
Realistic Image Synthesis

- Virtual Point Light (VPL) Methods -


Philipp Slusallek
Karol Myszkowski
Gurprit Singh

Overview of MC GI methods

- **General idea**
 - Generate samples from lights and camera
 - Connect them and transport illumination along paths
- **Path Tracing (with Next Event Estimation)**
 - For each pixel, generate random path from camera
 - At each vertex, generate one light sample on lights and connect
- **BiDir-Path Tracing**
 - For each pixel, generate random paths from camera and from lights
 - Connect them (in all possible ways) and transport light
- **Instant Radiosity/Global Illumination**
 - In preprocessing, generate fixed set of samples from lights (VPLs)
 - During rendering, directly connect to all of them and transport light
- **Lightcuts**
 - Do not connect to all VPLs, but use importance sampling
 - Create a hierarchical structure to efficiently select VPLs

Reuse of Light Paths

- **Bidirectional path tracing**
 - Starts a new light path for every eye sample
 - Many new path being traced
 - No correlation between samples → noise
- **Idea: Reuse light samples**
 - Generate random light samples in a preprocessing pass
 - Each light sample becomes a Virtual Point Light (VPL) illuminating the entire scene (and not just one pixel)
 - Significantly reduces light samples and required tracing of rays
 - Generates correlated errors across entire image
 - Not unbiased -- but *consistent*



Instant Radiosity [Siggraph97]

- **Trace few (dozens of) rays from light sources**
 - These corresponds to the light paths from BiDir path tracing
 - Each contains a fraction of the energy of a light source
- **Use them to generate ‘virtual point lights’ (VPLs)**
 - VPLs placed at every hit point along the path
 - Termination via Russian Roulette (or fixed path length (biased))
 - Trace shadow rays to all of them during rendering
 - Contains both direct and indirect diffuse illumination
- **Inherently smooth, except for sharp shadow boundaries**
 - Shadow artifacts in case of few VPLs
 - But converges consistently



Instant Radiosity: Remarks

- **Approximation of illumination**
 - VPLs provide an approximation of the light distribution in a scene
 - Converges to real distribution with larger number of VPLs
- **Dealing with non-diffuse surfaces**
 - Consider BRDF when reflecting “photons” and during illumination
 - OK for mostly diffuse: Highly glossy surfaces would reveal VPLs and would require very large numbers of VPLs for glossy interactions
- **Shadow ray can be traced coherently**
 - Select VPLs in a coherent way (e.g., by clustering)
 - Shoot packets of rays to VPL clusters





Instant Radiosity: $1/r^2$ Bias

- **Illumination Singularity**

- Scene would be way to bright near a VPL
 - Illumination can become arbitrarily large ($1/r^2$)
 - Would be averaged out eventually with more VPLs
 - But may need arbitrary large number of VPLs
- Possible Solution
 - Limit contribution such that it will not be too bright
 - Limit the $1/r^2$ term by some constant b
 - → Introduces bias: too dark, specifically in corners

- **Bias Compensation [Kollig&Keller'06]**

- Add back missing contribution through MC sampling
- Continue the eye paths randomly
 - Check bias: $1/r^2 - b$
 - If non-negative: Scale contribution by $(1/r^2 - b)/(1/r^2)$
- Optimization
 - Limit ray length to critical region: $r < 1/\sqrt{b}$



Also see: Simon Brown (<http://sjbrown.co.uk/2011/05/09/virtual-point-light-bias-compensation/>)

Instant Global Illumination

- **Idea: Use coherent form of bi-directional path tracing**
 - How can we cut corners without too many artifacts
 - Speed up the computation
 - Ensure that shadow rays are coherent for packet ray tracing
 - Require no communication between rays
 - Combine advantages of several different algorithms
 - Instant Radiosity: smooth diffuse lighting
 - Ray Tracing: reflections, refractions, visibility testing
 - Interleaved Sampling (ILS): better quality, easy to parallelize
 - Discontinuity Buffer: removes ILS artifacts
 - Limitations of compute clusters (but similar on GPUs)
 - Cannot communicate between nodes (too high latency)
 - Streaming computations
 - Master sends stream of jobs to clients
 - They eventually return the results – while already working on the next job(s)
 - Achieves almost perfect speedup (pipelining)


[Interactive Global Illumination, Wald, Kollig, et al., EGWR 2002]

More Samples, (Almost) For Free

- **IR: Way to few samples for good results**
 - Hard shadow borders
- **Idea**
 - Use different sets of VPLs for different pixels
 - E.g.: 9 sets of VPLs
 - Every 3x3th pixel uses same set of VPLs
- **Better quality:**
 - 9 times as many VPL *per image* than without
- **Easily parallelizable**
- Each node computes pixels with
 - same VPL id
 - Nicely scales with #nodes
- **But: Massive aliasing**
 - Can obviously see 2D grid ...
 - Could be avoided if all samples are used within a pixel (supersampling)
 - but we aim at speed here!

1	2	3	1	2
4	5	6	4	5
7	8	9	7	8
1	2	3	1	2
4	5	6	4	5

VPL sets used per pixel



Remove Aliasing

- **Idea: Discontinuity Buffer [Keller, Kollig]**
 - Filter irradiance among neighboring pixels
 - Use the same filter width (3x3)
 - Smoothing/removal of ILS-artifacts
 - Like irradiance caching, but more stable
- **Only filter in smooth regions**
 - Must detect discontinuities
 - Criterion: normal & distance
 - Ignore pixels that differ too much
- **Problem: Clients don't have access to neighboring pixels!**
 - Filtering has to run on the server
 - High server load
- **Server has to get additional data**
 - Normal, irradiance, distance
 - High network bandwidth !
- **A lesser issue on GPUs, but still ...**

1	2	3	1	2
4	5	6	4	5
7	8	9	7	8
1	2	3	1	2
4	5	6	4	5




Finally, Add QMC

- **Use Randomized Quasi Monte Carlo [Keller et al.]**
 - Used for generating the VPLs
 - Faster convergence, especially for small sampling rates
 - Can be combined easily with interleaved sampling
- **Plus: ‘Technical’ advantages of QMC**
 - Fast random number generation (table lookup + bit-ops)
 - Can reproduce any sequence of samples based on single seed value
 - Can easily synchronize different clients on same data
 - Each client can reproduce the sample set of any other client
 - Avoid ‘jumping’ of VPLs:
 - Just start with same seed every frame
 - For progressive convergence, just advance the seed value...
 - QMC sequences perfectly combine into the future...

Summary

- **Base ingredient**
 - Instant Radiosity + Ray Tracing
 - Plus fast caustic photon maps
- **Combine with Interleaved Sampling**
 - Better quality
 - Parallelizable
- **Remove artifacts with Disco-Buffer**
 - Faster convergence
 - Better parallelizability
- **Use randomized QMC**
 - Low sampling rates, parallelizability
- **Result: Definitely not perfect**
 - But not too bad for *only ~20 rays/pixel!*





Remaining Issues

- **Missing scalability**
 - Render nodes each compute some tiles of the entire image
 - Dynamically selected at runtime (load balancing)
 - All data must be send to master for filtering
 - Colors, normals, depths
 - Approach limited by network bandwidth
- **Filtering on the clients?**
 - Would require information from neighboring pixes
 - All clients would need to compute all VPLs
 - Not great, but doable
 - They often did compute multiple sets anyway
 - Because of dynamic load balancing
- **Filtering overhead was too costly**
 - Repeatedly test and sum up blocks of 3x3 pixels

Scalable IGI

- **Approach**

- Assign tiles of pixels to clients
 - Must include border for filtering
 - Overhead is ~10% for 40x40 pixels
- Trace ray in interleaved coherent packets
 - Use SIMD across pixels with the same VPLs
- Filtering of tile can be done on client
 - Constant time filtering using running sums
 - Add/subtract only at border of domain
 - Must only send final color
- Low-cost antialiasing
 - Nx supersampling: Assign VPLs into N groups
 - Trace different primary ray for every group
 - Connect this hit point to VPLs of group and average (again interleaved sampling)
- RQMC sampling of light sources
- SIMD shader interface




IGI with 50 Mio polygons

Results

- **Performance**

- Faster by 2.5-3x
- Almost perfect scalability (> 20 fps) plus good use of coherence



	Office	Conference	Power Plant
640x480	1.72 (1.41)	1.12 (0.85)	0.33 (0.28)
800x600	1.77 (1.45)	1.22 (0.94)	0.42 (0.29)
1000x1000	1.84 (1.49)	1.33 (1.03)	0.44 (0.31)
1600x1200	2.00 (1.63)	1.46 (1.09)	0.48 (0.34)

Table 3: Million rays per second on AthlonMP 1800+ CPUs at different resolutions with 16 VPLs, full shading and filtering. Using ray packet traversal the new system offers sub-linear costs in the number of pixels. The number in parenthesis are for dynamic environments, which impose a ray tracing overhead due to additional processing of scene changes.



Equal compute time images comparing old and new scalable approach


Bidirectional Instant Radiosity

- **Idea [Segovia, EGSR 06]**


- Generate VPLs where they matter
- Each VPL should have equal contribution to the *image*

- **Bidirectional Approach**

- Generate “VPLs” from light AND camera
 - N/2 samples each
- Estimate illumination at reverse VPLs using Instant Global Illumination
 - Shoot some paths each (e.g. ~5)
 - Gather light from forward/light VPLs
 - Reverse VPLs act as proxies
- Estimate importance of each VPL using M (e.g. 10) paths from camera (length 2)
- Resample VPLs (e.g. select 10%) according to contribution to camera with few samples (e.g. 5)
- Estimate accurate pdf for VPLs using more camera path (e.g. 50)
- Use selected VPLs during rendering with importance sampling



(a)



(b)



(c)



(d)


200 selected VPLs

Results

	standard <i>IR VPLs</i>	Reverse <i>IR VPLs</i>
Conference	8.7 %	1.3 %
Cruiser	6.2 %	3.8 %
U-Office	0%	10 %
Q3tourney1	2.1 %	7.9 %

Table 1: *Percentage of resampled VPLs* The resampling rate is equal to 1:10. According to the scene and the point of view, one method or the other provides the more relevant VPLs. For the conference room scene, most resampled VPLs are standard VPLs. For U-Office, the Reverse IR strategy is more efficient.

Results



(a)



(b)



(c)

Figure 6: Sampling Performances. Several pictures computed with the sampling techniques described in the article. Bidirectional sampling finds the relevant VPLs no matter the visibility layout. (a) A part of *q3tourney1* (courtesy of ID software) indirectly lit through a small corridor. Lights make at least 3 bounces before coming into the room. Most of the VPL are found with reverse IR strategy. (b) A simple office indirectly lit by a halogen lamp easily handled with Standard Instant Radiosity. (c) *The U Office* (raytraced). The scene is completely indirectly lit through a small corridor. Bidirectional IR quickly finds the relevant VPLs.

Dealing With Many Lights [EGSR'03]

- **Problem**

- Efficiency drops severely in highly occluded environments
 - Think: Large building with many (illuminated) rooms
- Probability of light being visible is low
 - Must generate many VPLs to get a good statistics, at all
 - Must send many shadow rays to VPLs that are almost certainly occluded

- **Idea**

- Ignore any lights that do not contribute illumination
- Avoid computing from lights, would load data for entire (huge) scene

- **Solution**


- Estimate importance of lights using path tracing (1 path per pixel)
 - From points to lights: Stops rays at walls, only touches visible geometry
 - Gives ~2 million samples (HD), could also average over last few frames
- Use importance sampling to distribute VPLs from lights

- **Issues**


- Average is over entire image (might miss lights illuminating small area)
- Can cause temporal aliasing (flickering) due to randomness of VPLs
 - Somewhat offset by deterministic QMC sampling (mostly same VPLs/light)

Estimate Importance of Lights

- Example: “Shirley10” – 10x10 rooms connected by doors



Estimate (1 sample/pix)



Real scene (10x10 rooms,
1 light each)

- Path traced image hardly recognizable ...
 - But: Estimate correct up to a few percent
- Used for importance sampling of lights

Lightcuts

- **Goals:**

- Efficient, accurate complex illumination
- In realistic and complex environments



Environment map lighting & indirect
Time 111s



Textured area lights & indirect
Time 98s


(640x480, anti-aliased, glossy materials)

Motivation

- **Hierarchies in Global Illumination**
 - Only used in FE methods so far
 - Can greatly improved performance
 - Take advantage of $1/N^2$ power fall-off
 - Group together light from distant objects & handles it together
 - Can reduce computational complexity from $O(N^2)$ to $O(N)$
- **Question: How to use them in MC-style algorithms**
 - Key idea: Sample points generated from lights and from camera
 - Could group them hierarchically, if generated in advance
 - Could handle illumination of a group as one sample
 - Allows adaptive/progressive refinement
 - Key issues:
 - How to group: Must have criteria for grouping (e.g. by “similarity”)
 - When to refine: Must have an efficient “oracle”


Lightcuts Problem

- Many light samples




Lightcuts Problem

- **Complex visibility**




Lightcuts Problem

- **Material properties with complex reflection**




Key Concepts

- **Light Cluster**
 - Approximate many lights by a single brighter light (the representative light)



Clustering of Light Samples

- **Sources of (many) light samples**
 - Point lights
 - Sampled area lights
 - Sampled HDR environment lighting
 - Generated secondary lighting samples (VPLs in IGI)
- **General idea**
 - Group light samples into binary tree
 - Leafs are the input light samples
 - Inner nodes combine illumination from their children
 - Choose a representative location from among children
 - Combine and bound attributes
 - Illumination uses a **cut** through the tree
 - Adaptively combines far away lights into one
 - Samples the integral evenly given bounds on power contribution, solid angle, visibility, and angular falloff



Criteria for Clustering

- **Contribution from a cluster**

- Given terms for material (M), geometry (G), visibility (V) and the intensity (I) of the (clustered) child light samples
- Illumination from the cluster is then given as

$$L_C = \sum_{i \in C} M_i(x_i, \omega_o) G_i(x_i) V_i(x_i) I_i$$

- **Approximation**

- However, this is too costly and is approximated as by a representative light sample j

$$\tilde{L}_C \approx M_j(x_j, \omega_o) G_j(x_j) V_j(x_j) \tilde{I}_j \quad \tilde{I}_j = \sum_{i \in C} I_i$$

- All properties are taken from representative, except light intensity
- Create a full cluster up to a single root node

- **Issue**

- Must have some way to bound the error of the approximation




Building the Light Tree

- Lights are split into types: Omni, oriented, and directional lights
 - Build a tree for each (but conceptually one big tree)
 - Directional lights are handled as point lights on a unit sphere
- Each cluster stores
 - Links to two children
 - Representative light (randomly chosen among children, \sim intensity)
 - Total intensity I_C (sum over all children)
 - Axis aligned bounding box
 - Oriented bounding cone (for oriented lights)
- Greedy bottom-up build:
 - In each step create cluster that minimizes total cost
- Cost model: $I_C(\alpha_C^2 + c^2(1 - \cos \beta_C)^2)$
 - α_C : Diagonal length of bounding box
 - β_C : Half angle of bounding cone (of light directions)
 - c : Constant for relative scaling of spatial/directional data
 - Set to half the scenes Bbox for oriented lights, zero otherwise

Choosing a Cut

- **General Approach**


- Set the cut to be the root node
- Choose the node from the cut with worst error
- Refine this node
 - Replacing it with its two children
- Terminate if relative error is below 1%
 - Can be computed because we have approximated illumination due to existing cut
 - Criterion due to Weber's law
 - Relative perception
 - In the paper they use 2% without artifacts



Illumination Equation

$$\text{result} = \sum_{\text{lights}} M_i G_i V_i I_i$$


Material term
Geometric term
Visibility term
Light intensity



Illumination Equation

$$\text{result} = \sum_{\text{lights}} M_i G_i V_i I_i$$


M_i Material term
 G_i Geometric term
 V_i Visibility term
 I_i Light intensity



Illumination Equation

$$\text{result} = \sum_{\text{lights}} M_i G_i V_i I_i$$


M_i | Material term
 G_i | Geometric term
 V_i | Visibility term
 I_i | Light intensity



Cluster Approximation


$$\text{result} = \sum_{\text{lights}} M_i G_i V_i I_i$$

Material term *Geometric term* *Visibility term* *Light intensity*



Cluster Error Bound

$$\text{error} \leq M_{\text{ub}} G_{\text{ub}} V_{\text{ub}} \sum_{\text{lights}} I_i$$




Bound each term

- Visibility ≤ 1 (trivial)
- Intensity is known
- Bound material and geometric terms using cluster bounding volume



Kitchen, 388K polygons, 4608 lights (72 area sources)



Kitchen, 388K polygons, 4608 lights (72 area sources)

Combined Illumination



Lightcuts 128s
4 608 Lights
(Area lights only)



Lightcuts 290s
59 672 Lights
(Area + Sun/sky + Indirect)

Combined Illumination



Lightcuts 128s
4 608 Lights
(Area lights only)
Avg. 259 shadow rays / pixel



Lightcuts 290s
59 672 Lights
(Area + Sun/sky + Indirect)
Avg. 478 shadow rays / pixel
(only 54 to area lights)

Extended Versions of Lightcuts

- **Reconstruction Cuts**

- Operates in image space
- Starts Lightcuts at coarse pixel grid (e.g. 16x16 pixels)
- Interpolates either colors or lighting info, or resamples
- Refines pixel grid where necessary (based on material, shadow info)

- **Multi-Dimensional Lightcuts**

- Realizes that antialiasing, motion blur, etc. require many samples per pixel
 - Inefficient if Lightcut is recomputed for each of them
- Instead build hierarchy of pixel samples and VPLs
 - Needs clever error bounds
- Traverse simultaneously, subdividing either cut based on cost function