Thanks and welcome.

Monte Carlo methods always, in theory, offer exact solutions to any subsurface scattering problem. In practice, however, the computation resources to achieve a low-noise solution are prohibitive. Thus, we continue to see new methods to reduce the variance of Monte Carlo approaches, like these three works presented at SIGGRAPH this week.

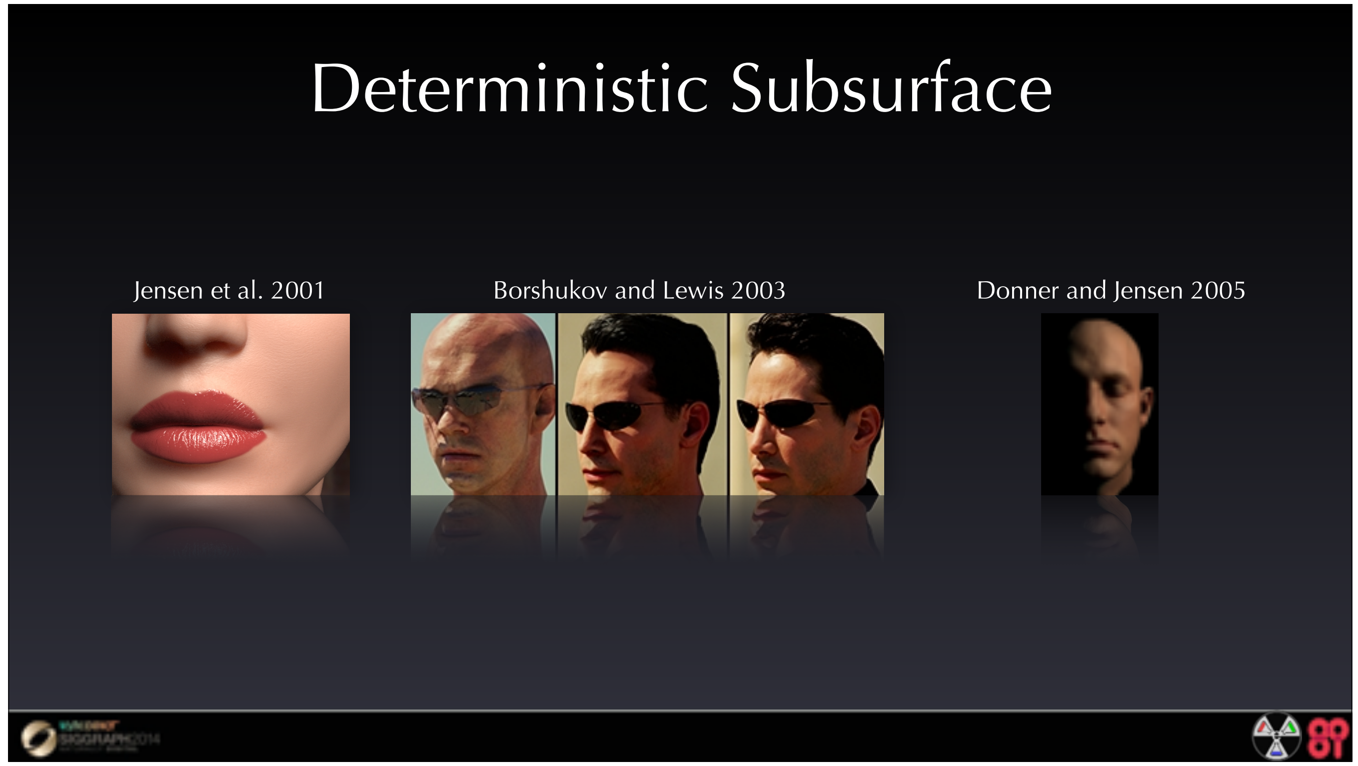

Deterministic approximations, such as these, remain highly relevant approaches to synthesising images of translucent media, such as human faces.

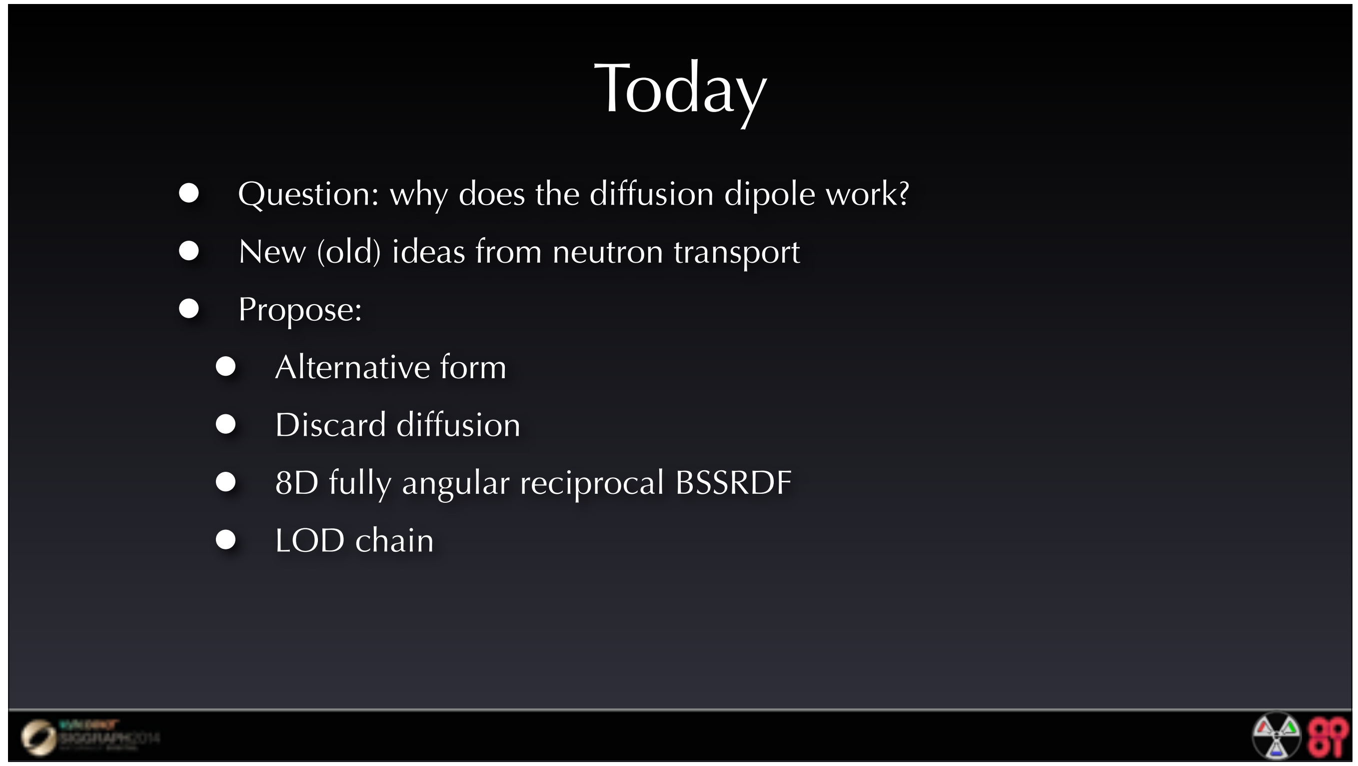

Today we describe a new approach to formulating an analytical approximate BSSRDF for rendering translucent materials that stems from asking the question: 'why does the diffusion dipole work?'. We propose that the positive and negative source configuration works because it is an approximation application of an exact statement from old neutron transport literature. Taking this connection further, we propose an alternative form for analytic BSSRDFs that discards diffusion and produces a fully angular, reciprocal BSSRDF that places the BSSRDF carefully in the LOD chain between explicit structure and BRDF.

We start by recalling a recent extension of the diffusion dipole, presented in 2011, called Quantized diffusion (QD).

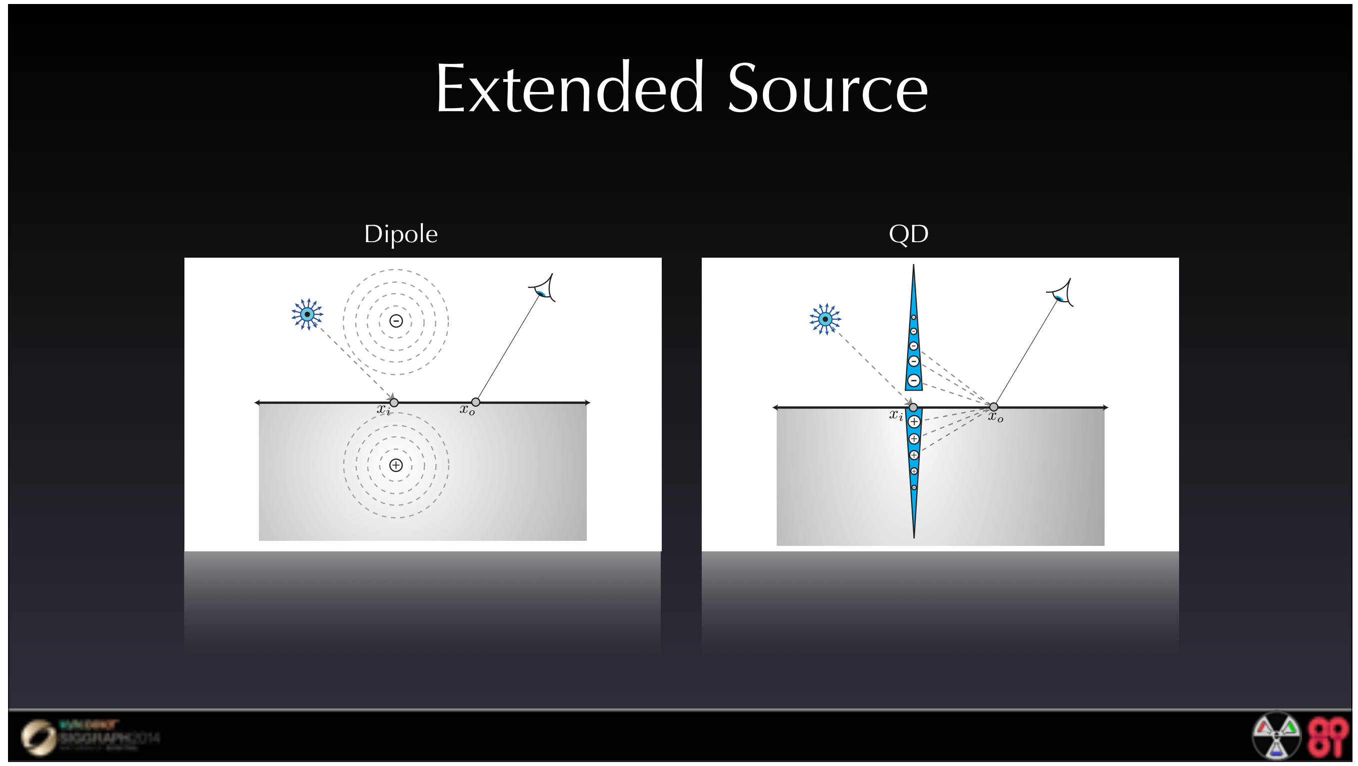

In QD, continuous distributions of positive and negative infinite-medium solutions combine to predict the existence flux at any surface location x_o where light is leaving the material due to illumination at x_i.

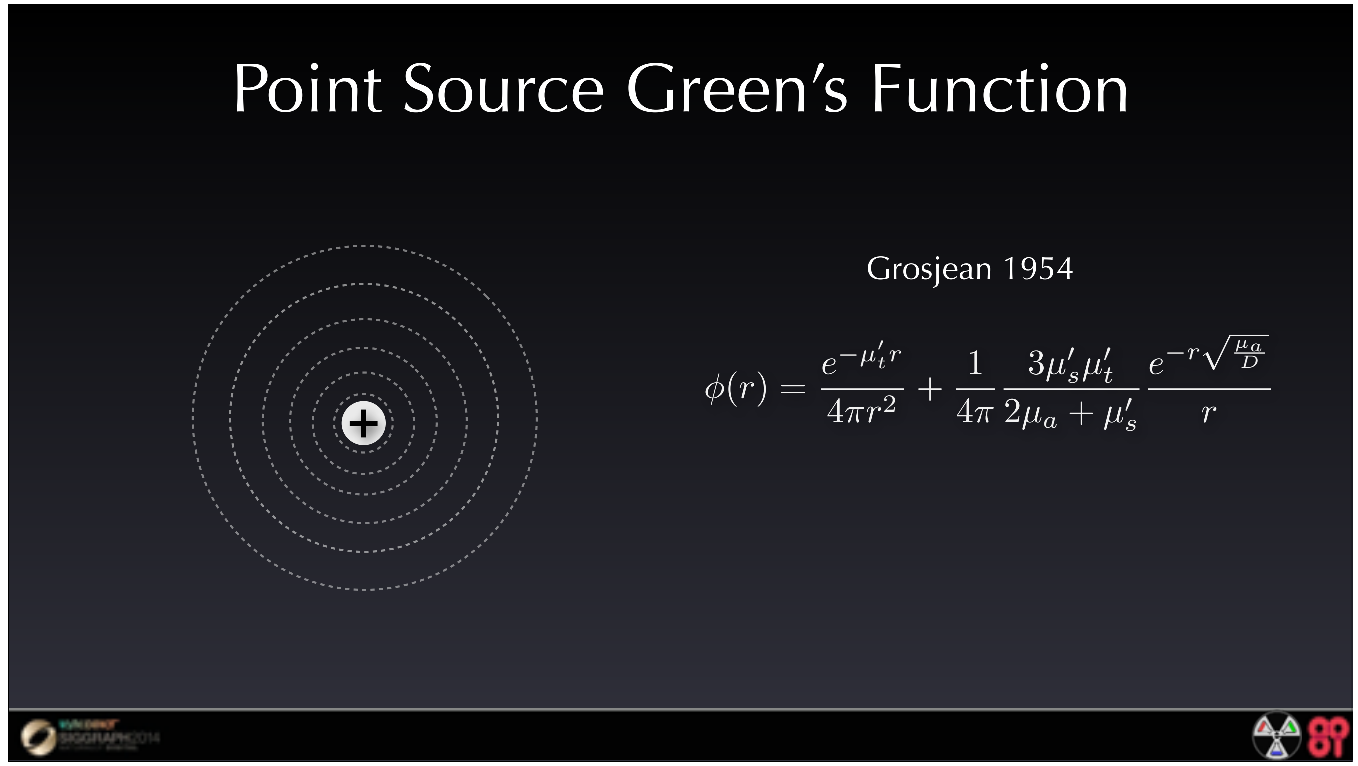

In the 2011 QD paper, Grosjean's 1954 analytic approximation to the fluence about a point source in infinite space was used as the foundation of the positive and negative source distributions defining the solutions for the semi-infinite medium.

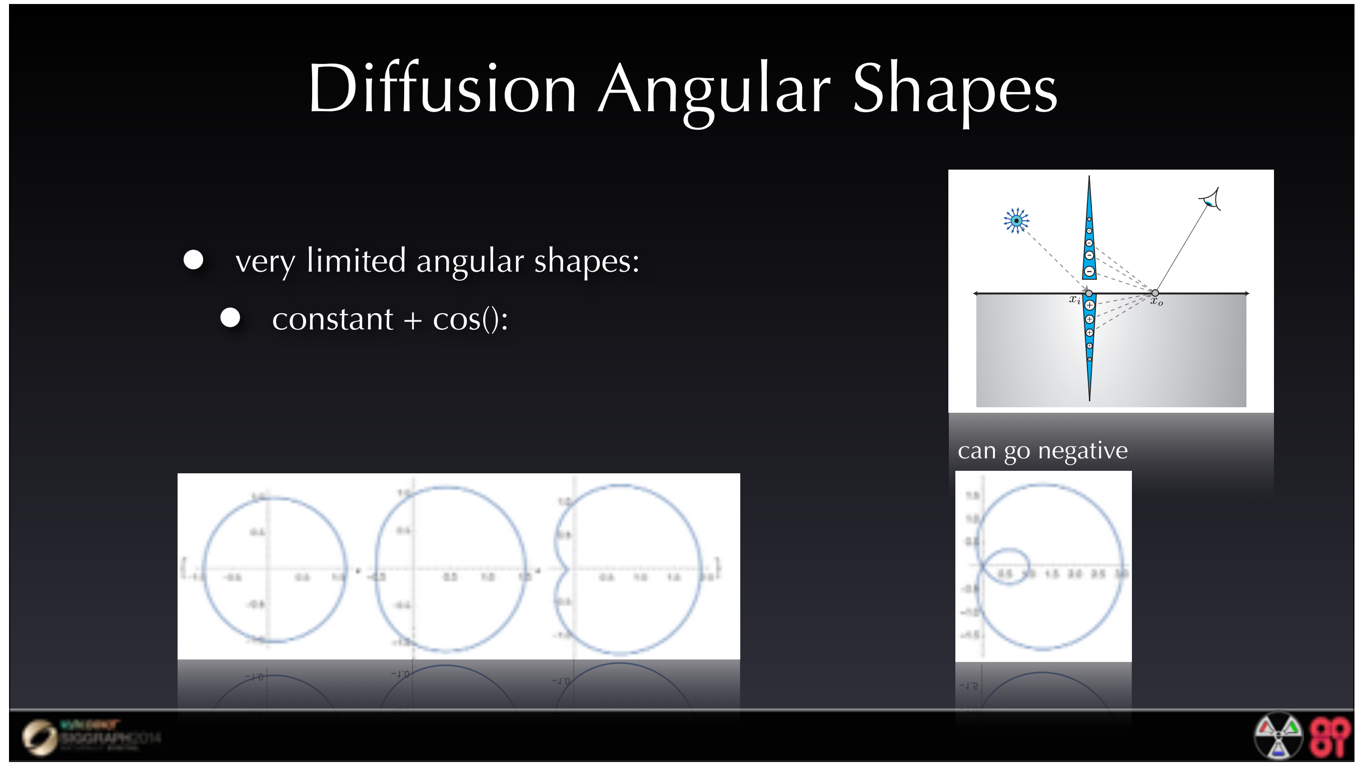

However, a particular issue with diffusion theory limits the angular accuracy of such approaches. This stems primarily from the limitation that diffusion solutions have a angular distribution everywhere that is restricted to a sum of a constant in angle and a cosine in angle. Worse, in some regions (such as very near a source) the radiance predicted by diffusion theory is actually negative, which is impossible physically.

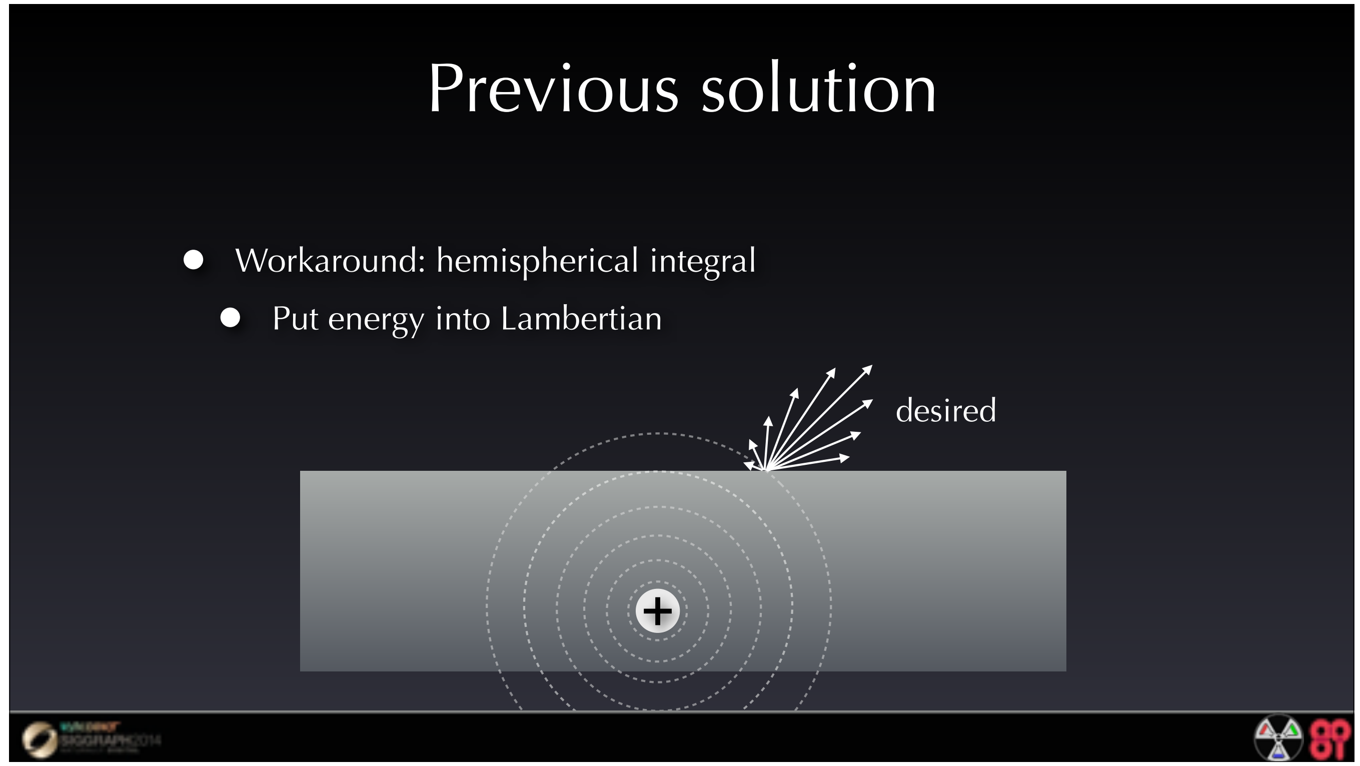

To avoid this issue in previous applications of diffusion theory at a surface, the workaround was to integrate the total energy predicted to leave the surface at some location (which is always positive), and redistribute that outgoing energy into a Lambertian shape. So instead of getting an angular varying distribution like this…

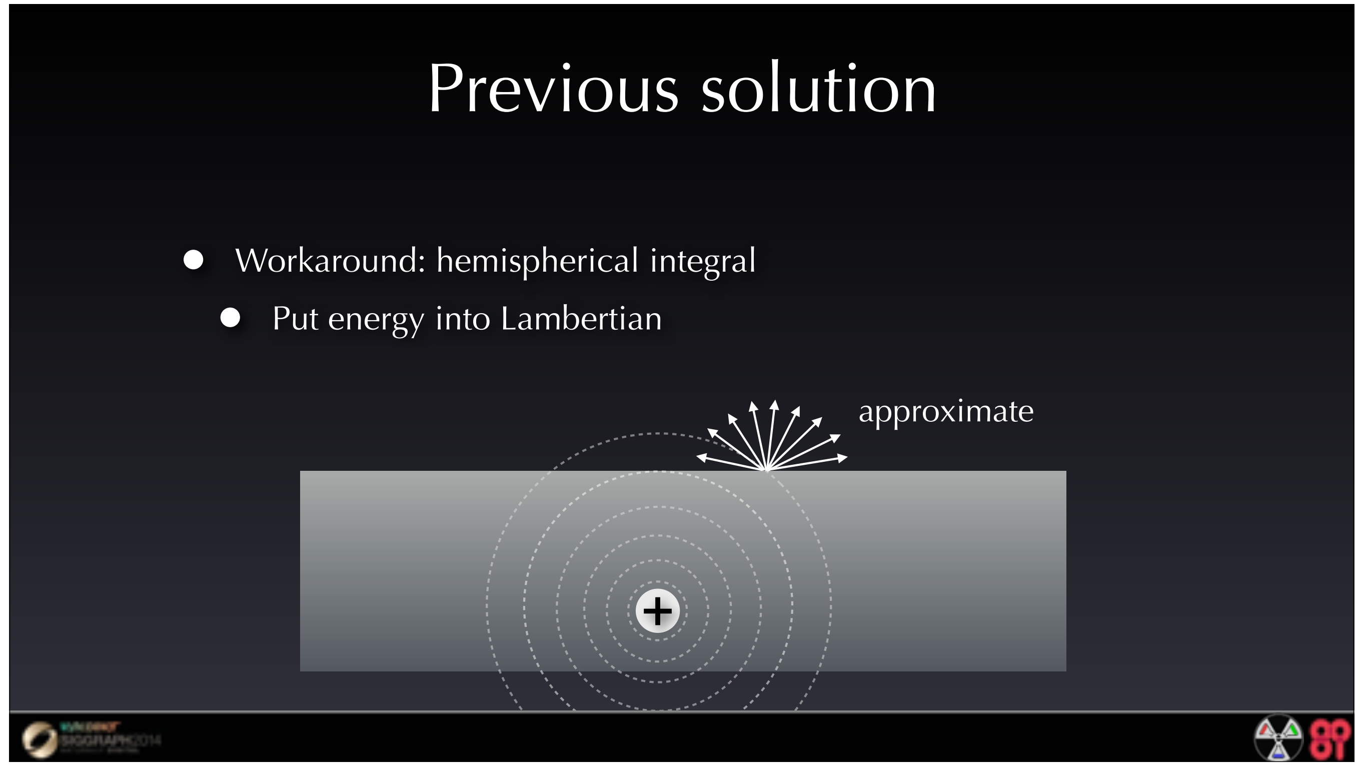

…we instead suffer the following approximation to avoid negativity.

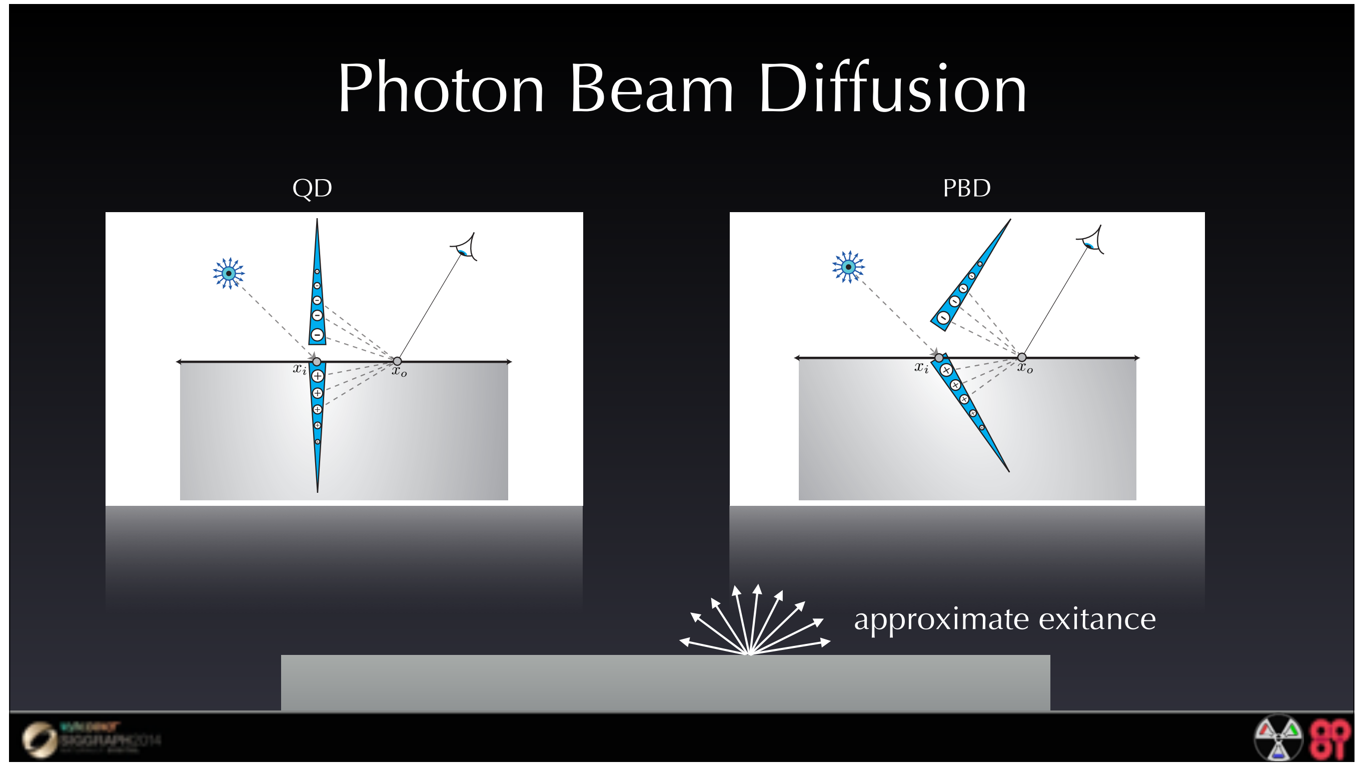

This limitation is inherent in the QD and Photon Beam Diffusion BSSRDFs. Both assume Lambertian existence from the surface to avoid the negativities predicted by the diffusion approximation.

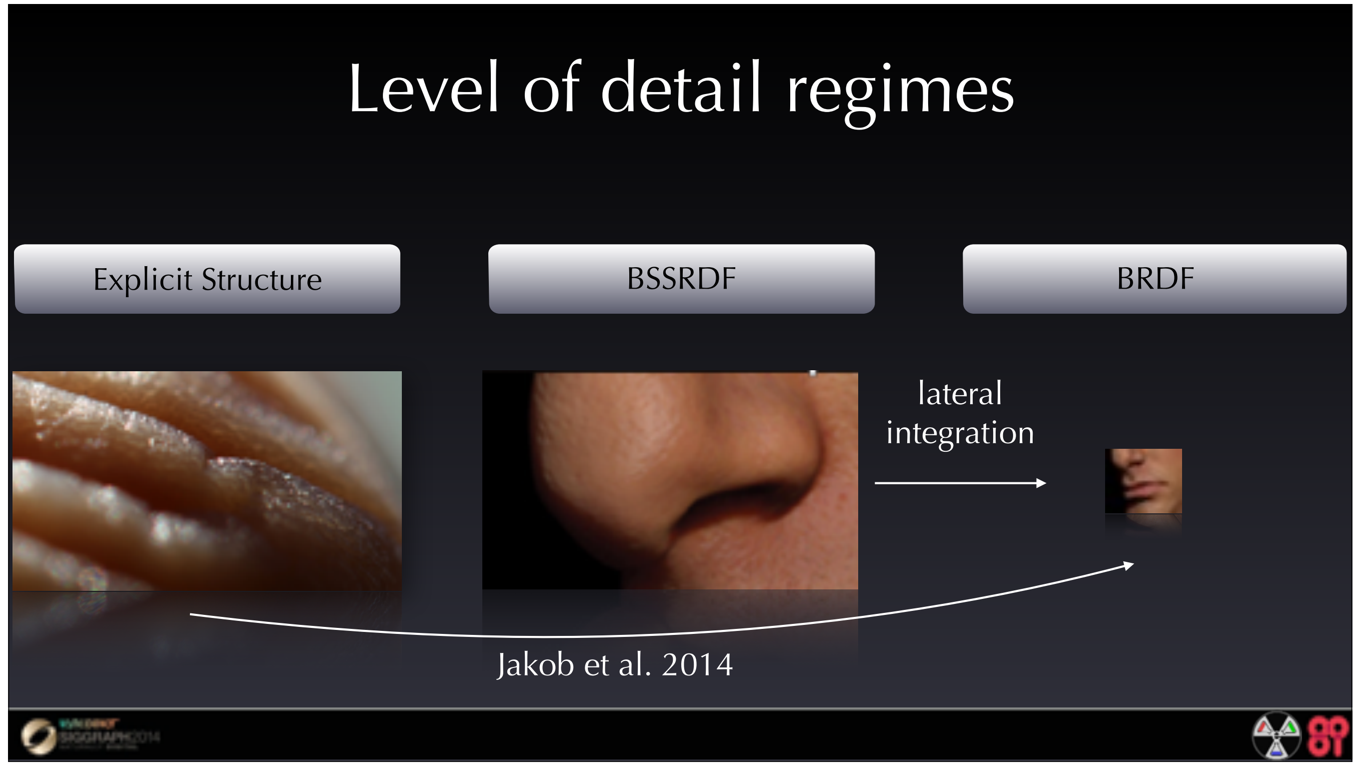

This is problematic when considering the BSSRDF and its neighbour, in the LOD chain, the BRDF, useful for rendering a surface where negligible lateral bleeding of light is visible (such as a very distant human face). We can imagine a continuous shot where a zoom from very close tissue (left) using explicit structure is blending into a BSSRDF model (middle) and finally to a BRDF model for efficiency in the distance (right). However, previous analytic BSSRDFs provide incredibly poor angular accuracy. The BRDF solution for multi-layered scattering layers is efficiently computable and highly accurate in angle (see, for example, this years paper by [Jakob et al. 2014]).

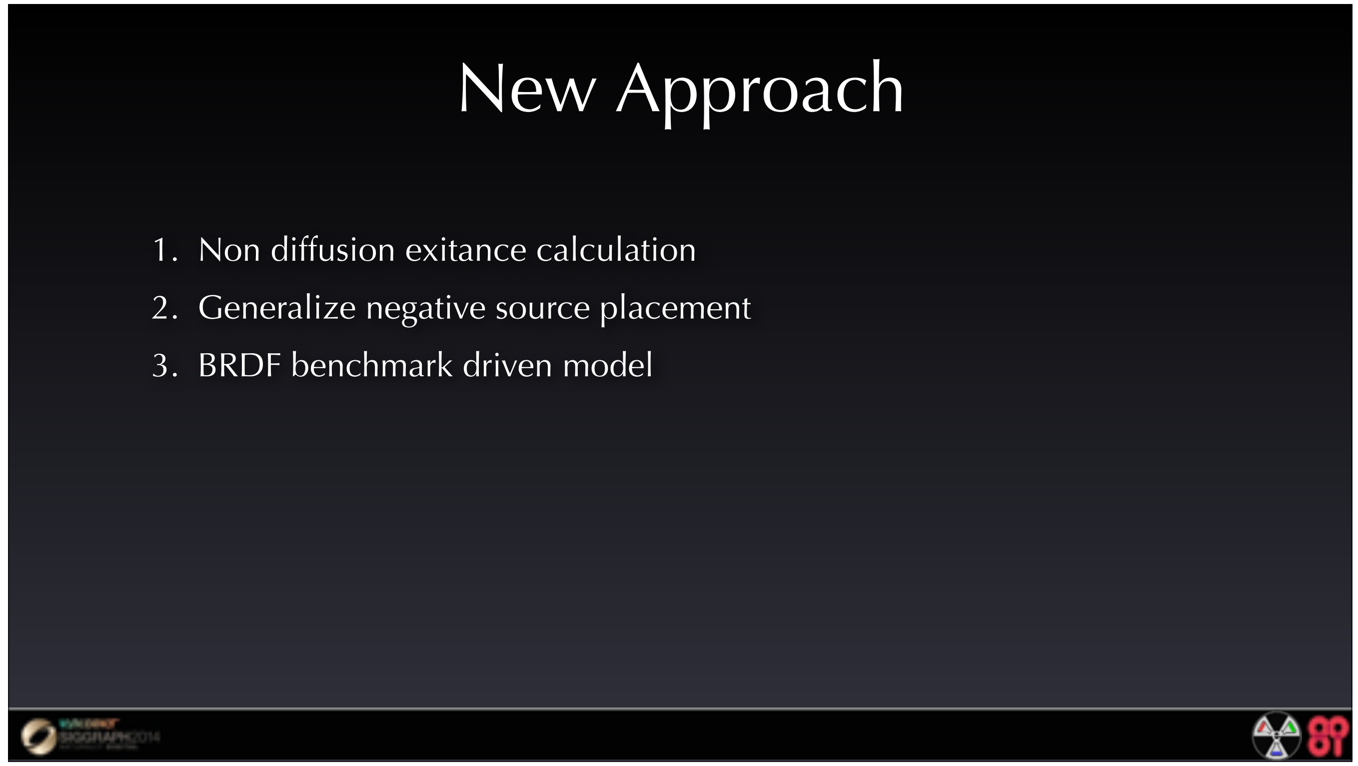

To produce a BSSRDF with improved angular accuracy, we propose to move away from a diffusion existence calculation. In addition, for better accuracy, we generalize previous placement of negative sources outside the medium. To ensure the most accurate match in the LOD chain to the BRDF, we drive step 2 by the known BRDF solution.

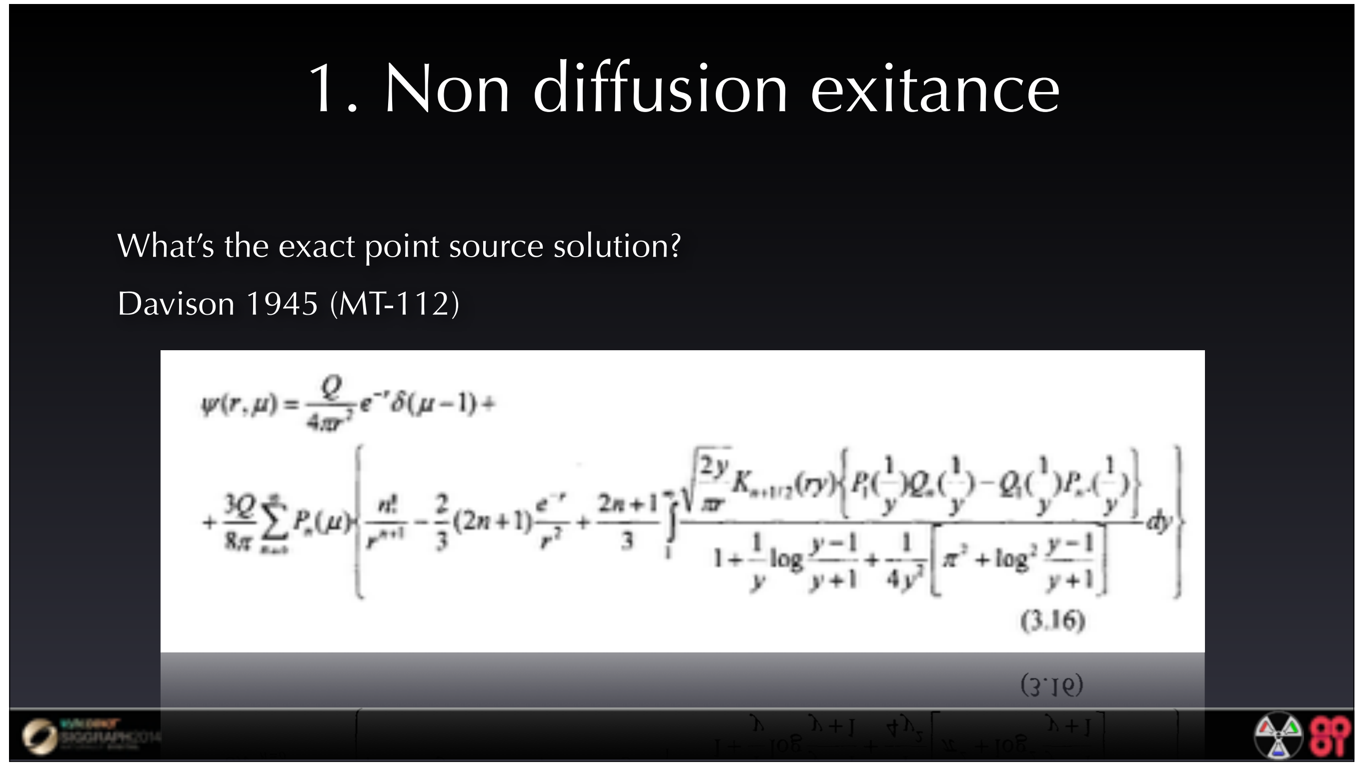

Step 1: if we move away from diffusion angular limitation, what, then, is the exact solution, in infinite media, for the angular distribution of radiance anywhere relative to an isotropic point source (assuming isotropic scattering)? The solution was first written down by Davison in 1945. Yikes! It involves a delta function in angle (necessary: this is the uncollided energy attenuated as it moves away from the point source, and directed in a singular direction—-directly away from the point source). The scattered solution is an infinite sum involving integrals of Bessel functions.

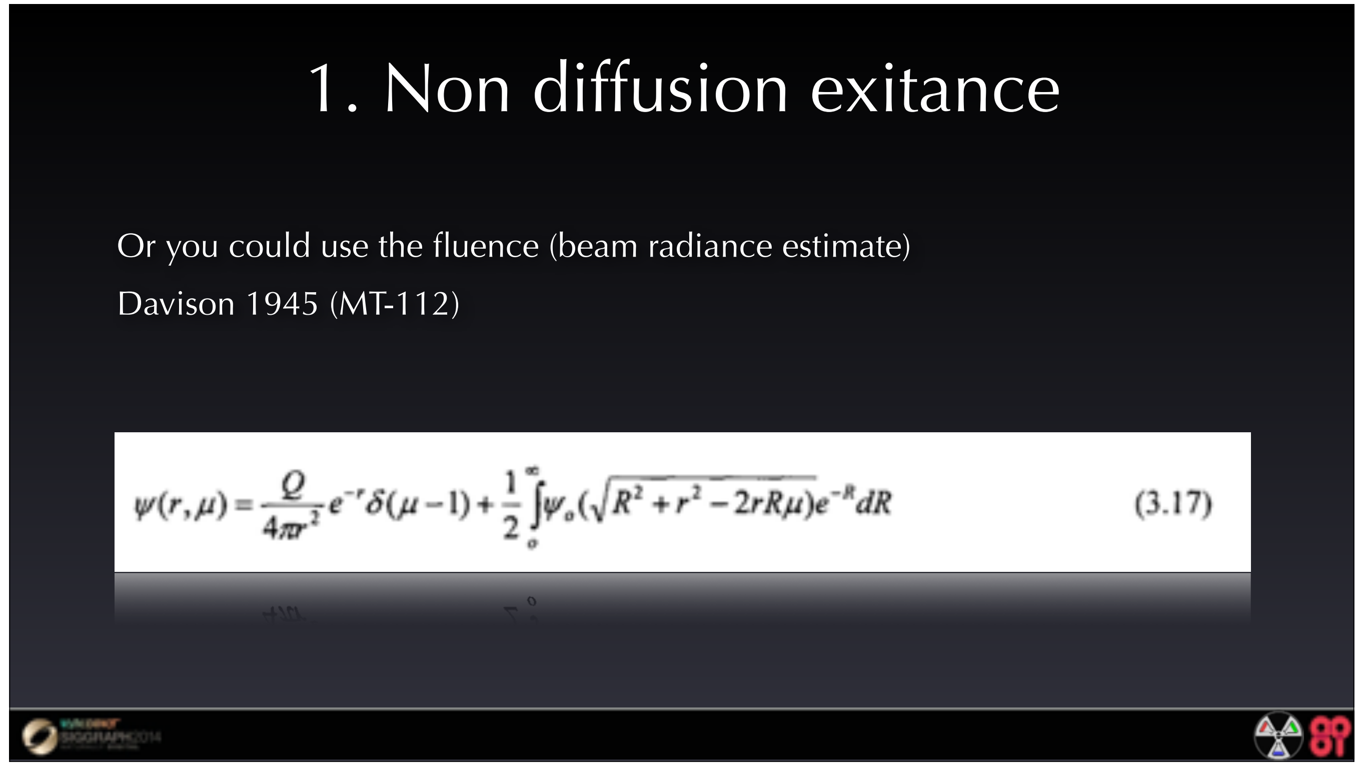

However, Davison noted in the same paper that the result is easily described (thanks to the assumption of isotropic scattering), as a 1D integral of the fluence (which we have a simple approximate form for from Grosjean).

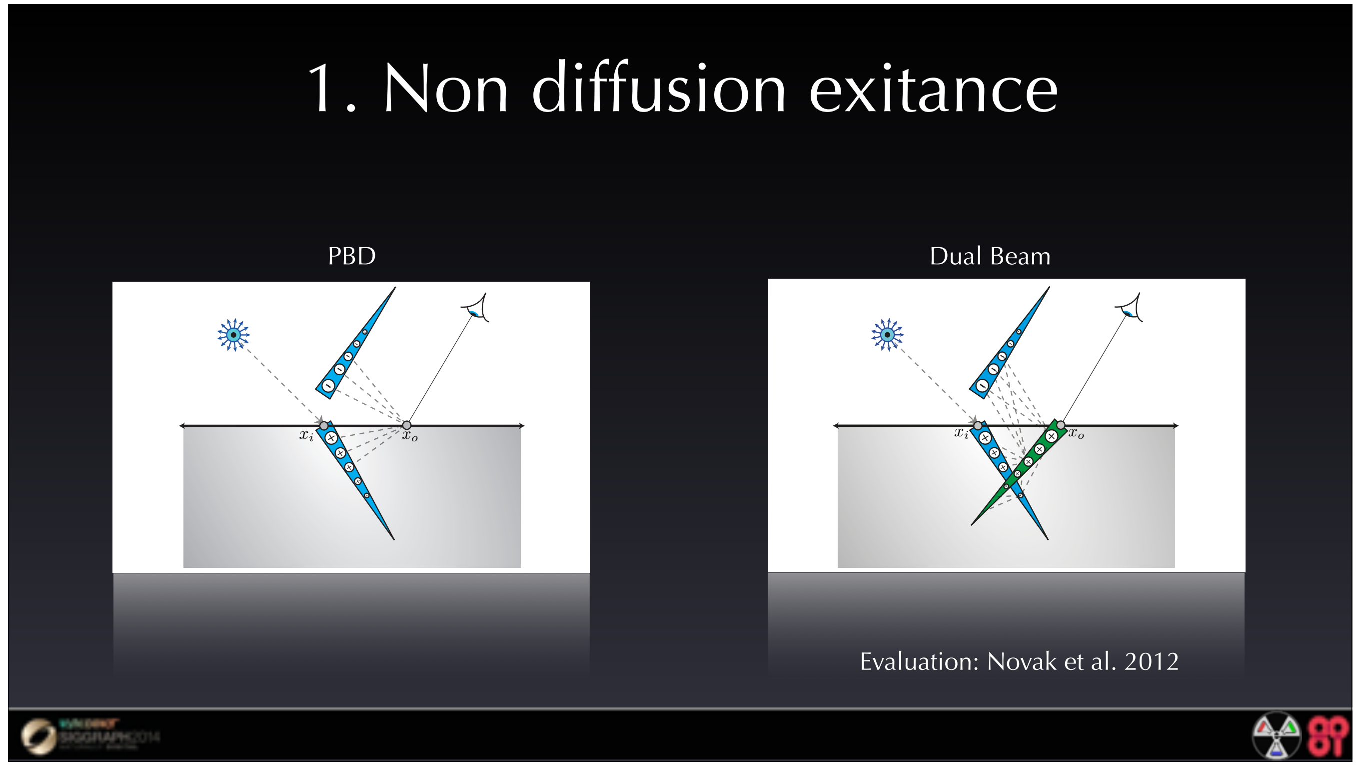

To apply this to the extended-source BSSRDF ideas of previous work, we replace the outgoing calculation of flux at x_o with Davison's line integral of the subsurface fluence. This is illustrated as the green continuum of positive 'detectors' inside the medium.

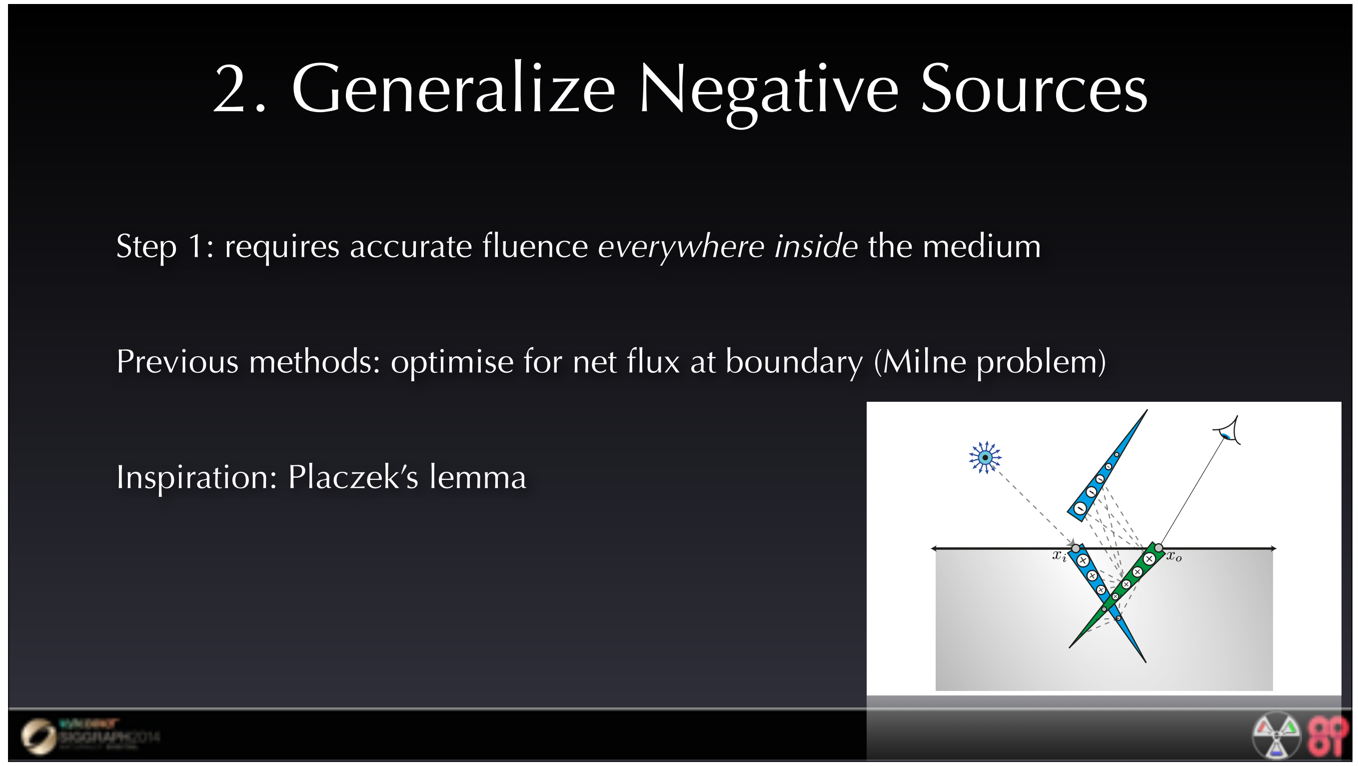

However, in contrast to previous work, we now require accurate fluence everywhere INSIDE the medium. This is where we leave diffusion theory behind, which imposed a boundary condition based on the Milne problem for solving for the negative source placement such that some property AT THE BOUNDARY was satisfied. To return somewhat to first principles, let's consider why the dipole idea works at all to begin with. Why is it that a negative copy of the internal source at just the right distance outside the medium does so well at predicting radial-exitant flux at the boundary (in particular, far from the incident location)? The first part of the answer is related to the linearity of the transport equation: if A and B are solutions to the equation, then so too is A+B. Thus, the sum of two infinite medium solutions, each satisfying the equation, are also solutions (except, of course, at boundaries, which must then reduce useful forms of their combination to only one or a few choices).

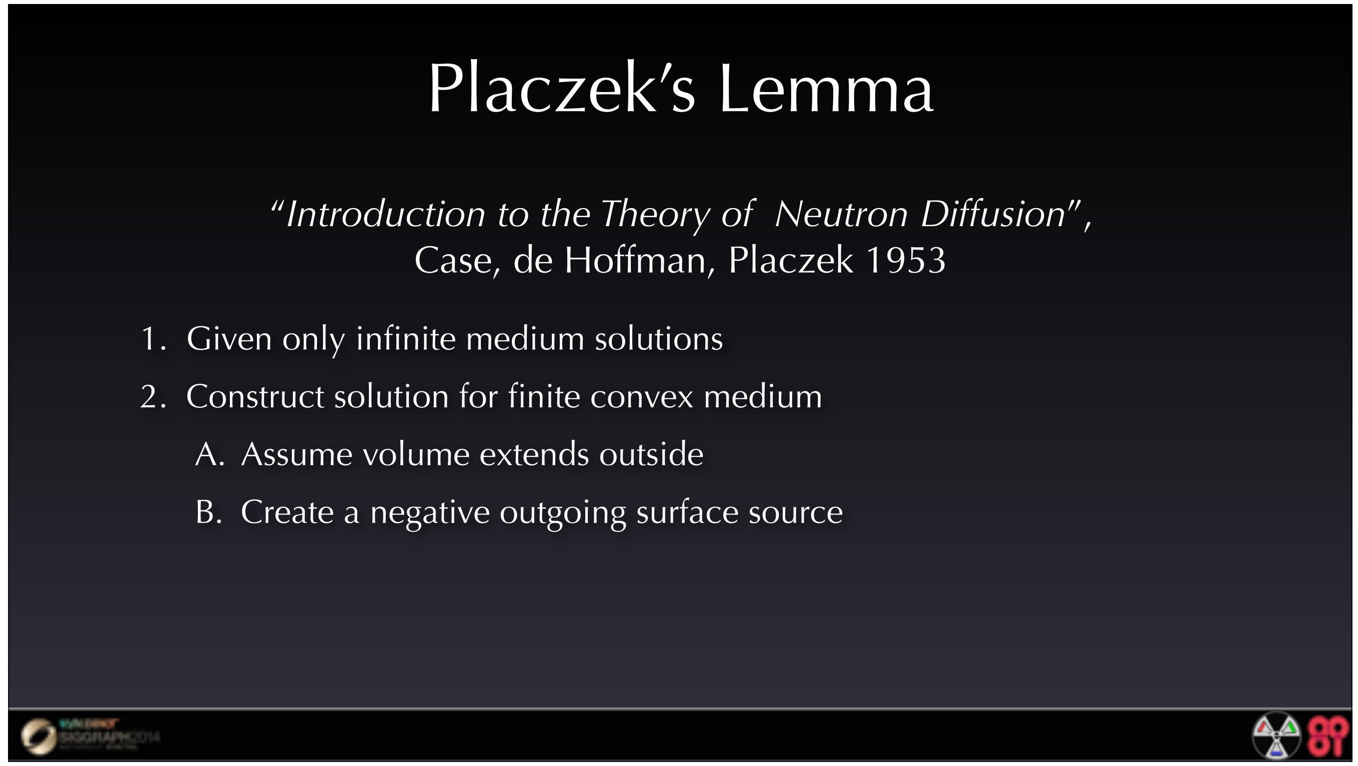

In 2011 we described the method of images (the placement of positive and negative infinite medium solutions to solve non-infinite problems) as an approximate application of Placezek's Lemma. Placezek's Lemma appears in a 1953 classic text on neutron transport (not easy to find) and indeed, sounds a lot like the method of images at first glance. However, it is an EXACT method for constructing solutions to finite problems (assuming a convex volume) using only infinite medium solutions.

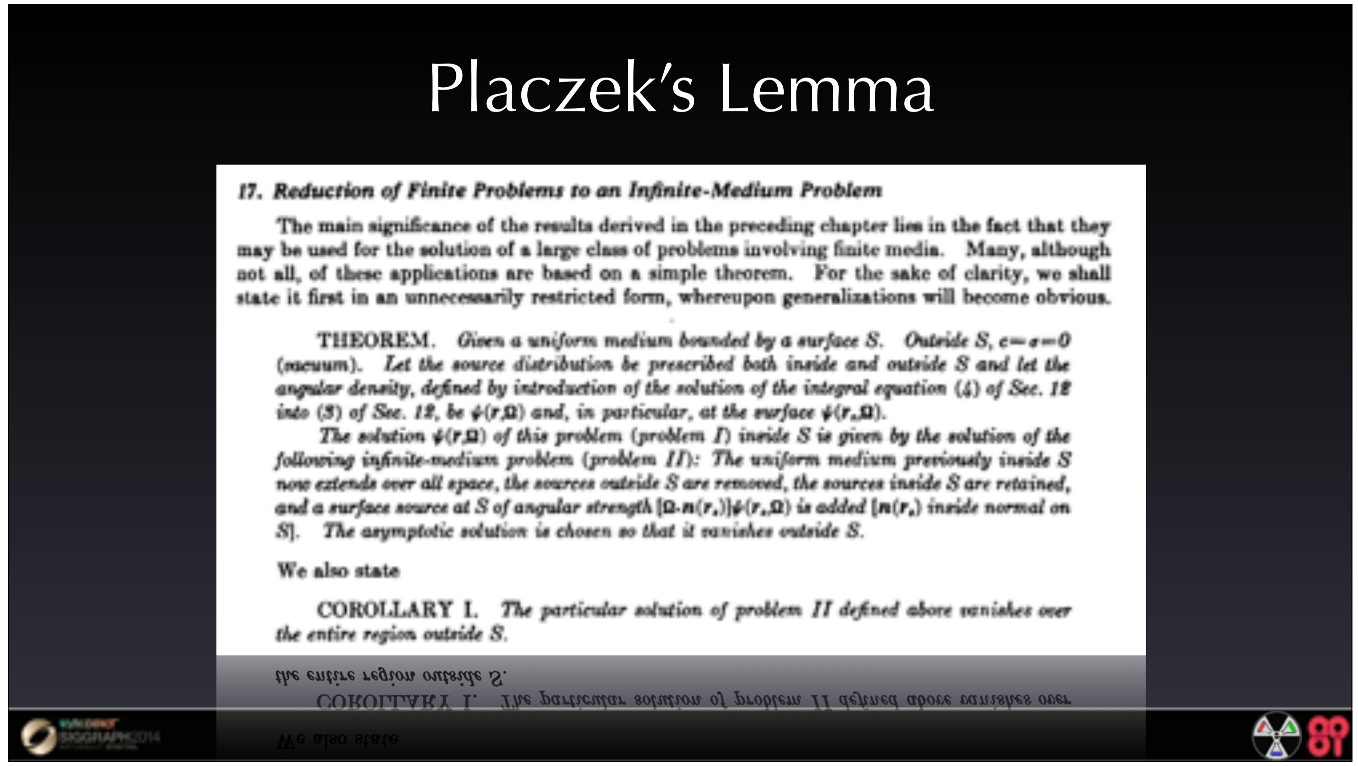

The statement of it is not terribly complex…

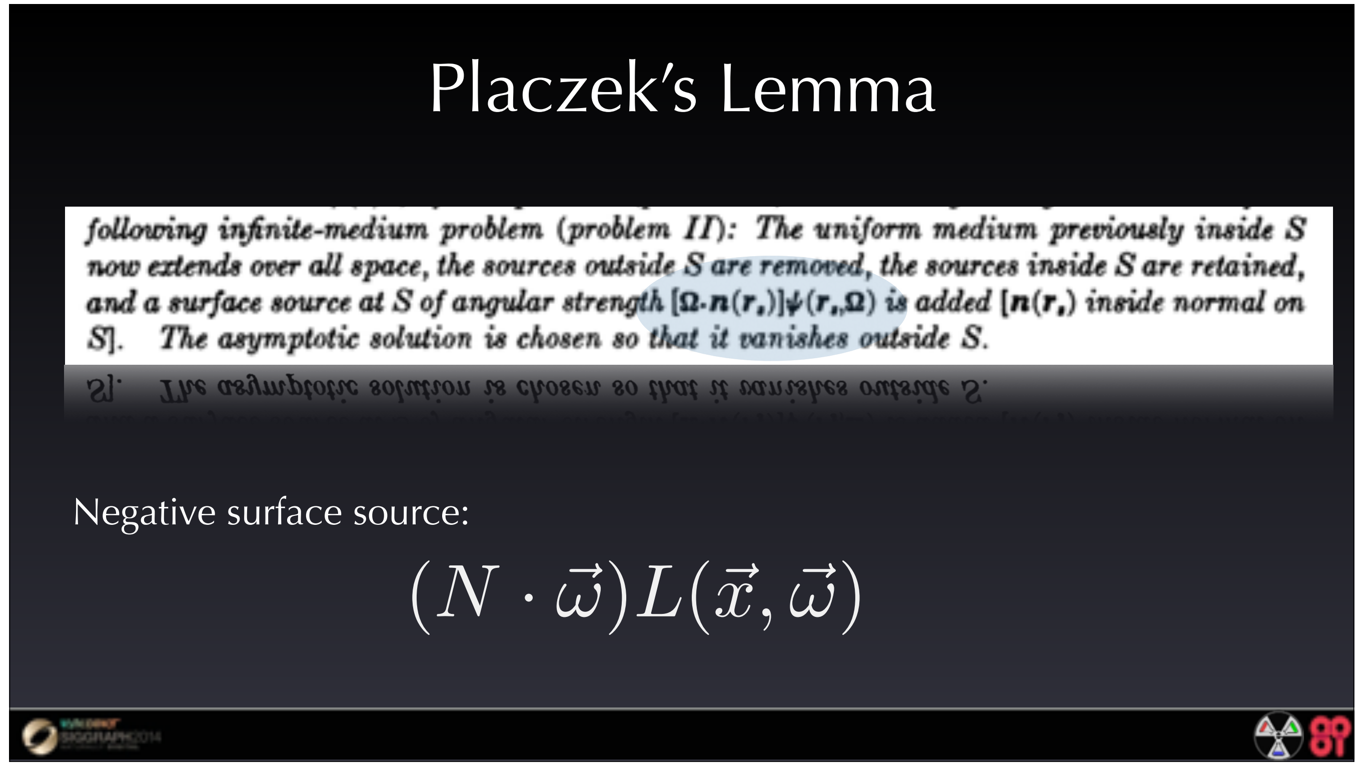

The construction is as follows: the scattering volume is now considered to extend everywhere in an infinite medium (i.e. outside the original volume). A new negative surface source is added to the system using this simple equation involving the normal dotted with the direction being considered. Now, the internal sources (the reduced-intensity beam in our case) and these new outward facing negative surface sources combine to produce the correct solution everywhere in the medium, using infinite medium green's functions. However, there is a chicken an egg problem: to know the magnitude of the negative surface source to produce this exact result requires knowing L(x,omega): the solution at the surface for the finite problem—-the very solution we seek in the first place! However, it is still worthwhile to consider this Lemma.

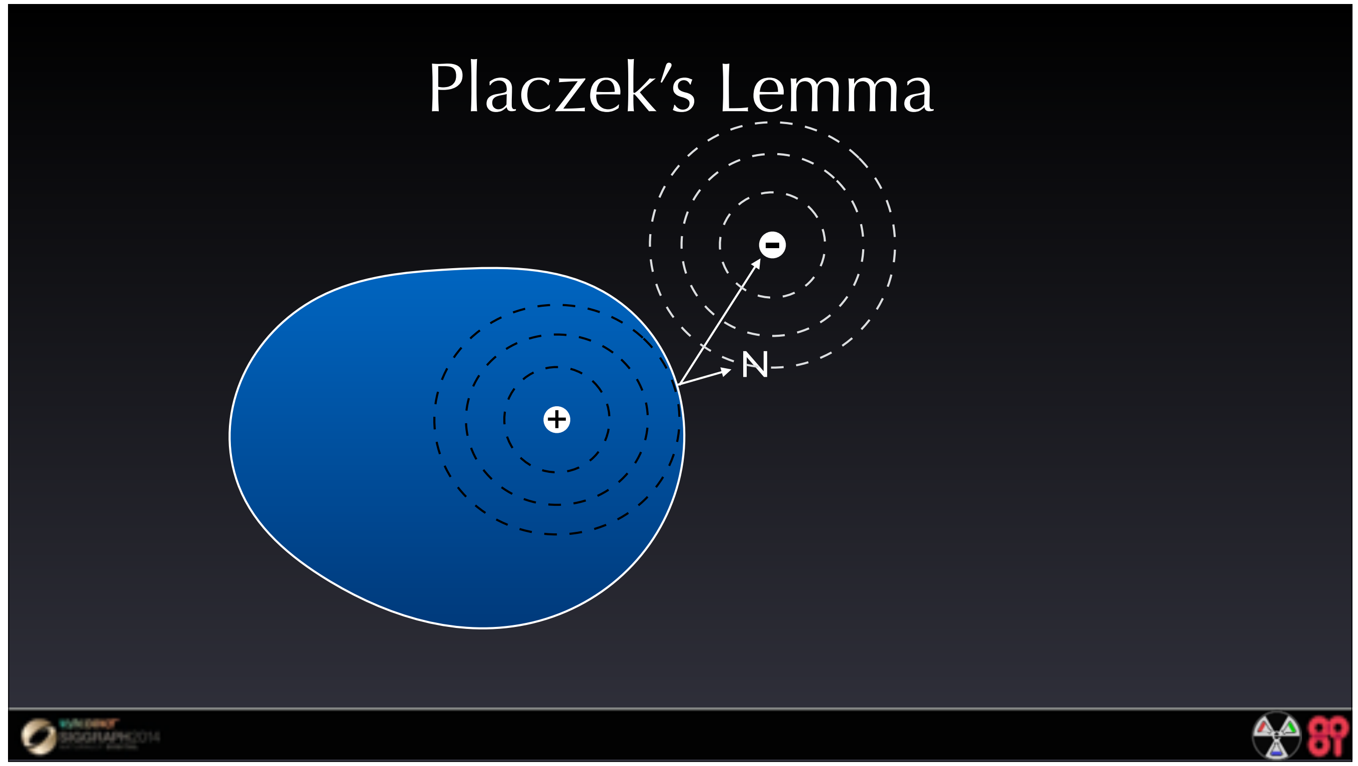

To see how this lemma works, consider a single point source inside the finite medium. (An extended reduced-intensity source is a superposition of these). For each surface location, and for each outgoing direction at that position, a negative source is added, which will first scatter at some point and place, effectively, a negative isotropic point source outside the volume.

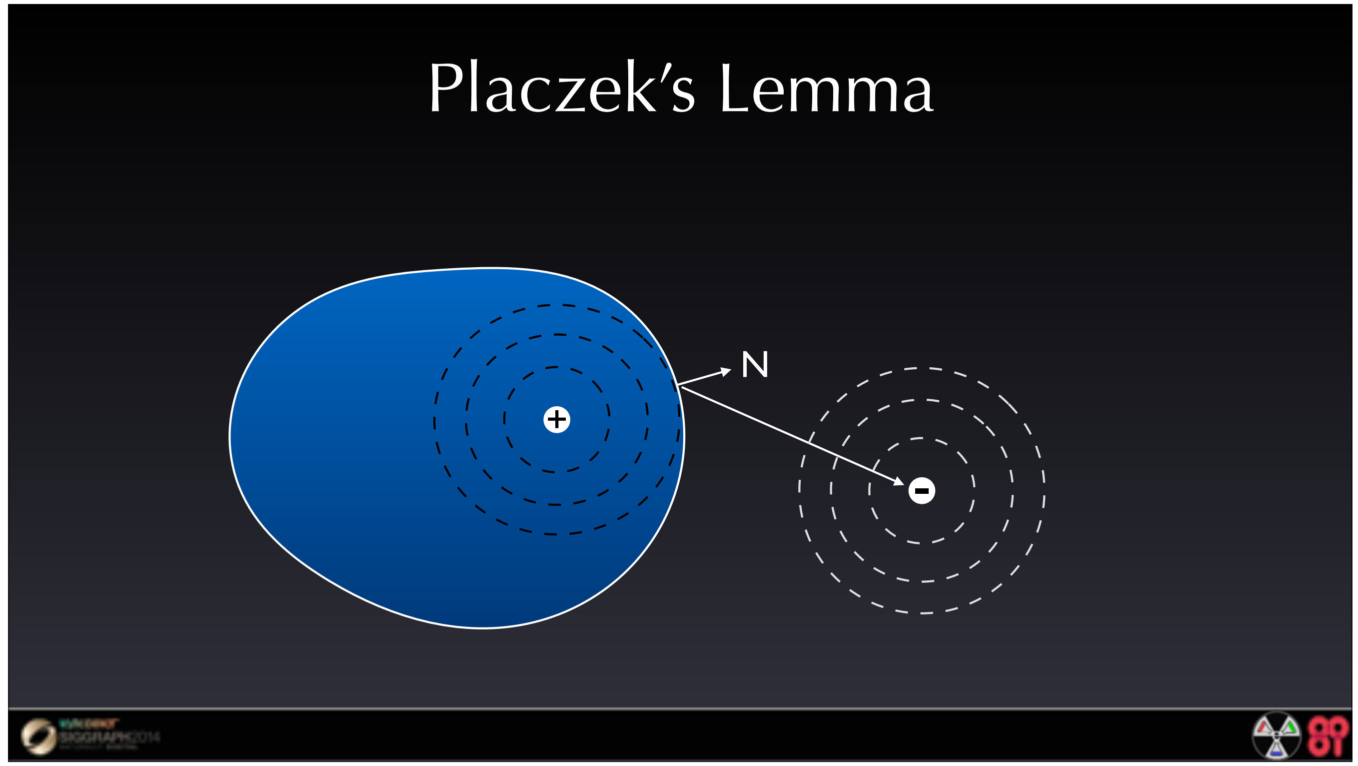

However, this happens for ALL outgoing directions, not just one, like in previous papers (like Jensen et al. 2011).

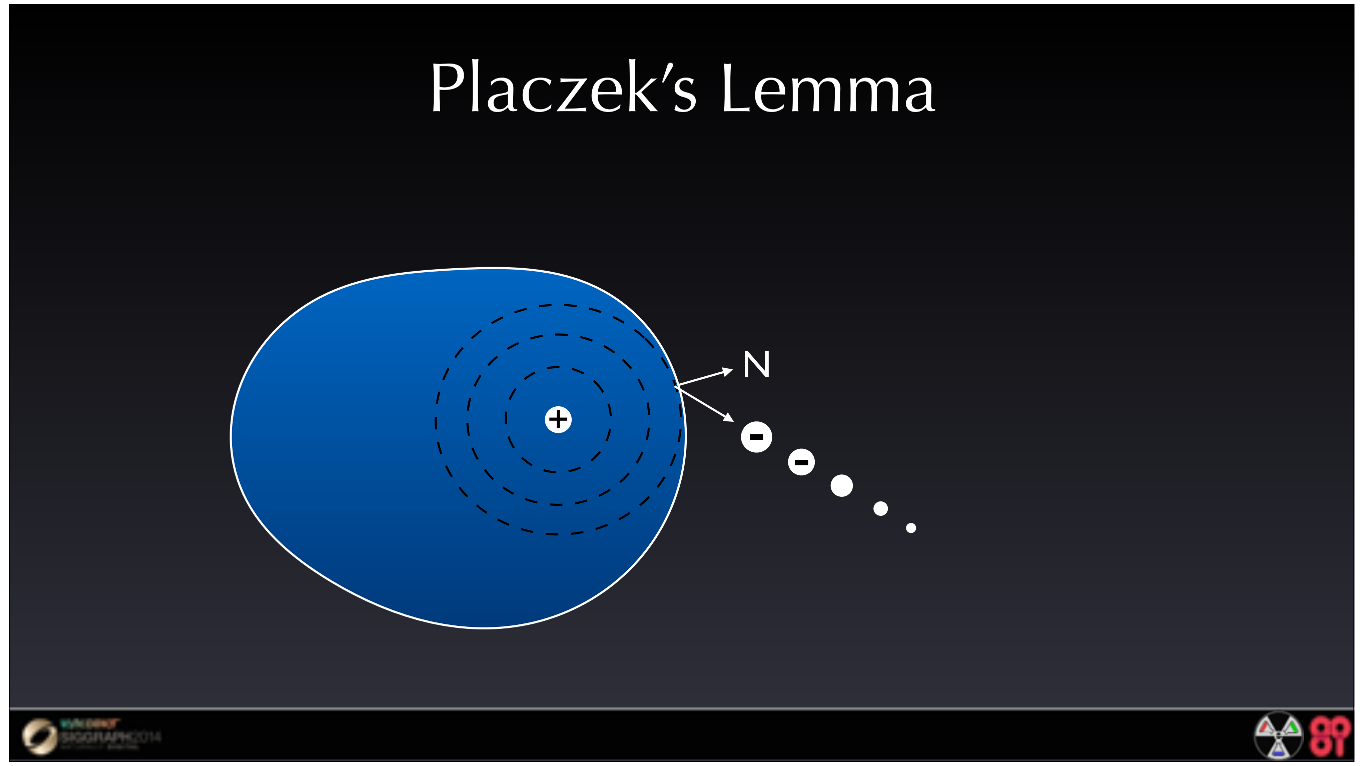

And since each of these surface sources begins at the surface, and the scattering medium extends outside the medium now, this effectively produces a continuous superposition of negative 'charges' outside the volume, due to the original point source inside.

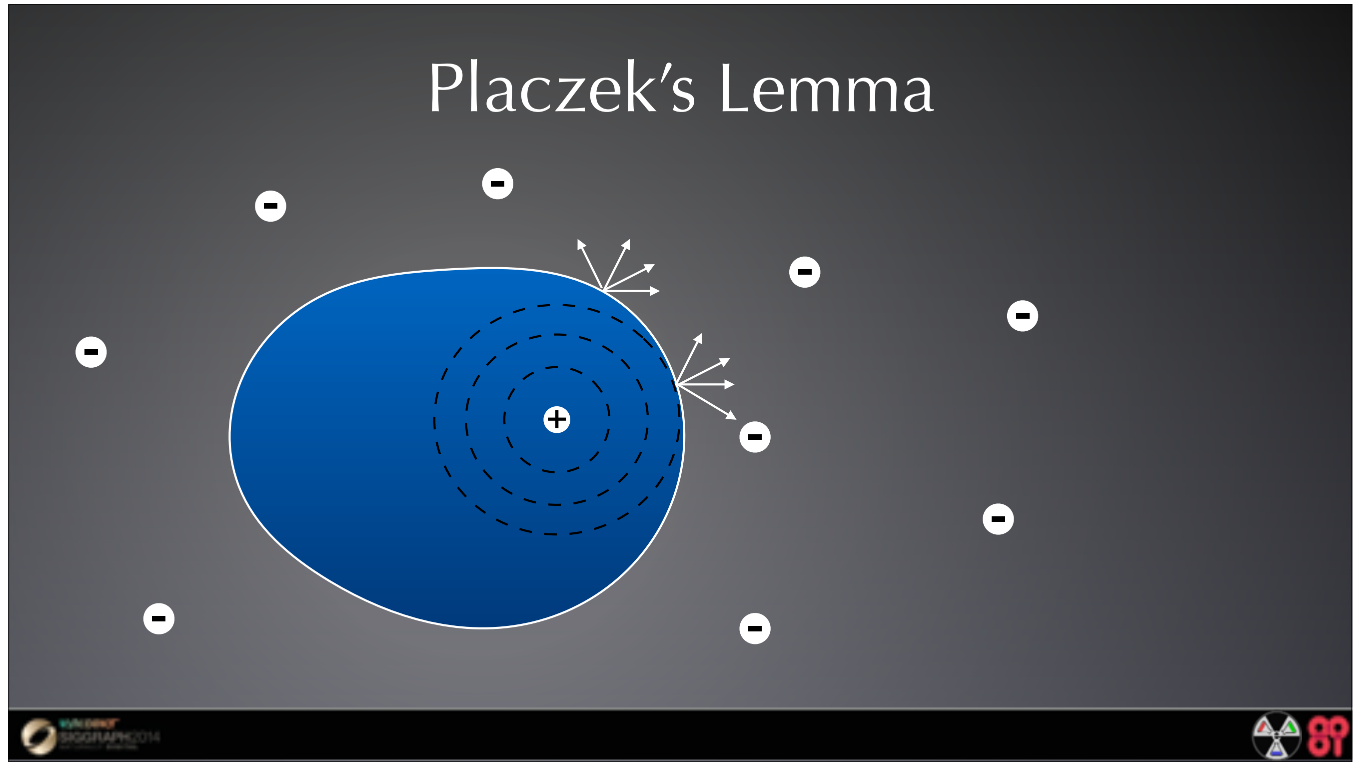

Considering all these extended beams, arising from all surface locations and angles, creates a continuous sea of negative energy outside the volume.

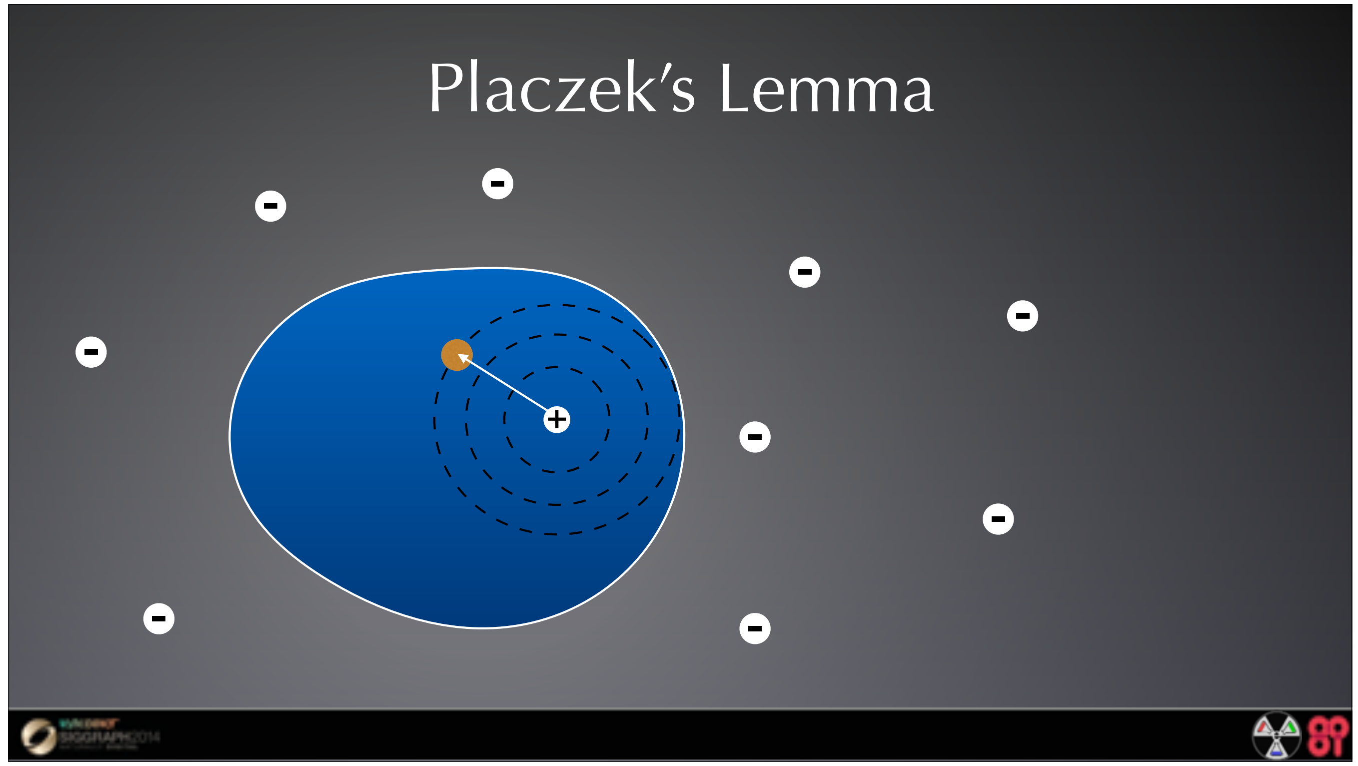

The exact solution anywhere in the medium is the sum of the original contribution from the interior positive source (computed with an infinite medium Green's function), and…

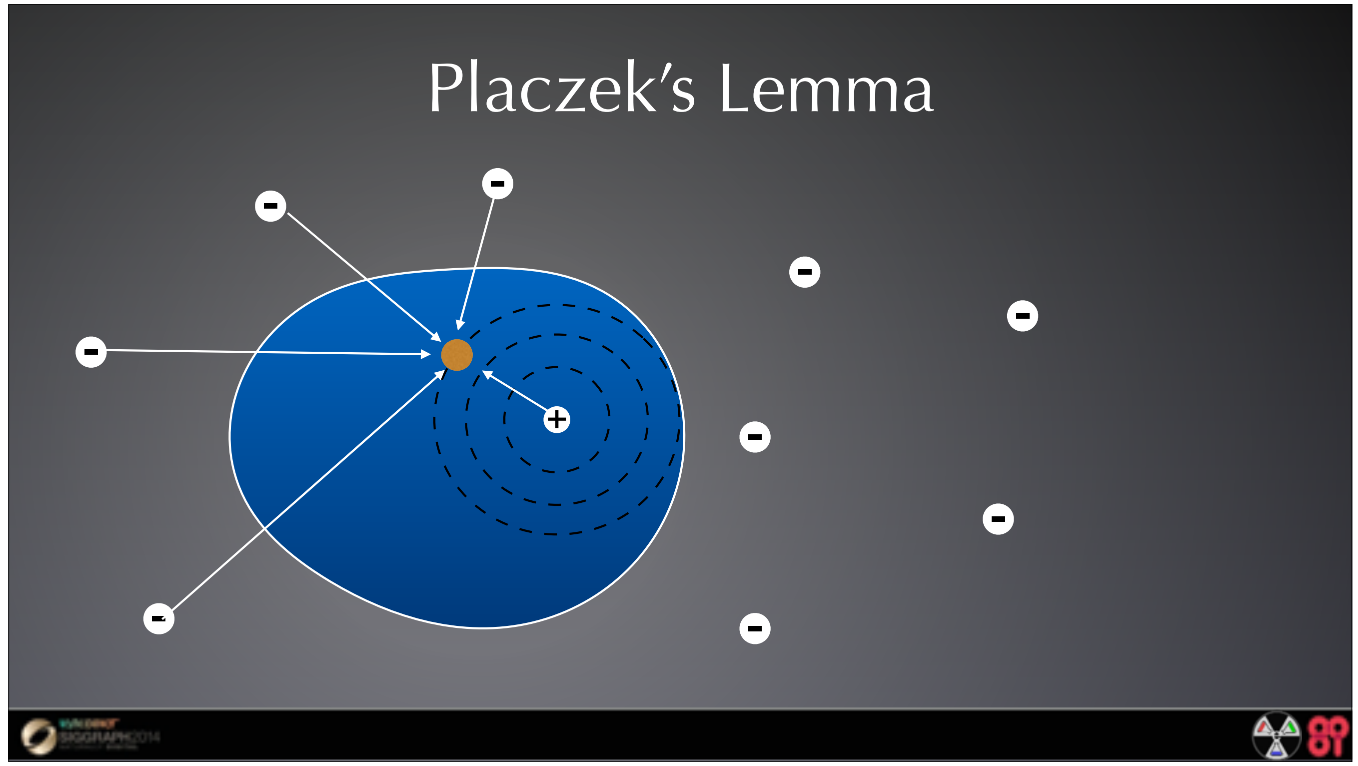

the integral of negative fluence from all of the exterior sources.

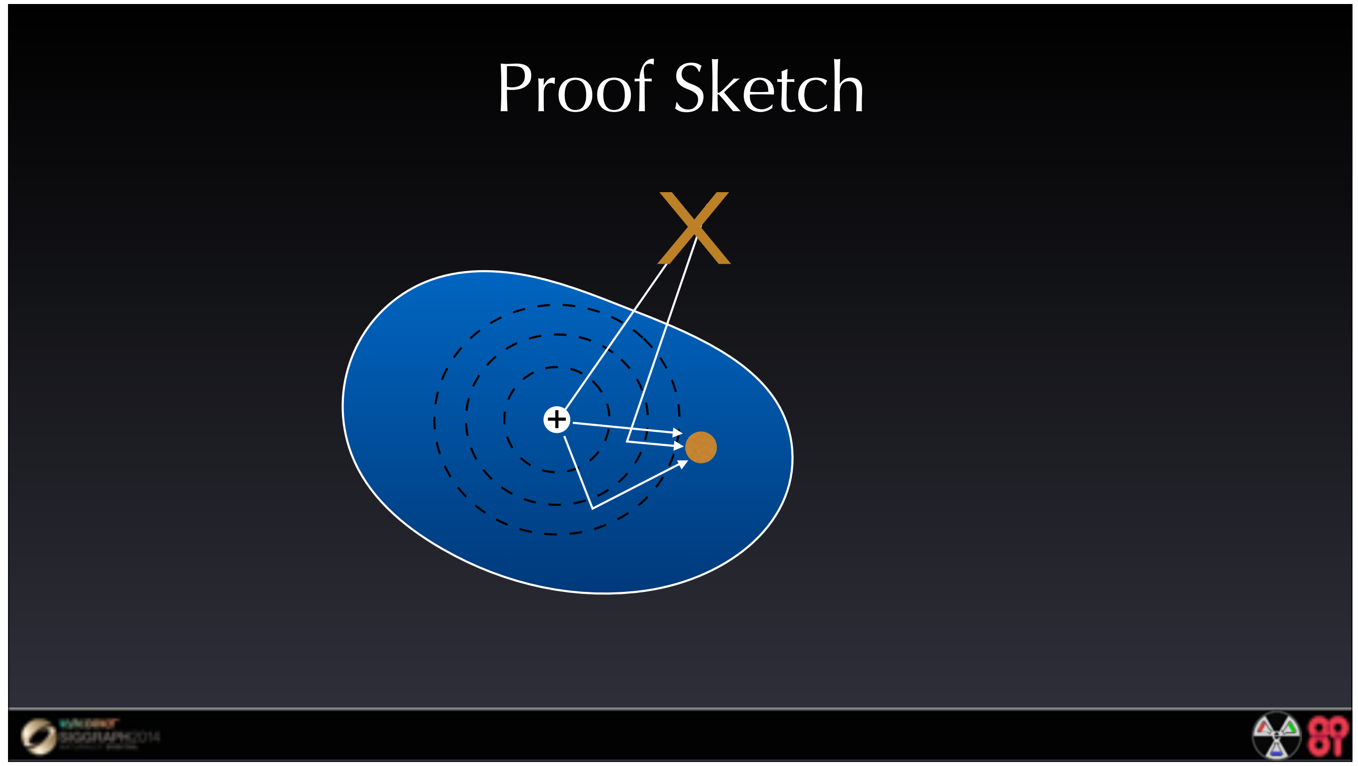

To convince yourself why something like this can work, consider that various paths described by the infinite solution used at the positive source will include some that exit and re-enter the original medium. As such, the infinite medium result will over predict the fluence at this orange location.

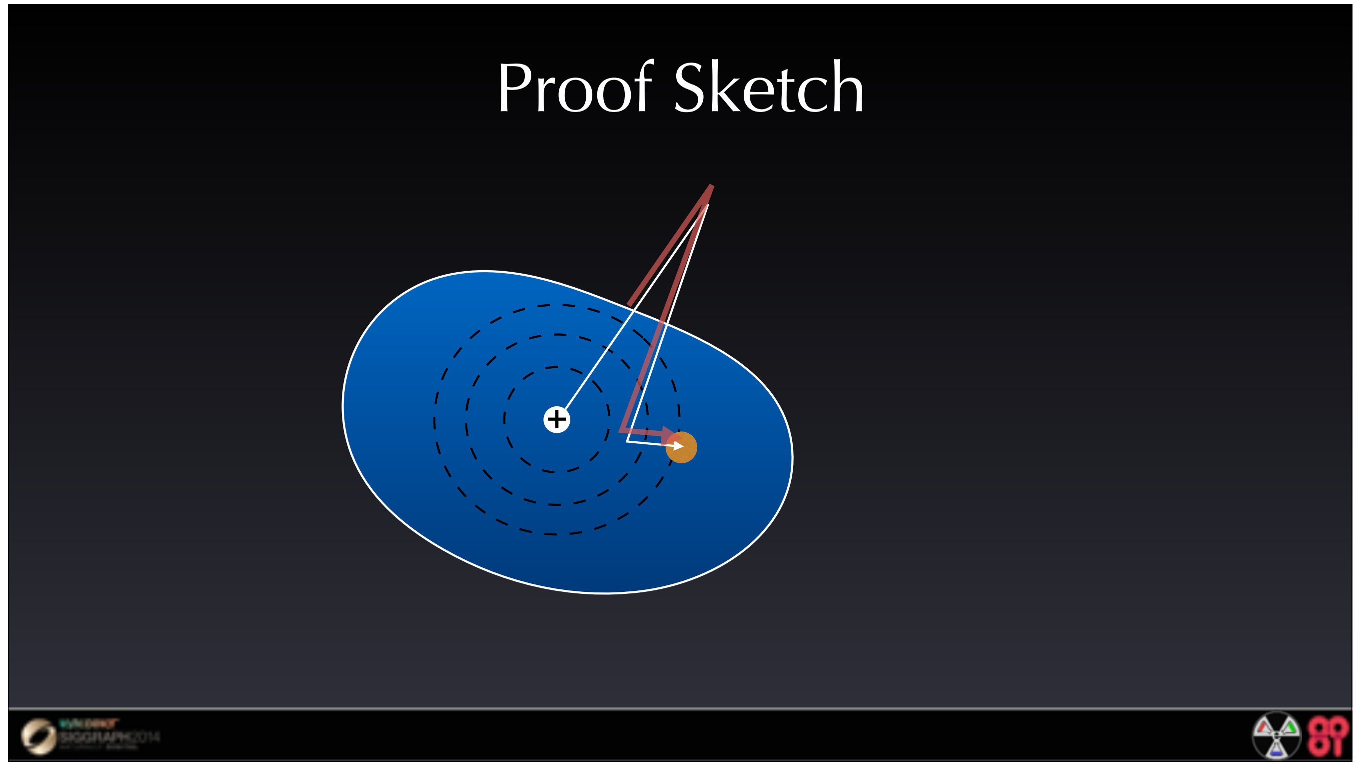

By adding a negative surface source just as this path leaves the volume the first time, this precisely cancels this extra energy, and the combination of positive source and negative external sources produces the finite medium desired answer.

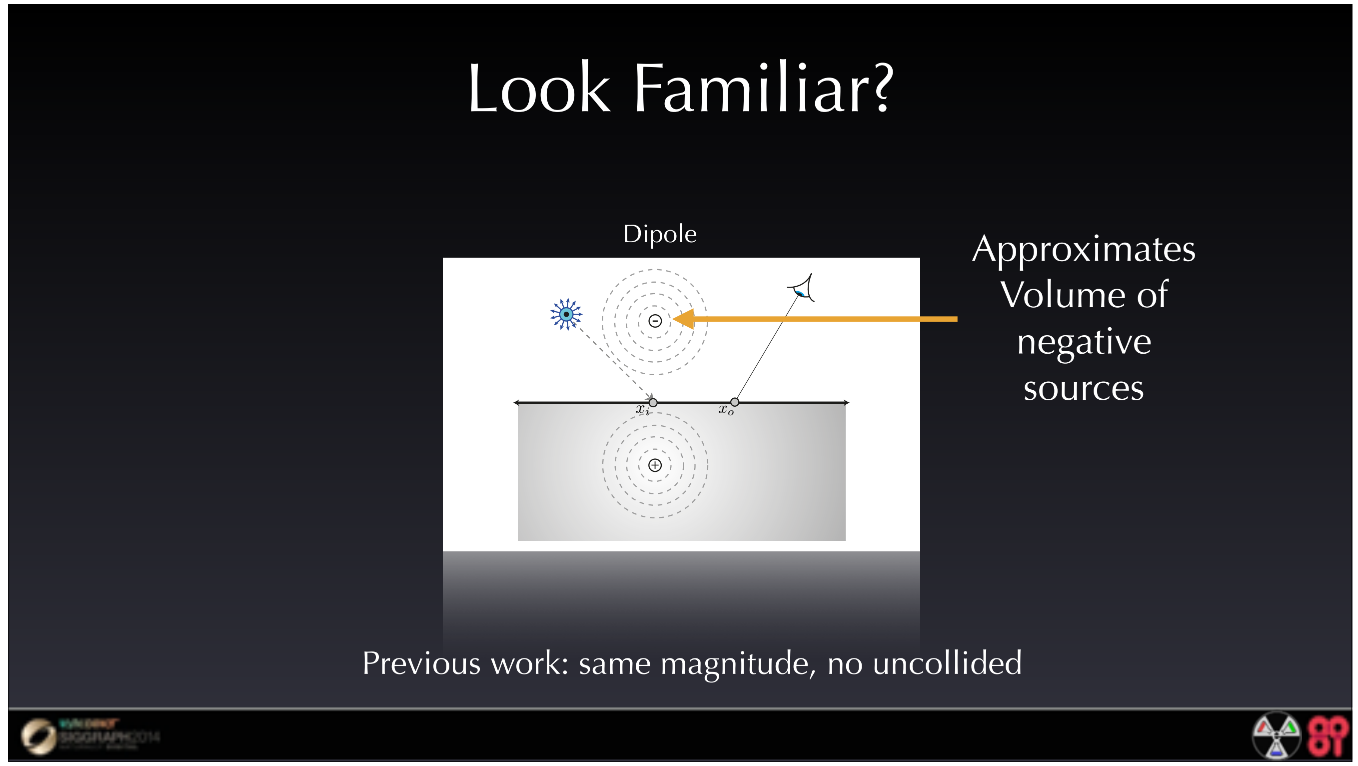

The original dipole is effectively an approximation of this idea in three ways: a) the continuous sea of negative energy is replaced by a single negative point source outside, b) the magnitude of that negative source is the same as the positive one, and b) the Green’s function for the negative sources is missing the uncollided term.

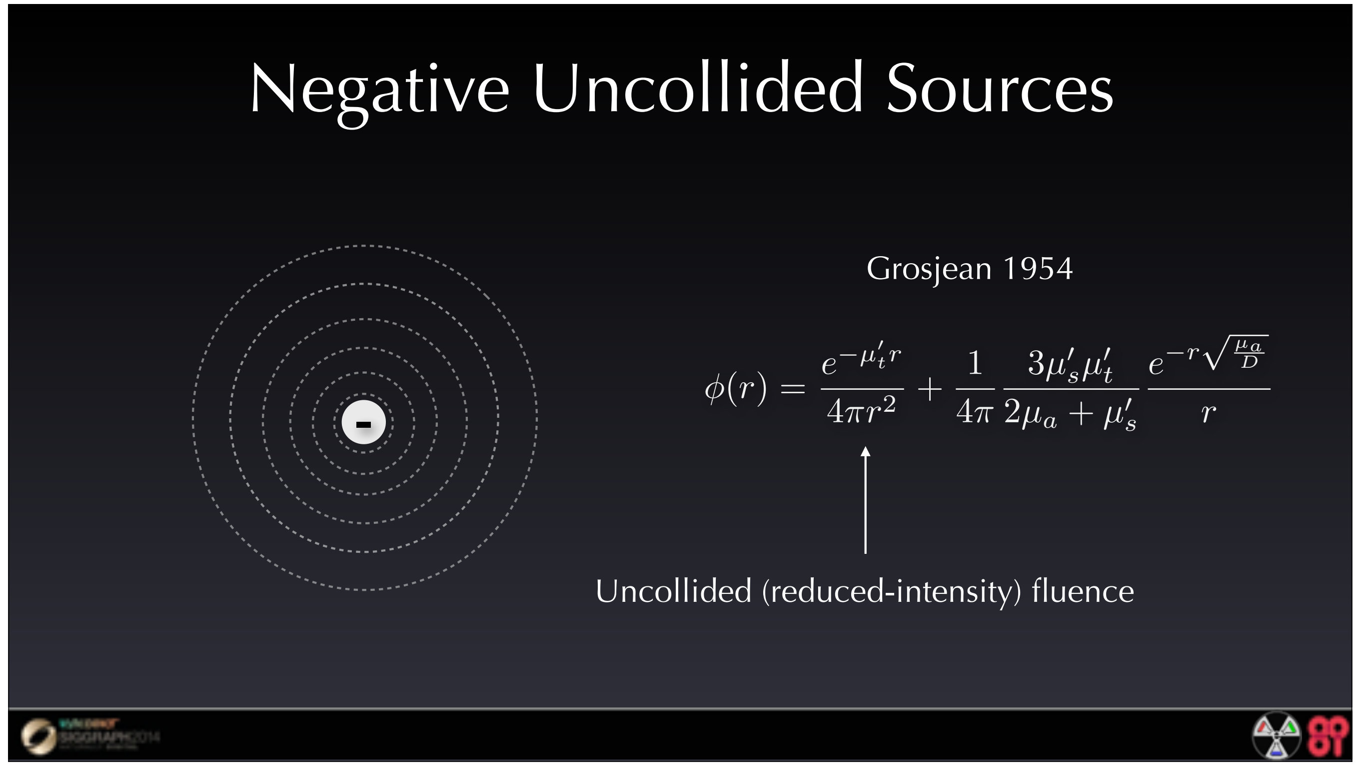

Placzek's lemma produces an exact solution by employing exact Green's functions. A highly accurate approximate Green's function contains an important uncollided term, which we can consider adding to the method of images for improved accuracy.

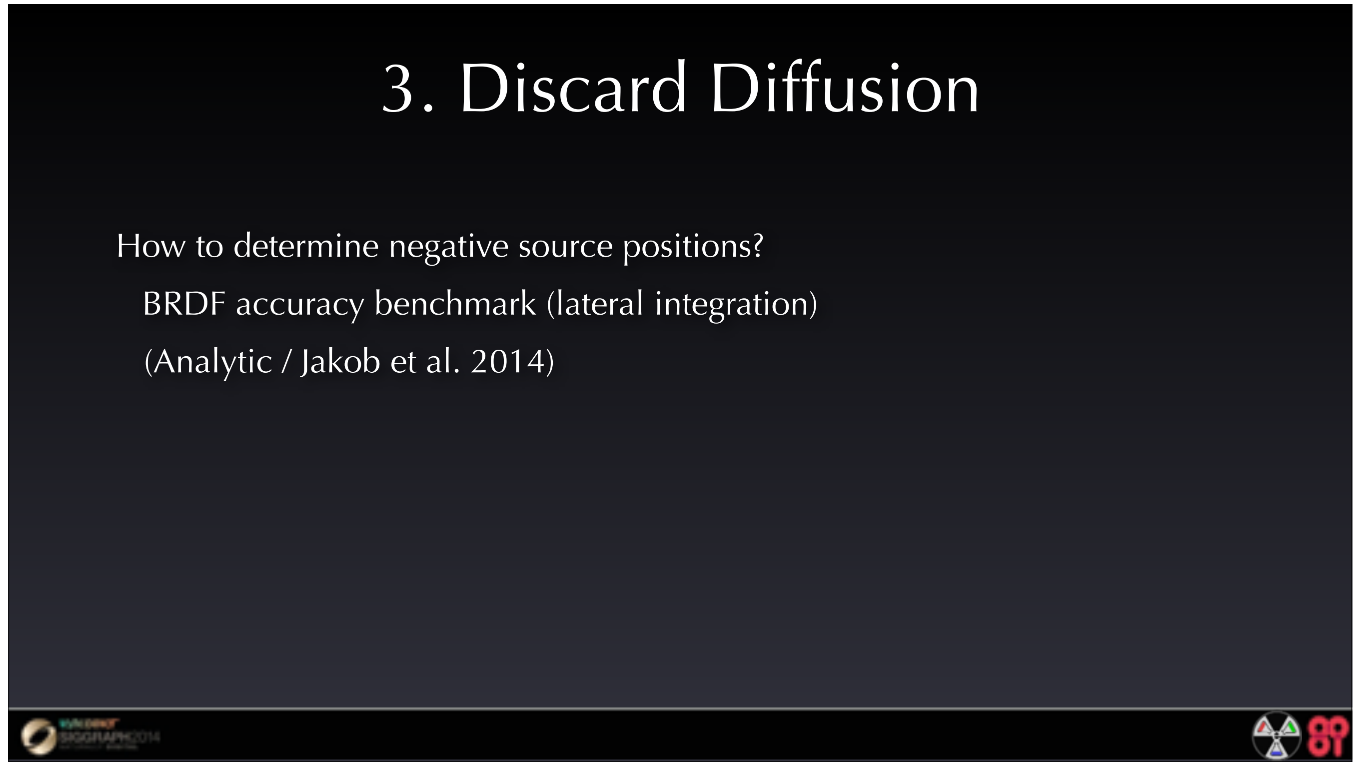

However, how do we maintain discrete negative approximate sources outside and place them in the most accurate way (without knowing the exact solution, as required by Placzek). To do this, we use the associated BRDF for the problem (computable using methods like [Jakob et al. 2014]).

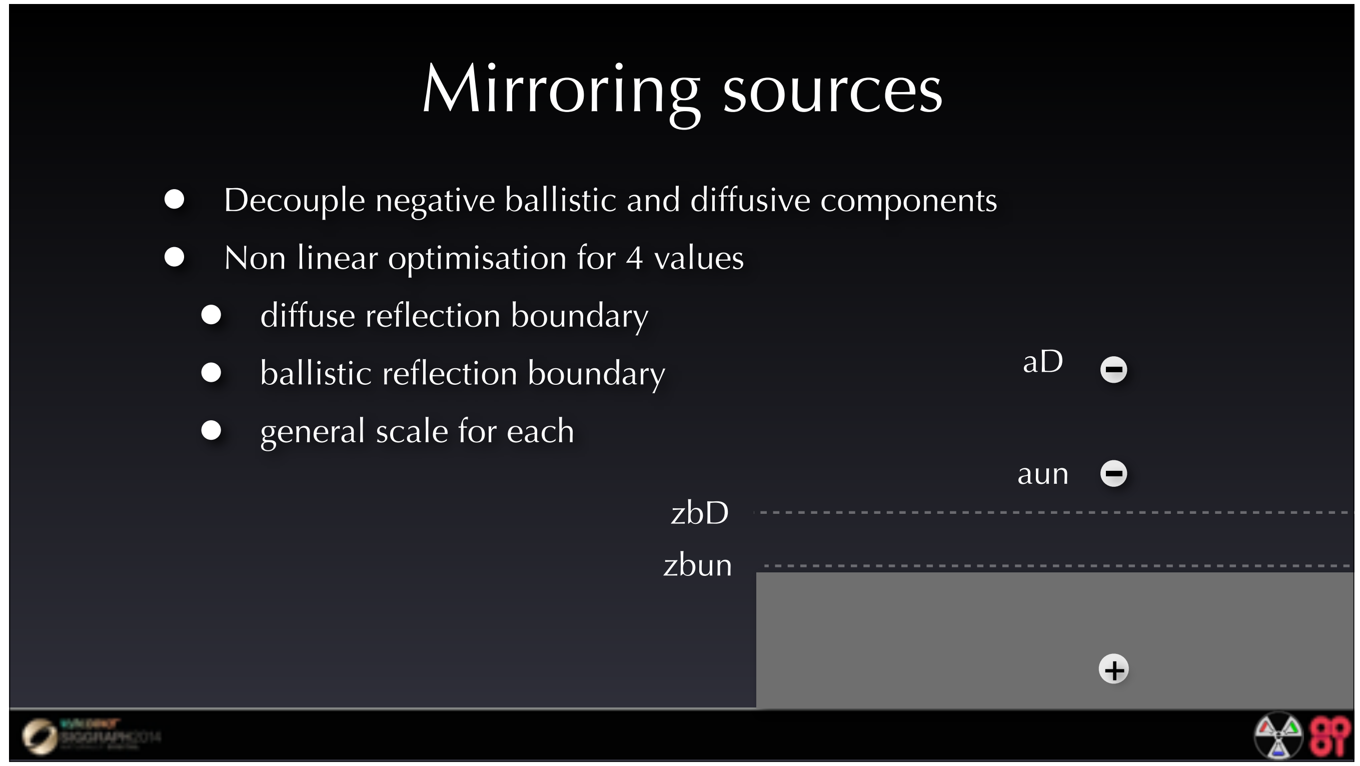

For each positive source in a semi-infinite medium we propose placing a negative diffusive and negative uncollided source outside the medium, where the reflection boundary, and intensities of the negative sources are all determined by a non linear optimisation for 4 parameters such that the BRDF derived from this BSSRDF is as accurate as possible (more details in the supplemental material).

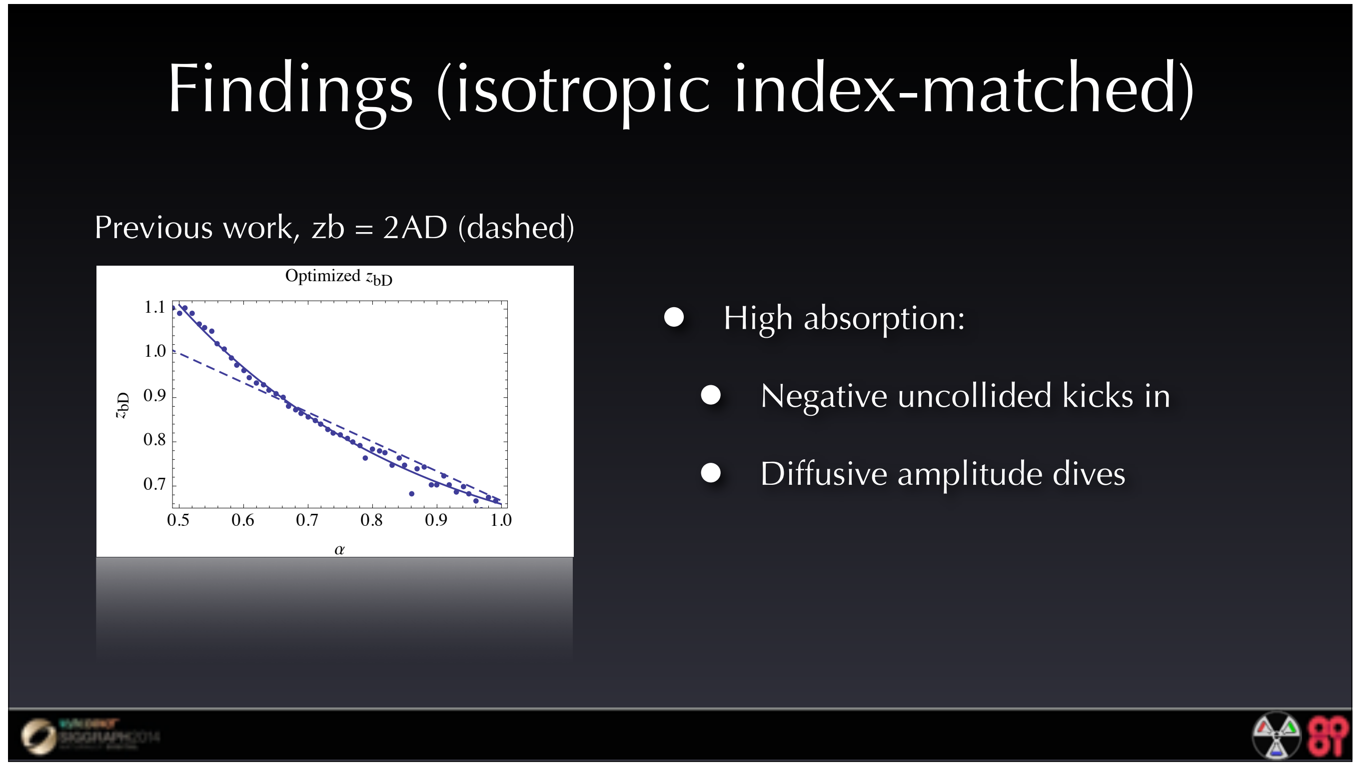

What we found when running the optimisation over a broad range of isotropic scattering media with single-scattering albedo alpha is: the optimal depth zb about which to mirror the diffusive portion of the negative sources is not quite that predicted by diffusion theory. This and other findings in more depth in the supp. material.

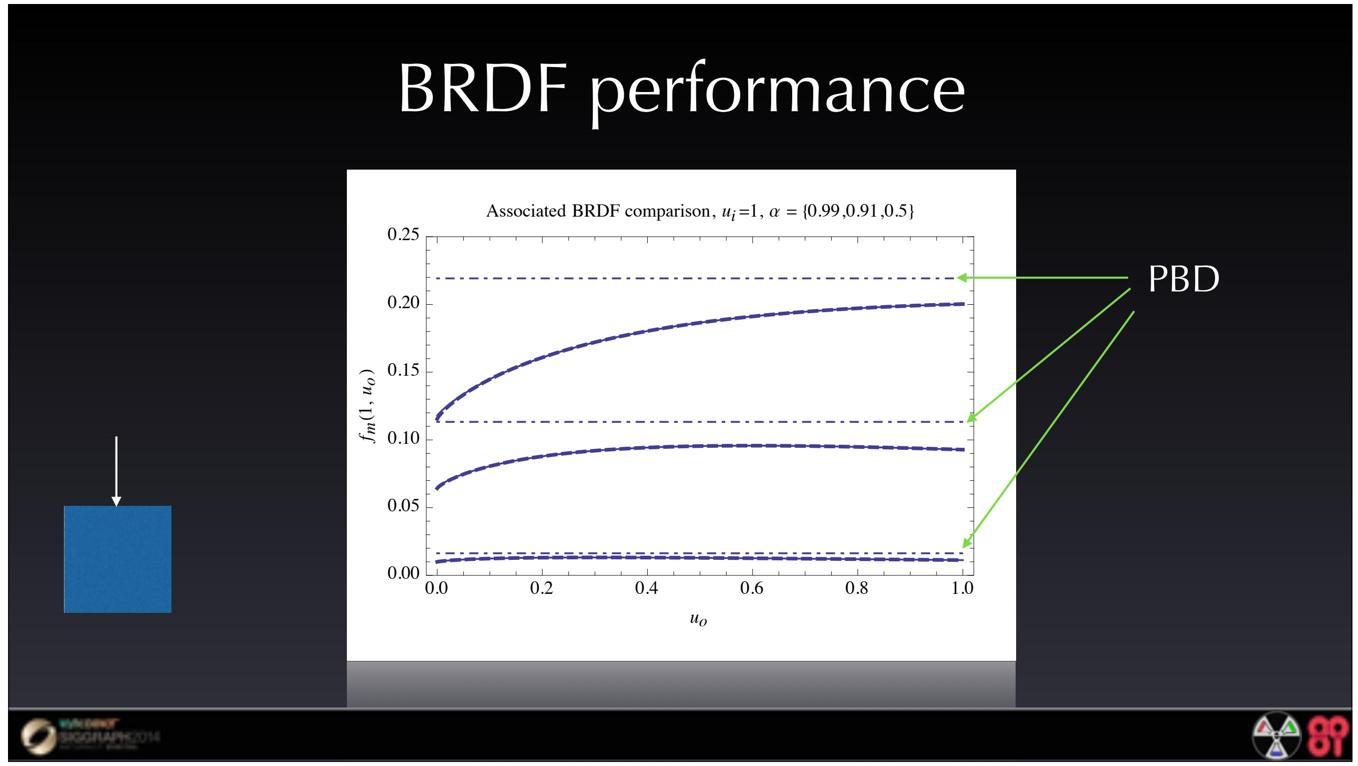

The BRDF of the new BSSRDF is much more accurate than previous Lambertian/diffusive approaches, like Photon Beam Diffusion.

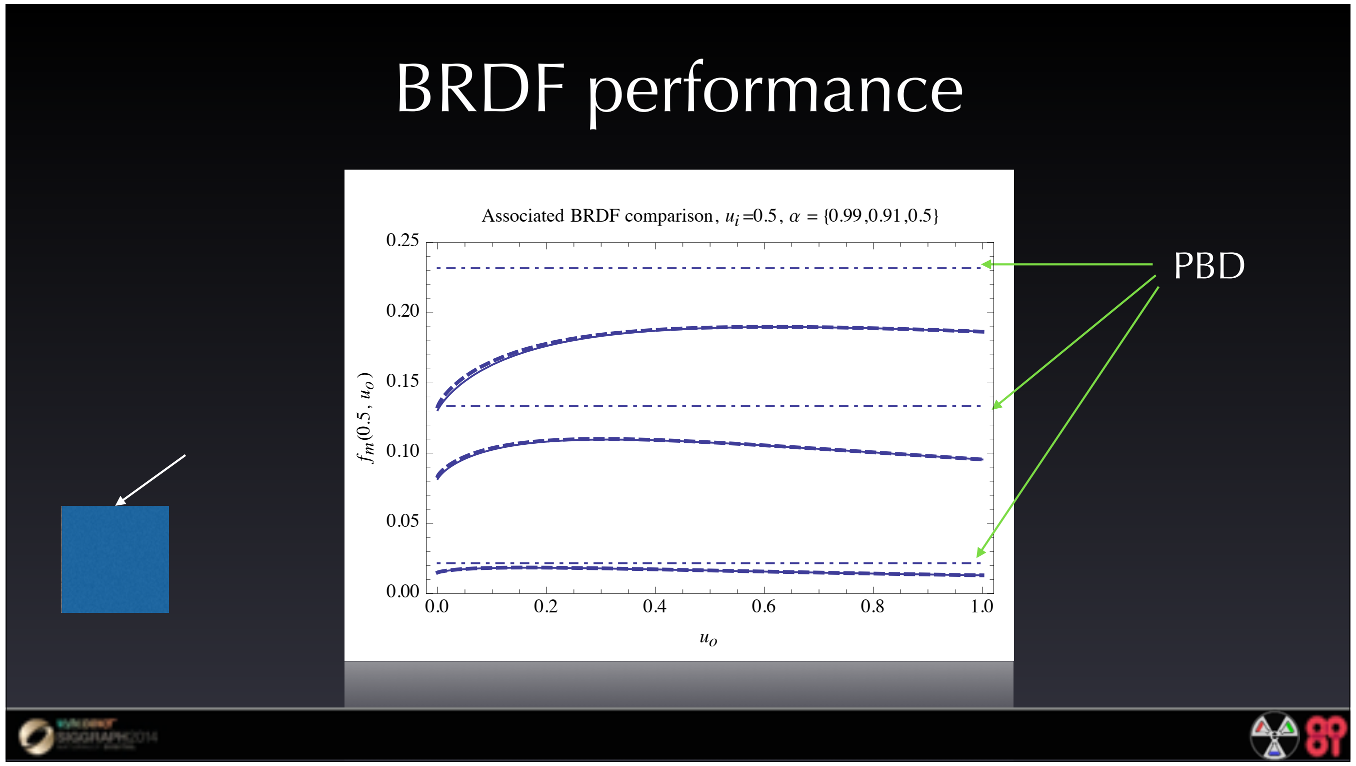

For a variety of incident angles…

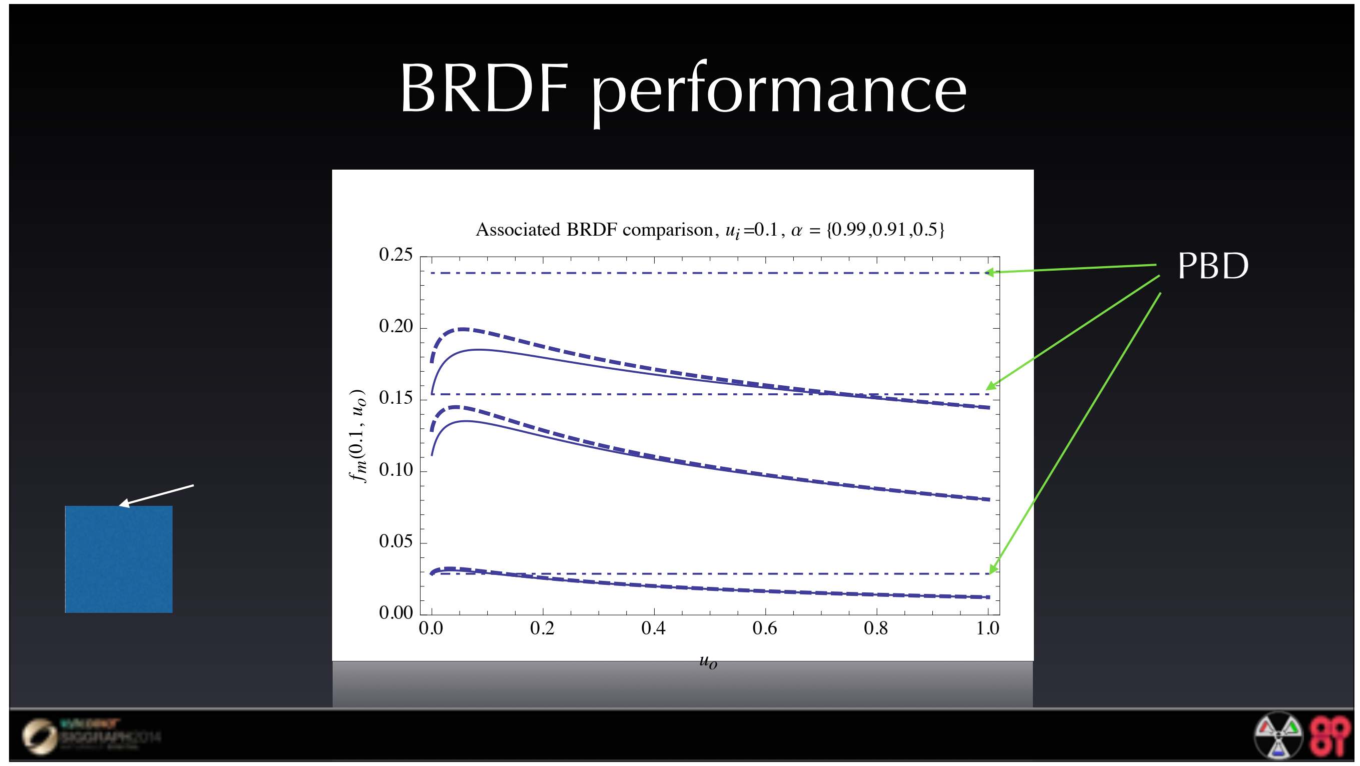
BRDF performance
Associated BRDF comparison, u_i=0.1, α = {0.99,0.91,0.5}
0.25
0.20
0.15
0.10
0.05
0.00
0.0
0.2
0.4
0.6
0.8
1.0
$f_m(0.1, u_o)$
u_o
PBD

…

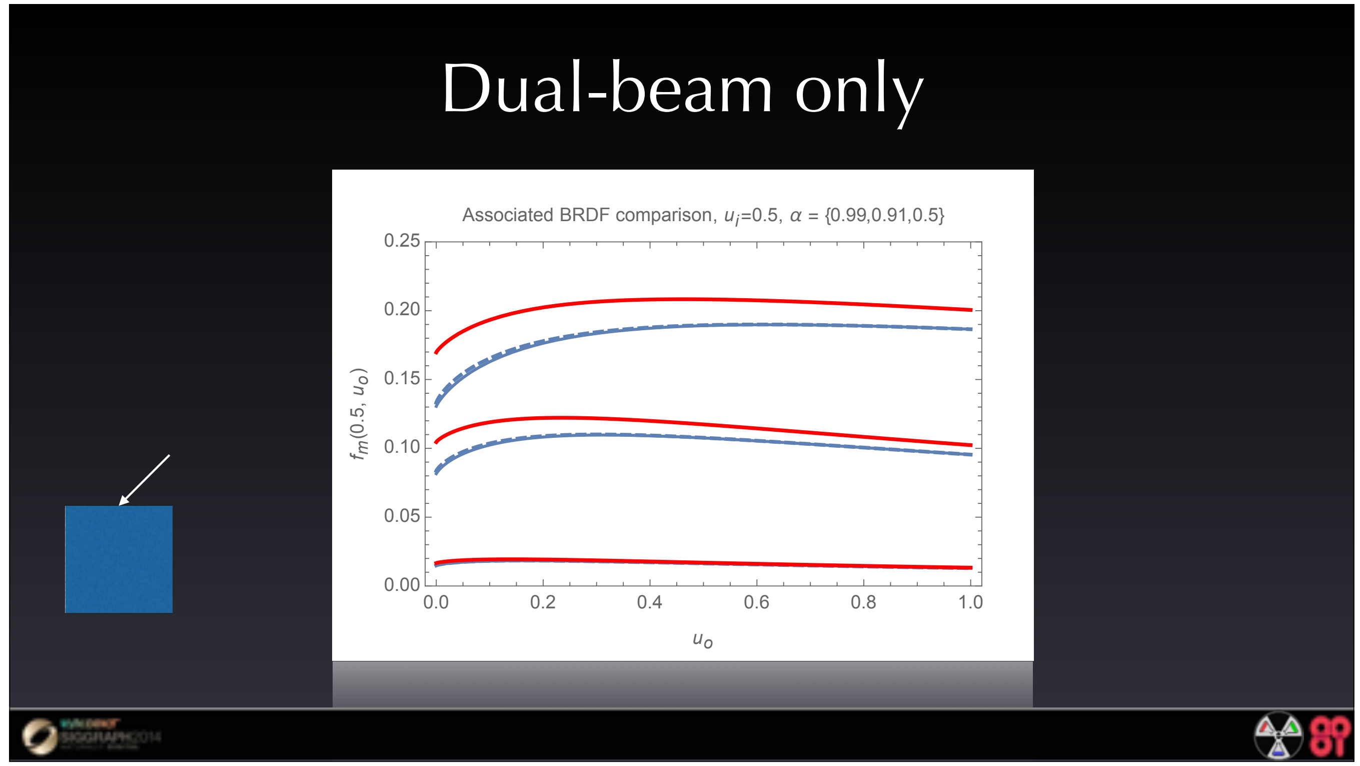

Here we show the improvement gained from only using the Davison-line-integral exitance calculation (but not the new source placements and uncollided term). We found both were required to get reasonable results over a variety of incident angles.

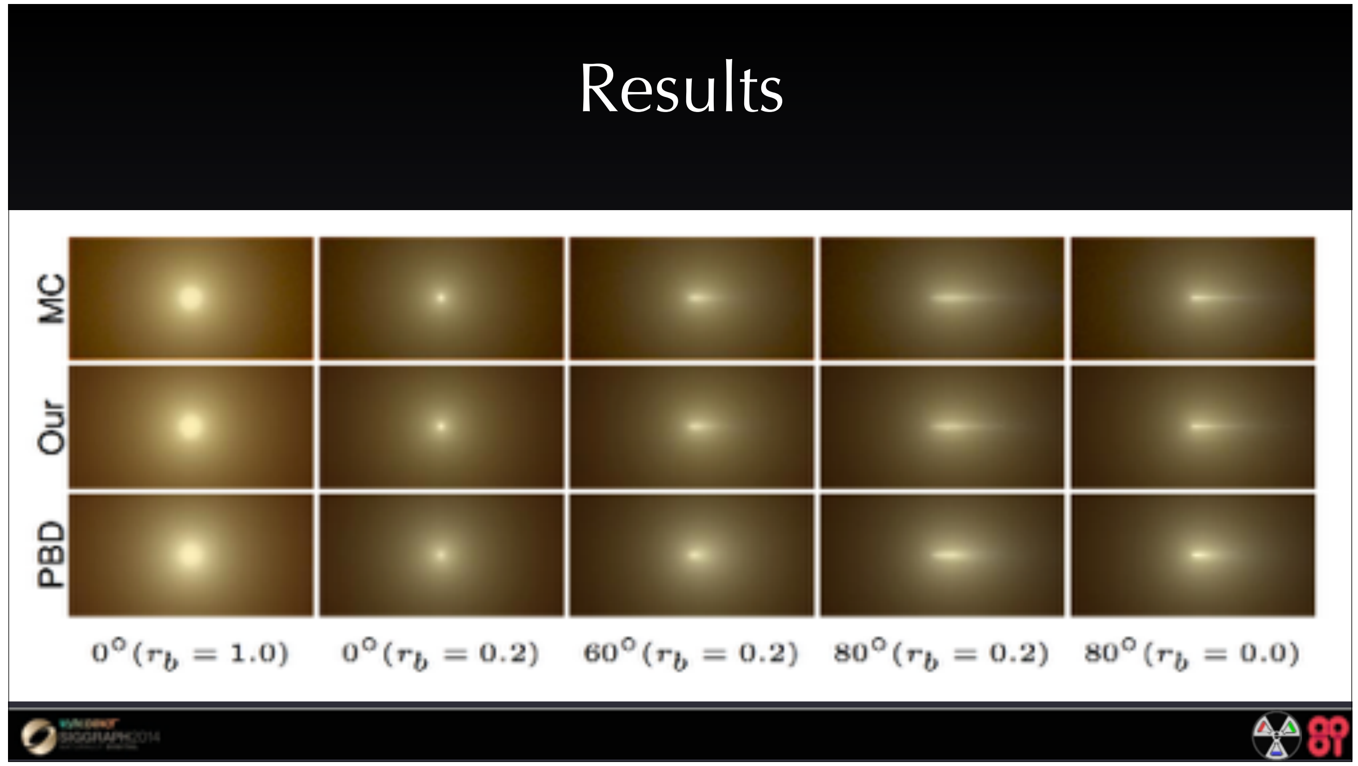

Our new BSSRDF also produces more accurate BSSRDF behaviour than Photon Beam Diffusion.

Conclusion

- Dual-beam BSSRDF formulation
  - Reciprocal 8D
  - BRDF benchmark guided
  - Discarded diffusion

Future Work

- TODO
  - Moment driven?
  - Fresnel
  - Anisotropic scattering
  - Placzek inspired solutions for MC or curved geometry?
- Other related searchlight solutions
  - [Liemert and Kienle 2012,2013, Gardner 2013, Machida 2013]

It may be possible to derive the 4 BSSRDF parameters from moments of the BRDF solution, avoiding the non-linear optimization. We have yet to include Fresnel effects and it is unclear how to extend method-of-images solutions to anisotropic scattering. Further investigations along these lines should include performance and accuracy relationships to recent searchlight solutions such as those listed here.

THANK YOU