Eurographics 2013
May 6-10, Girona (Spain)

# Path Space Regularization for Robust and Holistic Light Transport

Anton S. Kaplanyan and Carsten Dachsbacher
Karlsruhe Institute of Technology, Germany


Even the most recent existing methods are either not good **at** or not capable **of** handling complex illumination, such as reflected caustics on the floor.
In this work we will show how to combine the **strengths** of these methods together.
The goal of our work is to compute the complete light transport solution, including different phenomena, such as reflected caustics.


Even the most recent existing methods are either not good **at** or not capable **of** handling complex illumination, such as reflected caustics on the floor.
In this work we will show how to combine the **strengths** of these methods together.
The goal of our work is to compute the complete light transport solution, including different phenomena, such as reflected caustics.


Even the most recent existing methods are either not good **at** or not capable **of** handling complex illumination, such as reflected caustics on the floor.
In this work we will show how to combine the **strengths** of these methods together.
The goal of our work is to compute the complete light transport solution, including different phenomena, such as reflected caustics.

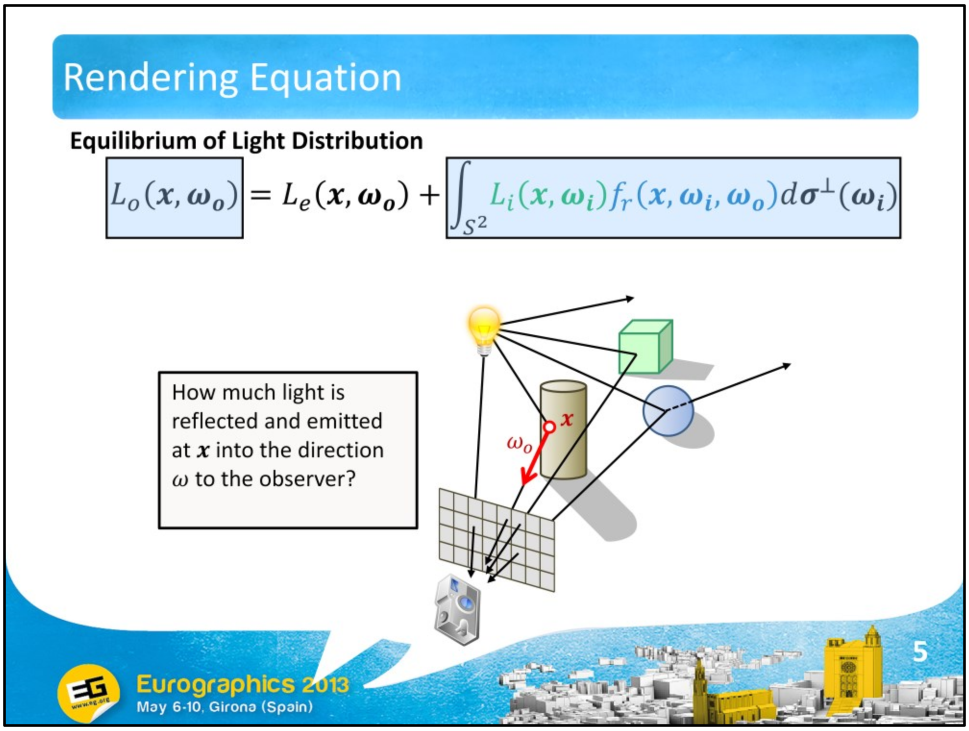


A short introduction to the source of the problem.

The equilibrium of light propagation is described using a Fredholm integral equation of the second kind.

This is a **rendering equation** and it describes the reflection distribution of light at the surface towards all directions.

The integral is the convolution of the **B**idirectional **R**eflectance **D**istribution **F**unction (BRDF) with the distribution of the light incident to the surface.

We are interested in the light coming towards the **camera** from **all** visible surfaces.

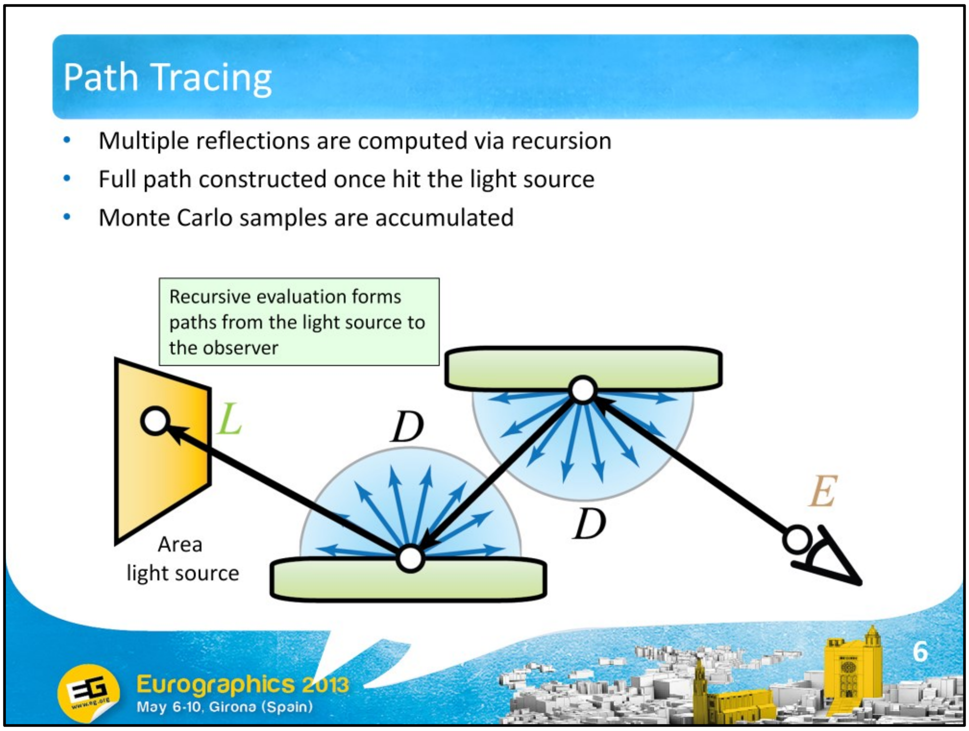


The first and the simplest method is called **path tracing**.
The rays are shot from the **camera** and bounced until **hit the light source**.
The complete paths are sampled **stochastically** and the rendering equation is solved using Monte Carlo integration.

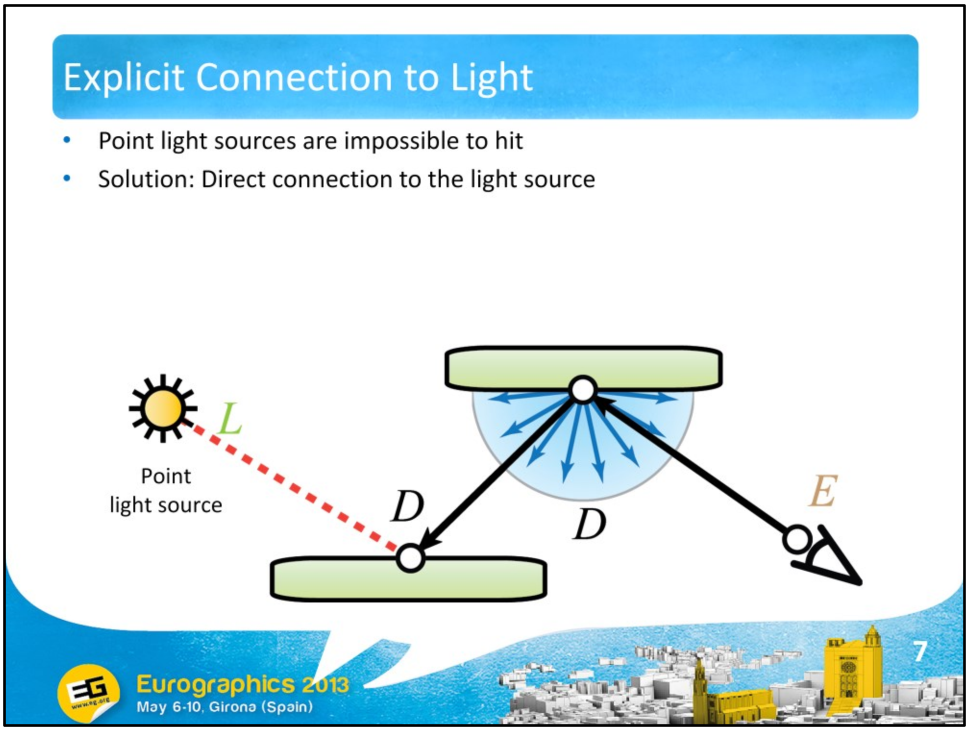


However it is impossible to directly **hit** a light source when rendering a scene in presence of **point lights** using path tracing.
The point light model is a common mathematical model used in graphics for many real-world emitters with tiny emissive area.
As a solution, such light sources are explicitly checked with **direct connection** at each interaction.

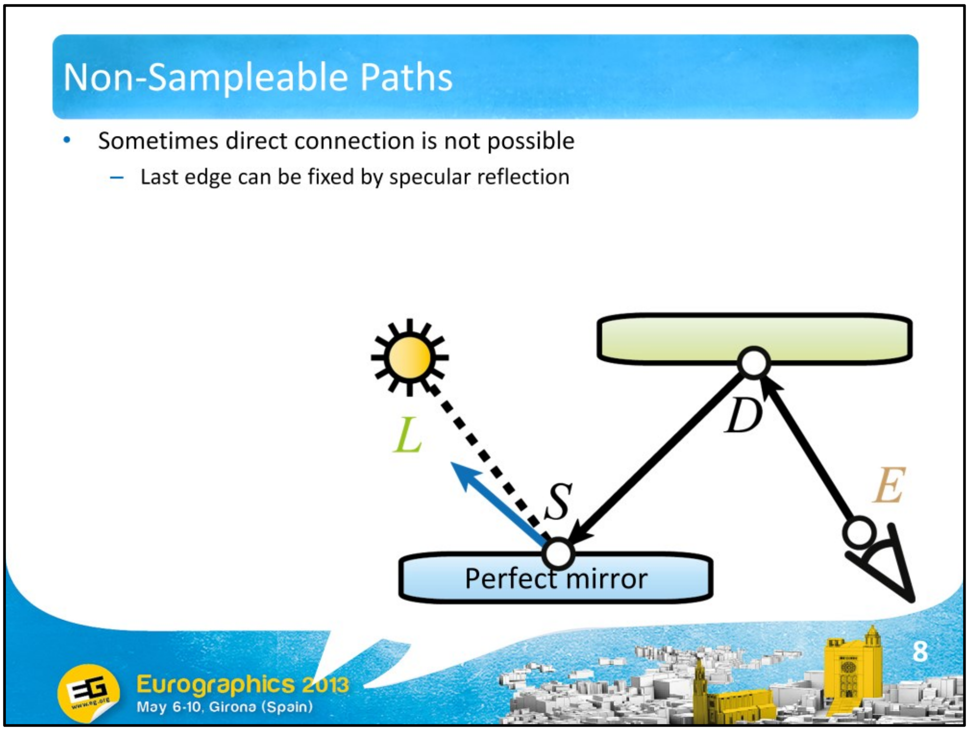


However in some cases even direct connection **does not** solve the problem. Imagine a **caustic**. Typically in such case the interaction next to the light source is **perfectly specular**, such as mirror reflection or glass refraction.
The direct connection **always fails**, since the connection direction should perfectly match the **stochastic** reflected direction.


There is a lot of previous work on how to construct difficult paths using different **connection methods** and **path construction** approaches, namely

1. Bidirectional Path Tracing (BDPT)
2,3. Photon mapping and its progressive extension
4. Vertex Connection and Merging as a combination of BDPT and PPM
5. And the recent Manifold Exploration mutation strategy for Metropolis light transport.

I will discuss each of these methods **later on** in my talk.

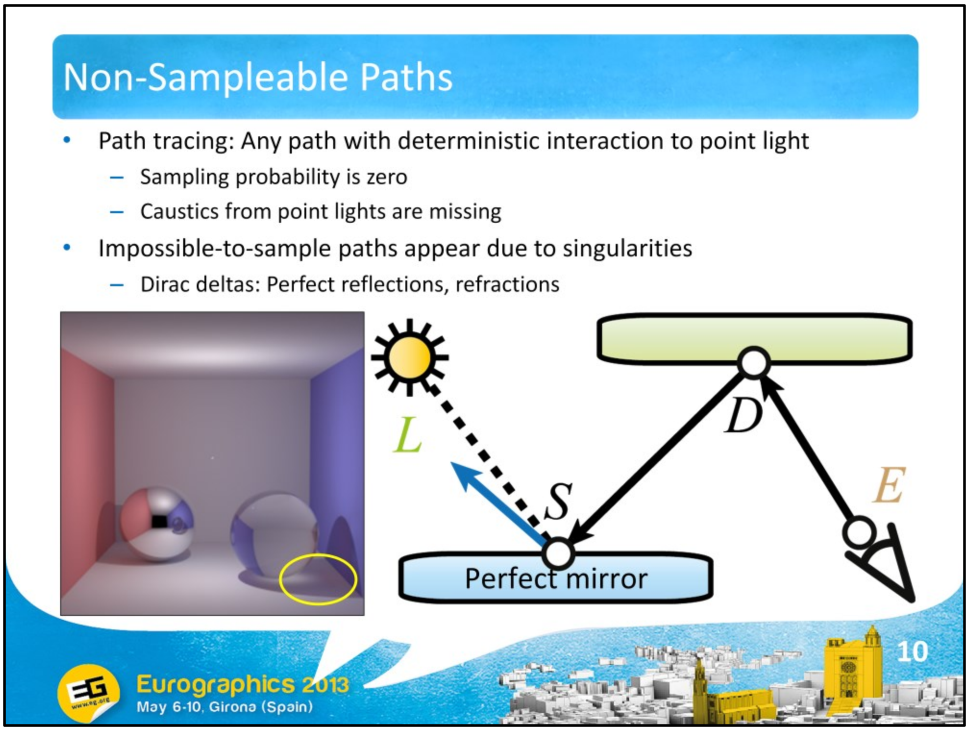


So the paths with the deterministic last interaction, such as caustics from point lights, **cannot be constructed** with path tracing.
For example, in this simple image on the left, rendered with path tracing, the caustics on the floor are **missing**.
This happens because the deterministic interaction is mathematically defined as a Dirac delta distribution, causing a **singularity** in the integrand.
However there are more advanced methods that **can** sample caustics from point lights.

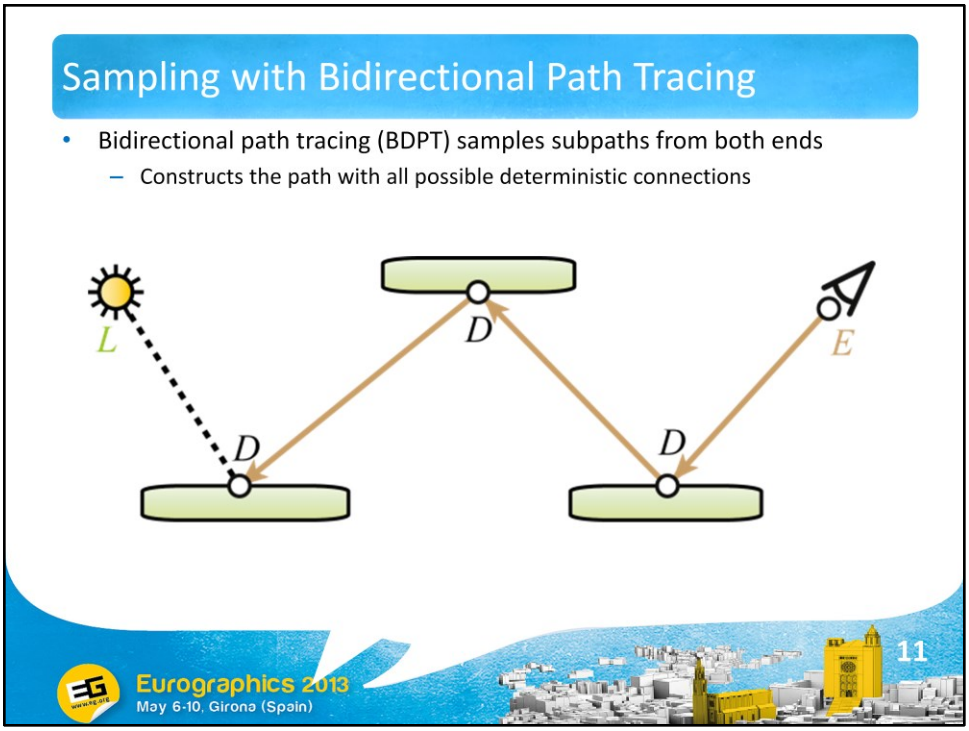


One of them is bidirectional path tracing.
It can potentially connect a given path at **every edge**, as illustrated on this path with four edges.
So the first option is to trace all the way to the light and then do a direct connection to the light source at the last edge.

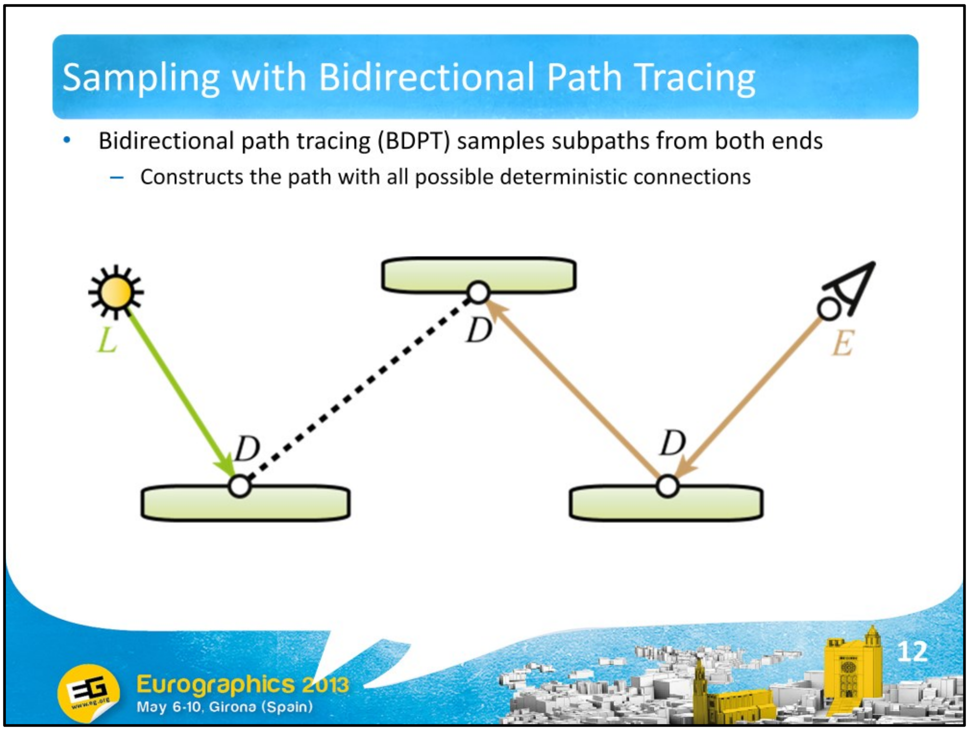


Or trace it further from the light and do a connection in the middle.

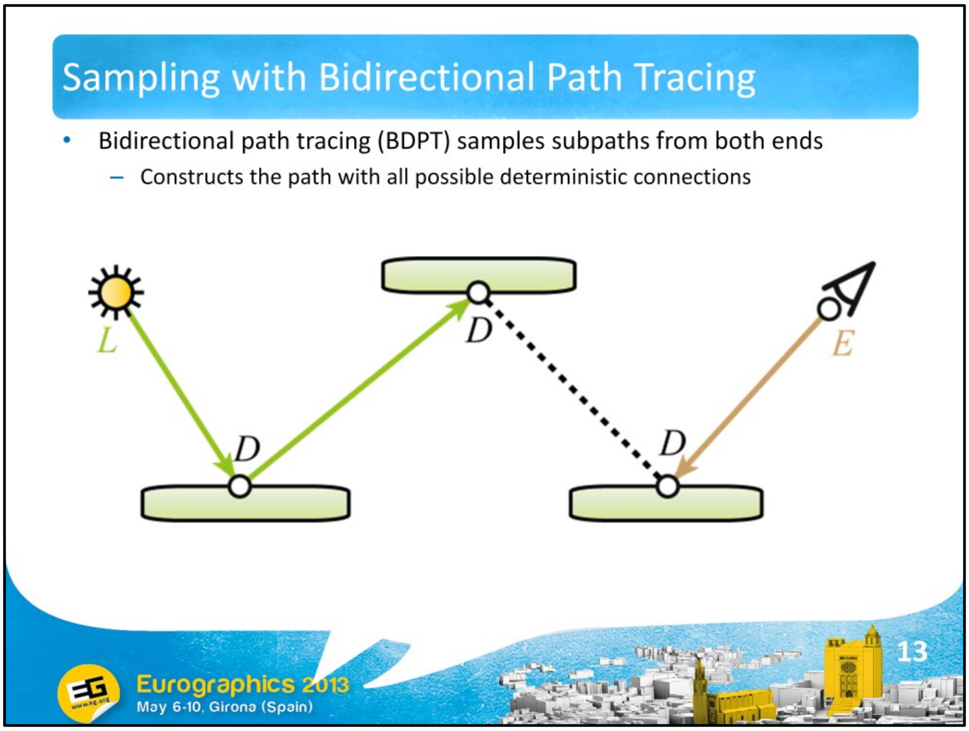


Or trace it further from the light and do a connection in the middle.

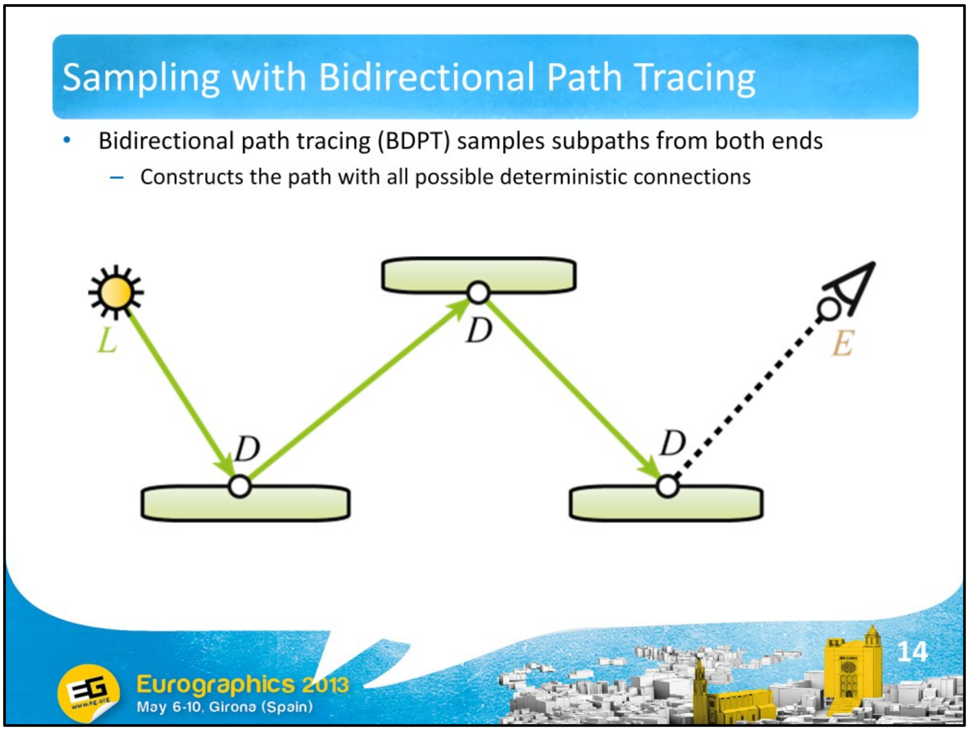


Or trace it all the way from the light source and do a direct connection to the camera sensor.

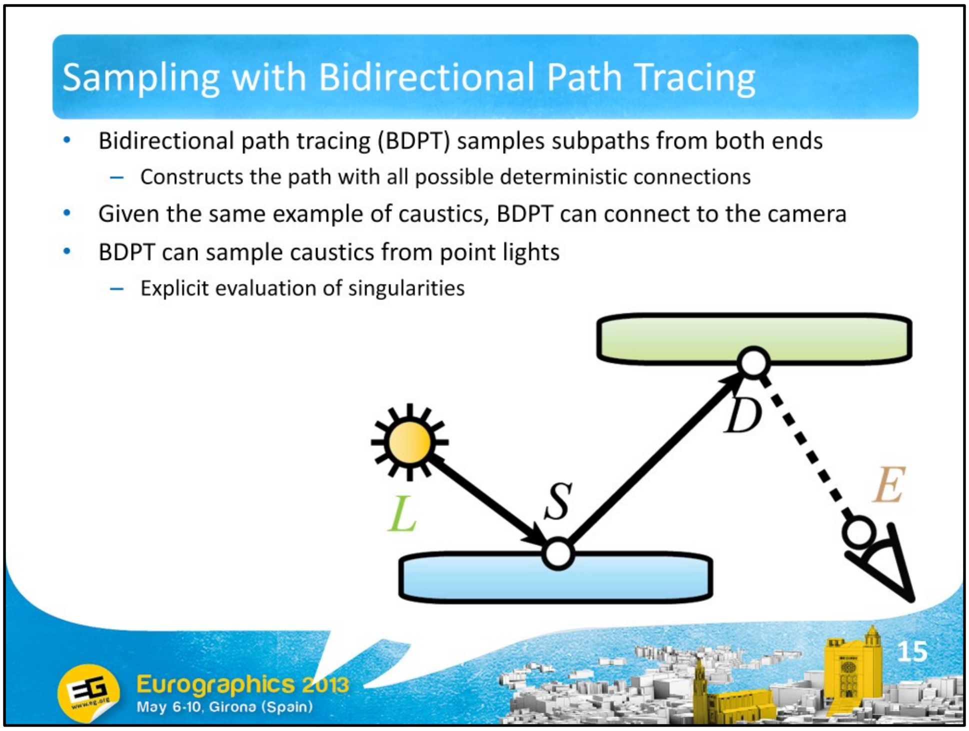


Consider the **previous example** of a caustic.
BDPT can construct such path starting from the light source by tracing up to the non-deterministic interaction, such as **diffuse surface**, and then connect to the camera.
The trick is that BDPT **bypasses** the singularities, instead of attempting to stochastically connect at them.

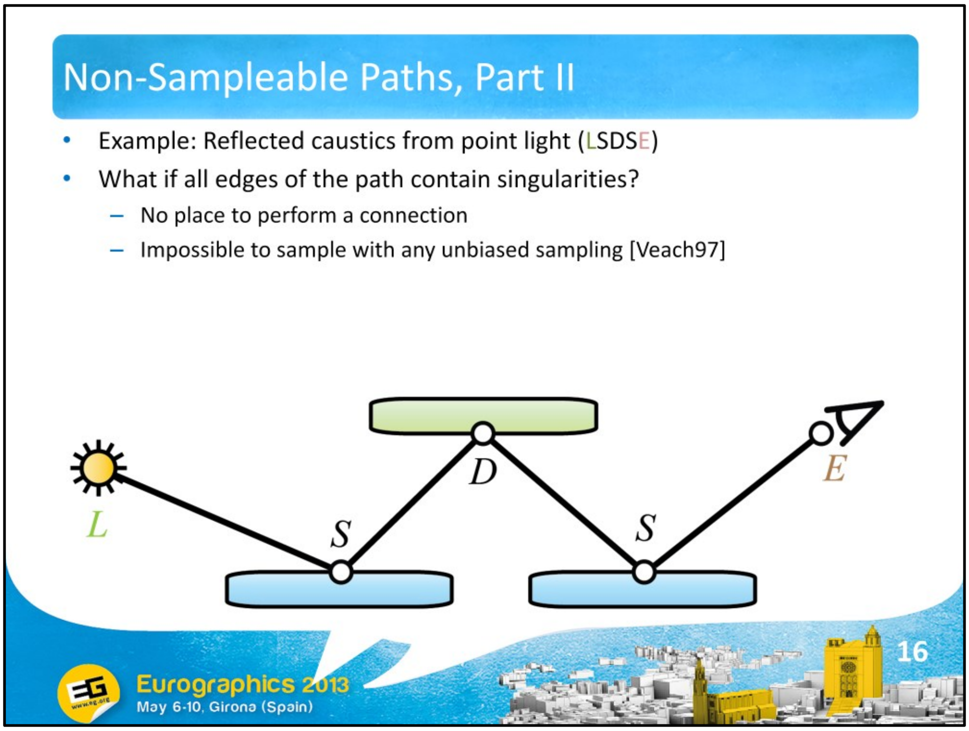


However, again, even such sophisticated methods like BDPT can **fail** to construct a path in some cases.
One of these cases is the **reflected caustics**.
The problem here is that **every edge** adjoin a singularity at one of its ends.
Such paths with no connectible edges are showed to be **non-sampleable with local unbiased methods** by Eric Veach.
And there is a more fundamental reason for that. Let’s take a brief **excursus**.

## Excursus: Undecidability of Ray Tracing

- Specular path tracing is *undecidable* [Reif et al. 1994]
  - Task: Find all specular paths from one fixed point to another ($DS^+D$)
  - Impossible to find all such paths on a Turing machine
  - Given arbitrary specular geometry

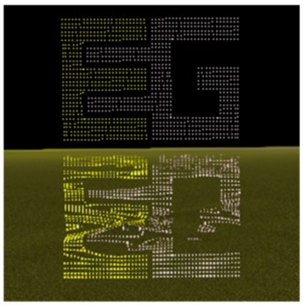

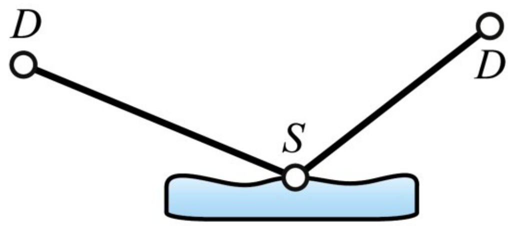




Interestingly, it is **not possible** to find all specular paths from one fixed point to another on a Turing machine given an arbitrary configuration of surfaces.
An example image on the left shows an EG logo modeled with **point light sources** placed over the **curved** water surface. The reflections of such light sources **pose** such an undecidable problem.

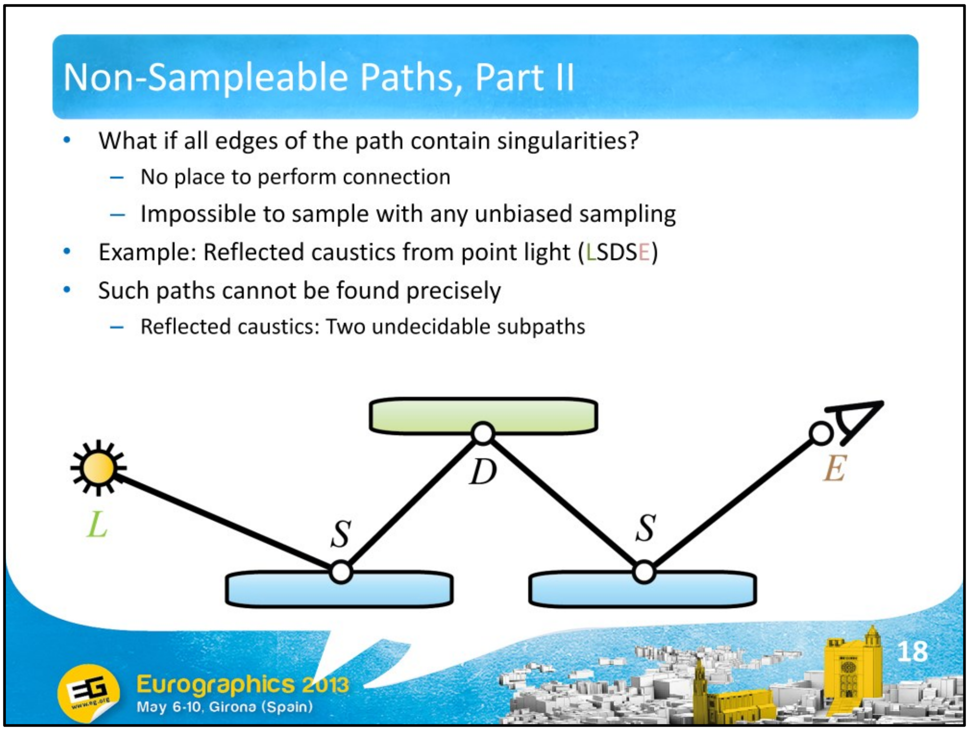


Coming back to an example of reflected caustics, we can see that this path consists of **two** undecidable subpaths.
Thus it does not matter from which direction to trace these subpaths, one of them will pose such problem during the connection.

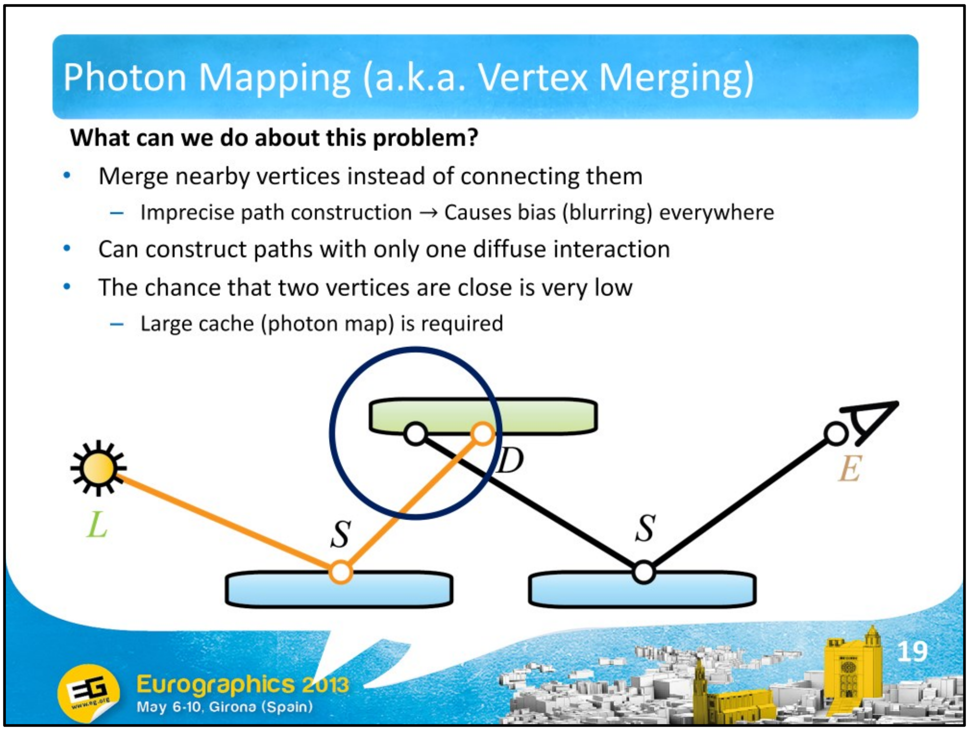


Not a secret that photon mapping can sample such paths, considered non-sampleable by local unbiased methods.
The reason is photon mapping constructs the full path by **merging** proximate vertices of different subpaths.
This is a **biased way**, causing blurring of image features. However this way **more difficult paths**, such as reflected caustics, can be sampled.

Note that the chance that two vertices happened to be located nearby is **very low**. Thus photon mapping requires a **large cache** of the light subpaths, called a photon map.

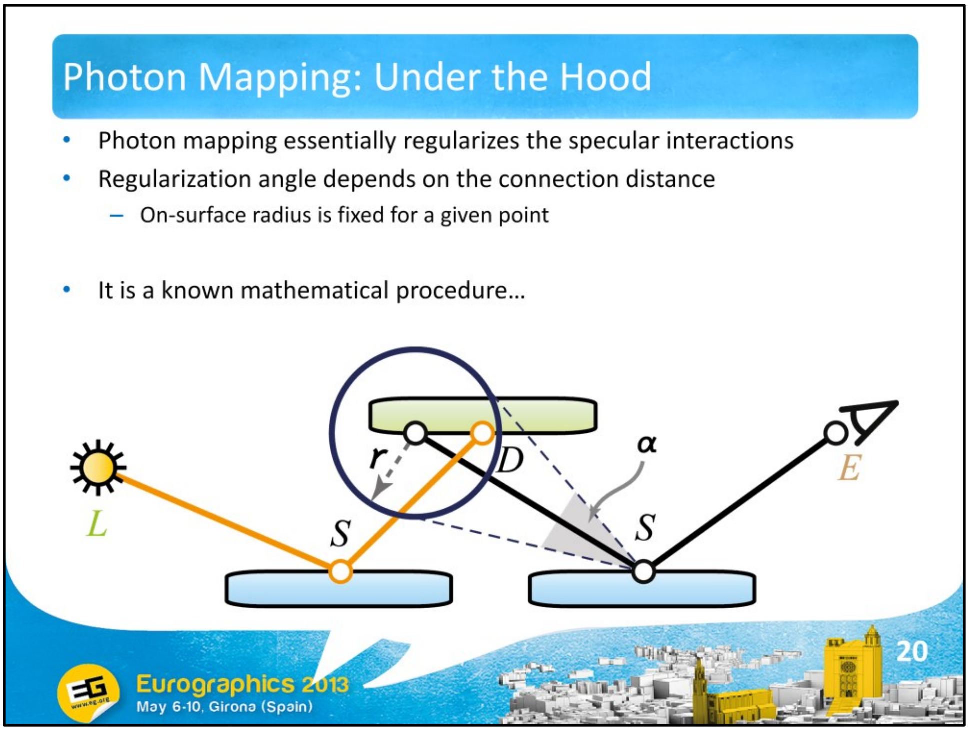


What photon mapping is going during density estimation is also known as regularization in mathematics.
Photon mapping **regularizes the interaction** by merging the path at the next vertex.
The **regularization angle**, *alpha*, can be derived from the fixed photon mapping radius *r* and depends on the connection distance.
It appears to be a standard mathematical procedure. We also use this procedure, however in a more smart way.

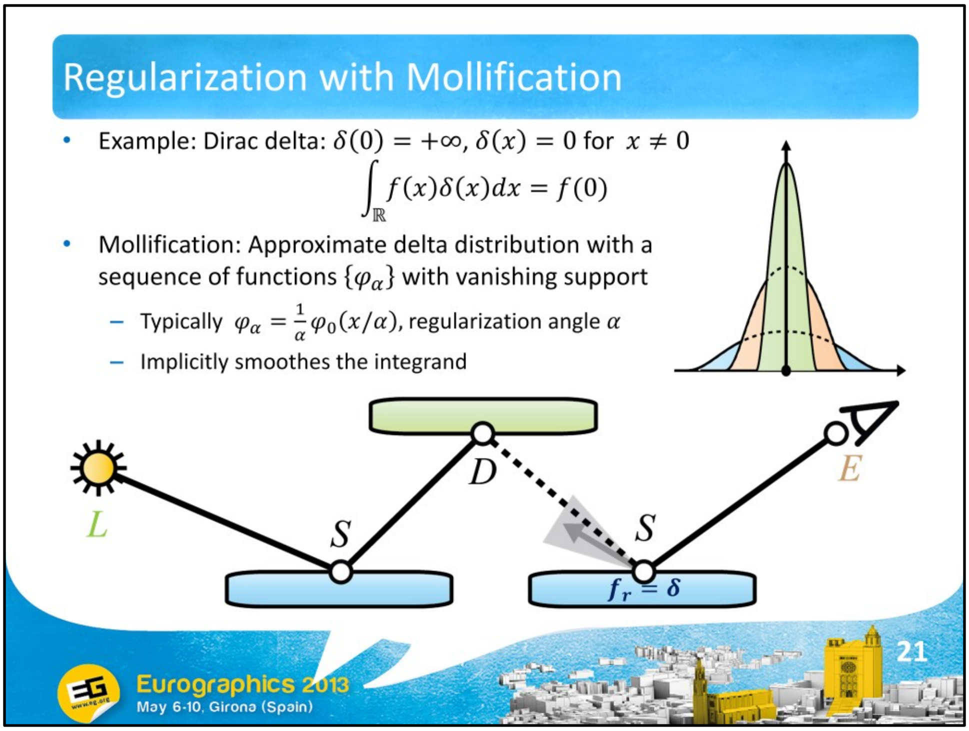


This procedure is called **mollification**.
Given a singularity caused by a **delta distribution**, like the one at the right specular vertex;
we construct a **sequence** of integrable smooth functions, that approach delta distribution in the limit.
Then we shrink the regularization angle during the integration, making the integrand less and less smooth.

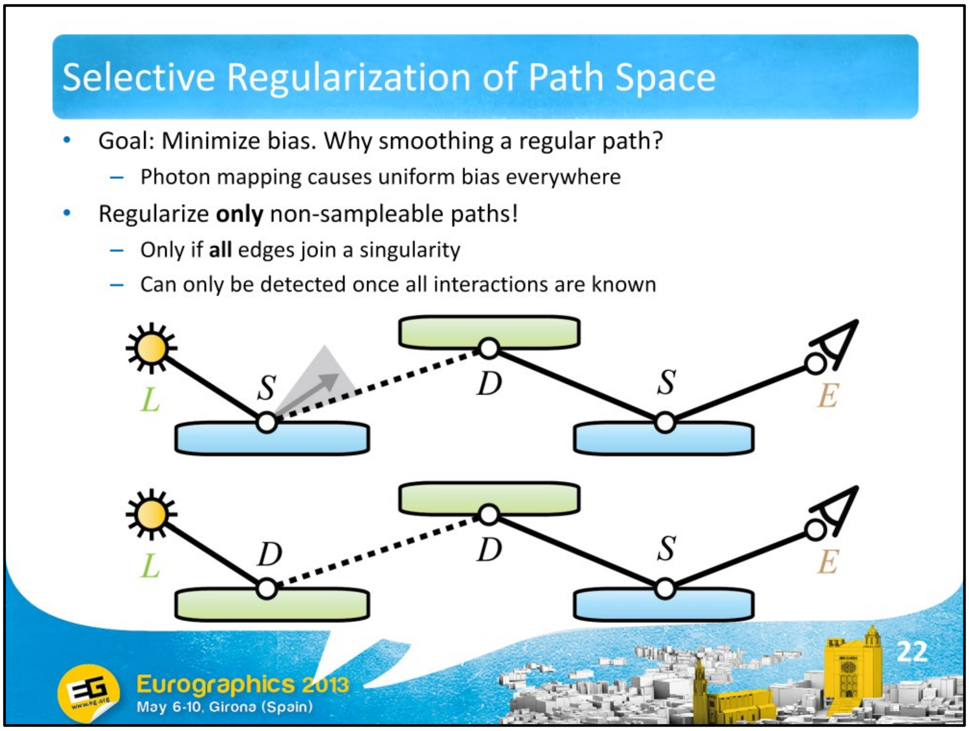


So, now that we forumlated what is going on, we can selectively regularize only when **necessary**.
Instead of regularizing all paths, including regular ones, we regularize **only non-sampleable paths (irregular for the sampling method)**, thus **minimizing** the amount of bias.
In case of BDPT, non-sampleable paths are detected if **all edges** of the path adjoin at least one singularity.
This situation can be recognized only once **all subpaths are traced** and all interactions are known.

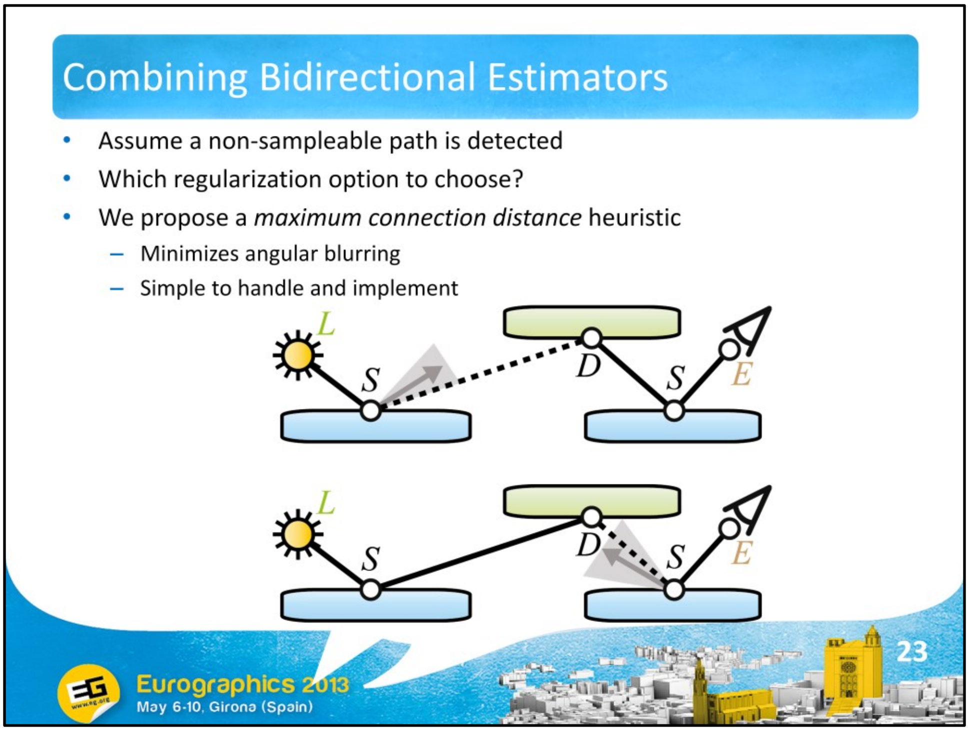


If the path generation method has multiple options of constructing the same path, such as a set of bidirectional estimators, then we need to choose between **several** different regularization options, as the two options for the same path in the bottom.

Here we propose to use the *maximum distance heuristic:* we regularize only if the connection edge is the longest among other path edges. Given that the on-surface radius is fixed, we minimize the angular smoothing caused by the regularization. Also it is very easy to practically handle.

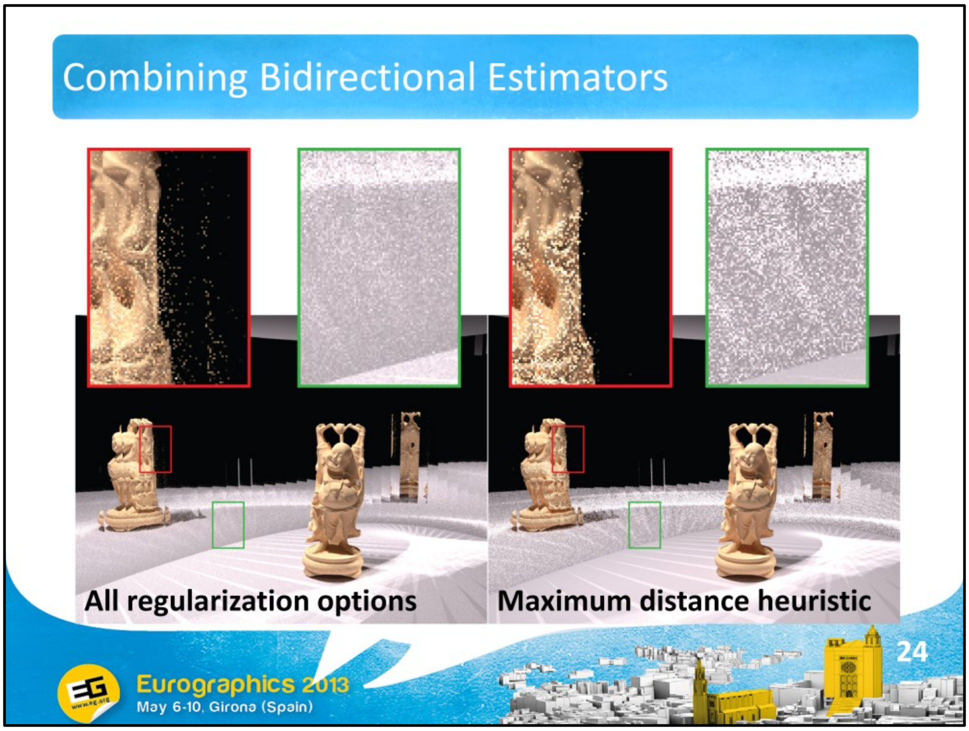


The example shows **all regularization options**, equally weighted, on the left; and the maximum distance heuristic on the right.
The bias caused by angular blurring is notably **smaller** on the right.

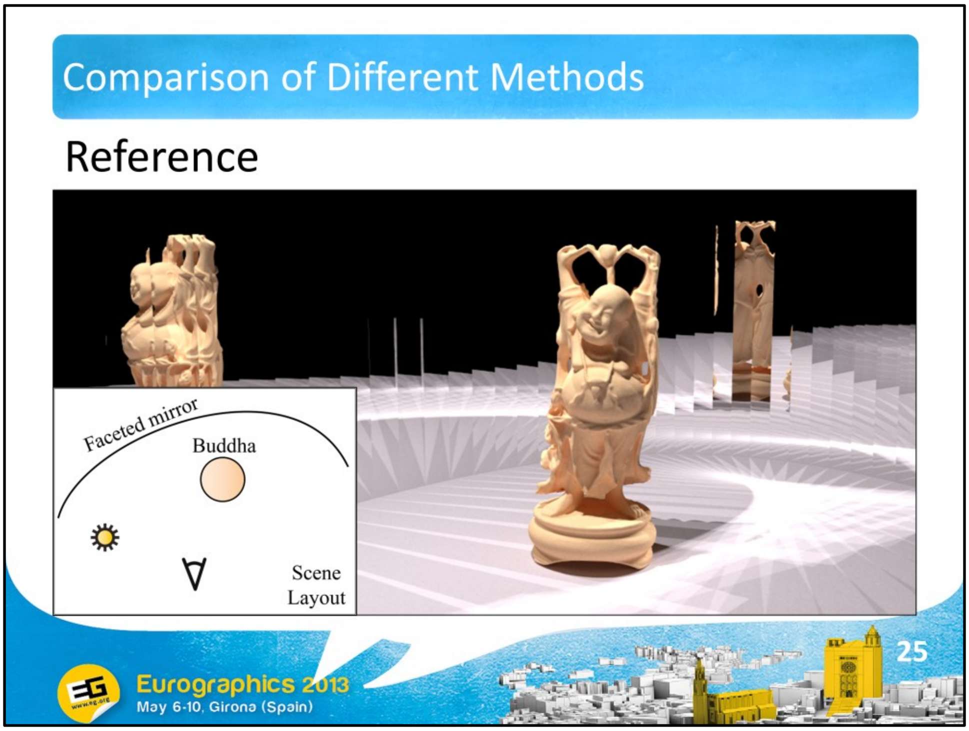


Now we will compare different popular methods with regularization.
This is a simple scene with caustics.

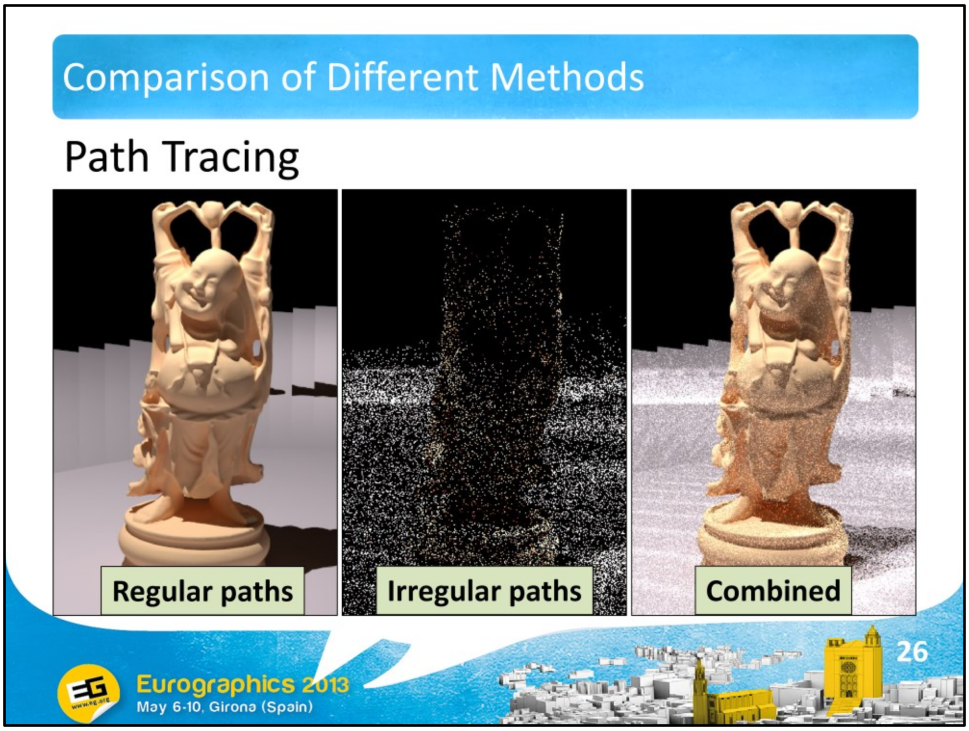


PT cannot sample **both caustics and reflected caustics**, thus requiring a lot of regularization.

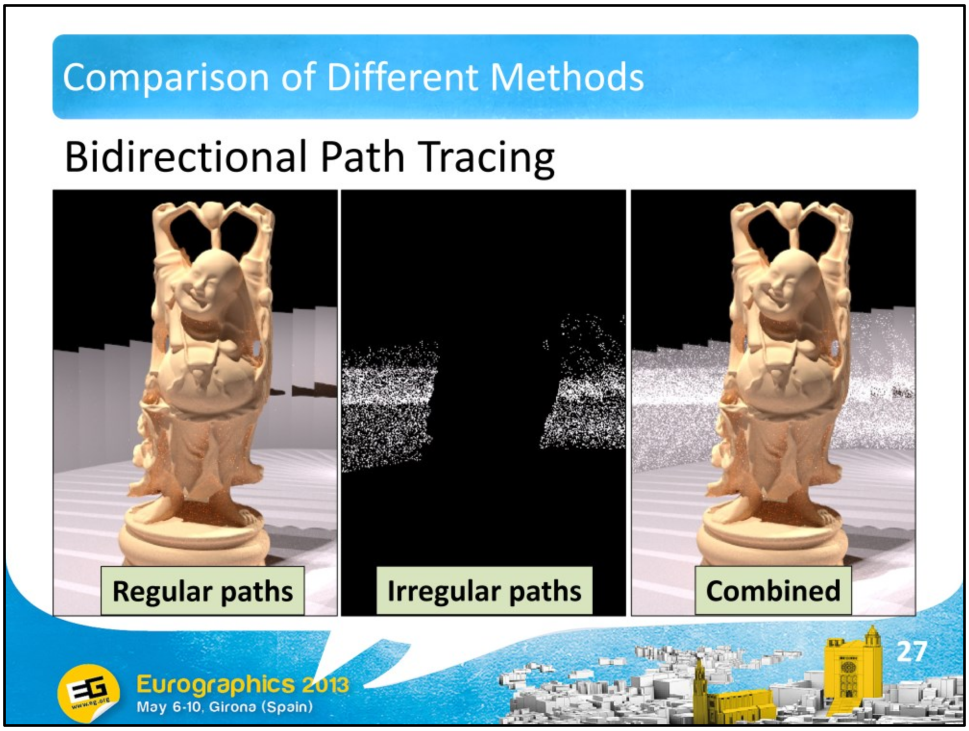


BDPT only requires regularization when the path cannot be constructed with edge connection, thus requiring to regularize only **reflected caustics**. Note that the amount of noise caused by irregular paths is **high** because the regularization angle is small.

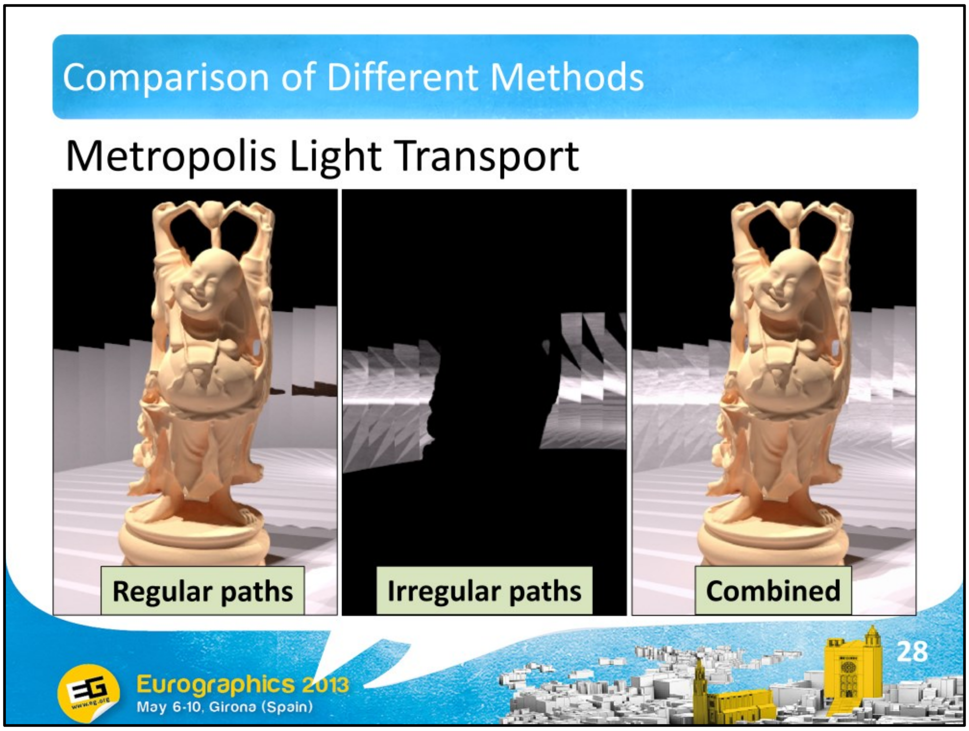


Finally, MLT solves the noise problem by implicitly **caching** the path it has already found before as a current state of Markov chain.

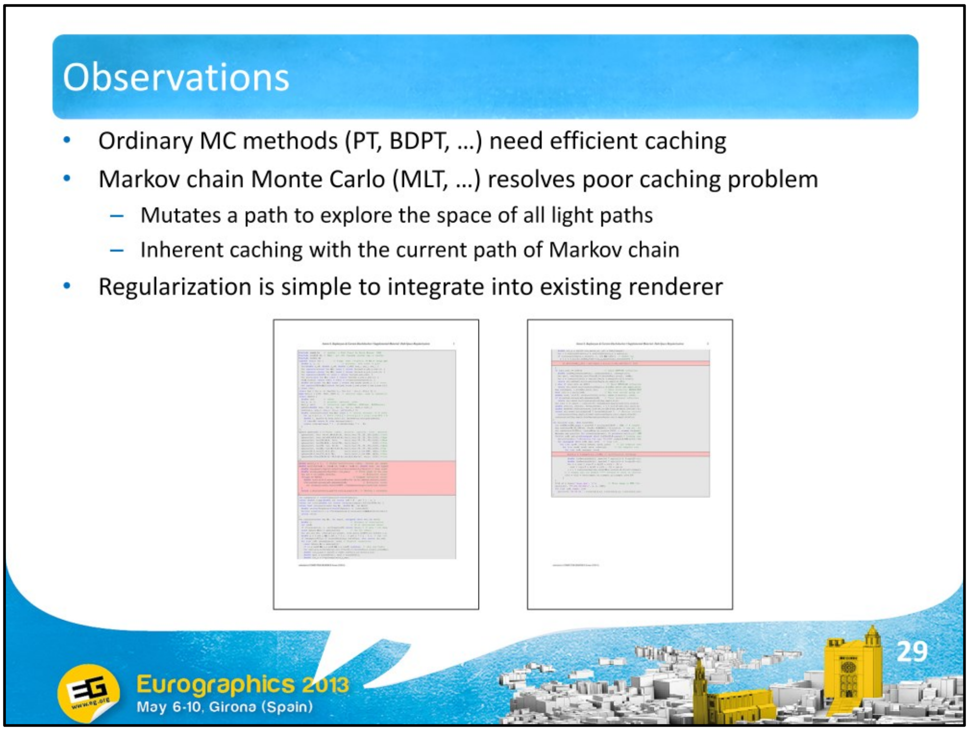


As we have seen, the ordinary Monte Carlo methods suffer from **noise**, thus are usually used with some efficient caching, such as photon map.
However Markov chain based methods, such as Metropolis light transport, naturally **resolve** this issue by caching the last important path as a current state of Markov chain.

Regularization is also **simple** to implement. Here is a code of a minimalistic path tracing. Regularization requires some changes only to the **evaluation routine** for specular BRDF. The required additions are marked in **red**.

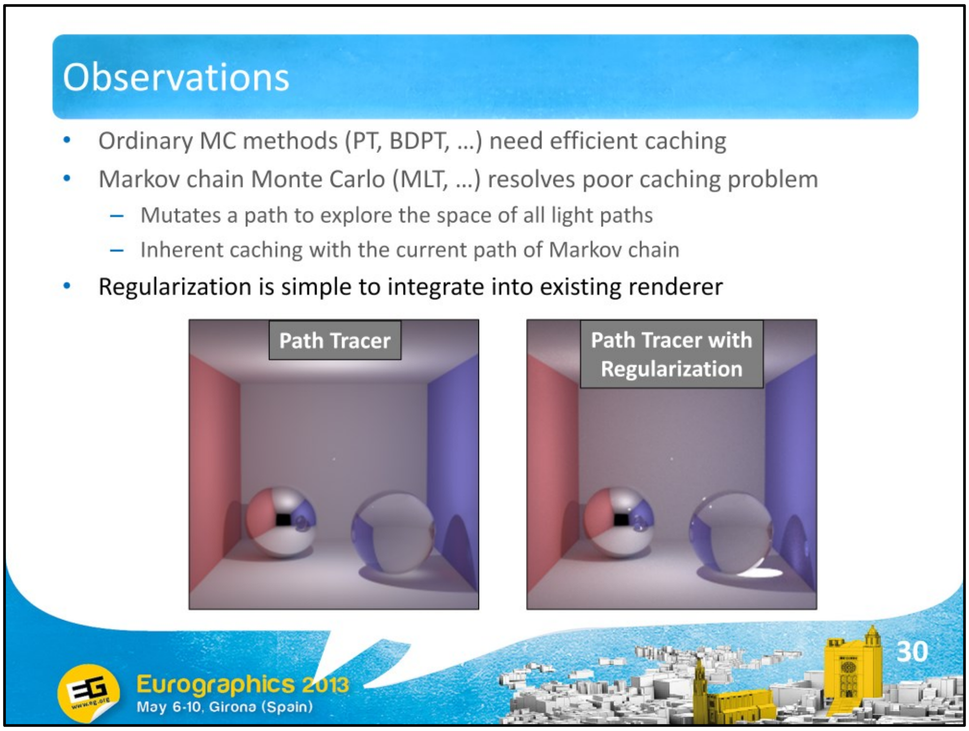


And you can see that with these small changes path tracing can already **handle caustics** of all kinds from point lights.

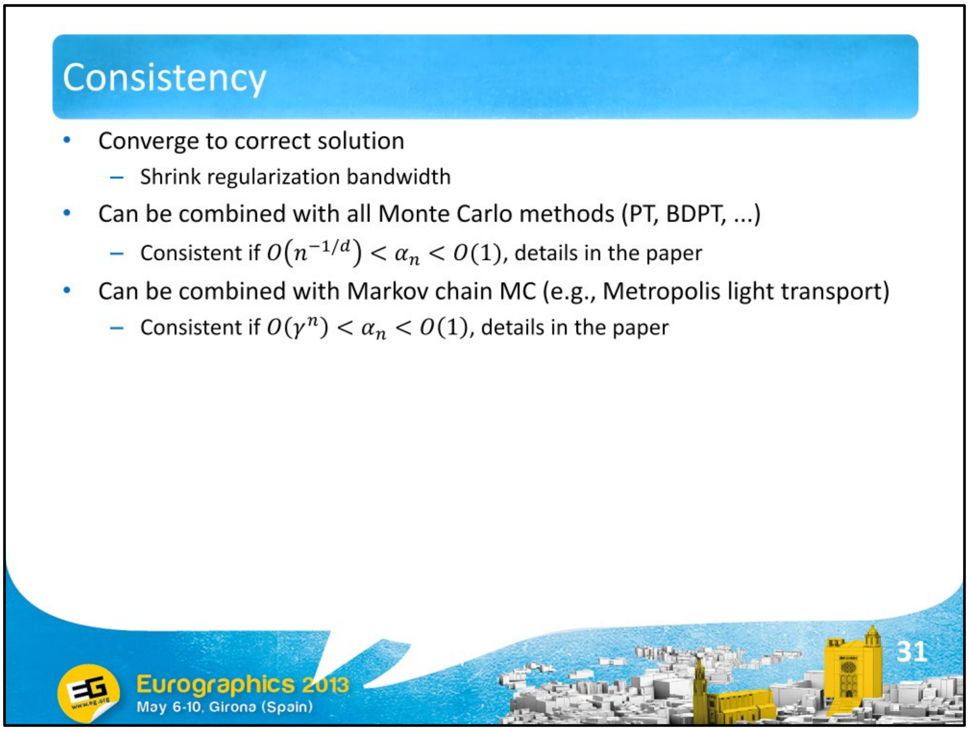


In order to make sure the regularization converges to **correct solution** in the limit, we need to **decrease** the regularization angle throughout the integration.
The Monte Carlo methods have the **same shrinkage conditions** as progressive photon mapping, which is not surprising.
However the MCMC methods, such as MLT, require slightly **different rate**. Please see the **details** in the paper.

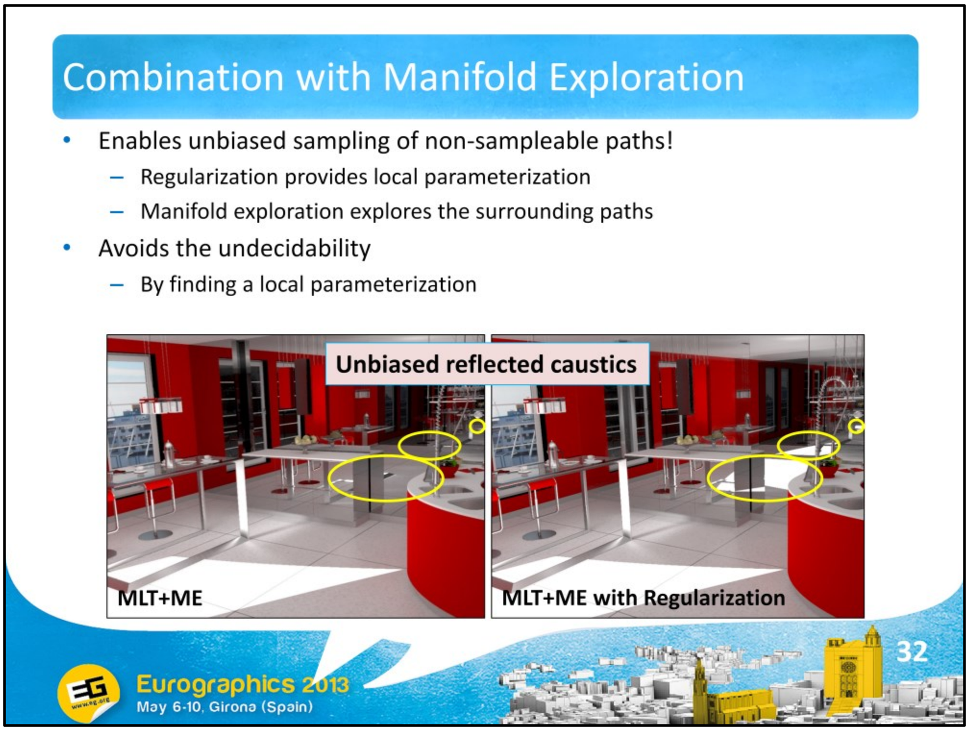


Recent manifold exploration mutation can connect two vertices through a chain of specular interactions, given a **valid local parameterization** for the connection.
We deliver a method, which **almost surely** provides the local parameterization for non-samplable paths, making the unbiased sampling of such paths **practical**.
This way we **avoid** stating undecidable problems.

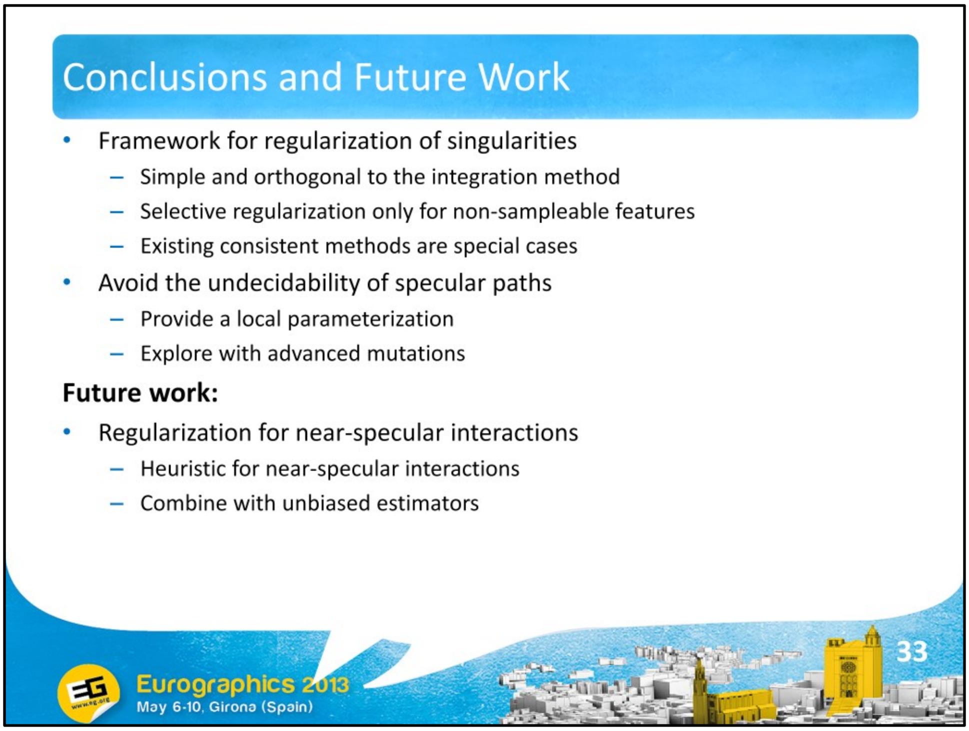


1. We presented a **selective regularization** framework for path space.
   a. This framework is **independent** on the integration method and can be used with any existing one.
   b. We have showed that the biased regularization is necessary **only** for irregular paths.
   c. We have explained why photon mapping **can** sample paths considered irregular by all local unbiased methods.
2. By providing proper seeds, we showed **for the first time** how to sample complex paths in unbiased way using the recent advanced methods.

We believe that our framework is just a **foundation** for many follow-up consistent methods, that can speed up the practical rendering.
For example, the interactions are usually not “black and white”, so **near-specular interactions** can be as difficult to handle as specular ones. We leave this problem for the future work, since it requires more **in-depth analysis** of interactions and the efficient combination with unbiased estimators.


Eurographics 2013
May 6-10, Girona (Spain)
Thank you for your attention
Questions?
Markov Chain PPM
Metropolis Light Transport
Metropolis Light Transport with Regularization