Layered Variance Shadow Maps

Andrew Lauritzen and Michael McCool
Computer Graphics Lab, David R. Cheriton School of Computer Science, University of Waterloo

Variance shadow maps (VSMs) provide an efficient way to produce high-quality, filtered shadows in real-time, but suffer from light bleeding artifacts. Layered variance shadow maps automatically partition the depth range into multiple layers, which alleviates or eliminates light bleeding artifacts.

Compared to VSMs, LVSMs:
- Reduce or eliminate light bleeding artifacts
- Do not require high-precision (32-bit float) texture filtering
- Can still be generated in one pass and need only one texture sample per pixel
- Allow shadow quality to be easily scaled up at the cost of more storage
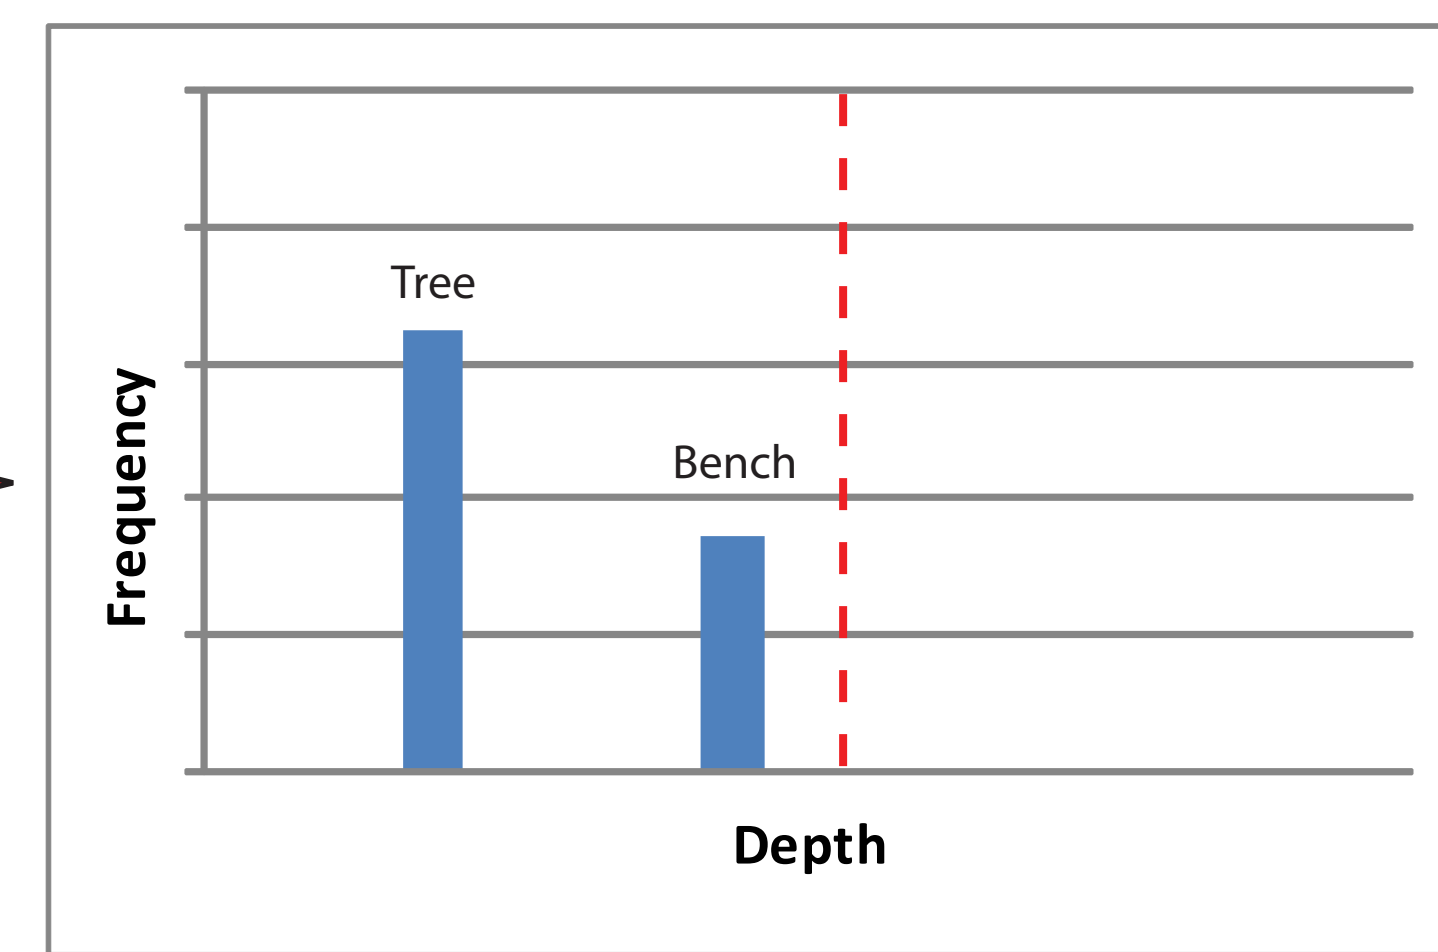
Variance Shadow Map

Layered Variance Shadow Map

Real-Time (60fps @ 1920x1200) Shadows Produced Using Our Algorithm
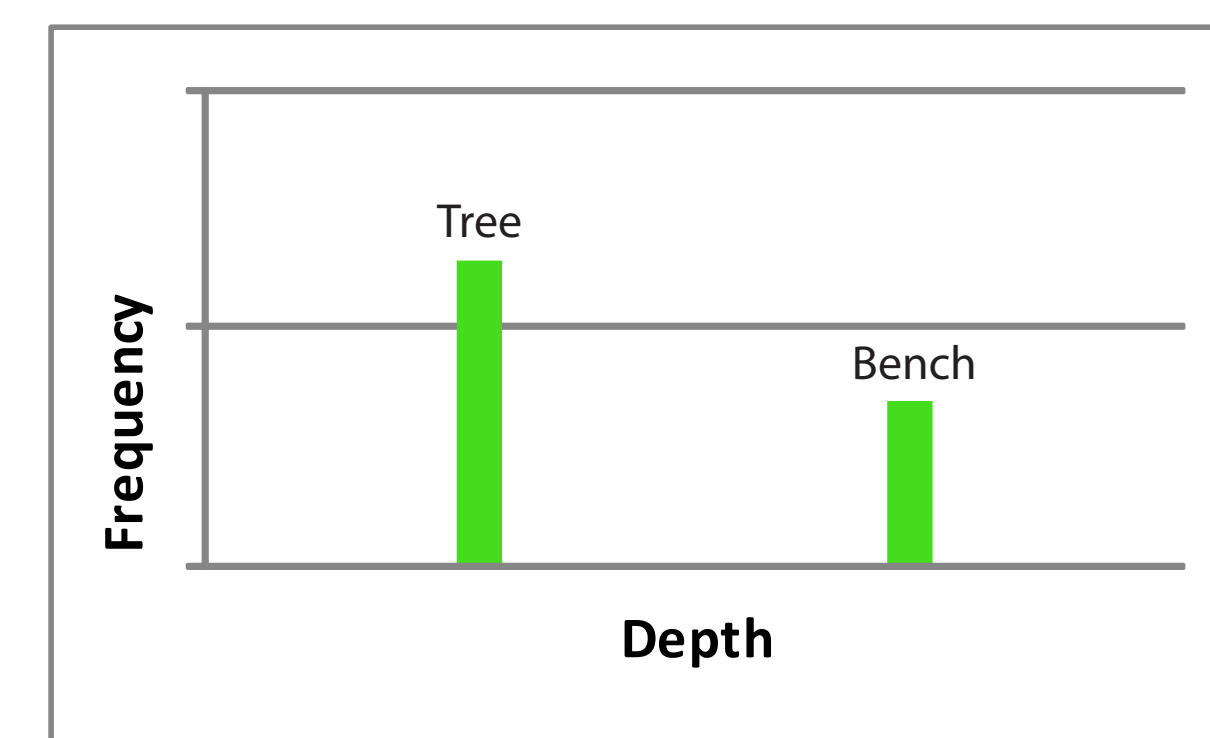
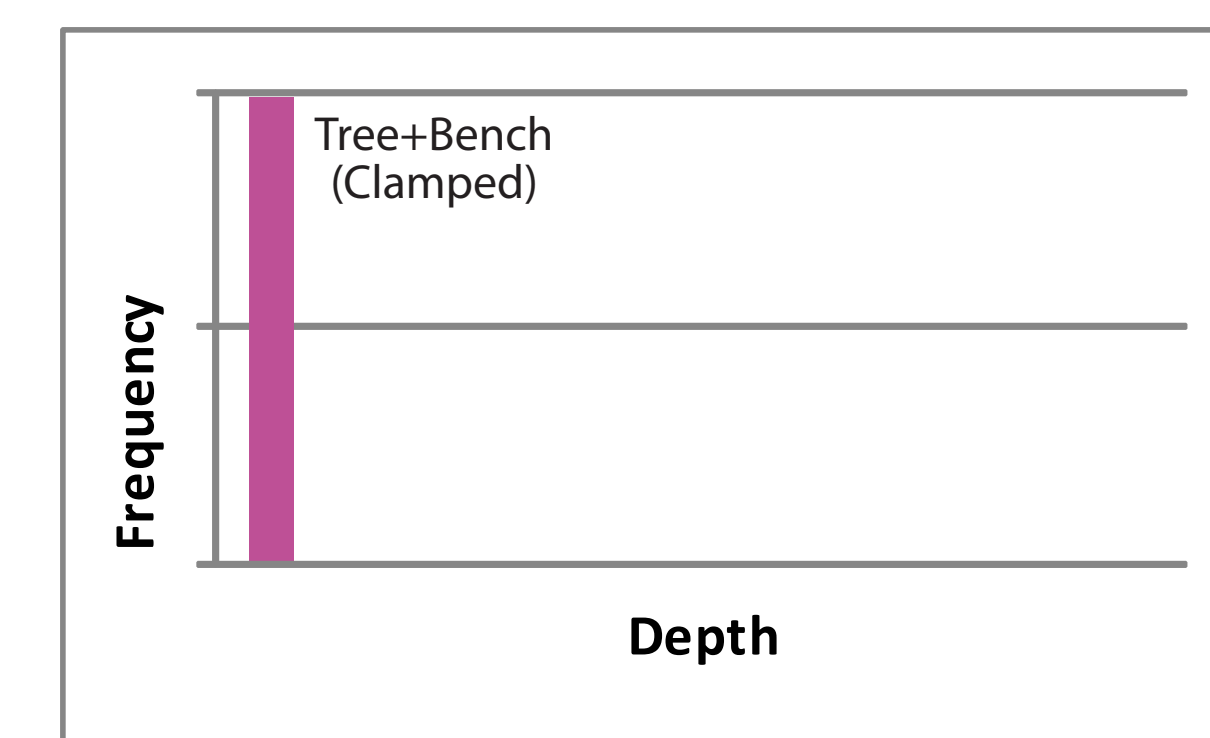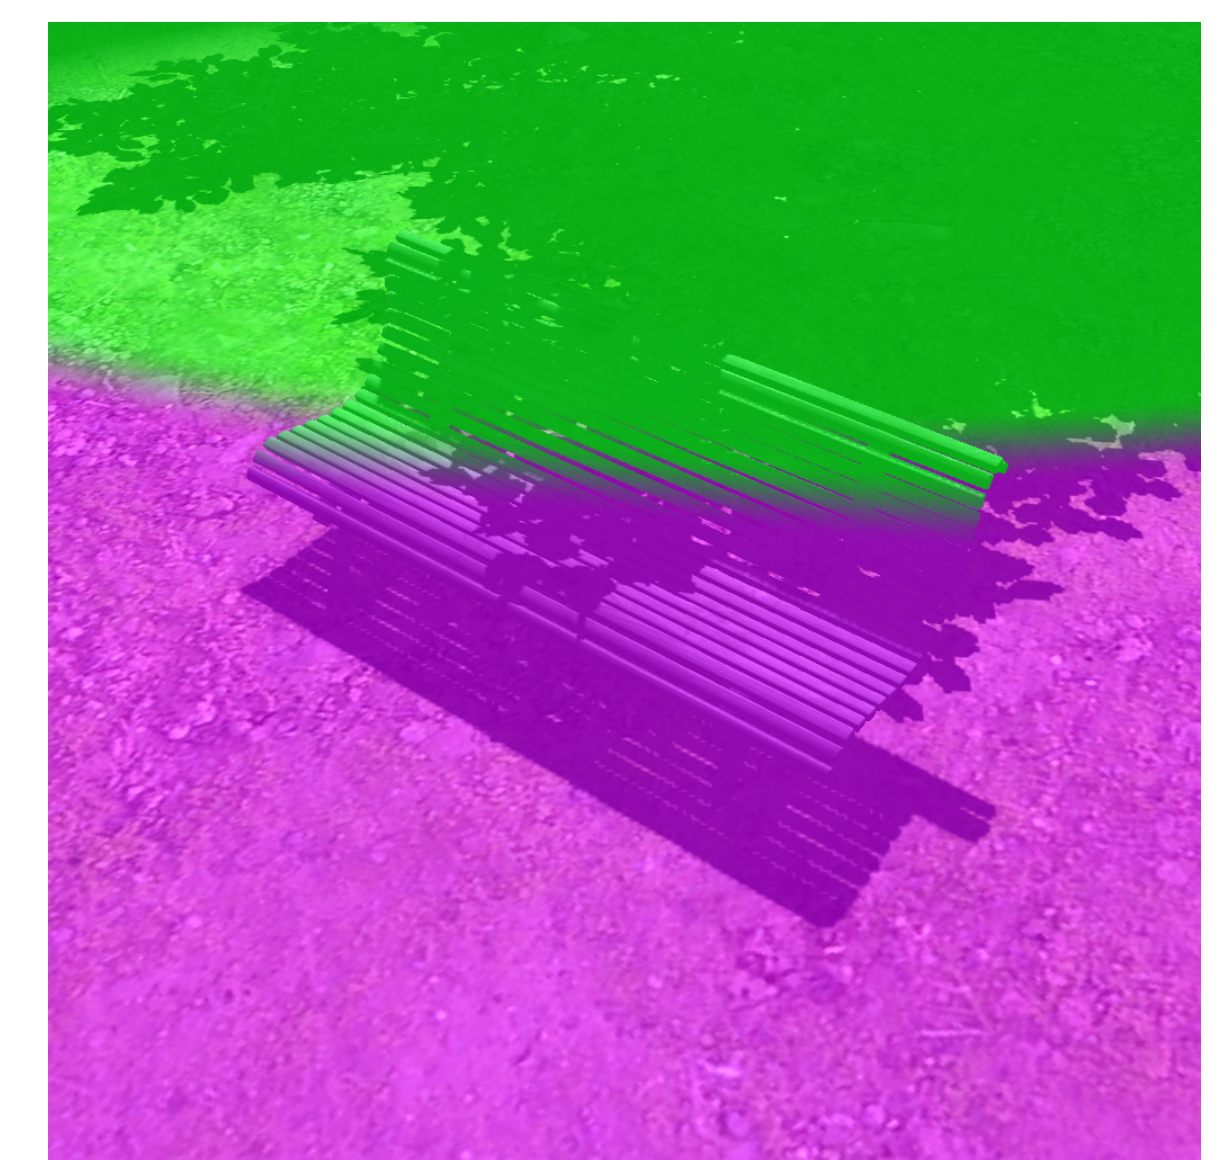Shadow Map Filter Region

Depth Distribution

Frequency

Tree

Bench

Depth

Layer 1

Frequency

Tree

Bench

Depth

Layer 2

Frequency

Tree+Bench (Clamped)

Depth

μ, σ^2

μ, σ^2

Chebyshev's Inequality

Chebyshev's Inequality

Reconstructed Visibility Function

Visibility

Depth

Tree   Bench   Ground

VSM
LVSM-1
LVSM-2
Ideal

Variance Shadow Map

Layers

Layered VSM