

# Scalability with Many Lights 1

## Lightcuts & Multidimensional Lightcuts

Bruce Walter, Cornell University

Course: Realistic Rendering with Many-Light Methods



Note: please see website for the latest slides.

Monday, August 6, 2012

These slides are for the SIGGRAPH 2012 course Realistic Rendering with Many-Light Methods. This part is titled Scalability with many lights, part one and will be presented by Bruce Walter from Cornell University. These slides are subject to change until the actual time of presentation, so please see the course website to get the latest version.

# Talk Overview

- Scalable solutions for many light rendering, Part I
  - Illumination at a single receiver: Lightcuts
    - » “Lightcuts: a Scalable Approach to Illumination” by Walter, Fernandez, Arbree, Bala, Donikian, Greenberg, SIGGRAPH2005
  - Illumination over a pixel: Multidimensional Lightcuts
    - » “Multidimensional Lightcuts” by Walter, Arbree, Bala, Greenberg, SIGGRAPH2006

# Why many lights?

- Simulate complex illumination using point lights
  - Area lights
  - HDR environment maps
  - Sun & sky light
  - Indirect illumination
- More lights → more accurate
  - And more expensive
  - Naive cost linear in lights



Area lights + Sun/sky + Indirect

# Accurate Approximation

- Do we really need to evaluate all lights?
  - Thousands to millions of lights
  - Average contribution miniscule
  - Below human perceptual limits
    - ▶ Weber's Law (roughly 2%)
- Evaluate a small subset
  - Chosen adaptively
  - Perceptually accurate approximation



Textured area lights & indirect

# Scalable Algorithm

- Scalable solution for many point lights
  - Sub-linear cost per light

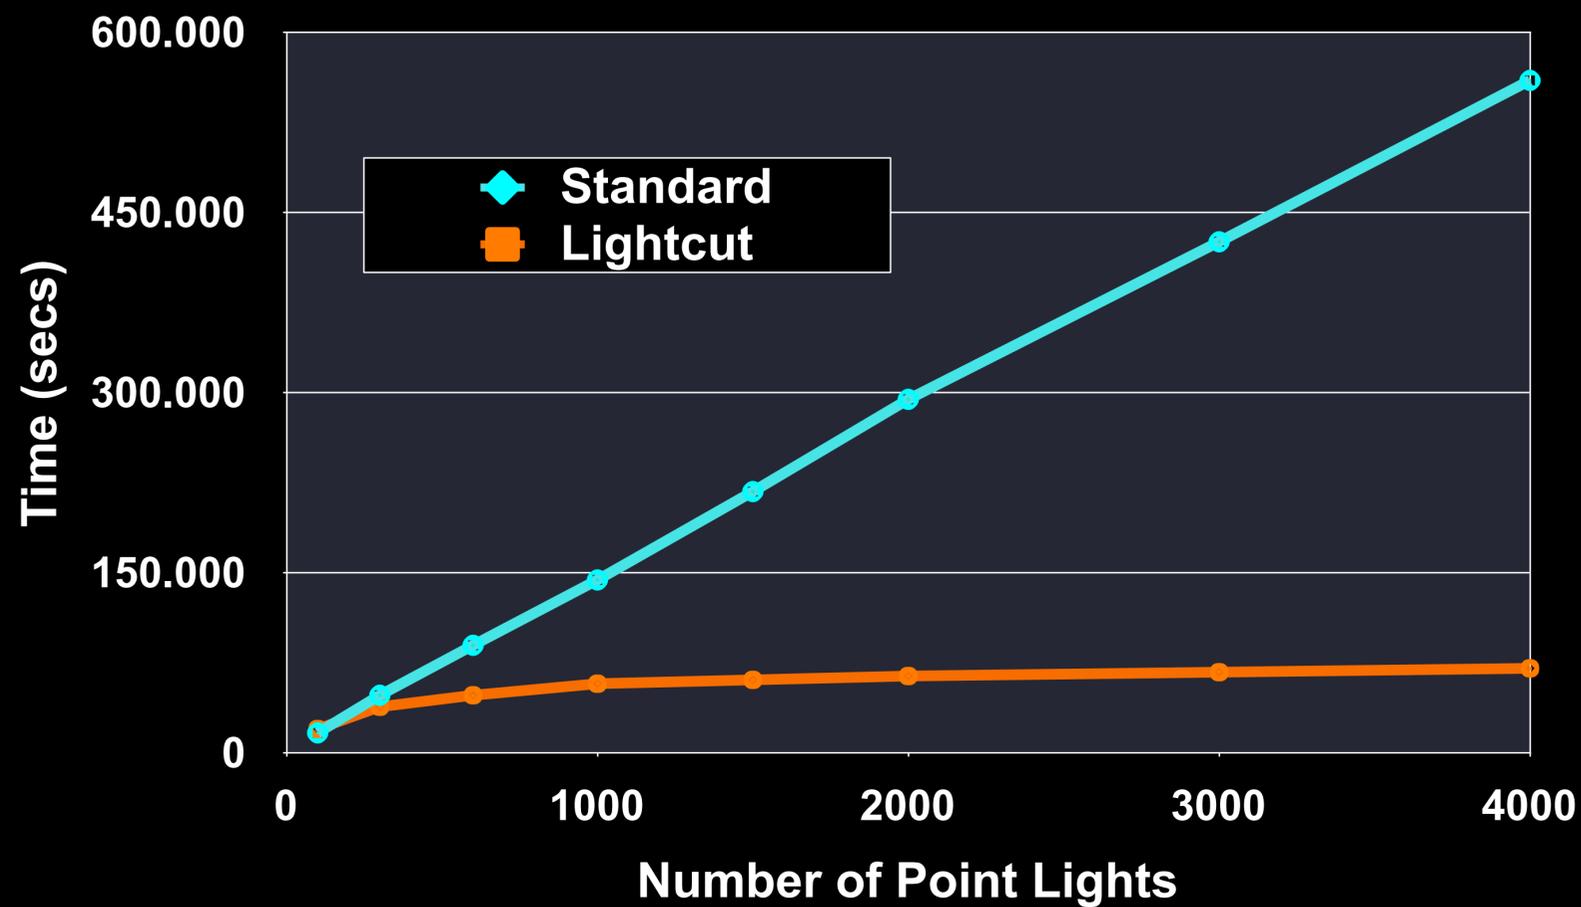
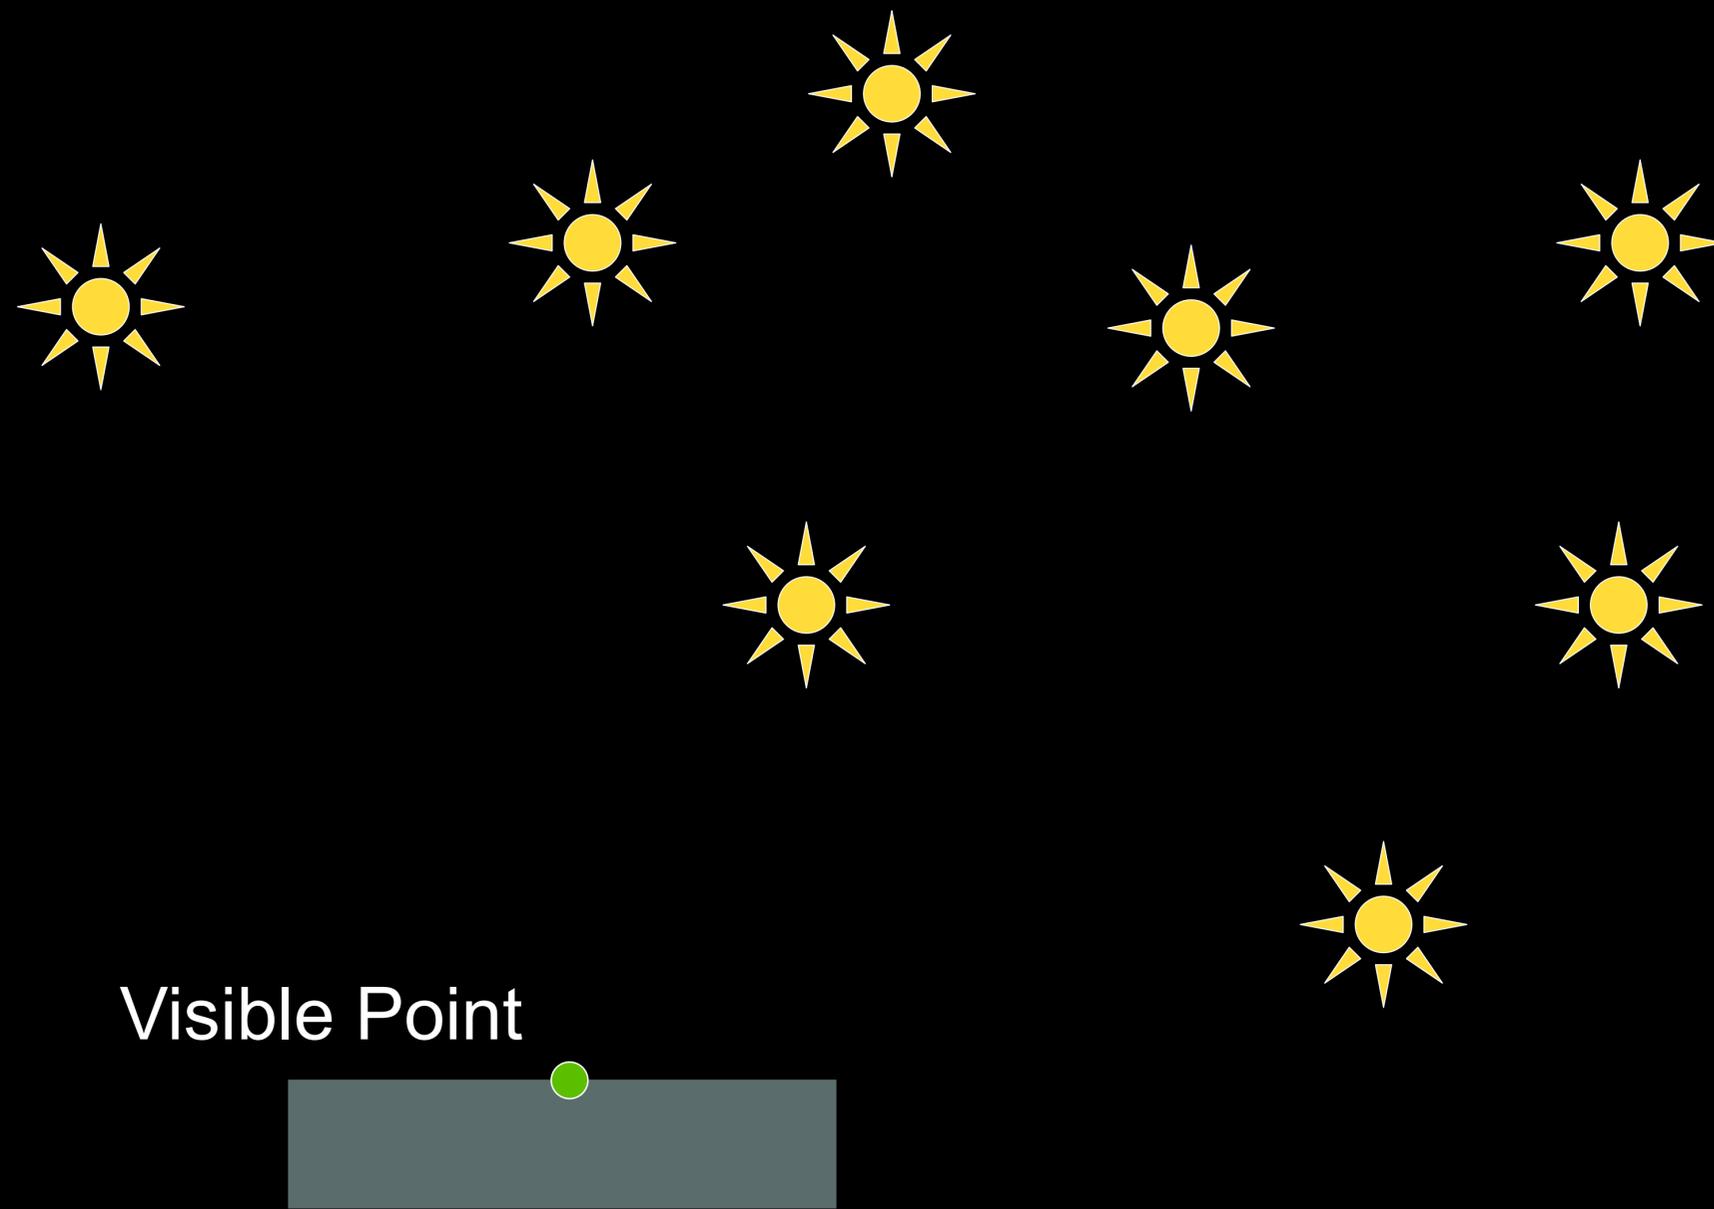


Tableau Scene  
Environment map lighting & indirect

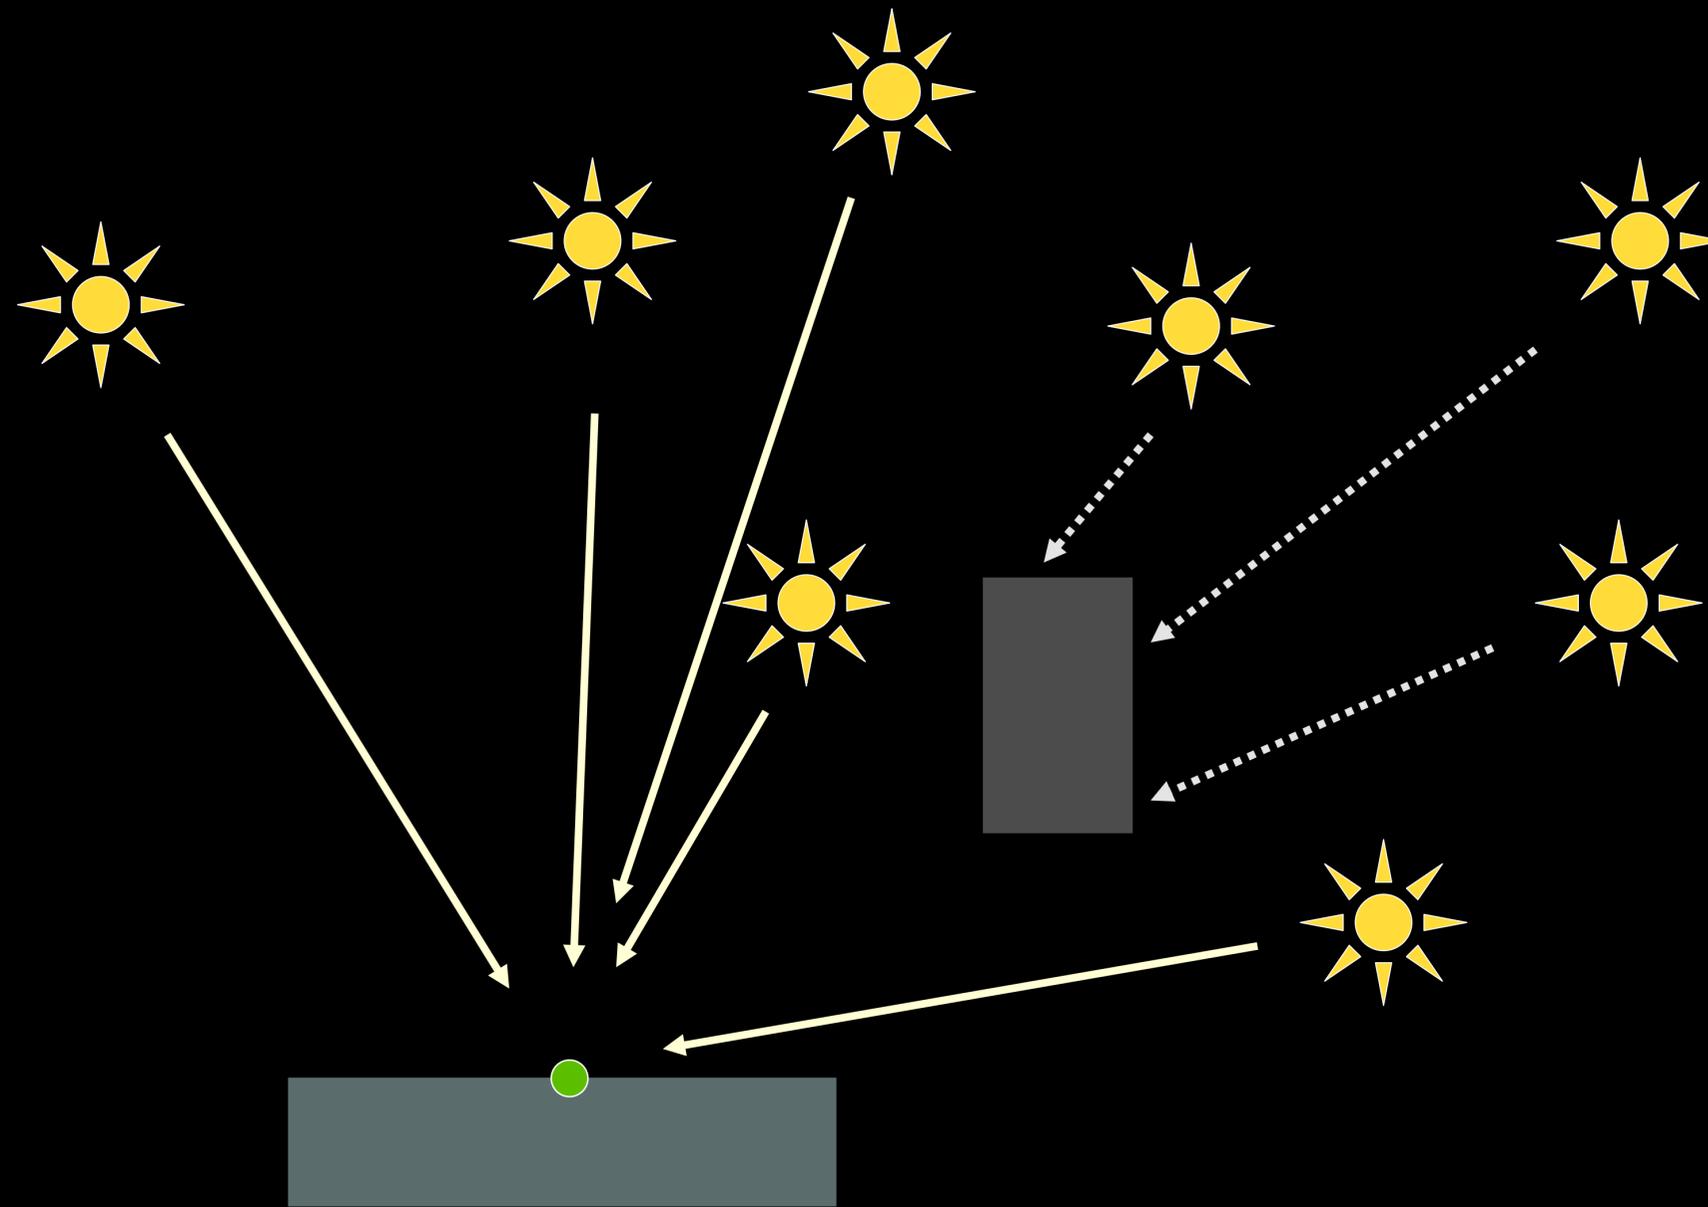
Monday, August 6, 2012  
 The core of lightcuts is a new scalable algorithm for efficiently approximating the light from many point lights. By many I mean thousands to millions. Here we show the time to compute this tableau scene with environment map illumination using varying numbers of lights. Evaluating each light individually gives a cost that increases linearly with the number of lights. Lightcuts' cost is strongly sub-linear and thus its advantage grows dramatically as the number of lights increases.

# Lightcuts Problem

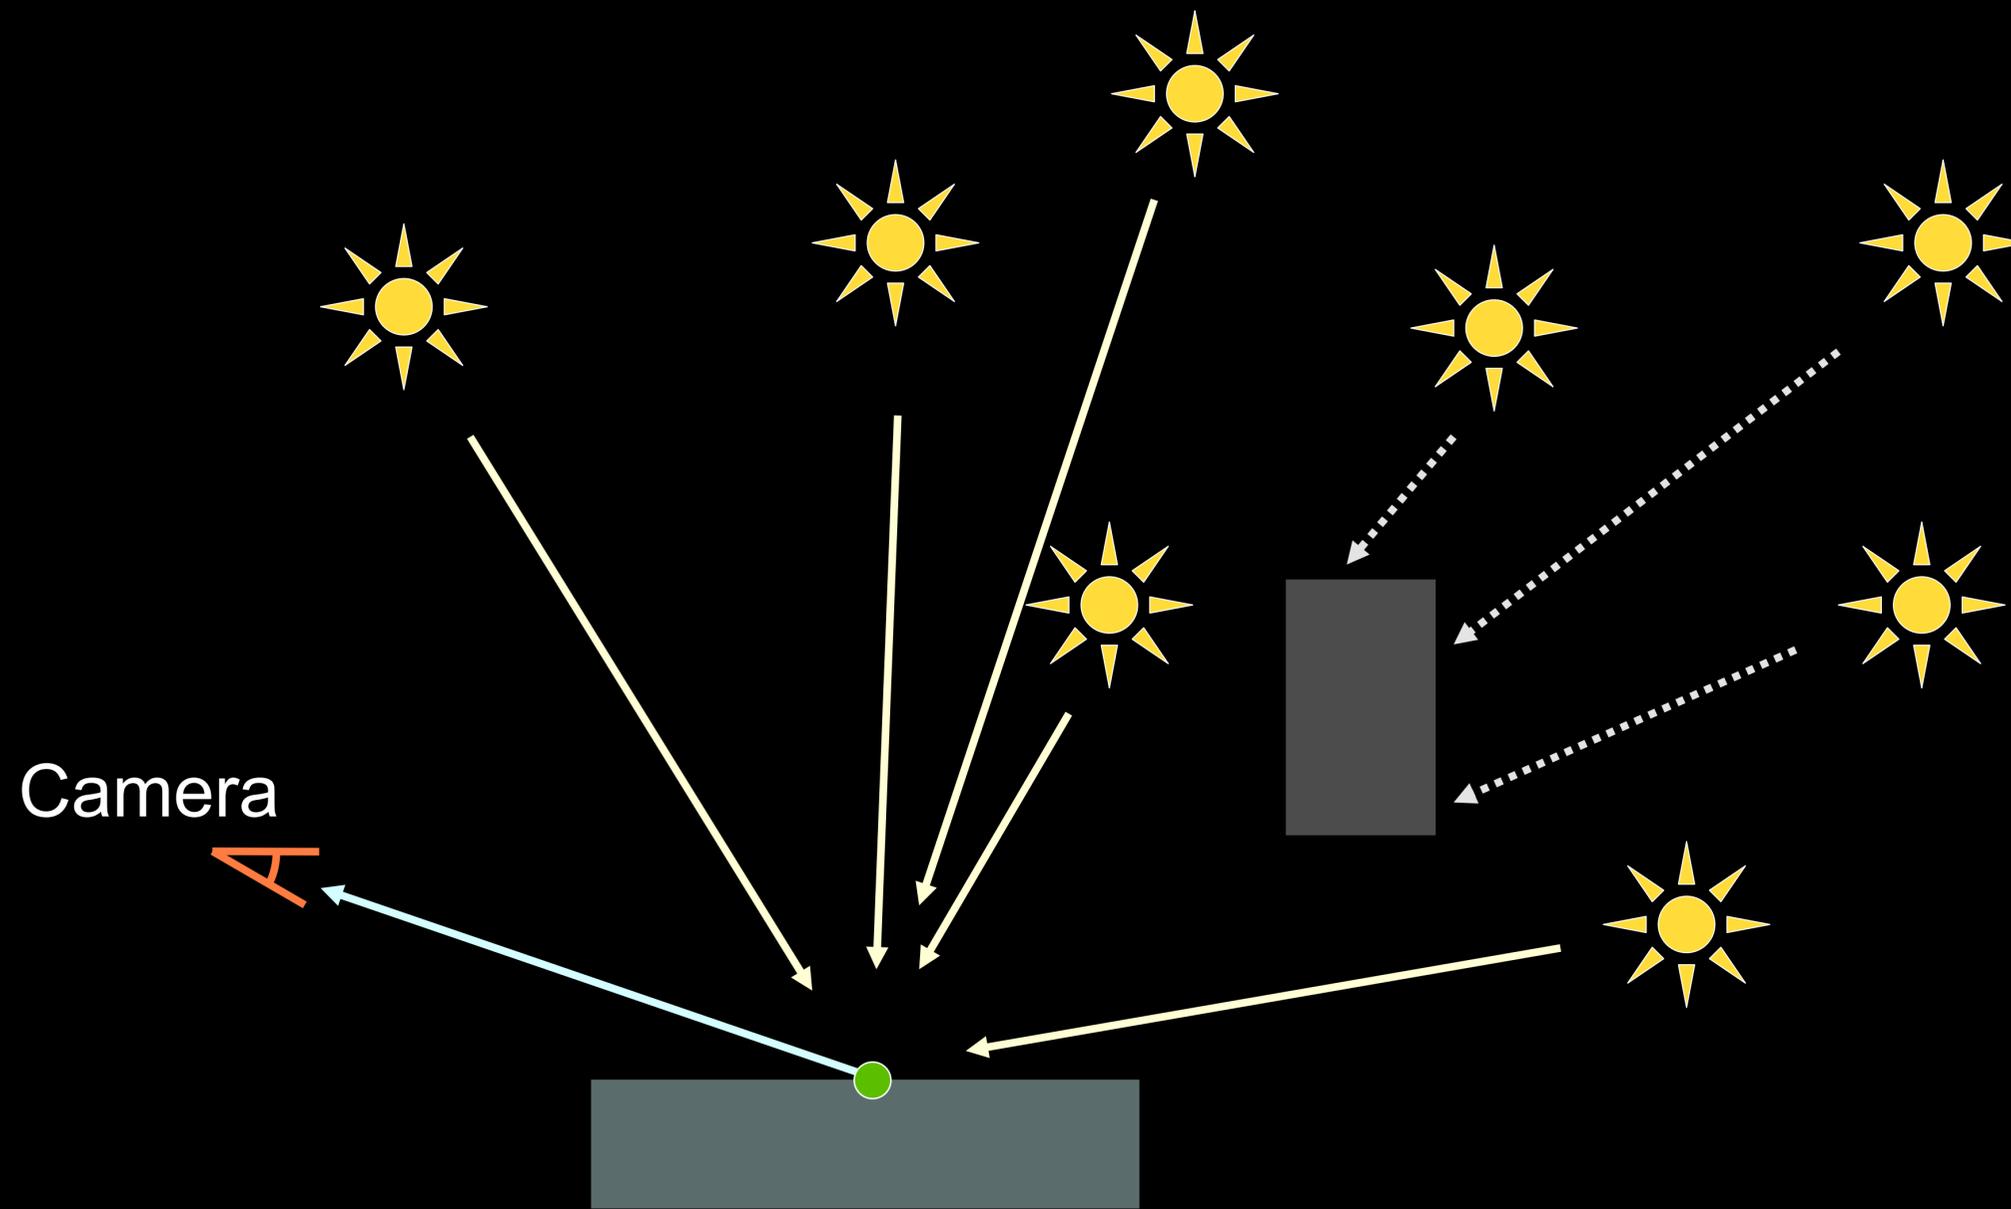


Monday, August 6, 2012  
Given a set of points we want to compute their contribution at some point of interest.

# Lightcuts Problem



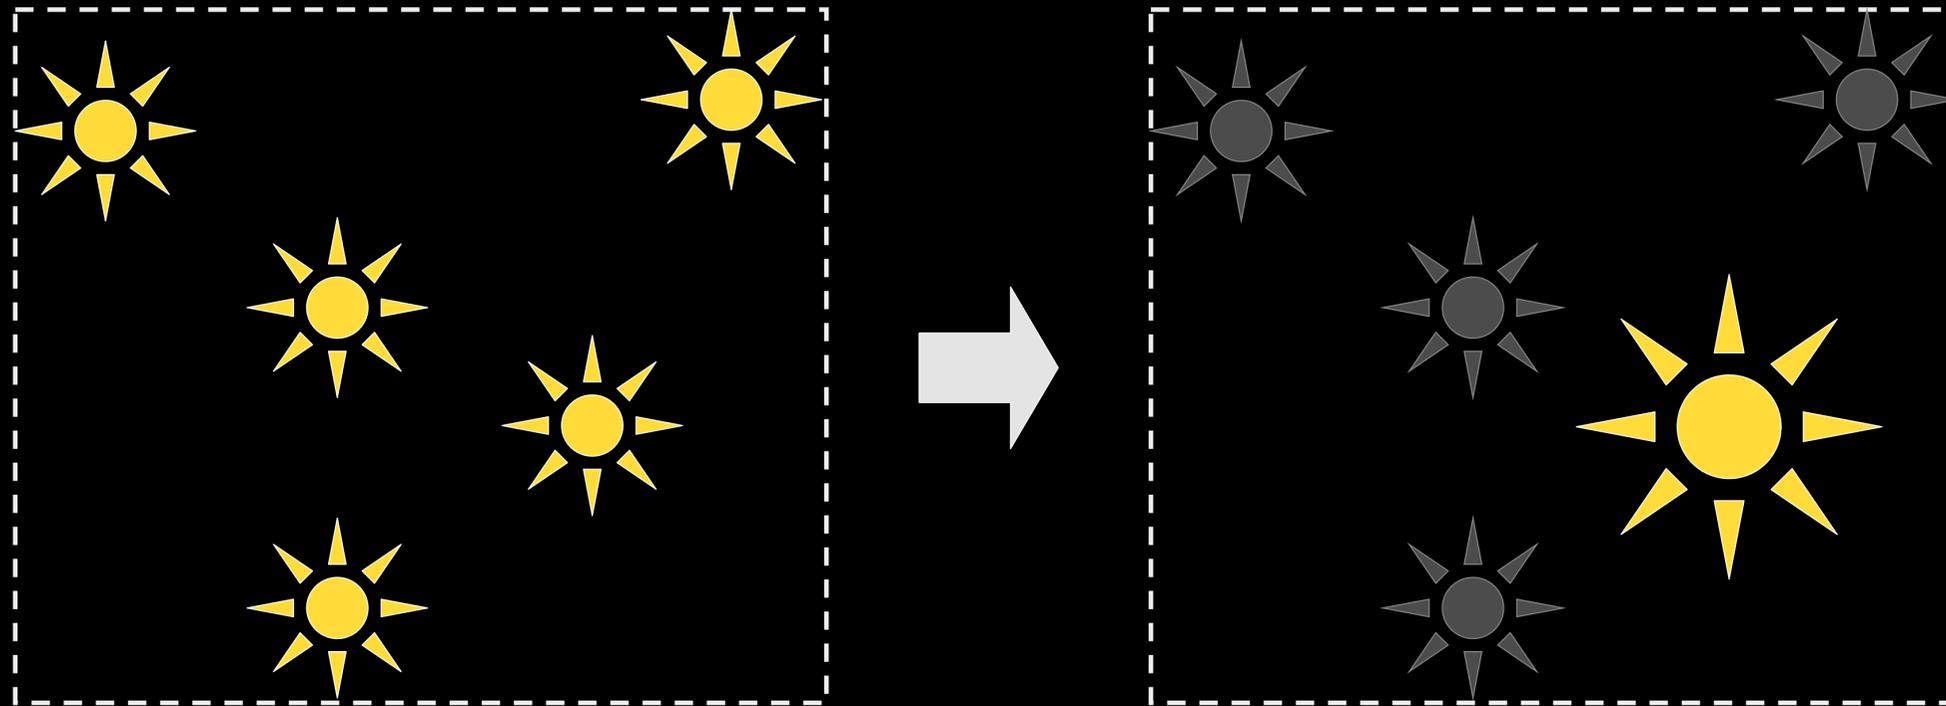
# Lightcuts Problem



# Key Concepts

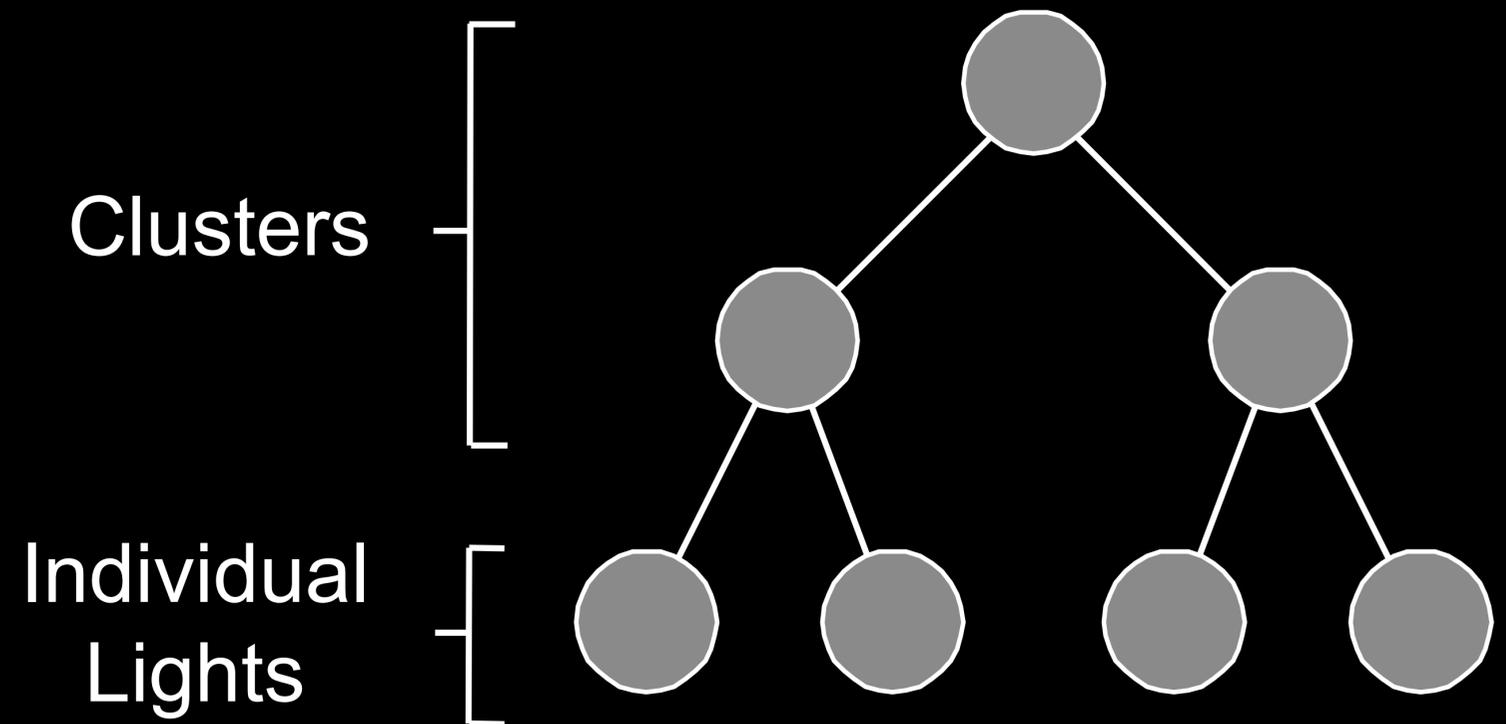
- Light Cluster

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



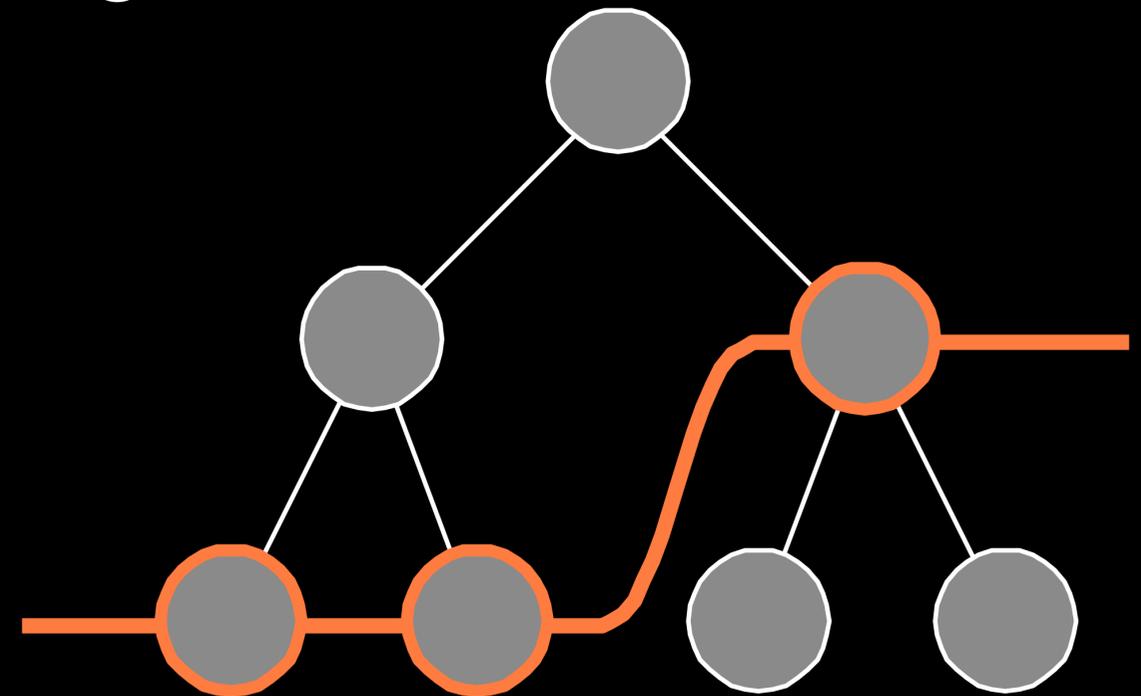
# Key Concepts

- Light Cluster
- Light Tree
  - Binary tree of lights and clusters

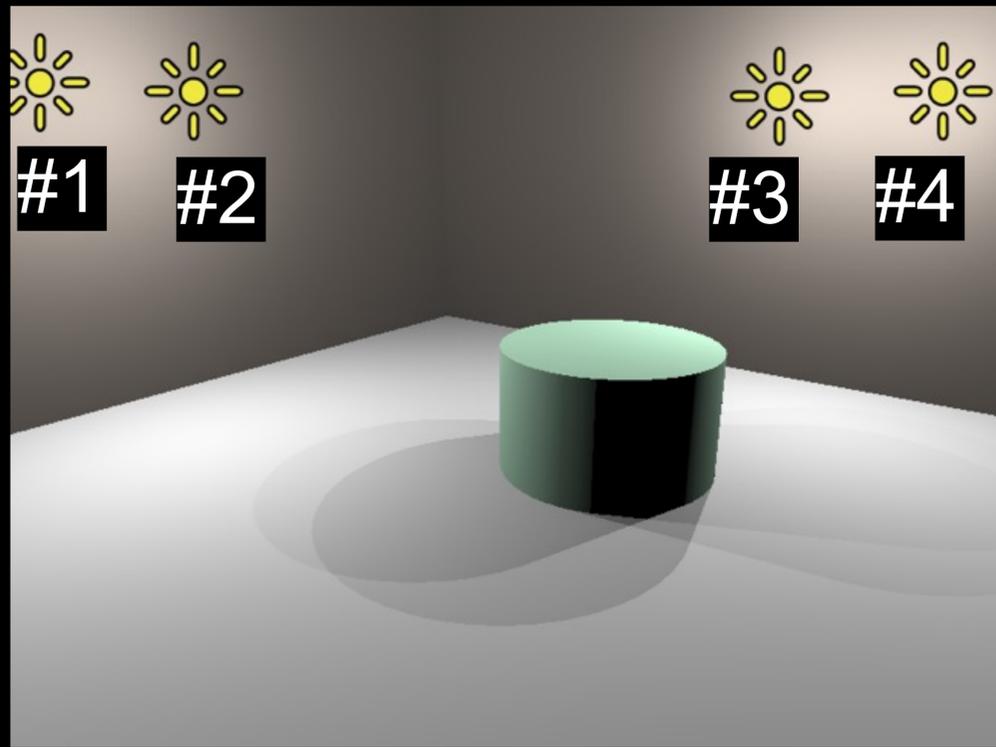


# Key Concepts

- Light Cluster
- Light Tree
- A Cut
  - A set of nodes that partitions the lights into clusters

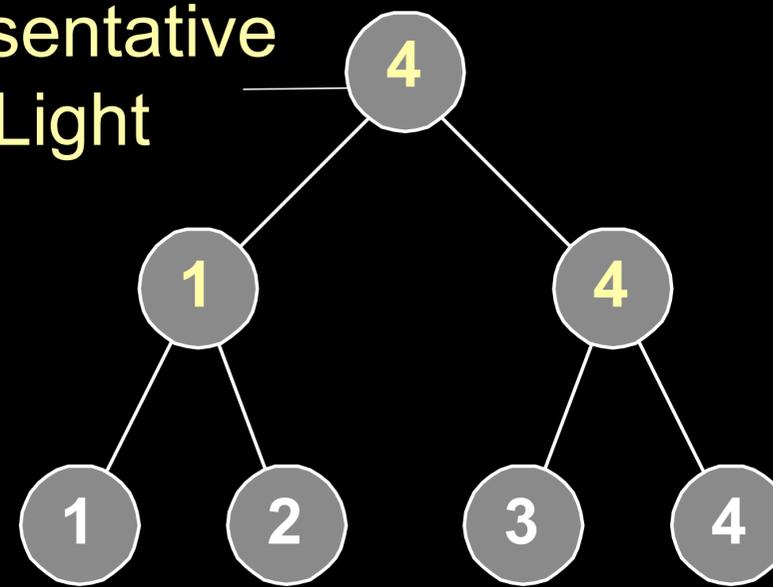


# Simple Example



## Light Tree

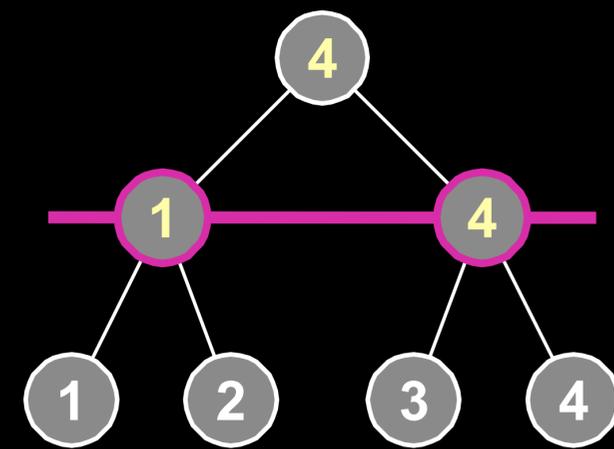
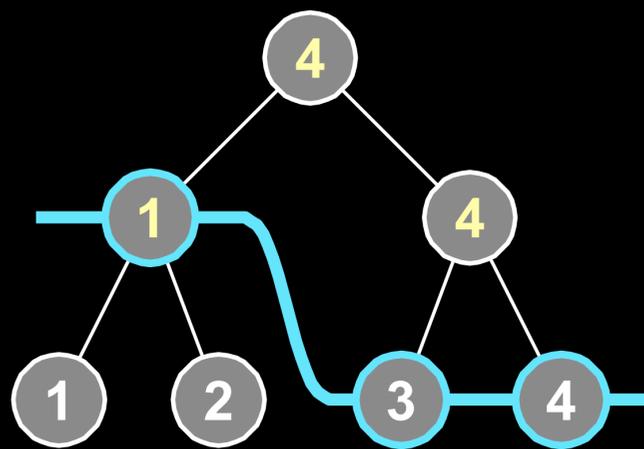
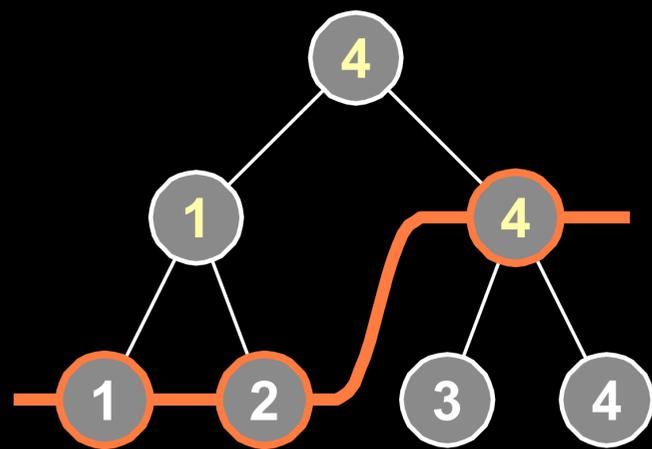
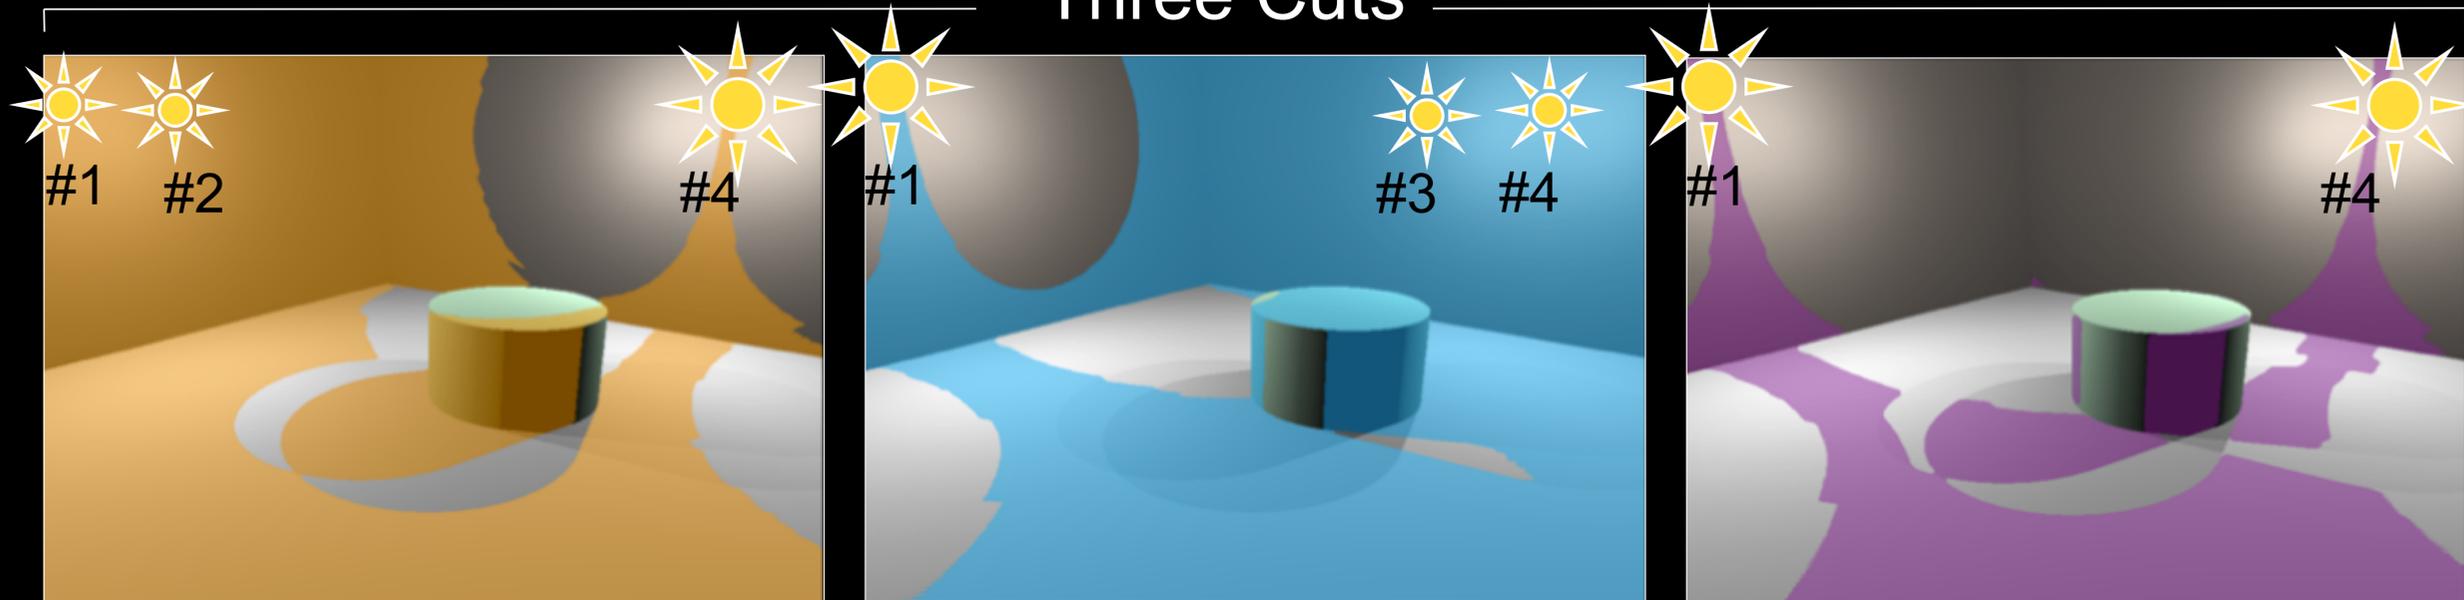
Representative Light



Clusters  
Individual Lights

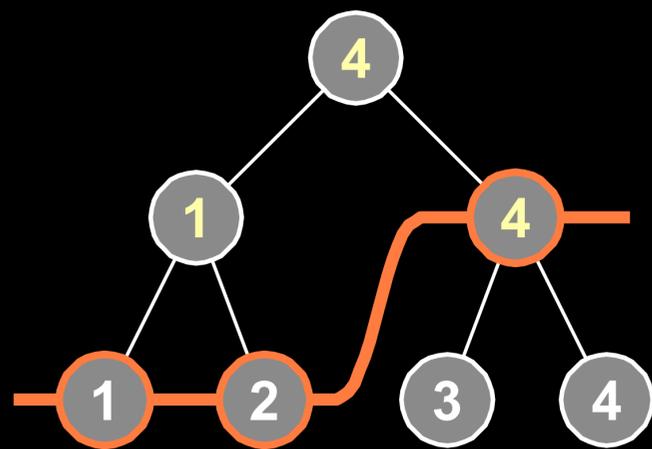
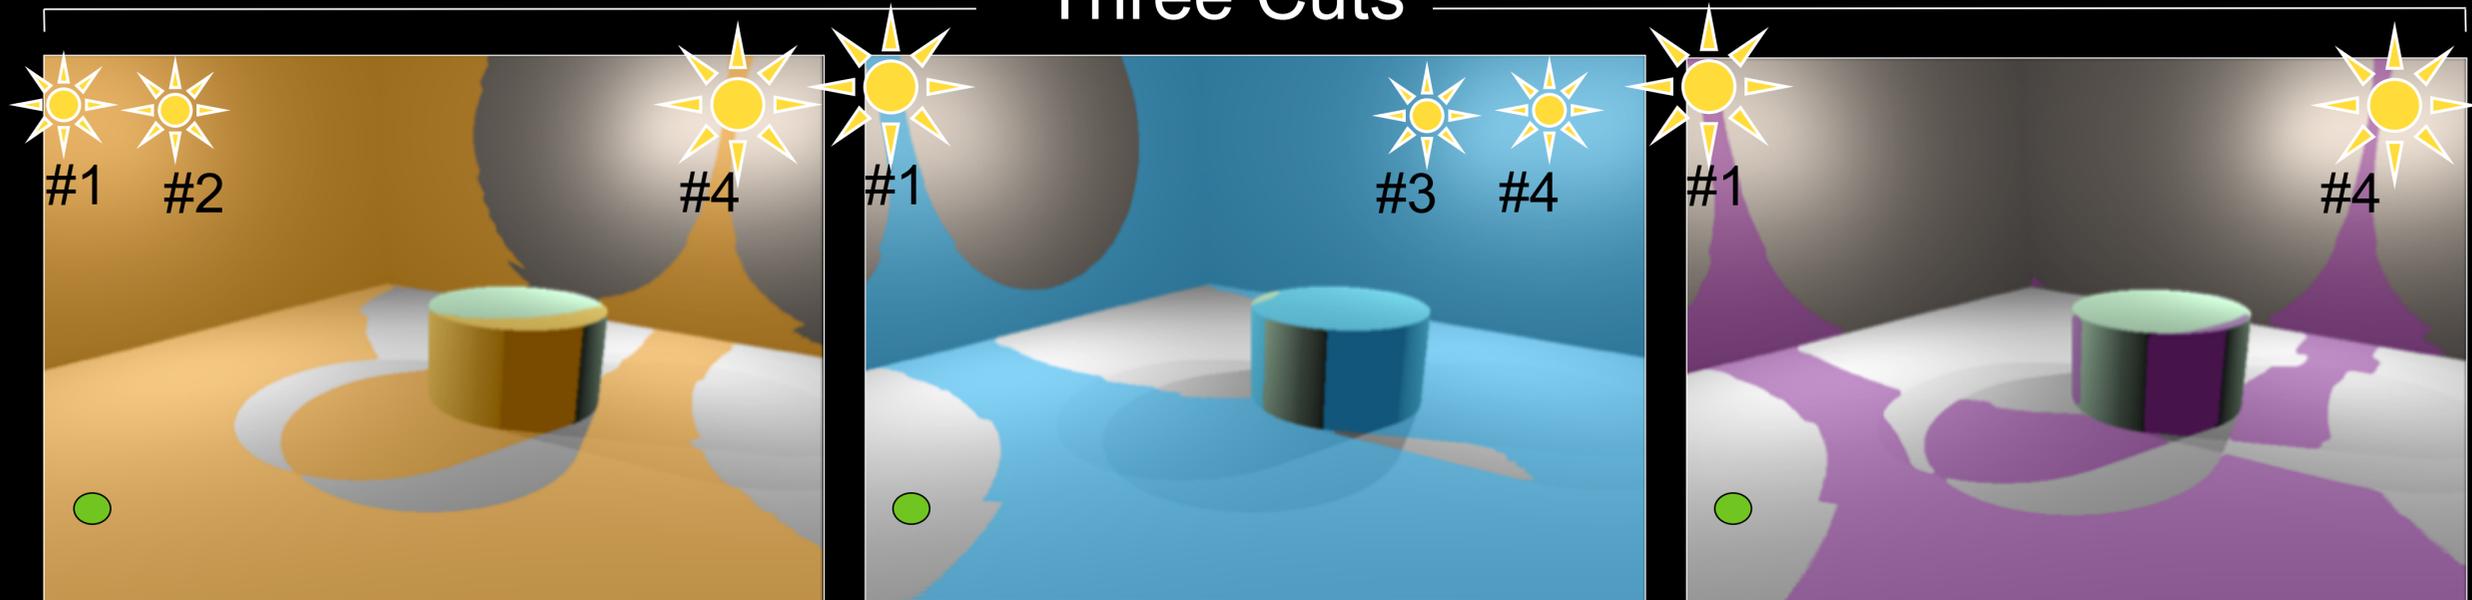
# Three Example Cuts

Three Cuts

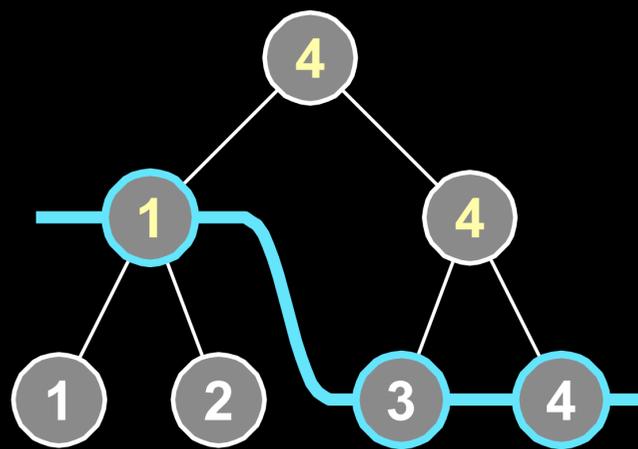


# Three Example Cuts

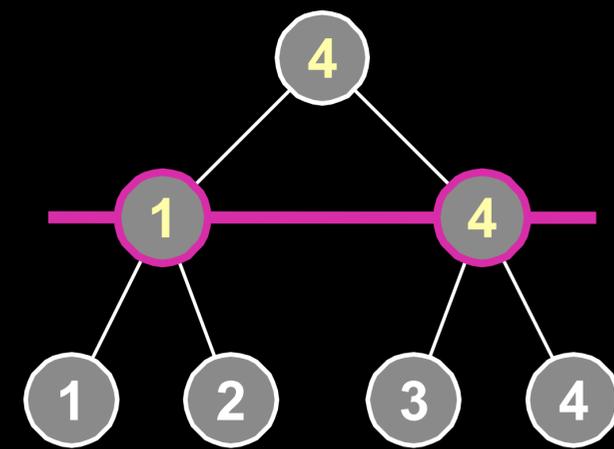
## Three Cuts



**Good**



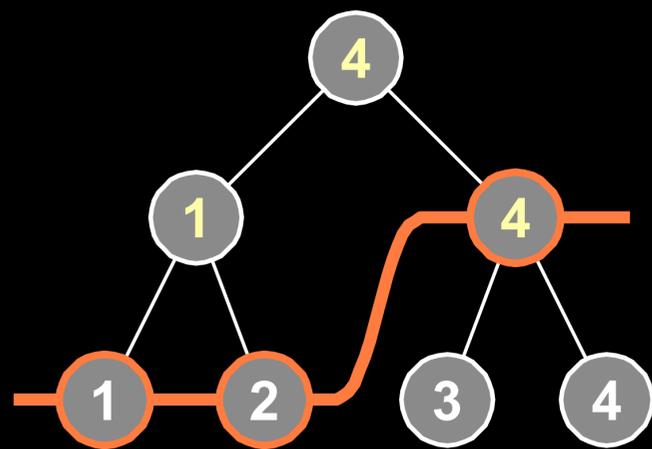
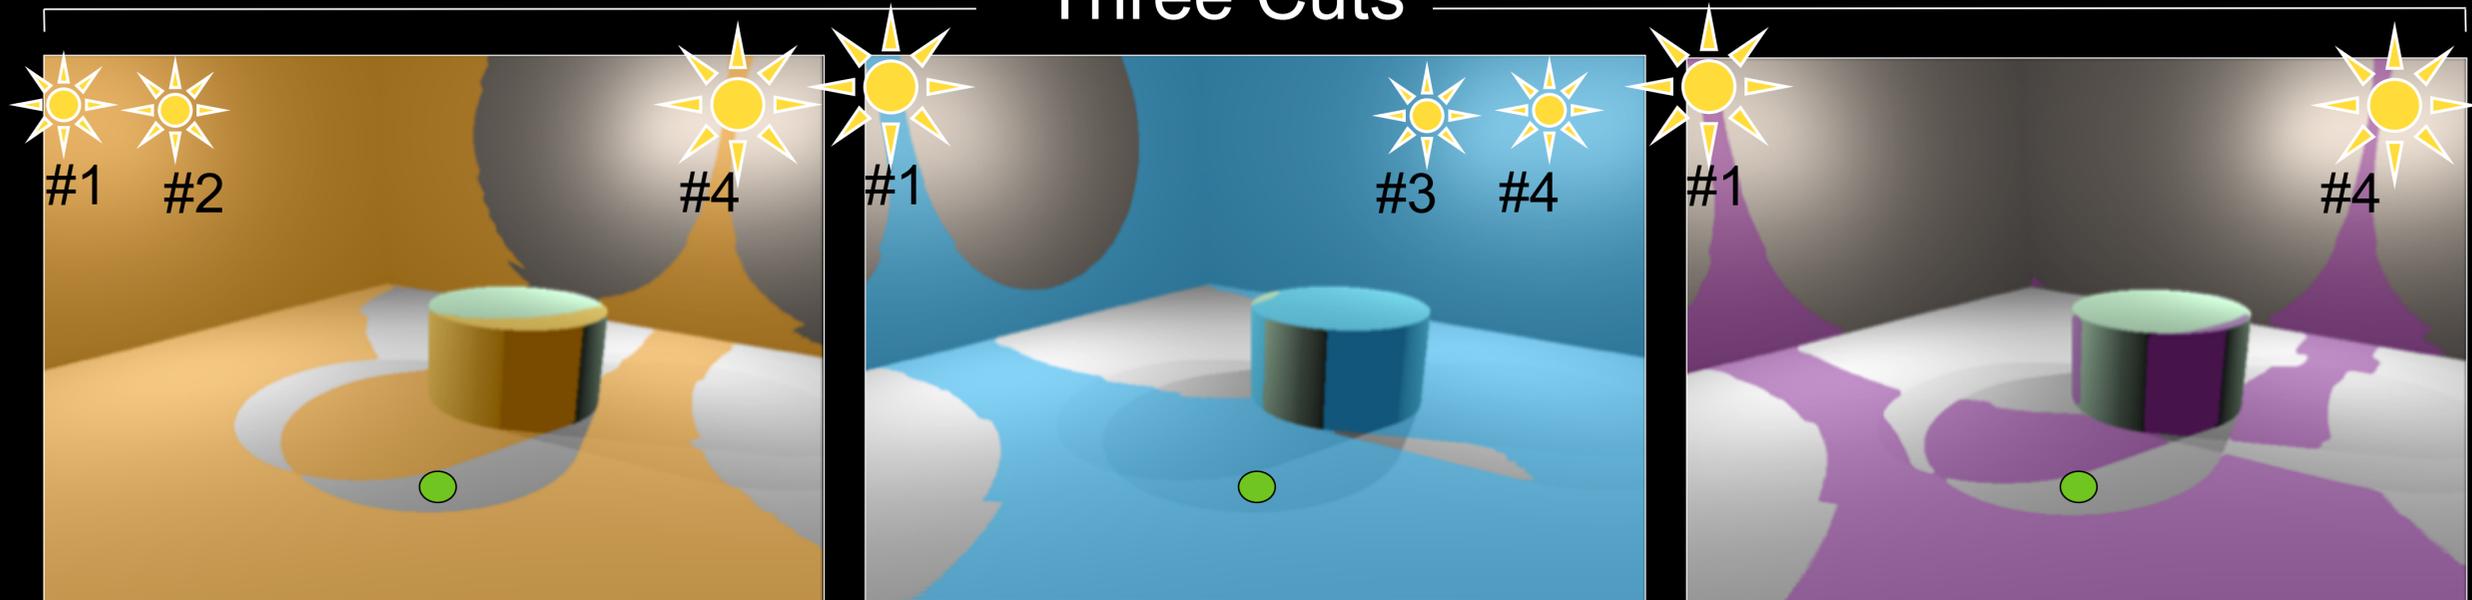
**Bad**



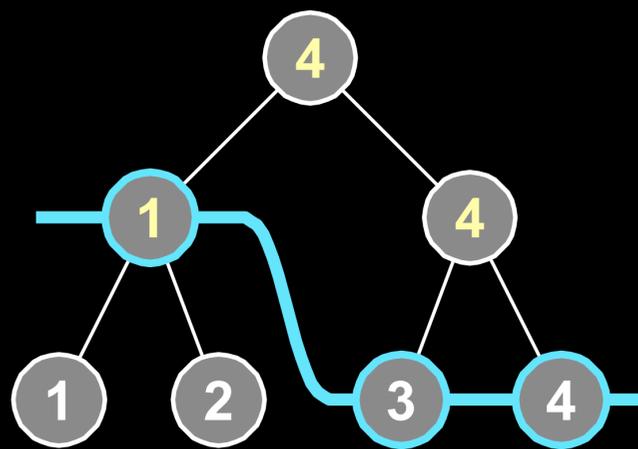
**Bad**

# Three Example Cuts

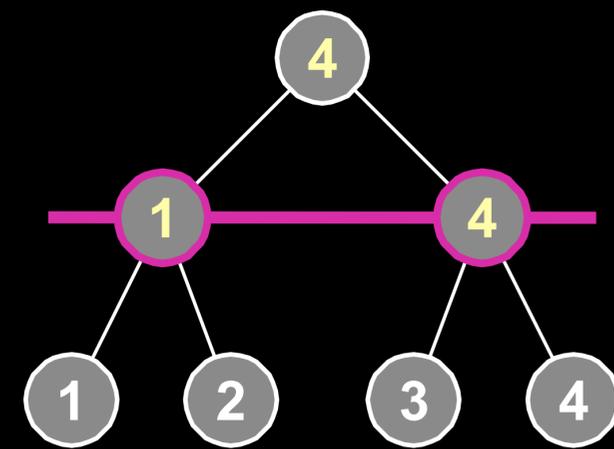
## Three Cuts



**Bad**



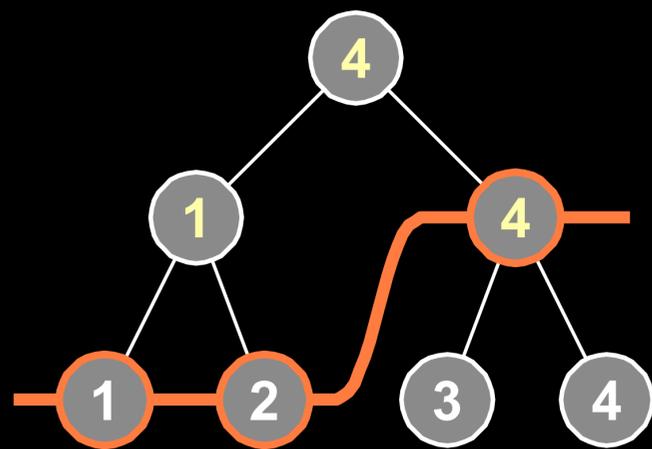
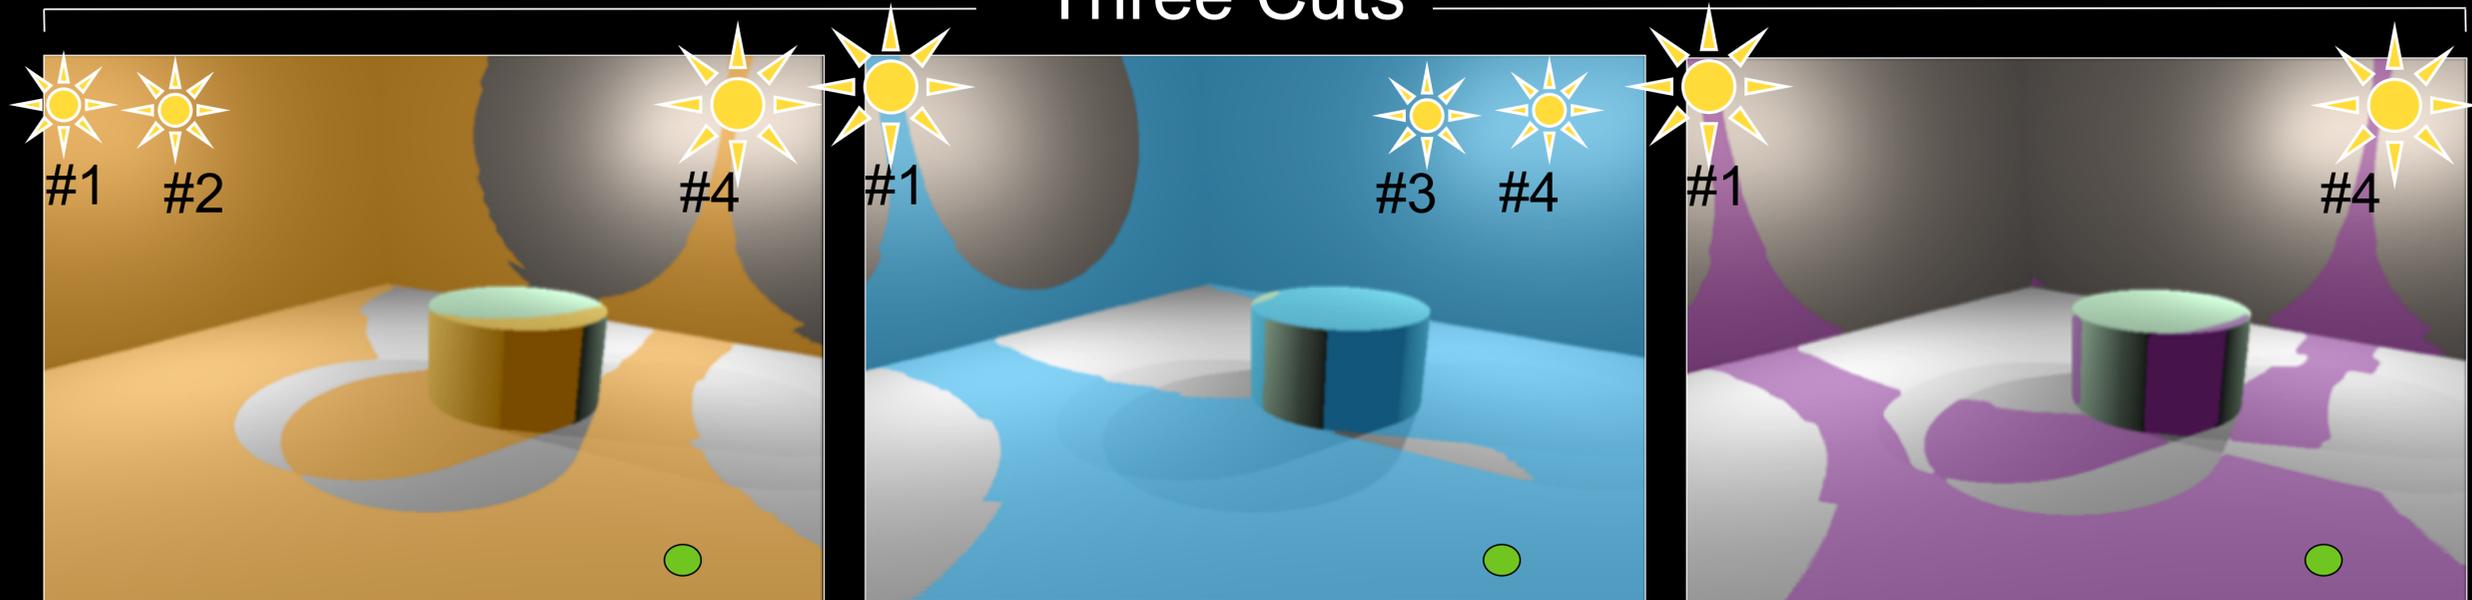
**Good**



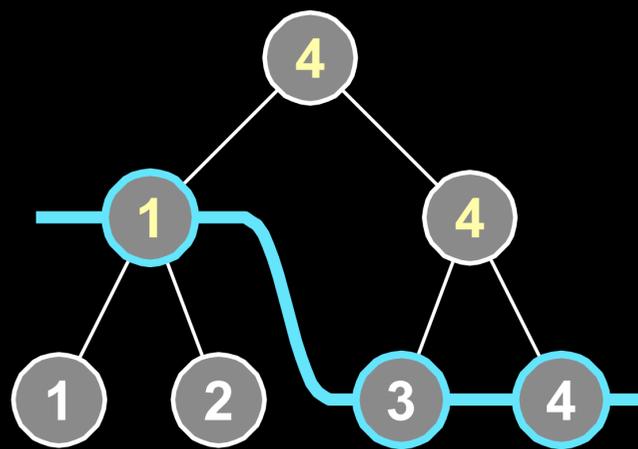
**Bad**

# Three Example Cuts

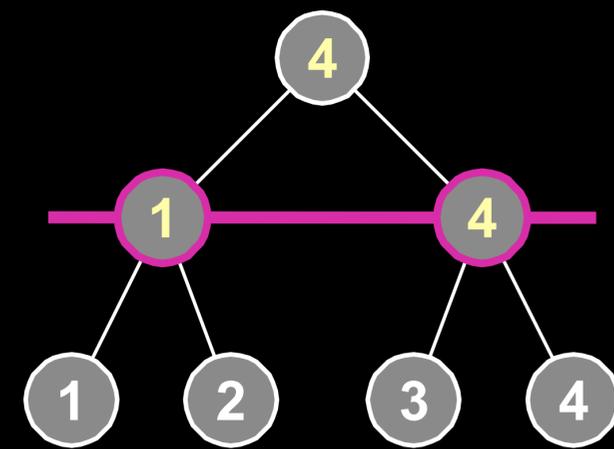
## Three Cuts



Good



Good



Good

Monday, August 6, 2012  
 For this point on the right side of the images, all three cuts usable. In this case the purple cut is the best choice because it will be the cheapest to compute as it contains the fewest nodes.

# Algorithm Overview

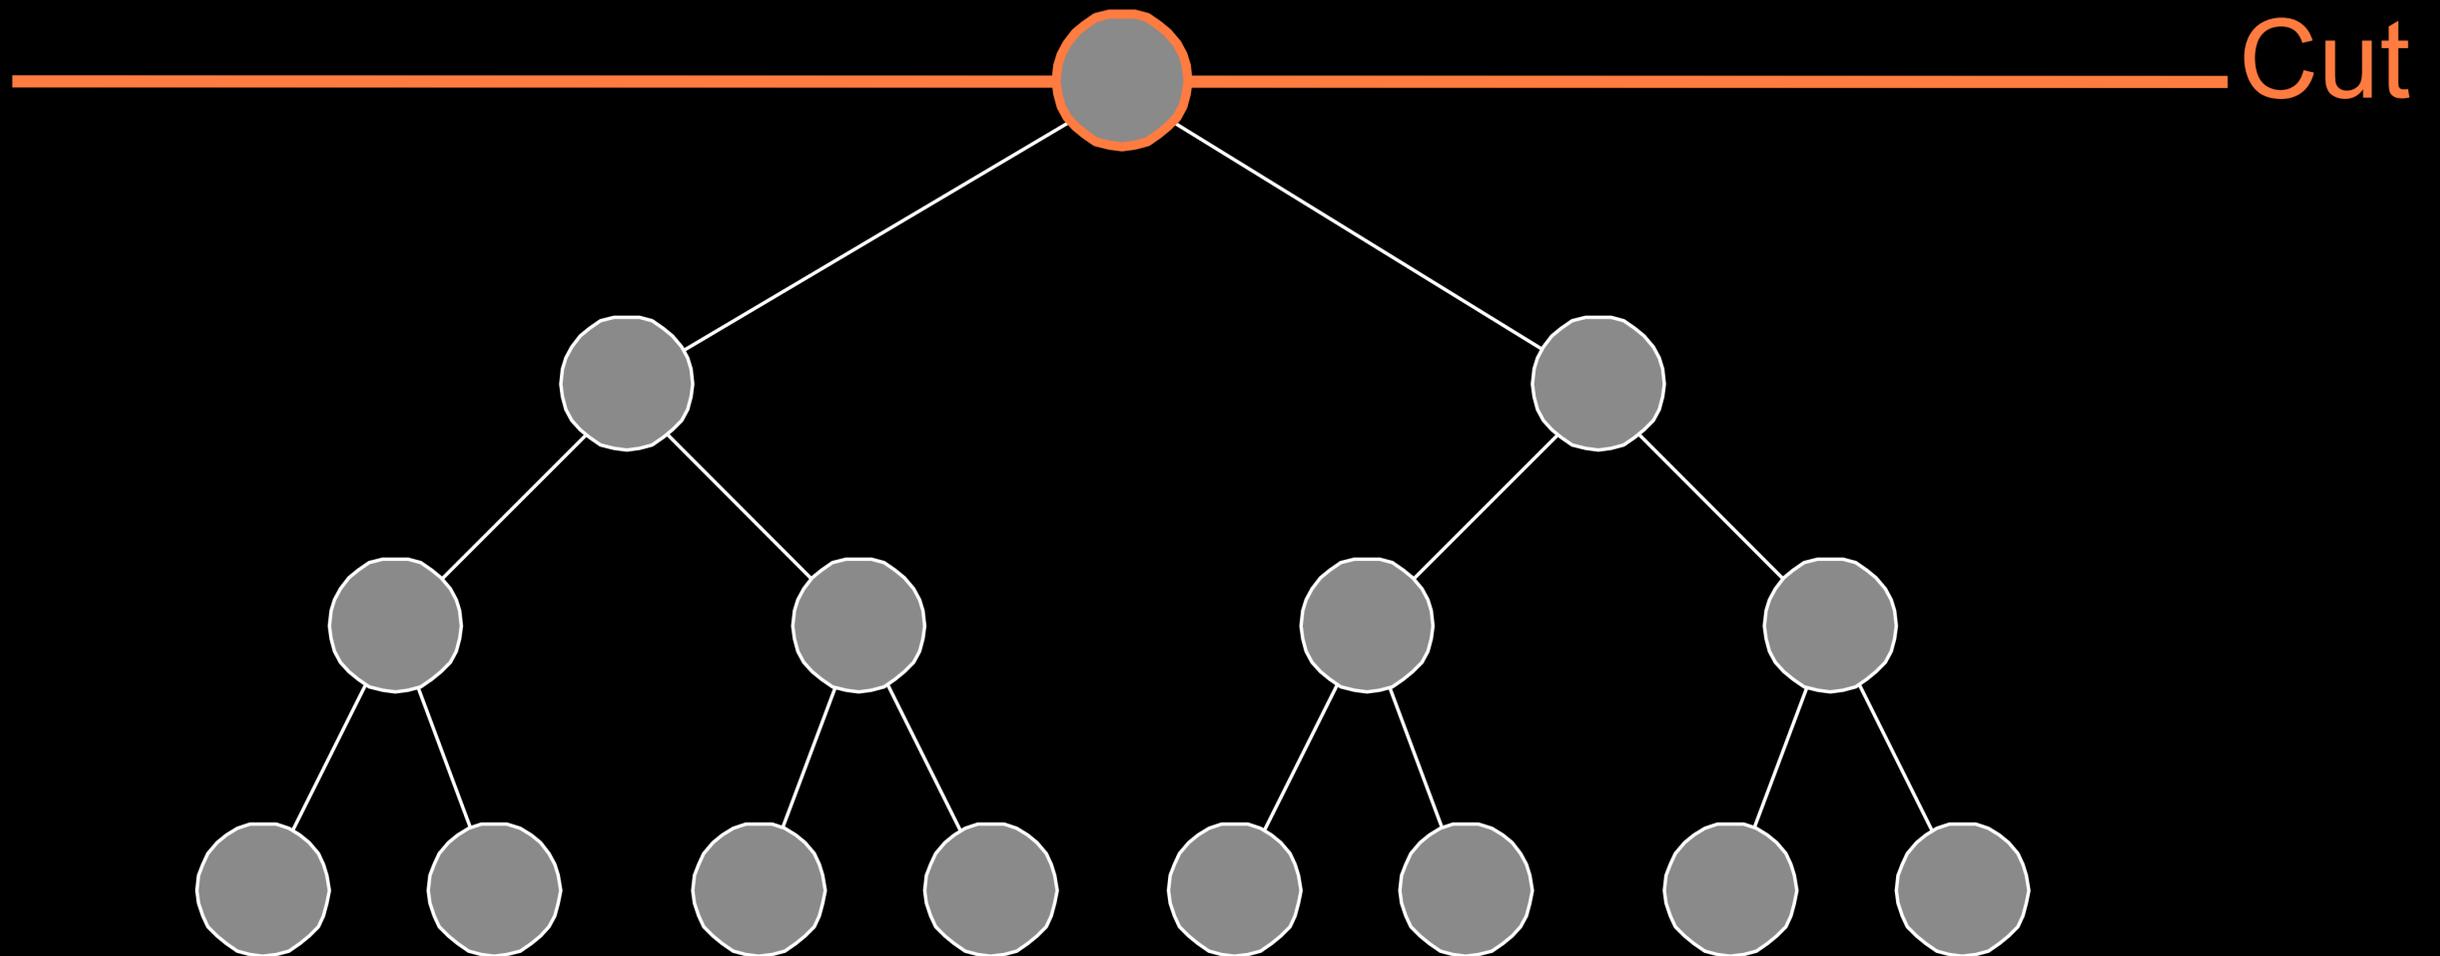
- Pre-process
  - Convert illumination to point lights
  - Build light tree
- For each visible point
  - Choose a cut to approximate the local illumination
    - ▶ Bound maximum error of cluster approximation
    - ▶ Refine cluster if error bound is too large

# Perceptual Metric

- Weber's Law
  - Contrast visibility threshold is fixed percentage of signal
  - Used 2% in our results
- Ensure each cluster's error  $<$  visibility threshold
  - Transitions will not be visible
  - Used to select cut

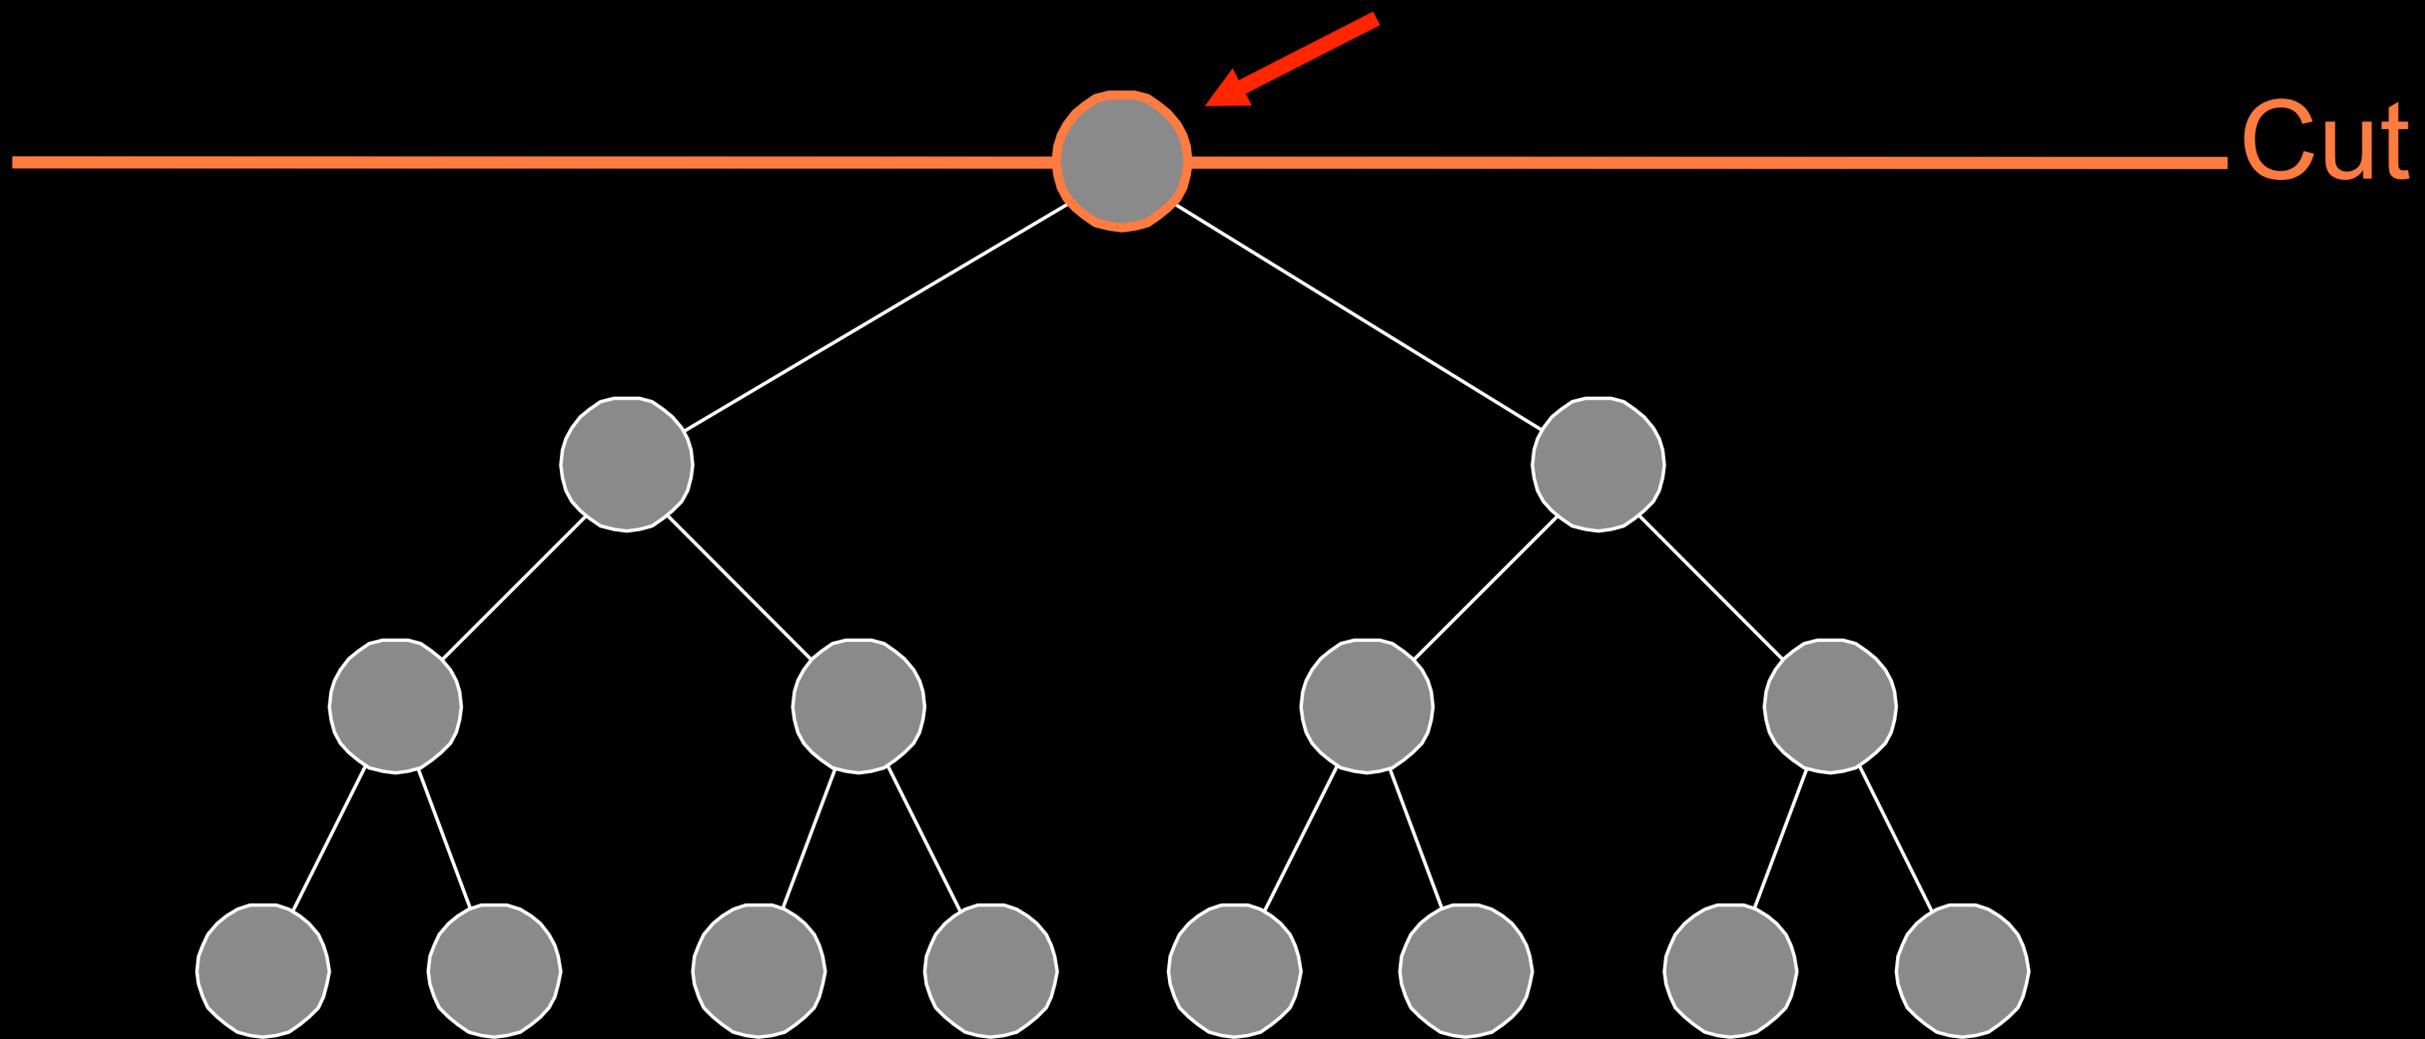
# Cut Selection Algorithm

- Start with coarse cut (eg, root node)



# Cut Selection Algorithm

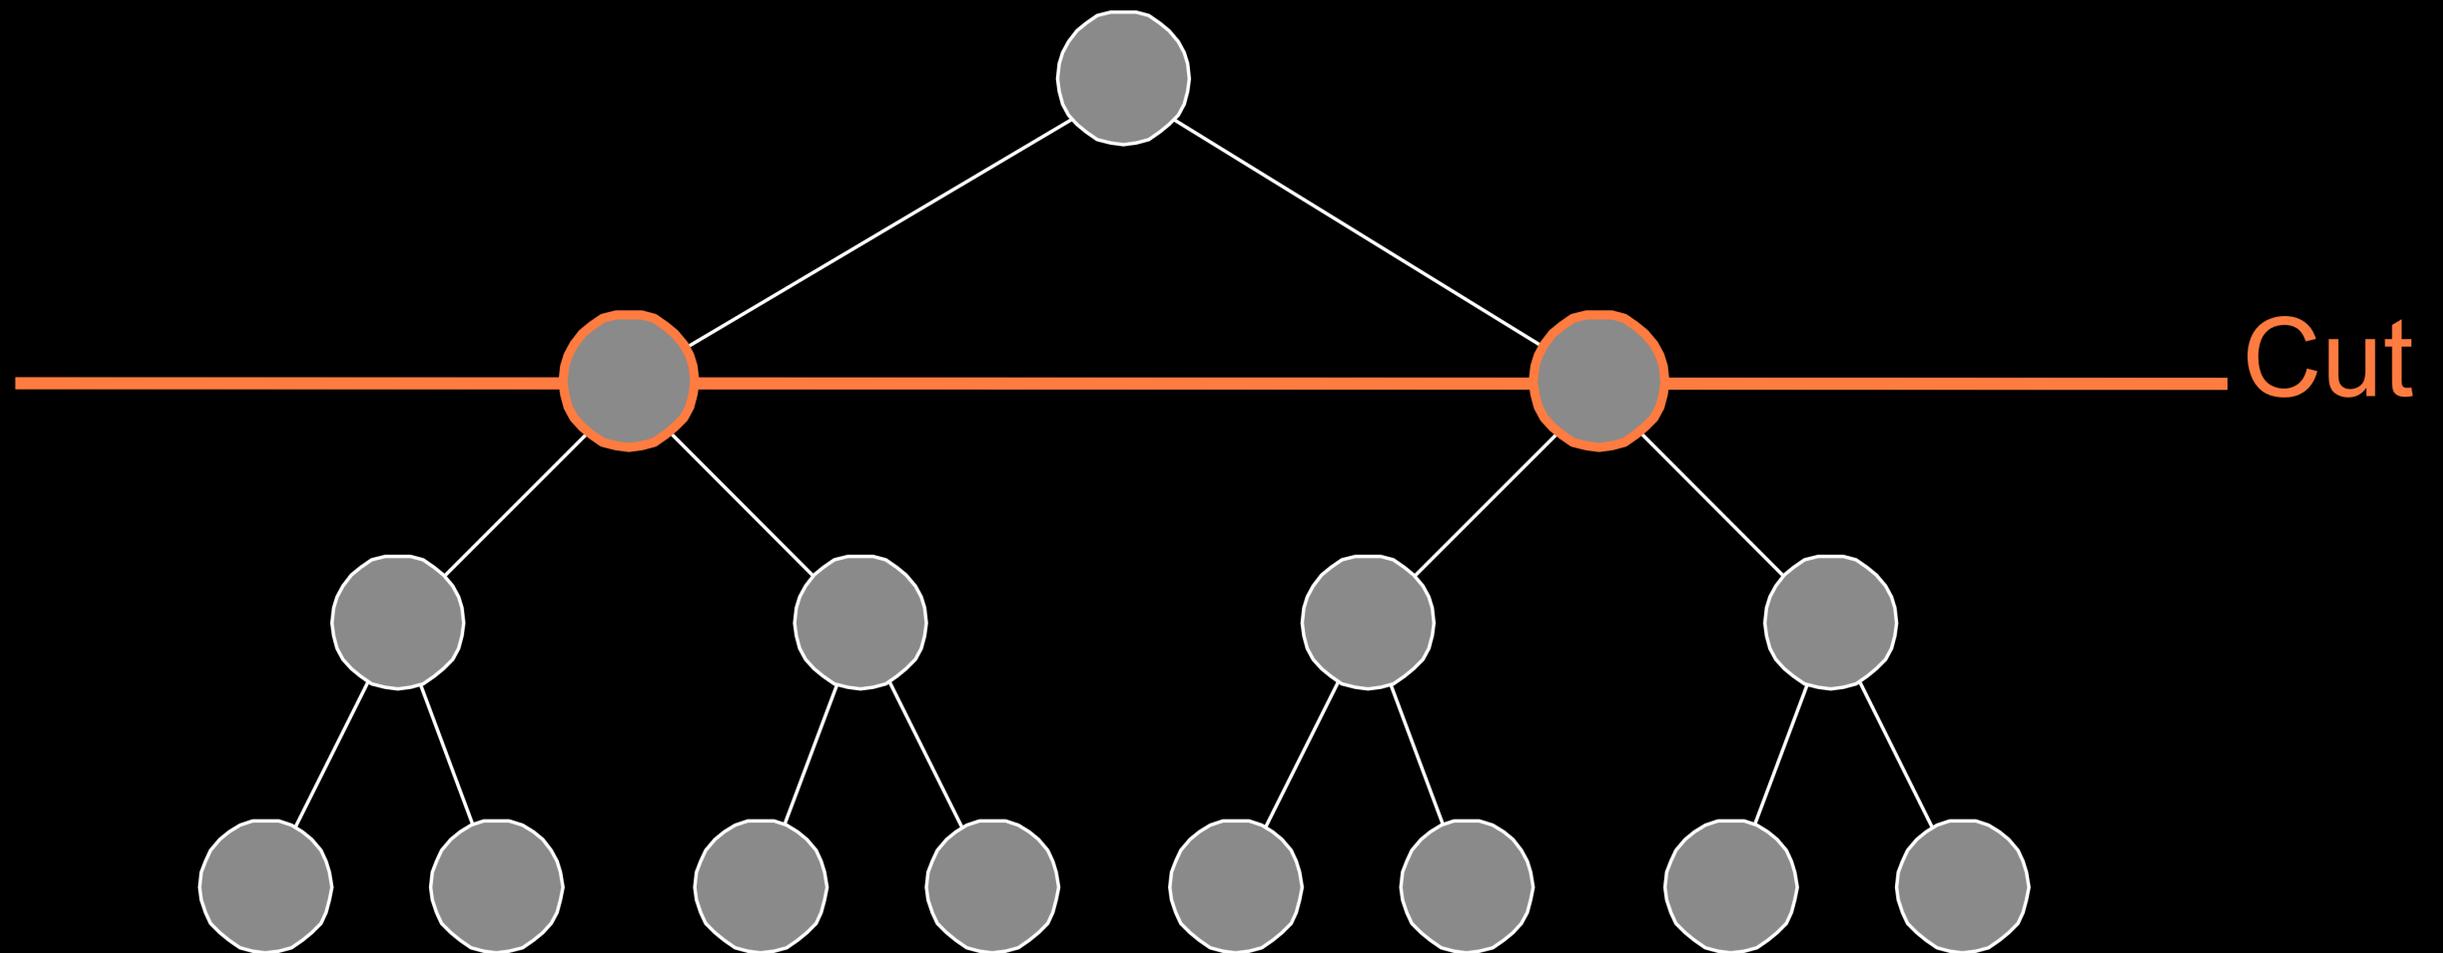
- Select cluster with largest error bound



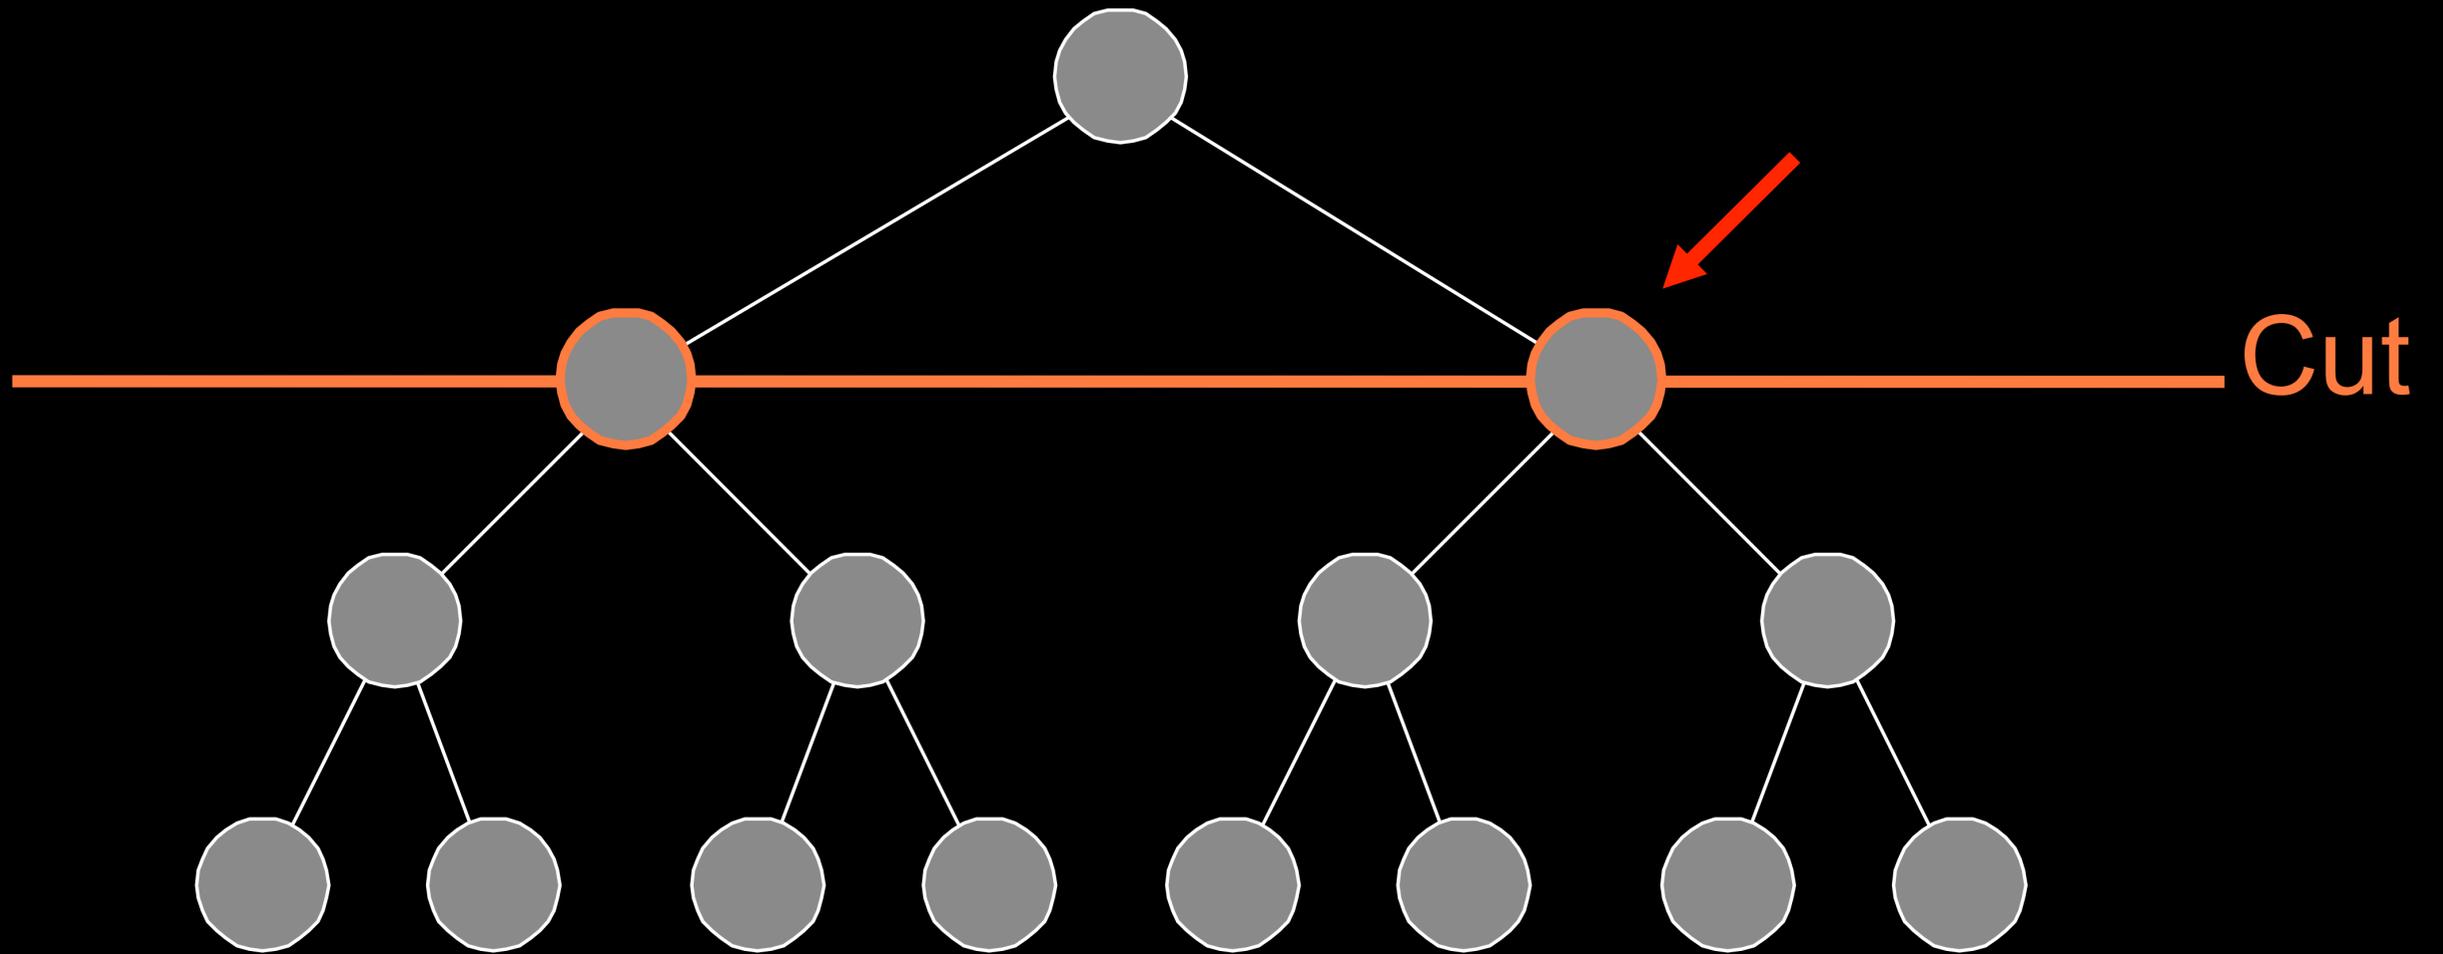
Monday, August 6, 2012  
 Then we select the node with the largest error bound. (the root node in this case).

# Cut Selection Algorithm

- Refine if error bound  $> 2\%$  of total

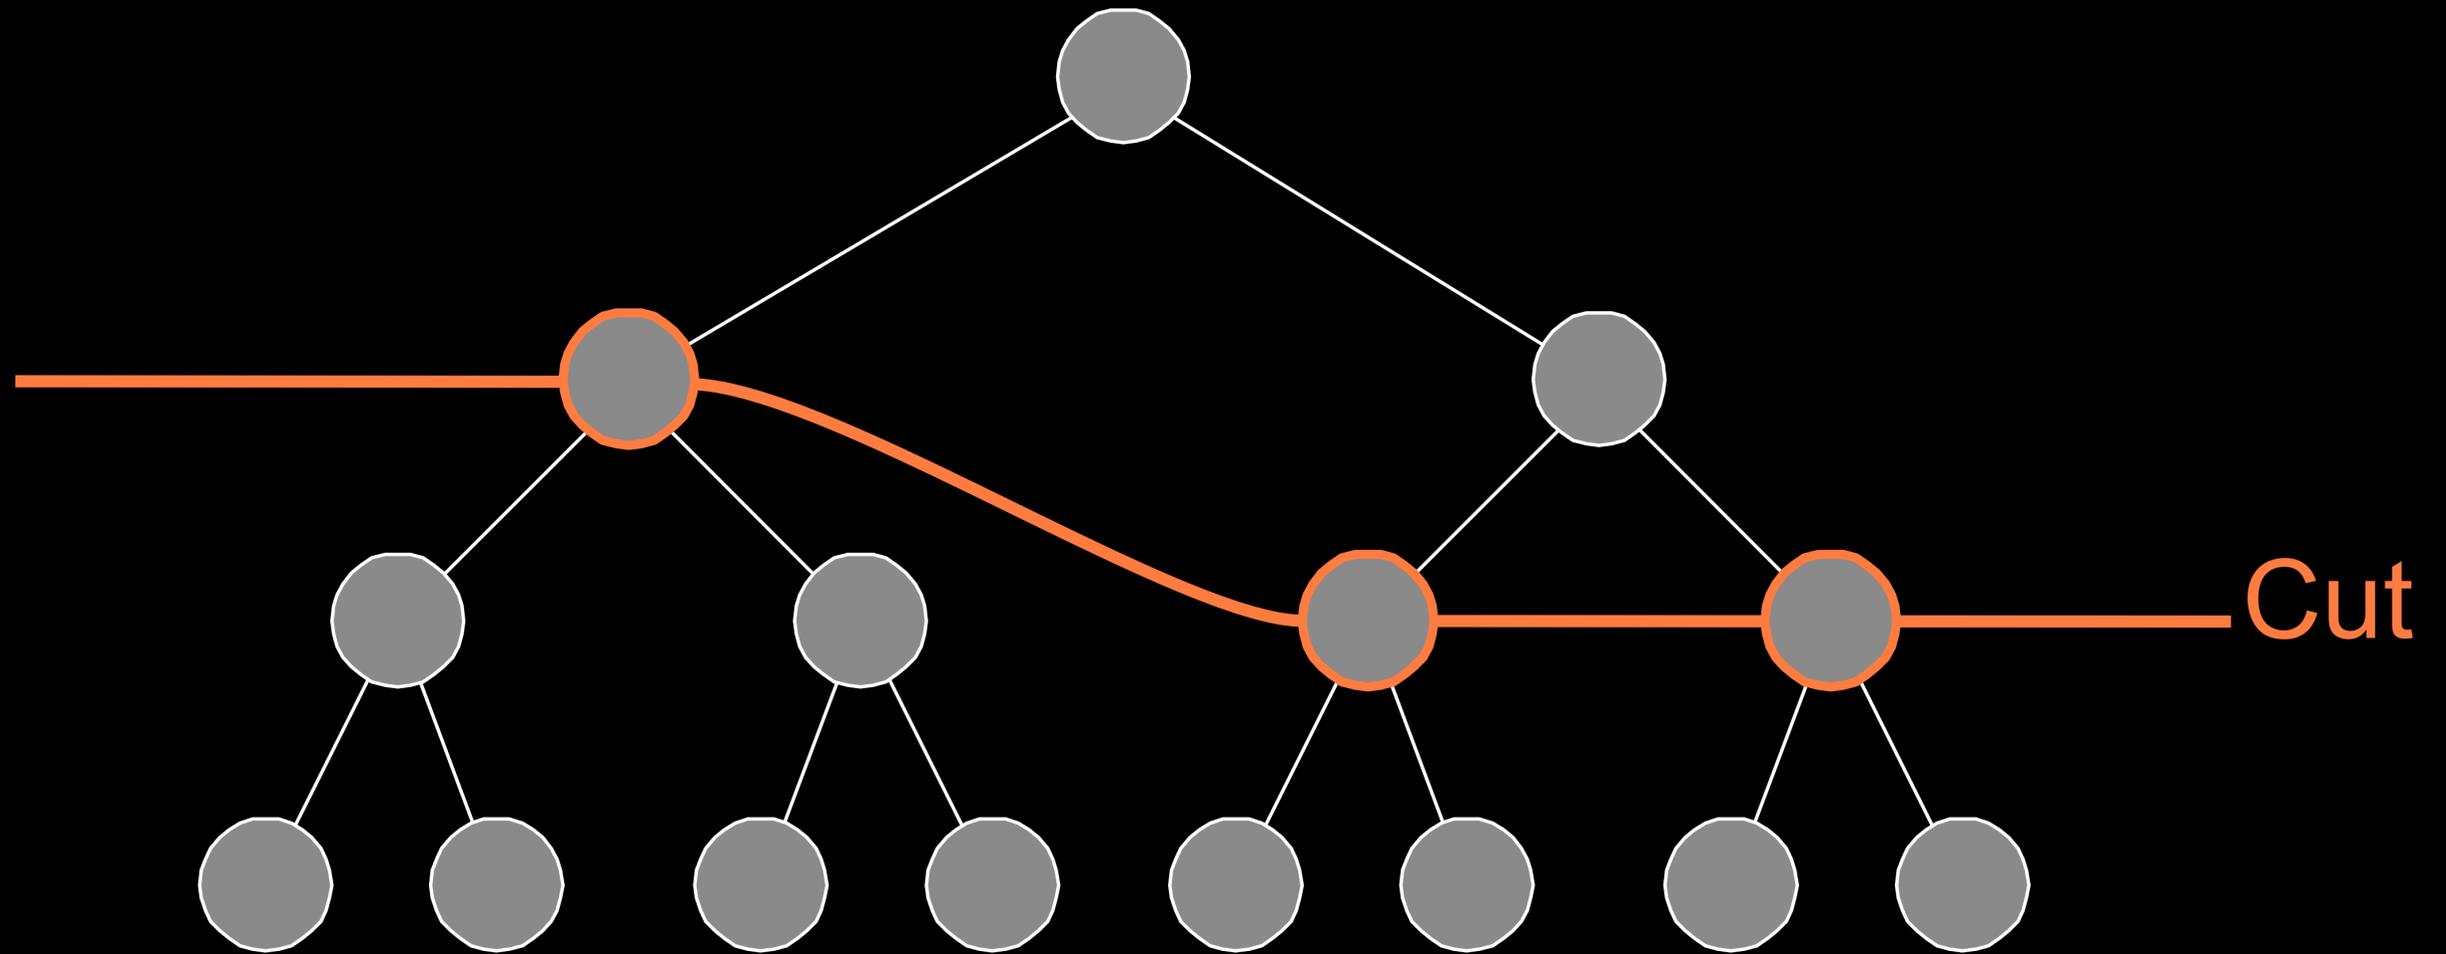


# Cut Selection Algorithm

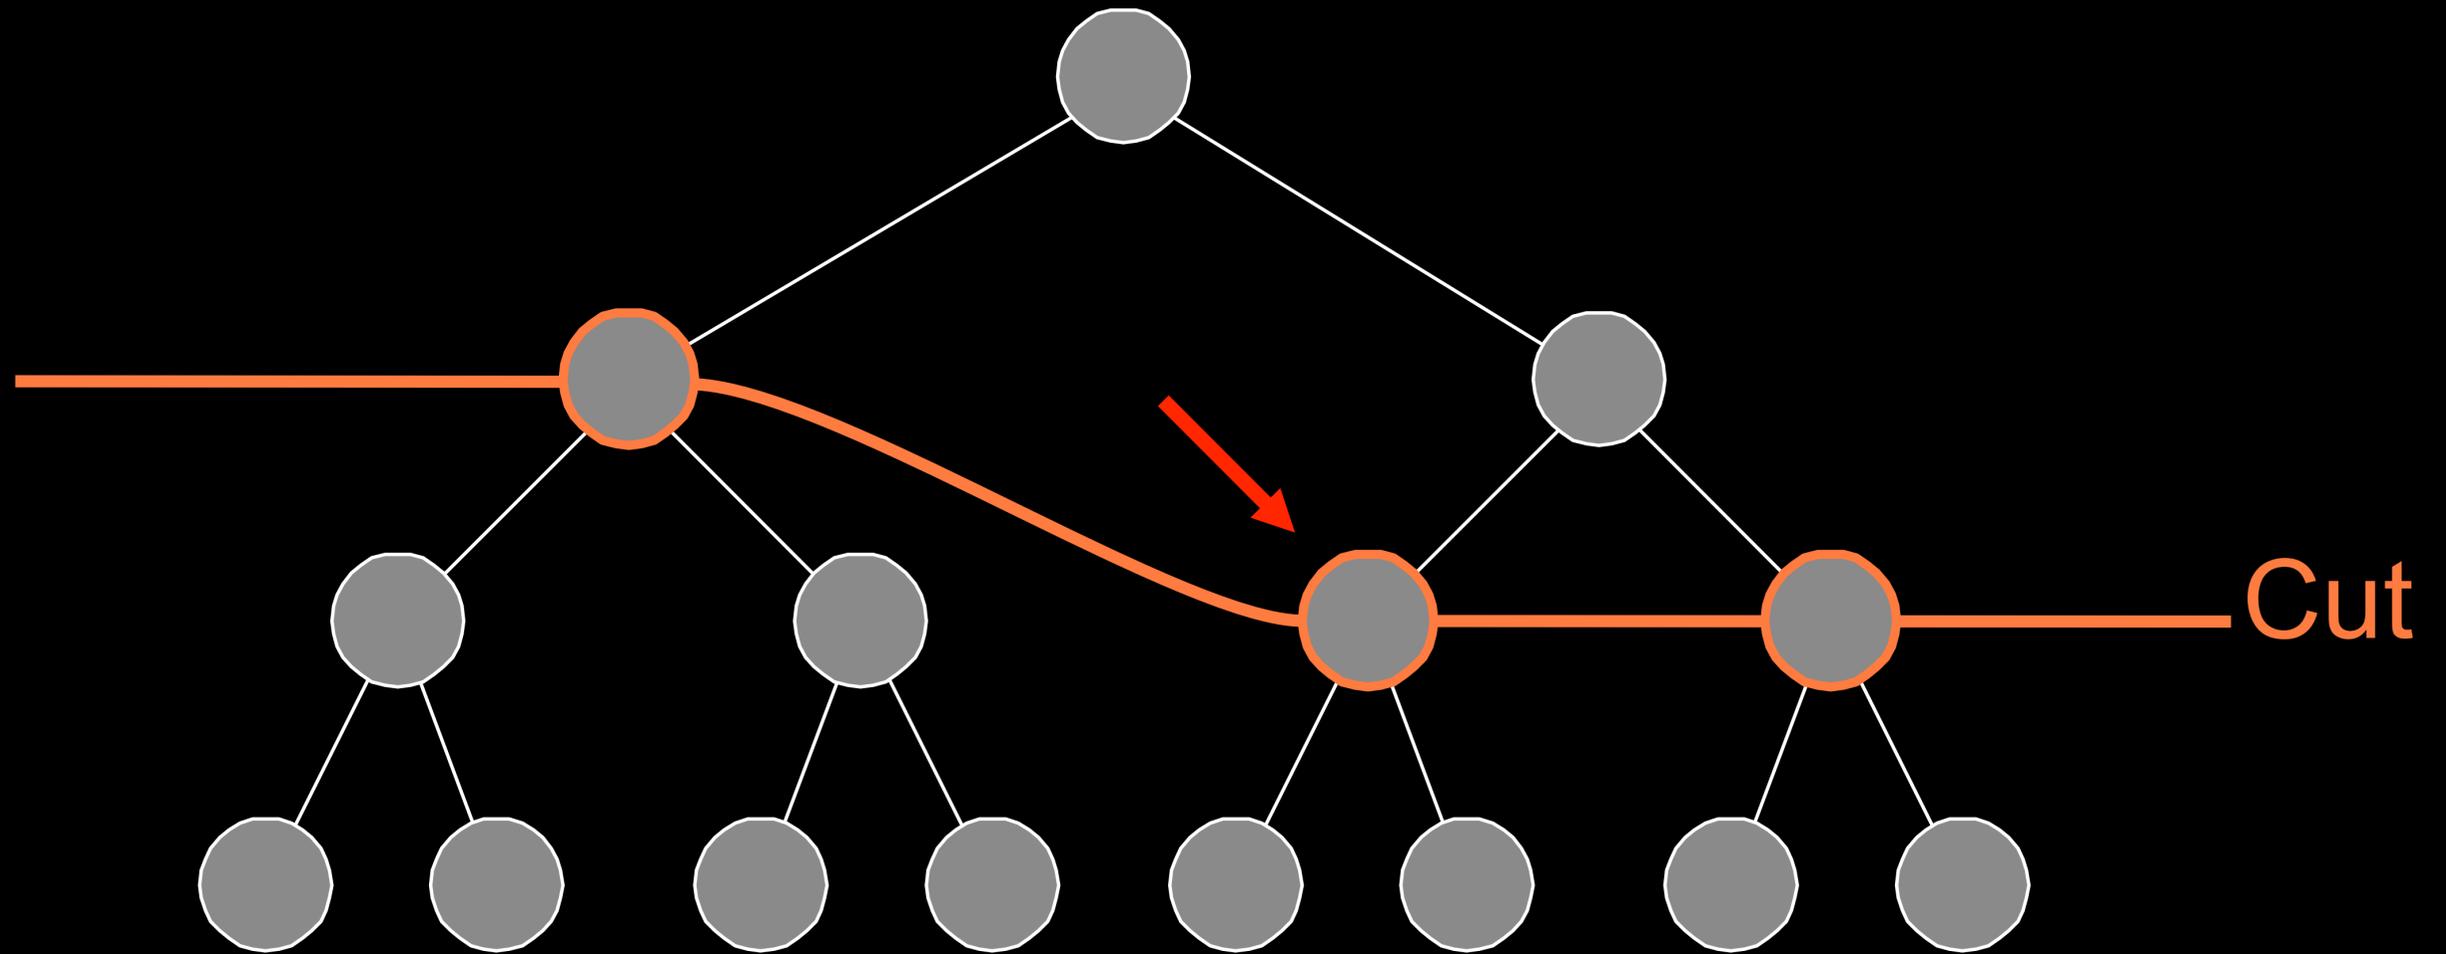


Monday, August 6, 2012  
Then we again select the cut node with the largest error bound

# Cut Selection Algorithm

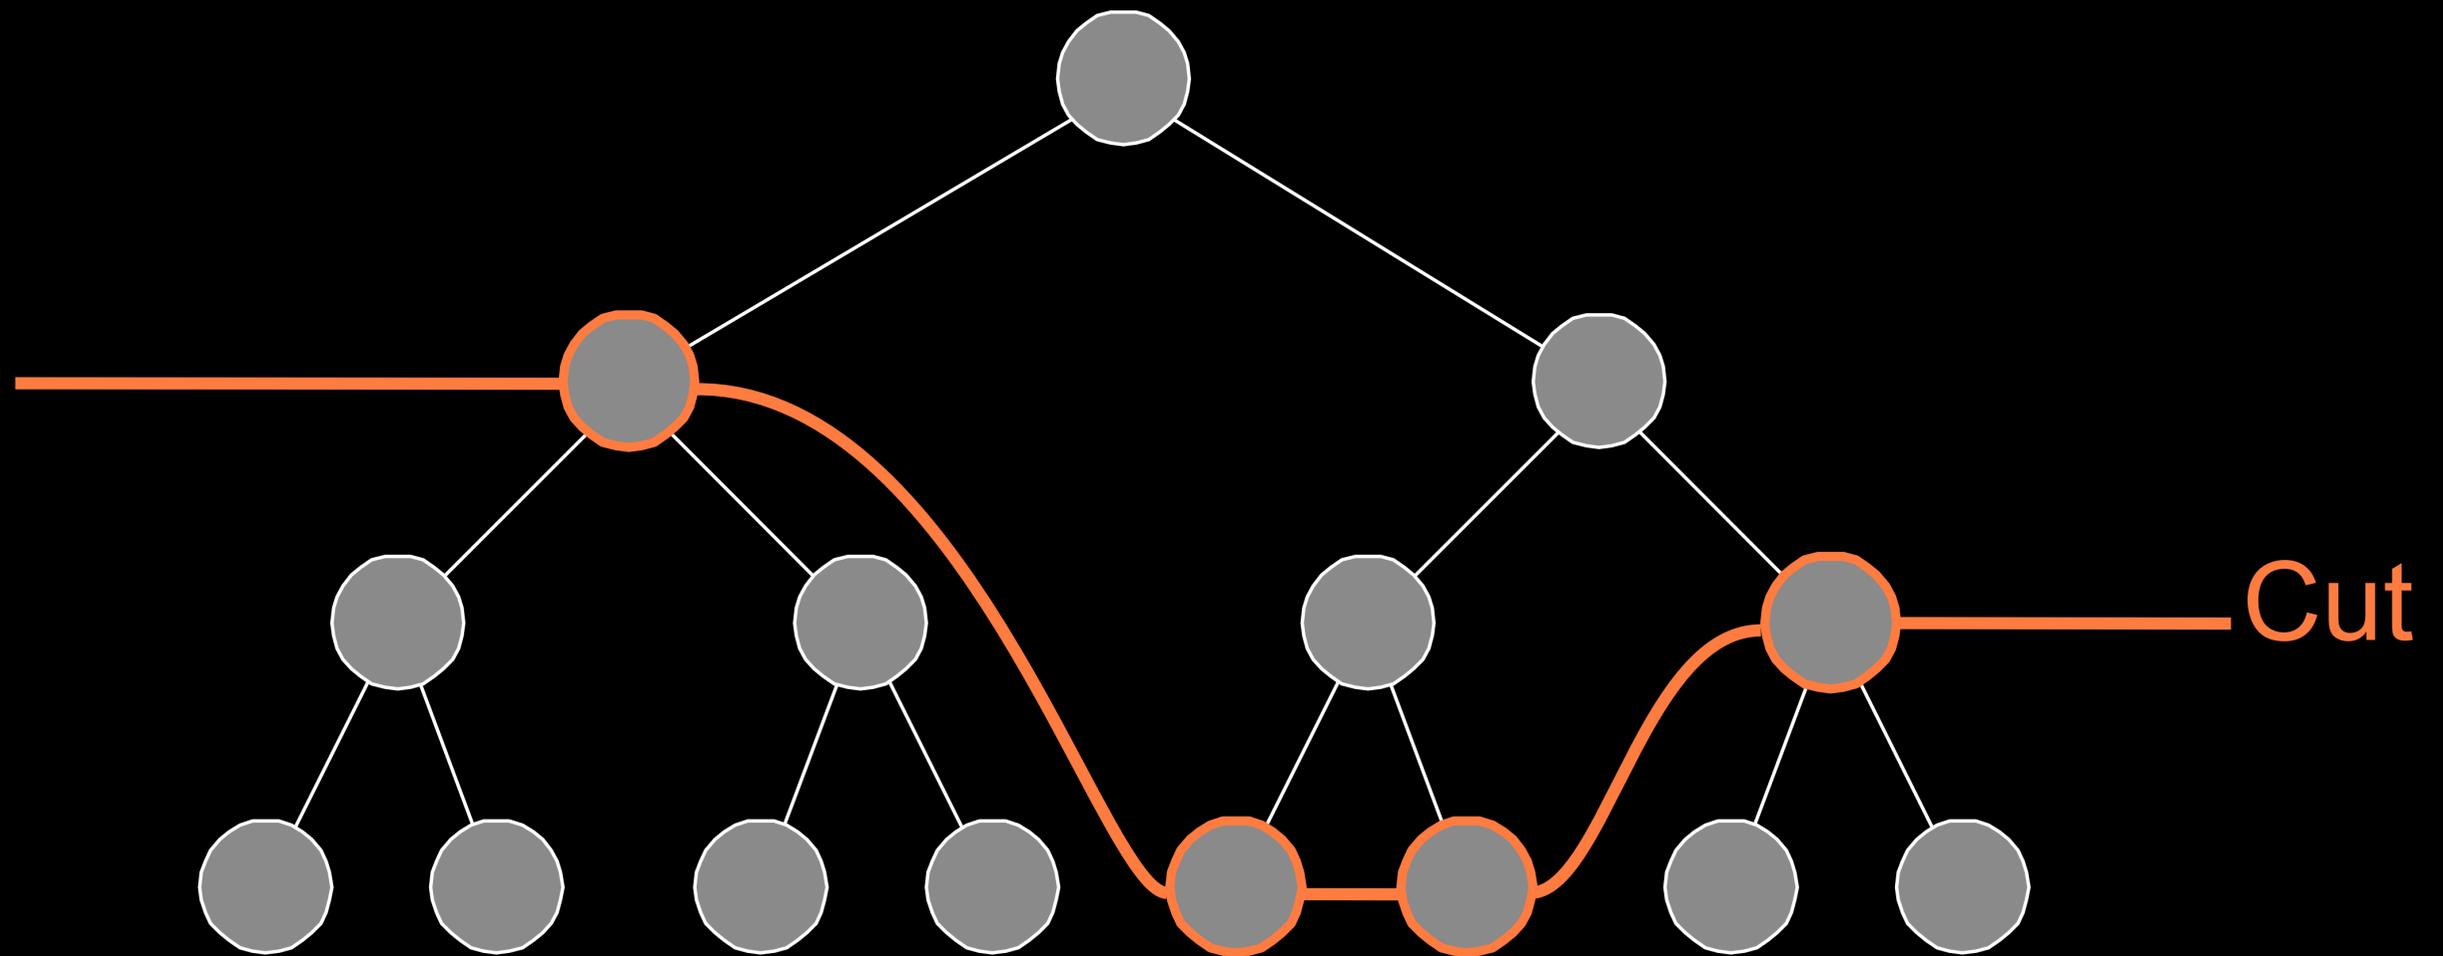


# Cut Selection Algorithm



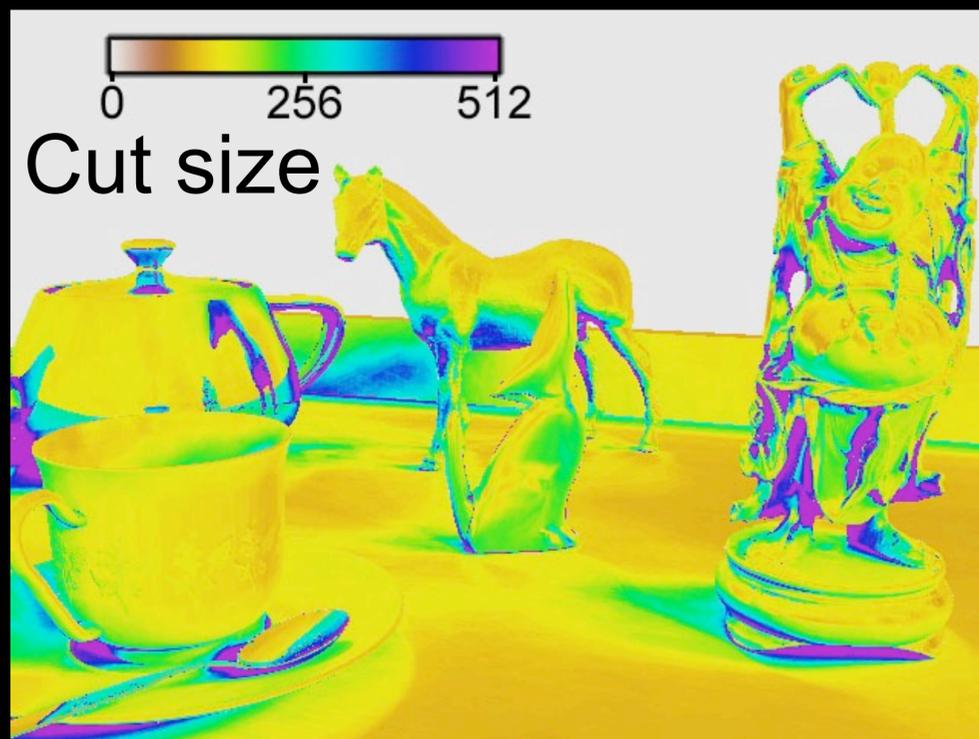
# Cut Selection Algorithm

- Repeat until cut obeys 2% threshold





Tableau, 630K polygons, 13 000 lights, (EnvMap+Indirect)



Tableau, 630K polygons, 13 000 lights, (EnvMap+Indirect)

Monday, August 6, 2012  
A false color image shows how the cutsize varies over the image. Note that it is much smaller than the 13000 lights. We also show an exaggerated error image which shows places where transitions occur and the problems they would cause if we allowed them to large enough to be visible.

# Lightcuts Recap

---

- Unified illumination handling
- Scalable solution for many lights
  - Locally adaptive representation (the cut)
- Analytic cluster error bounds
  - Most important lights always sampled
- Perceptual visibility metric

# Talk Overview

- Scalable solutions for many light rendering, Part I
  - Illumination at a single receiver: lightcuts
    - » “Lightcuts: a Scalable Approach to Illumination” by Walter, Fernandez, Arbree, Bala, Donikian, Greenberg, SIGGRAPH2005
  - Illumination over a pixel: Multidimensional Lightcuts
    - » “Multidimensional Lightcuts” by Walter, Arbree, Bala, Greenberg, SIGGRAPH2006

# Pixels are complex regions

- Pixels span regions over multiple dimensions
  - Complex illumination
  - Anti-aliasing
  - Motion blur
  - Participating media
  - Depth of field



$$\text{Pixel} = \int_{\text{Time}} \int_{\text{Pixel Area}} \int_{\text{Lights}} L(\mathbf{x}, \omega) \dots$$

Monday, August 6, 2012

Let me first describe the problem we're trying to solve. In order to compute high quality realistic images we need to be able to simulate multiple complex expensive phenomena. For instance the image on the right includes complex illumination, both from a captured environment map and indirect illumination, plus anti-aliasing and, on the right side of the image, motion blur for the spinning roulette wheel. Thus for each pixel we need to integrate the contribution over multiple domains.

# Pixels are complex regions

- Pixels span regions over multiple dimensions
  - Complex illumination
  - Anti-aliasing
  - Motion blur
  - Participating media
  - Depth of field



$$\text{Pixel} = \int_{\text{Volume}} \int_{\text{Time}} \int_{\text{Pixel Area}} \int_{\text{Lights}} L(\mathbf{x}, \omega) \dots$$

# Pixels are complex regions

- Pixels span regions over multiple dimensions
  - Complex illumination
  - Anti-aliasing
  - Motion blur
  - Participating media
  - Depth of field



$$\text{Pixel} = \int_{\text{Aperture}} \int_{\text{Volume}} \int_{\text{Time}} \int_{\text{Pixel Area}} \int_{\text{Lights}} L(\mathbf{x}, \omega) \dots$$

# Pixels are complex regions

- Pixels span regions over multiple dimensions
  - Complex illumination
  - Anti-aliasing
  - Motion blur
  - Participating media
  - Depth of field
  - Complex materials



Glossy

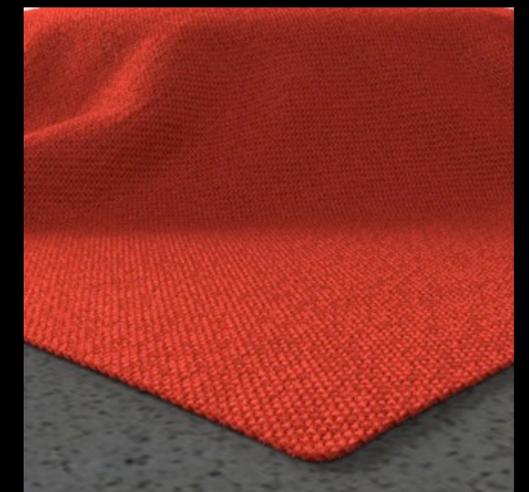


Subsurface



Volumetric

Before



After

["Bidirectional Lightcuts", Walter et al. 12]



# Bidirectional Lightcuts

*Walter Khungurn Bala*

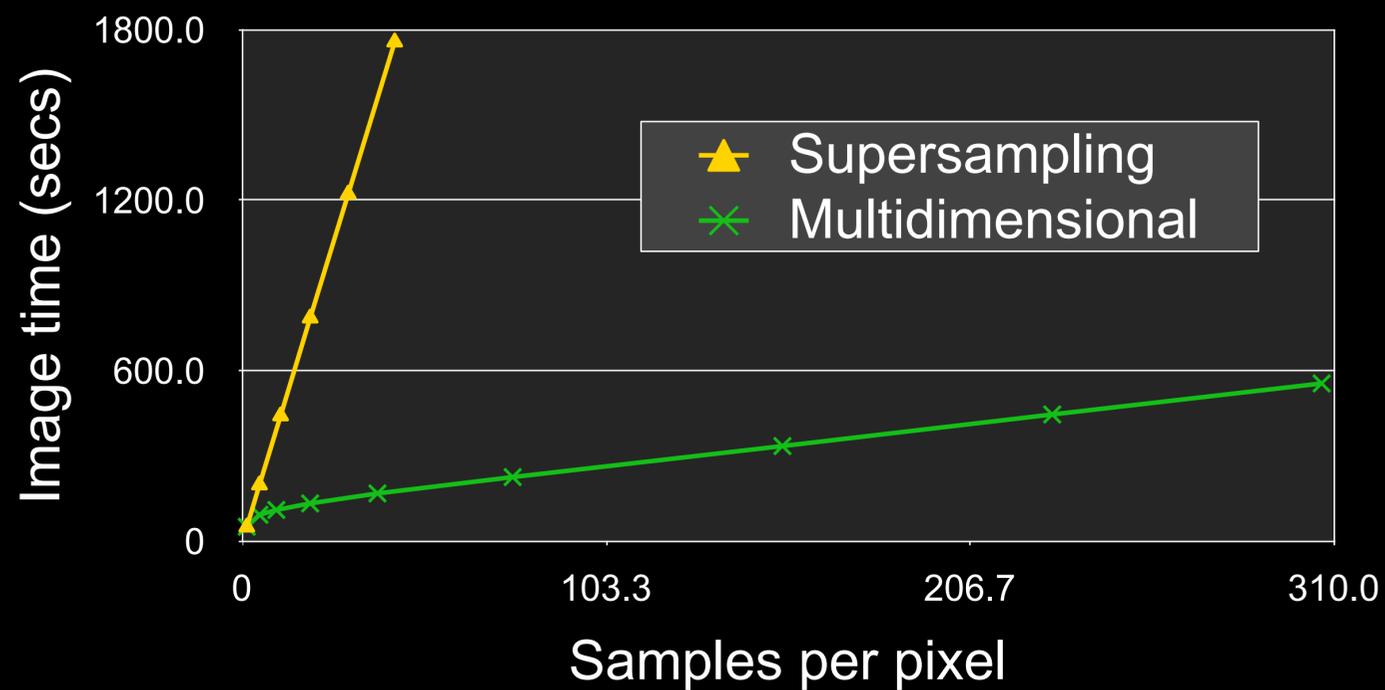
Tuesday, 2:00 - 3:30 PM  
Room 502

Monday, August 6, 2012

Come to the talk on Tuesday to hear more details.

# Insight

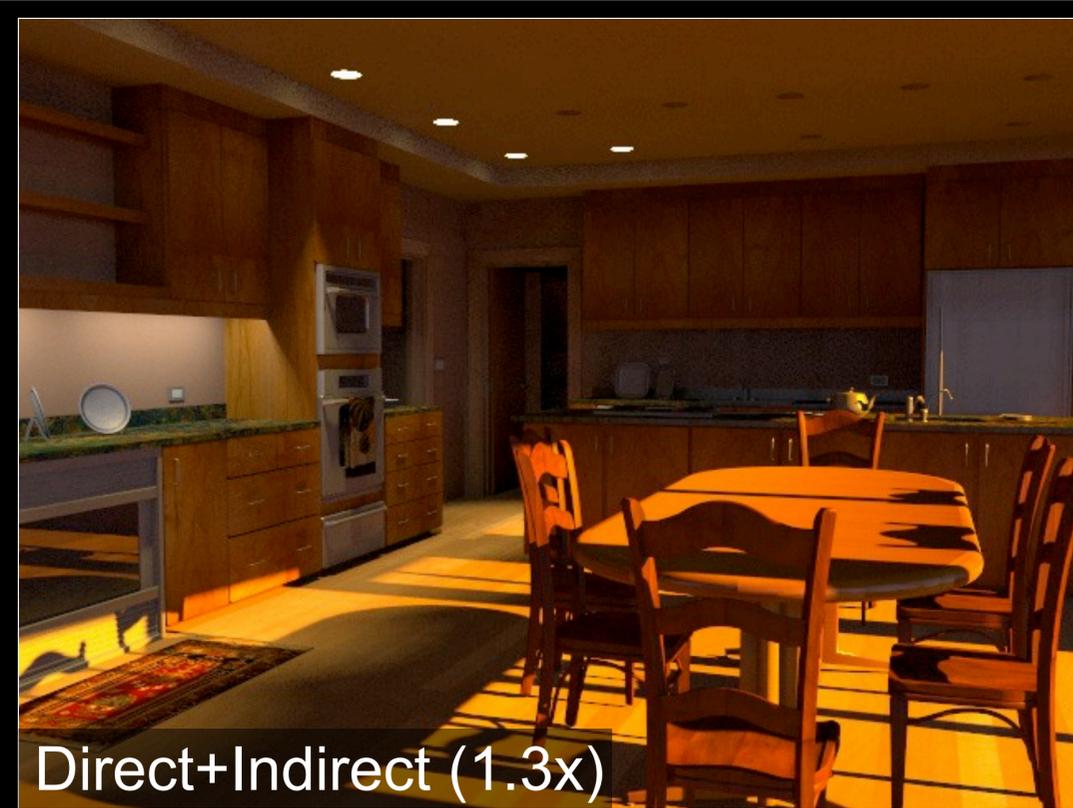
- Complex integrals over multiple dimensions
  - Requires many eye samples
- Holistic approach
  - Solve the complete pixel integral



Monday, August 6, 2012  
 What we want is compute total pixel values for our images which means multiple integrals. While many previous techniques have tried solving each integral separately but this does not result in a scalable solution. One of the key insights here is that solving the complete pixel integral is fundamental to achieving a scalable solution.



Direct only (relative cost 1x)



Direct+Indirect (1.3x)



Direct+Indirect+Volume (1.8x)



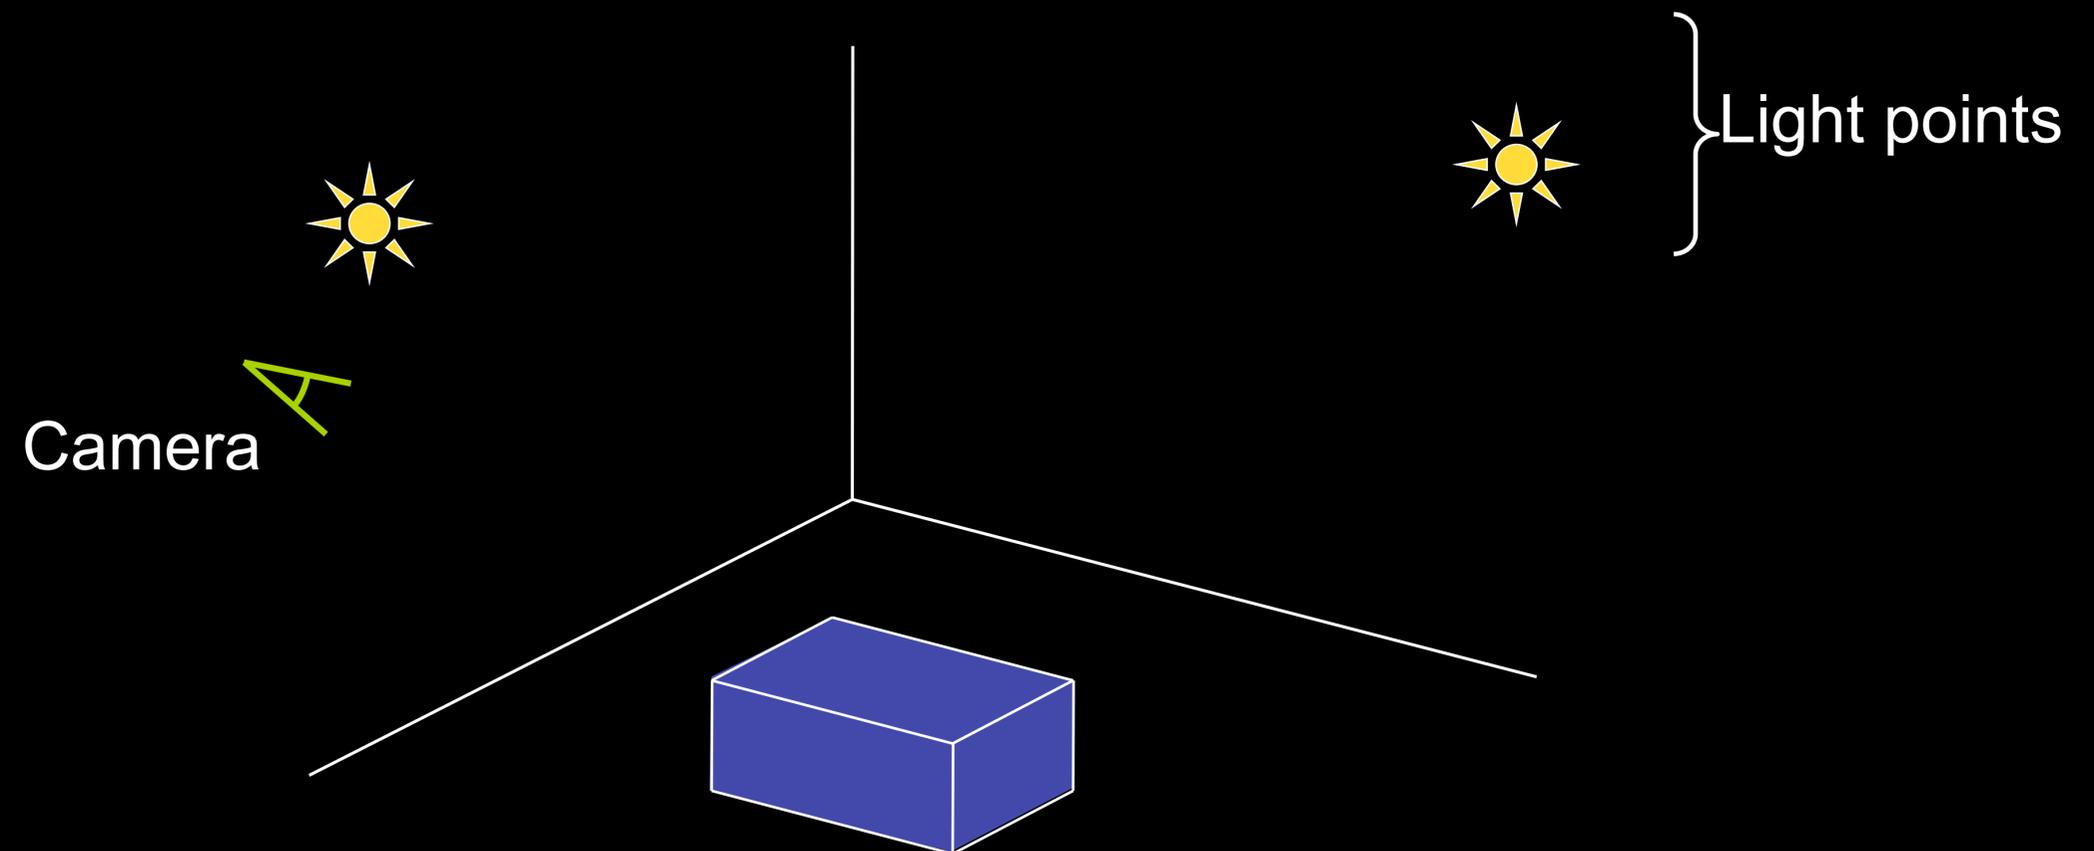
Direct+Indirect+Volume+Motion (2.2x)

Monday, August 6, 2012

This example shows another way in which the system is scalable. We take a single scene and progressively add more and more complex effects. Thus the image on the upper left is a direct only solution for the area lights and sun. The image on the upper right also include indirect illumination. The image on the lower left also added smoke or participating media to the scene and finally the image on the lower right further adds motion blur. Even though this image includes multiple complex effects its rendering time only increases by a factor of 2.2 compared to the direct only solution which is much better than in traditional approaches where adding each effect would greatly increase the image time.

# Point Sets

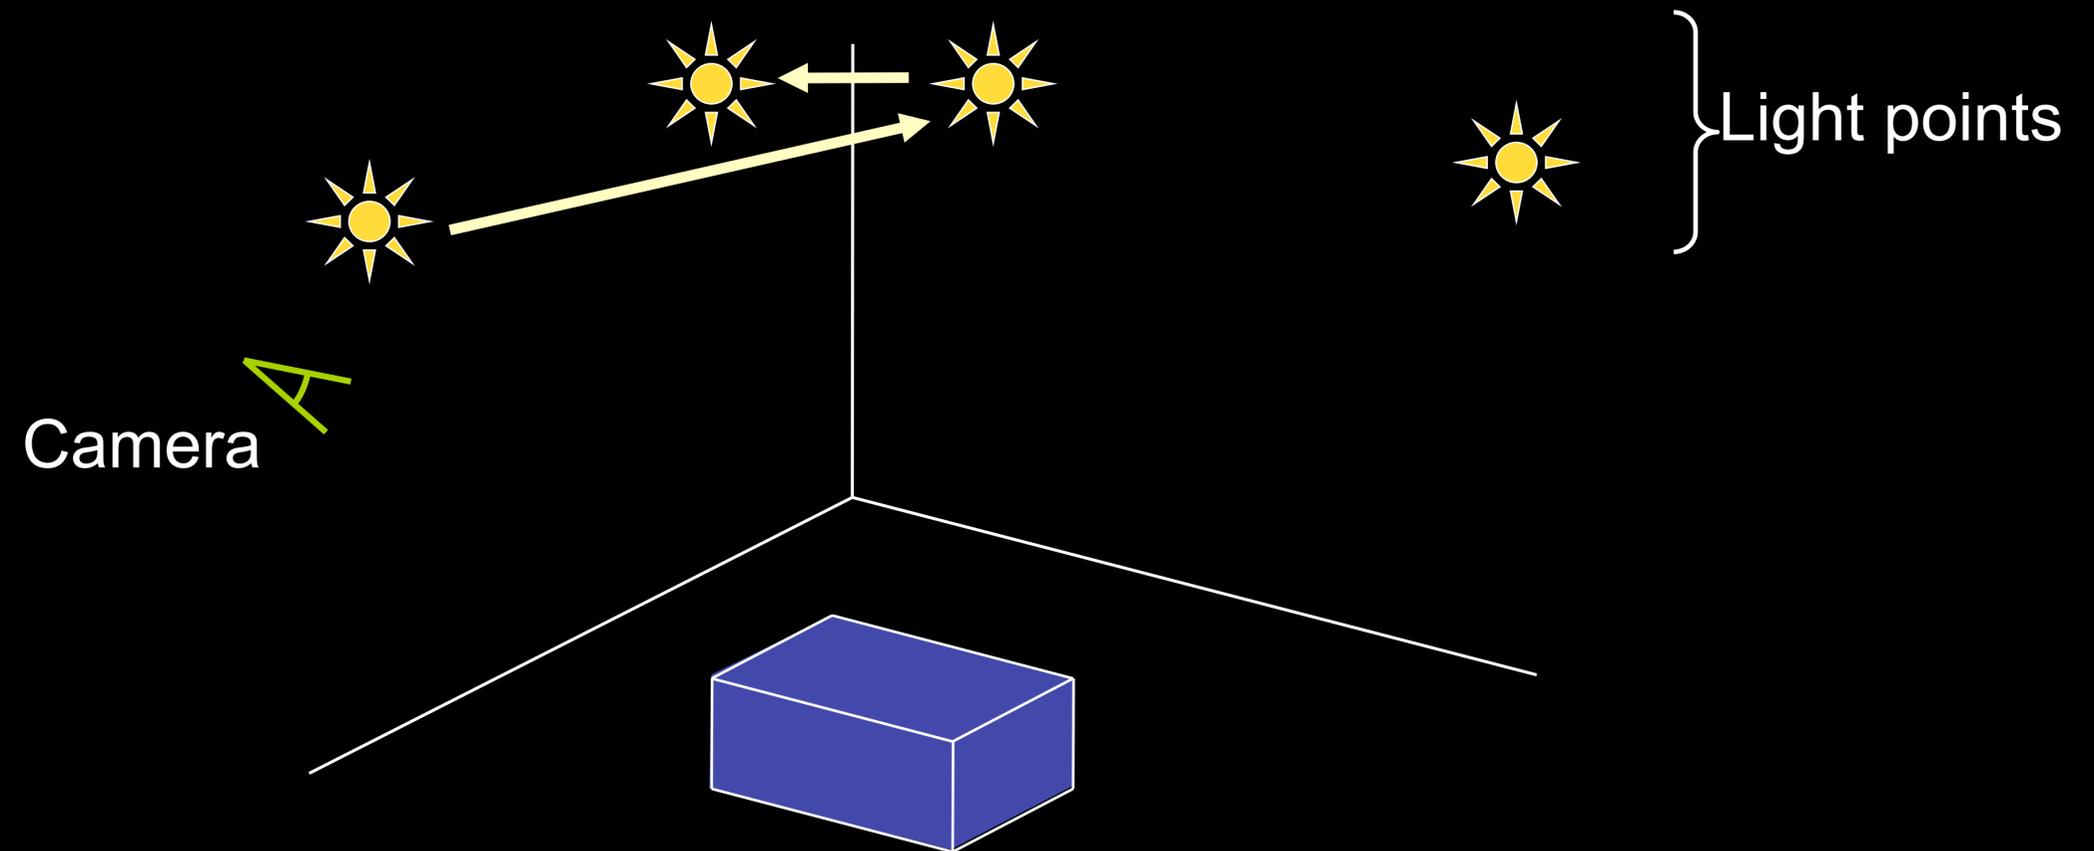
- Discretize full integral into 2 point sets
  - Light points (L)
  - Gather points (G)



Monday, August 6, 2012  
So now let's describe how the system works. We discretize the full integral equation using two point sets. First we convert all the light sources in the scene to point light approximations.

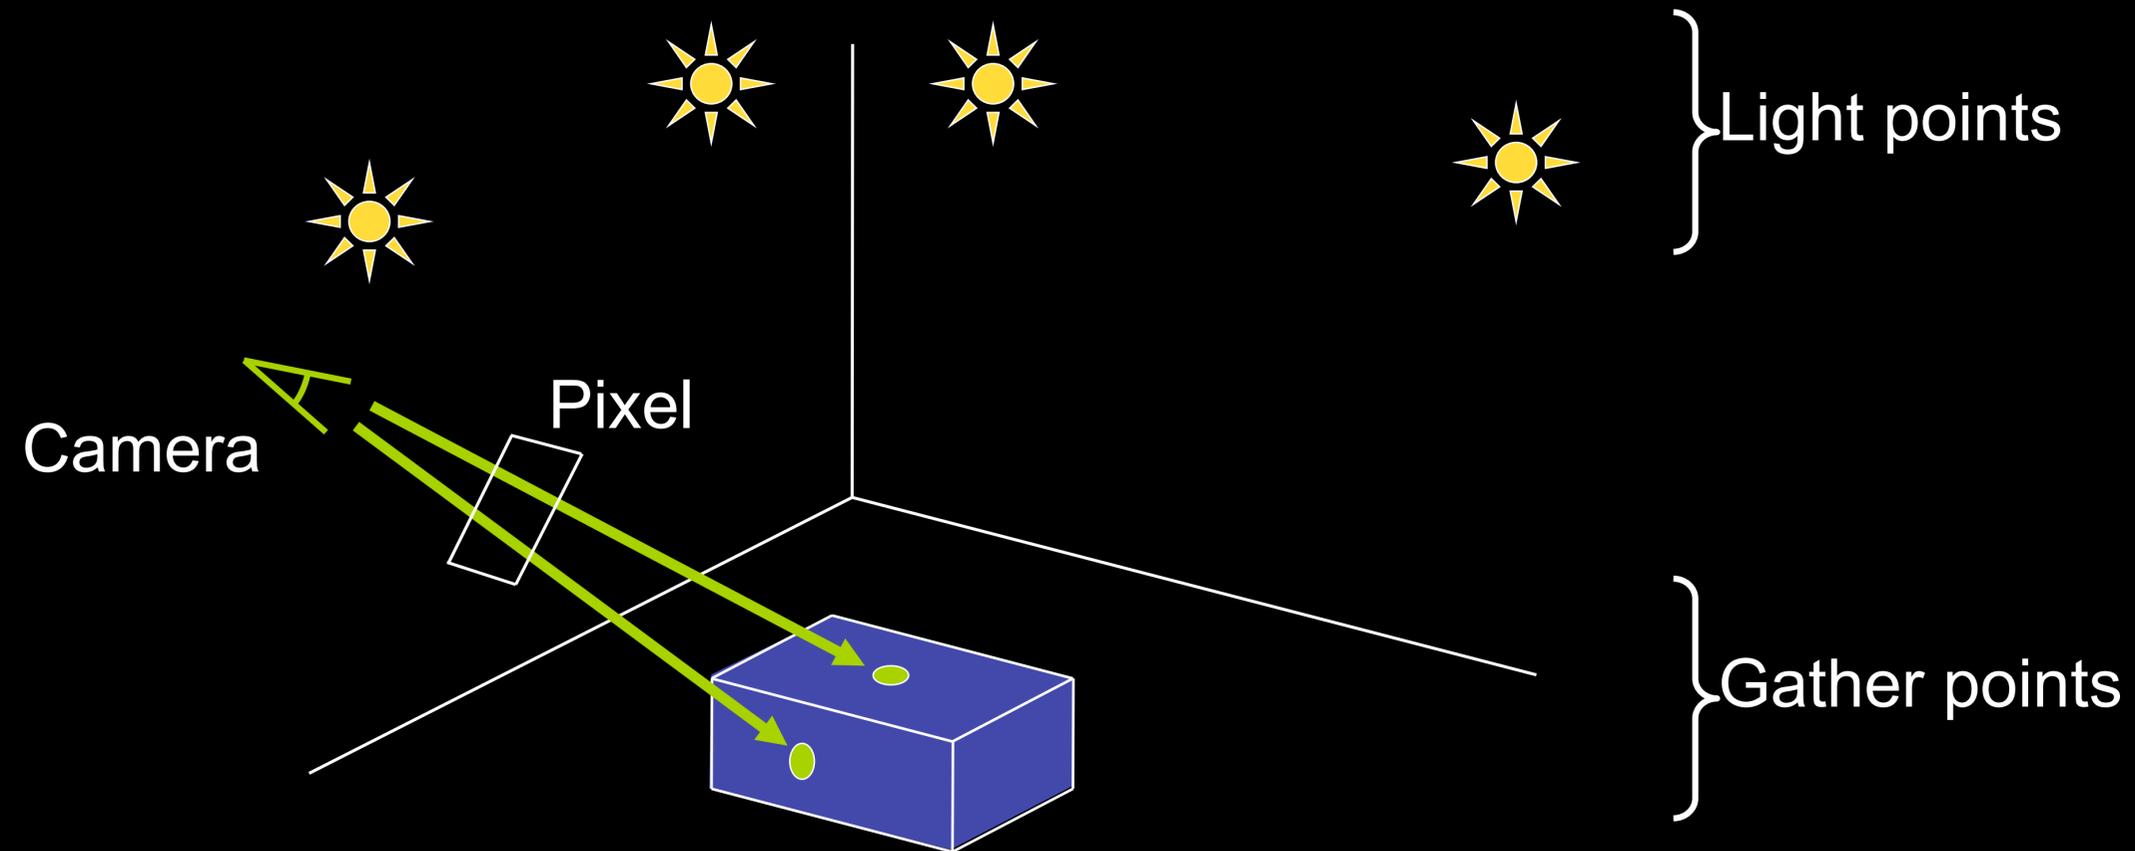
# Point Sets

- Discretize full integral into 2 point sets
  - Light points (L)
  - Gather points (G)



# Point Sets

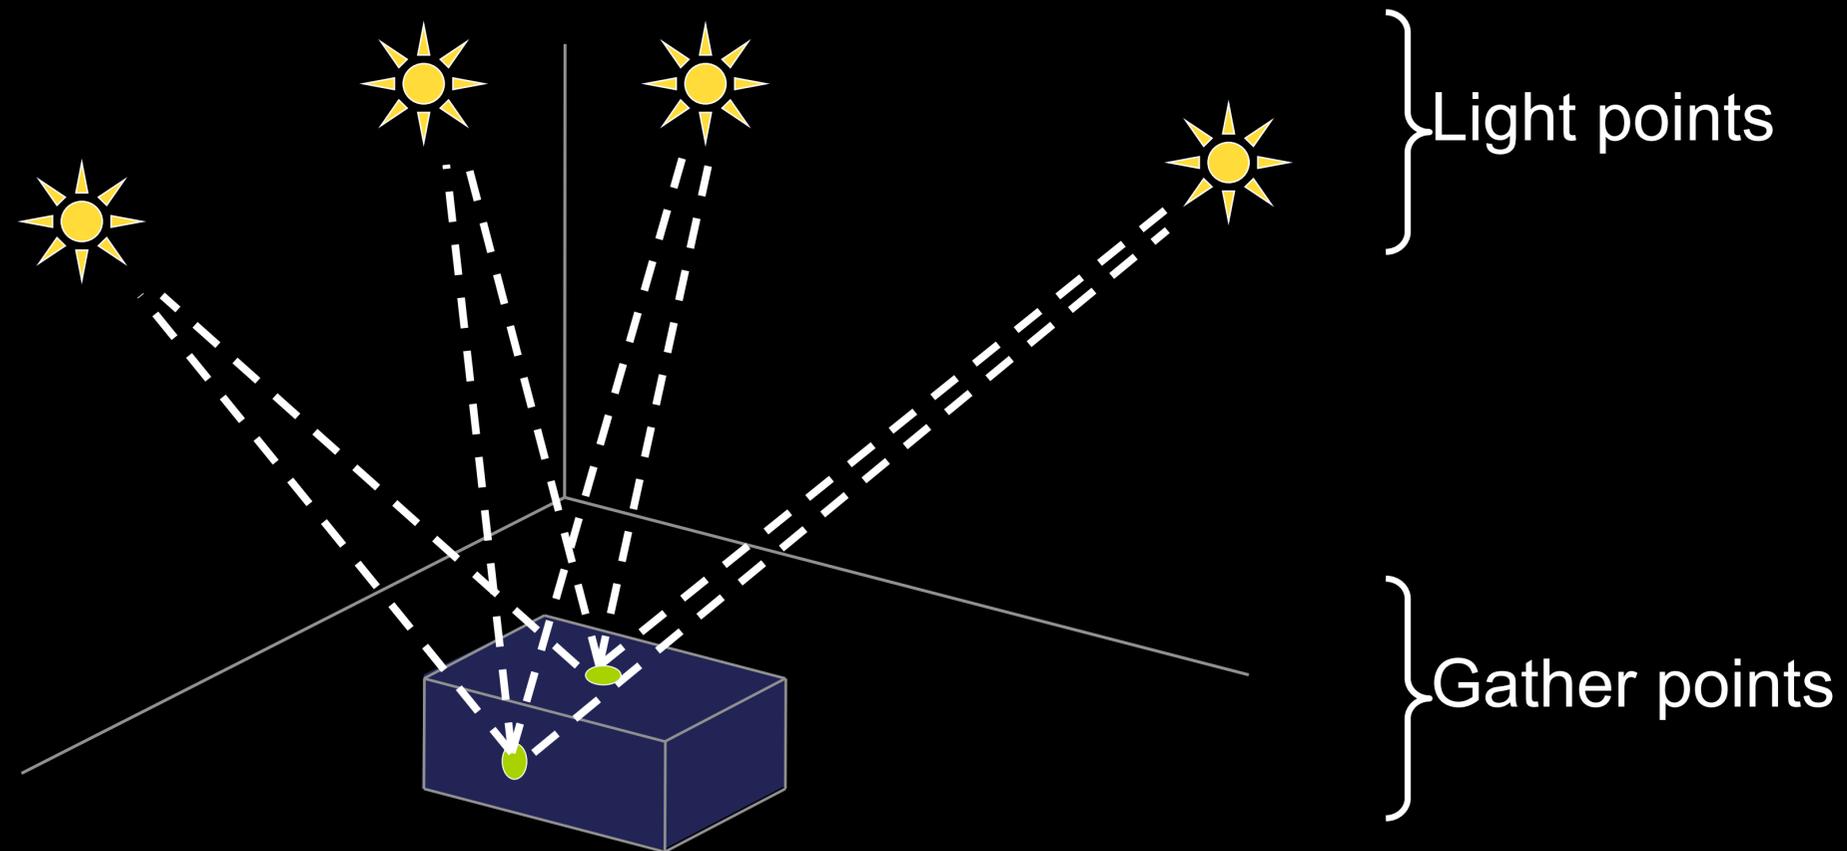
- Discretize full integral into 2 point sets
  - Light points (L)
  - Gather points (G)



Monday, August 6, 2012  
 Then for each pixel we trace rays from the camera through the pixel to generate gather points. Note that these point can be generated over time, over the volume, camera aperture, etc. to account for the different integral domains.

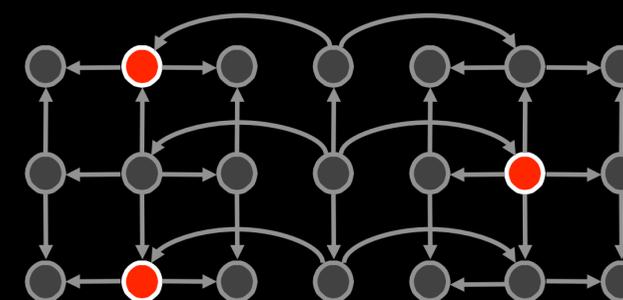
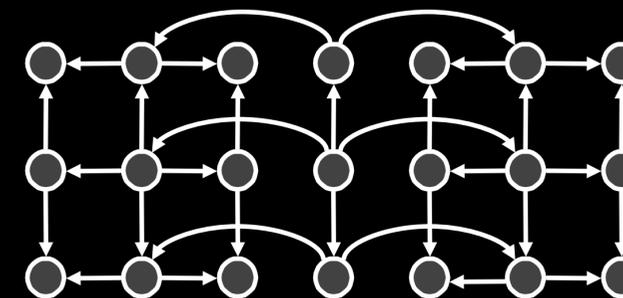
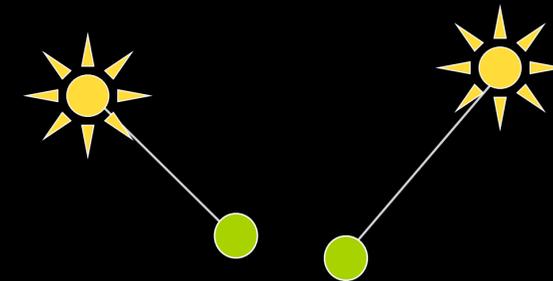
# Point Sets

- Discretize full integral into 2 point sets
  - Light points (L)
  - Gather points (G)



# Key Concepts

- Unified representation
  - Pairs of gather and light points
- Hierarchy of clusters
  - The product graph
- Adaptive partitioning (cut)
  - Cluster approximation
  - Cluster error bounds
  - Perceptual metric (Weber's law)



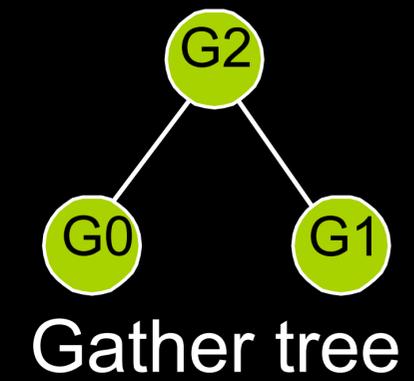
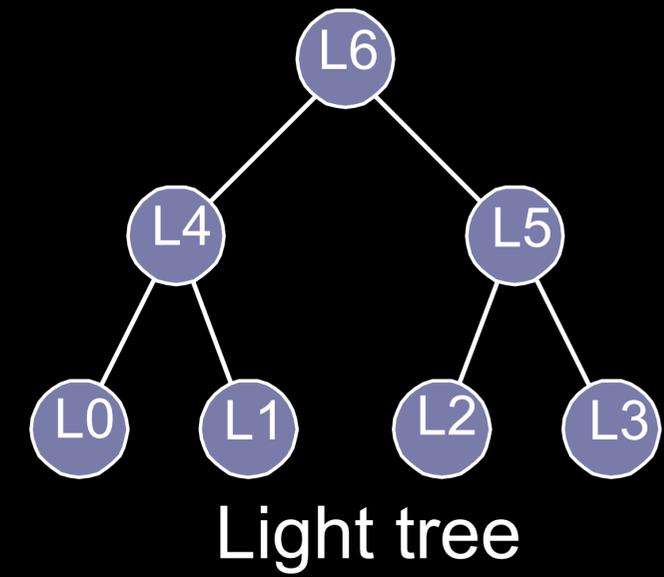
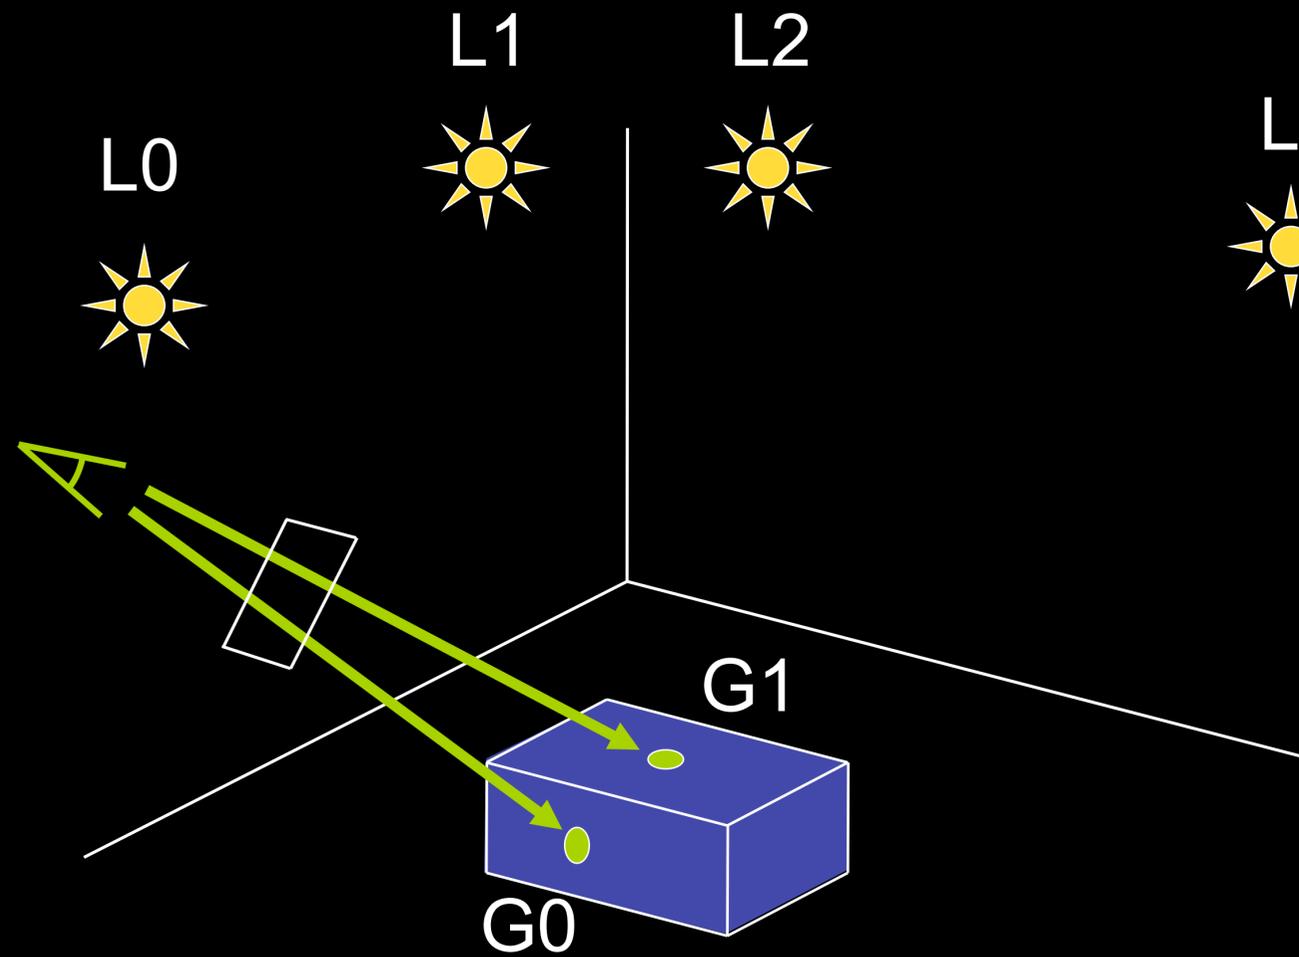
Monday, August 6, 2012  
 There are several key concepts we use for this scalable solution. First we reduce all the integrals to a unified representation, namely pairs of gather and light points. Then we form a hierarchy of clusters over the set of pairs which we call the product graph. Then we use this hierarchy to dynamically select a appropriate partitioning of the pairs into clusters which we call a cut. To achieve this we need an inexpensive way to approximate each cluster, to bound the error introduced by this approximation and an perceptual metric.

# Product Graph

---

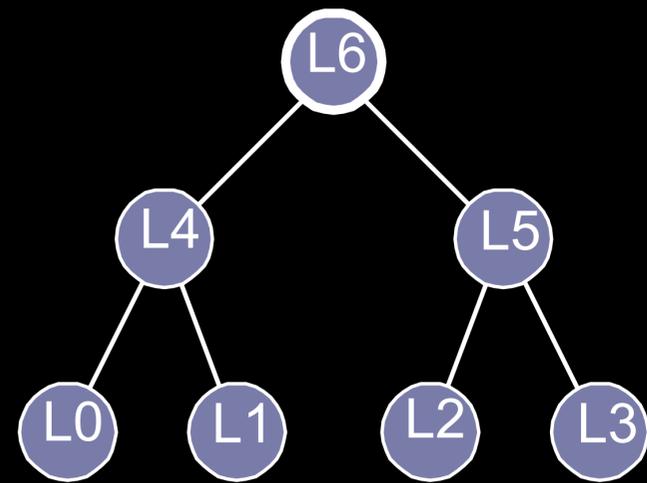
- Explicit hierarchy would be too expensive
  - Up to billions of pairs per pixel
- Use implicit hierarchy
  - Cartesian product of two trees (gather & light)

# Product Graph





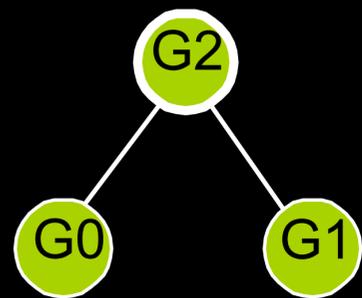
# Product Graph



Light tree

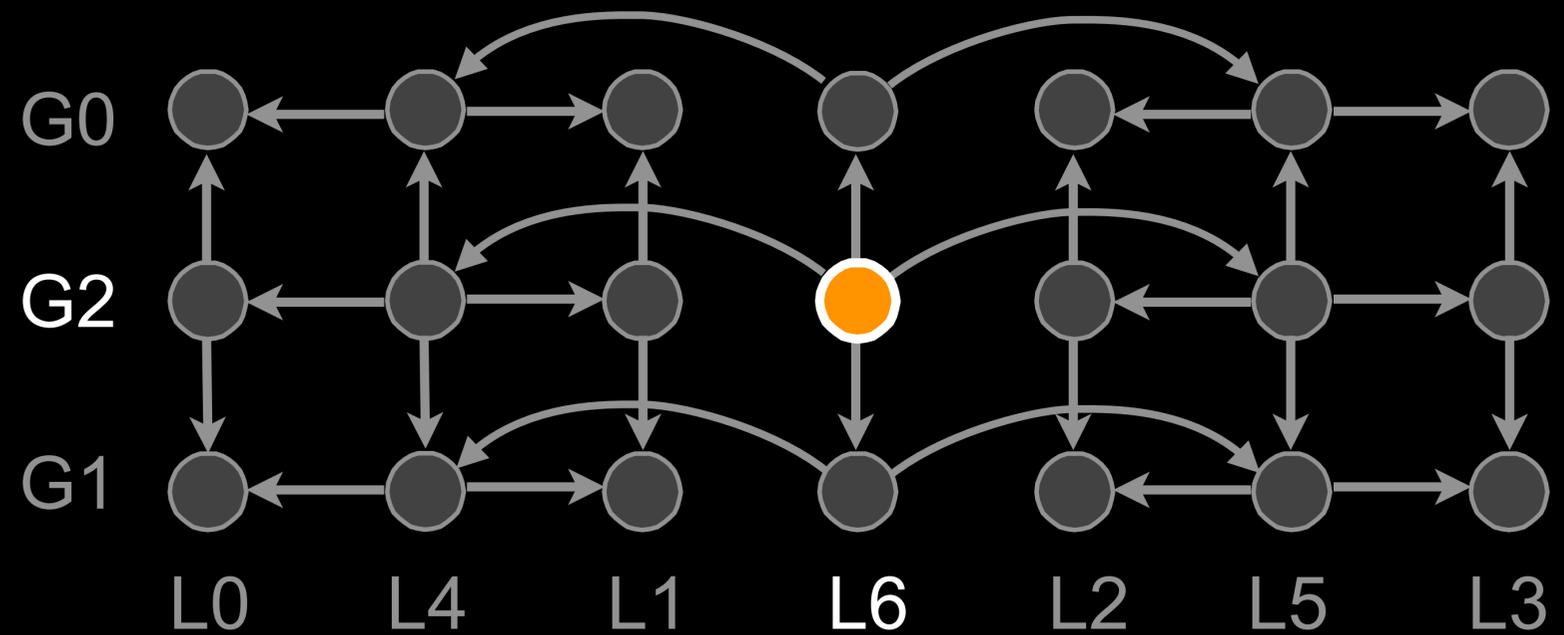
=

X

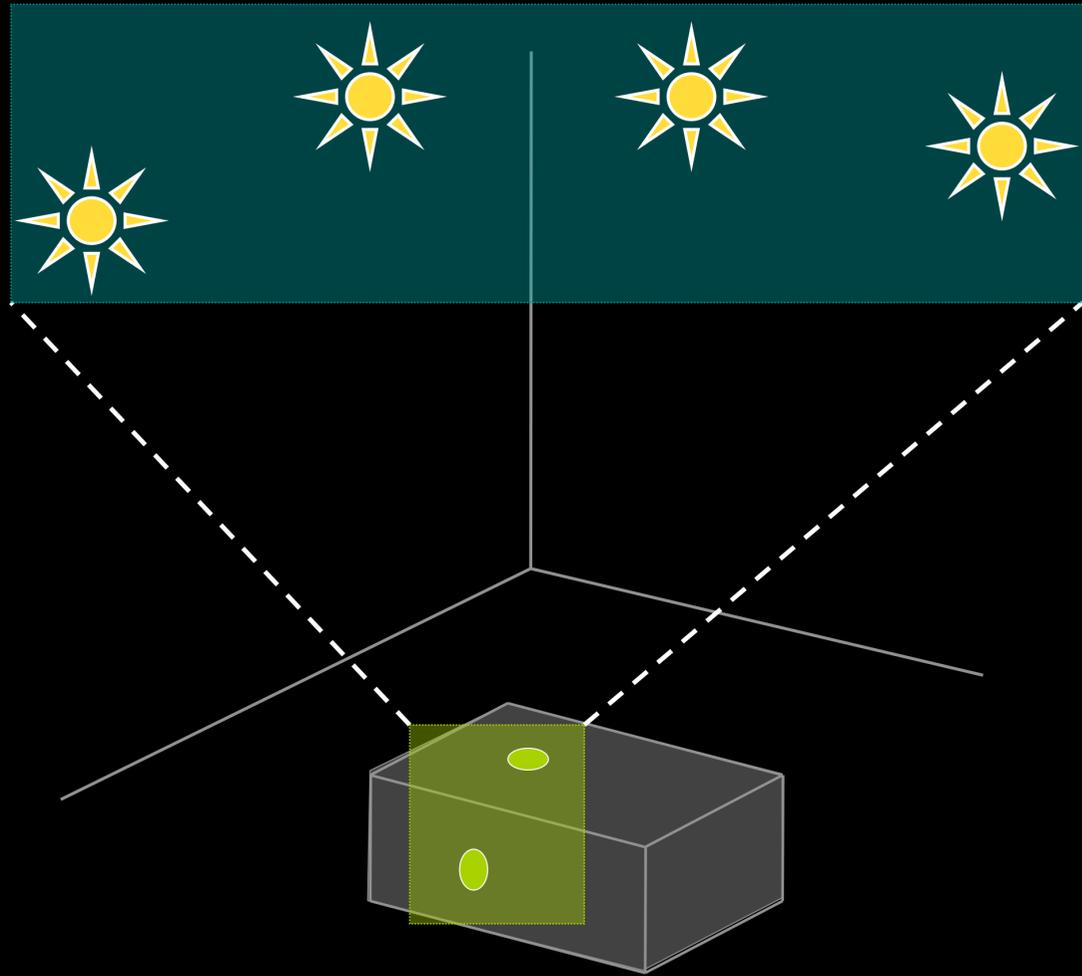


Gather tree

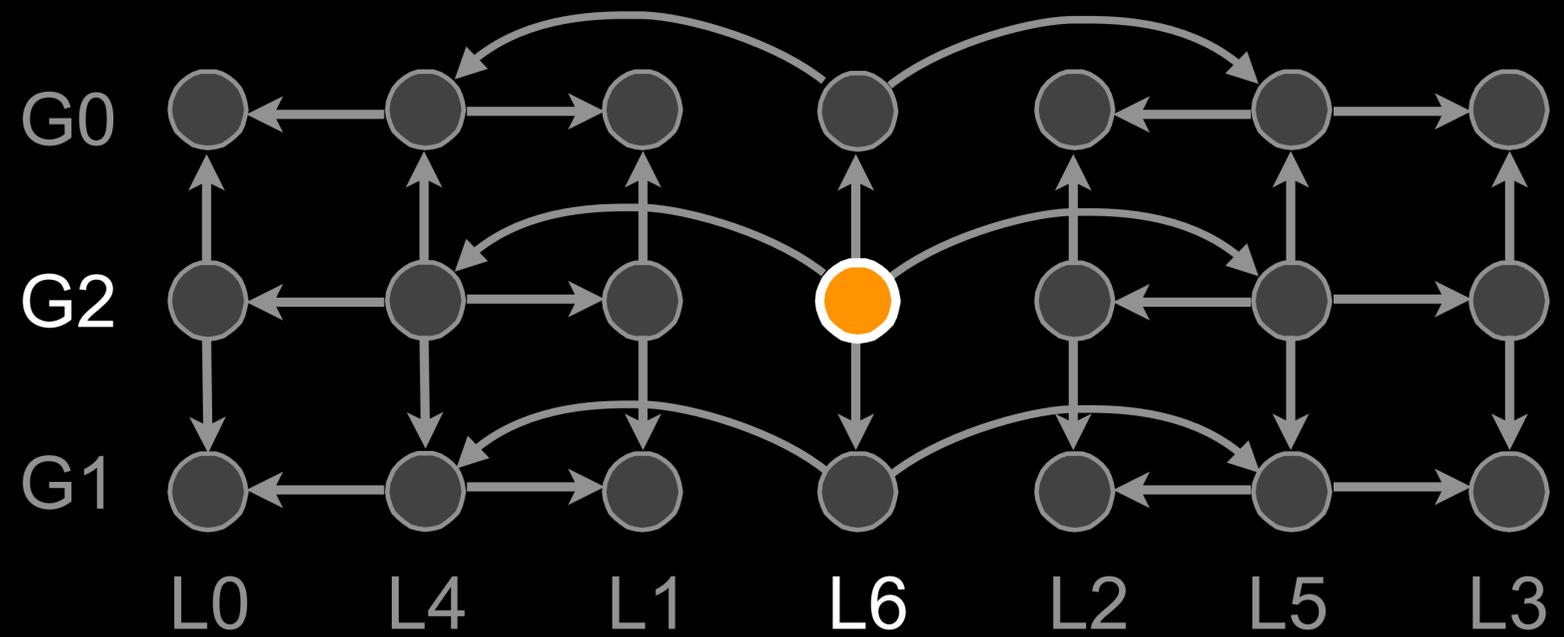
## Product Graph



# Product Graph



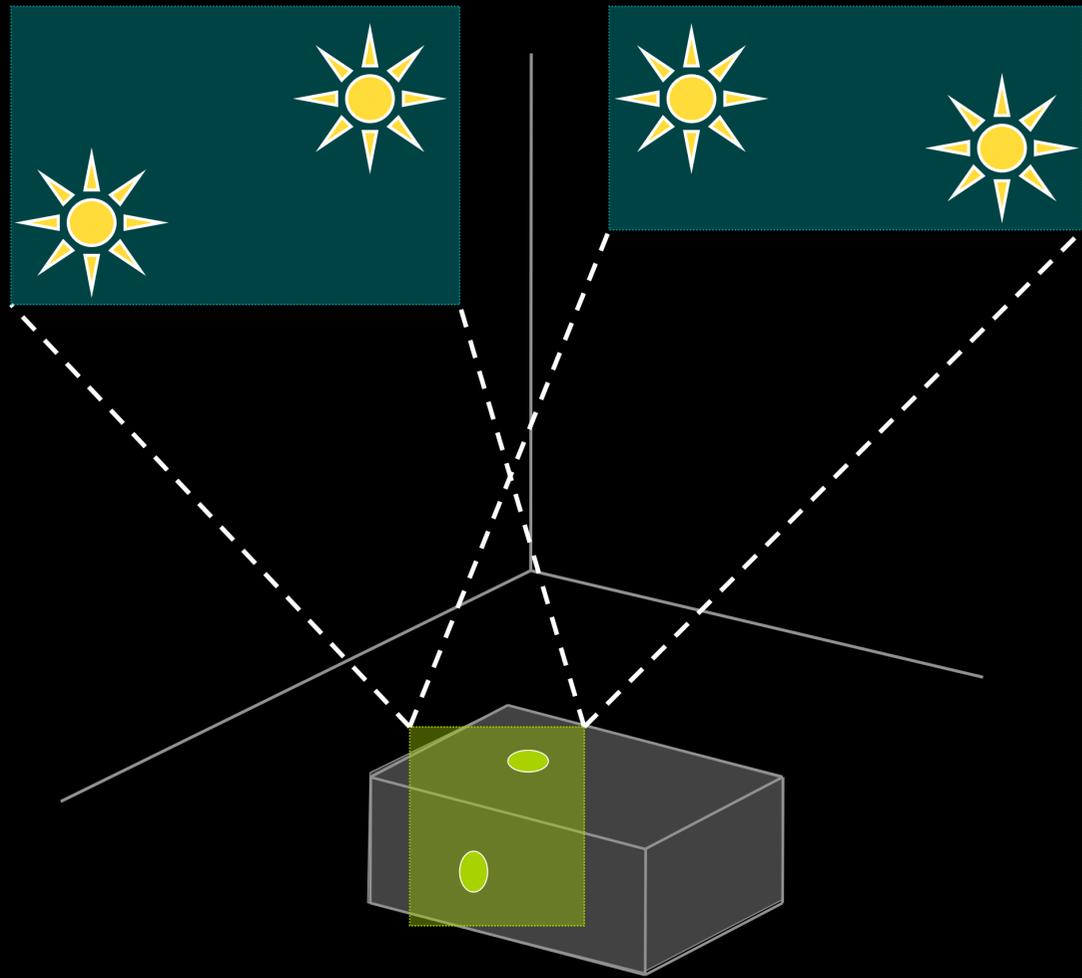
## Product Graph



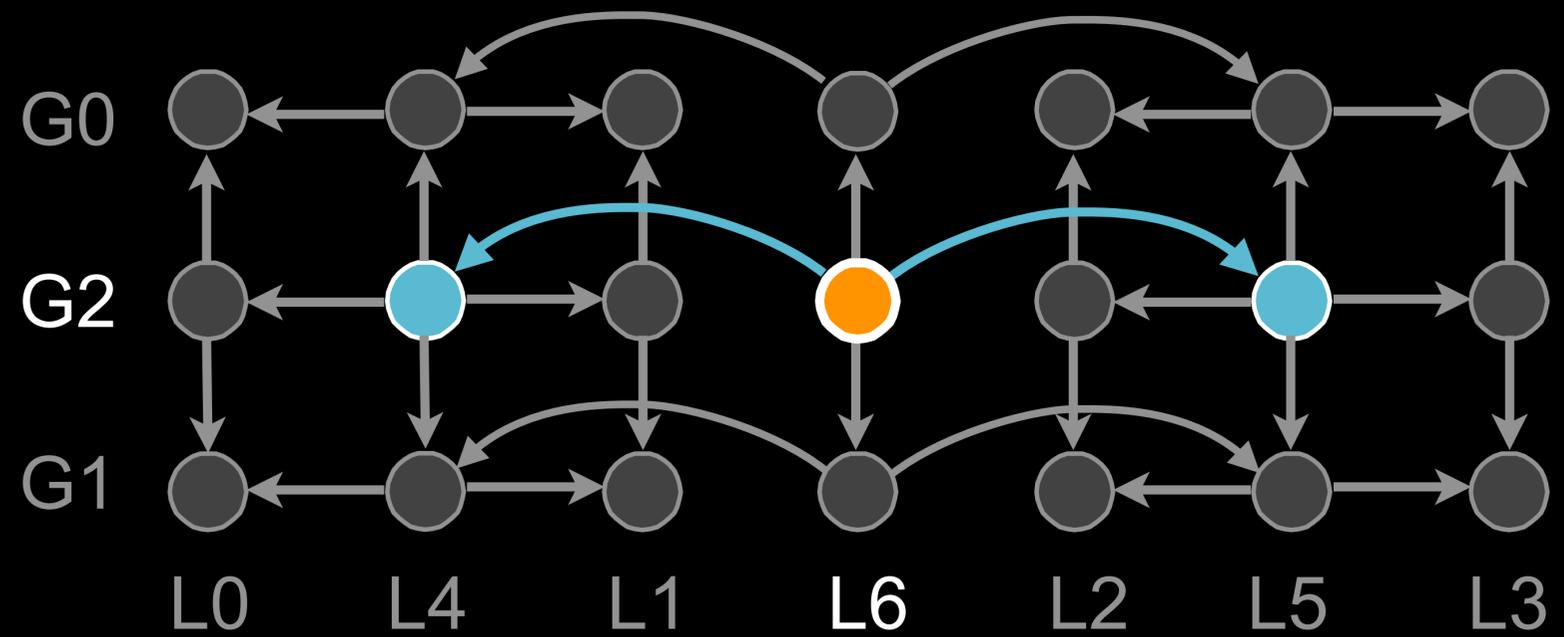
Monday, August 6, 2012

In the scene this corresponds to combining all the light points into a single cluster, combining all the gather points into a single cluster and then treating the set of all pairs as just a single interaction between these clusters. In general that is not going to be an accurate enough approximation of the pixel value, so we need a way to refine this approximation.

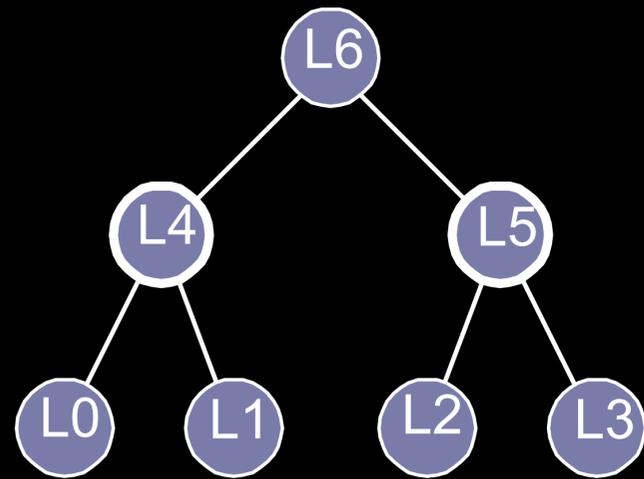
# Product Graph



## Product Graph



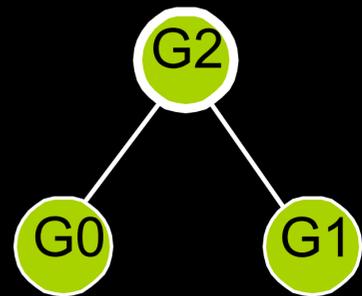
# Product Graph



Light tree

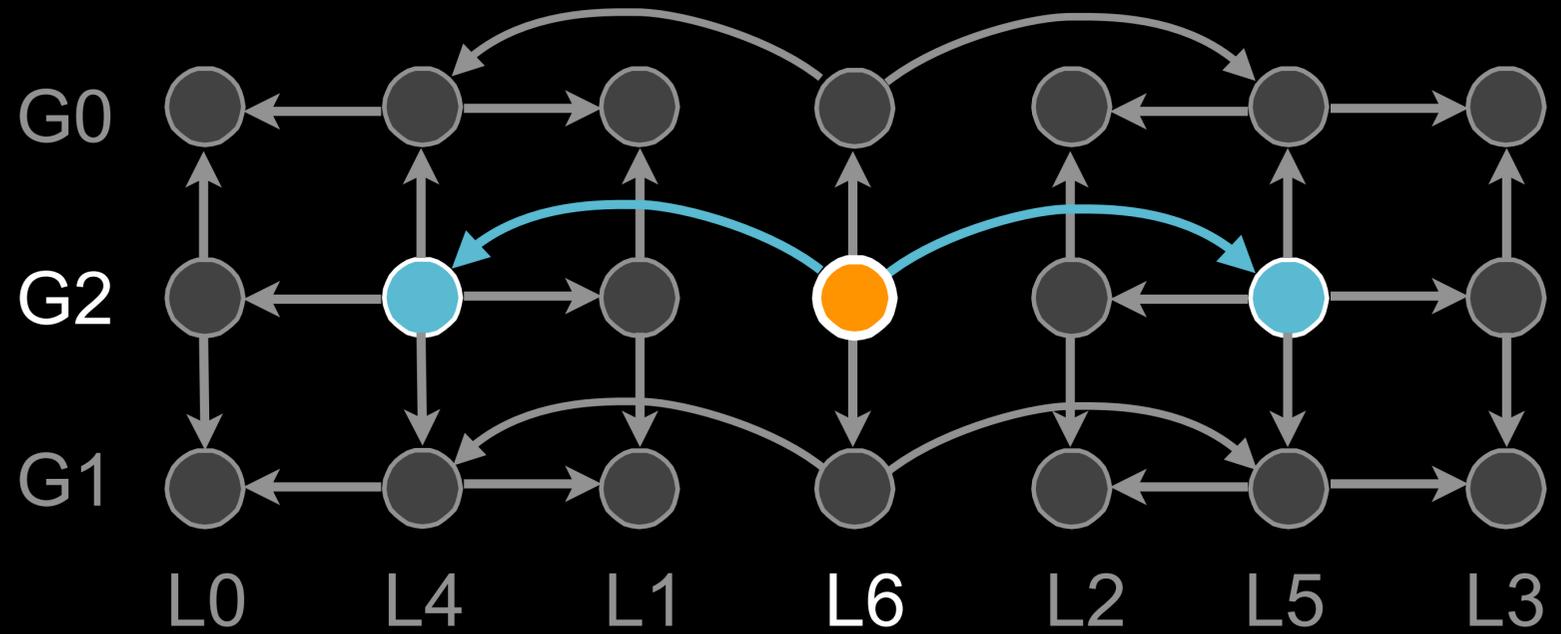
=

X

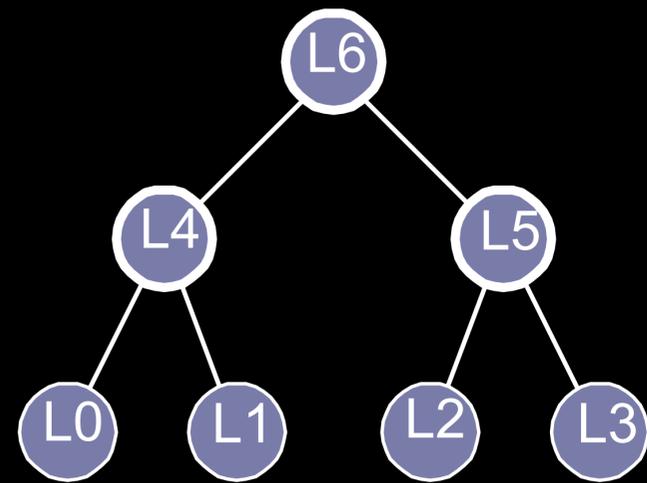


Gather tree

## Product Graph



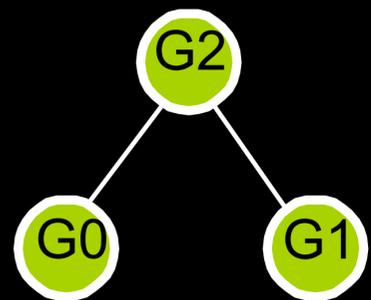
# Product Graph



Light tree

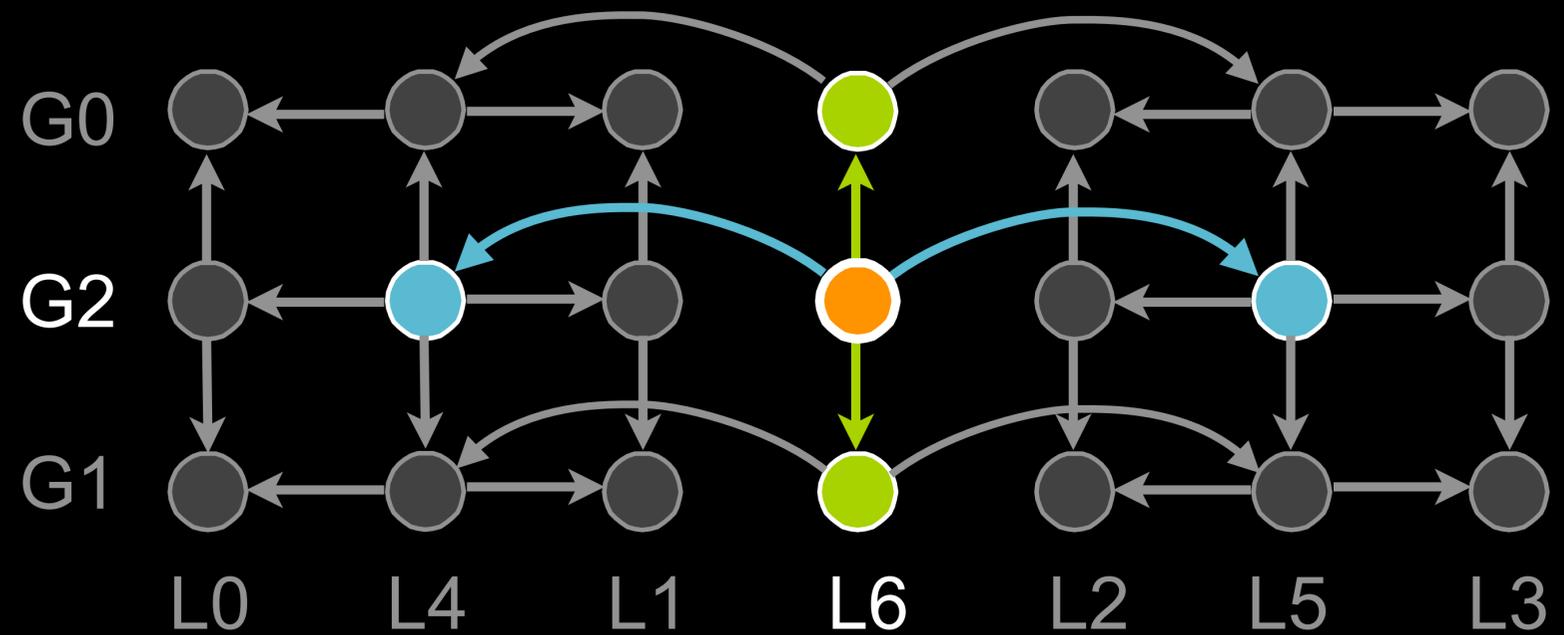
=

X



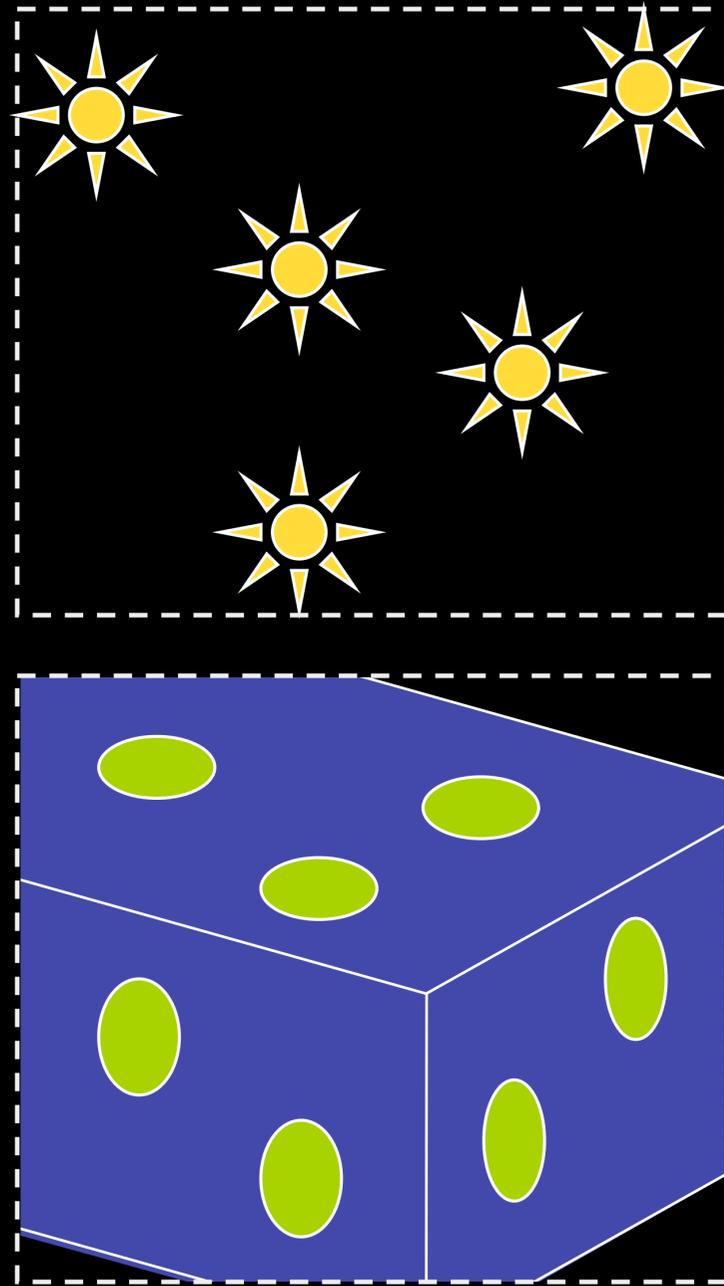
Gather tree

## Product Graph





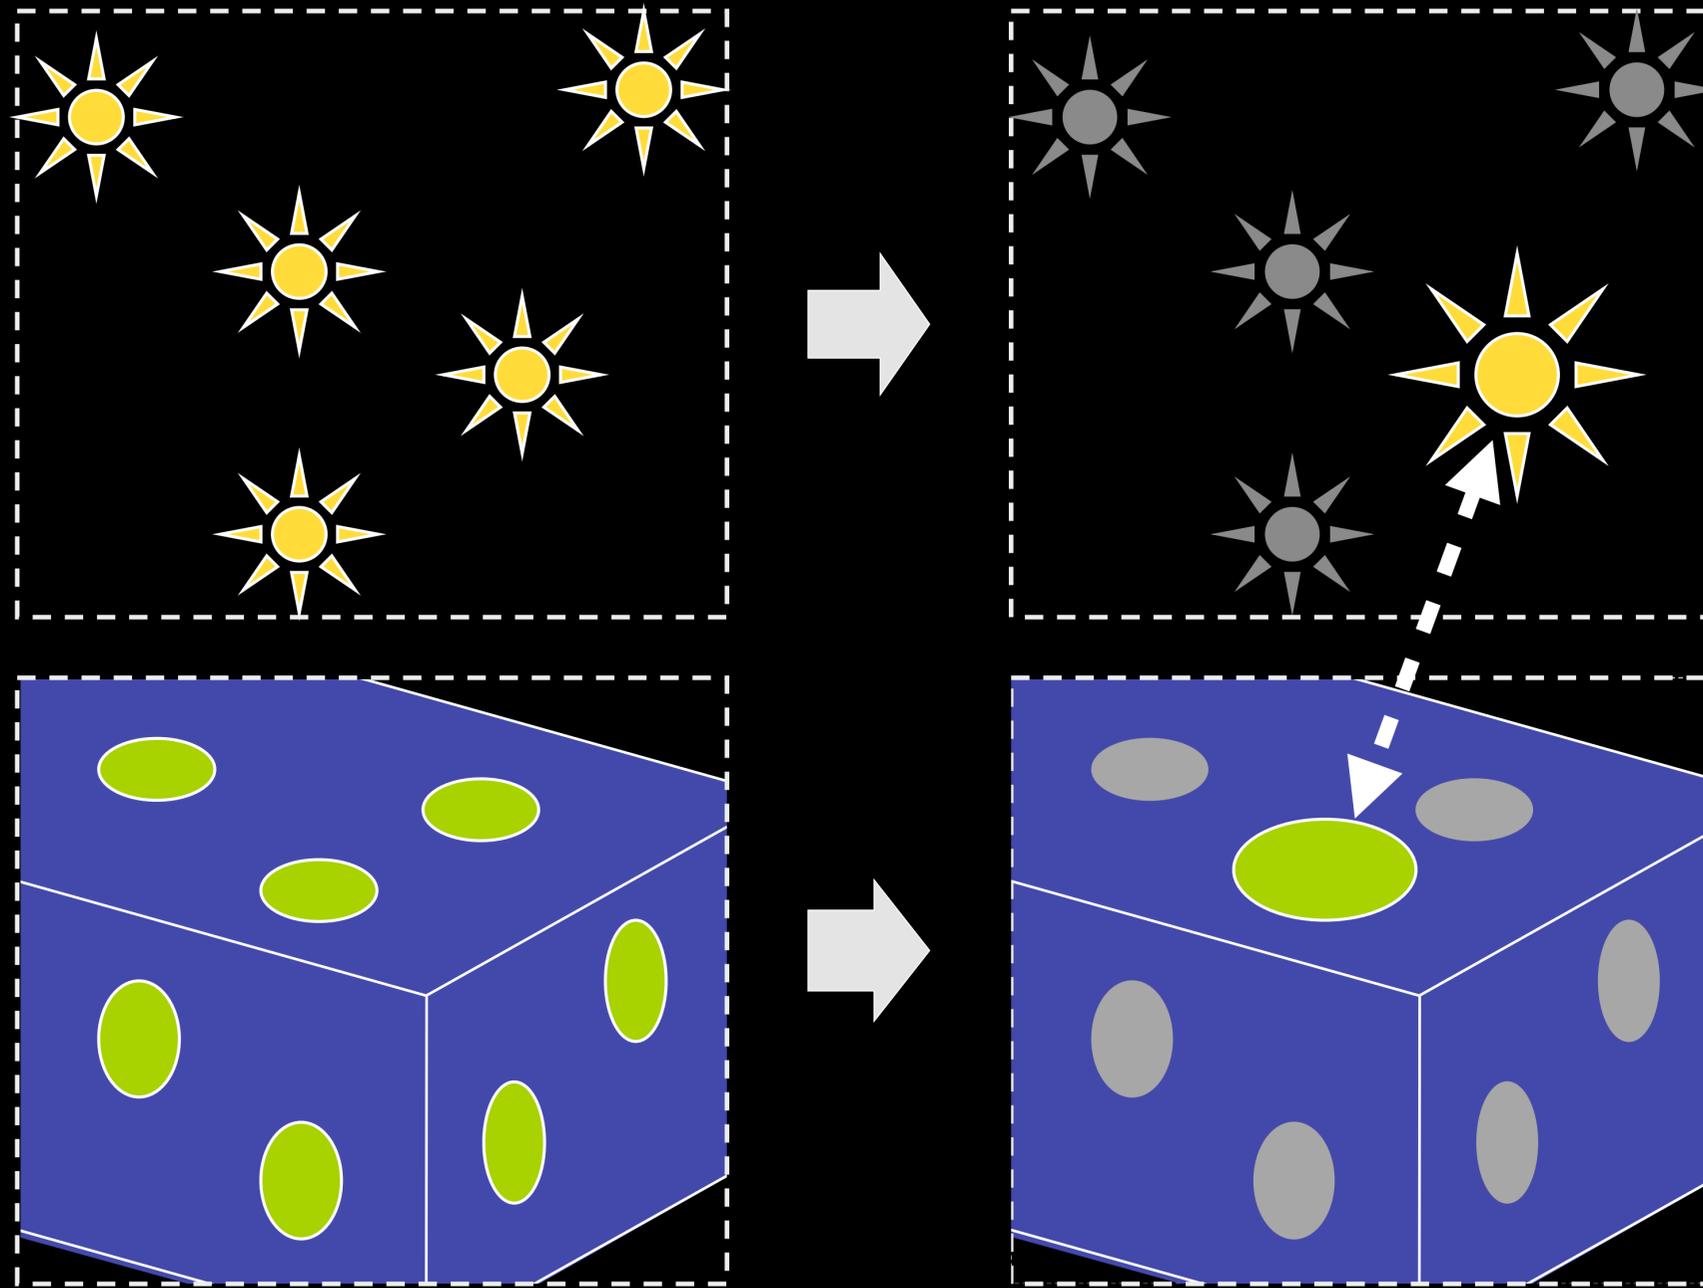
# Cluster Representatives



Monday, August 6, 2012

A cluster consists of a set of gather points and a set of lights and represents all possible pairings between the gather and light points. To cheaply approximate a cluster, we select a single pair and evaluate it, scaled by an appropriate factor based on the probability of picking that pair.

# Cluster Representatives

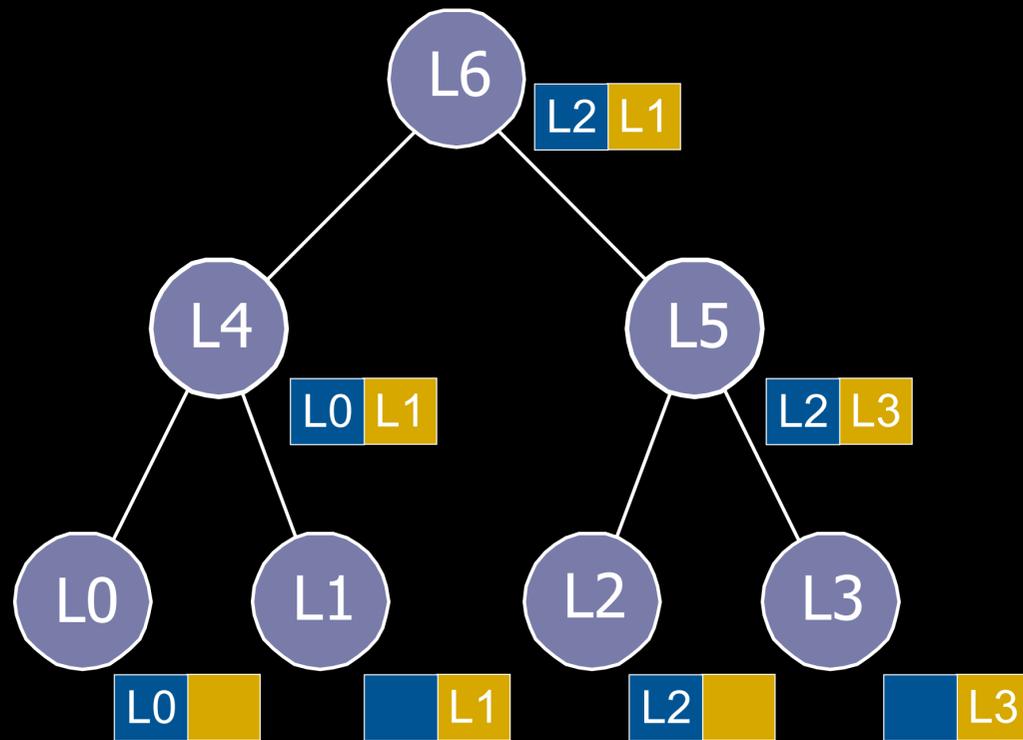


Monday, August 6, 2012

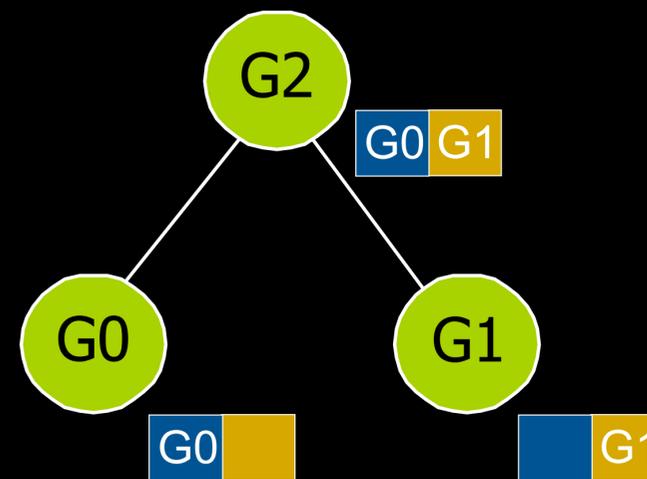
Each time we evaluate a cluster, we select one light as its representative. We pick one of the light to be the representative light and similarly one representative gather point. Then we approximate all the gather-light pairs using just the single pair of the representative gather point and the representative light point.

# Multiple Representatives

Light Tree



Gather Tree

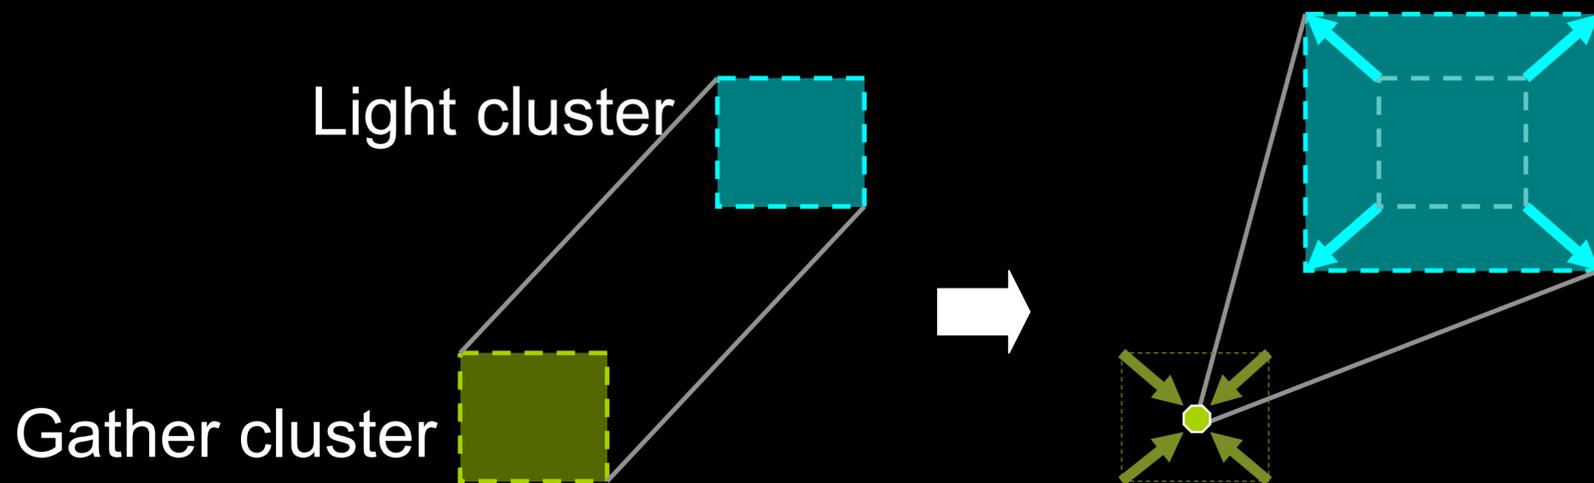


Exists at first or second time instant.

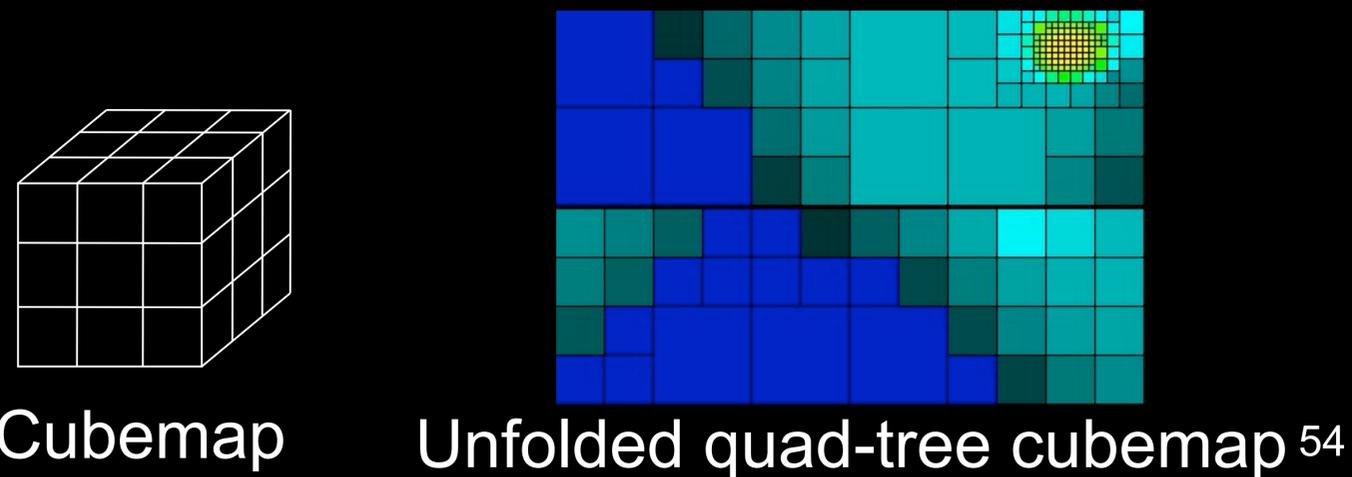
Monday, August 6, 2012  
 One important change from Lightcuts is that we store multiple representatives per cluster node. This is essential when integrating over time as only nodes at the same instant of time can be connected, but is also very useful even not integrating over time. It helps reduce correlation between pixels and reduce error regardless of which refinement type is chosen.

# Error Bounds

- Collapse cluster-cluster interactions to point-cluster
  - Use Minkowski sums



- Combine bounds over multiple points
  - Rasterize into cube-maps



Monday, August 6, 2012  
 We also need to generalize the error bounding to include gather cluster to light clusters rather than just the point to cluster bounds in the original Lightcuts. However we can collapse cluster-cluster interactions back into point-cluster interactions using Minkowski sums where we shrink one cluster to a point while simultaneously expanding the other cluster. We also need to combine the bounds from all points in a cluster into a common representation. We do this by rasterizing them into cubemaps that cover the sphere of directions. In the paper we used fixed resolution cubemaps but we have since switched to adaptive cubemaps (quad-tree on each face) because these better handle glossy BRDFs and other strongly directional distributions.

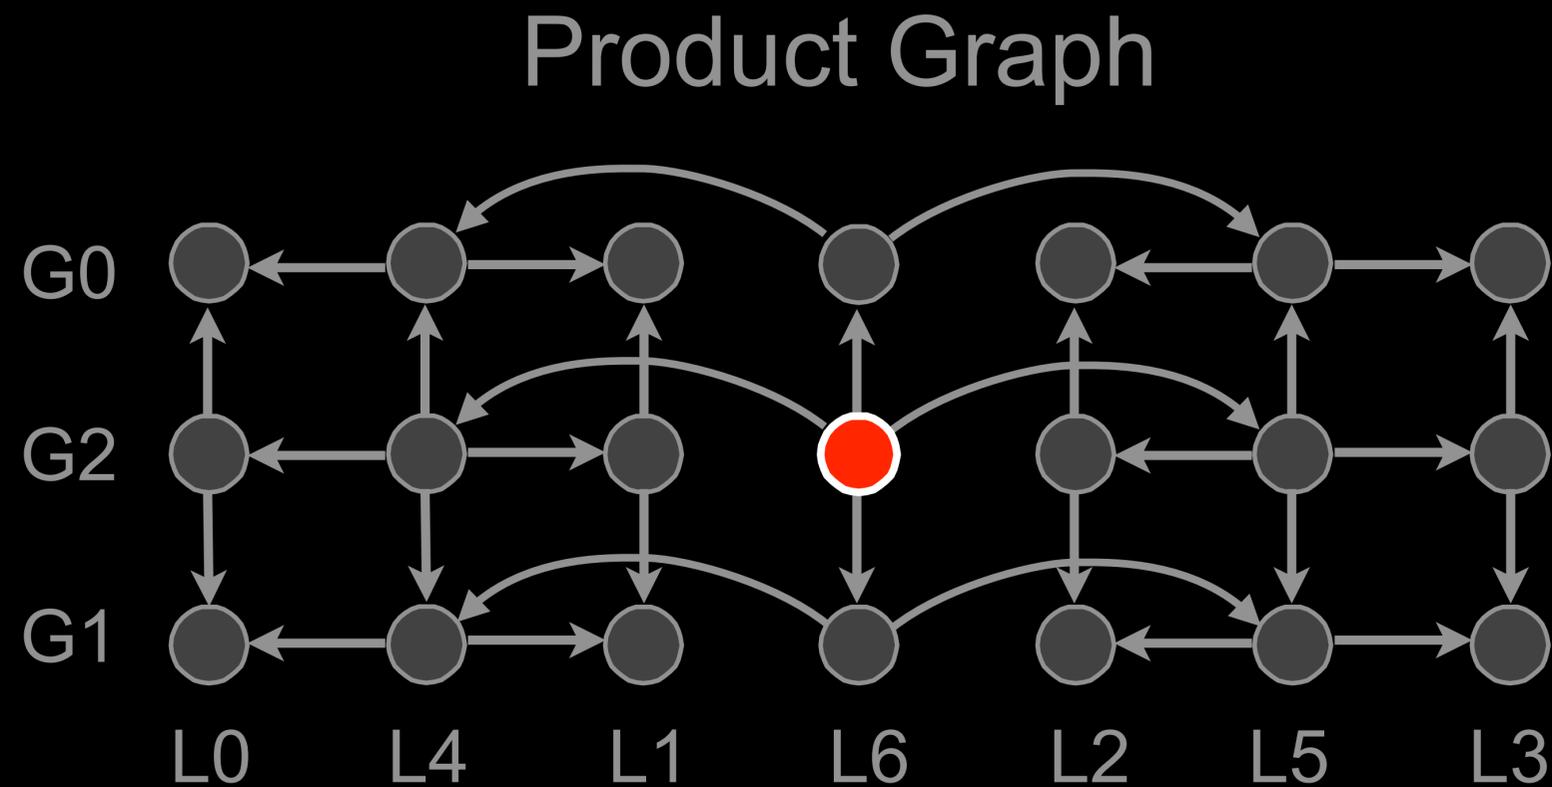
# Algorithm Summary

---

- Once per image
  - Create lights and light tree
- For each pixel
  - Create gather points and gather tree for pixel
  - Adaptively refine clusters in product graph until all cluster errors  $<$  perceptual metric

# Algorithm Summary

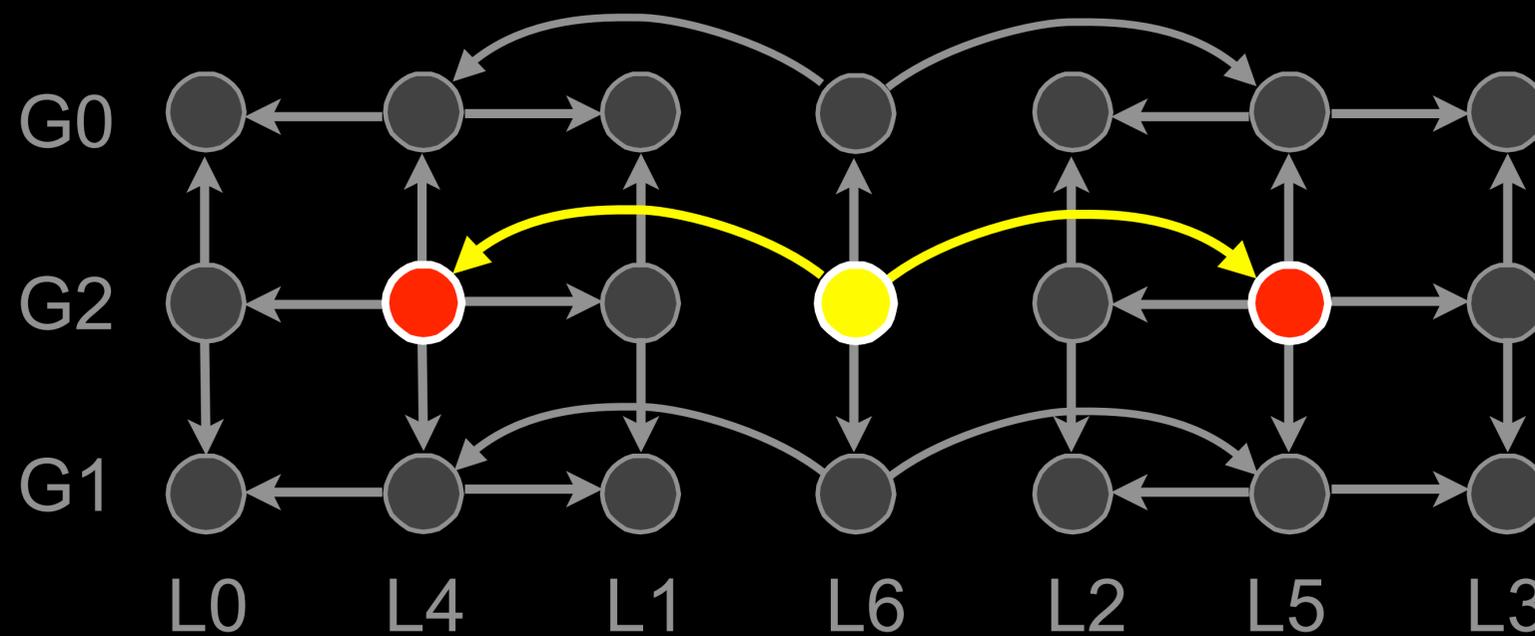
- Start with a coarse cut
  - Eg, source node of product graph



# Algorithm Summary

- Choose node with largest error bound & refine
  - In gather or light tree

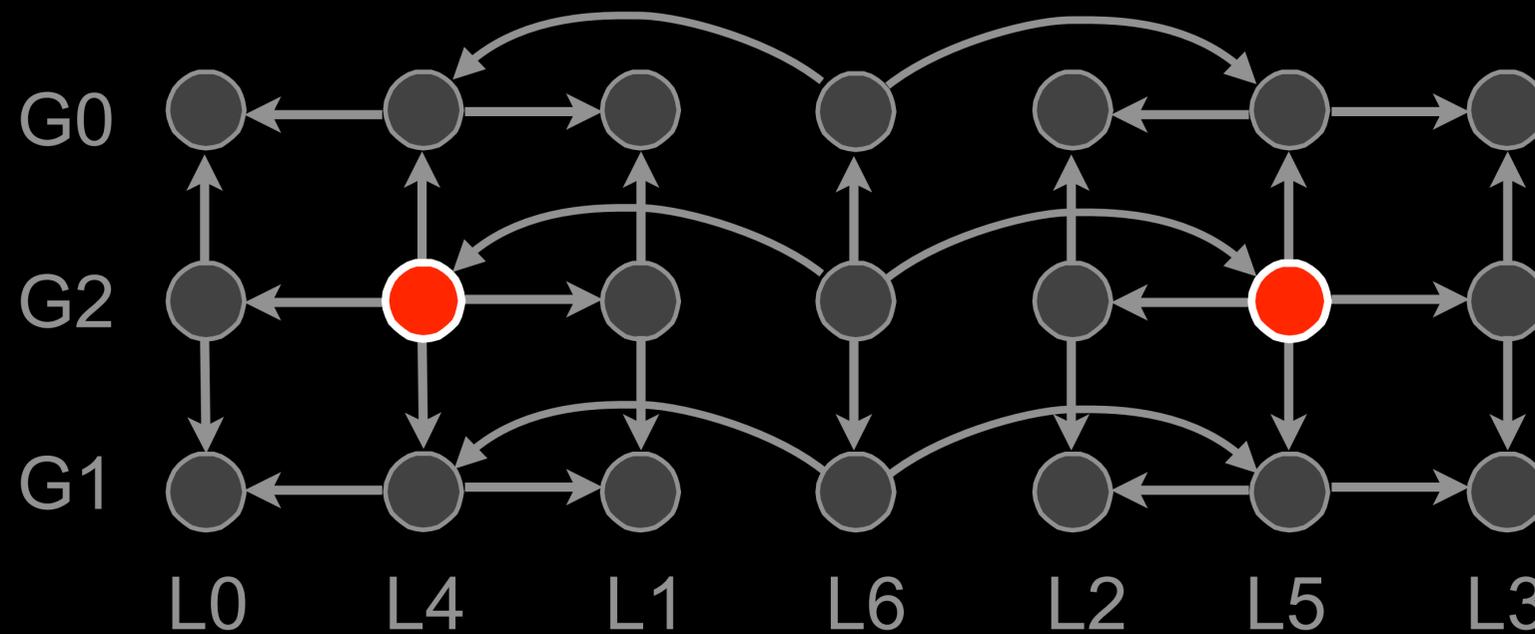
Product Graph



# Algorithm Summary

- Choose node with largest error bound & refine
  - In gather or light tree

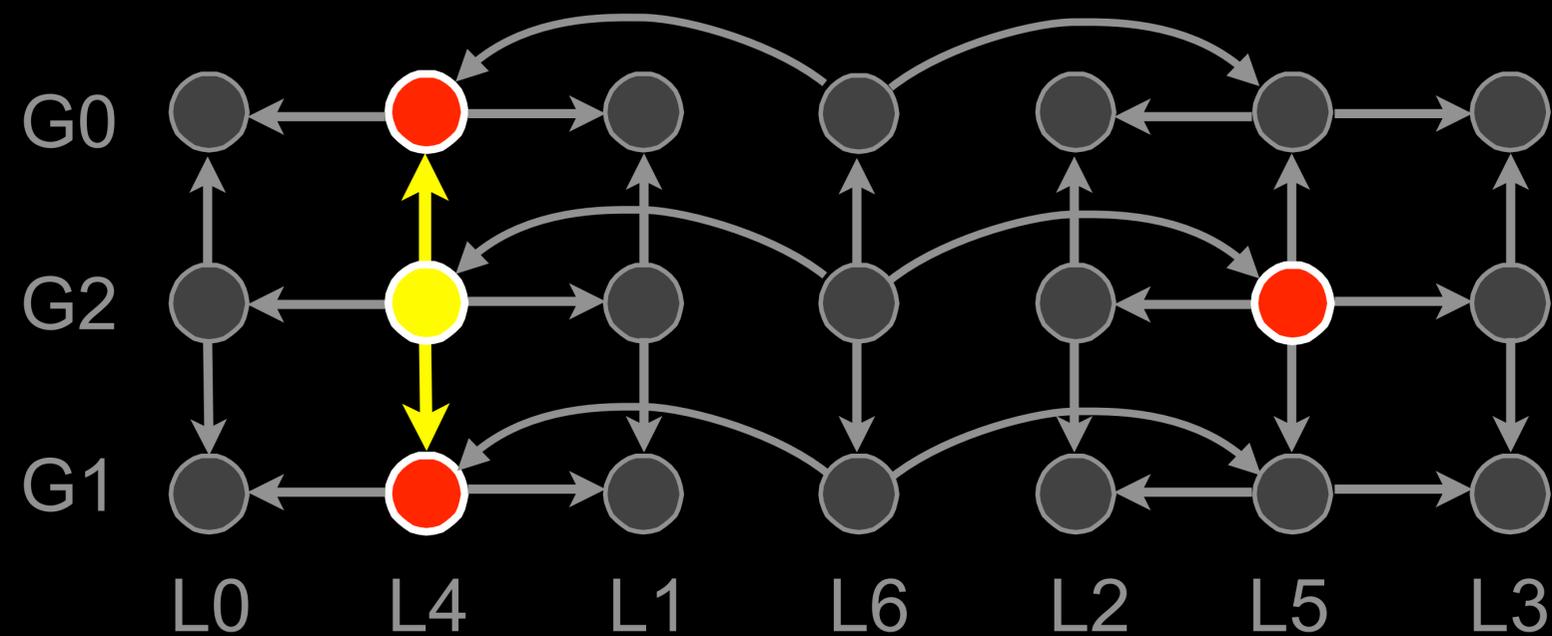
Product Graph



# Algorithm Summary

- Repeat process

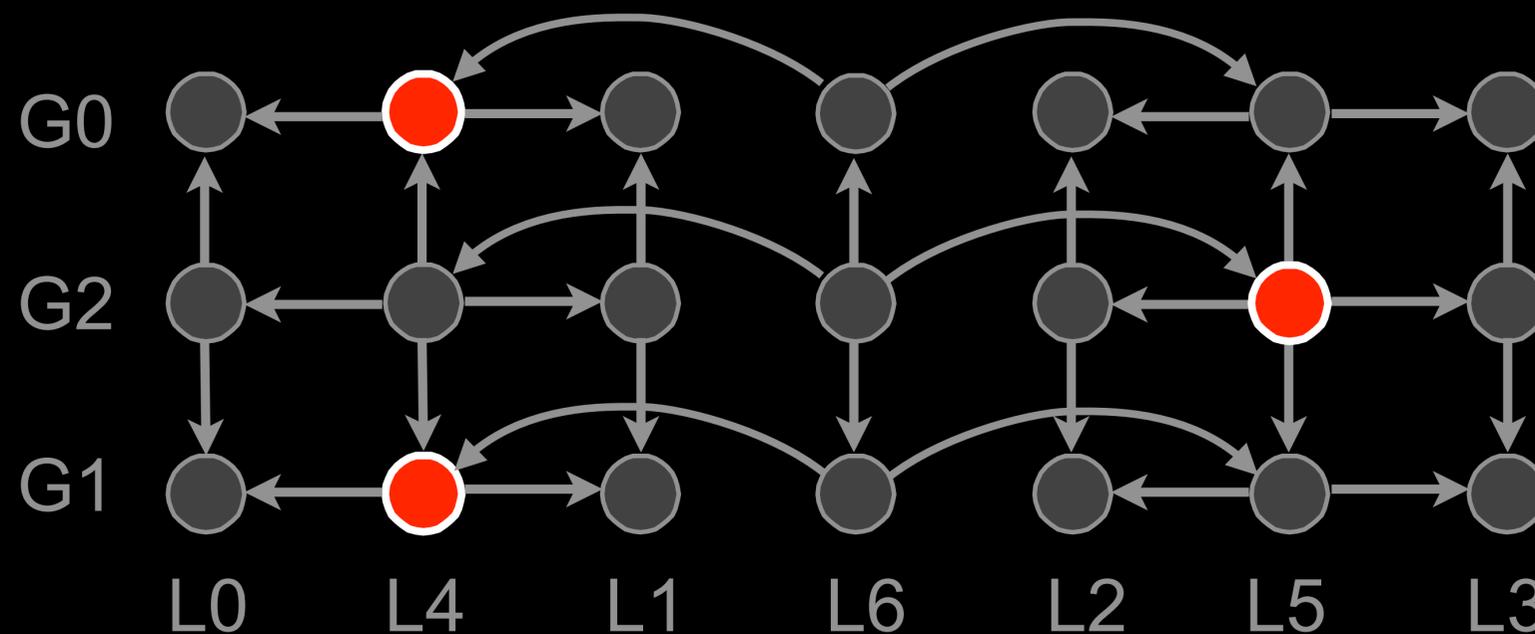
Product Graph



# Algorithm Summary

- Until all clusters errors < perceptual metric  
 –2% of pixel value (Weber’s law)

Product Graph



# Roulette



7,047,430 Pairs per pixel  
Avg cut size 174 (0.002%)

Time 590 secs

Monday, August 6, 2012

This is an example with a spinning roulette wheel which has strong motion blur. Here we are using 256 temporal samples per pixel for an average of over 7 million potential pairs per pixel. However the cut size is only 174 which means that we only actually evaluated 174 pairs which is only a tiny fraction of all the pairs.

# Conclusions

---

- Unified handling of complex effects
  - Motion blur, participating media, depth of field etc.
- Product graph
  - Implicit hierarchy over billions of pairs
- Scalable scene and illumination complexity
  - Millions to billions of point pairs per pixel

# The End



Monday, August 6, 2012

That concludes my talk and I'd be happy to answer any questions.

# Schedule

- 2:00 (05 min) ... Introduction & Welcome (*Křivánek*)
- 2:05 (30 min) ... Instant Radiosity (*Keller*)
- 2:35 (30 min) ... Handling difficult light paths (*Hašan, Křivánek*)
- 3:05 (25 min) ... Scalability with many lights I (*Walter*)
  
- 3:30 (15 min) ... Break
  
- 3:45 (20 min) ... **Scalability with many lights II** (*Hašan*)
- 4:05 (35 min) ... **Real-time many-light rendering** (*Dachsbacher*)
- 4:40 (30 min) ... **ML in Autodesk® 360 Rendering** (*Arbree*)
- 5:10 (05 min) ... **Conclusion - Q & A** (*All*)