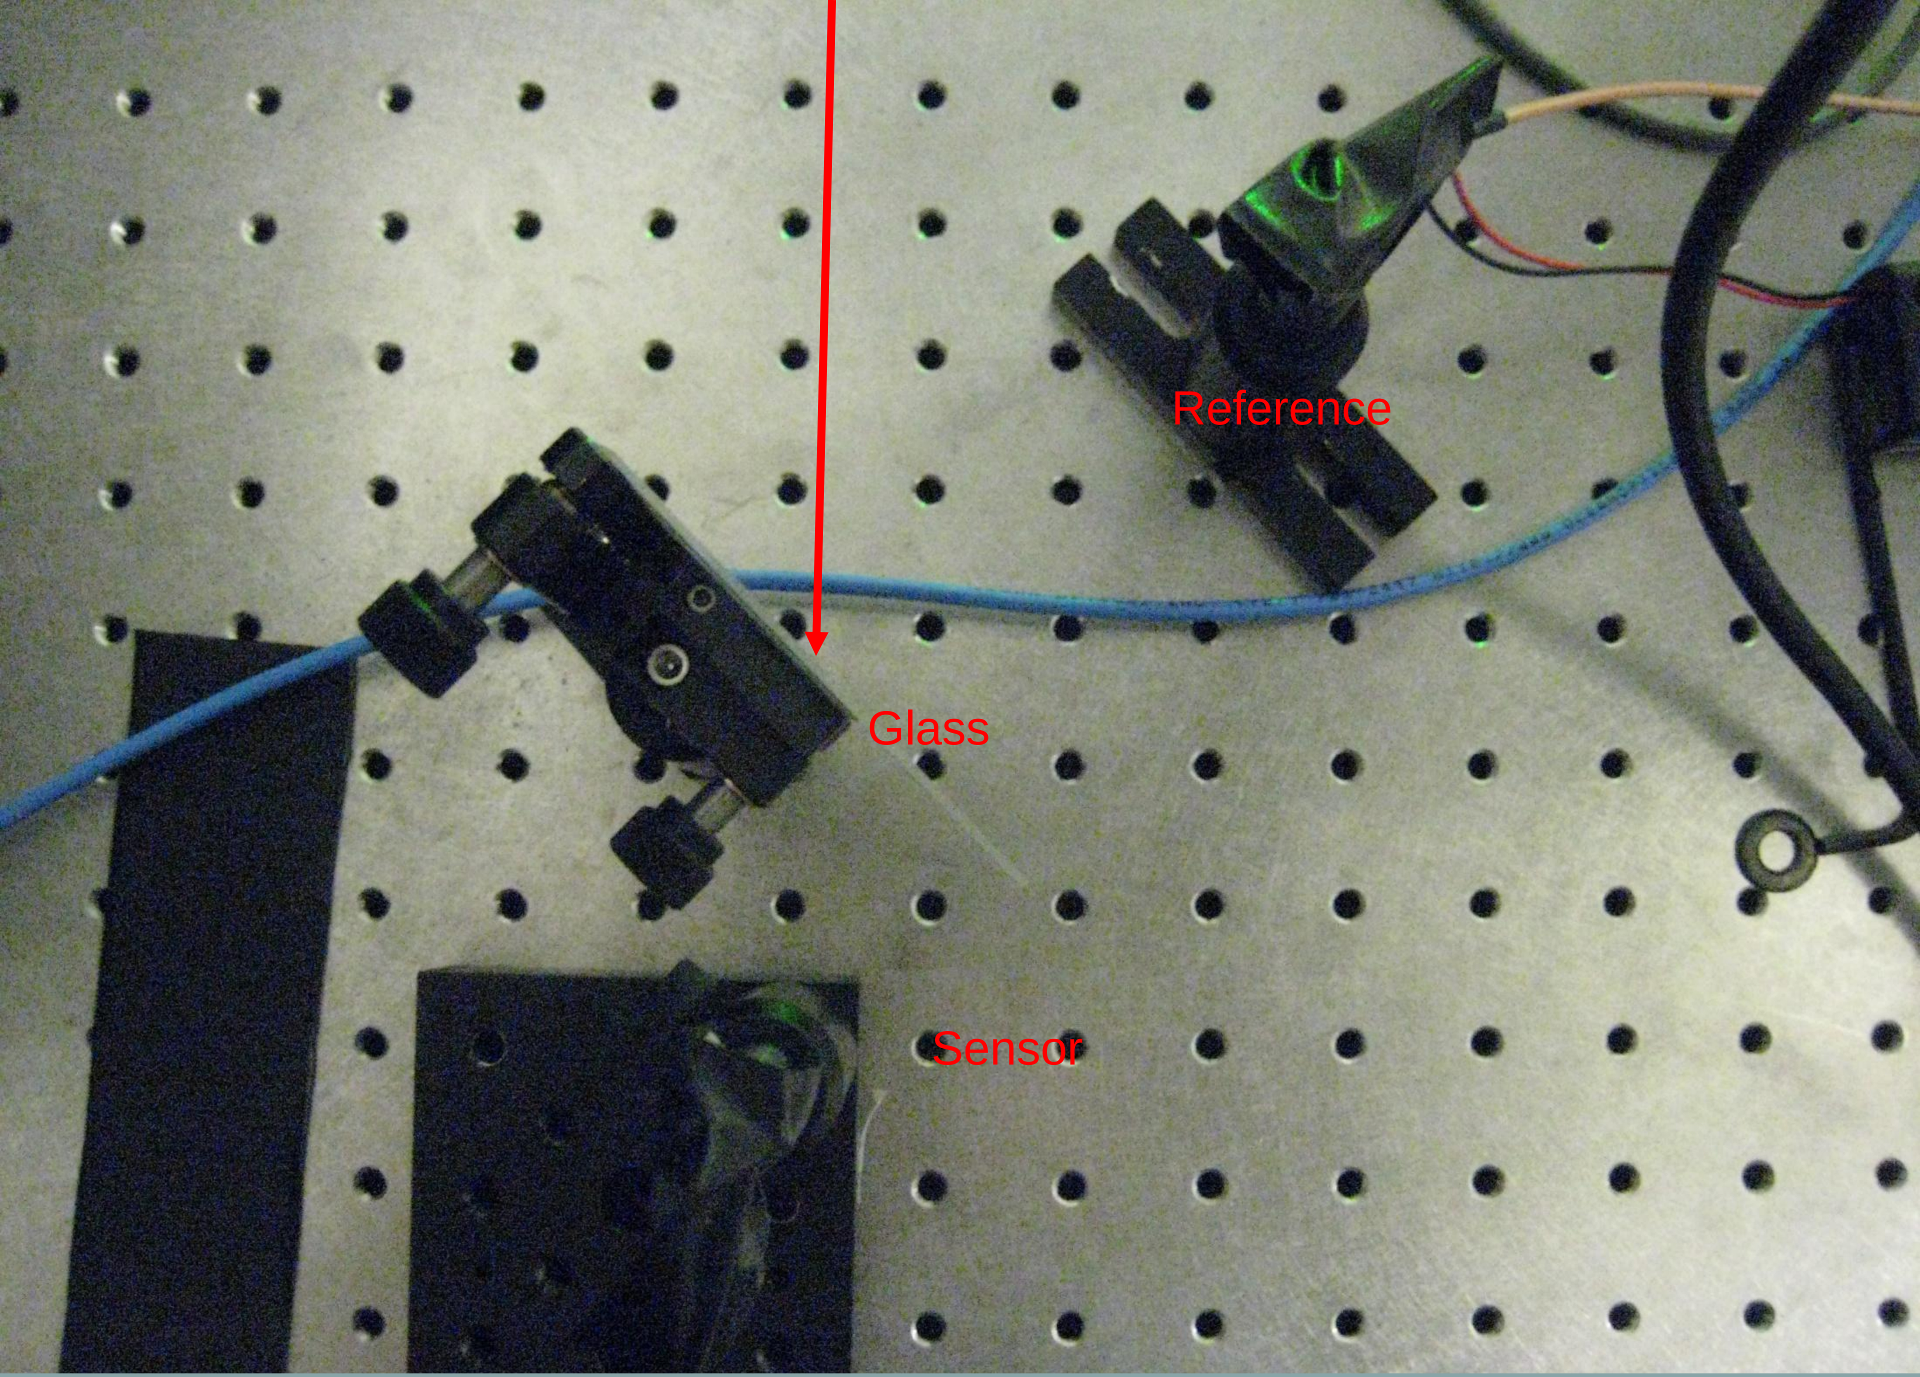

The laser enters the scene and strikes a glass surface.

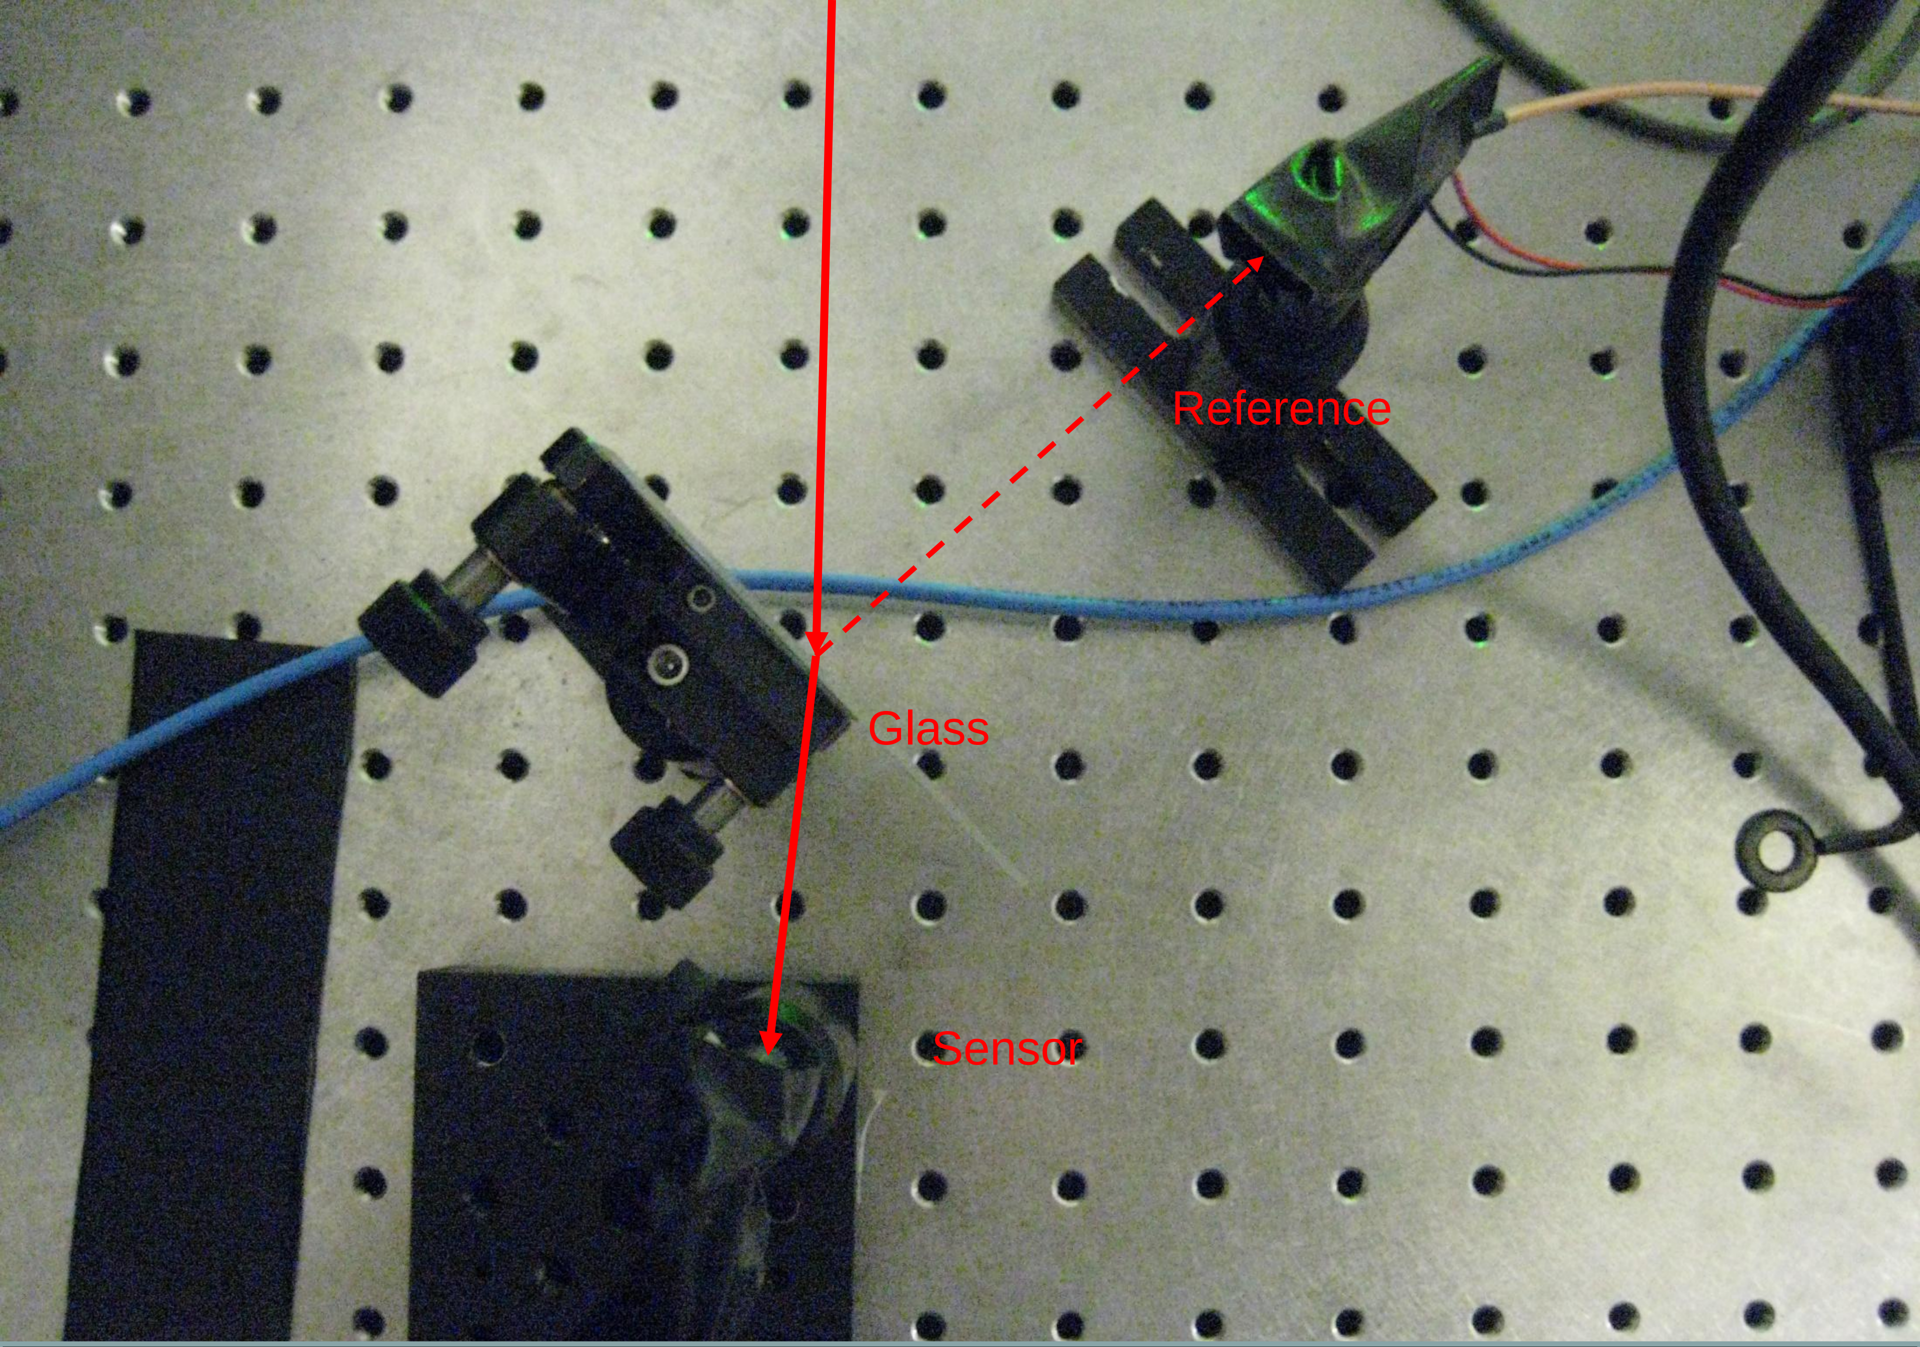

The glass reflects a small amount of light to the reference sensor, but the majority of the laser's energy falls on the main sensor. After capturing this data point, the sensor was moved .254 cm on a linear track.,

Synchronization

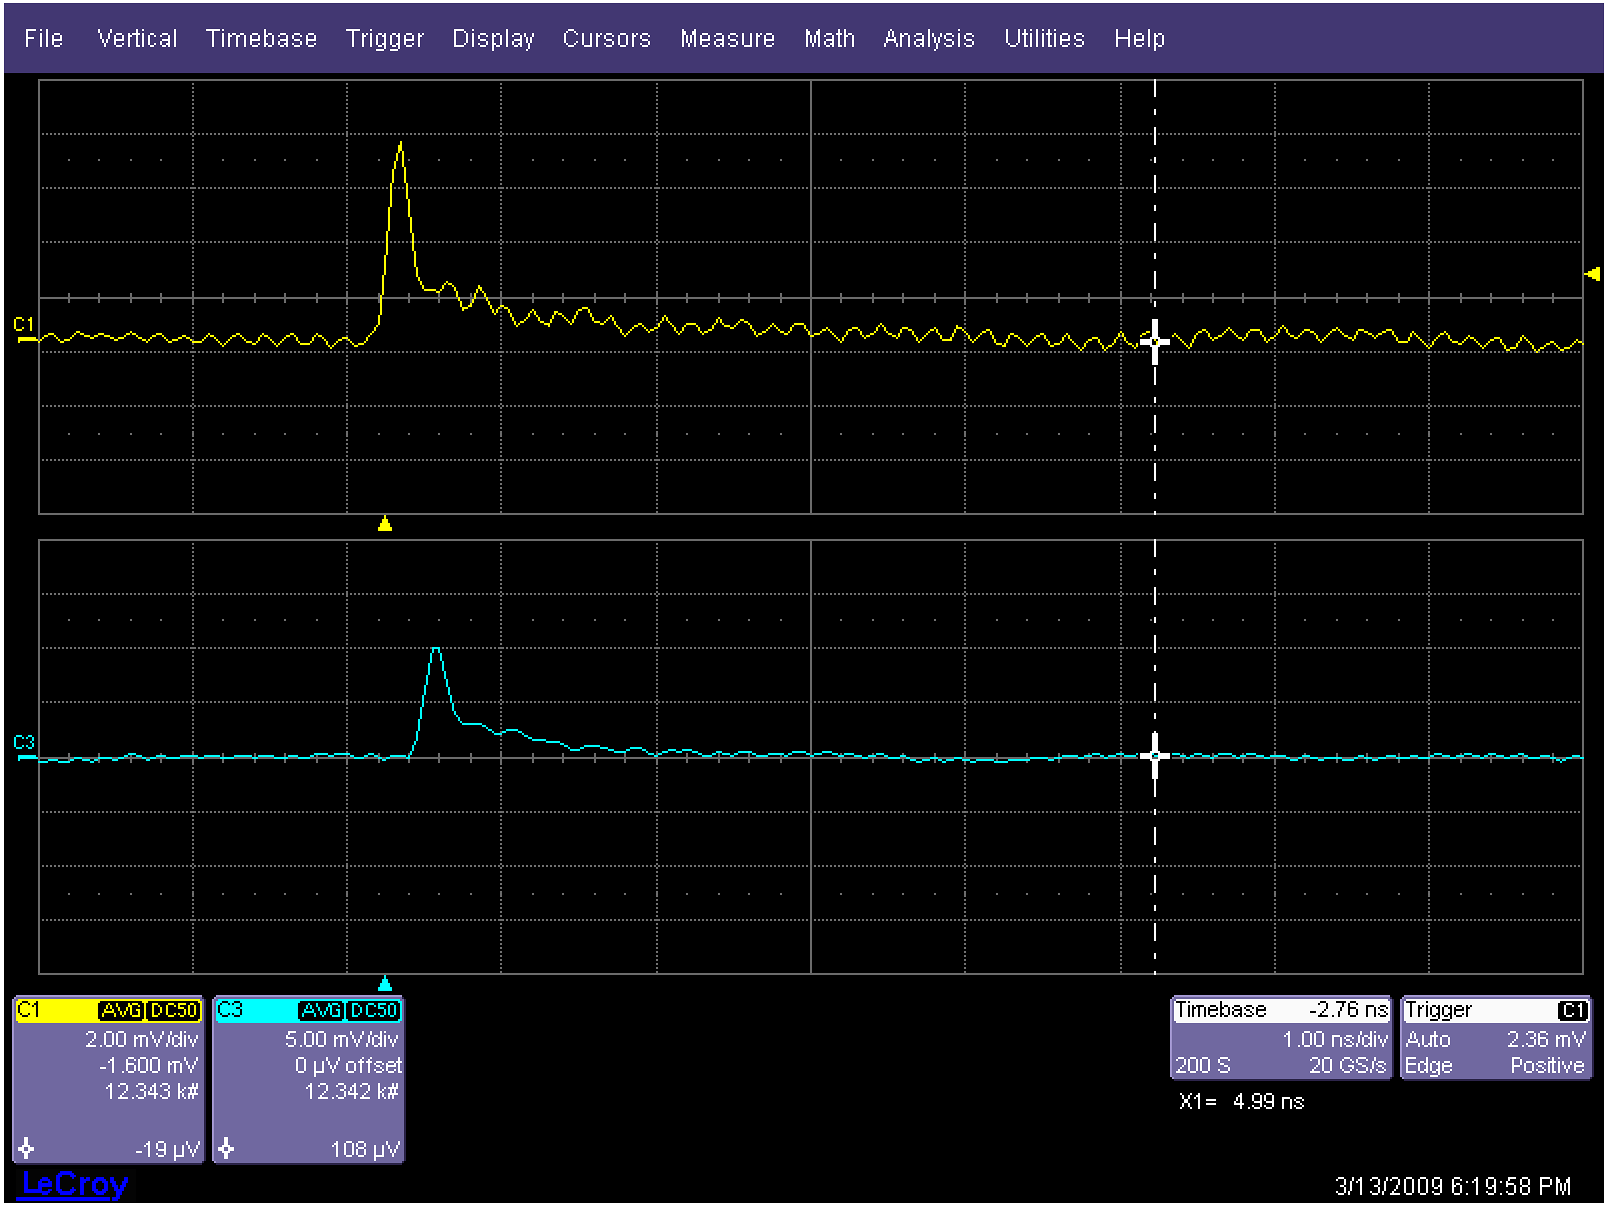

Screenshot from the Oscilloscope

Synchronization

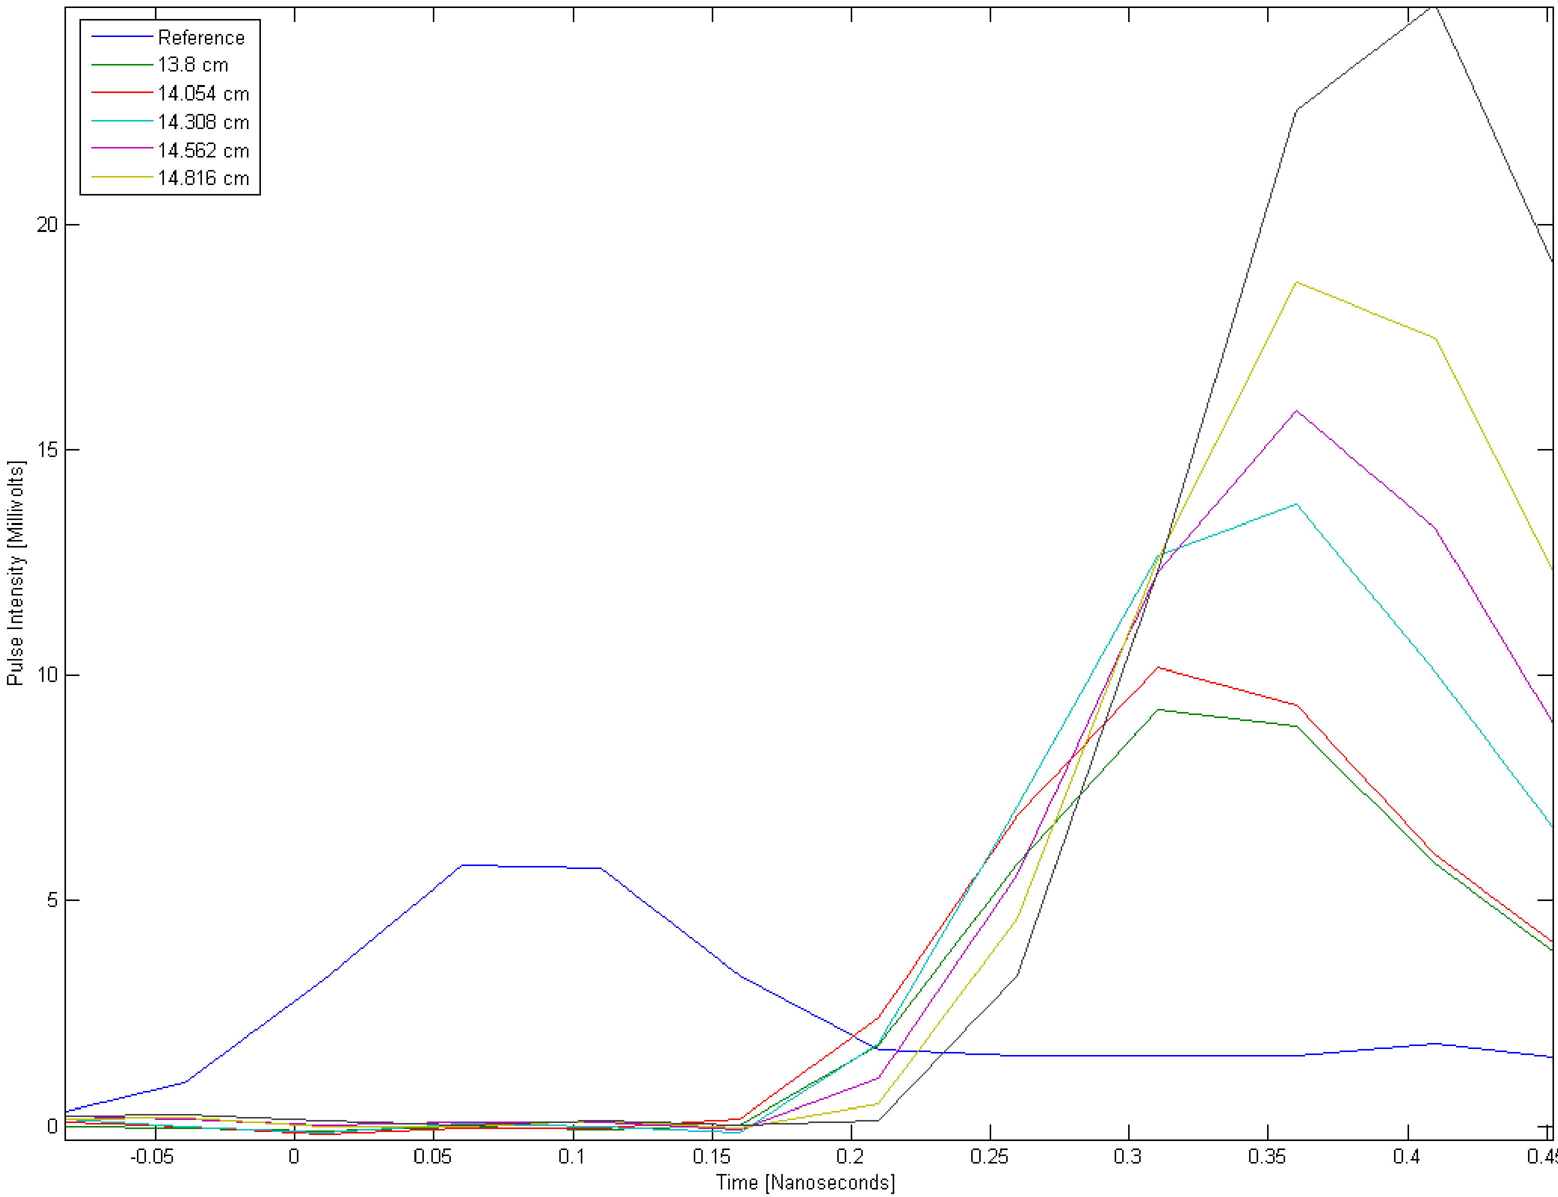

Time shifted pulses are observed as the sensor moves further from the reference trigger. Note that the

System Linearity

See video for basic experimental verification

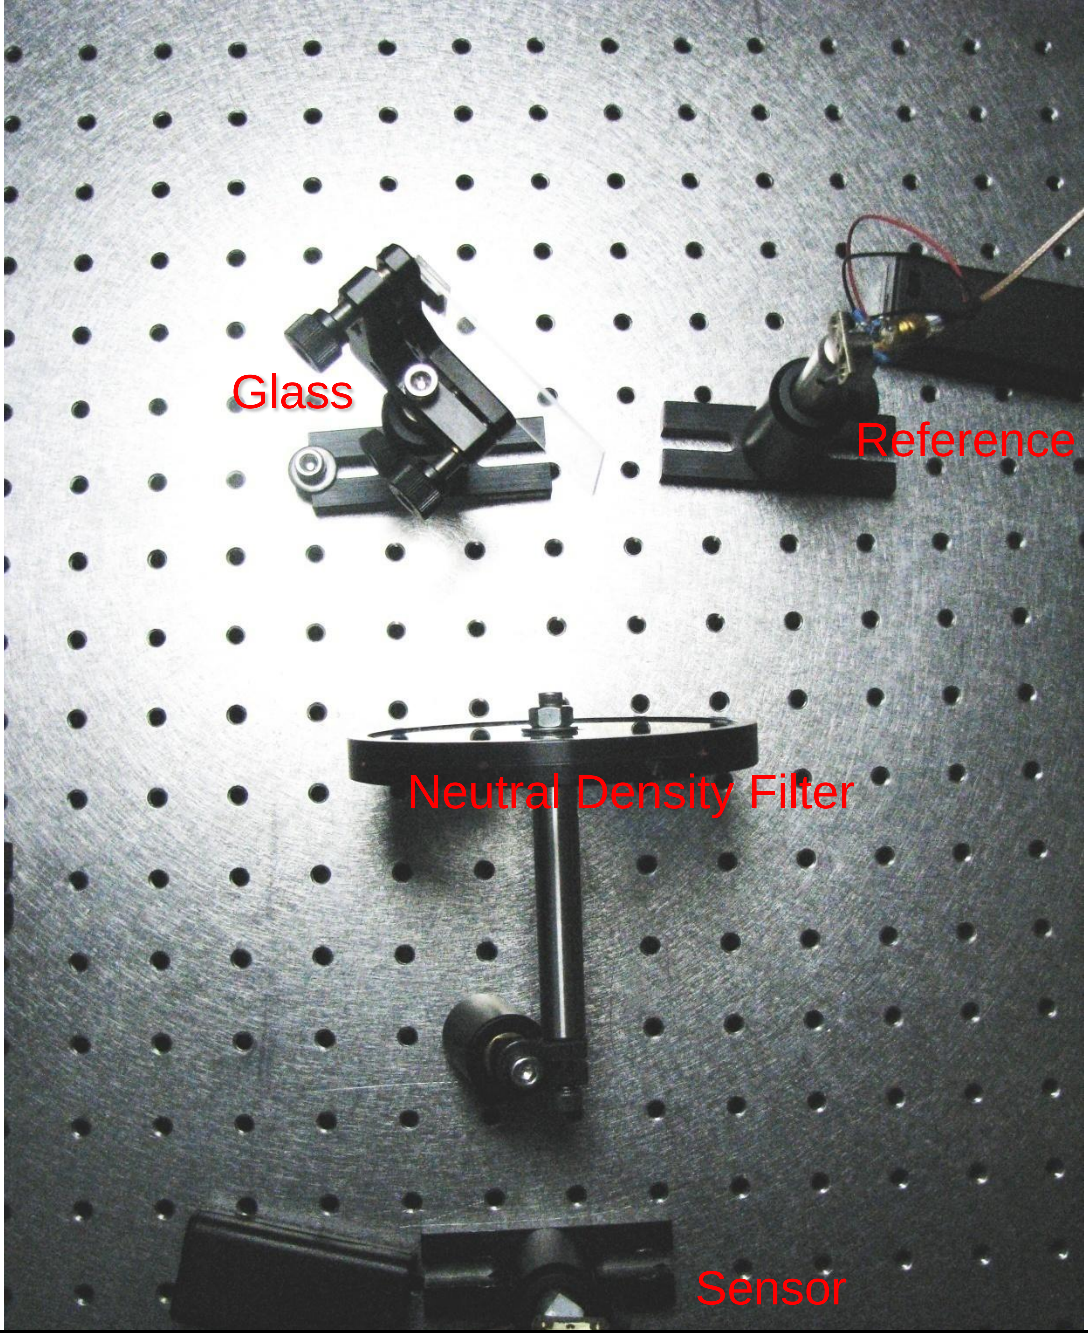

Here we show a setup used to demonstrate the linearity of the transient imaging camera system. The laser will be attenuated by a varying neutral density filter.

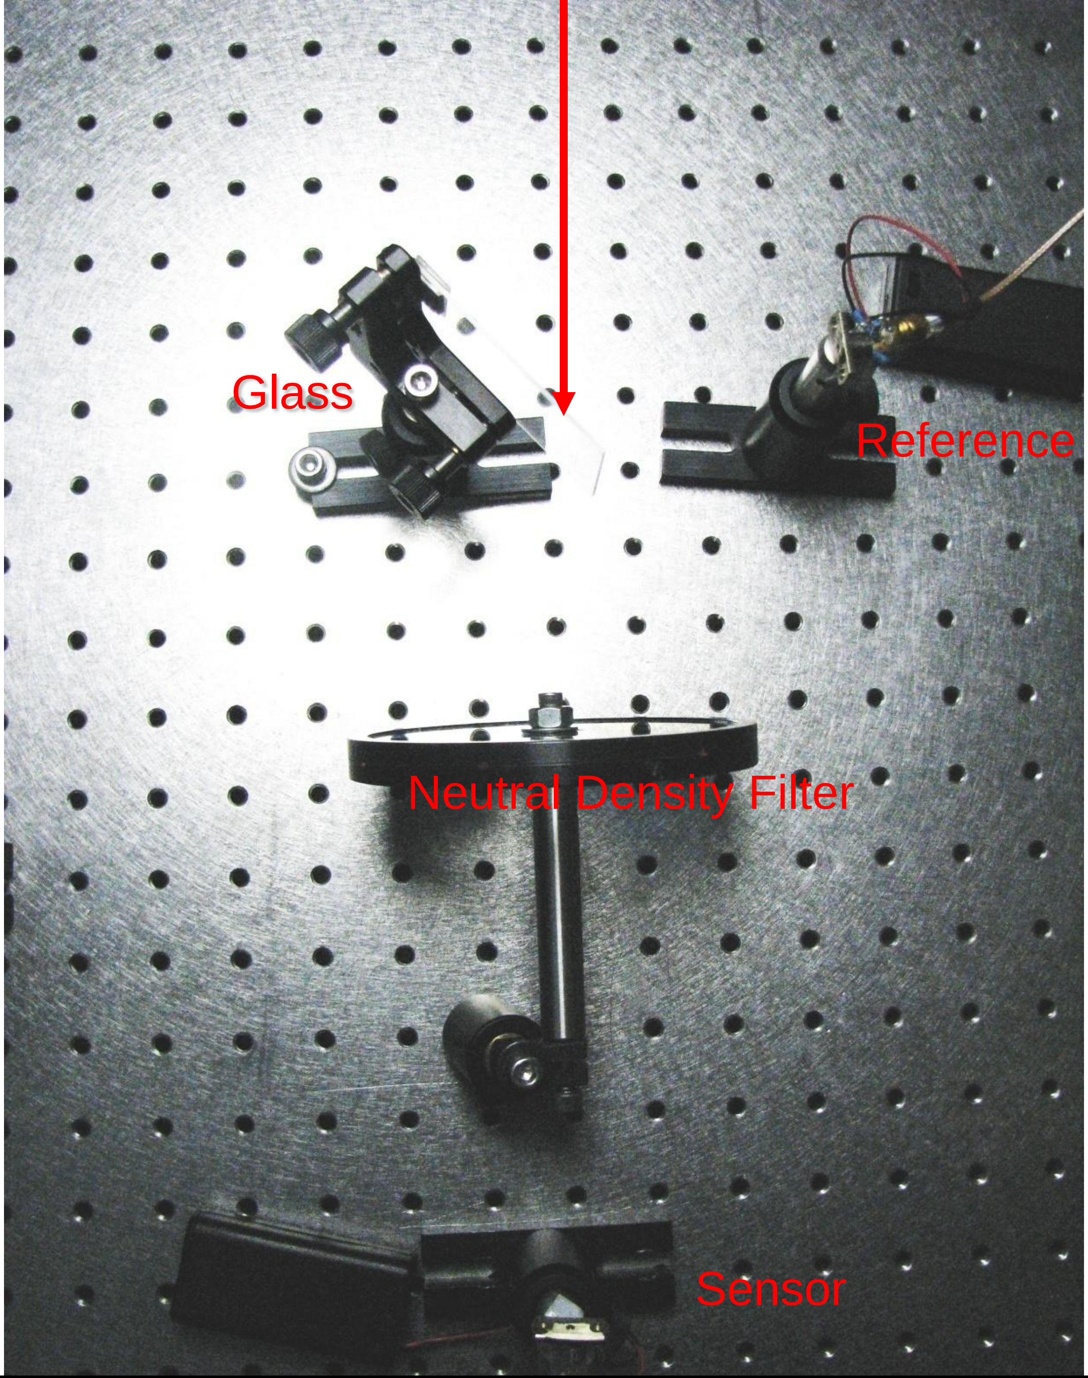

The laser enters the scene and strikes a glass surface.

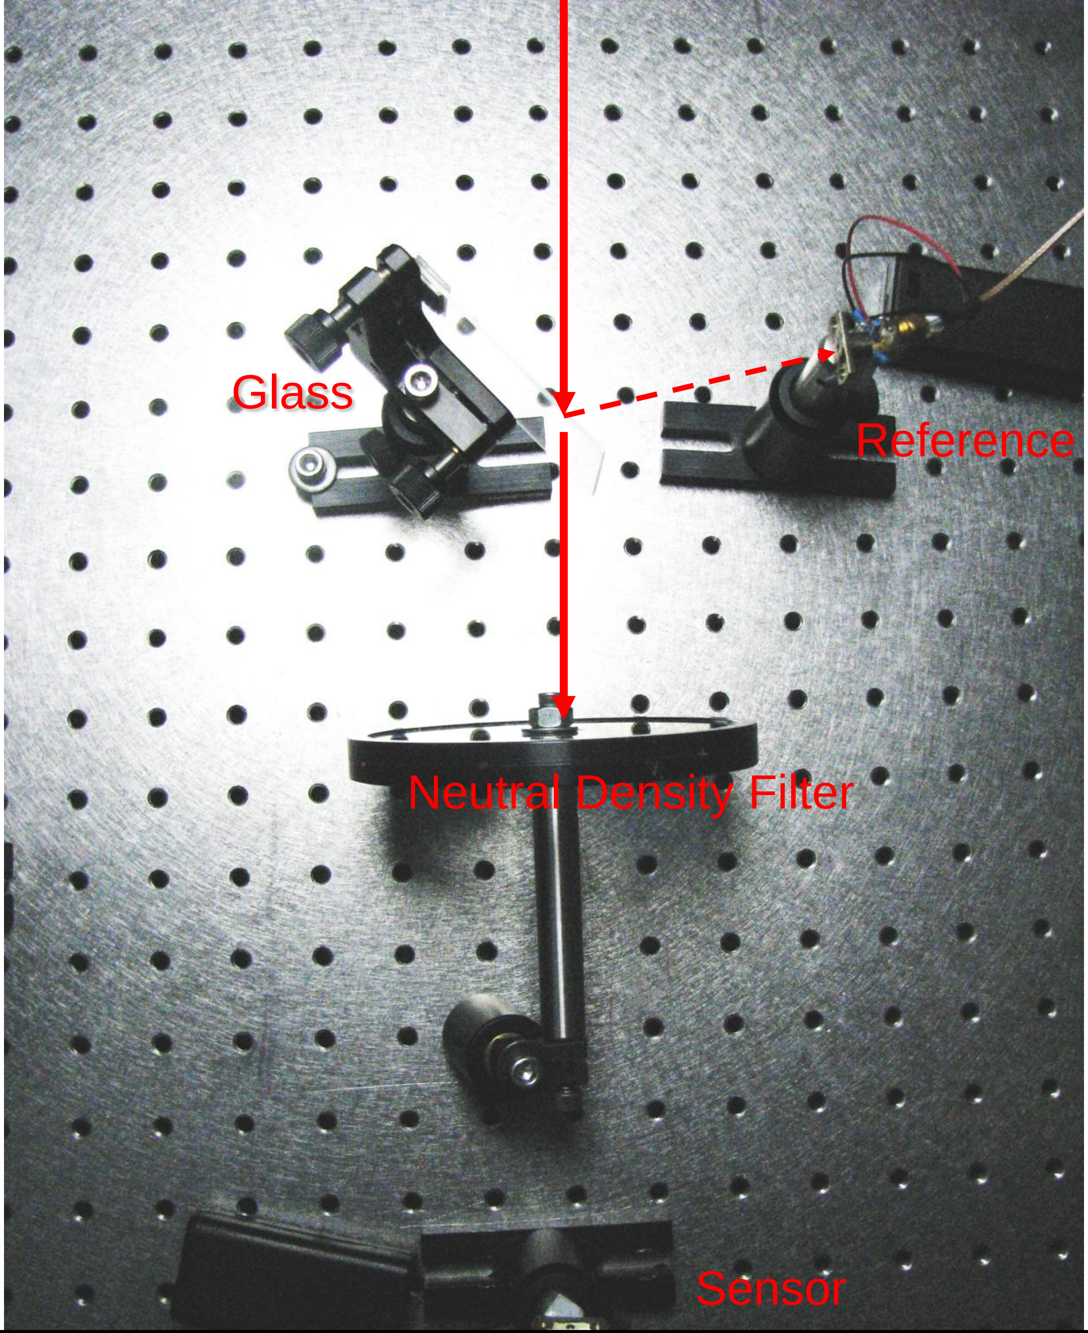

The glass reflects a small amount of the laser’s energy to the reference for oscilloscope triggering, but the majority of the energy continues to the neutral density filter.

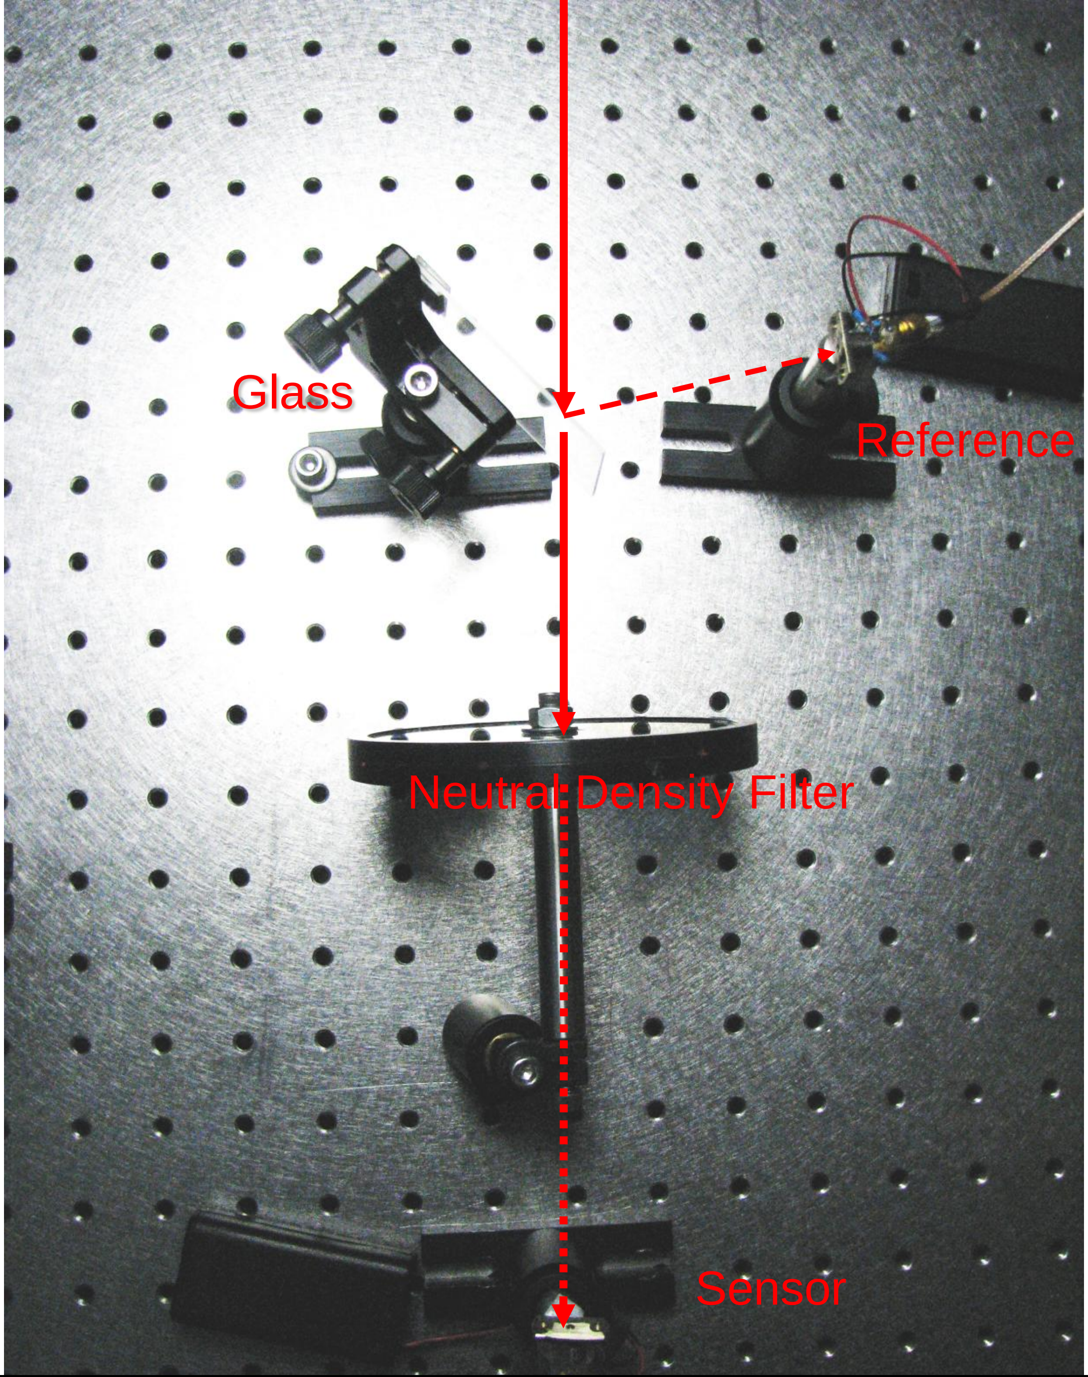

The light is attenuated by the neutral density filter and strikes the main sensor.

System Linearity

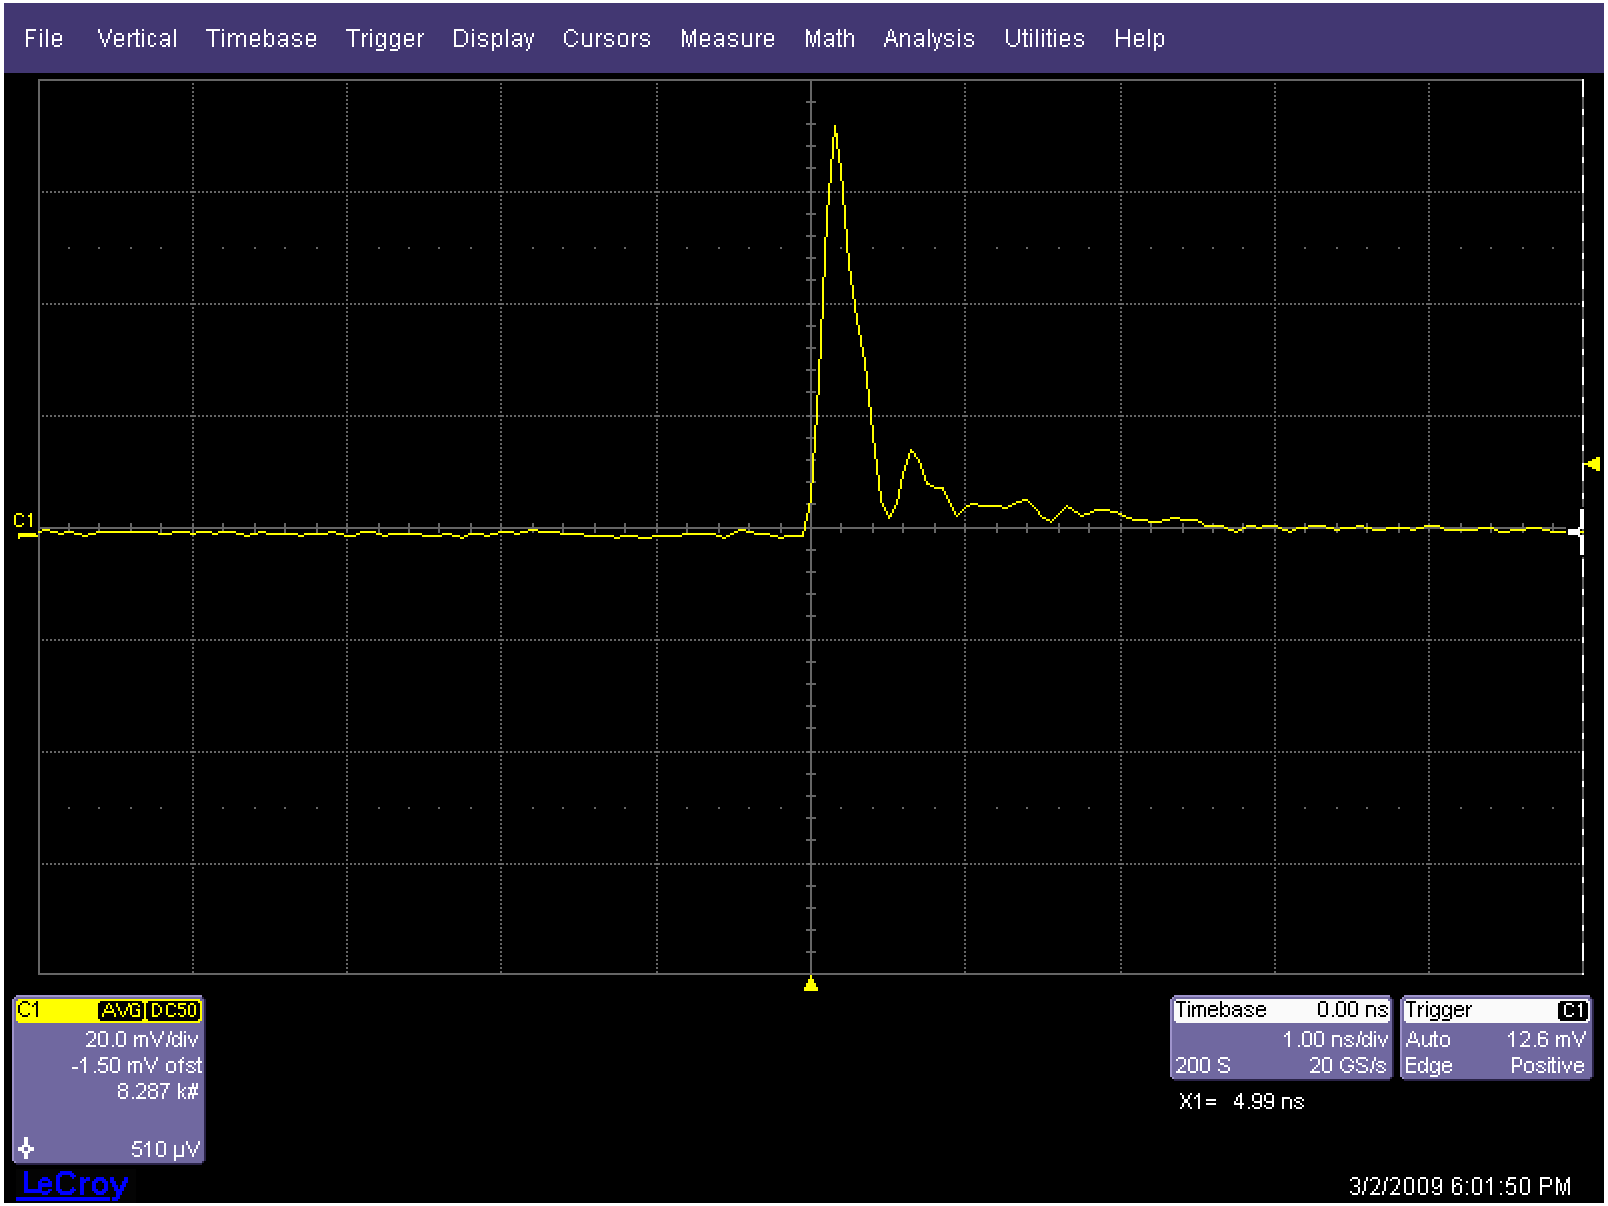

Screenshot from the oscilloscope of the least attenuated signal

System Linearity

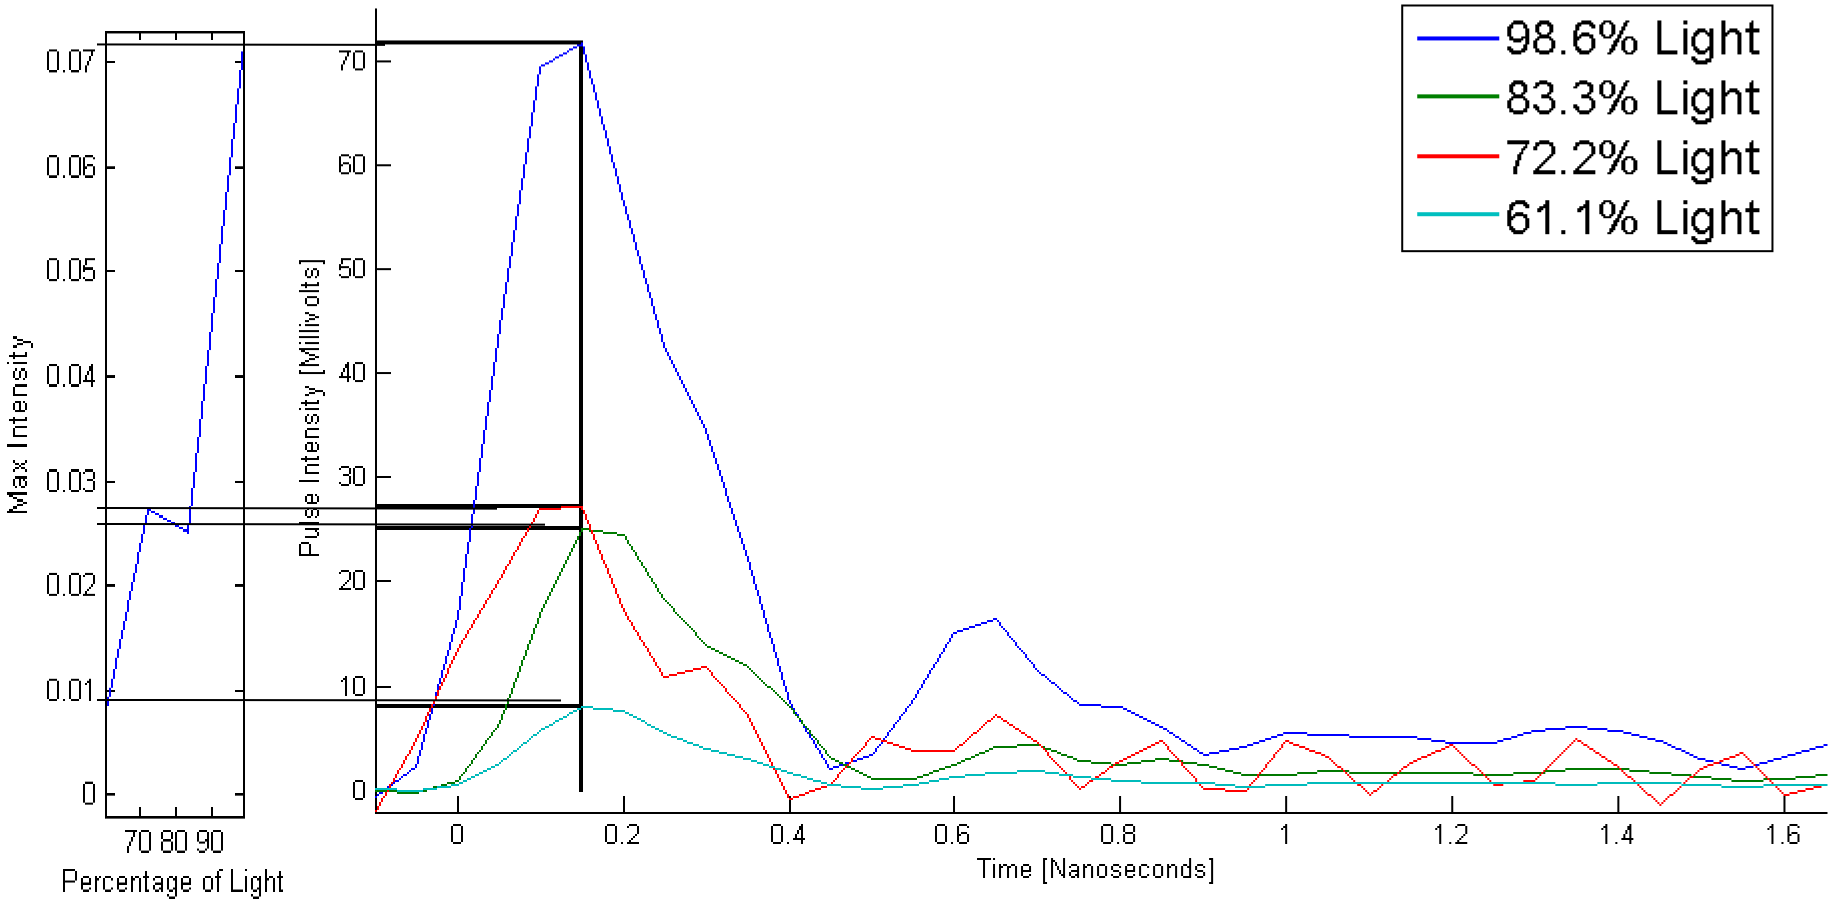

By varying the neutral density filter, we observed a linear system response.

Radial Falloff of Diffuse Light

See video for basic experimental verification

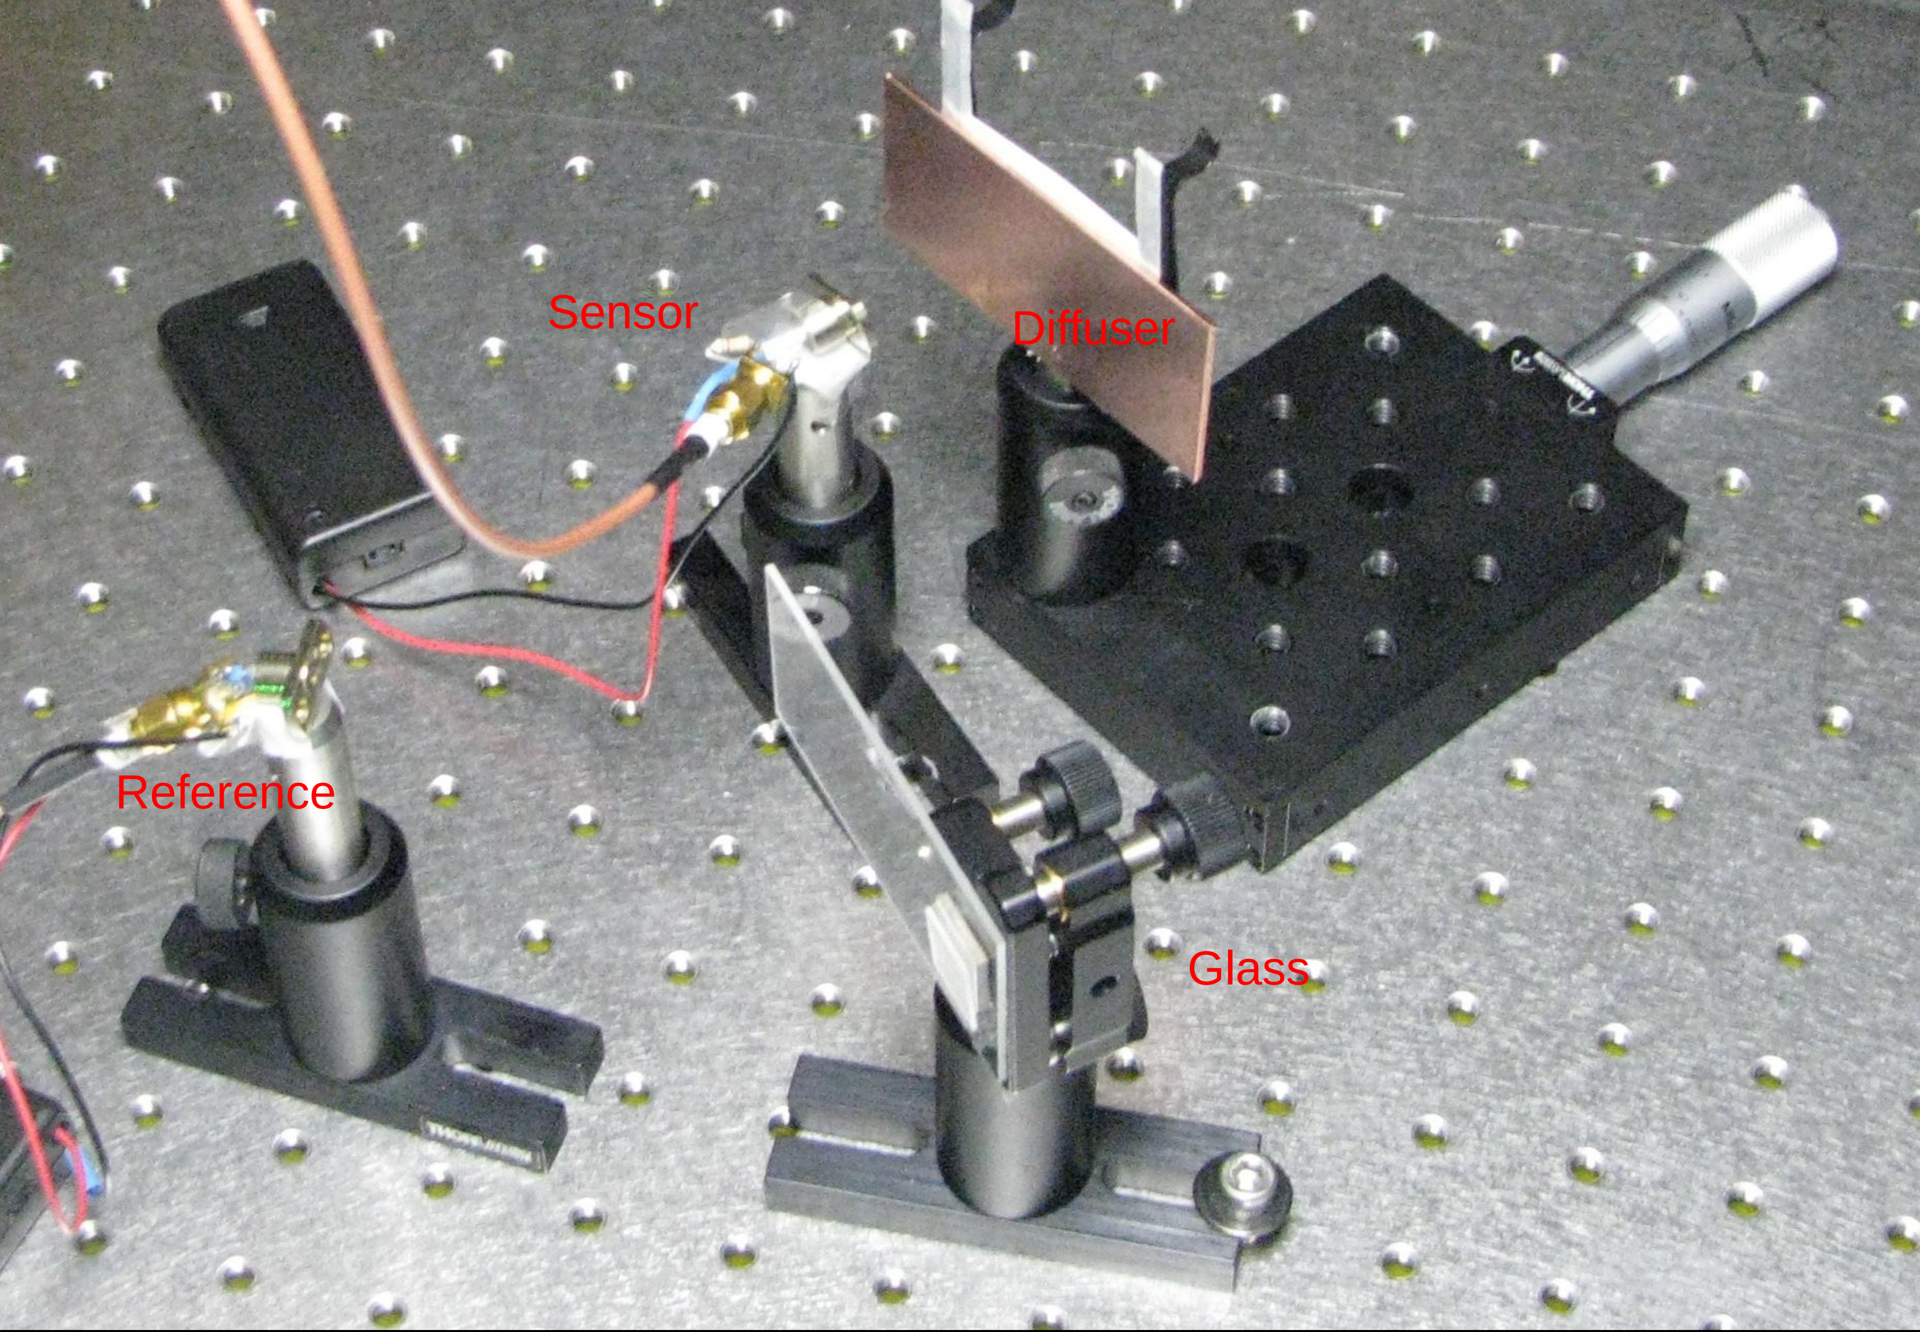

Here we show the experimental setup used to capture the first bounce from a diffuser to display the inverse square falloff pattern.

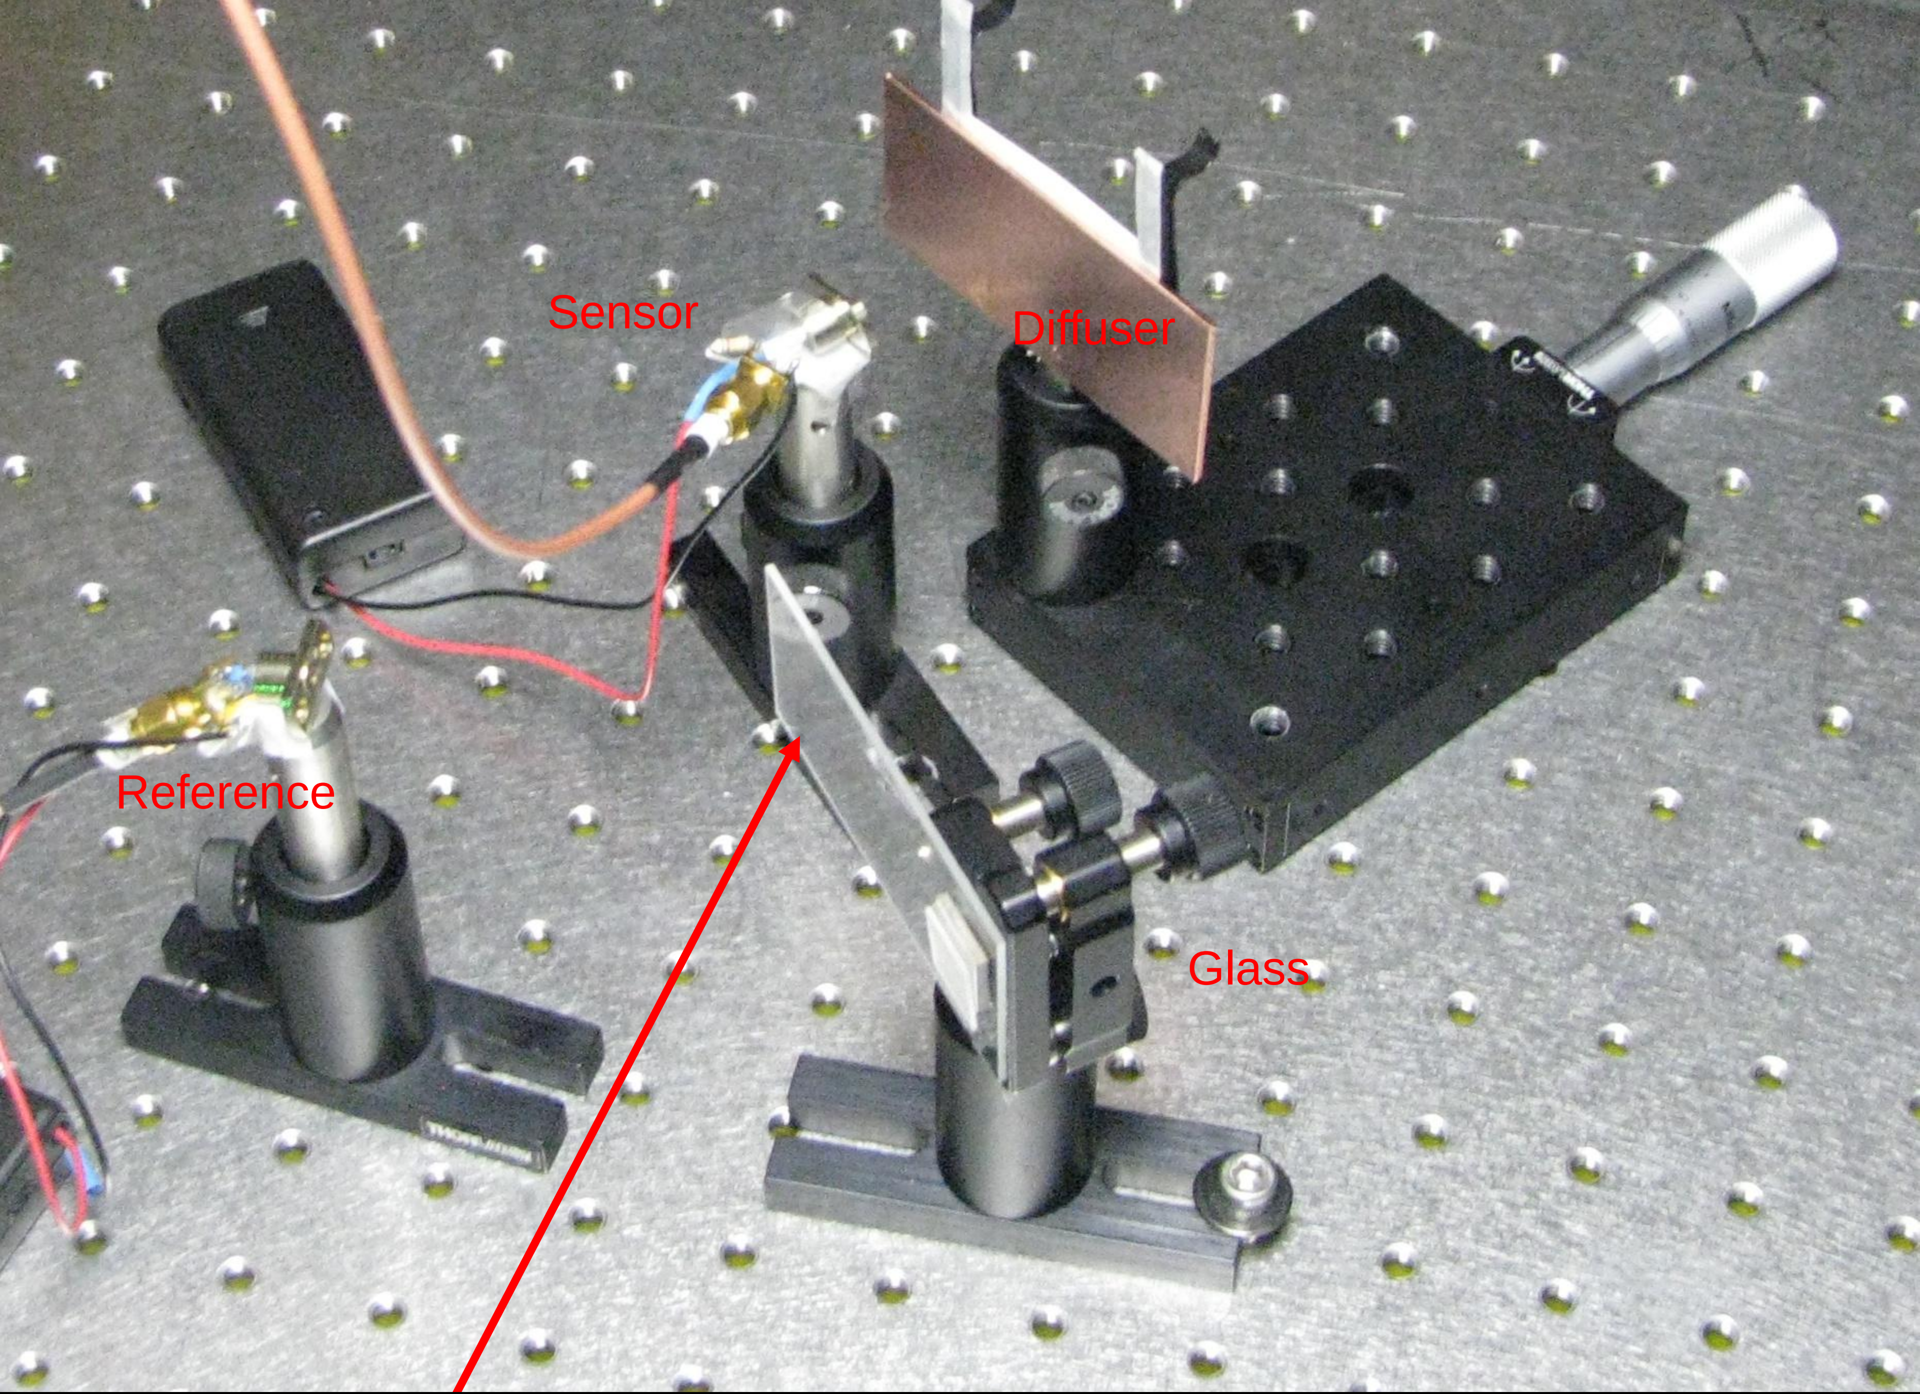

Light enters the scene and strikes a glass surface.

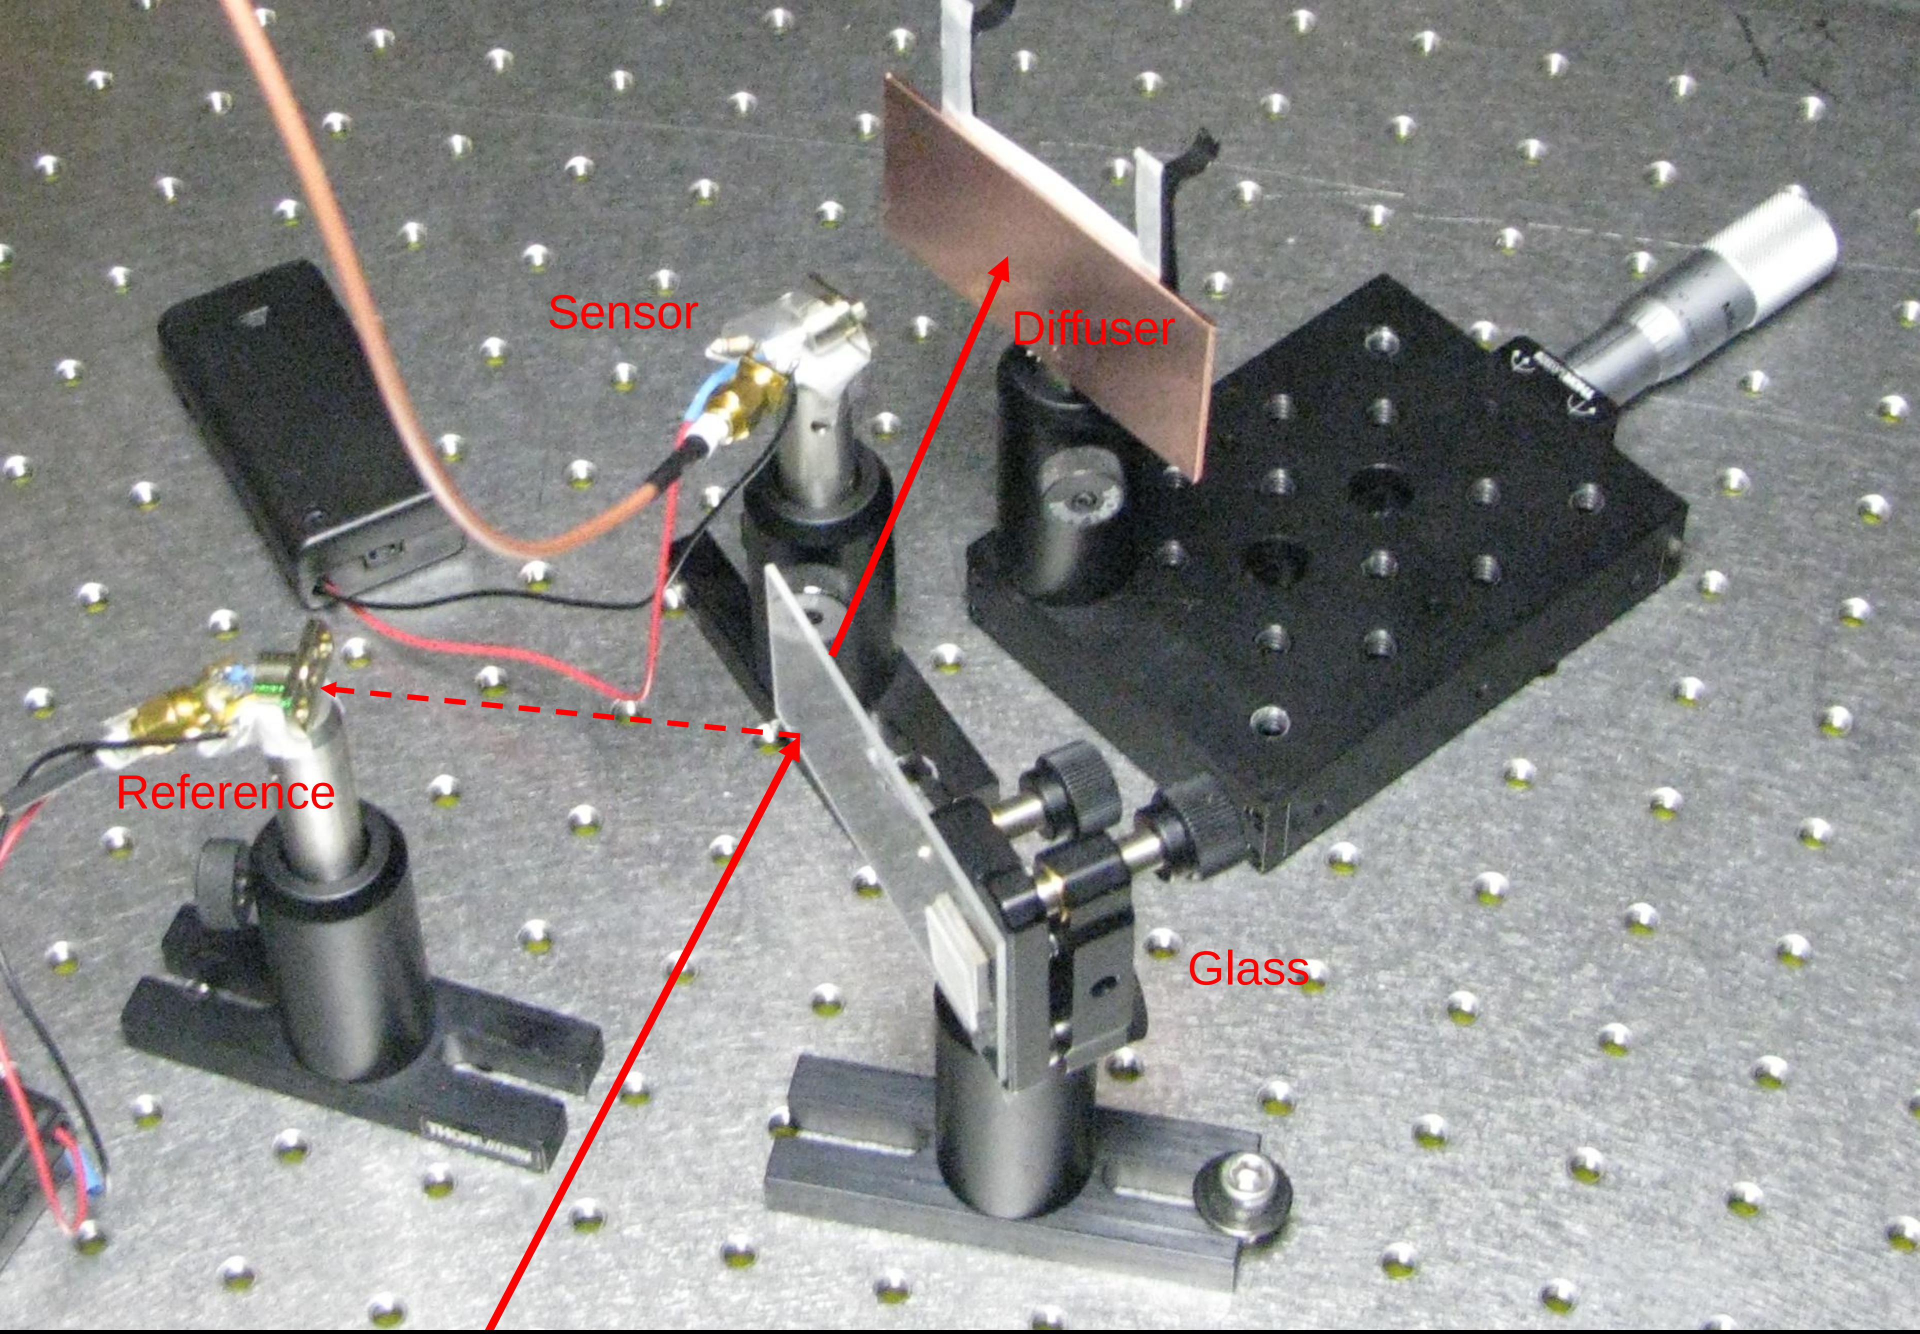

The light reflects off the glass surface and travels back to the reference which triggers the oscilloscope. The main energy in the laser continues on to the diffuser.

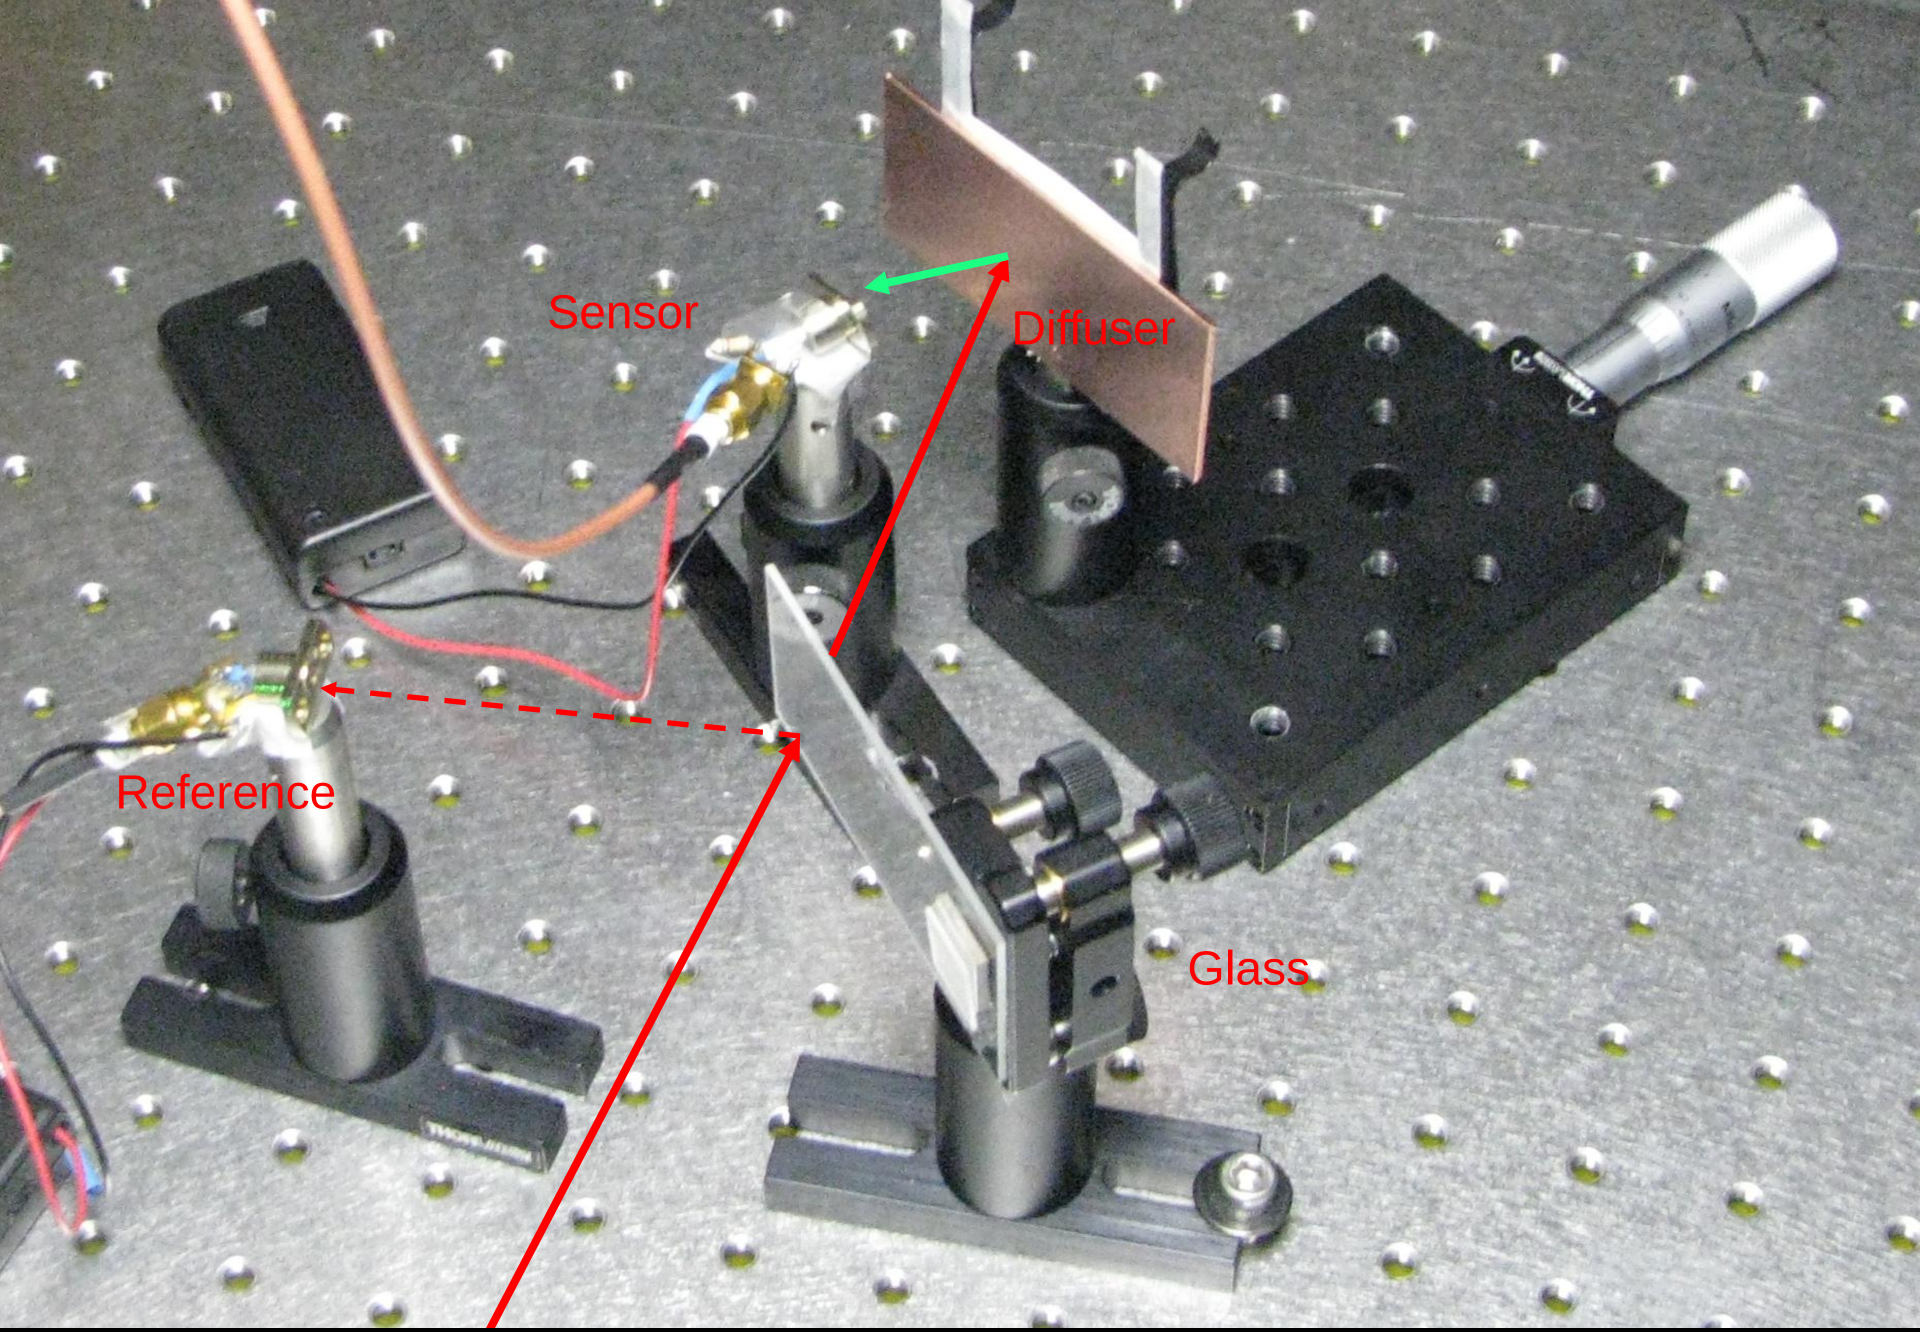

The laser strikes the diffuse surface and travels to the main sensor. The diffuser is mounted on a linear track which will move the diffuser further away from the sensor.

Radial Falloff

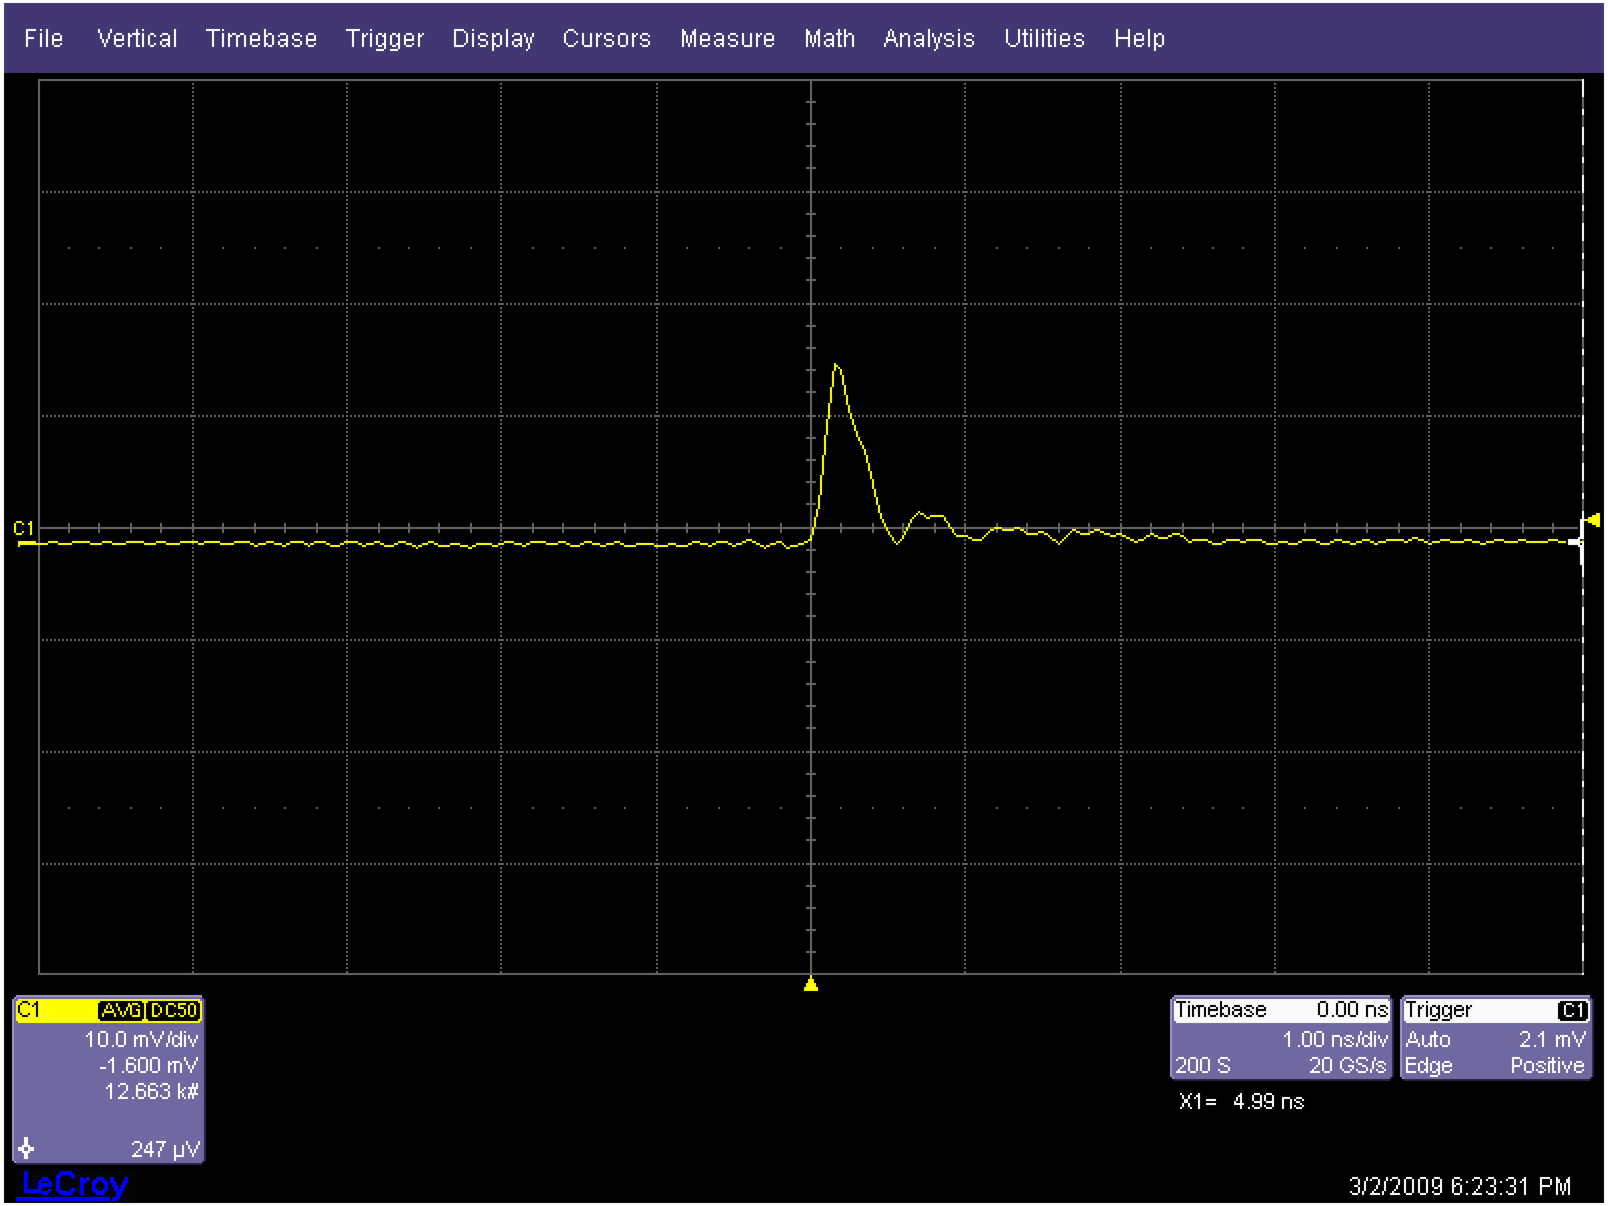

Screenshot of the signal on the oscilloscope when the diffuser is at 4.824 cm offset

Radial Falloff of Diffuse Light

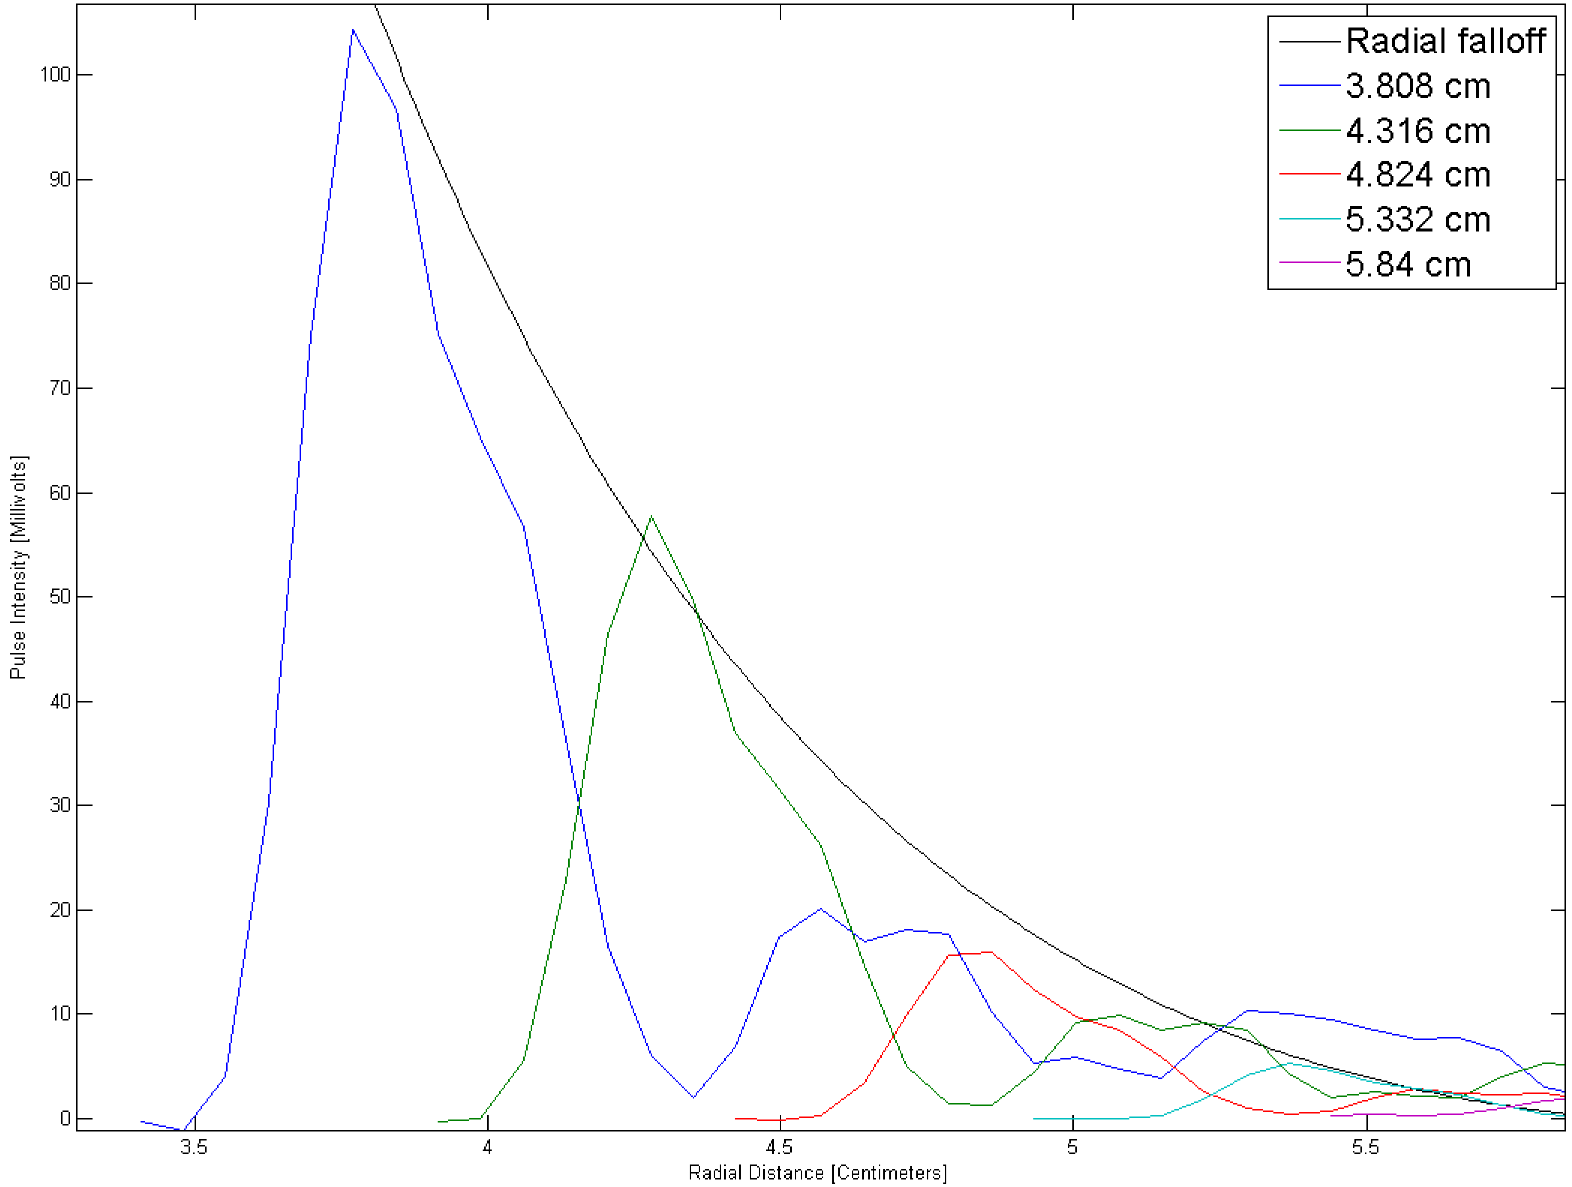

The light captured at different radial distances obeys an inverse radius squared relationship

Capturing Multiple Bounces

(Detecting third bounces)

See video for basic experimental verification

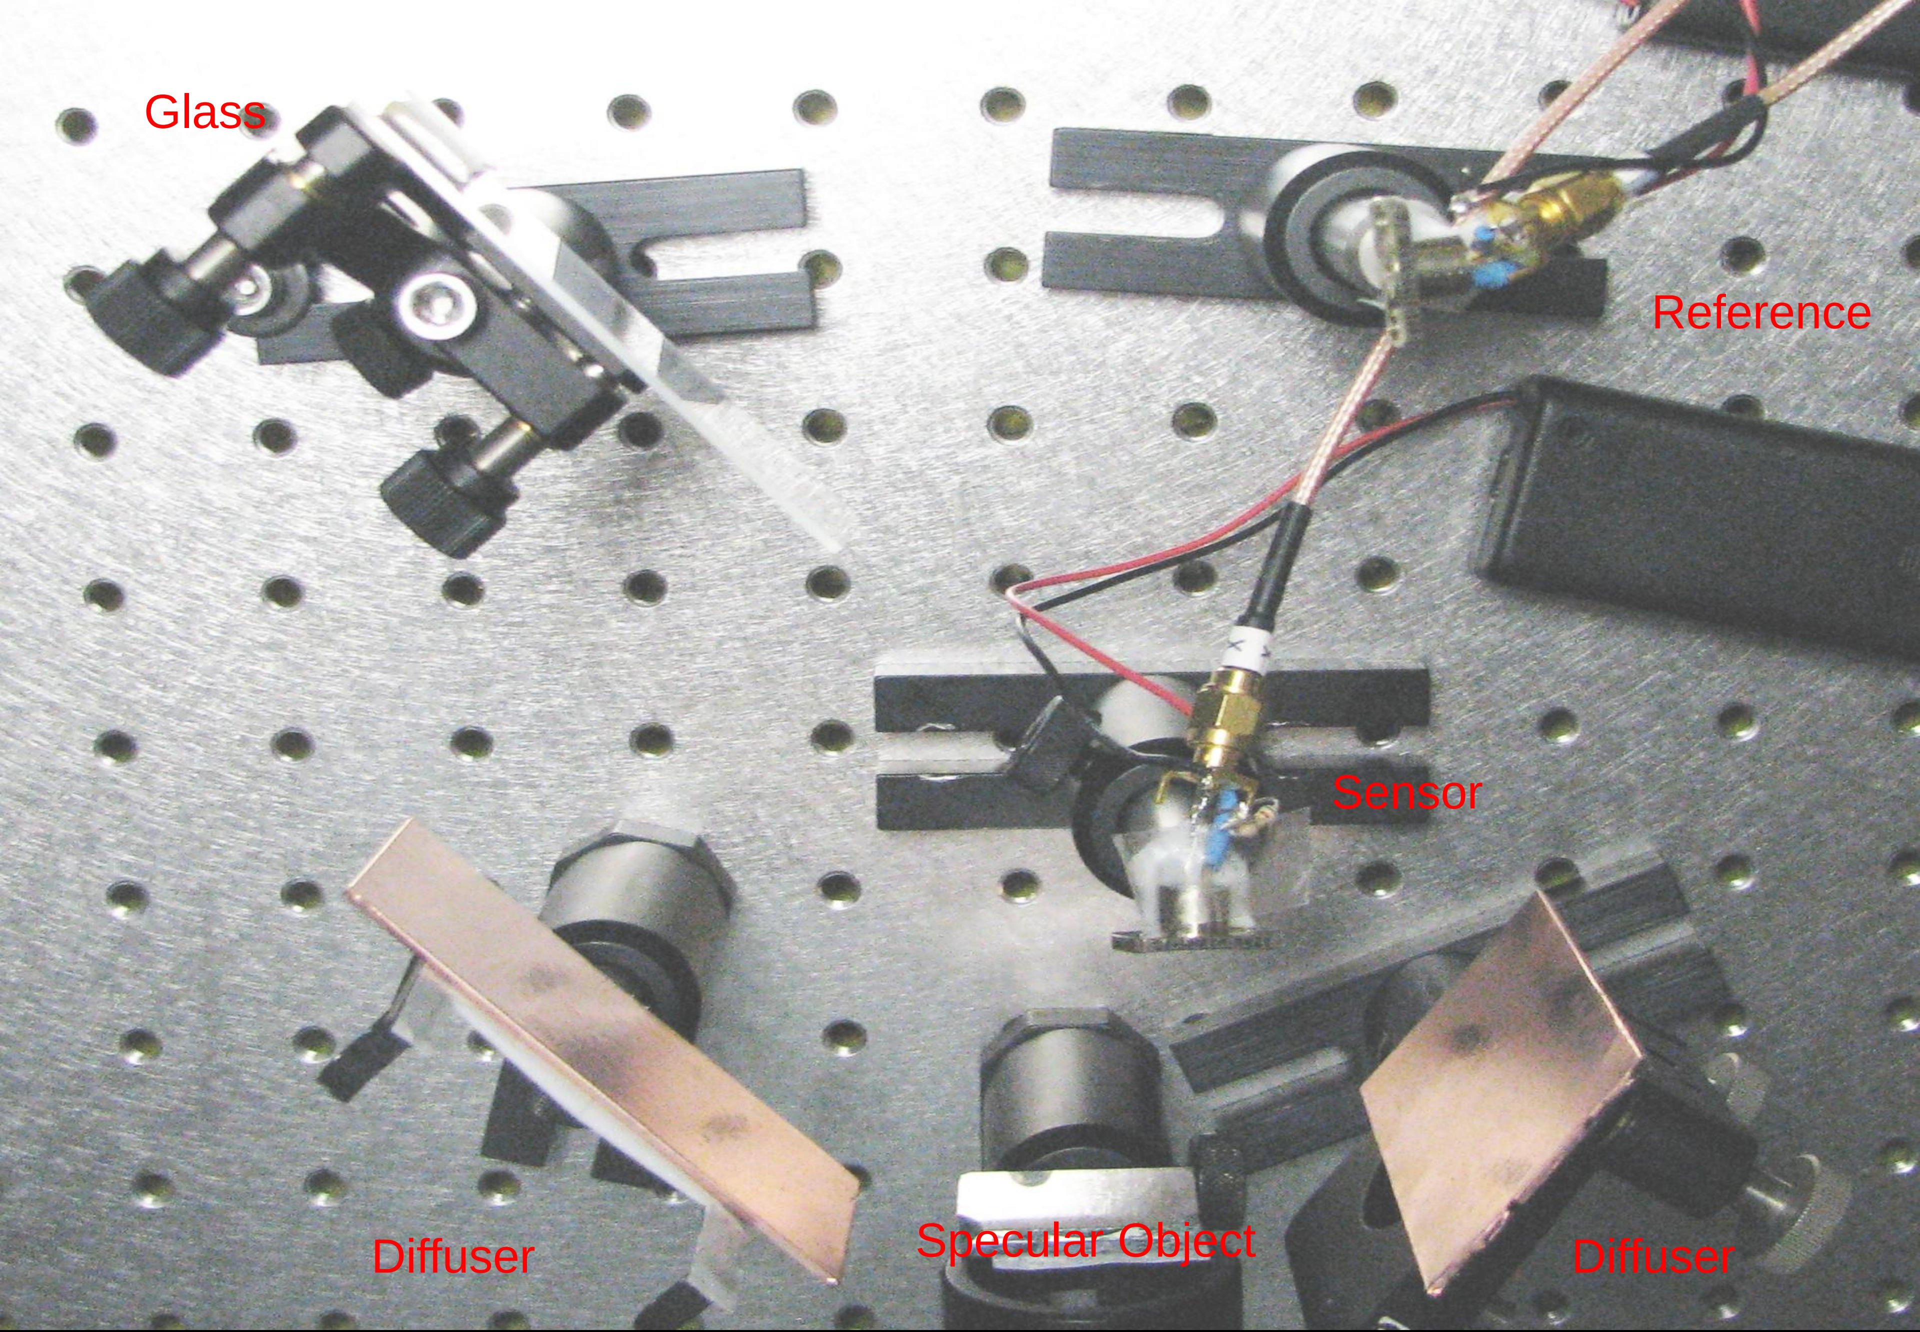

Here we show the experimental setup which shows that a signal will be discernable from the noise floor after the ray has reflected off of three patches.

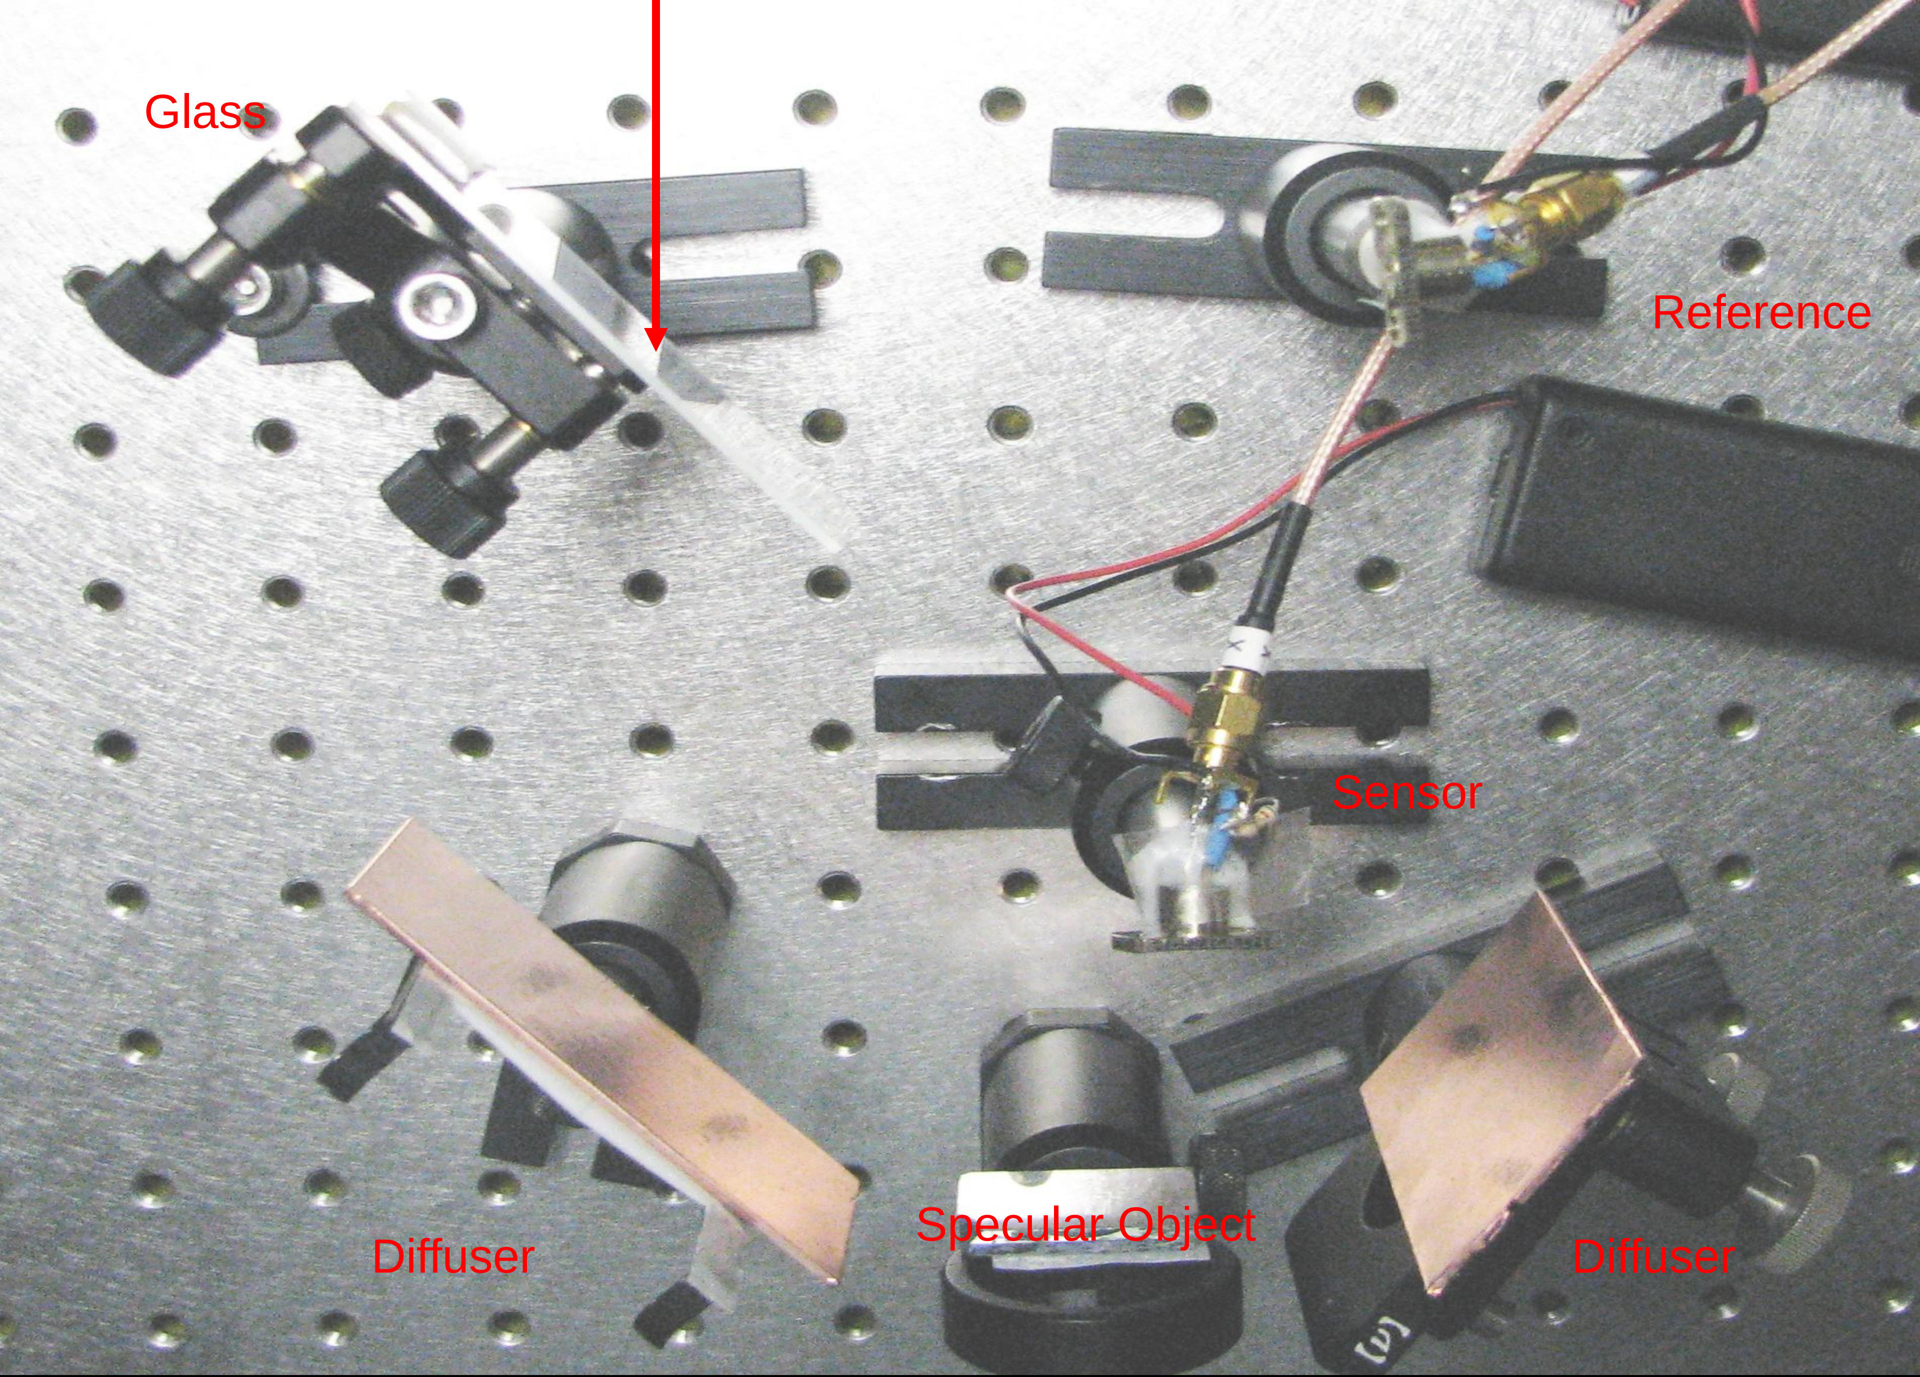

Light enters the scene and strikes a glass surface.

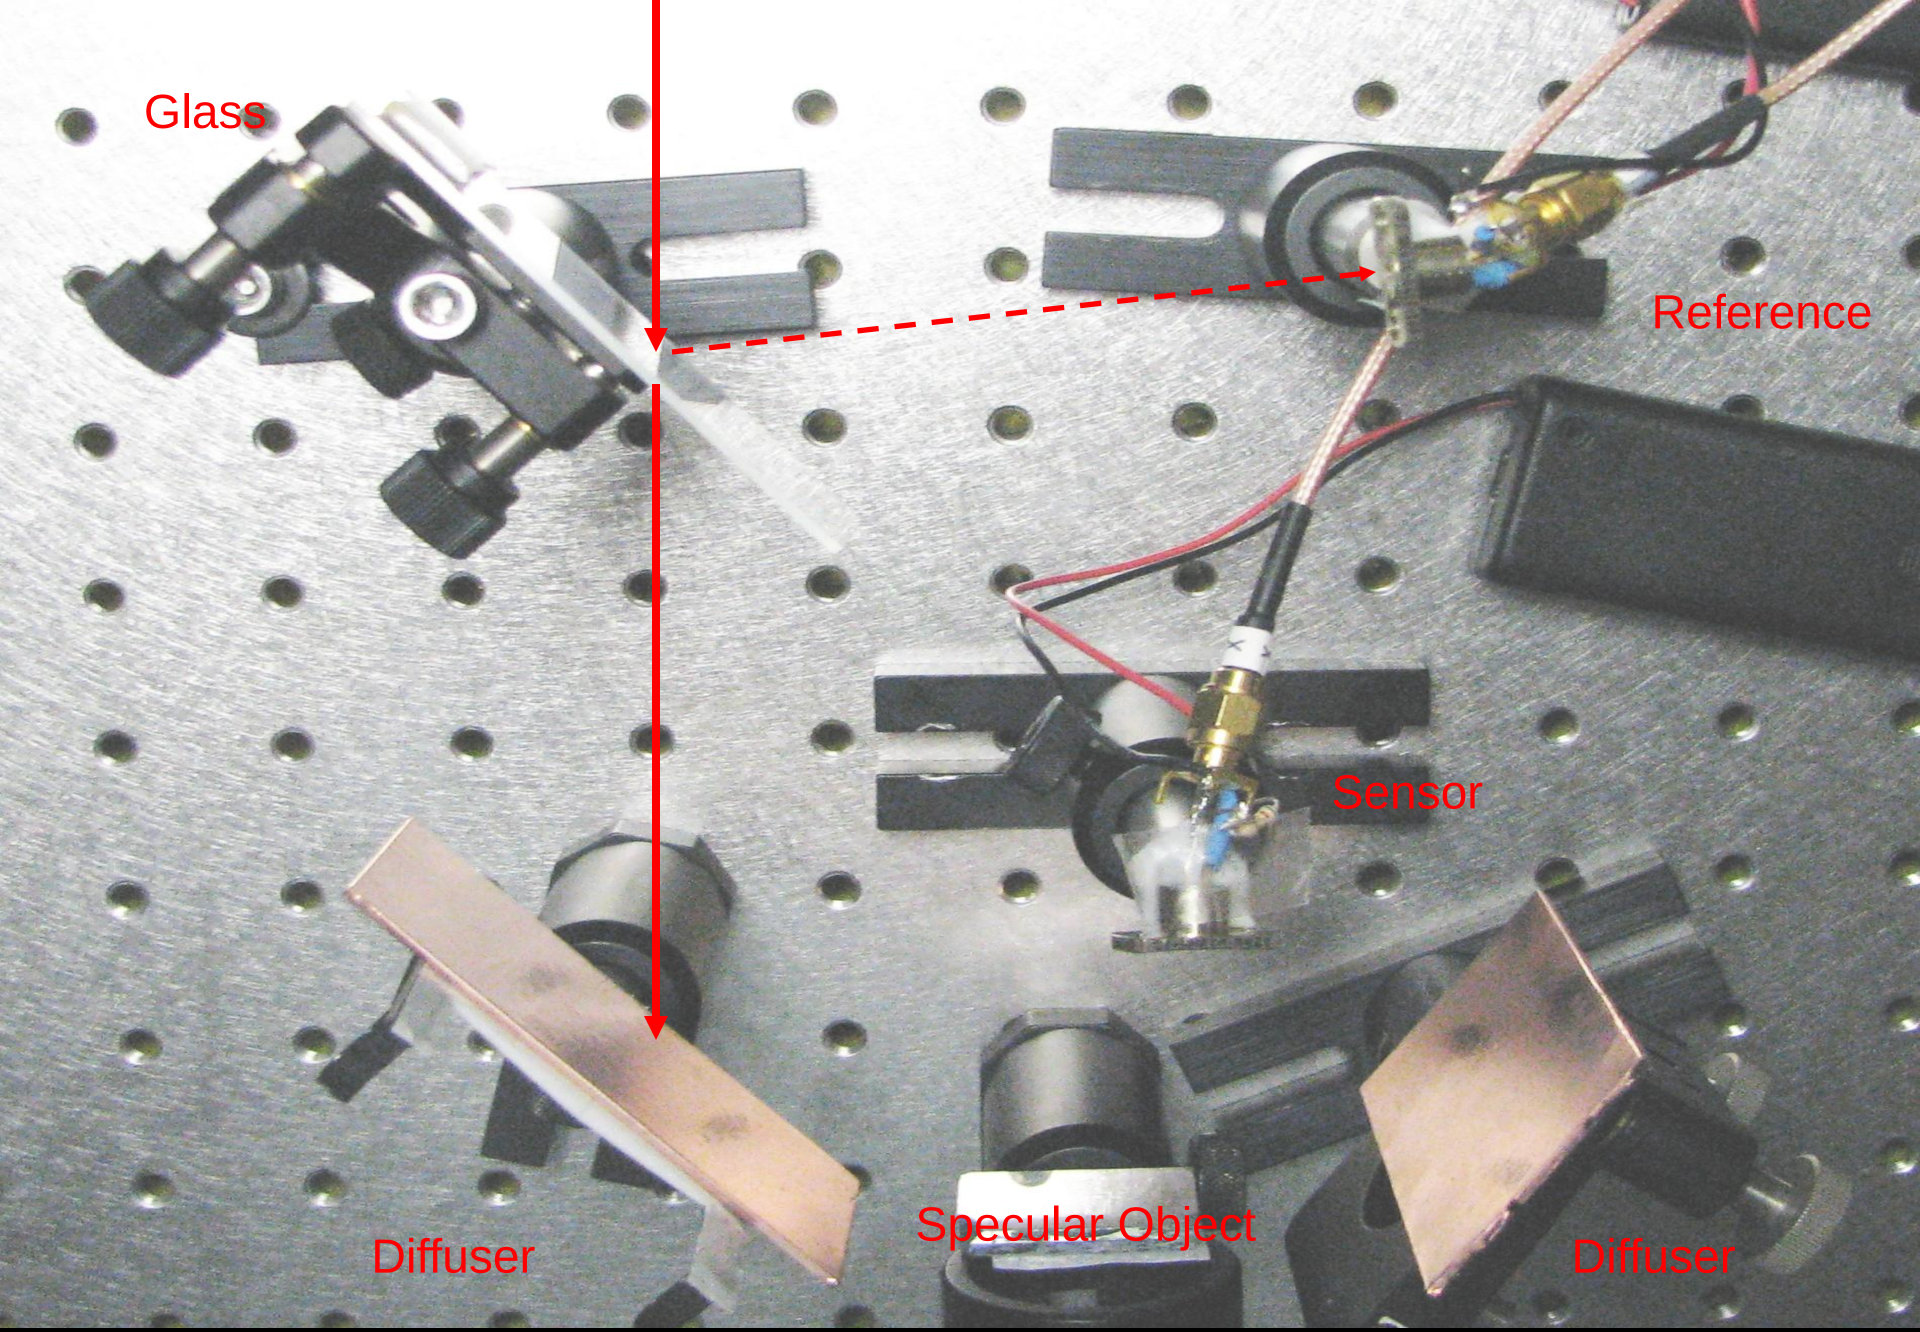

The light reflects off the glass surface and travels back to the reference. The main energy in the laser continues on to the diffuser.

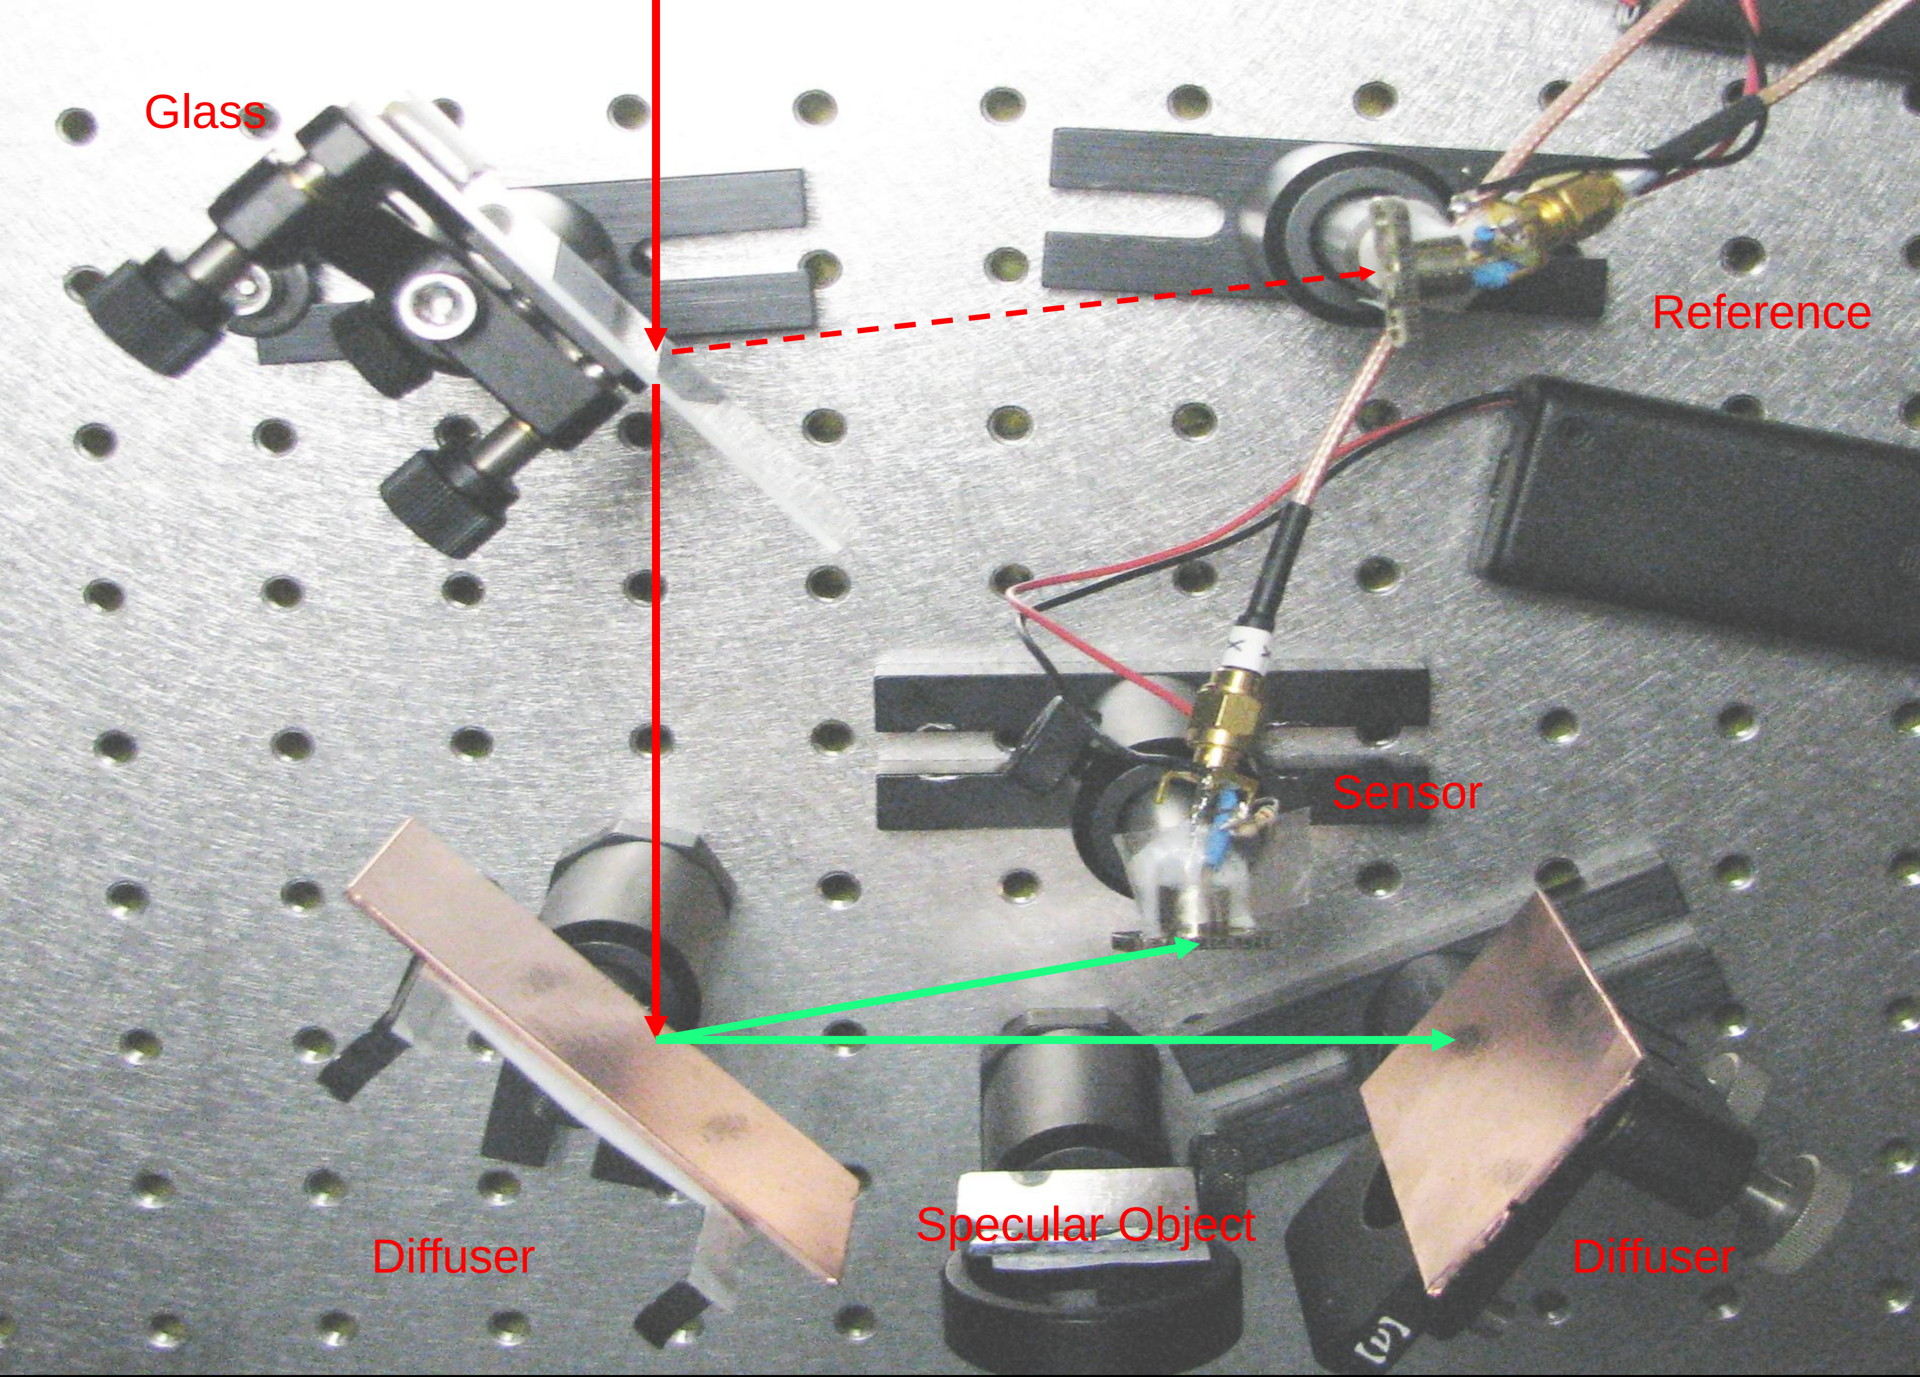

Light from the first diffuser can be captured by the sensor and also propagates to another diffuser.

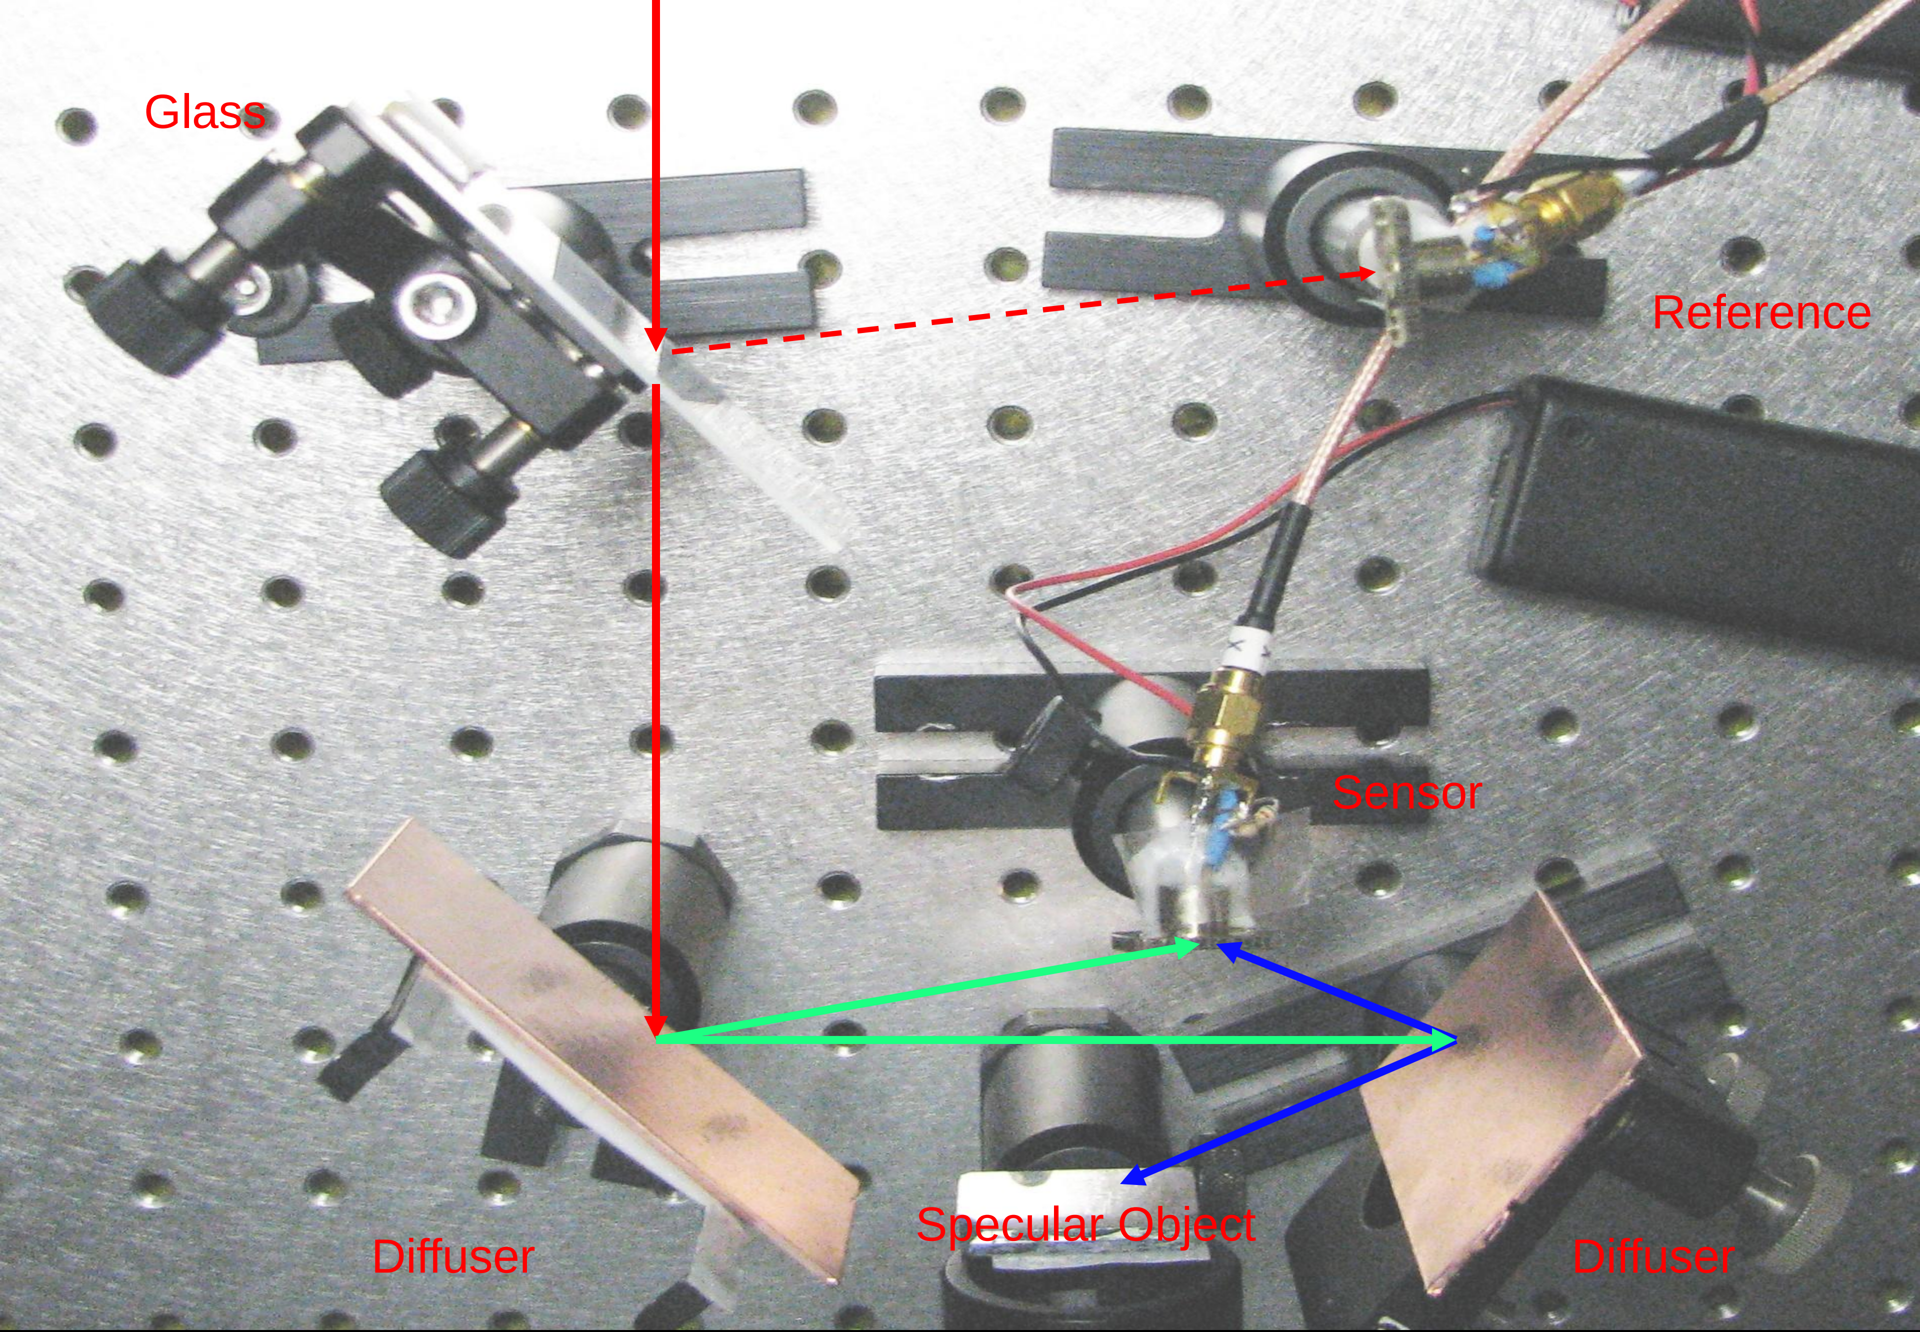

Light then reflects off of the second diffuser and can be captured by the sensor. Light from the second diffuser reflects towards a mirror which then is captured by the sensor.

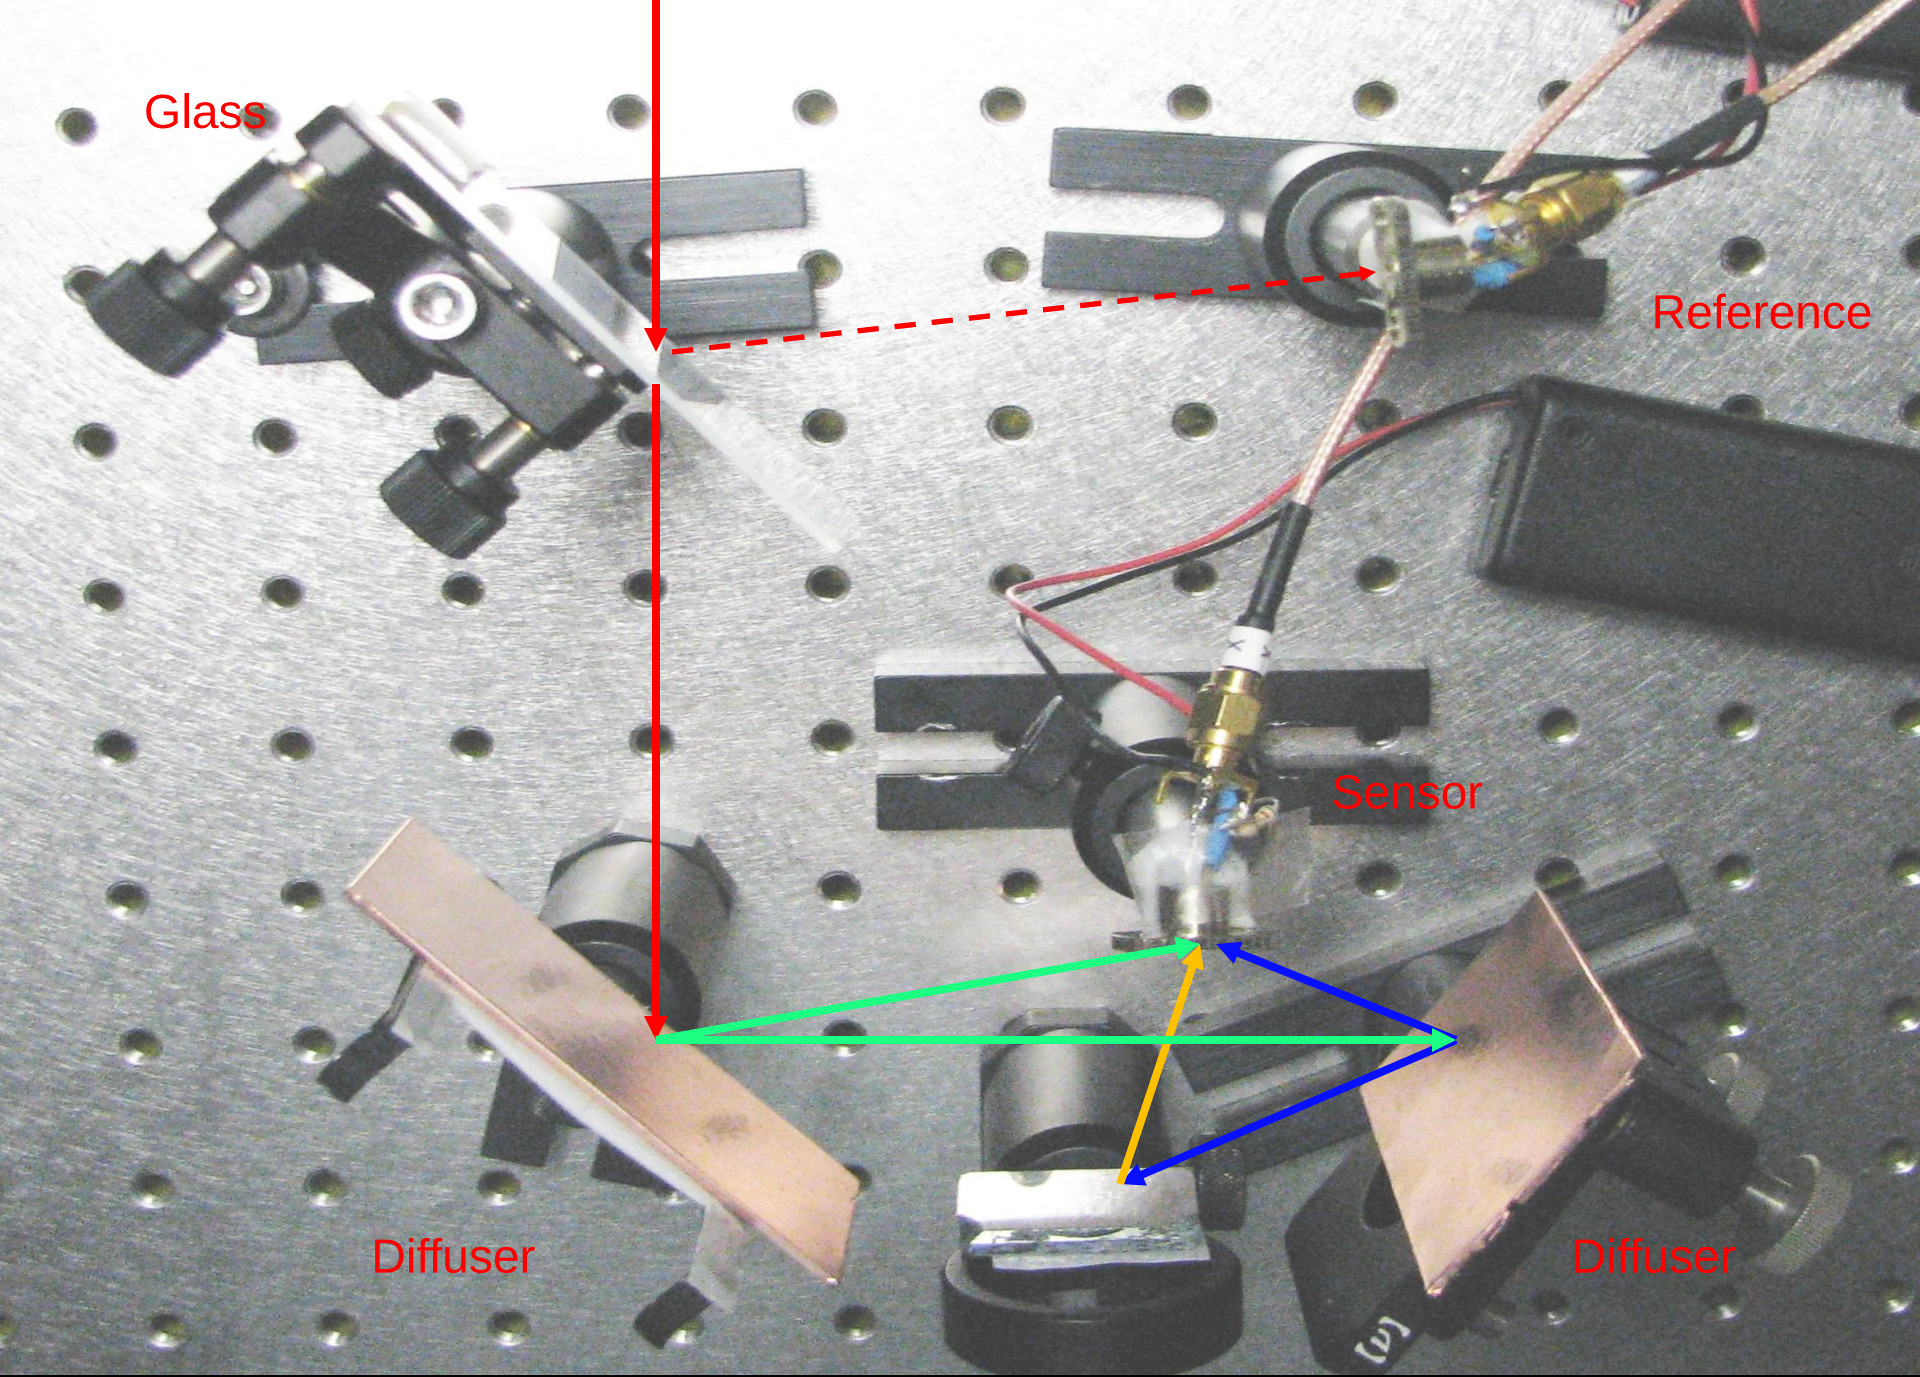

Finally, the remaining light reflects off the mirror to the sensor.

Capturing a Third Bounce

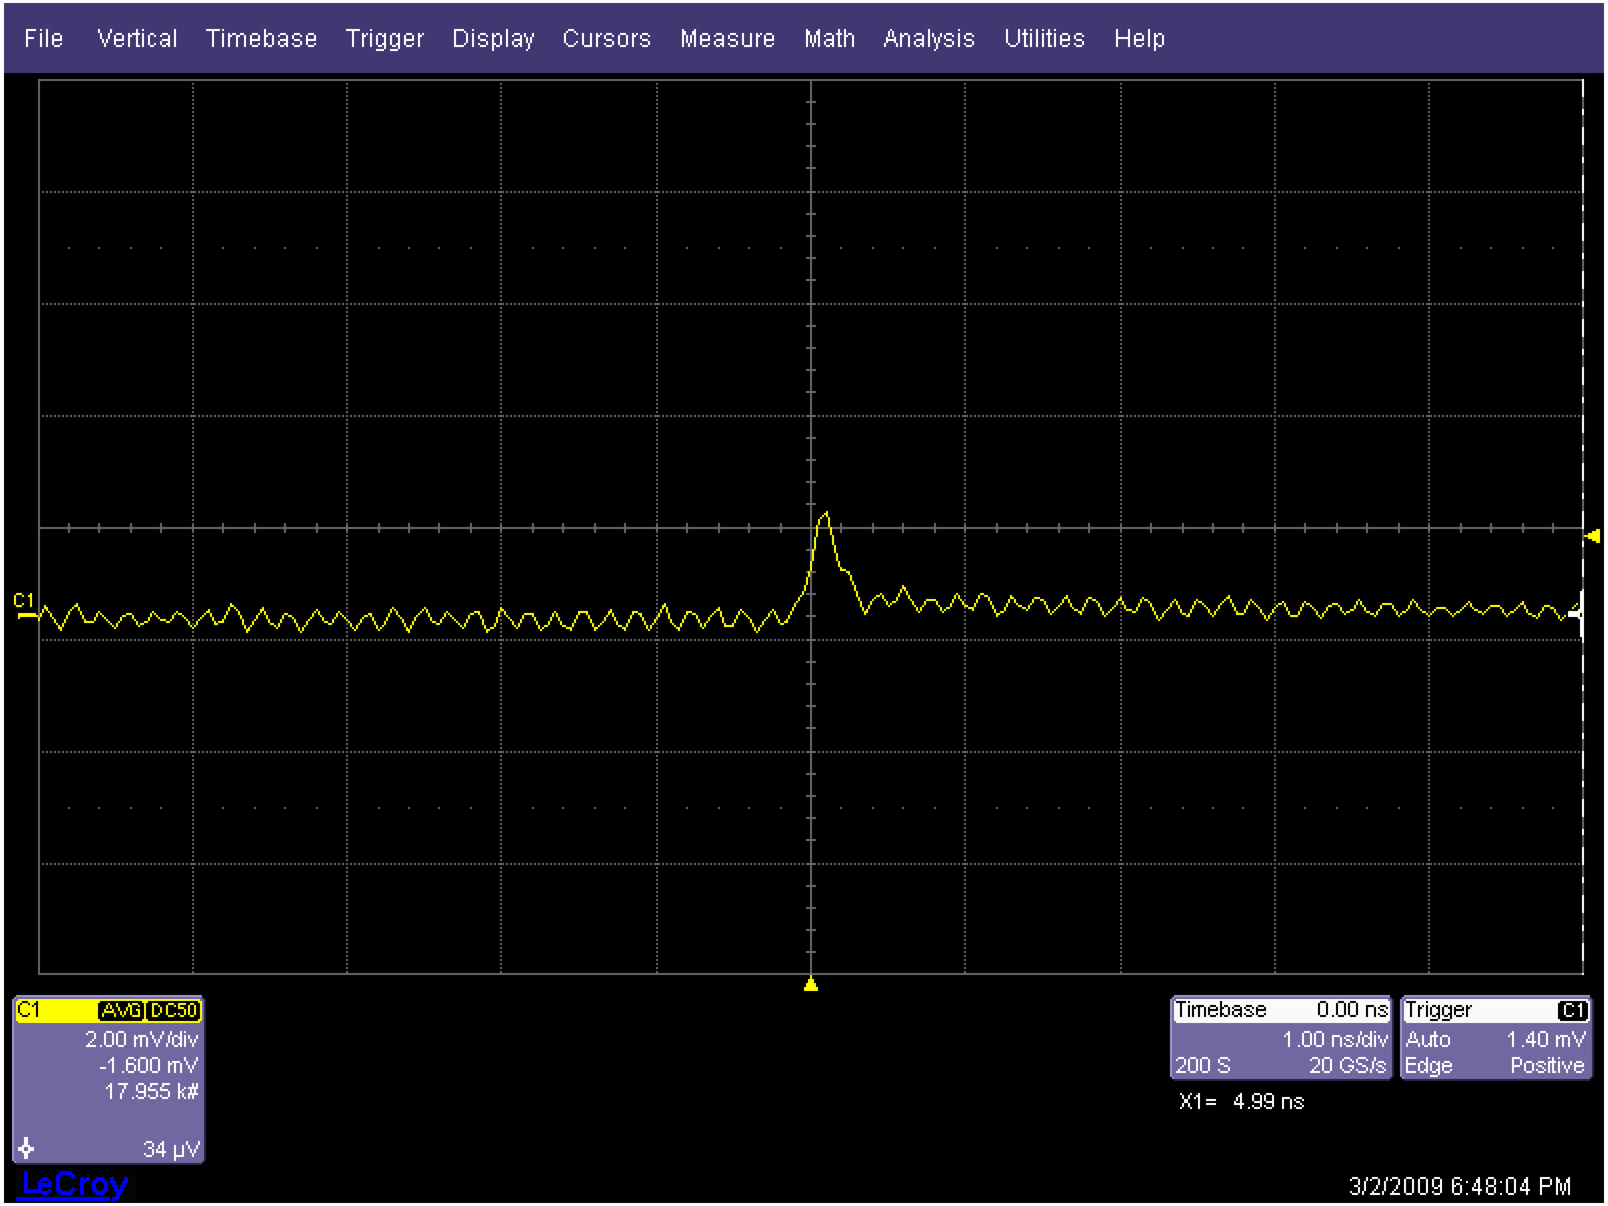

Screenshot from the oscilloscope when the second bounce was captured

Capturing a Third Bounce

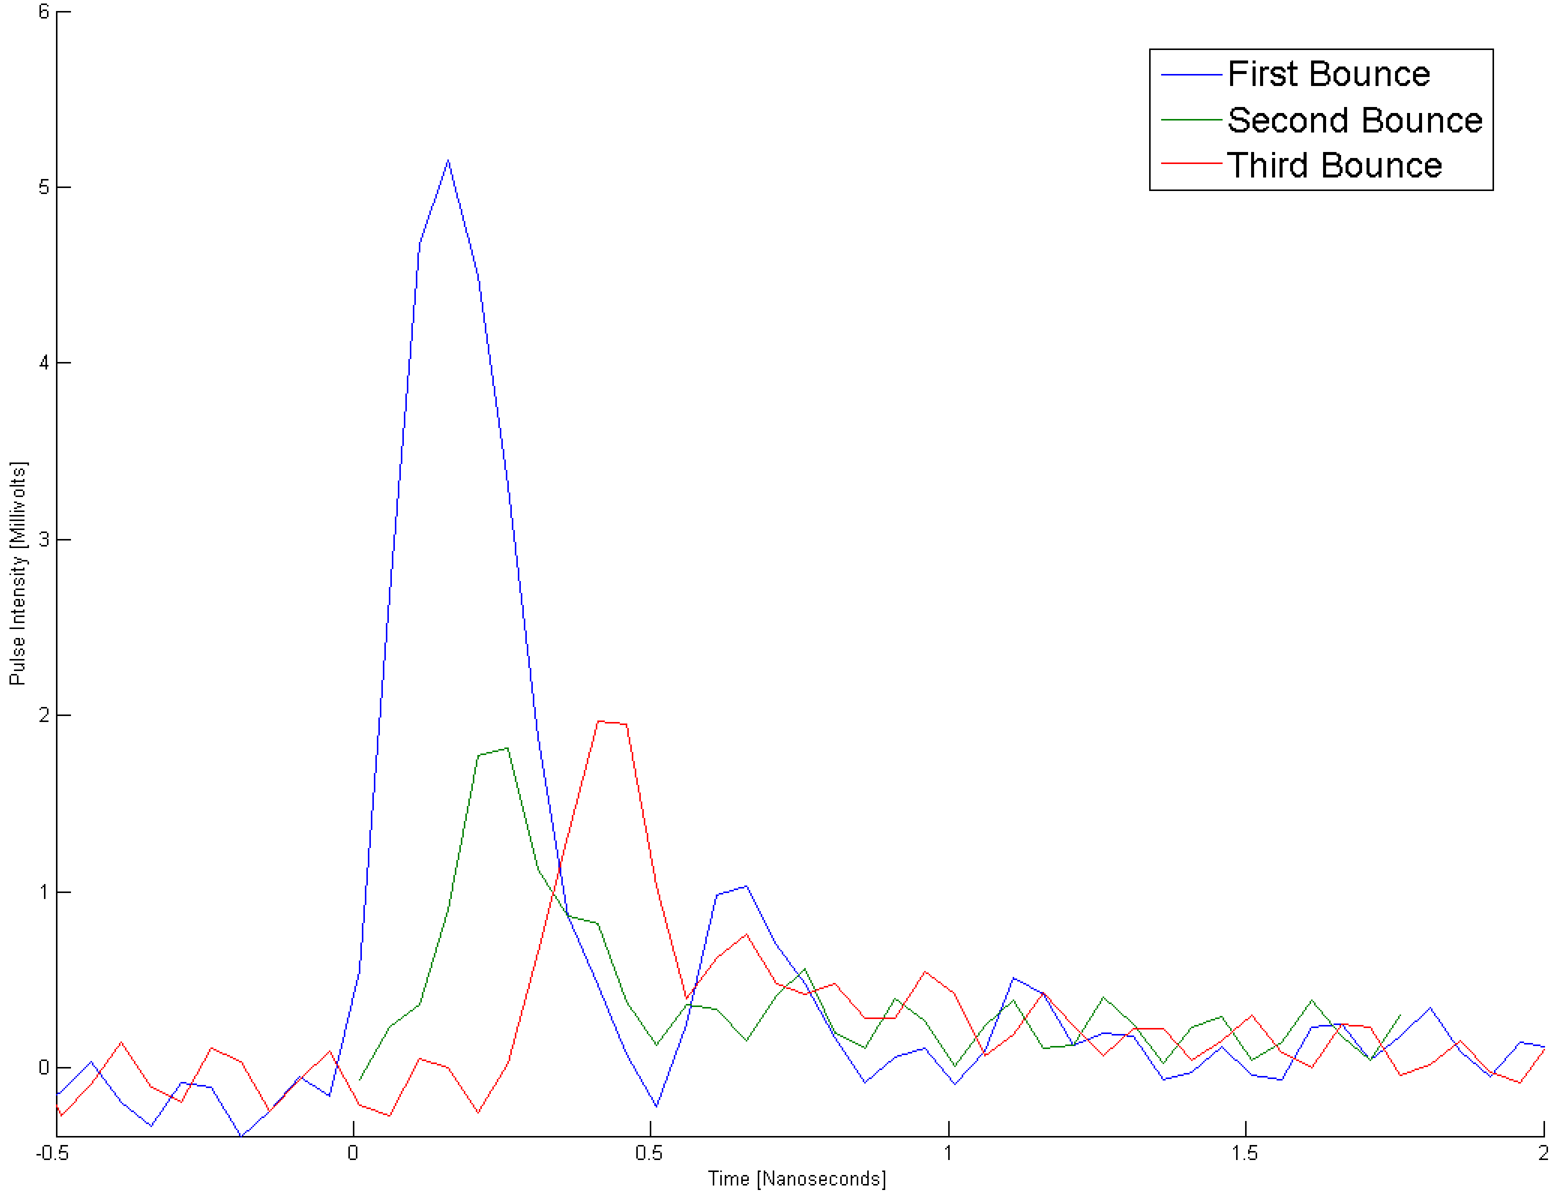

Even after two diffuse and one specular reflection, a third bounce is still significantly above the noise floor.

Missing Direct Reflection

Estimating hidden specular surfaces visible to the laser but invisible to the sensor

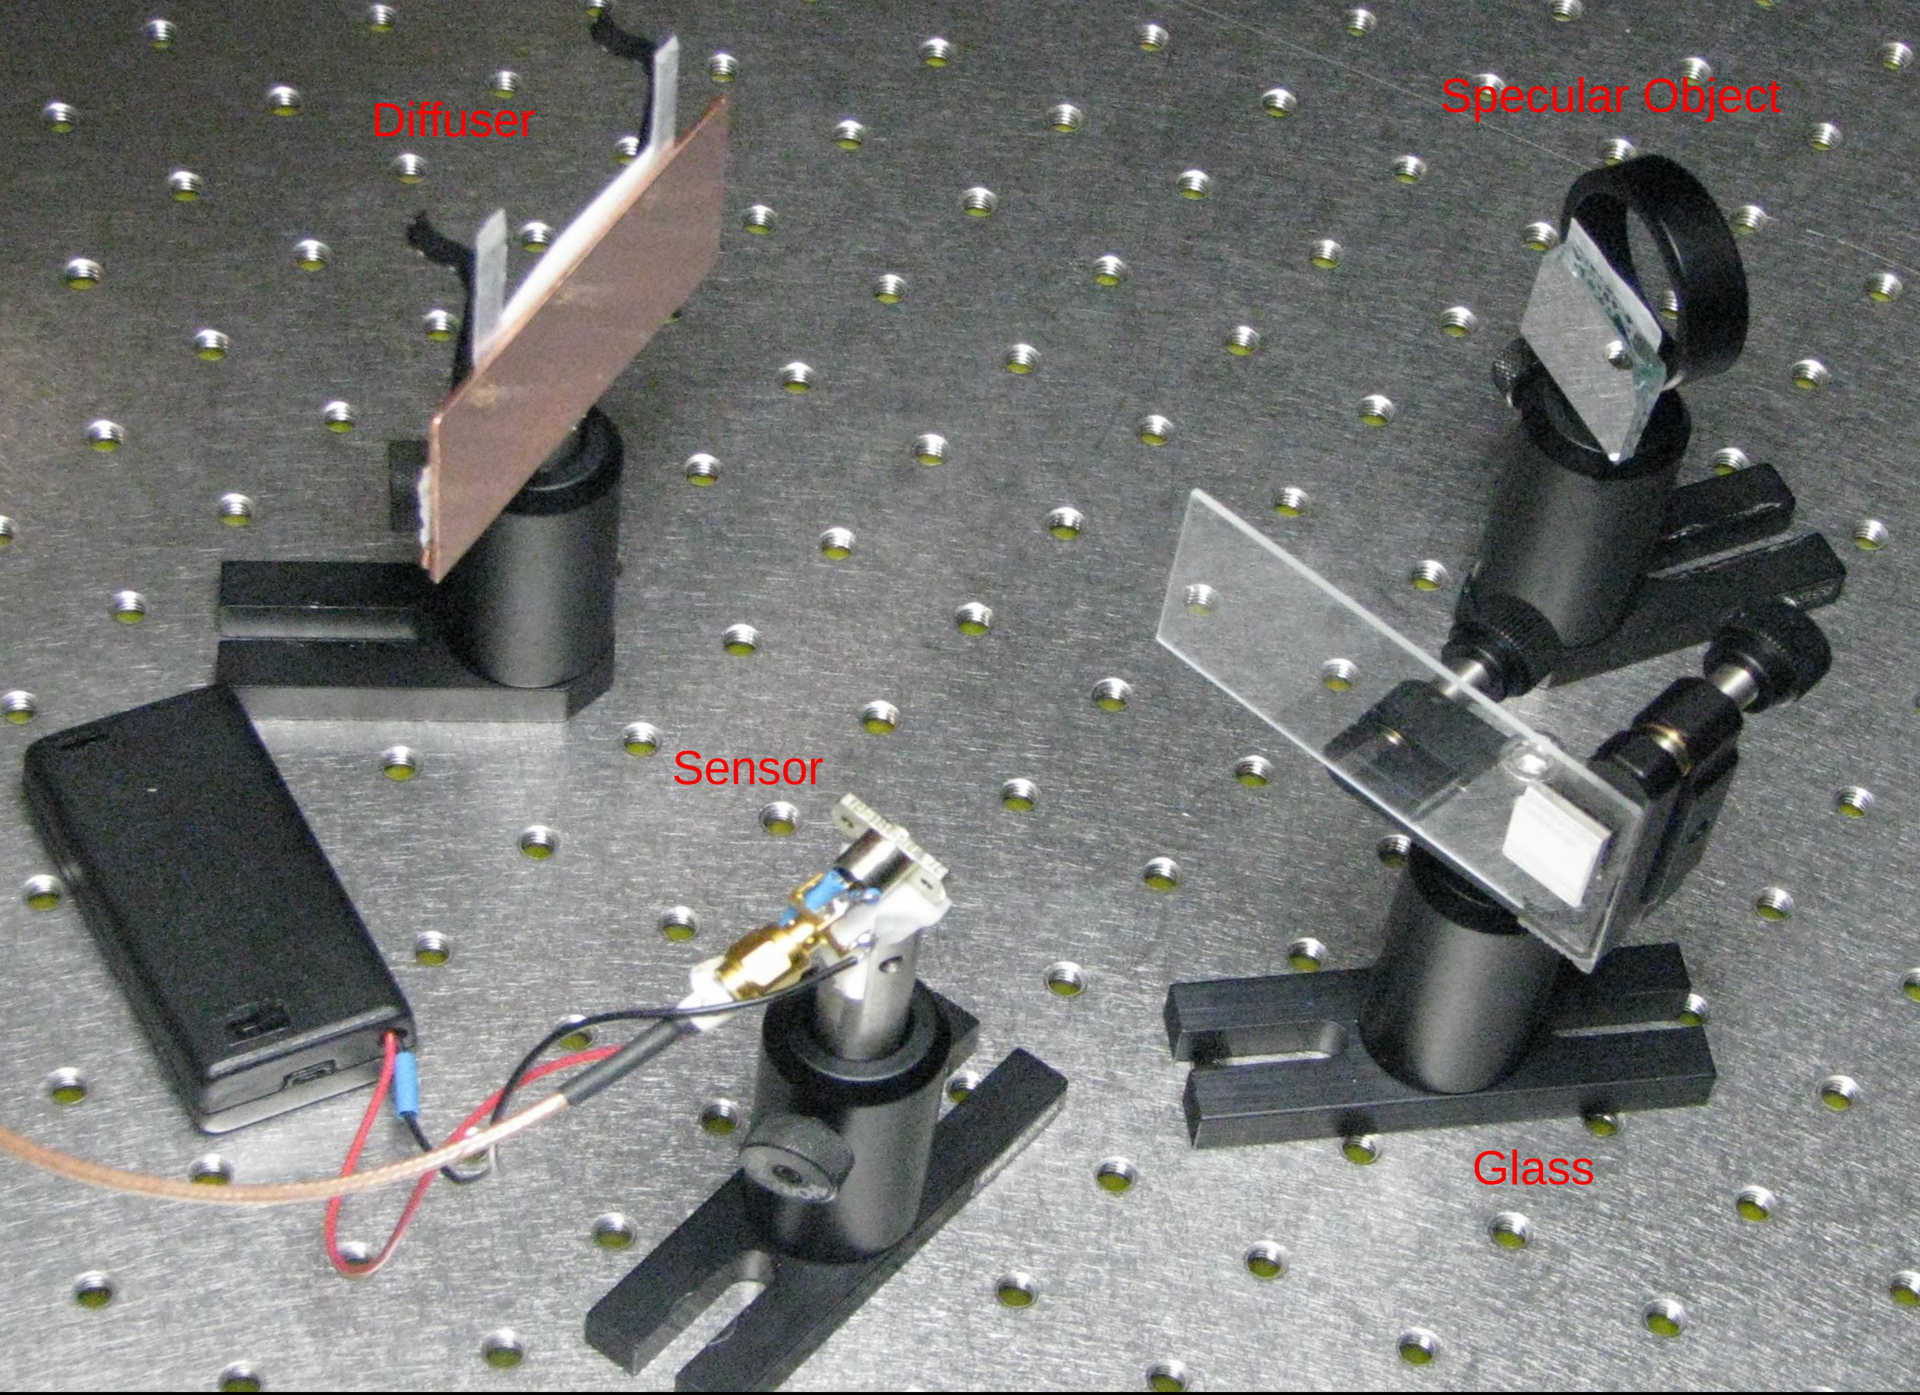

Here we show the setup for the missing direct reflection experiment.

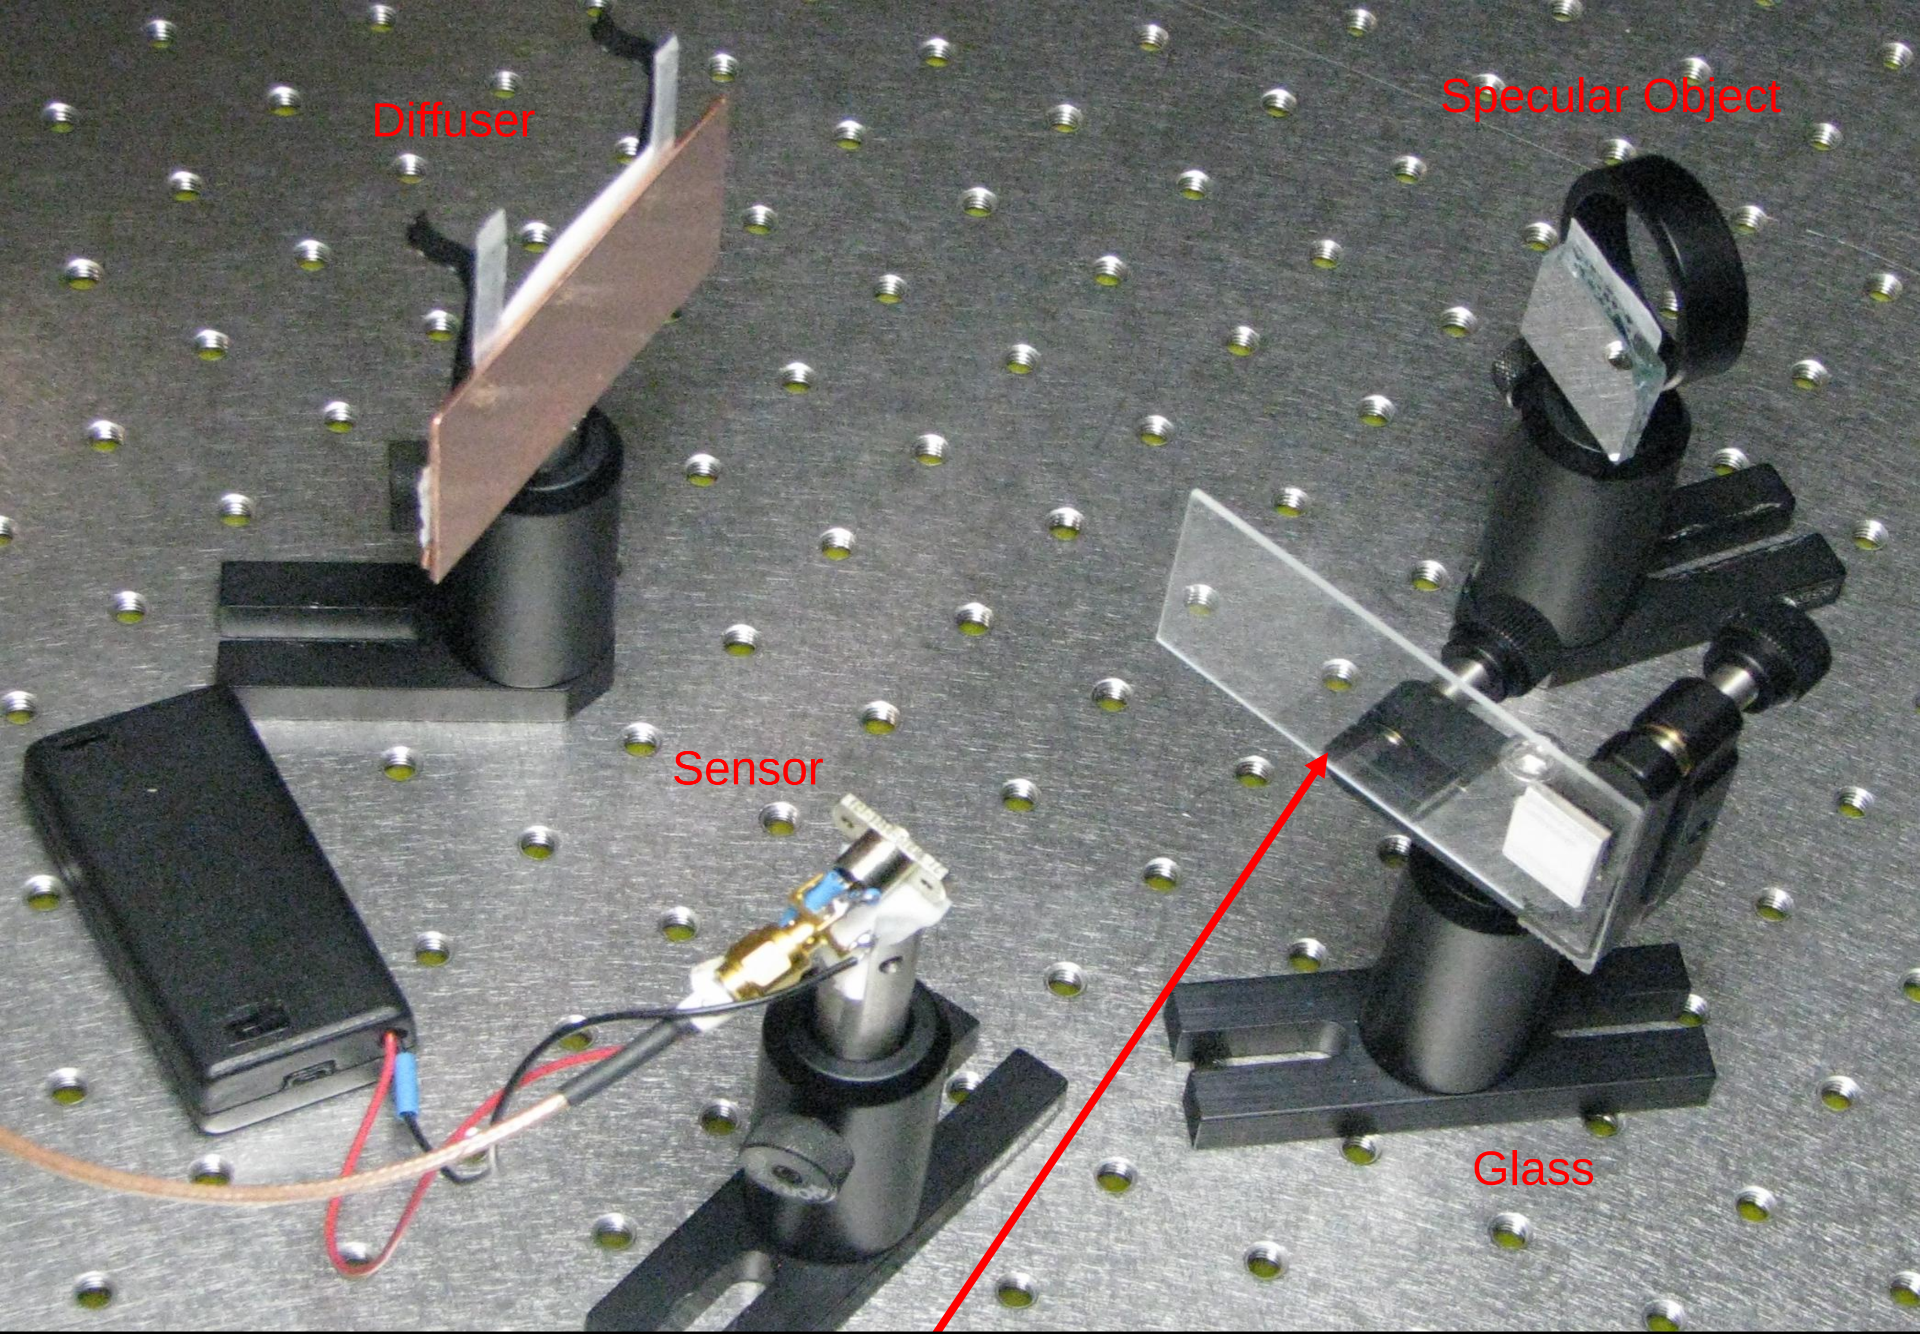

Light enters the scene and strikes a glass surface. Although it appears as if the sensor is looking at the specular object, it is only receiving the reference trigger from the glass reflector.

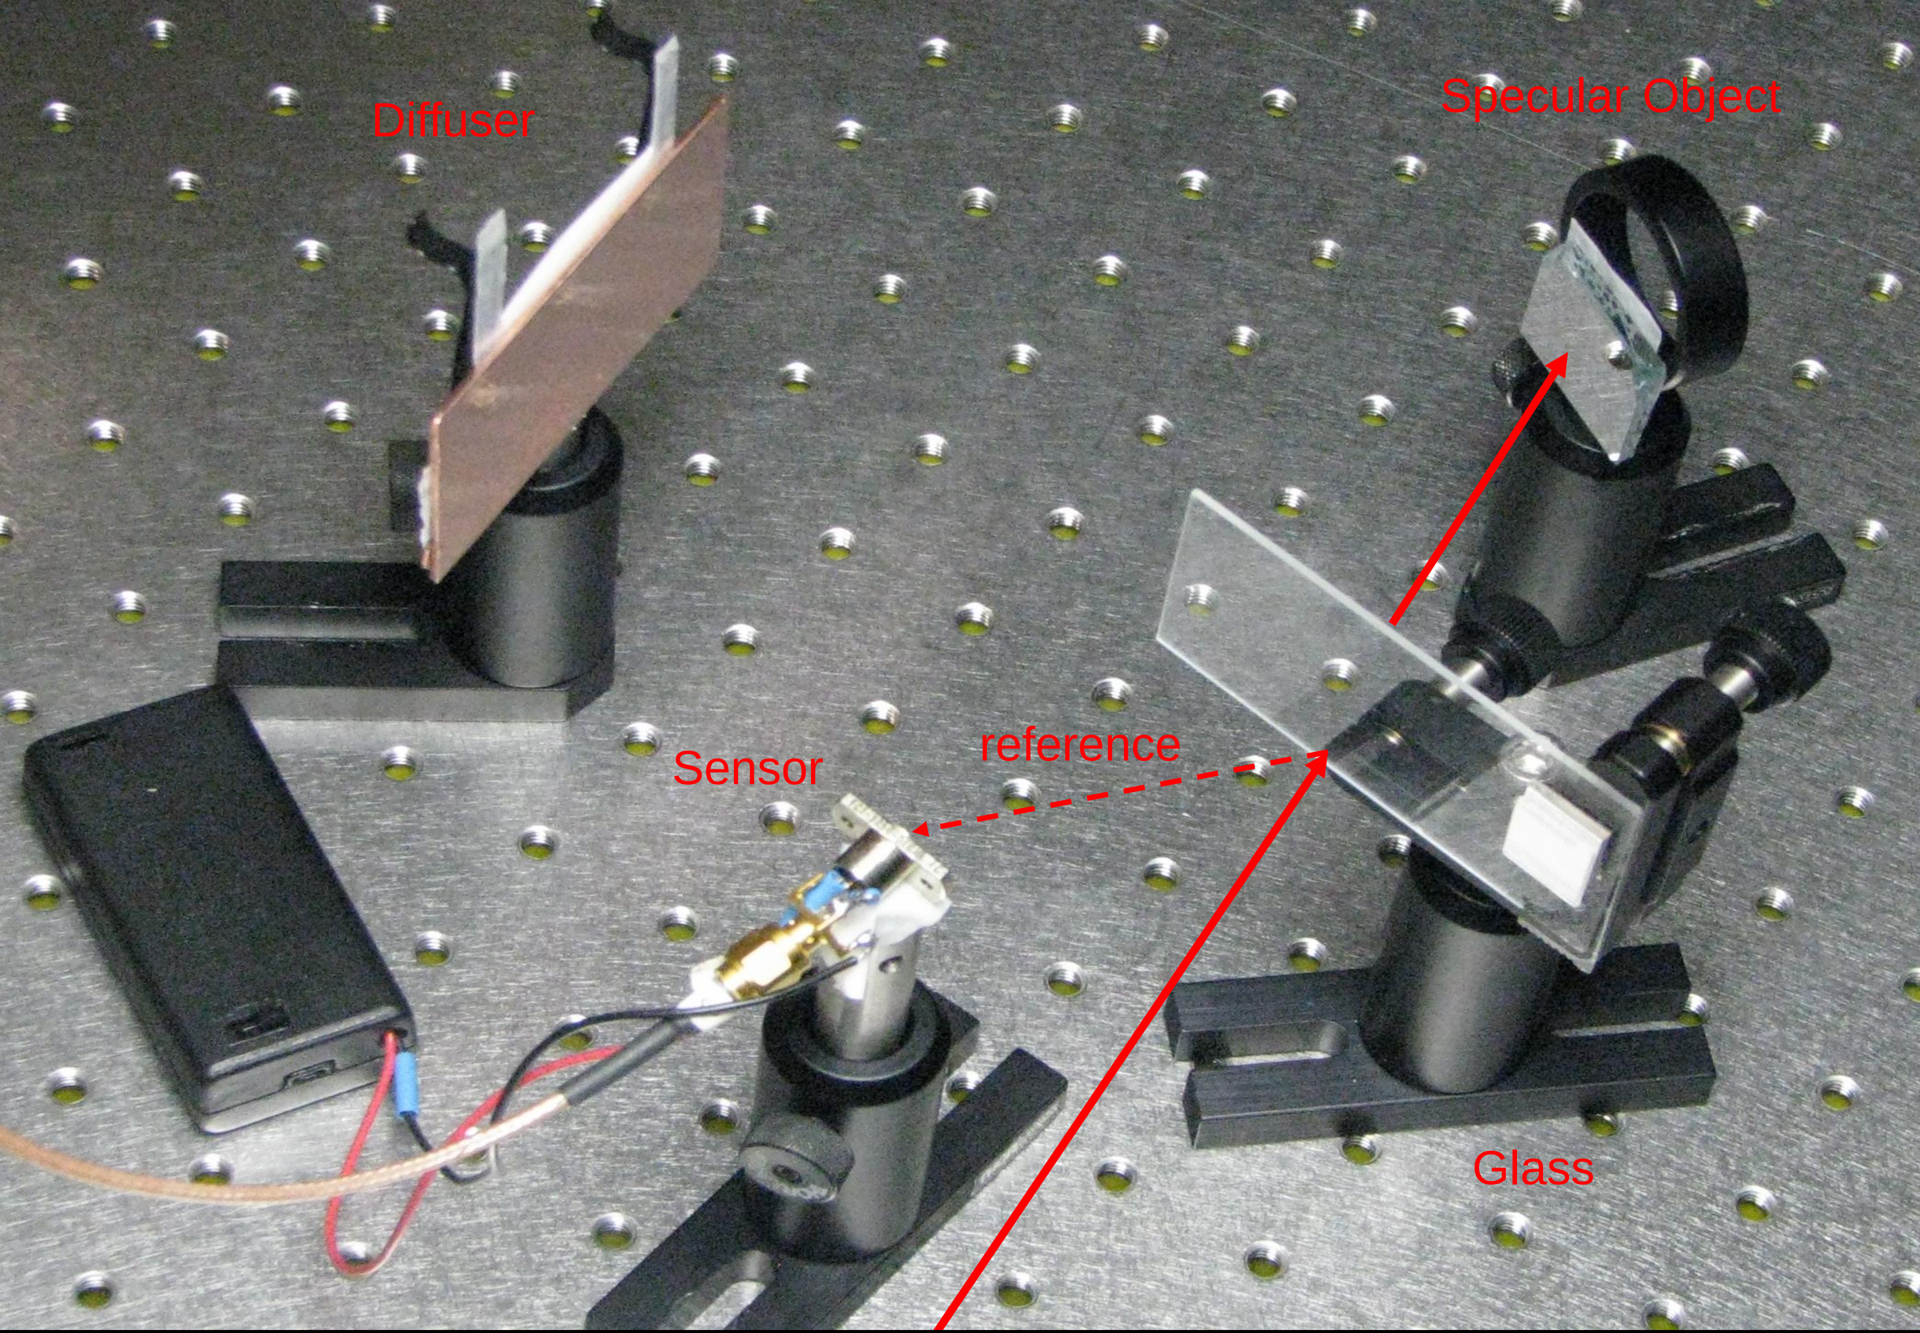

The light reflects off the glass surface and travels back to the reference. The main energy in the laser continues on to the diffuser.

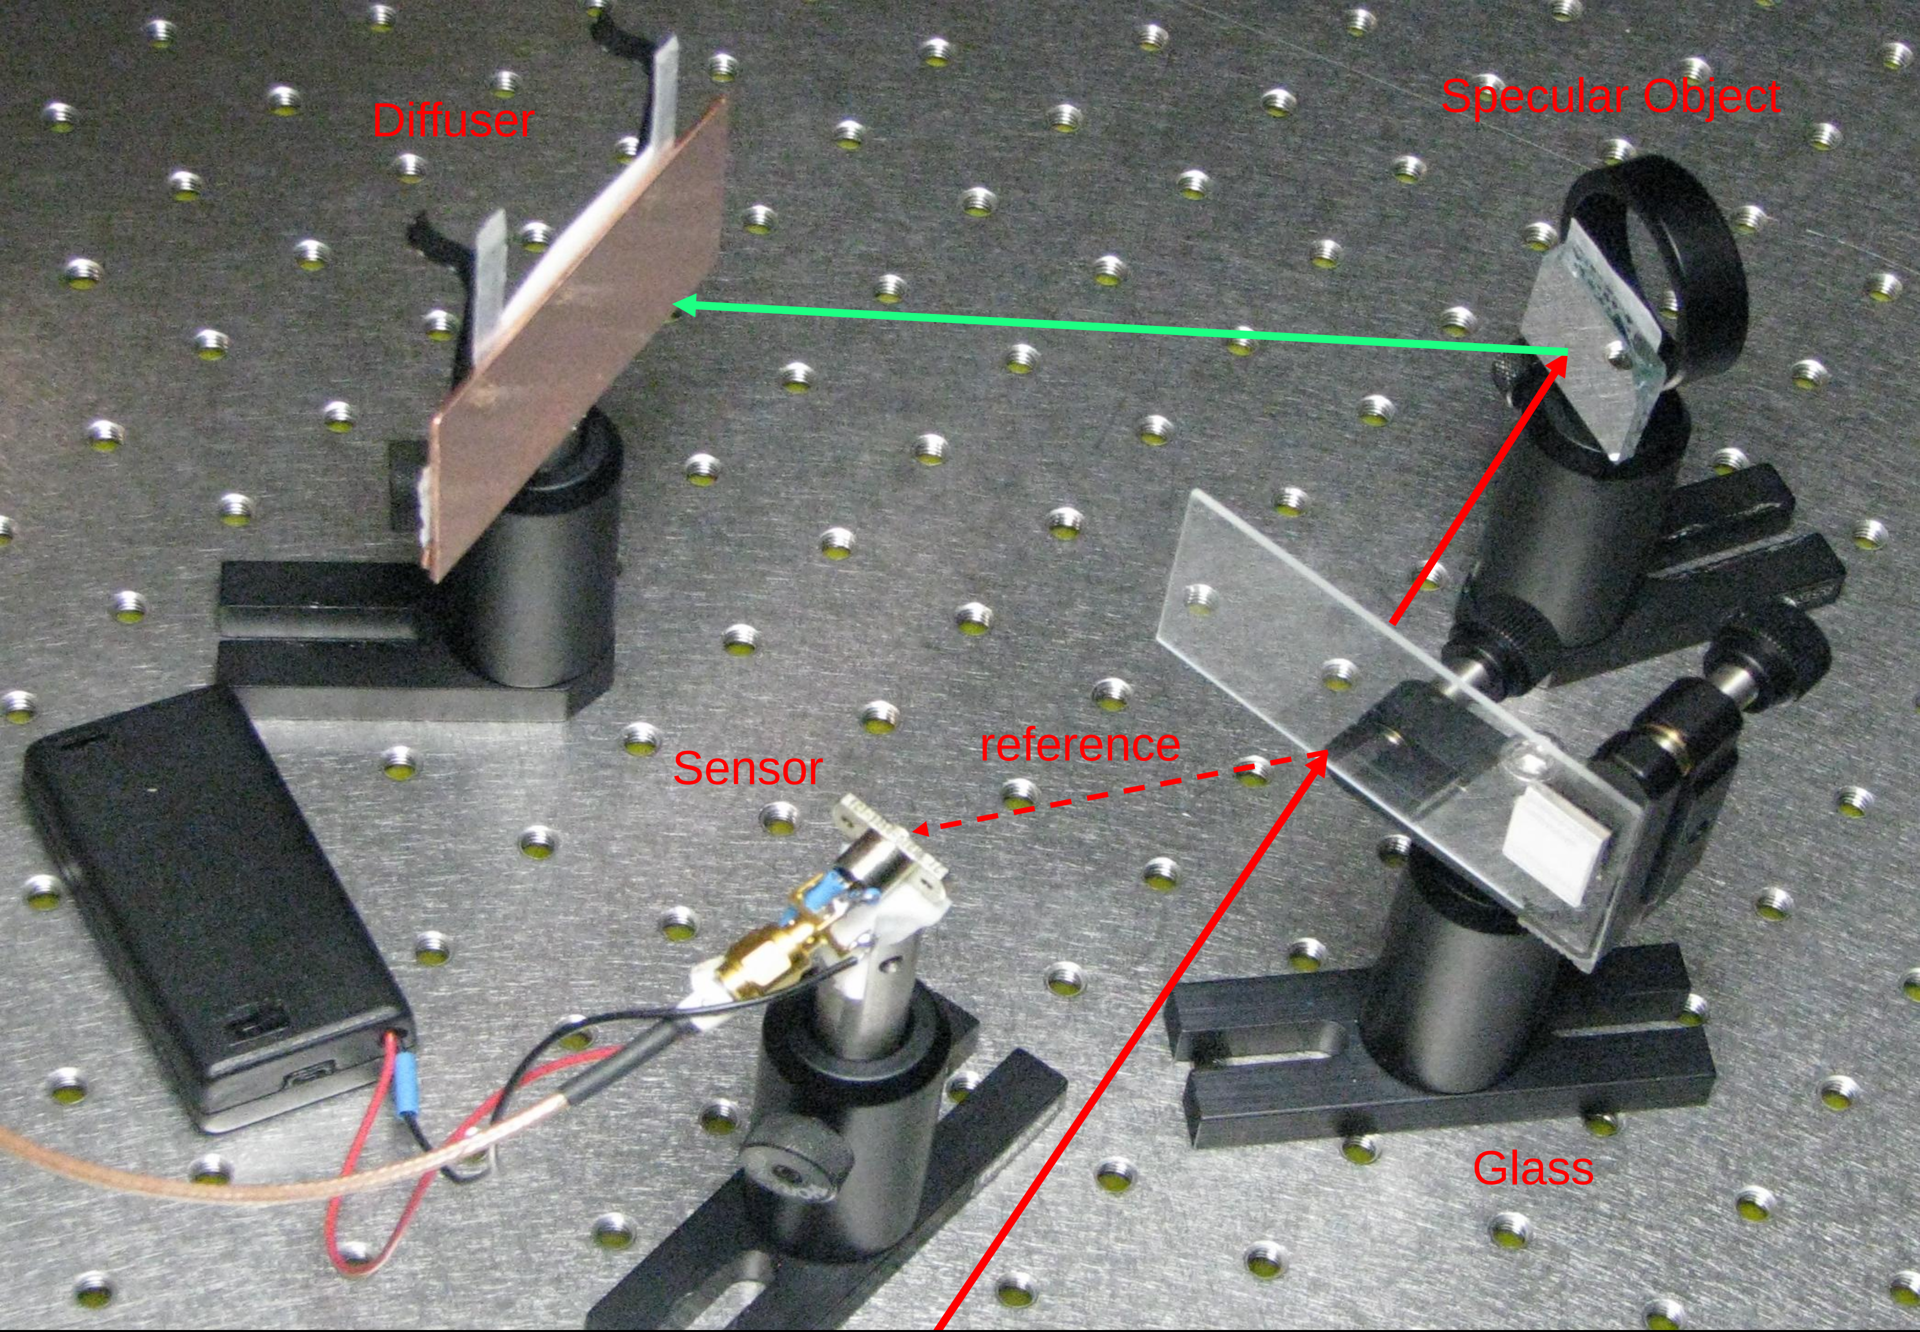

The remaining light strikes a specular object which reflects light towards a diffuser. Note that the sensor receives no light back from the specular object.

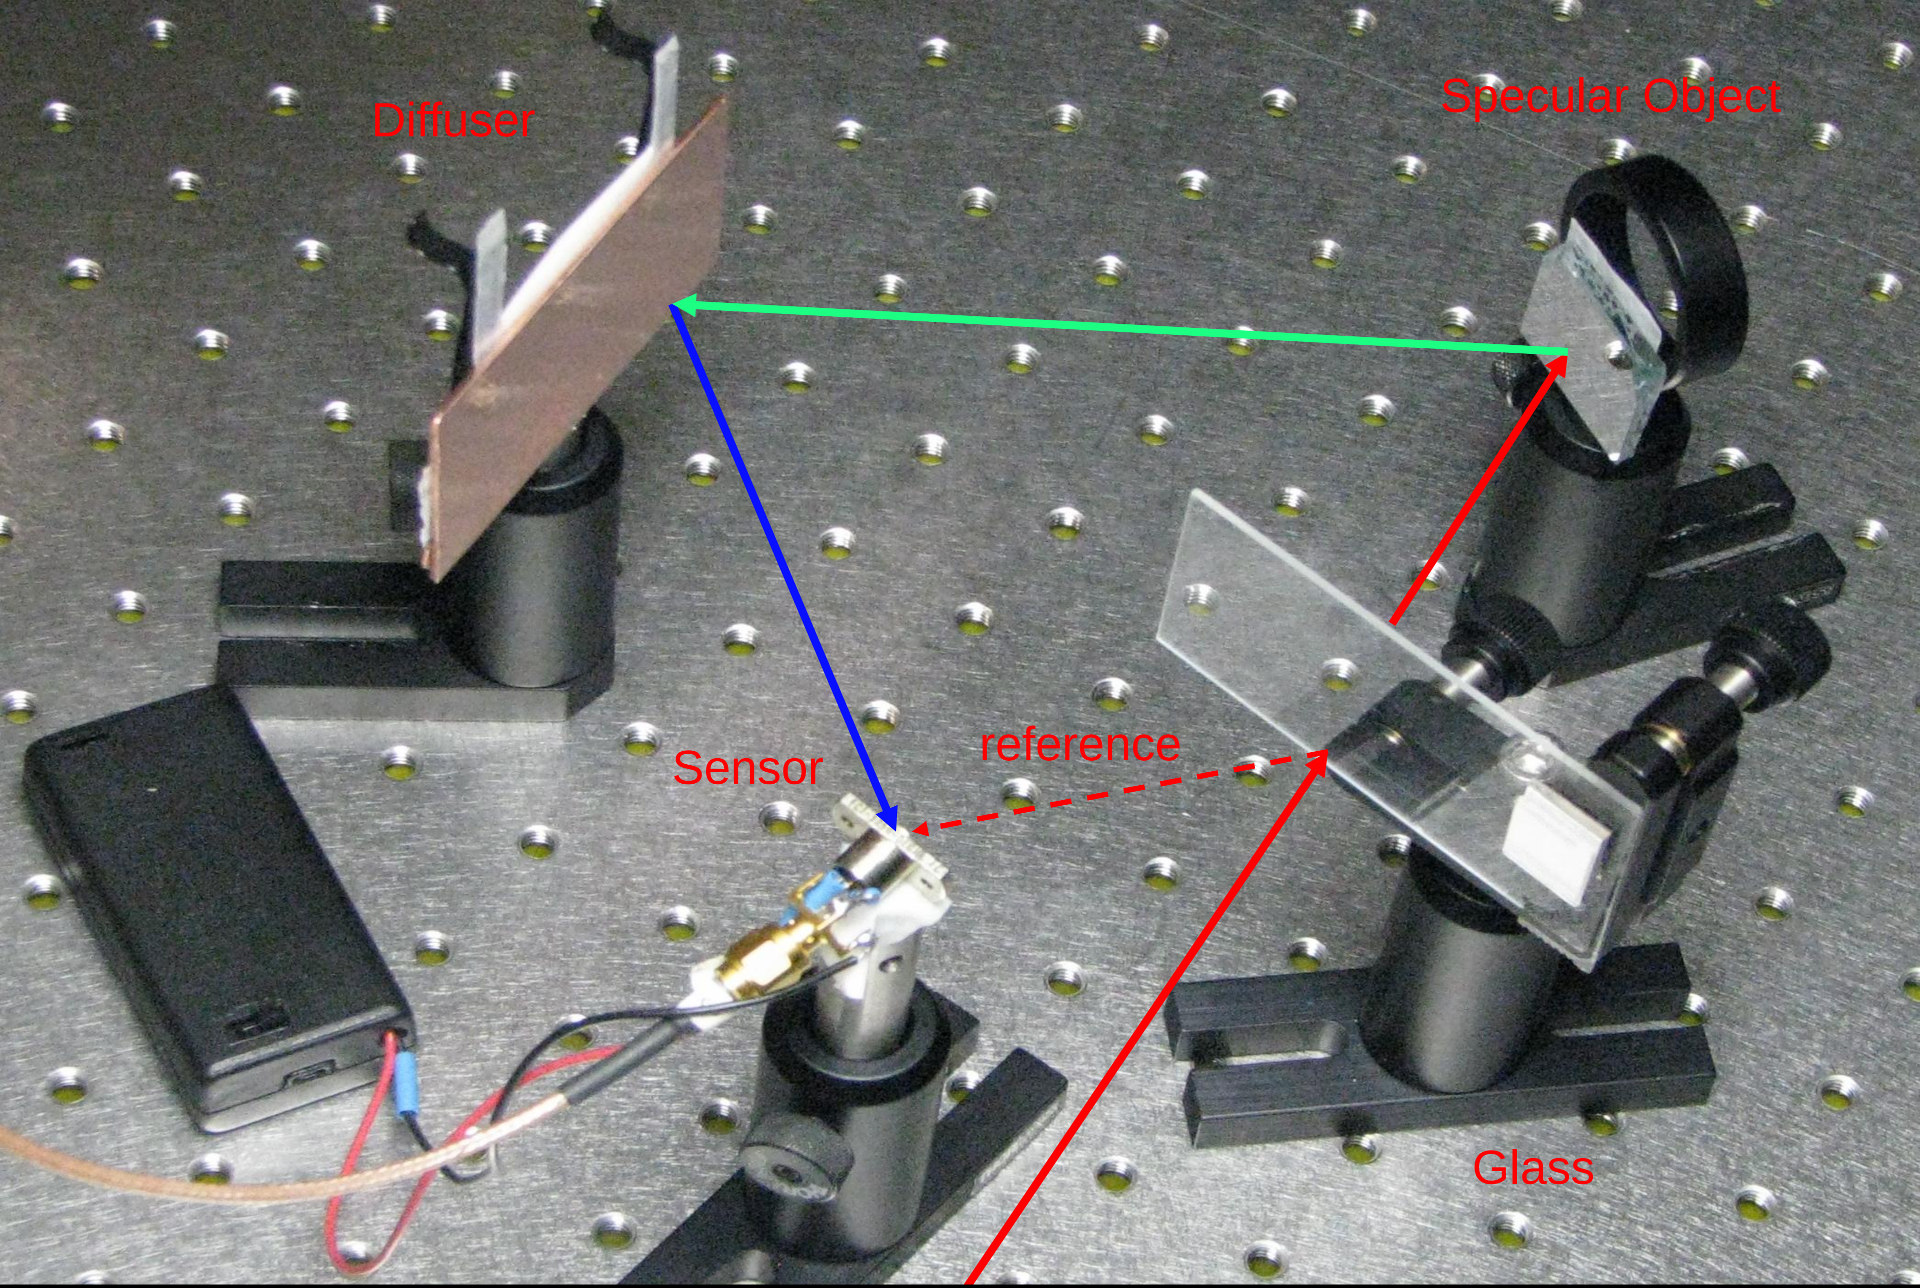

Light hits the diffuser which creates a diffuse hemisphere which can then be captured at the sensor. The experiment is repeated twice to obtain three raster scans based on the position of the light hitting the diffuser.

Missing Direct Reflection

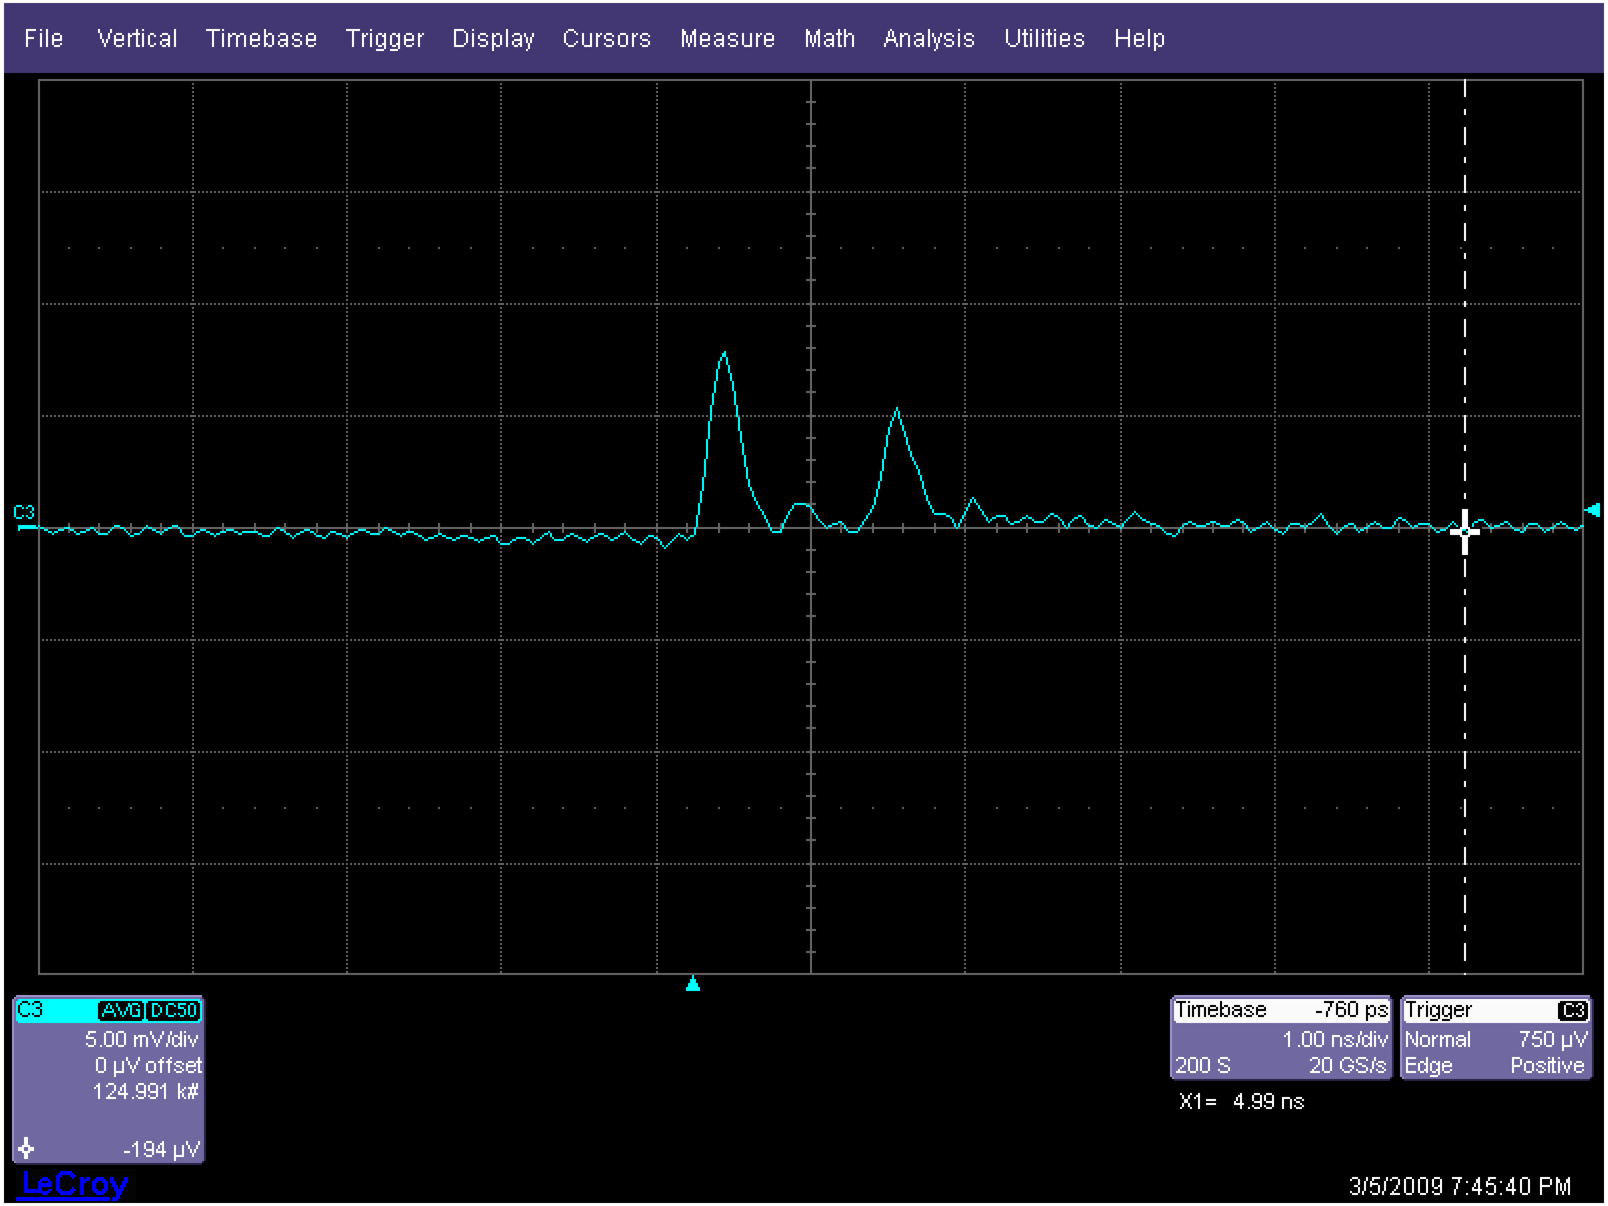

Screenshot of the captured signal on the oscilloscope during the right side raster scan

Missing Direct Reflection

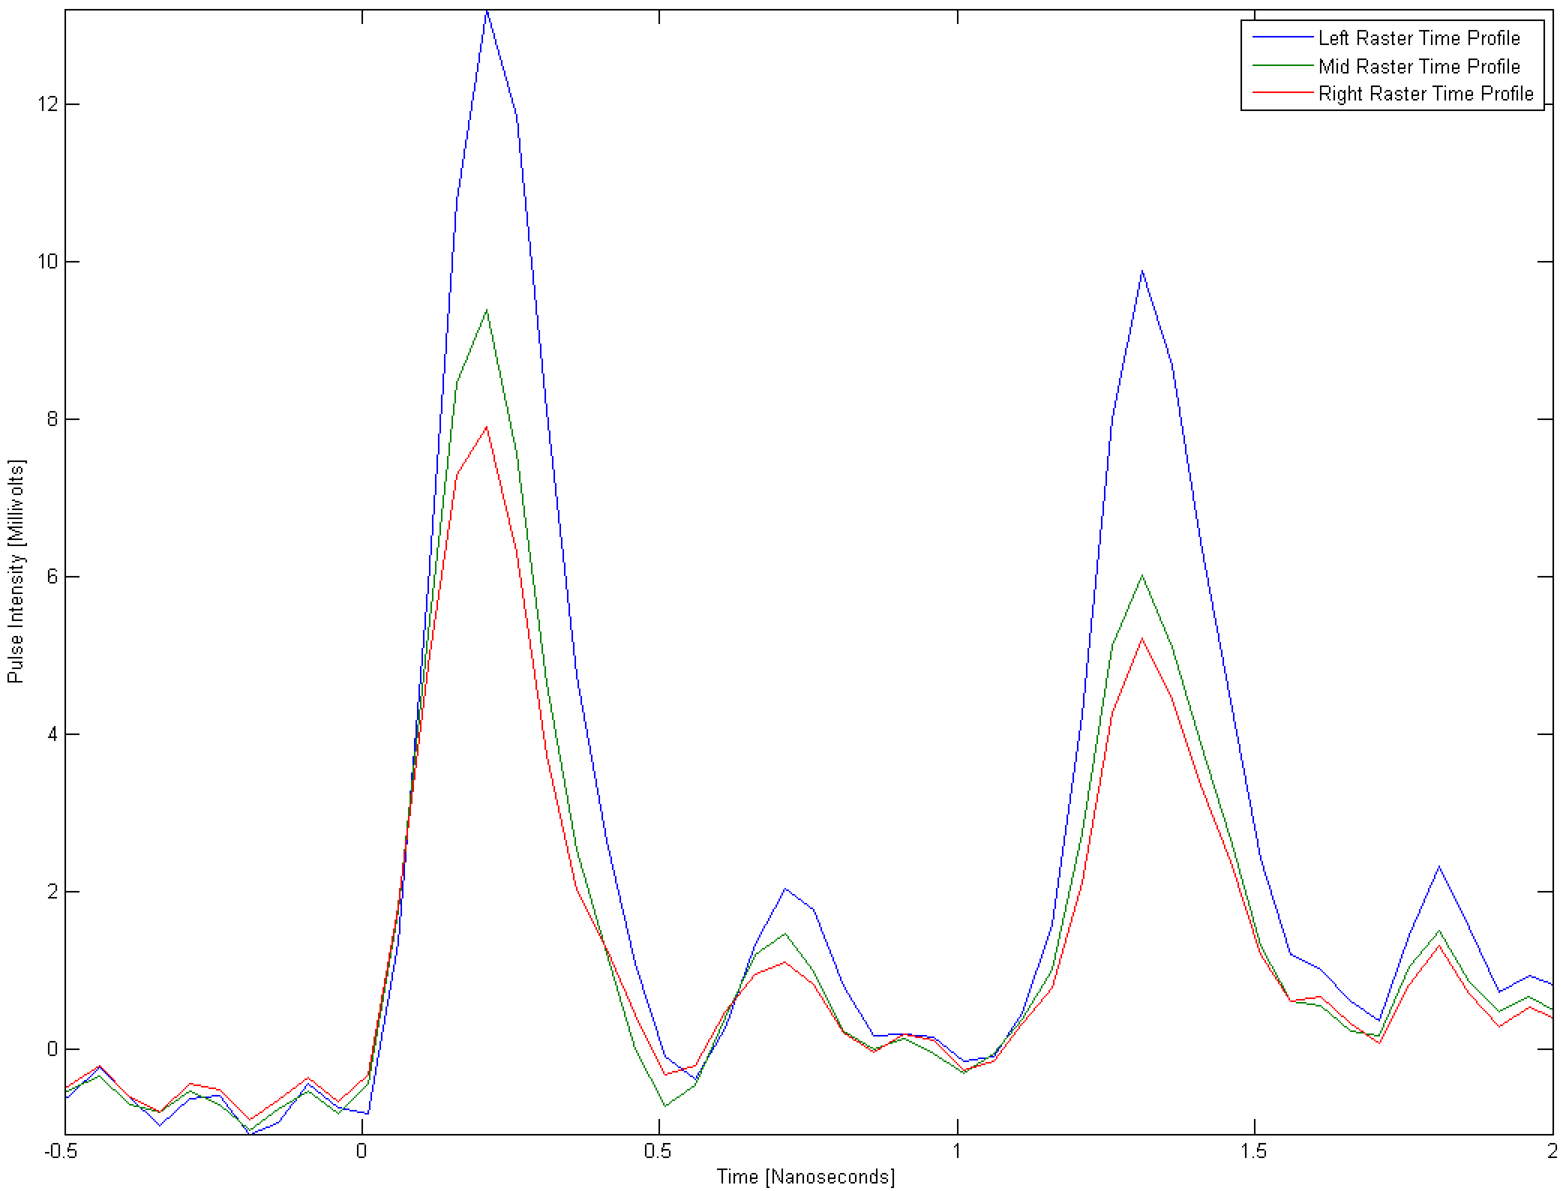

Plot showing TDOA for the 3 raster scans. Note that the plots are directly generated by the oscilloscope.

Missing Direct Reflection

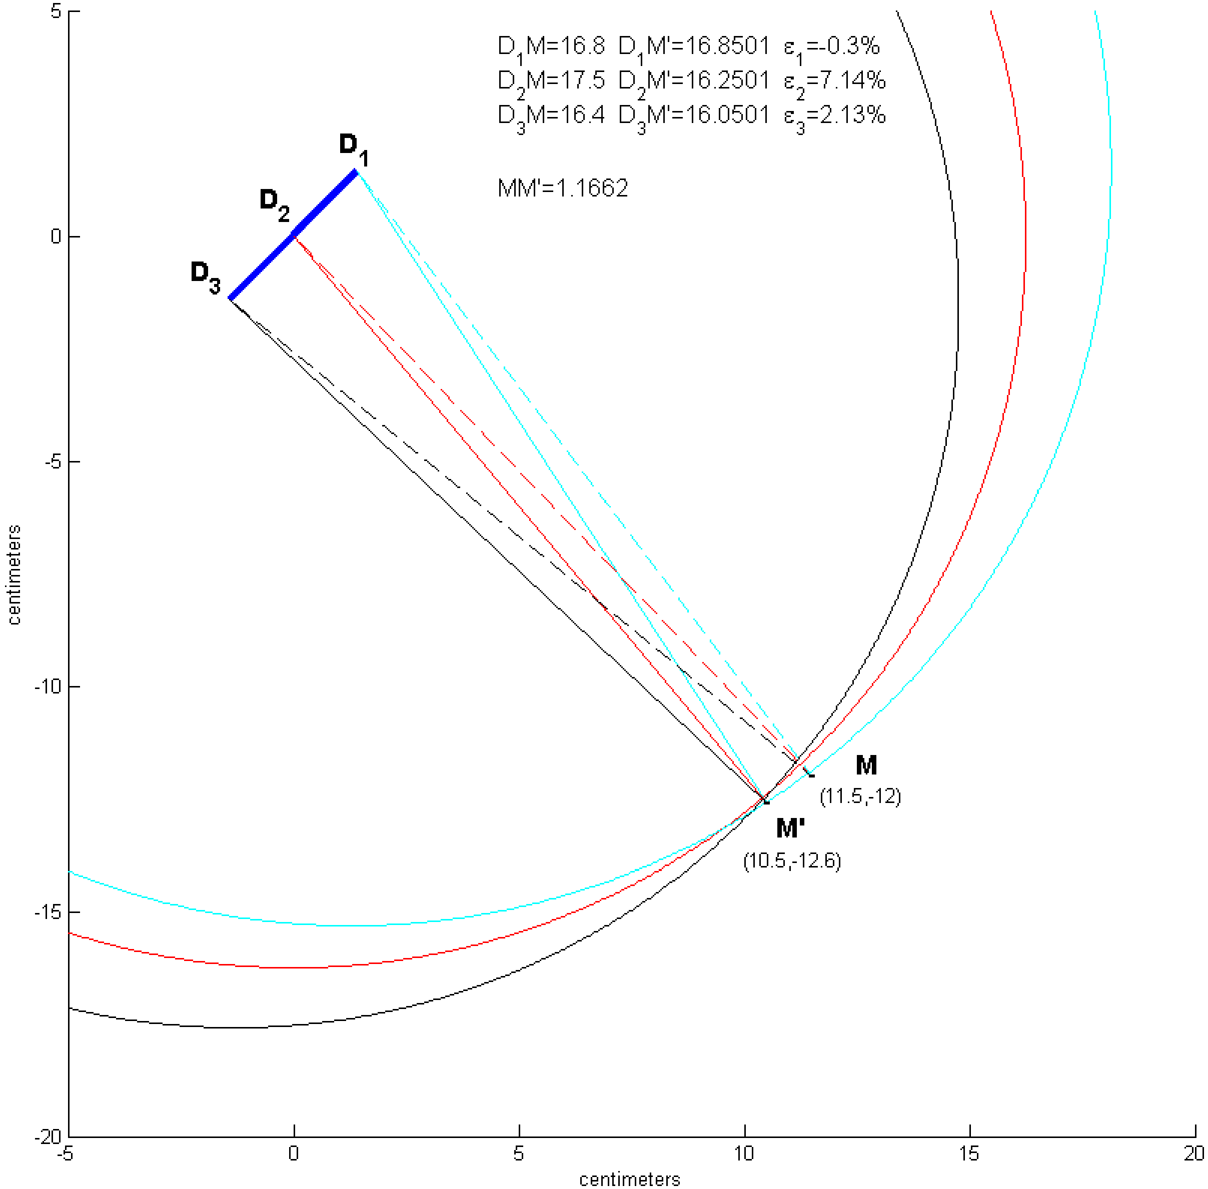

Original and reconstructed scene geometry. M is original position and M' is reconstructed

Looking Around the Corner : Single Patch

Note that the hidden surface is invisible to both the sensor AND the laser

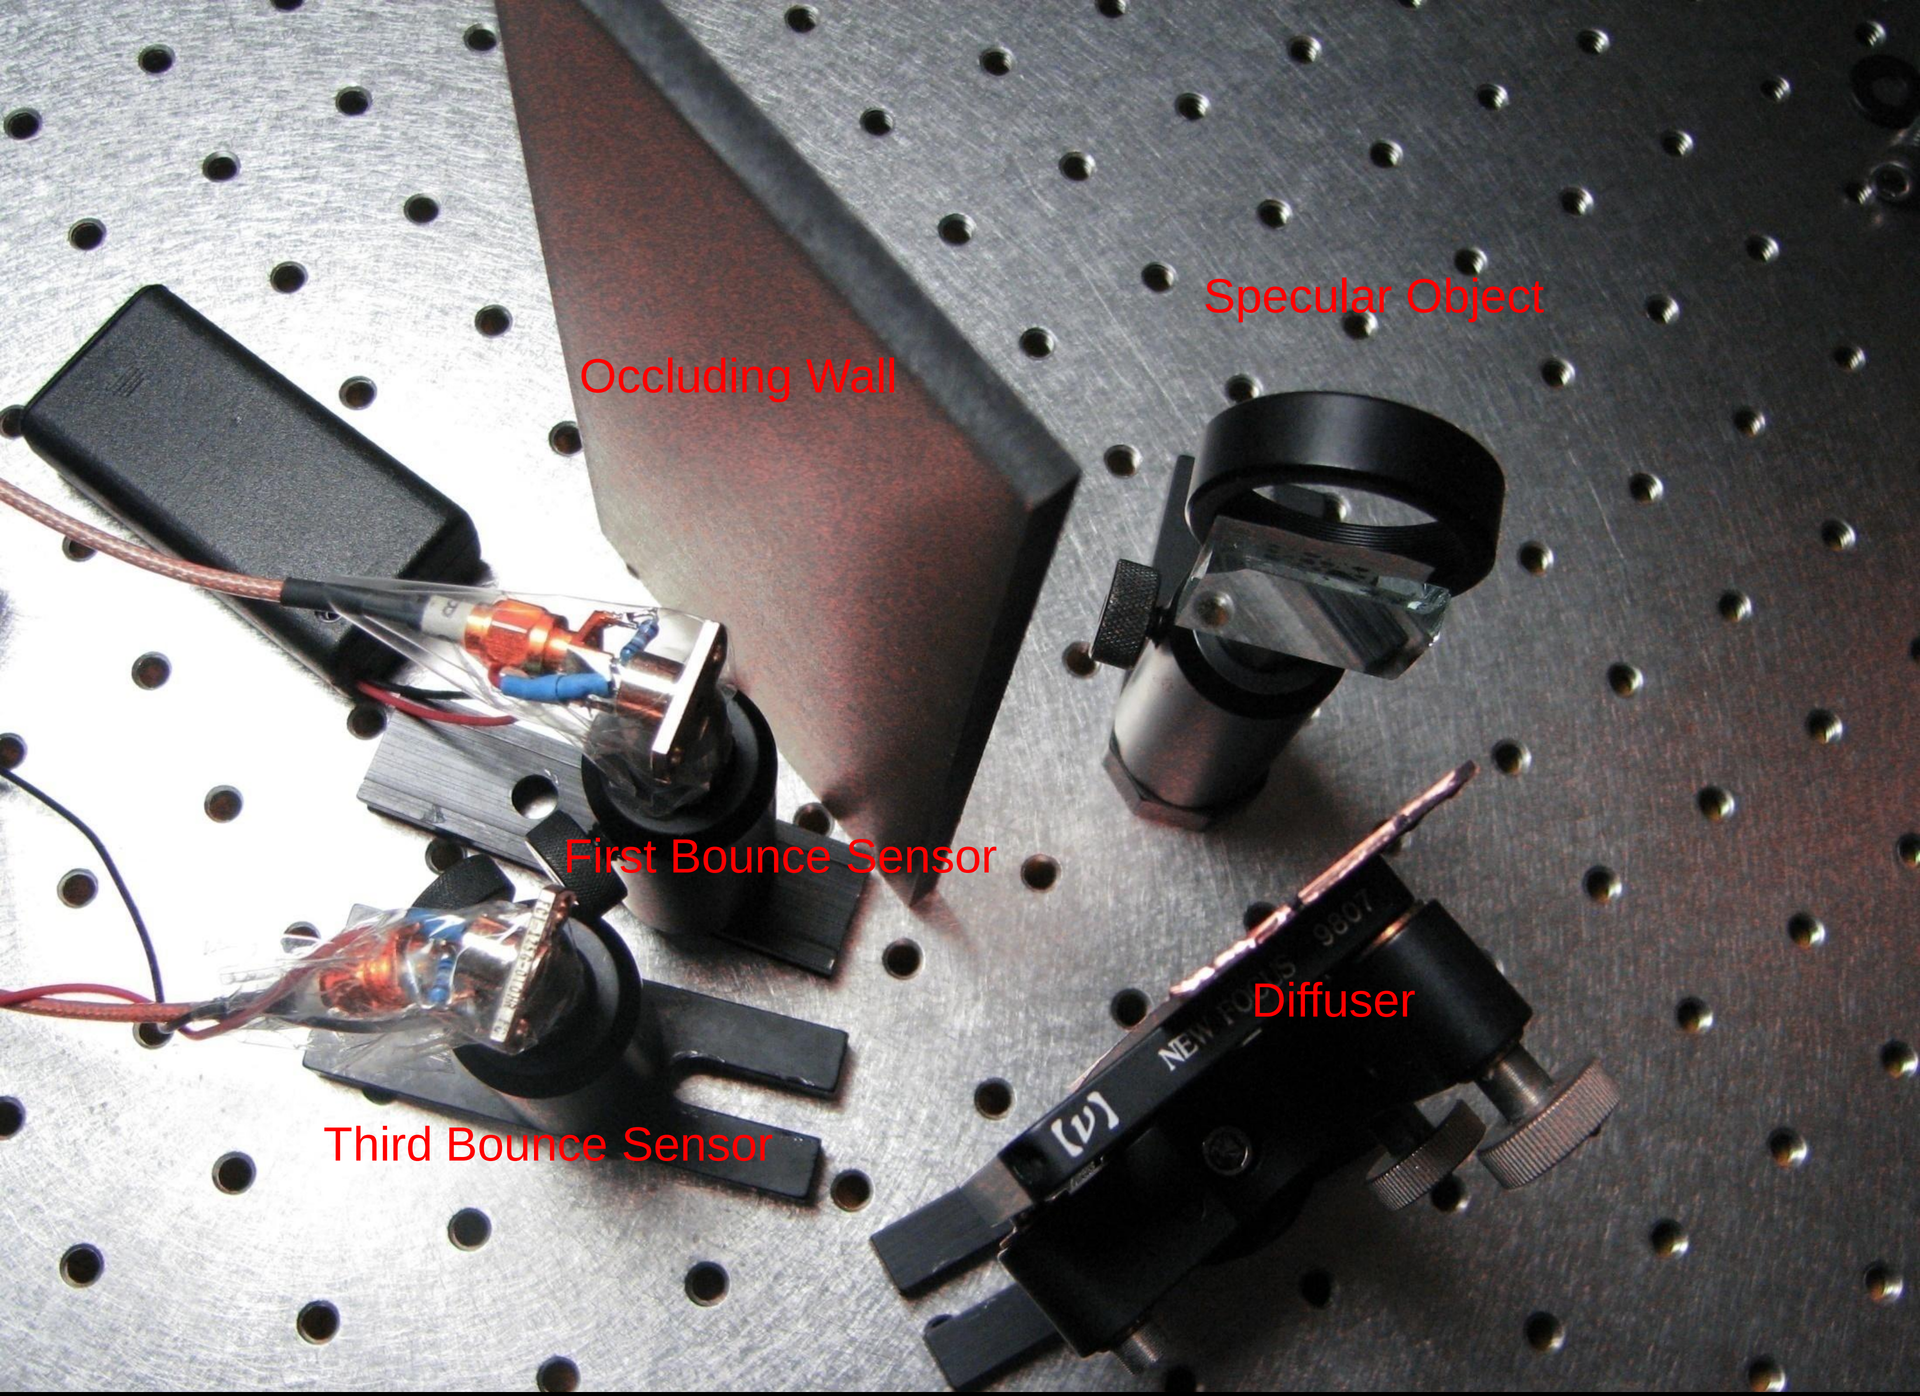

Here we show the setup for one example of a looking around the corner experiment using a single patch.

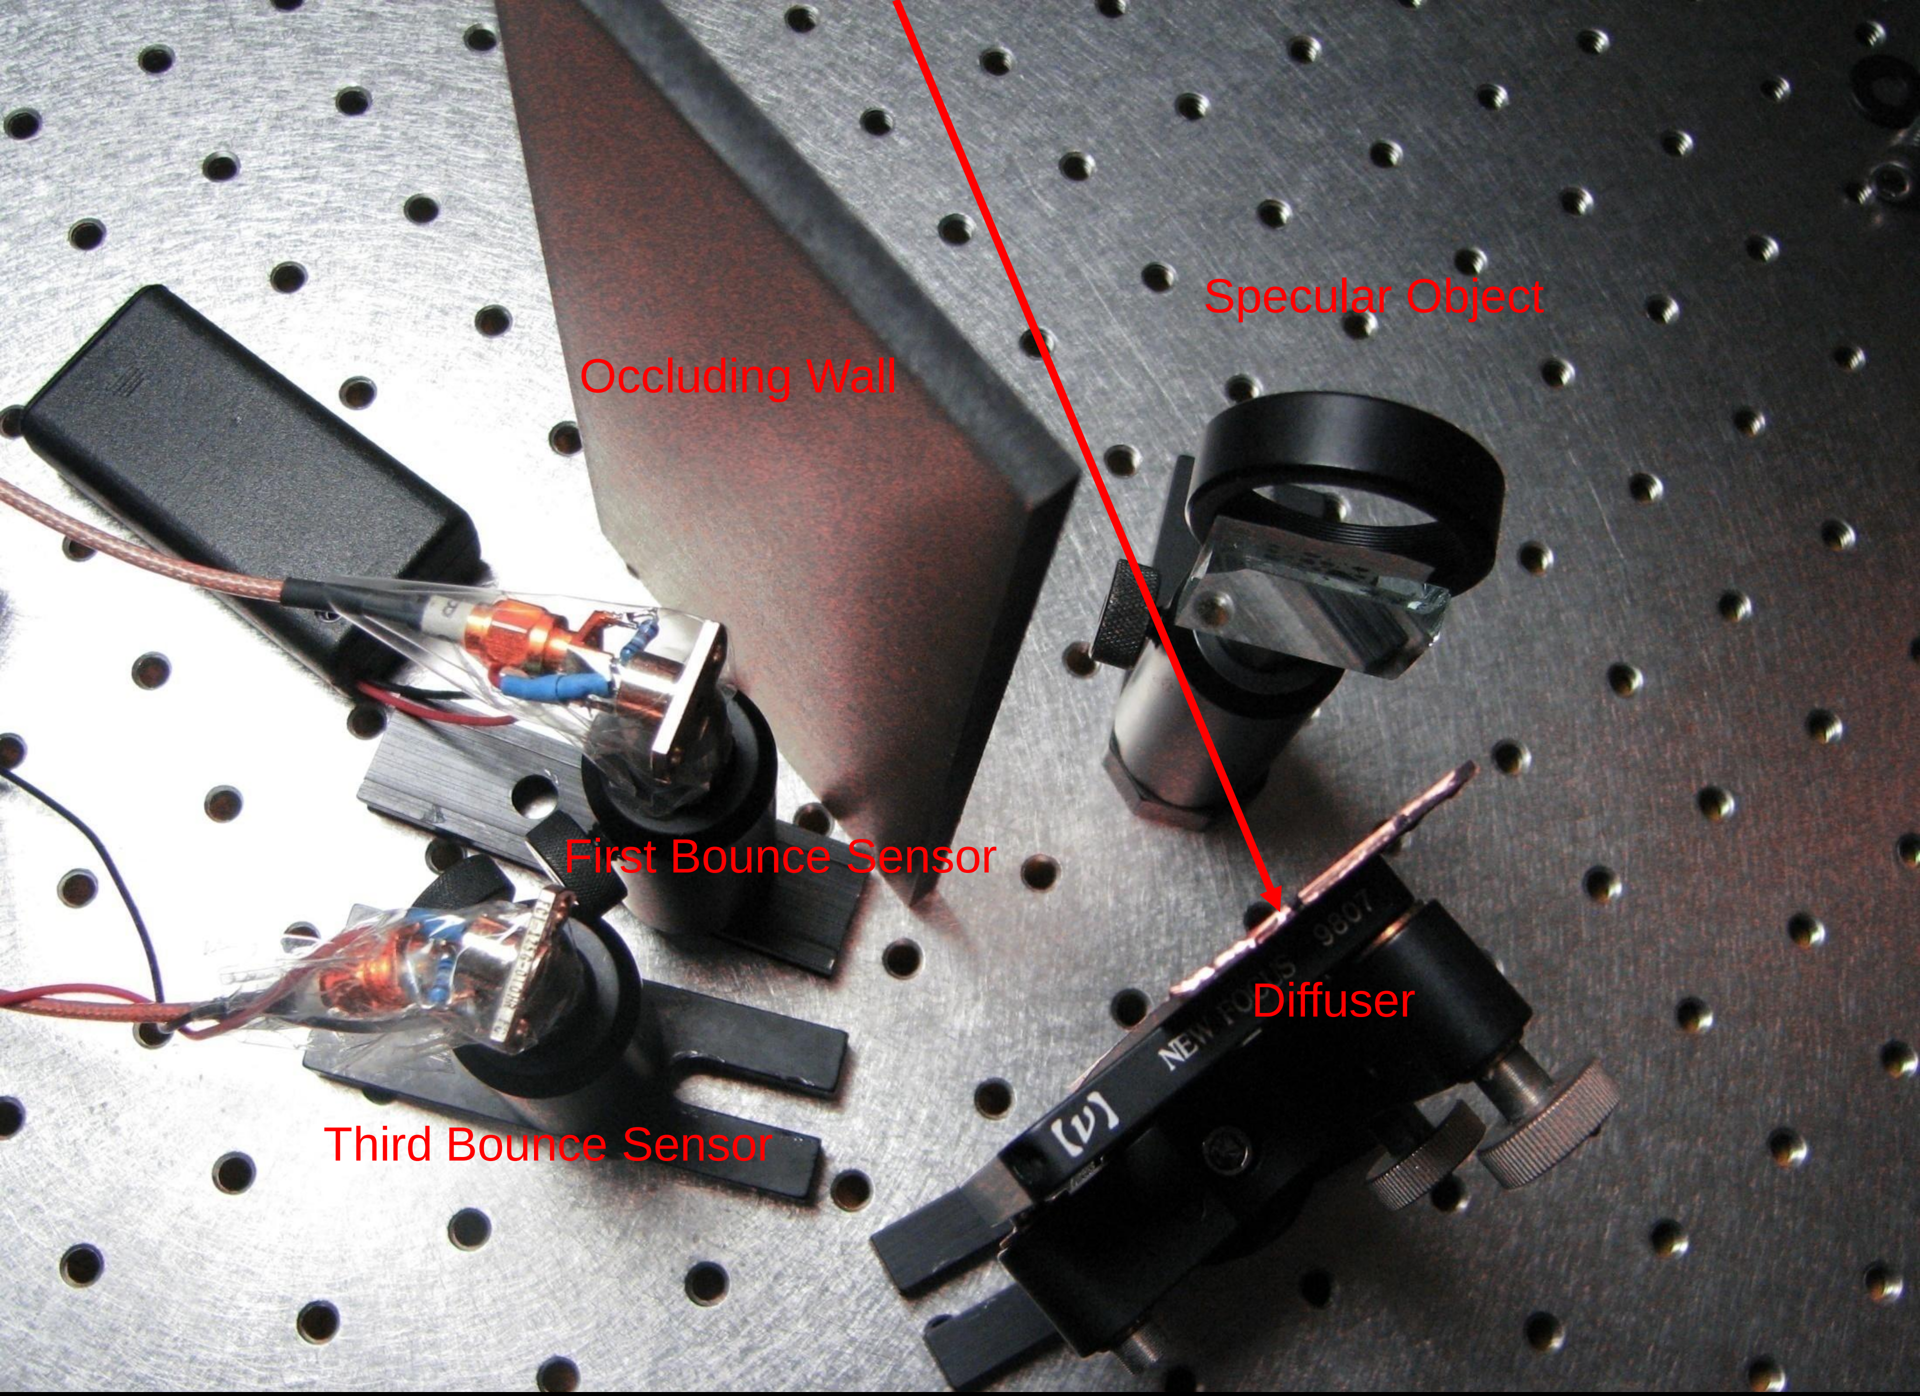

Light enters the scene and strikes a diffuser, which creates a hemisphere of diffuse light.

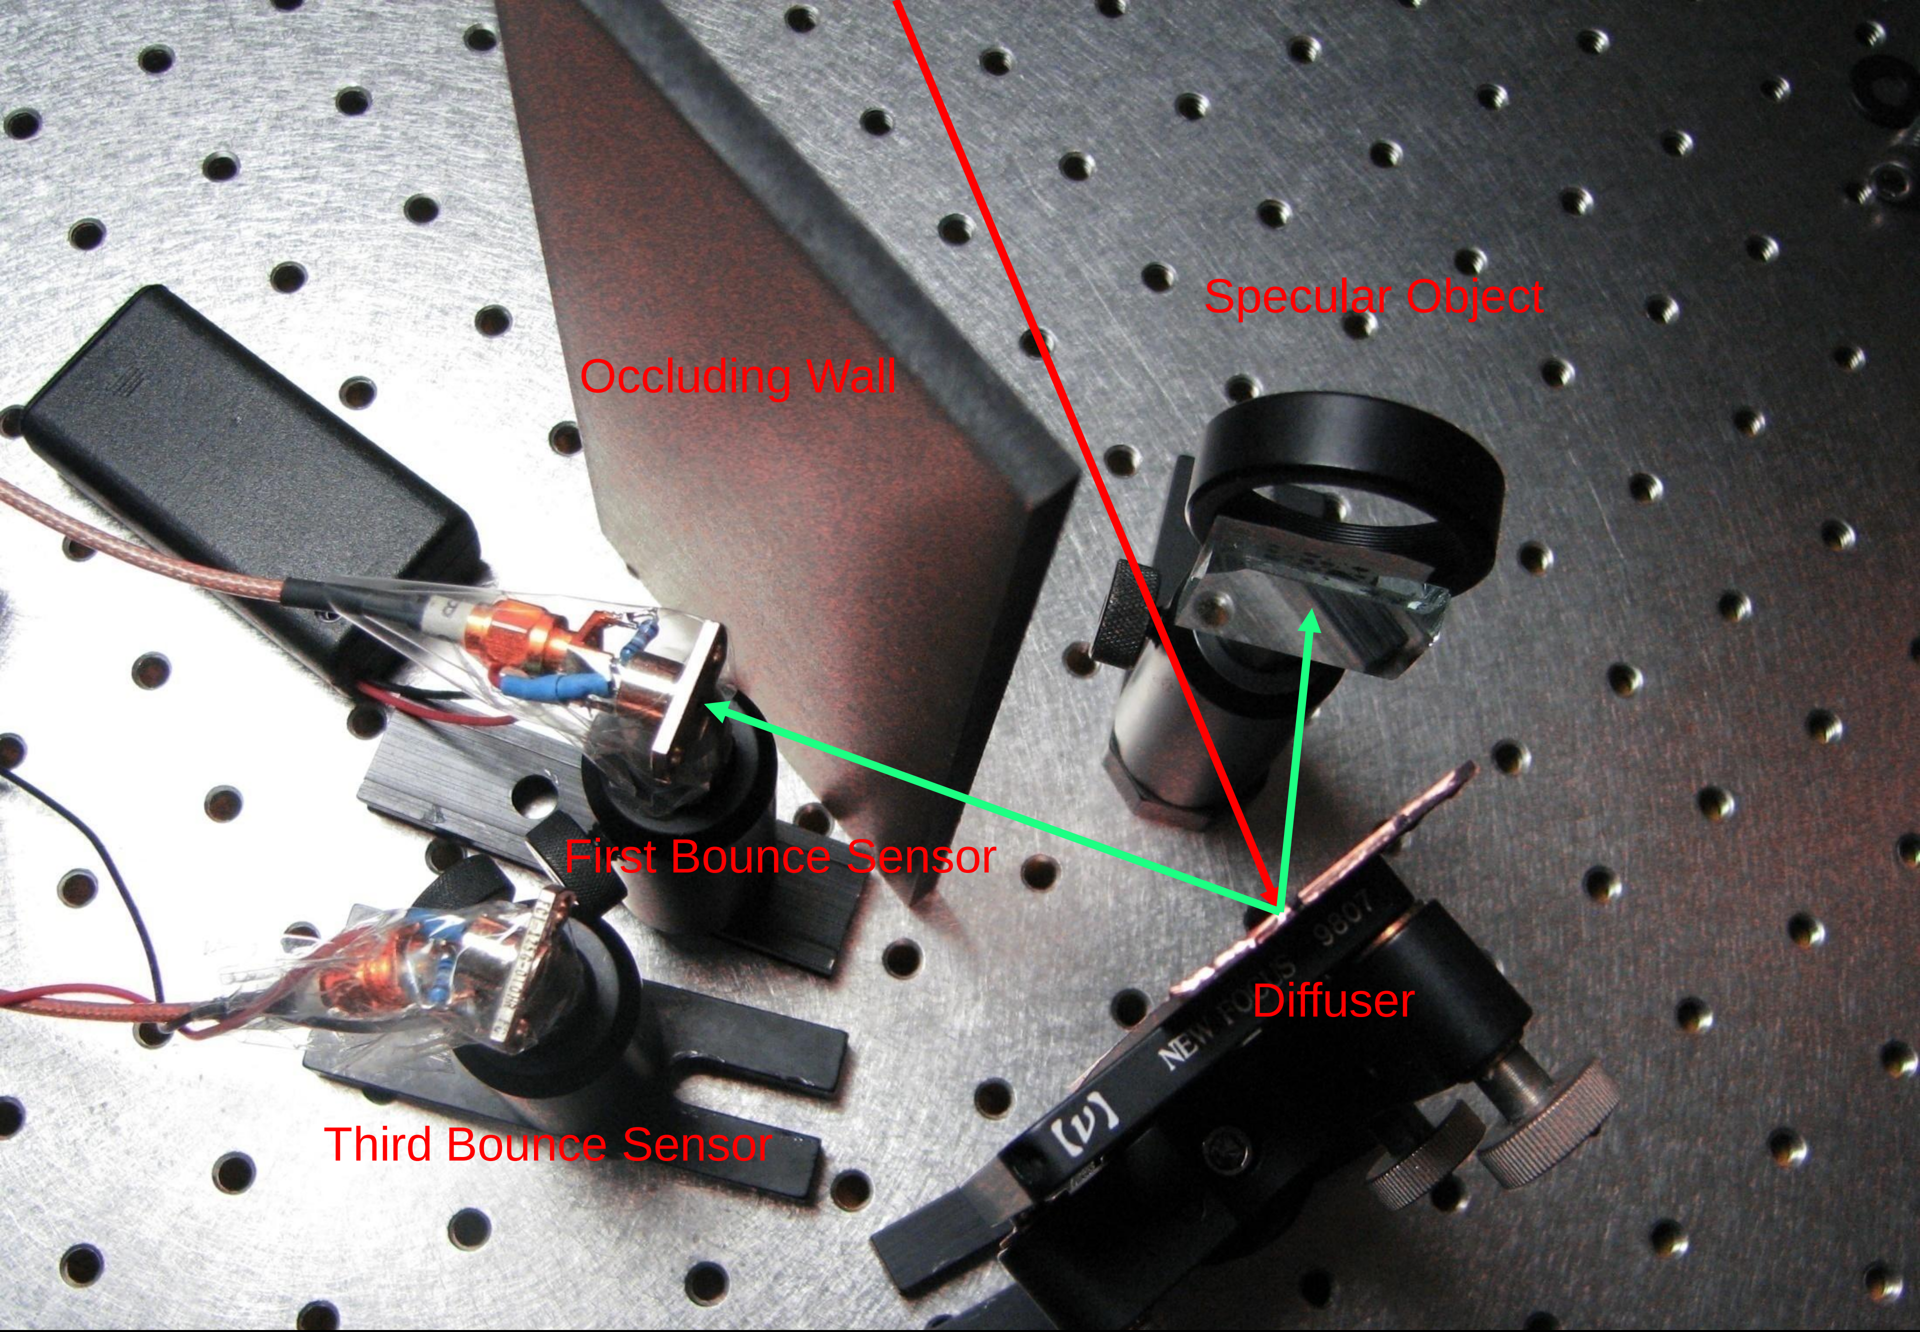

The hemisphere of light is captured by one of the sensors (first bounce shown in red) and also falls on a specular object.

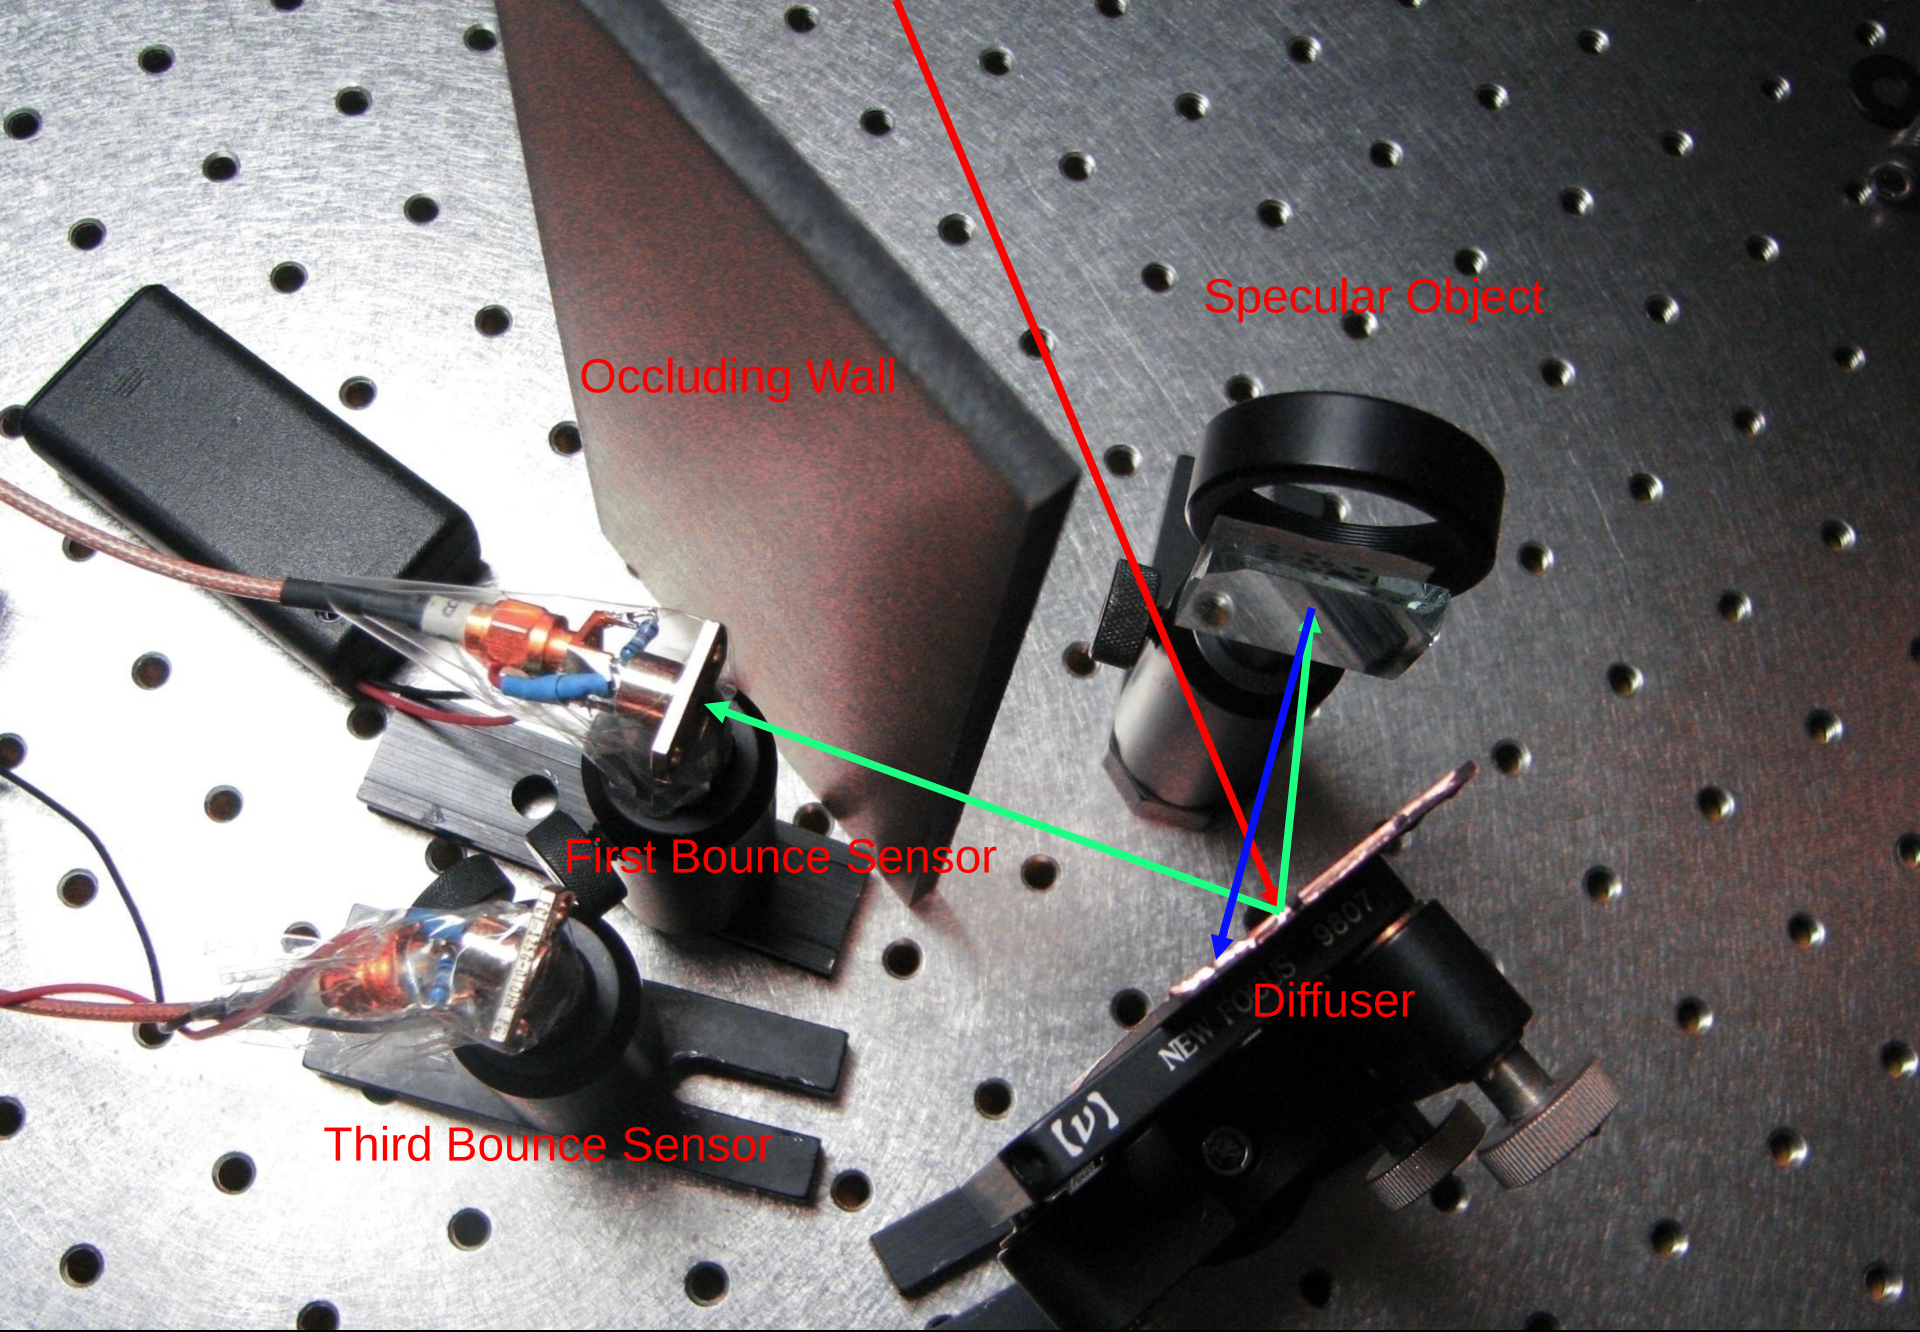

The specular object, hidden to both sensors by the occluding wall, reflects light back to the diffuser (second bounce shown in blue)

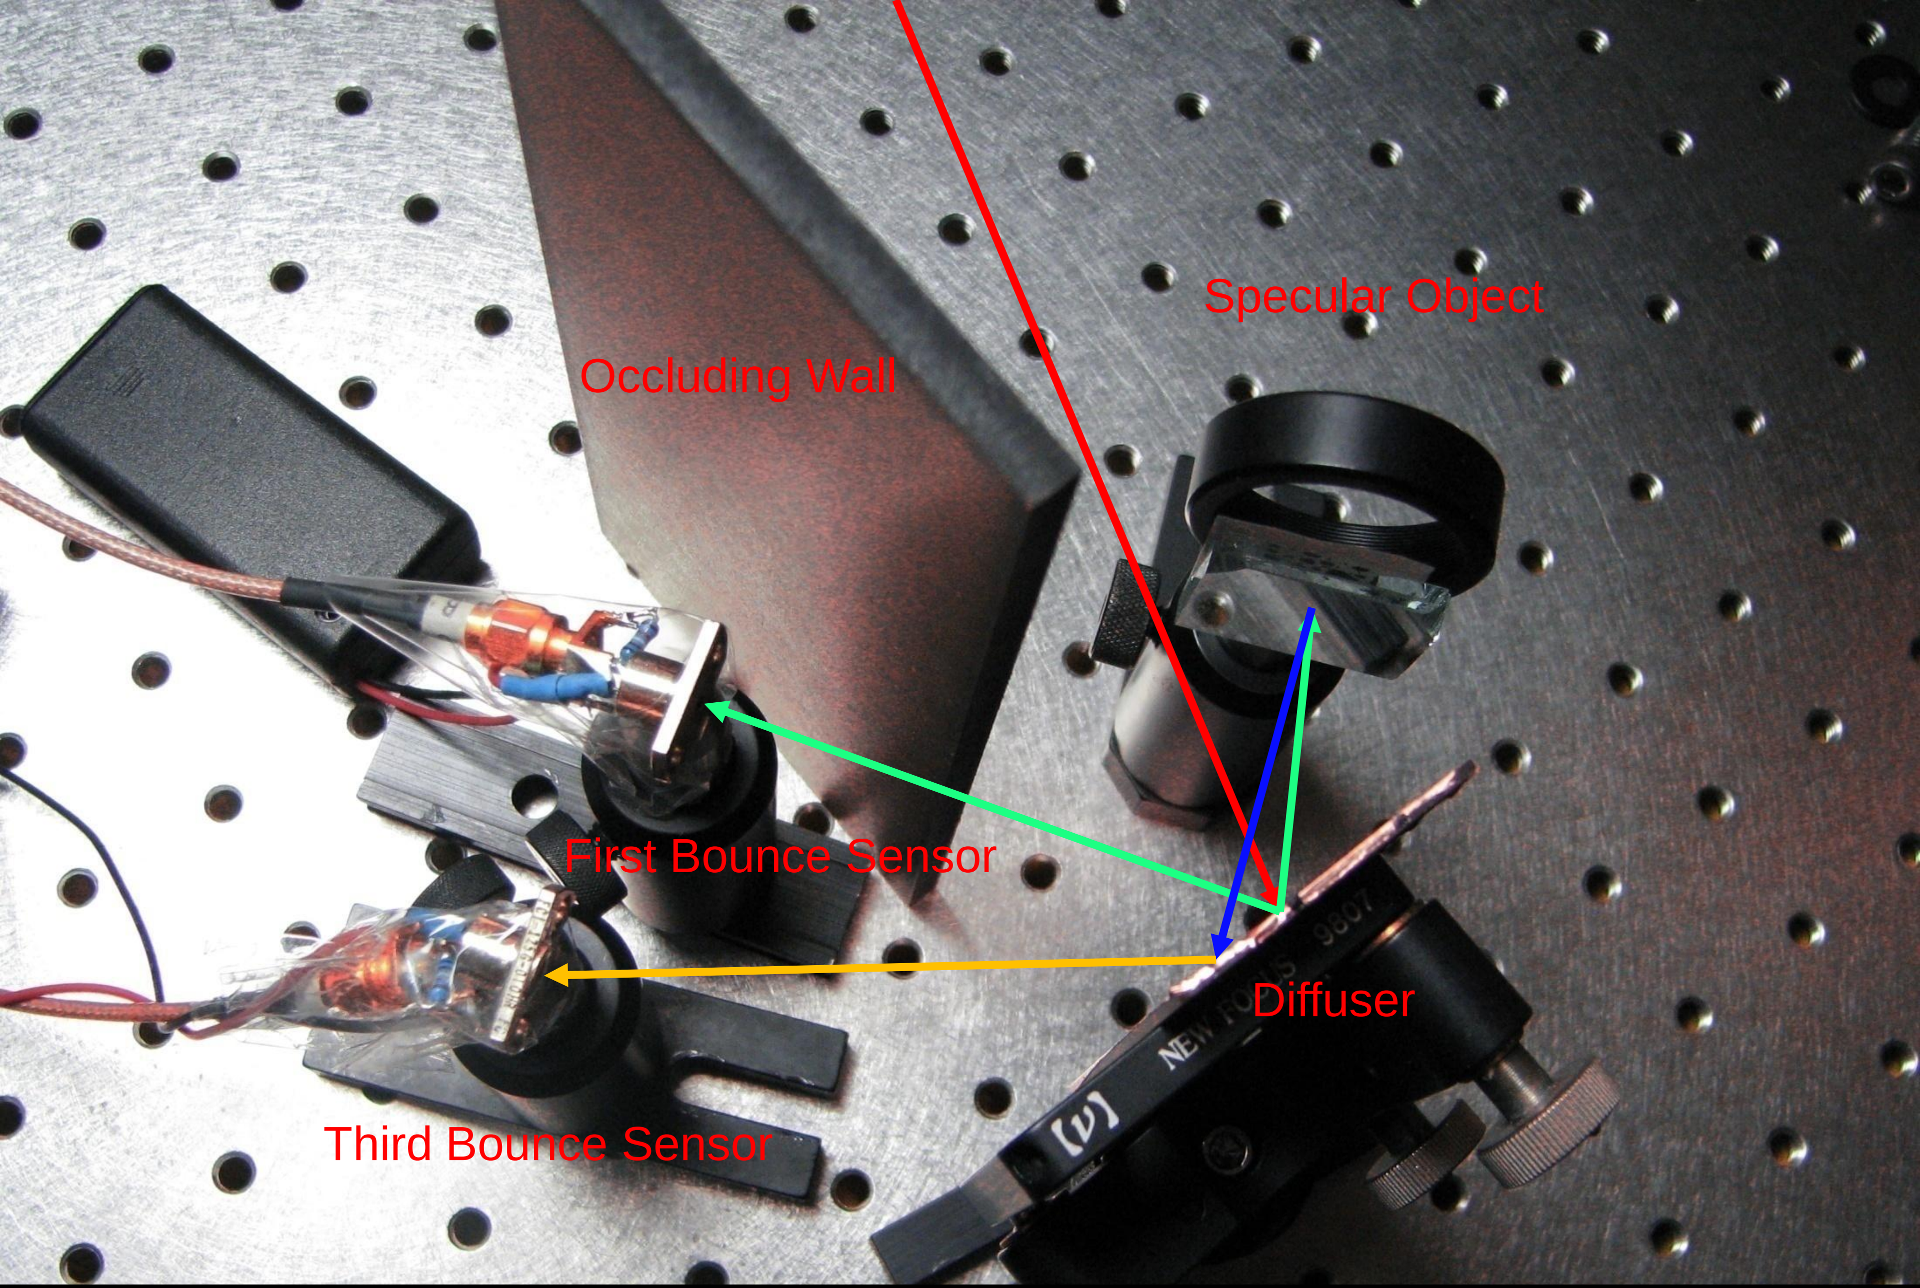

The remaining light which was reflected by the specular object, is finally captured by another sensor (third bounce shown in orange). Two more scans are then completed in the same fashion by adjusted the position of light which first hits the diffuser. This allows for scene geometry triangulation.

Looking Around the Corner: Single Patch

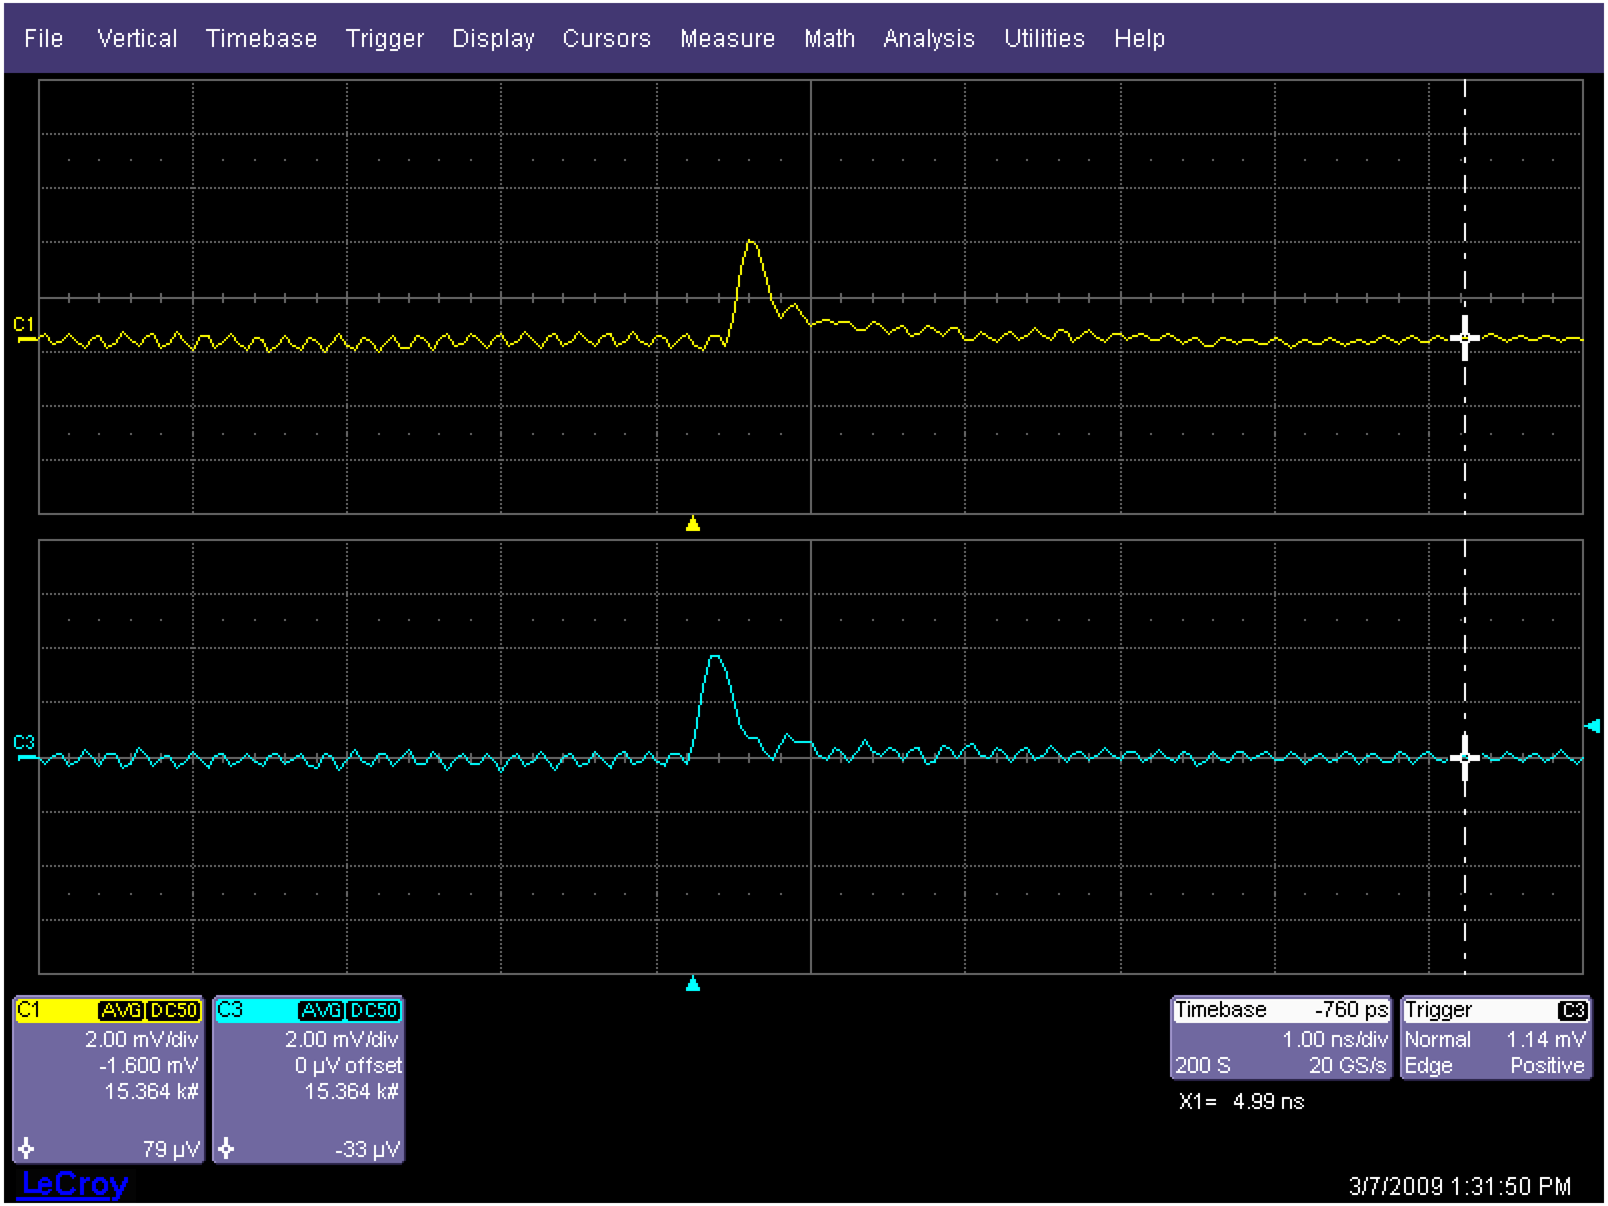

Screenshot of the captured signal on the oscilloscope for the middle raster scan

Looking Around the Corner: Single Patch Experiment

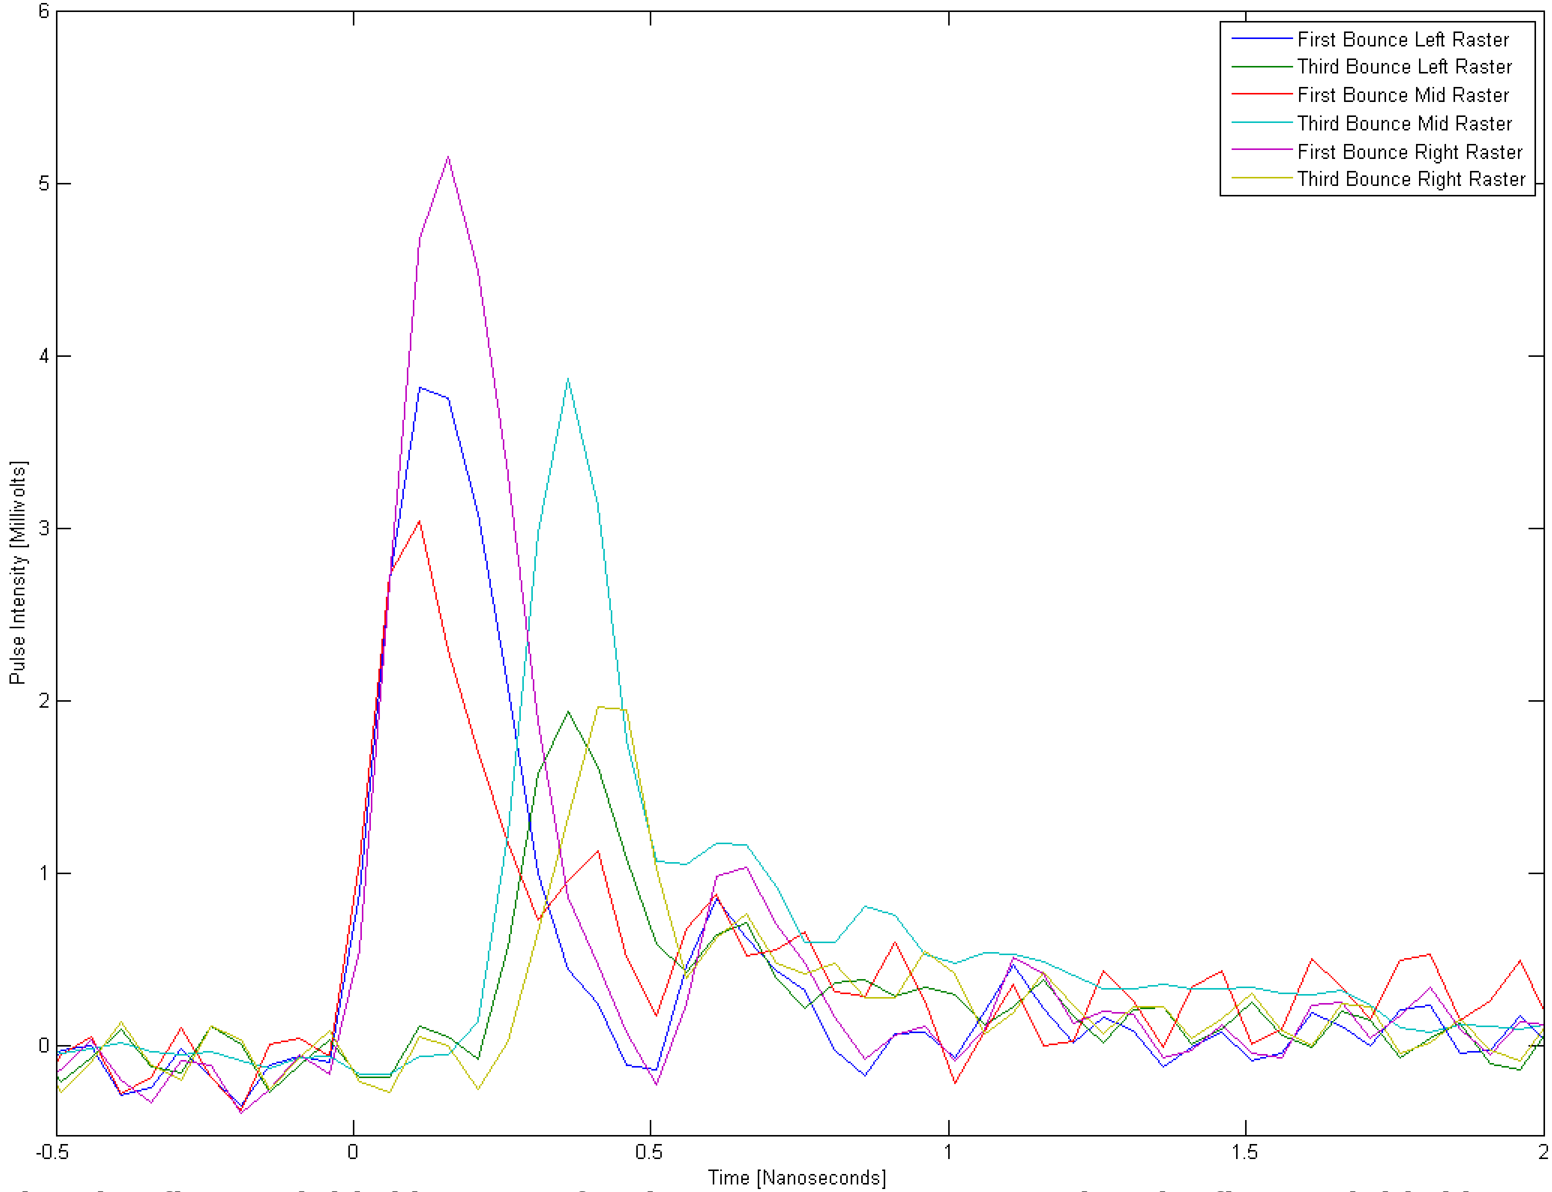

Plot showing first and third bounces for three raster scans. Note that the first and third bounces corresponding to the three different diffuser patches are shifted in time. This gives us the required distance information

Looking Around the Corner : 101 Barcode

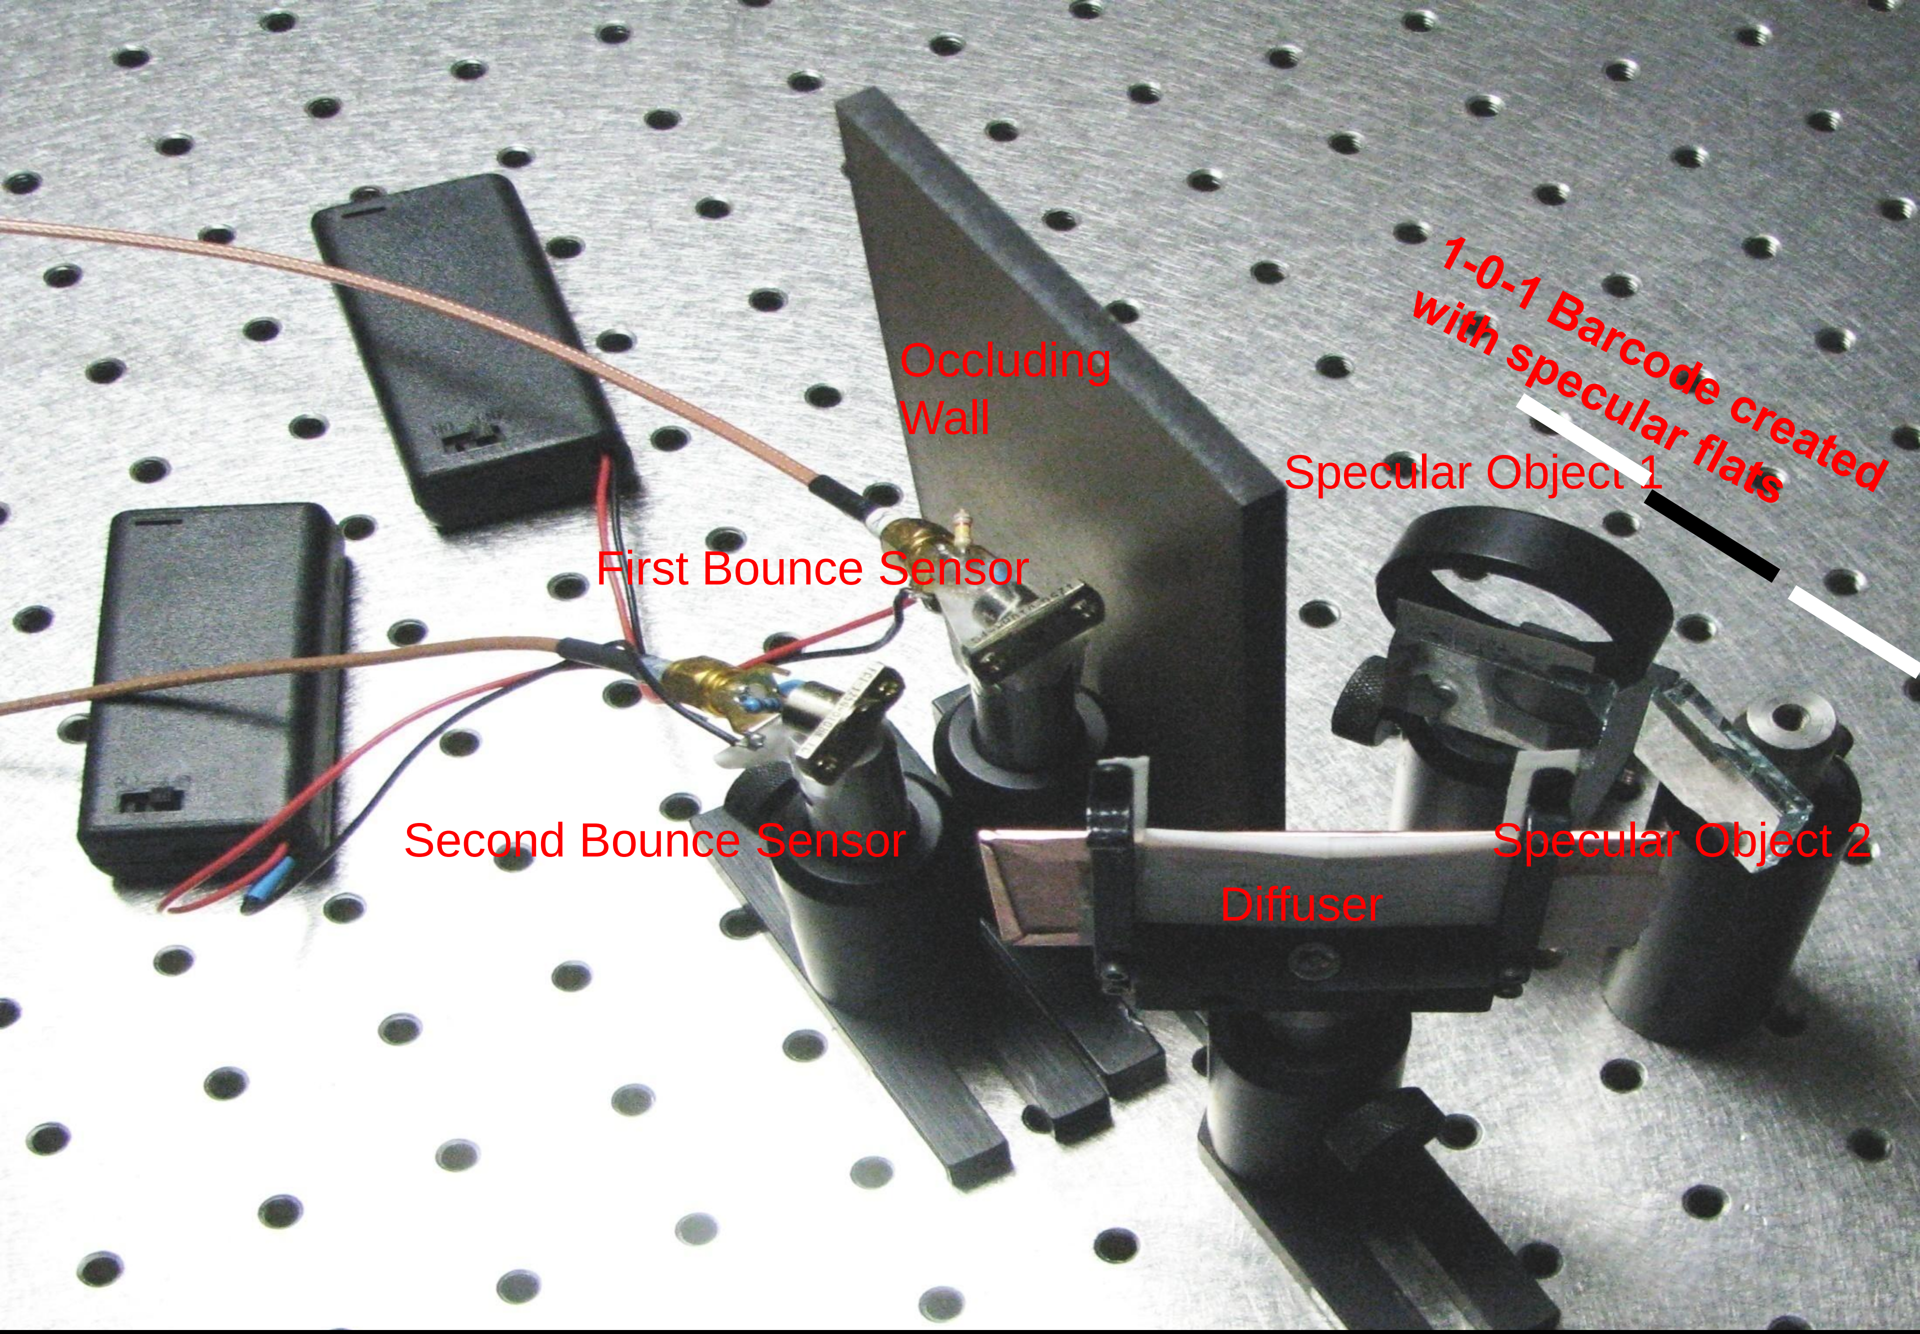

Here we show another looking around the corner experiment using two specular objects, which are representative of a 101 barcode

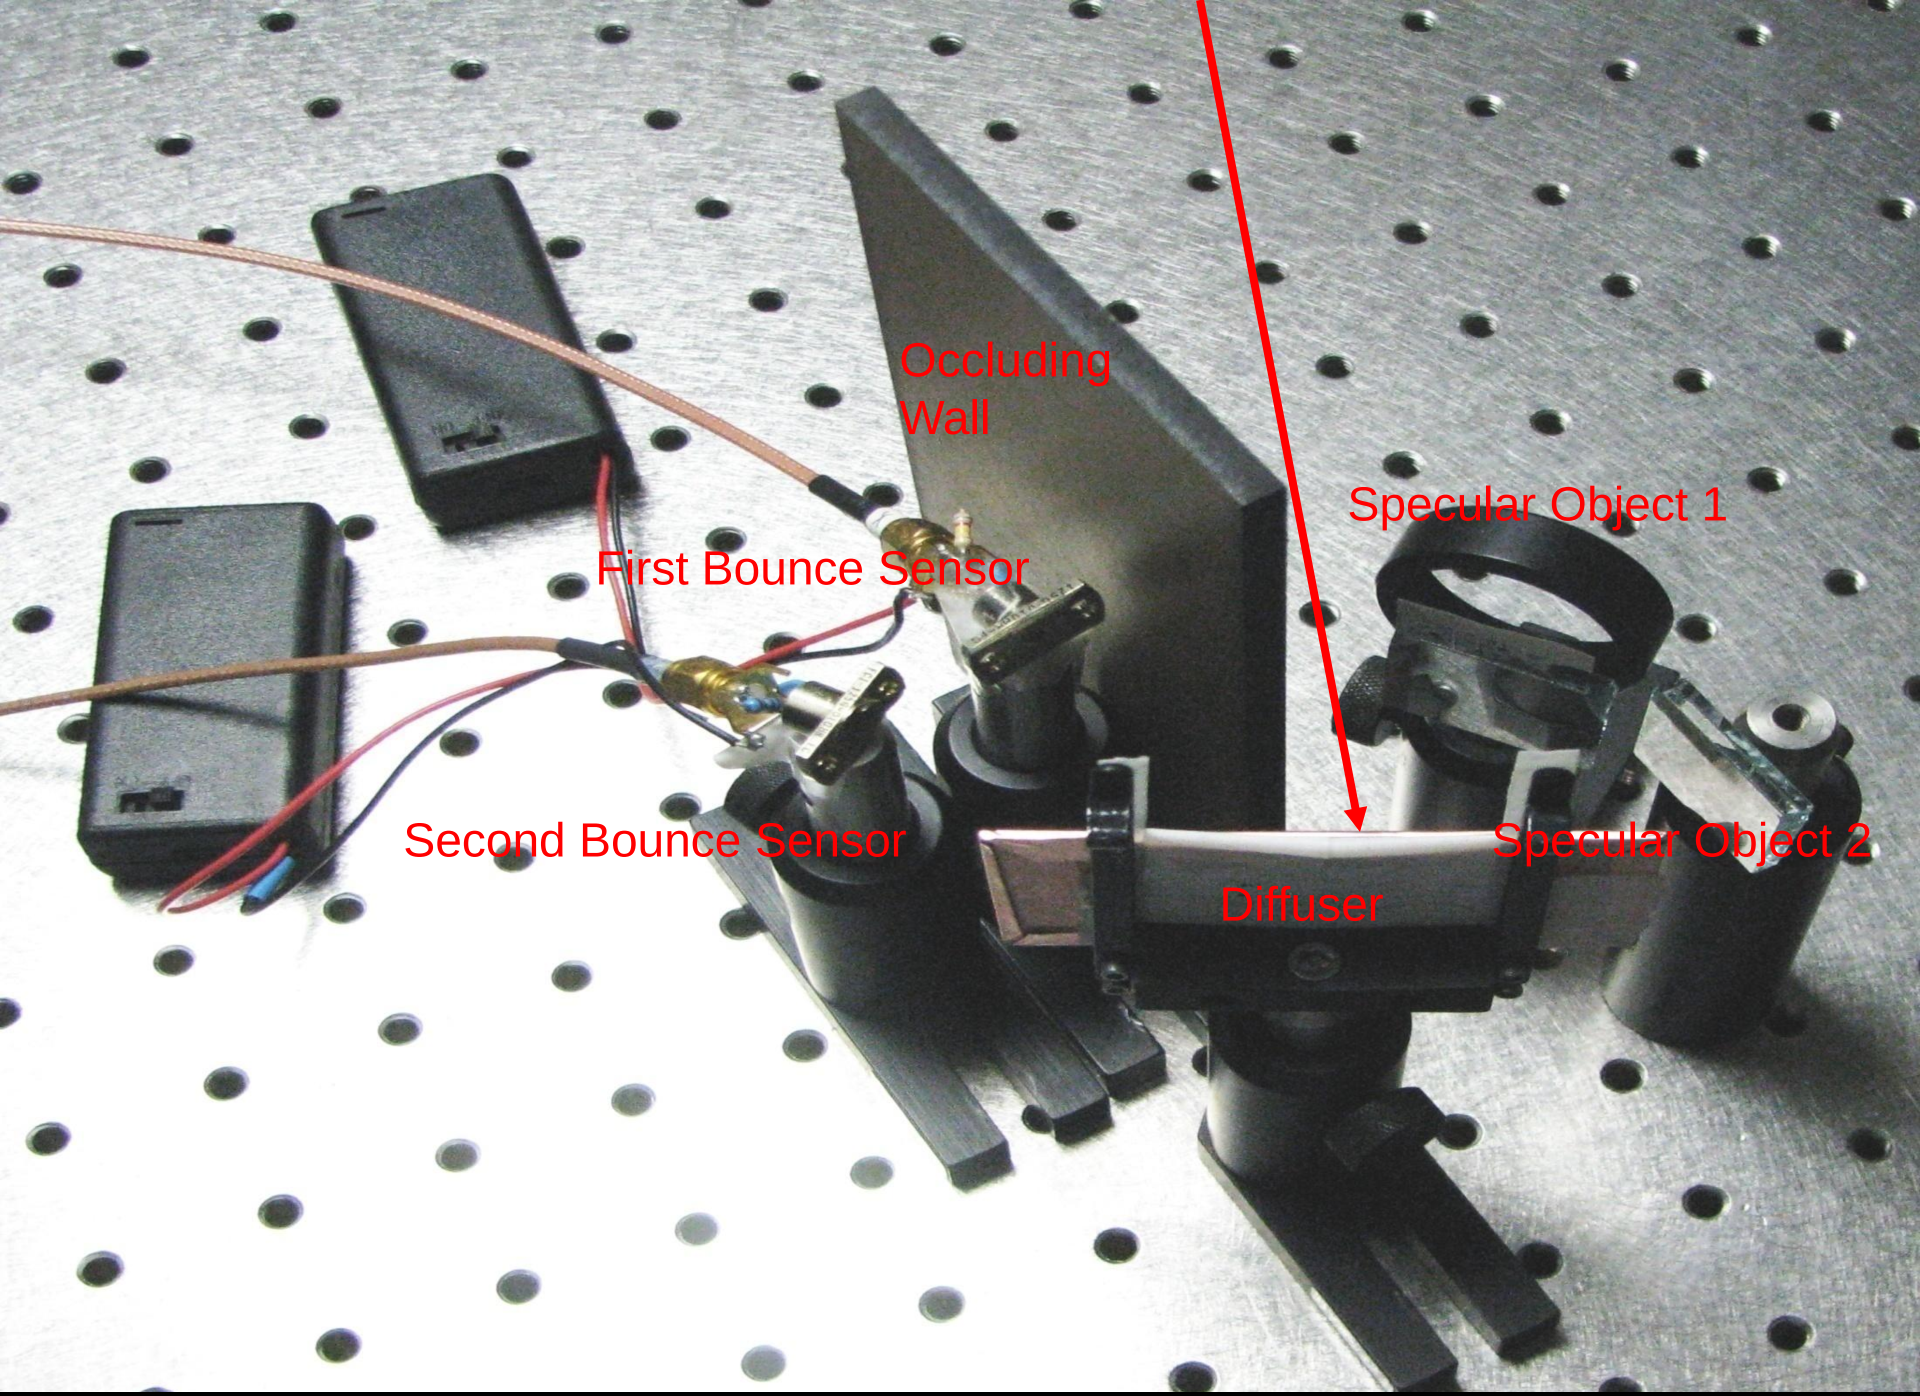

Light enters the scene and strikes a diffuser, which creates a hemisphere of diffuse light.

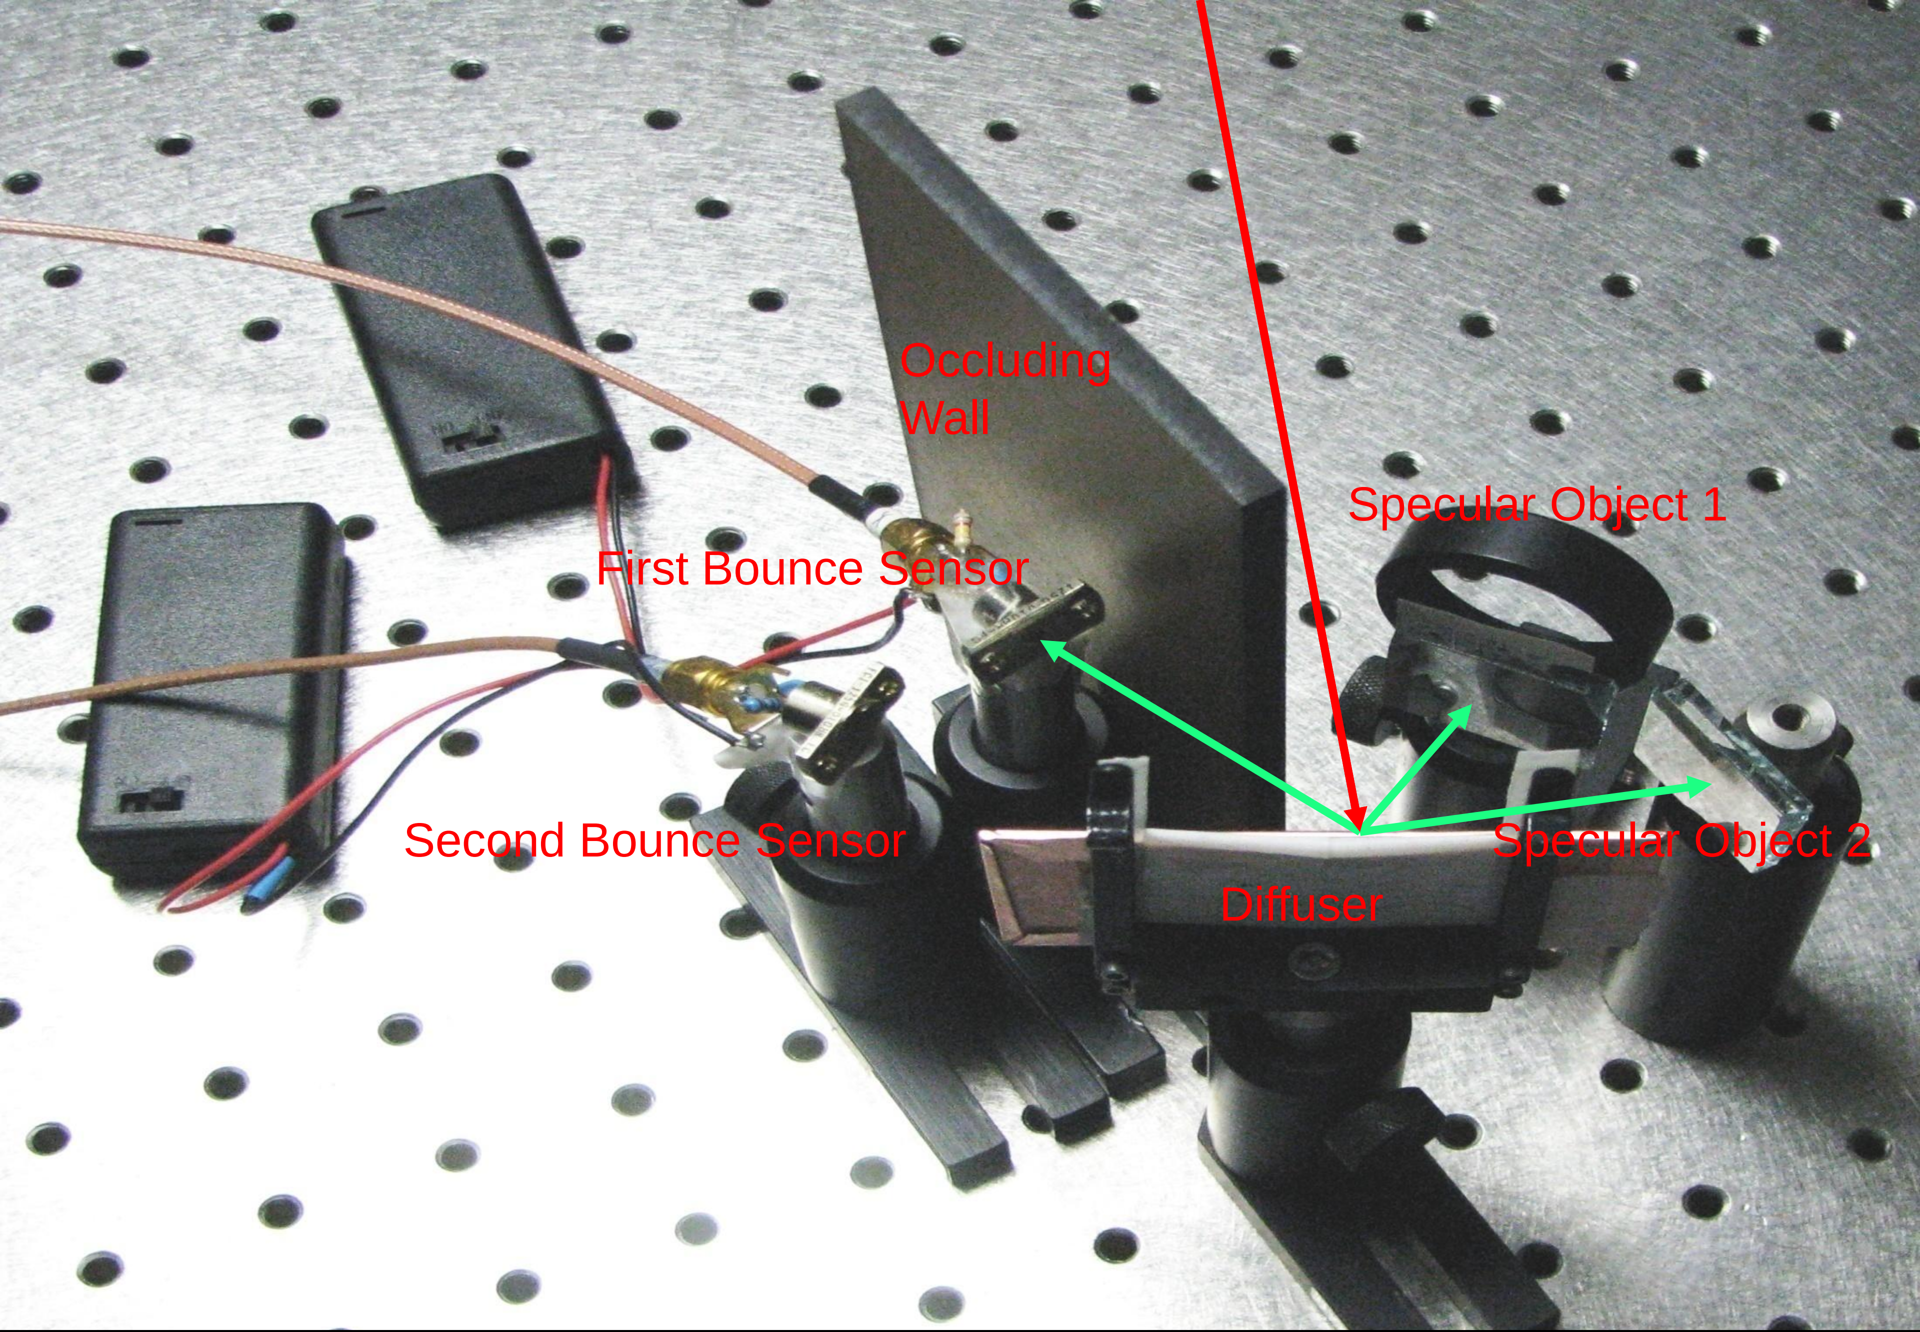

Light from the diffuser is captured by one of the sensors (first bounce shown in green). Light from the diffuser also falls upon two specular objects.

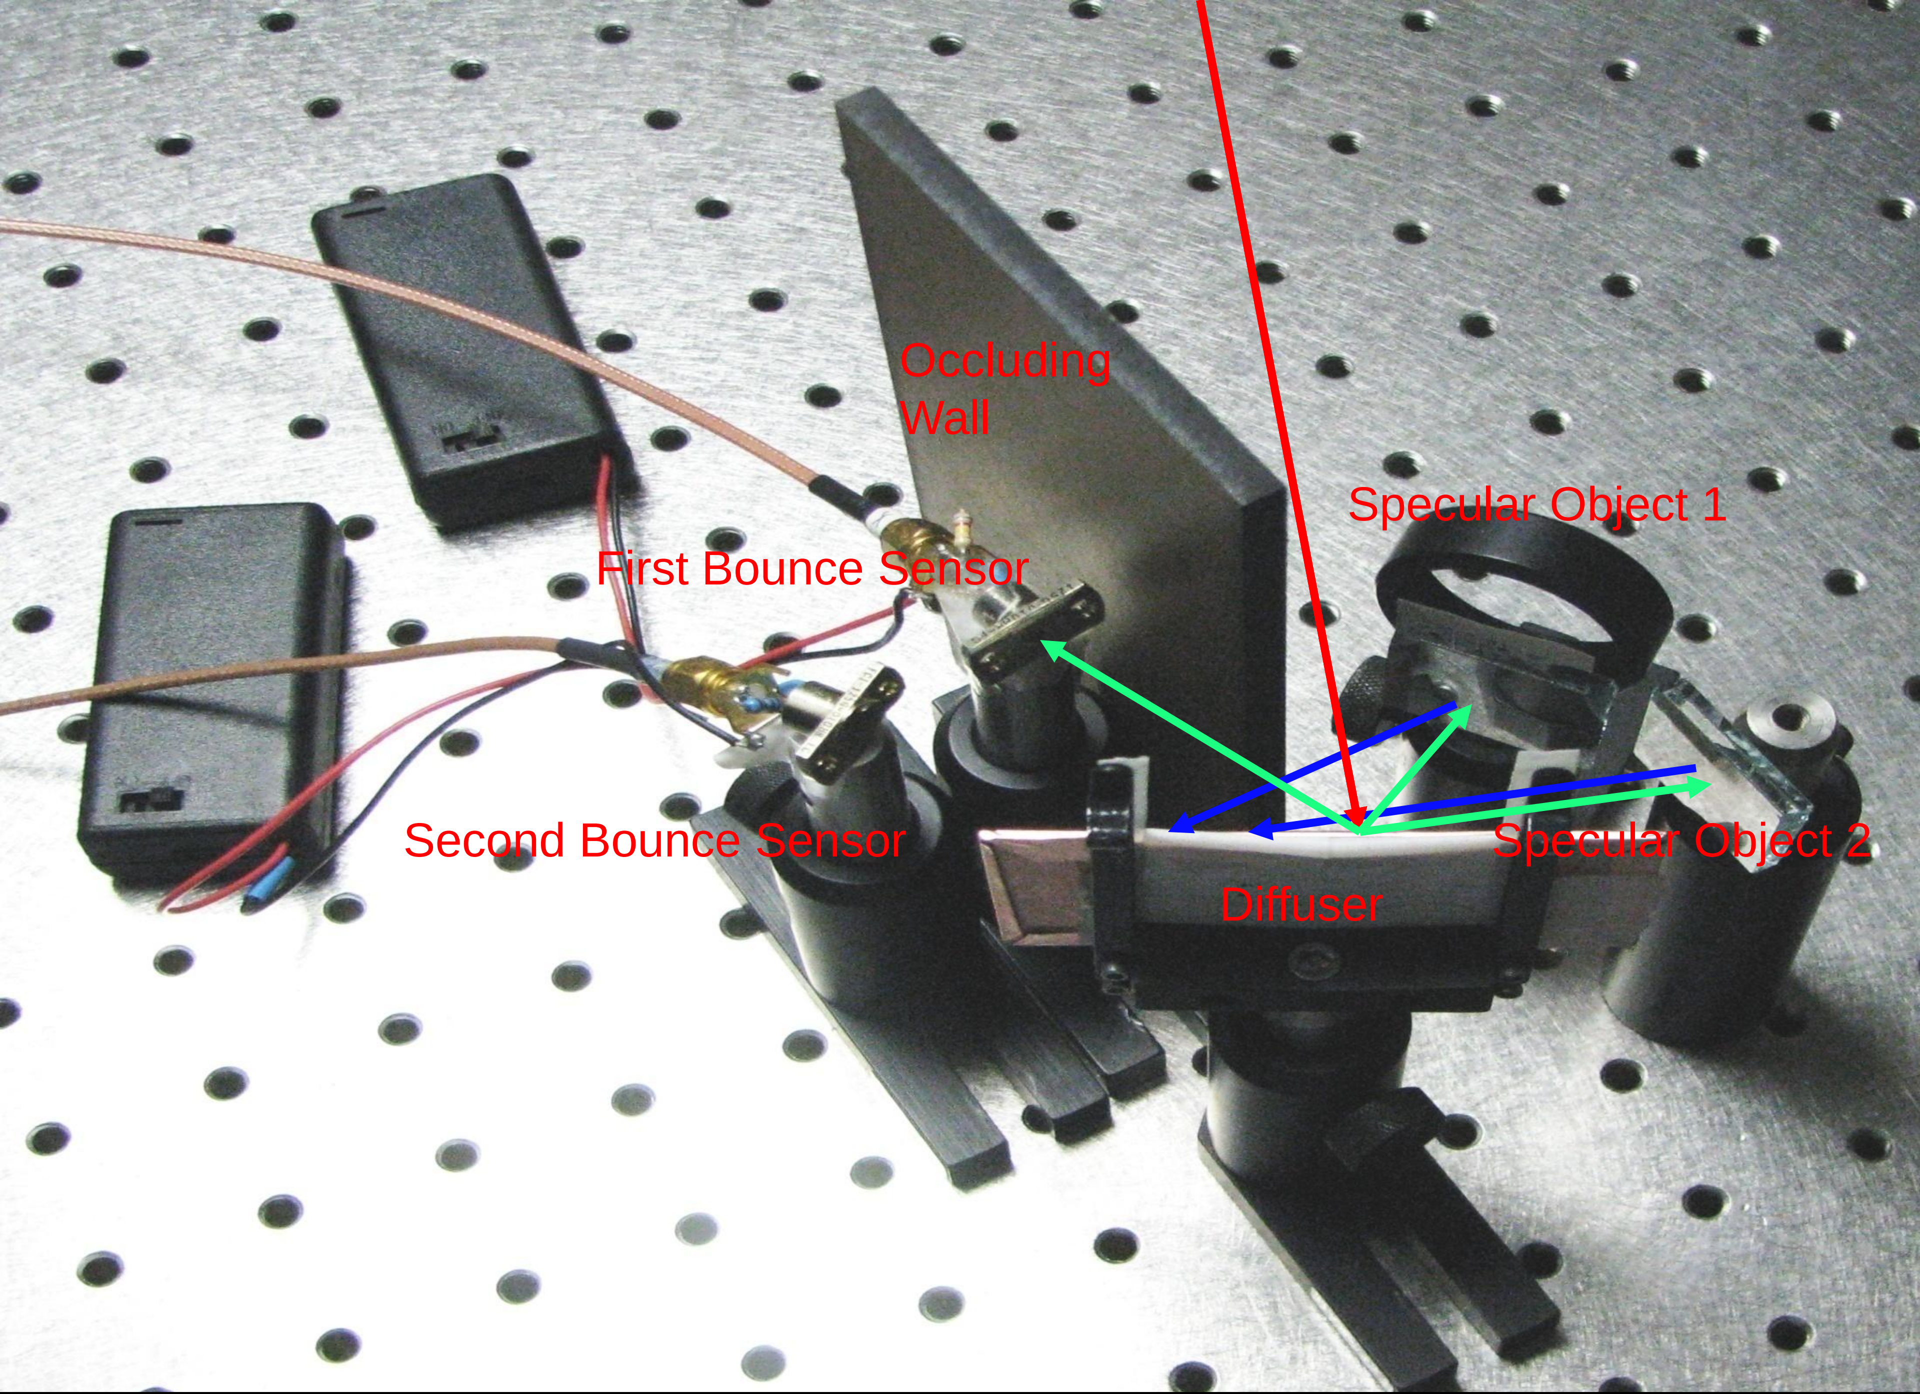

The specular objects reflect light back towards the diffuser (second bounce shown in blue).

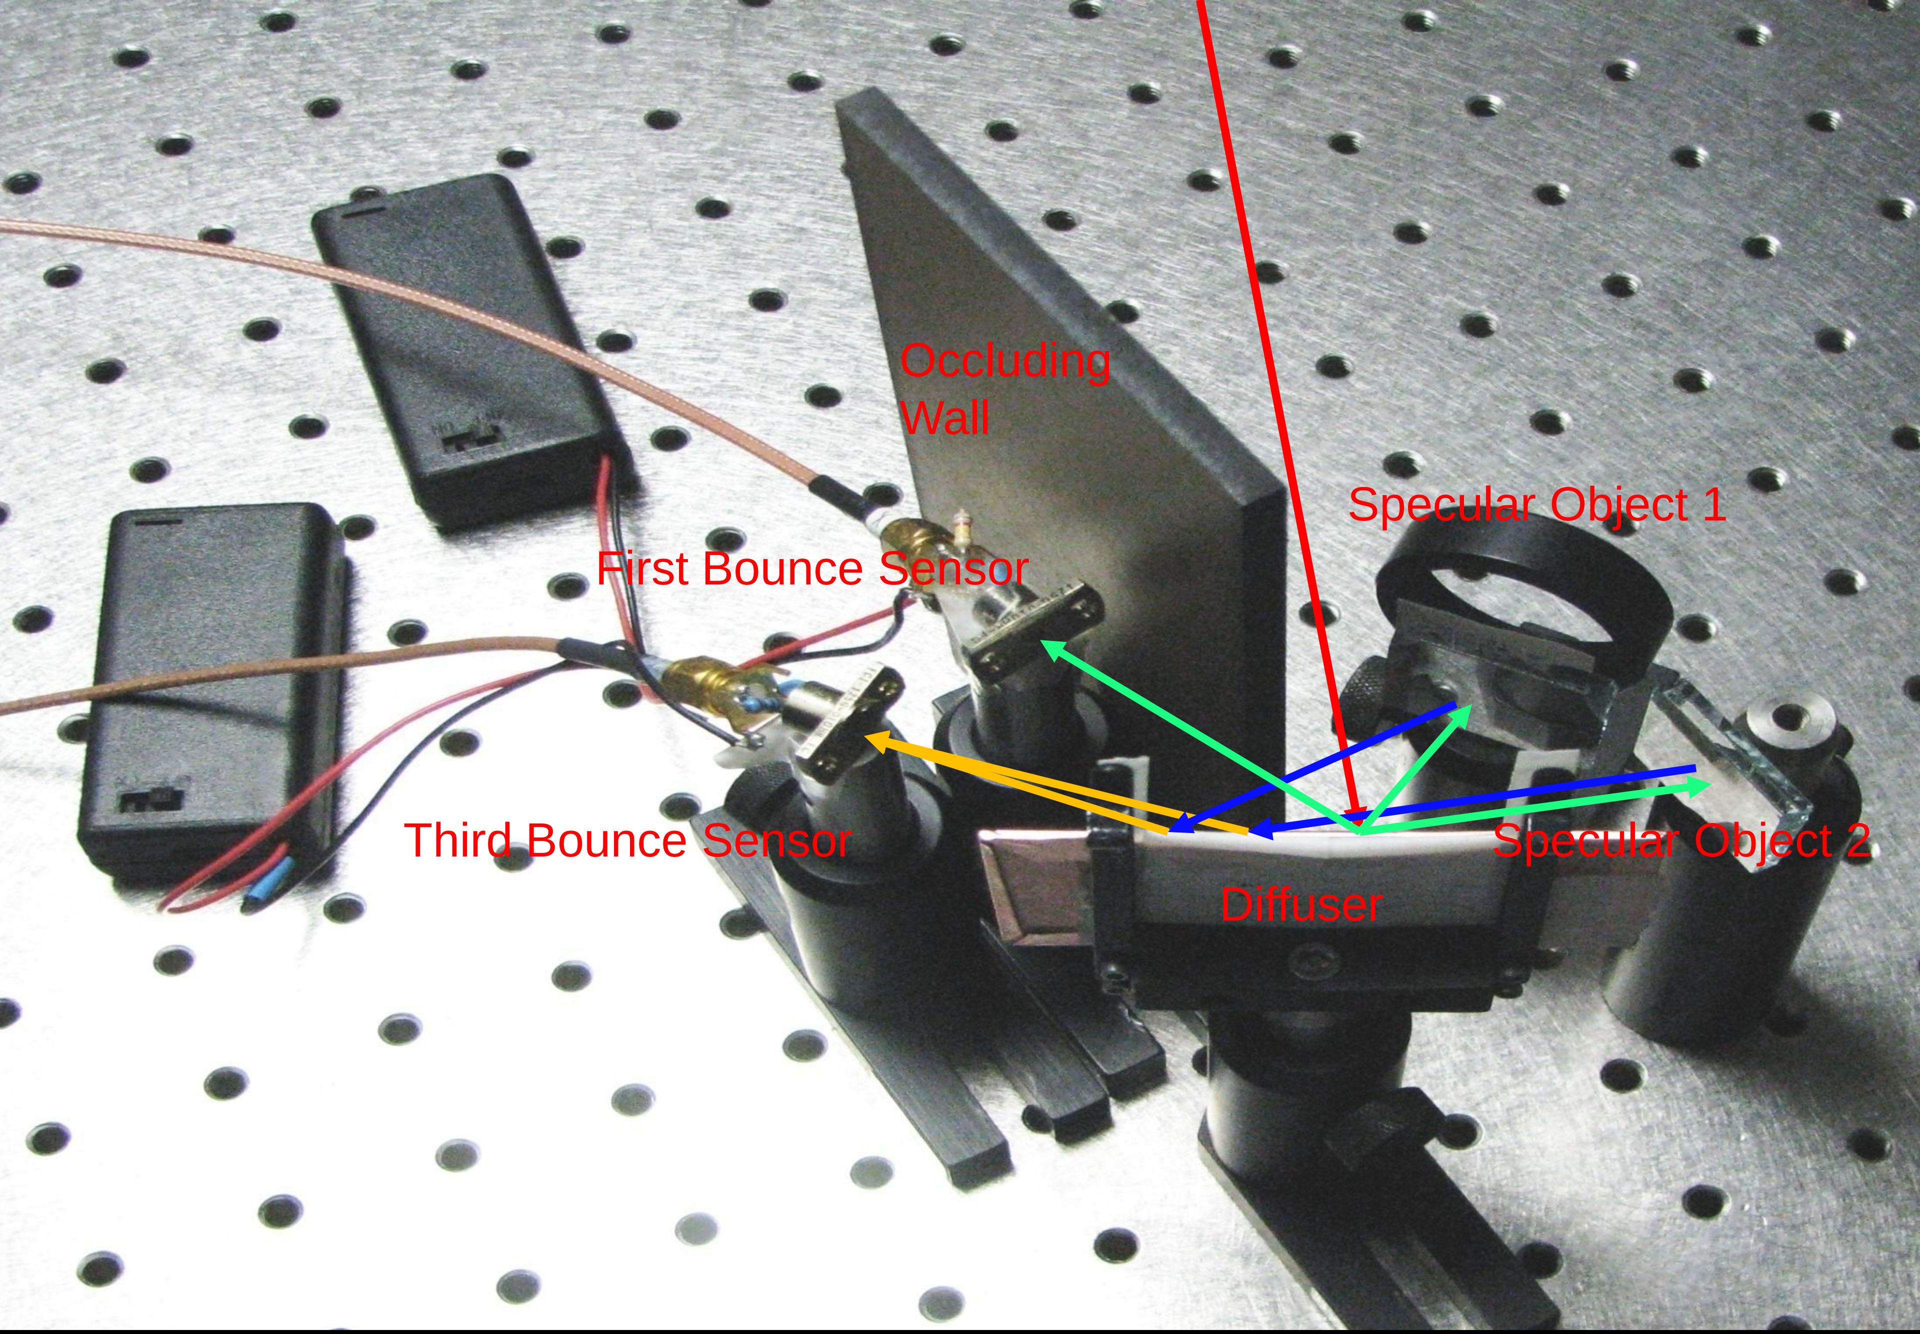

The distance between the two patches results in two time separated signals which are registered by the other sensor (third bounce shown in orange).

Looking Around the Corner: 101 Barcode Experiment

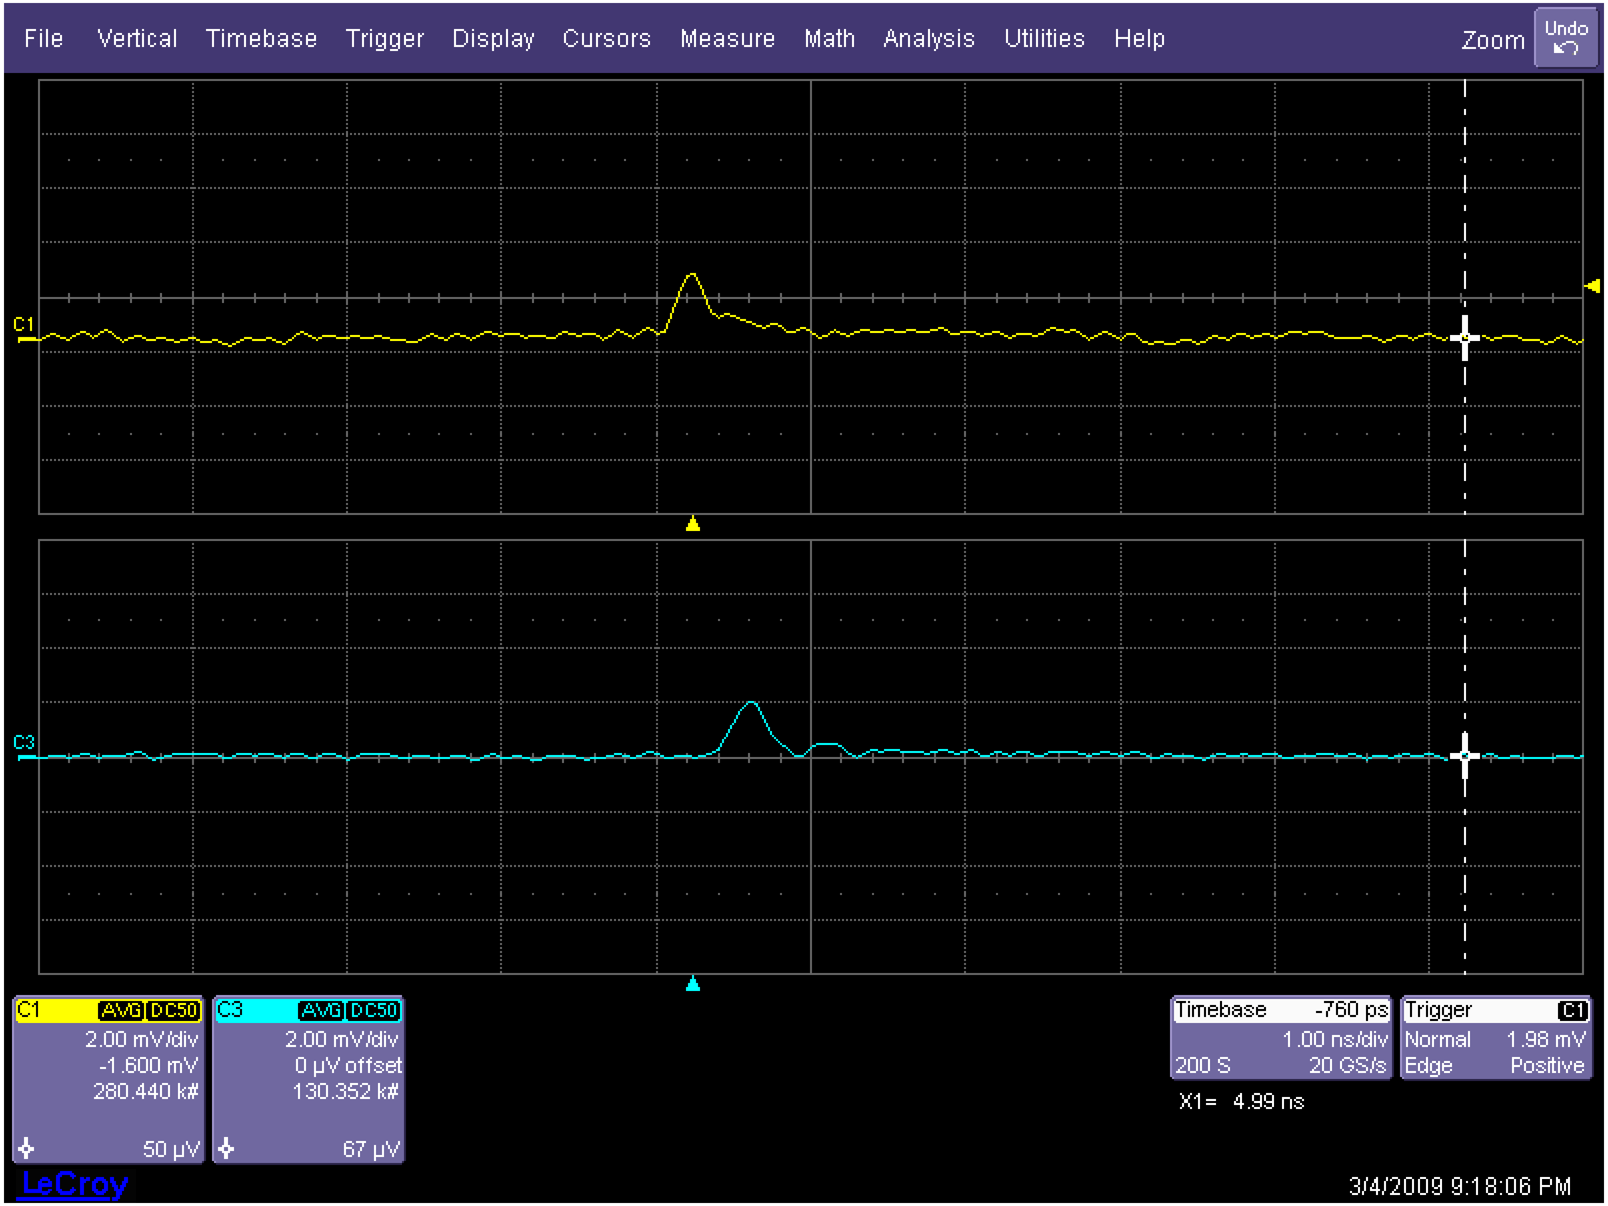

Screenshot of the captured signal on the oscilloscope for the middle raster scan.

Looking Around the Corner: 101 Barcode Experiment

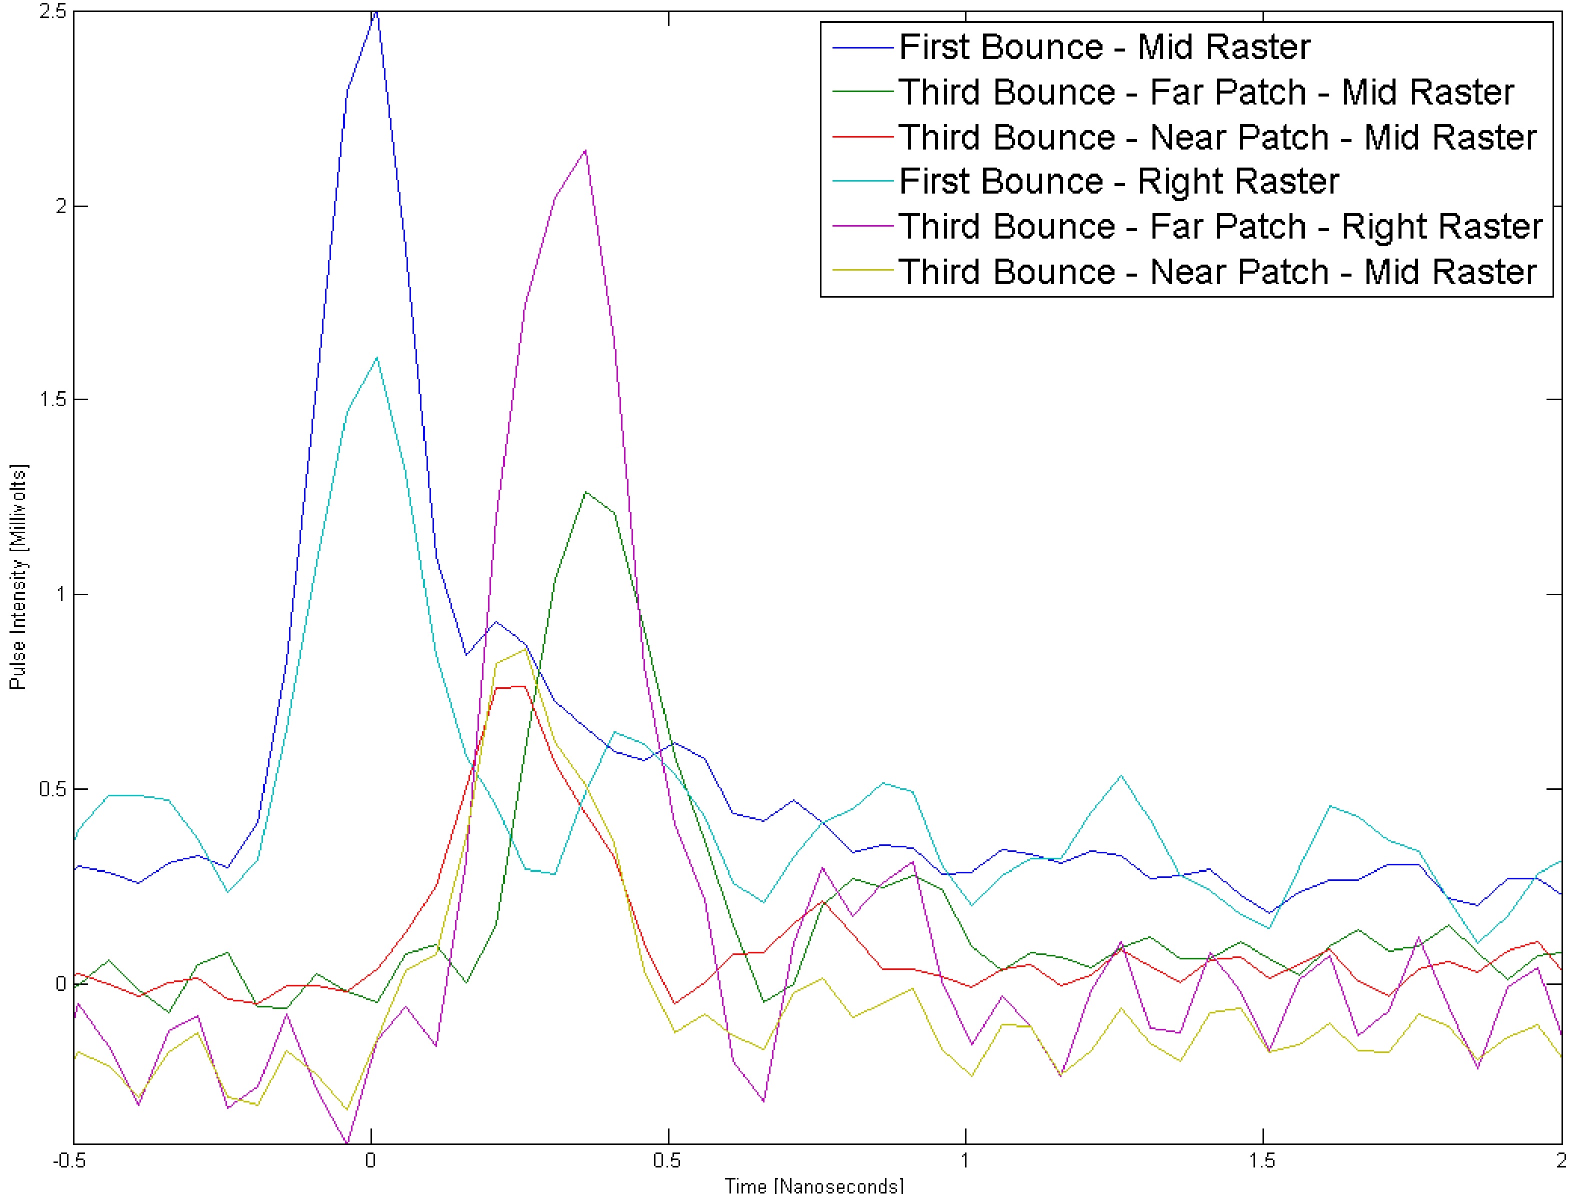

Plot showing the 1st bounce and the *two separately recorded 3rd bounces. Note* the very small delay *(≤ 200 ps) between two 3rd bounce arrivals. The plots include illumination of two separate diffuse patches, each one constrains the hidden surface geometry.*

Looking Around the Corner: 101 Barcode Experiment

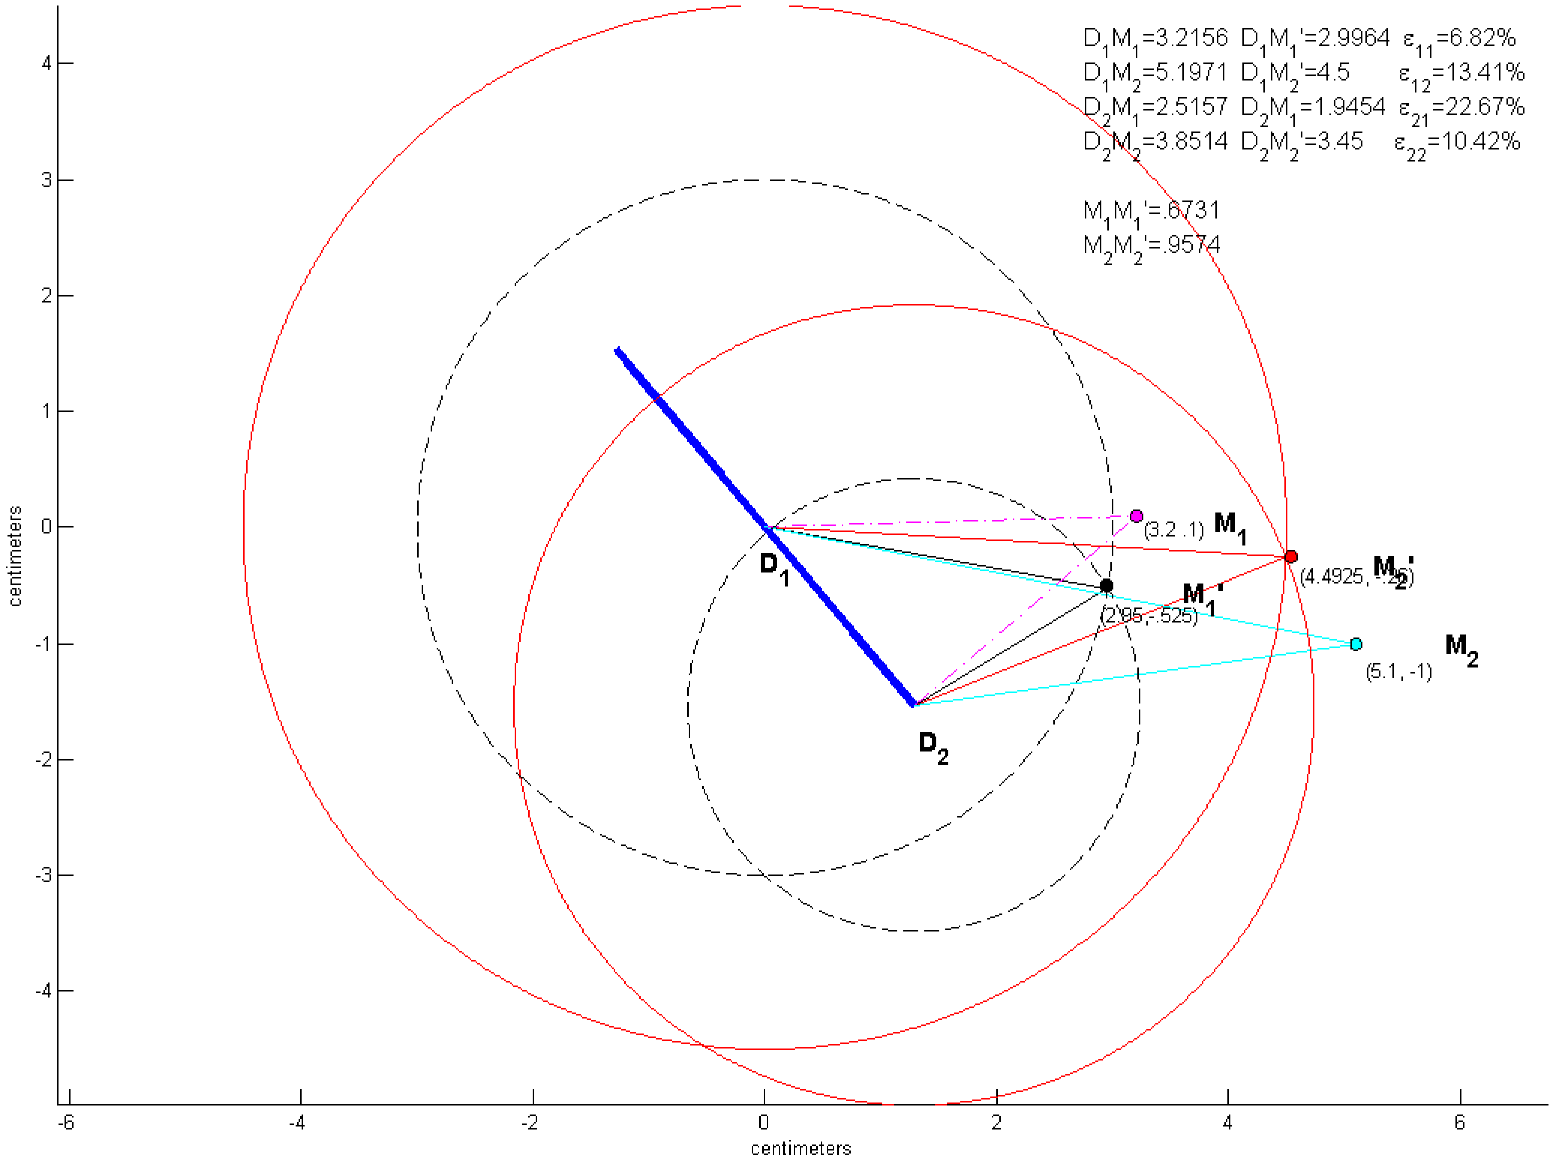

Scene Geometry Reconstruction

SUMMARY

Scene

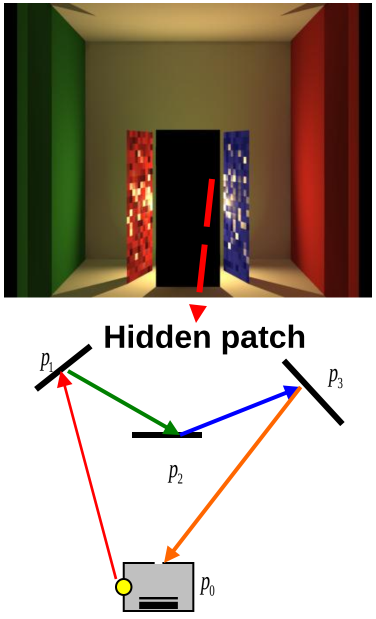

Consider a scene with 3 cards. The second card is hidden from the camera and the light source. We can only observe patch P_2 using third bounce information

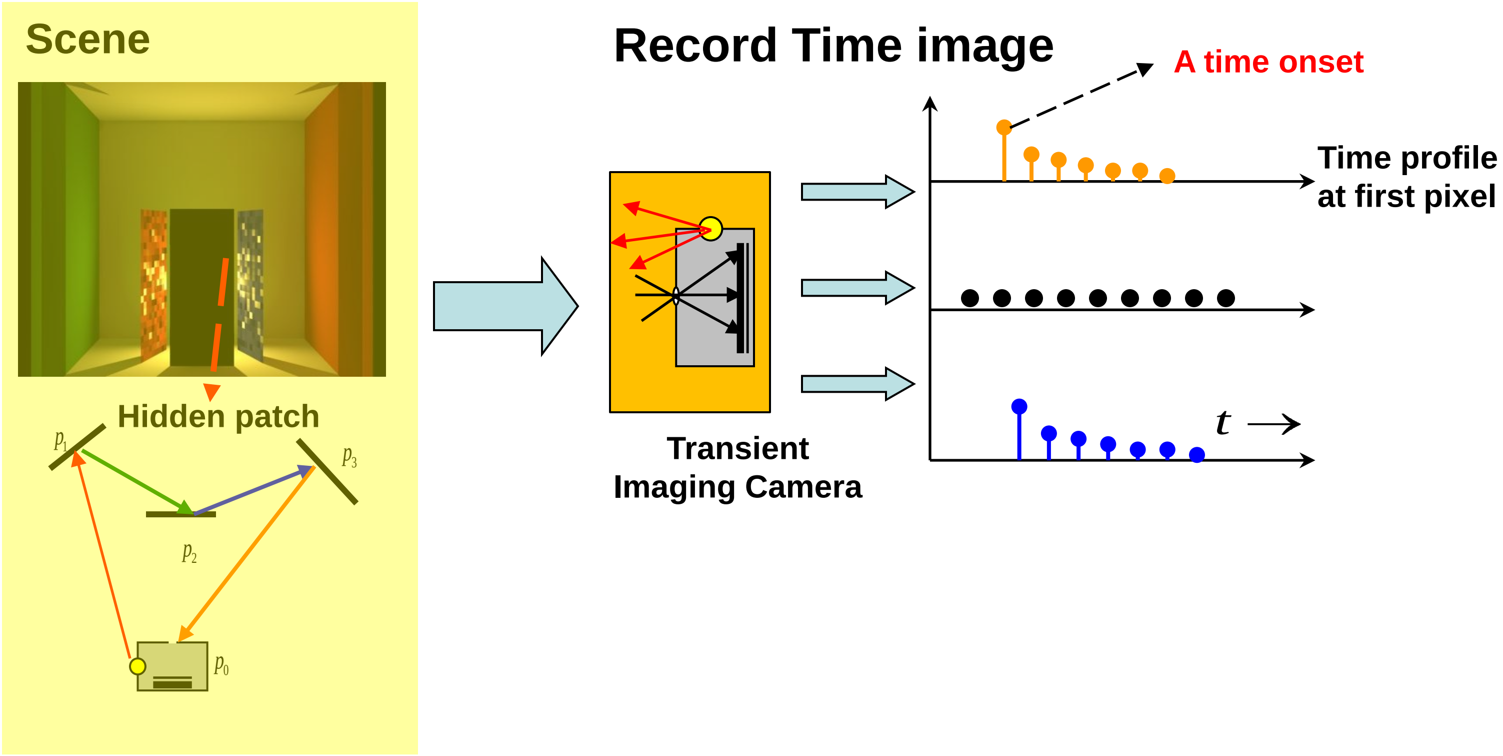

We capture a time image in response to every ray illumination. The time image is comprised of onsets. Note that we do not receive any light at the pixel observing P_2

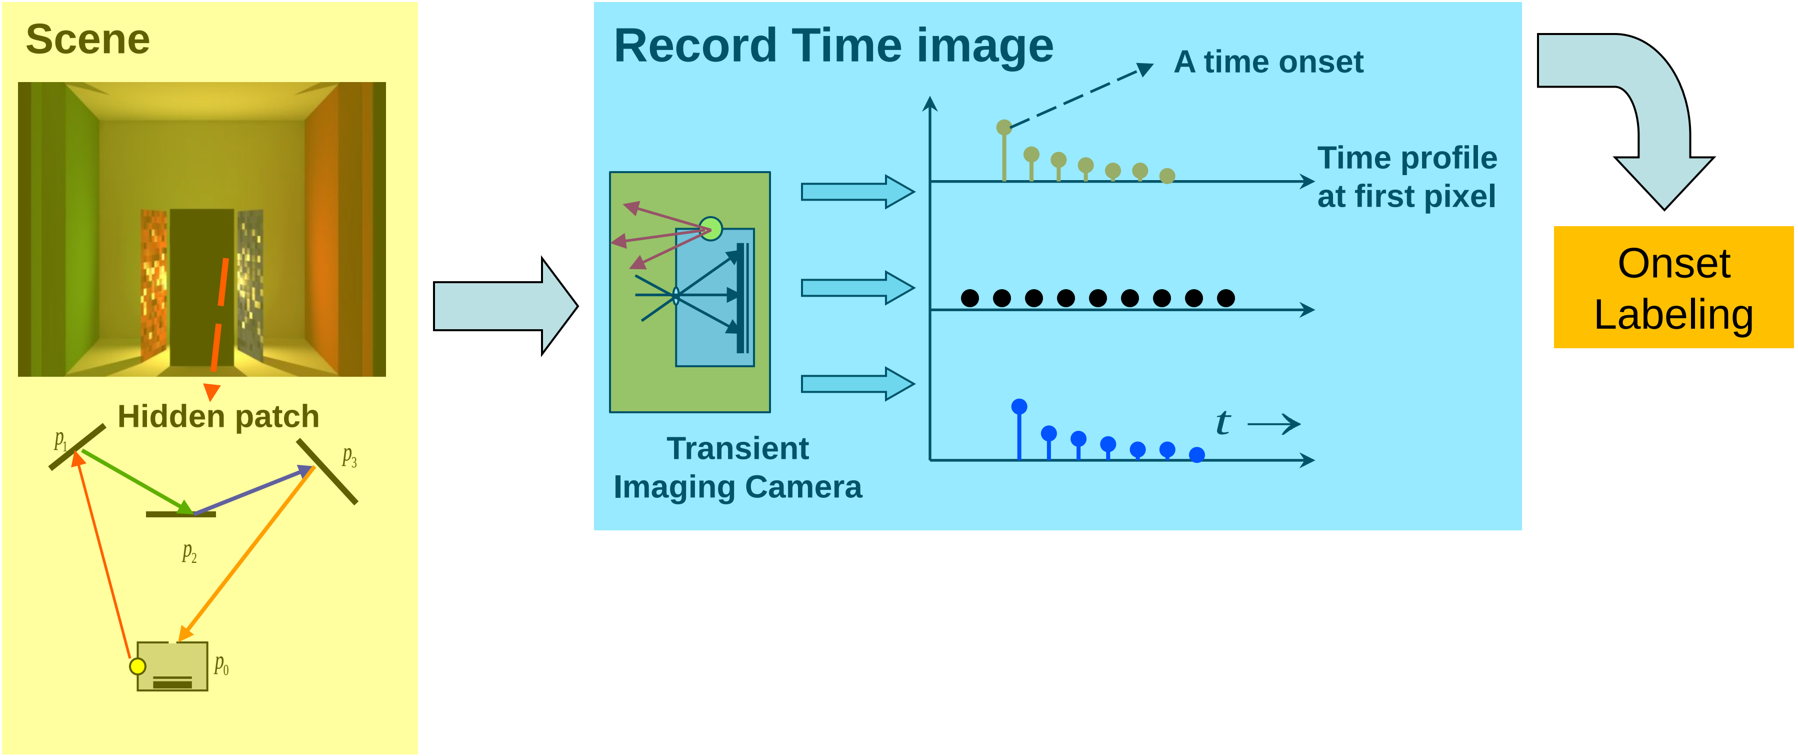

We then label the onsets corresponding to first and second bounces

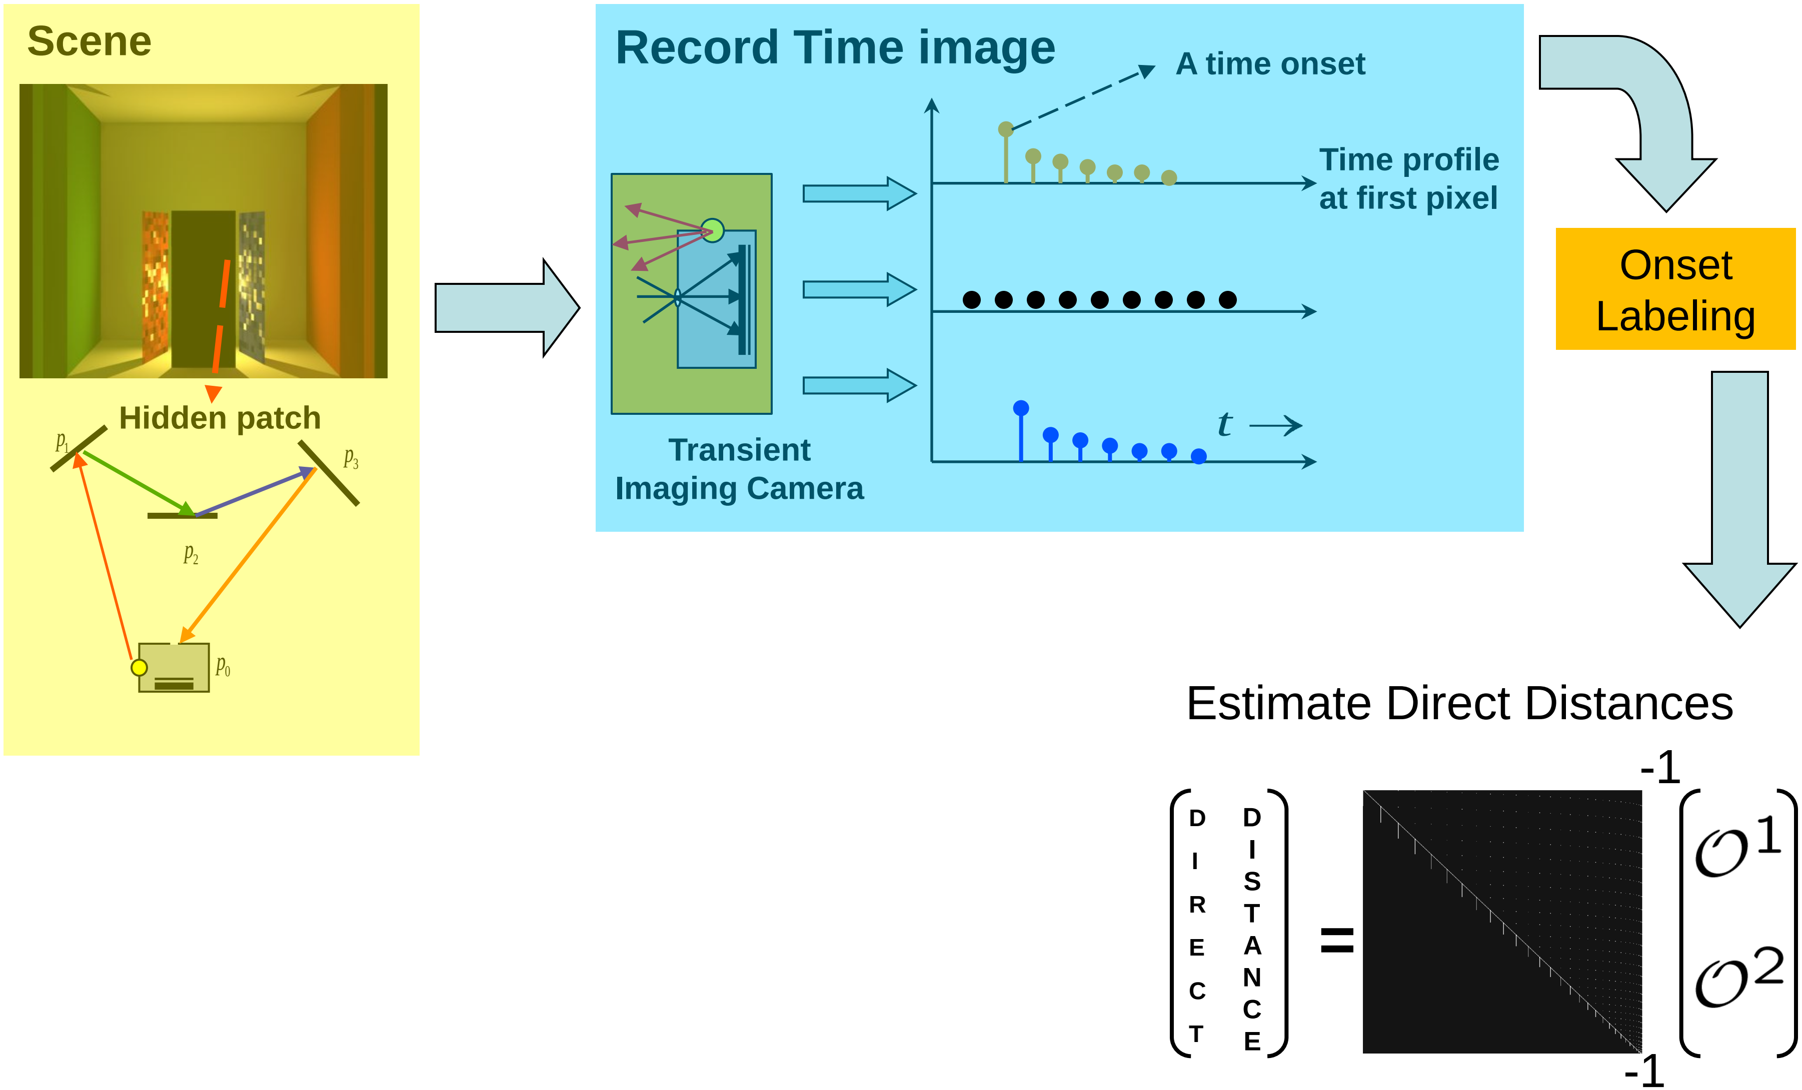

We use this information to infer distances to visible patches

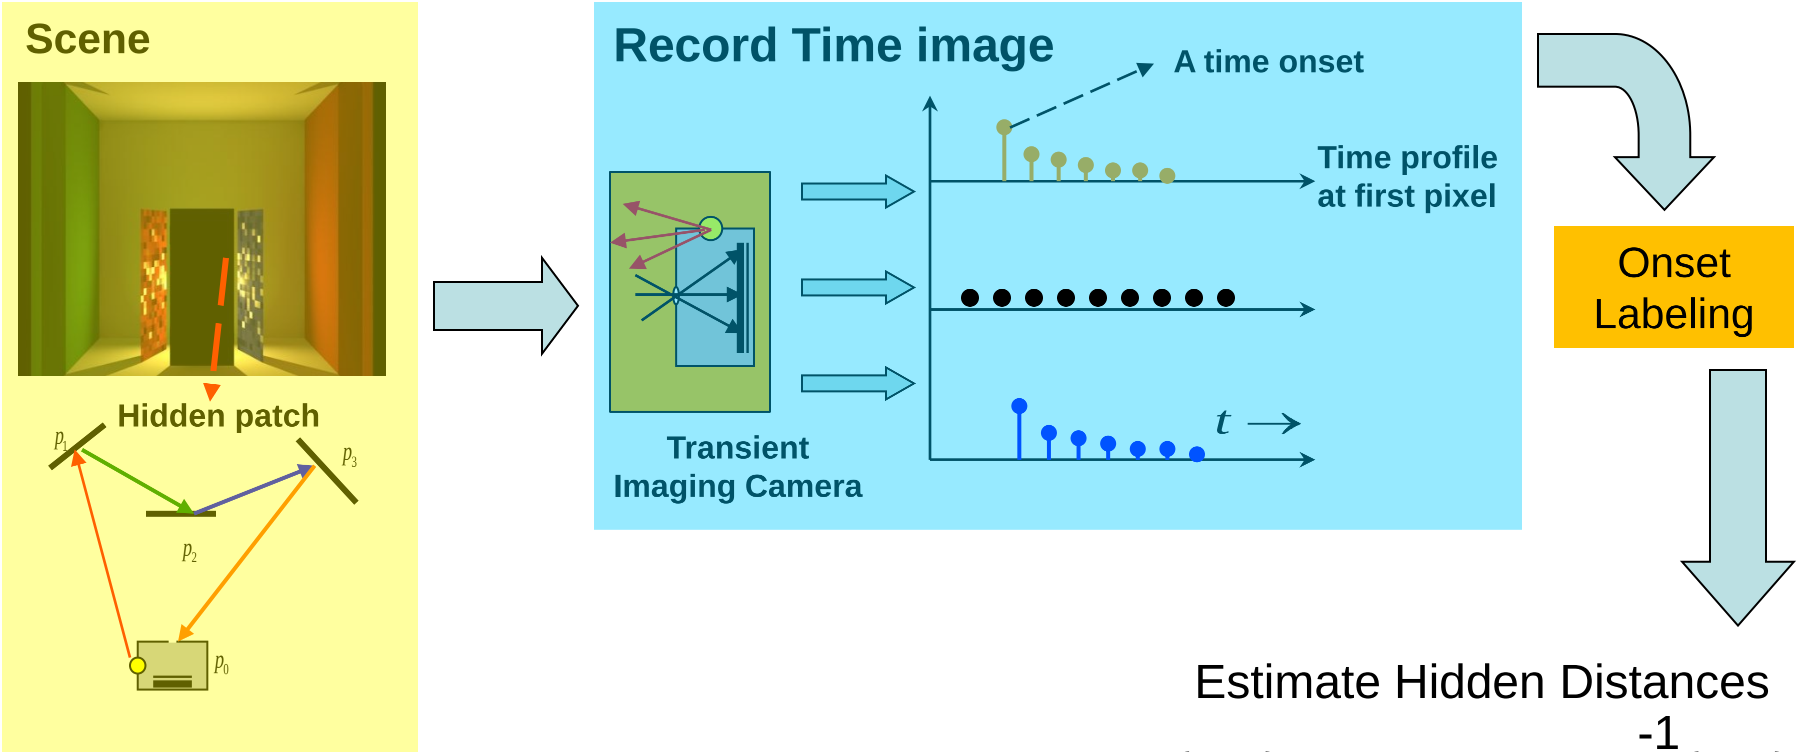

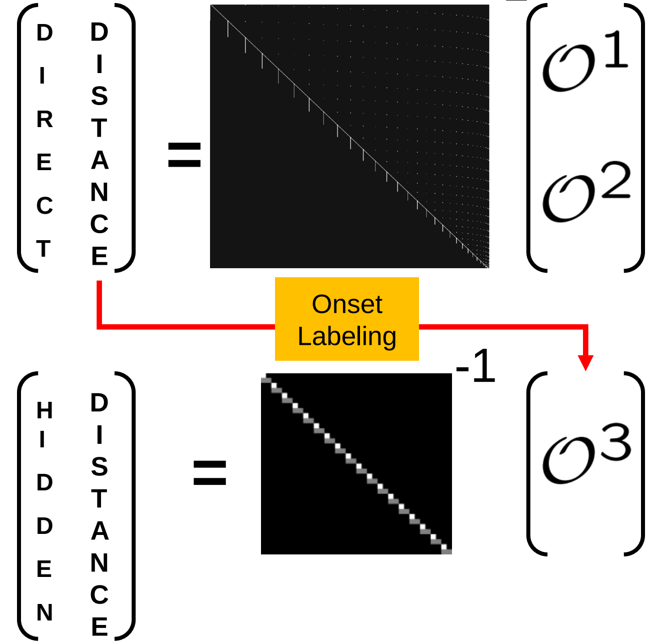

We then use direct distances to label the third bounce onsets

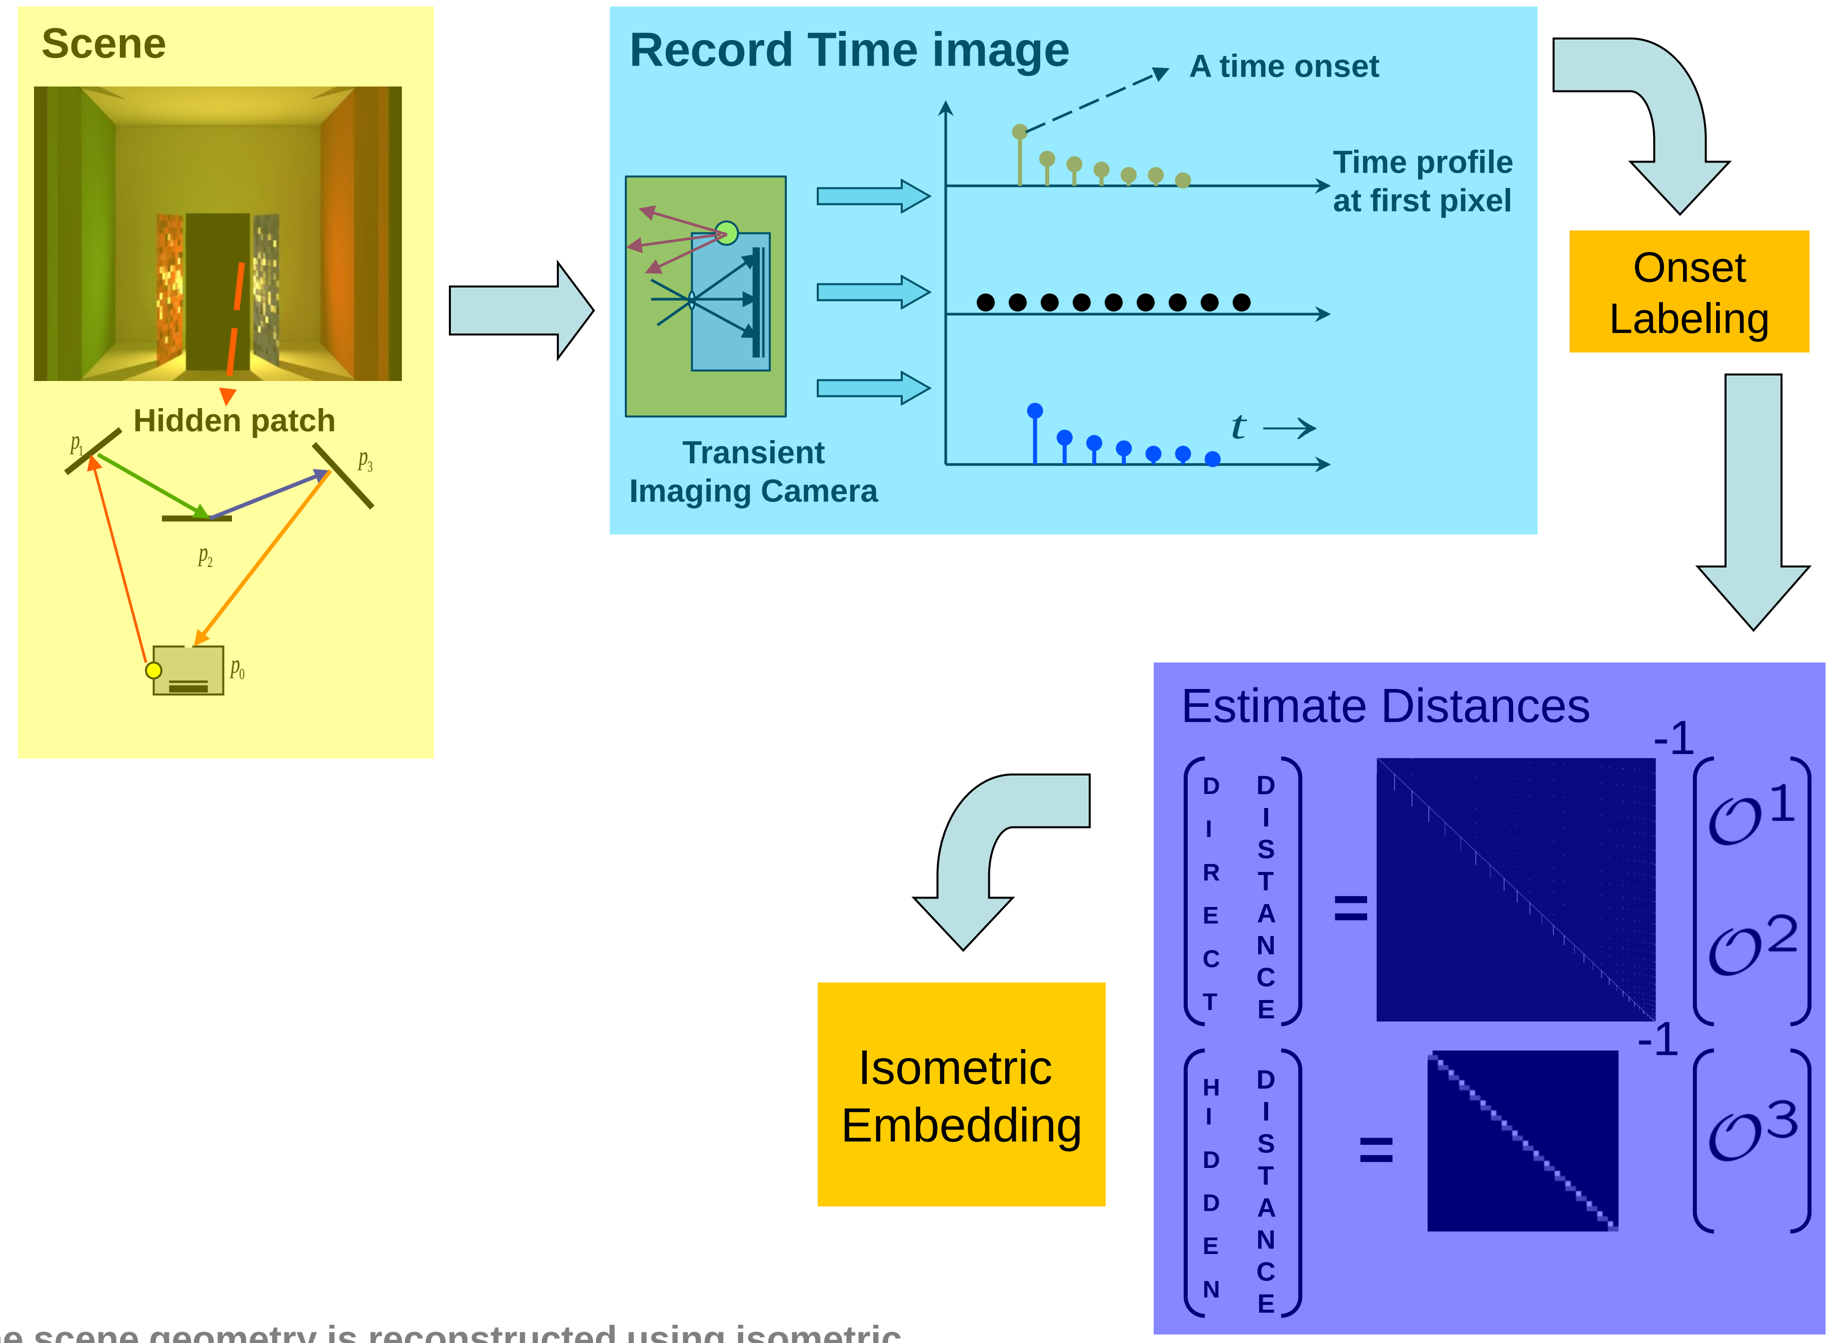

The scene geometry is reconstructed using isometric embedding and surface normal approximations

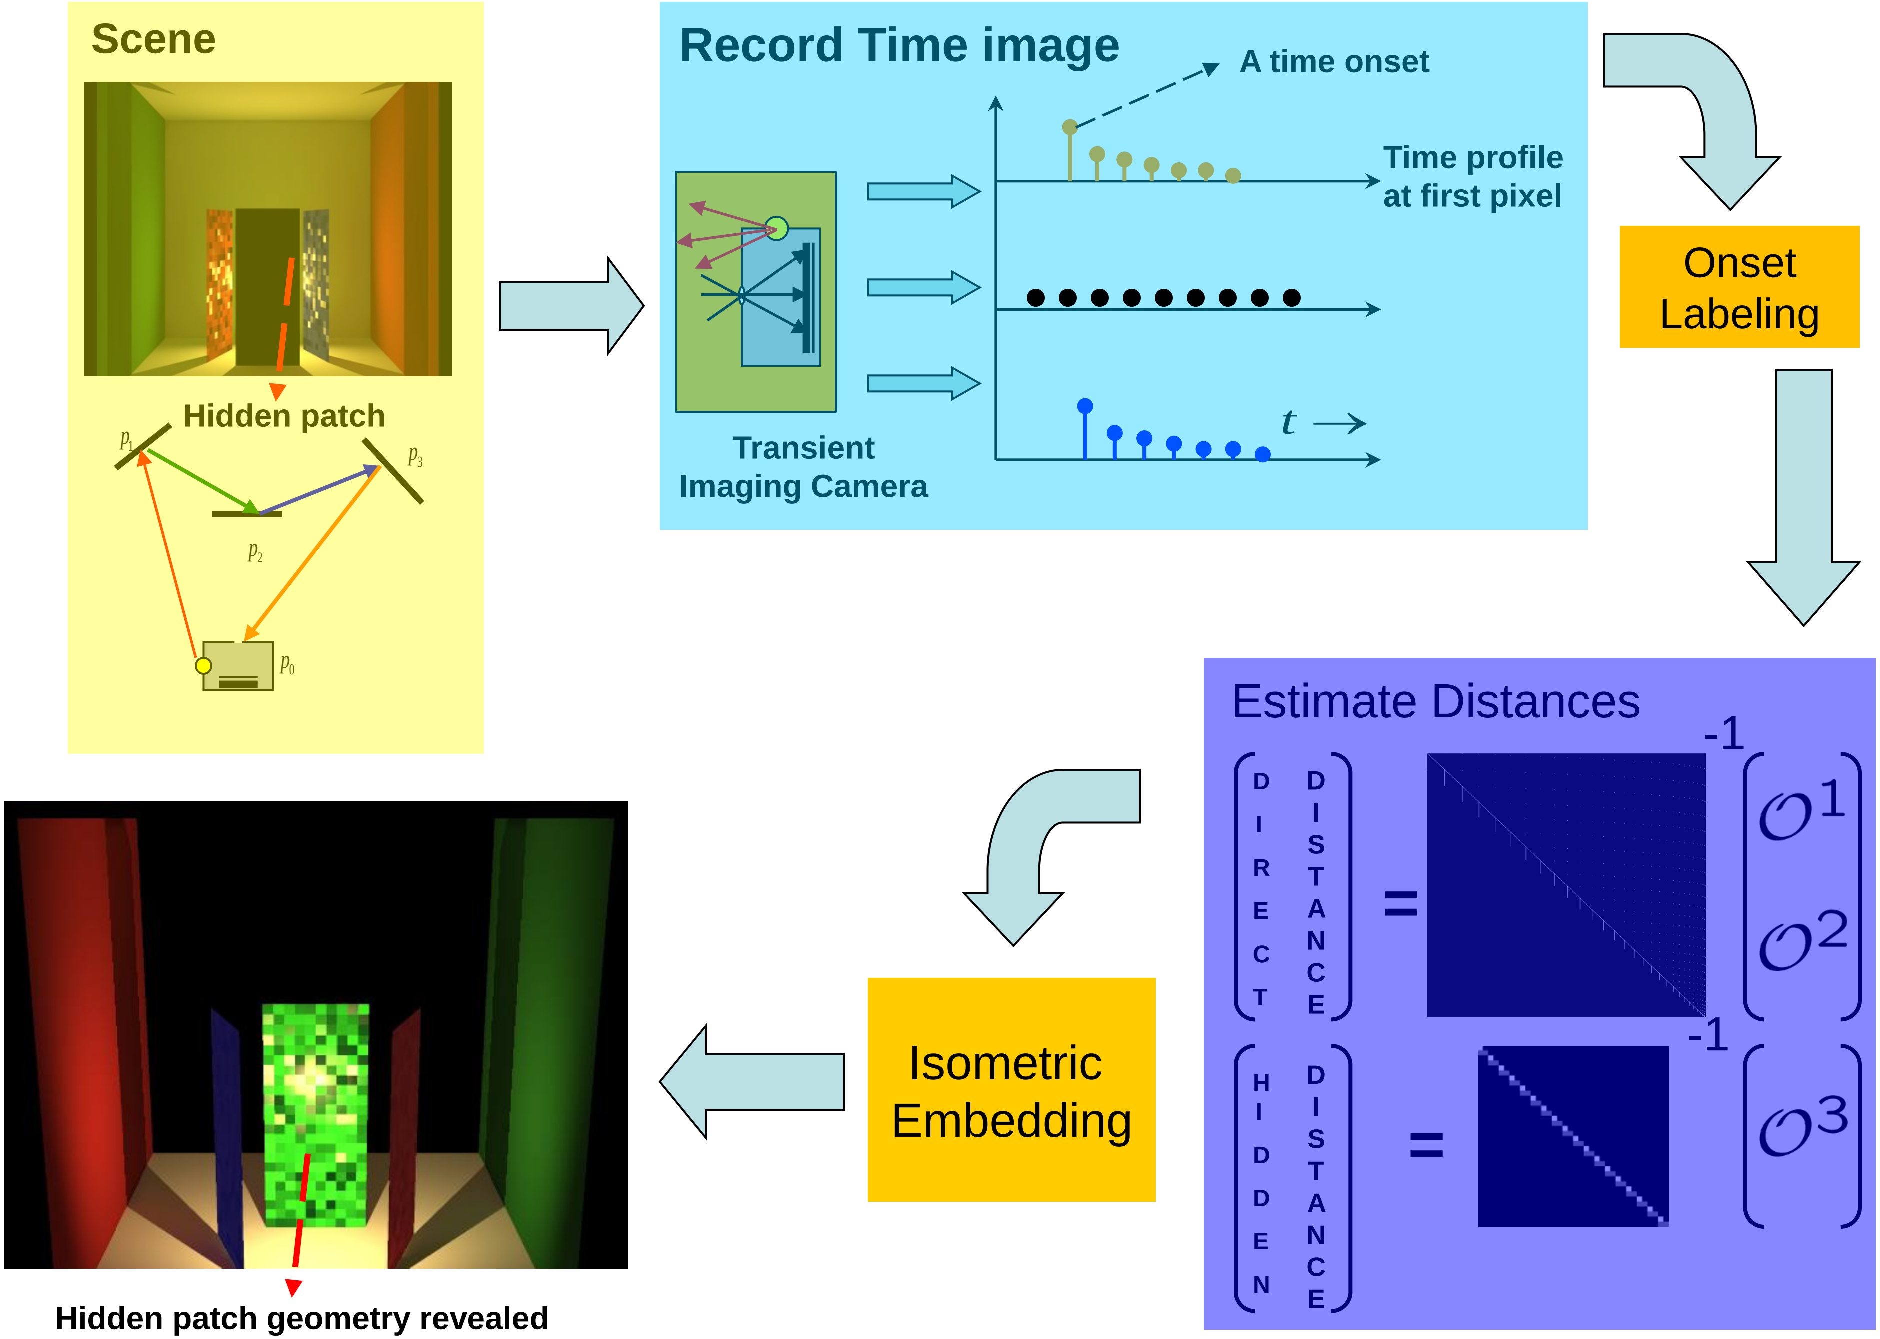
Scene
Hidden patch
p_1
p_2
p_3
p_0
Record Time image
A time onset
Time profile
at first pixel
$t \rightarrow$
Transient
Imaging Camera
Onset
Labeling
Estimate Distances
DIRECT DISTANCE
-1
$\mathcal{O}^1$
$\mathcal{O}^2$
HIDDEN DISTANCE
-1
$\mathcal{O}^3$
Isometric
Embedding
Hidden patch geometry revealed

This paper demonstrates that multipath analysis and higher dimensional sampling expands Computer Vision techniques for scene understanding.