A framework for Animation in Global Illumination Environments

I. Martin
X. Pueyo
D. Tost

Report LSI–98–32–R

A Framework for Animation in Global Illumination Environments

I. Martin, X. Pueyo

Institut d'Informàtica i Aplicacions
Universitat de Girona, Spain

D. Tost

Departament de Llenguatges i Sistemas Informàtics
Universitat Politècnica de Catalunya, Barcelona, Spain

Abstract: This paper proposes a framework for the production of animations of high quality radiosity environments which makes use of the a-priori knowledge of the dynamic properties of the scene in order to exploit temporal coherence. A discussion of the previous work leads to design a two-pass strategy that extends to the temporal dimension a related static method [MPT98] that computes a coarse global solution and then performs a final gathering step using hardware-graphics accelerators. The input data are a dynamic model of the environment through a period of time corresponding to the same camera recording. The strategy designed consists into two processes: a pre-processing stage and a production stage. The aim of the former one is to build a data structure able to store the eventual changes in the radiosity interactions throughout the sequence along with the predictable modifications in the image. The data structure is hierarchical so that maximum adaptability can be achieved in traversing it. The production stage computes all the frames by traversing the data structure and computing incrementally both the global and the local pass. A first implementation of the general framework is discussed along with future development lines.

1 Introduction

In the milestones track of realistic image synthesis, a step beyond is the production of animations using global illumination models, as dynamic behaviors greatly improve the perception of realism. Up to now, the advances in this field have been limited by the high computational cost of realistic animation synthesis. The reduction of the number of computations by taking advantage of frame-to-frame coherence is thus a challenging problem that have been often addressed in the bibliography. However, many previous works focus only at visibility computations and little has been done in global illumination. Several performing methods have been proposed for the interactive design of radiosity environments but they are not totally suitable for truly dynamic scenes. This paper presents a general framework for the production of efficient realistic animations using radiosity. The motion of the objects and the camera is supposed to be known a priori, an hypothesis which, although being restrictive, seems reasonable, taking into account that the design of an animation sequence is a slow, sometimes tedious process generally done using wireframe projections or simple local illumination models.

The proposed framework relies onto two basic assumptions: first, the energy interaction between any pair of elements of the environment should be treated at the appropriate level of detail, which may not be the same along the whole sequence, next the rendering of any image should also be computed at a suitable precision. Thus, the radiosity solution at any instant of a temporal sequence can be considered as a rough approximation of the next instant solution, which can be more less refined according to the relevance of the dynamic changes. Accordingly, any frame is an approximation of the next frame, which can be simply replicated from it, or interpolated or recomputed at relevant pixels.

These assumptions have led to the proposal of a general framework that applies a hierarchical sampling mechanism in both spaces: image-space and object-space. This allows to reconstruct with a given error level the continuous function of the global illumination of the environment over the period, as it is seen along the camera path. A two-pass general strategy is designed that allows two levels of treatment (global and local) and which exploits as much as possible all the levels of a hierarchical structuring of the dynamic environment. This general strategy extends to 4-D a previous method of the authors [MPT98].

The paper is structured into six sections. First the previous work is reviewed in section2, next, an overview of the strategy is presented in section 3. In section 4, the pre-processing stage is analyzed and the construction oracles are discussed. In section 5, the production stage is reviewed and the traversal oracles are commented. The first simulation results are analyzed in section 6 and the future work is presented along with the conclusions in section 6.

2 Previous Work

The main problem that faces the animation of global illumination environments is that the high computational cost of producing a single image is multiplied by the number of frames, typically at least 24 frames per second. Besides, successive frames of a single shot do not differ very much one to each other. This similarity, called frame-to-frame coherence or temporal coherence, can be used to accelerate different steps in the global illumination computations.

The use of frame-to-frame coherence has been early reported for the acceleration of hidden lines and surfaces removal [HZ82, Bad88, CCD90, GP91, Tos91, AH95]. A comparative analysis of these references can be found in [Mar94]. Herein, only the use of temporal coherence in global illumination is reviewed.

One of the most time consuming steps in global illumination algorithms is the visibility computations. In finite elements methods this process is embedded in the form factor computations. There has been several attempts to characterize the changes in the visibility produced by moving objects and incrementally compute the form factors. In [BWCG86] an a priori knowledge of the whole animation sequence is assumed, and the viewpoint is considered fixed. The form factors between patches that do not interact with moving objects are identified, computed in a preprocess and stored. In addition modified form factors are incrementally computed by using reprojection techniques. In [FYT94] the hierarchical radiosity is extended to dynamic environments by updating the hierarchy of links between static objects affected by moving occluders. Links can be refined, unrefined and temporally occluded. As the moving objects are known in advance the link events are easily predictable. However, links involving a dynamic patch are computed from scratch. The main limitation of this approach is that it cannot handle light sources movements. From a more theoretical point of view and restricted to 2-dimensions, in [ORDP96], a mechanism for the prediction of changes in form factors is provided based on geometrical relationships. A data structure called *visibility complex* is introduced that allows quick recomputations. In [DDP97] the work has been extended to 3D keeping memory and computation time reasonable.

Besides the visibility computations, the update in the energy transport has also been addressed. In [Che90, GSG90] an incremental computation of the radiosity solution is proposed. Both references are based on progressive radiosity. At each frame moving patches are first unshooted with their previous frame energy and next shooted with their current radiosity. This mechanism is called redistribution. For example, turning off a light source leads to shooting negative radiosity from the light source. This approach has proven to be valuable for interactive indoor design where the frame ratio can be low. A similar technique called re-propagation stores a dynamic array called Shadow Form Factor List (SFFL) which allows to exactly re-propagate radiosity only on the parts of the scene affected by the changes [MS94]. Although it allows near real-time animations it is restricted to a single change per frame. In addition, as it is based in a fixed initial meshing artifacts may appear at low patch resolutions. Finally, in [BS96] a global Monte Carlo method is used to compute the radiosity solution for a time interval, which avoids to perform repeated light transports.

Hierarchical radiosity has been the subject of research to find Dynamic solutions for The efficient recomputation of links in hierarchical radiosity is also an open subject of research. In [Sha97] links are classified in order to allow dynamic events to be easyly identified performing thus efficient computations. An approach towards interactive changes is presented in [DS97] that achieves fast recomputation of links and energy redistribution.

In [NDR96] the view space is sampled and then any camera path can be reconstructed from these samples. The view space is the set of all possible views in a 4D space (3D space + time). It allows to efficiently generating animation in generic global illumination environments but choosing the optimum set of samples is an open problem.

Synthesizing, leaving aside the use of frame-to-frame coherence for visibility computations and the interactive edition of radiosity environments, the temporally coherent computations of radiosity animations are still limited by strong initial hypotheses such as the number of changes per frame or the initial resolution of the meshing. The use of frame-to-frame coherence for dynamic radiosity scenes allowing both camera and actors movement and enabling to obtain any desired level of precision in the rendering is thus an open problem. It is the purpose of the framework next proposed.

3 Requirements of the framework

In order to compute a global illumination animation sequence the following requirements should be fulfilled:

- High degree of adaptativity. Computations can be performed at different levels so that a good tradeoff between error and computation time can be achieved.
- High quality images. The method should be able to compute images as accurate as reference images.
- Illumination should be computed and stored in a way that coherence is easily inferred.

The framework proposed herein tries to meet these requirements with two main features, first, by generalizing the concept of hierarchy, and next by splitting the computations into two pass: a coarse global one and a local refinement step. The need for these two features is next justified and the mean to develop them are next exposed.

Hierarchical sampling over time

The animation sequence can be viewed as a continuous function of the global illumination of the environment that must be sampled at a given rate in order to produce the frames needed. In the sampling process an approximation of the global illumination perceived by a virtual camera at a given instant is realized. Generally, this sampling is uniform: each frame is computed independently from the other ones, by applying the same strategy. However regular sampling presents obvious drawbacks, as the same level of precision in the computations is not required ever and everywhere. The same motivation that once led to switch from uniform environment meshing to hierarchical ones [HSA91] seems to apply to the temporal dimension. Not all the changes through time are equally relevant. Moreover some energy interactions do not even change.

Exploiting temporal coherence can thus be understood as the mean to perform the less computations needed along time in order to achieve the required degree of approximation of the global illumination function as it is seen by a moving observer. To guarantee this adaptability a hierarchical subdivision of the temporal dimension as well as of the spatial one seems to be necessary. Moreover, in order to make the computations as flexible as possible, the maximum hierarchization should be enabled, not only at the polygon level but also by using multiple levels of clusters [SAG94].

Assuming a previous knowledge of the animation, a hierarchical representation of the 4-D environment should thus be computed. Two different approaches exist depending on the primary criterion used for the subdivision: either time or space. Figure 1 illustrates these two possibilities. If the primary criterion is time, then the root level of the hierarchy contains a coarse representation of the whole sequence, which is recursively subdivided into smaller time intervals. The maximum level of subdivision would correspond to the frame rate. Each node

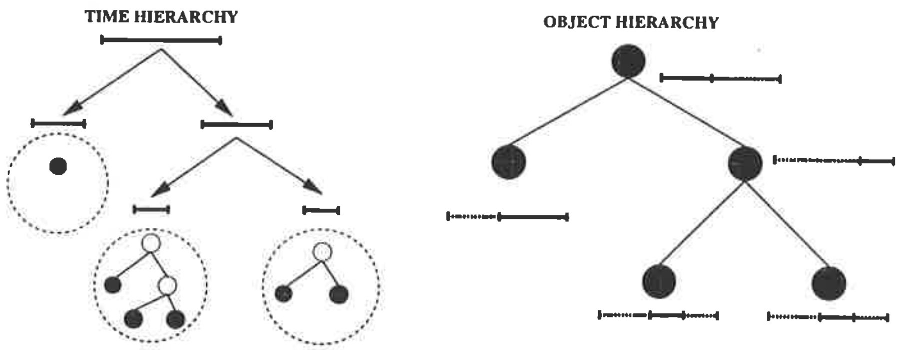

Fig. 1. Two options for storing temporal infomation. On the left, time is recursiveley subdivided and each interval stores all the objects information. On the right, objects are recursively subdivided and each object stores its time activity.

itself contains a spatial hierarchical subdivision of the environment. In the second approach the spatial subdivision instead is the primary criterion. Each node of the subdivision, contains a representation of the object (cluster, object, polygon , etc.) through time. The framework proposed herein is based on the latter approach although the analysis of the former one also seems to be interesting.

Two-pass strategy

In classical key-frame animation, two levels of production exist: the drawing of the key-frames which correspond to significant changes in the sequence, and the interpolation of the in-betweening frames. In global illumination animation these two levels could be interpreted as follows: the key-frames correspond to relevant changes of the radiance function whereas, inbetweenings are simple changes in the rendering stage. In order to allow maximum flexibility in the incremental computations, both levels of processing should be allowed the computation of the object-space radiosity interactions, as well as the image-space frame rendering. A two-pass strategy enabling to incrementally compute a global pass and a local pass seems to fit with this requirement.

In [MPT98] a two-pass radiosity algorithm for static scenes has been proposed whose structure fits with the strategy exposed above. It is herein extended to the temporal dimension.

4 Framework

4.1 The Hierarchical Sampling Description

The Hierarchical Sampling Description (HSD) is defined as a description of the sampling procedure at each instant of the sequence, for any object to be reconstructed through time. Each HSD contains a partition of the whole extent of the sequence. For each interval of the partition it stores the following information:

- The object is not active at this interval and an object at higher level of the hierarchy is used instead.
- The object is not active at this interval and a set of objects at lower levels are used instead.
- The object is active at this interval. A set of samples is stored to allow its reconstruction.

The hierarchy of HSD is constructed in a preprocessing step and it is then used to produce the sequence (Figure 2).

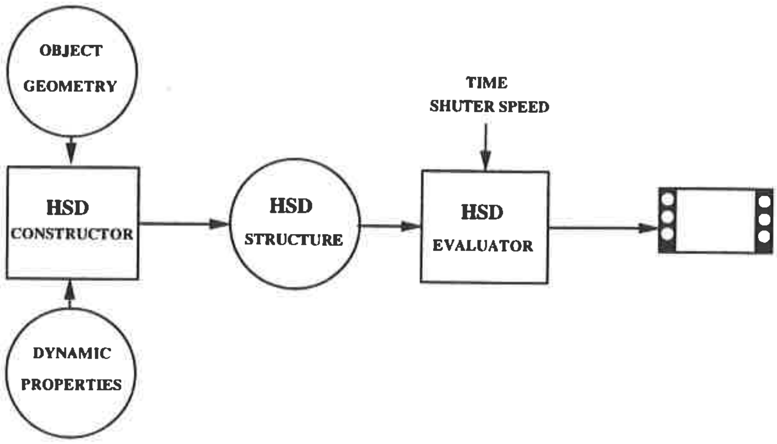

Fig. 2. The HSD pipeline. The main structure is created in a preprocessing step and then passed to the HSD evaluator that can reconstruct dynamic objects at any time.

In order to keep manageable data structures, only enough temporally coherent subsequences of the animation should be considered at a time. Otherwise, the memory usage to store temporal information in the pre-process would probably be prohibitive and the expected speed-up of the frames production would be very low. Different criteria could be given to split the sequence. A natural one is to divide it according to the different shots of the animation. A subsequence would thus correspond to one camera. In the navigation through compartmented buildings, the change of room should also be considered.

The construction algorithm is a recursive procedure based onto two main oracles. For each object, the procedure is first called with the top object and the whole sequence as an interval. The object is analyzed over time and a sequence of intervals is constructed that represents the temporal sampling strategy to be used. Active intervals are processed and samples are determined, while inactive ones are recursively passed to the procedure for all the descendents of the object being treated.

This procedure relies onto two oracles that drive the HSD creation:

- Interval Activity Oracle (IA): It decides if an object is active for a given time interval. This depends on the final impact (images) of using an object at a given level of the hierarchy.
- Interval Sampling Oracle (IS): It chooses the sampling set for an active interval. Once an interval has been chosen as active, a set of samples must be selected for the object reconstruction during this interval. No restriction for the reconstruction method is imposed in the mechanism herein proposed. The main factor to be taken into account is the coherence degree of the object value over the interval. In general, the greater the coherence, the smaller the sample set.

These two oracles have as a main source of information the dynamic properties of the scene. The a priori knowledge of these properties makes it possible to take decisions on the degree of coherence from frame-to-frame. However, the way this information is applied depends strongly on the type of object. In the next section the HSD strategy is applied to the special case of the objects of the two-pass strategy.

4.2 General strategy

The static strategy proposed in [MPT98] consists into two passes:

- A global pass which uses a clustering radiosity method to compute a coarse global solution. It is based on a hierarchical representation of the illumination, such that the energy exchange is represented by a top link which encloses all the energy interactions of the scene, enabling adaptive refinements according to a given error threshold.
- A local pass that performs a final gathering step using a hardware graphics accelerator and that stores illumination in textures. For each visible element of the cluster hierarchy a texture is computed that represents the illumination of the objects contained in this element. Thus, textures can be used to render at any resolution, from clusters (impostors) to subpatches of polygonal surfaces.

The extension of this method to an animation context would suppose computing the first frame from scratch and then performing the following recomputations for each subsequent frame:

- Global pass: form factor recomputation plus energy redistribution and refinement or unrefinement.
- Local pass: computation of textures for new visible objects and expand or collapse textures of size-changing objects.

These recomputations are very time consuming as they involve visibility changes. Most of the recomputation is concentrated into two objects: links (in the global pass) and textures (in the local pass). The hierarchical approach described above is applied to links as well as to textures in order to allow each recomputation to be performed as much adaptively as possible.

5 HSD on Links and Textures

5.1 Links

The energy exchange in the scene is represented by links. The energy carried by each link may change along the sequence, and its value must be reconstructed from a set of samples. Besides, a link can be used to transport energy during a given interval, but it may be disabled during another one because the energy transport is being performed at a different level of the link hierarchy. All this information is stored in HSD nodes, which are computed using the two HSD oracles described generically above:

- Interval Activity Oracle (IA) for links: It decides the level of the link hierarchy that must be used for a given interval. This should be decided using the refining criteria of the clustering method. Several refining criteria have been proposed that take into account some of these factors:
  - Energy. Classical hierarchical radiosity uses the absolute value of the energy transported by the link as a measure of the error. Energy variation can be due to visibility changes, geometric transformations of the shooter or the receiver, or energy change of the shooter.
  - Importance. Classical hierarchical radiosity can be extended to take into account the point of view to reduce computation, but obtaining a view dependent solution. Obiously changes in the importance will be related to camera movements.
  - Perceptual criteria. Some authors [GH97, MPT97] have proposed to refine links only when this refinement has a significant effect on the final image. These works are mainly based on the idea that energy values are not linearly mapped to pixel intensity values due to the human eye response to illumination. This mapping changes from frame to frame and it is especially noticeable in fast transitions from bright areas to dark ones and viceversa.

  A link will be active during a time interval if it passes the static refinement criteria used to compute the global solution. However this global solution will not be used for the final rendering in the method herein proposed but to generate the texture illumination for visible surfaces. Thus, at this stage, a coarse global solution suffices because details will be recomputed in the second pass. A simple energy plus importance criterion would be enough to produce the global pass.
- Interval Sampling Oracle for links: It chooses the sample set for any interval of activity previously determined. This strongly depends on the reconstruction mechanism chosen for the interval. In general, the temporal coherence degree will determine the way samples are chosen.

5.2 Textures

The final illumination for visible surfaces is computed in a second pass and it is stored at the hierarchy elements as textures. The oracles for the HSD structure construction are again predictions on the interval activity and the interval sampling:

- Interval Activity (IA) for textures: It decides the level of the cluster hierarchy that must be used for a given sequence. This information should be extracted from the camera used for the sequence. The size of the object projection on the screen can be used to determine at which level the object should be rendered.
- Interval Sampling (IS) for textures: Texture sampling along time is a complex issue because it depends on the perceptual response of human eye to image changes. Several factors can be considered:
  - Type of change. Human eye response is known to be more sensitive to contrast variations than to average changes on the illumination. Thus, a small moving shadow is usually more important than large uniform changes.
  - Speed of the texture in the image. This issue is related to motion blur. As frames correspond to the small time intervals in which the camera aperture is opened, fast moving objects will appear blurred in the final image. The textures of these objects may be computed at lower accuracy.

6 First Implementation

The framework herein proposed has just began to be implemented. Up to now only a subset of all the mechanisms has been implemented. A progressive implementation has been engaged starting from the brute force approach and continuing with limited enhancements allowing to analyze the impact of each new feature introduced.

The actual implementation only allows dynamic behavior to the camera. HSD is only applied to textures and Interval Sampling strategy is very simple as it uses one texture per interval. Interval Activity is determined analyzing each polygon projection in the viewport.

Camera movement has been modeled with A—W and translated to and extended MGF format to be processed by the implementation. Most of the camera parameters are animated (eye, center of pan and up vector).

Tables 1 and 2 show statistics for two animations. Sequence 1 corresponds to a scene of 85 polygons, 200 frames and x seconds. Sequence 2 corresponds to a scene of 261 polygons, 300 frames and x seconds. Both sequences have been rendered at 200x200 pixels. Plots show time per frames, number of textures computed per frame, number of textures resized per frame and number of textures reused per frame. Curves show that more computation is done when camera movement is fast and thus less coherence can be extracted.

MPEG files for tha animations can be found at:

`http://ima.udg.es/ imartin/egwr98/mpeg.html`

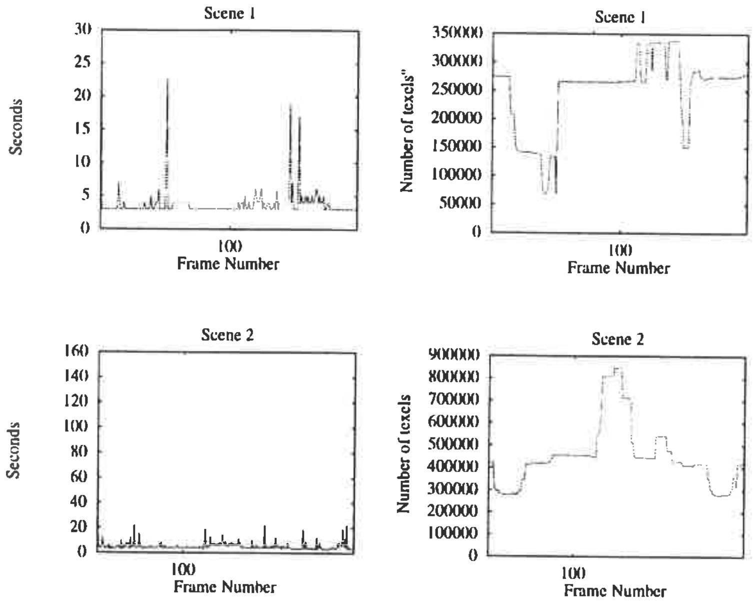

7 Conclusions and Future Work

In this paper a framework is proposed for the efficient production of animations in radiosity scenes with generic dynamic properties. It is based on the construction of a hierarchical representation of the 4-D environment (space+time). In addition, it makes an extensive use of clustering. Finally, it is a two-pass strategy, which enables the use of temporal coherence in the links recomptutation of the global pass as well as in the image textures update of the local pass. The framework allows full adaptativity in the dynamic changes description and it tries to optimize the recomputations in order to obtain an illumination reconstruction as accurate as possible.

This framework is currently being implemented in order to analyze comparatively several factors enabling accurate predictions on the best sampling and reconstruction strategy. The first results show that using hierarchical textures to produce high quality radiosity walkthroughs is very efficient.

Future developments include the implementation and analysis of the complete proposal. Specifically, a sampling strategy for textures is being developed based on temporal texture compression which should allow to reduce the number of textures to be computed without noticeable loss of quality while requiring lower memory.

Acknowledgements

This work has been partly supported by the project TIC 95-0630-C05-05 of the CICYT, and the grant 1995SGR-00232 of the DGR-Generalitat de Catalunya. Ignacio Martin has developed this work with a FI (Formacio de Investigadors) grant from the Generalitat de Catalunya.

References

[AH95] Stephen J. Adelson and Larry F. Hodges. Generating exact ray-traced animation frames by reprojection. *IEEE Computer Graphics and Applications*, 15(3):43–52, May 1995.

[Bad88] Sig Badt. Two algorithms for taking advantage of temporal coherence in ray tracing. *The Visual Computer*, 4(3):123–132, September 1988.

[BS96] Gonzalo Beseuievsky and Mateu Sbert. The Multi-Frame Lighting Method: A Monte Carlo Based Solution for Radiosity in Dynamic Environments. In *Rendering Techniques '96 (Proceedings of the Seventh Eurographics Workshop on Rendering)*, pages 185–194, New York, NY, 1996. Springer-Verlag/Wien.

[BWCG86] Daniel R. Baum, John R. Wallace, Michael F. Cohen, and Donald P. Greenberg. The Back-Buffer Algorithm: An Extension of the Radiosity Method to Dynamic Environments. *The Visual Computer*, 2(5):298–306, September 1986.

[CCD90] John Chapman, Thomas W. Calvert, and John Dill. Exploiting temporal coherence in ray tracing. In *Proceedings of Graphics Interface '90*, pages 196–204, Toronto, Ontario, May 1990. Canadian Information Processing Society.

[Che90] Shenchang Eric Chen. Incremental Radiosity: An Extension of Progressive Radiosity to an Interactive Image Synthesis System. In *Computer Graphics (ACM SIGGRAPH '90 Proceedings)*, volume 24, pages 135–144, August 1990.

[DDP97] Fredo Durand, Georges Drettakis, and Claude Puech. The Visibility Skeleton : A powerful and efficient multi-purpose Global Visibility Tool. In *Computer Graphics Proceedings, Annual Conference Series, 1997 (ACM SIGGRAPH '97 Proceedings)*, pages 89–100, 1997.

[DS97] George Drettakis and Francois X. Sillion. Interactive update of global illumination using a line-space hierarchy. In *Computer Graphics (ACM SIGGRAPH '97 Proceedings)*, volume 31, pages 57–64, 1997.

[FYT94] David A. Forsyth, Chien Yang, and Kim Teo. Efficient Radiosity in Dynamic Environments. In *Fifth Eurographics Workshop on Rendering*, pages 313–323, Darmstadt, Germany, June 1994.

[GH97] Simon Gibson and R. J. Hubbold. Perceptually driven radiosity. *Computer Graphics Forum*, 16(2):119–128, June 1997.

[GP91] Eduard Groller and Werner Purgathofer. Using temporal and spatial coherence for accelerating the calculation of animation sequences. In *Eurographics '91*, pages 103–13. Elsevier Science Publishers, Amsterdam, North-Holland, September 1991.

[GSG90] David W. George, Francois X. Sillion, and Donald P. Greenberg. Radiosity Redistribution for Dynamic Environments. *IEEE Computer Graphics and Applications*, 10(4):26–34, July 1990.

[HSA91] Pat Hanrahan, David Salzman, and Larry Aupperle. A Rapid Hierarchical Radiosity Algorithm. In *Computer Graphics (ACM SIGGRAPH '91 Proceedings)*, volume 25, pages 197–206, July 1991.

[HZ82] H. Hubshman and S.W. Zucker. Frame-to-frame coherence and the hidden surface computation: Constraints for a convex world. *ACM Transaction on Graphics*, 1(2):129–162, April 1982.

[Mar94] Ignacio Martín. Analisis de la coherencia temporal en escenas dinamicas iluminadas con radiosidad. Technical Report 94-18-T, LSI, Universitat Politecnica de Catalunya, October 1994.

[MPT97] Ignacio Martin, Xavier Pueyo, and Dani Tost. An image-space refinement criterion for linear hierarchical radiosity. In *Graphics Interface '97*, pages 26–36, San Francisco, CA, May 1997. Morgan Kaufmann. Available from http://www.dgp.toronto.edu/gi/gi97/proceedings/papers.

[MPT98] Ignacio Martin, Xavier Pueyo, and Dani Tost. A Two-Pass Hardware based method for Hierarchical Radiosity. Technical report, University of Girona, 1998. To be published in Eurographics 98.

[MS94] Stefan Muller and Frank Schoffel. Fast Radiosity Repropagation for Interactive Virtual Environments Using a Shadow-Form-Factor-List. In *Fifth Eurographics Workshop on Rendering*, pages 325–342, Darmstadt, Germany, June 1994.

[NDR96] Jeffry Nimeroff, Julie Dorsey, and Holly Rushmeier. Implementation and analysis of an image-based global illumination framework for animated environments. *IEEE Transactions on Visualization and Computer Graphics*, 2(4):283–298, December 1996.

[ORDP96] Rachel Orti, Stephane Riviere, Fredo Durand, and Claude Puech. Radiosity for Dynamic Scenes in Flatland with the Visibility Complex. In *Computer Graphics Forum*, volume 15, pages C237–C248, September 1996.

[SAG94] Brian Smits, James Arvo, and Donald Greenberg. A Clustering Algorithm for Radiosity in Complex Environments. In *Computer Graphics Proceedings, Annual Conference Series, 1994 (ACM SIGGRAPH '94 Proceedings)*, pages 435–442, 1994.

[Sha97] Erin Shaw. Hierarchical radiosity for dynamic environments. *Computer Graphics Forum*, 16(2):107–118, June 1997.

[Tos91] Dani Tost. An algorithm of hidden surface removal based on frame-to-frame coherence. In *Eurographics '91*, pages 261–273. Elsevier Science Publishers, Amsterdam, North-Holland, September 1991.

Departament de Llenguatges i Sistemes Informàtics
Universitat Politècnica de Catalunya

Research Reports – 1998

LSI–98–1–R "Optimal Sampling Strategies in Quicksort and Quickselect", Conrado Martinez, Salvador Roura.

LSI–98–2–R "Query, PACS and simple-PAC Learning", J. Castro and D. Guijarro.

LSI–98–3–R "Interval Analysis Applied to Constraint Feasibility in Geometric Constraint Solving", R. Joan-Arinyo and N. Mata.

LSI–98–4–R "BayesProfile: application of Bayesian Networks to website user tracking", Ramón Sangüesa and Ulises Cortés.

LSI–98–5–R "Some reflections on applying Workflow Technology to Software Process", Camilo Ocampo and Pere Botella.

LSI–98–6–R "Surface Fairing for Ship Hull Design Application", P. Brunet, A. Vinacua, M. Vivó, N. Pla, A. Rodríguez.

LSI–98–7–R "Trust Values for Agent Selection in Multiagent Systems", Karmelo Urzelai.

LSI–98–8–R "The use of SAREL to control the correspondence between Specification Documents", Núria Castell and Àngels Hernández.

LSI–98–9–R "Intervalizing colored graphs is NP-complete for caterpillars with hair length 2", C. Àlvarez, J. Diaz and M. Serna.

LSI–98–10–R "A unified approach to natural language treatment", Jordi Àlvarez.

LSI–98–11–R "Collision Detection: Models and Algorithms", Marta Franquesa and Pere Brunet.

LSI–98–12–R "Height-relaxed AVL rebalancing: A unified, fine-grained approach to concurrent dictionaries", Luc Bougé, Joaquim Gabarró, Xavier Messeguer and Nicolas Schabanel.

LSI–98–13–R "HyperChromatic trees: a fine-grained approach to distributed algorithms on RedBlack trees", Xavier Messeguer and Borja Valles.

LSI–98–14–R "Asynchronous Interface Specification, Analysis and Synthesis", Michael Kishinevsky, Jordi Cortadella, Alex Kondratyev and Luciano Lavagno.

LSI–98–15–R "Heuristics for the MinLA Problem: Some Theoretical and Empirical Considerations", Josep Diaz, Jordi Petit i Silvestre and Paul Spirakis.

LSI-98-16-R "Sampling matchings in parallel", Josep Diaz, J. Petit i Silvestre, María Jose Serna and Paul Spirakis.

LSI-98-17-R "The Parallel Approximability of the False and True Gates Problems for Nor Circuits", M. Serna and F. Xhafa.

LSI-98-18-R "Basic Geometric Operations in Ruler-and-Compass Constraint Solvers using Interval Arithmetic", R. Joan-Arinyo and N. Mata.

LSI-98-19-R "HDM: AN HETEROGENEOUS STRUCTURES DEFORMATION MODEL", Montse Bigordà and Dani Tost.

LSI-98-20-R "Visualization of Cerebral Blood Vessels", Anna Puig.

LSI-98-21-R "Cerebral Blood Vessels Modelling", Anna Puig.

LSI-98-22-R "Discrete Medial Axis Transform for Discrete Objects", Anna Puig.

LSI-98-23-R "Incorporating the Behavioural Information to the Schema Construction Processs of Federated Data Bases System", Luis Carlos Rodríguez G.

LSI-98-24-R "Del Texto a la Información", J. Atserias, N. Castell, N. Català, H. Rodríguez and J.Turmo.

LSI-98-25-R " Construcción automática de diccionarios de patrones de extracción de información", Neus Català and Núria Castell.

LSI-98-26-R "Syntactic Connectivity", Glyn Morrill.

LSI-98-27-R "Geometric Distance Constraint Satisfaction by Constraint-to-constraint relaxation", Lluis Solano Albajes and Pere Brunet Crosa.

LSI-98-28-R "Modelling Surfaces from Planar Irregular Meshes", Josep Cotrina Navau and Nuria Pla Garcia.

LSI-98-29-R "Computing Directional Constrained Delaunay Triangulations", Marc Vigo Anglada .

LSI-98-30-R "On the complexity of moving vertices in a graph", Antoni Lozano (Universitat Politècnica de Catalunya) and Vijay Raghavan (Vanderbilt University).

LSI-98-31-R "Defining and Translating Visual Schemas for Deductive Databases", Jordi Puigsegur, Joan A. Pastor, and Jaume Agustí.

LSI-98-32-R "A framework for Animation in Global Illumination Environments", I. Martin, X. Pueyo and D. Tost.

LSI-98-33-R "Fuzzy Heterogeneous Neurons for Imprecise Classification Problems", J. Valdés and Ll. Belanche.

LSI-98-34-R "Query Containment Checking as a View Updating Problem", C. Farr/e, E. Teniente and T. Urpí.

LSI-98-35-R "A Computational Characterization of Collective Chaos", Jordi Delgado and Ricard V. Solé.

LSI-98-37-R "Fringe analysis of synchronized parallel algorithms on 2–3 trees", R. Baeza-Yates, J. Gabarró and X. Messeguer.

LSI-98-38-R "Invocació explícita versus invocació implícita: Anàlisi comparativa de dos enfocaments de disseny de Sistemes d'Informació", Cristina Gómez, Luis Felipe Fernández, Juan Ramón López and Antoni Olivé.

Hardcopies of reports can be ordered from:

Núria Sánchez
Departament de Llenguatges i Sistemes Informàtics
Universitat Politècnica de Catalunya
Campus Nord, Mòdul C6
Jordi Girona Salgado, 1-3
08034 Barcelona, Spain
secrelsi@lsi.upc.es

See also the Department WWW pages, http://www-lsi.upc.es/