ReSTIR and Path Re-using

UCSD CSE 272
Advanced Image Synthesis
Tzu-Mao Li

with slides from Benedikt Bitterli

Motivation: can we do importance sampling of lights without complex data structures?

difficult to maintain & slow for real-time rendering

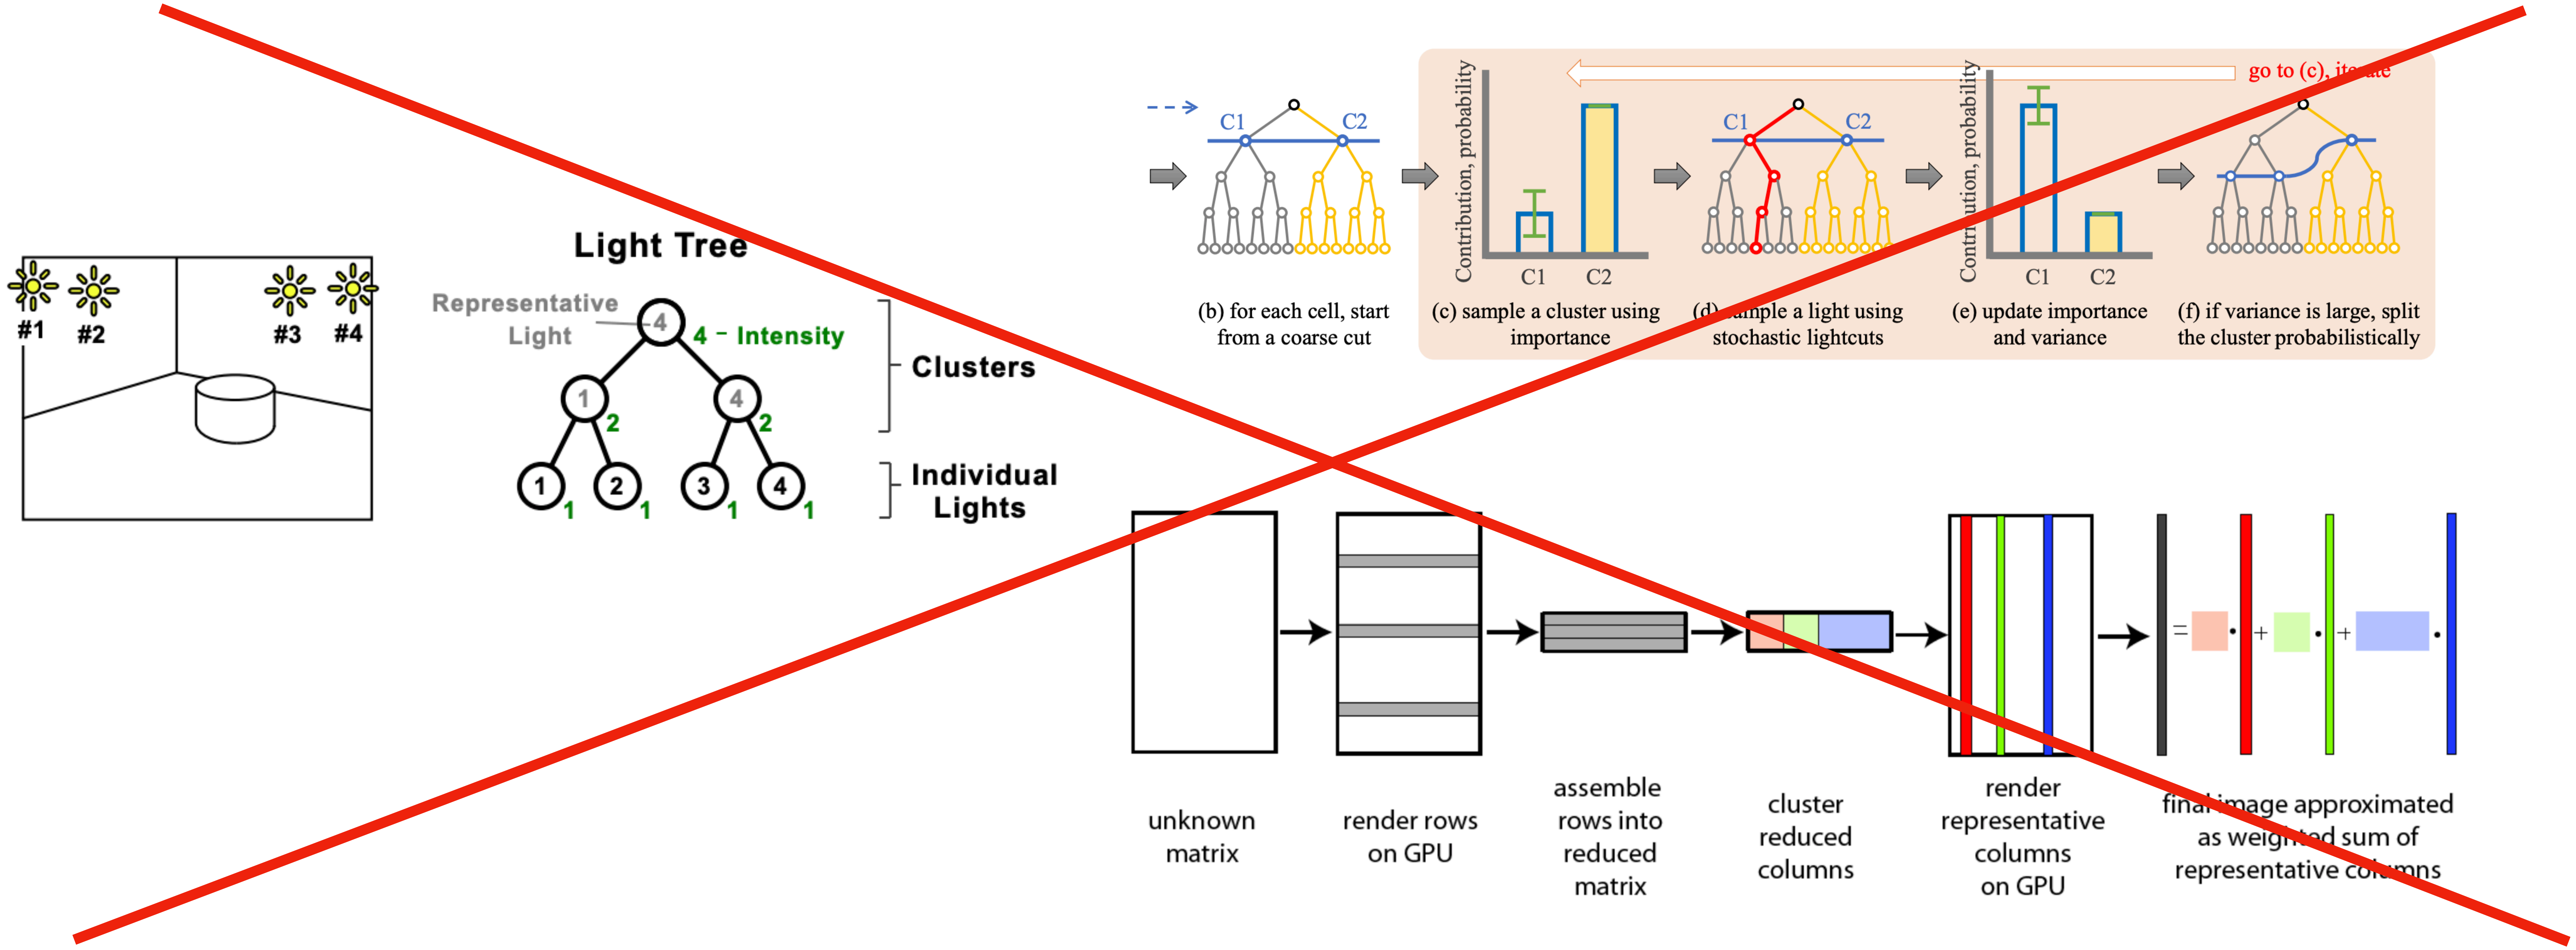

Eye candy

22.9 million triangles, 3.4 million emissive, dynamic triangles, rendered at interactive rates (20-40ms?)

ReSTIR: a general sampling algorithm

Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting

BENEDIKT BITTERLI, Dartmouth College
CHRIS WYMAN, NVIDIA
MATT PHARR, NVIDIA
PETER SHIRLEY, NVIDIA
AARON LEFOHN, NVIDIA
WOJCIECH JAROSZ, Dartmouth College

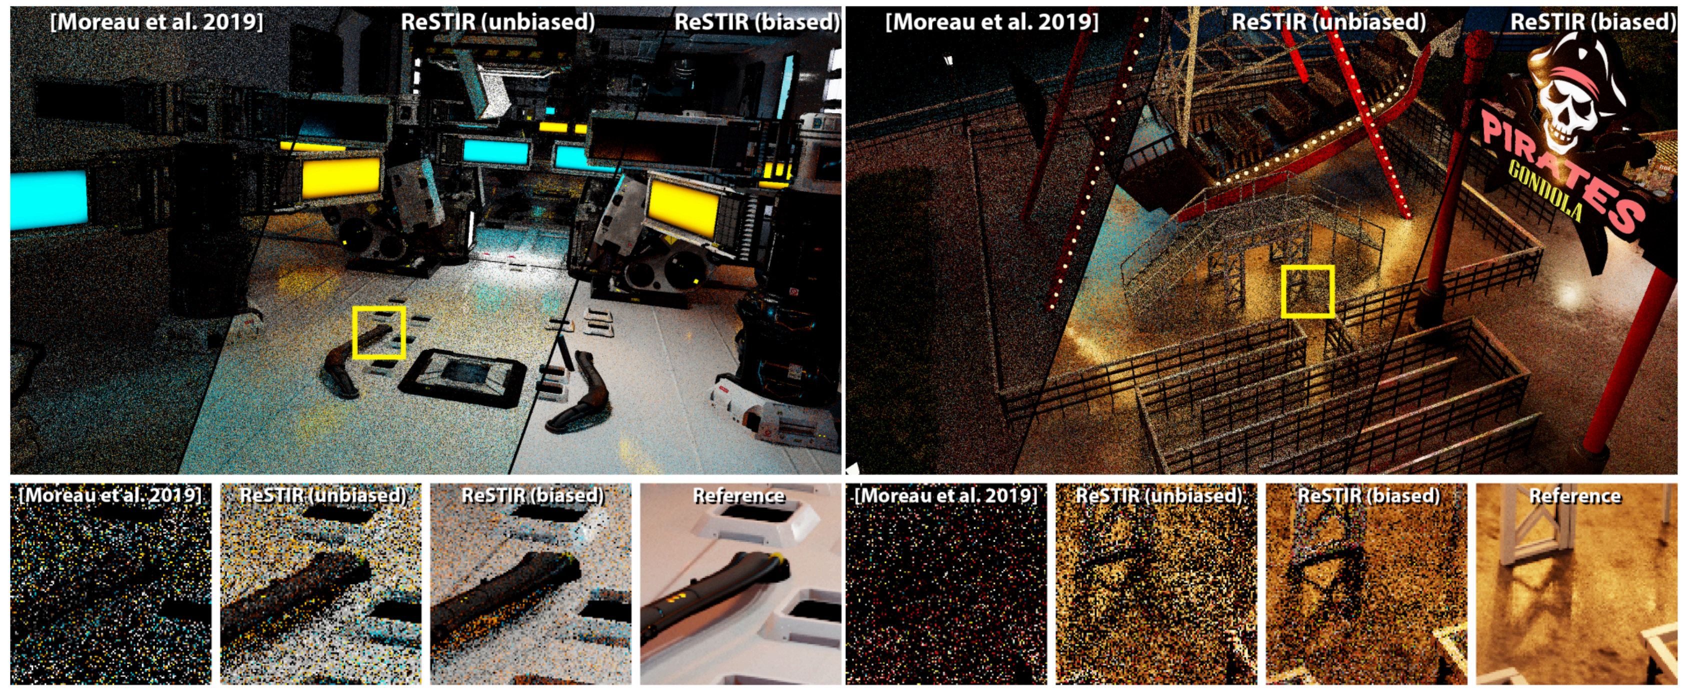

ReSTIR GI: Path Resampling for Real-Time Path Tracing

Y. Ouyang[1], S. Liu[1], M. Kettunen[1], M. Pharr[1], J. Pantaleoni[1]

[1]NVIDIA Corporation, Santa Clara, CA, USA

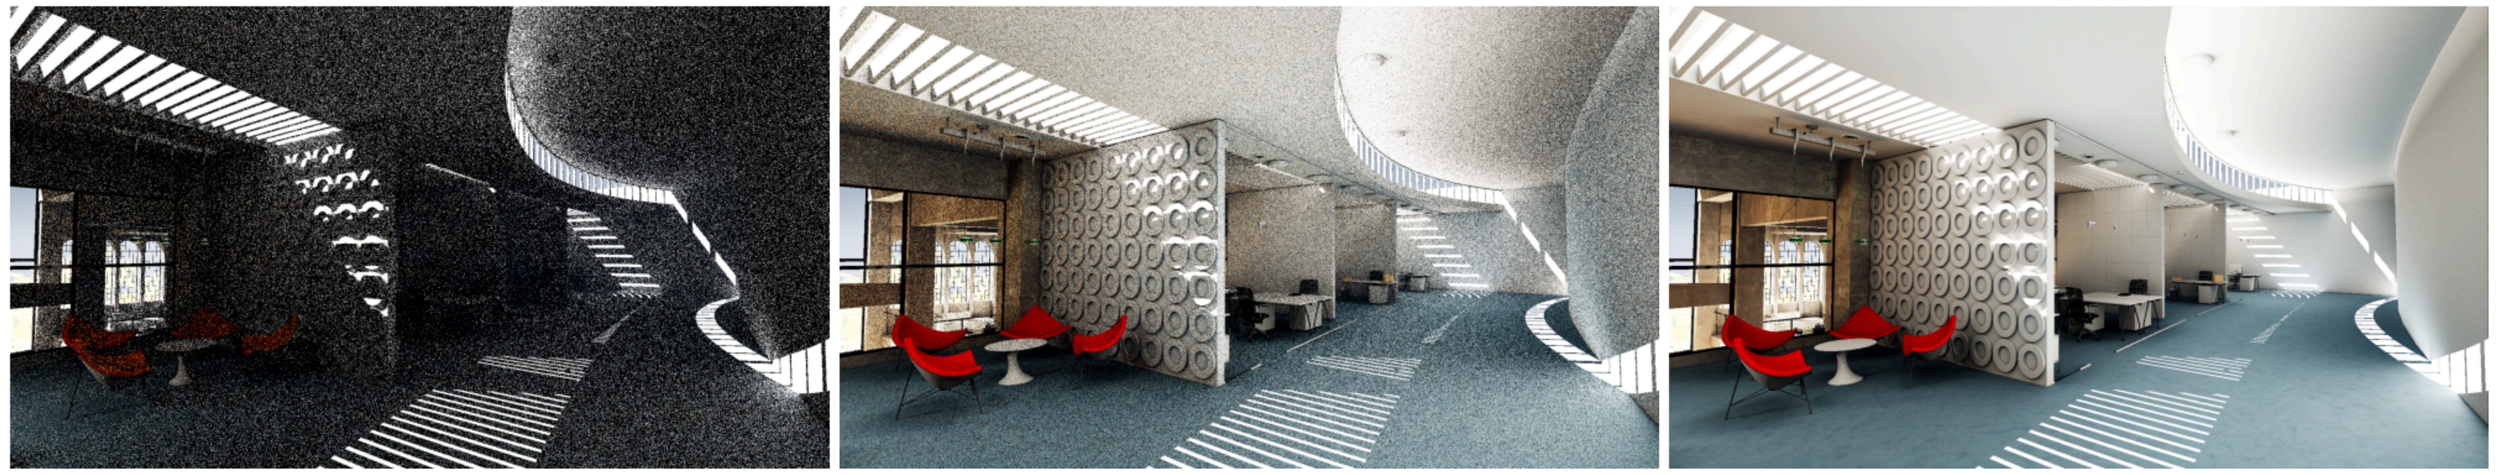

Fast Volume Rendering with Spatiotemporal Reservoir Resampling

DAQI LIN, University of Utah, USA
CHRIS WYMAN, NVIDIA, USA
CEM YUKSEL, University of Utah, USA

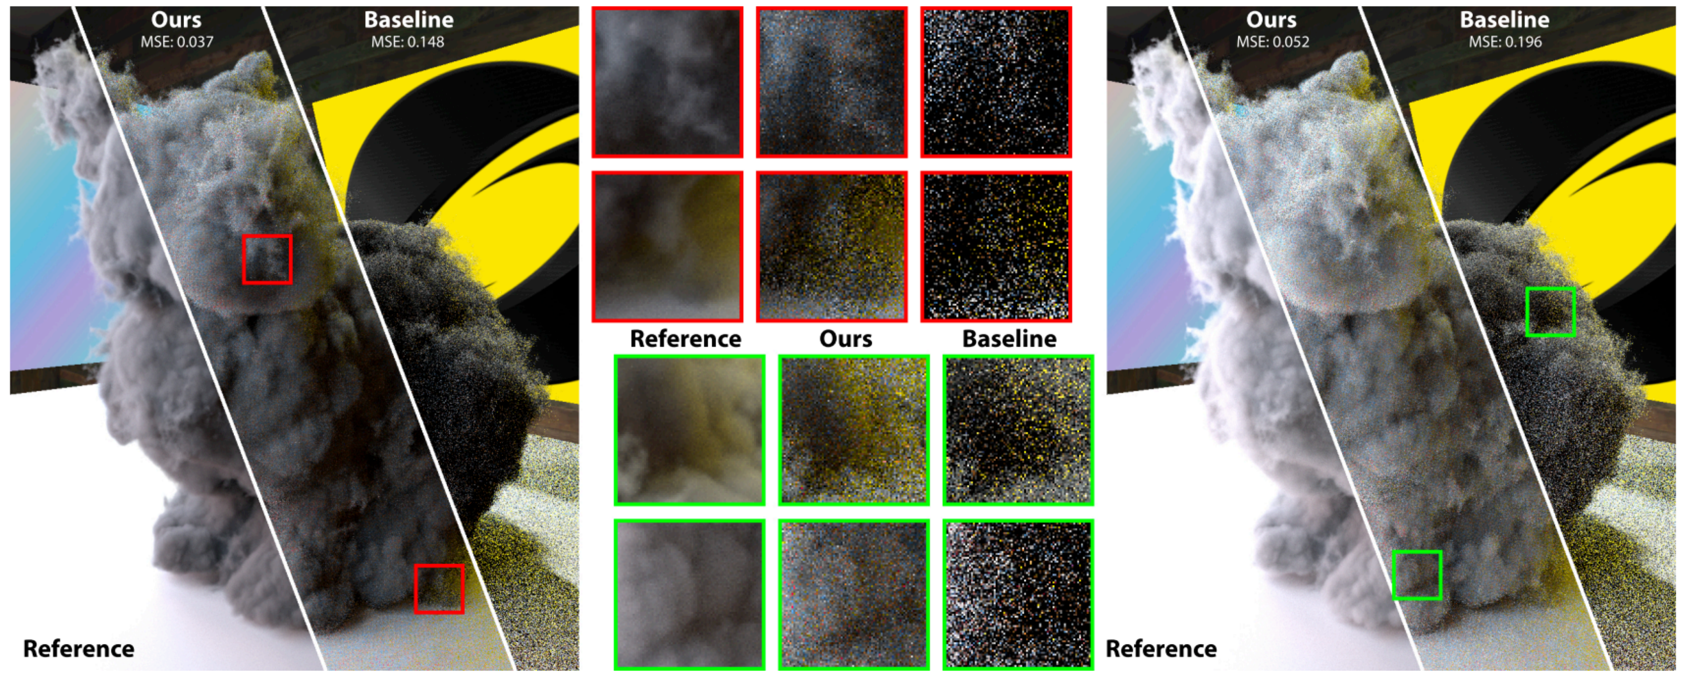

Generalized Resampled Importance Sampling: Foundations of ReSTIR

DAQI LIN*, University of Utah, USA
MARKUS KETTUNEN*, NVIDIA, Finland
BENEDIKT BITTERLI, NVIDIA, USA
JACOPO PANTALEONI, NVIDIA, Germany
CEM YUKSEL, University of Utah, USA
CHRIS WYMAN, NVIDIA, USA

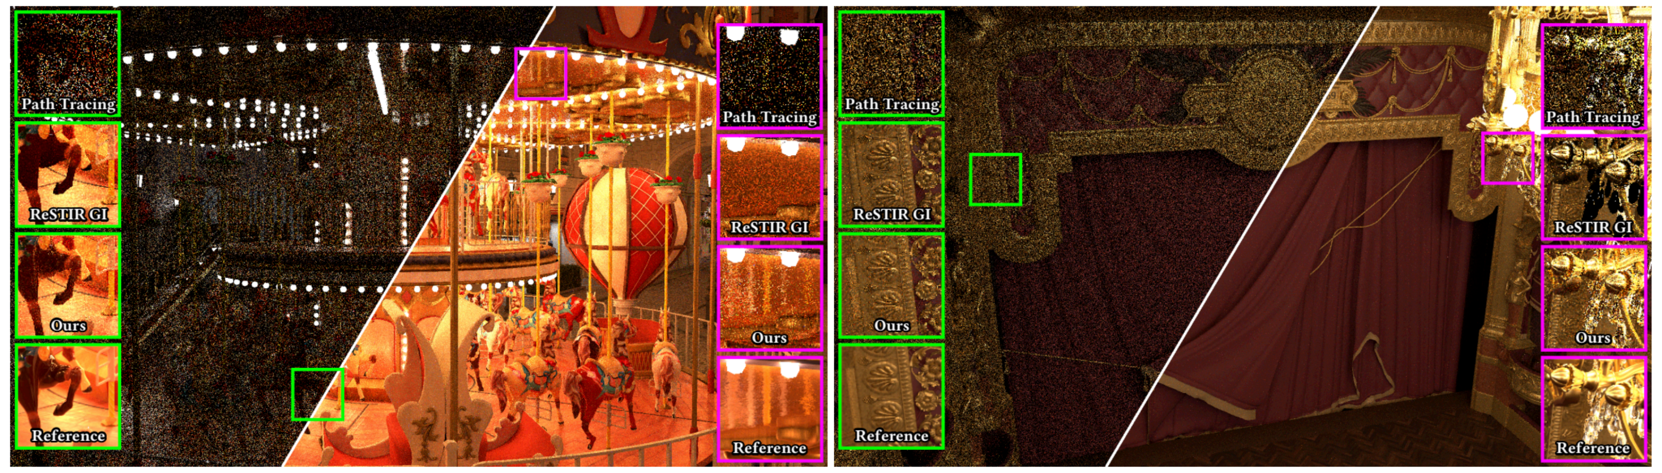

Let's start from the many-lights problem

intensity geometry term

$$\text{contribution} = L \cdot \rho \cdot G \cdot V$$

BSDF visibility

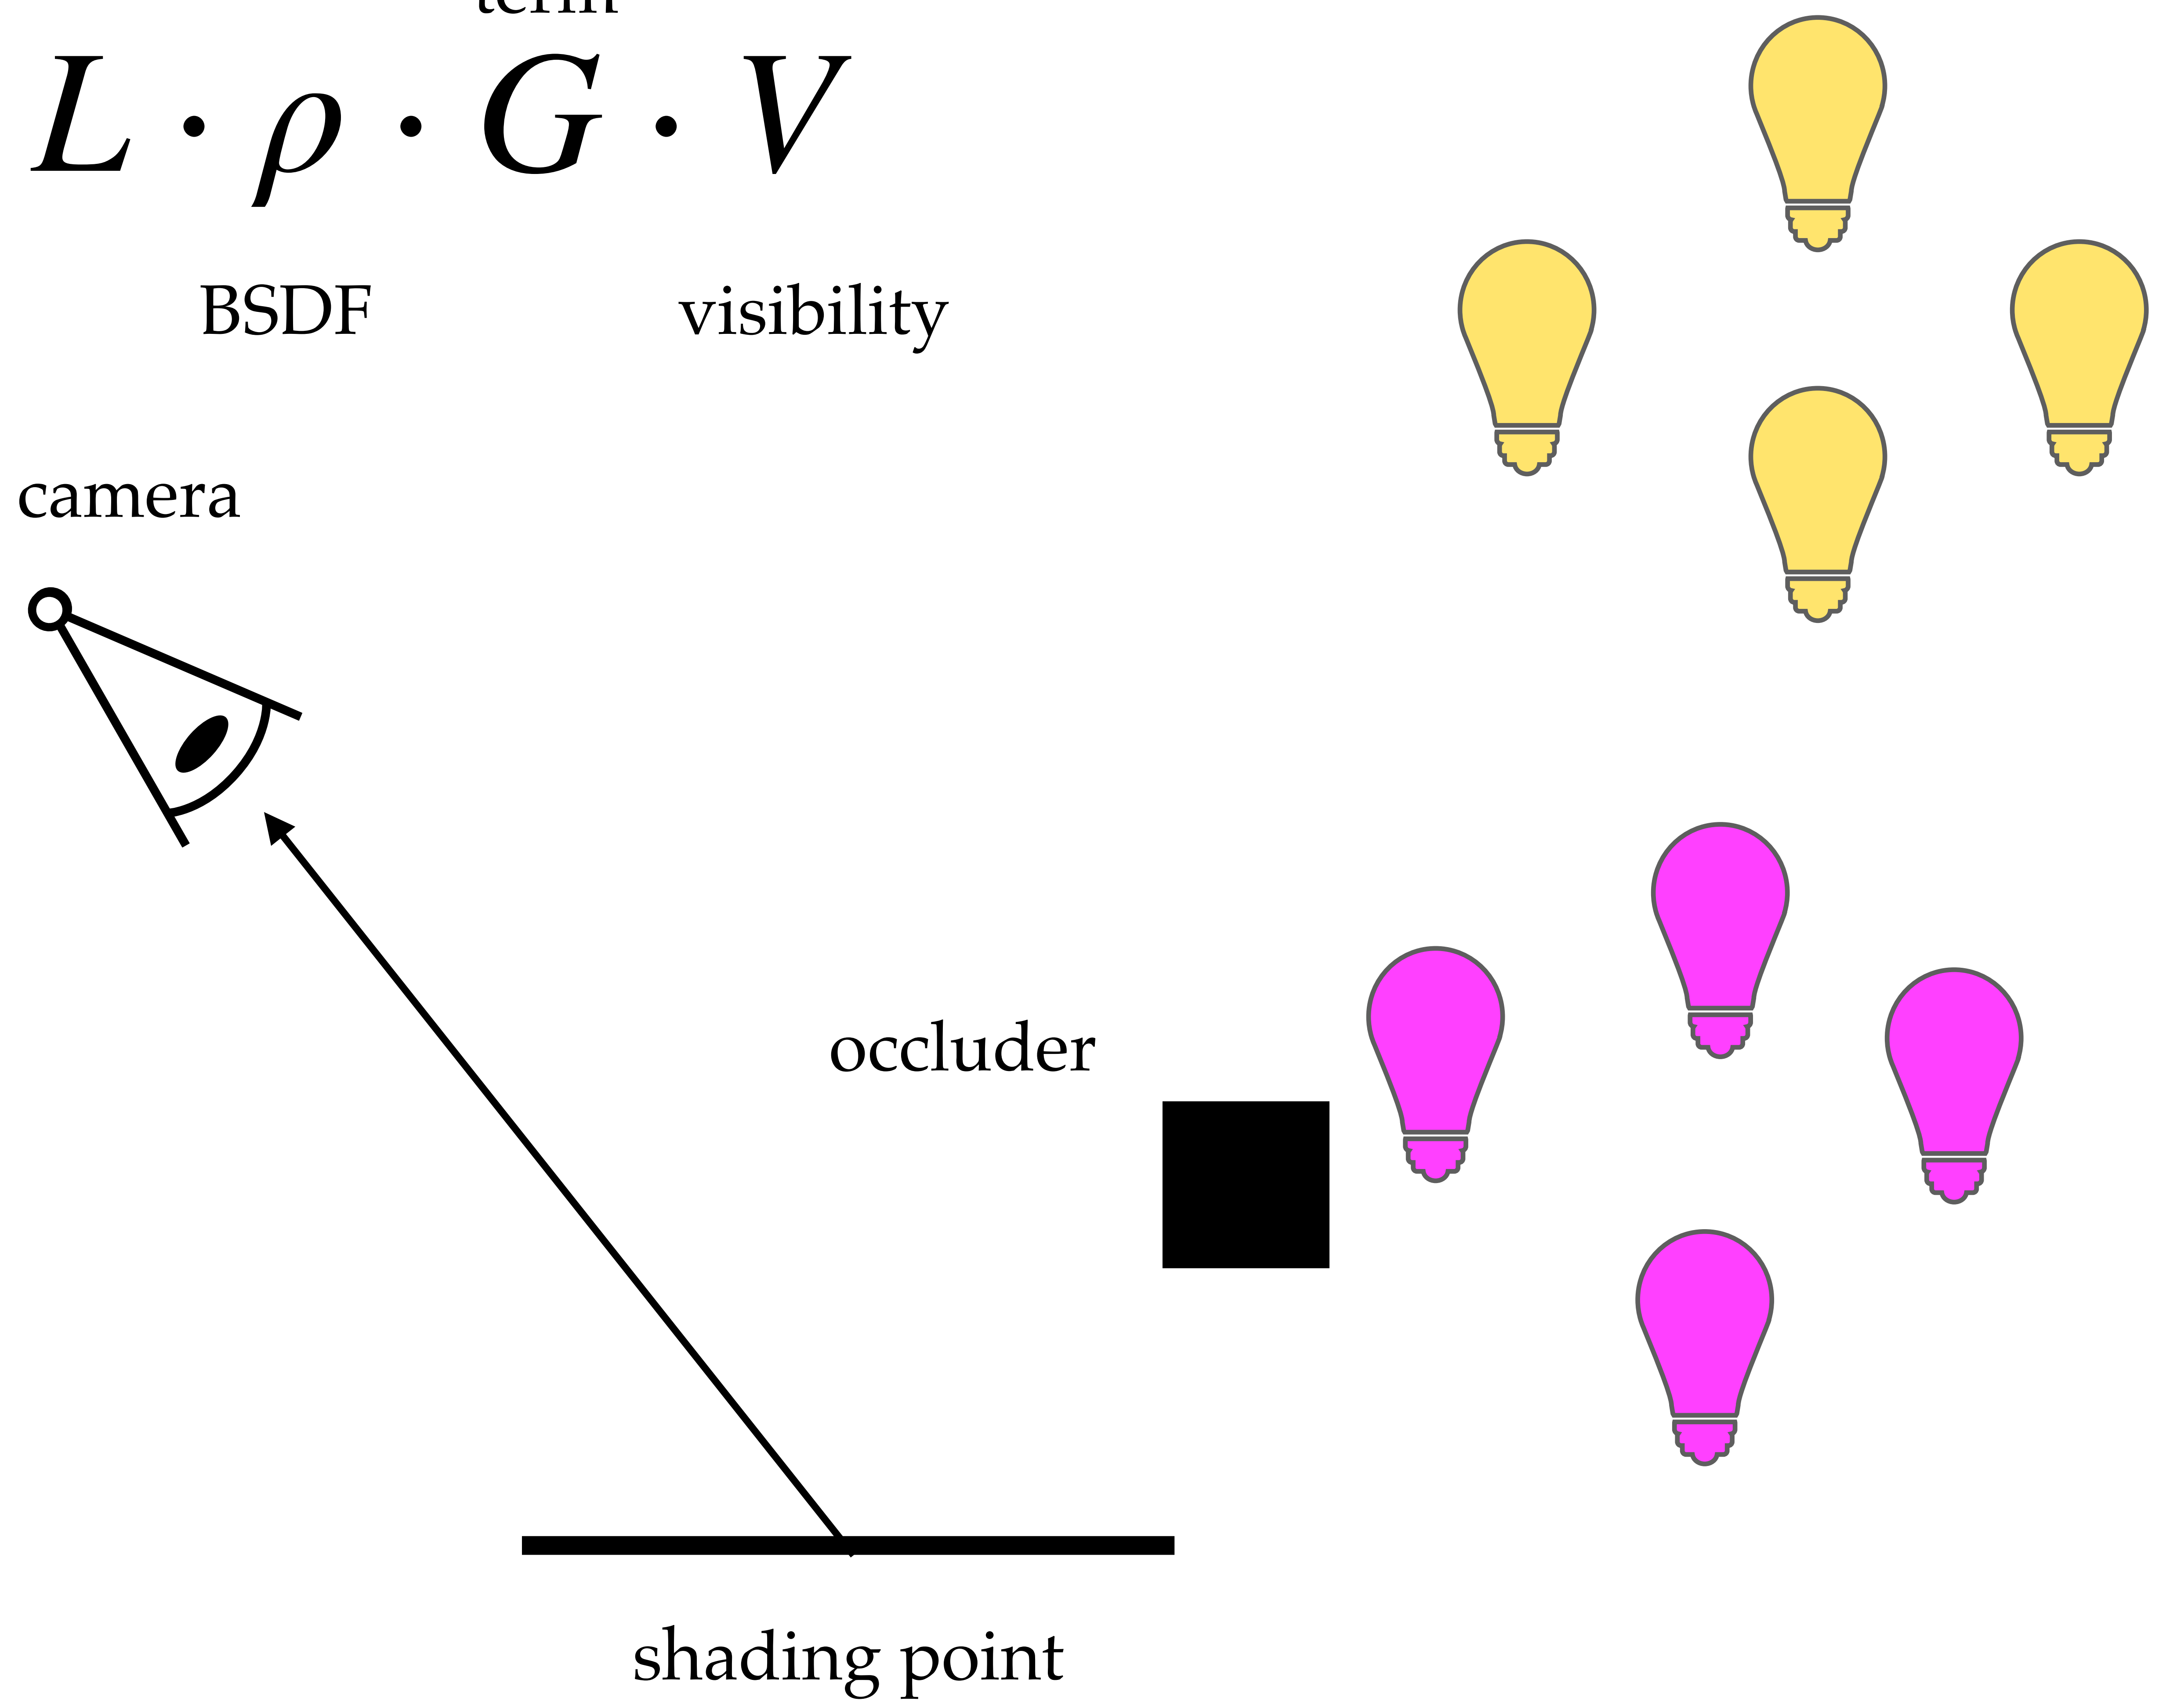

Idea: reuse neighboring pixels' sampling results

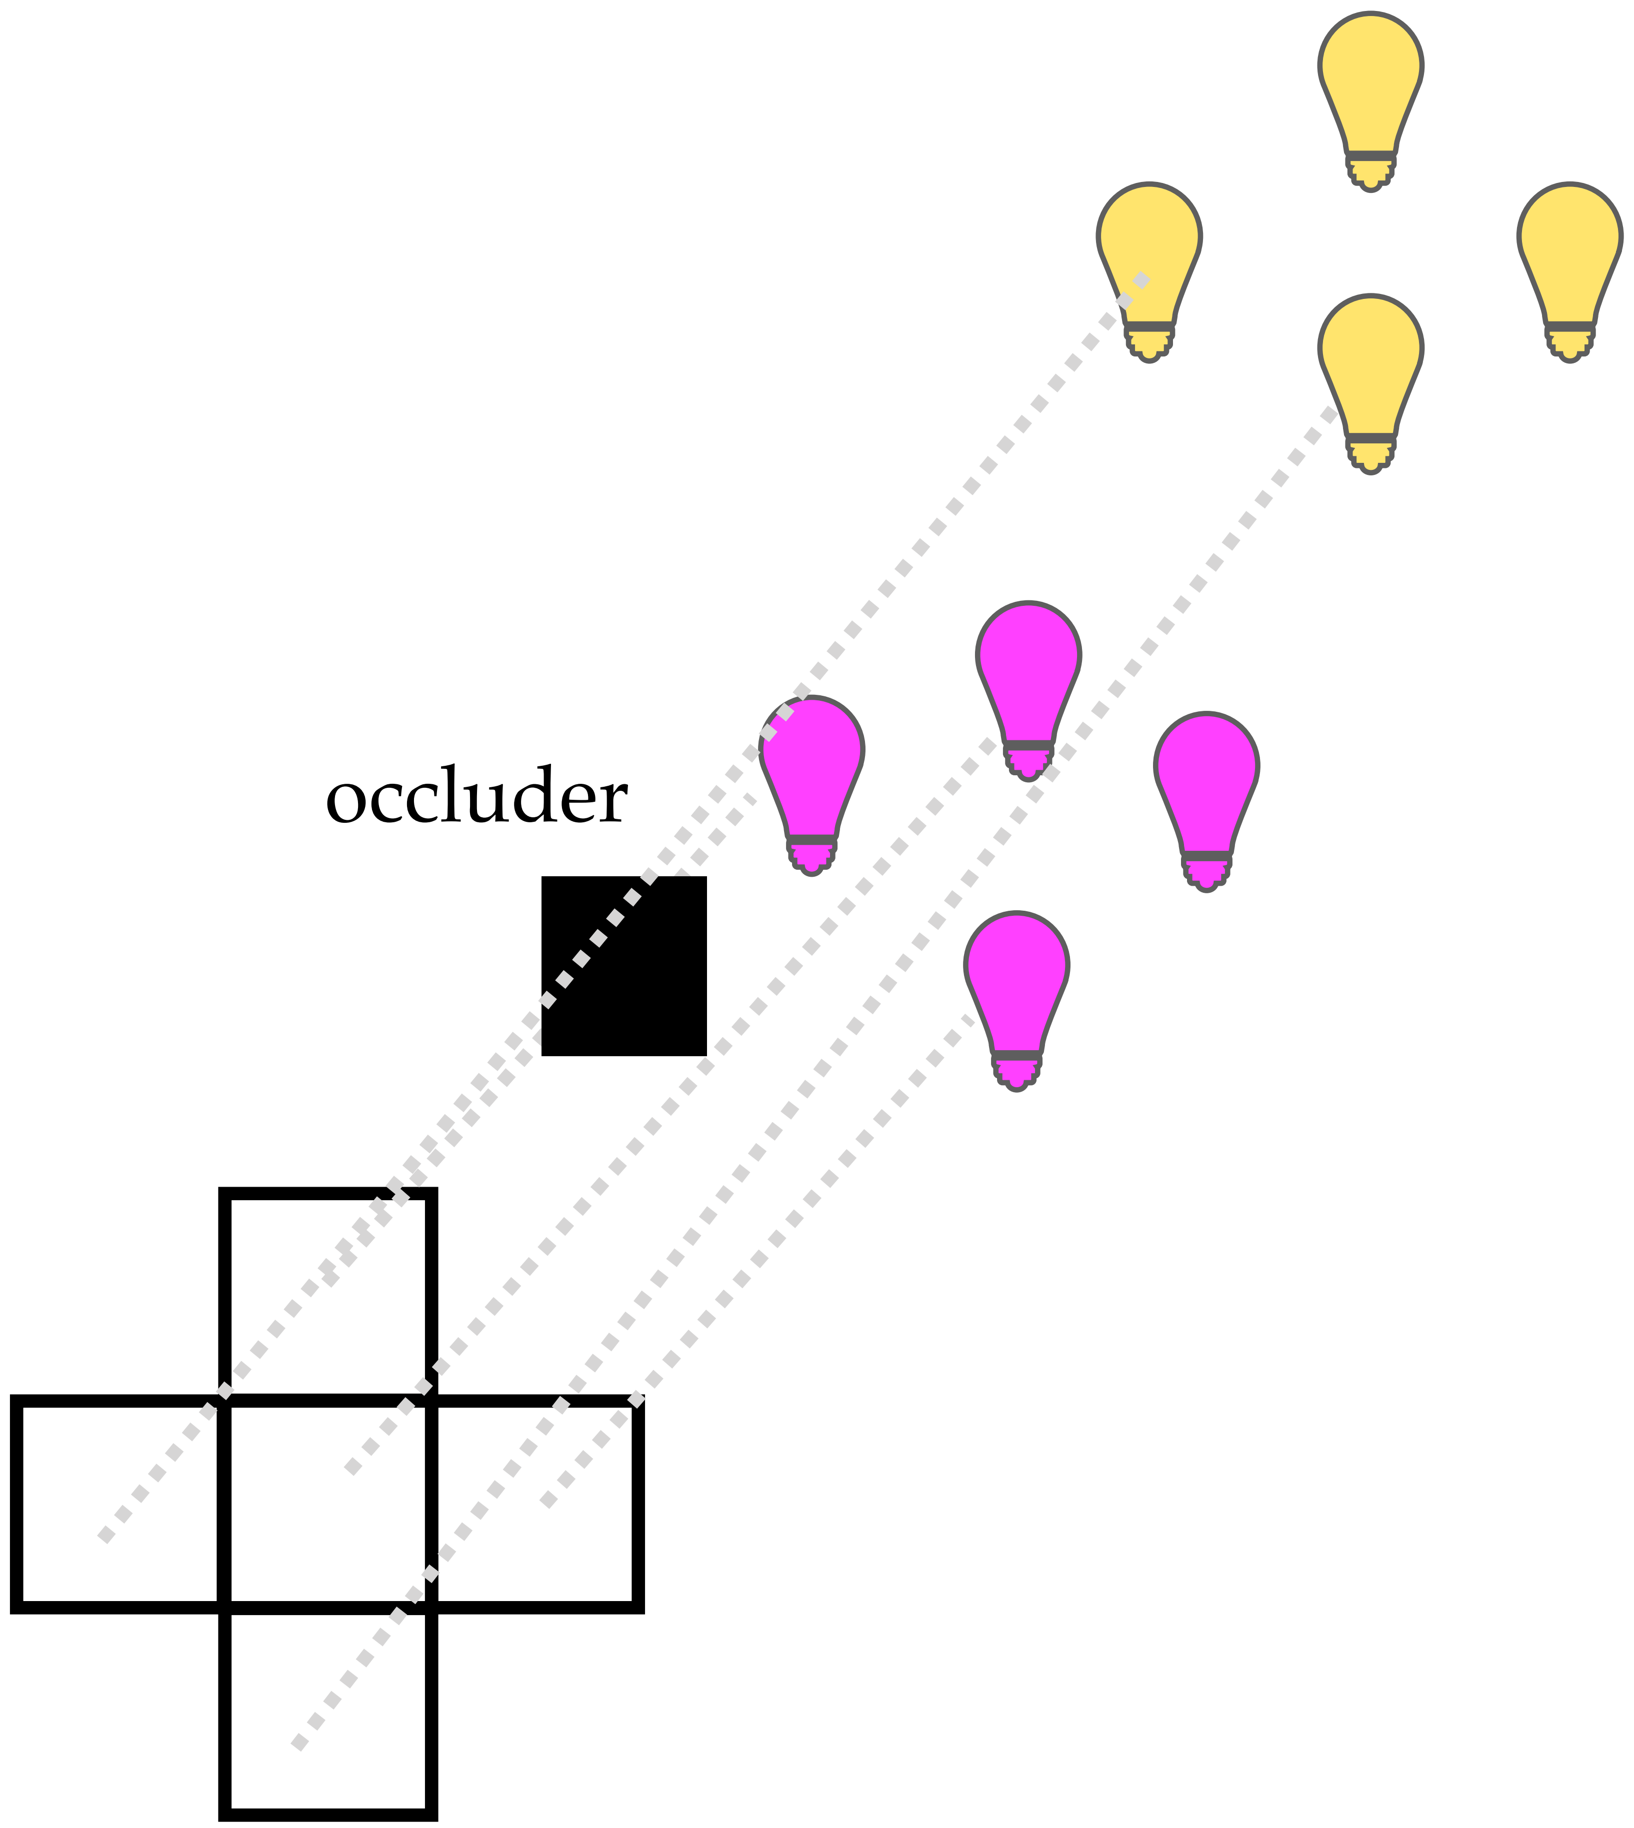

- each pixel starts with a single light sampled (e.g., uniform sampling)

Idea: reuse neighboring pixels' sampling results

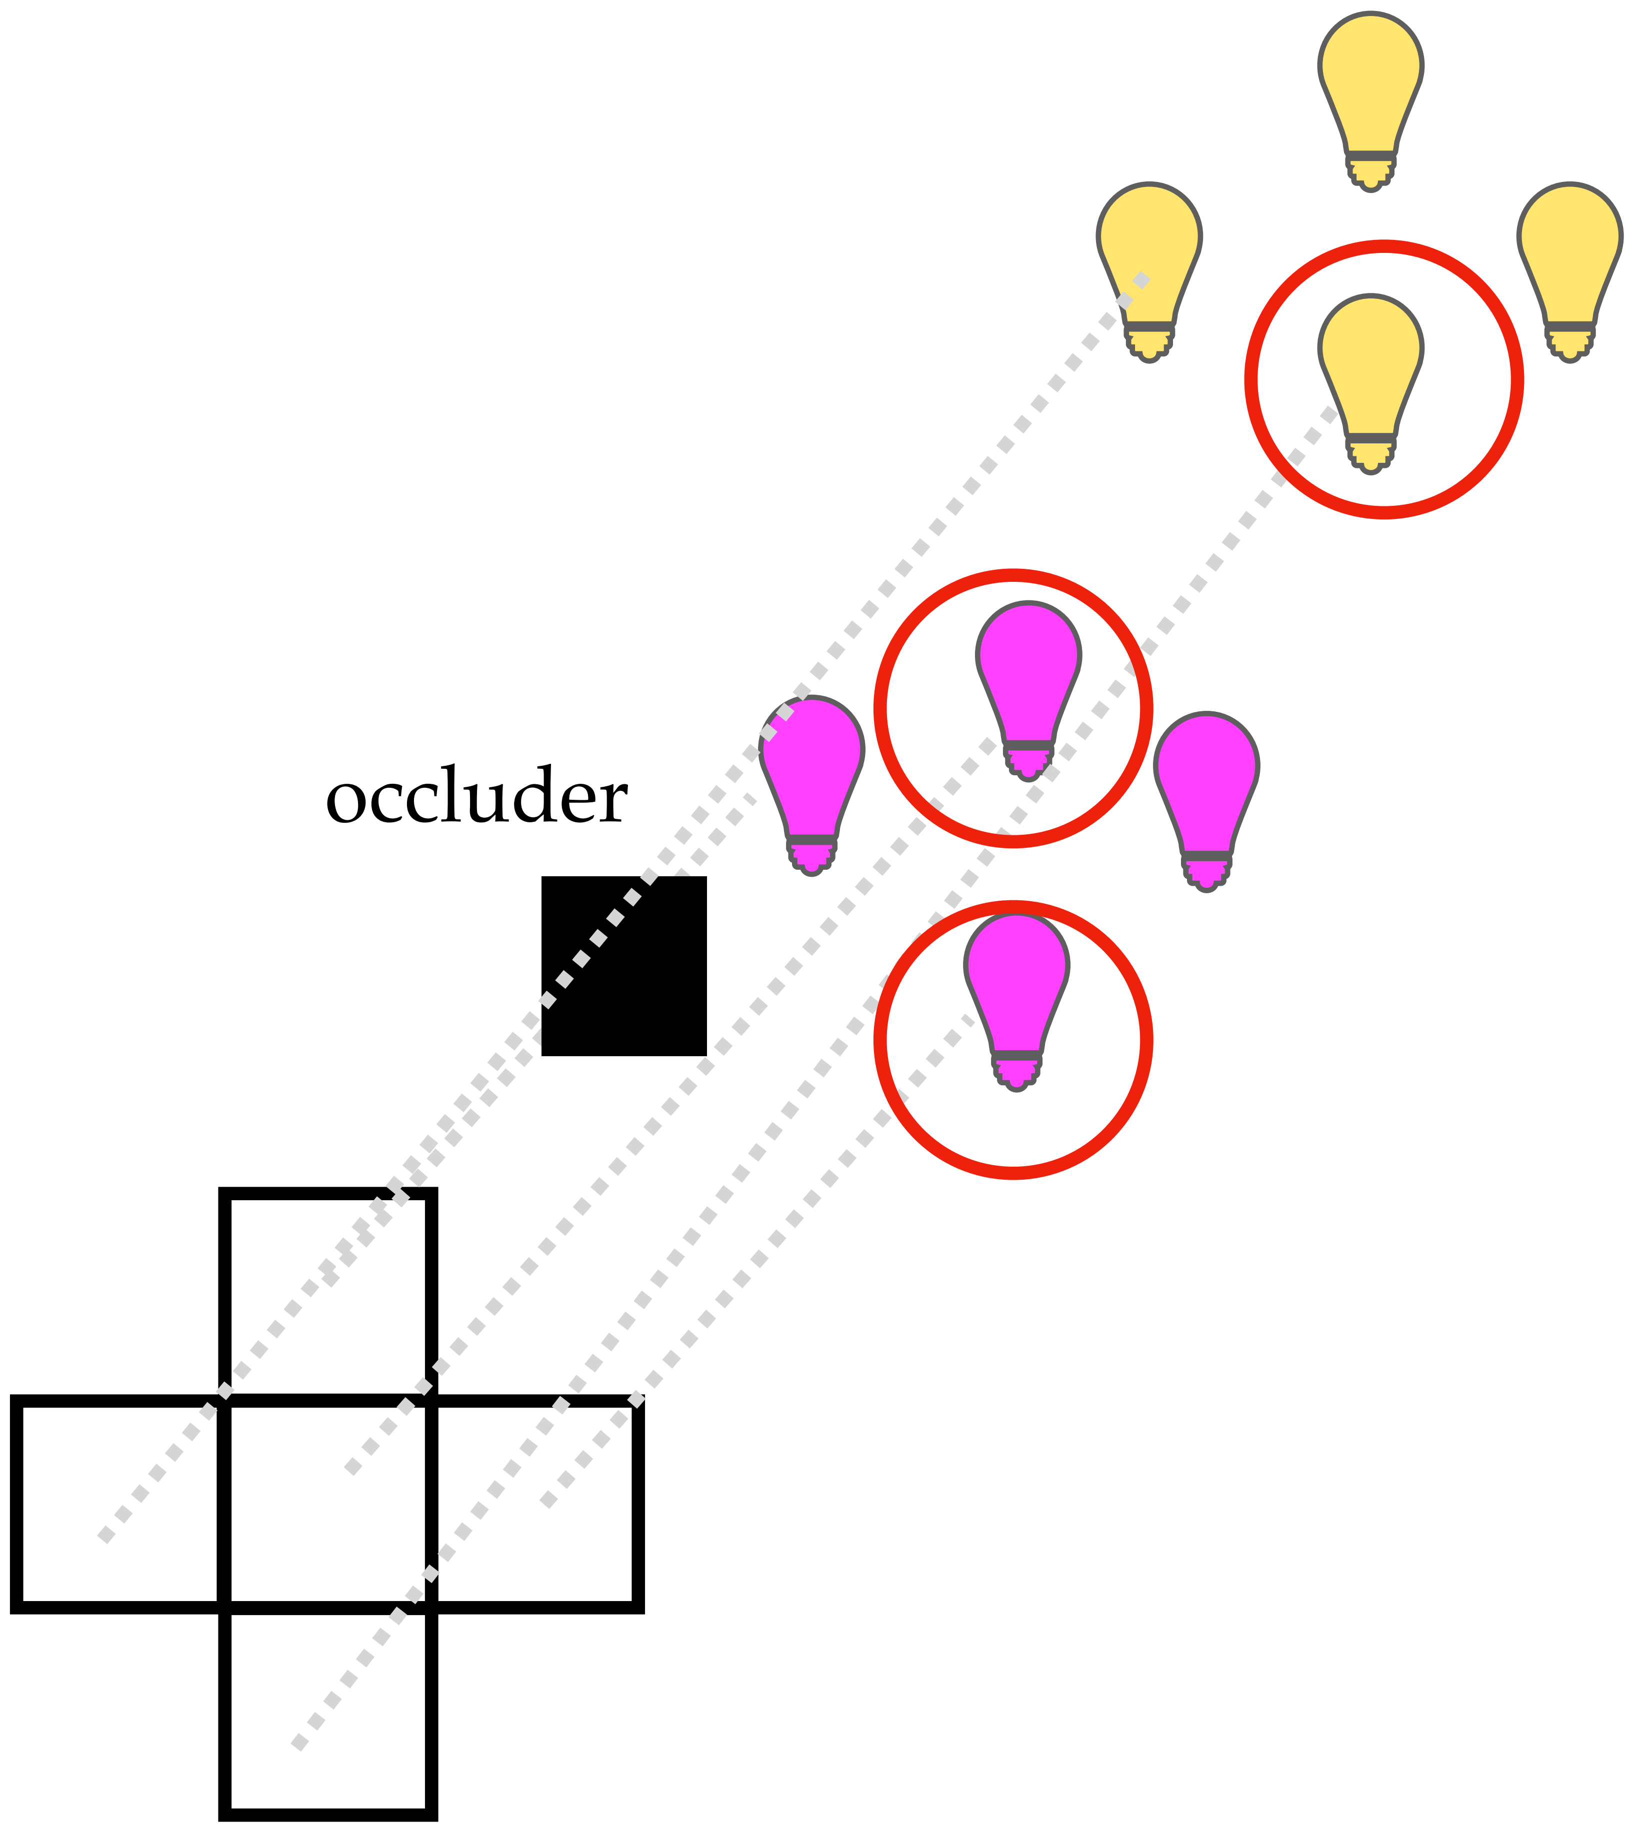

- each pixel starts with a single light sampled (e.g., uniform sampling)
- for the center pixel, pick the unoccluded lights from neighbor pixels

Idea: reuse neighboring pixels' sampling results

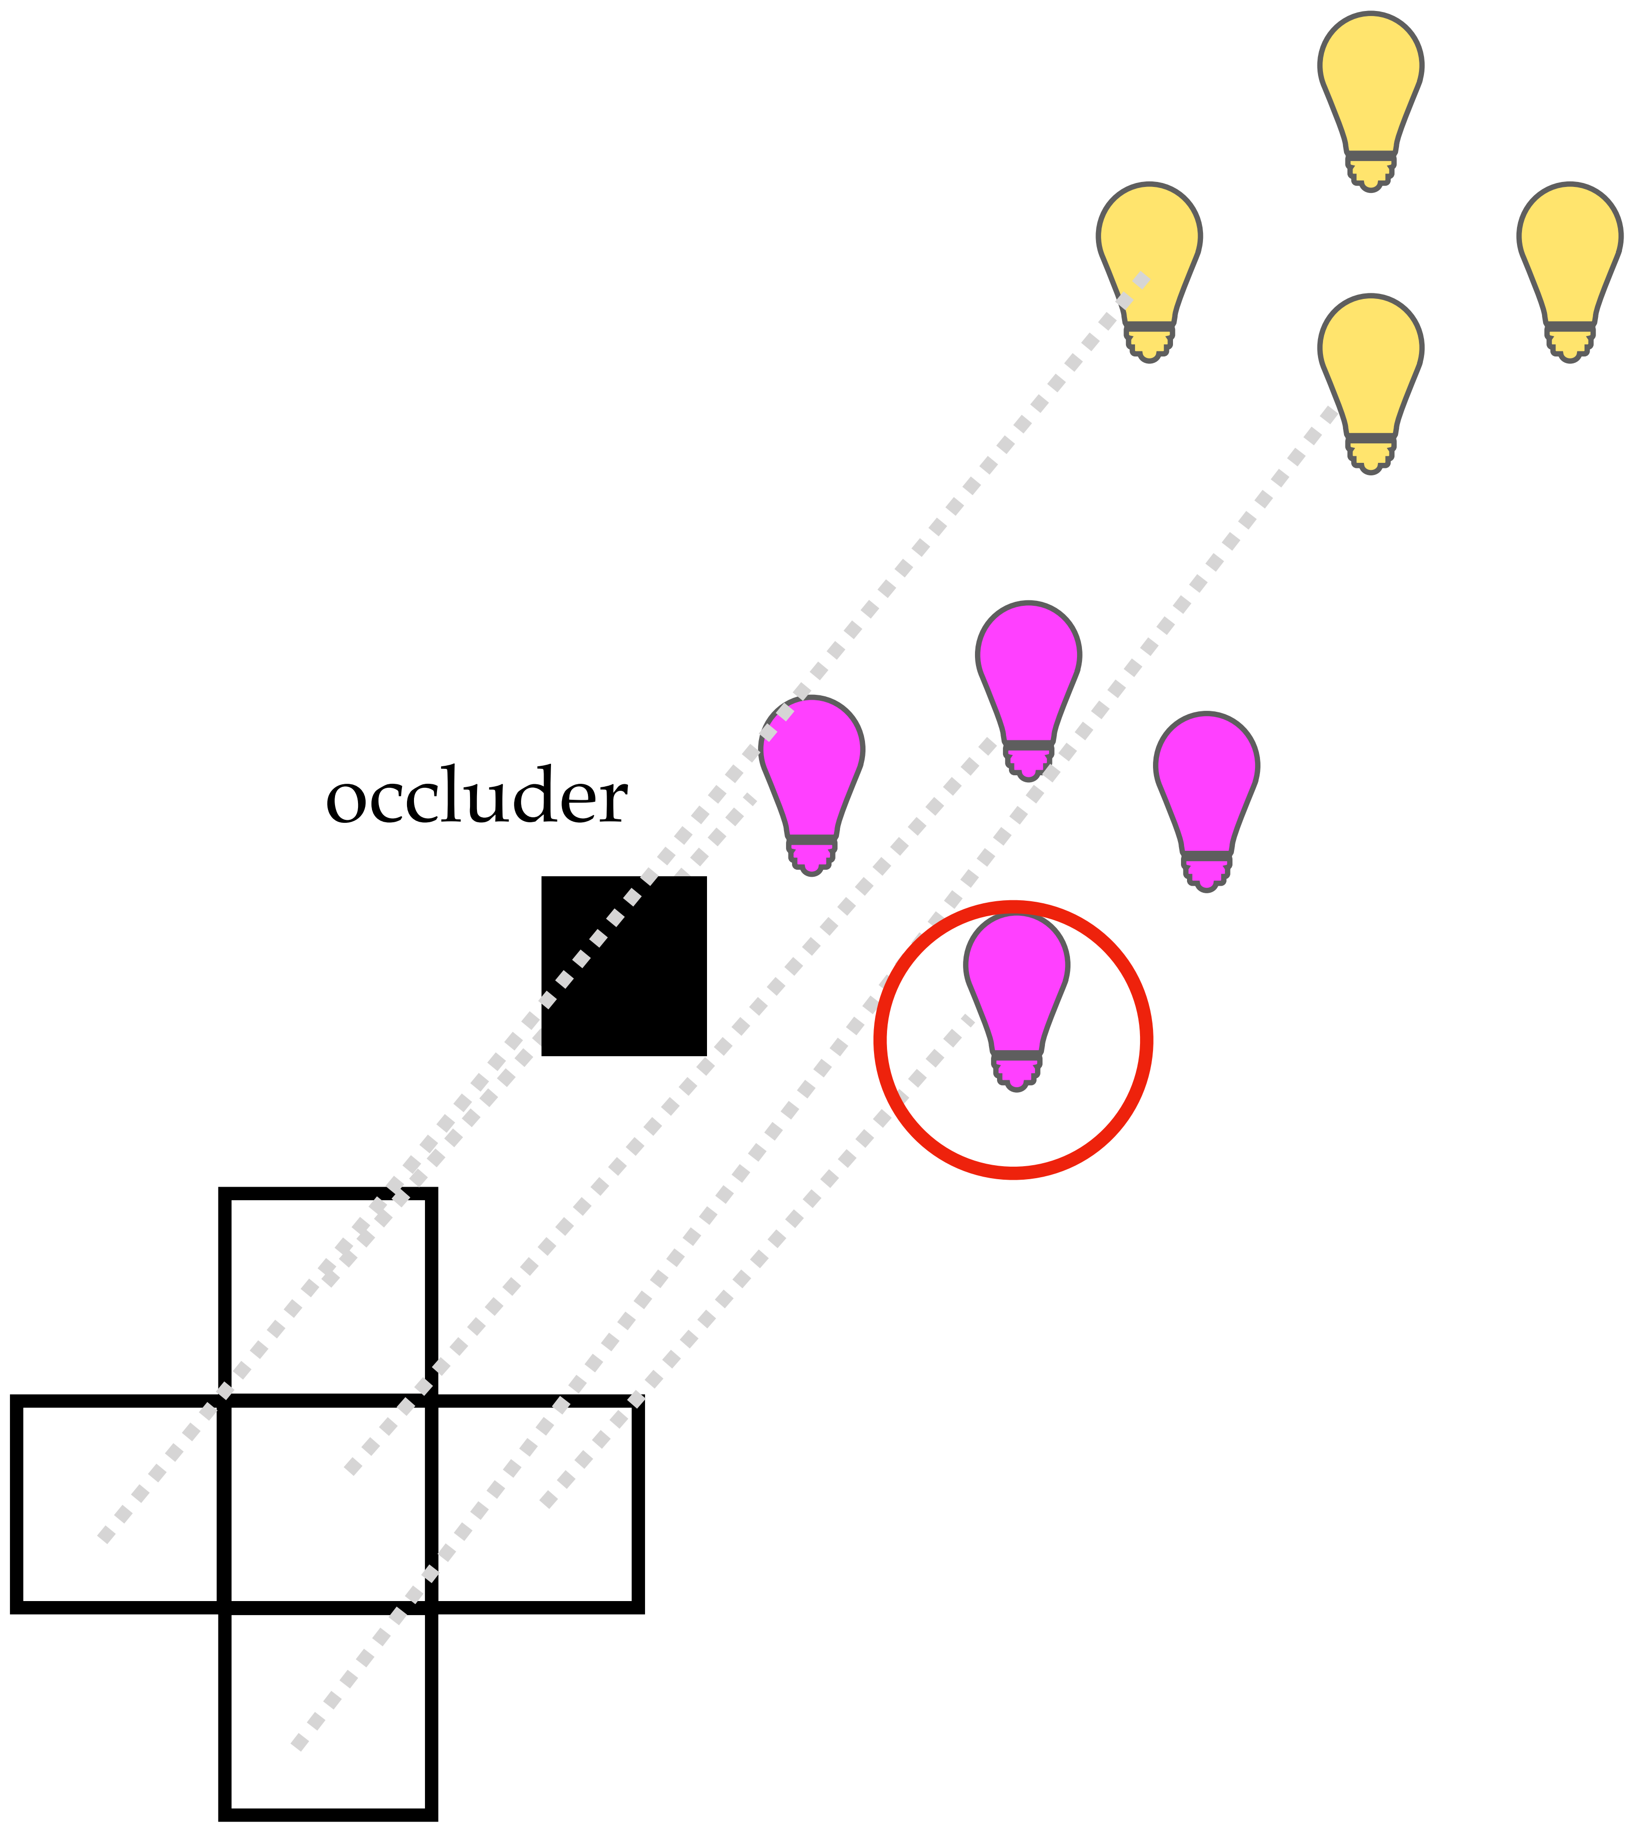

- each pixel starts with a single light sampled (e.g., uniform sampling)
- for the center pixel, pick the unoccluded lights from neighbor pixels
- sample from these lights using probability proportional to $L \cdot \rho \cdot G$

Idea: reuse neighboring pixels' sampling results

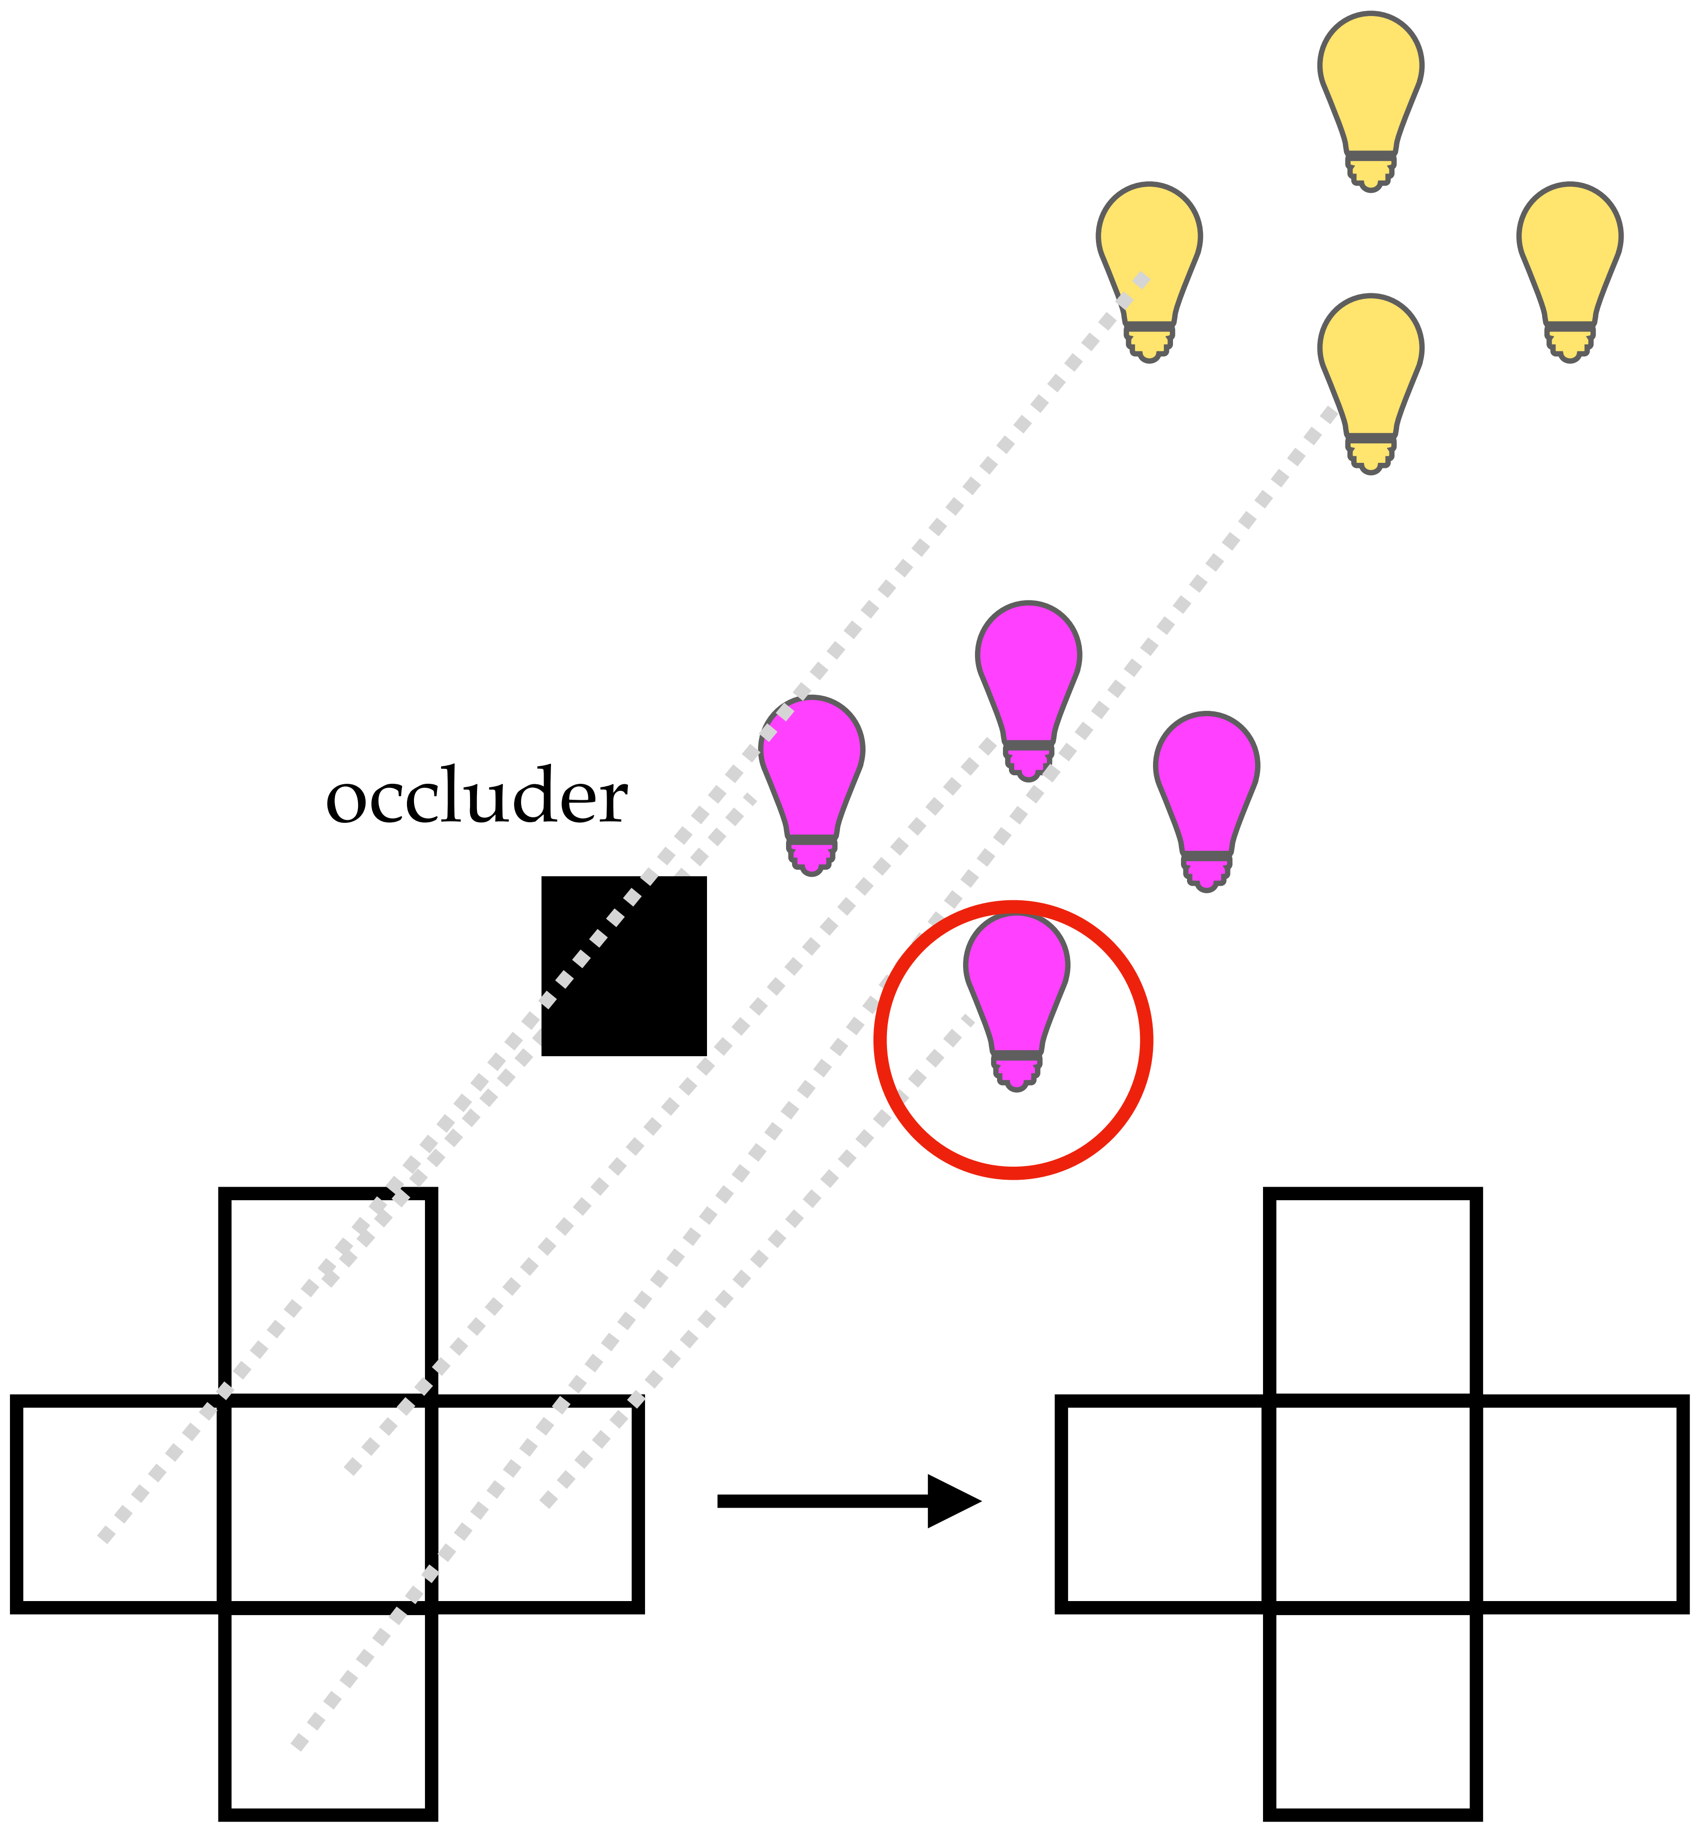

- each pixel starts with a single light sampled (e.g., uniform sampling)
- for the center pixel, pick the unoccluded lights from neighbor pixels
- sample from these lights using probability proportional to $L \cdot \rho \cdot G$
- propagate the information to the next frame

Idea: reuse neighboring pixels' sampling results

- benefits
  - occluded lights have low probability to be sampled
  - ρ & G are considered during reuse
  - sampling distribution is improved over time

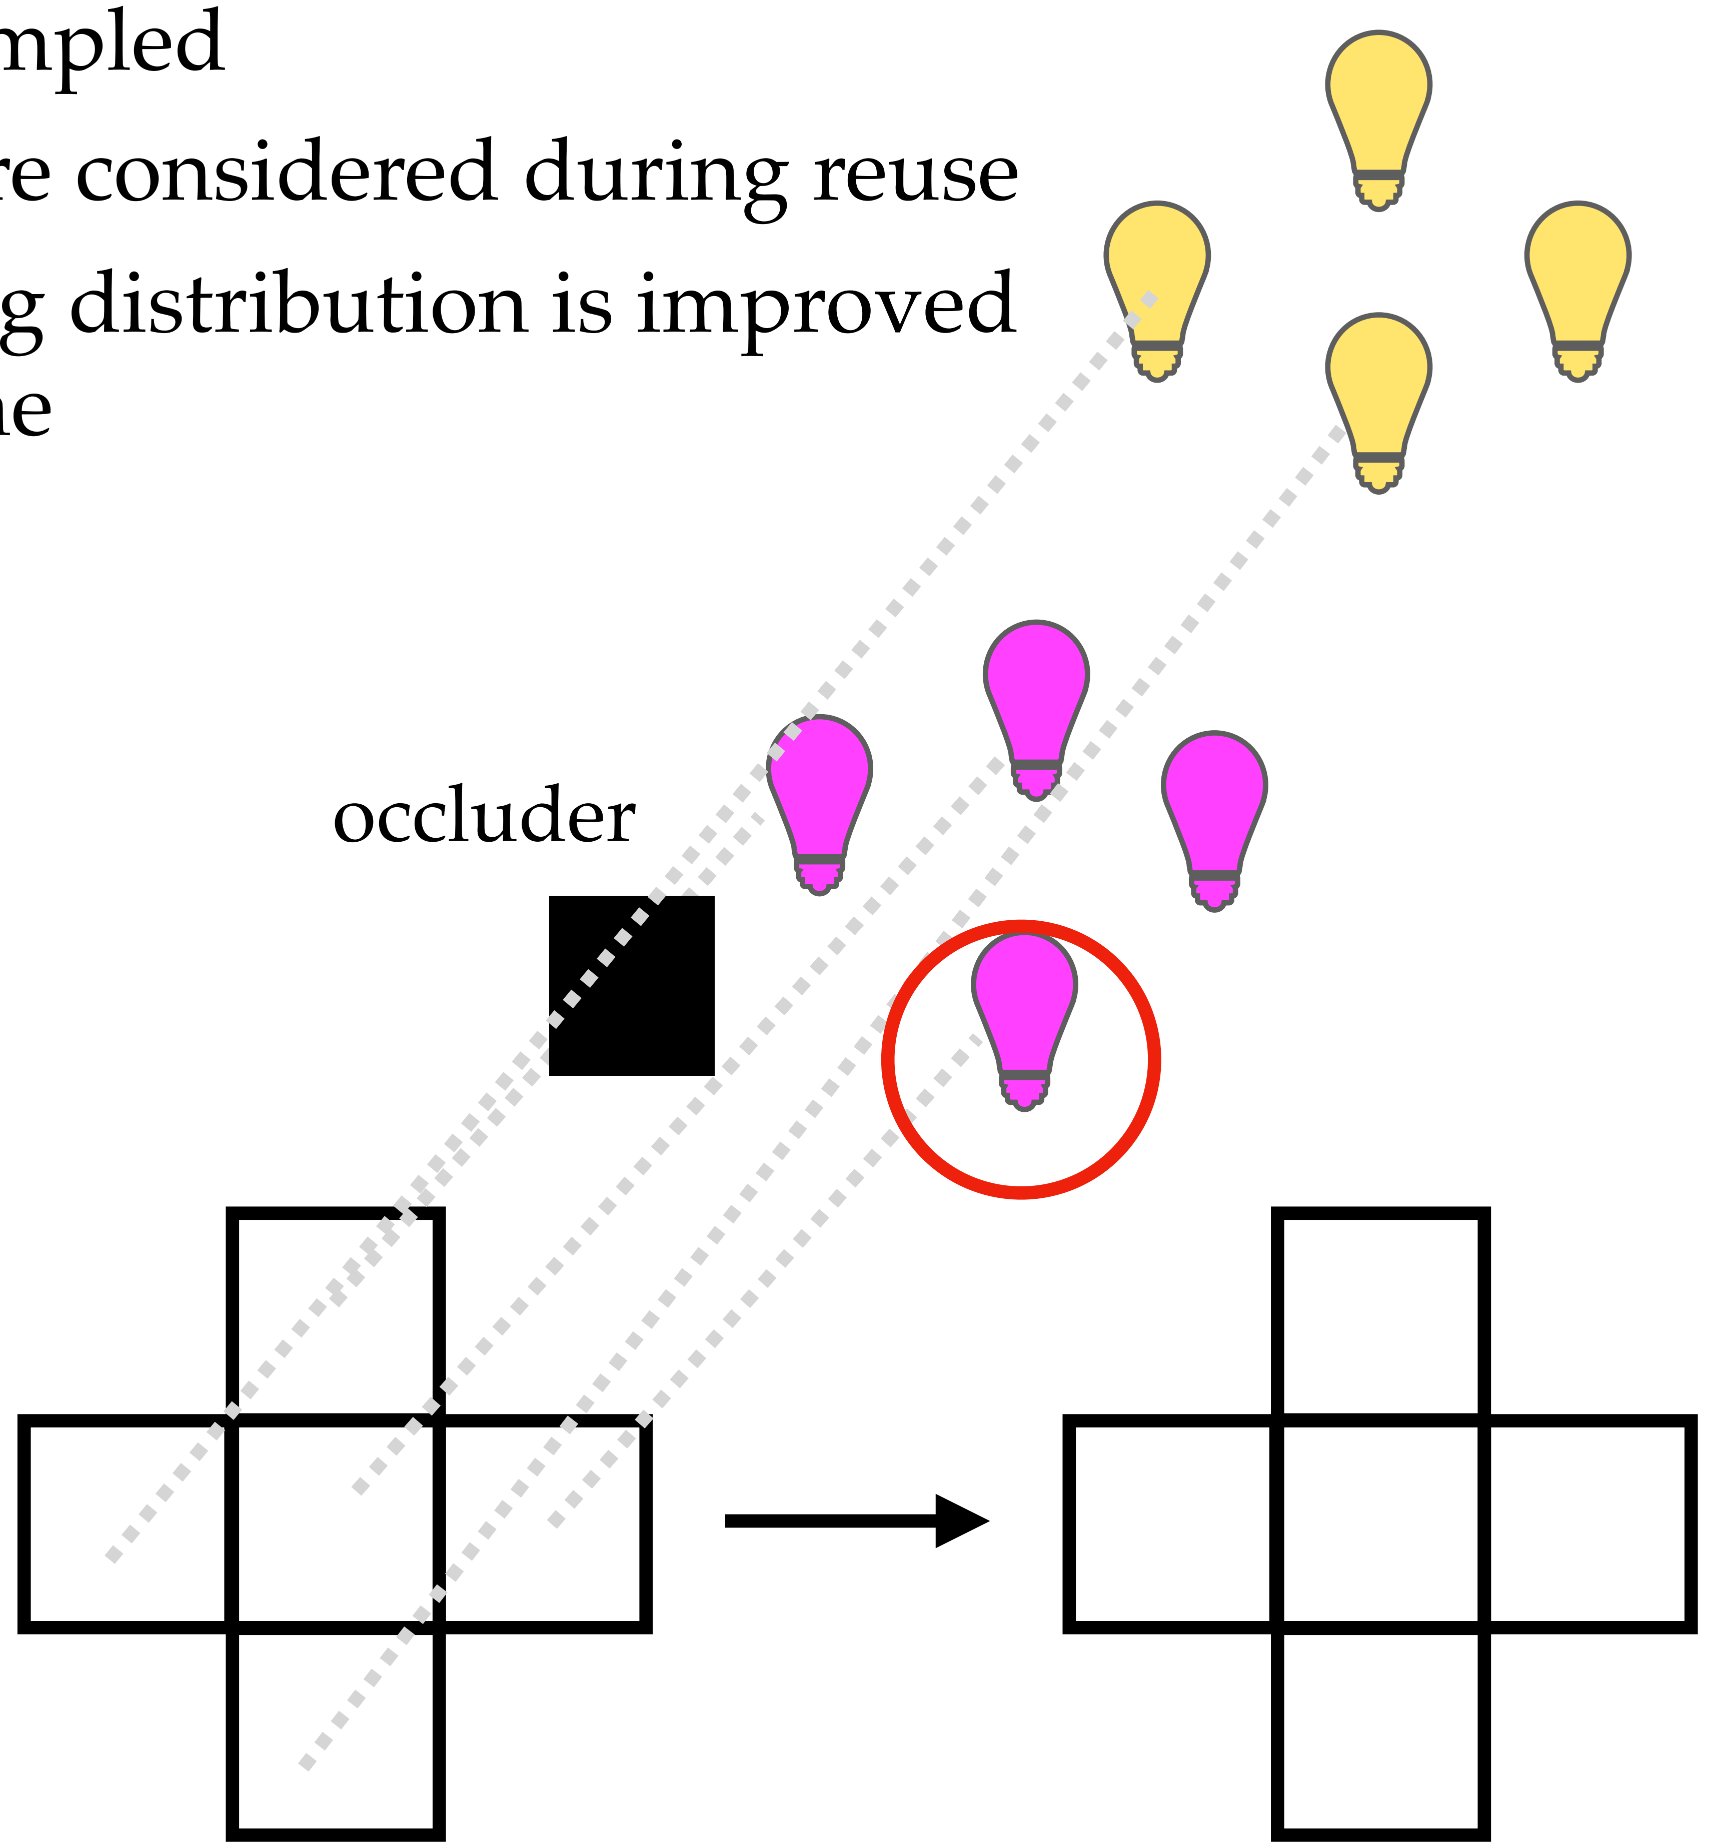

- each pixel starts with a single light sampled (e.g., uniform sampling)
- for the center pixel, pick the unoccluded lights from neighbor pixels
- sample from these lights using probability proportional to $L \cdot \rho \cdot G$
- propagate the information to the next frame

Math: resampled importance sampling

goal: approximately sampled arbitrary unnormalized target distribution $\hat{p}$

Importance Resampling for Global Illumination

Justin F. Talbot   David Cline   Parris Egbert

Brigham Young University

Math: resampled importance sampling

goal: approximately sampled arbitrary unnormalized target distribution $\hat{p}$

start with M samples with "candidate" distribution p

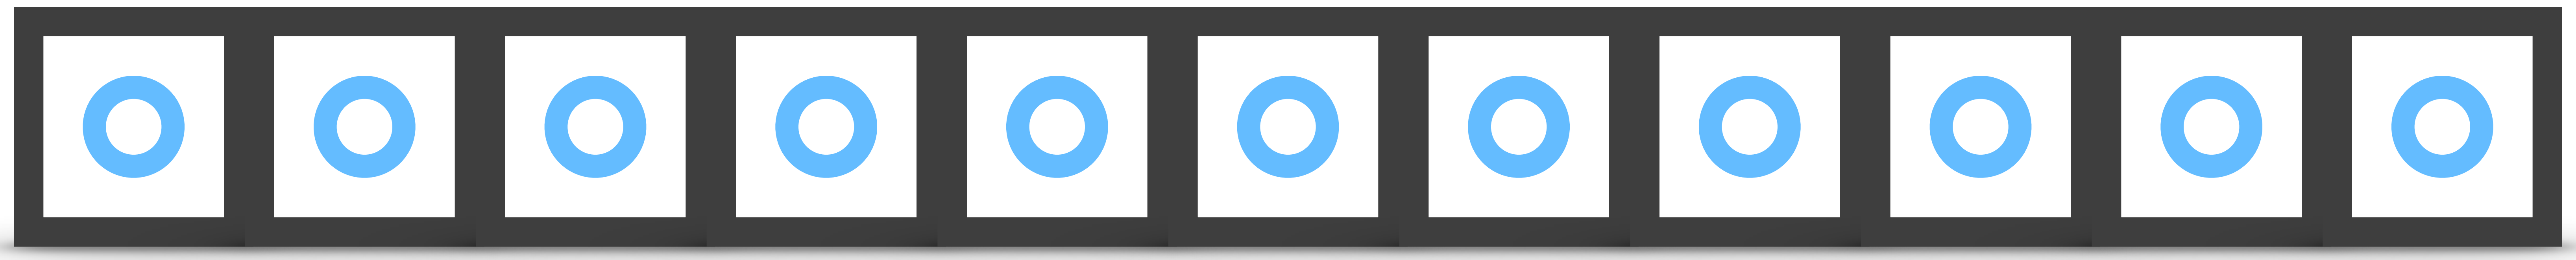

Math: resampled importance sampling

goal: approximately sampled arbitrary unnormalized target distribution $\hat{p}$

start with M samples with "candidate" distribution p

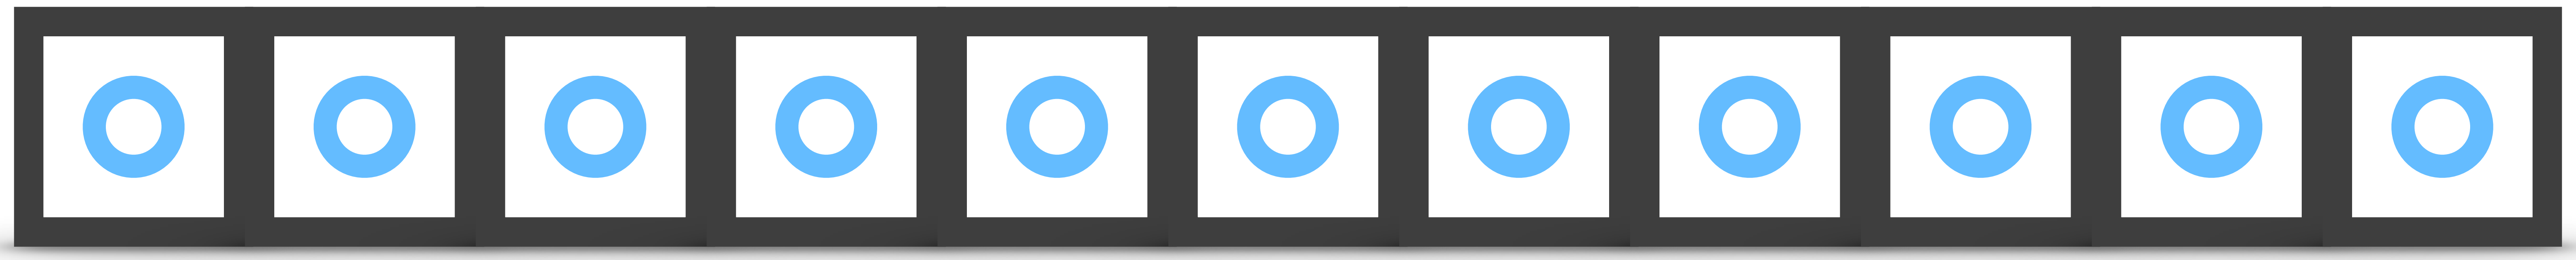

$$\frac{\hat{p}}{p} = \quad 5 \quad 8 \quad 0 \quad 2 \quad 3 \quad 5 \quad 20 \quad 4 \quad 6 \quad 1 \quad 7$$

evaluate $\frac{\hat{p}}{p}$ on all of them

Math: resampled importance sampling

goal: approximately sampled arbitrary unnormalized target distribution $\hat{p}$

start with M samples with "candidate" distribution p

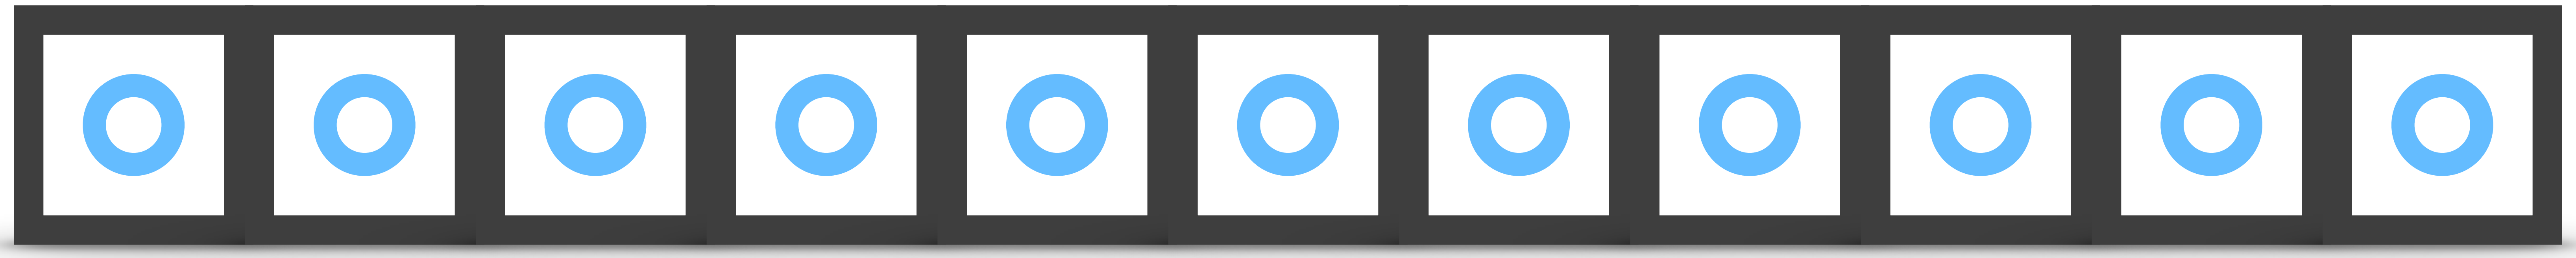

$$\frac{\hat{p}}{p} = 5 \quad 8 \quad 0 \quad 2 \quad 3 \quad 5 \quad 20 \quad 4 \quad 6 \quad 1 \quad 7$$

evaluate $\frac{\hat{p}}{p}$ on all of them

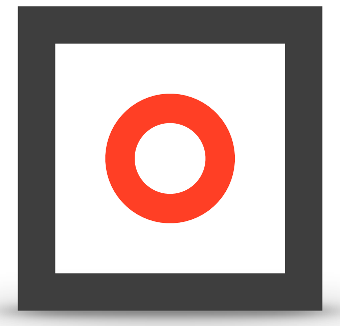

pick a sample with prob. proportional to $\frac{\hat{p}}{p}$

Math: resampled importance sampling

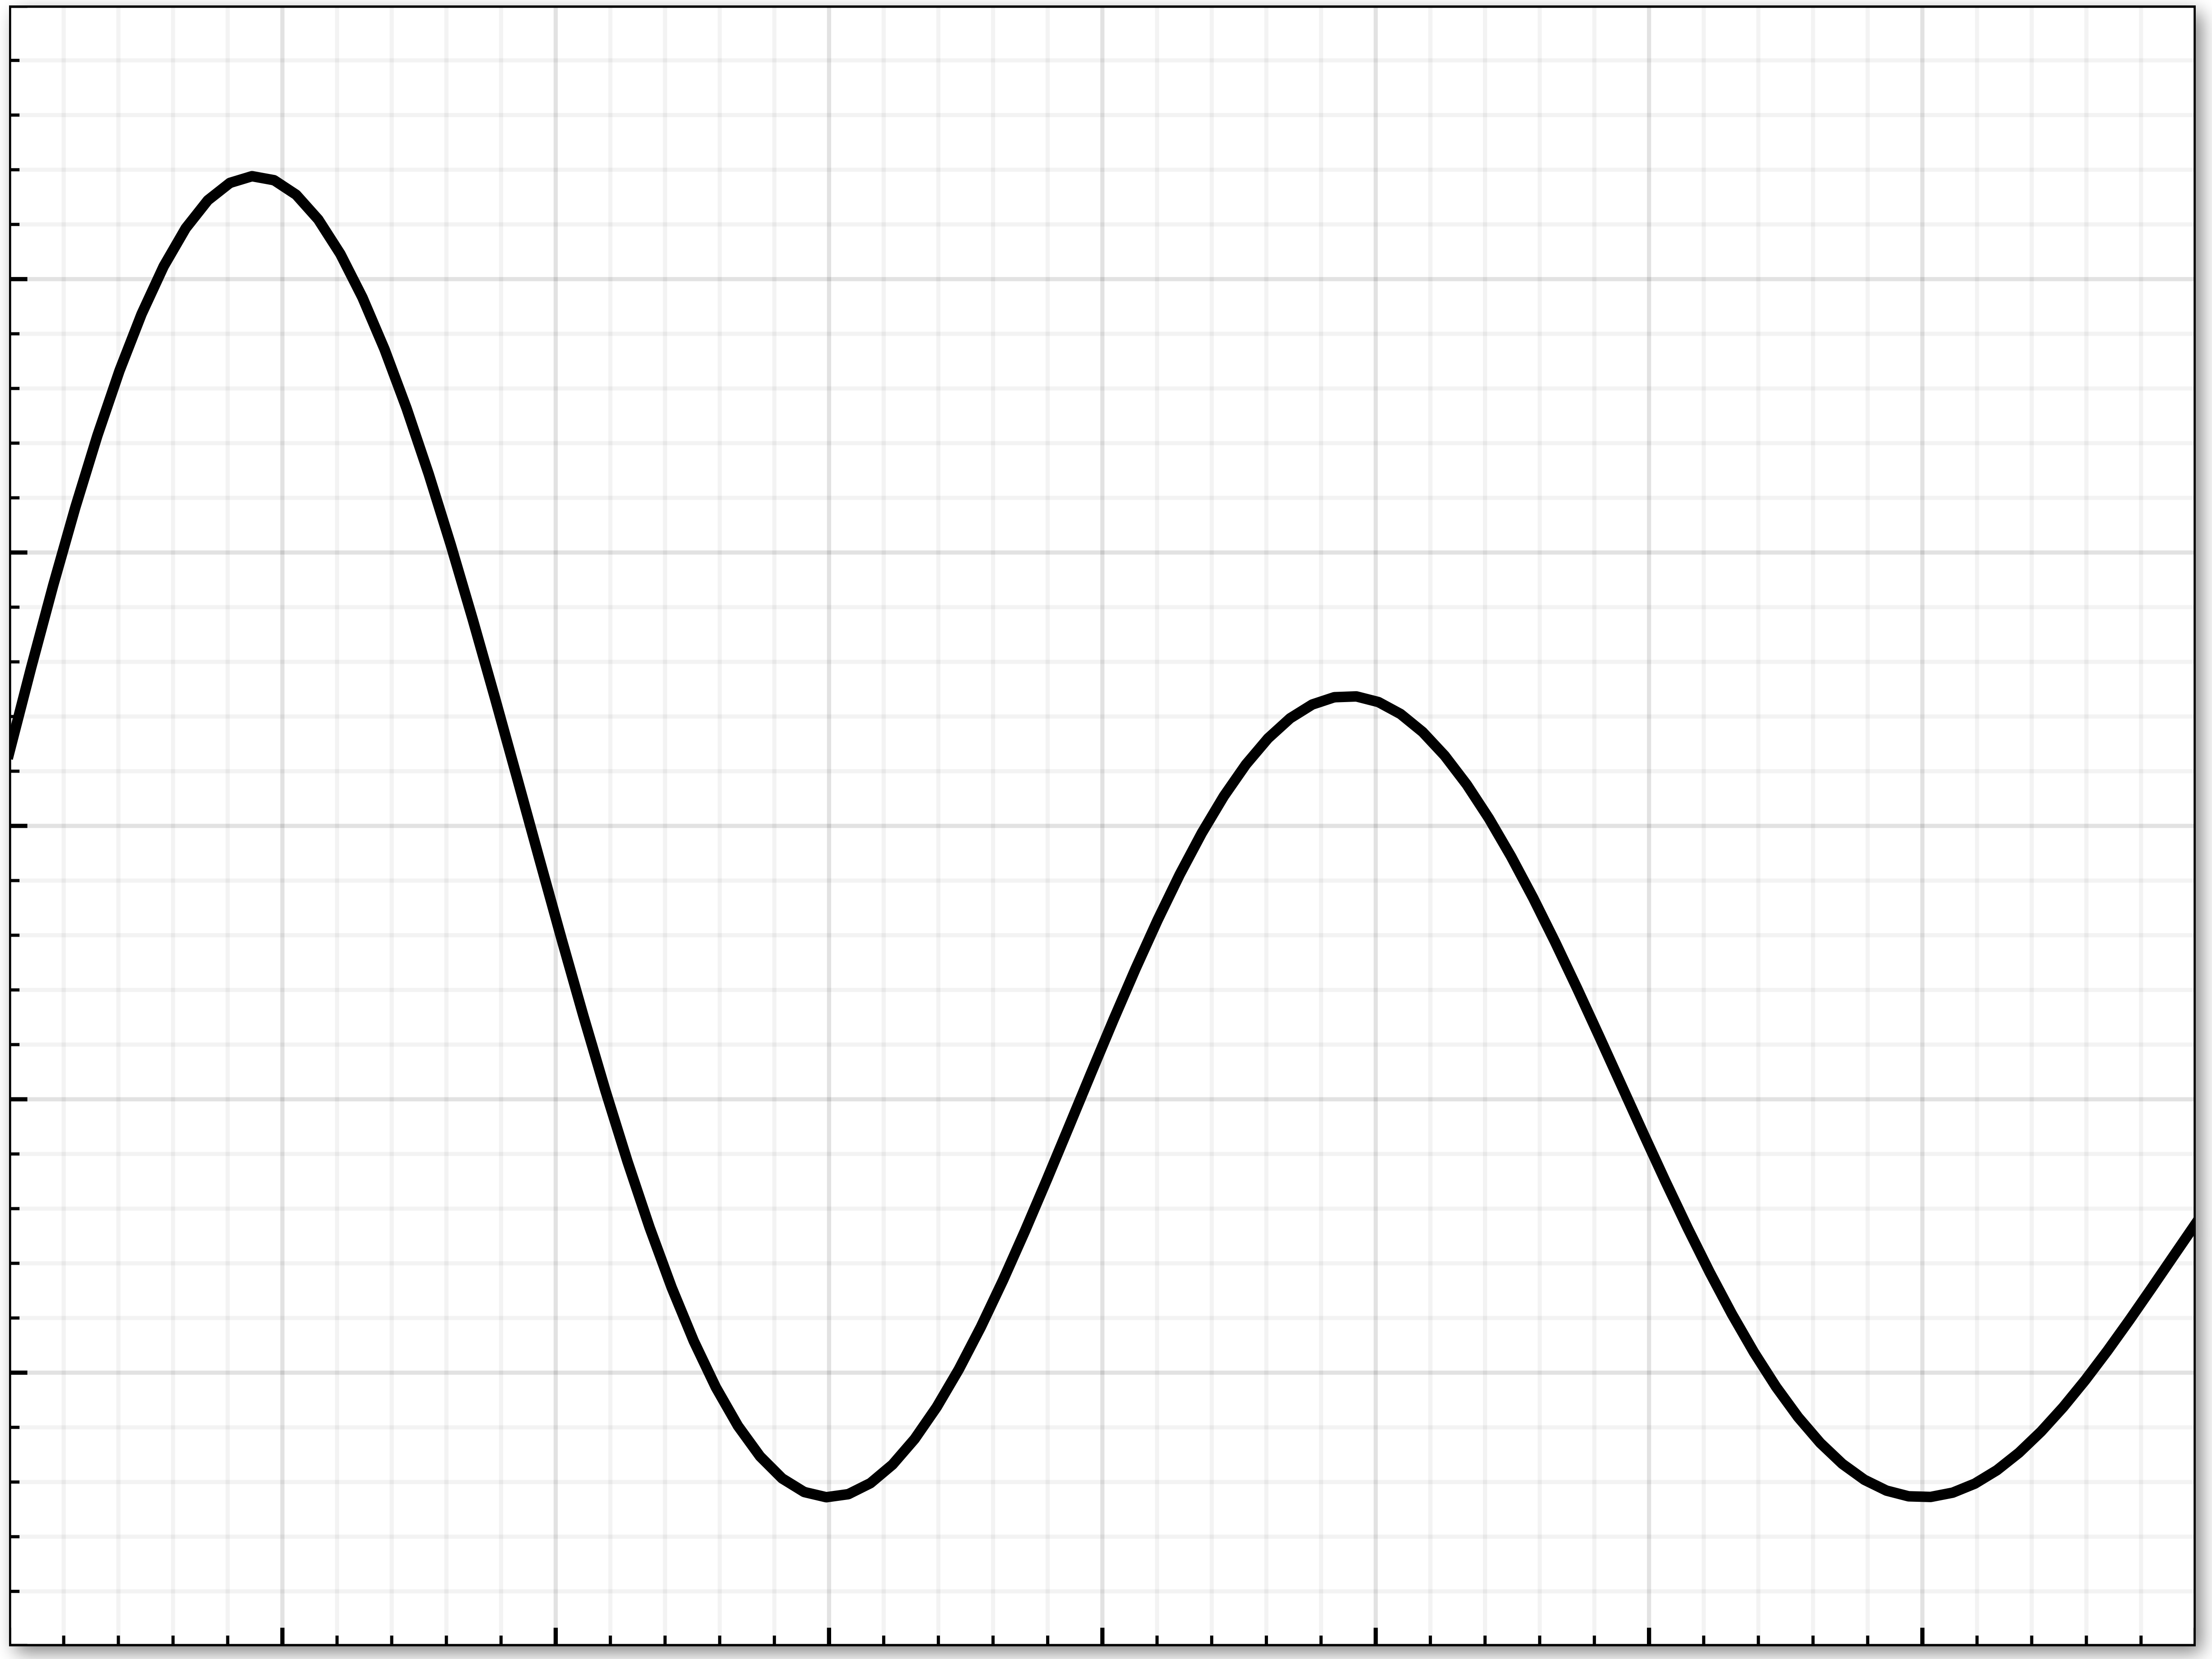

$\hat{p}$

Math: resampled importance sampling

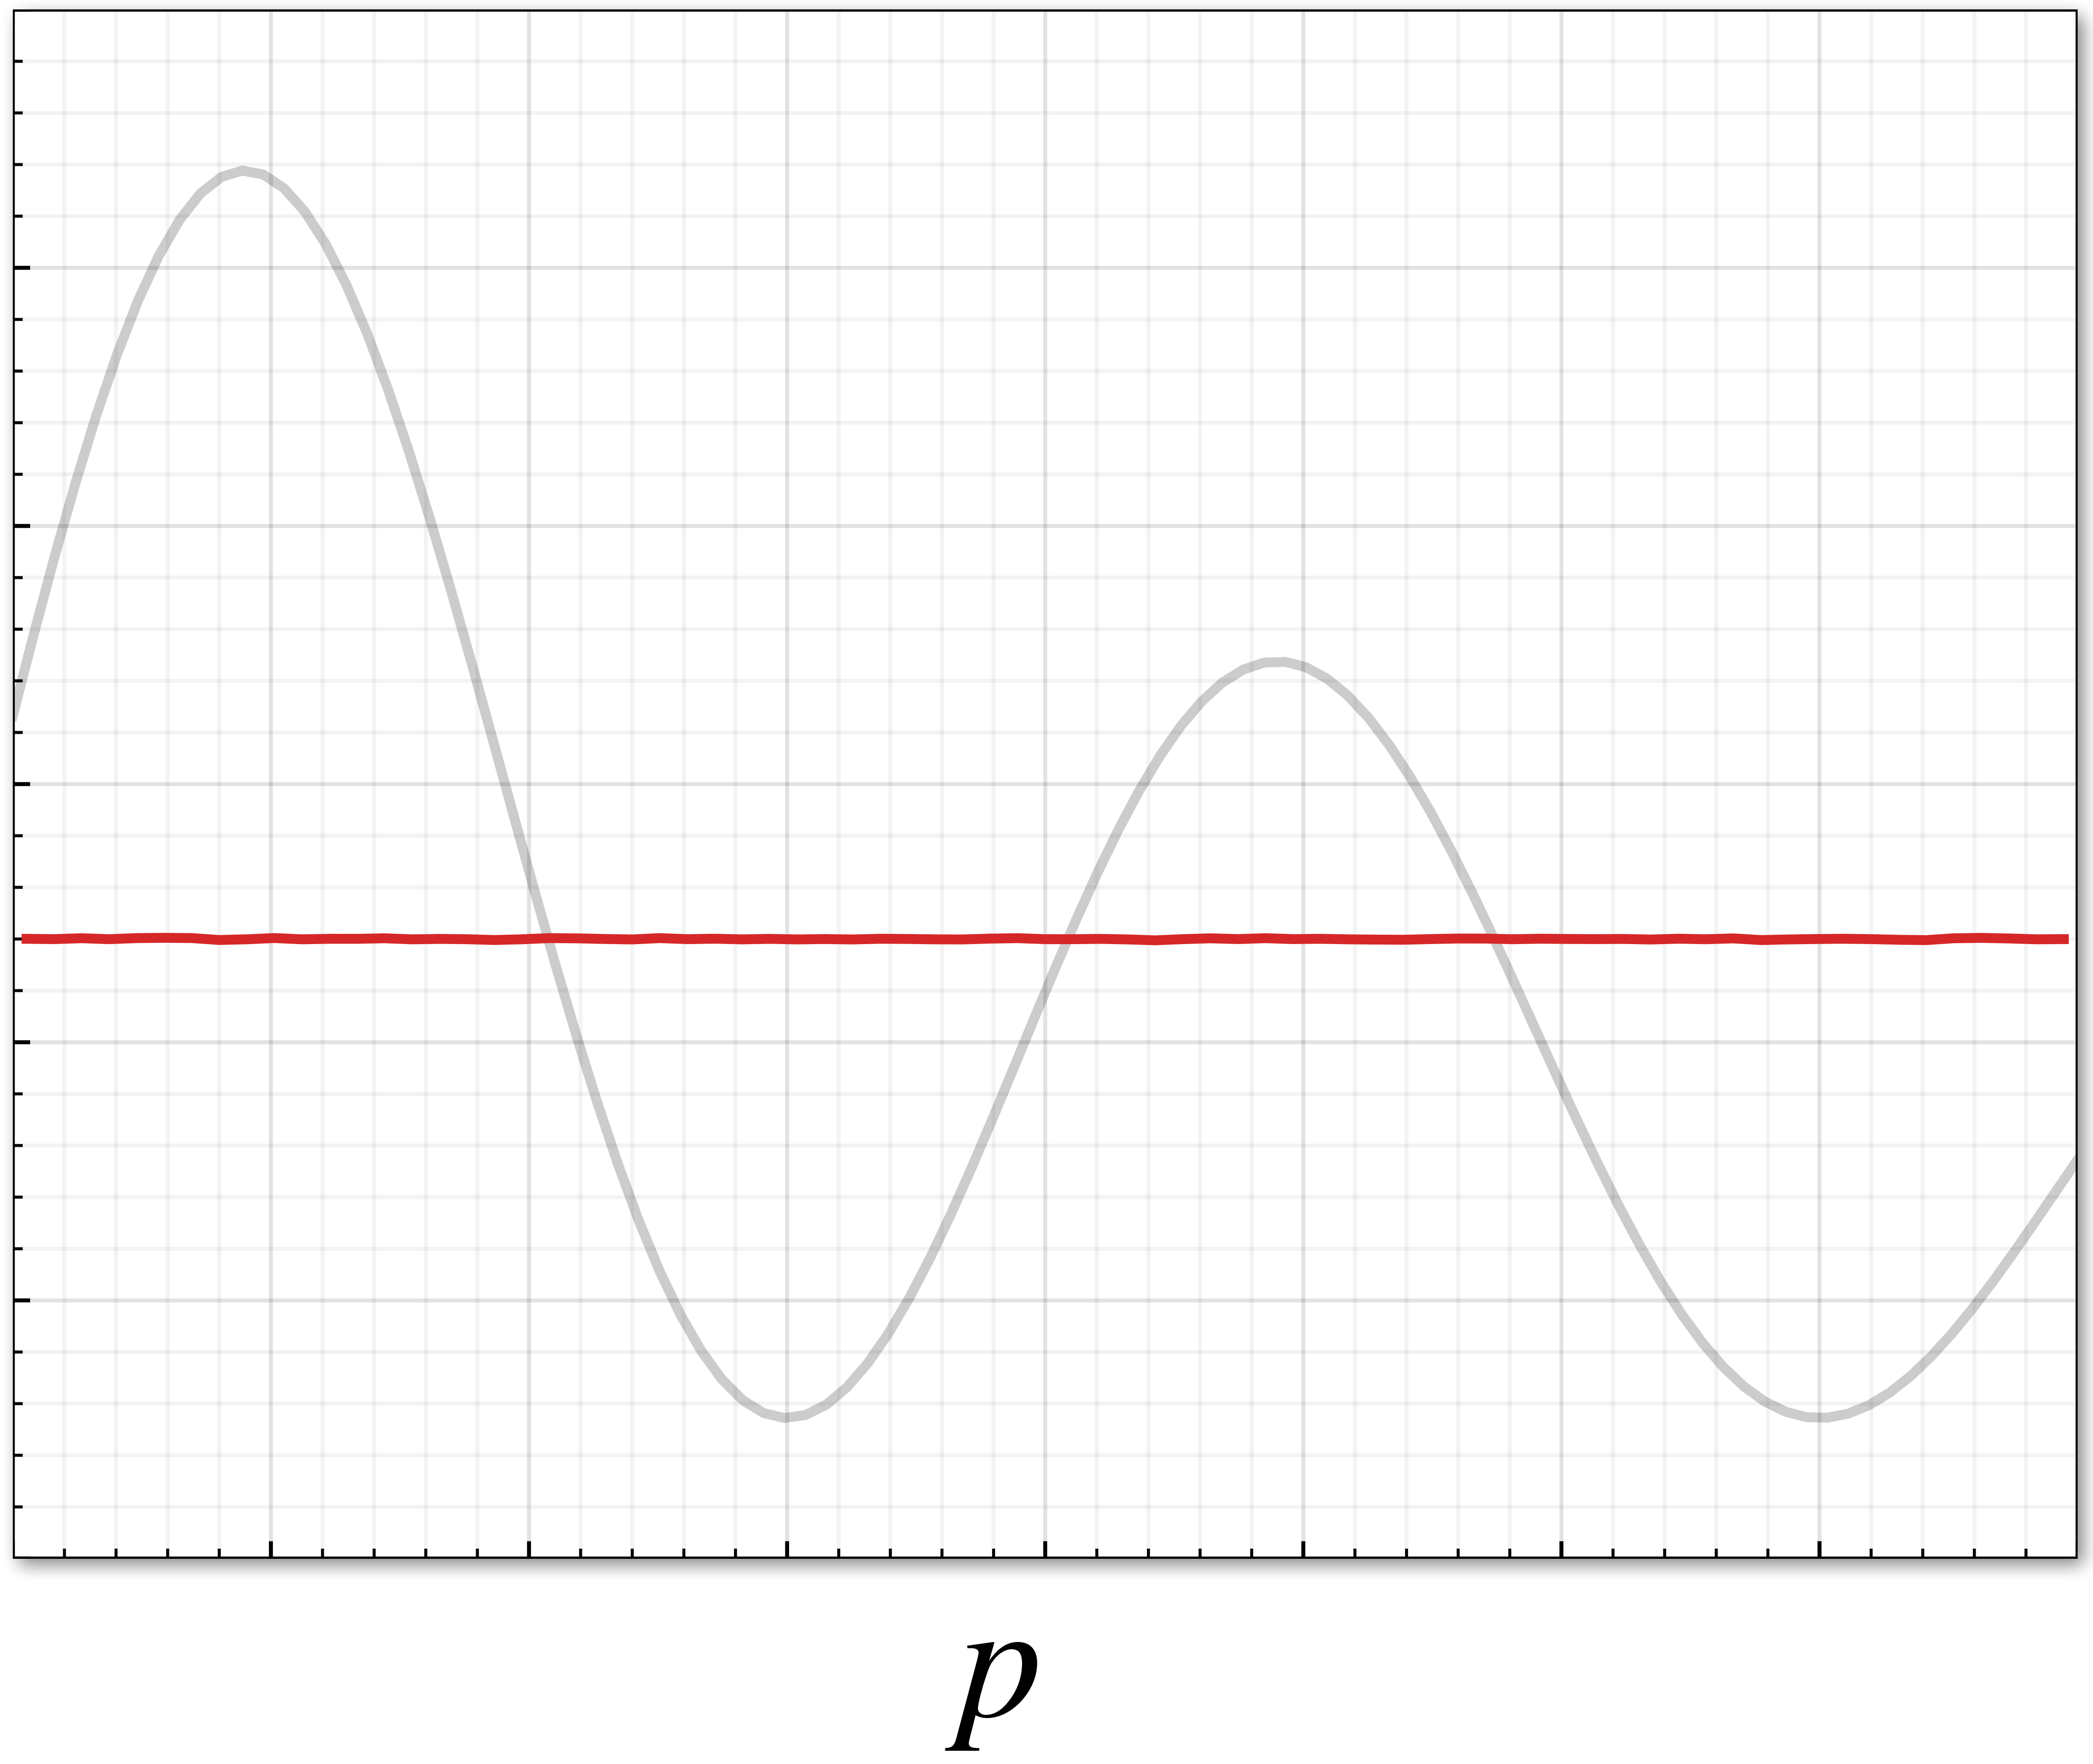

Math: resampled importance sampling

density of
the final sample

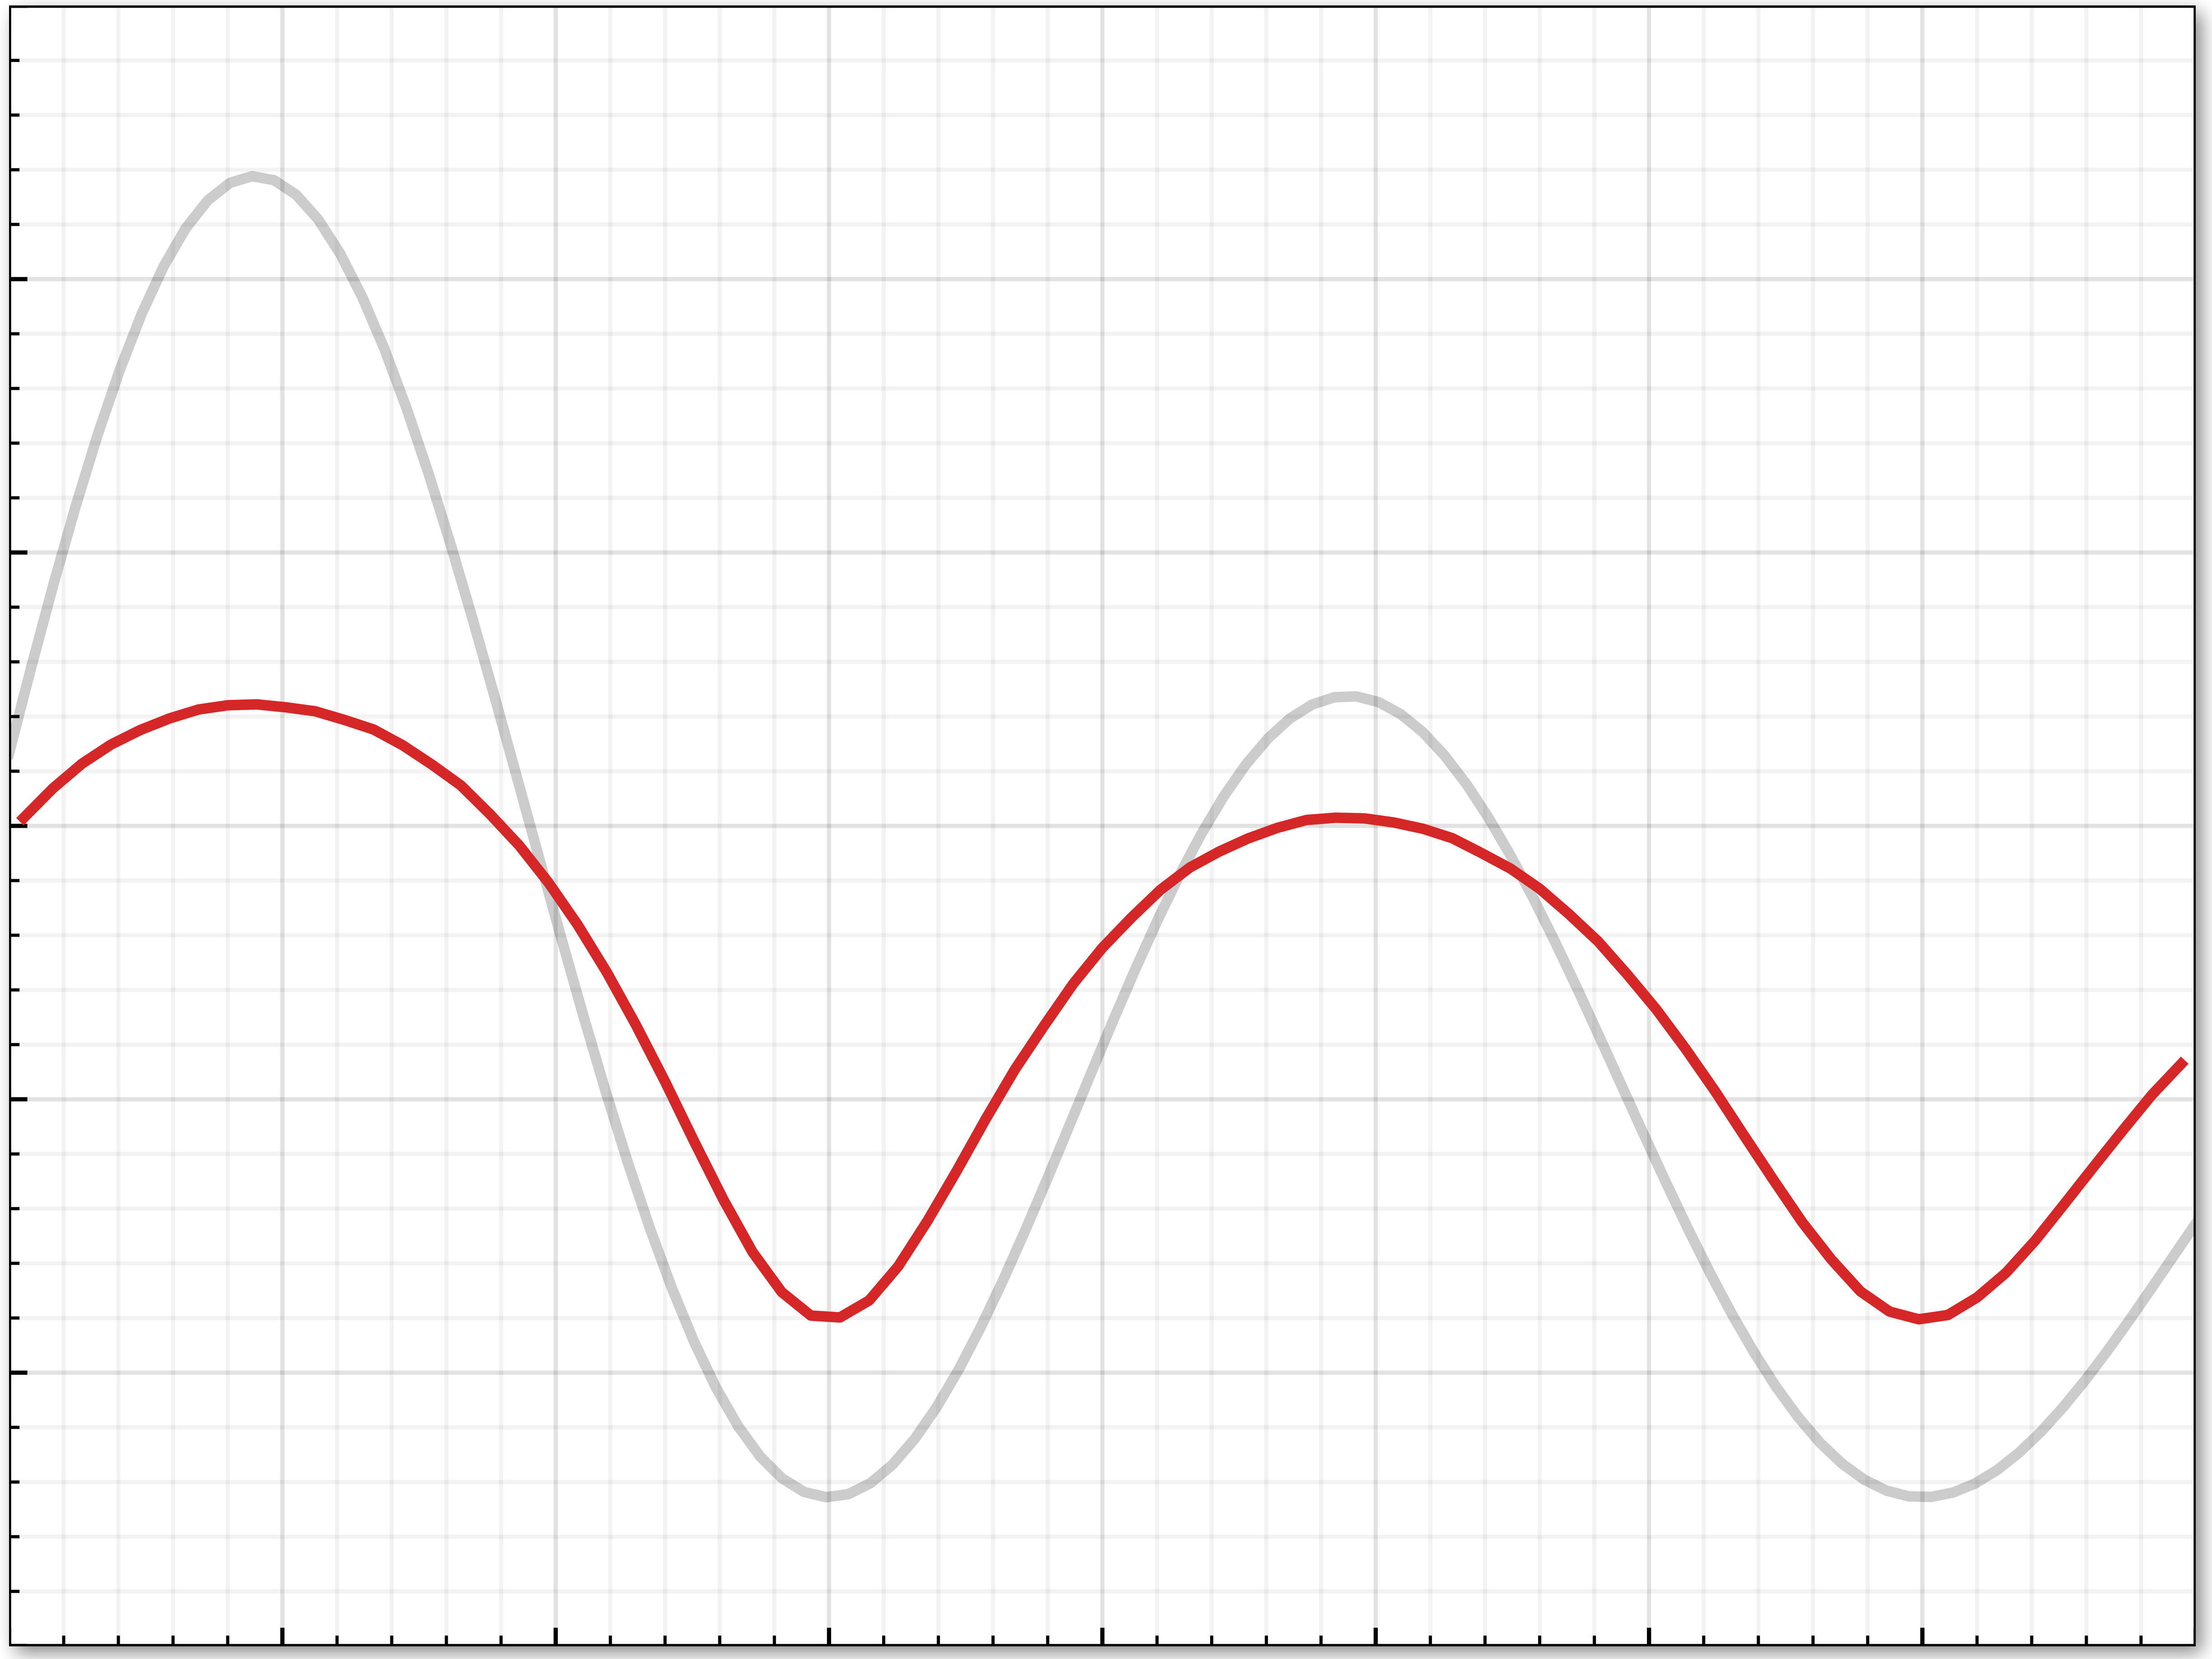

$M = 2$

Math: resampled importance sampling

density of
the final sample

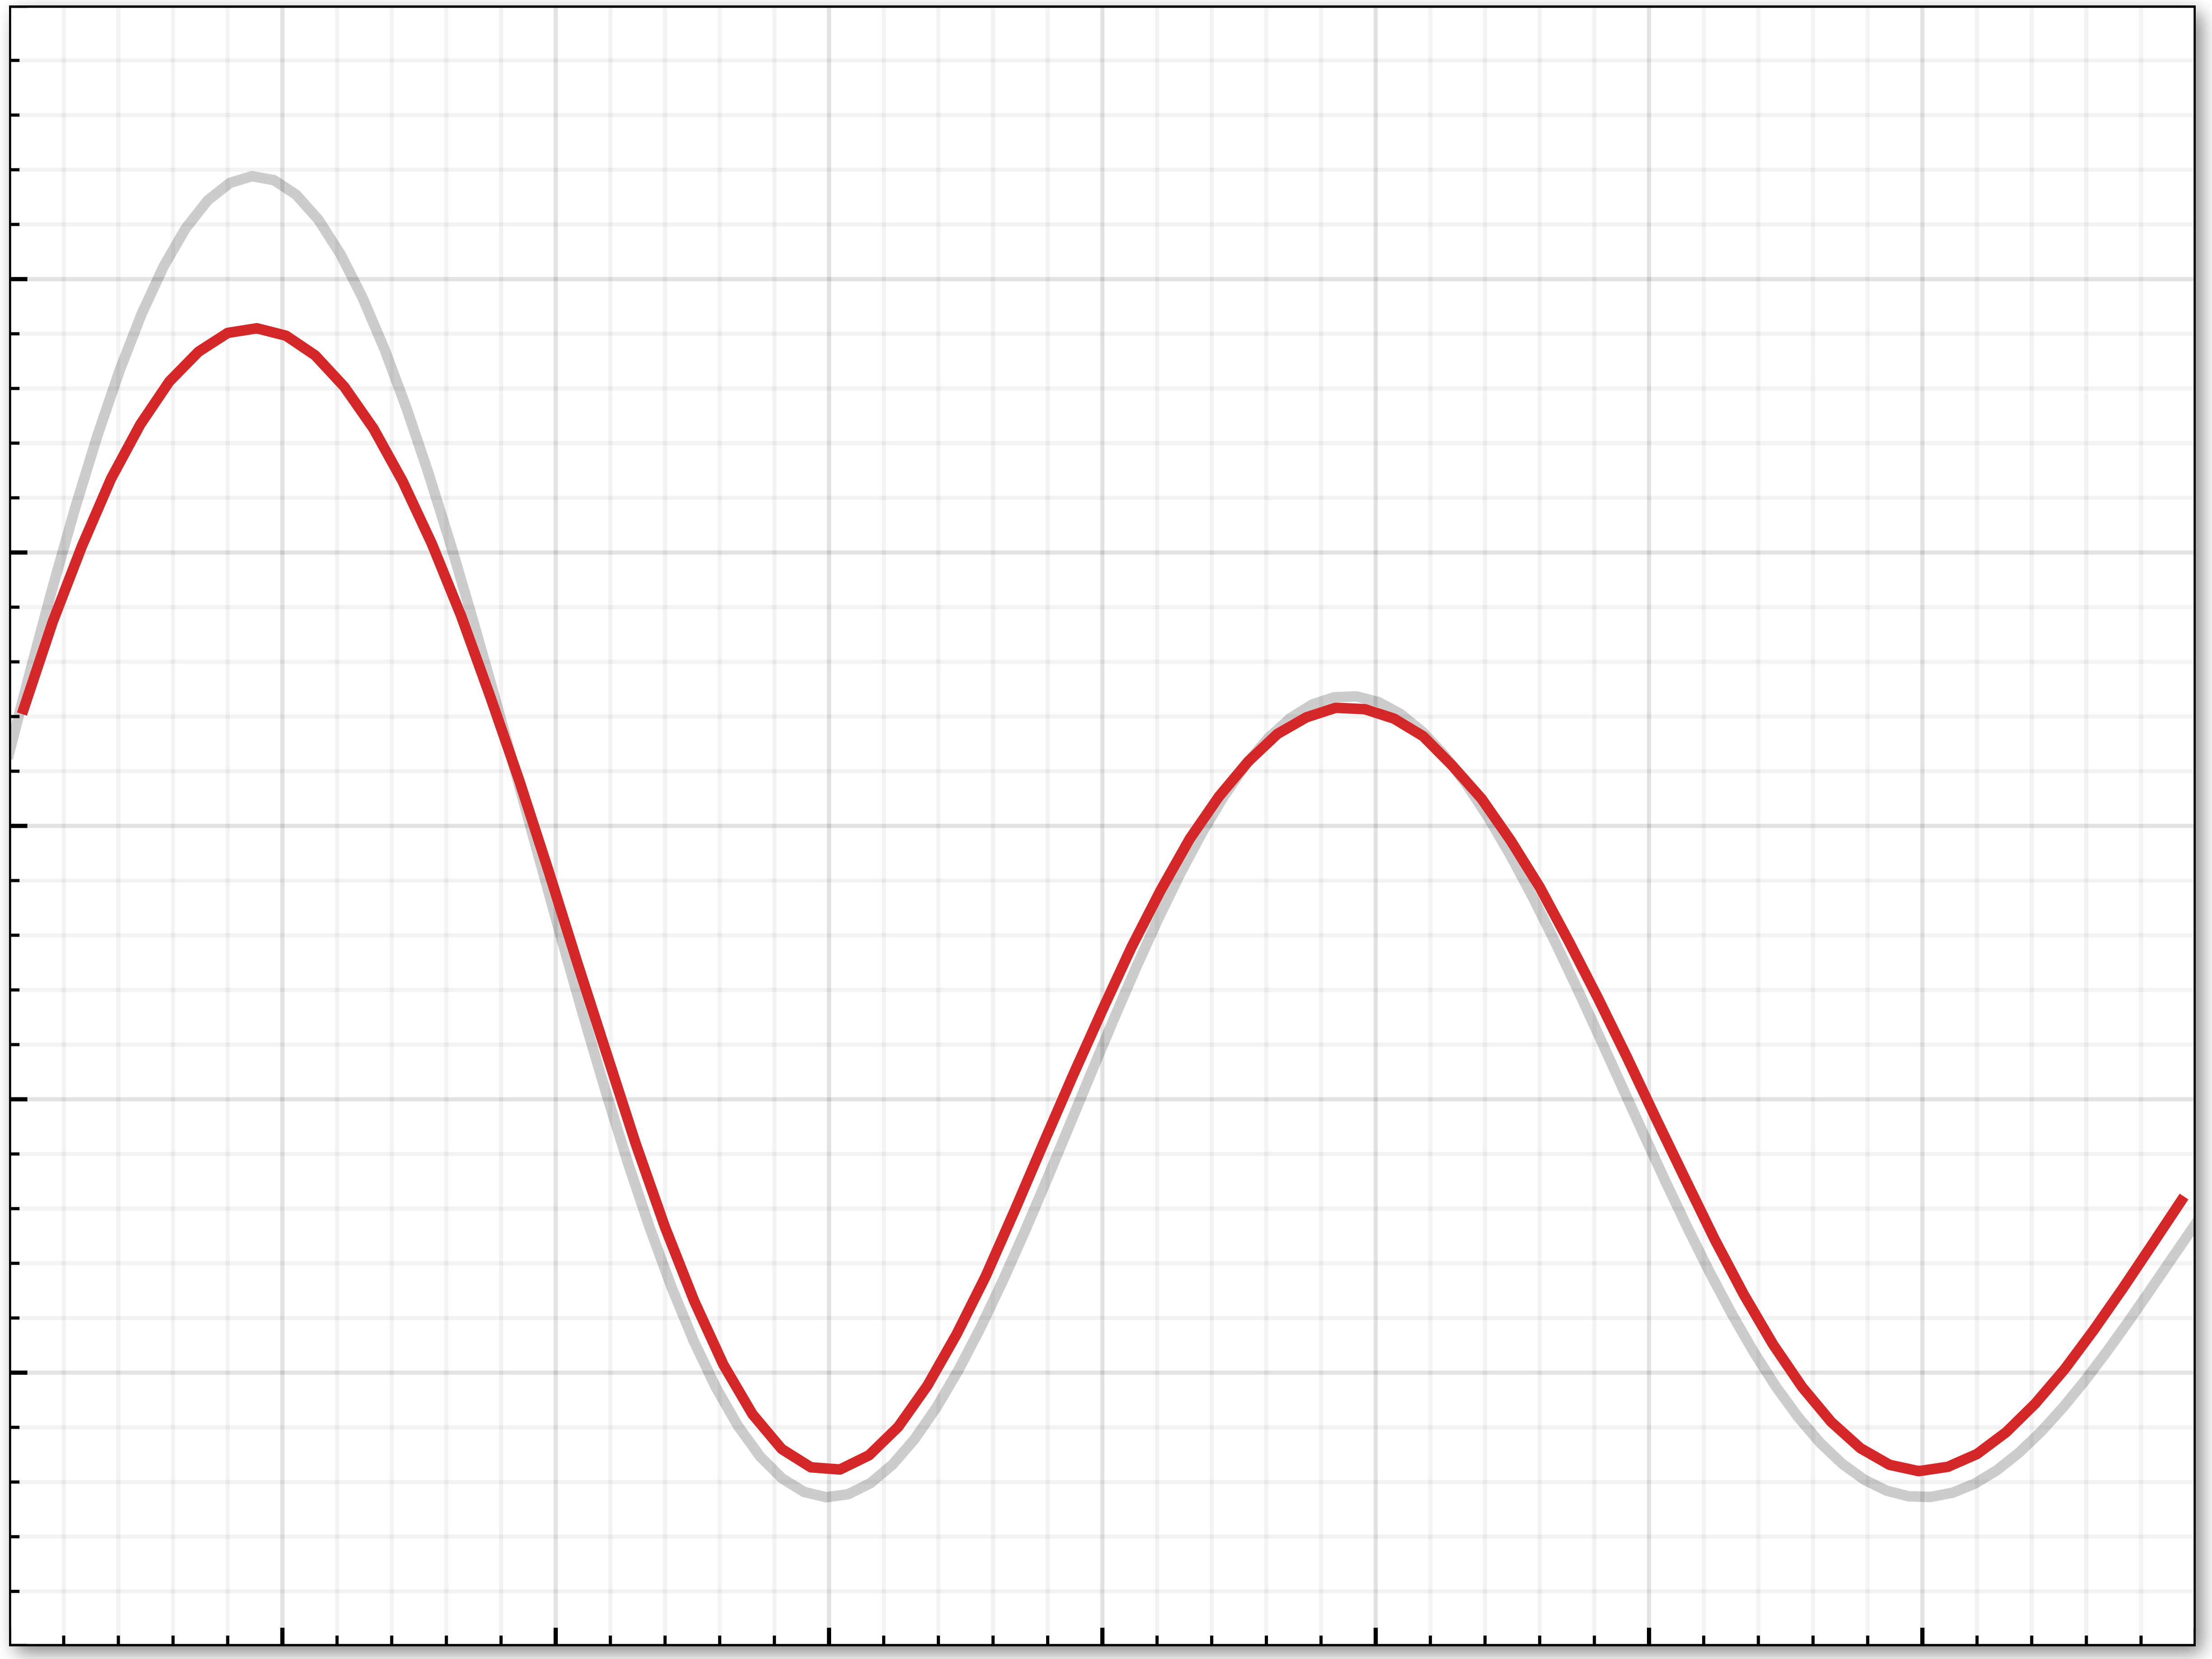

$M = 8$

Math: resampled importance sampling

density of
the final sample

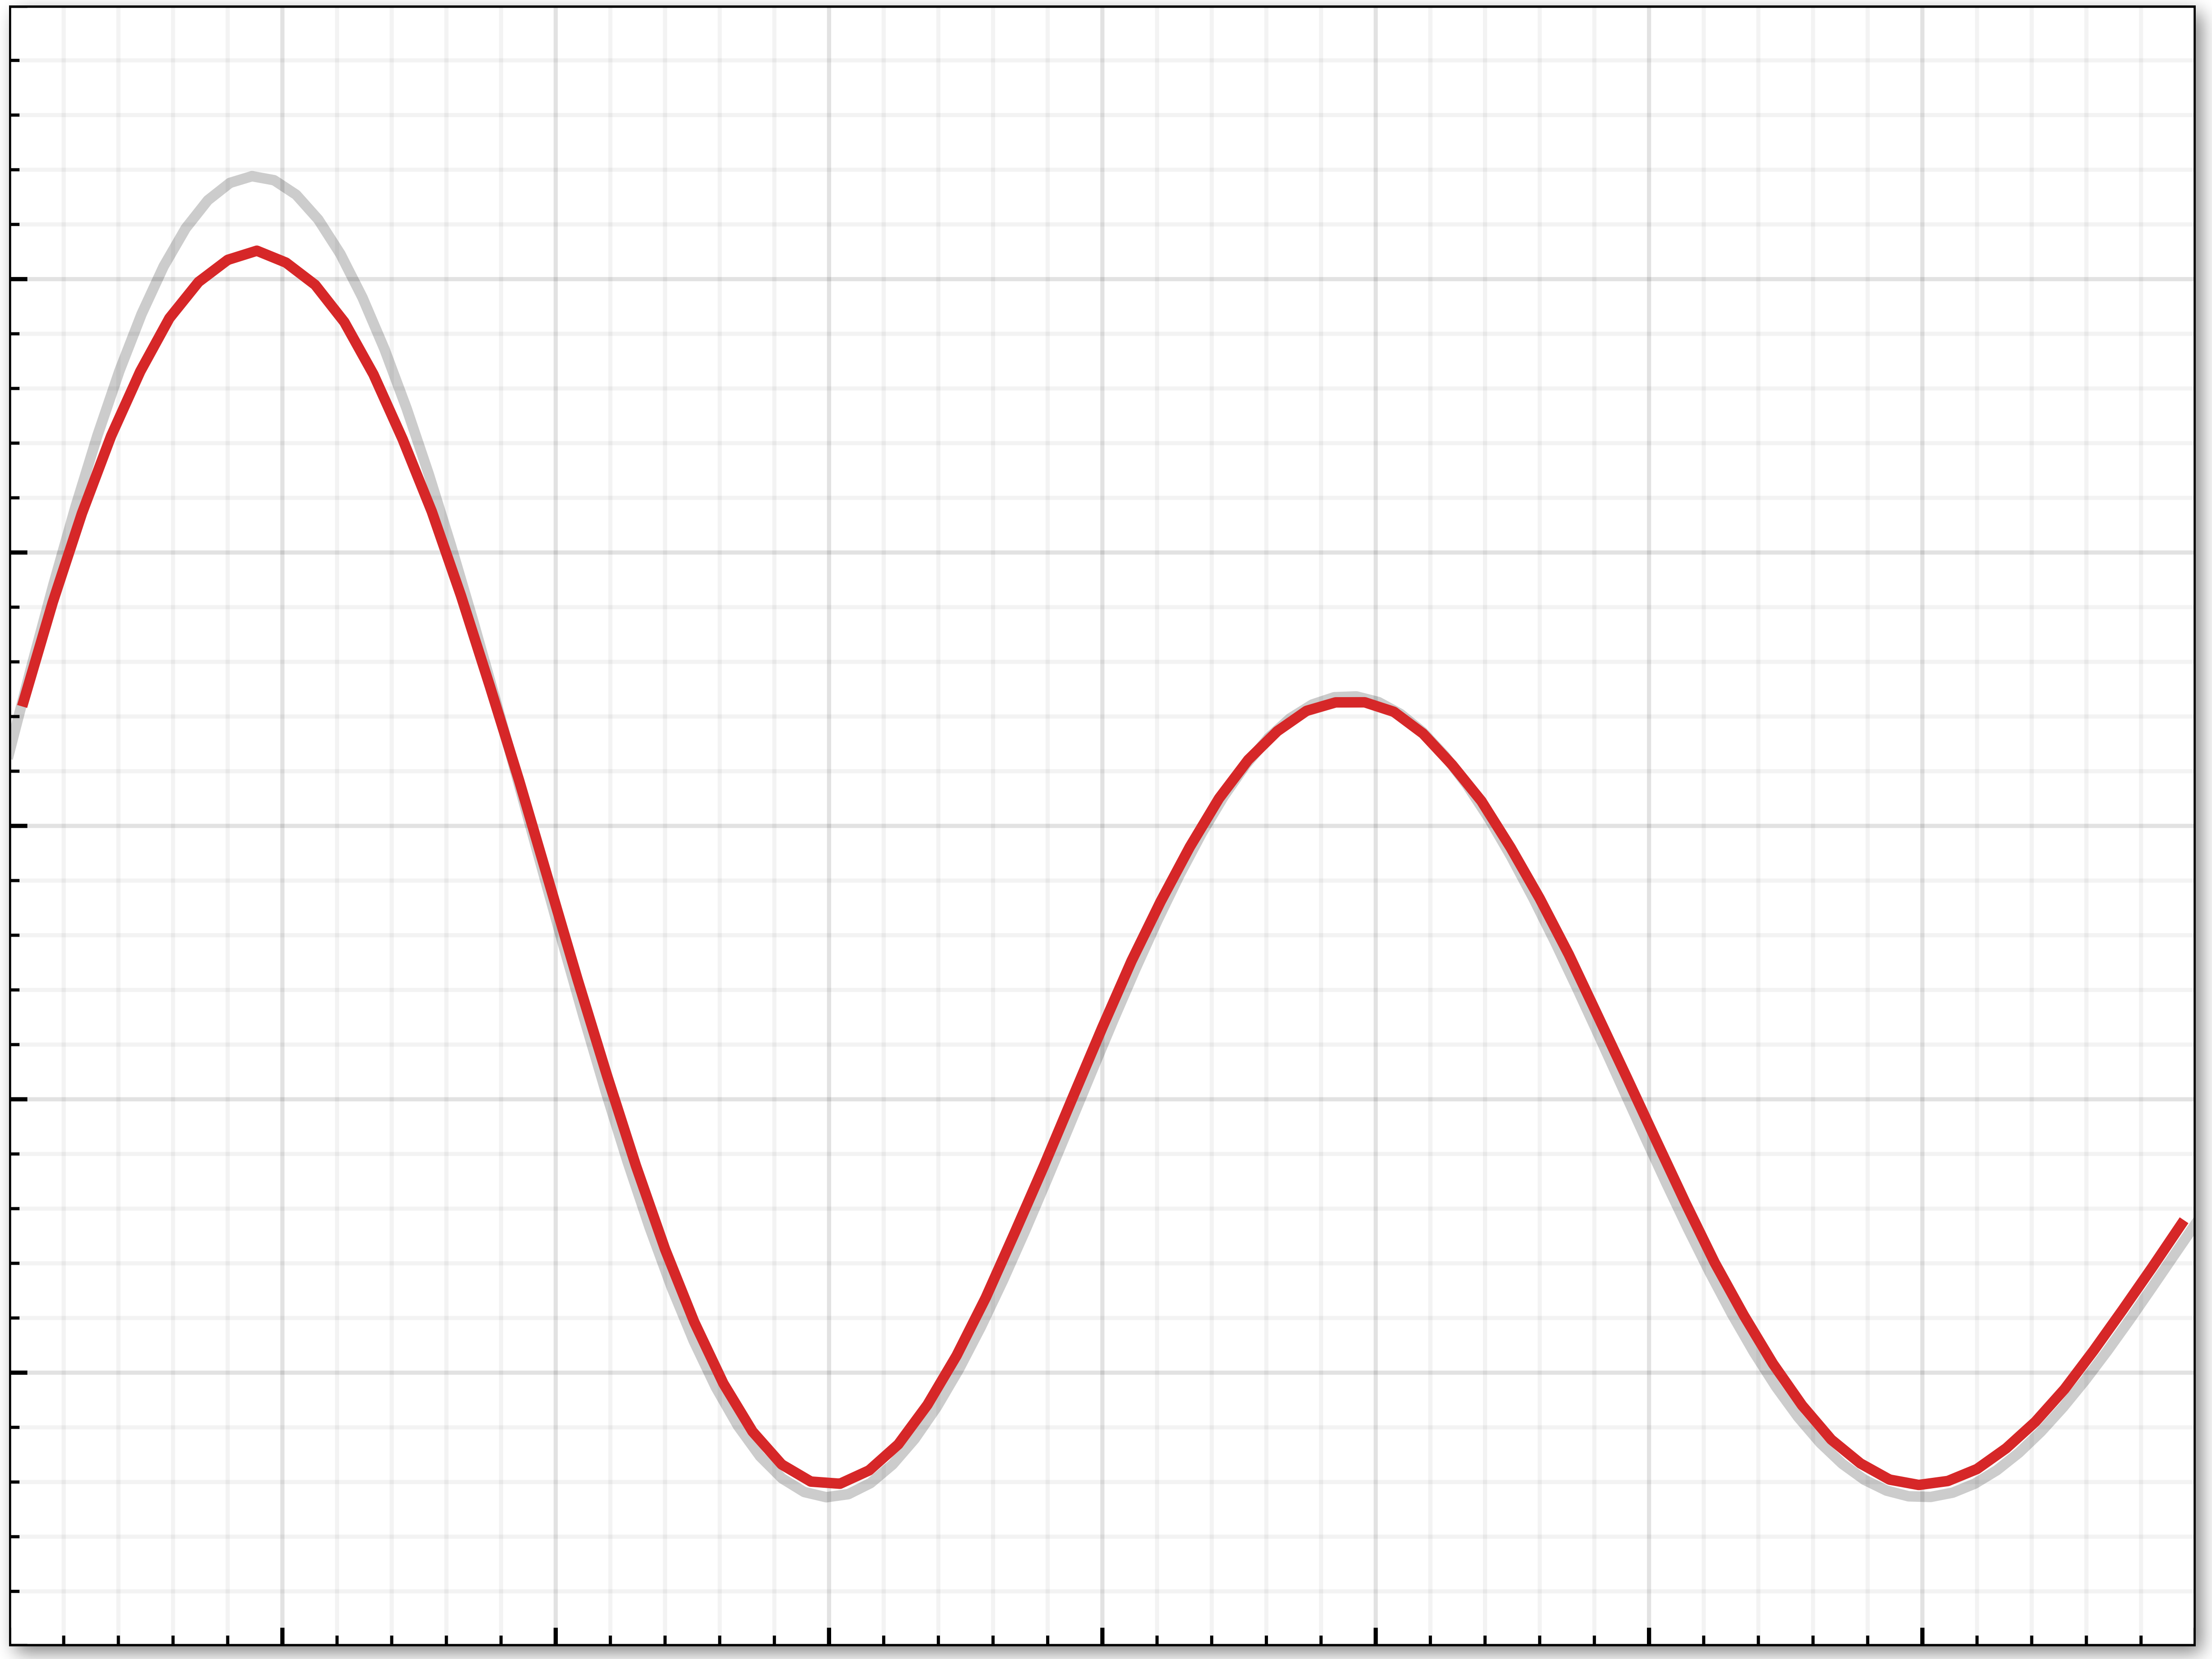

$M = 16$

Math: resampled importance sampling

approximate integrals using RIS samples y_j with unnormalized target $\hat{p}$

$$\int f(x)\mathrm{d}x \approx \frac{1}{N}\sum_{j}^{N} w_j \frac{f\left(y_j\right)}{\hat{p}\left(y_j\right)}$$

quiz: why do we need w_j?

Math: resampled importance sampling

approximate integrals using RIS samples y_j with unnormalized target $\hat{p}$

$$\int f(x)\mathrm{d}x \approx \frac{1}{N}\sum_{j}^{N} w_j \frac{f\left(y_j\right)}{\hat{p}\left(y_j\right)}$$

w_j is an unbiased approximation of the normalization factor of $\hat{p}$

$$w_j = \frac{1}{M}\sum_{i}^{M} \frac{\hat{p}(x_i)}{p(x_i)}$$

Math: resampled importance sampling

approximate integrals using RIS samples y_j with unnormalized target $\hat{p}$

$$\int f(x)\mathrm{d}x \approx \frac{1}{N}\sum_{j}^{N} W_j f\left(y_j\right) \qquad W_j = \frac{1}{M}\frac{1}{\hat{p}(y_j)}\sum_{i}^{M}\frac{\hat{p}(x_i)}{p(x_i)}$$

“unbiased contribution weight” (Lin/Kettunen 2022)
“properly weighted samples” (Liu 2001)

Example: resampled direct lighting

$$p \propto L$$

$$\hat{p} = L \cdot \rho \cdot G$$

Resampled importance sampling can be slow

need to build an array and compute CDF

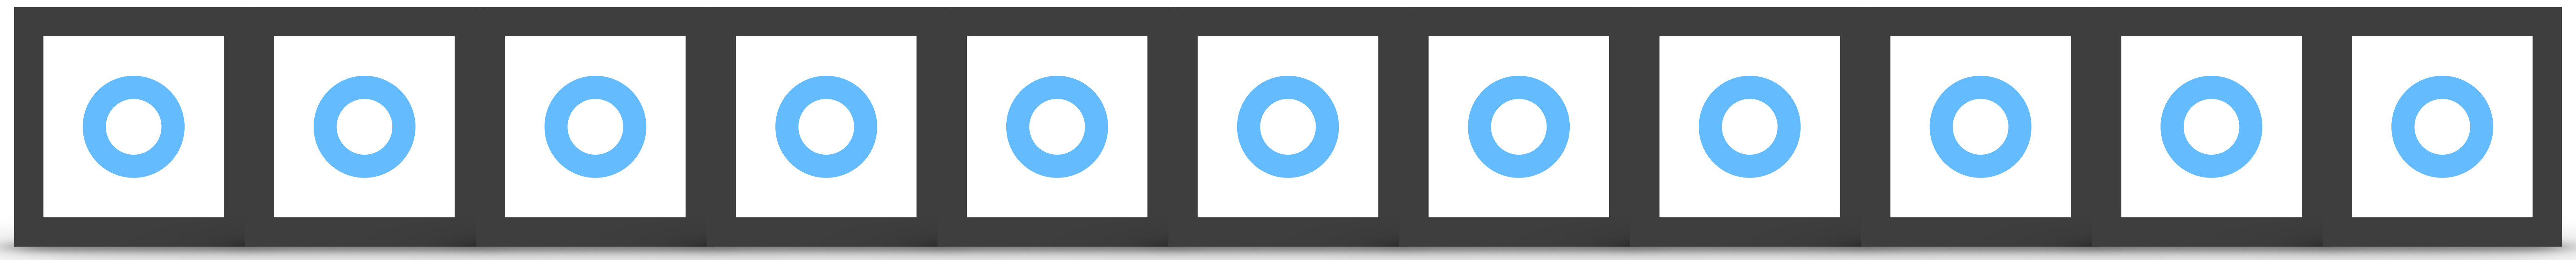

$$\frac{\hat{p}}{p} = \quad 5 \quad 8 \quad 0 \quad 2 \quad 3 \quad 5 \quad 20 \quad 4 \quad 6 \quad 1 \quad 7$$

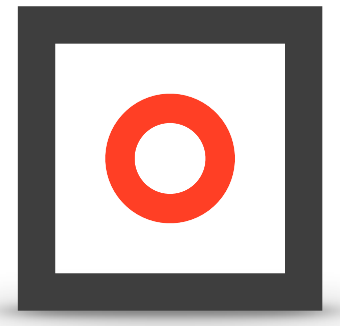

pick a sample with prob. proportional to $\frac{\hat{p}}{p}$

Solution: reservoir sampling

idea: streaming through the samples using rejection sampling

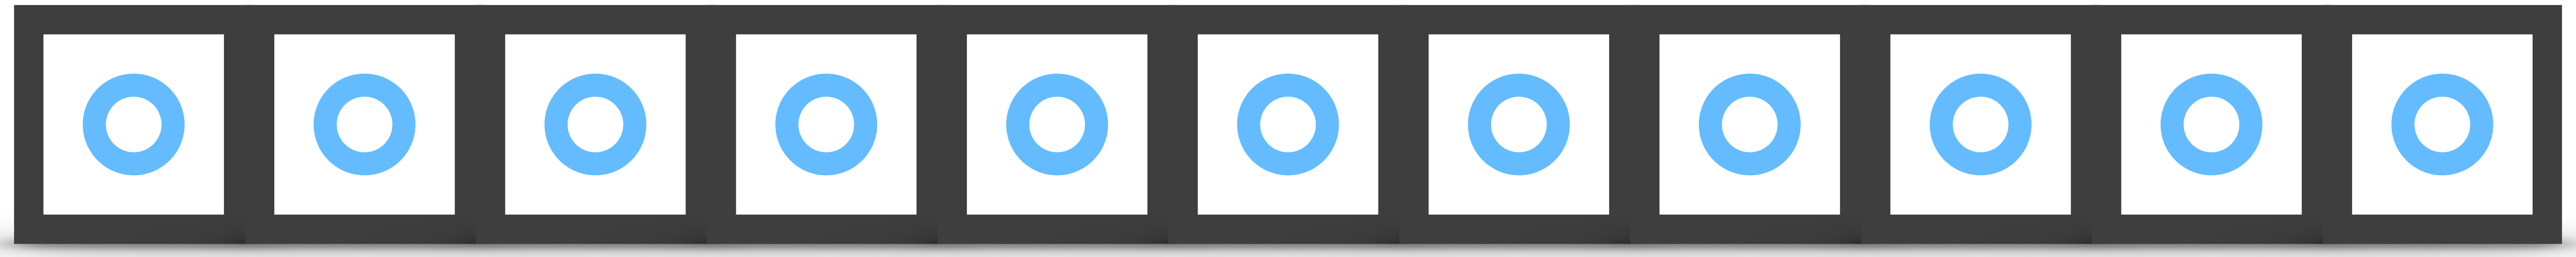

A general purpose unequal probability sampling plan

By M. T. CHAO

Institute of Statistics, Academia Sinica, Taipei, Taiwan

Solution: reservoir sampling

idea: streaming through the samples using rejection sampling

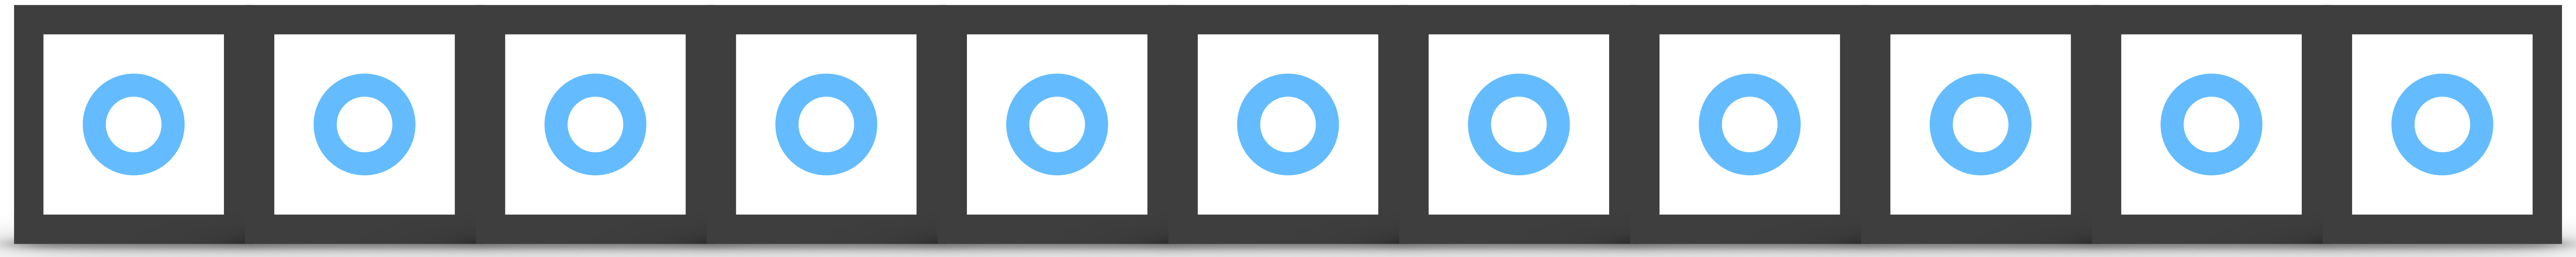

$y = ?$

$w_{\text{sum}} = 0$

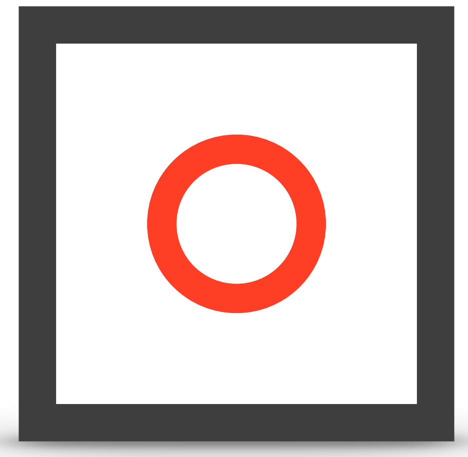

A general purpose unequal probability sampling plan

By M. T. CHAO

Institute of Statistics, Academia Sinica, Taipei, Taiwan

Solution: reservoir sampling

idea: streaming through the samples using rejection sampling

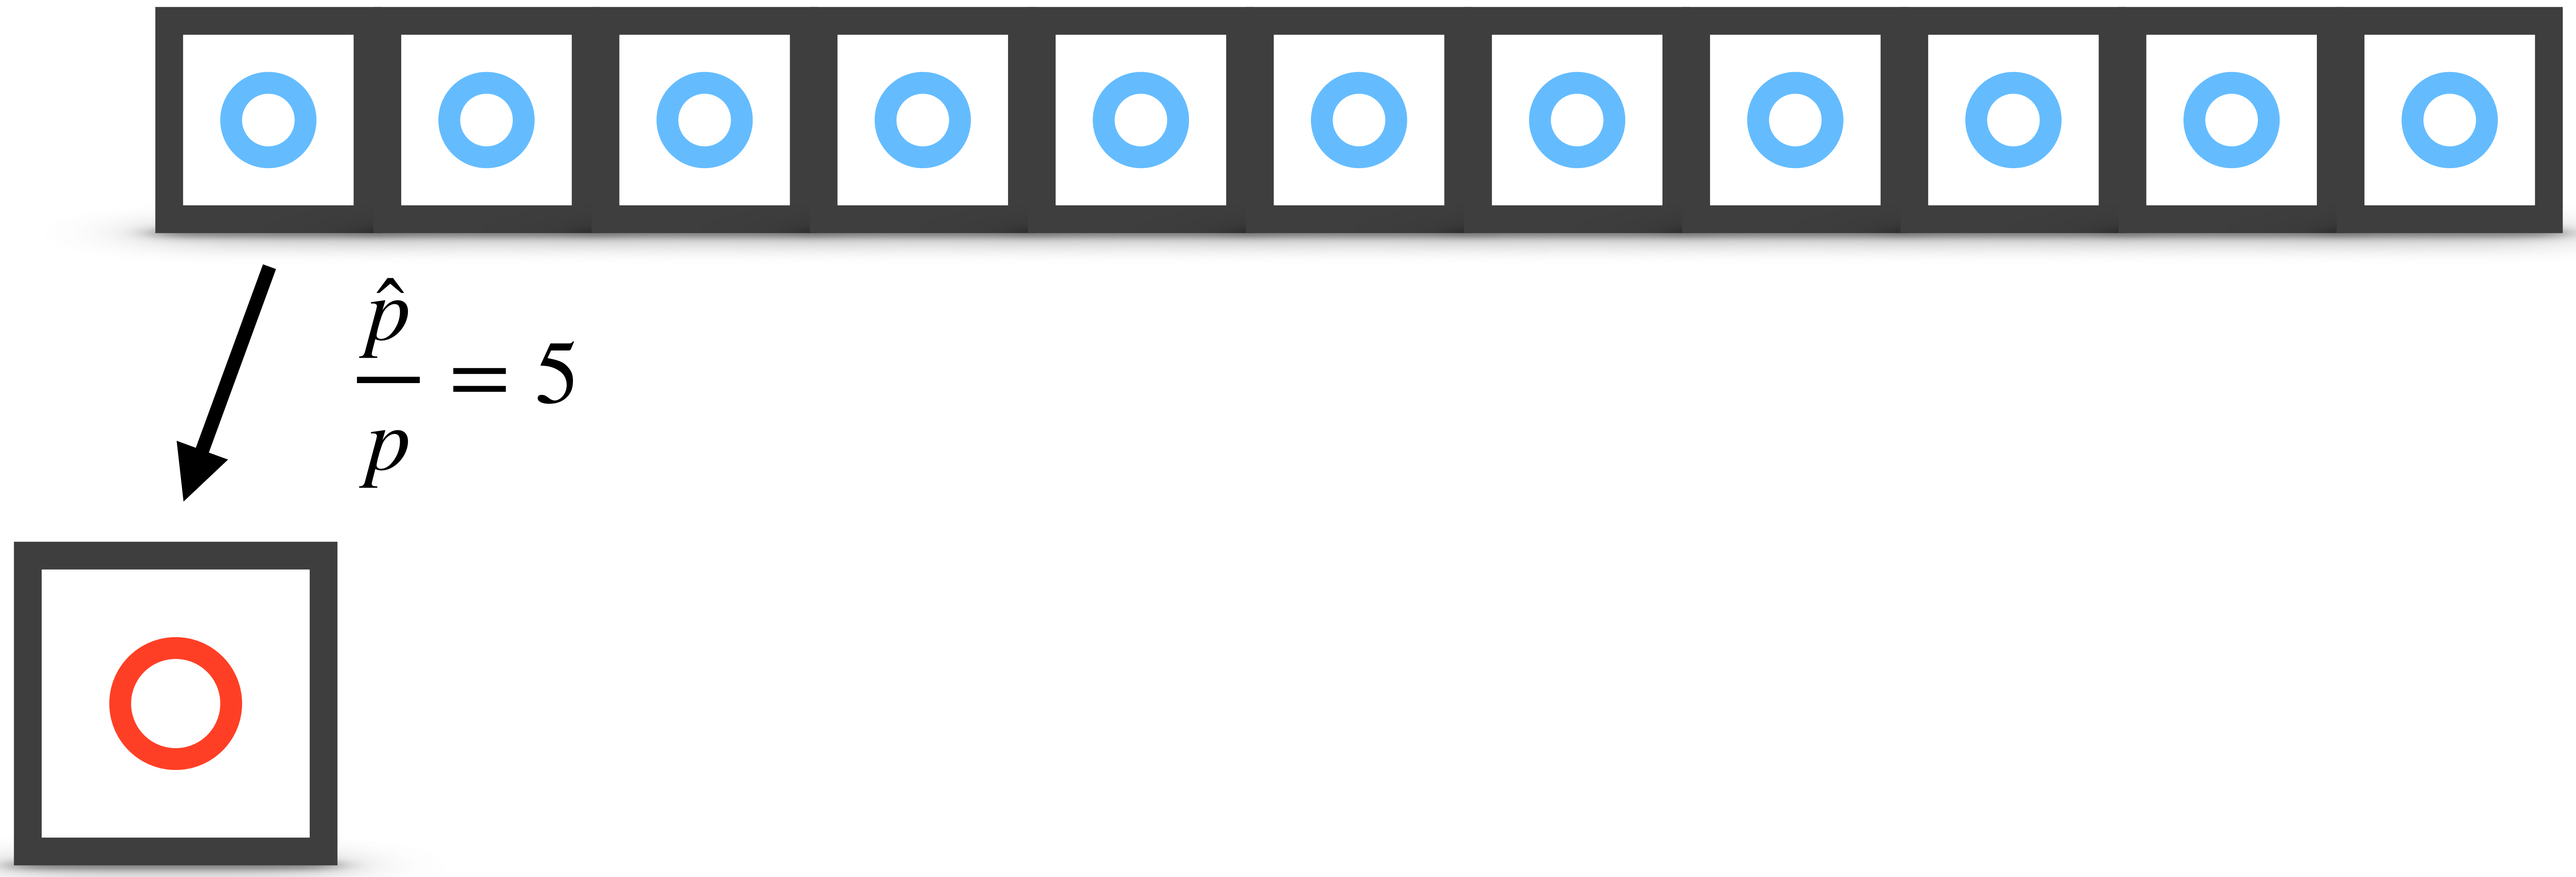

$\frac{\hat{p}}{p} = 5$

$y = 0$

$w_{\text{sum}} = 5$

A general purpose unequal probability sampling plan

By M. T. CHAO
Institute of Statistics, Academia Sinica, Taipei, Taiwan

Solution: reservoir sampling

idea: streaming through the samples using rejection sampling

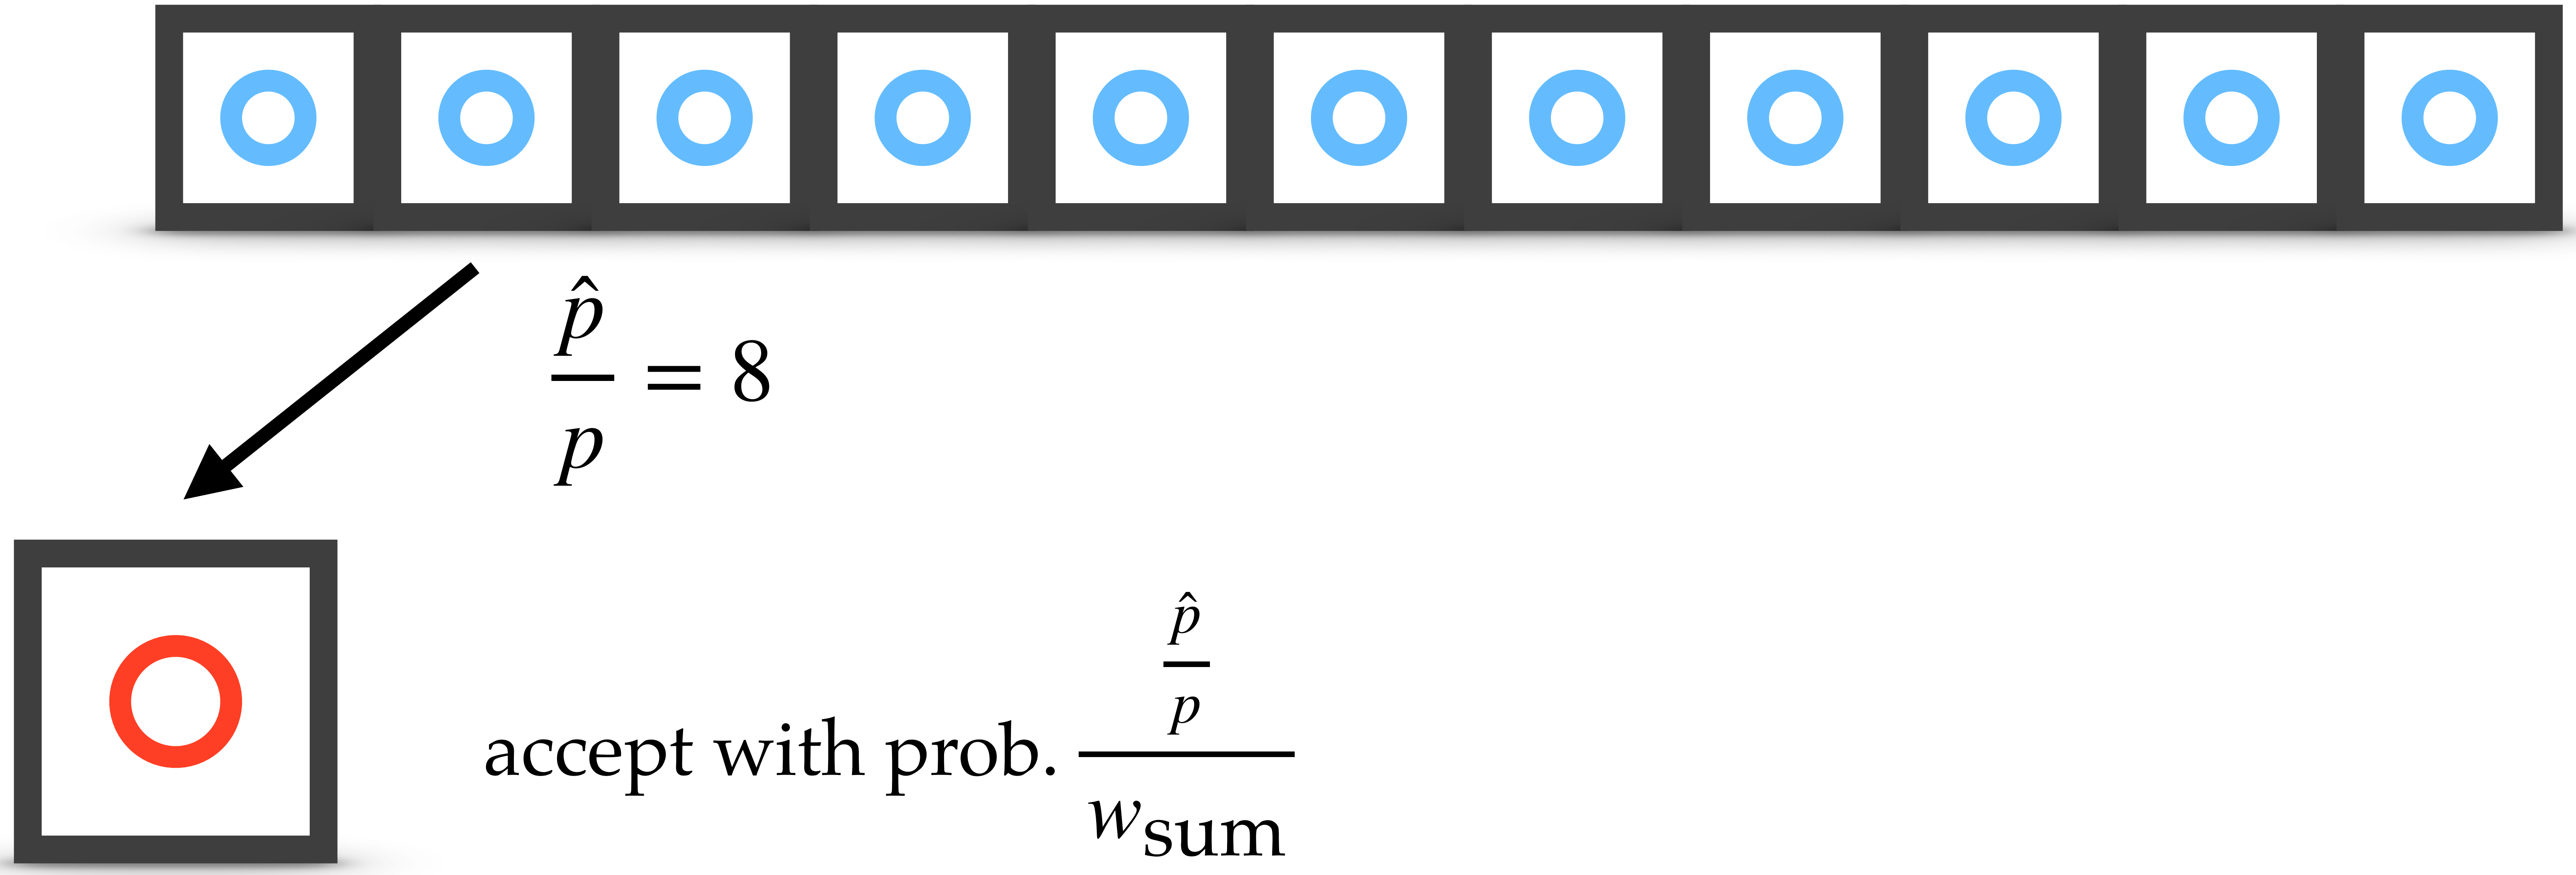

$y = 1$

$w_{\text{sum}} = 13$

accept with prob. $\frac{\hat{p}/p}{w_{\text{sum}}}$

A general purpose unequal probability sampling plan

By M. T. CHAO

Institute of Statistics, Academia Sinica, Taipei, Taiwan

Solution: reservoir sampling

idea: streaming through the samples using rejection sampling

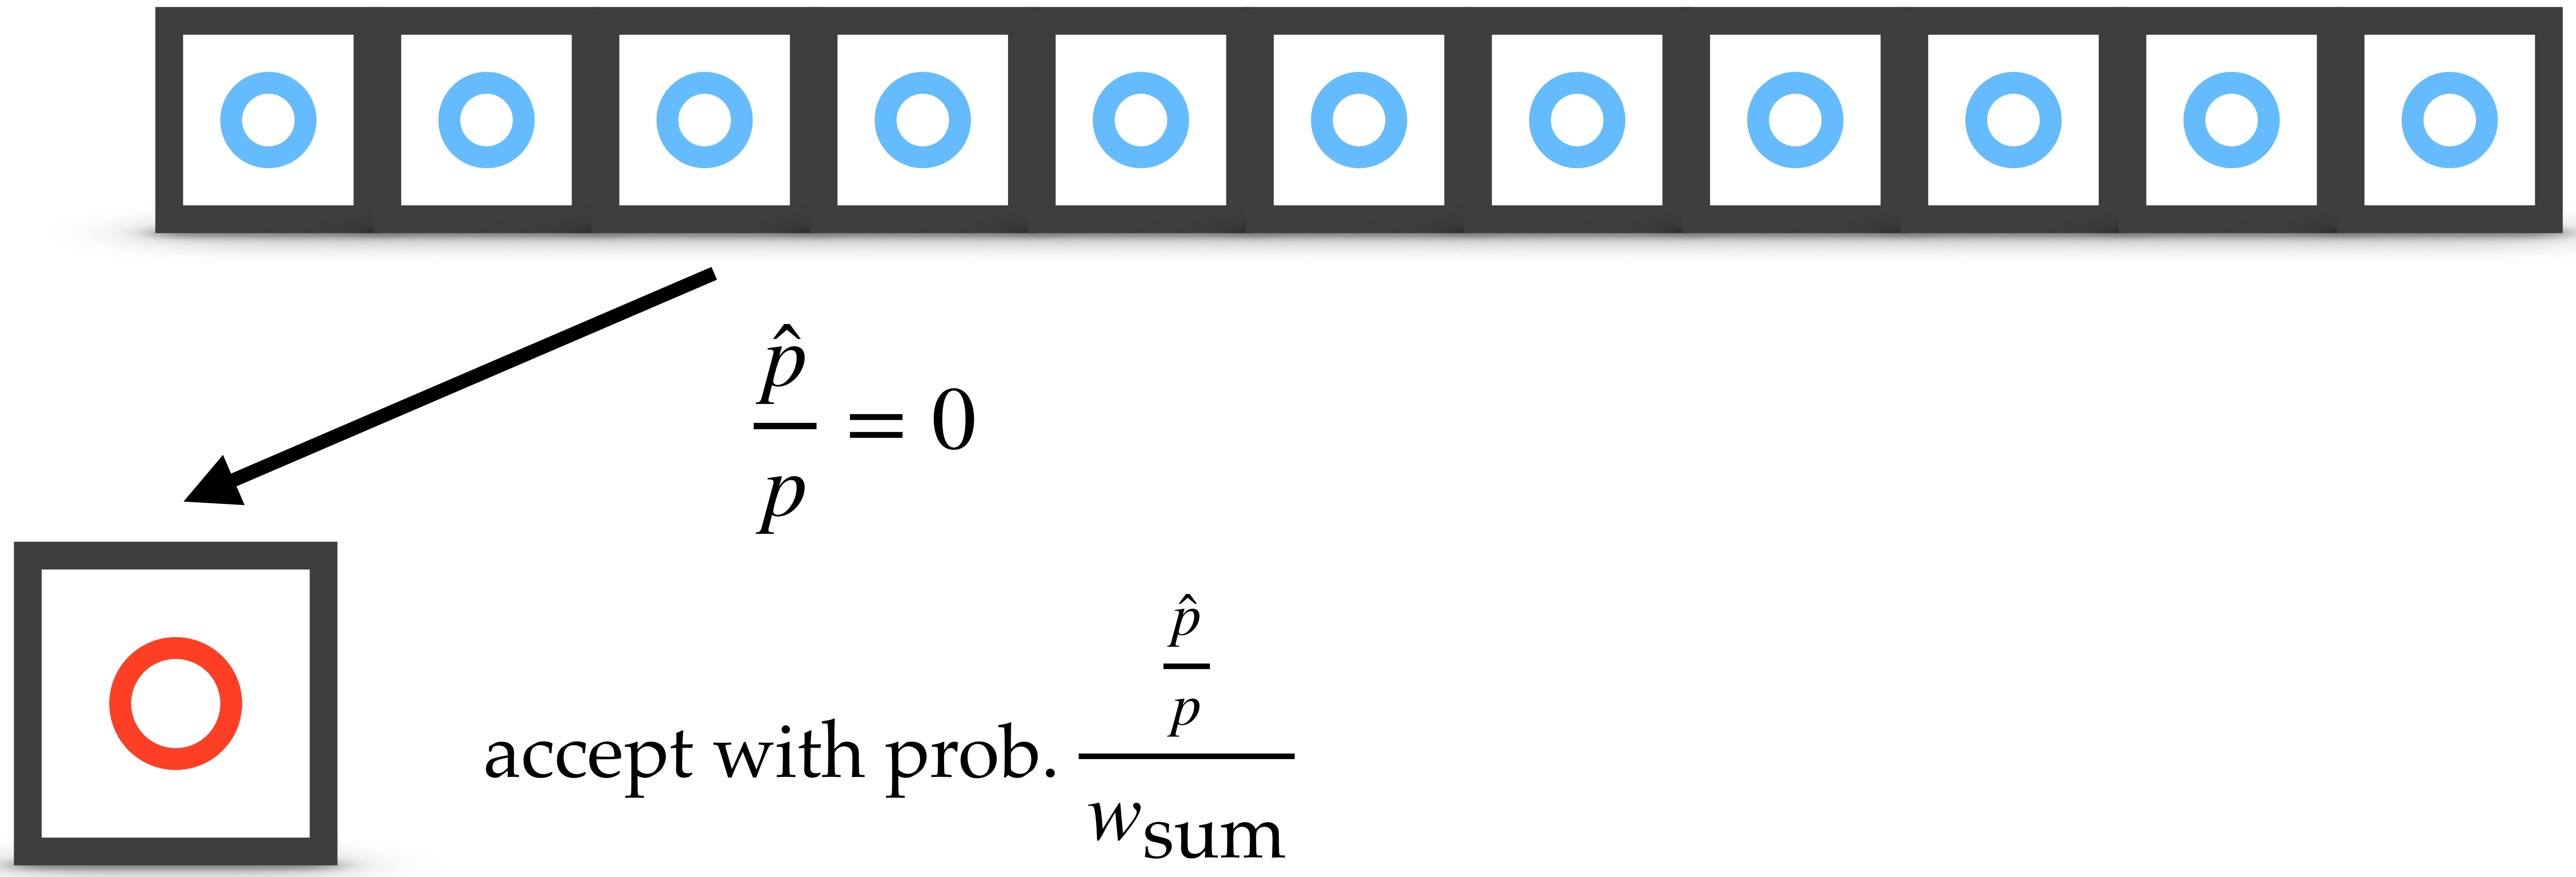

$$\frac{\hat{p}}{p} = 0$$

$y = 1$

$w_{\text{sum}} = 13$

accept with prob. $\dfrac{\hat{p}/p}{w_{\text{sum}}}$

A general purpose unequal probability sampling plan

By M. T. Chao

Institute of Statistics, Academia Sinica, Taipei, Taiwan

Reservoir sampling vs inverse transform sampling

reservoir sampling	**inverse transform sampling**
constant memory usage	O(M) memory usage
no precomputation	O(M) pre computation
O(M) computation per query	O(log(M)) computation per query
no stratification	can be stratified

Applying (reservoir) RIS to ReSTIR

recall: we want to reuse neighbor pixels' sampling results

each pixel stores a "reservoir" which is the result from RIS

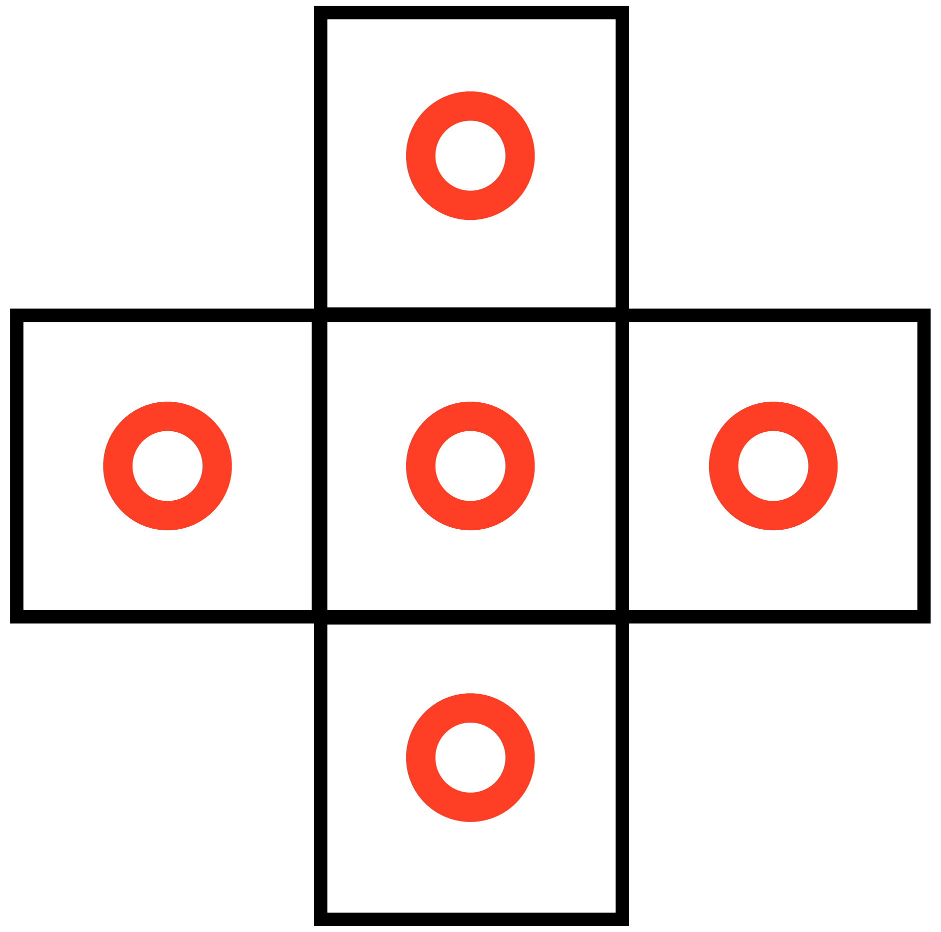

Applying (reservoir) RIS to ReSTIR

recall: we want to reuse neighbor pixels' sampling results

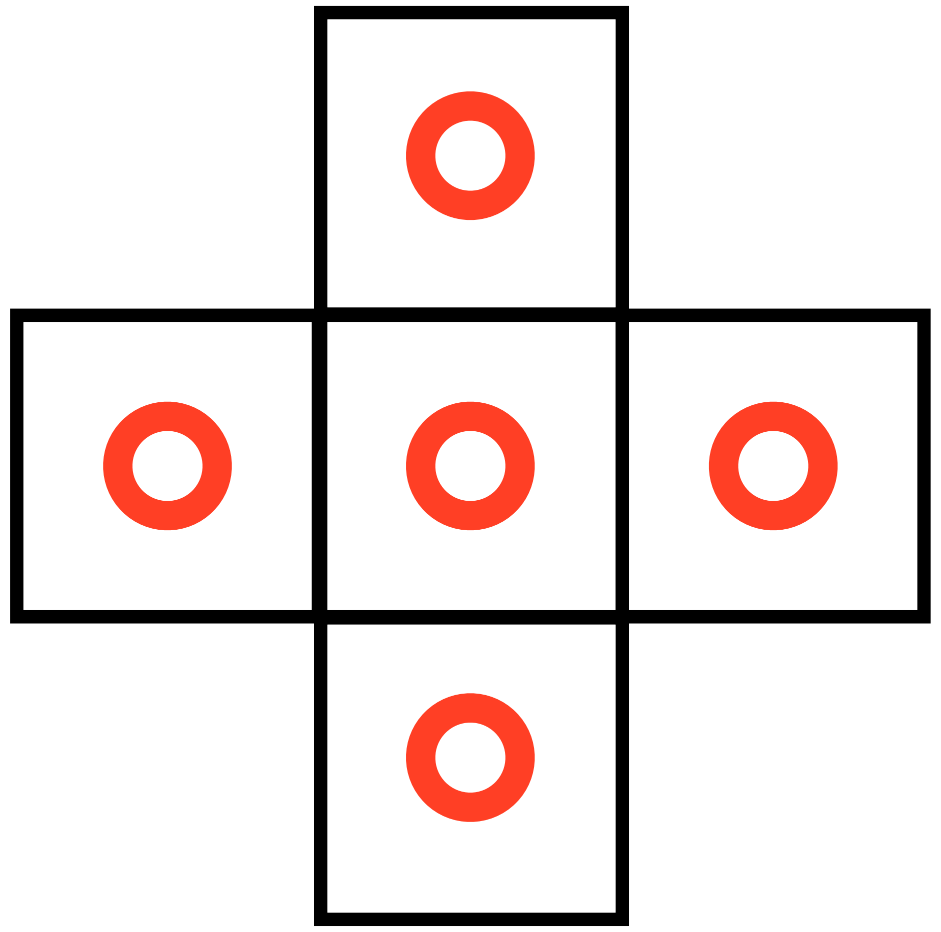

each pixel stores a "reservoir" which is the result from RIS

to reuse, we need to merge the reservoirs from two pixels

Applying (reservoir) RIS to ReSTIR

recall: we want to reuse neighbor pixels' sampling results

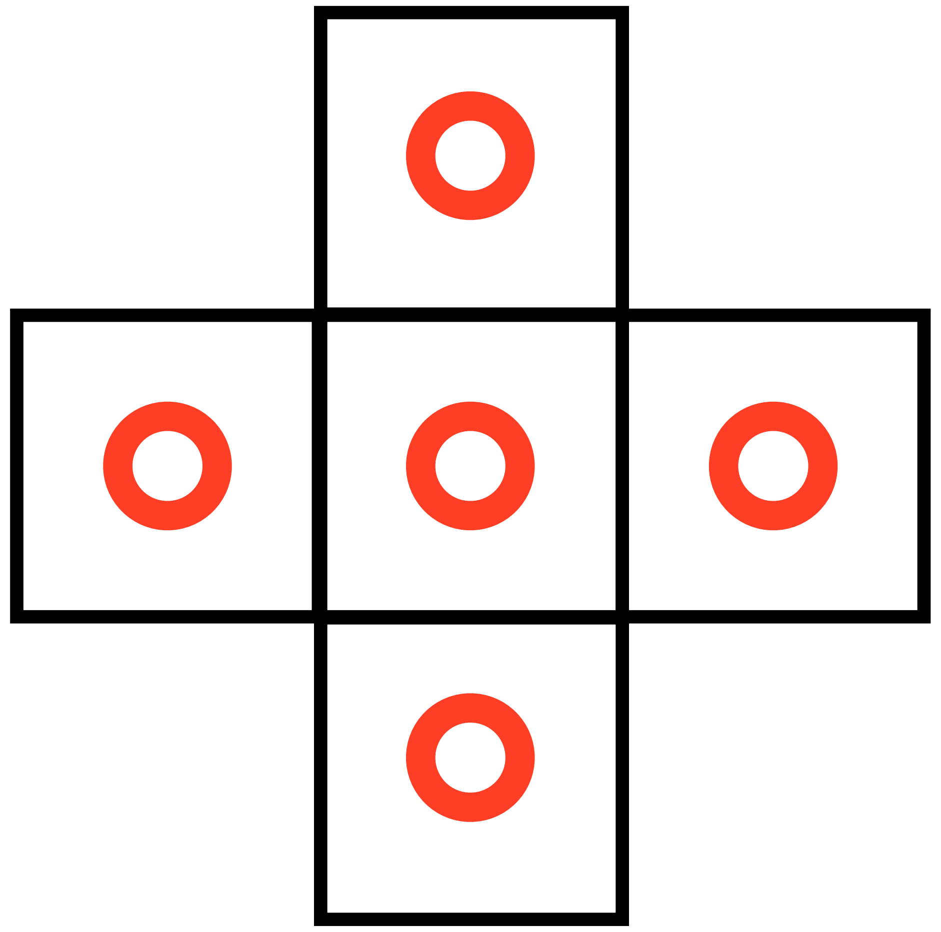

each pixel stores a "reservoir" which is the result from RIS

to reuse, we need to merge the reservoirs from two pixels

to merge, apply RIS to sample from the two reservoirs!

Applying (reservoir) RIS to ReSTIR

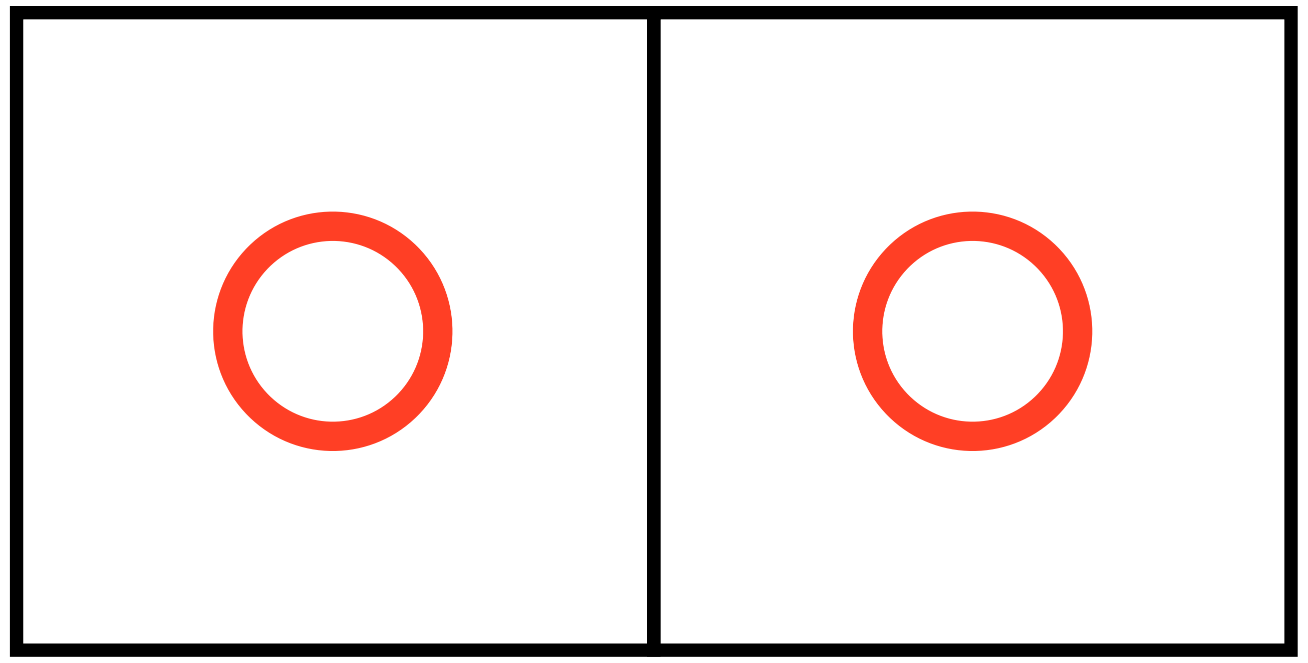

$$W_j = \frac{1}{M}\frac{1}{\hat{p}(y_j)}\sum_{i}^{M}\frac{\hat{p}(x_i)}{p(x_i)}$$

Applying (reservoir) RIS to ReSTIR

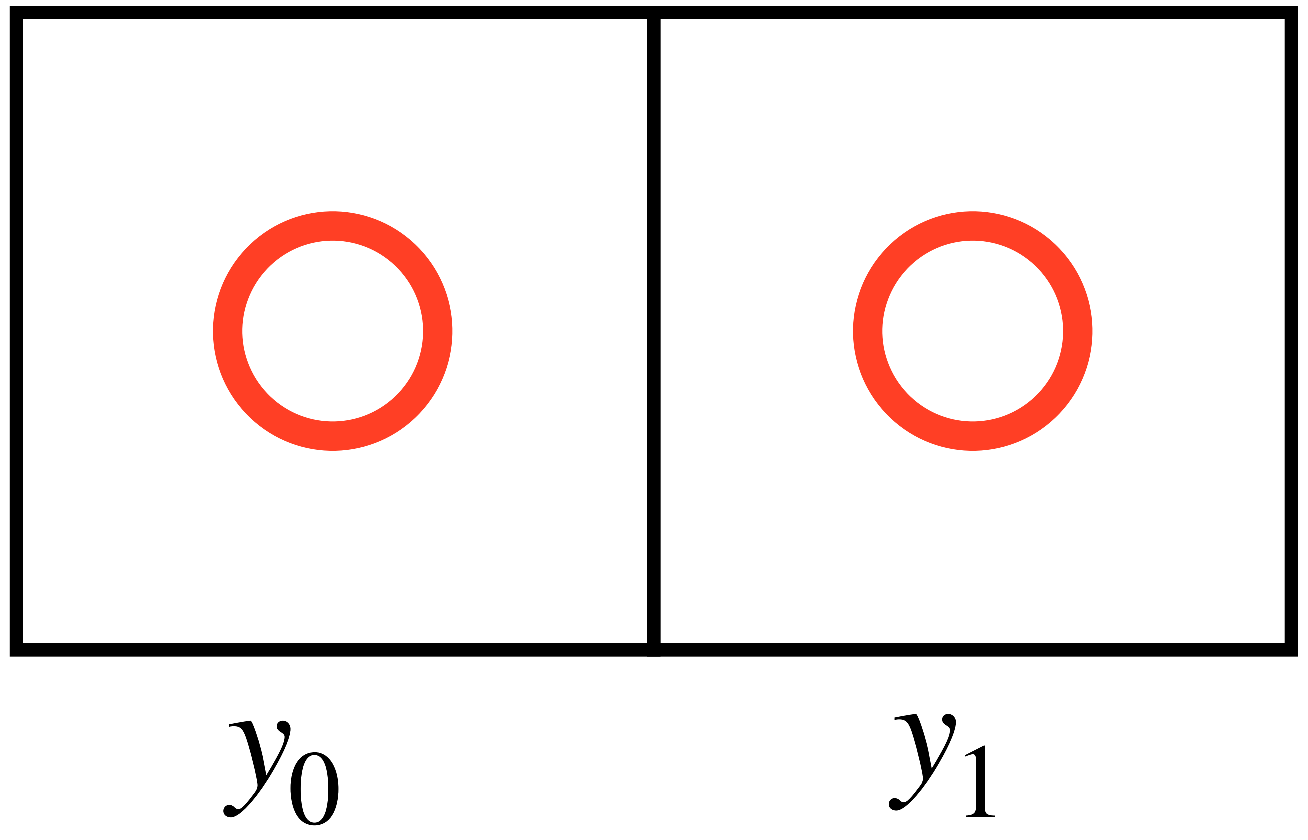

$W_0 = 10$, $M_0 = 5$

$W_1 = 15$, $M_1 = 6$

$$W_j = \frac{1}{M} \frac{1}{\hat{p}(y_j)} \sum_{i}^{M} \frac{\hat{p}(x_i)}{p(x_i)}$$

Applying (reservoir) RIS to ReSTIR

$$W_j = \frac{1}{M} \frac{1}{\hat{p}(y_j)} \sum_i^M \frac{\hat{p}(x_i)}{p(x_i)}$$

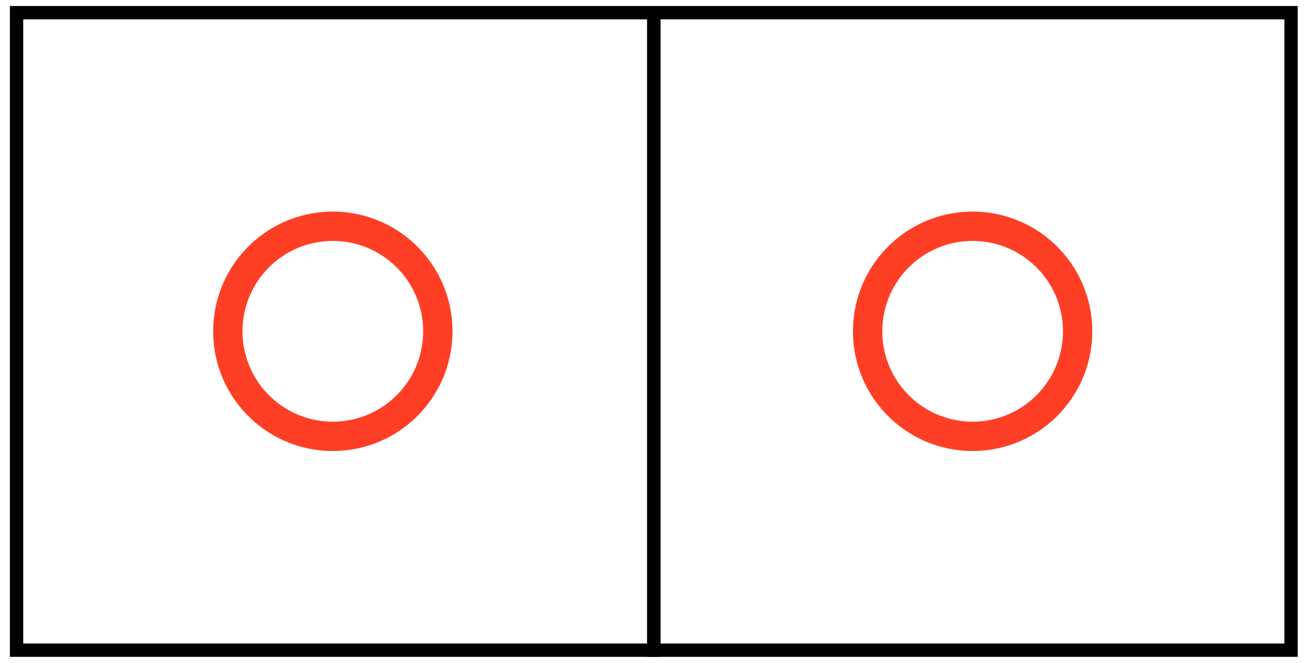

y_0

$W_0 = 10$

$M_0 = 5$

y_1

$W_1 = 15$

$M_1 = 6$

sample with prob. proportional to $W \cdot \hat{p}$

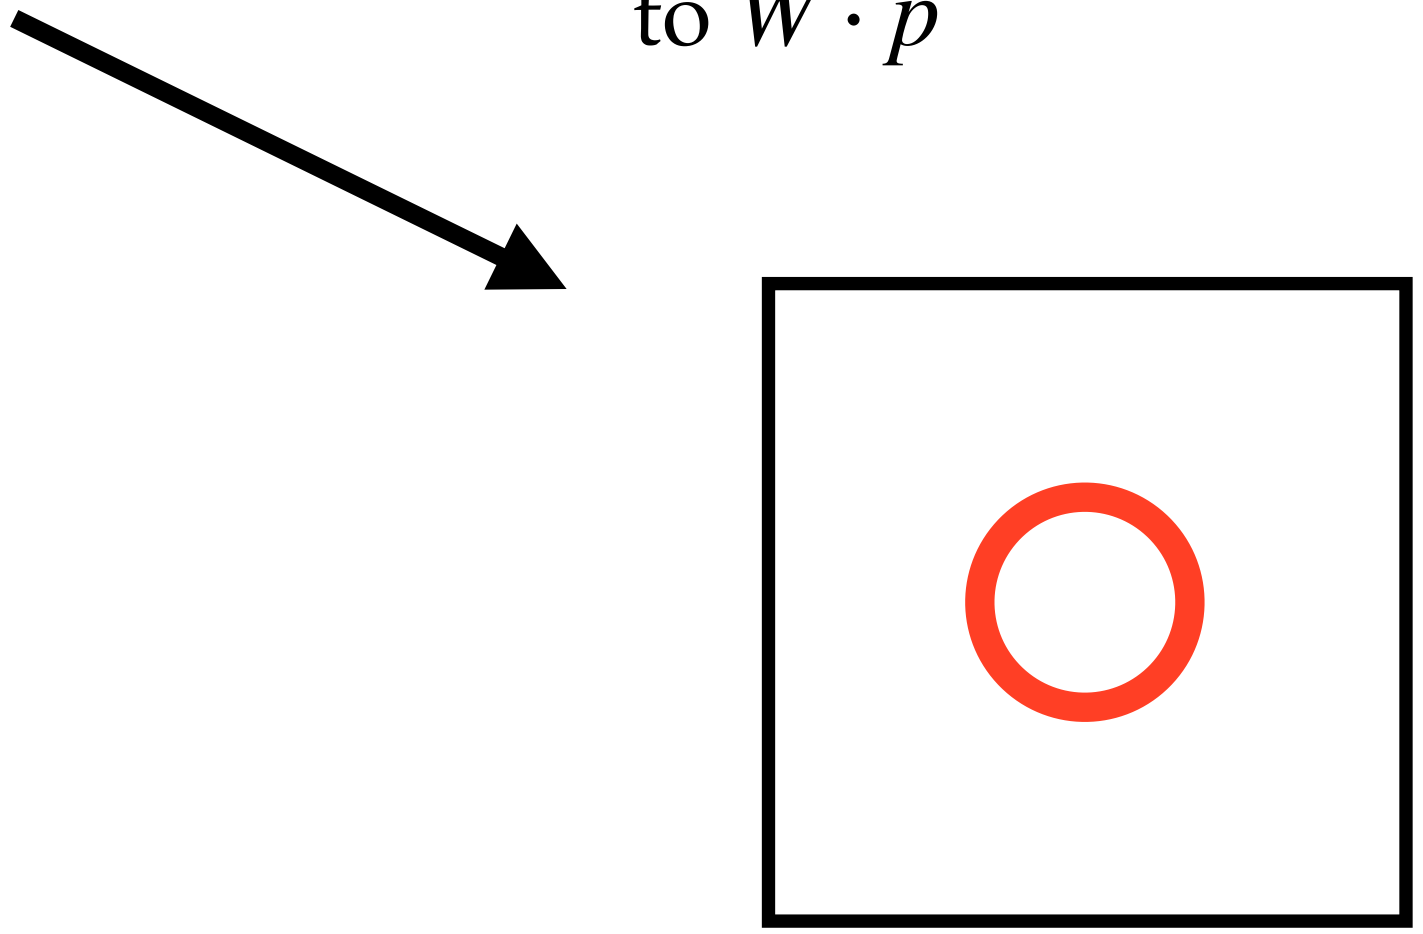

y_1

$$W = \frac{1}{\hat{p}(y_1)} \frac{M_0 W_0 + M_1 W_1}{M_0 + M_1}$$

$$M = M_0 + M_1 = 11$$

The ReSTIR algorithm

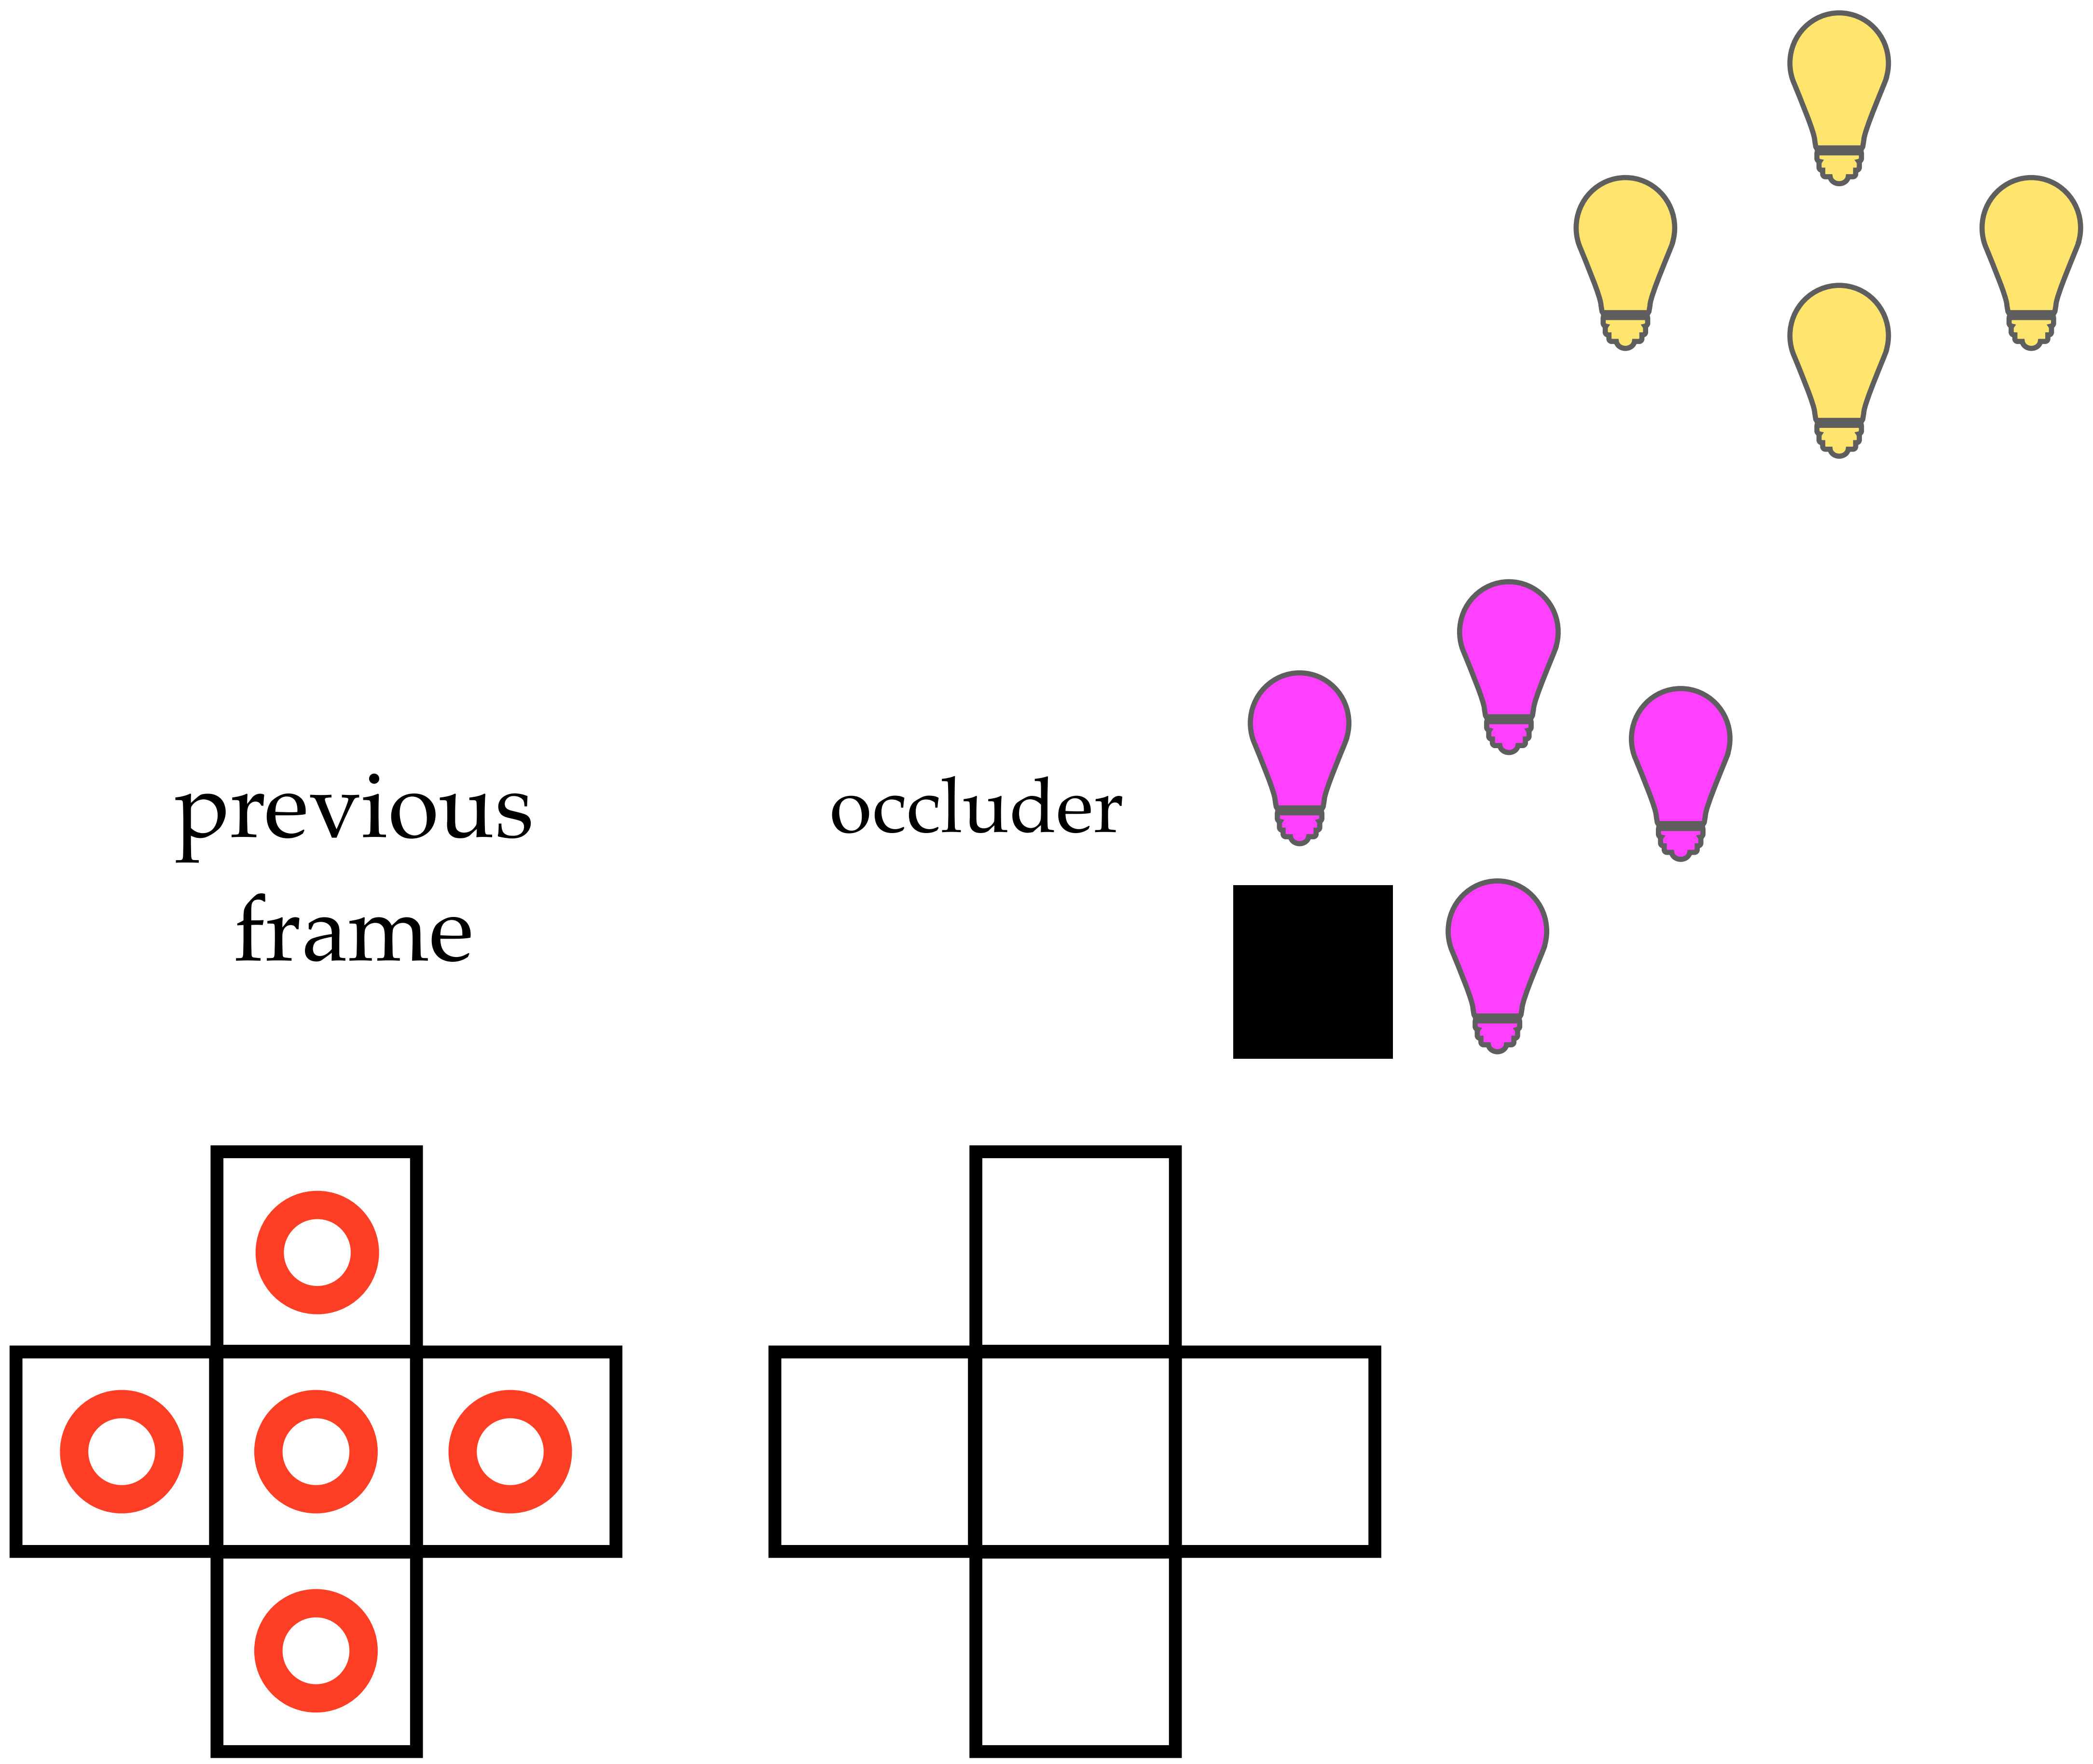

The ReSTIR algorithm

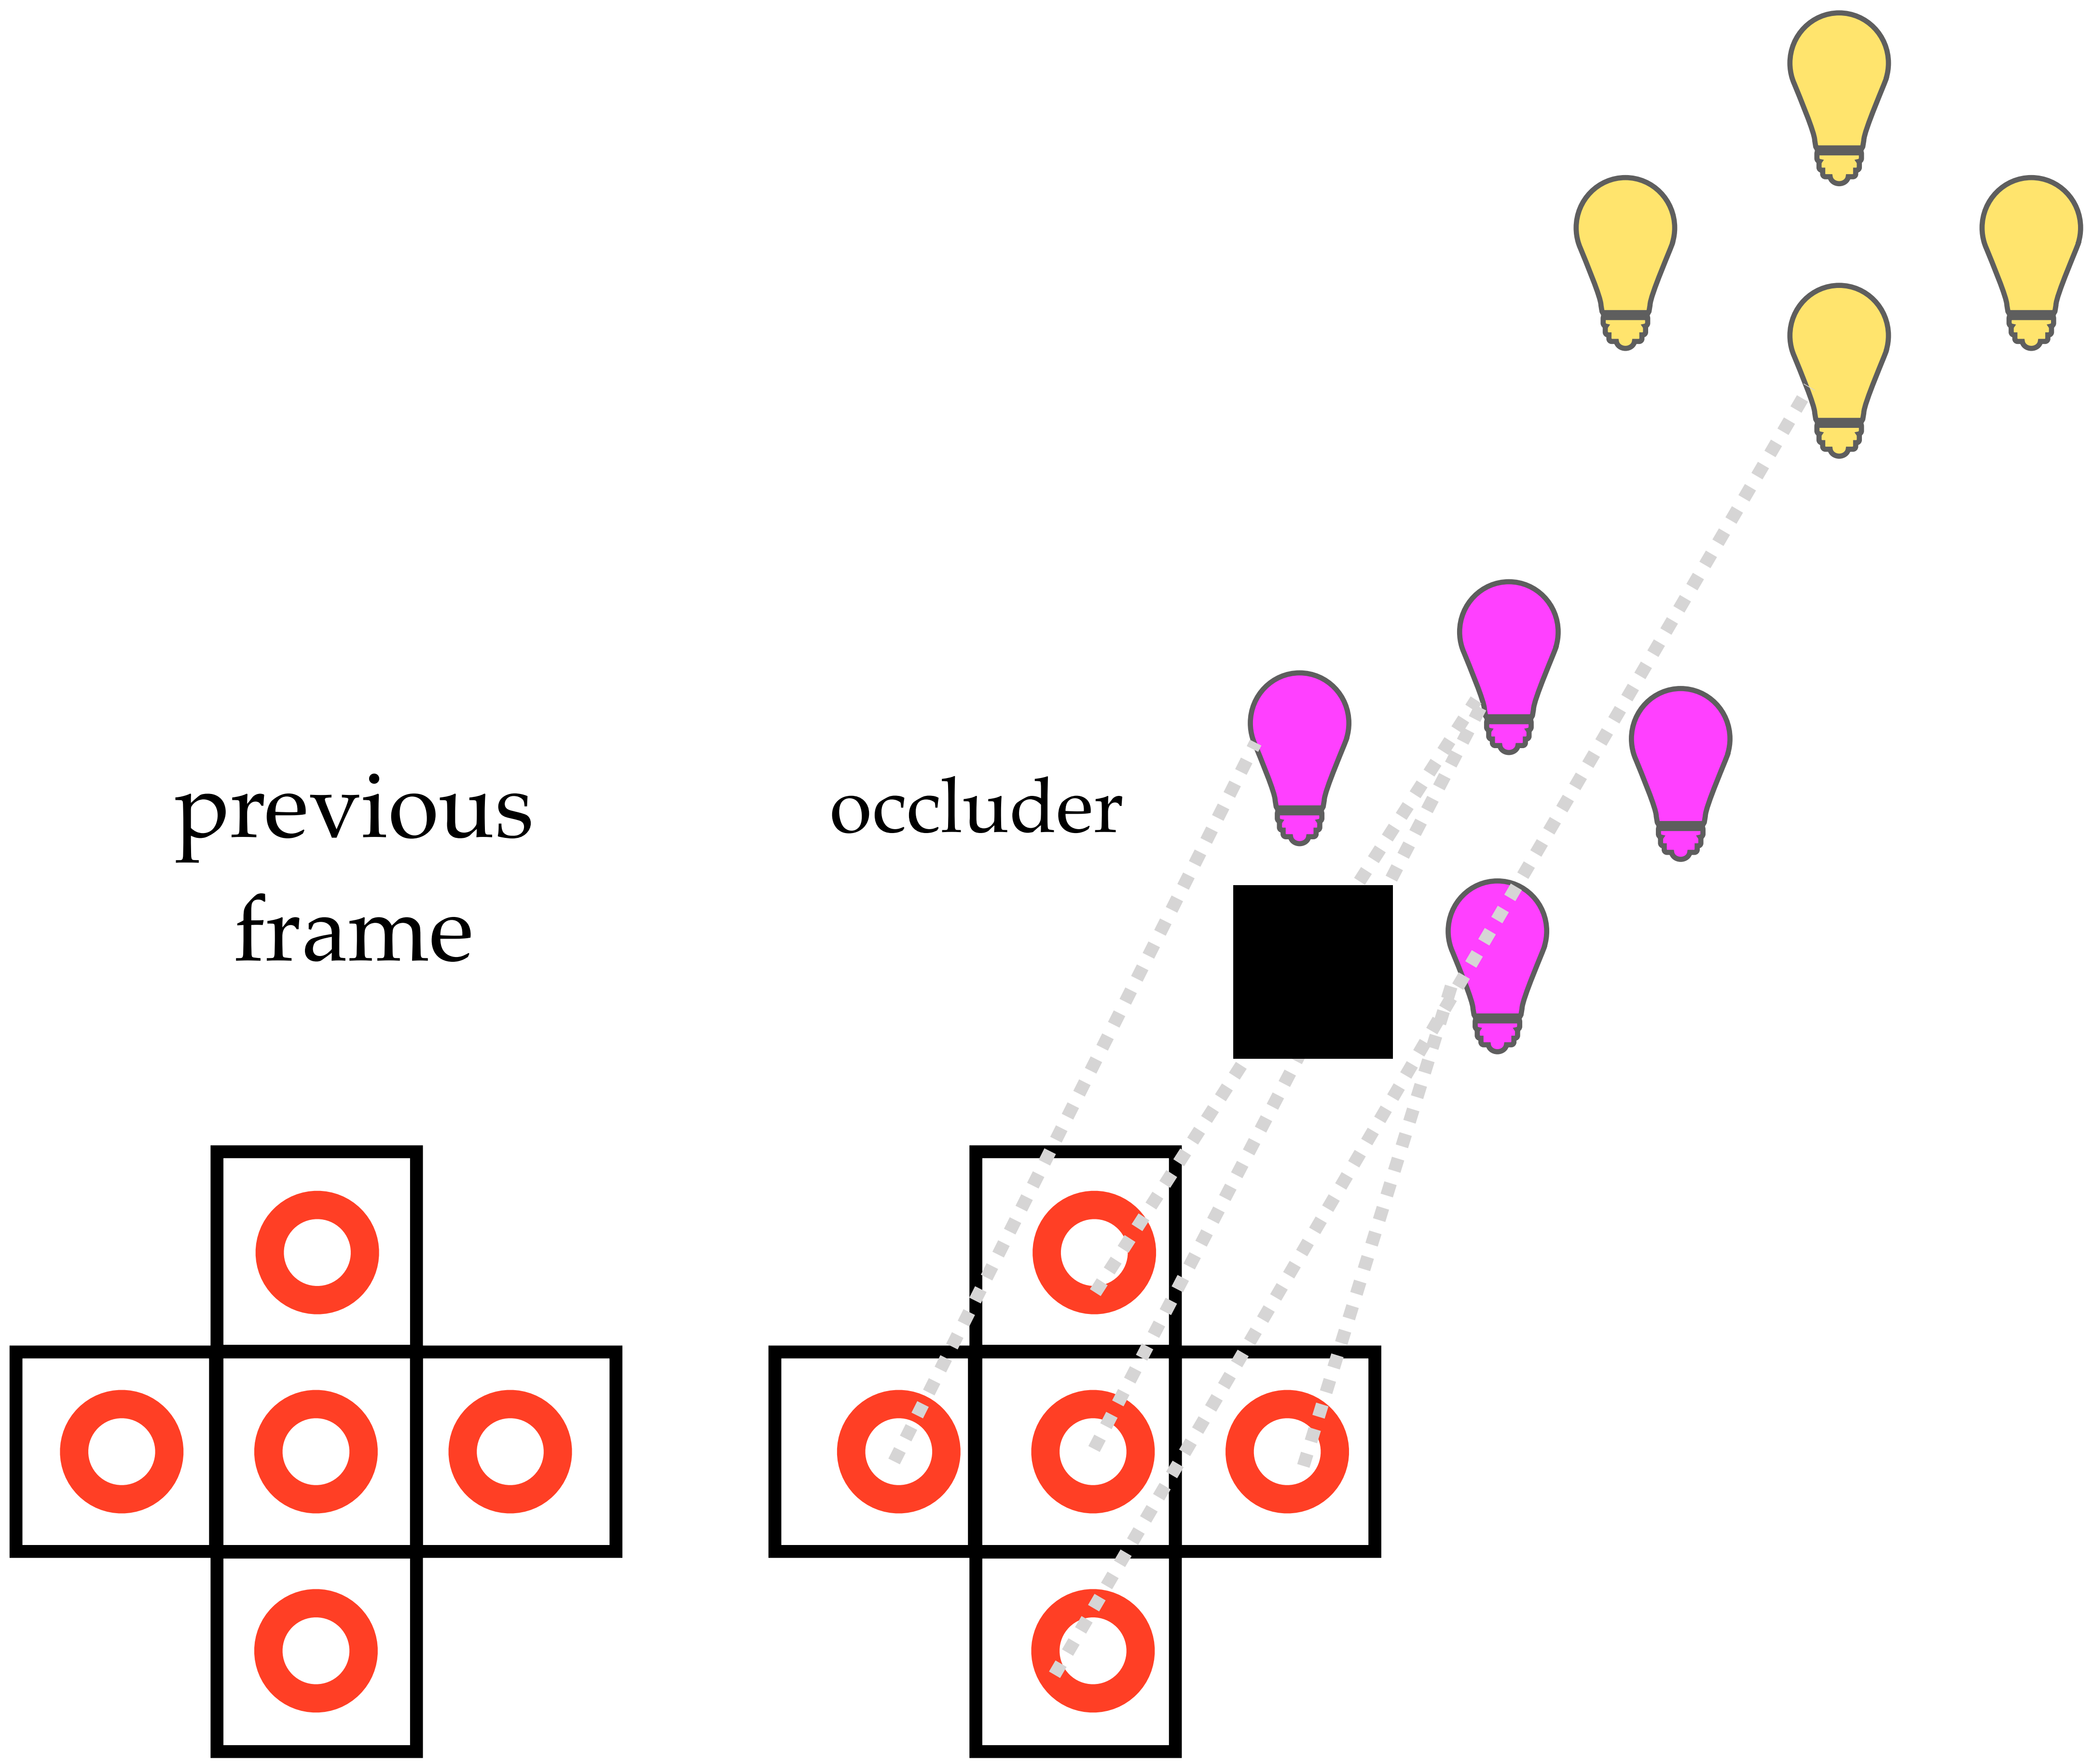

- each pixel sample a light using RIS (e.g., M = 32)

The ReSTIR algorithm

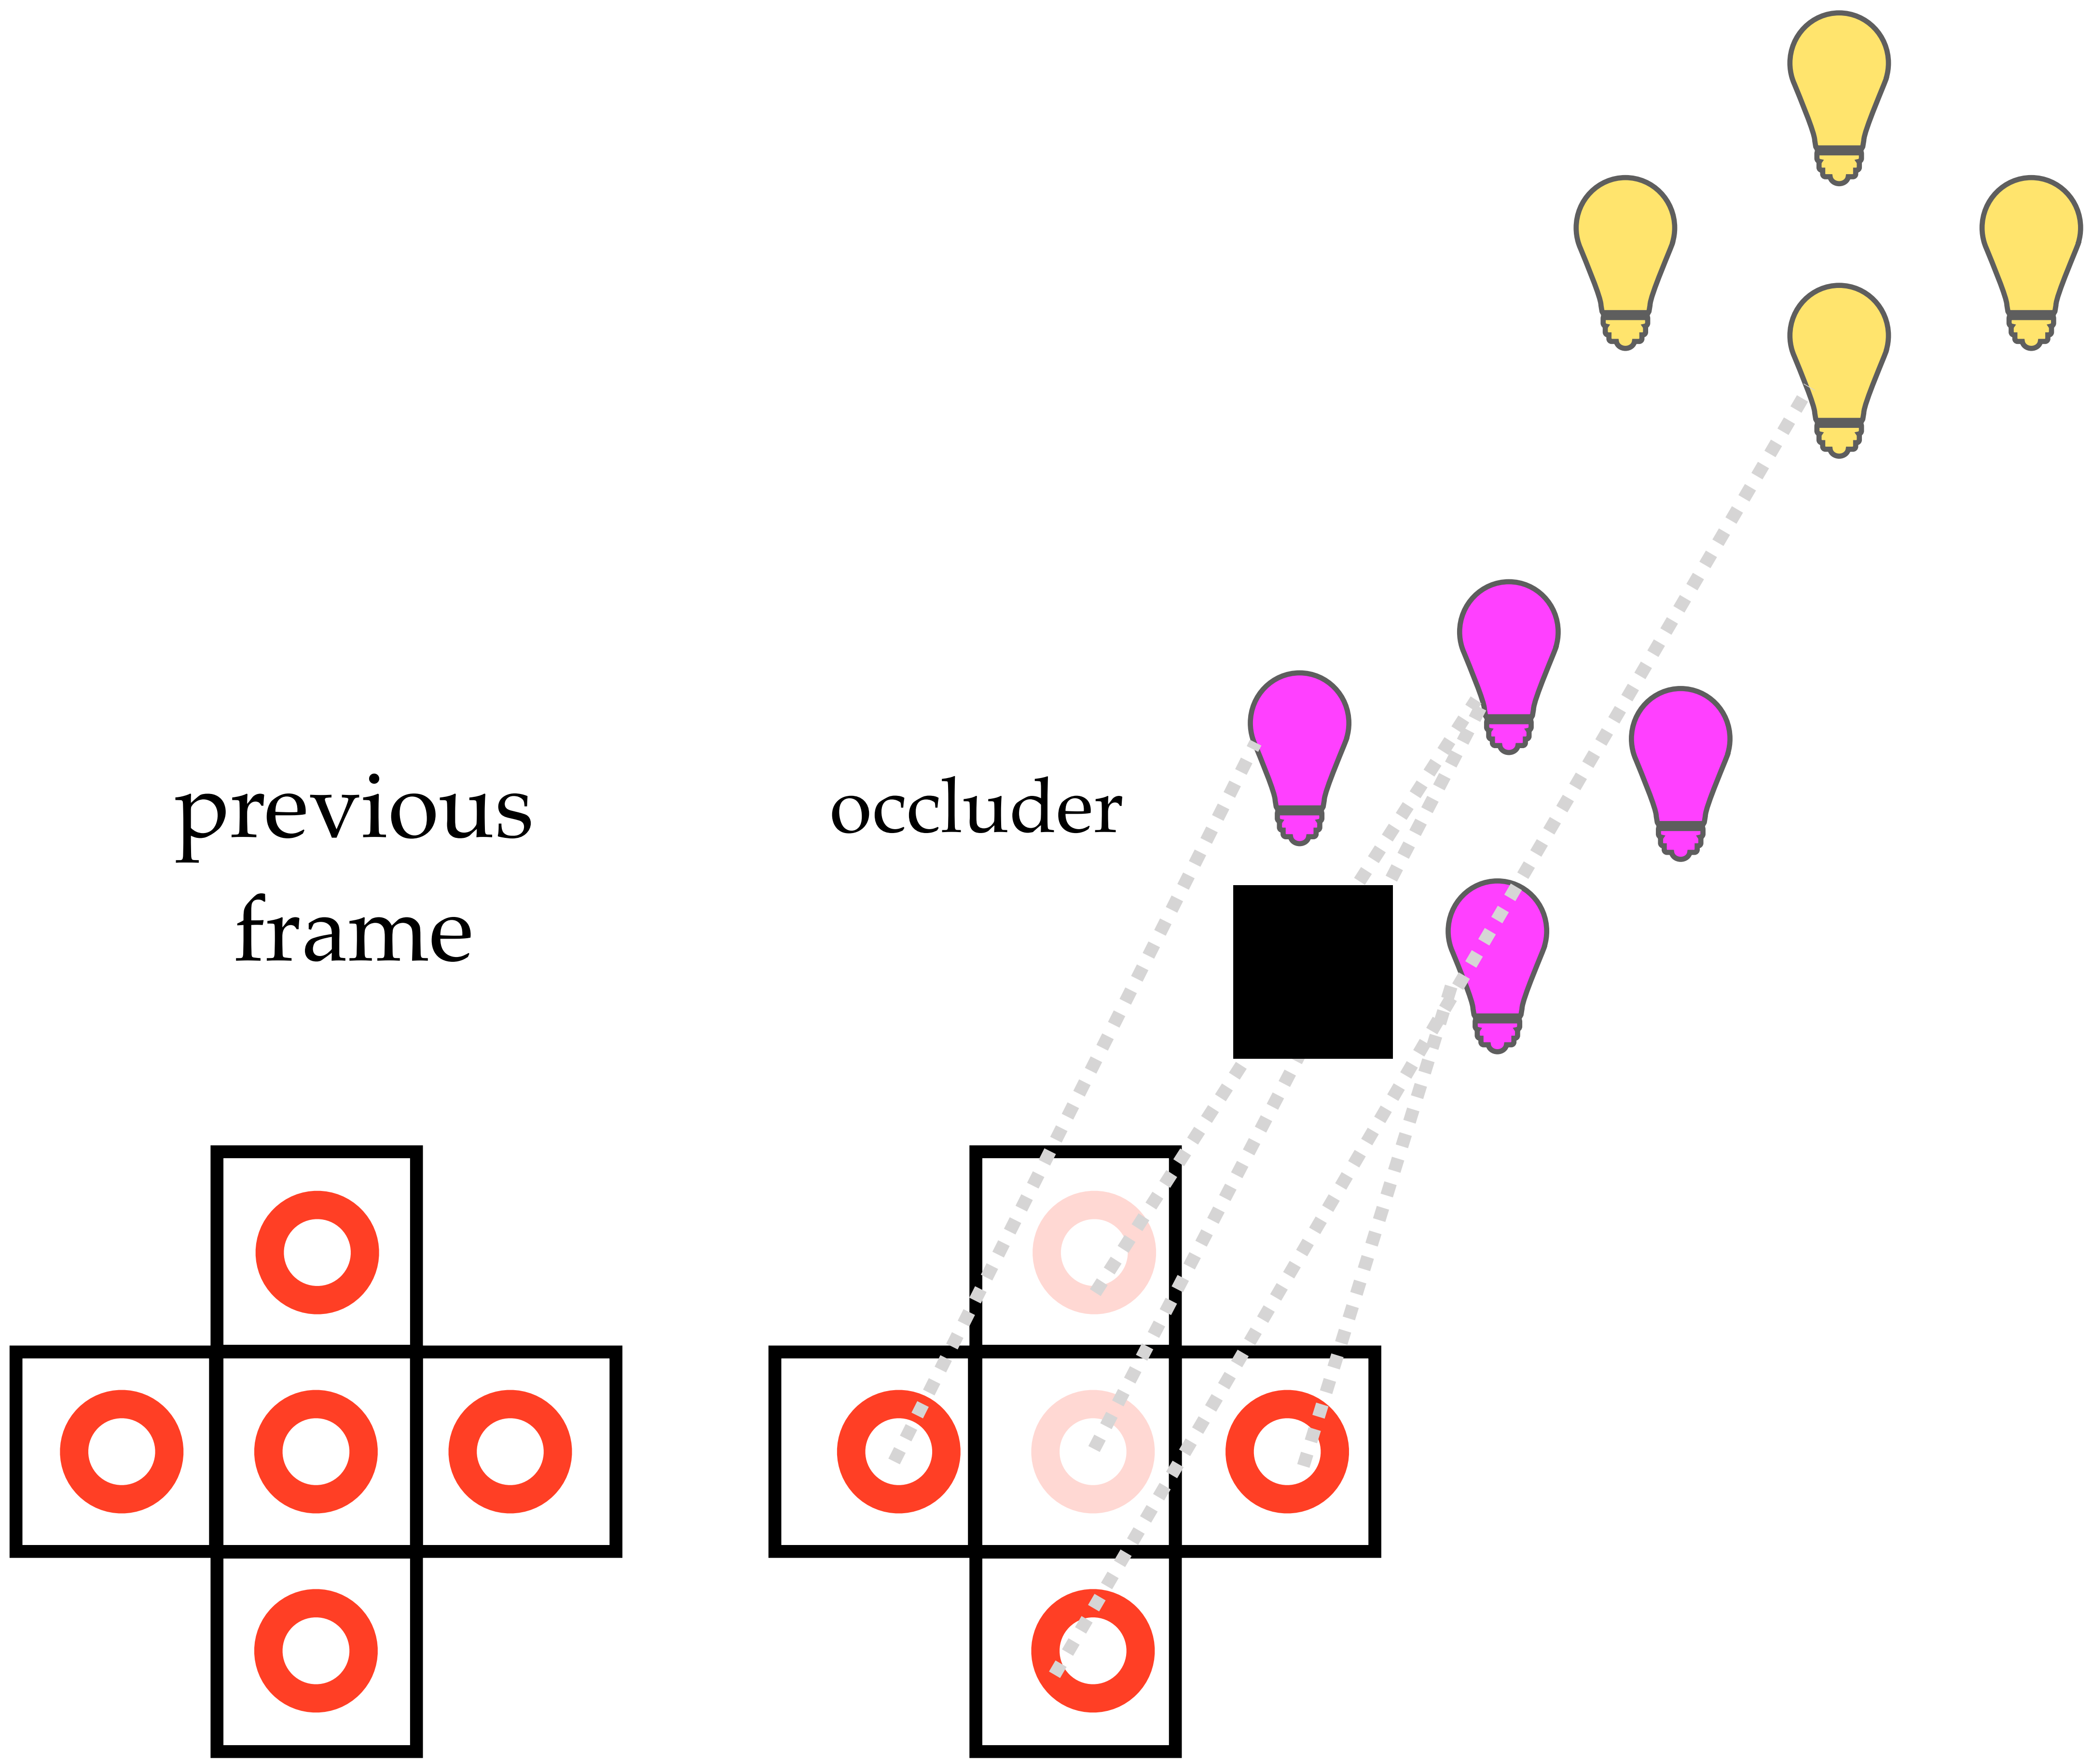

- each pixel sample a light using RIS (e.g., $M = 32$)
- evaluate visibility and set $W = 0$ for occluded pixels

The ReSTIR algorithm

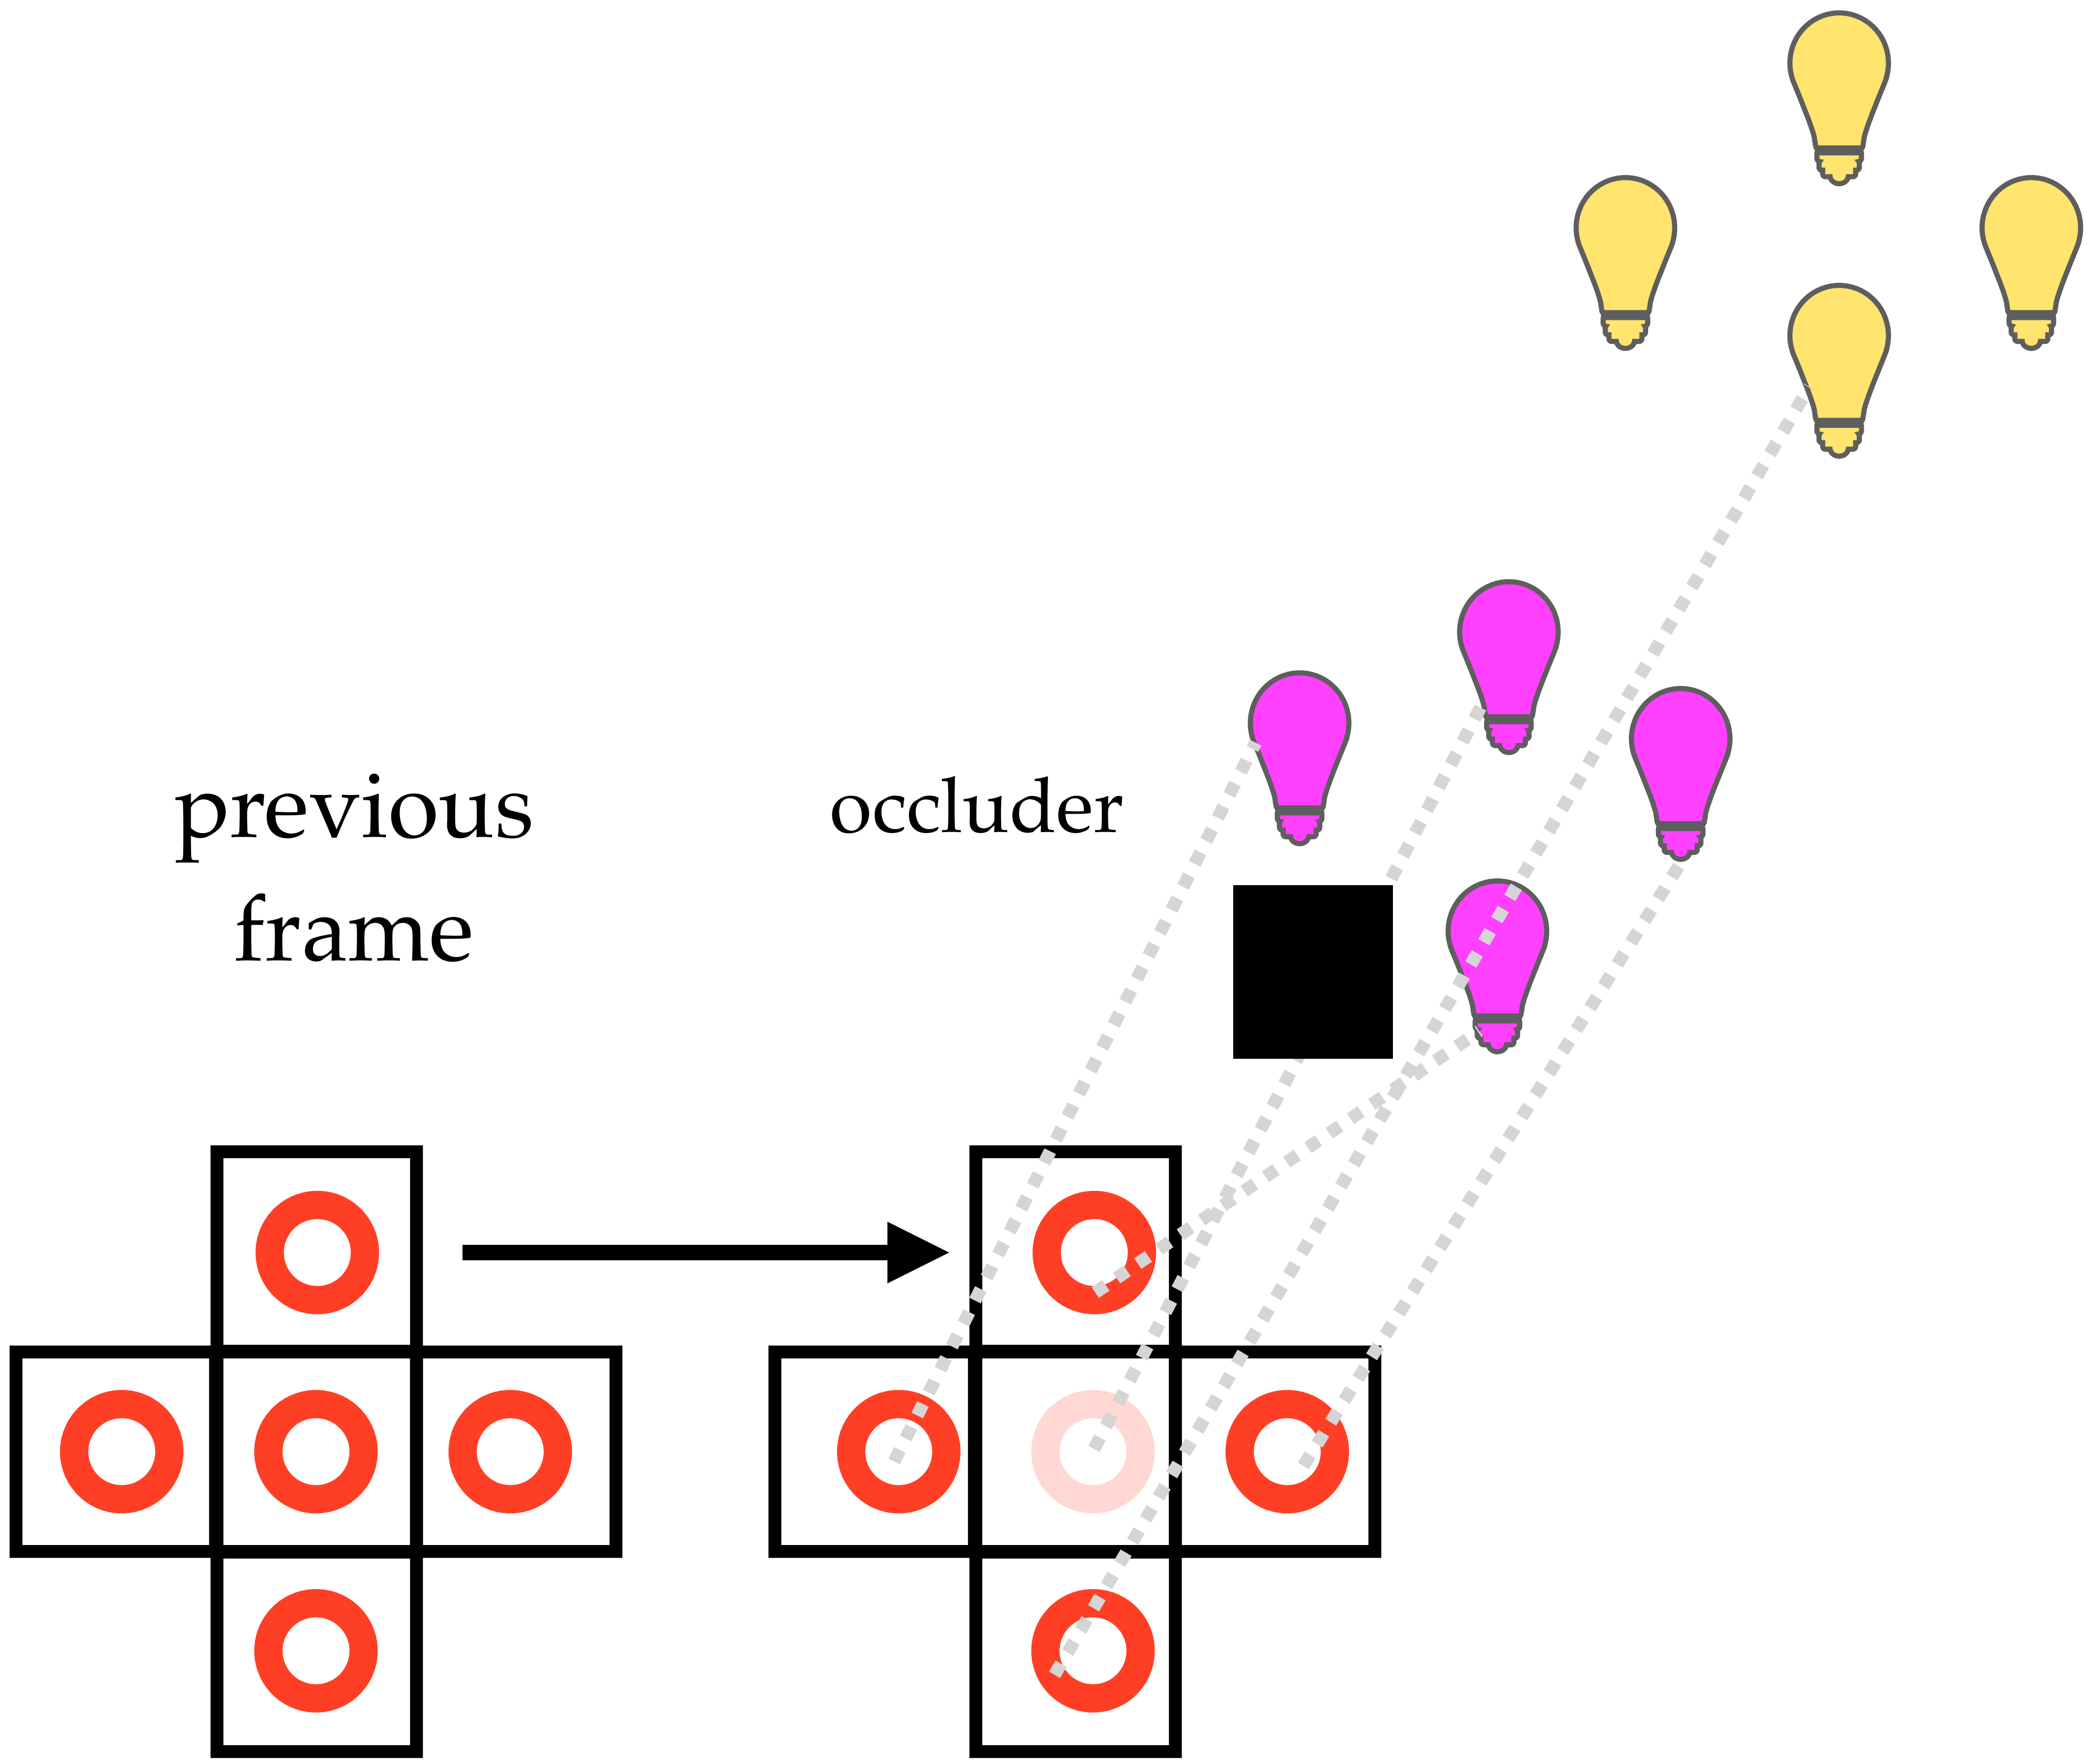

- each pixel sample a light using RIS (e.g., M = 32)
- evaluate visibility and set W = 0 for occluded pixels
- merge the reservoirs from previous frame

The ReSTIR algorithm

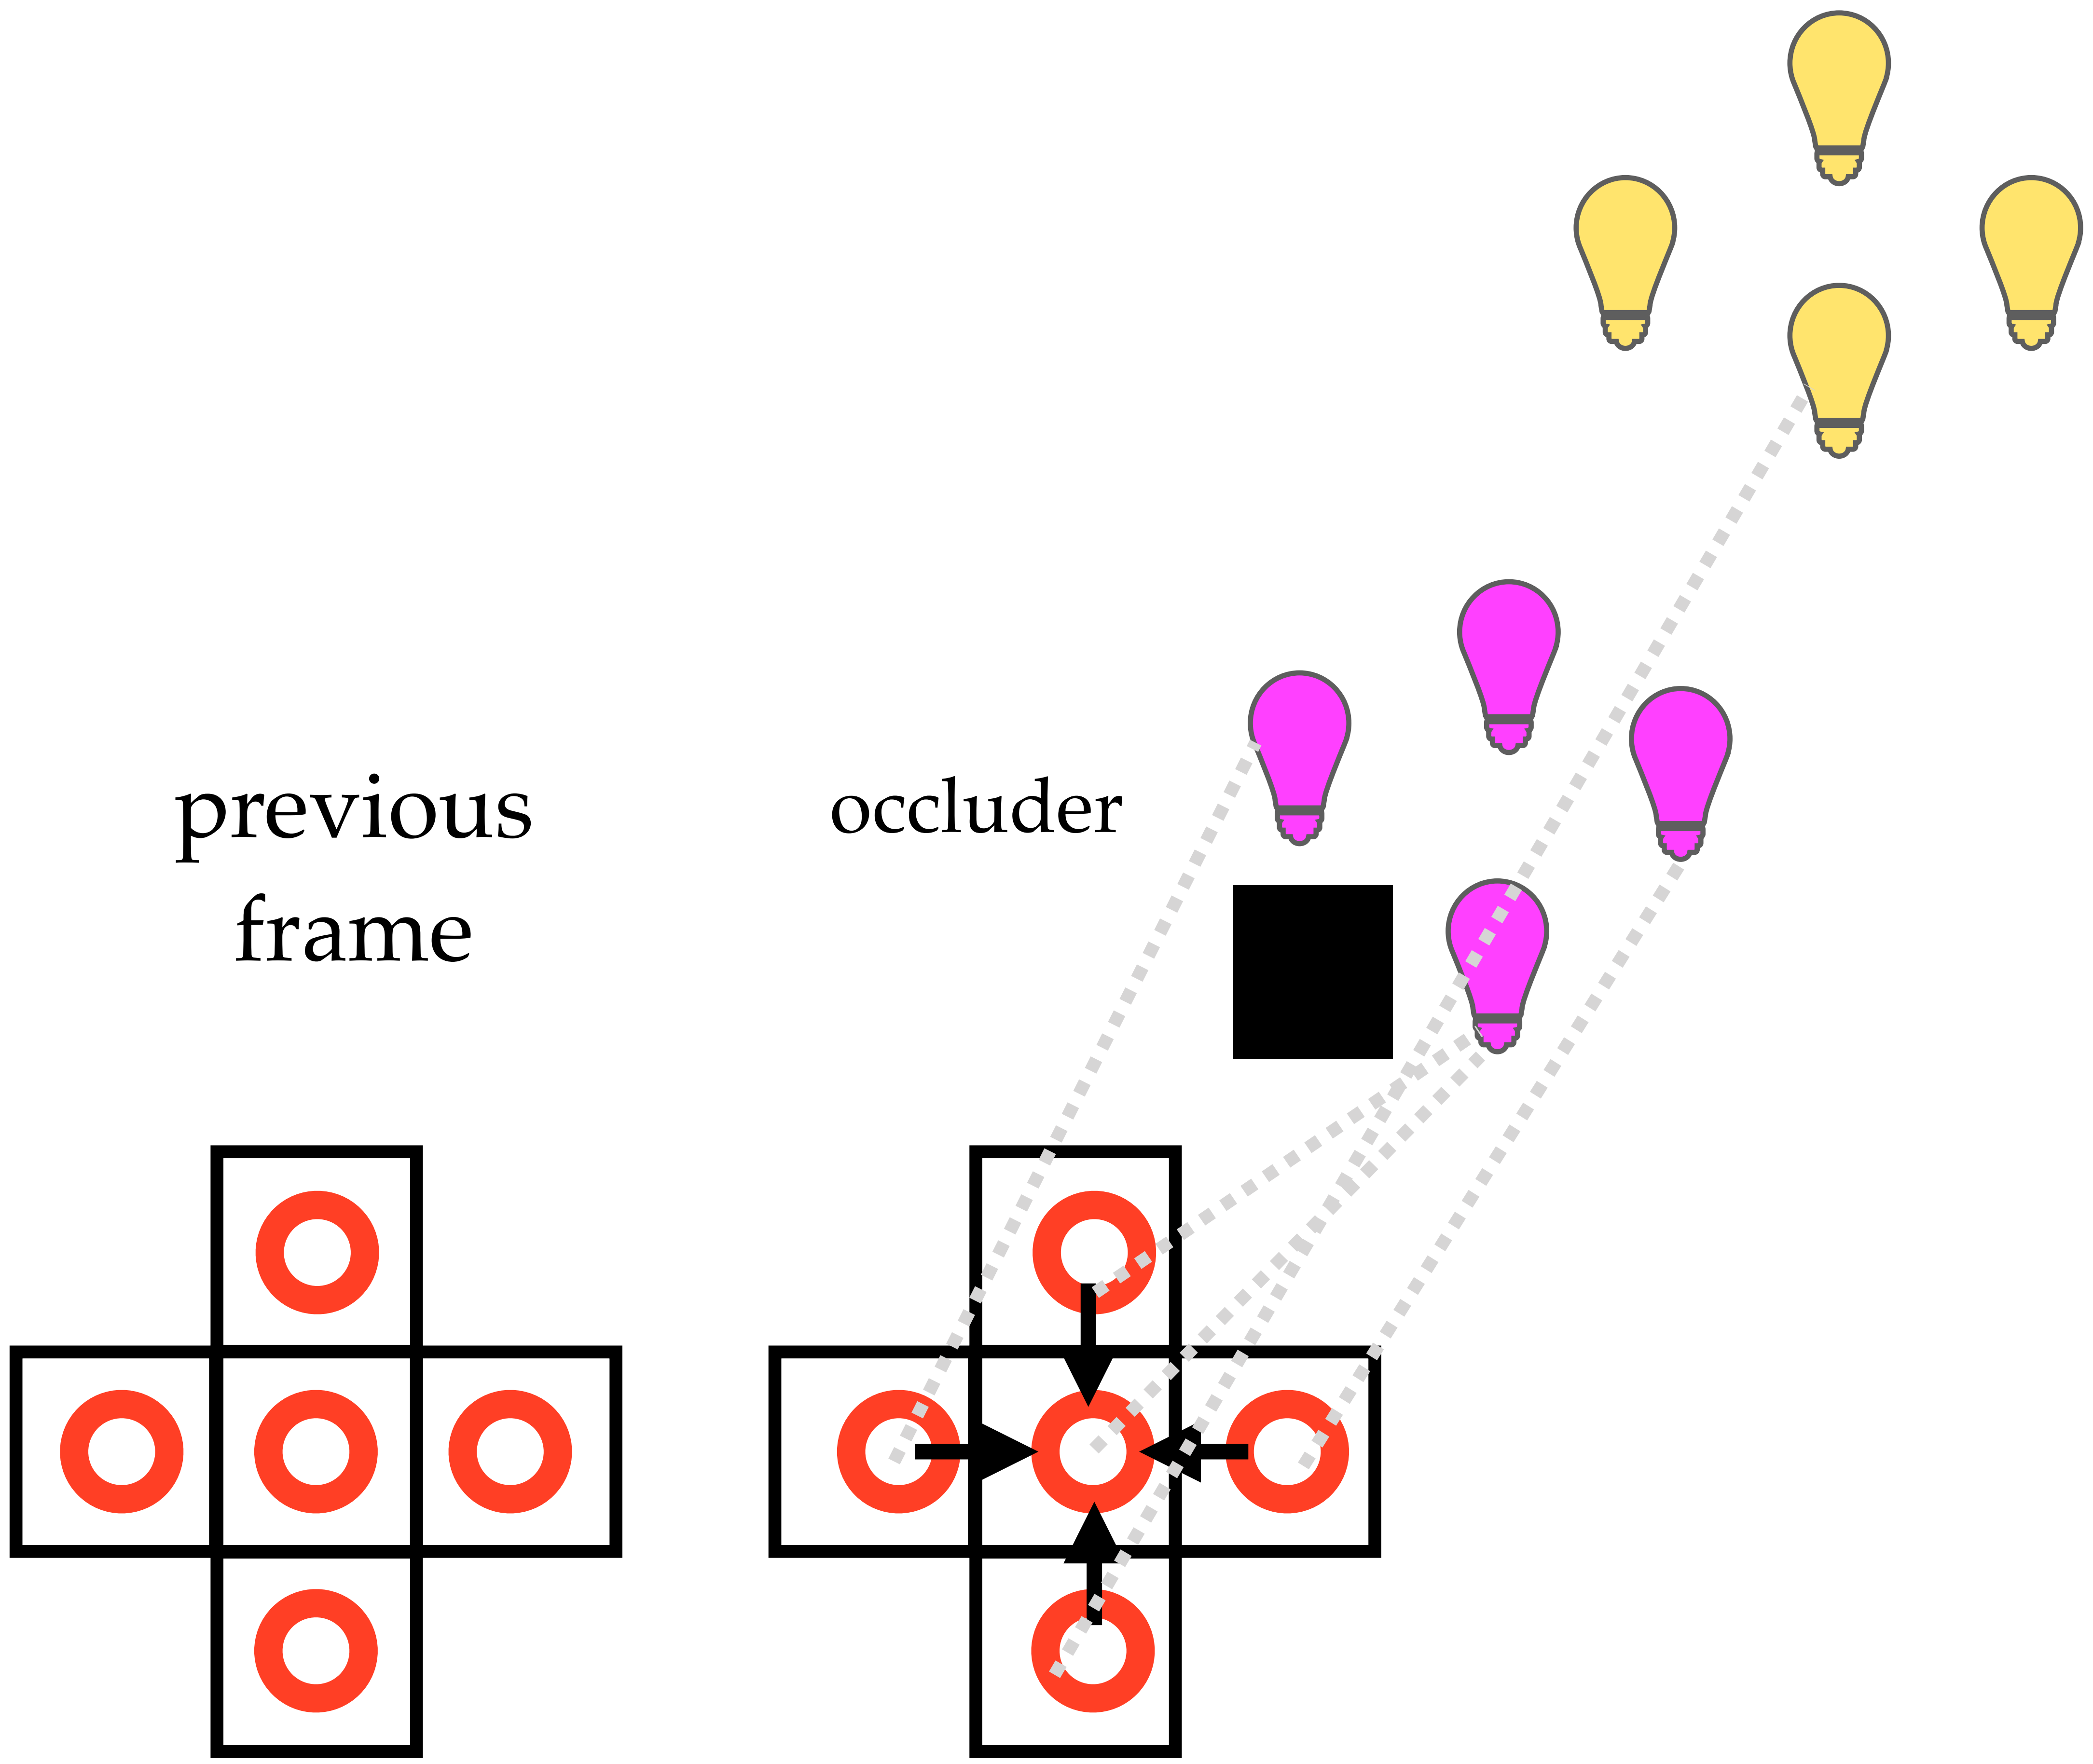

- each pixel sample a light using RIS (e.g., $M = 32$)
- evaluate visibility and set $W = 0$ for occluded pixels
- merge the reservoirs from previous frame
- merge the reservoirs from spatial neighbor pixels

Eye candy time


Reference
20’000 Emitters
Restaurant
Le Petit Coin


[Moreau et al., 2019], 34ms
20’000 Emitters


ReSTIR (unbiased), 30ms
20’000 Emitters


Reference
20'000 Emitters
Le Petit Coin
Le Petit Coin
MENU

[Moreau et al., 2019], 30ms
20'000 Emitters


ReSTIR (unbiased), 26ms
20’000 Emitters
MENU


Reference
23’000 Emitters
S
City
Subway


[Moreau et al., 2019], 29ms
23’000 Emitters


ReSTIR (unbiased), 17ms
23’000 Emitters

ReSTIR vs MCMC

both maintain a chain of samples and reuse previous ones

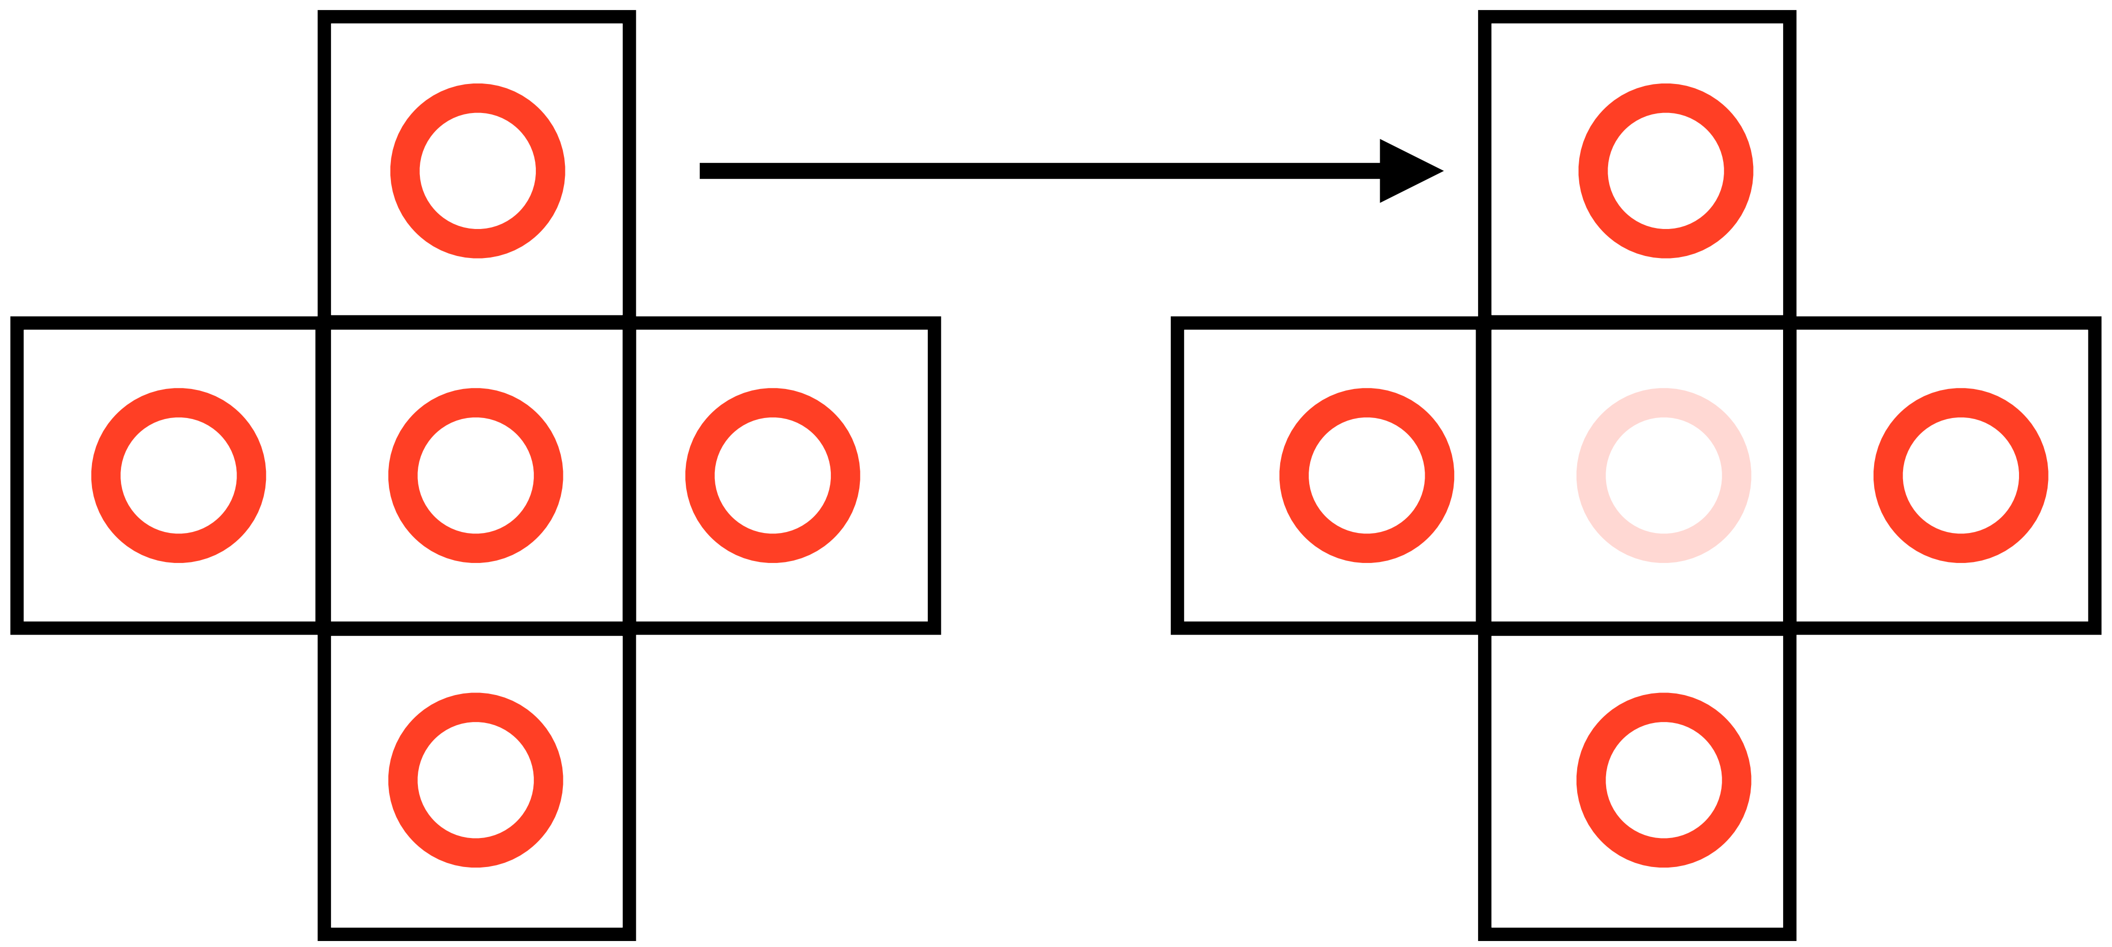

Extending ReSTIR to handle global illumination

want to reuse paths instead of lights

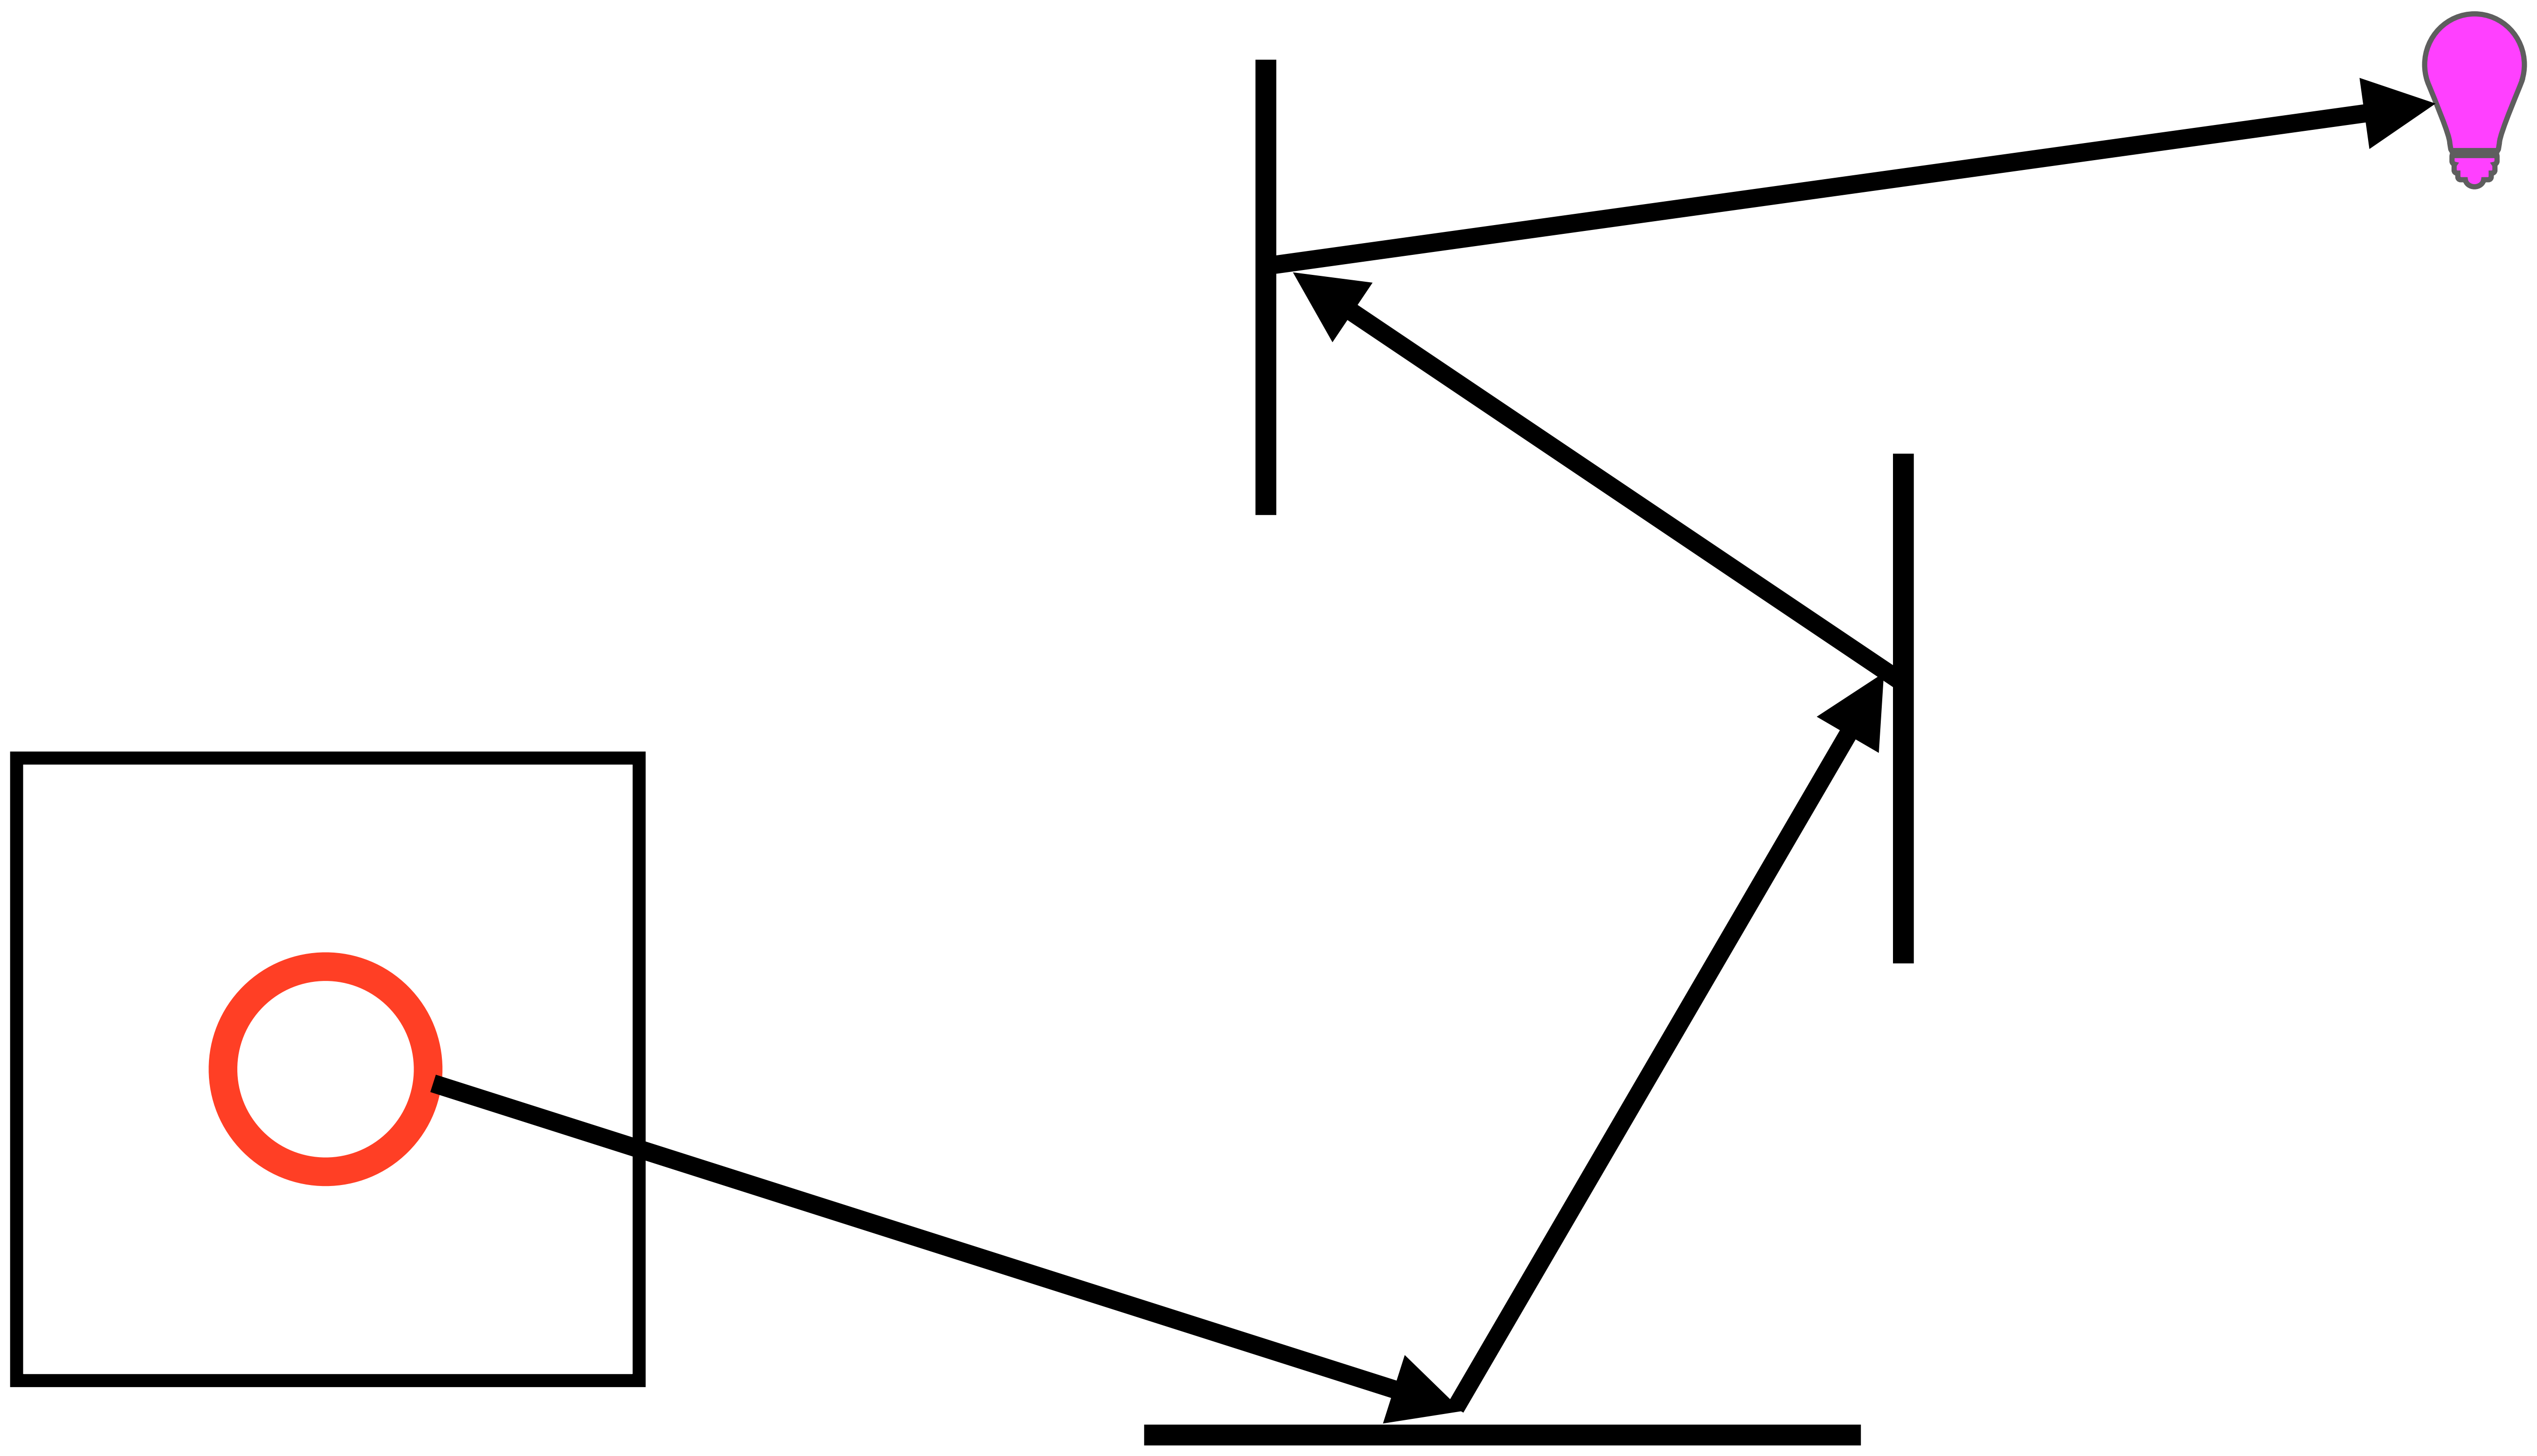

Extending ReSTIR to handle global illumination

want to reuse paths instead of lights

idea: treat the second path vertex as virtual point light

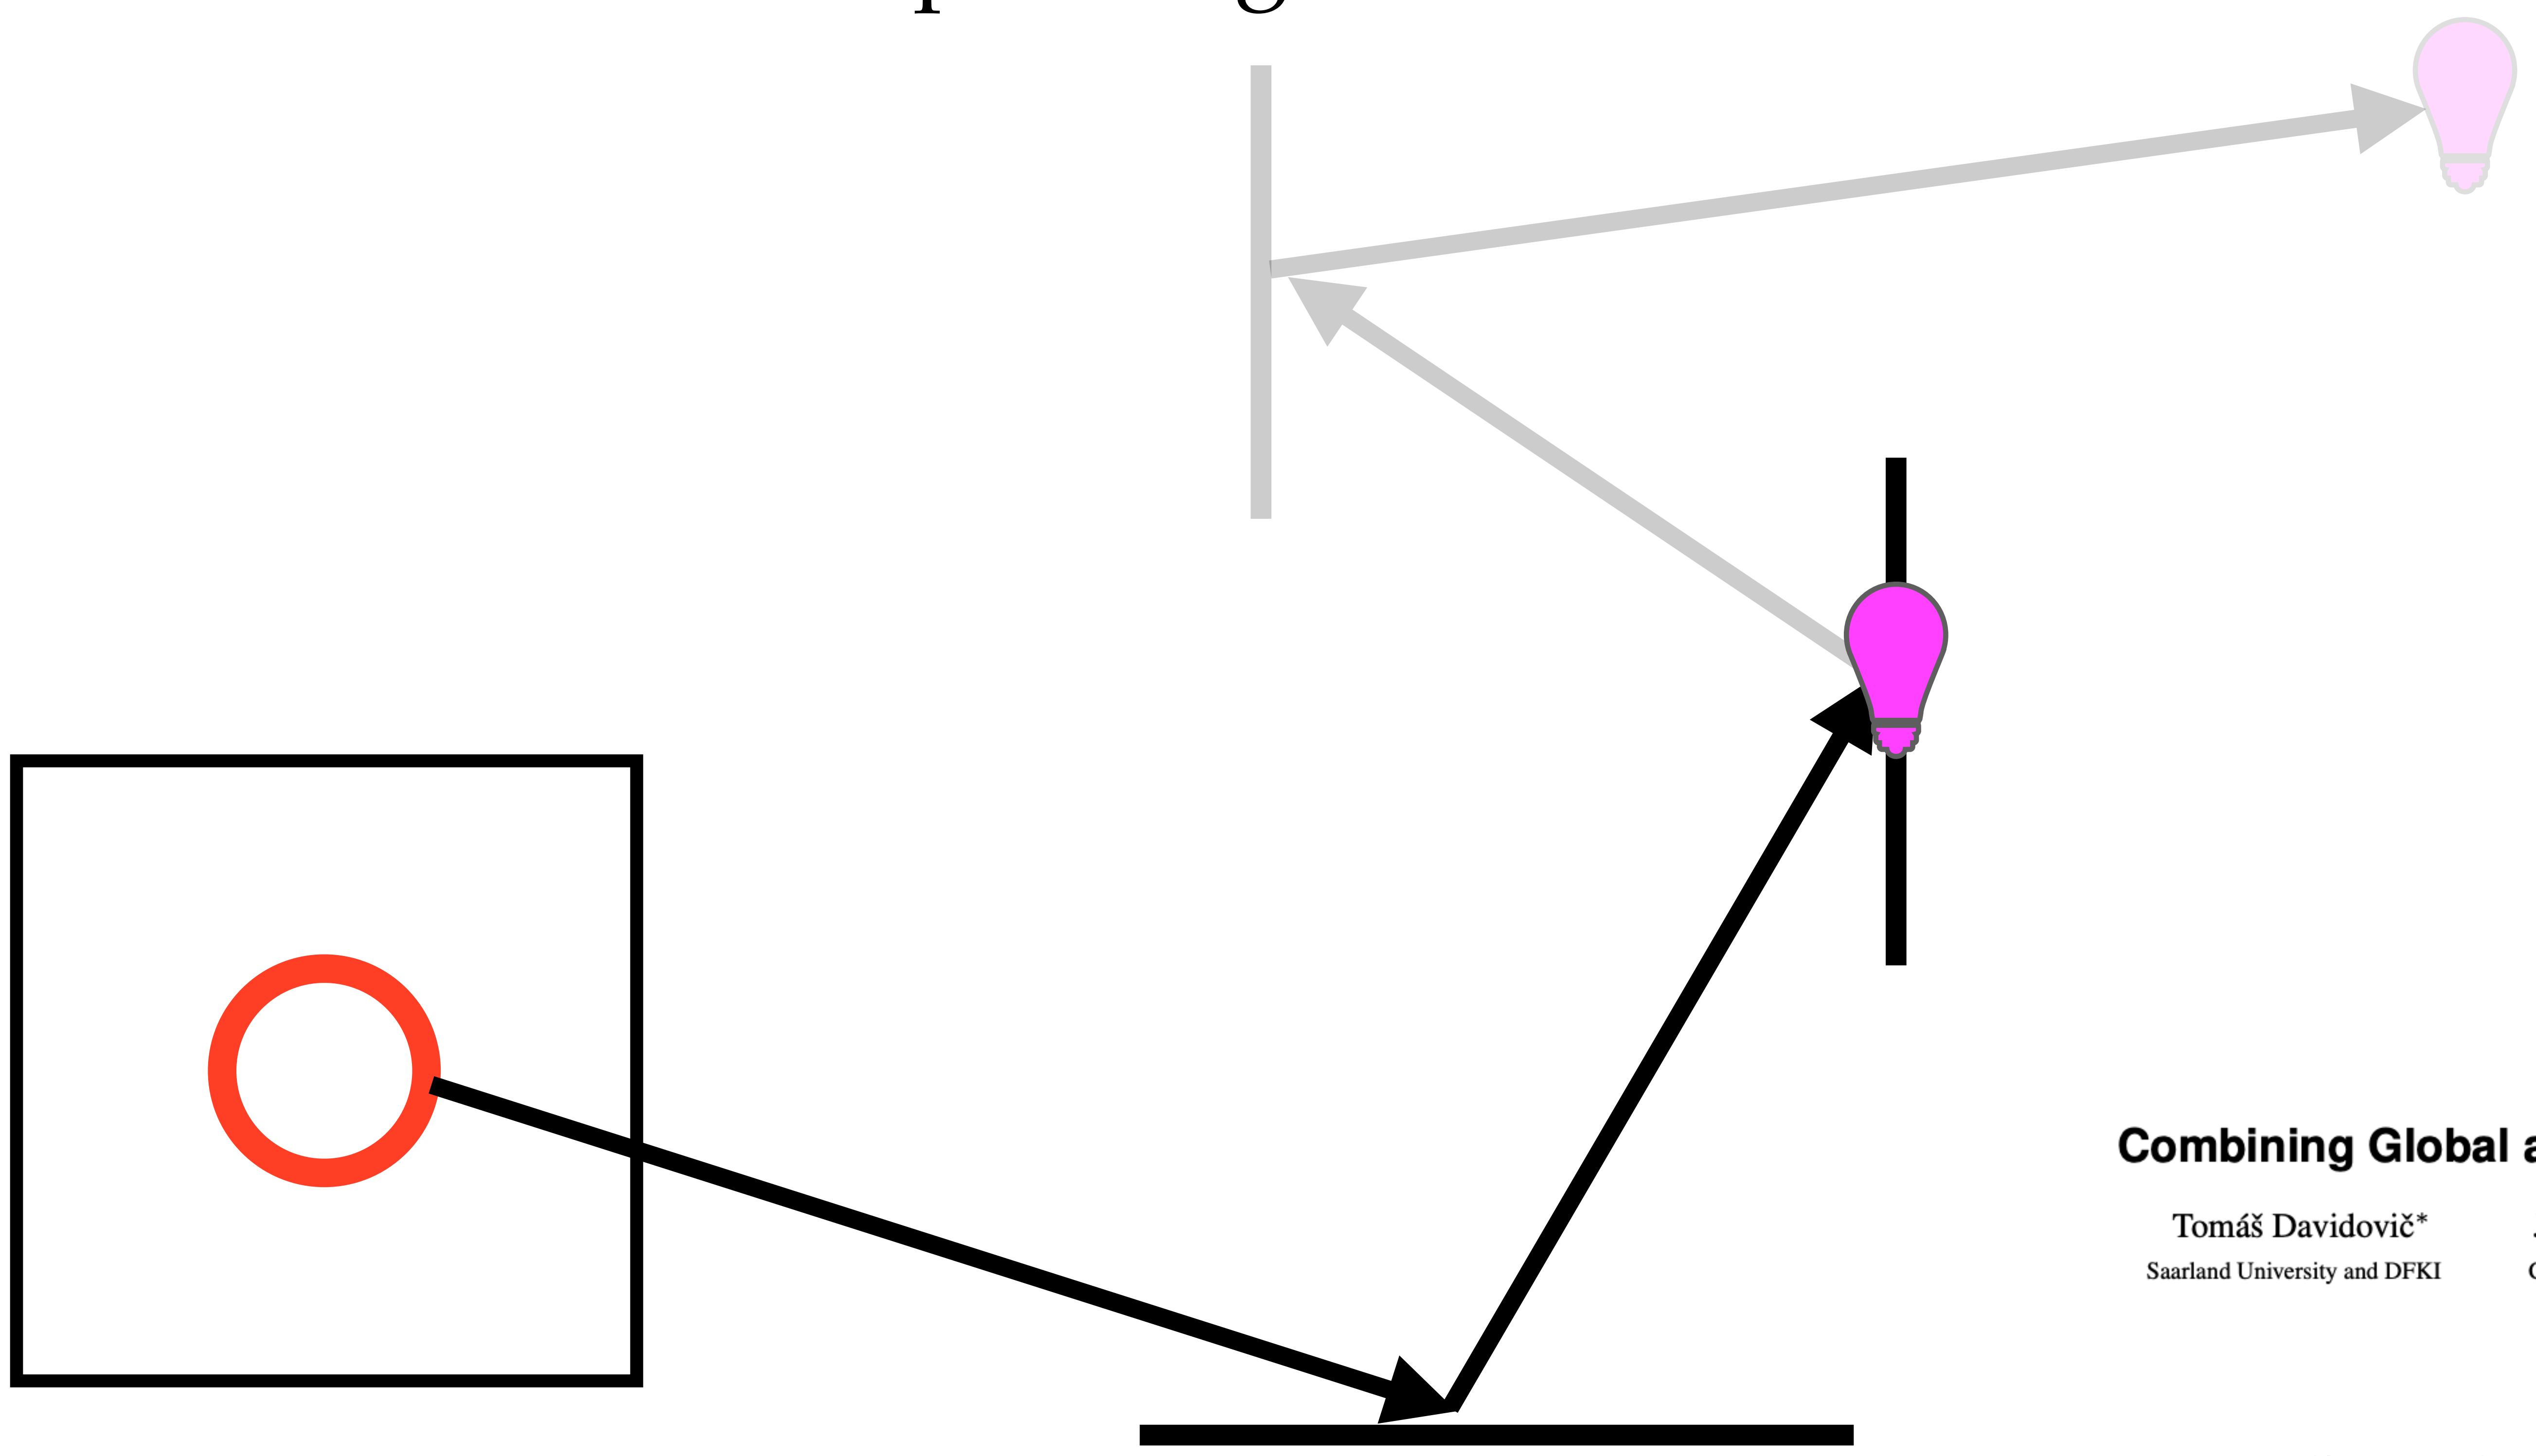

Accelerating Path Tracing by Re-Using Paths

Philippe Bekaert
Max-Planck-Institut für Informatik, Saarbrücken, Germany.
Philippe.Bekaert@mpi-sb.mpg.de

Mateu Sbert
Institut d'Informàtica i Aplicacions, Universitat de Girona, Spain.
mateu@ima.udg.es

John Halton
University of North Carolina at Chapel Hill.
halton@cs.unc.edu

Combining Global and Local Virtual Lights for Detailed Glossy Illumination

Tomáš Davidovič* — Saarland University and DFKI
Jaroslav Křivánek — Charles University, Prague; Cornell University
Miloš Hašan — Harvard University
Philipp Slusallek — Saarland University and DFKI
Kavita Bala — Cornell University

ReSTIR GI: Path Resampling for Real-Time Path Tracing

Y. Ouyang[1], S. Liu[1], M. Kettunen[1], M. Pharr[1], J. Pantaleoni[1]

[1]NVIDIA Corporation, Santa Clara, CA, USA

Extending ReSTIR to handle global illumination

want to reuse paths instead of lights

idea: treat the second path vertex as virtual point light

during reservoir merging, connect to the VPL

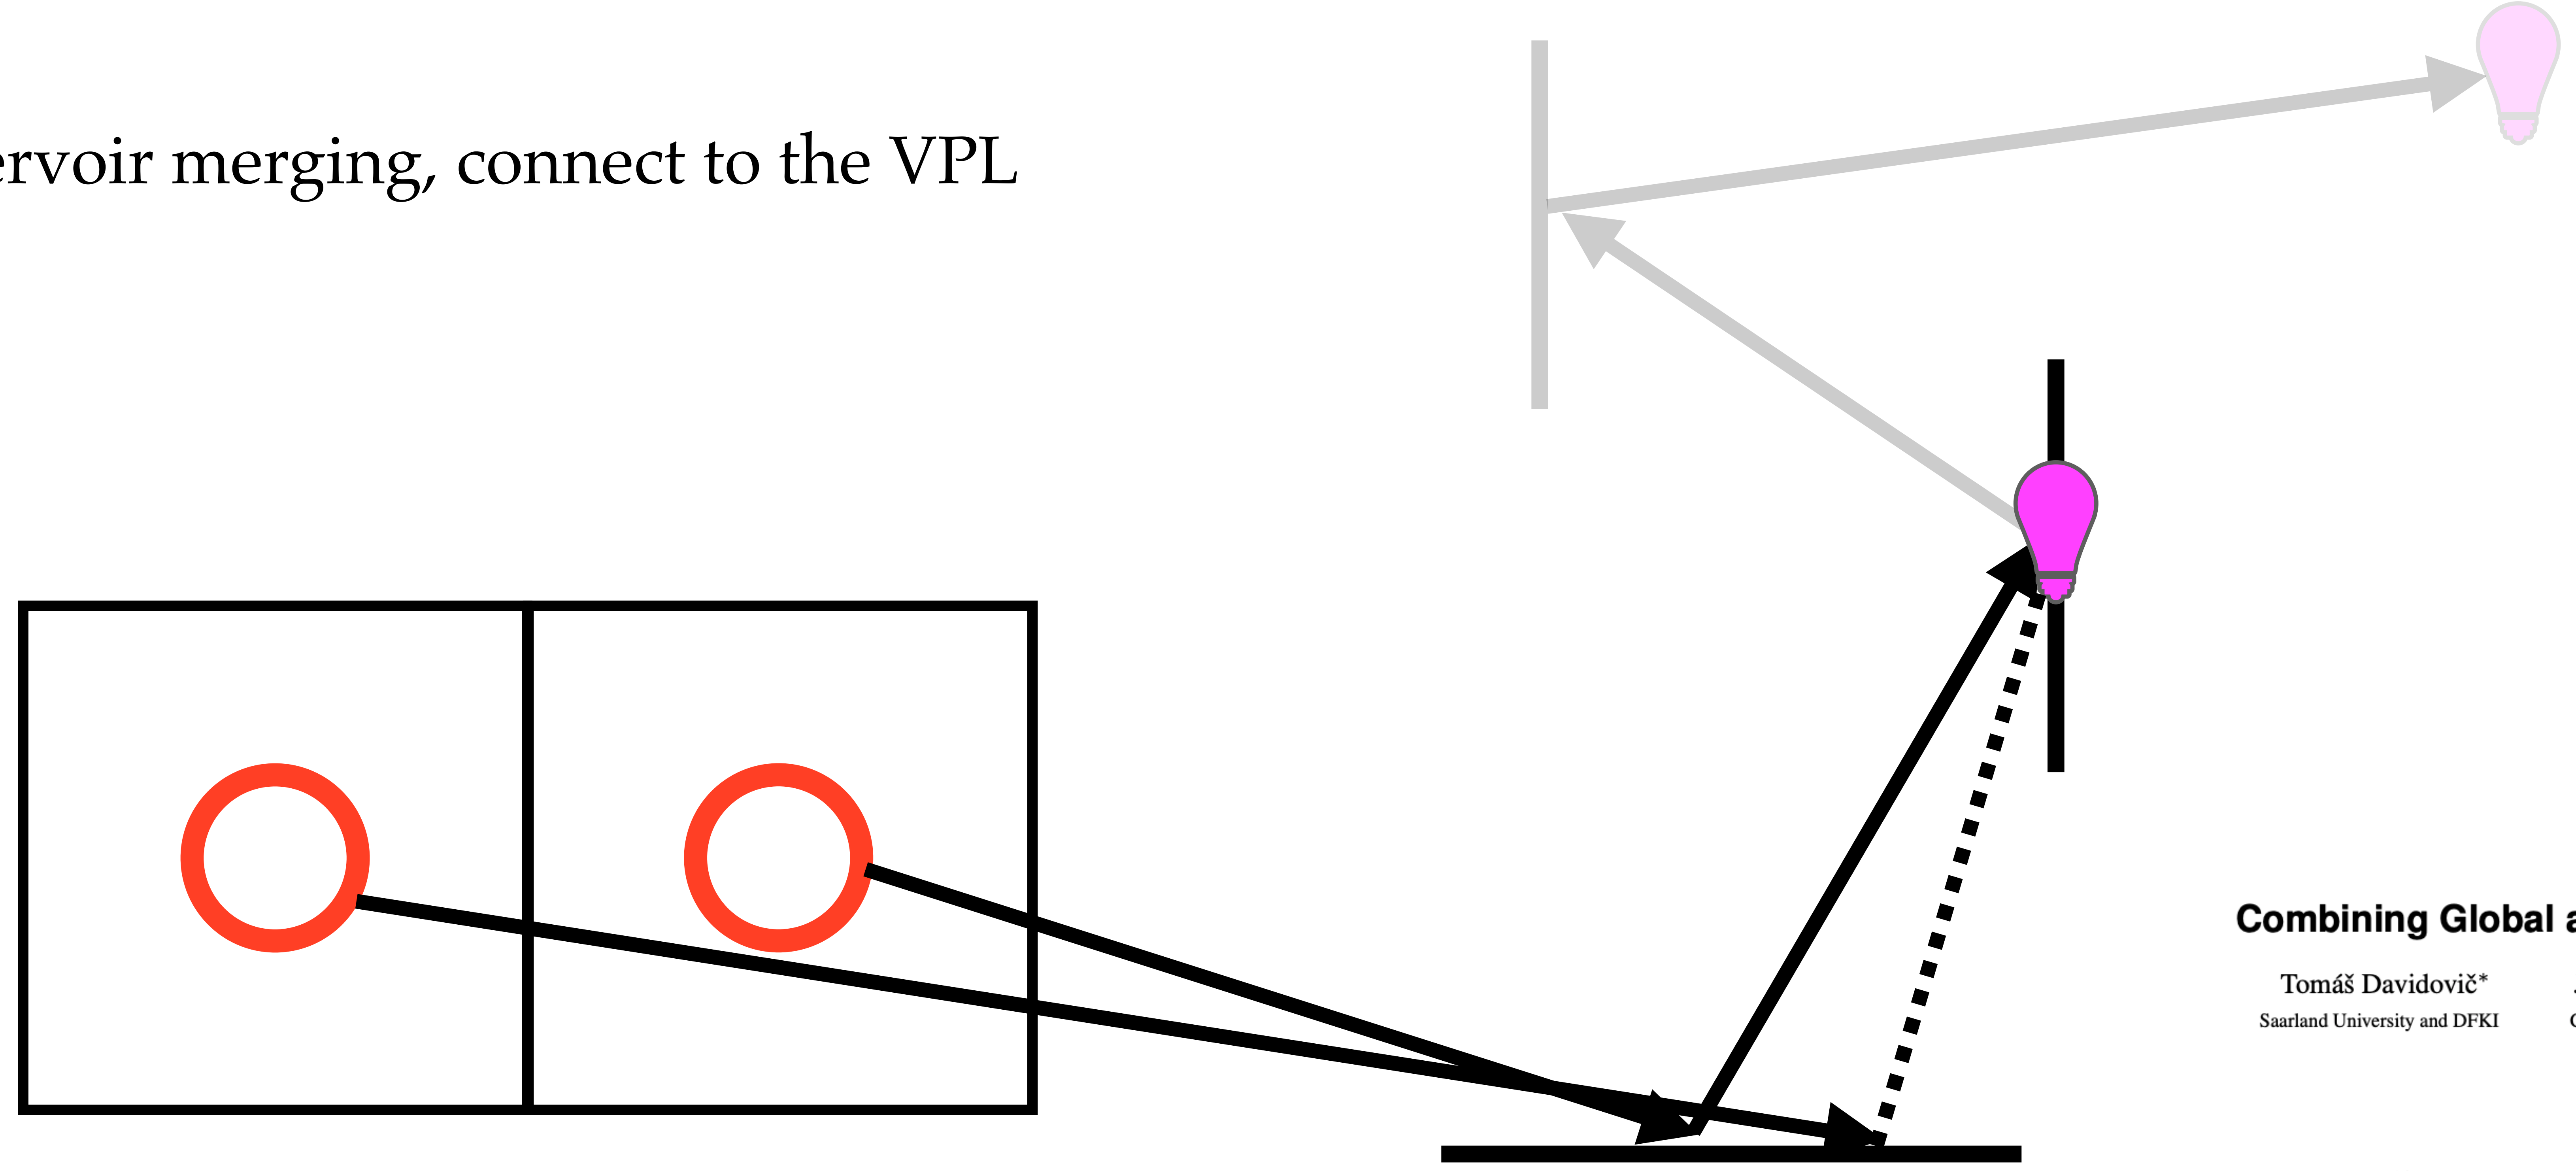

Accelerating Path Tracing by Re-Using Paths

Philippe Bekaert
Max-Planck-Institut für Informatik, Saarbrücken, Germany.
Philippe.Bekaert@mpi-sb.mpg.de

Mateu Sbert
Institut d'Informàtica i Aplicacions, Universitat de Girona, Spain.
mateu@ima.udg.es

John Halton
University of North Carolina at Chapel Hill.
halton@cs.unc.edu

Combining Global and Local Virtual Lights for Detailed Glossy Illumination

Tomáš Davidovič* — Saarland University and DFKI
Jaroslav Křivánek — Charles University, Prague; Cornell University
Miloš Hašan — Harvard University
Philipp Slusallek — Saarland University and DFKI
Kavita Bala — Cornell University

ReSTIR GI: Path Resampling for Real-Time Path Tracing

Y. Ouyang[1], S. Liu[1], M. Kettunen[1], M. Pharr[1], J. Pantaleoni[1]

[1]NVIDIA Corporation, Santa Clara, CA, USA

ReSTIR GI

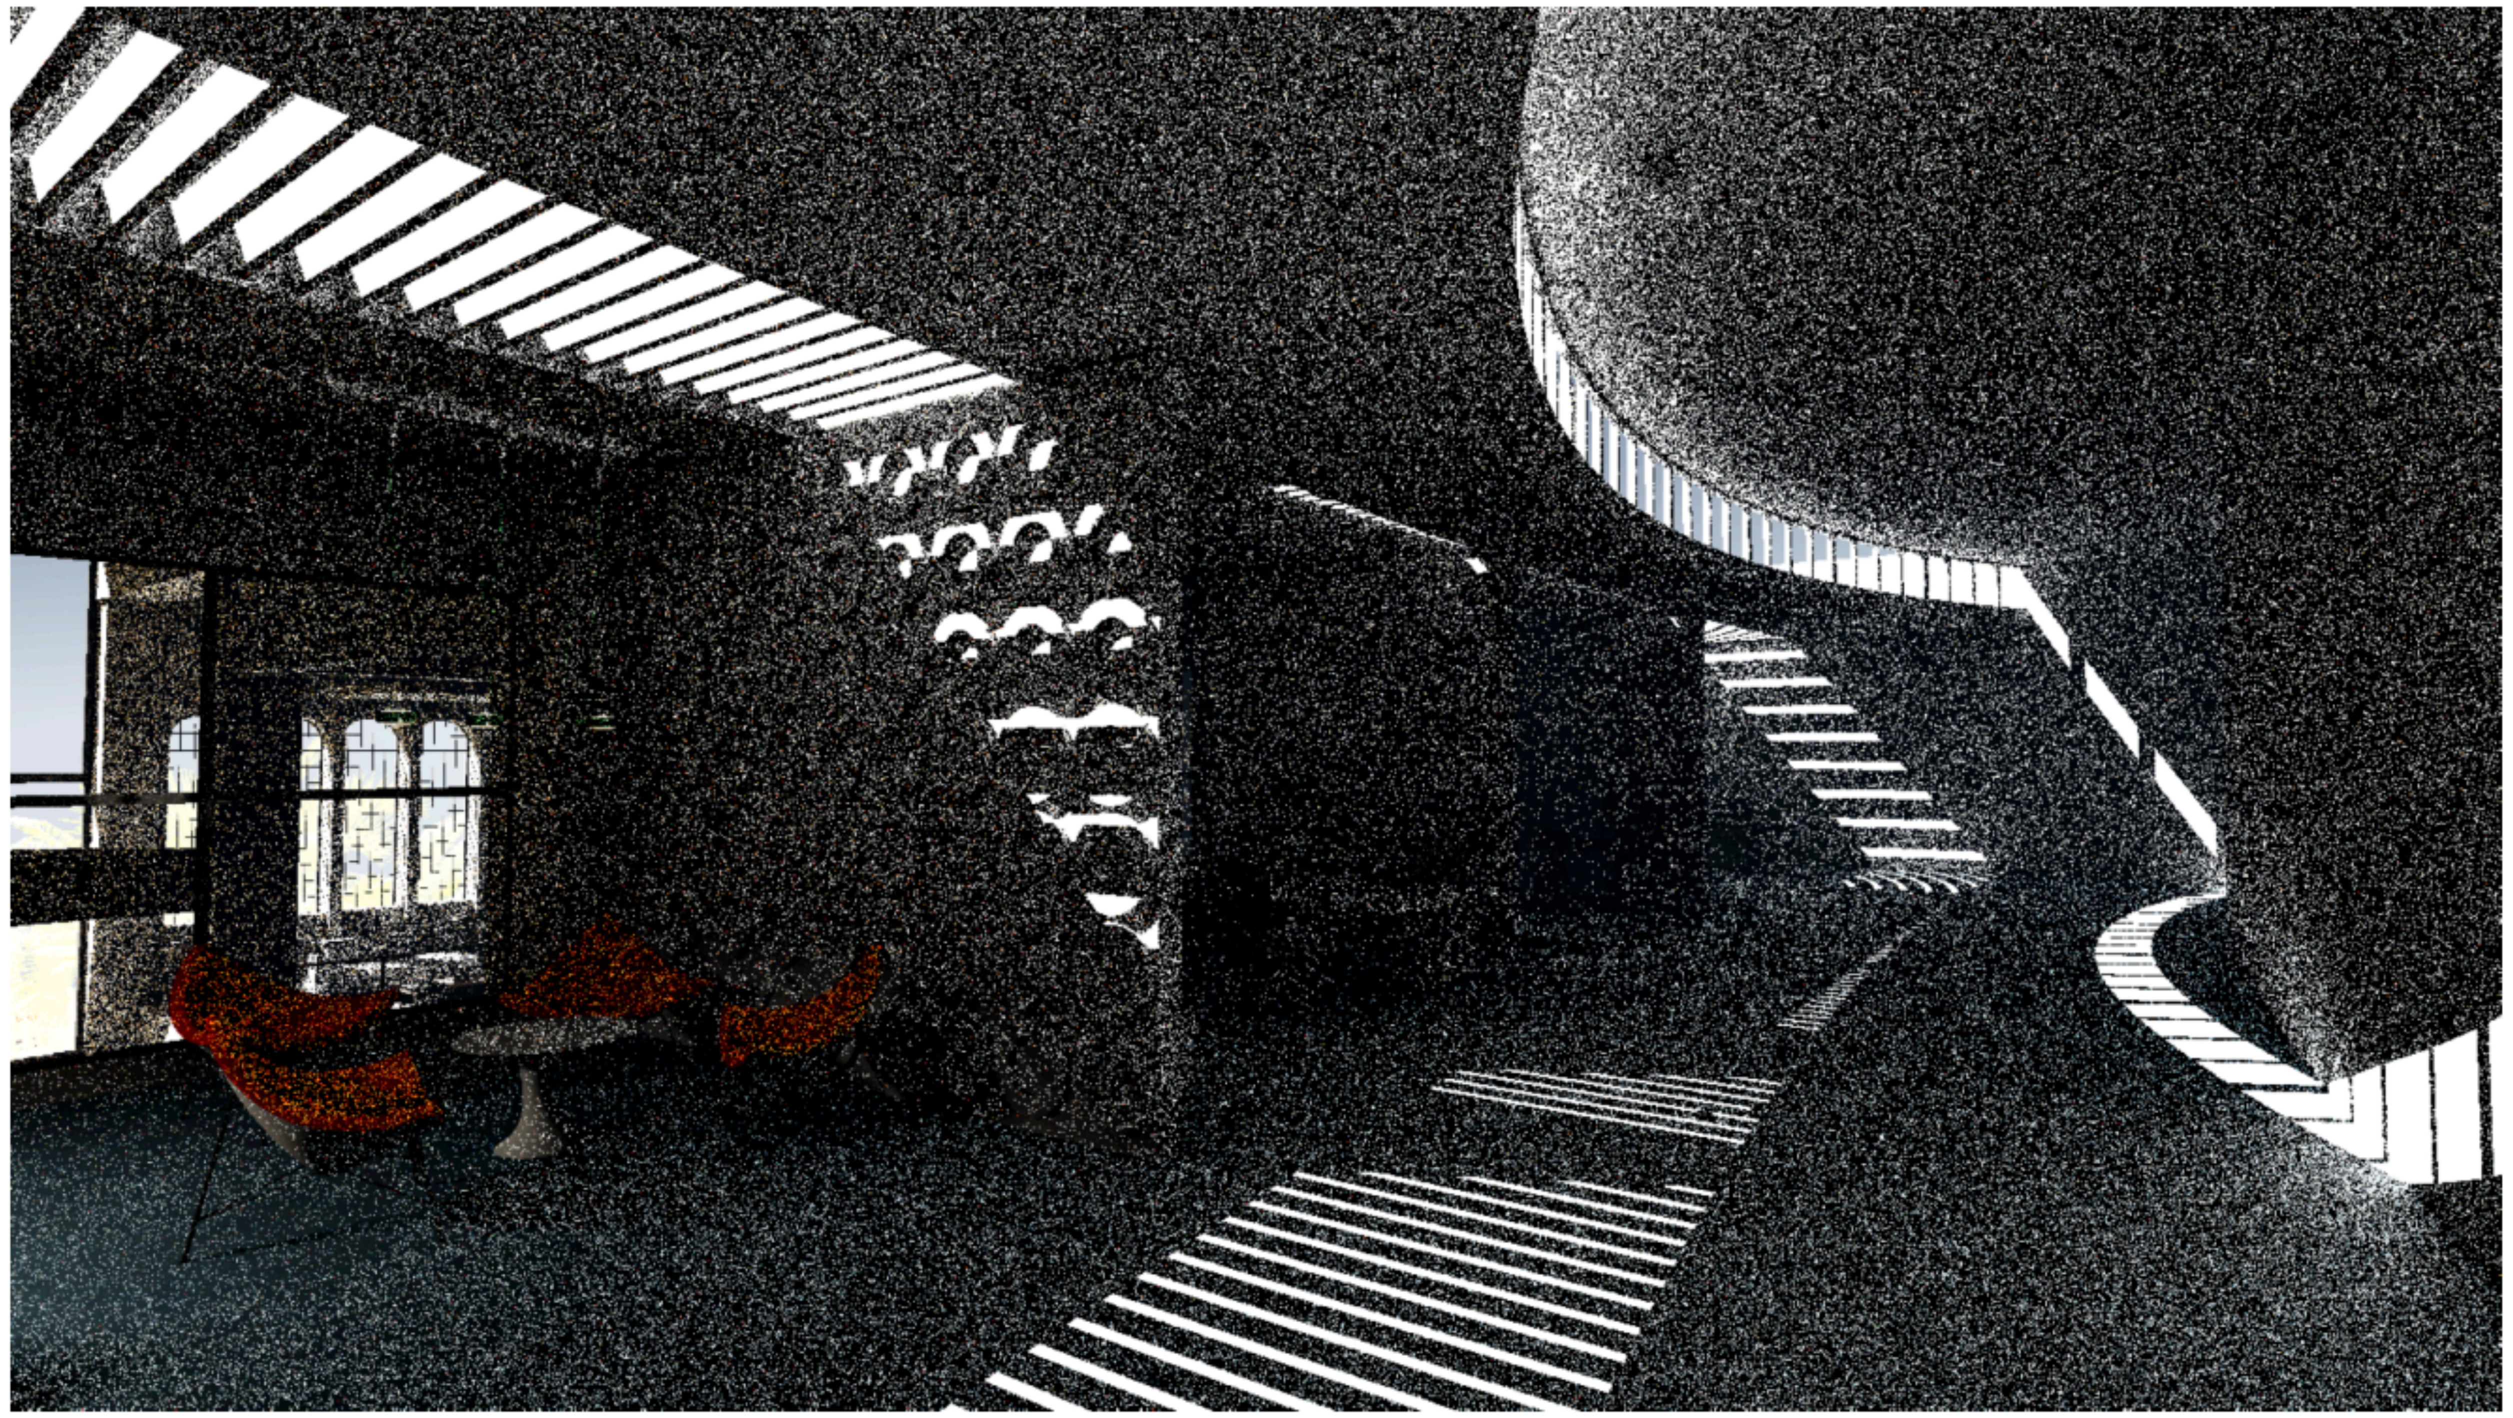

path tracing (8 ms)

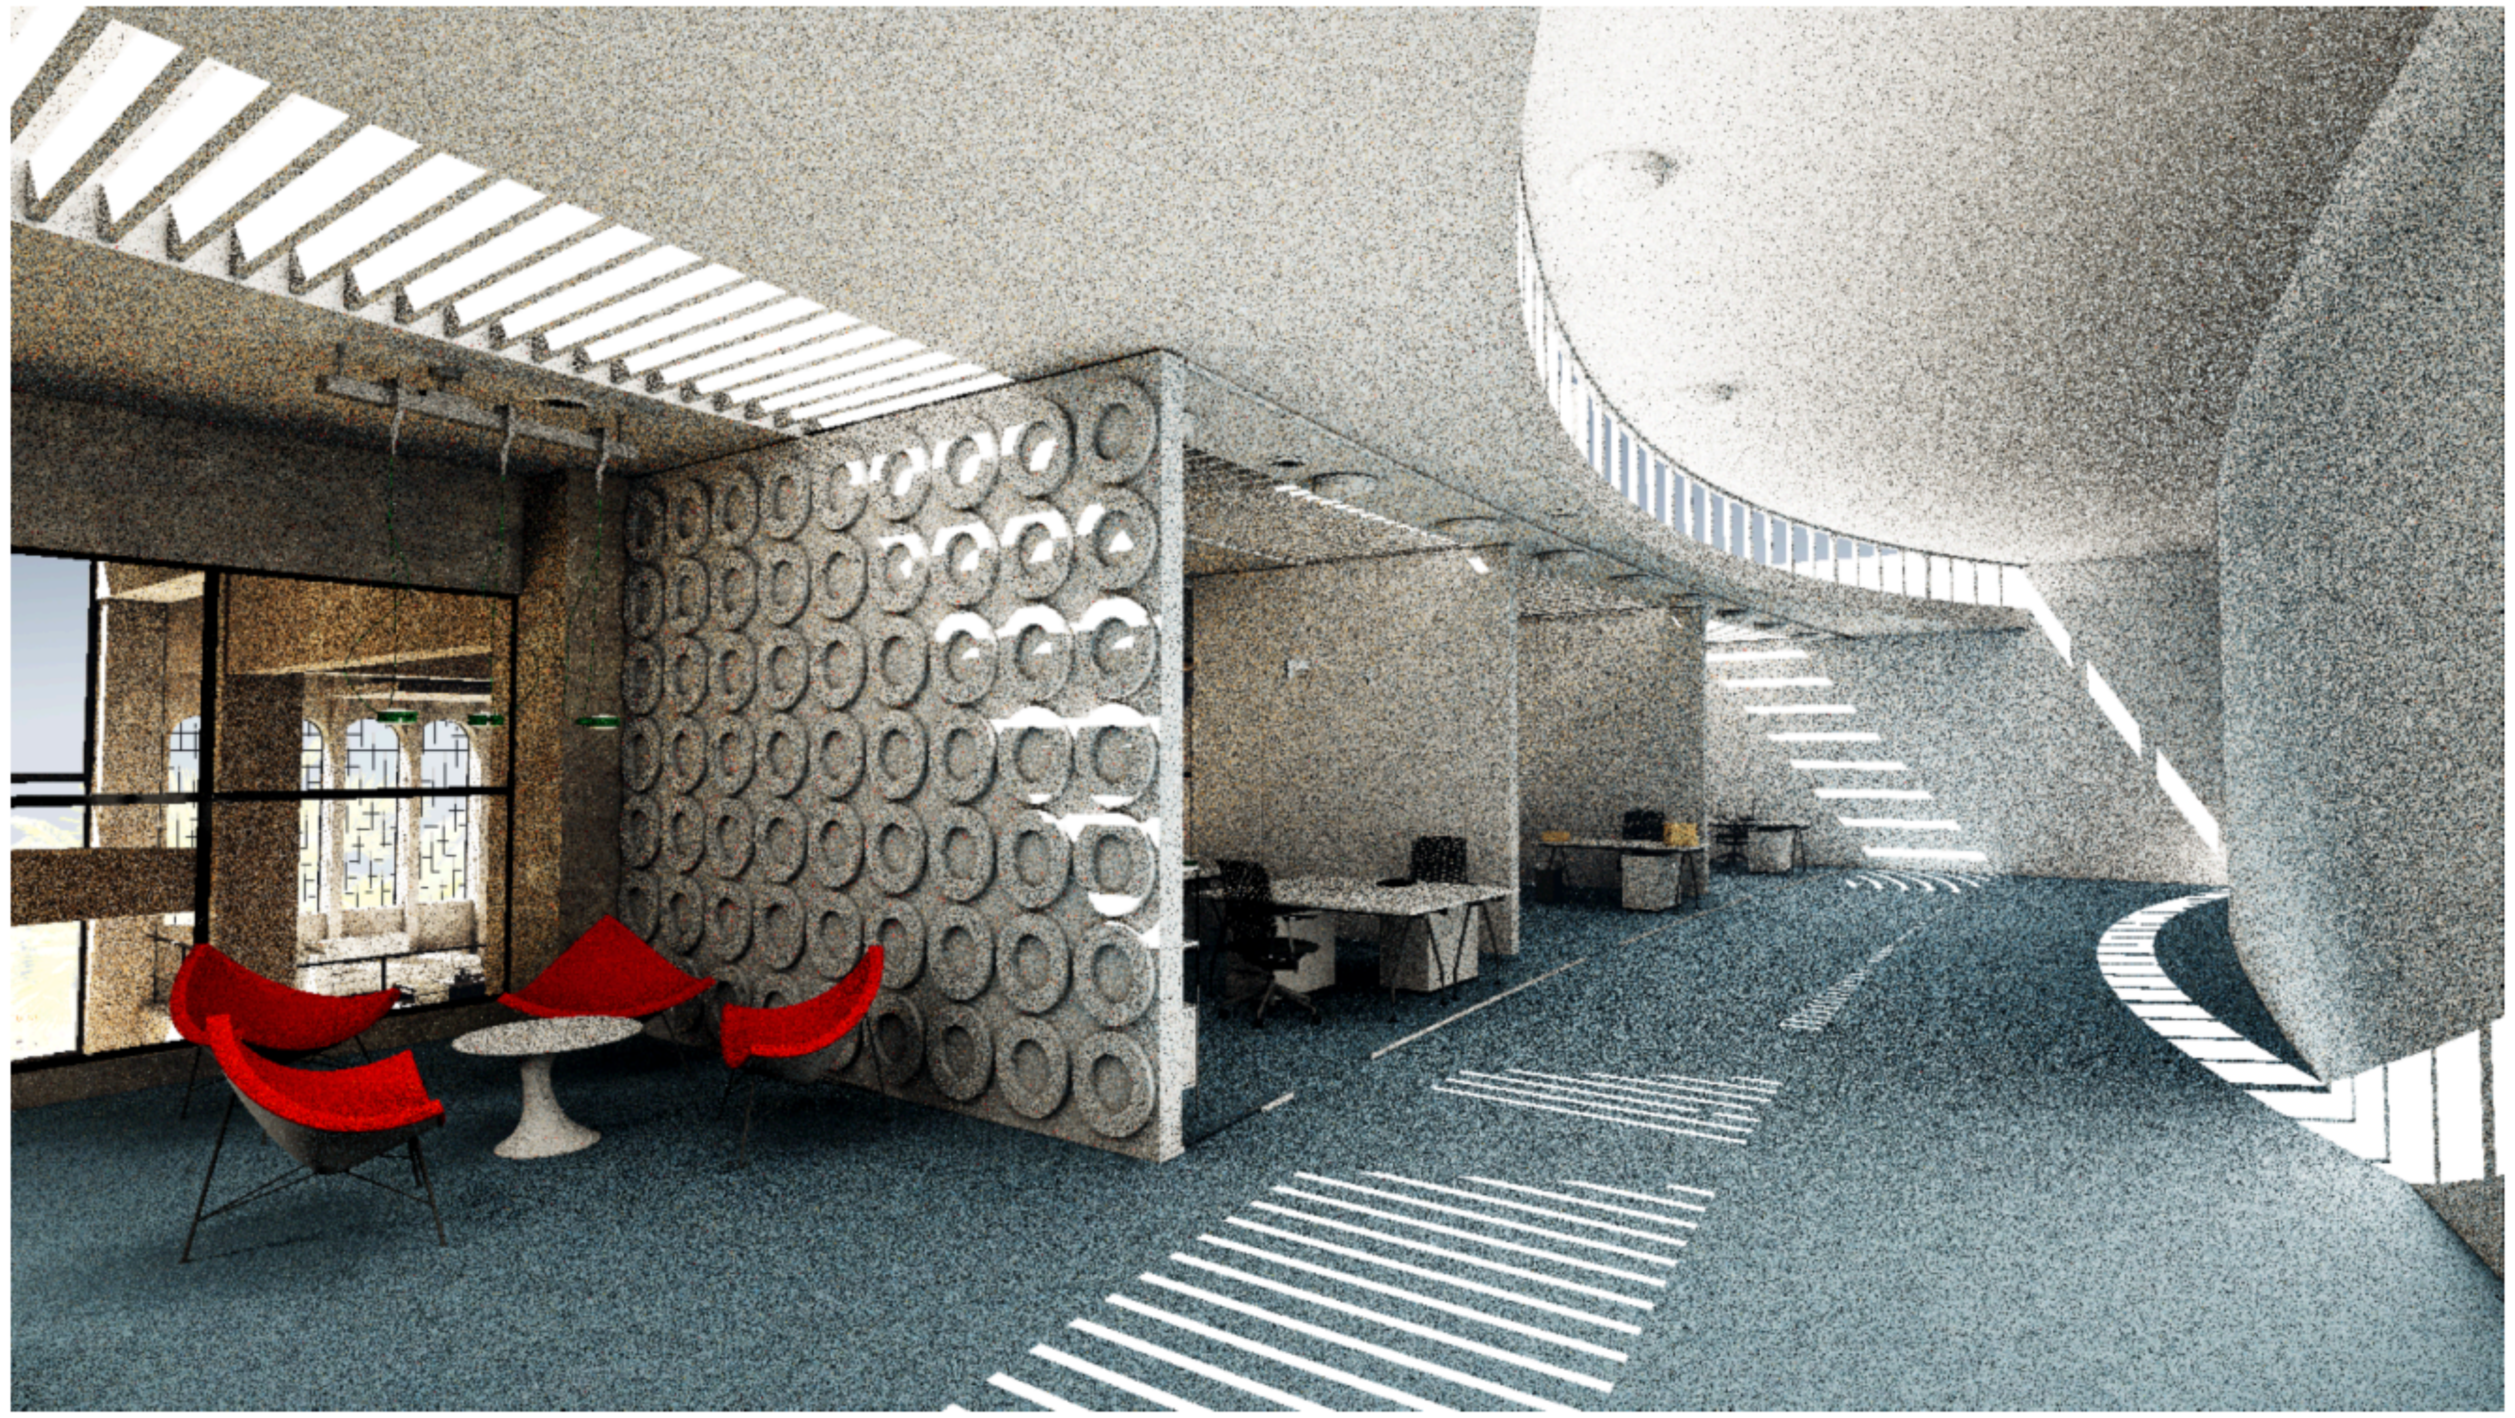

ReSTIR GI (8.9 ms)

Extending ReSTIR to handle GI

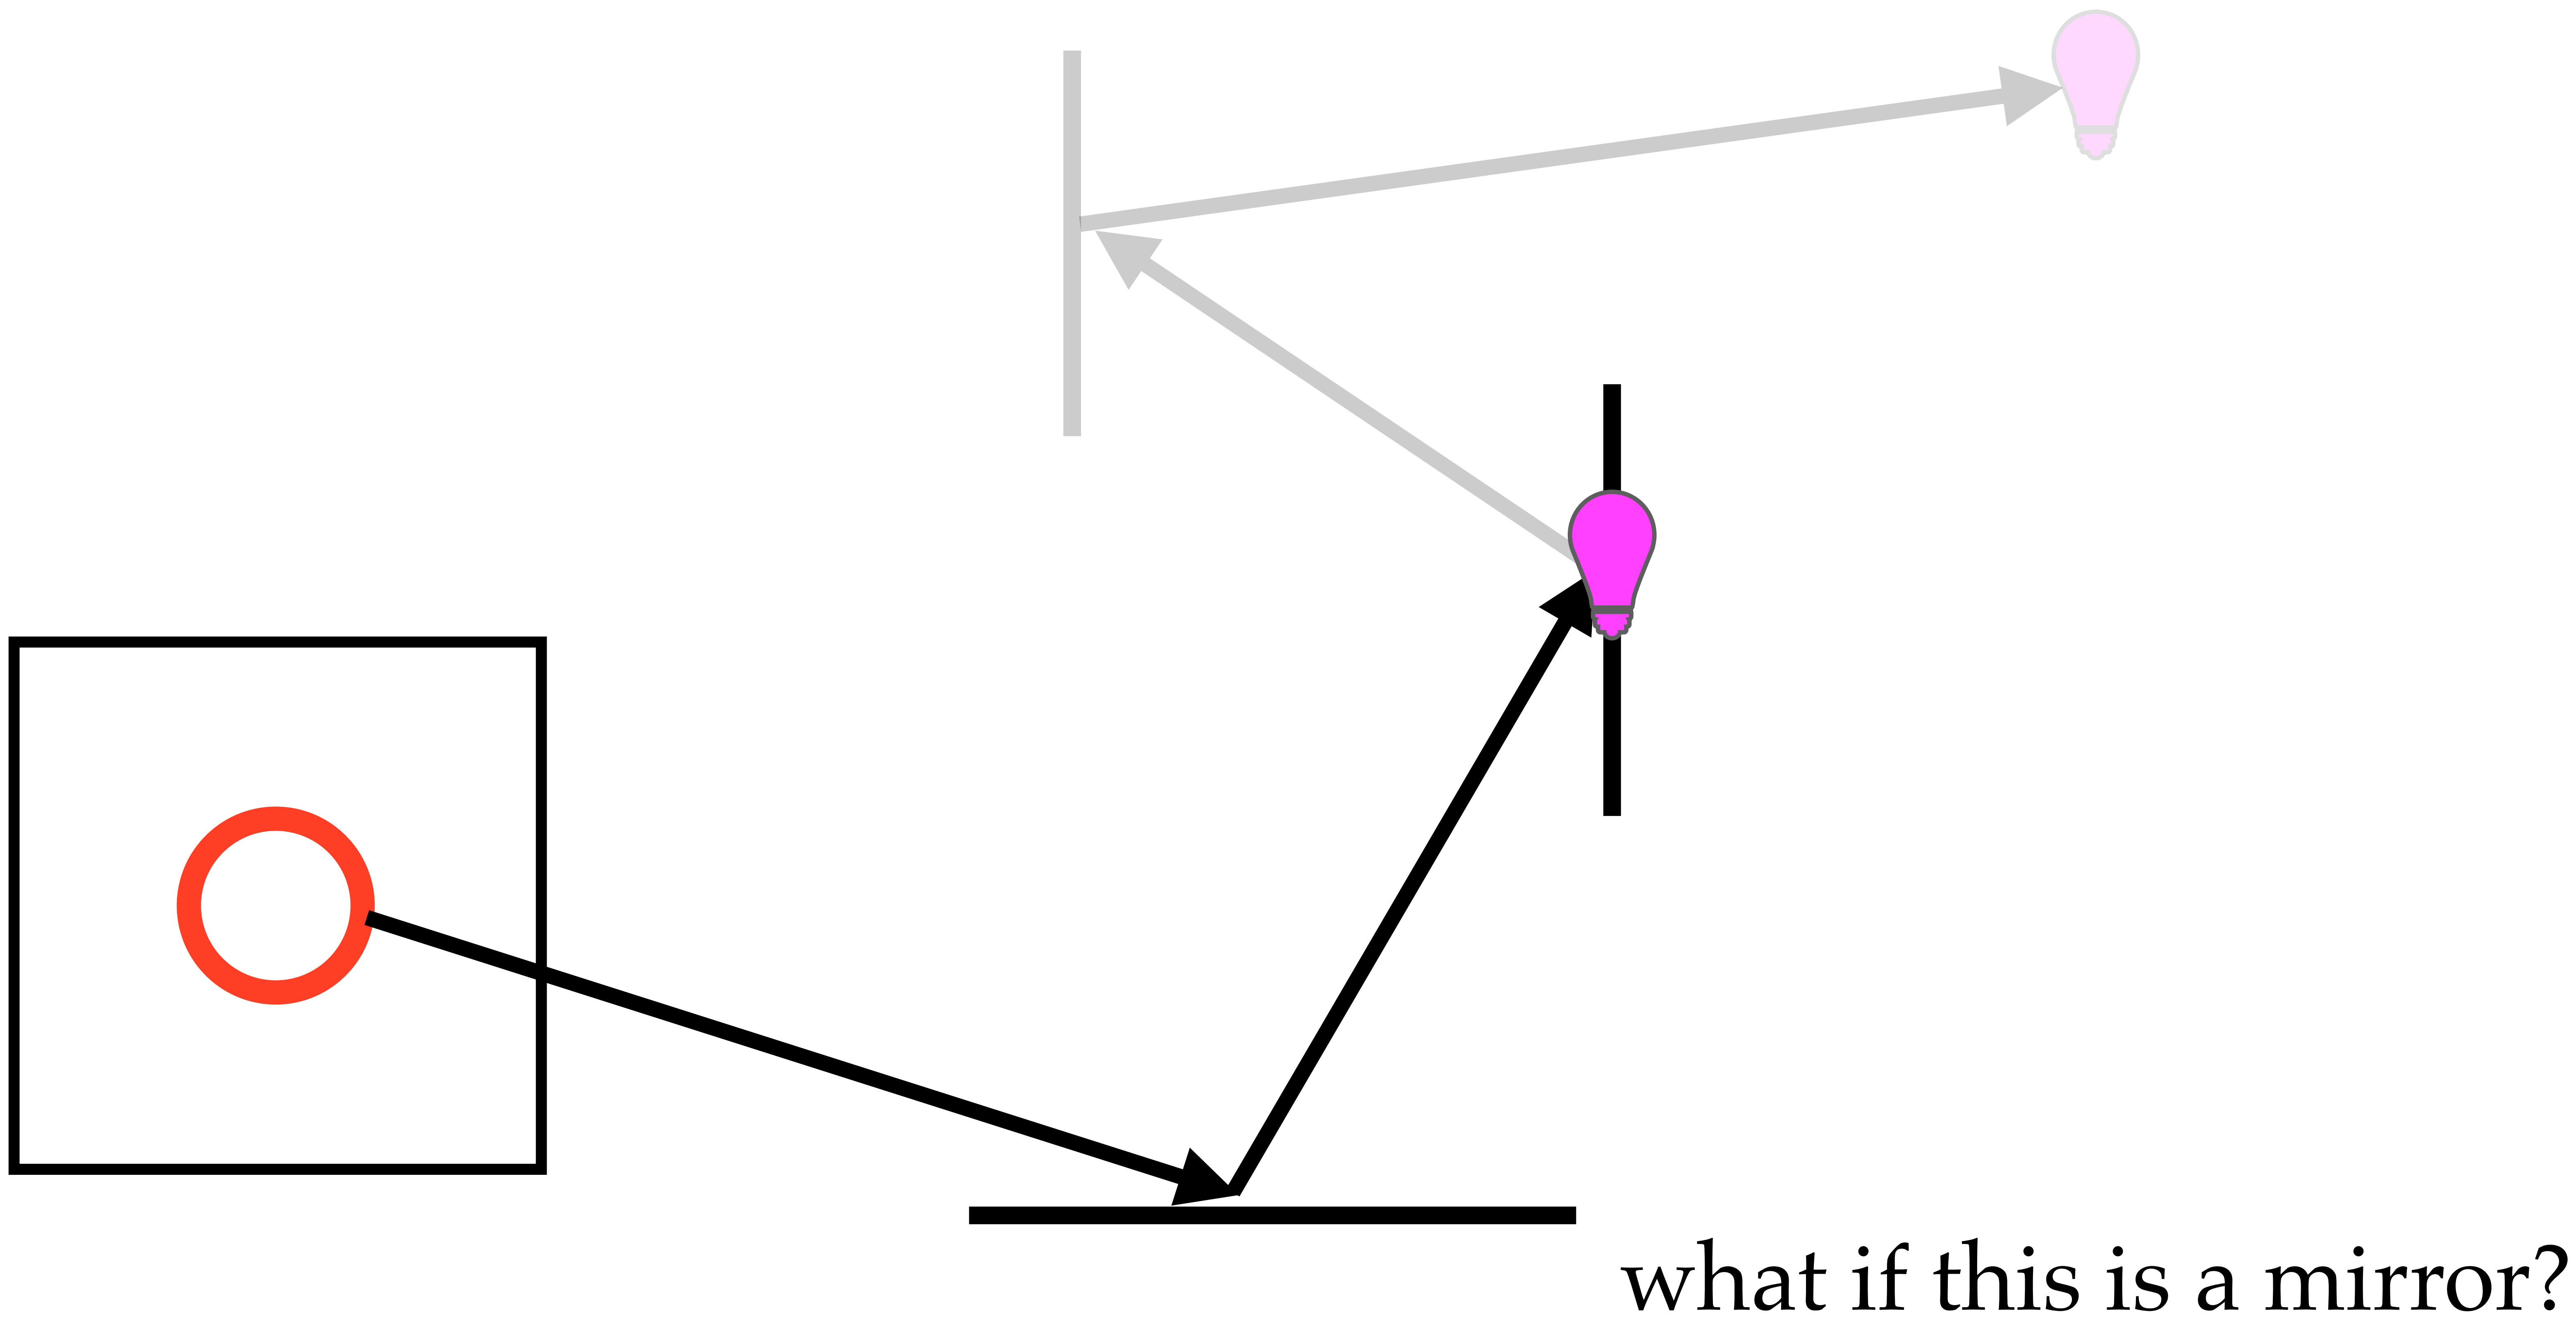

Extending ReSTIR to handle GI

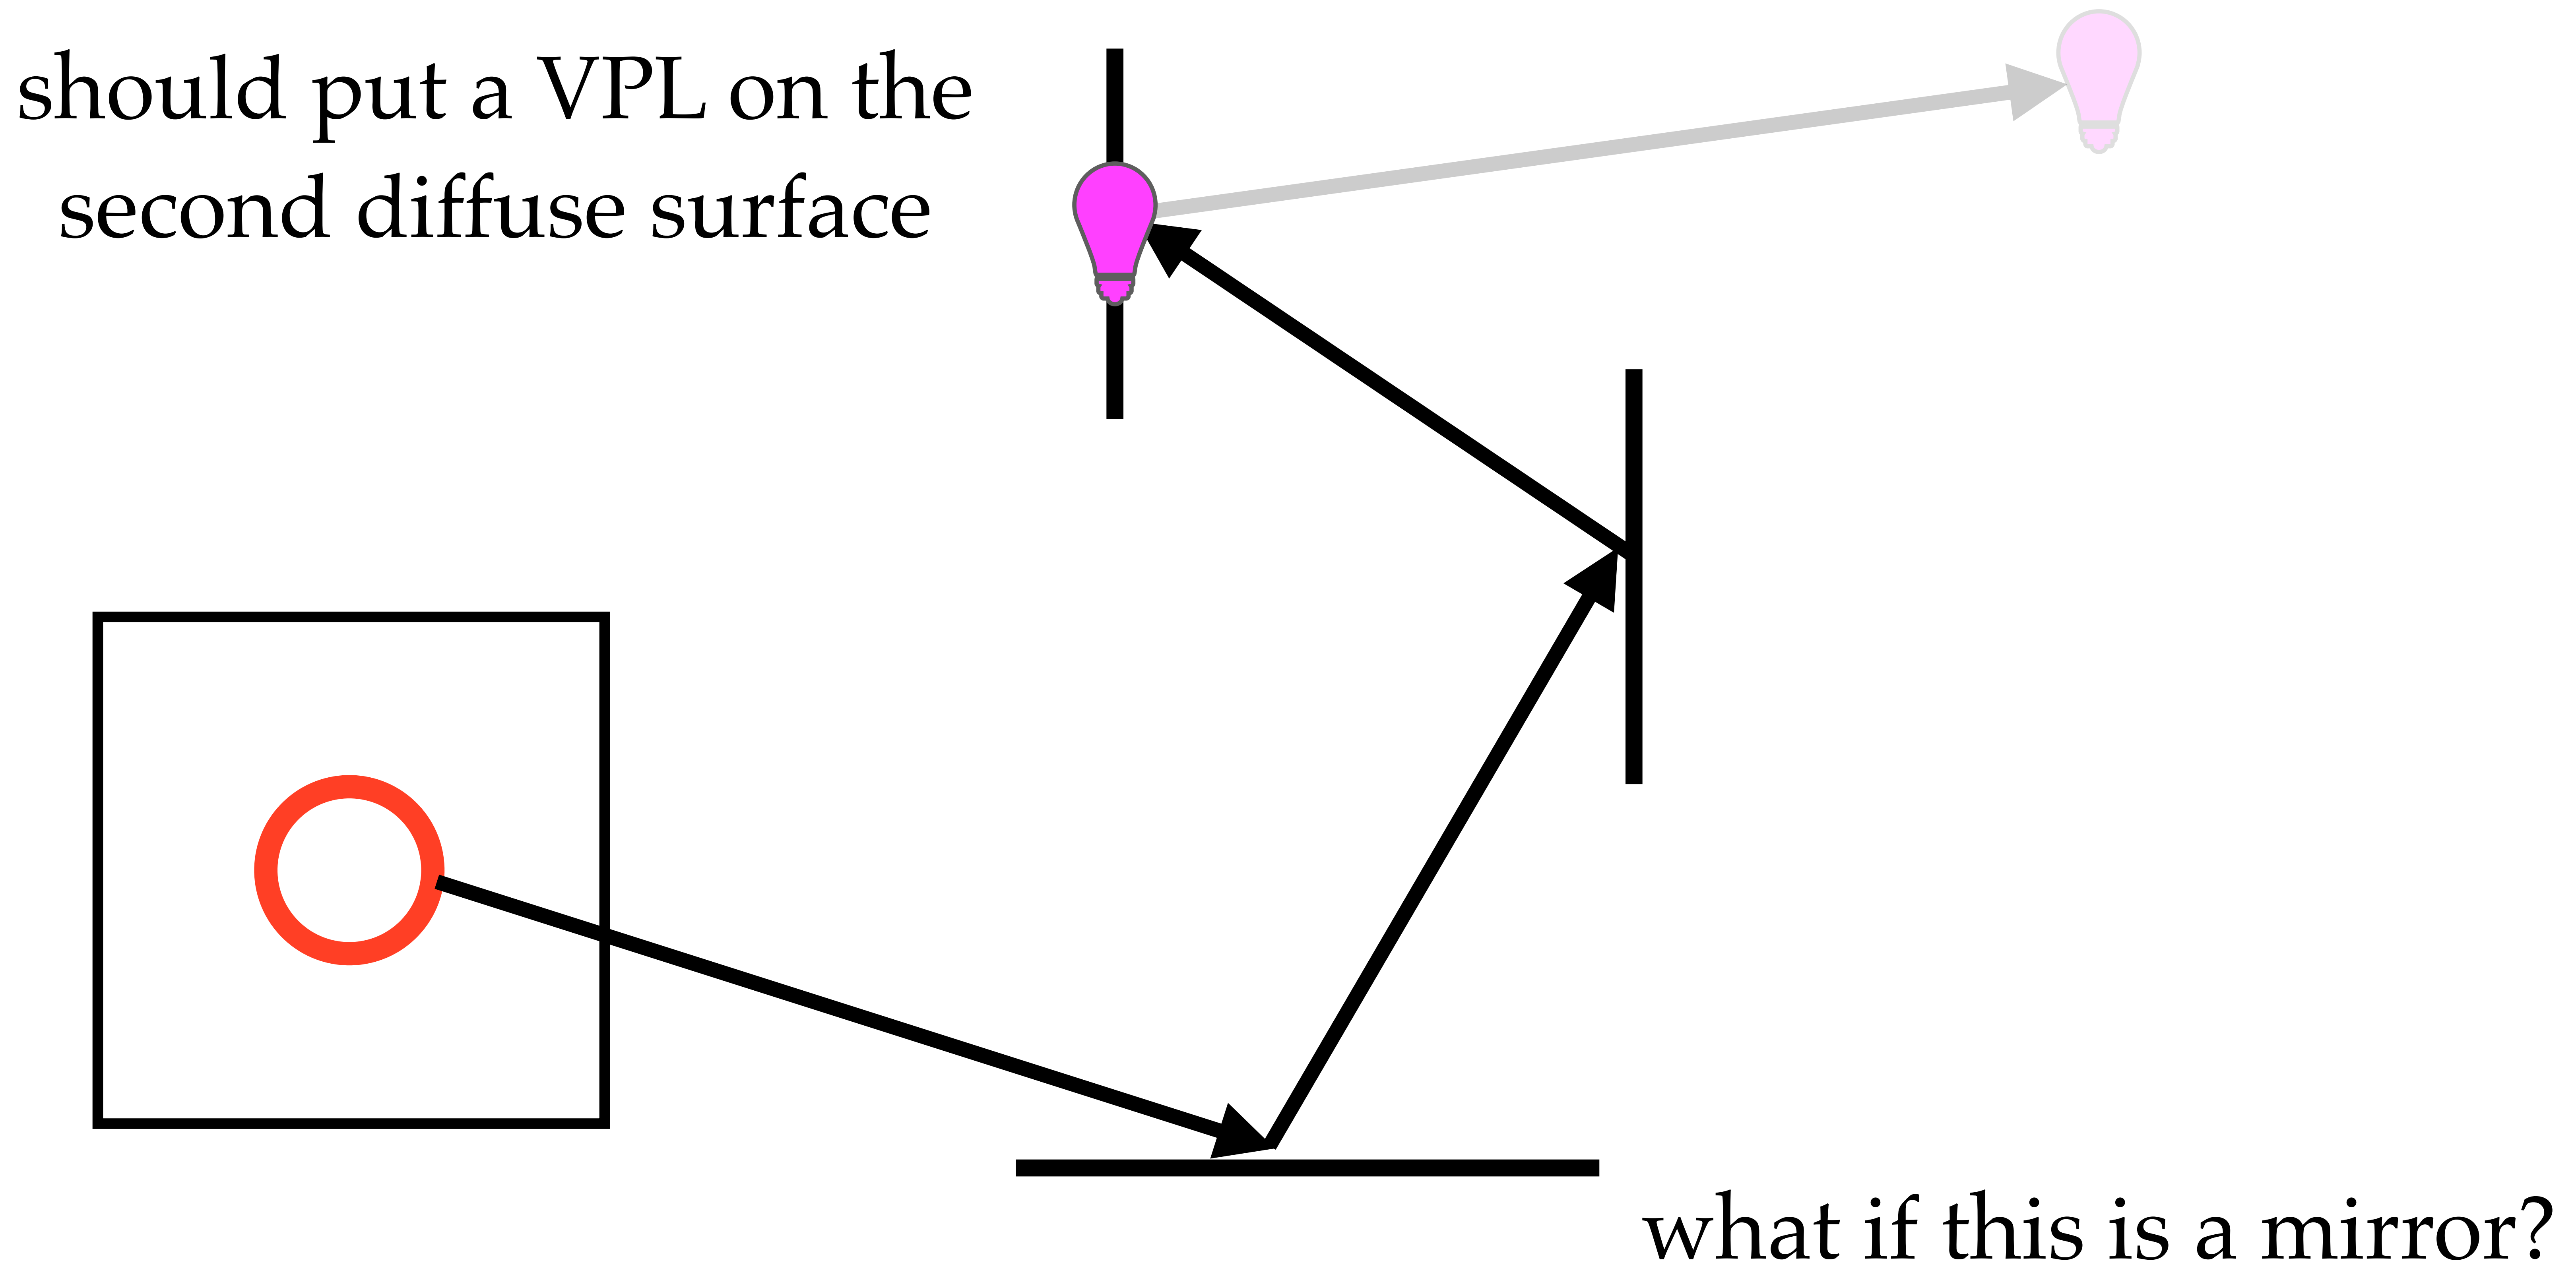

Extending ReSTIR to handle GI

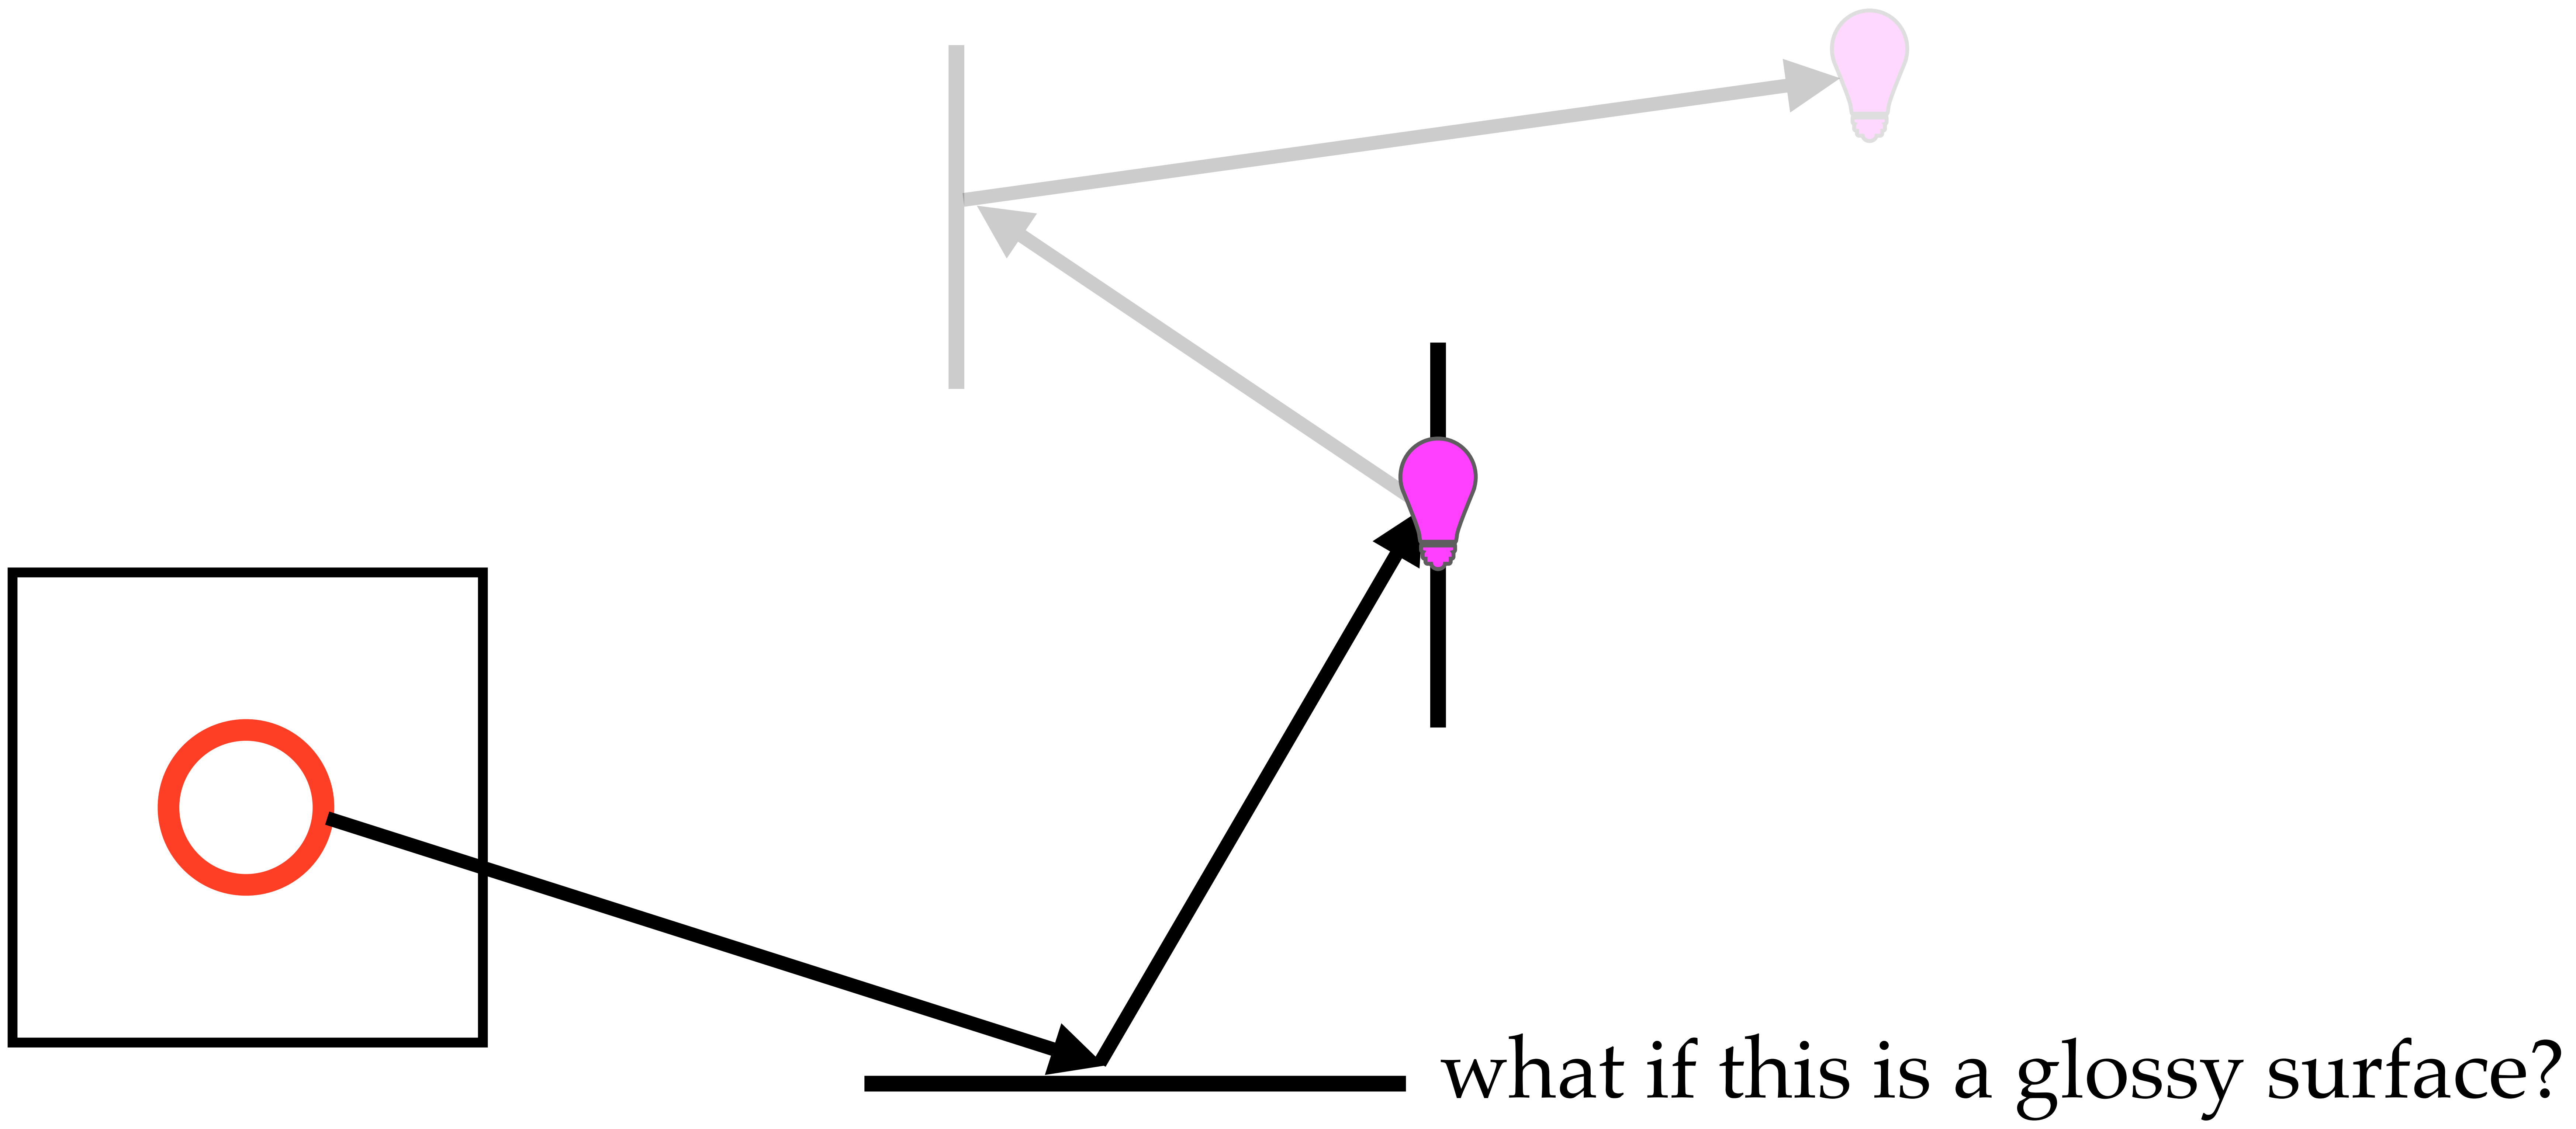

Extending ReSTIR to handle GI

general formulation: we want to find a “shift mapping” to transfer paths between pixels

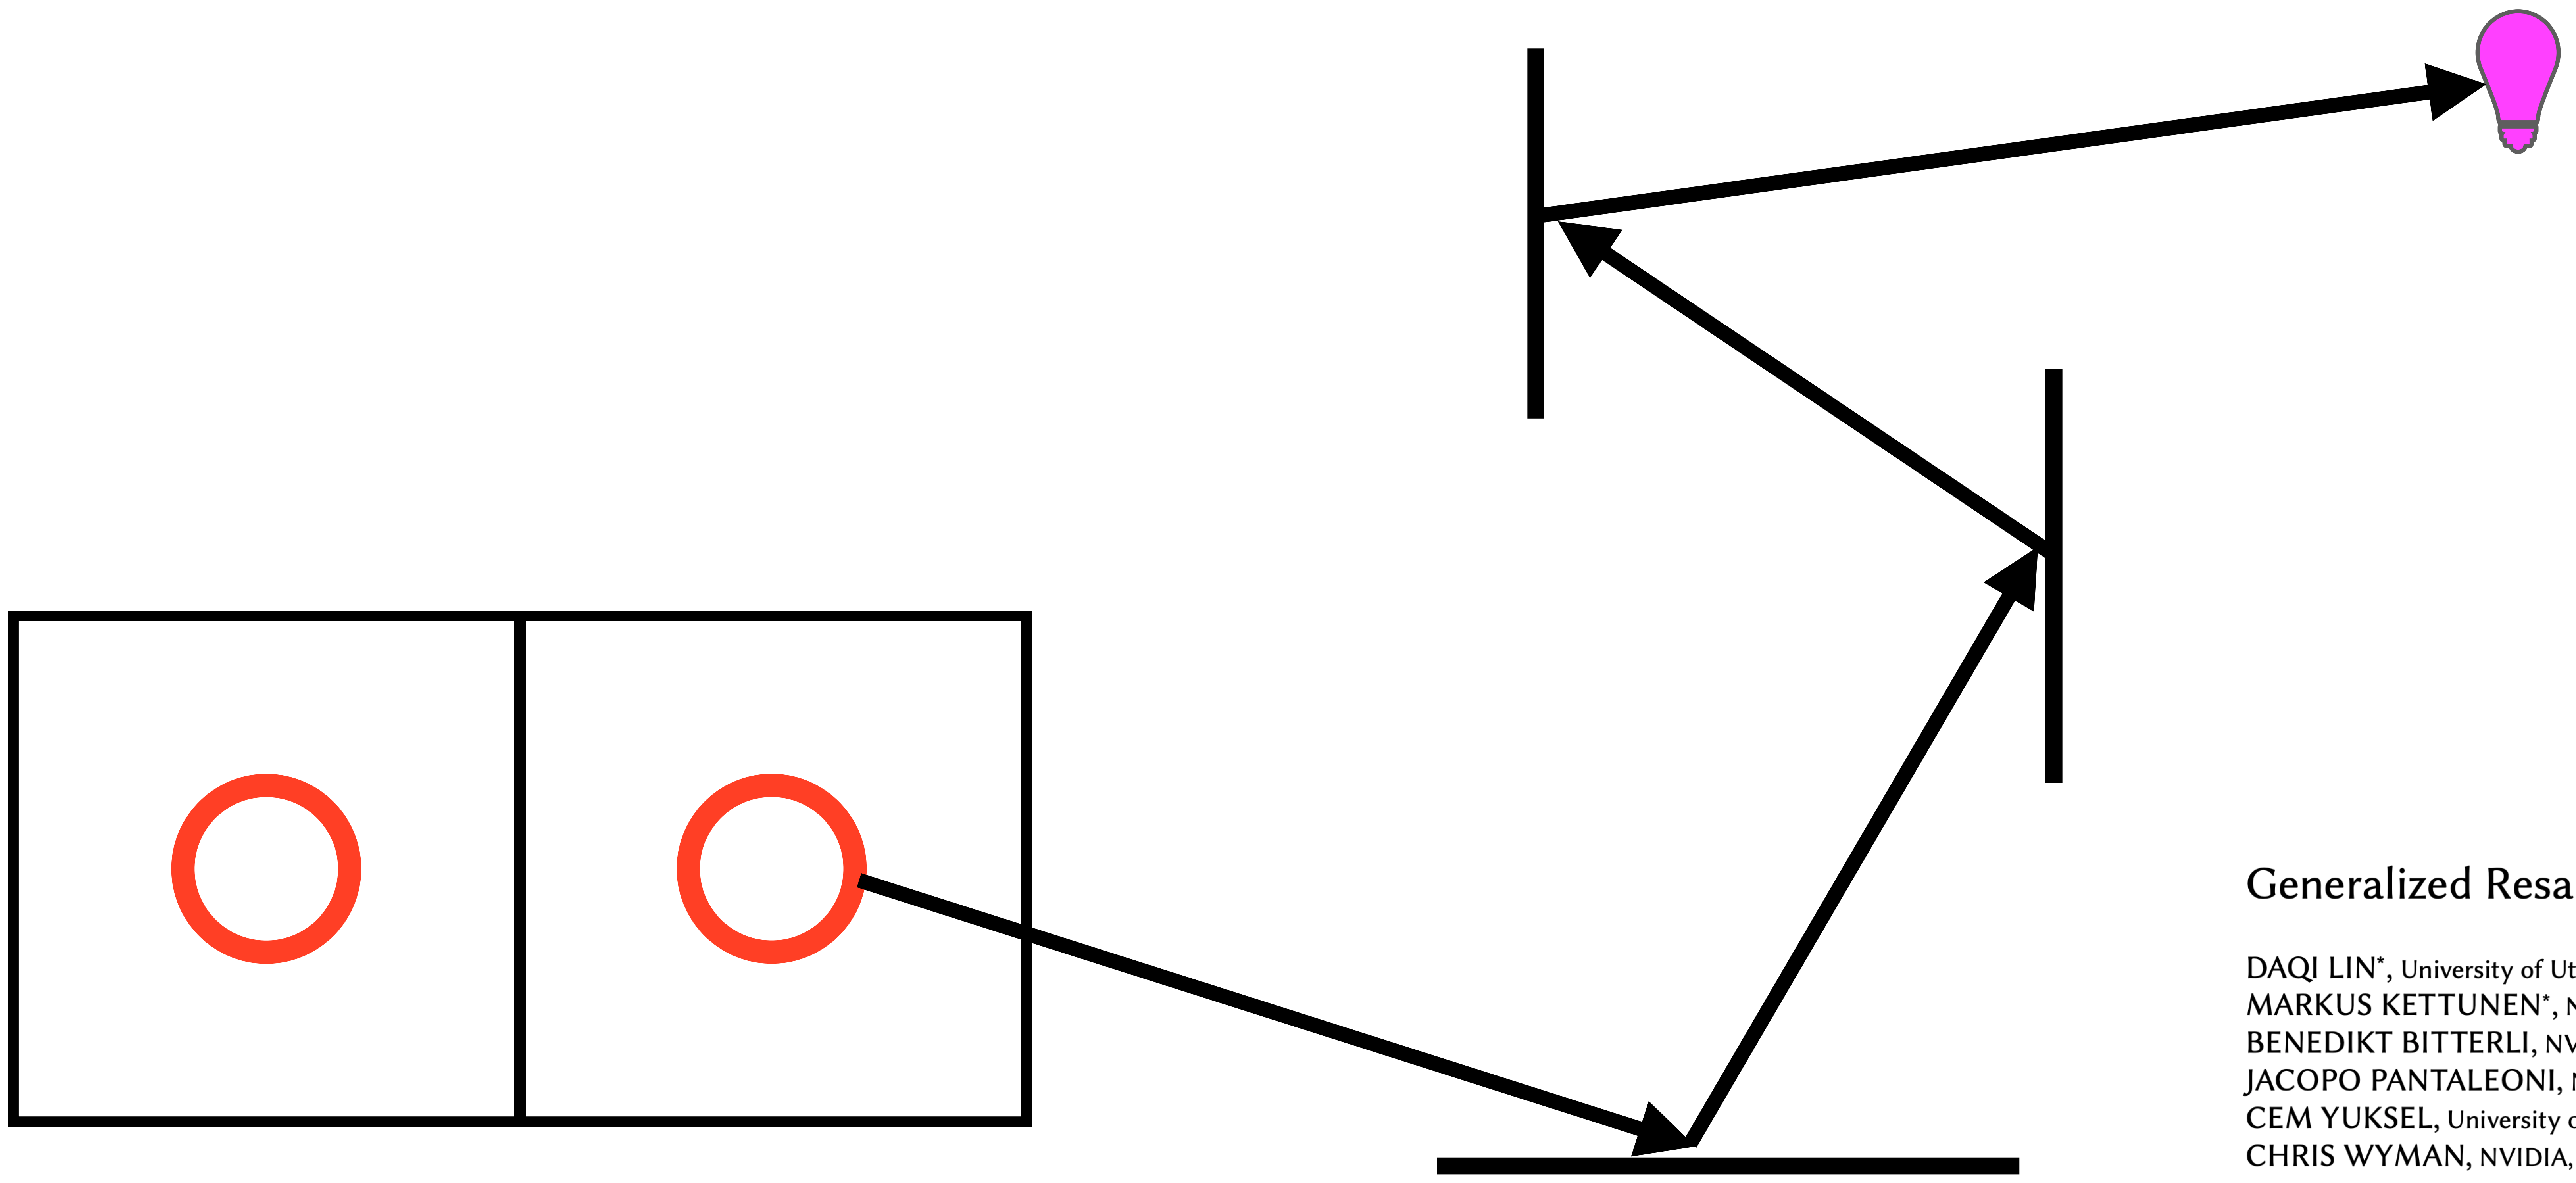

Generalized Resampled Importance Sampling: Foundations of ReSTIR

DAQI LIN*, University of Utah, USA
MARKUS KETTUNEN*, NVIDIA, Finland
BENEDIKT BITTERLI, NVIDIA, USA
JACOPO PANTALEONI, NVIDIA, Germany
CEM YUKSEL, University of Utah, USA
CHRIS WYMAN, NVIDIA, USA

Gradient-Domain Metropolis Light Transport

Jaakko Lehtinen[1,2] Tero Karras[1] Samuli Laine[1] Miika Aittala[2,1] Frédo Durand[3] Timo Aila[1]

[1]NVIDIA Research [2]Aalto University [3]MIT CSAIL

Extending ReSTIR to handle GI

general formulation: we want to find a “shift mapping” to transfer paths between pixels

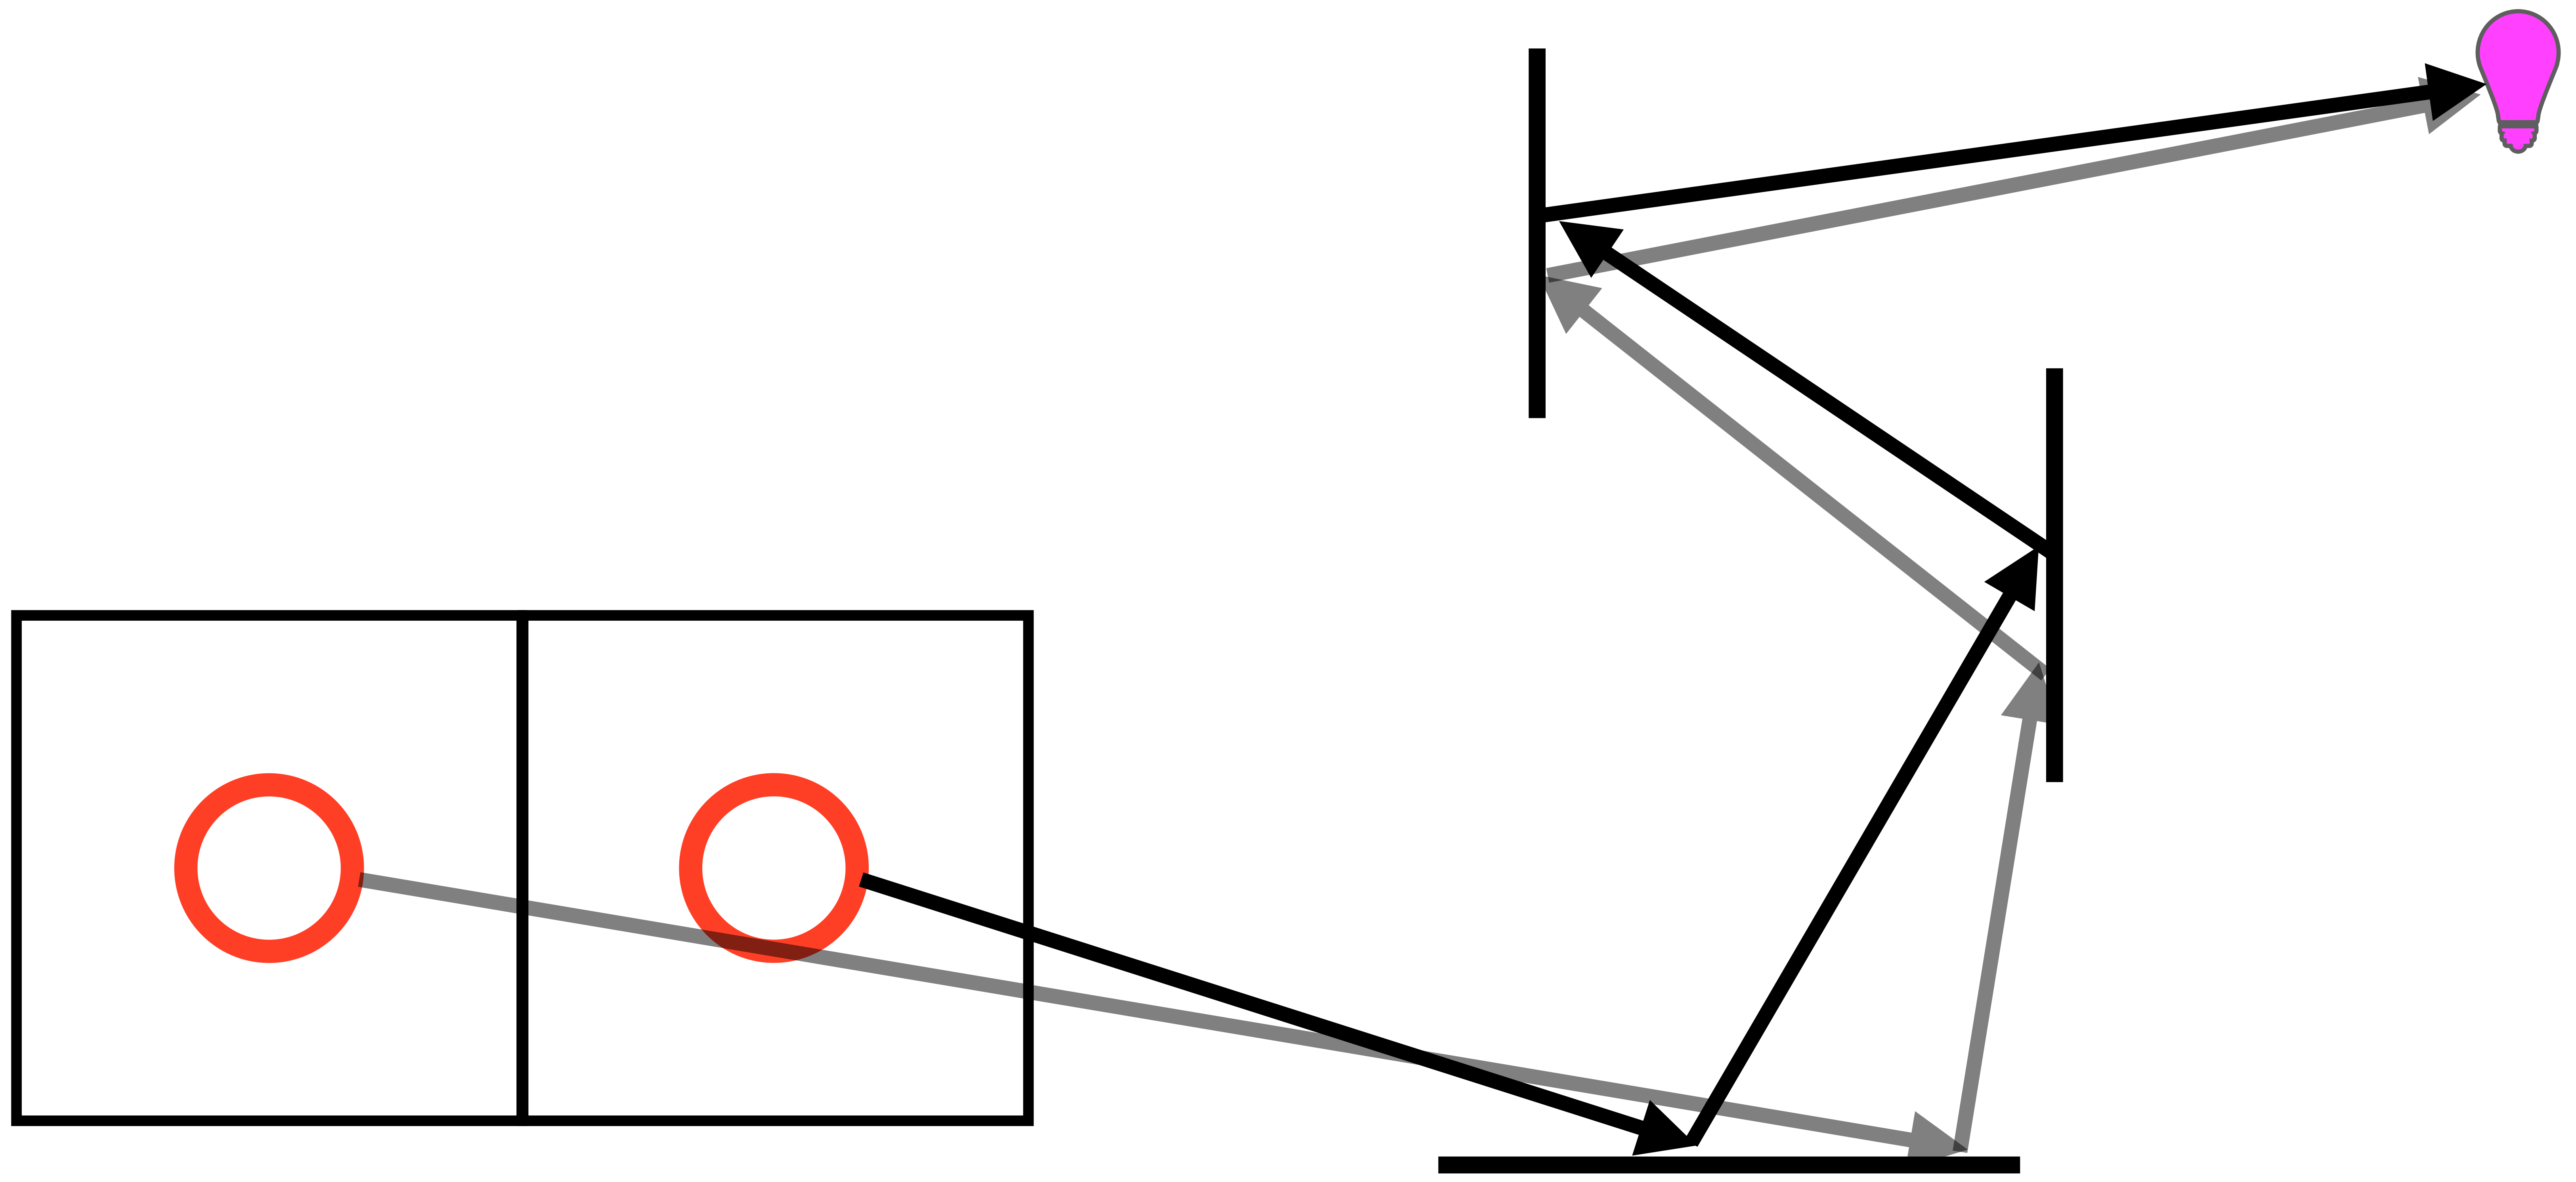

Extending ReSTIR to handle GI

general formulation: we want to find a “shift mapping” to transfer paths between pixels

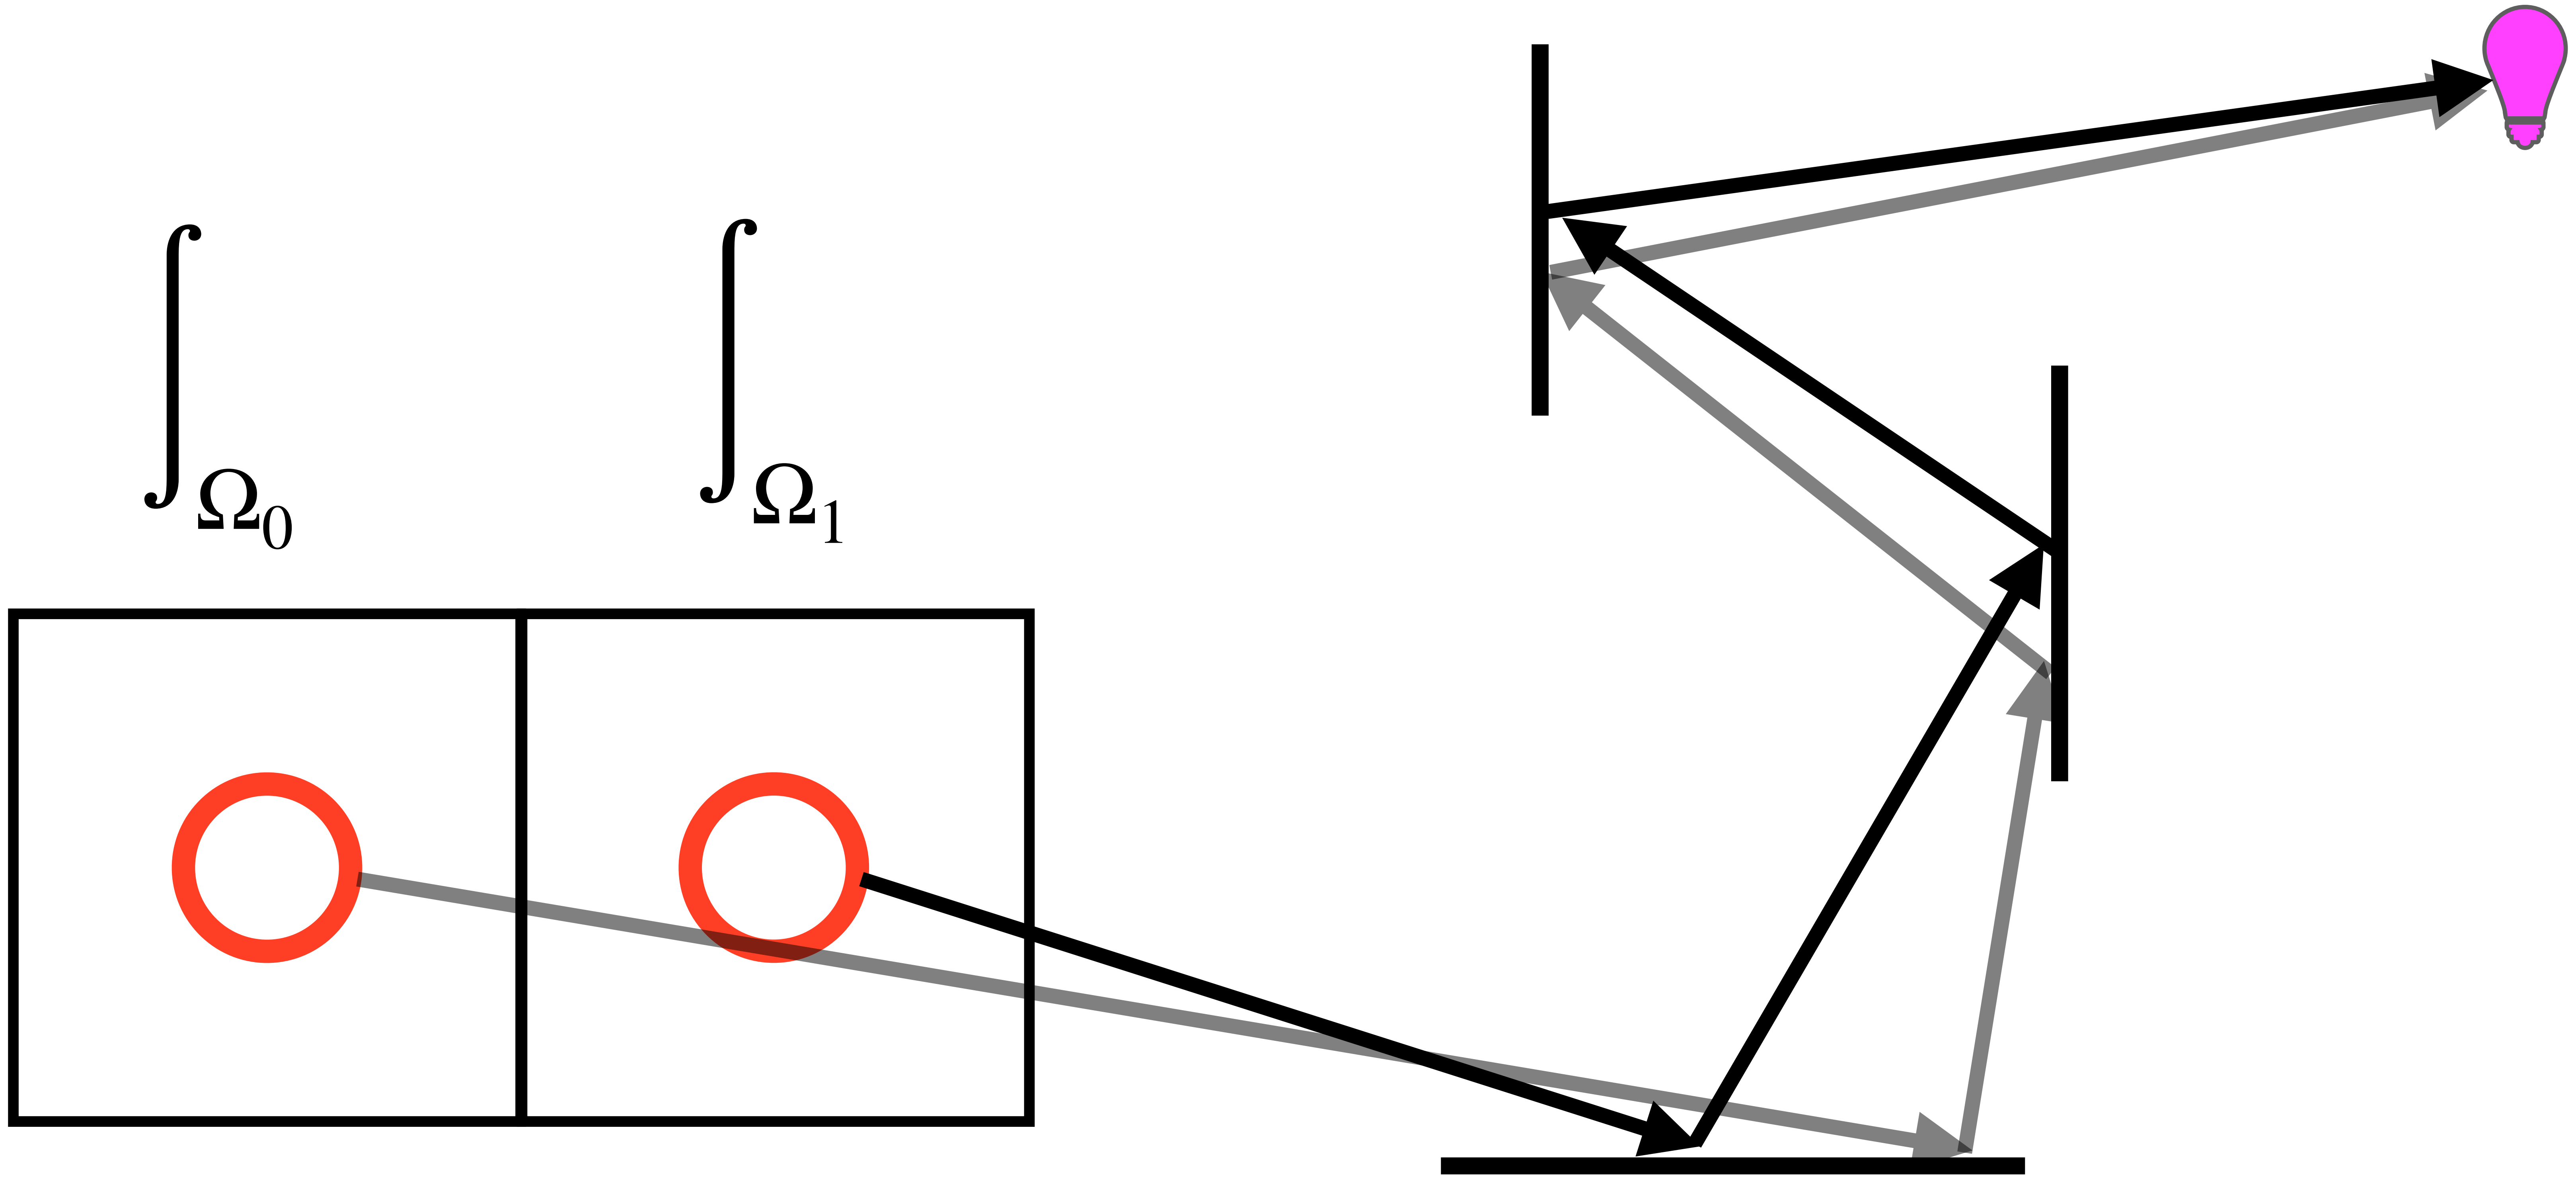

Extending ReSTIR to handle GI

general formulation: we want to find a “shift mapping” to transfer paths between pixels

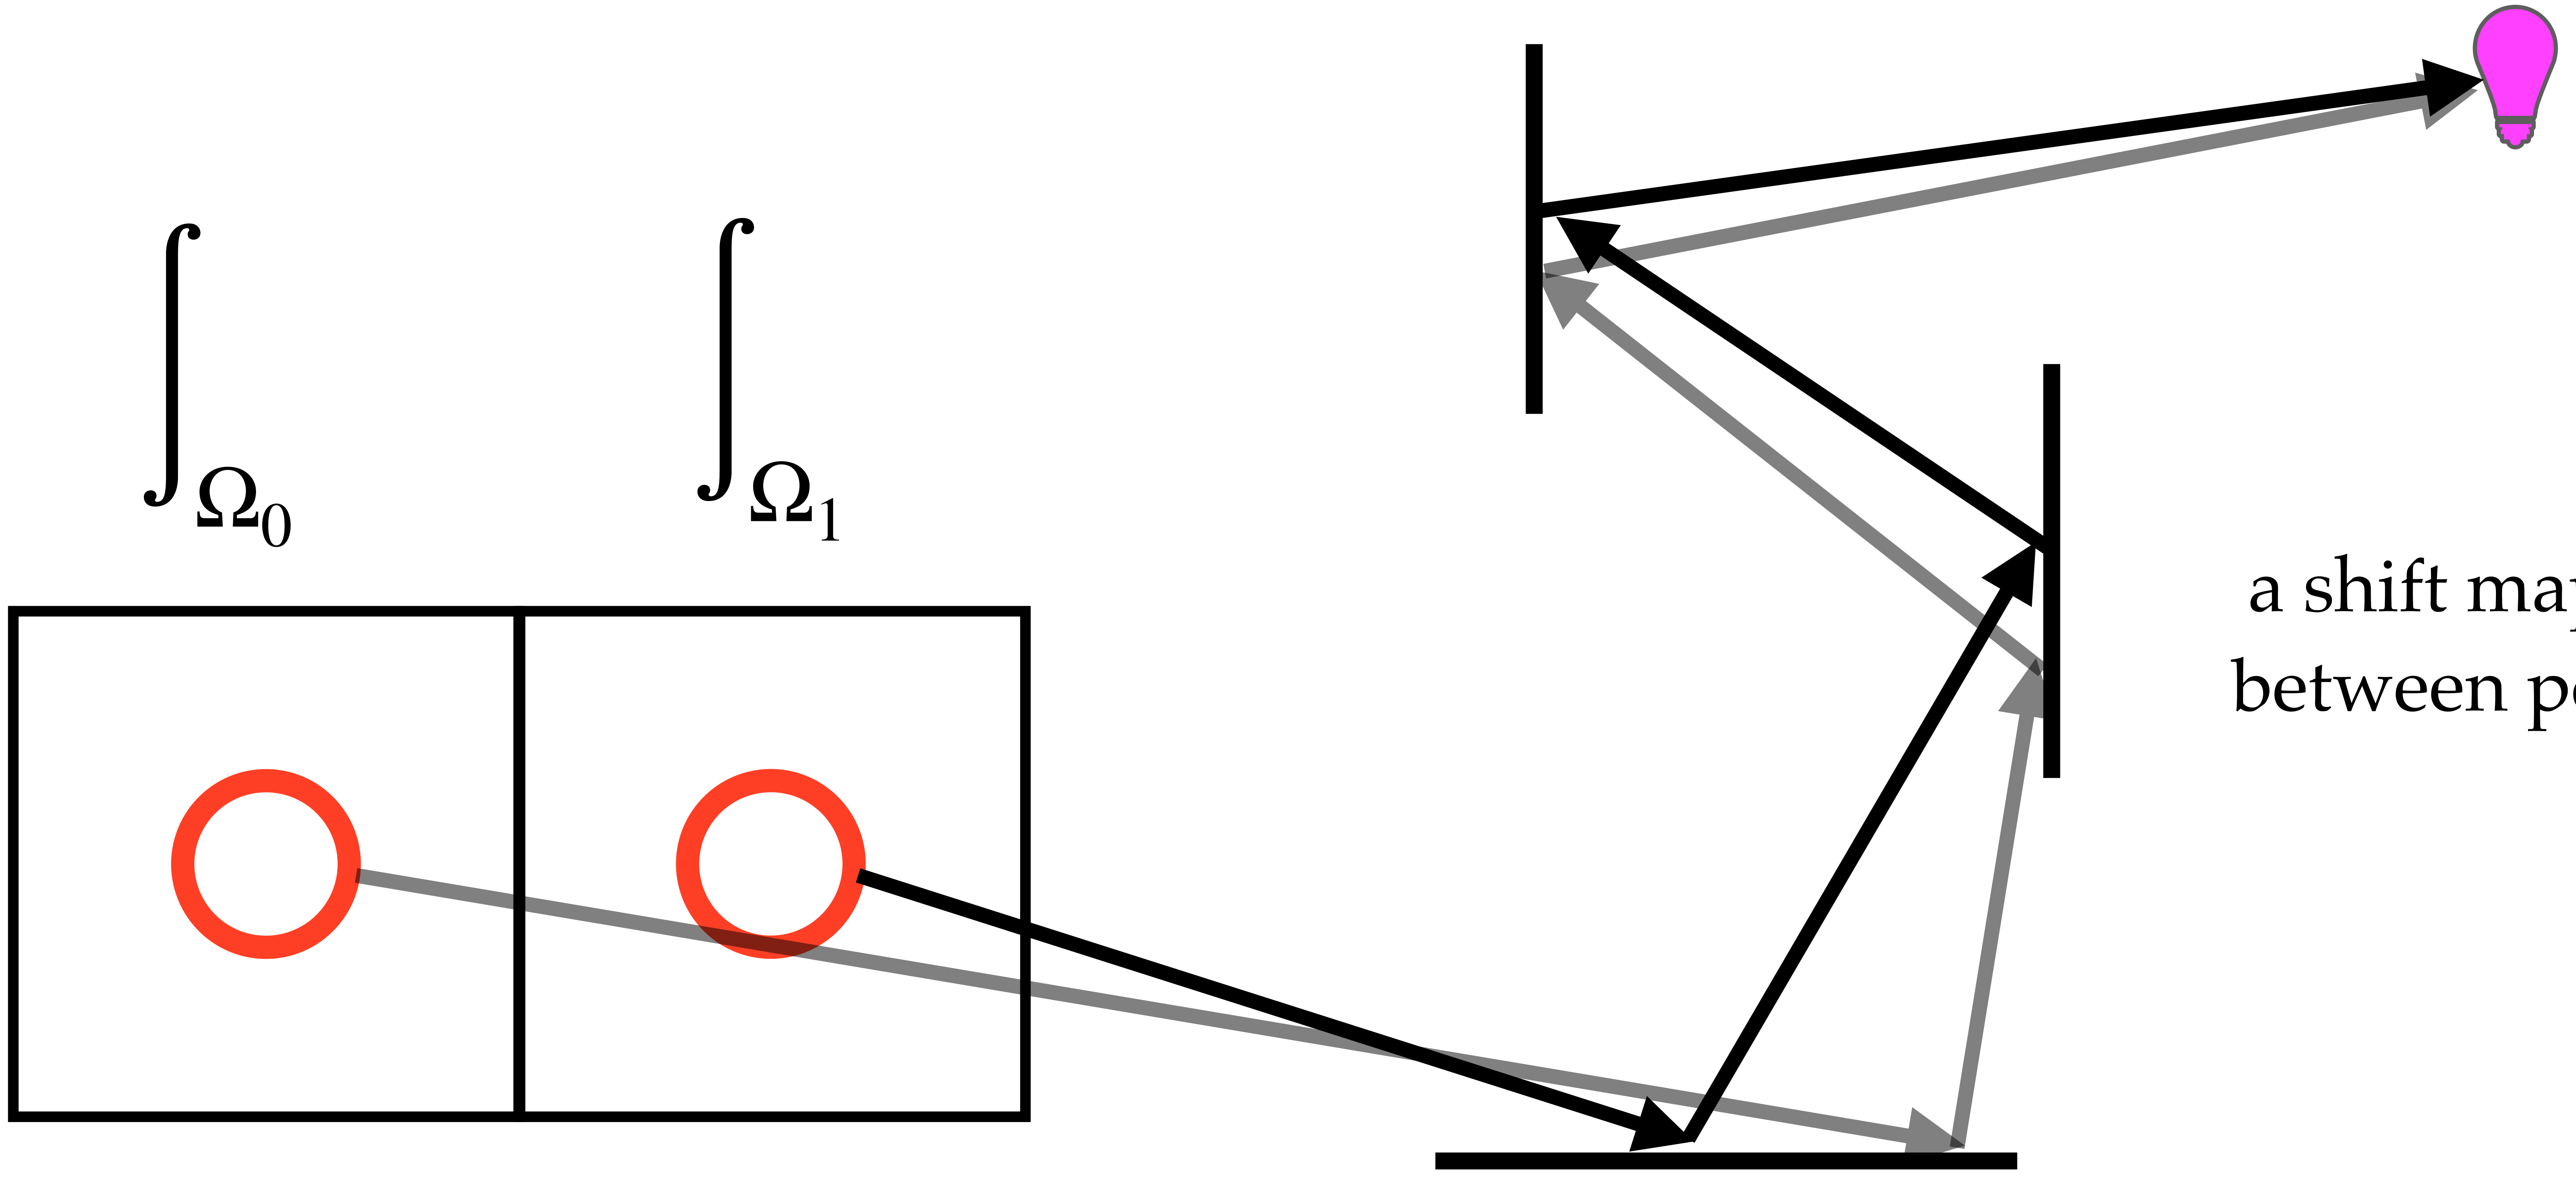

a shift mapping is an invertible mapping between points in the domains Ω_0 and Ω_1

ReSTIR for general GI

path tracing (70 ms)

ReSTIR PT (70 ms)

ReSTIR for general GI

path tracing (80 ms)

ReSTIR PT (80 ms)

Volumetric ReSTIR

- basically the same idea, with shift mapping designed for volume rendering
- requires very careful engineering for high performance

Fast Volume Rendering with Spatiotemporal Reservoir Resampling

DAQI LIN, University of Utah
CHRIS WYMAN, NVIDIA
CEM YUKSEL, University of Utah

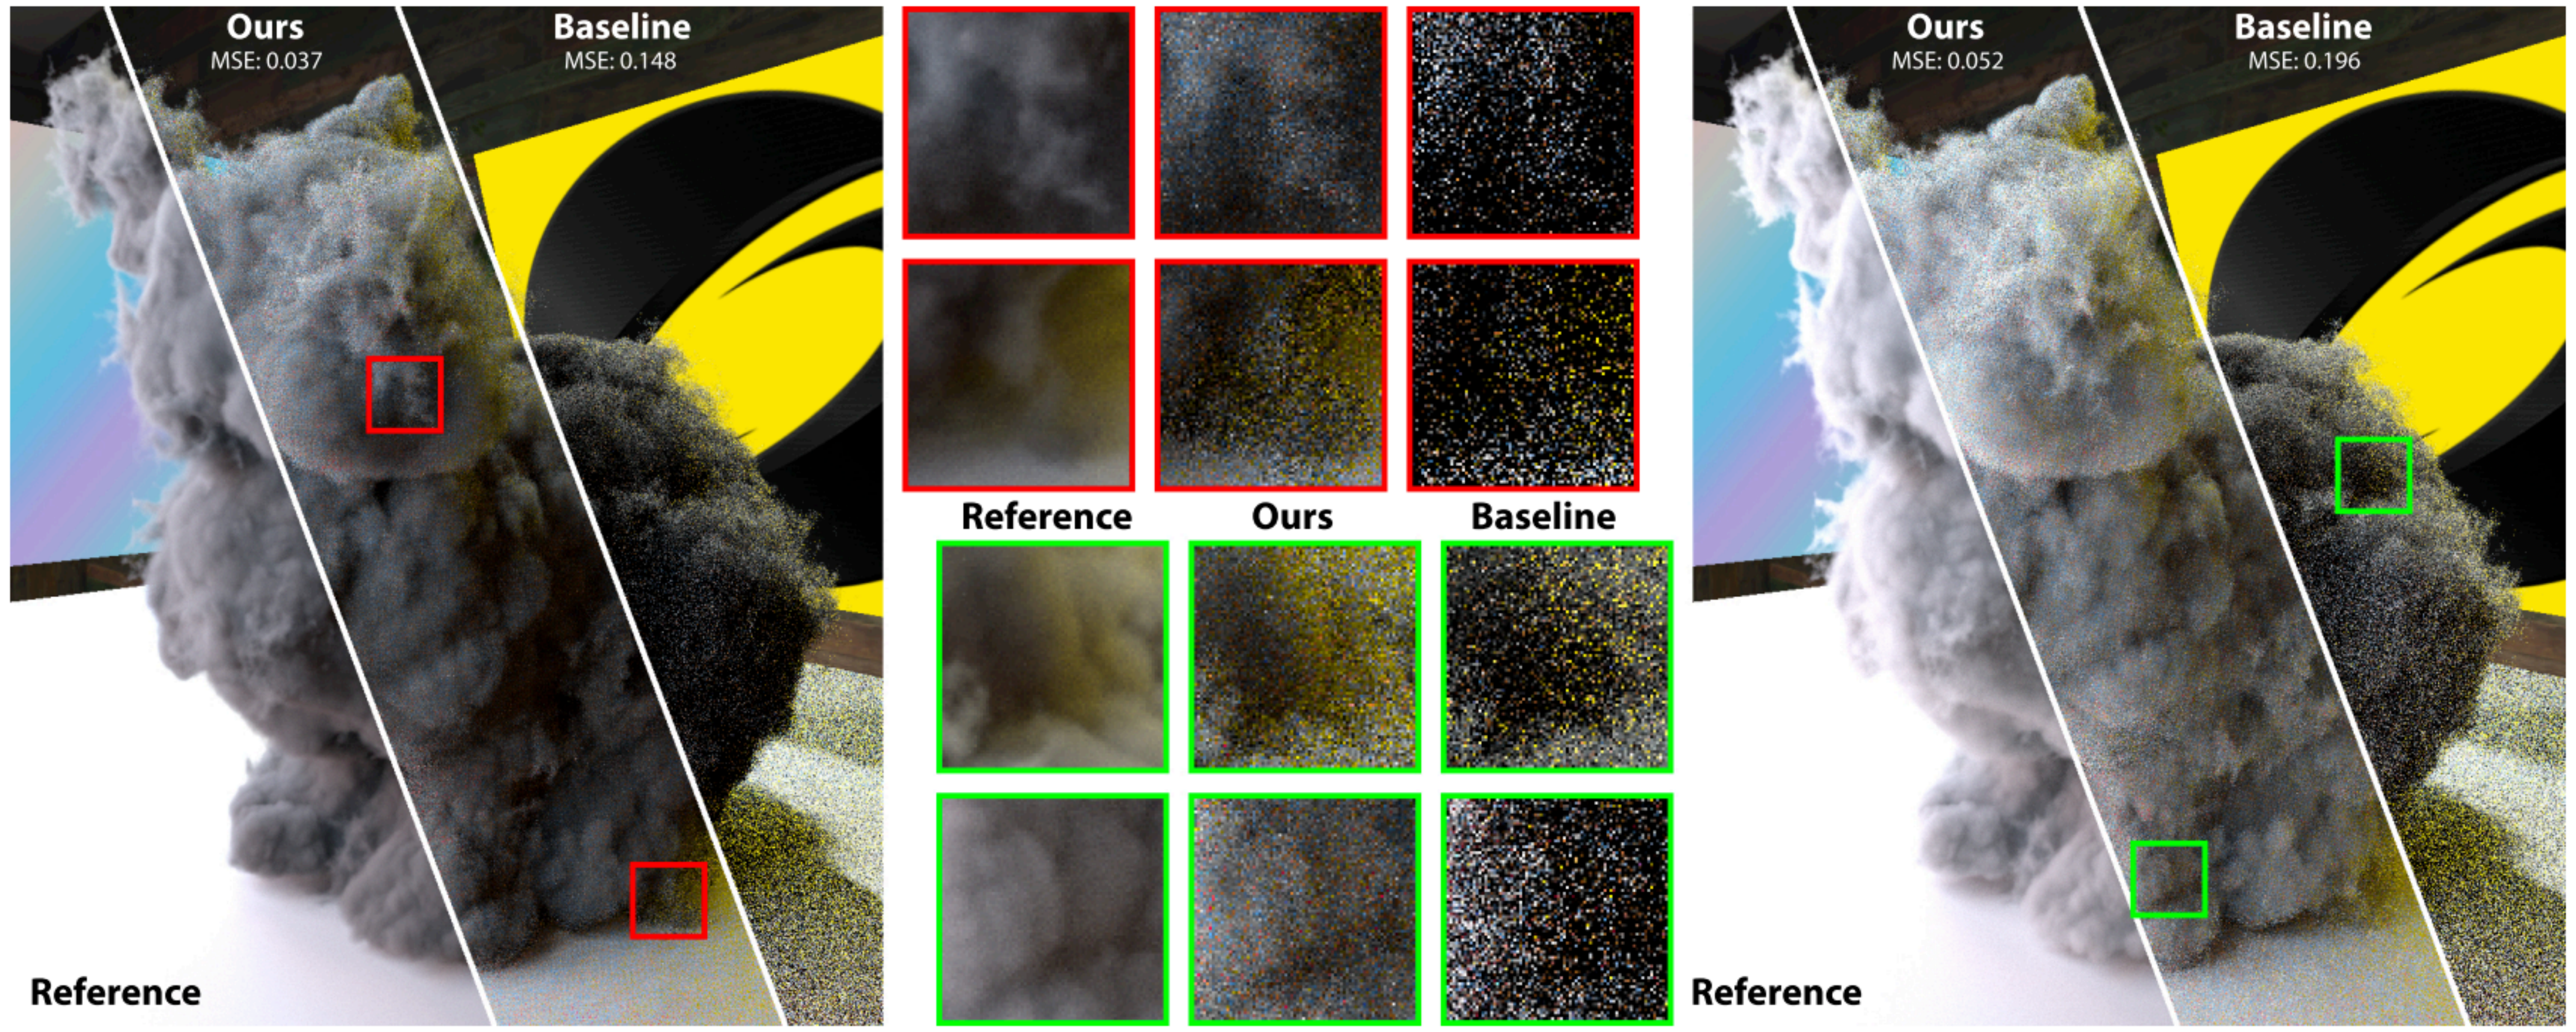

Volumetric ReSTIR

Some cool theories from Lin 2022

most important message: you should “cap” the M count when merging reservoirs!

$$M = \min\left(M_0 + M_1, M_{\text{max}}\right)$$

Generalized Resampled Importance Sampling: Foundations of ReSTIR

DAQI LIN*, University of Utah, USA
MARKUS KETTUNEN*, NVIDIA, Finland
BENEDIKT BITTERLI, NVIDIA, USA
JACOPO PANTALEONI, NVIDIA, Germany
CEM YUKSEL, University of Utah, USA
CHRIS WYMAN, NVIDIA, USA

Discussion: ReSTIR vs path guiding vs MCMC

Next: production rendering for visual effects

The Reyes Image Rendering Architecture

Robert L. Cook
Loren Carpenter
Edwin Catmull

Pixar
P. O. Box 13719
San Rafael, CA 94913