

An Efficient Denoising Algorithm for Global Illumination



Michael Mara, Morgan McGuire,
Benedikt Bitterli, and Wojciech Jarosz

Dartmouth

Path Tracing Produces Compelling Results

16 paths/pix



256 paths/pix



+ denoising



Reference




Direct

Traced Indirect

Denoised Indirect

Final



2.5ms

0.5ms

0.5ms

3ms

2.5ms

2ms


Proposed Path Traced VR Pipeline

Direct

Traced Indirect

Denoised Indirect

Final



Shadowmaps

G-Buffer

Deferred Shading

Trace Indirect

Denoise Indirect

Composite and Post

2.5ms

0.5ms

0.5ms

3ms

2.5ms

2ms

Proposed Path Traced VR Pipeline



Noise




Smearing




Flicker



Blur



Boiling



Ghosting

Related Work


Offline Denoisers

- Survey papers:
 - Gautron et al. [2014]
 - Zwicker et al. [2015]
- Small sampling of recent work:
 - Dammertz et al. [2010]
 - Rousselle et al. [2012]
 - Moon et al. [2014]
 - Kalantari et al. [2015]
 - Bitterli et al. [2016] →
 - Bako and Vogels et al. [2017]



Interactive Denoiser

- Joint tracing/filtering approach
- ~80ms
- 4 paths/pixel
- Matte surfaces only
- Simple scenes





HPG/SIGGRAPH 2017 Realtime Denoisers

- Schied et al. “Spatiotemporal Variance-Guided Filtering: Real-Time Reconstruction for Path Traced Global Illumination.” **HPG 2017**
- Chaitanya et al. “Interactive Reconstruction of Monte Carlo Image Sequences using a Recurrent Denoising Autoencoder.” **SIGGRAPH 2017**
- Mara et al. “An Efficient Denoising Algorithm for Global Illumination.” **HPG 2017**

Robust Extension to Previous Ad-Hoc Approaches

[McGuire, HPG 2012]

[Mara, I3D 2013]



[Mara, HPG 2016]



264 *Low bandwidth,
low-power data flow cut*

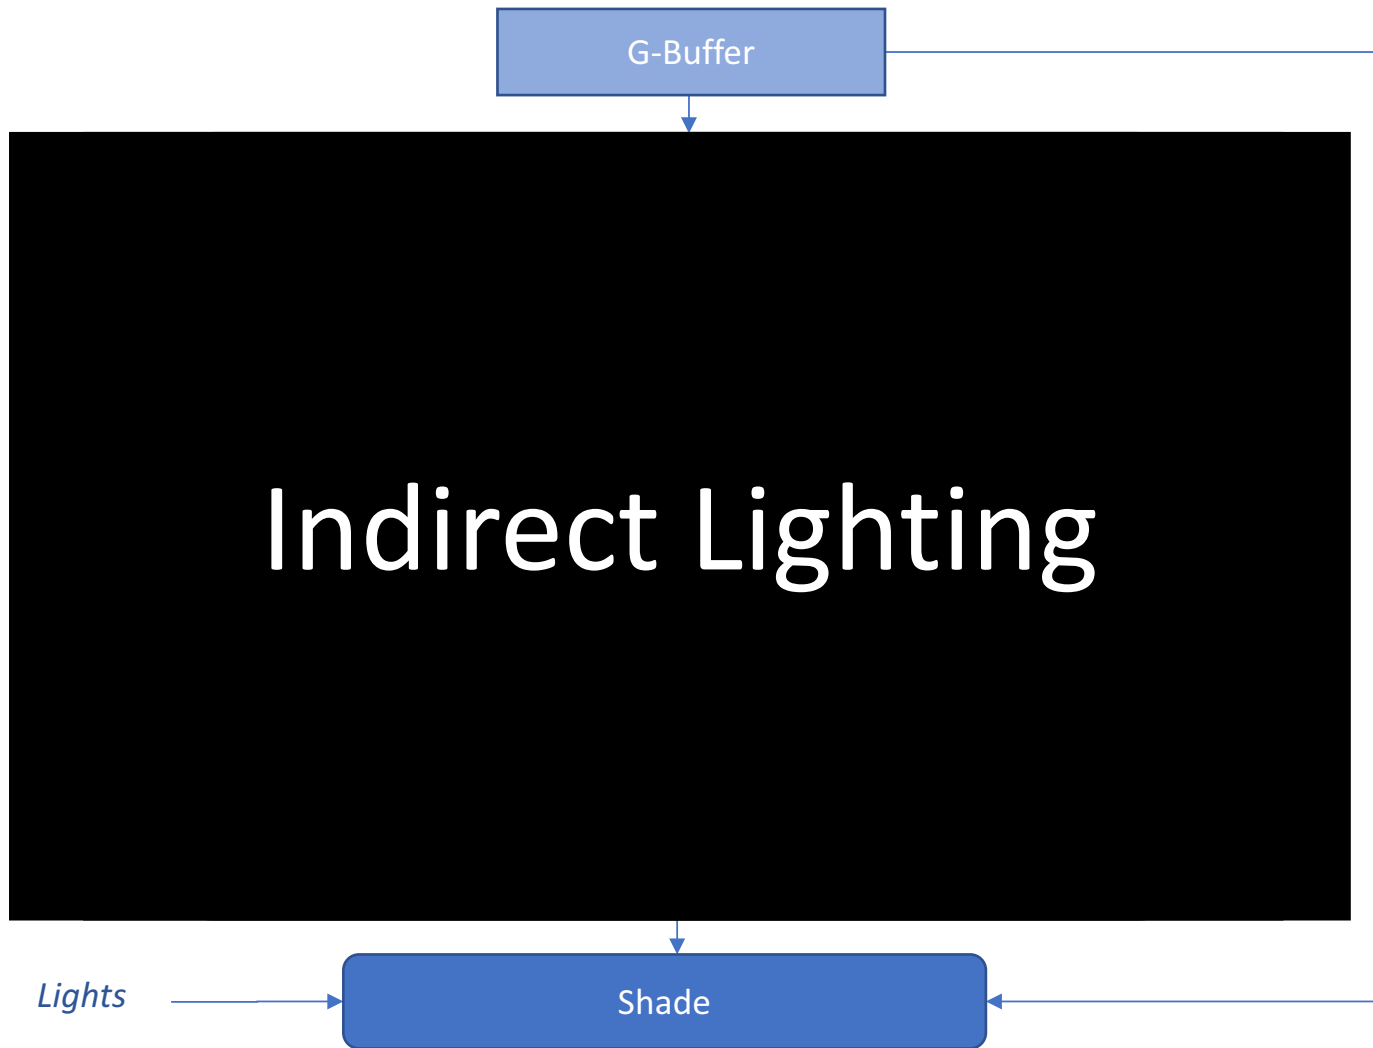


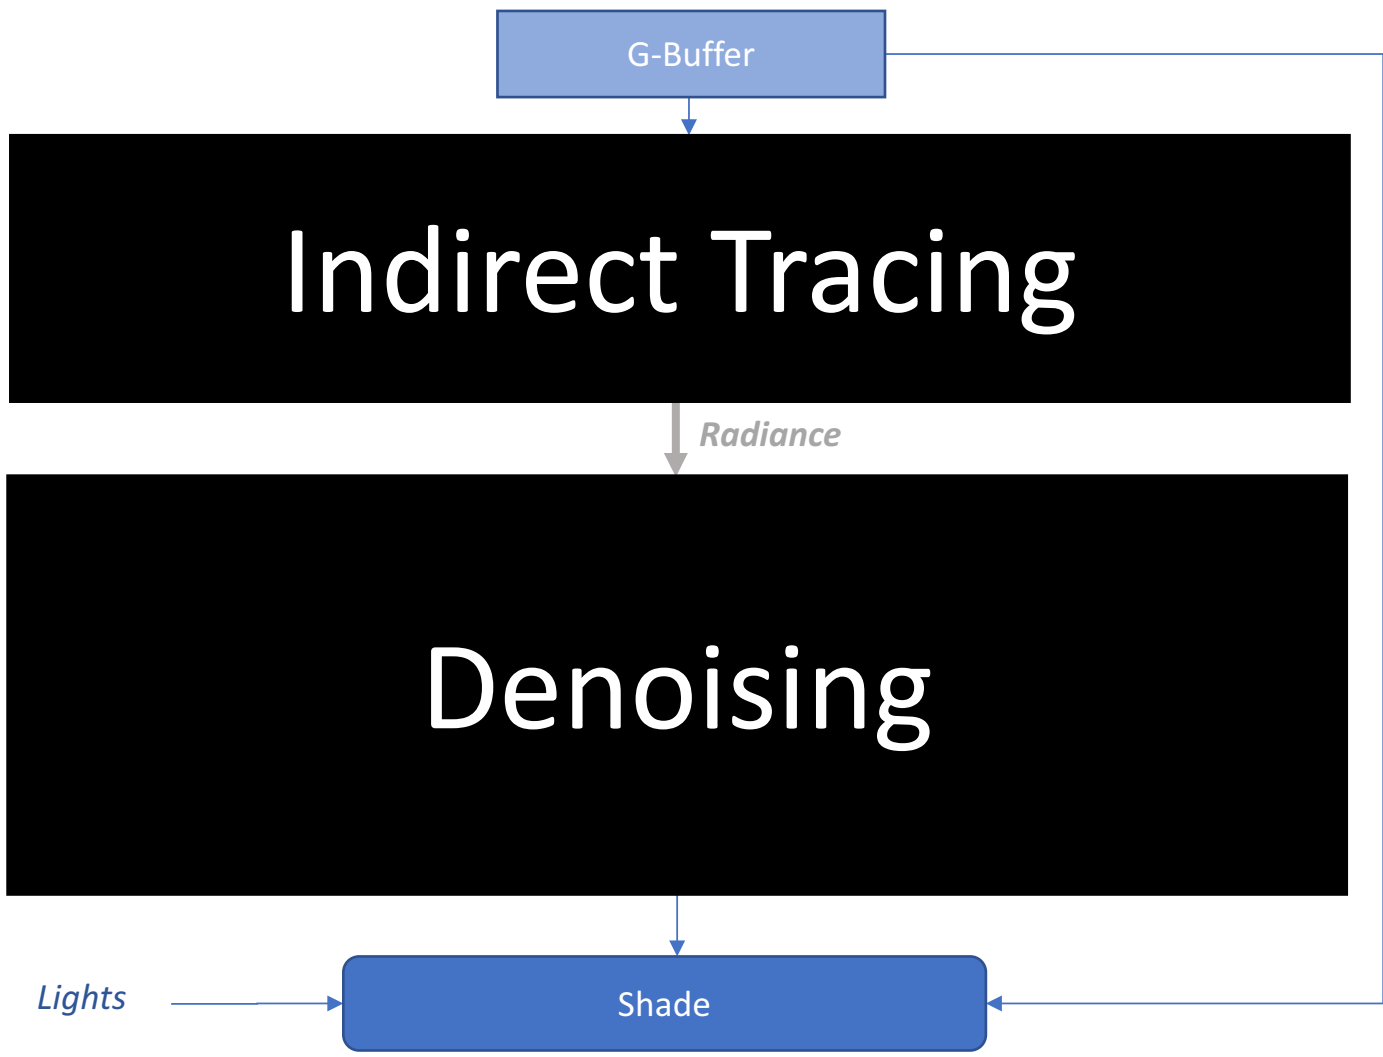
Indirect Radiance

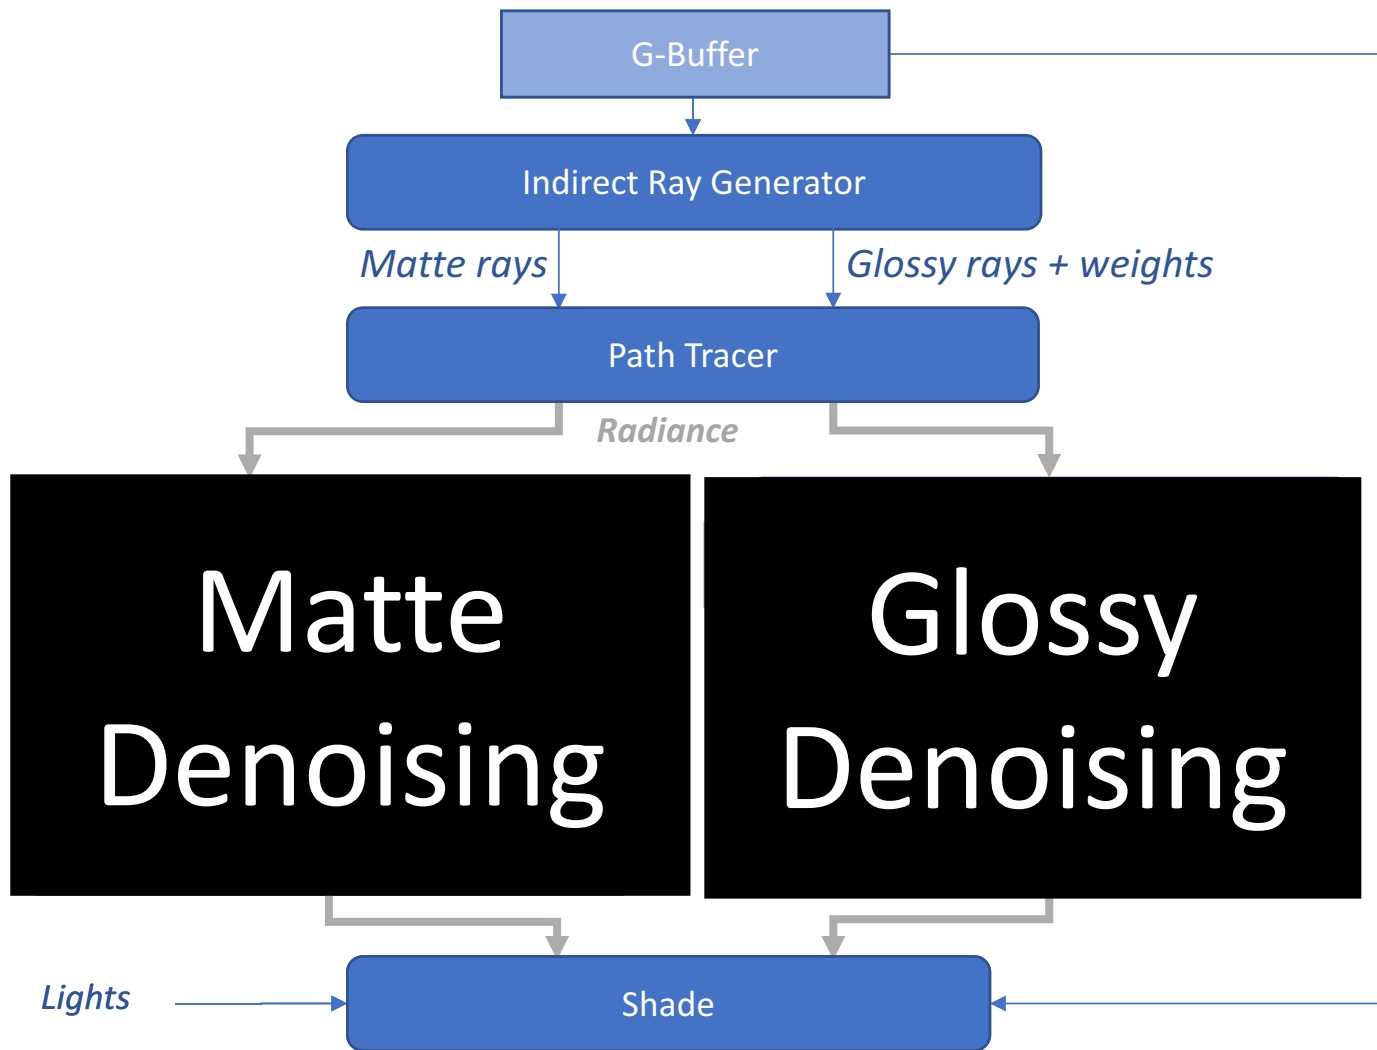
Net Radiance

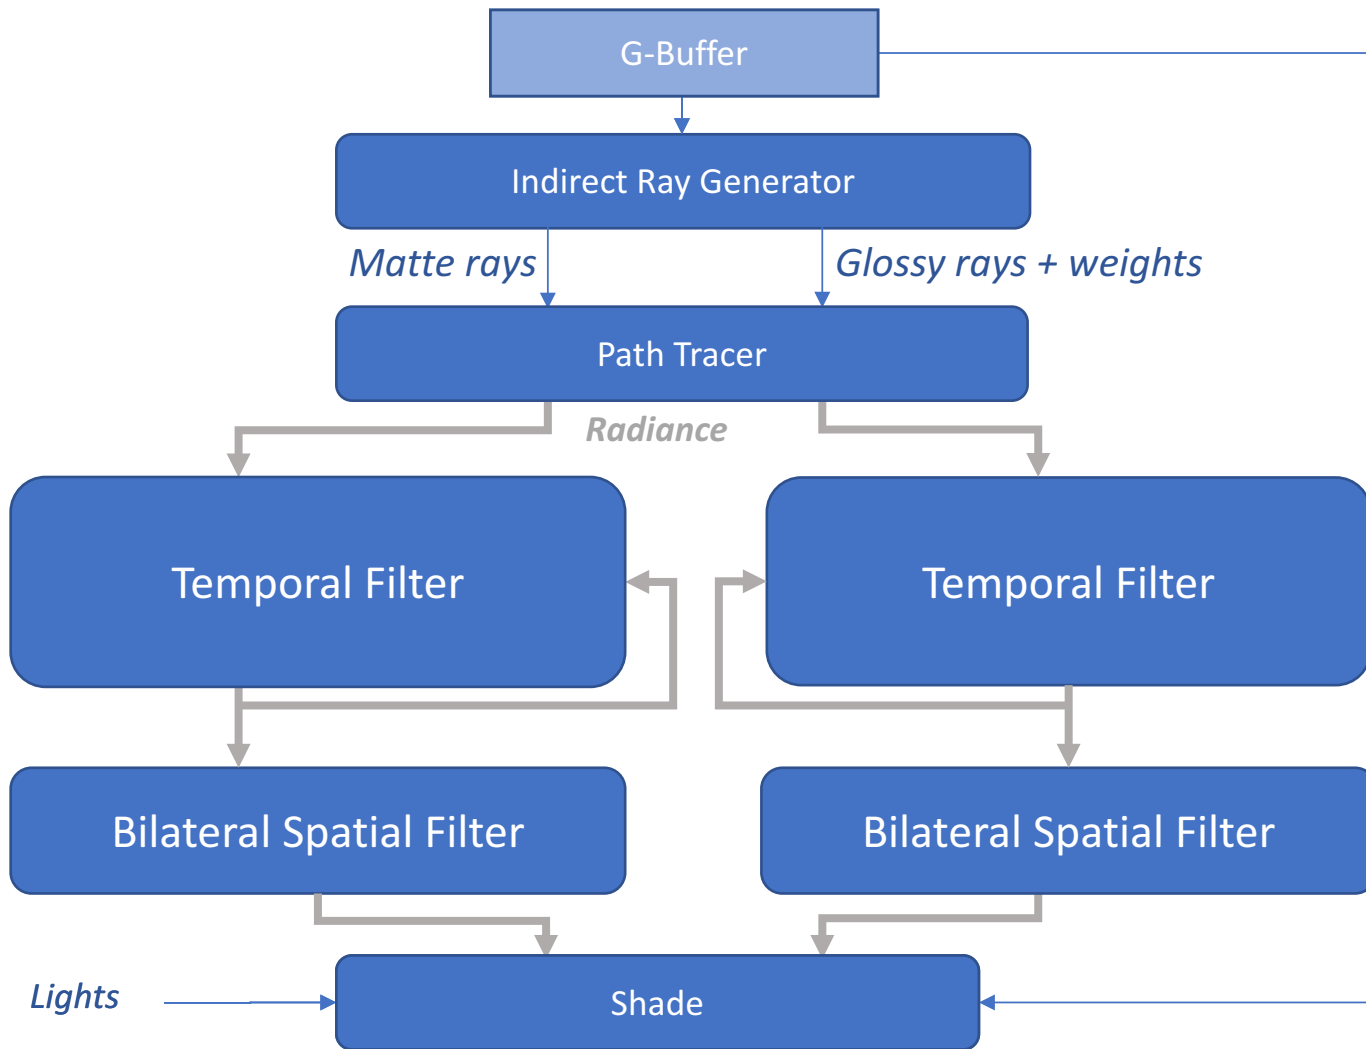
[Crassin, JCGT 2015]

Our Algorithm









Factored BRDF

$$f(\hat{\omega}_i, \hat{\omega}_o) = m(\hat{\omega}_i, \hat{\omega}_o) (1 - F(\hat{\omega}_i, \hat{\omega}_o))^2 + g(\hat{\omega}_i, \hat{\omega}_o) F(\hat{\omega}_i, \hat{\omega}_o)$$

Factored BRDF

$$f(\hat{\omega}_i, \hat{\omega}_o) = \underbrace{m(\hat{\omega}_i, \hat{\omega}_o)}_{\text{matte}} (1 - F(\hat{\omega}_i, \hat{\omega}_o))^2 + \underbrace{g(\hat{\omega}_i, \hat{\omega}_o)}_{\text{glossy}} F(\hat{\omega}_i, \hat{\omega}_o)$$

Factored Monte Carlo Integration

$$L_o(X, \hat{\omega}_o) = \frac{1}{2N} \sum_i^{2N} L_i(X, \hat{\omega}_i) \frac{f(\hat{\omega}_i, \hat{\omega}_o) |\hat{n} \cdot \hat{\omega}_i|}{p(\hat{\omega}_i)}$$

$$L_o(X, \hat{\omega}_o) = \frac{\pi}{N} \sum_i^N L_i(X, \hat{\omega}_i) (1 - F(\hat{\omega}_i, \hat{\omega}_o))^2 m(\hat{\omega}_i, \hat{\omega}_o) \\ + \frac{1}{N} \sum_j^N L_i(X, \hat{\omega}_j) \frac{F(\hat{\omega}_j, \hat{\omega}_o) g(\hat{\omega}_j, \hat{\omega}_o) |\hat{n} \cdot \hat{\omega}_j|}{p_g(\hat{\omega}_j)}$$

Approximation of Matte Estimator

$$\frac{\pi}{N} \sum_i^N L_i(X, \hat{\omega}_i) (1 - F(\hat{\omega}_i, \hat{\omega}_0))^2 m(\hat{\omega}_i, \hat{\omega}_0)$$




$$\frac{(1 - F(\hat{\omega}_s, \hat{\omega}_0))^2 m(\hat{\omega}_s, \hat{\omega}_0) \pi}{N} \sum_i^N L_i(X, \hat{\omega}_i)$$

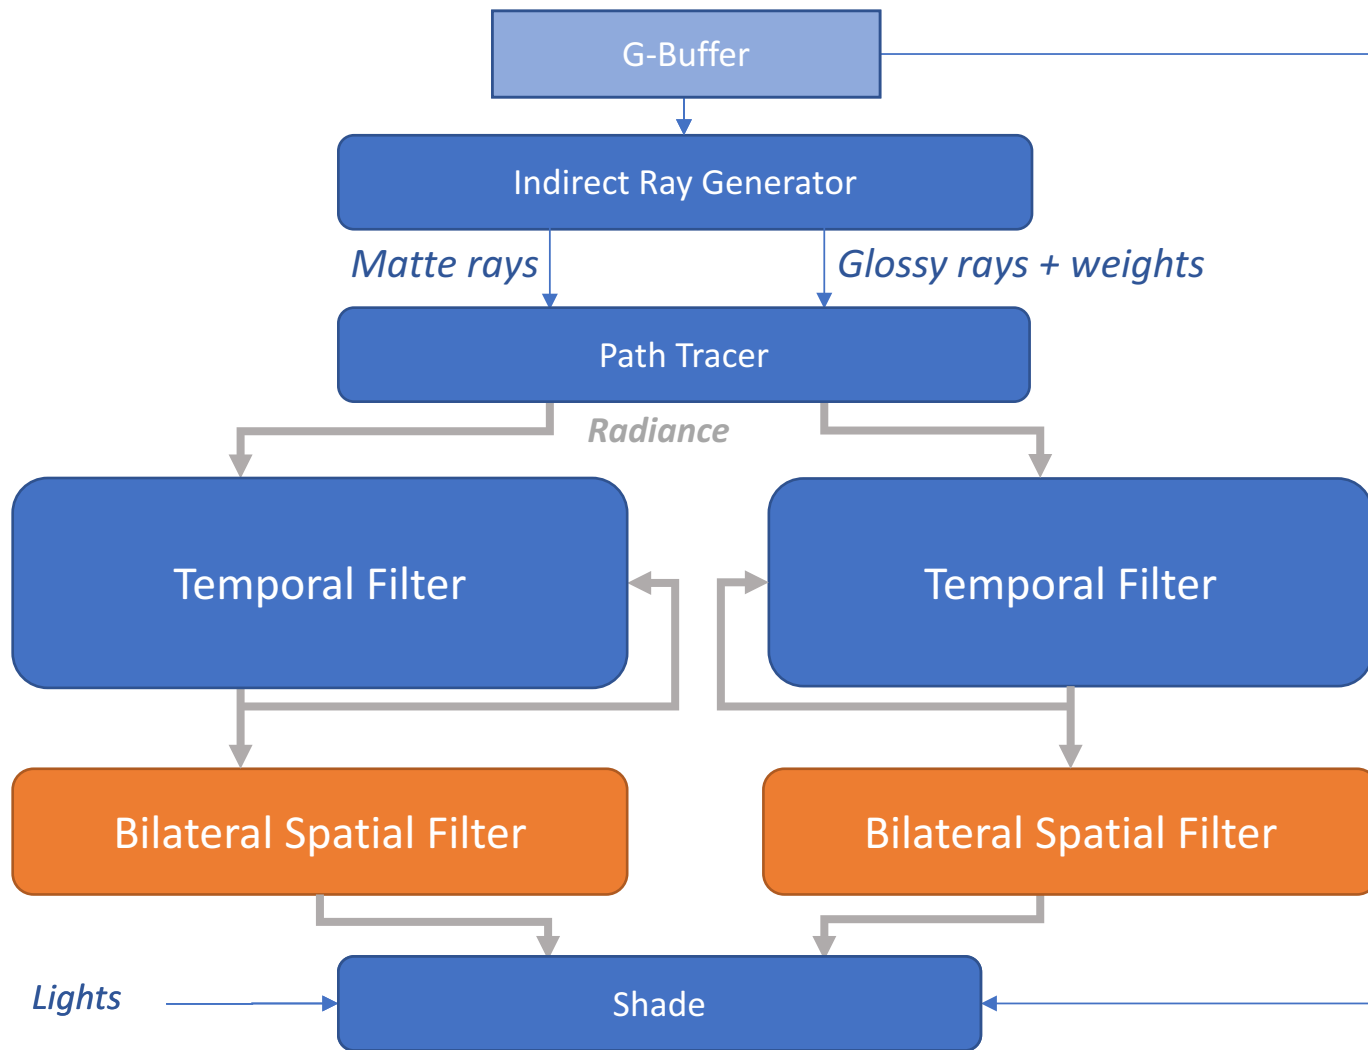
Approximation of Glossy Estimator

$$\frac{1}{N} \sum_j^N L_i(X, \hat{\omega}_j) \frac{F(\hat{\omega}_j, \hat{\omega}_0) g(\hat{\omega}_j, \hat{\omega}_0) |\hat{n} \cdot \hat{\omega}_j|}{p_g(\hat{\omega}_j)}$$




$$\frac{F(\hat{\omega}_s, \hat{\omega}_0)}{N} \sum_j^N L_i(X, \hat{\omega}_j) \frac{g(\hat{\omega}_j, \hat{\omega}_0) |\hat{n} \cdot \hat{\omega}_j|}{p_g(\hat{\omega}_j)}$$





Spatial Filter Taps

Matte Bilateral Filter




Sparse

Glossy Bilateral Filter

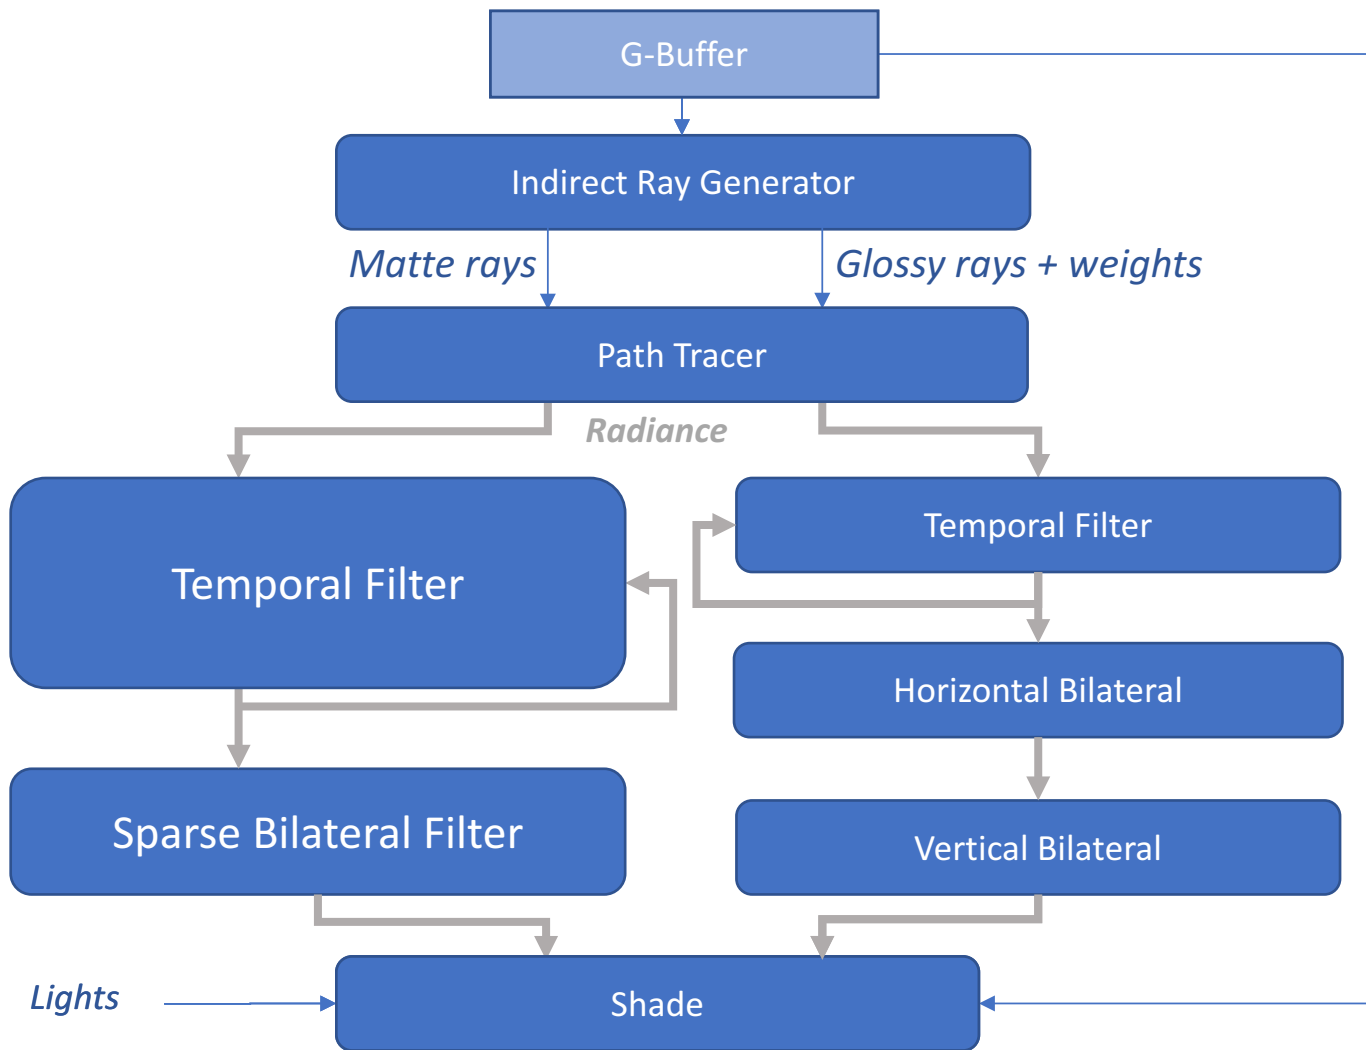


Horizontal



Vertical

Two-pass





WITHOUT DENOISING

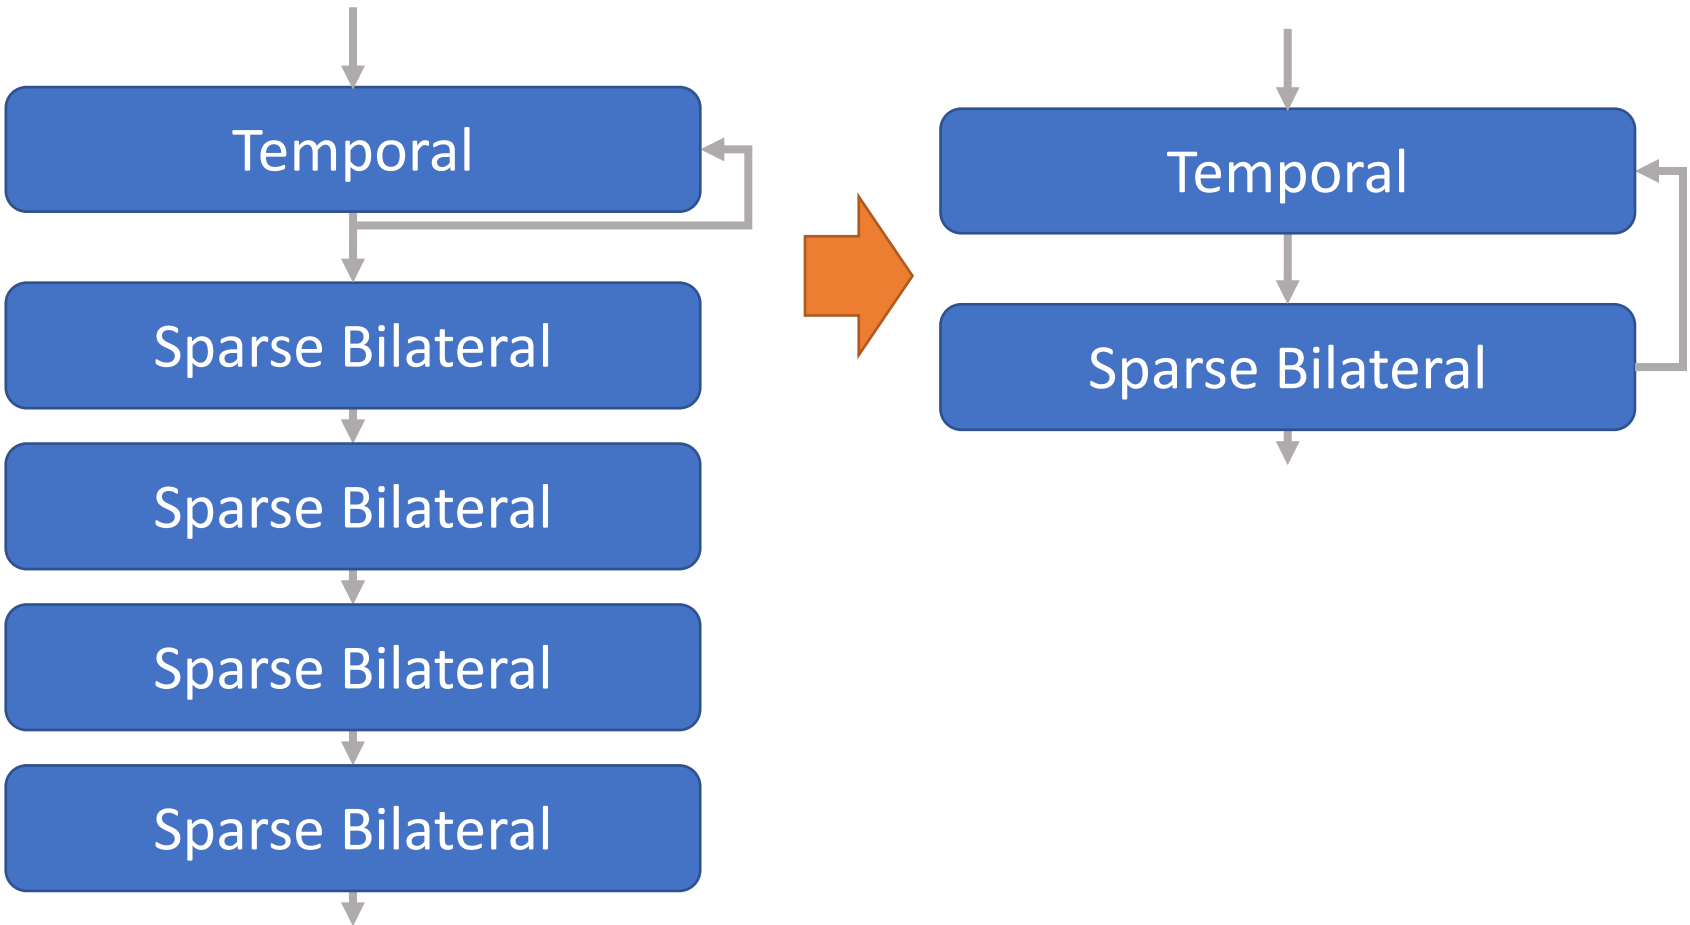


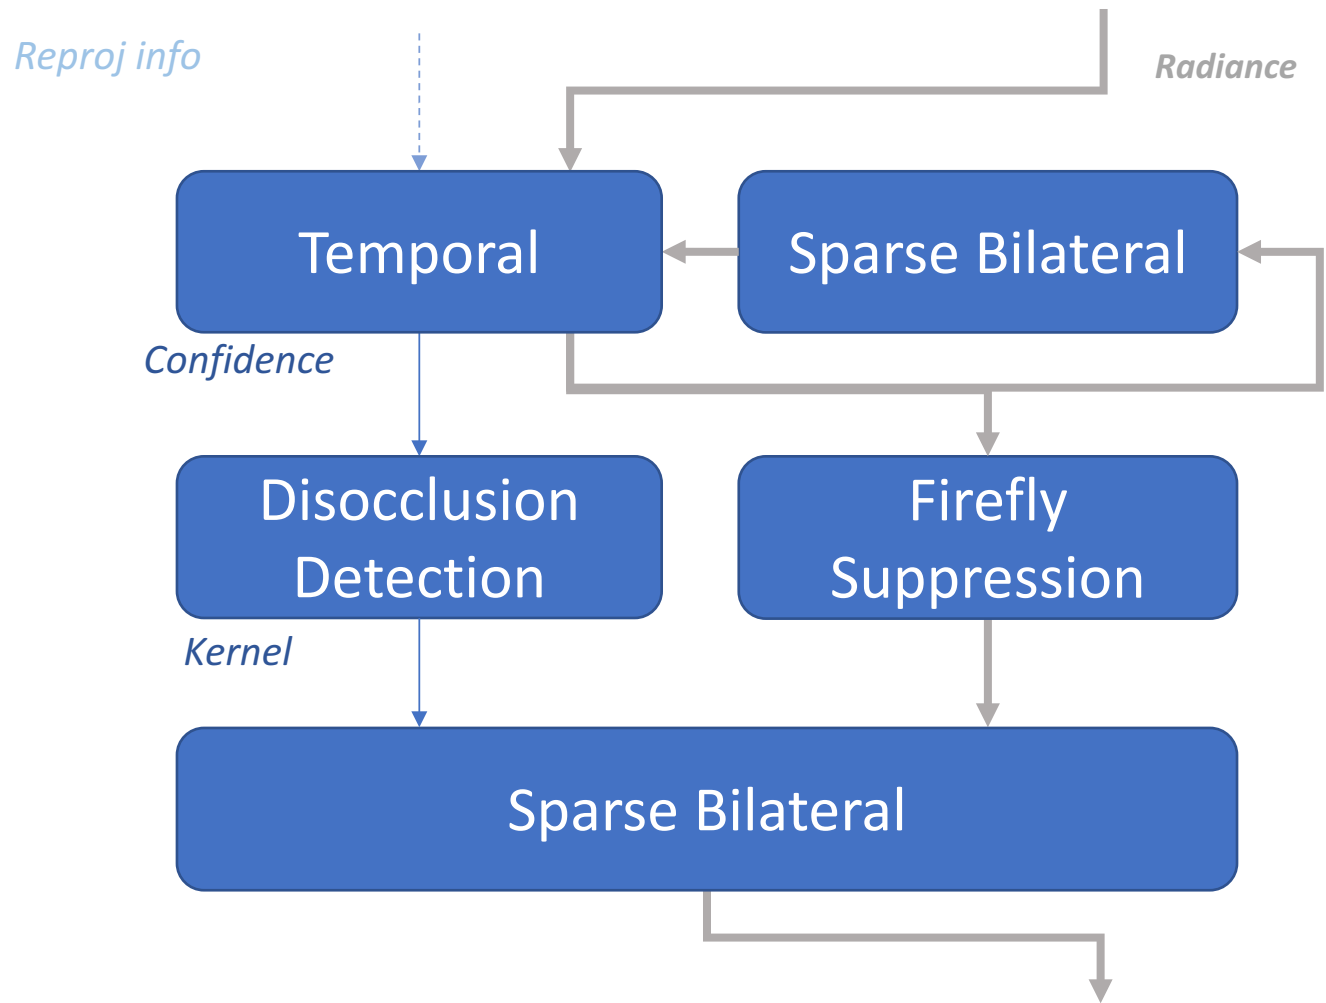
NAÏVE RESULTS

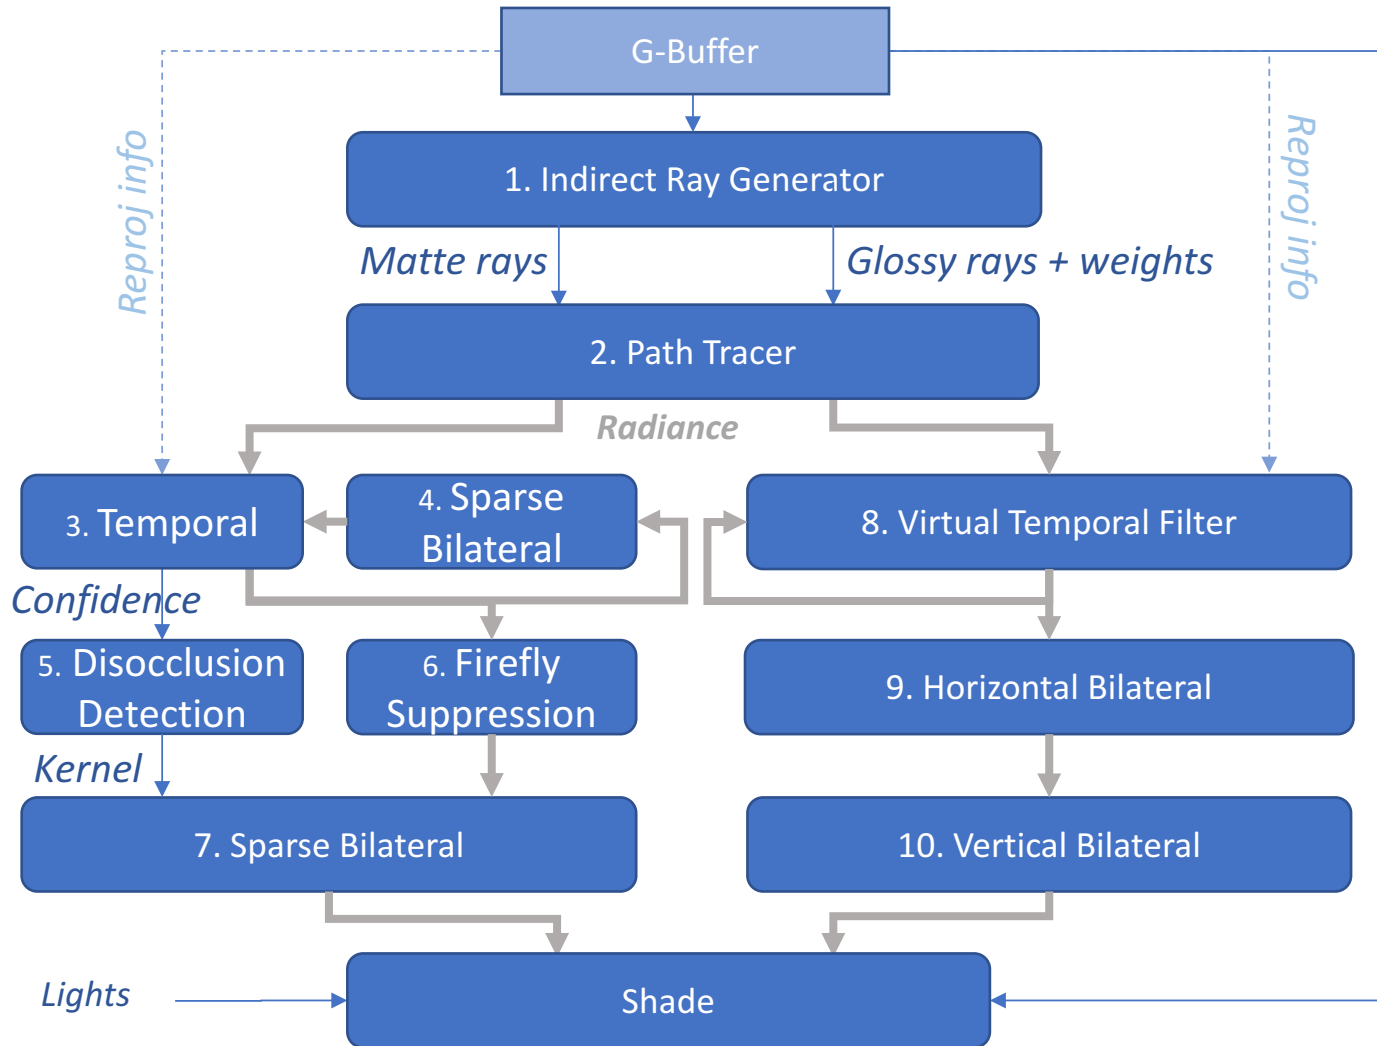
Jointly Design Spatial and Temporal Filters



Spatio-Temporal Feedback









Final Results

Offline Results



None



NLM




WLR



NFOR



Ours



Offline Results








None

NLM

WLR

NFOR

Ours



Offline Results












None

NLM

WLR

NFOR

Ours



Offline Results










None

NLM

WLR

NFOR

Ours



Offline Results







None

NLM

WLR

NFOR

Ours



Offline Results


Scene	NLM	WLR	NFOR	Ours
<i>Cornell Box</i>	55s	54s	110s	0.16s
<i>Glossy Sponza</i>	63s	54s	138s	0.28s
<i>Horse Room</i>	60s	52s	129s	0.31s
<i>Bathroom</i>	56s	50s	126s	0.32s
<i>Living Room</i>	58s	50s	119s	0.31s

5120×2880 \approx 14.7Mpix

Realtime Results

Table 1: Run-time in milliseconds for real-time filter pipeline stages at 1280×720.

Stage	Time	Stage	Time
3. Matte Temporal	0.18	7. Matte Bilateral	3.36
4. Matte Pre-Bilateral	3.32	8. Glossy Temporal	0.23
5. Disocclusion Detection	0.76	9. Glossy H. Bilateral	0.34
6. Firefly Suppression	0.55	10. Glossy V. Bilateral	0.43












Ground Truth





















FULL ALGORITHM

3-BOUNCE ILLUMINATION, 2 PATHS PER PIXEL

Real-time Demo

Conclusions

- Fast 2 paths/pixel denoising via factoring
- Joint spatio-temporal filters
- Real-time denoised pathtracing is coming

Future Work

- In future: 6-10 rays/frame
- Direct Followup: Lower bandwidth

Before



After

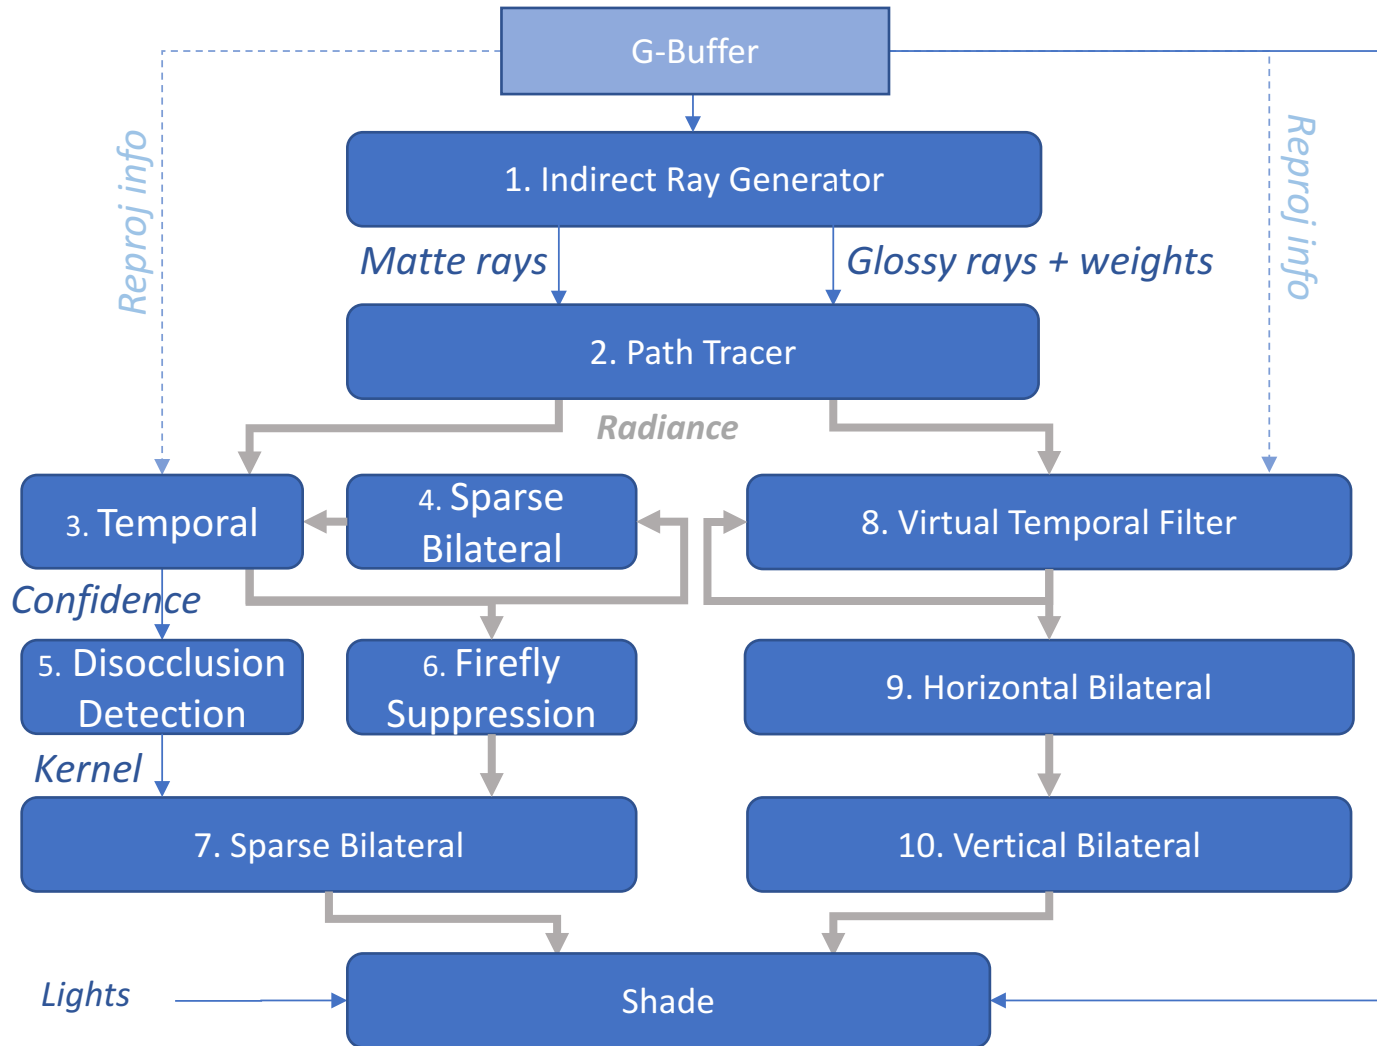
Thank you!

casual-effects.com/research/Mara2017Denoise/

Acknowledgements

Carsten Benthin and Ingo Wald at Intel and the NVIDIA OptiX team for their technical and financial support of this research; and the artists Guillermo M. Leal Llaguno, Marko Dabrovic, Nacimus Ait Cherif, “Jay Artist”, Nic Hull, and Guedis Cardenas for their scenes.

Extra Slides



WITHOUT SPATIAL FILTERS



WITHOUT TEMPORAL FILTERS

Differences in final images from paper

- Slightly increased fresnel
- Expanded clip neighborhood for temporal to 2 pixel radius
- TAA on direct term

Temporal Filters

- Re-use computation from past frames
- Most often done with reverse-reprojection
- Used for everything from smoothing AO to full antialiasing
- Helps reduce noise
- Leaning too heavily introduces “ghosting”

Cross-Bilateral Filters

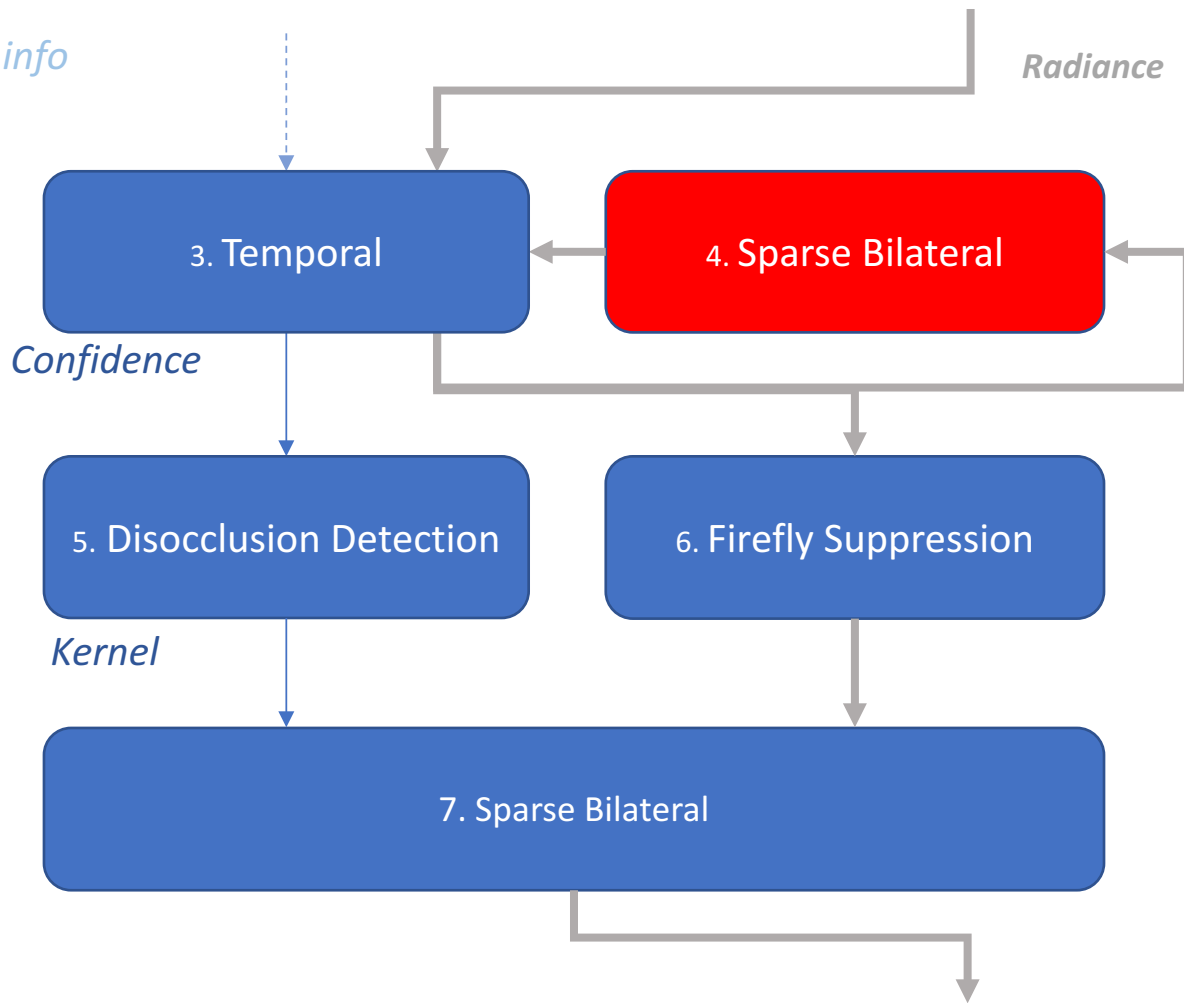
- Trades detail for smoothness
- Combination of too weak spatial + temporal filters:
“boiling”

Success

- Denoising <16ms
- No obvious noise
- No obvious “boiling”
- No obvious “ghosting”

Reproj info

Radiance



Sparse Bilateral Filter

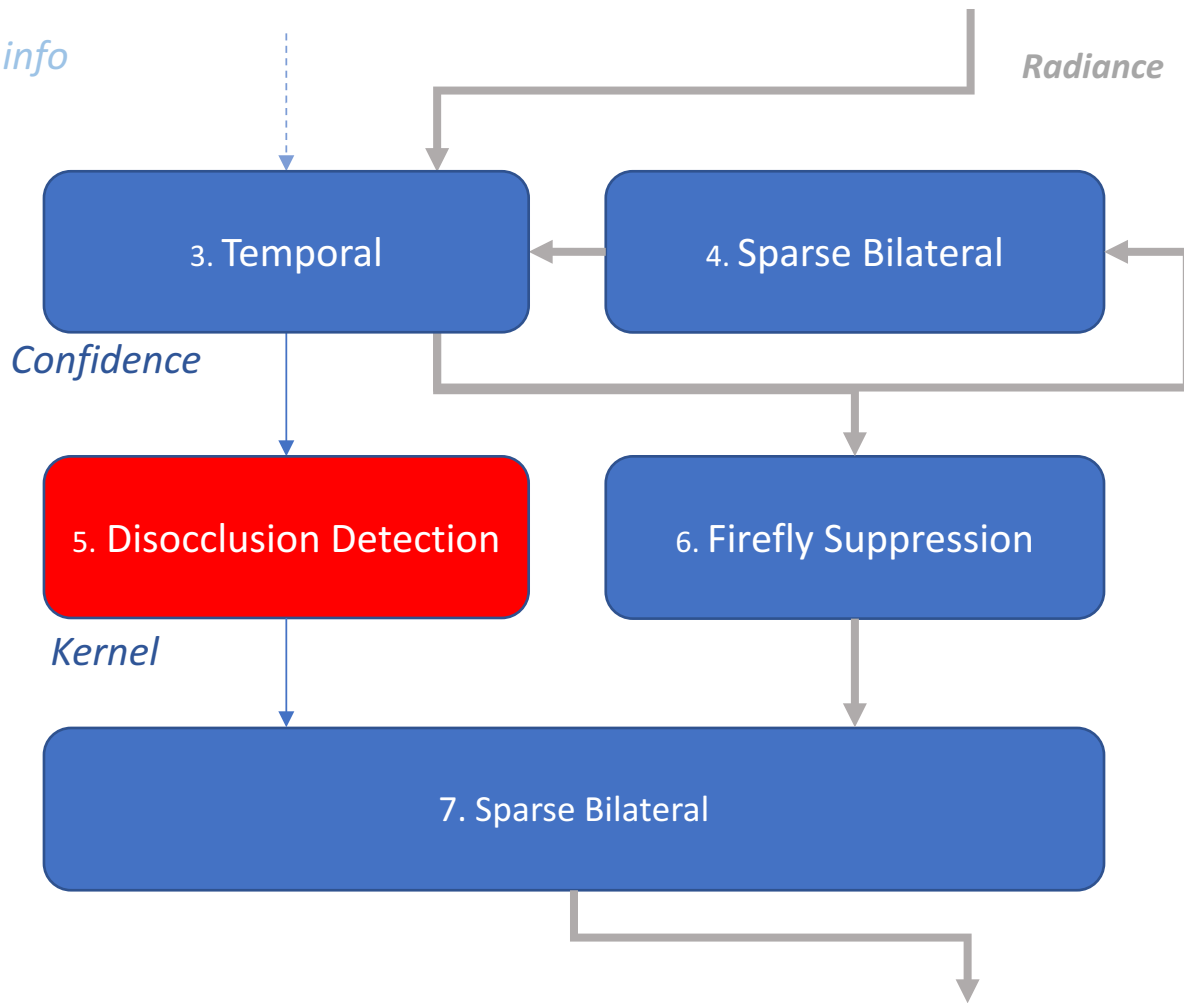
- Cross-Bilateral Filter
- Depth + normals to adjust Gaussian filter weight
- Too big for full filter
- Too big to separate without extreme artifacts
- Use sparse taps

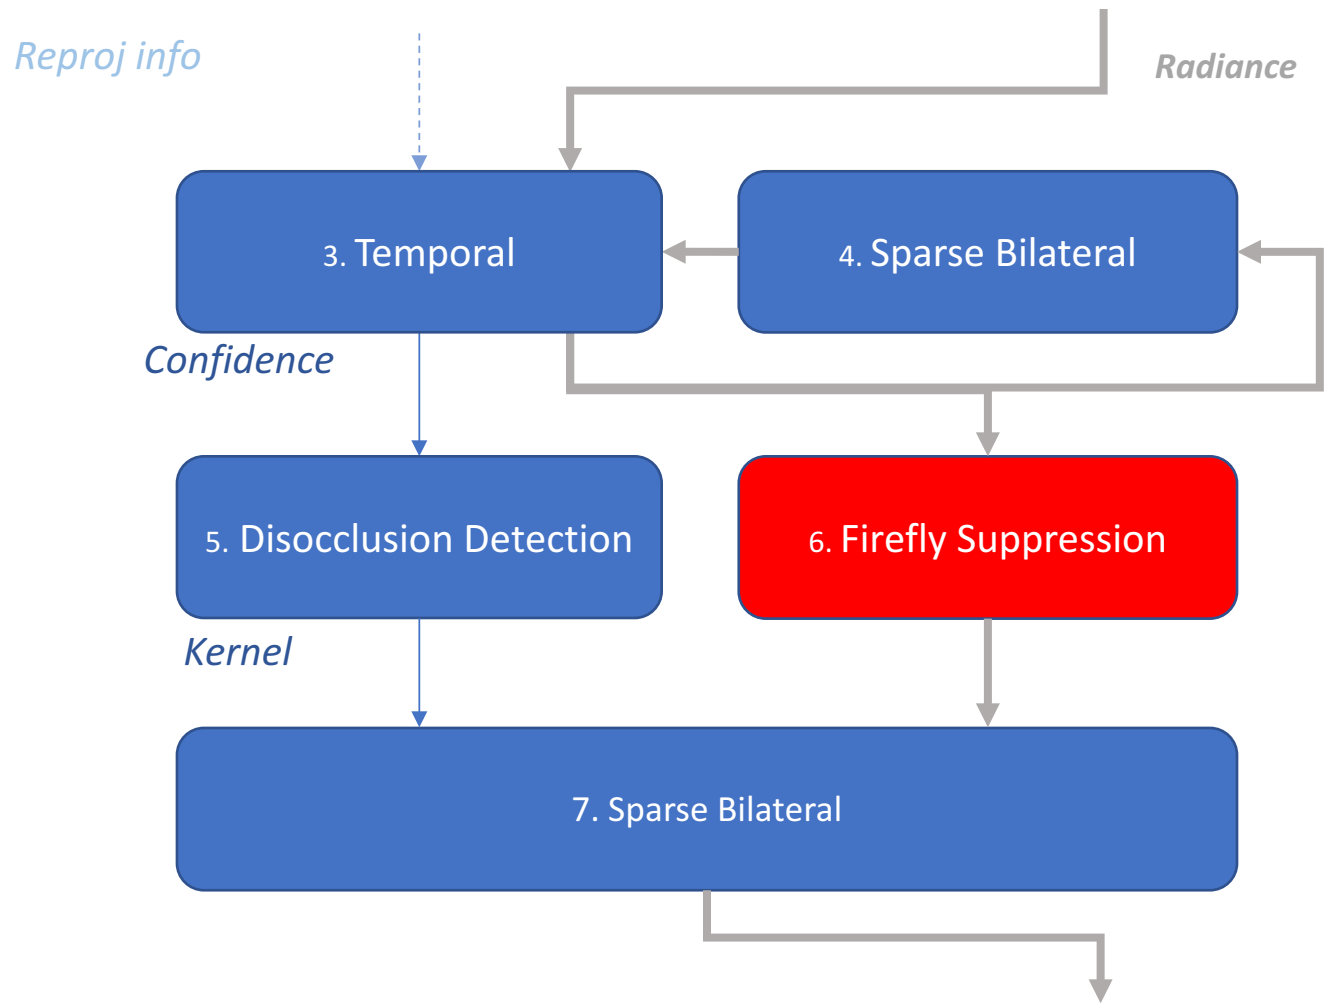
Sparse Bilateral Filter Taps

- Borrowed from [Mara2016]


Reproj info

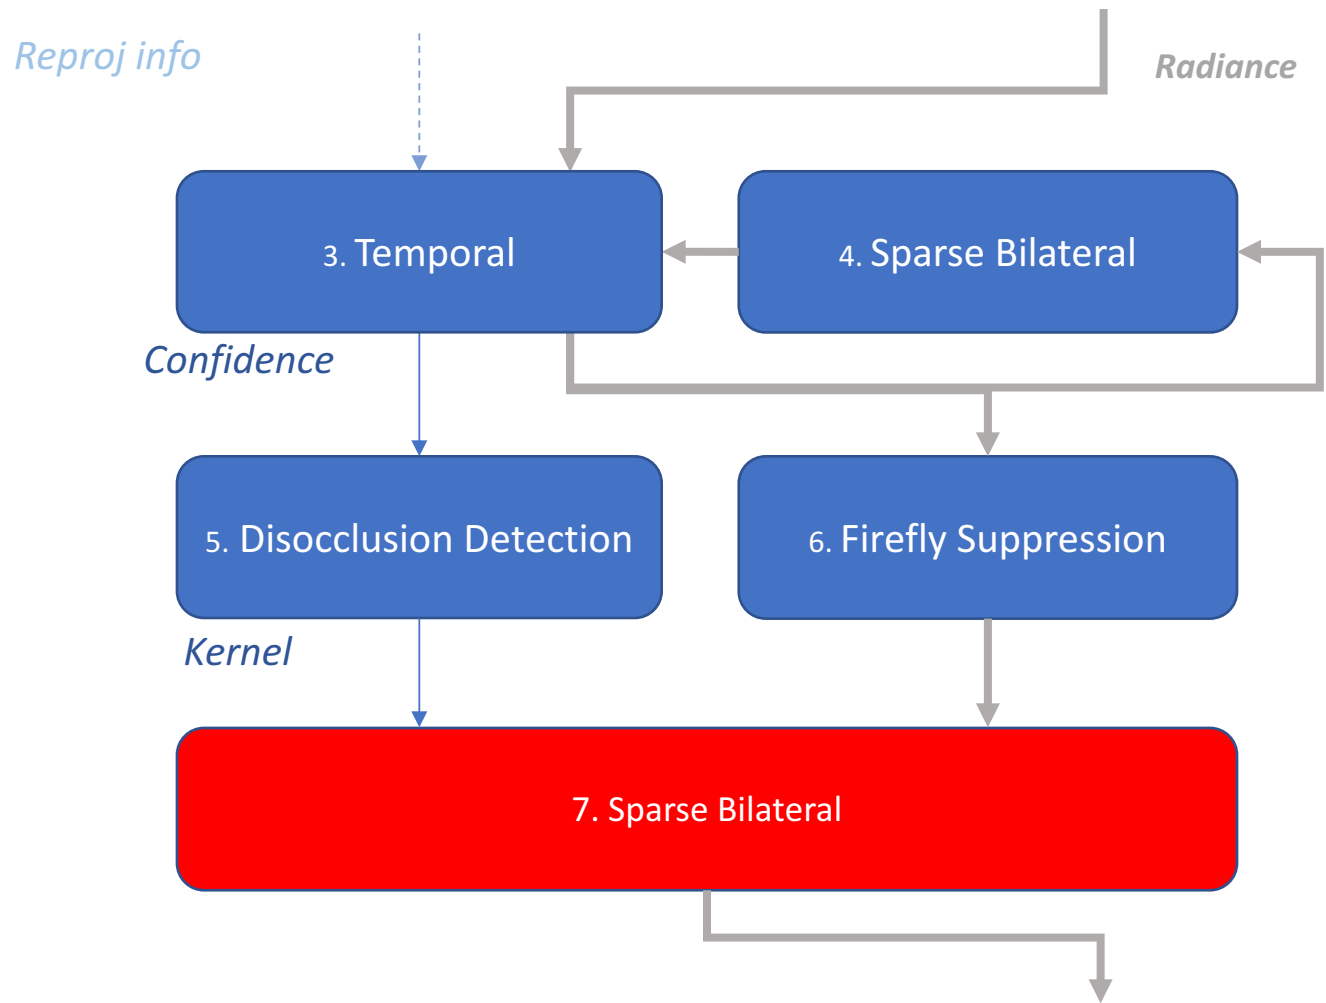
Radiance





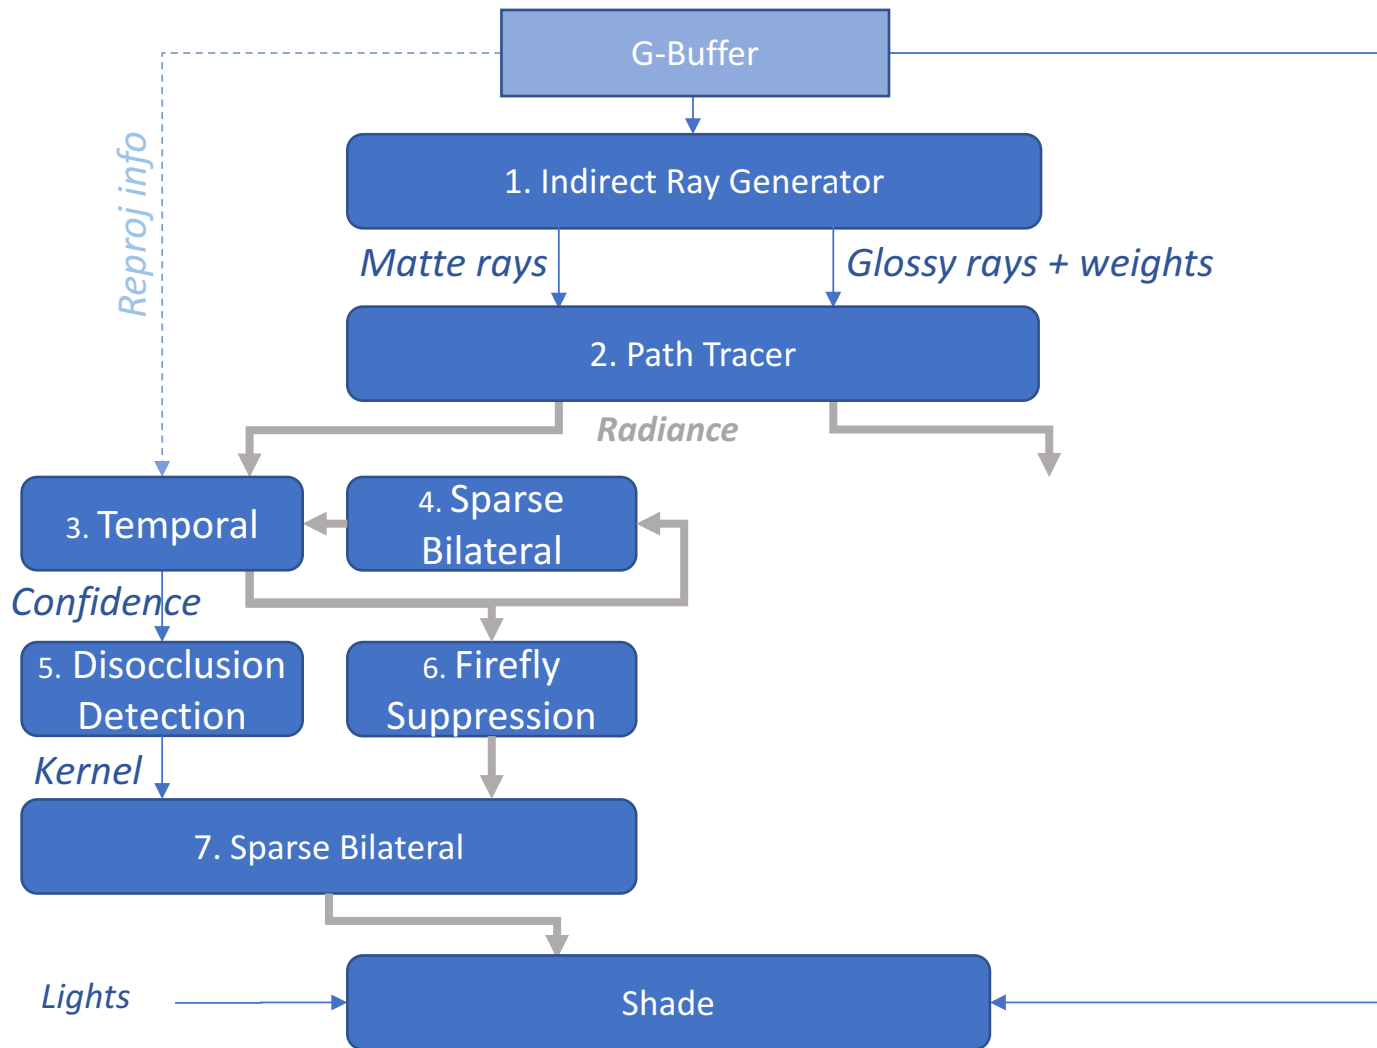
Firefly Suppression using 3x3 Median Filter

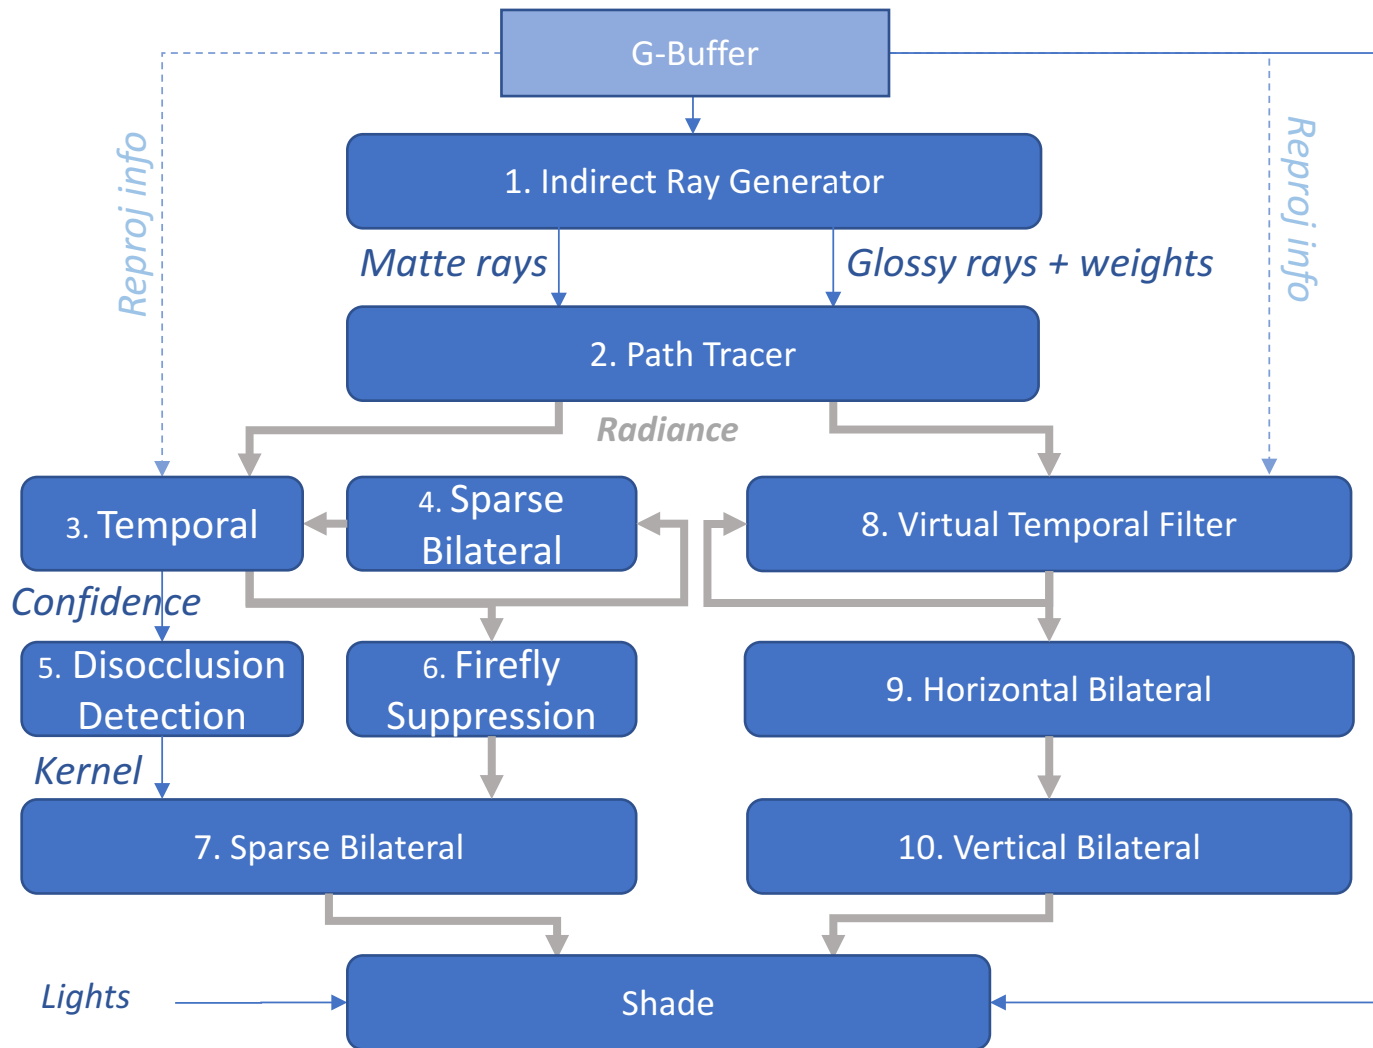






WITHOUT DISOCCLUSION FILTER





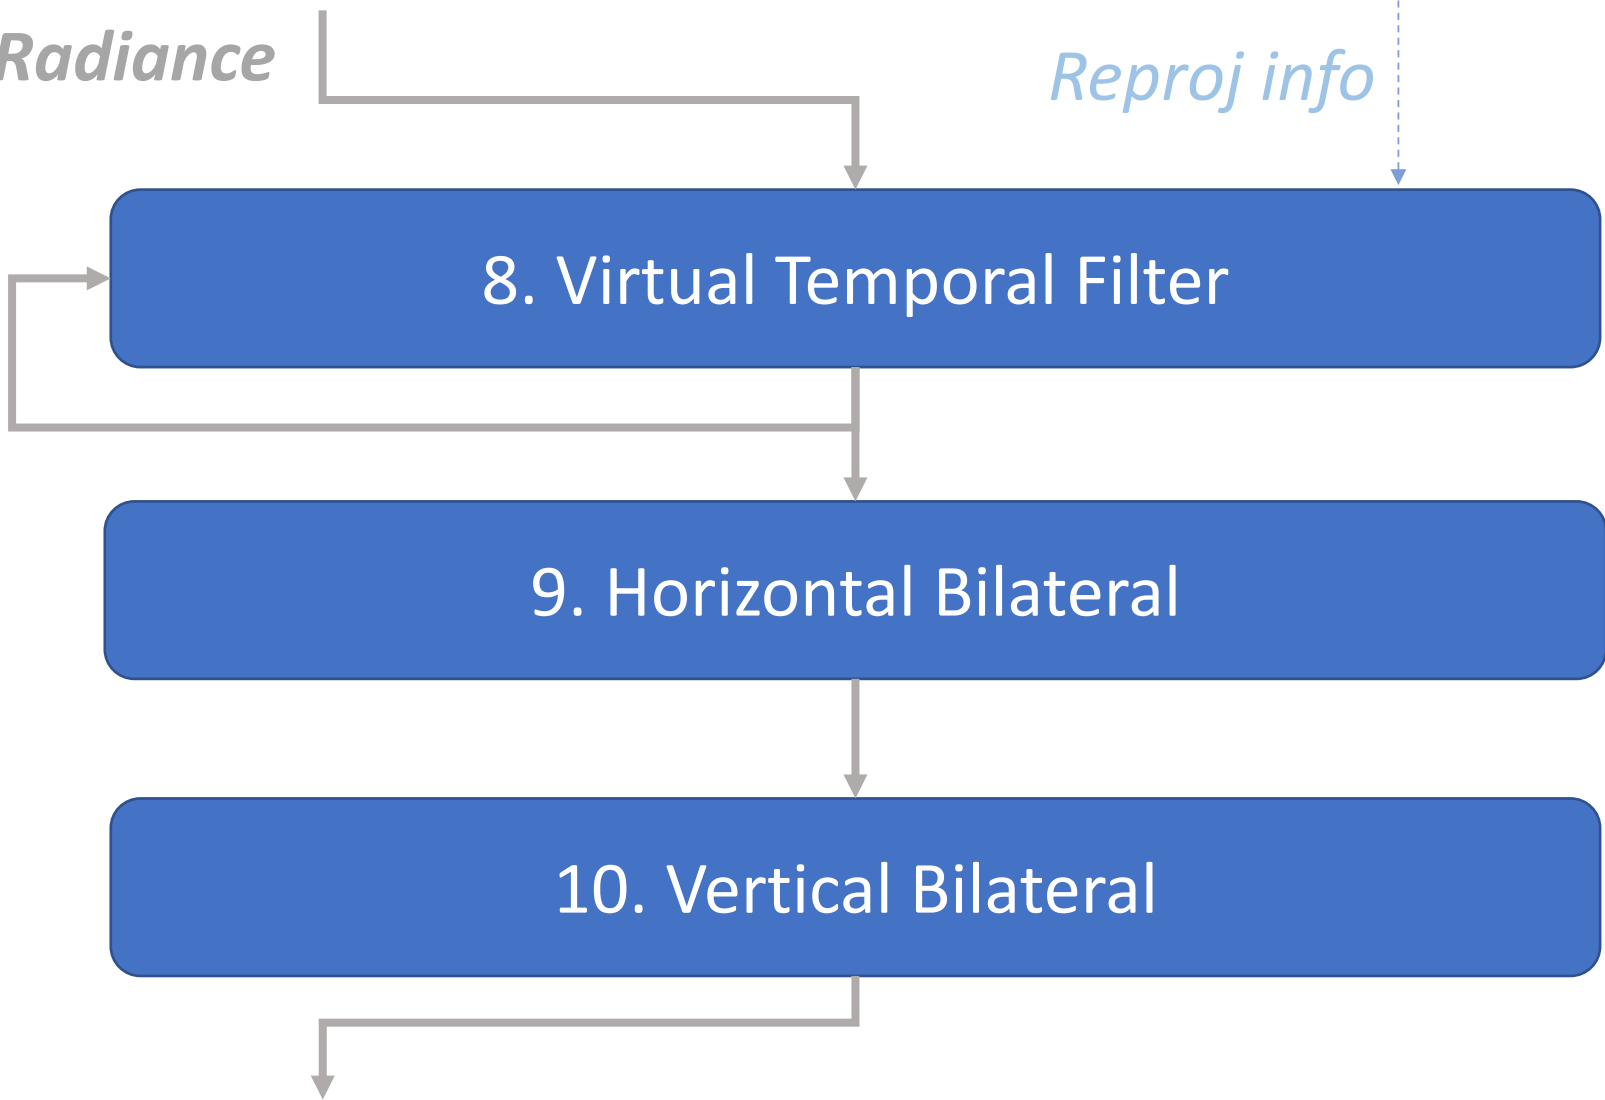
Radiance

Reproj info

8. Virtual Temporal Filter

9. Horizontal Bilateral

10. Vertical Bilateral



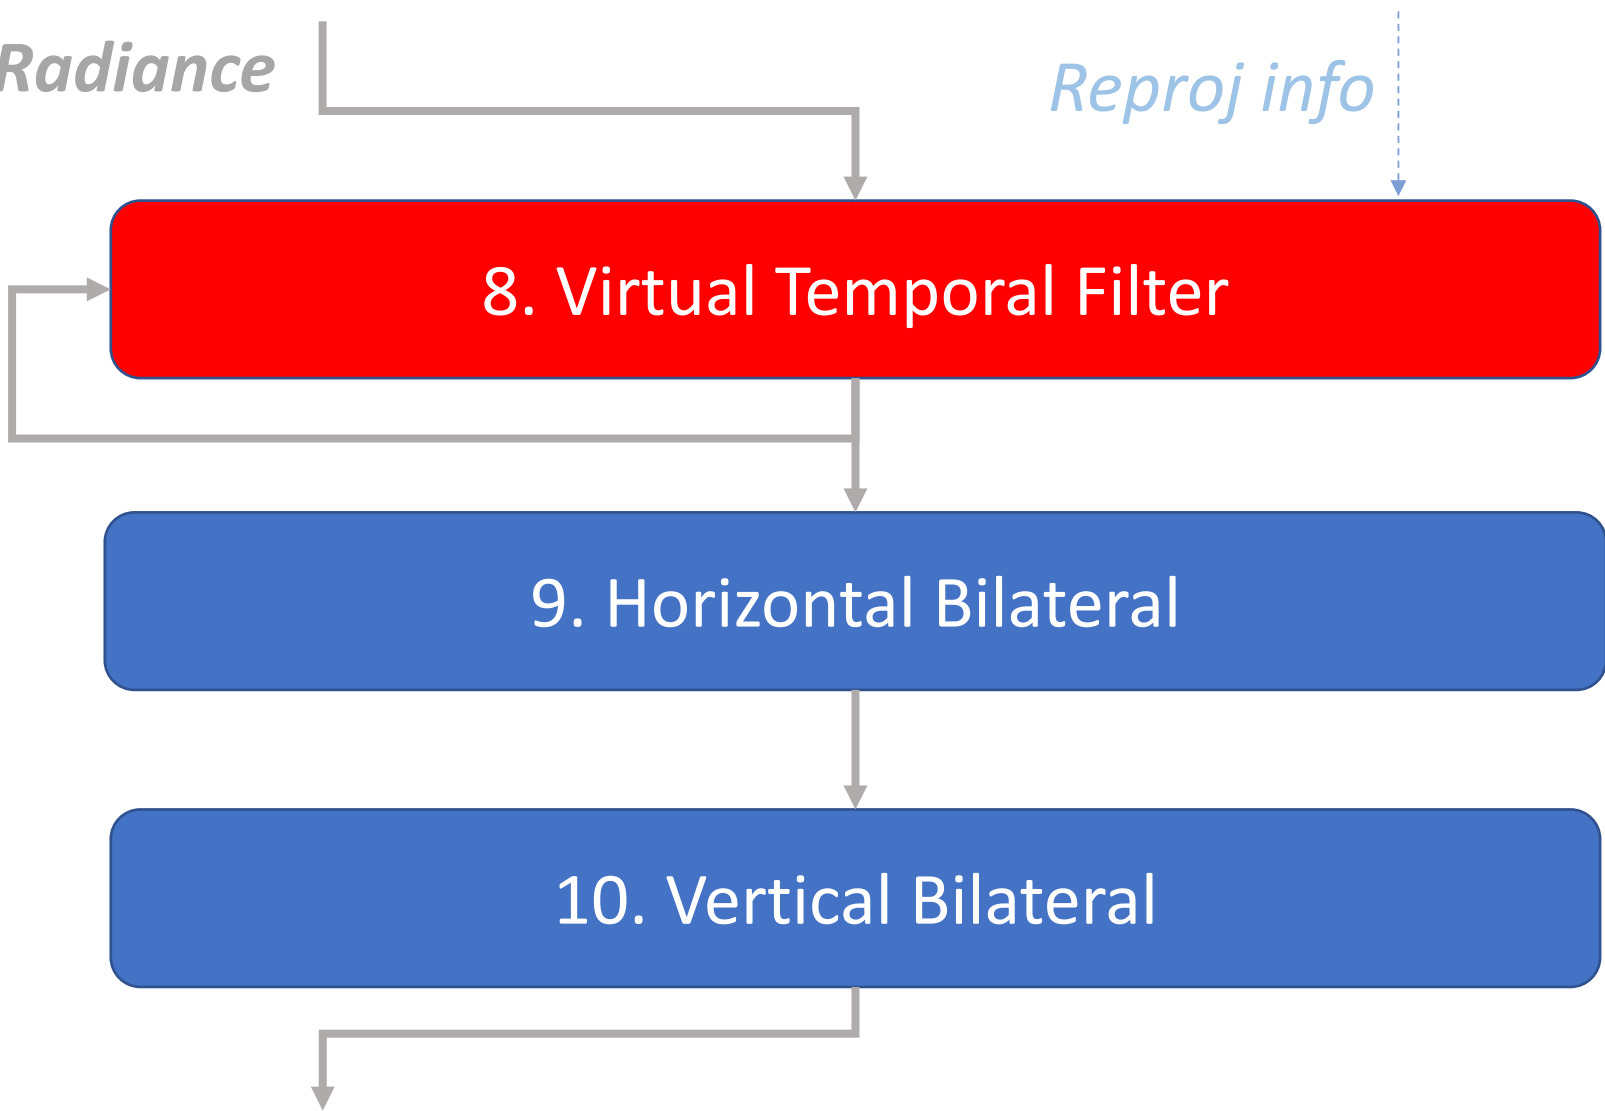
Radiance

Reproj info

8. Virtual Temporal Filter

9. Horizontal Bilateral

10. Vertical Bilateral





WITHOUT VIRTUAL GLOSSY REPROJECTION