Model reduction of radiative heat transfer by a hierarchical radiosity method

Mickaël Le Bohec, Denis Lemonnier, Didier Saury

Model reduction of radiative heat transfer by a hierarchical radiosity method

M. Le Bohec[a,b,*], D. Lemonnier[a,*], D. Saury[a]

[a] *Institut Pprime UPR CNRS 3346 - CNRS / ISAE-ENSMA / Université de Poitiers, 1 avenue Clément Ader, B.P. 40109, F-86961 Futuroscope Chasseneuil CEDEX, FRANCE*

[b] *Aix Marseille Université, CNRS, IUSTI UMR 7343, 13453, Marseille, France*

Abstract

Radiative heat transfer in buildings can be determined through the resolution of the radiosity equations. This technique requires the evaluation of geometric couplings between the elements of the scene called view factors. Obtaining them is particularly challenging, especially when obstructions are involved. In addition, the system to be solved can be cumbersome, since each surface is likely to interact with all the others and the number of nodes involved in the description of a complex scene is large. We present a solution method that refines the mesh of a scene while simultaneously constructing a representation at different scales of the view factors between the contributing elements, in order to avoid calculating all the transfers at the highest resolution. It thus considerably reduces the calculation time and makes the tool suitable with a routine use in a design perspective.

Nomenclature

A	surface, [m^2]
E	irradiation, [$W \cdot m^{-2}$]
F	view factor, [−]
J	radiosity, [$W \cdot m^{-2}$]
k	number of initial elements, [−]
m	number of created elements, [−]
M	emissive power, [$W \cdot m^{-2}$]
M°	blackbody emissive power, [$W \cdot m^{-2}$]
n	number of elements in a scene, [−]
n_F	number of view factors, [−]
r	distance, [m]
v	visibility between two points, [−]
V	visibility between two elements, [−]
$\mathbf{x}$	a point in space, [−]

Greek symbols

ε	emissivity, [−]
ϵ	error, [−]

*Corresponding authors: mickael.lebohec@gmail.com, denis.lemonnier@ensma.fr

θ	angle, [rad]
ρ	reflectivity, [−]
τ_f	filling ratio, [−]
τ_o	obstruction ratio, [−]
τ_r	reduction ratio, [−]
φ	net heat flux, [$\mathrm{W \cdot m^{-2}}$]

Subscripts and superscripts

ε	related to a threshold
p, q	elements indices
pq	from element p to element q

Accents

$\hat{}$	non dimensional quantity
$\check{}$	corrected quantity

1. Introduction

A common way to model radiative heat transfer between solid elements is to resort to the radiosity method [1, 2]. In its usual form, it assumes opaque, grey and diffuse surfaces separated by a transparent medium. Once the domain is discretized into n elements over which the temperature, the reflectivity and the irradiation are assumed to be uniform, the set of discrete equations to be solved for the radiosity J_p on element p reads as:

$$J_p = M_p + \rho_p \sum_{q=1}^{n} F_{pq} J_q \tag{1}$$

The view factor F_{pq} expresses the geometric coupling between the elements p and q. It represents the portion of radiant energy leaving the surface p and coming directly onto q. Its evaluation may be a complex task and will not be discussed in this paper. We can refer to [3] for a comprehensive introduction to the subject. We simply remind that this factor is defined by the expression:

$$F_{pq} = \frac{1}{A_p} \int_{A_p} \int_{A_q} v_{pq} \frac{\cos\theta_p \cos\theta_q}{\pi r_{pq}^2} \mathrm{d}A_q \mathrm{d}A_p \tag{2}$$

where v_{pq} characterizes the visibility between two points located on the surfaces elements A_p and A_q. This parameter is 1 if these points can see each other and 0 otherwise.

In thermal modeling, net heat fluxes are often more useful than radiosities. They are deduced, for

a closed enclosure, in the following way:

$$\begin{aligned}\varphi_p &= J_p - E_p \\ &= J_p - \sum_{q=1}^{n} F_{pq} J_q \\ &= \sum_{q=1}^{n} F_{pq}(J_p - J_q) \\ &= \sum_{q=1}^{n} \varphi_{pq} \end{aligned} \tag{3}$$

where E_p denotes the irradiation of element p, that is the incident radiant heat flux on this element.

To solve the radiosity system, one must first evaluate the n^2 potential interactions between the elements and, for each pair, test whether the $n-2$ other elements interfere or not, totally or partially. Obtaining the matrix of the system has thus a cost of order $O(n^2)$ if the obstructions do not have to be tested and of order $O(n^3)$ if they have to.

We introduce in this paper a multiresolution method allowing reducing this complexity to an order $O(k^2+m)$ without obstructions and $O(k^3+km)$ with obstructions. Here, k is the number of elements of the initial mesh (expected to be rather coarse) and m the number of elements created by successive refinement based on the hierarchical method outlined below.

This technique comes from the photorealistic rendering community where it was extensively used and extended [4, 5, 6, 7]. In heat transfer, to the authors' knowledge, only the work by Bindick [8] can be mentioned, but this reference fails to provide some of the key parameters that strongly impact the performance of the method.

2. Hierarchical radiosity method

2.1. Presentation

The hierarchical radiosity method [4] is based on the fact that the radiative transfer between two elements can be evaluated at different levels of detail, depending on the importance of this particular exchange on the global solution. The element-to-element transfers with a modest impact can be assessed at a low level of detail, while those having a high impact need to be calculated on a finer grid.

This method therefore builds a multiresolution mesh of the scene allowing the different sub-meshes to interact at adapted scales. The different levels consist of successive subdivisions of the initial grid. Any refinement of a facet (or mesh element) creates a new level of detail on it. Thus, each initial element is the root of a tree in which each floor corresponds to a different resolution as shown in Figure 1. The complete mesh of the scene is therefore the union of the trees of all the initial facets.

Geometrical subdivisions are intended to refine the calculation of energy exchanges. Beforehand, the radiative transfer between each pair of elements is evaluated using an *estimator* whose value is compared to a given user-defined *threshold*. Facets targeted by transfers whose intensity exceeds the

threshold are subdivided. This results in using a denser mesh where the higher amounts of energy are involved.

Since the exchanges must be known prior to determining if any refinement is necessary, transfers at the coarser scale must be evaluated to initiate the method. Then, if the scene is initially composed of k elements, and if there are no obstructions, the cost of this first evaluation is of order $O(k^2)$. The rest of the work consists in calculating some of these transfers at higher levels of accuracy by using a finer mesh. The overall complexity is actually $O(k^2 + m)$, where m is the number of created elements to refine the spatial grid [4]. Thus, the coarser the initial mesh (the smaller k), the more efficient the hierarchical method is.

Refinement, when necessary, is performed for each given pair of facets. Therefore, one element can interact with the others at different levels of detail. In the example depicted by figure 1, element 2 interacts with element 3 at a higher level than with element 1. These interactions are represented by *links*. The power received by a facet can therefore come from multiple levels. The solution step will have to connect these different scales to access the final solution.

2.2. Estimator

Several estimator can be used to estimate the transfers between the elements to guide the refinement process. Their choice influences the dynamics of refinement and the sparsity of the resulting algebraic system. Whatever the estimator, the evaluation of view factors is required. In the context of the hierarchical method, the point-to-disk approximation [9, 10, 11, 12] is usually used to that end. In the present work, we avoid this approximation and use a classical surface integration method [1, 2].

2.2.1. View factor

A first possibility is to simply use the view factor F_{pq} as an estimator [4] since it is related to the fraction of power leaving an element p that reaches another element q. The intensity of the transfer is therefore proportional to this purely geometrical factor. Then, the refinement of a scene for a given threshold F_ε is independent of the thermal boundary conditions of the problem. This refinement can be performed once for all and used in several subsequent simulations based on the same geometry without any additional cost.

However, the refinement is not optimal for heat transfer problems since it does not take into account the actual power transferred between elements. Some facets can therefore be unnecessarily subdivided, if they involve a significant view factor, but are at a relatively low temperature compared to the rest of the scene, for instance.

2.2.2. Irradiation

Another option is to choose as an estimator the irradiation $E_{pq} = F_{pq}J_p$ of an element q by an element p [4, 11, 13, 14]. Then, all the important transfers will be refined even if the involved facets are distant. Conversely, weak transfers will no longer be refined even if the facets are close.

With this estimator, an initial evaluation of E_{pq} is needed and therefore J_p must be known on the starting mesh. The process is then *progressive*, alternating resolution and refinement until convergence. This progressivity can also be exploited by decreasing the value of the irradiation threshold at each step. This logic allows to quickly solve the system in the early steps and get a better starting solution for the subsequent steps. In practice, the threshold at each step s is halved, that is $E_{\varepsilon,s+1} = E_{\varepsilon,s}/2$. It is therefore the smallest threshold of this sequence, $E_{\varepsilon,min}$, that characterizes the final precision of the computation.

2.3. Refinement

The refinement procedure is now described focusing on the use of the irradiation based estimator.

There are two different ways of presenting the process. The most common relies on the geometric description of the tree structure [3, 8, 15]. The second is more theoretically based on the finite element framework and wavelets theory [5, 16, 17]. Here we follow the first approach and introduce new illustrations allowing to graphically and accurately describe the different steps of the algorithm.

2.3.1. General case

The refinement algorithm operates recursively on each pair of elements. For two facets p and q, it compares the estimators E_{pq} and E_{qp} to the threshold E_ε. If at least one of these two quantities is greater than E_ε, a local refinement is needed. If $E_{pq} > E_{qp}$, then q is subdivided, and p otherwise. Subdivisions are referred hereafter as the *children* of their *parent* element. On the other end, if the two estimators are below the threshold, there is no need to refine any facet : a link is created between p and q and the procedure stops for this pair. After each refinement step, the same process is reproduced between the unrefined element and the children of the refined element.

In order to limit the number of refinements (it could be infinite, especially for close neighbors), a minimum facet size A_ε is prescribed, below which the process is stopped.

Figure 2 illustrates this procedure over some iterations. The geometric configuration is shown at the top, followed by the profile of the estimator E_{pq} of the transfer from p to q and its corrected equivalent $\check{E}_{pq}$. The latter is expressed as:

$$\check{E}_{pq_i} = \frac{A_q}{A_{q_i}} E_{pq_i} \tag{4}$$

where q_i is one of the descendant facets of q. The area weighting ratio is simply used to compensate the variation of E_{pq} due to the size of the targeted surface and allows a graphical comparison of the estimator at different mesh levels (see figure 2, step 3). The continuous profile of the real transfer is drawn with a plain line, the discrete profile of the estimator is represented by a bold line and the energy passing between elements is represented by triangles formed by dotted lines. For sake of simplicity, the element p is here supposed to be indivisible ($A_p < A_\varepsilon$). Subdivisions only take place on element q.

At the onset of the first step, there is no subdivision and the transfer is always roughly represented by a single value. Since, in the present example, it is below the threshold $E_{\varepsilon,1}$, no refinement is needed and the process stops for this couple of elements. After going through all the other pairs of facets and having solved the whole system, if $E_{\varepsilon,min}$ is not reached, a new step begins with a lower threshold, $E_{\varepsilon,2}$. Now, still referring to figure 2, we observe that E_{pq} is greater than $E_{\varepsilon,2}$ and, consequently, the element q is subdivided (assuming that $E_{pq} > E_{qp}$). But, as the new fluxes remain above the threshold, one more refining step is needed. At the third iteration, all the irradiations are below the prescribed limit and the refinement stops. At the beginning of the last step, the edge values are less than $E_{\varepsilon,3}$ and do not require further modifications. But the values in the center remain too large and the procedure continues for them only until reaching $E_{\varepsilon,min}$. The profile of $\check{E}_{pq}$ shows that the number of subdivisions is greater on the peak of the continuous profile.

2.3.2. Scene with obstructions

When some elements obstruct to the exchange between two facets, the relevant visibility coefficient v_{pq} has to be assessed. This may be computationally intensive because every facet in the scene has to be tested to determine its potential shadowing effect. However, such an evaluation may take advantage of the hierarchical logic of refinement. Let us define V_{pq} the global visibility between two elements p and q as:

$$V_{pq} = \frac{1}{A_p A_q} \int_{\mathbf{x} \in A_p} \int_{\mathbf{x}' \in A_q} v_{pq}(\mathbf{x}, \mathbf{x}') \mathrm{d}A_q \mathrm{d}A_p \tag{5}$$

The progressive process inherent to the hierarchical method allows exploiting the information obtained at previous levels and avoids unnecessary computations. In particular, if two facets fully see each other ($V_{pq} = 1$), it is unnecessary to test the visibility between the sub-elements resulting from a possible refinement. Similarly, if two facets cannot see each other ($V_{pq} = 0$), then the refinement procedure stops. The only configurations that require repeated calculations of the visibility are found when the elements can partially see each other, i.e. when $V_{pq} \in]0, 1[$.

Figure 3 illustrates the progression of the refinement process for an exchange perturbed by an obstacle. We here set a null final threshold in order to focus on the impact of the obstruction. Initially, the visibility is partial and is therefore computed again for the two elementary transfers of the second iteration. Since it remains partial after this step, it is once more evaluated in the third iteration. At this point, the visibility no longer needs to be calculated because it is now total on the edges ($V_{pq} = 1$) and null in the center ($V_{pq} = 0$). If the maximum resolution is not reached yet, the transfers on the edges can be refined without further calculation of the visibility and those in the center no longer need to be refined (they are null).

In the example illustrated in figure 3, the threshold is set to zero and, despite this, the problem is actually reduced. Only six transfers are evaluated at the finest resolution instead of the eight possible. With the presence of obstruction, choosing a zero threshold allows a lossless reduction of the problem which avoids assessing exchanges that are actually null.

In addition, the multiresolution mesh of the hierarchical method limits the number of elements to be tested for computing the visibility, since any facet includes its own subdivisions. Then, to evaluate a possible occlusion between two elements p and q, it is enough to test the k facets of the initial mesh. Finally, in the presence of obstructions, the method has a complexity of order $O(k^3 + km)$.

2.4. Solving the hierarchical system

Once the scene is completely refined for a given threshold and the hierarchy fully constructed, the radiosity system needs to be solved. To that end, each element must *gather*, through the links that have been created, the energy received from other elements, before *passing* it to its hierarchy (parents and children). This is equivalent to a Gauss-Seidel iteration [3]. These steps are repeated until convergence of the solution.

For any element p, the first step computes the gathered radiosity J_p^{gath} as:

$$J_p^{gath} = \sum_{q \in \mathcal{L}_p} (1 - \varepsilon_p) F_{pq} J_q \tag{6}$$

where $\mathcal{L}_p$ is the set of elements linked to p by a transfer. This quantity actually represents the amount of reflection in the radiosity.

At the end of the gathering step, the radiosity of each facet still ignores the influence of the transfers taking place at the other levels of the hierarchy. The passing step then completes the process by making the elements of the same hierarchy communicate with each other. It also provides a solution of the system for *each* level of detail of the mesh.

To this purpose, the lower elements (leaves) of the hierarchy are first completed by passing them the energy density of their ancestry. Since each level entirely overlaps the parent (coarser) one, the upper levels are then assigned the area-weighted average of the energy density of their children.

In the example shown in figure 4, each descendant of facet 2 must gather the energy received from the environment: from facet 3 for the first subdivision level, and from facet 1 for the second level (2.1.1 to 2.2.2). By itself, element 2 does not interact with any facet of the scene. Its radiosity will be known after its children have transmitted theirs. Each of them must then pass its energy to its offspring (since what is received by 2.1 is also received by 2.1.1 and 2.1.2) and to its ancestry (since what is received by 2.1 is also received, partly, by 2). Focusing on the left part of the tree, we first get the radiosity of facet 2.1.1 by *pushing* down the gathered radiosity and adding the emissive power:

$$J_{2.1.1} = M_{2.1.1} + J_{2.1.1}^{gath} + J_{2.1}^{gath} + J_2^{gath} \tag{7}$$

The rest of the hierarchy will then be solved by *pulling* up the radiosities of the extremities:

$$J_{2.1} = \frac{1}{2} J_{2.1.1} + \frac{1}{2} J_{2.1.2} \tag{8}$$

$$J_2 = \frac{1}{2}J_{2.1} + \frac{1}{2}J_{2.2} \tag{9}$$

For this reason, the above process is often referred as the *push-pull* phase [4, 13, 14] and is repeated until convergence of the solution. In the present example, the resolution for facets 1 and 3 is straightforward: since they have no hierarchy, all the irradiations are received by a single element.

3. Model reduction

The performance of the hierarchical method (hm) can be assessed by comparison against the conventional method (cm). To make sense, this comparison must be conducted at the same *maximal* resolution (figure 5) here defined as the inverse of the minimal facet size $1/A_\varepsilon$.

3.1. Definitions

The filling ratio τ_f is defined as the number of view factors calculated with the hierarchical method $n_{F,hm}$, compared to that needed by the classical method, $n_{F,cm}$:

$$\tau_f = \frac{n_{F,hm}}{n_{F,cm}} \tag{10}$$

The reduction ratio is further defined as $\tau_r = 1-\tau_f$. It characterizes the model reduction capability of the method. The number $n_{F,hm}$ depends on the selected threshold value of the estimator while $n_{F,cm}$ is simply :

$$n_{F,cm} = \frac{n(n-1)}{2} \tag{11}$$

thanks to the reciprocity rule $A_iF_{ij} = A_jF_{ji}$ and since all mutual view factors are null (facets are convex elements).

The rate of obstruction (13) is the proportion of point pairs that do not see each other. This is an intrinsic geometric property of a given scene. It is defined as:

$$\tau_o = \frac{\displaystyle\int_{\mathbf{x}\in A}\int_{\mathbf{x}'\in A}\left[1 - v(\mathbf{x},\mathbf{x}')\right]\mathrm{d}\mathbf{x}'\mathrm{d}\mathbf{x}}{\displaystyle\int_{\mathbf{x}\in A}\int_{\mathbf{x}'\in A}\mathrm{d}\mathbf{x}'\mathrm{d}\mathbf{x}} = 1 - \frac{1}{A^2}\int_{\mathbf{x}\in A}\int_{\mathbf{x}'\in A} v(\mathbf{x},\mathbf{x}')\mathrm{d}\mathbf{x}'\mathrm{d}\mathbf{x} \tag{12}$$

For a discretized scene, τ_o can be approximated as the ratio of the number of null view factors $n_{F=0,cm}$ to the total number of view factors $n_{F,cm}$ in the full matrix of the conventional method :

$$\tau_o \approx \frac{n_{F=0,cm}}{n_{F,cm}} \tag{13}$$

With this definition, τ_o is clearly correlated to the potential of reduction of the model since the computation of null view factors tends to be eliminated by the hierarchical method.

However, $n_{F=0,cm}$ cannot be known before the calculation of every view factor, which is precisely what the method tries to avoid. The rate of obstruction is rather evaluated stochastically picking a

large number N of point pairs $(\mathbf{x}_i, \mathbf{x}_j)$ following a uniform probability distribution. For each of them, the visibility $v(\mathbf{x}_i, \mathbf{x}_j) = v_{ij}$ is tested. The reduction rate is then estimated as:

$$\tau_o \approx 1 - \frac{1}{N} \sum_{k=0}^{N} v_{ij,k} \tag{14}$$

A null threshold leads to the calculation of all transfers at the finest resolution (except when a complete obstruction occurs). The solution obtained this way is identical to that of the conventional method since null view factors have no impact on the results. This provides a reference against which the relative L_1-approximation error on radiosities and radiant fluxes is calculated as :

$$\epsilon(\mathbf{J}) = \frac{\|\mathbf{J}_{hm} - \mathbf{J}_{cm}\|_1}{\|\mathbf{J}_{cm}\|_1} \qquad \text{and} \qquad \epsilon(\boldsymbol{\varphi}) = \frac{\|\boldsymbol{\varphi}_{hm} - \boldsymbol{\varphi}_{cm}\|_1}{\|\boldsymbol{\varphi}_{cm}\|_1} \tag{15}$$

3.2. Analysis on ideal configurations

The performances of the hierarchical method are first assessed on two idealized configurations (figure 6). The first one is a unit cube with all faces at 20 °C, except the left one set at 100 °C. Three different uniform emissivities are considered, from 0.8 to 1.0. The radiosity and net heat flux distributions are presented over the back (a) or right (b) face (target). The second configuration simply adds an obstacle kept at 20 °C in the center of the cavity.

3.2.1. Without obstacle

Figure 7 shows the evolution of the error on the prediction of radiosities and net heat fluxes when the threshold decreases down to zero during the computation. In practice, the idea is to stop the refinement before the entire view factor matrix is filled.

It is noticeable that radiosity and net heat flux converge together at the same rate. However, the radiosity predictions are more accurate of approximately one order of magnitude for a given threshold. In other words, a model aiming at approximating radiosities requires fewer view factors, and is therefore lighter, than when it comes to approximating net heat fluxes.

The evaluation of radiosities tends to be more accurate when the emissivity increases. Actually, in the limiting case of a black cube, this quantity is uniformly equal to the emissive power on each (isothermal) face. The error is then null because, in this specific configuration, radiosities are available without any view factor computation. This is no longer true for gray surfaces, and the initial error on J is then maximal. In short, the higher the emissivity, the fewer radiosity variations have to be captured by refinement.

On the contrary, the net heat flux varies on each face even for a black cube as stated by equation 16 since in this special situation $J_p = M_p^\circ$.

$$\varphi_p = J_p - \sum_{q=1}^{n} F_{pq} J_q = M_p^\circ - \sum_{q=1}^{n} F_{pq} M_q^\circ \tag{16}$$

Figure 7 shows that the emissivity has an opposite effect on the net heat flux since it increases the error. This is however a consequence of the normalization based on the maximal emissive power of the scene M_{max} as shown by equation 17. The normalized threshold $\hat{E}_\varepsilon$ increases when the emissivity decreases and fewer view factors are computed for a given value of this parameter.

$$\hat{E}_\varepsilon = \frac{E_\varepsilon}{M_{max}} = \frac{E_\varepsilon}{(\varepsilon M^\circ)_{max}} = \frac{E_\varepsilon}{\varepsilon M^\circ_{max}} \tag{17}$$

Figure 8 exhibits the typical profiles of the evolution of the reduction rate as a function of the threshold. It involves three stages. At the early stages of the calculation, only a few new links are created at each iteration. Then, in a second step, the number of subdivisions sharply increases and necessitates much more view factors to be calculated. Finally, an improvement of the solution may still be reached without creating many new links but involving further iterations. The last two steps are necessary if one seeks a very high level of accuracy and should be avoided in the view of model reduction.

The results demonstrate that a good approximation can be obtained with only a few percents of the view factors being calculated (the first step of the three-part process described above). For instance, errors on net heat flux remain between 2.0 % and 2.3 % with only 5.1 % and 7.0 % of the view factors to be evaluated (for a threshold of $\hat{E}_{\varepsilon,min} \approx 0.001$).

The reduction rate starts with a value of 100 % because, on the coarsest mesh, the number of view factors needed to obtain a solution is negligible compared to that finally required for the desired resolution. The process does not converge to a null reduction rate, which means that even if the solution produced by the hierarchical method is identical to the one obtained with the full matrix, as illustrated on figure 7, the reduction technique prevents the view factors linking two elements from the same face of the cube from being calculated.

The asymptotic value of the reduction ratio is therefore geometry-dependent and can be determined analytically in this simple configuration. Indeed, a cube composed of n elements has $n/6$ elements per face. A cube meshed with n uniformly distributed elements has $n/6$ elements per face. Since the elements of a same face do not see each other, there are $6(n/6)(n/6-1)/2$ null view factors (from equation (11)) out of $n(n-1)/2$. The proportion of null values is therefore $(n/6-1)/(n-1)$ which tends to $1/6$ for large values of n. This corresponds to the obstruction ratio observed earlier on figure (8).

Figure 9 displays the error on the net flux φ for different maximal resolutions $1/A_\varepsilon$. The error is computed with respect to a solution at the largest resolution of 4096 elements/m^2. As expected, the accuracy increases with the maximal resolution, but it is worth noting that all the curves display the same convergence rate until they reach their asymptote. These asymptotic values differ because once a given maximum resolution is prescribed, no further refinement can occur.

Figure 10 displays the net flux field on a face adjacent to the heated surface, that is, on the target face of figure 6a. Three distributions are plotted for three different maximal resolutions but with a

same threshold of $\hat{E}_{\varepsilon,min} = 1/2^9 \approx 0.002$. This value corresponds to the last threshold for which errors are similar for the three simulations (see figure 9).

The numerical values associated with figure 10 are presented in table 1. The solution displayed in figure 10a is, as expected, the closest to convergence, the error with respect to a full solution of the same resolution (64 elements/m^2 in its case) is 0.8 % and the reduction ratio of 31.6 % is the closest to $1/6 \approx 16.7$ %, the asymptotic intrinsic value of the geometry. Conversely, the solution of figure 10c is the most distant from a full solution of the same resolution (1024 elements/m^2 in its case). Despite this, their errors, compared to a same solution with a maximum resolution of 4096, are equivalent.

Moving to the lowest resolution (64 elements/m^2) to the highest (1024 elements/m^2) necessitates calculating 10 times more view factors with the hierarchical method. With the conventional approach, it would have cost 257 times more view factors evaluations (from equation (11)).

Finally, figure 11 shows 2D distributions of the net heat flux received on a face adjacent to the heated one at different stages of an entire progression of the method. The emissivity of the cube inner surfaces is 0.8 and the maximum resolution is 1024 elements/m^2. As the threshold $\hat{E}_{\varepsilon,min}$ decreases, the mesh becomes denser and the profile is more accuralety drawn. At the lowest values of the threshold, the maximum resolution is reached everywhere. Even when this is established, additional transfers are evaluated and the error falls from 1.2 to 0 % between $\hat{E}_{\varepsilon,min} = 1/2^{11}$ and $1/2^{39}$ (see figure 7).

The first step (figure 11a) starts with only one element (per face) which is enough to describe the geometry. The net heat flux is therefore represented by a single value. Then (figure 11b) the refinement process adds new elements close to the hot face, where the irradiation (the estimator) assumes the largest values. The maximal resolution is even rapidly reached in this area. Then (figure 11c) the refinement proceeds especially at the periphery because of the proximity to the neighboring elements of the other surfaces which involves large view factors.

Within a few steps (basically from $\hat{E}_{\varepsilon,min} = 1/2^9$), the general shape of the net flux field is accurately reproduced and there is no significant difference between the two solutions at $\hat{E}_{\varepsilon,min} = 1/2^{11}$ and $\hat{E}_{\varepsilon,min} = 1/2^{39}$.

3.2.2. With obstacle

A second ideal case is considered to investigate the impact of an independent occluder added to the configuration (see figure 6b).

Figures 12 and 13 show respectively the errors and the reduction and obstruction ratio for this case. They are the equivalents of figures 7 and 8 for the unoccluded case. They exhibit very similar profiles showing that the occlusion does not affect the typical progression of the method. In particular, the convergence rate is unaffected. All remarks concerning figures 7 and 8 remain valid here.

Due to the presence of the obstacle, the asymptotic reduction ratio increases to 27.8 % instead of 16.7 % in the unoccluded case. However, the accuracy of the reduced model is preserved. Still with the threshold $\hat{E}_{\varepsilon,min} \approx 0.001$, the error on the net heat flux remains between 1.9 % and 2.2 % with only 4.2 % to 7.2 % of the view factors to be evaluated. These values are similar to what was observed

in the unoccluded case.

Figure 14 represents the asymptotic reduction rates when varying the obstruction ratio. The scenes remain similar to the configuration drawn in figure 6b but the size of the internal cube changes. The larger the obstacle is, the higher the obstruction ratio. The asymptotic reduction ratio is the minimum reduction that can be expected from the method. In practice, the reduction will be much more substantial, as it can be seen on figure 13, since the refinement will stop before reaching this asymptotic value using an appropriate threshold.

The curve with round-shaped symbols in figure 14 represents the reduction ratio obtained with the hierarchical method and the curve with no symbols refers to the ideal situation of a reduction ratio equal to the obstruction ratio. This plot proves that the obstruction ratio τ_o is a good estimator of potential reduction due to obstruction. Large obstruction ratios allow the strongest reduction possibilities. For a mesh of infinitely small facets, the asymptotic reduction ratio would be equal to the obstruction ratio. Here a slightly different is observed due to the finite discretization, since a portion of the view factor is still involved in visibility computations.

Figure 15 compares the computation times of the hierarchical method (hm) and the conventional full matrix method (cm). For the former, the data are obtained with a final threshold of $\hat{E}_{\varepsilon,min} \approx 0$. It can be viewed as the worst-case computational time, since allowing a larger tolerance would shorten the computations.

A significant gain is observed when using the hierarchical method : it reaches several orders of magnitude as the desired resolution increases. This is mainly due to the fact the hierarchical method retains a constant number of elements to be tested for visibility (those of the initial mesh) whatever the resolution. Stopping certain branches of the hierarchy, when full obstruction is detected, or switching to a view factor calculation that does not imply visibility testing when no obstructions are found, also helps to reduce the cost of the simulation. Without this information, the conventional method must calculate all the view factors and needs to test the visibility for each one.

Figure 16 illustrates the net heat flux distribution over a face adjacent to the heated surface one for different steps of an entire progression of the method.

The emissivity of the walls is 0.8 and the maximum resolution is 1024 elements/m^2 as in figure 11. It is worth pointing out that, in the early stage, the thermal signature of the occluder is not visible and it appears as the method proceeds.

3.3. Simulation of a realistic case

The proposed scene is a living room occupied by two people and equipped with a coffee table and a chair, as depicted in figure 19. The source is a fireplace whose core is at 275 °C. The surface temperature of the occupants is 30 °C while the rest of the room (including the furniture) is at 19 °C. All surfaces have an emissivity of 0.95. The initial mesh represents a total area of 39.27 m^2 distributed into 188 elements, including the ceiling and the bounding walls which are not represented. After refinement to a maximum resolution of about 10 cm^2, the finest level of detail is composed of 27 116

elements. Without a hierarchical method it would be necessary to evaluate 368 million view factors (see equation (11)). Flux results are not plotted on the chimney because it would degenerate the color scale by adding too high a value.

Figure 17 and 18 show the evolution of the error on the net heat flux (evaluated with equation (15)) and of the reduction ratio as functions of the minimum threshold. The profiles are similar to those obtained in the ideal cubic configurations, which means that the dynamic of the method does not vary with the complexity of the scene. Here, however, errors are larger in the early stages, probably because of the higher number and increased complexity of the occluders.

In this configuration, the asymptotic reduction ratio is quite high (44.6 %) though the full potential of the method is not exploited. Therefore, the reference solution can be obtained while computing only 44.6 % of view factors. The entire processing of the scene took 2 days and 4 hours (on an Intel® Core™ i5 processor of 2.7 GHz), which is long but affordable. A conventional method could not access this solution in a reasonable time.

In this scene, the transition in the reduction ratio is found for a threshold $\hat{E}_{\varepsilon,min}$ around 10^{-3} and 10^{-4} instead of 10^{-4} and 10^{-6} in previous cases. This is due to the larger dimensions of this scene, involving lower values of view factors for a given maximal resolution, since the distance between elements is larger.

Figure 19 shows the net flux distributions received by all the elements in the living room for two different minimum thresholds: the reference solution is obtained by setting $\hat{E}_{\varepsilon,min} \approx 0$ and the further reduced solution with $\hat{E}_{\varepsilon,min} = 1/2^{14} \approx 6.1 \times 10^{-5}$. The difference between these two solutions does not exceed 2.7 % and is almost imperceptible in the graphs. In the second case ($\hat{E}_{\varepsilon,min} = 1/2^{14}$) the reduction rate is 96.2 % and the computation time is drastically reduced to 1 h 30 min (instead of 52 h with $\hat{E}_{\varepsilon,min} \approx 0$). Allowing a 3 % error on the results divides the computation time by 35 and the number of view factors to be evaluated by 15.

4. Conclusions

We provide an original implementation of the hierarchical method in the context of thermal simulations in buildings. This technique allows reducing the algorithmic complexity of computation from $O(n^2)$ to $O(k^2 + m)$ without obstruction and from $O(n^3)$ to $O(k^3 + km)$ with obstructions (with $k \ll n$). It makes it possible to tackle more complex radiative scenes as the one presented in figure 19.

In graphic rendering, from where the method is derived, the quantity of interest is the radiosity, which determines the color of a lightened object. In heat transfer problems, radiosity but also net heat flux are sought depending on the situation. It is observed that the latter converges more slowly than the former and, therefore, needs more view factors to be calculated.

Our analysis highlights the obstruction ratio as a critical parameter. It is defined as the ratio between the number of null view factors to the total number of view factors of a scene when computed

with the conventional full matrix method. This parameter is linked to the potential of reduction of a scene model because the hierarchical method uses regularities on large scales to guide the computation on smaller ones and uses in particular the obstruction to avoid further expensive visibility test. We have suggested a stochastic method to allowing to cheaply estimate the asymptotic reduction potential through the obstruction ratio (14).

The hierarchical method proceeds by refining recursively a scene until every heat transfers are below some minimum threshold. This progression was obtained by successively halving the threshold until reaching the desired minimum. When the minimum threshold tends to zero, the method can be seen as lossless: it converges to a solution exactly identical to the one which would have been obtained with a conventional method. But it still provides some optimisations by avoiding to calculate a large portion of null view factors.

Allowing a larger uncertainty in the solution helps to increase the reduction rate of the model. In this context, a reliable solution was obtained on a complex scene in a significantly reduced time (divided by 35) with only few percents of error.

Future works will explore the way to use the hierarchical method on challenging general geometries with curved surfaces. Such scenes need a lot of small elements to fit the geometry, and few refinements are possible. Consequently, the hierarchical method works almost like a classical full matrix method. Many authors proposed to use different kinds of clustering to tackle this problem [13, 14, 18, 19, 20]. Extending the hierarchical methodology to semitransparent medium through the zonal method is another perspective.

References

[1] M. F. Modest, Radiative heat transfer, Academic press, 2013.

[2] J. R. Howell, M. P. Menguc, R. Siegel, Thermal radiation heat transfer, CRC press, 2010.

[3] M. F. Cohen, J. Wallace, P. Hanrahan, Radiosity and realistic image synthesis, Academic Press Professional, 1993.

[4] P. Hanrahan, D. Salzman, L. Aupperle, A rapid hierarchical radiosity algorithm, ACM SIGGRAPH Computer Graphics 25 (4) (1991) 197–206. doi:http://doi.org/10.1145/127719.122740.

[5] S. J. Gortler, P. Schröder, M. F. Cohen, P. Hanrahan, Wavelet radiosity, in: Proceedings of the 20th annual conference on Computer graphics and interactive techniques, 1993, pp. 221–230. doi:http://doi.org/10.1145/166117.166146.

[6] L. Aupperle, P. Hanrahan, A hierarchical illumination algorithm for surfaces with glossy reflection, in: Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques, 1993, pp. 155–162. doi:10.1145/166117.166137.

[7] F. Sillion, A unified hierarchical algorithm for global illumination with scattering volumes and object clusters, IEEE Transactions on visualization and computer graphics 1 (3) (1995) 240–254. doi:10.1109/2945.466719.

[8] S. Bindick, M. Stiebler, M. Krafczyk, Fast kd-tree-based hierarchical radiosity for radiative heat transport problems, International Journal for Numerical Methods in Engineering 86 (9) (2011) 1082–1100. doi:https://doi.org/10.1002/nme.3091.

[9] J. R. Wallace, K. A. Elmquist, E. A. Haines, A ray tracing algorithm for progressive radiosity, ACM SIGGRAPH Computer Graphics 23 (3) (1989) 315–324. doi:10.1145/74334.74366.

[10] G. Coombe, M. J. Harris, A. Lastra, Radiosity on graphics hardware, in: Proceedings of Graphics Interface 2004, 2004, pp. 161–168. doi:10.1145/1198555.1198782.

[11] B. E. Smits, J. R. Arvo, D. H. Salesin, An importance-driven radiosity algorithm, ACM SIGGRAPH Computer Graphics 26 (2) (1992) 173–282. doi:10.1145/142920.134080.

[12] P. Schröder, Numerical integration for radiosity in the presence of singularities, Fourth Eurographics Workshop on Rendering (1993).

[13] A. J. Willmott, P. S. Heckbert, M. Garland, Face cluster radiosity, in: Rendering Techniques '99, 1999, pp. 293–304.

[14] E. Gobbetti, L. Spanò, M. Agus, Hierarchical higher order face cluster radiosity for global illumination walkthroughs of complex non-diffuse environments, Computer Graphics Forum 22 (3) (2003) 563–572. doi:https://doi.org/10.1111/1467-8659.t01-1-00704.

[15] M. Stamminger, H. Schirmacher, P. Slusallek, H.-P. Seidel, Getting rid of links in hierarchical radiosity, Computer Graphics Forum 17 (3) (1998) 165–174. doi:https://doi.org/10.1111/1467-8659.00264.

[16] P. Schröder, S. J. Gortler, M. F. Cohen, P. Hanrahan, Wavelet projections for radiosity, Computer Graphics Forum 13 (1994) 141–151. doi:https://doi.org/10.1111/1467-8659.1320141.

[17] P. Schröder, Wavelet radiosity : Wavelet methods for integral equations, ACM SIGGRAPH '96 (1996) 143–165.

[18] L. Alonso, F. Cuny, S. P. Jean, J.-C. Paul, S. Lazard, E. Wies, The virtual mesh : A geometric abstraction for efficiently computing radiosity, ACM Transactions on Graphics 20 (3) (2001) 169–201. doi:http://doi.org/10.1145/501786.501789.

[19] B. Smits, J. Arvo, D. Greenberg, A clustering algorithm for radiosity in complex environments, Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques (1994) 435–442doi:http://doi.org/10.1145/192161.192277.

[20] P. H. Christensen, D. Lischinski, E. J. Stollnitz, D. H. Salesin, Clustering for glossy global illumination, ACM Transactions on Graphics 16 (1) (1997) 3–33. doi:http://doi.org/10.1145/237748.237749.

Figures

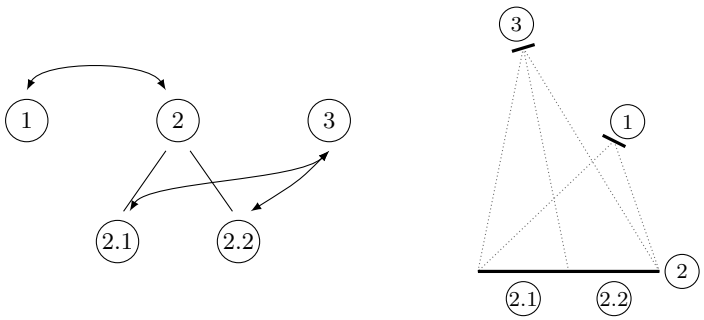

Figure 1: Example configuration where transfers (represented on the right by triangles) exist at different levels of detail.

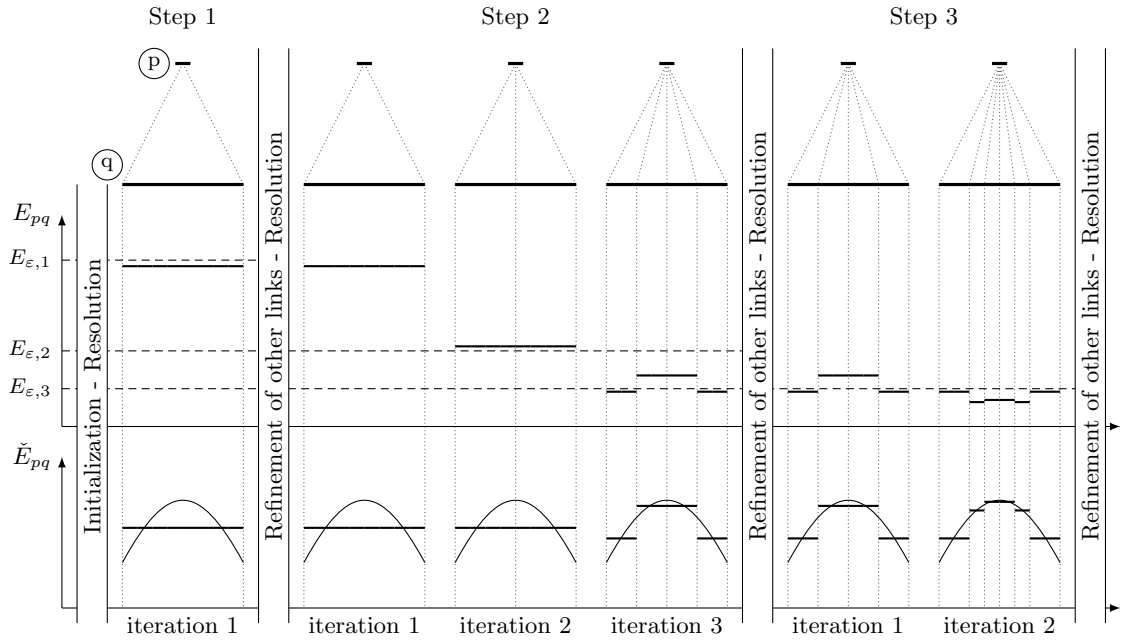

Figure 2: Implementation of the refinement procedure. Each step is characterized by a threshold and consists in a refinement phase and a solution phase.

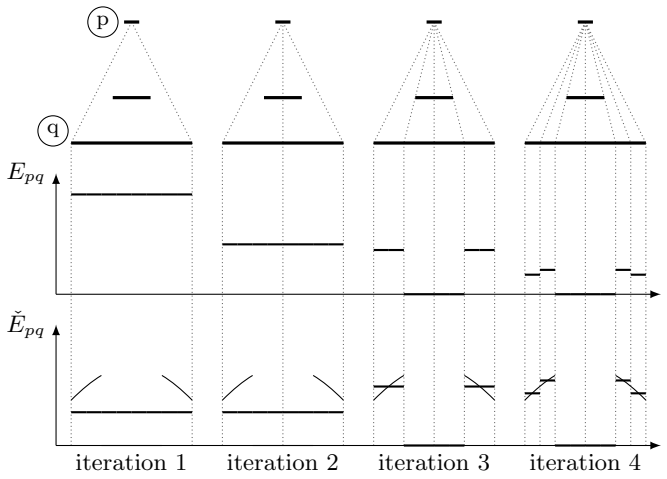

Figure 3: Implementation of the refinement procedure in a case involving an obstacle.

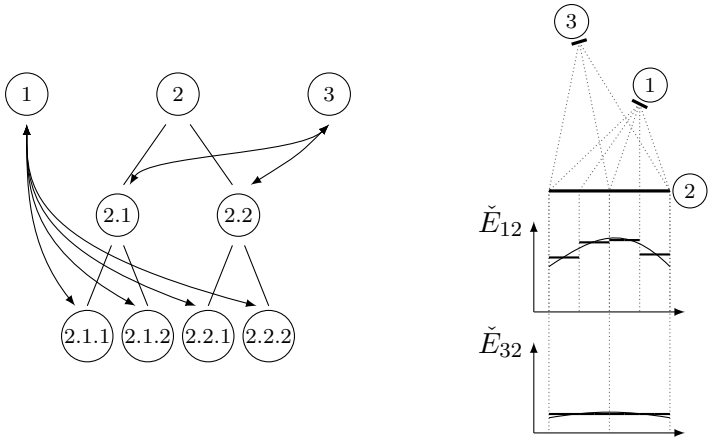

Figure 4: Situation illustrating the two steps involved in the solution phase.

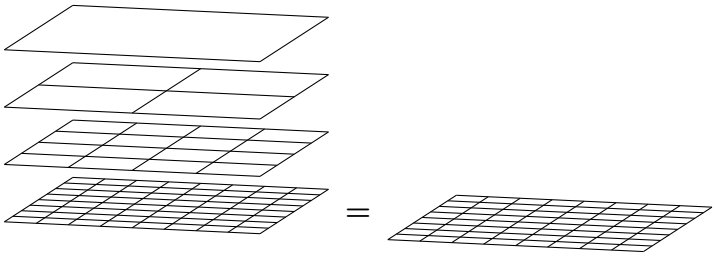

Figure 5: Example of meshes allowing a comparison (at maximal resolution) between the hierarchical method (left) and the conventional full matrix method (right) in terms of the number of view factors involved.

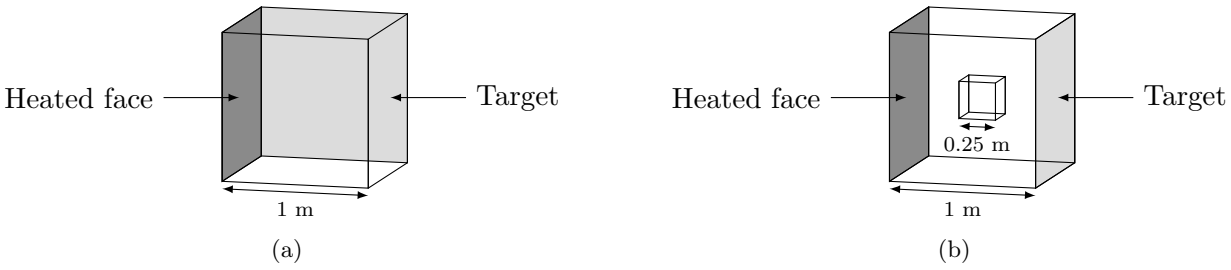

Figure 6: Ideal configurations. The left face is at 100 °C, the other ones at 20 °C.

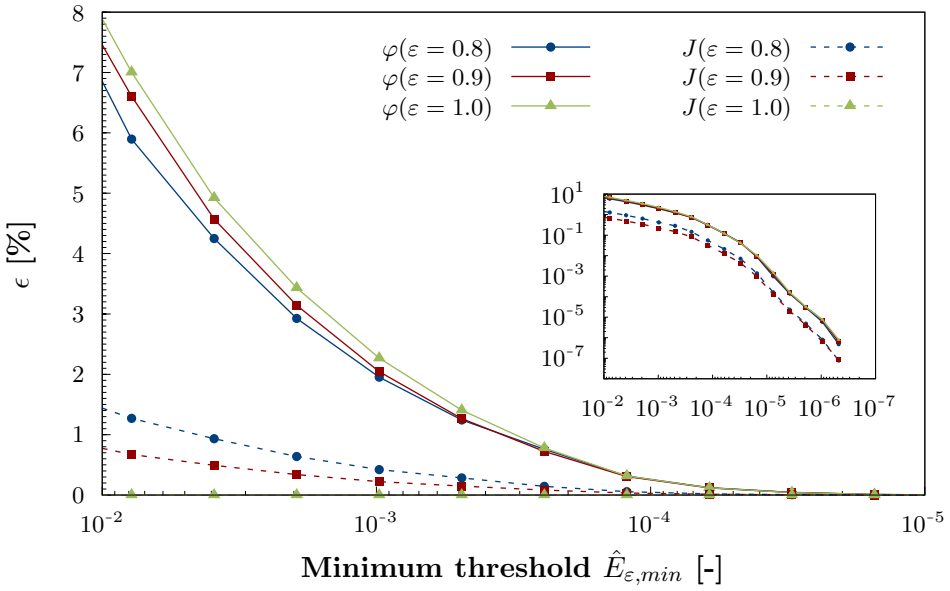

Figure 7: Evolution of the error on the net fluxes φ (full line) and the radiosities J (dotted line) using the irradiation estimator for different emissivities in the unoccluded case. Transfers are nondimensionalized by the maximal emissive power M_{max}.

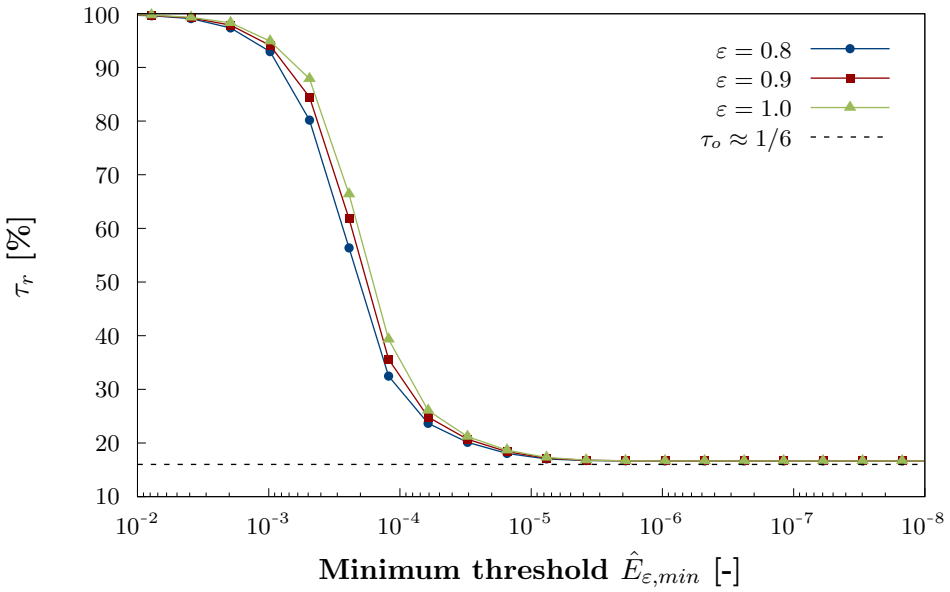

Figure 8: Evolution of the reduction ratio τ_r during refinement using the irradiation estimator for different emissivities in the unoccluded case.

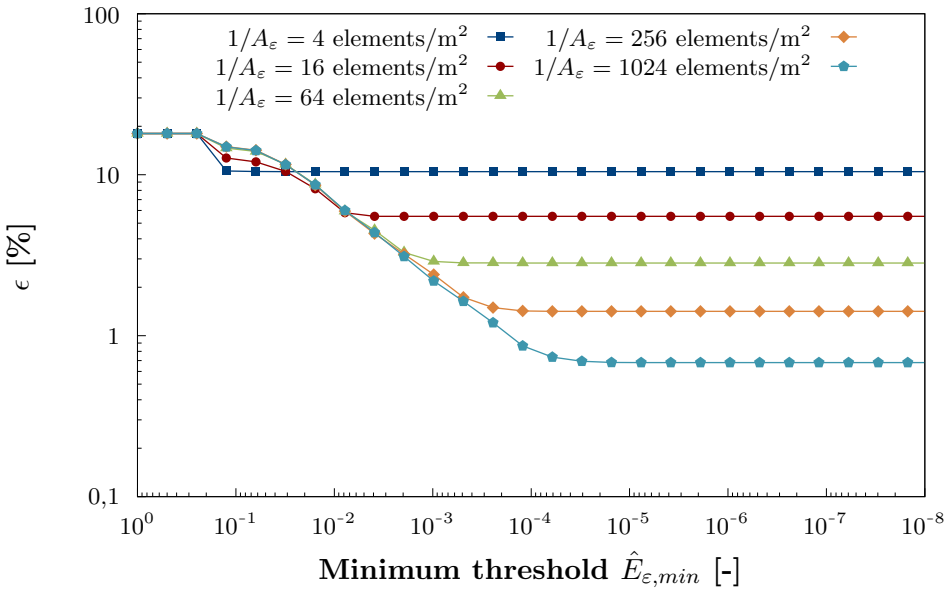

Figure 9: Evolution of the error on the net flux φ for different maximal resolution $1/A_\varepsilon$. The error is computed with respect to a solution at a resolution of $1/A_\varepsilon = 4096$ elements/m^2.

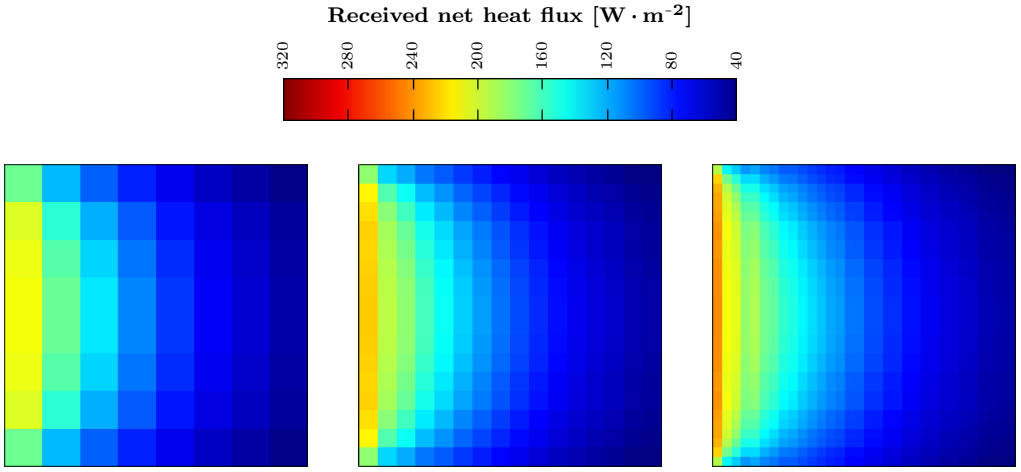

(a) $1/A_\varepsilon = 64$ elements/m^2 (b) $1/A_\varepsilon = 256$ elements/m^2 (c) $1/A_\varepsilon = 1024$ elements/m^2

Figure 10: Distribution of the net flux φ on the target surface for a same threshold $\hat{E}_{\varepsilon,min} = 1/2^9$ and for different maximal resolutions $1/A_\varepsilon$ with an emissivity of 0.8.

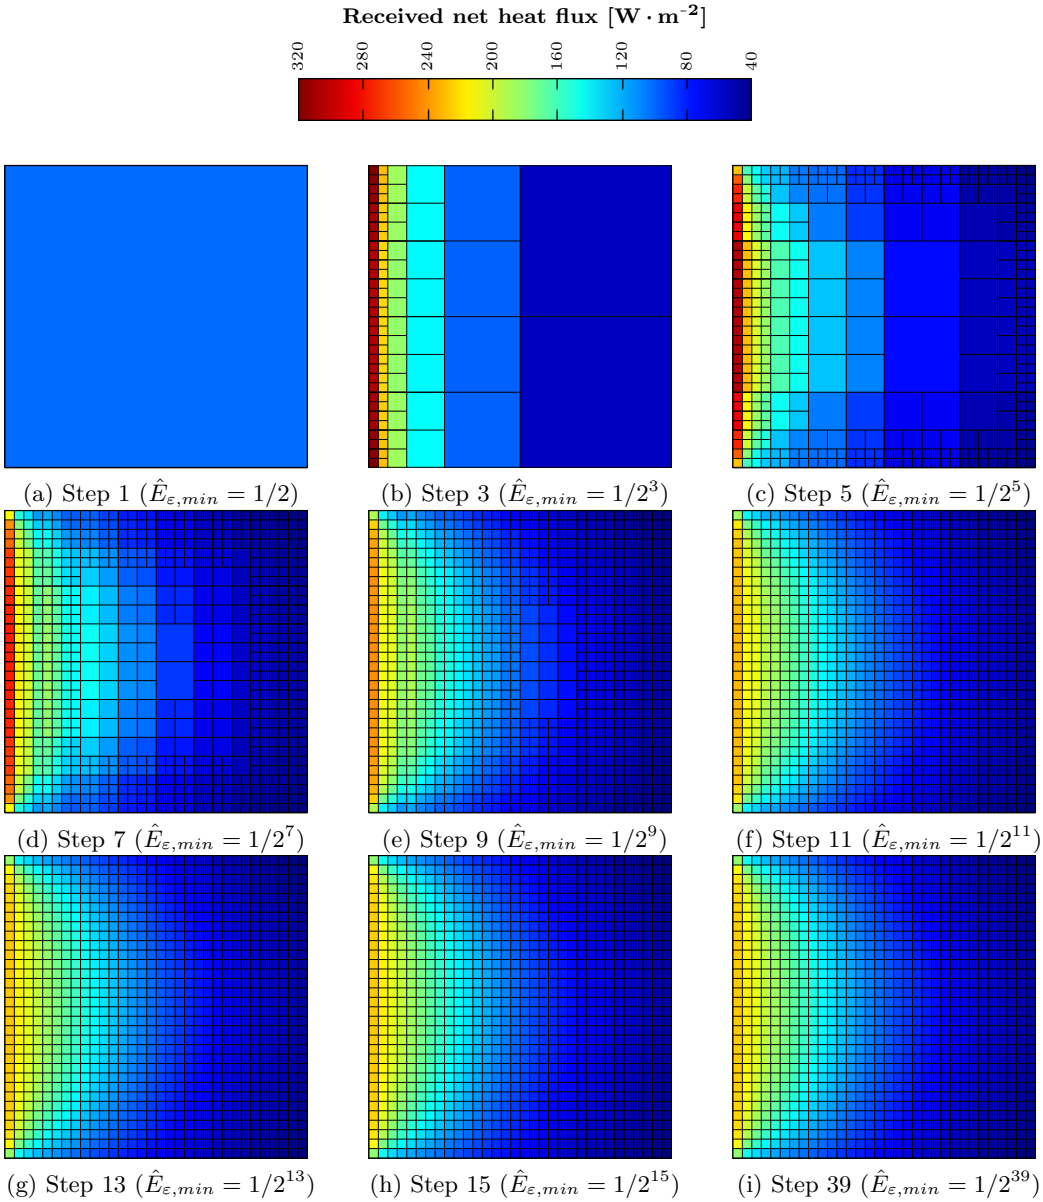

(a) Step 1 ($\hat{E}_{\varepsilon,min} = 1/2$) (b) Step 3 ($\hat{E}_{\varepsilon,min} = 1/2^3$) (c) Step 5 ($\hat{E}_{\varepsilon,min} = 1/2^5$)

(d) Step 7 ($\hat{E}_{\varepsilon,min} = 1/2^7$) (e) Step 9 ($\hat{E}_{\varepsilon,min} = 1/2^9$) (f) Step 11 ($\hat{E}_{\varepsilon,min} = 1/2^{11}$)

(g) Step 13 ($\hat{E}_{\varepsilon,min} = 1/2^{13}$) (h) Step 15 ($\hat{E}_{\varepsilon,min} = 1/2^{15}$) (i) Step 39 ($\hat{E}_{\varepsilon,min} = 1/2^{39}$)

Figure 11: Distribution of net flux on a face adjacent to the hot surface at different steps of the reduction process ($\varepsilon = 0.8$; $1/A_\varepsilon = 1024\,\text{elements/m}^2$).

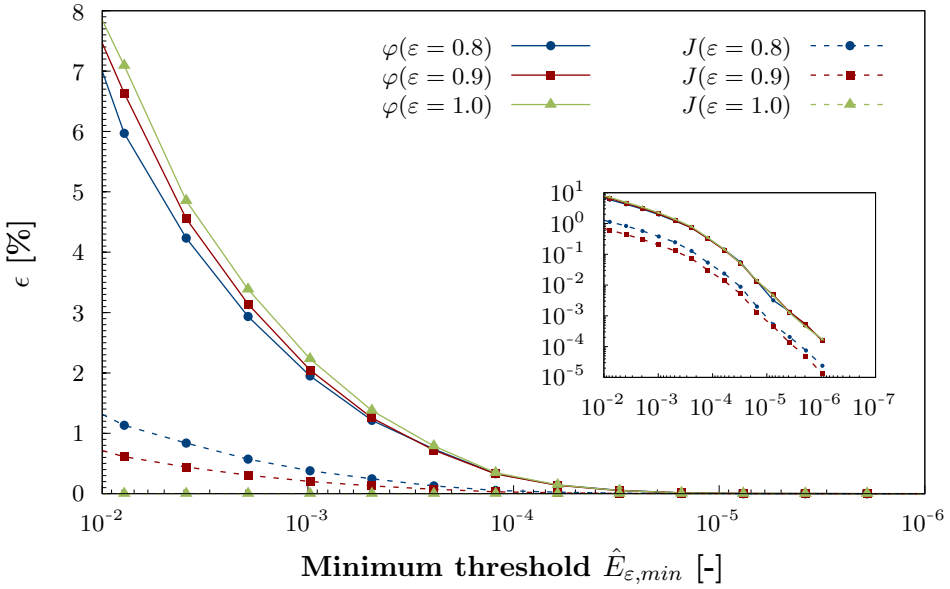

Figure 12: Evolution of the error on the net fluxes φ (full line) and the radiosities J (dotted line) using the irradiation estimator for different emissivities in the occluded configuration with an obstruction ratio of $\tau_o \approx 27.8\%$.

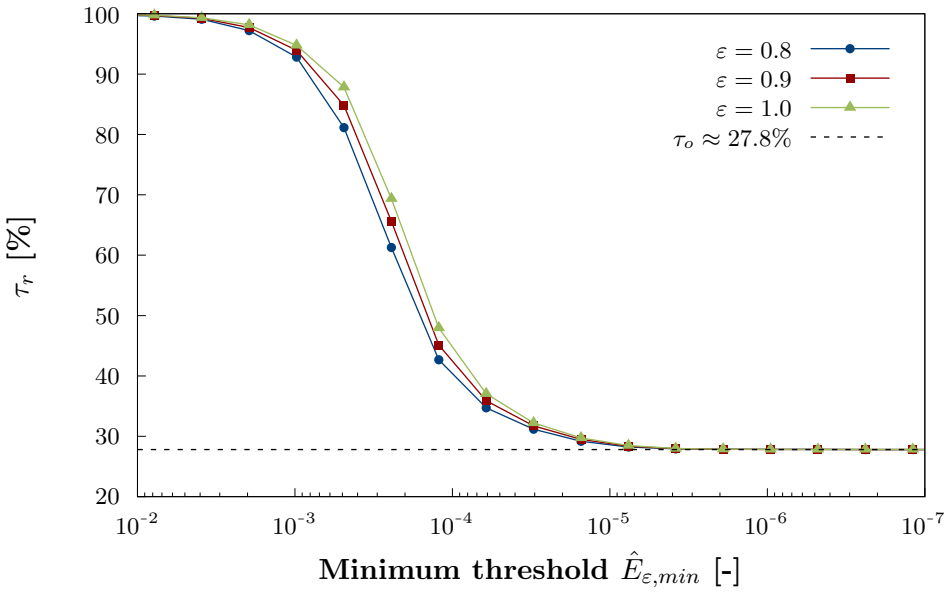

Figure 13: Evolution of the reduction ratio τ_r using the irradiation estimator in the occluded configuration with an obstruction ratio of $\tau_o \approx 27.8\%$.

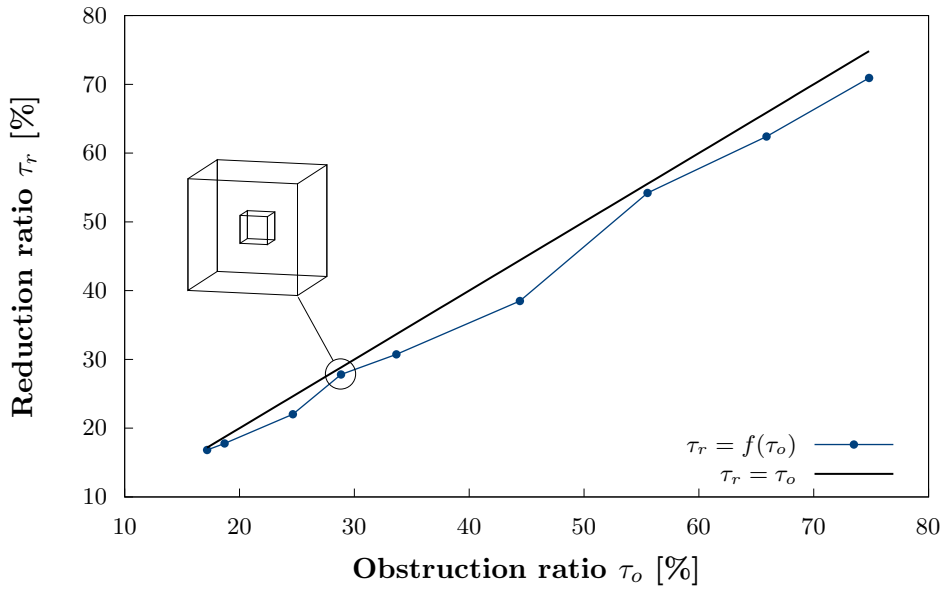

Figure 14: Evolution of the asymptotic reduction ratio τ_r for different occluded scenes of various obstruction ratio. Those were obtained in different concentric cubes scenes with different sizes of the internal cube.

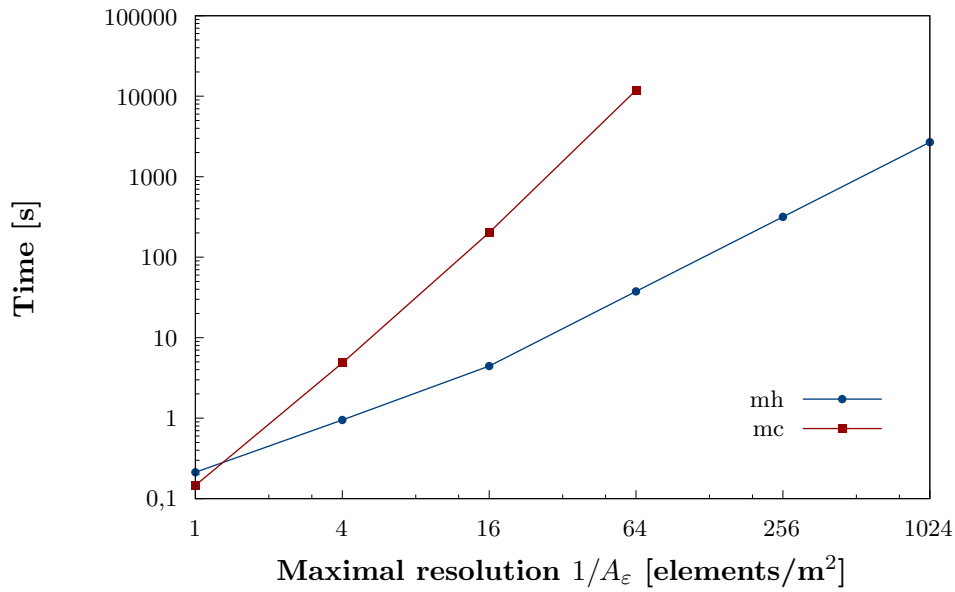

Figure 15: Computation time for various resolutions with the hierarchical method (round-shaped symbols) and the conventional full matrix method (square-shaped symbols).

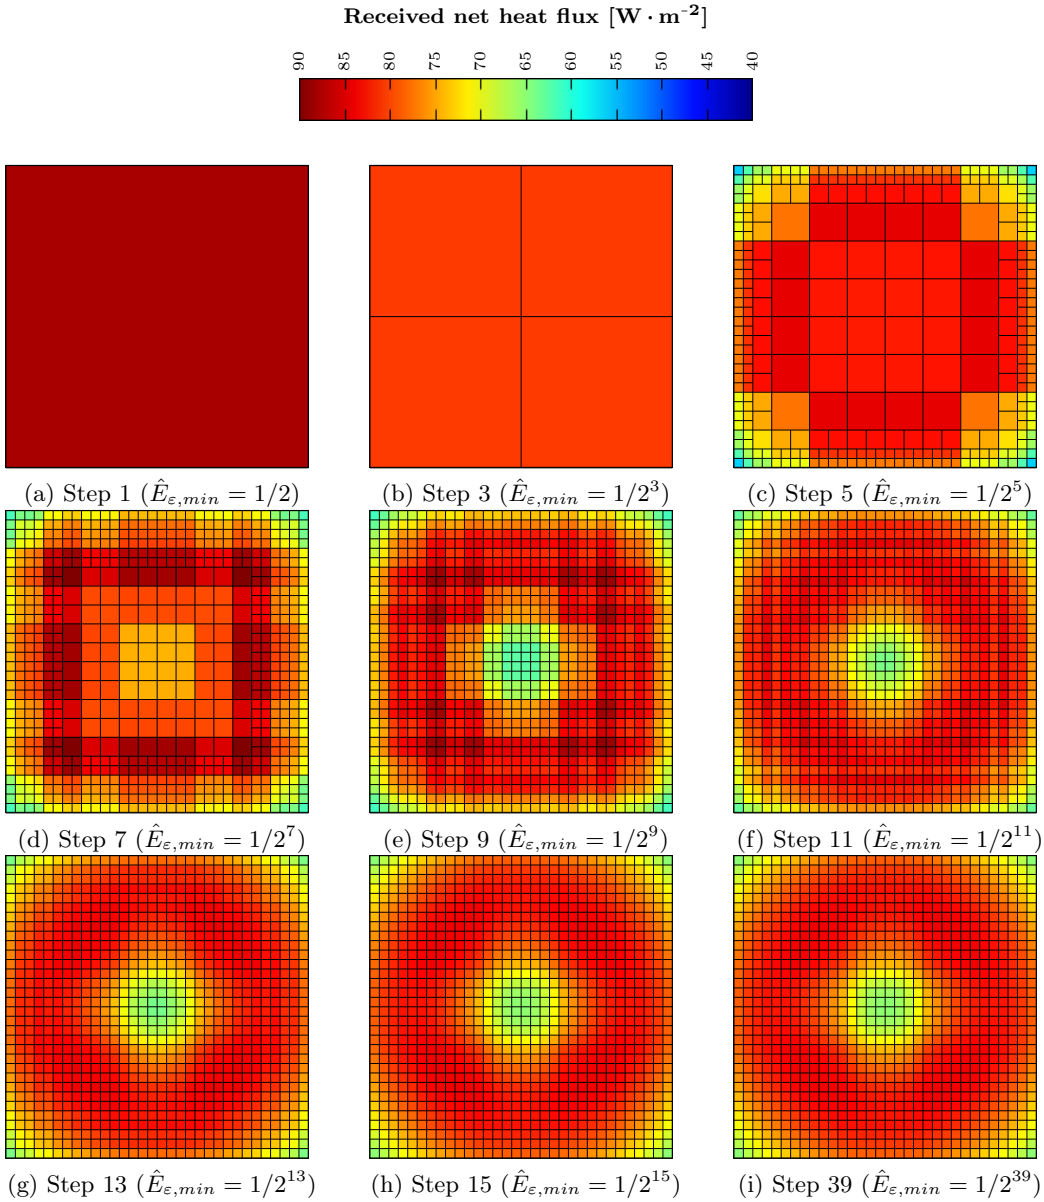

(a) Step 1 ($\hat{E}_{\varepsilon,min} = 1/2$) (b) Step 3 ($\hat{E}_{\varepsilon,min} = 1/2^3$) (c) Step 5 ($\hat{E}_{\varepsilon,min} = 1/2^5$)

(d) Step 7 ($\hat{E}_{\varepsilon,min} = 1/2^7$) (e) Step 9 ($\hat{E}_{\varepsilon,min} = 1/2^9$) (f) Step 11 ($\hat{E}_{\varepsilon,min} = 1/2^{11}$)

(g) Step 13 ($\hat{E}_{\varepsilon,min} = 1/2^{13}$) (h) Step 15 ($\hat{E}_{\varepsilon,min} = 1/2^{15}$) (i) Step 39 ($\hat{E}_{\varepsilon,min} = 1/2^{39}$)

Figure 16: Distribution of net flux on a face adjacent to the hot surface for different threshold values ($\varepsilon = 0.8$; $1/A_\varepsilon = 1024\,\text{elements/m}^2$).

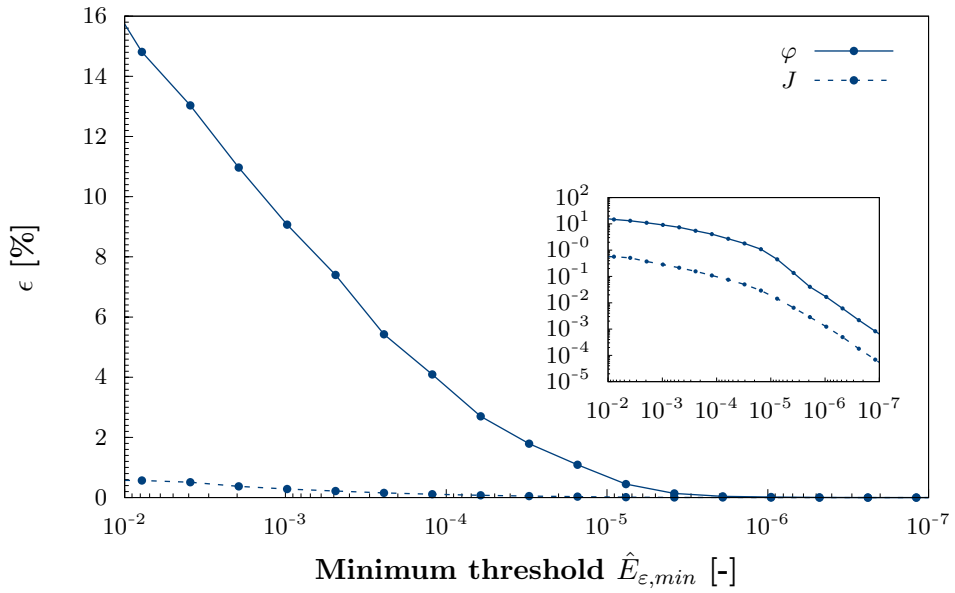

Figure 17: Evolution of the error of φ using the irradiation estimator in the chimney configuration.

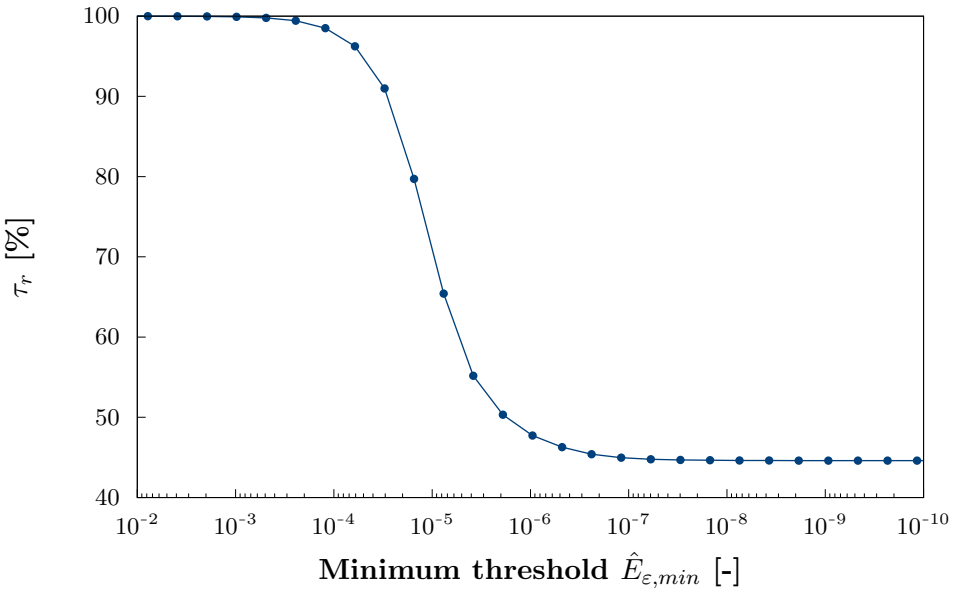

Figure 18: Evolution of the reduction ratio τ_r using the irradiation estimator in the chimney configuration.

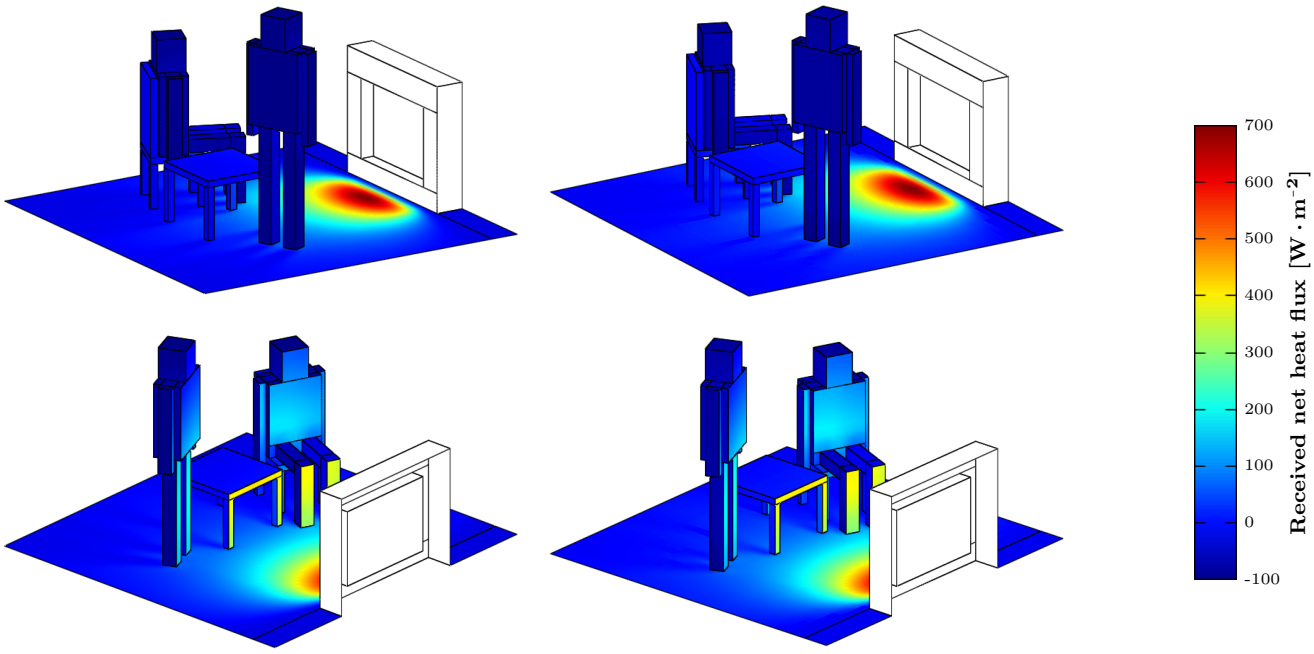

Figure 19: Net heat flux distribution in the elements of a furnished living room, computed with a minimum threshold of $\hat{E}_{\varepsilon,min} \approx 0$ (left) and $\hat{E}_{\varepsilon,min} = 1/2^{14} \approx 6.1 \times 10^{-5}$ (right).

Tables

$1/A_\varepsilon$ [elements/m^2]	Error with respect to the same resolution [%]	Error with respect to a resolution of 4096 [%]	τ_r [%]	$n_{F,hm}$ [-]
64	0.8	3.3	31.6	50316
256	2.3	3.2	80.4	230976
1024	2.9	3.1	97.4	498540

Table 1: Numerical values of errors, reduction ratio and number of view factor needed for computation associated with figure 10.