EIGEN-ANALYSIS FOR RADIOSITY SYSTEMS

Gladimir V. G. BARANOSKI[1]

Randall BRAMLEY[2]

Jon G. ROKNE[1]

Department of Computer Science, The University of Calgary, Calgary, AB, CANADA[1]
Computer Science Department, Indiana University, Bloomington, IN, USA[2]
gbaranos@cpsc.ucalgary.ca

ABSTRACT

The convergence of iterative methods used to solve the linear equations arising from radiosity systems depends on the distribution of the eigenvalues of the radiosity coefficient matrix. In this paper we prove that all eigenvalues of the radiosity coefficient matrix are real and positive. This fact may allow us to obtain fast radiosity solutions using the knowledge about the spectrum of the matrix. Moreover, the physical meaning of the eigenvectors in global illumination applications is an open problem in graphics. In order to contribute to the clarification of this question, we present some experiments based on the theory of matrices, in which we show interesting features that arise when the eigenvectors are used as solution vectors in graphics settings.

Key Words: eigenvalue, eigenvector, spectral analysis, radiosity.

INTRODUCTION

Determining the radiant exitance (or radiosity) of each patch in a closed environment involves solving, either implicitly (through radiosity-specific methods[1]) or explicitly (through general matrix methods like Gauss-Seidel [Cohen-Greenberg85]) the linear system $Gb = e$, in which G represents the radiosity coefficient matrix, b represents the vector of unknowns and e represents the vector of emittances. Because of the large size of these linear systems and the relatively low accuracies required in the solutions b, iterative methods are commonly used.

Why is it relevant to search for faster iterative methods to solve the radiosity system of equations when the most expensive stage of the radiosity pipeline is the calculation of form factors? The reason is that there are many practical applications in which a radiosity system has to be solved repeatedly, due to changes either to the vector of reflectances or to the vector of emittances, while the form factors remain unchanged. An example of such an application would be the simulation of the incidence of sun light over a building during a day. In this case the form factors would have to be computed only once, while the system would have to be solved 720 times to account for solar positions at every minute.

The convergence properties of iterative methods rely extensively on the spectrum or set of eigenvalues of the coefficient matrix G. An *eigenvector* ν of a matrix G is a nonzero vector that does not rotate when G is applied to it. In other words, there is some scalar constant λ, an *eigenvalue* of G, such that $G\nu = \lambda\nu$. Every square matrix G of order n has n possibly nondistinct complex eigenvalues $\lambda_1, \lambda_2, ..., \lambda_n$. When G is symmetric the eigenvalues are real-valued. The

[1] We use the expression *radiosity-specific methods* to group methods specifically developed to solve the radiosity problem, although those methods can be considered variations of numerical methods such as Southwell Iteration [Goertler *et al.* 94] or SOR [Hageman-Young81].

set $\sigma(G)$ of eigenvalues of a matrix is called its spectrum.

It has already been shown [Baranoski *et al.* 95a] [Neumann94] that we can obtain relatively inexpensive estimates of the eigenvalues of G using the Gerschgorin Circle Theorem [Burden-Faires93]. Moreover, it has also been shown [Baranoski *et al.* 95a] that high reflectivities cause a larger number of interreflections, causing the eigenvalues to become more spread out, which in turn slows down the convergence of the iterative methods. However, it has not yet been shown analytically that all the eigenvalues of the radiosity coefficient matrix G are *real* and *positive*, which we intend to do in next section. In theory one could apply any iterative method to solve a radiosity system of equations. However, this proof guarantees that fast methods, like Chebyshev [Baranoski *et al.* 95b] and Conjugate Gradient [Neumann94, Shewchuck94] among others, can be applied with confidence of their convergence.

Furthermore, Arvo [Arvo95a, Arvo95b] has shown that several fundamental operators that arise in global illumination can be uniformly approximated by matrices. Then, if one can determine what the eigenvectors of a global illumination matrix, like the radiosity matrix G, represent in terms of the physical application, it may be possible to obtain accurate approximations of them. These eigenvectors in turn could be used to obtain low rank approximations of those matrices using SVD type approaches [Golub-Loan89].

The problem is that the physical meaning of the eigenvectors and eigenvalues is unknown in global illumination to date. In other fields, on the other hand, researchers know their meaning, and use them not only in analysis of problems, but also in practical applications. For instance, as mentioned by Schewchuk [Shewchuck94], the eigenvectors of the stiffness matrix associated with a discretized structure of uniform density represent the natural modes of vibration of the structure being studied, and the corresponding eigenvalues define the natural frequencies of vibration [Lancaster-Tismenetsky85].

Since for physically meaningful problems the inverse of the matrix G has all positive entries, the Perron-Frobenius theorem [Blackwell61, Lancaster-Tismenetsky85] implies that all components of the eigenvector corresponding to the smallest eigenvalue of G are strictly of the same sign. Thus this eigenvector can be scaled and interpreted as a vector of radiosities. We explore some of the implications of this fact, and show interesting features of using eigenvectors as solution vectors in graphics settings in order to determine their physical meaning in the radiosity context.

EIGENVALUES OF THE RADIOSITY COEFFICIENT MATRIX

The transpose of an $n \times n$ matrix $G = g(ji)$ is the matrix $G^t = g(ij)$. A square matrix G is said to be symmetric if $G = G^t$. Initially, to prove that all eigenvalues of the radiosity matrix G are real and positive, consider that it can be made symmetric by scaling its rows:

$$G^s = DG \tag{1}$$

where D is the diagonal matrix in which the diagonal entry d_{ii} is the quotient of the area and the reflectivity of patch i.

Since DG is symmetric, its eigenvalues are real-valued. Moreover, by applying the Gerschgorin Circle Theorem [Burden-Faires93], one can verify that they are also positive. Hence DG is a positive definite matrix [Burden-Faires93]. The definition of positive definite means that

$$x^H DGx > 0 \tag{2}$$

for all $x \in C$, where C is the complex plane and x^H is the Hermitian transpose of the vector x.

Let x be an eigenvector of G and λ be an eigenvalue. Then

$$Gx = \lambda x, \tag{3}$$

where λ and $x \neq 0$ are possibly complex.

Then:

$$DGx = \lambda Dx \tag{4}$$

and

$$x^H DGx = \lambda x^H Dx \tag{5}$$

The left side of Equation (5) is necessarily real and positive from Equation (2). Furthermore, the definition of an eigenvector x implies that it is nonzero, hence

$$x^H Dx = \sum \overline{x}_i x_i d_i \geq d_{min} \sum x_i \overline{x}_i = d_{min} x^H x = d_{min} \parallel x \parallel > 0 \quad (6)$$

Equations (5) and (6) imply that λ is real and $\lambda > 0$. Therefore all of the eigenvalues of G are real and positive.

Although it may seem that this proof applies directly to the continuous radiosity operator, more details need to be considered. The critical point is that the diagonal entry i of the matrix D is the ratio of the area and reflectivity of the i-th patch. For the continuous case, the area is zero and the above argument cannot be used directly. However, the continuous operator is a compact operator [Arvo95b], sometimes also called a completely continuous operator in functional analysis. Because of this we can construct a sequence of finite dimensional operators G_k that converge uniformly to the continuous operator G_∞, each with a spectrum consisting only of positive real eigenvalues. That sequence can be constructed using a sequence of uniformly refined discretizations of the scene, for example. The limit G_∞ will necessarily have a spectrum that is real and nonnegative. Because G_∞ is nonsingular [Kajiya85], it cannot have zero as an eigenvalue. Its compactness also implies that it has only a point spectrum, which implies that G_∞ has only positive real eigenvalues in its spectrum.

IMPLICATIONS OF THE PERRON-FROBENIUS THEOREM

Any matrix that can be expressed in the form:

$$A = sI - B, \quad s > 0, \quad B \geq 0 \quad (7)$$

for which $s \geq \rho(B)$, the spectral radius [2] of B, is called an M-matrix [Berman-Plemmons87].

Recall that the matrix G can be represented by:

$$G = I - F \quad (8)$$

where F is the scaled form factor matrix. When a scene has no concave patches, the diagonal entries of the matrix G are ones and the off-diagonal entries are the negatives of products of the form factors and the corresponding reflectivities. When reflectivities are less than 1.0 in value, the summation of those products in any row is necessarily less than one. This implies that the spectral radius of F is less than one. Therefore G is an M-matrix.

Since G is a nonsingular M-matrix and $\rho(F) < 1$ (which can be proved by applying the Gerschgorin Circle Theorem [Burden-Faires93] to F), the matrix version of the Neumann lemma [Berman-Plemmons87] for convergent series gives:

$$G^{-1} = (I - F)^{-1} = \sum_{k=0}^{\infty} F^k = I + F + F^2 + F^3 + \cdots \quad (9)$$

This implies that the inverse G^{-1} of G is thus a positive matrix, having all positive components. In other words, G^{-1} is a nonnegative matrix. Furthermore, since G is irreducible [Blackwell61, Berman-Plemmons87, Lancaster-Tismenetsky85] (because the visibility of patch i from patch j implies the reverse), G^{-1} is also irreducible.

The Perron-Frobenius theorem, which concerns square irreducible nonnegative matrices, implies that the largest eigenvalue $\zeta_{\max}$ of G^{-1} is real and positive (as we have already established), and the corresponding eigenvector ν of G^{-1} has all components strictly of the same sign [Blackwell61, Lancaster-Tismenetsky85]. Moreover, if we multiply the relationship:

$$G^{-1}\nu = \zeta_{\max}\nu \quad (10)$$

through by G we get $\nu = \zeta_{\max} G\nu$, or

$$G\nu = \lambda_{\min}\nu \quad (11)$$

where $\lambda_{\min} = 1/\zeta_{\max}$. Hence the eigenvector of G corresponding to its smallest eigenvalue can be scaled in such way that all its components are positive, and so can be physically interpreted as a radiosity vector. In the next section we show some images obtained this way, and try to give a

[2] The spectral radius $\rho(C)$ of a matrix C is defined by $\rho(C) = max \mid \lambda \mid$, where λ is an eigenvalue of C [Burden-Faires93].

physical interpretation to them.

The Perron-Frobenius theorem has further implications. Let $x > 0$ mean the vector x is nonnegative and at least one component is positive. Then the largest eigenvalue of G^{-1} is equal to the maximum over all $x > 0$, of the minimum over all positive components x_i, of $\frac{(G^{-1}x)_i}{x_i}$. Here $(G^{-1}x)_i$ means the i-th component of the vector $G^{-1} * x$. The minimum and maximum can be reversed in order.

G is a monotone operator, that is, $Gx \geq 0$ implies $x \geq 0$. This means if the right hand side vector of the linear system is nonnegative, then the solution must be nonnegative. Of course this also follows directly from the physical application! Using the above, it may be possible to provide an *upper* bound on the smallest eigenvalue of G, thereby analytically proving the spreading of eigenvalues as the average reflectance of the environment increases.

EXPERIMENTS WITH EIGENVECTORS

Experiments Set-Up

The test model used in our experiments is shown in Figure 1a, and consists of a sphere centered in a cube (r=1.0, d=3.0, h=6.0). The sphere is divided into 128 patches and each of the faces of the surrounding cube are divided into 144 patches summing up a total of 992 patches. In our experiments we use a classical radiosity approach [Glassner95] in which all surfaces are assumed to be perfect diffuse reflectors.

We use two different sets of parameters to allow a wider range of observations. In set 1 we use one light source corresponding to 16 patches, with an emittance equal to 10.0. In set 2 we use two light sources, each corresponding to 16 patches and with emittance equal to 5.0. The reflectivities of the light sources are set to 0.1, and the reflectivities of the other surfaces are presented in Figure 1b. The form factors are computed using the PDM method described in [Baranoski92]. The eigenvalues and eigenvectors are computed using *MATLAB* [MathWorks94] through

a)

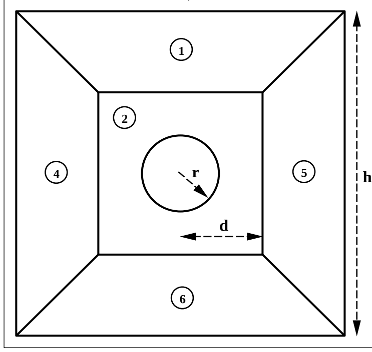

b)

surface	set 1	set 2
face 1	0.90	0.90
face 2	0.75	0.75
face 3	0.75	0.75
face 4	0.85	0.85
face 5	0.85	0.85
face 6	0.65	0.65
sphere	0.65	0.55

Figure 1: a) Sketch of the test model. b) Sets of reflectivities used in the experiments.

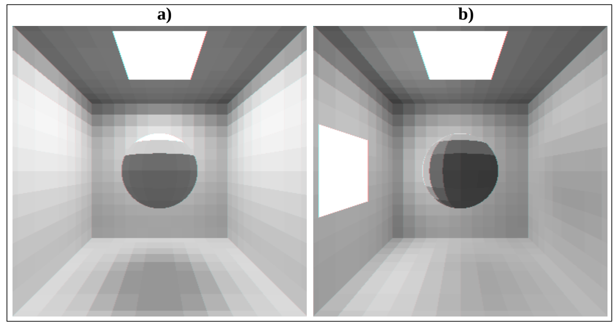

Figure 2: Images corresponding to the solution (radiosity) vectors of the linear systems regarding: a) set 1 and b) set 2.

the QR method [Burden-Faires93] whose algorithms are provided by the *EISPACK* routines [Smith *et al.* 76].

The images presented in this paper are rendered using flat shading and greyscale to allow a better detection of the features associated with the eigenvector components. Figure 2 shows the images corresponding to the solution (radiosity) vectors of the radiosity systems of linear equations regarding the two sets of parameters, which are solved using the Chebyshev method [Baranoski *et al.* 95a]. These images are used in our experiments as reference images to be compared with images obtained using eigenvectors as solution vectors.

Eigenvectors Corresponding to the Smallest Eigenvalue of the Radiosity Matrices

Figure 3 shows the eigenvectors ν_S corresponding to the smallest eigenvalues of the radiosity matrices G for two different choices of config-

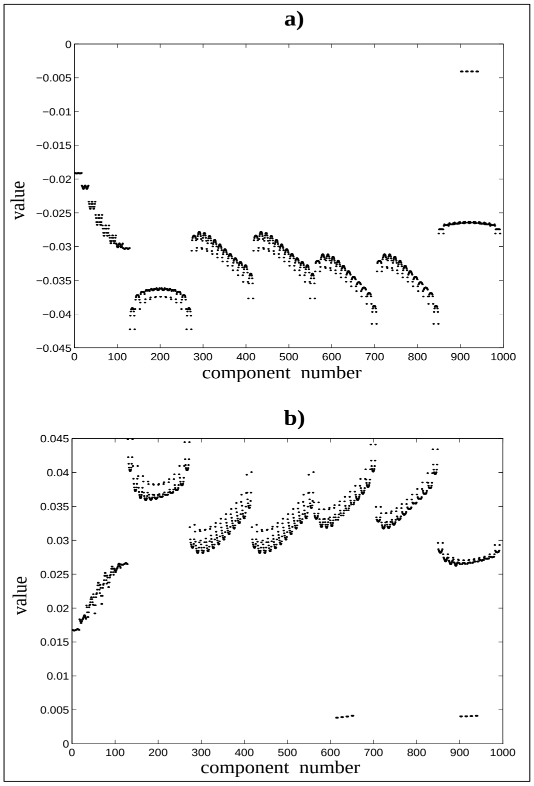

Figure 3: Eigenvectors ν_S of the matrices G associated with: a) set 1 and b) set 2.

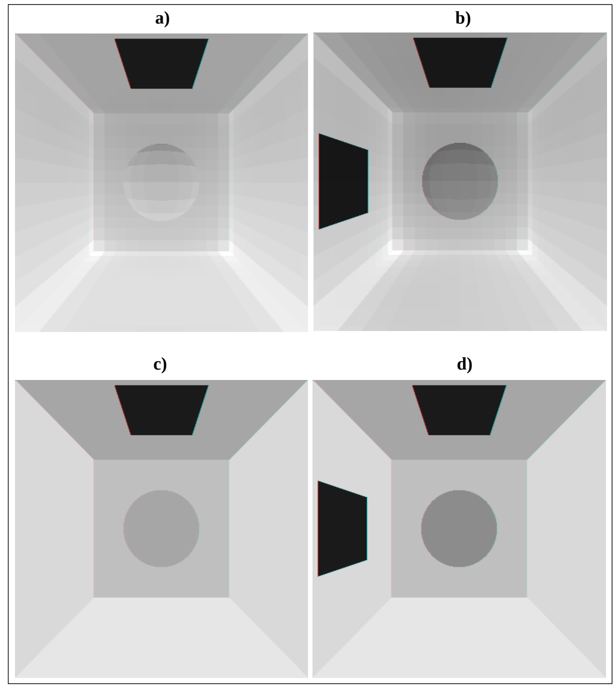

Figure 4: Images obtained using as solution vectors the eigenvectors ν_S of the matrices G associated with: a) set 1 and b) set 2; and images obtained using as solution vectors the vectors of reflectivities associated with: c) set 1 and d) set 2.

uration and/or parameters. As expected from the Perron-Frobenius theorem, their components have all the same sign. We can notice seven distinguished groups of points which are associated, from left to right, with the sphere and the six faces. The eigenvector components corresponding to the patches with lowest reflectivities, which in our experiments correspond to the emitter patches, are represented by the points with lowest absolute values.

After taking the absolute values and normalizing these eigenvectors (by dividing all their components by the absolute value of their largest component), we use them as solution vectors to display the images of the scenes (Figure 4, top row). The features presented in these images seem to be associated with the distribution of the reflectivities in the scenes (Figure 4, bottom row). In Figure 4a we can notice that the top of the sphere, which is closer to an area with low reflectivity, is darker than its bottom. In Figure 4b, areas of the scene that are directly exposed to the luminaires are also darker, possibly due to the low reflectivities assigned to the emitter patches. These features seem to indicate that the absolute values of components of the eigenvectors ν_S are directly proportional to the reflectivity values of the corresponding patches. Furthermore, they are also associated with the direct interaction of reflectivities. In other words, patches directly exposed to areas with low reflectivity, shown as darker areas in the images, correspond to components of the eigenvectors ν_S with low absolute values (assuming normalized eigenvectors).

Eigenvectors Corresponding to the Largest Eigenvalue of the Symmetric Radiosity Matrices

A SVD (single value decomposition) type approach [Golub-Loan89] may be used to provide a low rank approximation for the symmetric matrices G^s. A low rank approximation of G^s is given by:

$$G^s_p = \lambda_n \nu_n \nu_n^{'} + \lambda_{n-1}\nu_{n-1}\nu_{n-1}^{'} + ... + \lambda_{n-p}\nu_{n-p}\nu_{n-p}^{'} \quad for \quad p \leq n-2 \quad (12)$$

where $\nu_1', \nu_2', ..., \nu_n'$ correspond to the transposes of the eigenvectors of G^s. For a symmetric

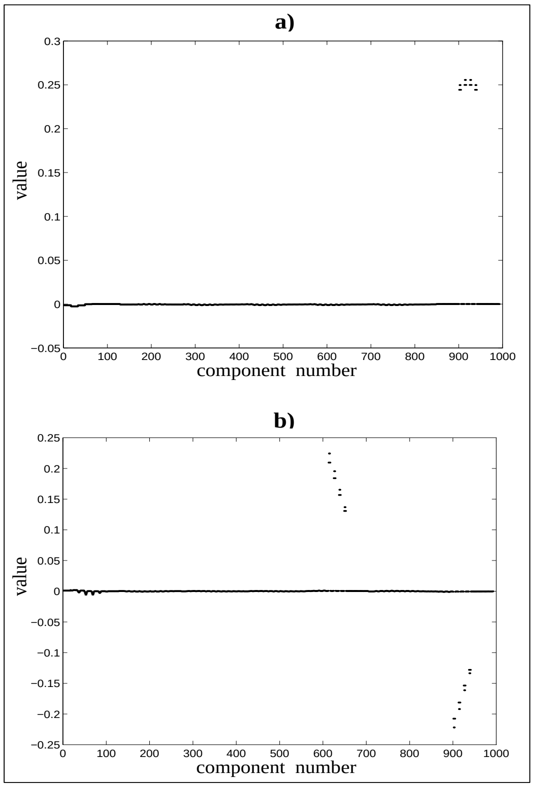

Figure 5: Eigenvectors ν_L of the matrices G^s associated with: a) set 1 and b) set 2.

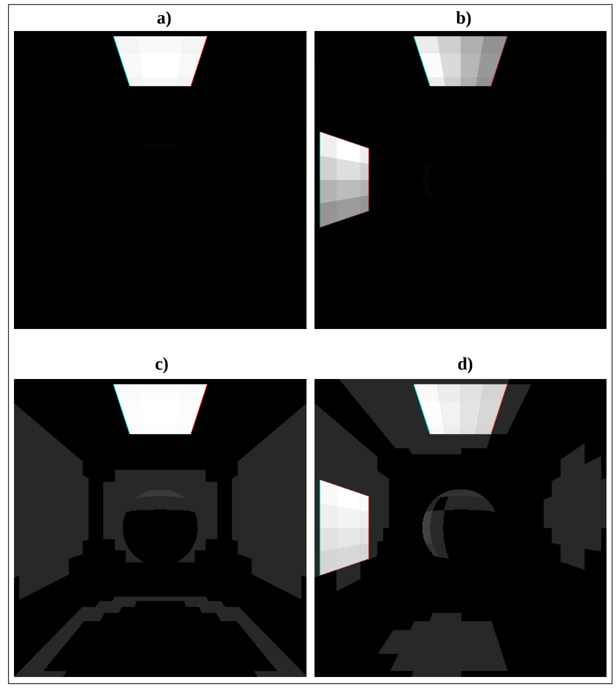

Figure 6: Images obtained using as solution vectors the eigenvectors ν_L of the matrices G^s associated with: a) set 1 (with $\gamma = 1.0$), b) set 2 (with $\gamma = 1.0$), c) set 1 (with $\gamma = 2.2$) and d) set 2 (with $\gamma = 2.2$).

matrix the SVD is the same as the eigenvalue-eigenvector decomposition. The principal components are then the eigenvectors, corresponding to the largest eigenvalues, ν_L. Because of this, we decided to extend our investigation to the eigenvectors ν_L of the matrices G^s (Figure 5).

After taking the absolute values and normalizing the eigenvectors ν_L, we use them as solution vectors to display the images regarding the two sets of parameters. As one would expect looking at the plots of eigenvectors (Figure 5a and 5b), the images (Figure 6a and 6b) are almost completely dark, with the exception of the emitter patches. These images were displayed using a Gamma correction function [Foley *et al.* 90] provided by *XV* [Bradley94] in which the Gamma value, γ, is set to 1.0. When we increase this value to $\gamma = 2.2$, some interesting features appeared (Figure 6c and 6d). These features seem to be related with the paths of direct light propagation. They also show that there is more useful information associated with the components of the eigenvectors ν_L than the almost straight lines in the plots of Figure 5 indicate. Furthermore, on face 2 (see Figure 1) of the image presented in Figure 6d there is no sign of any feature associated with the paths of direct light propagation, as one would expect since that surface is also exposed to the luminaires.

Figure 7 presents a zoom in of the components of the eigenvectors ν_L associated with the non emitter patches, and reveals the patterns associated with the features presented in Figures 6c and 6d. To analyze the physical meaning of these patterns more closely, we set the components of the eigenvectors ν_L associated with the emitter patches to 1.0, and took the absolute values and normalized the remaining ones. The resulting images, presented in Figures 8a and 8b, are very close to the solution images (Figures 2a and 2b). Increasing the value of γ from 1.0 to 2.2, which has the effect of increasing the brightness of the scenes, we can notice that the similarities with the solution images become even more evident.

Where does the association with the paths of

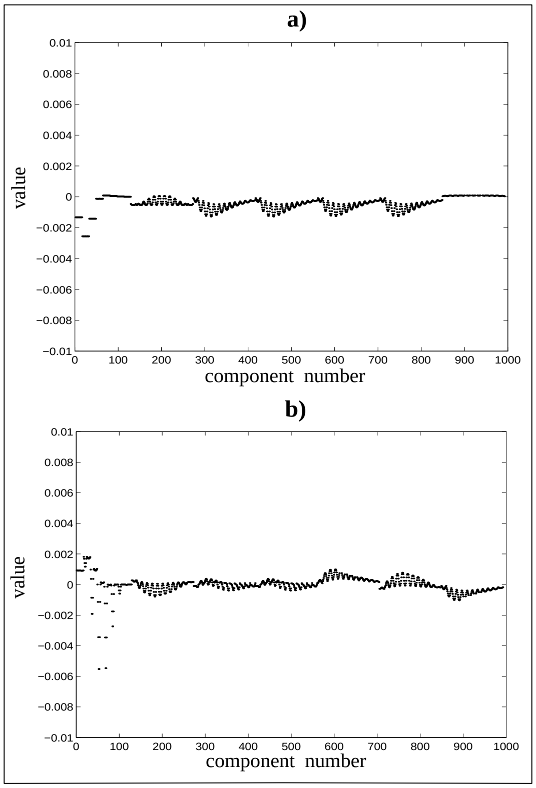

Figure 7: Zoom in of the components of the eigenvectors ν_L regarding the non emitter patches and associated with a) set 1 and b) set 2.

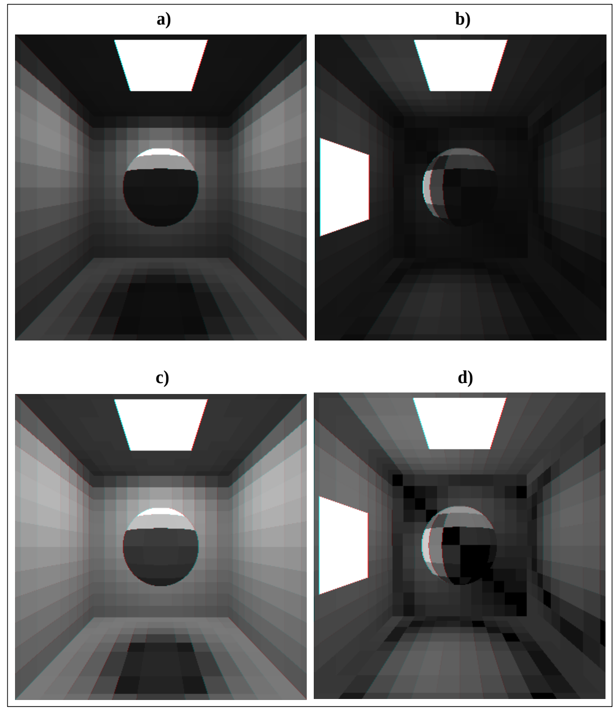

Figure 8: Images obtained using as solution vectors the adjusted versions of the eigenvectors ν_L regarding the matrices G^s associated with: a) set 1 (with $\gamma = 1.0$), b) set 2 (with $\gamma = 1.0$), c) set 1 (with $\gamma = 2.2$) and d) set 2 (with $\gamma = 2.2$).

direct light propagation comes from? Looking at the graphs presented in Figures 5 and 7, we can notice that the absolute values of the components of the eigenvectors ν_L of G^s are inversely proportional to the reflectivity values of the corresponding patches. Moreover, the components with the highest absolute values, henceforth called dominants, are associated with the patches with the lowest reflectivity in the environment, which correspond in our experiments to the emitter patches. The next components with high absolute values are those whose corresponding patches are directly exposed to the patches associated with the dominant components (Figures 6c and 6d).

However, if we assign different reflectivity values to the emitter patches, such that they no longer correspond to the dominant components of the eigenvectors ν_L of G^s (Figure 9), the association with the paths of direct light propagation can not be established. This aspect is illustrated in the images presented in Figure 10 where we set the reflectivity of the emitter patches to 0.9.

In this case, the dominant components will correspond to the sphere patches, since they now present the lowest reflectivities in both scenes. Furthermore, the scene in which we assign a lower reflectivity value for the sphere (Figure 10b) presents a higher brightness than the other one (Figure 10a).

To further investigate the relationship between the components of the eigenvectors ν_L and the reflectivities of the patches, we define vectors Q, in which the entry q_i corresponds to the diagonal entry d_i of D. After normalizing these vectors we use them to display the images regarding the two sets of parameters. Comparing these images (Figure 11) with the previous ones (Figure 10), we can notice a similar color gradation on the spheres, especially the dark spots on the top and on the bottom. This aspect suggests that the components of the eigenvectors ν_S of the matrices G^s are not only associated with the direct interaction of reflectivities, but they are also associated with the areas of the patches. This additional dependency comes from the symmetriza-

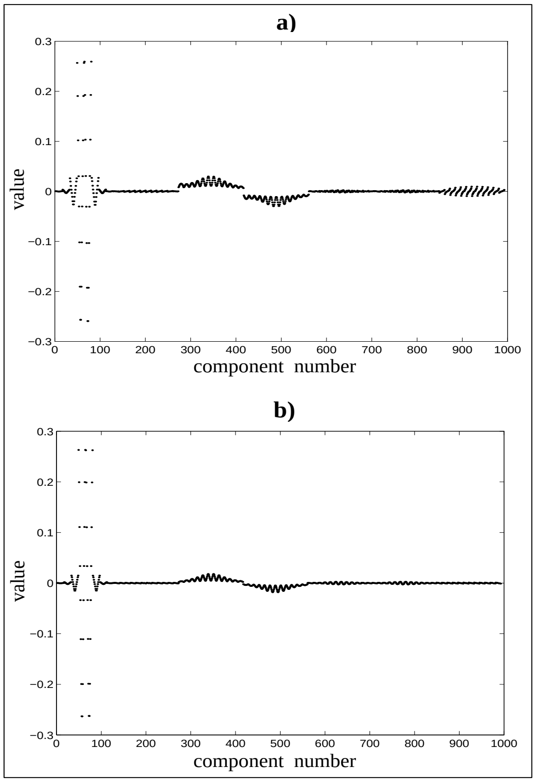

Figure 9: Eigenvectors ν_L of the matrices G^s (with the reflectivity of the emitter patches set to 0.9) associated with: a) set 1 and b) set 2.

tion process which scales the diagonal elements of G by the diagonal entries of D.

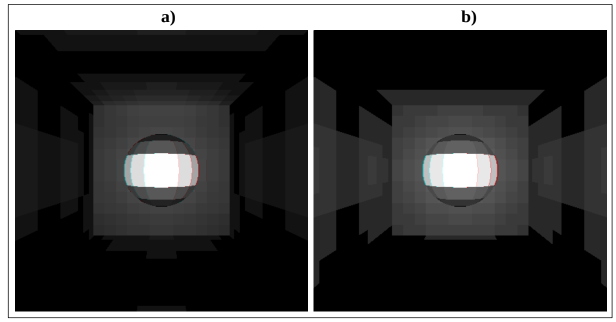

Figure 10: Images obtained using as solution vectors the eigenvectors ν_L of the matrices G^s (with the reflectivity of the emitter patches set to 0.9) associated with: a) set 1 (with $\gamma = 2.2$) and b) set 2(with $\gamma = 2.2$).

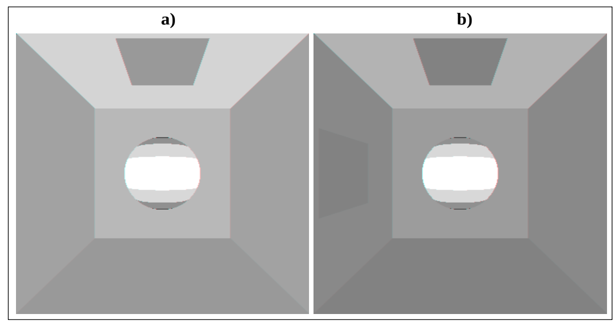

Figure 11: Images obtained using as solution vectors the vectors Q associated with: a) set 1 and b) set 2.

SUMMARY AND FUTURE WORK

Arvo [Arvo95b] has suggested that functional analysis might be a useful tool for providing a better theoretical foundation for global illumination. This suggestion is followed in this paper in the form of an investigation of the spectral properties of the radiosity matrix. We prove that all the eigenvalues of this matrix are real and positive. We believe that this proof, which is directly related to the spectral analysis of the equations governing the transport of radiant energy in global illumination, is essential for the application of fast iterative solvers to the radiosity systems. We also point out that the Perron-Frobenius theorem may be used to prove analytically the spreading of the eigenvalues as the brightness of the scene increases. This issue will be addressed in the next stage of our research.

In this paper we also show some interesting features of using the eigenvectors corresponding to the smallest and the largest eigenvalues of radiosity matrices and their symmetric versions as solution vectors in graphics settings. While these features provide evidence that there is potentially useful information related to these eigenvectors, more research is needed to gain a fuller understanding of their physical meaning. We intend to proceed with this investigation that, we believe, may lead us to faster global illumination solutions.

ACKNOWLEDGMENTS

We would like to thank Dr. Peter Shirley for his valuable support during the early stages of this research, Dr. James R. Arvo and Dr. Wayne Eberly for their comments regarding the proof presented in this paper, Dr. Richard M. Levy for his feedback regarding possible applications for our work and the anonymous reviewers for their helpful suggestions. The work presented in this paper was supported in part by the Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq, Proc. 200876/93-7, Brazil) and by the Natural Sciences and Engineering Research Council of Canada (NSERC, Grant 69-1266, Canada).

REFERENCES

Arvo, J. 1995 *Analytic Methods for Simulated Light Transport* PhD thesis, Yale University.

Arvo, J. 1995 The role of functional analysis in global illumination In *Rendering Techniques'95 (Proceedings of the Sixth Eurographics Rendering Workshop)*, pages 115–126.

Baranoski, G., Bramley, R., Shirley, P. 1995 Fast radiosity solutions for environments with high average reflectance In *Rendering Techniques'95 (Proceedings of the Sixth Eurographics Rendering Workshop)*, pages 345–356.

Baranoski, G., Bramley, R., Shirley, P. 1995 Iterative methods for fast radiosity solutions Technical report, Indiana University.

Baranoski, G. 1992 The parametric differential method: An alternative to the calculation of form factors, *Computer Graphics Forum (Eurographics'92 Proceedings)*, Vol. 11, no 3, pp 193–204.

Berman, A., Plemmons, R. 1987 *Nonnegative Matrices in the Mathematical Sciences* Academic Press, New York.

Blackwell, D. 1961 Minimax and irreducible matrices, *Journal of Mathematical Analysis and Applications*, Vol. 3, pp 37–39.

Bradley, J. 1994 *XV - Interactive Image Display for the X Window System* Bryn Mawr, PA, USA Version 3.10a.

Burden, R., Faires, J. 1993 *Numerical Analysis* PWS-KENT Publishing Company, Boston, 5 edition.

Cohen, M., Greenberg, D. 1985 The hemicube: A radiosity solution for complex environments, *Computer Graphics (SIGGRAPH Proceedings)*, Vol. 19, no 3, pp 31–40.

Foley, J., Dam, A.van, Feiner, S., Hughes, J. 1990 *Computer Graphics: Principles and Practice* Addison-Wesley Publishing Company, Reading, Massachusetts, 2 edition.

Glassner, A. 1995 *Principles of Digital Image Synthesis* Morgan Kaufmann, San Francisco, 1 edition.

Goertler, S., Cohen, M., Slusallek, P. 1994 Radiosity and relaxation methods, *IEEE Computer Graphics and Applications*, Vol. 14, no 6, pp 48–58.

Golub, G., Loan, C. V. 1989 *Matrix Computations* John Hopkins University Press, Baltimore, 2 edition.

Hageman, L., Young, D. 1981 *Applied Iterative Methods* Academic Press, New York, 1 edition.

Kajiya, J. 1985 The rendering equation, *Computer Graphics (SIGGRAPH Proceedings)*, Vol. 20, no 4, pp 143–150.

Lancaster, P., Tismenetsky, M. 1985 *The Theory of Matrices* Academic Press, San Diego, 2 edition.

MathWorks, 1994 *MATLAB User's Guide* Natick, MA, USA Version 4.2c.

Neumann, L. 1994 New efficient algorithms with positive definite radiosity matrix In *Proc. of the Fifth Eurographics Rendering Workshop*, pages 219–237.

Shewchuck, J. 1994 An introduction to the conjugate gradient method without the agonizing pain Technical report, Scholl of Computer Science, Carnegie Mellon University.

Smith, B., Boyle, J., Dongarra, J., Garbow, B., Ikebe, Y., Klema, V., Moler, C. 1976 *Matrix Eigensystem Routines - EISPACK Guide*, volume 6 of *Lecture Notes in Computer Science* Springer-Verlag, 2 edition.