Speeding Up Progressive Radiosity by Overshooting

Martin Feda
Technical University of Vienna
Institute of Computer Graphics
Karlsplatz 13/186-2
Vienna, Austria
feda@cg.tuwien.ac.at

Abstract

Overshooting techniques have proven to significantly speed up the convergence of radiosity computation. Similar methods (overrelaxation) have also been widely used in numerical mathematics to speed up iterative solution methods for linear equation systems, e.g. Gauss-Seidel iteration. This paper gives a comprehensive description of all overshooting techniques. Advantages and disadvantages of the different techniques are also described.

1 Introduction

Overshooting is a simple technique to accelerate the progressive refinement radiosity convergence. Standard progressive refinement radiosity distributes the self-emitted or received unshot radiosity ΔB_i of shooting patches to the environment [Cohe 88]. ΔB_i is the difference between the radiosity already shot from the patch i in earlier iterations, and the current value of the radiosity B_i. After shooting its radiosity, a patch will receive additional light from other patches, which has to be distributed again. Furthermore, a part of its own radiosity is reflected back from the environment to itself. The standard algorithm shoots this new radiosity in a later iteration step.

The basic idea of the overshooting method is to estimate the radiosity that will be received later on and to shoot it to the environment together with the actual unshot radiosity ΔB_i [Feda 92]. In other words, the difference $\Delta B_i'$ between previously shot radiosity and the current estimate of the final radiosity of patch i is distributed to the environment instead of ΔB_i. $\Delta B_i'$ is called the estimated unshot radiosity. If the estimate is good, no further shooting is necessary from this patch.

Shooting more radiosity (estimated unshot radiosity $\Delta B_i'$) than available (actual unshot radiosity ΔB_i) results in a negative unshot radiosity ΔB_i^{new} of the shooting patch i after this iteration:

$$\Delta B_i^{new} = \Delta B_i - \Delta B_i' \tag{1}$$

$$\Rightarrow \qquad \Delta B_i^{new} \leq 0$$

In case $\Delta B_i'$ was overestimated, patch i will never receive enough radiosity later on to compensate for its negative unshot radiosity. In this case, the remaining negative unshot radiosity ΔB_i has to be corrected in a later iteration step by shooting negative radiosity.

Therefore, the overshooting method is guaranteed to converge to the same radiosity solution as the standard algorithm, but it converges faster, if the shootable radiosities $\Delta B_i'$ are well estimated.

There are several overshooting methods, which use different techniques to estimate $\Delta B_i'$ and to handle negative unshot radiosity. This paper summarizes these techniques and their distinguishing features. All these methods speed up the convergence of progressive refinement radiosity significantly at very low additional cost.

Note that estimating the unshot radiosity of a shooting patch does not directly affect its radiosity B_i. That means that patches hidden from all shooting patches show no improvement compared to pure progressive refinement radiosity. Ambient light [Cohe 88] can be used to brighten intermediate images even more than overshooting alone.

2 Overrelaxation

Overrelaxation [Gort 94] simply uses a fixed overrelaxation factor ω:

$$\Delta B_i' = \omega \cdot \Delta B_i \qquad \text{where} \quad 1 \leq \omega < 2 \tag{2}$$

This method is equivalent to the method well-known in numerical mathematics, which can be thought of an extrapolation of radiosity values before and after shooting ΔB_i. It typically works best for overrelaxation factors in the range 1.2 to 1.5, depending on the scene. The shooting patch selection is based on the absolute value of unshot radiosity, that means the patch that maximizes $|\Delta B_i \cdot A_i|$ is selected.

The advantage of overrelaxation is simplicity. It can be applied to any iterative solution method, even to the gathering approach. The disadvantage is that the speed-up depends on the experience of the user to select a good overrelaxation factor.

3 Ambient Overshooting

In [Cohe 88], the ambient light is described, which is used to improve the radiosity values B_i for display purposes. This ambient light is based on the current unshot radiosities of all patches and on the average reflectance of the environment. It proved to give a good estimate of the final radiosity result.

$$\rho_{av} = \frac{\sum_{i=1}^{N} \rho_i \cdot A_i}{\sum_{i=1}^{N} A_i}$$

$$Ambient = \frac{1}{1-\rho_{av}} \cdot \frac{\sum_{i=1}^{N} \Delta B_i \cdot A_i}{\sum_{i=1}^{N} A_i} \qquad (3)$$

A_i denotes the area of patch i, and ρ_i its reflection coefficient. See [Cohe 88] for details on the estimated ambient light.

This ambient light can also be used for overshooting [Feda 92]. The estimated unshot radiosity can then be computed by:

$$\Delta B_i' = \Delta B_i + \rho_i \cdot Ambient \qquad (4)$$

The shooting patch selection is based on the absolute value of estimated unshot radiosity, that means the patch that maximizes $|\Delta B_i' \cdot A_i|$ is selected.

The *Ambient* term must never become negative, because in this case the solution would not converge [Feda 92]. That means, the sum of unshot powers $\Delta B_i \cdot A_i$ of all patches i must always be positive. This can be accomplished if no shooting patch shoots more than the total unshot power. Therefore formula (4) has to be modified in the following way:

$$\Delta B_i' = \min\left(\Delta B_i + \rho_i \cdot Ambient \ , \ \sum_{j=1}^{N} \Delta B_j \cdot A_j / A_i\right) \qquad (5)$$

Ambient overshooting typically achieves significantly better speed-up than overrelaxation. However, it depends on the quality of estimation of ambient light, and therefore on the scene geometry and reflectance. Its advantage compared to other overshooting methods is its independence from any particular form factor method. It can be used for radiosity with ray-traced form factors [Wall 89], Monte Carlo radiosity [Shir 91], and other radiosity methods where no patch-to-patch form factors are explicitly computed.

4 Positive Overshooting (Gathering and Shooting)

Shooting from patch i requires to compute the form factors F_{ji} to all receiving patches j, that is the column i in the form factor matrix. Using the reciprocity relationship $F_{ij} \cdot A_i = F_{ji} \cdot A_j$ you get the corresponding F_{ij} almost for free.

With these form factors F_{ij} it is possible to predict the minimum amount of radiosity that will be shot towards patch i in subsequent iterations by gathering the actual unshot radiosities from all patches j. The gathered unshot radiosities are used for overshooting [Shao 93].

With this method, it is not necessary to shoot negative radiosity, since the negative radiosity resulting from overshooting is guaranteed to be compensated in later iterations. Therefore, all negative patches can be ignored for shooting patch selection before they get positive. This implies also that negative unshot radiosities are not considered for gathering, because they will become positive or at least zero anyway.

Thus, the formula to evaluate $\Delta B_i'$ is:

$$\Delta B_i' = \Delta B_i + \rho_i \cdot \sum_{j=1}^{N} \Delta B_j^+ \cdot F_{ij}$$

$$\text{where } \Delta B_j^+ = \begin{cases} \Delta B_j & \text{if } \Delta B_j > 0 \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

The main advantage of this method is that no negative shots and no heuristics are required, since $\Delta B_i'$ is based on known form factors. The main disadvantage is that in most cases $\Delta B_i'$ is substantially underestimated, because no interreflections are considered in the estimate. Furthermore, this method requires patch-to-patch form factors, and therefore it cannot be used e.g. for ray-traced form factors (vertex-to-patch) [Wall 89] and Monte Carlo radiosity [Shir 91].

5 Super-Shoot-Gather

This estimation method exploits the idea of gathering unshot radiosities. Additionally, it computes the amount of radiosity shot to the environment which will eventually be reflected back to the shooting patch, from there again reflected to the environment and back to the shooting patch, and so on ad infinitum [Gort 94].

The ratio of radiosity shot from patch i that comes back from the environment in the first step is given by

$$\sum_{j=1}^{N} \rho_i \cdot \rho_j \cdot F_{ij} \cdot F_{ji} \tag{7}$$

The radiosity that comes back from the environment to the shooting patch i after any number of bounces is given by a geometric series with the ratio given in (7). Combining the limes of this geometrical series with formula (6), $\Delta B_i'$ can be estimated by

$$\Delta B_i' = \frac{\Delta B_i + \rho_i \cdot \sum_{j=1}^{N} \Delta B_j^+ \cdot F_{ij}}{1 - \sum_{j=1}^{N} \rho_i \cdot \rho_j \cdot F_{ij} \cdot F_{ji}} \tag{8}$$

Negative patches can be ignored before they get positive, because the negative radiosity resulting from overshooting is guaranteed to be compensated in later iterations.

In its original form, super-shoot-gather requires a NxN matrix of unshot radiosity values, which is used to account for the pre-shot radiosity from a patch j to shooting patch i, when later on patch j gets its turn [Gort 94]. The simplified version presented here eliminates these high memory costs, because the gathered radiosity is not added to B_i, but only used for shooting. After this shooting step it is represented by negative unshot radiosity ΔB_i, in the same way as the additional radiosity estimated by the other overshooting methods.

The advantage of this method is that radiosities can never be overestimated. No negative shots and no heuristics are required, since $\Delta B_i'$ is based on known form factors. The super-shoot-gather estimate accounts for interreflections between the shooting patch and the environment, thus leading to faster convergence compared to positive overshooting. However, the influence of interactions between patches in the environment on the shooting patch are not considered. As positive overshooting, this method requires patch-to-patch form factors, and therefore it cannot be combined with ray-traced form factors or Monte Carlo radiosity, as ambient overshooting can.

6 Enhanced Super-Shoot-Gather with Ambient Light

This algorithm is a combination of ambient overshooting and super-shoot-gather. The ambient light is not used to determine the estimated unshot radiosity $\Delta B_i'$ of the shooting patch, but to enhance the unshot radiosities of all other patches in the scene. These estimated unshot radiosities are then used for super-shoot-gather [Xu 94].

The estimated unshot radiosity is given as follows:

$$\Delta B_i' = \frac{\Delta B_i + \rho_i \cdot \sum_{j=1}^{N} \left(\Delta B_j + \rho_j \cdot Ambient \right) \cdot F_{ij}}{1 - \sum_{j=1}^{N} \rho_i \cdot \rho_j \cdot F_{ij} \cdot F_{ji}} \tag{9}$$

The shooting patch selection cannot be done depending on $\Delta B_i'$, because the estimation of requires form factors. Therefore the shooting patch selection is done in the same way as for ambient overshooting, that means the patch that maximizes $|(\Delta B_i + \rho_i \cdot Ambient) \cdot A_i|$ is selected.

The *Ambient* term must never become negative, because in this case the solution would not converge [Feda 92]. This problem is the same as in ambient overshooting and can be solved in the same way: the energy that is shot from a shooting patch must be limited to the total unshot energy. This problem, however, was not considered in [Xu 94].

This method achieves the highest convergence rate of all methods. However, it requires patch-to-patch form factors, and therefore it cannot be combined with ray-traced form factors or Monte Carlo radiosity, as ambient overshooting can.

7 Summary

Several overshooting algorithms have been described that offer different advantages, but also have different disadvantages. For comparisons of the speed-up, a general ranking is impossible, because the speed-up of ambient overshooting and enhanced super-shoot-gather depends on the quality of the ambient estimate. Typically, enhanced super-shoot-gather achieves fastest convergence, the second-fastest algorithm is typically ambient overshooting, then normal super-shoot-gather, and the slowest one is positive overshooting. The speed-up of overrelaxation depends on the relaxation coefficient ω. If ray-traced form factors or other methods are used, where no patch-to-patch form factors are computed, ambient overshooting should be preferred.

References

[Cohe 88] M. Cohen, S. Chen, J. Wallace, D. Greenberg
A Progressive Refinement Approach for Fast Radiosity Image Generation
Computer Graphics, Vol. 22, Nº 4, 1988

[Feda 92] M. Feda, W. Purgathofer
Accelerating Radiosity by Overshooting
Proceedings of the 3rd Eurographics Workshop on Rendering, Bristol, UK, 1992

[Gort 94] S. Gortler, M. Cohen, P. Slussalek
Radiosity and Relaxation Techniques
IEEE Computer Graphics and Applications, Vol. 14, Nº 6, 1994

[Shao 93] M.-Z. Shao, N. Badler
Analysis and Acceleration of Progressive Refinement Radiosity Method
Proceedings of the 4th Eurographics Workshop on Rendering, Paris, France, 1993

[Shir 91] P. Shirley
Radiosity via Ray Tracing
Graphics Gems II, J. Arvo (Editor), Academic Press Inc., London-San Diego, 1991

[Wall 89] J. Wallace, K. Elmquist, E. Haines
A Ray Tracing Algorithm for Progressive Radiosity
Computer Graphics, Vol. 23, Nº 3, 1989

[Xu 94] W. Xu, D. Fussel
Constructiong Solvers for Radiosity Equation Systems
Proceedings of the 5th Eurographics Workshop on Rendering, Darmstadt, Germany, 1994