C. Project Description
Transient Signal Processing for Realistic Imagery

Capturing and reproducing the richness of our visual world is a central goal of computer graphics. While image processing has had a large impact on antialiasing [Cro77] and local shading [RH01c, BJ01], and despite the importance of Fourier analysis in optics, [Hec01, Goo96] there has been little work in computer graphics on frequency analysis of global lighting. Filling this gap would be a significant step, especially given the importance of sampling the space of rays for rendering [DBB03, Jen01] and material acquisition [War92, MWL+99, MPBM03a, MPN+02]. Frequency analysis of global light transport is challenging for two reasons. First, the domain of the signal is intricate because the space of light rays is five-dimensional: three dimensions for the position and two for the direction. Second, light transport is linear in the light emitted by the sources, but the resulting signal is non-stationary and has discontinuities, which prevents the use of traditional global Fourier analysis. As the geometric and visual complexity of 3D scenes continues to grow, a solid signal-processing perspective is becoming not only relevant but crucial for efficient lighting simulation and material acquisition.

As illustrated in Fig. 1, our goal is to derive a signal-processing framework for all mechanisms of light transport. We propose research in one core theoretical area and two application areas. These areas are summarized in Fig. 2 and all focus on signal characteristics in real imagery.

Theoretical signal-processing framework for light transport. We will study the *local frequency content* as well as the *discontinuities* of the radiance function and how they are affected by phenomena such as shading, occlusion, and travel in free space. We will characterize how the radiance signal is transformed as light propagates and interacts with objects. This theoretical work will serve both as conceptual inspiration and will provide analytical formulas for other parts of our work.

Exploiting signal characteristics for rendering applications. We will develop practical algorithms for lighting simulation and real-time rendering. These algorithms will exploit knowledge of the characteristics of the radiance signal. Our acceleration strategies will rely on lower sampling rates when appropriate and perceptual masking due to high-frequency patterns.

Acquisition and characterization of scene properties. We will develop techniques to acquire real-world material properties. We will first rely on brute force to acquire and characterize ground-truth data. This will allow us to validate assumptions made in other projects, and also exploit these characteristics for faster acquisition using lower sampling rates and statistical characterization.

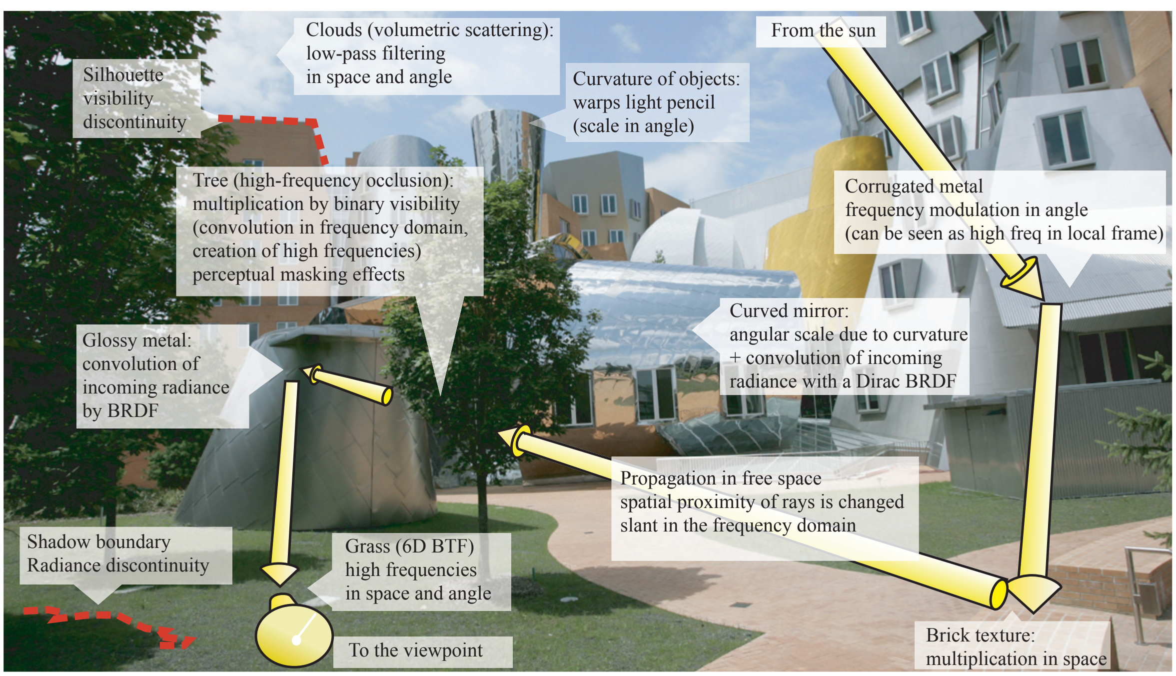

Figure 1: *Signal processing perspective on light transport.*

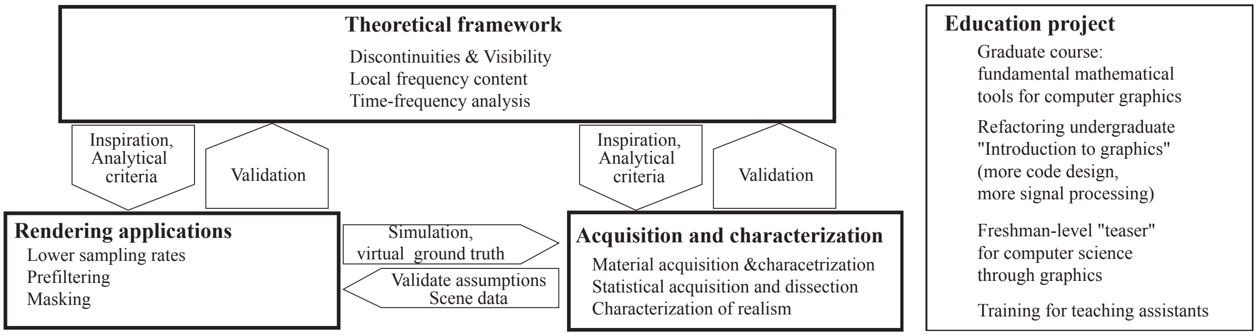

Figure 2: *Our research includes three interwoven fronts. They will all share a conceptual framework and focus on frequency content and discontinuities of the radiance function and the image. Our education project builds on firm fundamentals and implementation.*

This proposal builds on the PI's extensive contributions on the space of light rays and visibility [DDP96, DDP97b, DDP02, DDP97c, Dur99]. This will provides us with a framework for discontinuities as well as important tools to work in the space of light rays The application areas extend the PI's work on lighting simulation [DDP99, JDS $^+$01], real-time rendering [DDTP00, DDSD03, CD03, AKDS04, CD04] materials appearance [NDM04], and edge-preserving filtering [OCDD01, DD02, JDD03, ED04]. This work will provide firm signal-processing foundations for realistic image synthesis and material acquisition.

Beyond five years. We believe that beyond the five-year period, this project will open new areas of research at the interface between computer graphics and perception. Signal processing is crucial to understanding low-level vision as well as mid-level aspects such as texture perception and visual attention. The recent field of the statistics of natural images [SO01] studies regularities in visual stimuli and how they affect perception. This suggests that a solid signal-processing framework for light transport will provide important insights about ecological optics [BGG96], the study of how the characteristics of our visual world explain the mechanisms of our perception. We are convinced that important aspects of non-photorealistic pictorial styles can also be studied through signal processing. We believe that a better understanding of the signal-processing aspects of realistic visual phenomena and how our perception processes them is crucial for creating compelling depiction [Dur02].

Technological impact on society. Urban and architectural lighting is an important component of our daily environment. With the development of light sources with better power and spectral characteristics, a significant number of cities will rethink their urban lighting. Our interaction with lighting architects has shown that they want to use lighting simulation tools but are fustrated by the long rendering time and complexity of the currently-available software. Our lighting simulation methods (Section 3) will address these two points by greatly accelerating lighting simulation and recasting all algorithmic parameters in a unified framework affording a direct control over the time-quality tradeoff. We will work with students in the nearby architecture department to develop state-of-the-art simulation techniques well-suited to the needs of architects. Hollywood productions have similarly been slow at adopting global illumination, in part because the behavior of the algorithms is harder to predict. Industrial Lights and Magic is only starting to use irradiance caching for one-bounce indirect diffuse lighting, not even full global illumination. Our work will provide better control of the quality-time tradeoff and substantial acceleration (Section 3). In addition, discussions with industry players in realistic rendering indicate that material acquisition is one of their top priorities. The work we describe in Section 4 will dramatically facilitate this process. Games and simulations will also benefit from the real-time rendering techniques enabled by our research for display acceleration (Section 3.2) and pre-computed light transport (Section 3.1.4). As members of MIT's Oxygen alliance, we are in contact with partners such as Hewlett Packard and Nokia who are interested in the applications of our real-time and realistic rendering work. We have developed a close relationship with the MIT Deshpande Center which promotes innovation and provides invaluable contacts and opportunities for commercialization of technology from research labs.

1 Introduction: Relevance of signal processing for realistic imagery.

Light in a scene is transported, occluded and filtered by complex materials. By the time it reaches our eyes, the visual array is an intricate function, and analyzing it or reproducing it is a task that is challenging at best.

Low-level human vision is about transient image processing. What we know of low-level human vision suggests that signal processing is a crucial aspect of scene perception. Our low-level vision appears to perform local frequency analysis and edge detection using receptor fields similar to Gabor wavelets. Non-linearities such as lateral inhibitions emphasize transients such as visual boundaries. We have the ability to both characterize the stationary part of a stimulus and to detect and focus on transients.

Realistic graphics remains costly and the time-quality tradeoff is difficult to control. The equations of light transport are well-understood, but despite efficient heuristics and acceleration, realistic rendering remains costly. Phenomena such as interreflection or complex materials are costly because they involve integration over large domains. This is frustrating because their effect is subtle, although crucial. More troublesome, rendering time and quality can be hard to control. Algorithms are sensitive to internal parameters that vary from scene to scene. The frequency content of visual phenomena is implicitly exploited through efficient heuristics. For example, indirect diffuse illumination varies slowly because it integrates the illumination over the visible hemisphere. Techniques such as irradiance caching [WRC88] compute this component only at sparse sample points and interpolate. In signal-processing terms, the integral of diffuse lighting is a low-pass filter. This leads to a band-limited signal that can be sampled at a lower rate and reconstructed. Irradiance caching is unfortunately limited to diffuse reflection, but we believe that a general signal-processing perspective on light transport will allow us to generalize it. In summary, the state of the art relies on a skillful yet hard-to-control blend of brute force and clever heuristics. We need a better understanding of the behavior of light transport. We need techniques that relate phenomena and their simulation to final image quality.

Acquisition is challenging, yet our vision can recognize materials under unknown lighting. For inverse problems such as scene property acquisition, the space of possible solutions and the ambiguity present in images often make the problem intractable without severe assumptions or additional input. In addition, the dimensionality can be large, making direct costly in time and storage. This is frustrating because our visual system can, for example, recognize materials from single images under unknown illumination. The Holy Grail is to have acquisition systems that can capture material appearance with the same ease. As pointed out by vision research, the solution of this ill-posed problem must rely on knowledge about the characteristics of these visual components and their image projection, e.g. [SO01, DAW01, FP99, LZW03, TFA03]. A better understanding of these characteristics is needed.

Summary: We need to characterize and exploit the frequency content and discontinuities of different components of light transport. We have seen that for image synthesis and scene property acquisition, it is important to consider the characteristics of the signals involved, be they in the scene or image domains. Frequency content and discontinuities must be characterized and exploited.

2 Theoretical signal-processing framework for light transport

Light transport can be described as a big linear system that takes as input the emitted light and outputs a radiance distribution. However, characterizing this system is challenging. The domain is high-dimensional, and the signal is not stationary and presents discontinuities. We cannot rely directly on Fourier analysis because it focuses on space-invariant systems. Instead, we propose to analyze the signals in terms of *local* frequency content and *discontinuities*. This theoretical work will both serve as inspiration and provide analytical criteria for rendering applications. However, a number of applications only share conceptual ideas with this theory, and we will be able to develop them independently.

Background on radiance and the domain of light rays. We first present background on the domain of the light signal. This background will prove important to understanding our proposed theory in the rest of the section. The 5D radiance (or plenoptic [AB91]) function is defined on the domain of rays. A ray $\vec{r}$ is defined by its origin (3D) and its direction (2D on S^2). One spatial direction is special; It is collinear to the propagation direction. Recall that we want to characterize signal modification when light is propagated. We therefore consider the 4D slice of radiance orthogonal to the propagation direction. We call such a 4D slice a *local light field*. In the absence of scattering, rays propagate in straight lines, which puts an important emphasis on line geometry. In 3D, the space of lines is 4D and corresponds to the four degrees of freedom of local light fields. In our research, we will make some of the preliminary studies in 2D, since previous studies of light transport in flatland have proved fruitful [Hec92, PV96, ORDP96, RH01a]. In 2D, line space is two-dimensional and ray space is three-dimensional.

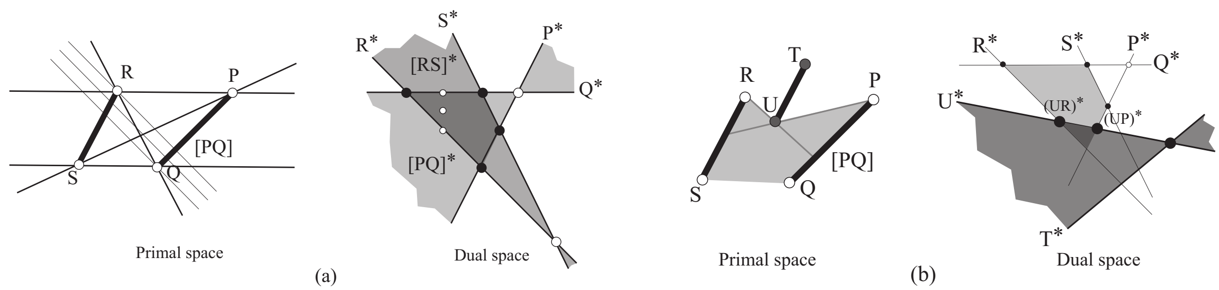

Figure 3: *(a) The set of lines piercing two segments is a quadrilateral in the dual. The four dual vertices correspond to the four lines joining pairs of segment endpoints. (b) Occlusion by $[TU]$ corresponds to intersection in the dual.*

Background on line duality. Since line space is hard to comprehend, we will use a dual space where lines are mapped to points. In 2D, a line $y = ax + b$ can be mapped to a point (a, b) (see Fig. 3). Formally, a projective space must be used to account for vertical lines [Sto91]. We note objects in the dual space with a star *. In the dual, parallel lines become vertically-aligned points. The set of lines going through a point becomes a line (e.g. R^* in Fig. 3). The lines going through a segment in the primal become a double wedge ($[PQ]^*$ and $[RS]^*$ in Fig. 3). In the case where the segments are vertical, the wedge degenerates and its boundaries are parallel. This point will prove important in Section 2.2. The lines piercing the two segments correspond to the quadrilateral intersection of the two wedges. Similar parameter spaces can be used in 3D. We will rely on the 2-plane parameterization [LH96, GGSC96]. Most of the notions we discuss in Section 2.2 in 2D translate to 3D using an affine transform of this parameterization. We avoid the singularities of 3D line space because we only consider local neighborhoods.

2.1 Discontinuities of the radiance function and 3D visibility

We can now turn to the study of discontinuities involved in light transport. Most radiance discontinuities are due to *visibility* changes. Visibility problems are central in image synthesis. In previous research, we have shown that line-space studies of visibility yield crucial insights. Consider the example of Fig. 3, right. Occlusion between the two segments $[RS]$ and $[PQ]$ by segment $[TU]$ corresponds in the dual to removing the wedge $[TU]^*$ from the quadrilateral of the two segments. The vertices of the resulting polygon correspond to lines such as (RU) that are visibility discontinuities, called *visibility events*. For example, if $[RS]$ is a light source, then $(RU)^*$ is a penumbra boundary. We have generalized these ideas and developed a framework, the *3D visibility complex*, which describes all visibility properties of a three dimensional scene [DDP96, DDP97b, DDP02, Dur99].

Our theoretical study of the 3D visibility complex has resulted in important insights and on the implementation of a simplified version, the *visibility skeleton* [DDP97c, DDP97a], which we have applied to lighting simulation using the radiosity method [DDP99]. The skeleton permits fast visibility queries and it permits to match the elements shape to discontinuities of radiance. The theoretical insights afforded by the 3D visibility complex also allowed us to develop a novel visibility preprocessing approach [Dur99, DDTP00].

We are continuing our research on visibility in line space, in particular on temporal aspects and on the topological properties of the visibility complex. Space prevents us from elaborating on these aspects, and we instead choose to focus on frequency issues. When the scale and number of discontinuities goes beyond an appropriate resolution, we need to be able to compute band-limited representation and assess the frequency content due to occlusion. For example, occlusions created by the leaves of a distant tree should be treated in a statistical or spectral manner, as discussed below.

2.2 Local frequency analysis of light transport

Related work. We build on seminal work on texture antialiasing [Hec89], light-field analysis [IMG00, CTCS00], and local shading [RH01c]. These approaches use Fourier analysis to assess how frequency content is affected by various visual phenomena. They have led to crucial practical tools for texture mapping [Hec89] and new diffuse lighting representations [RH01b]. They have also benefited inverse problems such as shape from texture [Alo86, MR97] and material/lighting recovery [RH01c]. They use Fourier analysis in a broad sense and do not strictly require stationarity. For example, Heckbert [Hec89] locally derives filtering bandwidth based on local perspective mapping. Since our main goal is to derive sampling requirements and interpolation for light transport, we will also assume local stationarity. We will however later challenge this assumption and use *time-frequency analysis*, as we discuss at the end of the section.

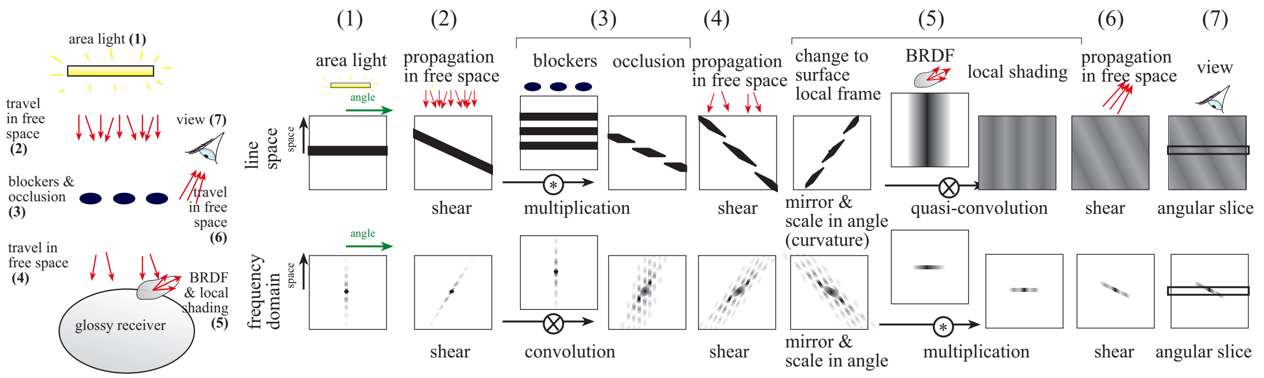

Figure 4: *Illustration of our proposed framework for light transport in the frequency domain. For a simple scene, we show the linear operators that affect frequency content in the space and angular domain.*

A flavor of time-frequency analysis, wavelets [Mal98], has received much attention in graphics and image analysis. *Discrete* wavelets have been commonly used as basis functions and analysis filters. Wavelets have not been used in the context we suggest, which is to characterize light transport as a linear system. Of course, wavelet applications exploit similar characteristics but in a data-driven manner and in the context of projection methods. In contrast, we will characterize the frequency response of the *equations* of light transport in the spirit of linear systems.

Shinya et al.'s pencil tracing [STN87] is similar to our goal. They characterize the spread (or footprint) of a 4D bundle of rays for adaptive ray tracing. However, they only consider propagation in free space, reflection and refraction. They do not consider the frequency content of radiance. Furthermore, we will study additional phenomena such as occlusion and local shading.

Local frequency content. We are interested in local signal content, so we only consider a local pencil of rays. In particular, we do not consider the full 2-sphere $S^{\in}$. We can decompose the four degrees of freedom into two time 2D: the 2D *spatial* dimensions on the plane orthogonal to the ray direction, and the 2D *angular* dimensions. This is important because many previous approaches are restricted to spatial or angular variations. We use a parameterization [STN87] similar to the two-plane parameterization but which better separates spatial and angular domains. It is the equivalent of the 2D duality presented above, and which we use to illustrate our ideas.

We have derived the qualitative aspects of fundamental processes such as transport in free space, occlusion, and shading. Some of these phenomena have already received a signal-processing treatment in the literature, often limited to either angular or spatial variations. We are now in the process of formalizing analytical formulas and studying more phenomena. In the rest of this section, we discuss the frequency aspects of the main mechanisms of light transport. We consider the scene in Fig. 4 as inspiration and follow light as it travels from the lamp. It first travels through free space, is occluded by a number of blockers, and is reflected by a glossy smooth object before reaching the eye. Each of these steps corresponds to precise mechanisms with important modifications of the frequency content of radiance.

2.2.1 Input signals: lamps and projectors.

The input signal is the radiance emitted by light sources. A point light is the Cartesian product of a constant in angle and a Dirac in space. An extended light sourced such as the one in Fig. 4d is a box in space times a constant in angle. Sources such as a projector consist of a Dirac in space and an image signal in the angular domain.

2.2.2 Travel in free space: shear of ray space

After leaving the source, light travels in free space (Fig. 4 step 2). Travel in free space is a crucial operation because the angular variation also turns into spatial variation. Consider a slide projector: At the source, we have a Dirac in space and the image in the angular domain. At the receiver, the signal of the image is present in both space and angle.

This aspect can be better described in a dual space. In Fig. 5, we show spatial propagation between two neighborhoods, O and P. We denote rays in the neighborhood of O by their supporting line and we add a prime for rays after propagation (OP propagates to OP'). The duality is the same as in Fig 3, but for better parameterization, we consider AB vertical so that both AB and CD are mapped to the line at infinity. The points on AB map to parallel lines in the dual space. The points on CD have duals that are also parallel but with a different direction. Rays that are spatially adjacent

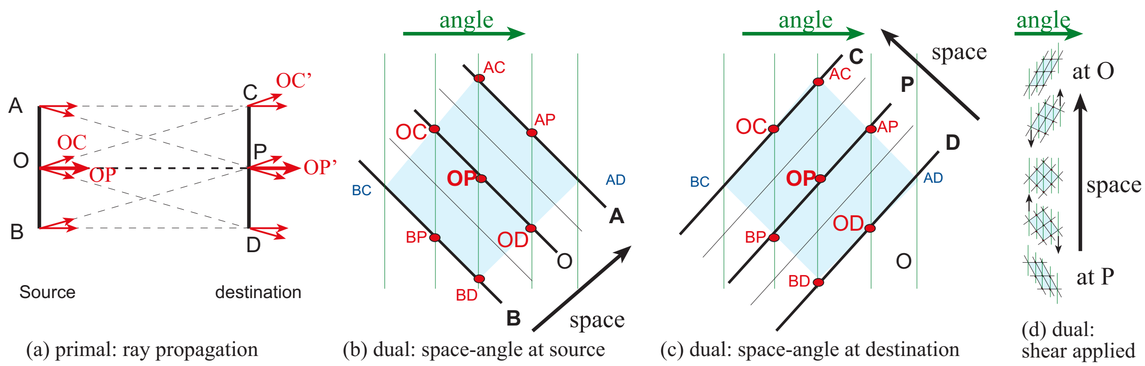

Figure 5: *Ray propagation changes the notion of spatial adjacency. Rays such as OC and OP are not adjacent anymore after propagation to OC' and OP'. This can be visualized in two ways: The direction corresponding to space changes (b and c) or the dual space is sheared vertically as in (d).*

at O propagate to rays that are not adjacent anymore (e.g. OC and OP). This is because the notion of spatial dimension has changed. At O, the space dimension goes along segment AB in primal space, which corresponds to the direction orthogonal to $A^\star$ in the dual. In contrast, space at P goes along CD and corresponds to the direction orthogonal to $D^\star$ in the dual. Consider the rays originating at O. We see in Fig. 5(b) that their duals are parallel to $A^\star$ and that they share the same spatial coordinate. After propagation, the change of spatial dimension puts them at different spatial locations (Fig. 5(c)). However, the angular dimension is not altered; The direction of rays does not change during propagation.

The change of spatial dimension during propagation is a linear shear operation in line space (Fig. 5(d)). The frequency content of the local light field is also sheared by the same amount but in the perpendicular direction ([Bra00], p. 333). The length of propagation corresponds directly to the amount of shearing. The longer the travel, the more pronounced the shear, and therefore the more spatial neighborhood is distorted. This explains the slide projector effect. This also corresponds to the important notion of *footprint* exploited by ray-tracing techniques such as cone-tracing [Ama84], pencil tracing [STN87] or ray differentials [Ige99]. The further away a ray bundle goes, the more pronounced the spatial variations. Our approach also generalizes Heckbert's texture mapping study [Hec89] and plenoptic sampling [CTCS00]. Both previous works only study the spatial content of the input signal (Lambertian assumption).

2.2.3 Visibility: creation of high frequencies by convolution in the frequency domain

In step 3 of Fig. 4, the rays are occluded by blockers. Occlusion is responsible for intricate frequency effects. As discussed above, these can be considered as discontinuities and visibility events. However, at a larger scale, a crucial contribution of our research will be to understand the effect of occlusion on the frequency content.

Fig. 6 illustrates the effect of visibility. At the location of occlusion, radiance is multiplied by the binary function of the blockers, where $\emptyset$ denotes occlusion. According to the multiplication theorem, this means that occlusion corresponds to a *convolution in the frequency domain*. At the location of occlusion, high frequencies are created only in the spatial domain since the blocker's binary occlusion is a constant in angle and the shape of the blockers in space (Fig. 6(b)). The frequency content is the Cartesian product of the 2D Fourier transform of the blocker's binary projection and a 2D Dirac in angle (Fig. 6(c)). However, after propagation, the shear linear transform described in the previous section results in high frequencies in both space and angle (Fig. 6(d) and (e)).

This perspective will allow us to recast and extend important results such as convolution shadows [SS98], feature-based visibility [SD95], and our convolution approach to volumetric visibility [DDTP00]. These approaches do not consider both space and angle and do not explicitly provide frequency modification information. We will also extend recent work on frequency content of shadows of microgeometry [RKB04] that only considers local interaction for one blocker's edge. This will allow us to take into account the effect of occlusion for plenoptic sampling [CTCS00].

2.2.4 Local interaction: quasi-convolution

In Fig. 4 step 5, the light reaches a smooth glossy object. Surface-light interactions are central to the appearance of materials. For these phenomena, we will build on Ramamoorthi and Hanrahan's work [RH01c] that shows that local illumination primarily corresponds to a convolution in angle. It is the integral of the incoming radiance mulitplied by the BRDF rotated to the surface normal. This rotation corresponds to the kernel shift in convolution. We will adapt

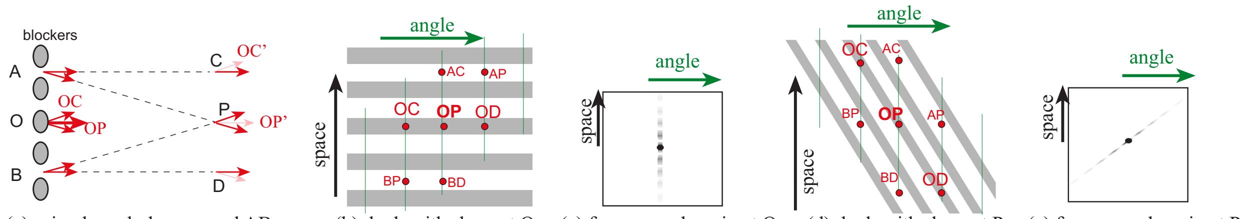

(a) primal: occluders around AB (b) dual: with shear at O (c) frequency domain at O (d) dual: with shear at P (e) frequency domain at P

Figure 6: *Occlusion around AB corresponds to multiplication by binary bands in the dual space. In (b-e) we apply the shear operations that make space and angle orthogonal for different locations. Around O, the binary visibility is constant along the angular dimension, while around P, high frequencies occur both in spatial and in angular domains.*

their work in the context of our line parameterization and will extend it to consider not only angular but also spatial content.

In a nutshell, the ray bundle needs to be rotated to the surface local frame and mirrored to account for reflection. We also need to correct for the distortion of the linear parameterization when the incident angle is not normal. The angular content is scaled by the curvature of the object, and radiance is then convolved with the 2D slice of the BRDF corresponding to the incident angle. Note that, in our parameterization, rotation is only approximated to a first order, in contrast to the spherical harmonics used by Ramamoorthi. This is the price we pay for a representation that naturally combines space and angle. However, since we are interested in local signal content, this approximation is acceptable.

This signal-processing perspective has important consequences for rendering as well as for material perception and acquisition. As pointed out by Ramamoorthi [RH01c], the outgoing signal is band-limited by the BRDF. This explains results by Dror et al. on illumination-independent material recognition [DAW01]. Moreover, in 3D, the angular scaling due to curvature is anisotropic according to the principal curvatures. This explains Fleming's hypothesis that our perception uses anisotropy in the image to deduce curvature anisotropy [Fle03].

2.2.5 Summary and other phenomena

As shown in Fig. 4 and described in this section, light undergoes a number of modifications as it propagates in a scene. Although the individual mechanisms are simple, their compound effect leads to complex signals. A number of domain distortions occur because of travel in free space (step 2, 4 and 6) and local curvature (step 5, left). The spectrum itself is affected by multiplication-convolution operations due to occlusion (convolution in frequency, creation of high frequencies) and shading (multiplication in frequency, band-limited filtering). As discussed through the section, our unified signal perspective affords important insights on rendering techniques as well as human and machine vision.

We will formalize the mechanisms described above and will derive analytical formulae. We will also study a number of other important phenomena which we only list informally: texture mapping (multiplication by a function of space), volumetric scattering (low-pass, e.g. [JMLH01, PA04]), bump mapping (frequency modulation), depth of field (low-pass), diffraction [Sta99]. Fig. 1 gives a flavor of these effects in a real scene. In the long term, we will also study the temporal dimension and motion blur effects.

2.3 Longer-term: Time-frequency analysis

So far, we have assumed local stationarity to analyze frequency content. It is important to challenge this assumption, evaluate where it breaks, and use more appropriate tools when necessary. In particular, note that local shading is not a true convolution because the BRDF slice and the incoming radiance vary over space. Similarly, occlusion boundaries correspond to non-stationary events.

To address these challenges, we will build on time-frequency analysis or transient signal processing [Fla99, FS98, Mal98], which studies systems that are linear but vary in time. In this case, the convolution theorem does not hold. The frequency spectrum can be characterized over time using different versions of the short-time-Fourier-transform, Gabor wavelets or power distributions similar to spectrograms used in speech analysis. In these representations, the stationary case corresponds to a spectrum that does not vary over time. Fortunately, for a number of practical cases, the spectrum varies slowly over time. These systems are called *underspread* [Mat00]. In this case, approximate convolution and multiplication theorems exist. The spread of a system, as well as the error due to approximate convolution theorems, can be calculated [Mat00]. We will apply these tools to evaluate and extend our work on local frequency analysis.

In the general case (normal spread), time-frequency analysis can still provide important information about light propagation. We will use new tools termed *input Wigner distribution (IWD)* and *output Wigner distribution (OWD)* [HM01, Mat00]. They characterize from where frequency content comes and where it goes in a linear time-varying system. From a light transport perspective, this means that we will be able to estimate the potential outgoing frequency content with the OWD, as well as which part of the incoming frequency content are important with the IWD. These two pieces of information are crucial for analysis and synthesis problems, and they closely relate to factorizations of BRDF [KM99, MAA01, RH02, LRR04] that separate components that depend on incoming and outgoing directions.

3 Exploiting signal characteristics for rendering applications

The theoretical concepts described above will lead to practical algorithms in lighting simulation and real-time rendering. There are two different ways we want to exploit signal characteristics:

Sparser sampling of low-frequency components. If we know that a component is low-frequency, use a low sampling rate and interpolate. A heuristic based on this approach is irradiance caching [WRC88].

Masking. High-frequency patterns can perceptually mask other artifacts. For example, shadow artifacts are harder to detect on high-frequency textures [FPSG97, BM98, Mys98].

Related work. Frequency or multiscale content has been exploited in the context of wavelet radiosity [HSA91, GSCH93, DDP99]. However, these approaches do not have the flexibility of sampling approaches. In Monte-Carlo integration, the integral rendering equation is approximated by finite sums of random paths. Heuristics such as irradiance caching [WRC88] or photon mapping [Jen01] implicitly take advantage of the frequency content of various components of light propagation; in our research, we will make this explicit and will generalize it. A number of techniques use notions related to bandwidth in a ray's neighborhood. They propagate a footprint to estimate adaptive refinement [Ama84, STN87, Ige99, SW01, CLF+03]. Our work is complementary and will exploit this notion to propagate frequency requirements. In addition, authors have exploited the frequency content of the image to make better use of rays [BM98, Mys98, RPG99, Kel01]. We will generalize this by considering intermediate radiance frequency content and by decoupling components of the image to isolate signals of different bandwidth.

3.1 Lower sampling rate

3.1.1 Acceleration techniques revisited and unified

A number of rendering techniques implicitly exploit knowledge of the required sampling rate. These techniques are extremely powerful, but they involve a number of parameters that may need to be adjusted for each scene. We will revisit these techniques and recast them in our signal-processing framework for light transport. Footprint approaches [Ama84, STN87, Ige99, SW01, CLF+03] correspond to propagation in free space (Section 2.2.2). Feature-based visibility [SD95] and our occlusion sweep [DDTP00] deal with occlusion (2.2.3) and travel (2.2.2). Texture prefiltering [Hec89] is a combination of travel in free space (2.2.2) and local-frame rotation (2.2.4). Masking-based global illumination [BM98, Mys98, RPG99] exploits multiplication by texture (2.2.5). Irradiance caching [WRC88]and photon mapping [Jen01] exploit the low-pass filtering due to diffuse shading (2.2.4). This list is not exhaustive and we will express all the criteria involved in these techniques and others into our framework. It will allow us to make all the parameters directly relevant to visual quality. Our ultimate goal is to have a single parameter that controls the quality-time tradeoff.

3.1.2 Glossy rendering

The general case of glossy materials remains costly to simulate. For each visible point, the shading equation has to be computed by integrating over the hemisphere using a large number of rays. Importance sampling of the environment [ARBJ03, ODJ04] or the BRDF [LRR04] dramatically improves convergence, but a couple of hundred of rays are still necessary for each pixel. The situation is frustrating because glossy reflection is essentially blurry. In our work, we will exploit this blurriness by trading sampling rate in the angular domain for sampling in the image.

A blurry highlight is costly to integrate because it has a wide angular support. Fortunately, it results in a blurry image (Fig. 7), and image sampling can be reduced and interpolated (Fig. 7). Reciprocally, a narrower highlight results in a sharper image that requires a higher image sampling rate. But since the angular spread of the highlight is narrower,

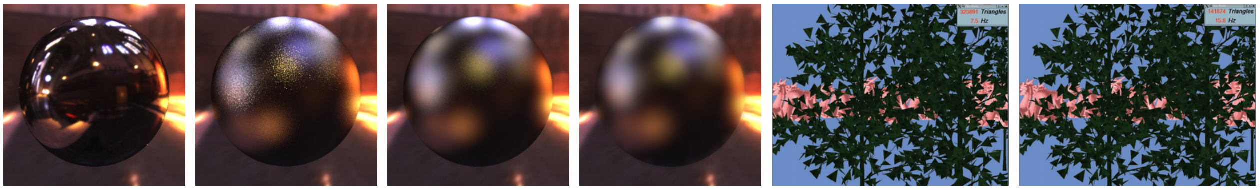

(a) 20 rays (b) 20 rays (c) 1280 rays (d) subsampled, same cost as (b) (e) fine LOD (f) coarse LOD

Figure 7: **BRDF rendering:** *We propose to trade image sampling for angular sampling (manual simulation using Photoshop). (a) 20 rays per pixel are enough for a shiny BRDF. (b) and (c) For a rough objects, more rays are needed because the specular lobe is wider. (d) Since the resulting image is blurry, we can compute reflection at every* 8×8 *pixel and interpolate. A layer mask was used to ensure a sharp silhouette.* **LOD and masking** *(e) Scene rendered with high-resolution 3D models. (f) Based on our masking estimation, we use coarser levels of details for the dragon. The complex pattern of tree leaves masks the artifacts.*

fewer rays need to be used for each pixel. This is a typical case of frequency vs. time (or angle) localization. For smooth functions such as Gaussians, the spread in space is inversely proportional to the bandwidth in frequency, and BRDF functions tend to behave like Gaussians. The validation of this assumption will be discussed in Section 4.1. We can expect to perform BRDF rendering at a given quality for a constant cost, with an image sampling rate that is inversely proportional to the ray (angular) sampling. This idea exploits similar properties as frequency-based ray tracing [BM95] or hierarchical Monte-Carlo ray tracing [Kel01]. However, they work in a data-driven fashion, based on image values, while we deduce blurriness from the shading equation and the material parameters.

This idea is simple to exploit for environment-map rendering of smooth objects without occlusion. We must take into account the BRDF frequency content as well as the curvature of the object (Section 2.2.4). Exploiting the convolution interpretation of local shading provides the required sampling rate. Image discontinuities caused by occluding contours must be respected by interpolation (see below).

The general case (3D environment) is more involved because of high spatial frequencies or discontinuities in the illumination. Fortunately, point lights are excluded from hemispherical Monte-Carlo integration (they have measure null) and the remaining illumination is C1-continuous in the spatial domain [Arv94, DF94]. We must estimate the spatial radiance variation due to parallax of nearby objects and soft shadows. This is related to the derivatives of irradiance as studied by Ward and Heckbert [WH92], Arvo [Arv94], or in our recent contribution [AKDS04]. In the case of extended sources, variations due to occlusion cannot be ignored. We will treat this component separately and will compute a simplified version of Arvo's Jacobian [Arv94] similar to our shadow refinement oracle [DDP99].

We expect the visual quality to increase even faster than the numerical accuracy, because we will enforce the appropriate amount of smoothness in the image (Fig. 7(d) is smoother than Fig. 7(c)). We will use perceptual metrics to further assess this hypothesis. This approach will be an important step for time-quality control, since it will use the image-angular tradeoff to achieve roughly constant quality for a given time.

Image interpolation When using lower sampling rates for rendering, we need to interpolate at the appropriate bandwidth while respecting image discontinuities. We will build on the bilateral filter [TM98, DD02] to address these issues. Similar to McCool's use of anisotropic diffusion [McC99], we will use auxiliary maps such as z-buffer and object ID to influence the edge-preserving notion of similarity in the bilateral filter. We will use deferred shading in order to use information about the visible surfaces to drive adaptive sampling. The image will be rendered at full resolution, but some components of shading will be sub-sampled. In contrast to traditional adaptive schemes, each pixel will eventually be considered for shading. However, if we find that a given component of the pixel's shading can be interpolated from higher levels of the image, we will save computation. Interpolation is difficult at object silhouettes, especially for the background object since occluded information is missing. We will build on layered-depth images [SGwHS98] to provide information about the shading of occluded parts close to visible pixels.

Extension to bump mapping, soft shadows, depth of field, and motion blur. For bumped-map blurry BRDFS, we can still exploit limited bandwidth in the image if we restrict ourselves to similar normals. This simply adds a term based on normals to the bilateral similarity. In addition, bump mapping creates high frequencies in the image that cause masking effects. We can assess them and exploit them to subsample in angle. We will also extend our technique to soft shadows (area of the source vs. penumbra sharpness) and depth-of-field effects (size of aperture vs. depth of field).

This will require appropiate bandwidth estimates for these effects, proper decoupling of the visual components (e.g. interpolate only the incoming light, before texture mapping) and appropriate similarity terms for interpolation.

3.1.3 Frequency-based Metropolis light transport

Frequency-based path mutation. The Metropolis light transport algorithm developed by Veach and Guibas [VG97] is a powerful unbiased method to treat complex light paths. It exploits coherence in path space by mutating light paths that contribute strongly to the image. This way, even if a path has low probability, once the algorithm has found it, it correctly explores the complete neighborhood. Although the idea is appealing, the technique is difficult to control because path-mutation rules must be implemented and they strongly influence the efficiency of the algorithm. Once a "good" path is found, one has to decide how to mutate the path and by how much.

We will use our frequency framework to estimate the local complexity of the light-transport operator. We will implement the analytical criteria discussed in Section 2 to compute the frequency effect of each step in a light path. This will provide us with crucial information to drive mutation. The distance of mutation will be driven by spatial and angular local bandwidth. We are hopeful that it will greatly improve the stability and efficiency of Metropolis light transport, at the cost of bias. However, bias is not an important issue for most visual applications.

Visibility preprocessing. Visibility will raise the biggest challenges for frequency evaluation. We need an efficient data structure to query the frequency content of a set of blockers that occlude a ray pencil. We will adapt volumetric techniques such as feature-based visibility [SD95, SS96] and our occlusion sweep [DDTP00]. For each cell of a spatial hierarchy, we will store directional frequency information for a small set of directions (6 to 20). For each leaf node and each direction, we will compute the projection of the blockers and their 2D Fourier transform, for which we will encode a simple summary (subsampled radial power spectrum for example). For higher nodes in the spatial hierarchy, we will use combination operators based on the travel in free space mechanism (2.2.2). Note, however, that this approach does not take into account the phase information and only provides an approximation. Correlation between successive blockers is not taken into account. In addition, the angular resolution will need to be low for storage reasons. However, we believe that the information provided by this data structure will be a large improvement to control algorithms such as Metropolis light transport. Furthermore, this frequency visibility information can be used to drive the construction of the hierarchy in the context of a cost-driven construction [ABCC03]

3.1.4 Precomputed radiance transfer

Precomputed radiance transfer (PRT) preprocesses the appearance of an object under any light direction [SKS02]. Light is assumed to come from an infinite distance (no spatial variation). For a point on the object, we store how light is reflected as a function of incoming light. Incoming and outgoing light are encoded on the sphere and usually represented using spherical harmonics (SH). This is a typical case of frequency-based light transport, since SH correspond to Fourier bases on the sphere. The object is sampled spatially and a transfer matrix is computed for each sample.

Spatial vs. angular sampling PRT methods focus on the angular content of radiance. Incoming light is assumed to come from infinity and is usually sampled at the center of an object. Sloan et al. [SKS02] consider spatially-varying lighting by using multiple samples, but they do not provide criteria for spatial sampling. Similarly, all methods compute the transfer function at points on the object but the actual required sampling rate is unclear. As a first step, we have recently studied the spatial variations of incoming radiance projected onto SH [AKDS04]. We have computed the gradient of coefficients and performed extrapolation and interpolation based on a Taylor series expansion. In our research, we will use our theoretical results to derive sampling criteria both in space and angle.

Non-linear approximation of transport kernels for full-frequency PRT PRT approaches use projection on linear basis, potentially truncated for non-linear approximations [NRH03]. Handling high-frequency effects requires large sets of basis functions. We will use mixtures of Gaussians to provide more compact representations of the transport function and smoother interpolation. Consider a given view direction. Instead of projecting the outgoing light onto linear basis functions, we will optimize a small set of Gaussians to approximate the data. Gaussians have several advantages. They are compact and can represent arbitrary frequency content by varying their variance. They can be integrated quickly against an environment map using mipmaps. In addition, when one varies the viewpoint, Gaussians provide more natural interpolation of the outgoing-light kernel, avoiding the fade-in fade-out artifacts of linear bases.

We have realized a preliminary implementation of this idea, using a single box function instead of Gaussians. The results are promising but the performances are not optimal yet because boxes are not as smooth as Gaussians and are not as efficient as mipmaps. In addition, we must develop an interpolation technique for multiple Gaussians. For adjacent view directions, we must match the closest Gaussians to ensure that interpolation will be meaningful.

3.2 Visibility and masking for real-time and offline rendering

As discussed above, occlusion creates discontinuities and high frequencies. In our recent work, we have developed a number of techniques for fast exact or conservative visibility computation. We must now take into account occlusion by complex blockers such as trees.

Real-time rendering acceleration using partial visibility and masking Acceleration algorithms such as visibility culling or levels of detail (LOD) are crucial to reach high frame rate. Visibility culling permits the rapid elimination of completely hidden objects, while levels of detail replace distant objects with coarser models. Visibility is usually a binary decision: each object is determined visible or not. In contrast, we propose to detect and take advantage of *partial* visibility to improve levels of detail selection, as suggested by Andújar et al. [ASVNB00]. However, a simple partial visibility information such as the percentage of occluded area is hardly enough. *Perceptual masking* must be accounted for: a complex pattern overlaid on a signal can *mask* it [FPSG97, BM98, Mys98]. Taking into account the frequency content of the occluding pattern is essential for appropriate handling of partial visibility.

We are developing computational tools to estimate masking from the viewpoint. We use an occlusion mask [ZMHH97] and propose three approaches to estimate masking from the map: morphological dilation-erosion operators to estimate the radius of the locally-biggest hole; a non-linear construction of a pyramids of the map that compiles a notion of hole size; and a frequency approach using DCT coefficients. Fig. 7 (c-d) shows a prototype implementation using the DCT frequency content. We will compare the three approaches and integrate them in a real-time walkthrough system, using perceptual models of LOD error [LH01]. While we describe our masking technique in the context of real-time rendering, it also has applications to offline lighting simulation.

Volumetric visibility and masking The technique we just described considers masking from the current viewpoint using image-space computation. We will extend this work to precompute masking information from a volumetric perspective. This will relieve load at runtime for real-time applications. We will use a strategy similar to that described in Section 3.1.3 for Metropolis light transport. The data structure can then be used in two ways. We can precompute cell-to-cell masking information and store the amount of masking for each pair in the spirit of from-region visibility. We can also query the data structure online by casting a few rays to get masking information about moving objects.

4 Acquisition and characterization of scene properties

In the last two sections, we have presented projects that derive or make assumptions on the characteristics of real signals. It is crucial to validate these characteristics and compare them against real data. In addition, the acquisition of data such as material properties has strong applications. Knowledge of the characteristics of such signals will allow us to design faster acquisition devices that perform appropriate sampling and take advantage of limited bandwidth and regularities. As a first step, we will focus on material data.

4.1 BRDF acquisition

High-resolution acquisition of anisotropic BRDFs Unfortunately, material property data such as the BRDF are difficult to get. Most available measurements have a relatively low number of samples (less than 10,000 for a 4D function). A notable exception is the dataset by our former post-doc and colleague Wojciech Matusik [MPBM03a]. He has measured more than a hundred real materials at a high sampling rate (4M samples), using a design inspired by Marschner's approach [MWL$^+$99]. In this setup, a sphere is photographed under a given illumination direction. The normal variation on the sphere allows each image to capture a 2D slice of the BRDF. Unfortunately, the sphere design limits acquisition to materials that can take a spherical shape and to isotropic BRDFs.

We are modifying this procedure to capture full BRDFs for anisotropic materials (Fig. 8). A cylinder is replacing the sphere, which will facilitate wrapping materials around it. The cylinder provides only one degree of freedom for the normal, but we use the second dimension to capture anisotropy; A large number (30 to 40) of stripes of the material cut at different angles will be wrapped onto the cylinder. The stripes will provide the rotation of the material with respect to the incident angle, the crucial aspect of anisotropy. In contrast to the sphere case, we will have to not only rotate the source along a 1D circle, but also rotate the object to acquire the full 4D data.

We are currently building the device and we are hoping to start the measurement in the next three months. We have estimated that each material will take about two days to acquire and will result in about 180 million samples. Although the acquisition time is significant, this will be the first time that high-resolution anisotropic BRDF data are available. We are committed to releasing the data to the scientific community to promote research on material properties.

Experimental validation of BRDF models A variety of BRDF models have been proposed. However, validation of these models have been scarce due to the lack of high-resolution data. Building on our dataset, we will evaluate the performance of popular BRDF models. While previous work in BRDF modeling have validated their models with some measurements, our work will be the first to quantitatively compare different models based on a sizable dataset. We have performed a preliminary study on the isotropic materials acquired by Matusik [MPBM03a]. Our first finding highlights the profound difference between two popular formulations of the specular lobe: 1. defined around the mirror direction ($V \cdot R$), and 2. defined by the half-vector ($H \cdot N$), where V is the view, N is the normal, and R is the mirror direction. Most previous discussions of reflectance have focused on the intensity at grazing angles. We found that the shape of the lobe also has important implications. Our experimental data show that the half-angle $H \cdot N$ lobe is substantially more faithful than the mirror direction lobe, which is confirmed visually in rendered images [NDM04]. This finding is important, because most researchers currently use the Lafortune model [LFTG97] to fit their data, although it is built on the $V \cdot R$ lobe. Using an $H \cdot N$ lobe would significantly improve their results.

We will pursue this experimental validation, assessing the importance of different components such as Fresnel, masking, or shadowing terms. We will validate anisotropic models once we acquire the data. We will compute the frequency content of the BRDF data to evaluate the assumptions made in our proposed glossy-rendering technique (Section 3.1.2). We will also further study the error and uncertainty present in the measurements to establish our results on a firm basis.

Efficient acquisition that exploits known characteristics of BRDF signal The estimated acquisition time of two days for anisotropic materials is acceptable for scientific research, but not for practical applications. We plan to exploit the characteristics of the BRDF signal as measured using our brute-force approach to derive more efficient acquisition techniques. First, we want to compute the frequency content along the anisotropy dimension for typical materials. We expect the signal to be smooth along this dimension, thereby allowing for coarser sampling rates. In general, we want to exploit the local frequency content of BRDFs to adapt sampling rates. This project is related to Matusik's recent work in which he uses wavelets to derive efficient material acquisition [MPBM03b]. We want to extend this work in several ways. First, he assumes that BRDFs are measured point-by-point, while efficient methods rely on 2D slices observed by a camera (see above). Second, he considers that all samples are perfectly accurate. We will address these limitations and will derive optimal slices of the BRDF based on our measured data and their signal characteristics. We will take into account measurement uncertainty to guarantee high-quality acquisition.

4.2 Statistical material acquisition

The above strategies rely on slow measurements of the BRDF in controlled environments. In the long term, we will devise acquisition techniques that work in uncontrolled environments with fewer inputs. They will rely on knowledge of light transport effects as well as a-priori knowledge about the characteristics of the material property's signal.

Capture of materials from a single image under unknown illumination Humans have the fascinating ability to recognize materials under unkown illumination, despite the fundamental ambiguities present in the image. As shown by Ramamoorthi et al. [RH01c] and our framework (Section 2.2.4), inverse lighting is a deconvolution problem with an ambiguity on illumination and BRDF. Inspired by recent work on illumination-independent material recognition [DAW01] we believe that we can resolve the deconvolution ambiguity using a-priori knowledge about the illumination. Natural illumination statsitics have been characterized numerically [DLAW01] and our framework will provide further insights on their frequency properties. The BRDF measurements and characterization described above will provide

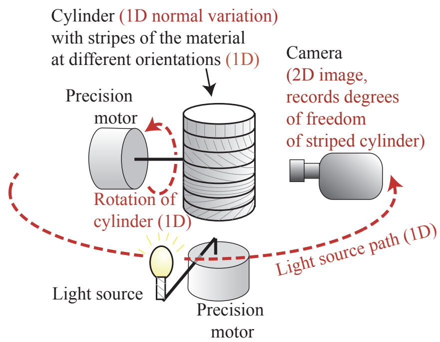

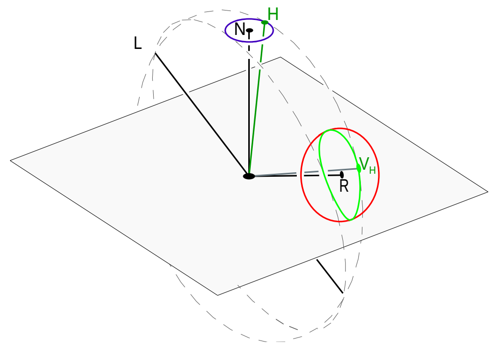

Figure 8: **Left: Anisotropic BRDF acquisition.** *An image captures variations in normal and angle of anisotropy. The light and the cylinder rotate to provide the remaining two degrees of freedom of a BRDF.* **Specular lobe** *The $V \cdot R$ (mirror direction) lobe compared with the $H \cdot N$ (half-angle). Note how the $H \cdot N$ lobe in green is narrower than the $V \cdot R$ lobe in red. We showed that real BRDF tend to have narrow lobes at grazing angles, validating the $H \cdot N$ lobe.*

information about the subspace of possible BRDFs. Equipped with this a-priori knowledge, we can greatly reduce ambiguity in material acquisition from images under unknown illumination.

Fur, hair and complex BTF acquisition We want to develop techniques that can capture the appearance of complex materials such as fur with a small number of images under natural lighting. We note that the appearance of these materials is precisely characterized by complexity, and that we should therefore capture complexity at a mid-level rather than individual samples of the function. Following the tradition of statistical analytical models of BRDFs, e.g. [CT82, HTSG91], we will use our framework to study analytically the lighting phenomena that occurs at a meso-scale for a complex BTF: occlusion, shading, multiple scattering, shadowing, etc. This will provide us with important insights about the frequency content in both the spatial and angular domains. In particular, occlusion between the geometry makes BTFs quite different from spatially-varying BRDFs: BTFs contain high frequencies in the angular domain. We will first restrict our attention to fur [Gol97] to narrow down the microgeometry and BRDF [MJC⁺03]. We are hopeful that the frequency analysis will reveal that different components of light transport and different characteristics of the fur material will result in distinct observed frequency effects in images. Video images will undoubtably be necessary to capture micro-parallax effects and high frequencies due to occlusion. We will then generalize this work to other complex materials such as hair, grass, and fabric.

5 Summary of proposed research and plans beyond five years

We have highlighted the strong potential a signal-processing perspective on global light transport has to impact the field of computer graphics.

Theoretical framework for global light transport. We have described how light transport involves spectrum distortion due to transport and curvature, as well as convolutions in the angular domain (local shading) and convolutions in the frequency domain due to occlusion. In summary, shading by blurry BRDFs simplifies the signal while occlusion makes it more complex. In our research, we will formalize this framework and study other visual phenomena such as texture mapping, and participating media. We will also evaluate when the local stationarity assumption is valid and use tools from time-frequency analysis to extend our work to space-varying mechanisms.

Exploiting signal characteristics for rendering applications. We will exploit knowledge of the frequency content of mechanisms of light transport in two ways. If we know that the signal is low-frequency, we will sample sparsely and interpolate. If the signal has high frequencies, we will exploit perceptual masking and the reduced sensitivity of the human visual system. In other words, if there is coherence, we will exploit it, and if there is no coherence at all, the human visual system will not notice artifacts. We will exploit these strategies for lighting simulation and real-time rendering.

Acquisition and characterization of scene properties. Material measurement is an important application in its own, but it will also allow us to validate results from our theoretical and rendering research. In addition, the analysis of typical real-world materials will allow us to develop new techniques for faster material acquisition.

The proposed research will establish strong signal-processing fundamentals to realistic graphics and will enable faster and easier creation of high-quality computer graphics imagery.

Beyond five years. The potential of signal processing for computer graphics go beyond the five-year scope of the proposed research. Perception mechanisms must be studied further and exploited. Non-photorealistic styles also rely on frequency content (e.g. Impressionists, Pointillists) and discontinuities (line drawing). The signal-processing-level difference and similarities of realistic and non-realistics depictions must contain important insights about the way we perceive and depict the world. In the long term, our goal is to characterize realistic images from a signal-processing standpoint [FL03, CH03] and to design a "realism filter" that takes an image and makes it look more realistic.

Our long-term philosophy is that all images have *constant visual complexity* because our visual system can only cope with a given visual complexity. The bandwidth of our perception is deceivingly limited, e.g. [Mil56]. Therefore, all scenes and phenomena should all take the same computation time to render. The angular-image tradeoff discussed in Section 3.1.2 is a first step towards this goal. Our statistical acquisition of fur is another step. In general, when a phenomenon becomes too complex, our perception abstracts it and "summarizes" it at a higher level. This is for example the case of texture information. If we find the appropriate way to evaluate the notion of complexity and to summarize a phenomenon, we can dramatically improve image synthesis and acqusition.

In the very long term, we want create a multi-disciplinary center or consortium on pictorial sciences. Inspired by the development of interdisciplinary linguistic and vision communities, we want to promote a cross-disciplinary approach to pictures to bring together computer scientists, psychologists, art historians and imaging professionals.

6 Education and broader impact

We will focus on two complementary educational goals: depth in the fundamentals of computer graphics and interdisciplinary breadth. We are introducing a new course to implement this vision and are redesigning the traditional "introduction to computer graphics" undergraduate class.

Undergraduate teaching: hands-on introduction to graphics. We believe that computer graphics can motivate students for computer science. In our teaching career, research in computer graphics has provided a way to attract students' attention and to anchor fundamental notions of computer science with examples drawn from their interests in games, movies, or simulation. We are redesigning our introduction to computer graphics course in order to make it more attractive, more hands on, and more accessible to younger students. The course will have a series of assignments for students to experiment with important concepts of computer graphics and to build a simple ray tracer from scratch. The course will culminate with a final project and a rendering competition to stimulate students' creativity. We are also committed to making our slides available to the community through MIT's OpenCourseWare initiative. Even before this initiative, we have put our current slides, and scholars at a number of institution have asked to use them for their teaching. See http://graphics.lcs.mit.edu/classes/6.837/F03/.

Graduate course: mathematical and numerical fundamentals. Together with colleague Jovan Popović, we introduce this year a new graduate and upperclassmen course in the core curriculum. This course introduces fundamental computational and mathematical tools for computer graphics. We cover techniques such as Monte-Carlo integration, finite elements, light transport, signal processing, and perception models. While we study applications in computer graphics, we focus on the underlying tools, which will make the course appealing to a broad audience. We firmly believe that strong mathematical fundamentals in these areas are crucial for students in many fields of computer science such as robotics, vision, and computer graphics.

Teaser activity for first-year students. Throughout January, MIT students can attend non-traditional courses. We will develop a one-week "teaser" for computer science through computer graphics, and provide them with a fun hands-on experience in computer science. We hope that this teaser will inspire a large number of students to do their undergraduate degree in computer science. Constrained and stimulated by the heavy research load in graphics in January, we plan to also use this teaser as an opportunity to involve upperclassmen and Master's of Engineering student in teaching activities.

Teaching Assistant training We will develop a program to help graduate students get a deeper teaching experience. When we were in graduate school in France, as a part of teaching duties, TAs had to attend each year 10 days of seminars and work on teaching. The emphasis was not on ready-made solutions, but on the introduction and discussion of important issues. This is why we wish to return the favor, and introduce a similar program at MIT. We believe that teaching assistants will greatly benefit from the discussions with peer fellow TAs and the feedback from experienced professors. We hope that it will convince some of them to pursue an academic career.

Involvement of undergraduates in research. MIT has a well-known research program for undergraduates (Undergraduate Research Opportunities Program, UROP) which we expect to utilize to include several undergraduates in this research project. We will put a particular emphasis on first-year students by proposing projects that do not require extensive programming experience. We hope that an experience in the research lab will convince these students to major in computer science. We are, for example, currently working with a female undergraduate Tiffany Wang, who is performing acquisition high-quality acquisition of textured 3D models.

Promotion of the participation of females and minorities. The principal investigator is committed to promoting the participation of females and underrepresented minorities in science and engineering, both at the undergraduate and graduate levels. We will promote the participation of undergraduate students to our research as described above. Our team currently includes three female graduate students (Sara Su, Soonmin Bae and Tilke Judd), a female post-doc (Barbara Cutler) and an underrepresented-minority student (Paul Green).

International impact through teaching and research collaborations. Our advanced graphics course will be offered in the context of the Singapore-MIT Alliance. The lectures will be broadcast to Singapore and we will have weekly review sessions in video-conference. This allows talented students from all Asia to benefit from our course development and research experience. We are also engaged in tight collaborations with the Artis team of Francois Sillion in Grenoble, France. We are co-supervising a student, Stéphane Grabli, and have been working together on a number of projects such as Billboard Clouds [DDSD03] and flash photography [ED04]. A number of Artis and MIT students have made multi-month stays in the other team. We are committed to strengthening this collaboration through joint projects and exchange of students and post-docs.

Local collaborations and promotion of a Cambridge graphics community. Cambridge has a substantial number of computer graphics researchers at MIT, Harvarad and Mitsubishi Electric Research Lab. We are in close contact with these researchers through collaborations (with Hanspeter Pfister at MERL), serving on PhD committees, and regular informal visits. We are also promoting synergies between these teams by organizing informal and formal events that gather researchers, interns and graduate students. We are starting this year a yearly mini Symposium that will gather graphics researchers from Cambridge for one afternoon. Each scientist or student will have three minutes to give an overview of their research. The goal is to present researchers and students, particularly undergraduates, with a diverse and exciting overview of the research done in the neighborhood. We believe that it will then facilitate informal discussions and collaborations.

Interdisciplinary and education impact through conference organization We are planning to organize a number of interdisciplinary events during the five-year period of this research. In May 2005, we are hosting at MIT a *Symposium on Computational photography* together with colleagues Marc Levoy (Stanford) and Rick Szeliski (Microsoft), http://photo.csail.mit.edu. This event will gather 200 scientists and professionals from fields such graphics, vision, photography and optics.

In the third year of this project, we hope to organize, together with Bill Thompson and Pete Shirley from the University of Utah, an informal workshop on visual signal statistics that gathers vision scientists, psychologists and computer graphics researchers. This will probably take the form of an ACM Siggraph/Eurographics Campfire.

The PI is also on the advisory board of an exciting event on scientific illustration, *Image and Meaning 2*. In this workshop, two hundred people from diverse areas of science, illustration, and education will participate in interdisciplinary activities to study how one learns from images and learns by producing images. We hope that this event will have a strong impact on the role of images in scientific and engineering education. See http://web.mit.edu/i-m/.

7 Result from Prior NSF support

The principal investigator has not received NSF funding in the past.

References

[AB91] E. H. Adelson and J. R. Bergen. The plenoptic function and the elements of early vision. *M. Landy and J. A. Movshon, (eds) Computational Models of Visual Processing*, 1991.

[ABCC03] Boris Aronov, Hervé Brönnimann, Allen Y. Chang, and Yi-Jen Chiang. Cost-driven octree construction schemes: An experimental study. In *Annu. ACM Symp. Comput. Geom*, 2003.

[AKDS04] T. Annen, J. Kautz, F. Durand, and H.P. Seidel. Spherical harmonic gradients for mid-range illumination. In *Eurographics Symposium on Rendering*, 2004.

[Alo86] Y. Aloimonos. Detection of surface orientation from texture. i. the case of planes. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 584–593, 1986.

[Ama84] John Amanatides. Ray tracing with cones. *Computer Graphics*, 18(3):129–135, July 1984.

[ARBJ03] Sameer Agarwal, Ravi Ramamoorthi, Serge Belongie, and Henrik Wann Jensen. Structured importance sampling of environment maps. In Jessica Hodgins and John C. Hart, editors, *Proceedings of ACM SIGGRAPH 2003*, volume 22(3) of *ACM Transactions on Graphics*, pages 605–612. ACM Press, 2003.

[Arv94] James Arvo. The irradiance Jacobian for partially occluded polyhedral sources. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 343–350. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.

[ASVNB00] Carlos Andújar, Carlos Saona-Vázquez, Isabel Navazo, and Pere Brunet. Integrating occlusion culling and levels of detail through hardly-visible sets. In *Eurographics 2000*, 2000.

[BGG96] Vicki Bruce, P.R. Green, and M.A. Georgeson. *Visual Perception : Physiology, Psychology and Ecology*. Psychology Pr, 3rd edition, 1996.

[BJ01] Ronen Basri and David Jacobs. Lambertian reflectance and linear subspaces. In *Proceedings of the Eighth International Conference On Computer Vision (ICCV-01)*, pages 383–390, Los Alamitos, CA, July 9–12 2001. IEEE Computer Society.

[BM95] Mark R. Bolin and Gary W. Meyer. A frequency based ray tracer. In *Proceedings of SIGGRAPH 95*, Computer Graphics Proceedings, Annual Conference Series, pages 409–418, August 1995.

[BM98] Mark R. Bolin and Gary W. Meyer. A perceptually based adaptive sampling algorithm. In Michael F. Cohen, editor, *Computer graphics: proceedings: SIGGRAPH 98 Conference proceedings, July 19–24, 1998*, Computer Graphics -proceedings- 1998, pages 299–310, New York, NY 10036, USA and Reading, MA, USA, 1998. ACM Press and Addison-Wesley.

[Bra00] Ronald Bracewell. *The Fourier Transform and its Applications*. Mc Graw Hill, 3rd edition, 2000.

[CD03] Eric Chan and Fredo Durand. Rendering fake soft shadows with smoothies. In *Proceedings of the 14th Eurographics workshop on Rendering*, pages 208–218. Eurographics Association, 2003.

[CD04] Eric Chan and Fredo Durand. An efficient hybrid shadow rendering algorithm. In *Proceedings of the 15th Eurographics workshop on Rendering*, 2004.

[CH03] F. Cutzu and R. Hammoud. Estimating the photorealism of images: Distinguishing paintings from photographs. In *Proc. of CVPR*, 2003.

[CLF+03] P. H. Christensen, D. M. Laur, J. Fong, W. L. Wooten, and D. Batali. Ray differentials and multiresolution geometry caching for distribution ray tracing in complex scenes. In P. Brunet and D. Fellner, editors, *Proceedings of the 24th Annual Conference of the European Association for Computer Graphics (EG-03)*, volume 22, 3 of *Computer Graphics forum*, pages 543–552, Oxford, UK, September 1–6 2003. Blackwell Publishing Ltd.

[Cro77] Franklin C. Crow. The aliasing problem in computer-generated shaded images. *Communications of the ACM*, 20(11):799–805, November 1977.

[CT82] R. Cook and K. Torrance. A reflection model for computer graphics. *ACM Transactions On Graphics*, 1(1):7–24, 1982.

[CTCS00] Jin-Xiang Chai, Xin Tong, Shing-Chow Chan, and Heung-Yeung Shum. Plenoptic sampling. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, pages 307–318, July 2000.

[DAW01] Ron O. Dror, Edward H. Adelson, and Alan S. Willsky. Recognition of surface reflectance properties from a single image under unknown real-world illumination. In *Proc. IEEE workshop on IOAV in Lighting : Psychophysics and Computation*, December 2001.

[DBB03] Philip Dutre, Philippe Bekaert, and Kavita Bala. *Advanced Global Illumination*. AK Peters Limited, 2003.

[DD02] Frédo Durand and Julie Dorsey. Fast bilateral filtering for the display of high dynamic range image. In John Hughes, editor, *SIGGRAPH 2002 Conference Graphics Proceedings*, Annual Conference Series, pages 257–265. ACM Press/ACM SIGGRAPH, 2002.

[DDP96] Frédo Durand, George Drettakis, and Claude Puech. The 3D visibility complex: A new approach to the problems of accurate visibility. In Xavier Pueyo and Peter Schröder, editors, *Eurographics Rendering Workshop 1996*, pages 245–256, New York City, NY, June 1996. Eurographics, Springer Wein. ISBN 3-211-82883-4, http://www-imagis.imag.fr/Publications/.

[DDP97a] F. Durand, G. Drettakis, and C. Puech. 3d visibility made visibly simple. In *video 13th Annu. ACM Sympos. Comput. Geom.*, 1997.

[DDP97b] Frédo Durand, George Drettakis, and Claude Puech. The 3d visibility complex: a unified data-structure for global visibility of scenes of polygons and smooth objects. In *Canadian Conference on Computational Geometry*, August 1997. http://www-imagis.imag.fr/˜Fredo.Durand.

[DDP97c] Fredo Durand, George Drettakis, and Claude Puech. The visibility skeleton: a powerful and efficient multi-purpose global visibility tool. *Computer Graphics*, 31(3A):89–100, August 1997. http://www-imagis.imag.fr/Publications/.

[DDP99] Frédo Durand, George Drettakis, and Claude Puech. Fast and accurate hierarchical radiosity using global visibility. *ACM Transactions on Graphics*, April 1999. to appear. http://www-imagis.imag.fr/˜Fredo.Durand.

[DDP02] Frédo Durand, George Drettakis, and Claude Puech. The 3d visibility complex. *ACM Transactions on Graphics*, 21(2):176–206, April 2002.

[DDSD03] Xavier Décoret, Frédo Durand, François X. Sillion, and Julie Dorsey. Billboard clouds for extreme model simplification. In Jessica Hodgins and John C. Hart, editors, *Proceedings of ACM SIGGRAPH 2003*, volume 22(3) of *ACM Transactions on Graphics*, pages 689–696, 2003.

[DDTP00] Frédo Durand, George Drettakis, Joëlle Thollot, and Claude Puech. Conservative visibility preprocessing using extended projections. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, pages 239–248, July 2000.

[DF94] G. Drettakis and E. Fiume. A fast shadow algorithm for area light sources using backprojection. *Computer Graphics*, 28(Annual Conference Series):223–230, July 1994. http://www-imagis.imag.fr/˜George.Drettakis/pub.html.

[DLAW01] Ron O. Dror, Thomas K. Leung, Edward H. Adelson, and Alan S. Willsky. Statistics of real-world illumination. In *Proceedings of CVPR*, December 2001.

[Dur99] Frédo Durand. *3D Visibility, analysis and applications*. PhD thesis, U. Joseph Fourier, Grenoble, 1999. http://www-imagis.imag.fr.

[Dur02] F. Durand. An invitation to discuss computer depiction. In *Proc. of the ACM/Eurographics Symposium on Non-Photorealistic Animation and Rendering (NPAR)*, 2002.

[ED04] Elmar Eisemann and Frédo Durand. Flash photography enhancement via intrinsic relighting. In *SIGGRAPH*, 2004.

[FL03] Hany Farid and Siwei Lyu. Higher-order wavelet statistics and their application to digital forensics. In *IEEE Workshop on Statistical Analysis in Computer Vision (in conjunction with CVPR)*, 2003.

[Fla99] Patrick Flandrin. *Time-Frequency/Time-scale Analysis*. Academic Press, 1999.

[Fle03] R. Fleming. *Human visual perception under real-world illumination*. PhD thesis, Massachusetts Institute of Technology, 2003.

[FP99] W. T. Freeman and E. Pasztor. Learning low-level vision. In *7th International Conference on Computer Vision*, pages 1182–1189, 1999. See also http://www.merl.com/reports/TR99-12/.

[FPSG97] J. Ferwerda, Pattanaik, Shirley, and Greenberg. A model of visual masking for computer graphics. In *Computer Graphics (Proc. Siggraph)*, 1997.

[FS98] Hans G. Feichtinger and Thomas Strohmer, editors. *Gabor Analysis and Algorithms: Theory and Applications*. Birkhauser Boston, 1998.

[GGSC96] S.. Gortler, R. Grzeszczuk, R. Szeliski, and M. Cohen. The lumigraph. In *Proc. of SIGGRAPH*, 1996.

[Gol97] Dan B. Goldman. Fake fur rendering. In Turner Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 127–134. ACM SIGGRAPH, Addison Wesley, August 1997. ISBN 0-89791-896-7.

[Goo96] Joseph Goodman. *Introduction To Fourier Optics*. McGraw-Hill, 1996.

[GSCH93] Steven J. Gortler, Peter Schroder, Michael F. Cohen, and Pat Hanrahan. Wavelet radiosity. In *Computer Graphics Proceedings, Annual Conference Series, 1993*, pages 221–230, 1993.

[Hec89] P. Heckbert. Fundamentals of texture mapping and image warping. Master's thesis, University of California at Berkeley, Department of Electrical Engineering and Computer Science, 1989.

[Hec92] Paul S. Heckbert. Radiosity in flatland. In A. Kilgour and L. Kjelldahl, editors, *Computer Graphics Forum (EUROGRAPHICS '92 Proceedings)*, volume 11(3), pages 181–192, September 1992.

[Hec01] Eugene Hecht. *Optics*. Pearson Addison Wesley, 4th edition, 2001.

[HM01] F. Hlawatsch and G. Matz. Quadratic time-frequency analysis of linear, time-varying systems. In L. Debnath, editor, *Wavelet Transforms and Time-Frequency Analysis*, chapter 9, pages 235–287. Boston: Birkhäuser, 2001.

[HSA91] Pat Hanrahan, David Salzman, and Larry Aupperle. A rapid hierarchical radiosity algorithm. In Thomas W. Sederberg, editor, *Computer Graphics (SIGGRAPH '91 Proceedings)*, volume 25, pages 197–206, July 1991.

[HTSG91] Xiao D. He, Kenneth E. Torrance, François X. Sillion, and Donald P. Greenberg. A comprehensive physical model for light reflection. In *Computer Graphics (Proceedings of SIGGRAPH 91)*, volume 25, pages 175–186, July 1991.

[Ige99] Homan Igehy. Tracing ray differentials. In *Proceedings of SIGGRAPH 99*, Computer Graphics Proceedings, Annual Conference Series, pages 179–186, August 1999.

[IMG00] Aaron Isaksen, Leonard McMillan, and Steven J. Gortler. Dynamically reparameterized light fields. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, pages 297–306, July 2000.

[JDD03] Thouis R. Jones, Frédo Durand, and Mathieu Desbrun. Non-iterative, feature-preserving mesh smoothing. In *ACM Transactions on Graphics, SIGGRAPH'03 proceedings*, August 2003. http://www-grail.usc.edu/pubs.html.

[JDS+01] Henrik Wann Jensen, Frédo Durand, Michael M. Stark, Simon Premoze, Julie Dorsey, and Peter Shirley. A physically-based night sky model. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, pages 399–408, August 2001.

[Jen01] Henrik Wann Jensen. *Realistic Image Synthesis Using Photon Mapping*. AK Peters, 2001.

[JMLH01] Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. A practical model for subsurface light transport. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, pages 511–518, August 2001.

[Kel01] Alexander Keller. Hierarchical monte carlo image synthesis. *Mathematics and Computers in Simulation*, 55(1–3):79–92, February 2001.

[KM99] Jan Kautz and Michael D. McCool. Interactive rendering with arbitrary BRDFs using separable approximations. In ACM, editor, *SIGGRAPH 99. Proceedings of the 1999 SIGGRAPH annual conference: Conference abstracts and applications*, Computer Graphics, pages 253–253, New York, NY 10036, USA, 1999. ACM Press.

[LFTG97] E. Lafortune, S.-C. Foo, K. Torrance, and D. Greenberg. Non-linear approximation of reflectance functions. *Computer Graphics*, 31(Annual Conference Series):117–126, 1997.

[LH96] M. Levoy and P. Hanrahan. Light field rendering. In *Proc. of SIGGRAPH*, 1996.

[LH01] David Luebke and Benjamin Hallen. Perceptually-Driven simplification for interactive rendering. pages 223–234, 2001.

[LRR04] Jason Lawrence, Szymon Rusinkiewicz, and Ravi Ramamoorthi. Efficient BRDF importance sampling using a factored representation. In *ACM SIGGRAPH 2004 Full Conference DVD-ROM*, August 2004. To appear.

[LZW03] A. Levin, A. Zomet, and Y. Weiss. Learning to perceive transparency from the statistics of natural scenes. In *Adv. in Neural Info. Proc. Systems*, volume 15. MIT Press, 2003.

[MAA01] Michael D. McCool, Jason Ang, and Anis Ahmad. Homomorphic factorization of BRDFs for high-performance rendering. In Eugene Fiume, editor, *SIGGRAPH 2001, Computer Graphics Proceedings*, Annual Conference Series, pages 185–194. ACM Press / ACM SIGGRAPH, 2001.

[Mal98] Stéphane Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, San Diego, 1998.

[Mat00] Gerald Matz. *A Time-Frequency Calculus for Time-Varying Systems and Nonstationary Processes with Applications*. PhD thesis, Technischen Universität Wien, 2000.

[McC99] Michael D. McCool. Anisotropic diffusion for monte carlo noise reduction. In *ACM Transactions on Graphics*, volume 18 (2), pages 171–194, 1999.

[Mil56] George A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63:81–97, 1956. http://www.well.com/user/smalin/miller.html.

[MJC+03] Stephen R. Marschner, Henrik Wann Jensen, Mike Cammarano, Steve Worley, and Pat Hanrahan. Light scattering from human hair fibers. In Jessica Hodgins and John C. Hart, editors, *Proceedings of ACM SIGGRAPH 2003*, volume 22(3) of *ACM Transactions on Graphics*, pages 780–791, 2003.

[MPBM03a] Wojciech Matusik, Hanspeter Pfister, Matthew Brand, and Leonard McMillan. A data-driven reflectance model. *ACM Transactions on Graphics*, 22(3):759–769, July 2003.

[MPBM03b] Wojciech Matusik, Hanspeter Pfister, Matthew Brand, and Leonard McMillan. Efficient isotropic BRDF measurement. In *Proceedings of the 14th Eurographics workshop on Rendering*, pages 241–248. Eurographics Association, 2003.

[MPN+02] Wojciech Matusik, Hanspeter Pfister, Addy Ngan, Paul Beardsley, Remo Ziegler, and Leonard McMillan. Image-based 3D photography using opacity hulls. In *SIGGRAPH 2002 Conference Proceedings*, Annual Conference Series, pages 427–437, 2002.

[MR97] J. Malik and R. Rosenholtz. Computing local surface orientation and shape from texture for curved surfaces. *Int. J. Computer Vision*, pages 149–168, 1997.

[MWL+99] Stephen R. Marschner, Stephen H. Westin, Eric P. F. Lafortune, Kenneth E. Torrance, and Donald P. Greenberg. Image-based BRDF measurement including human skin. In *Eurographics Workshop on Rendering*, 1999.

[Mys98] Karol Myszkowski. The visible differences predictor: applications to global illumination problems. In *Eurographics Workshop on Rendering*, Vienna, Austria, June 1998.

[NDM04] Addy Ngan, Fredo Durand, and Wojciech Matusik. Experimental validation of analytical brdf models. In *ACM Siggraph technical sketch*, 2004.

[NRH03] Ren Ng, Ravi Ramamoorthi, and Pat Hanrahan. All-frequency shadows using non-linear wavelet lighting approximation. In *ACM SIGGRAPH 2003 Full Conference DVD-ROM*, 2003.

[OCDD01] B. M. Oh, M. Chen, J. Dorsey, and F. Durand. Image-based modeling and photo editing. *Proc. SIGGRAPH*, 2001.

[ODJ04] Victor Ostromoukhov, Charles Donohue, and Pierre-Marc Jodoin. Fast hierarchical importance sampling with blue noise properties. In *ACM SIGGRAPH 2004 Full Conference DVD-ROM*, August 2004. To appear.

[ORDP96] Rachel Orti, Stéphane Rivière, Frédo Durand, and Claude Puech. Radiosity for dynamic scenes in flatland with the visibility complex. In Jarek Rossignac and François Sillion, editors, *Computer Graphics Forum (Proc. of Eurographics '96)*, volume 16(3), pages 237–249, Poitiers, France, September 1996.

[PA04] S. Premoze and M. Ashikhmin. Efficient rendering of multiple scattering effects in participating media. In *Eurographics Symposium on Rendering*, 2004.

[PV96] Pocchiola and Vegter. Topologically sweeping visibility complexes via pseudotriangulations. *GEOMETRY: Discrete & Computational Geometry*, 16, 1996.

[RH01a] R. Ramamoorthi and P. Hanrahan. Analysis of planar light fields from homogeneous convex curved surfaces under distant illumination. In *HVEI*, pages 195–208, 2001.

[RH01b] Ravi Ramamoorthi and Pat Hanrahan. An efficient representation for irradiance environment maps. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, pages 497–500, August 2001.

[RH01c] Ravi Ramamoorthi and Pat Hanrahan. A signal-processing framework for inverse rendering. *Proceedings of SIGGRAPH 2001*, pages 117–128, August 2001. ISBN 1-58113-292-1.

[RH02] Ravi Ramamoorthi and Pat Hanrahan. Frequency space environment map rendering. *ACM Transactions on Graphics*, 21(3):517–526, July 2002.

[RKB04] Ravi Ramamoorthi, Melissa Koudelka, and Peter Belhumeur. A fourier theory for cast shadows. In *ECCV*, pages I 146–162, 2004.

[RPG99] Mahesh Ramasubramanian, Sumanta N. Pattanaik, and Donald P. Greenberg. A perceptually based physical error metric for realistic image synthesis. In Alyn Rockwood, editor, *Siggraph 1999, Computer Graphics Proceedings*, Annual Conference Series, pages 73–82, Los Angeles, 1999. ACM Siggraph, Addison Wesley Longman.

[SD95] François Sillion and George Drettakis. Feature-based control of visibility error: A multi-resolution clustering algorithm for global illumination. In Robert Cook, editor, *SIGGRAPH 95 Conference Proceedings*, Annual Conference Series, pages 145–152. ACM SIGGRAPH, Addison Wesley, August 1995. held in Los Angeles, California, 06-11 August 1995, http://www-imagis.imag.fr/˜Francois.Sillion/.

[SGwHS98] Jonathan Shade, Steven J. Gortler, Li wei He, and Richard Szeliski. Layered depth images. In Michael Cohen, editor, *Proceedings of SIGGRAPH 98*, Annual Conference Series, Addison Wesley, pages 231–242. Addison Wesley, 1998.

[SKS02] Peter-Pike Sloan, Jan Kautz, and John Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics*, 21(3):527–536, July 2002.

[SO01] E P Simoncelli and B A Olshausen. Natural image statistics and neural representation. *Annual Review of Neuroscience*, 24:1193–1216, May 2001.

[SS96] Cyril Soler and François Sillion. Accurate error bounds for multi-resolution visibility. In Xavier Pueyo and Peter Schröder, editors, *Eurographics Rendering Workshop 1996*, pages 133–142, New York City, NY, June 1996. Eurographics, Springer Wein. ISBN 3-211-82883-4, http://www-imagis.imag.fr/˜Cyril.Soler/csoler.gb.html.

[SS98] Cyril Soler and François Sillion. Fast calculation of soft shadow textures using convolution. In *Computer Graphics Proceedings, Annual Conference Series:* SIGGRAPH '98 (Orlando, FL), page ?? ACM SIGGRAPH, New York, July 1998. http://www-imagis.imag.fr/˜Cyril.Soler/csoler.gb.html.

[Sta99] Jos Stam. Diffraction shaders. In Alyn Rockwood, editor, *Siggraph 1999, Computer Graphics Proceedings*, Annual Conference Series, pages 101–110, Los Angeles, 1999. ACM Siggraph, Addison Wesley Longman.

[STN87] Mikio Shinya, Tokiichiro Takahashi, and Seiichiro Naito. Principles and applications of pencil tracing. In Maureen C. Stone, editor, *Computer Graphics (SIGGRAPH '87 Proceedings)*, volume 21(4), pages 45–54, July 1987.

[Sto91] J. Stolfi. *Oriented Projective Geometry: A Framework for Geometric Computations*. Academic Press, 1991.

[SW01] Frank Suykens and Yves Willems. Path differentials and applications. In *Eurographics Workshop on Rendering*, pages 257–268, 2001.

[TFA03] M. F. Tappen, W. T. Freeman, and E. H. Adelson. Recovering intrinsic images from a single image. volume 15. MIT Press, 2003.

[TM98] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proc. IEEE Int. Conf. on Computer Vision*, pages 836–846, 1998.

[VG97] Eric Veach and Leonidas J. Guibas. Metropolis light transport. In Turner Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 65–76. ACM SIGGRAPH, Addison Wesley, August 1997. ISBN 0-89791-896-7.

[War92] Greg Ward. Measuring and modeling anisotropic reflection. In *Computer Graphics (Proceedings of SIGGRAPH 92)*, volume 26, pages 265–272, July 1992.

[WH92] Gregory J. Ward and Paul Heckbert. Irradiance gradients. *Third Eurographics Workshop on Rendering*, pages 85–98, May 1992.

[WRC88] Gregory J. Ward, Francis M. Rubinstein, and Robert D. Clear. A ray tracing solution for diffuse interreflection. *Computer Graphics*, 22(4):85–92, August 1988.

[ZMHH97] Hansong Zhang, Dinesh Manocha, Thomas Hudson, and Kenneth E. Hoff III. Visibility culling using hierarchical occlusion maps. In Turner Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 77–88. ACM SIGGRAPH, Addison Wesley, August 1997. ISBN 0-89791-896-7, http://www.cs.unc.edu/˜zhangh/hom.html.