



HAL
open science

Coupling Conduction, Convection and Radiative Transfer in a Single Path-Space: Application to Infrared Rendering

Mégane Bati, Stéphane Blanco, Christophe Coustet, Vincent Eymet, Vincent Forest, Richard Fournier, Jacques Gautrais, Nicolas Mellado, Mathias Paulin, Benjamin Piaud

► **To cite this version:**

Mégane Bati, Stéphane Blanco, Christophe Coustet, Vincent Eymet, Vincent Forest, et al.. Coupling Conduction, Convection and Radiative Transfer in a Single Path-Space: Application to Infrared Rendering. ACM Transactions on Graphics, 2023, 42 (4), pp.1-20. 10.1145/3592121 . hal-04090428v2

HAL Id: hal-04090428

<https://hal.science/hal-04090428v2>

Submitted on 11 May 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NoDerivatives 4.0 International License

Coupling Conduction, Convection and Radiative Transfer in a Single Path-Space: Application to Infrared Rendering

MÉGANE BATI, IRIT, Université de Toulouse, UPS, CNRS, France
 STÉPHANE BLANCO, LAPLACE, Université de Toulouse, CNRS, INPT, UPS, France
 CHRISTOPHE COUSTET, Meso-Star, France
 VINCENT EYMET, Meso-Star, France
 VINCENT FOREST, Meso-Star, France
 RICHARD FOURNIER, LAPLACE, Université de Toulouse, CNRS, INPT, UPS, France
 JACQUES GAUTRAIS, CRCA, CBI, Université de Toulouse, CNRS, UPS, France
 NICOLAS MELLADO, CNRS, IRIT, Université de Toulouse, UPS, France
 MATHIAS PAULIN, IRIT, Université de Toulouse, UPS, CNRS, France
 BENJAMIN PIAUD, Meso-Star, France

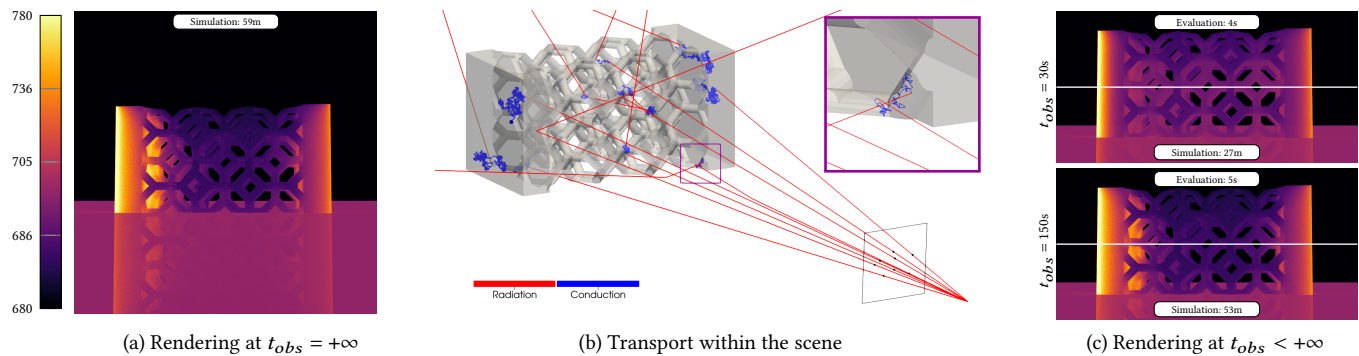


Fig. 1. We propose a Monte Carlo approach to tackle multiple physics with a single algorithm, translating their coupling into a single path-space composed of randomly chained sub-paths for each physics. Application is exemplified with heat transfer. (a) An infrared image of a steady state thermal exchanger, with temperature imposed on the left and right walls. (b) Monte Carlo paths alternate between heat-transfer modes (here conduction and radiation). (c) A huge benefit is the fast production of transient simulations, at any time, using the information gathered in (a), i.e. from only one Monte Carlo run at steady state.

In the past decades, Monte Carlo methods have shown their ability to solve PDEs, independently of the dimensionality of the integration domain and for different use-cases (e.g. light transport, geometry processing, physics simulation). Specifically, the path-space formulation of transport equations is a key ingredient to define tractable and scalable solvers, and we observe nowadays a strong interest in the definition of simulation systems based on Monte Carlo algorithms. We also observe that, when simulating combined physics (e.g. thermal rendering from a heat transfer simulation), there is a

Authors' addresses: Mégane Bati, IRIT, Université de Toulouse, UPS, CNRS, Toulouse, France, megane.bati@irit.fr; Stéphane Blanco, LAPLACE, Université de Toulouse, CNRS, INPT, UPS, Toulouse, France, stephane.blanco@laplace.univ-tlse.fr; Christophe Coustet, Meso-Star, Toulouse, France, christophe.coustet@meso-star.com; Vincent Eymet, Meso-Star, Toulouse, France, vincent.eymet@meso-star.com; Vincent Forest, Meso-Star, Toulouse, France, vincent.forest@meso-star.com; Richard Fournier, LAPLACE, Université de Toulouse, CNRS, INPT, UPS, Toulouse, France, richard.fournier@laplace.univ-tlse.fr; Jacques Gautrais, CRCA, CBI, Université de Toulouse, CNRS, UPS, Toulouse, France, jacques.gautrais@univ-tlse3.fr; Nicolas Mellado, CNRS, IRIT, Université de Toulouse, UPS, Toulouse, France, nicolas.mellado@irit.fr; Mathias Paulin, IRIT, Université de Toulouse, UPS, CNRS, Toulouse, France, mathias.paulin@irit.fr; Benjamin Piaud, Meso-Star, Toulouse, France, benjamin.piaud@meso-star.com.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in , <https://doi.org/10.1145/3592121>.

lack of coupled Monte Carlo algorithms allowing to solve all the physics at once, in the same path space, rather than combining several independent MC estimators, a combination that would make the global solver critically sensitive to the complexity of each simulation space. This brings to our proposal: a coupled, single path-space, Monte Carlo algorithm for efficient multi-physics problems solving.

In this work, we combine our understanding and knowledge of Physics and Computer Graphics to demonstrate how to formulate and arrange different simulation spaces into a single path space. We define a tractable formalism for coupled heat transfer simulation using Monte Carlo, and we leverage the path-space construction to interactively compute multiple simulations with different conditions in the same scene, in terms of boundary conditions and observation time. We validate our proposal in the context of infrared rendering with different thermal simulation scenarios: e.g., room temperature simulation, visualization of heat paths within materials (detection of thermal bridges), heat diffusion capacity of thermal exchanger. We expect that our theoretical framework will foster collaboration and multidisciplinary studies. The perspectives this framework opens are detailed and we suggest a research agenda towards the resolution of coupled PDEs at the interface of Physics and Computer Graphics.

CCS Concepts: • **Computing methodologies** → **Simulation types and techniques**; **Rendering**; **Physical simulation**; • **Applied computing** → **Physics**; **Mathematics and statistics**.

ACM Reference Format:

Mégane Bati, Stéphane Blanco, Christophe Coustet, Vincent Eymet, Vincent Forest, Richard Fournier, Jacques Gautrais, Nicolas Mellado, Mathias Paulin, and Benjamin Piaud. 2023. Coupling Conduction, Convection and Radiative Transfer in a Single Path-Space: Application to Infrared Rendering. 42, 4 (August 2023), 20 pages. <https://doi.org/10.1145/3592121>

1 INTRODUCTION

In the past decades, the simulation of radiative transfer has been extensively studied in Computer Graphics in the context of Physically Based Rendering (PBR), leading to the development of powerful models, algorithms and rendering systems [Jakob et al. 2022; Pharr et al. 2016]. Such systems are nowadays tailored to cope with the ever-increasing complexity of the scenes (e.g. in terms of geometry and materials) and the simulation needs (e.g. spectral complexity, coupling several physics). This is achieved by leveraging efficient data structures and data models to manage and access the scene data, and also by exploiting the scalability of Monte Carlo methods to solve the Radiative Transfer Equation in path-space.

In the present work, we advocate that these rendering approaches, which are originally motivated by movies and games use-cases, will lead to potential breakthroughs in other fundamental fields, as illustrated in climate science by Villefranque et al. [2022]. We propose to demonstrate this potential for the simulation of heat transfer, a process that involves *radiative transfer* (similarly to standard rendering), but also *conduction* and *convection*. In order to *couple* these three phenomena, we build upon the proposal of Tregan et al. [2023], which introduces a *single coupled path-space* for modeling the heat-transfer, and which is expected to be solved by Monte Carlo. Starting from this theoretical model, we formulate and demonstrate a practicable, tractable, and scalable Monte Carlo simulation framework. Unlike discretized methods or naive coupling, our framework allows to render infrared images *without having to resolve the temperature field first*. To generate a thermal image (see Figure 1(b)), rendering paths start from the camera, traverse the physics of coupled heat transfer in all its temporal dimensions, and lead back to the "sources": not the objects surface temperature, but the heating sources within the scene, or the boundary conditions.

A notable property of our approach is its scalability: it provides the capability to simulate coupled physical processes at very large scale with complex geometry, for instance a thermal simulation at the city scale (see Figure 4). In addition, the paths information can be collected and stored by running one *single* steady-state Monte Carlo simulation (Figure 1(a)), and can then be reused to simulate *any* transient state at *no cost*, e.g. to produce transient infrared movies rendering the energy propagation in capacitive objects (Figure 1(c)). The same path information can also be reused for instant rendering of any other boundary and initial conditions, providing interactive feedback (see the supplemental video). We believe that, combined with the visual analysis of the thermal paths throughout the scene, *interactive thermal simulation* is a game changer for use cases related to energy-transition, e.g. the rehabilitation of buildings.

From the very origin, the computer graphics community is strongly connected with physics and physicists have always heavily benefited of computer graphics. However, both communities have been historically structured upon their own agendas, and decisive advances had to first pervade one community through domain-specific communications before they became visible by the other. In this sense, even when they address closely similar questions, the two communities advance separately. For instance, physicists could import in their own field the usage of geometrical libraries and acceleration grids [Villefranque et al. 2019] only once they became a major reference in computer graphics [Wald et al. 2014]. In mirror, computer graphics researchers have imported null-collision algorithms only recently [Raab et al. 2008] despite very early references in particle transport [Coleman 1968]. Even more recent was the convergence towards common theoretical foundations of differential light transport in computer graphics and domain deformation sensitivities in transport physics [Howell and Daun 2021; Lapeyre et al. 2020; Li et al. 2018; Zhang et al. 2019; Zhao et al. 2020].

Our approach focuses on *coupling* multiple physics in a single path space. The idea of coupling is precisely one of these interface questions that has recently emerged in both Computer Graphics (multiscale coupling [Bitterli and d'Eon 2022; Guo et al. 2018; Heitz et al. 2016; Wang et al. 2022]) and Physics, in fields such as electromagnetic, solid physics, photochemistry [Dauchet et al. 2018, 2013; Galtier et al. 2016; Gattepaille et al. 2018]. We believe that computer graphics and physics research would strongly benefit from a common theoretical framework to foster collaboration on that topic. This leads us to share our understanding and knowledge of Physics and Computer Graphics to propose the following contributions:

- building upon the proposal of Tregan et al. [2023], we define a tractable formalism for heat transfer simulation using Monte Carlo, and present how we leverage the path-space information to interactively compute multiple simulations with different conditions in the same scene, in terms of boundary conditions and observation time (see Section 4),
- we present some results that highlight the practical significance of this proposal (see Section 5),
- the perspectives we expect from this framework are detailed in Section 6 and, beyond heat transfer, we explain how its core principles open avenues towards the resolution of *coupled* PDEs in a broader sense.

2 RELATED WORKS

Radiative transfer and Physically Based Rendering. In Computer Graphics, the radiative transfer is modeled either using the rendering equation [Kajiya 1986] on surfaces, or by the Radiative Transfer Equation [Chandrasekar 1960] (RTE) in volumes. In order to efficiently render images, these equations have been reformulated to model the simulation as a path space sampling problem [Veach 1998], which can be accelerated using dedicated data structures [Pharr et al. 2016]. The transfer on surfaces is modeled using Bidirectional Reflectance Distribution Functions (BRDF), e.g. modeling the micro-geometry of rough surfaces [Bitterli and d'Eon 2022; Wang et al. 2022] or diffractive ones [Holzschuch and Pacanowski 2017], glint paints [Chermain et al. 2020], or volumic micro-structures when

the medium is diffuse and/or multi-layer [Randrianandrasana et al. 2021]. These models might require a spectral [Wilkie et al. 2014] or wave-optics [Steinberg et al. 2022] rendering. In case of volumes, macroscopic transport can be modeled in participating media with spatial and spectral heterogeneities [Galtier et al. 2013; Kutz et al. 2017] efficiently stored in dedicated data structures [Villefranche et al. 2019], and can be computed using Monte Carlo, even if being non-linear. Considering that light has a finite speed, movies of light propagation can be generated using transient rendering [Jarabo et al. 2014; Marco et al. 2019].

Monte Carlo beyond radiative transfer. In order to model realistic scenes, the light transfer needs to be simulated, but it also requires some phenomena to be described by other physical rules, e.g. fluids dynamics for clouds or smokes [Rioux-Lavoie et al. 2022]. Hence, even if they focus on radiative transfer, contributions proposed by the computer graphics community are not restricted to this unique physic. For instance, the Walk-on-Sphere (WoS) algorithm, introduced in physics by Muller [1956], has been recently implemented by Sawhney and Crane [2020] to solve PDE-based geometry processing tasks. This algorithm has been extended to spatially varying coefficients [Sawhney et al. 2022] and bi-directional resolution from the sources [Qi et al. 2022]. These works open new perspective to process a wide range of physical models using Monte Carlo. However, in these works, when simulating infrared pictures, the WoS is used to compute temperatures using Monte Carlo, and only then, this information is used to render the final image, in a post-processing step [Sawhney et al. 2022]. By contrast, we focus in the present paper on the idea of *coupling* several physics, pushing the idea that multiple models can be simulated at once, processing physics and rendering in a single step, using a single path space.

Coupling several physics. Rather than relying on a precomputation, we focus on solving the *physical simulation at the same time as rendering*. This idea of coupling has already been introduced for light transport in rendering, even if not explicitly stated. Specifically, several works have proposed to model light interaction at *different scales* for materials with micro-geometries, e.g. rough [Bitterli and d'Eon 2022; Heitz et al. 2016; Wang et al. 2022], multi-layer [Guillén et al. 2020; Guo et al. 2018] of granular diffuse [Meng et al. 2015] materials. This scale coupling appears when evaluating the BRDF in the scene (macroscopic scale) using a microscopic Monte Carlo estimation. Usually, the goal is not to evaluate this microscopic estimation per se, but rather to build a model in pre-process, e.g., to minimize variance. For these reasons, most models use position-free approximations and ignore lateral transport in the material.

Except this coupling of the RTE at different scales, we are not aware of other work solving multiple physics coupling using Monte Carlo. For solving problems involving different physics, the definition of a single path space seems to be promising: there is no need then to use different Monte Carlo algorithm for going from one physical model to another, but rather to follow a unique path going from one space to another.

Paper organization. Coupled heat simulation (i.e. radiative transfer, convection and conduction, as reminded in Sec. 3) can indeed be expressed in a single path space (Sec 4.1), as demonstrated in

theoretical physics by Tregan et al. [2023]. We build upon this result to adapt its formulation to computational and practical constraints (Sec. 4.2). Our proposed coupled model is not only spatial: it is also temporal so that it handles unstationary states (Sec. 4.3). We then show how a single simulation allows to build propagation animations, as well as to change boundary conditions, using *path-replay* ideas [Mul 2016; Vicini et al. 2021] (Sec. 4.4).

3 THE PHYSICAL MODEL AND ITS DIDACTIC VERSION

In this section, we review the physical model of heat-transfer, i.e. convection, conduction and radiative transfer. First, we present the model in its general formulation. We recall the RTE with some thermal sources associated to matter in a state of local thermodynamic equilibrium (Sec. 3.1) and the advecto-diffusion model (Sec. 3.2) (i.e. conduction and convection), and we explain their coupling (Sec. 3.3). Since our concern lies essentially on the coupling, we then present a simplified version of the model (Sec. 3.4), which will be used hereafter. A synthetic summary of the fundamentals of thermodynamics is given in App. A.1. Heat-transfer specialists can directly jump to Section 4.

3.1 Thermal RTE

The model describing the transport of radiative energy is classically given by a) a linear single-velocity Boltzmann equation in the volume (Eq. 1), and b) the Rendering Equation on the opaque boundaries of the system (Eq. 2):

$$\vec{\omega} \cdot \vec{\nabla} L_\lambda = -\kappa_a [L_\lambda - L_\lambda^{eq}(\theta)] - \kappa_s L_\lambda + \int_{4\pi} \kappa_s p_s(\vec{\omega}|\vec{\omega}_i) L_{\lambda,i} d\vec{\omega}_i \quad (1)$$

$$L_\lambda = \epsilon L_\lambda^{eq}(\theta) + (1 - \epsilon) \int_{2\pi} p_r(\vec{\omega}|\vec{\omega}_i) L_{\lambda,i} d\vec{\omega}_i \quad (2)$$

where $L_\lambda \equiv L_\lambda(\vec{x}, \vec{\omega}, t)$ is the radiance (or specific intensity), $\vec{x}, \vec{\omega}, t, \lambda$ denote respectively the position, direction, time and wavelength, with $L_{\lambda,i} \equiv L_\lambda(\vec{x}, -\vec{\omega}_i, t)$. The absorption and scattering coefficients (κ_a and κ_s), the emissivity ϵ as well as the volumic and surface phase functions (p_s and p_r) are all quantities that potentially depend on $\vec{x}, \vec{\omega}, \lambda$. The Planck function $L_\lambda^{eq}(\theta)$ describes the equilibrium radiance, and depends on the temperature θ and the wavelength.

Note that radiative transfer is considered in its stationary version, due to the time scales of photon propagation compared to the time scales of other transfer modes. Nevertheless, radiance keep time dependency as soon as we consider coupled situations where the temperature θ evolves over time. Also note that, the variation of the index of refraction with the temperature being negligible in our context, the refractive form of the RTE [Ament et al. 2014] is not considered.

In radiative heat transfer, Equations 1 and 2 are strictly compatible with the equilibrium constraints. This imposes the specific form of the source terms (involving L_λ^{eq}) in the equations, regarding the emission and absorption parts.

3.2 Advection-diffusion

The conductive and convective processes that we find respectively in solid media and fluids are typically modeled by an advecto-diffusive

description of the temperature θ :

$$\rho c \frac{\partial \theta}{\partial t} = -\vec{\nabla} \cdot \left[-k \vec{\nabla} \theta + \rho c \vec{V} \theta \right] + p$$

which can be written as:

$$\frac{\partial \theta}{\partial t} = -\vec{\nabla} \cdot \left[-\alpha \vec{\nabla} \theta + \vec{V} \theta \right] + S \quad (3)$$

where $\theta \equiv \theta(\vec{x}, t)$, $\vec{V} \equiv \vec{V}(\vec{x}, t)$, $\alpha \equiv \alpha(\vec{x}) = \frac{k(\vec{x})}{\rho c}$ and $S \equiv S(\vec{x}, t) = \frac{p(\vec{x}, t)}{\rho c}$. The properties k , ρ and c are respectively the thermal conductivity, the density and the heat capacity of the material. The term $-k \vec{\nabla} \theta$ represents the surface density vector of conductive flux given by Fourier's law. The term $\rho c \vec{V} \theta$ is the temperature advection flux (transport of energy by the fluid movement). It vanishes in solids ($\vec{V} = \vec{0}$). For fluids, the velocity field \vec{V} is prescribed, and how it can be computed by fluid mechanics models is not discussed here (see Perspectives Sec. 6). Finally, in S , p is a source term of voluminal power.

Advecto-diffusive phenomenology appears in numerous fields of physics, well beyond thermal transfers. Classically, numerical grid methods are employed to obtain the solutions. These have the usual limitation in terms of scalability [Sawhney and Crane 2020] and statistical alternatives have emerged in the physics community as early as in the 1960s. In line with the theoretical work of Feynman-Kac [Feynman and Brown 2005; Itô et al. 1996; Kac 1947; Kac et al. 1951], who writes the solution of these equations as an expectation on drifted stochastic Brownian processes, several Monte Carlo computational alternatives have emerged. The seminal work on walk-on-sphere proposal [Muller 1956] introduces an elegant algorithm with outstanding convergence properties. Recent work by Sawhney et al. [2023] extends the initial WoS to cases with Neumann-like boundary conditions. Nevertheless, handling heterogeneous situations in diffusivity and velocity field remains an open question and this research area is active, even for stationary cases [Lejay and Maire 2010, 2013; Sawhney et al. 2022]. The instationarity brings an additional difficulty for the management of the events ending inside the sphere. Some bridges have been established between these proposals and the literature of first-pass Green methods [Deaconu et al. 2017; Deaconu and Lejay 2006, 2007; Hwang and Mascagni 2003; Karl 2019; Sabelfeld and Popov 2020], which open interesting perspectives but, until now, do not allow to scale up to practical geometric and temporal complexity.

3.3 Coupling

In most cases, the energy transfers within a system are based on the coupling of the aforementioned modes. The radiative transfer model is naturally coupled to the temperature field through the Planck function, which appears in both the volume and boundaries of the system (Eq. 1 and 2). The temperature field is given by the advecto-diffusive heat transfer equation (Eq. 3) which is affected itself by radiative exchanges. The coupling with radiation occurs in two different ways: through sources in the volume and at the system boundaries. Regarding the first point, the radiative source $S = \frac{p}{\rho c}$ in the volume (see Eq. 3) is no longer prescribed but expressed from

the radiative flux:

$$p = -\vec{\nabla} \cdot \left(\int_{4\pi} d\vec{\omega} \int_0^{+\infty} d\lambda L_\lambda \vec{\omega} \right), \quad (4)$$

where the dependency on the radiance L_λ becomes explicit. Regarding the second point, exchanges through system boundaries, we use the flux continuity property. The connection condition between a fluid domain and an opaque solid is expressed by the continuity of the temperature (θ is unique at the interface) and the continuity of energy fluxes. The conductive energy flux on the solid side of the interface $\varphi_{cond,S}$ is strictly equal to the sum of the radiative and convective fluxes on the fluid side. Since the velocity of the fluid is null at the boundary (no-slip condition), the convective flux is restricted to the conductive part. Hence:

$$\varphi_{cond,S}[\theta] = \varphi_{cond,F}[\theta] + \varphi_{rad}[L_\lambda], \quad (5)$$

where $[\cdot]$ denotes the operator dependence on temperature and radiance. Equation 5 shows the coupling at the surface between the variables θ and L_λ (the temperature field and the radiance at the boundary).

3.4 A simplified version of the model to enlighten the coupling

Now, we expose the simplified version of the heat transfer model that will be used in Sec. 4 to design a single path-sampling Monte Carlo algorithm for the coupling of conduction, convection and radiation. For didactic reasons, we will stick to the very basis of coupled heat transfer, making strong simplifications of the initial model presented above. Each simplification made here will be closely reconsidered in Perspectives (Section 6).

Opaque diffuse-gray solids and transparent fluids. The scene is considered to be only made of either opaque solids (no radiation inside the solids, only conduction) or transparent fluids. Radiative transfer reduces therefore to straight line propagation across fluids (neither diffusion nor absorption/emission in the fluid) so it is driven by the standard Rendering Equation (Eq. 2) at all locations along the solid-fluid interfaces.

The radiative properties of solid surfaces are independent of temperature. Their emissivity ϵ is independent of direction and wavelength (gray), and reflection can be either diffusive (cosine distributed) or specular (or a mix of both). Integrating over all frequencies, in terms of $L = \int_0^{+\infty} L_\lambda d\lambda$ and using $\int_0^{+\infty} \pi L_\lambda^{eq}(\theta_S) d\lambda = \sigma \theta_S^4$ (Stefan-Boltzmann law of black body emittance), Rendering Equation (2) writes:

$$L = \epsilon \frac{\sigma \theta_S^4}{\pi} + (1 - \epsilon) \int_{2\pi} L_i p_r(\vec{\omega}|\vec{\omega}_i) d\vec{\omega}_i \quad (6)$$

where σ is the Boltzmann constant and θ_S the solid surface temperature. The net radiative flux density at a solid surface of unit normal \vec{n} is:

$$\varphi_{rad} = \int_{2\pi} \left(\underbrace{\epsilon \frac{\sigma \theta_S^4}{\pi}}_{\text{emitted}} - \underbrace{\epsilon L_i}_{\text{absorbed}} \right) \vec{\omega}_i \cdot \vec{n} d\vec{\omega}_i \quad (7)$$




Fig. 2. *Construction of a coupled path.* (a) The path starts with a radiative sub-path γ initiated at the camera, potentially reflecting at the solid surface until absorption at \vec{x}_γ . The Monte Carlo weight is $\theta_S(\vec{x}_\gamma)$ if the solid surface temperature is known (Eq. 19). (b) When the solid surface temperature is not known, one of the three heat transfer modes is randomly selected (Eq. 23). If conduction is selected, the Monte Carlo weight is $\theta_S(\vec{x}_\gamma - \delta\vec{n})$, i.e. the temperature of the solid at an injection location δ -away from the interface. If convection is selected, the Monte Carlo weight is the fluid temperature θ_F . If radiation is selected, a direction $\vec{\omega}$ is sampled and the Monte Carlo weight is the incoming radiance temperature $\theta_R(\vec{x}_\gamma, \vec{\omega})$. If this radiance temperature is unknown, a new radiative sub-path is initiated as when starting from the camera. (c) In the conductive case, after injection within the solid, when the solid temperature is unknown, a walk-on-sphere path is followed until the solid border is found at a location \vec{x}_{wos} (Eq. 24). The Monte Carlo weight is then $\theta_S(\vec{x}_{wos})$. (d) When the reached location is at the interface between two solids and when the temperature of this interface is unknown, one of the two solids is randomly selected and the Monte Carlo weight is the temperature of the selected solid δ -away from the interface, i.e. $\theta_S(\vec{x}_\gamma - \delta_1\vec{n})$ or $\theta_S(\vec{x}_\gamma + \delta_2\vec{n})$ (Eq. 26). (e) In the convective case, when θ_F is unknown, a location \vec{x}_F is sampled at the surface of the solid surrounding the fluid cell. The Monte Carlo weight is the solid temperature $\theta_S(\vec{x}_F)$ at this location (Eq. 28). (f) The algorithm alternates between different sub-path modes until it reaches a known temperature.

Uniform solids by parts. Solids are rigid (no deformation velocity field) and are divided into sub-parts where conductivity and specific heat capacity are uniform. For each sub-part, Eq. 3 therefore turns into pure diffusion:

$$\frac{\partial \theta_S}{\partial t} = \alpha \Delta \theta_S, \quad (8)$$

where $\theta_S \equiv \theta_S(\vec{x}, t)$ is the solid temperature. At any location at an interface of unit normal \vec{n} between two solids (denoted $s1$ and $s2$), the conductive flux density on both sides is defined as

$$\varphi_{cond} = -k\vec{\nabla}\theta_S \cdot \vec{n}, \quad (9)$$

while flux continuity implies that

$$\varphi_{cond,s1}(\vec{x}, t) = \varphi_{cond,s2}(\vec{x}, t). \quad (10)$$

Perfectly mixed fluid cells. All connected fluid cells are assumed isothermal except for thin boundary layers along the solid-fluid interfaces. Across these layers, the heat flux density φ_{conv} is modeled as proportional to the difference between the temperature θ_F of the isothermal fluid core and the surface temperature θ_S of the facing solid (Newton law):

$$\varphi_{conv} = h(\theta_S - \theta_F). \quad (11)$$

The proportionality coefficient h (the convective exchange coefficient), via the Nusselt number, is a direct translation of the fluid flux inside the boundary layer and is assumed to be independent of the temperature. Under this assumption, as detailed in App. A.1, Equation 3 can be rewritten:

$$\rho c \mathcal{V} \frac{d\theta_F}{dt} = \int_{\mathcal{S}} h(\vec{x}) (\theta_S - \theta_F) d\vec{x}, \quad (12)$$

where $\theta_F \equiv \theta_F(t)$, \mathcal{V} is the volume of the fluid cell and \mathcal{S} is its boundary. At any locations along a solid-fluid interface, there is conduction on the solid side and both convection and radiation on the fluid side. Here, energy conservation implies that

$$\varphi_{cond}(\vec{x}, t) = \varphi_{conv}(\vec{x}, t) + \varphi_{rad}(\vec{x}, t). \quad (13)$$

We note that in spite of the simplifications made here, this model is the one that is most commonly used in many application areas.

4 COUPLING THE THREE HEAT TRANSFER MODES

In this section, we detail how to extend the model described in Sec. 3.4 in order to make it tractable for real-world scenarios. The key physical interactions modeled in our proposal (i.e. transport in medium and interactions at boundaries) define the main steps of our algorithm. The following section is structured around these interactions, the corresponding equations are not only highlighted but also illustrated in Figure 2 in order to simplify the reading. The pseudocode is provided in App. B.

Overview. In Section 4.1, we define how we model *backward* paths (from camera to sources) propagating through the entire system. These paths combine the three heat transfer modes, and transport temperature information from locations where temperature is imposed (sources) to locations where temperature is sought (camera pixels). Section 4.2 shows how to address the geometric/physical information using only line-surface intersection operators. For this, we replace the standard walk-on-sphere algorithm with a delta-sphere random walk algorithm for thermal diffusion inside solids. In Section 4.3 we go beyond the stationary case and extend our model to transient heat-transfer. Finally, we analyze in Section 4.4 the physical pictures associated to the resulting path statistics, and how this information can be used to reconstruct new infrared images very rapidly when changing the sources. The question of further using this propagation information is addressed in Section 5.

4.1 Paths combining the three heat transfer modes

Let us start with pure radiative transfer, typically addressing the Monte Carlo estimation of the spectrally integrated radiance $L \equiv L(\vec{x}, \vec{\omega})$ at the location \vec{x} of an infrared camera inside the fluid, as in Fig. 2. A standard russian roulette algorithm leads to the backward sampling of a multiple reflection random path Γ ending at location

\vec{x}_Γ on the solid-fluid interface, the Monte Carlo weight being the infrared source at this location, i.e. $W_\Gamma = \frac{\sigma\theta_S^4(\vec{x}_\Gamma)}{\pi}$. If we temporarily assume that the surface temperature of the solid is known, then the Monte Carlo algorithm can be summarized as the following expectation:

$$L = \mathbb{E} \left[\frac{\sigma\theta_S^4(\vec{x}_\Gamma)}{\pi} \right]. \quad (14)$$

Defining the radiance temperature as

$$\theta_R = \left(\frac{1}{\sigma} \pi L \right)^{1/4}, \quad (15)$$

and assuming that all temperatures are close to a known reference temperature θ_{ref} , it is standard practice to linearize radiance as function of radiance temperature using the 1-order development

$$L = \frac{\sigma\theta_R^4}{\pi} \simeq \frac{\sigma\theta_{ref}^4}{\pi} + \frac{4\sigma\theta_{ref}^3}{\pi} (\theta_R - \theta_{ref}). \quad (16)$$

In the Monte Carlo context, this linearization can be repeated inside the expectation in Eq. 14:

$$L \simeq \mathbb{E} \left[\frac{\sigma\theta_{ref}^4}{\pi} + \frac{4\sigma\theta_{ref}^3}{\pi} (\theta_S(\vec{x}_\Gamma) - \theta_{ref}) \right] \quad (17)$$

$$= \frac{\sigma\theta_{ref}^4}{\pi} + \frac{4\sigma\theta_{ref}^3}{\pi} (\mathbb{E}[\theta_S(\vec{x}_\Gamma)] - \theta_{ref}) \quad (18)$$

and the right sides of Eq. 16 and 18 can be equaled to give the radiance temperature expressed itself as an expectation:

radiance temperature

$$\theta_R = \mathbb{E} [\theta_S(\vec{x}_\Gamma)]. \quad (19)$$

As detailed in supplementary material (Sec. 3), the linearization approximation can be removed, in which case Equation 19 expresses θ_R as the expectation of a random variable whose definition involves θ_R . We denote this situation as *non-linear coupling*, further discussed in Section 6.

With these equations in hand, we have a Monte Carlo algorithm where the radiance temperature is modeled as an expectation of solid-surface temperatures. But in a coupled heat transfer problem, there are surfaces where the temperature is not known. Instead, it results from the combined presence of conduction on the solid side, and convection plus radiation on the fluid side. Equation 13 translates this combination ; combined with the previous definitions and the same linearization of $(\sigma\theta_S^4 - \sigma\theta_R^4)$ as for Eq. 16, it leads to

$$-k\vec{\nabla}\theta_S \cdot \vec{n} = h(\theta_S - \theta_F) + \int_{2\pi} h_R(\theta_S - \theta_R) \frac{\vec{\omega} \cdot \vec{n}}{\pi} d\vec{\omega} \quad (20)$$

where $h_R = 4\sigma\epsilon\theta_{ref}^3$.

The term $\vec{\nabla}\theta_S \cdot \vec{n}$ can be approximated using finite difference:

$$-\vec{\nabla}\theta_S \cdot \vec{n} \approx \frac{\theta_S(\vec{x} - \delta\vec{n}) - \theta_S(\vec{x})}{\delta}, \quad (21)$$

where δ must be small compared to the solid thickness. Thanks to this formulation, we can express θ_S as an expectation:

$$\theta_S(\vec{x}) = P_{cond} \theta_S(\vec{x} - \delta\vec{n}) + P_{conv} \theta_F + P_{rad} \int_{2\pi} \theta_R \frac{\vec{\omega} \cdot \vec{n}}{\pi} d\vec{\omega}, \quad (22)$$

with $P_{conv} = \frac{h}{\frac{k}{\delta} + h + h_R}$, $P_{rad} = \frac{h_R}{\frac{k}{\delta} + h + h_R}$ and $P_{cond} = 1 - (P_{conv} + P_{rad})$. Note that, contrary to the others simplifications and approximations we made previously, approximation (21) is required to express θ_S as an expectation, and can not be released.

To represent the test with three branches, we define the random vector $(\mathcal{H}_{cond}, \mathcal{H}_{conv}, \mathcal{H}_{rad})$ on support $\{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$ with respective probabilities P_{cond} , P_{conv} and P_{rad} , and defining $\vec{\Omega}$ as a random direction cosine-distributed around \vec{n} , then

solid temperature at a solid-fluid interface

$$\theta_S(\vec{x}) = \mathbb{E} \left[\mathcal{H}_{cond} \theta_S(\vec{x} - \delta\vec{n}) + \mathcal{H}_{conv} \theta_F + \mathcal{H}_{rad} \theta_R(\vec{x}, \vec{\Omega}) \right]. \quad (23)$$

Reporting Equation 19 into Equation 23 we can use double randomization [Maire and Nguyen 2016]: expectation of an expectation is an expectation and therefore linearly combining two path-sampling Monte Carlo algorithm leads to one single path-sampling Monte Carlo.

So we can now estimate the radiance temperature at an infrared camera location by sampling a multiple reflection radiative path until it stops at a solid-fluid interface. Then, in case the solid surface temperature is not known, the algorithm simply continue the same path by

- randomly choosing the heat transfer mode according to P_{cond} , P_{conv} and P_{rad} ;
- in case conduction is selected, retaining the solid temperature $\theta_S(\vec{x} - \delta\vec{n})$ at a distance δ away from the interface;
- in case convection is selected, retaining the fluid temperature θ_F ;
- in case radiation is selected, first sampling the direction $\vec{\omega}$ and retaining the radiance temperature $\theta_R(\vec{x}, \vec{\omega})$.

The same logic of double randomization is used to continue the algorithm for the selected heat transfer mode when the temperature to be retained is unknown.

The conduction term in Equation 23 requires computing $\theta_S(\vec{x} - \delta\vec{n})$. This can be translated as the sampling of a walk-on-sphere random path WOS within the corresponding uniform solid sub-part, the path ending at a location \vec{x}_{WOS} along the sub-part boundary:

walk-on-sphere

$$\theta_S(\vec{x} - \delta\vec{n}) = \mathbb{E} [\theta_S(\vec{x}_{WOS})]. \quad (24)$$

If at \vec{x}_{WOS} the solid temperature is known, the algorithm stops retaining this temperature as the Monte Carlo weight. If the solid temperature is not known there are two cases. Either the sub-part boundary is a solid-fluid interface, in which case we come back to Equation 23 and continue the algorithm by selecting again one heat transfer mode. Or the sub-part boundary is a solid-solid interface, in which case we need to make a finite difference approximation on

both sides. Equation 10 becomes

$$k_1 \frac{\theta_S(\vec{x} - \delta_1 \vec{n}) - \theta_S(\vec{x})}{\delta_1} \approx k_2 \frac{\theta_S(\vec{x}) - \theta_S(\vec{x} + \delta_2 \vec{n})}{\delta_2}, \quad (25)$$

and

solid temperature at a solid-solid interface

$$\theta_S(\vec{x}) = \mathbb{E} \left[\mathcal{H}_{cond,1} \theta_S(\vec{x} - \delta_1 \vec{n}) + \mathcal{H}_{cond,2} \theta_S(\vec{x} + \delta_2 \vec{n}) \right], \quad (26)$$

where the Bernoulli random variables $\mathcal{H}_{cond,1}$ has probability $P_{cond,1} = \frac{k_1}{\delta_1} / (\frac{k_1}{\delta_1} + \frac{k_2}{\delta_2})$ and $\mathcal{H}_{cond,2} = 1 - \mathcal{H}_{cond,1}$. So the algorithm continues with either a δ_1 jump inside solid 1 or a δ_2 jump inside solid 2, and starts a new walk-on-sphere algorithm from the reached location.

The convection term in Equation 23 requires computing θ_F . The stationary version of Equation 12 (i.e. $\frac{d\theta_F}{dt} = 0$) leads to

$$\theta_F = \frac{\int_S h(\vec{x}) \theta_S(\vec{x}) d\vec{x}}{\int_S h(\vec{x}) d\vec{x}}, \quad (27)$$

and therefore

fluid temperature

$$\theta_F = \mathbb{E} \left[\theta_S(\vec{X}_S) \right], \quad (28)$$

where \vec{X}_S is a random location along the fluid cell boundary with probability density proportional to the convective exchange coefficient. So when θ_F is needed, a location is sampled on the boundary, with location probability as a function of the intensity of convective exchanges, and retains the solid temperature at this location.

Finally, the radiation term in Equation 23 requires computing $\theta_R(\vec{x}, \vec{\omega})$, and this is identical to our initial question of estimating θ_R at a camera location.

Altogether, as summarized in Figure 2, Equations 19, 23, 26 and 28, combined with a walk-on-sphere algorithm for pure conduction within the uniform solid sub-parts, provide a closed Monte Carlo algorithm. This algorithm samples paths across the system, switches between heat transfer modes at the solid-fluid interfaces, ends when reaching a location where the temperature is known and retains this temperature value as the Monte Carlo weight. At this stage, although the core of the algorithm is defined, its efficient implementation impacts the choices made on the physical model itself in order to benefit from complex geometry scalability.

4.2 Walk-on-delta-sphere

According to our hypothesis, the solids are divided into uniform sub-parts and the fluids are transparent, which means that all the parts properties can be accessed from their envelop. Same thing for the heat transfer parameters: conductivity and diffusivity on both sides for solid-solid interfaces; conductivity and diffusivity on the solid side, emissivity and convective heat transfer coefficient on the fluid side for solid-fluid interfaces. Fast access to these data boils down to the fast sampling of surface locations for each of the three heat transfer modes.

As far as radiation is concerned, standard ray-tracing strategies are immediately available to optimize the computational effort associated to line-surface intersections [Pharr et al. 2016]. For convection, sampling \vec{X}_S is trivial: there is no dependence on location and cumulated distributions can be pre-constructed for the geometrical primitives (e.g. triangles) of the envelope, weighting them by their size and the local value of the convective exchange coefficient. For conduction, walk-on-sphere algorithms open a new set of questions. If the discussion was strictly restricted to the basic heat transfer model of the present section, then we could simply use the acceleration strategies that have recently been proposed for finding the triangle corresponding to the smallest tangent sphere as function of the sphere-center location [Shellshear and Ytterlid 2014; Wald et al. 2014]. But our question is broader: we want to set up path-sampling algorithms flexible enough for extension to heterogeneous diffusivity, heterogeneous velocity (when advection is considered), volume-coupling with radiation within semi-transparent media (fluids or solids), and nonlinear coupling with additional physical models (e.g., diffusivity and velocity depend on temperature). These questions are at least partially addressed in the walk-on-sphere literature [Sabelfeld 2019; Sabelfeld and Smirnov 2021], but they remain very open. We then take a practical decision: considering that we already made the choice of using a finite difference approximation for the handling of solid-solid and solid-fluid interfaces, we choose to replace the walk-on-sphere algorithm with a walk-on- δ -sphere algorithm [Tregan et al. 2023], i.e. sampling the next location on a sphere of fixed radius δ instead of the tangent sphere. For stationary diffusion inside uniform solids, exactly like walk-on-sphere, the algorithm is exact far from the boundary. Starting from a location \vec{x} , we handle the presence of the boundary with the following procedure:

- We sample an isotropically distributed random direction $\vec{\Omega}$.
- We test if a triangle is intersected before the distance δ in direction $\vec{\Omega}$. If an intersection is found we store the intersection distance s^+ , otherwise $s^+ = \delta$.
- We do the same, defining s^- for the opposite direction $-\vec{\Omega}$.
- We define $\tilde{\delta} = \min(s^-, s^+)$, and we jump at $\vec{x} + \tilde{\delta} \vec{\Omega}$, i.e. on the sphere of radius $\tilde{\delta}$ (which will be of radius δ if no triangle was intersected).

This is a first order spatial approximation (i.e. exact when the temperature profile is trilinear). The pseudo-code of this algorithm, as well as the full Monte Carlo procedure are given in App. B.

4.3 Extension to transient heat transfer

We model transient conduction and convection by keeping the temporal terms in Equations 8 and 12, respectively. During the path sampling, we also track time, such that the backward construction of the path may stop because the initial time t_I is reached. As the initial conditions of the system are known, all temperatures are known at t_I , whatever the reached location of the Monte Carlo weight.

Convection. Because of the kinetic nature of Eq. 12, time tracking can be directly computed for convection. The formal solution writes

$$\theta_F(t) = \exp\left(-\frac{t-t_I}{\tau_{conv}}\right)\theta_F(t_I) + \int_0^{t-t_I} \frac{1}{\tau_{conv}} \exp\left(-\frac{t_b}{\tau_{conv}}\right)\theta_F^*(t-t_b) dt_b \quad (29)$$

with $\tau_{conv} = \frac{\rho c V}{h_S}$ and $\bar{h} = \frac{1}{S} \int_S h(\vec{x}) d\vec{x}$, and where $\theta_F^*(t)$ is the same average surface temperature as in the stationary case. It was already formulated as an expectation in Eq. 28, leading to:

$$\theta_F^*(t) = \frac{1}{h_S} \int_S h(\vec{x}) \theta_S(\vec{x}, t) d\vec{x} = \mathbb{E} \left[\theta_S(\vec{X}_S, t) \right]. \quad (30)$$

Observing that $\exp\left(-\frac{t-t_I}{\tau_{conv}}\right) = \int_{t-t_I}^{+\infty} \frac{1}{\tau_{conv}} \exp\left(-\frac{t_b}{\tau_{conv}}\right) dt_b$, Equation 29 can be written as an expectation by defining an exponentially distributed random variable $T_{b,conv}$ of mean τ_{conv} , itself defining $\mathcal{H}_{I,conv} = H(T_{b,conv} - t + t_I)$, where H is the Heaviside function. $\mathcal{H}_{I,conv}$ is 1 if the initial condition is retained, 0 otherwise:

$$\begin{aligned} \theta_F(t) &= \int_0^{+\infty} \frac{1}{\tau_{conv}} \exp\left(-\frac{t_b}{\tau_{conv}}\right) \\ &\quad \left[H(t_b - t + t_I) \theta_F(t_I) + H(-t_b + t - t_I) \theta_F^*(\vec{x}, t - t_b) \right] dt_b \\ &= \mathbb{E} \left[\mathcal{H}_{I,conv} \theta_F(t_I) + (1 - \mathcal{H}_{I,conv}) \mathbb{E} \left[\theta_S(\vec{X}_S, t - T_{b,conv}) \right] \right] \\ &= \mathbb{E} \left[\mathcal{H}_{I,conv} \theta_F(t_I) + (1 - \mathcal{H}_{I,conv}) \theta_S(\vec{X}_S, t - T_{b,conv}) \right] \end{aligned} \quad (31)$$

The resulting transient convective sub-path is sampled the following way:

- $T_{b,conv}$ is sampled and is interpreted as an exponentially distributed backward jump in time;
- when this jump leads to a time $t - T_{b,conv}$ before t_I , the initial fluid temperature is retained;
- otherwise \vec{X}_S is sampled to get a position at the surface of the solid surrounding the fluid (as for the stationary path) and the solid temperature is retained, at this location at time $t - T_{b,conv}$.

Conduction. It is handled exactly the same way as convection, as detailed below. Up to now we used a first order approximation in space for diffusion when estimating the temperature of a solid-solid or a solid-fluid interface (re-injection inside the solid at distance δ when conduction is selected) and when adapting the walk-on- δ -spheres algorithm in the vicinity of the boundary. We use the same level of spatial approximation to write the following approximate version of transient diffusion equation (see Eq. 8). As detailed in App. A.2, this is done by using a finite difference to transform the Laplace operator in an arbitrarily chosen Cartesian coordinate system and then averaging over all the possible orientations of the coordinate system:

$$\frac{\partial \theta_S}{\partial t} = \alpha \left(-\frac{6}{\delta^2} \theta_S + \frac{6}{\delta^2} \theta_S^* \right), \quad (32)$$

where $\theta_S^*(\vec{x}, t)$ is the same average temperature over the δ -sphere as for the stationary case. It can be translated as an expectation using




Fig. 3. A stationary Monte Carlo estimation of a probe (triangle) temperature (left) is a weighted sum of the sources, here the known exterior (dot) and heater (cross) temperatures. The propagator enables to evaluate the samples under different source temperatures (top) shown by different colors, without performing a new simulation. The observation time can also be changed (bottom) for an extra-memory cost, then the initial temperature also becomes a source (squares).

the isotropic random direction $\vec{\Omega}$:

$$\theta_S^*(\vec{x}, t) = \mathbb{E} \left[\theta_S(\vec{x} + \delta \vec{\Omega}, t) \right]. \quad (33)$$

The formal solution of Equation 32 is

$$\begin{aligned} \theta_S(\vec{x}, t) &= \exp\left(-\frac{t-t_I}{\tau_{cond}}\right)\theta_S(\vec{x}, t_I) + \\ &\quad \int_0^{t-t_I} \frac{1}{\tau_{cond}} \exp\left(-\frac{t_b}{\tau_{cond}}\right)\theta_S^*(\vec{x}, t-t_b) dt_b \end{aligned} \quad (34)$$

which is consistent with the solution of Eq. 8 as detailed in supplementary material (Sec. 2). As for Eq. 29, we directly write

$$\theta_S(\vec{x}, t) = \mathbb{E} \left[\mathcal{H}_{I,cond} \theta_S(\vec{x}, t_I) + (1 - \mathcal{H}_{I,cond}) \theta_S(\vec{x} + \delta \vec{\Omega}, t - T_{b,cond}) \right], \quad (35)$$

and the corresponding sampling algorithm is exactly the same as for transient convection.

4.4 Storing and displaying path information

For each camera pixel, the above described algorithm is used to sample n paths starting from the camera at the observation time. The path traverses the system backward in time, switching from one heat transfer mode to the other, until either the initial time is reached (initial condition) or it crosses a location where the temperature is known (imposed temperature). These are the two only ways for a path to find its end, which means that the Monte Carlo weight is always the value of the temperature at the end of the path. As for standard rendering, generating the image does not require to store the complete history of the path in addition to the Monte Carlo weights. But in the context of thermal simulation, there are several practical use-cases where storing the entire path space can have a very strong impact. For instance, in the context of energy transition,

engineers and citizens can use infrared cameras to understand the influences of a fluctuating environment, get insight about their heat control choices, or about the insulation properties of building. From this perspective, statistically analyzing the heat transfer paths gives more information than a simple average of the associated weight (i.e. the radiative flux at the pixel surface). Displaying the paths within any scientific viewer has already a strong practical significance by itself. The same statement is made to draw an analysis from the visualization of light paths [Kopylov and Dmitriev 2000], we stick here to simple visualizations although more advanced methods exist (e.g. [Zirr et al. 2015]).

We illustrate this point in Section 5, together with two other useful practices that emerge as straightforward consequences of the present algorithmic proposal due to the linearity of the physical model. First, storing only *the end of the path* allows to very rapidly replay the rendering for new initial conditions or new imposed temperatures (see Figure 3). This idea is similar to relighting [Mul 2016] where sources can be modified after the simulation. Note that the term *path-replay* comes from the work of Vicini et al. [2021]. Second, storing *all surface encounters* allows to change the observation time. Indeed, when replaying the rendering at any date between the initial condition and the observation time for which the paths were constructed, the initial condition will be found somewhere between two known surface encounters. Specifically, it will be located inside an identified uniform solid sub-part or an identified fluid cell. Using this information, the Monte Carlo weight can be directly recomputed with a known initial temperature value. In practice, this means that a single Monte Carlo run computes all the required information to generate a sequence of infrared images showing the temporal response evolving over time, for a fixed camera location and a fixed geometry (see Figure 1 and supplementary video).

5 RESULTS

In this section, we show the usefulness of coupling several physics and manipulating the propagator on some concrete applications, mainly related to building construction and rehabilitation. The theoretical tools enable to analyze infrared renderings in the thermal context. Here, we show some results on infrared images, but it is much more general: a flux, an electric consumption, etc. could also be computed the same way.

Implementation details. We developed our prototype implementation using *Stardis* [Sta 2018] for Monte Carlo simulations and evaluations, restricting ourselves to the didactic model (i.e. having uniform conditions). This assumption is further discussed in supplementary material (Sec. 4). We implemented the path-replay of an infrared image and the time propagator, the later being restricted to a uniform initial temperature. The surfacic phase function p_r is composed of a specular part and a Lambertian part, sampled accordingly to the specular fraction parameter s (the probability that reflection is specular otherwise diffuse). We set the reinjection distance δ about 1/20 of the object's characteristic size. For all the presented infrared images, we provide in table 1 either their simulation or evaluation times, performed on an Intel Xeon Silver 4208 machine with 32 CPUs. As our implementation is not optimized, these timings might

Table 1. Scene configurations and associated computation times, for either full Monte Carlo simulation (s) or replay (r). For teaser, replay times are only two orders of magnitude smaller than MC times. This is because the replay code anticipates for further internal sources (see supplemental). For farm simulations, the replay code has been simplified for engineering needs and the replay to Monte Carlo ratio is of the order of $5 \cdot 10^{-7}$ (Fig. 6 replay times estimated on average over a large number of run replications).

| Scene | triangles | size | spp | time |
|-------------------|-----------|-----------|-------|------------------------|
| Fig. 1 – propag. | 13k | 1024x1024 | 256 | s: 59m |
| Fig. 1 (a) | 13k | 1024x1024 | 256 | r: 6s |
| Fig. 1 (c) – top | 13k | 1024x1024 | 256 | r: 4s |
| Fig. 1 (c) – top | 13k | 1024x1024 | 256 | s: 27m |
| Fig. 1 (c) – bot. | 13k | 1024x1024 | 256 | r: 5s |
| Fig. 1 (c) – bot. | 13k | 1024x1024 | 256 | s: 53m |
| Fig. 4 (a) | 9,97k | 1024x768 | 1024 | s: 10h20m |
| Fig. 4 (b) | 89,6k | 1024x768 | 1024 | s: 9h44m |
| Fig. 4 (c) | 806k | 1024x768 | 1024 | s: 9h33m |
| Fig. 4 (d) | 9,6M | 1024x768 | 1024 | s: 9h14m |
| Fig. 5 int. | 380 | 1024x1024 | 3000 | s: 25h02m |
| Fig. 5 ext. | 380 | 1024x1024 | 3000 | s: 13h58m |
| Fig. 6 int. | 380 | - | 10000 | s: 15s |
| Fig. 6 int. | 380 | - | 10000 | r: $8 \cdot 10^{-6}$ s |
| Fig. 6 ext. | 380 | - | 10000 | s: 2m04s |
| Fig. 6 ext. | 380 | - | 10000 | r: $8 \cdot 10^{-6}$ s |

be further reduced, (e.g., using vectorization). In any case, the purpose of our prototype implementation is to demonstrate that the didactic model already enables interesting results and analysis of the coupled heat transfer.

Validation. As there is no other code against which to validate the infrared images, the *Stardis* code is validated on temperature field computation and all the references are available at the webpage <https://www.meso-star.com/projects/stardis/validation.html>. These validations are made against analytical models on simple geometries and cross-comparisons with commercial finite difference codes (*ANSYS Fluent* [Flu 2017] and *Comsol* [Com 2020]) on more complex geometries. Note that the open access *Syrthes* [Syr 2021] code embeds *Stardis*, which allows systematic comparisons on the *exact* same CAD material. These validations are provided for cases that remain tractable by deterministic approaches (requiring the discretization of the whole volume) but cannot be made for complex geometries such as cities where solving on mesh is impractical.

The propagator implementation is validated against the classical Monte Carlo simulation, both on a temperature computation (see Fig. 7) and infrared renderings (see Fig. 1 – right).

5.1 Infrared imaging using coupled-paths : the case of the Heat exchanger

We first illustrate how we use coupled-paths to generate an infrared image, here of a heat exchanger (in Fig. 1) made of a Kelvin cells foam ($k = 237 \text{ W m}^{-1} \text{ K}^{-1}$, $\rho = 2700 \text{ kg m}^{-3}$, $c = 890 \text{ J K}^{-1} \text{ kg}^{-1}$) with emissivity $\epsilon_e = 0.9$ and specular fraction $s_e = 0.5$ for the reflective part, placed on a floor with emissivity $\epsilon_f = 0.8$ and a specular fraction $s_f = 0.8$. The whole is plunged within a uniform radiative background maintained at 680K. For the sake of clarity, we only

consider coupled conductive and radiative paths in this illustration. At initial time, the temperature of the system is set at $\theta_I = 700\text{K}$ everywhere, the left and right walls of the exchanger are set at 850K and 750K respectively, and keep being imposed. The temperature of the floor remains imposed at $\theta_I = 700\text{K}$. Left panel in Fig. 1 shows the infrared rendering at infinite time ($t_{obs} = +\infty$), i.e. the stationary state: imposed temperatures at the walls propagate towards the core of the exchanger, as expected, but not very deeply (for the given set of parameters) because the central part is conversely cooled by radiative exchange with background and floor.

Middle panel illustrates how this rendering has been built using coupling between conductive and radiative paths. Starting from pixels of the virtual camera, some radiative paths propagate back to the structure up to an intersection point. There, they can switch to conduction within the material or be reflected. Then, they chain conductive paths through the material (blue subpaths) and radiative paths between internal faces (red subpaths), until they either meet a left/right wall or the floor where the temperature is known (e.g. blue dot on the left wall), or exit to radiative background. Using such paths, each pixel value can be estimated in a single procedure integrating the fully coupled thermal transfer, without having to compute the temperature field prior the rendering.

Visualizing paths helps to interpret the results. For instance, the probability that a path ends at the known temperature of a wall is higher for paths that first intersect the structure near either wall. Hence such paths will convey more often the wall information (850K or 750K) than the background radiative value (680K) or floor (700K). Conversely, paths that meet the structure near the core will more often end by radiative exits to background or floor, and meet a wall only rarely so that the exchanger core remains darker. This contrast between core and walls will depend on the probability exiting to background or floor before meeting either wall, a probability which is driven by the material properties (e.g. the higher the conductivity, the lower this probability) and the foam density (the higher the density, the lower this probability). We also note that the floor having a high specular part, paths that hit it as first intersection from the camera can bounce back to the structure and then retrieve its reflected image, albeit attenuated because the floor emissivity is high so that its own imposed temperature prevails: the floor here acts as an attenuated infrared mirror.

As mentioned in Section 4, storing information about the paths that have been realized in this first rendering can be reused to hugely accelerate rendering when other conditions apply (see the supplementary video). We illustrate this point with the right panel of Fig. 1, where this information is used to render transient states at $t_{obs} = 30\text{s}$ and $t_{obs} = 150\text{s}$. Contrasting with stationary state, paths can now meet the initial condition back in time (in addition to ending at walls, floor or background), so that some of the paths that would have ended at the wall at stationary state now meet the initial condition before reaching the wall. As a consequence, they retrieve $\theta_I = 700\text{K}$ instead of the higher wall temperature, resulting overall into a lower temperature. Similarly, some of the paths exiting to background from within the core at stationary state (retrieving then 680K) now meet the initial condition before exiting, and retrieve 700K , resulting overall into a higher temperature. Of course, the probability of meeting the initial condition before intersecting a

wall / floor or exiting to background vanishes as time grows and the full history to the stationary state can be rebuilt interactively (see the supplementary video) with no need to sample again the full path space for each snapshot.

5.2 Scalability

In Figure 4 we reproduce an image that was displayed in Villefranque et al. [2022] to illustrate the ability to synthesize urban infrared images. It includes heat transfer within the buildings at each scale, as well as transfer between buildings (via the ground or via radiative exchanges), a feature typically required for the modeling of urban heat islands. In our scene, the environment is cold and each room within each building is heated with an air conditioning system. As a typical output, this image highlights the thermal bridges associated to concrete slabs and balconies, as well as a strong heat loss due to an insufficient thermal insulation of the roof.

Figure 4 shows that our objective of a full geometric scalability is reached: computation times are independent of the number of building in the scene. With the choices we made in Section 4, data-access acceleration is exactly the same as for today's standard rendering and we can indeed expect that computation times are insensitive to the number of geometric primitives. But in this scene, interactions are even more complicated, as the buildings are interacting with each other. Despite this, the computation times are the same when imaging 1, 9, 81, or 961 buildings. Under the hood, it means that the average computation time required to construct the paths is nearly the same if the building is surrounded by a known environment (the path stops when it exits the building), or if path can reach the neighbor buildings and continue their way inside these buildings. Note that inside a building, a path can have a wide range of interactions, e.g., cross the walls by conduction, interact with internal air by convection inside the rooms, or jump from one wall location to another by radiation. We believe that the stable simulation times are due to the fact that cities are dominantly extended in two dimensions. Making the scene more complex does not change significantly the average time a heat transfer path travels through the housing before it exits by convection toward the external surrounding air or by radiation toward the upper atmosphere, both here assumed at known temperatures. Computation times would on the contrary increase when the geometry becomes more complex if this increases the length that paths must travel before exit, i.e. before encountering a known temperature or an initial condition (e.g. an extended three-dimensional structure at stationary regime with known temperatures only at the outside). So computation times may change when changing the geometry because of heat transfer features changing the structure of the paths (as it is the case for multiple reflection and multiple scattering in standard rendering), but not because more triangles were used to accurately describe the scene.

5.3 Application: a thermal insulation study

We now use a case study of thermal insulation to illustrate how the Monte Carlo resolution of the coupled heat-transfer model may be used. We provide a deeper description in supplementary material (Sec. 1).




Fig. 4. Computation times insensitive to geometric complexity. Images of buildings in a cold environment with air-conditioning in each room. Although buildings thermally interact with each other via the ground or because of infrared exchanges between external walls, computation time is the same for 1, 9, 81 and 961 buildings (small differences are only due to the number of building pixels).

Analysis of an infrared image. How do we read an infrared image to thermally understand a building and optimize its design (esp. to minimize energy loss)? We take as an example a simplified building inspired from a typical farm from central France made of a room for living topped by an empty hayloft, and we question the difference of its thermal behavior when insulation is put inside versus outside. A wall-mounted heater is turned on in the room and we take, at the stationary state, a nightly infrared picture of the whole building from outside (Fig. 5, upper panel). At first sight, we obviously diagnose the thermal bridge at the ground level in both cases (which calls for primary action to prevent this). To go further, interpreting the shallow temperature difference in external surfaces temperature at hayloft and room levels is less intuitive. To this end, we can take a closer look at the paths themselves, starting here from a probe point located on the inside of the room (Fig. 5, lower panels). Note that the 3D paths are projected on the plane.

To illustrate how reading paths can help, we first focus on three of them, each one starting in convective mode from the probe point (gray paths starting from the gray sphere in Fig. 5, middle panel). With internal insulation (left), we see one of them ending at the outside soil temperature, a second one ending at the heater, and the third one ending at the external air temperature near the roof. Changing the configuration to external insulation (right), these three paths start the same but undergo very different histories: one reaches external air by conduction through the gable, and the two other ones end at the heater by convection. Most importantly, we visually note that paths now chain conduction inside the walls more often. Confirming these "statistics with the eyes" with 60 paths (Fig. 5, lower panel), we understand that a larger fraction of paths can reach the hayloft, thereby transporting heat upside, such heat being then lost for the living room. Reading paths then shows that the inside walls are forefront: when unconstrained under external insulation, they act as thermal bridges towards the empty hayloft and homogenize temperature over the whole building.

Varying conditions. As a consequence of the computational cost, storing as much information as we can about the paths statistics can be of a huge interest since they are unaffected by time or boundaries conditions (such as different heating commands, alternative weather conditions, or even for integrating day/night alternation over days, ...). Computing the propagator only once for each insulation choice and camera settings is information enough to explore sensitivity




Fig. 5. (top) Infrared renderings of a farm being insulated by (left) the interior or (right) the exterior. The boundary conditions are the heater at 300K inside the lower room, the ground at 286K at 3m depth, the surrounding air temperature at 280K and the incident radiance temperature at 280K for downward directions (coming from the upper atmosphere assuming a foggy environment i.e. equal to the air temperature). (below) Example paths: starting with convection from a probe inside the lower room, 3 paths in first picture, 60 in second picture. For each path, radiative subpaths are in red, conductive subpaths in blue and convective subpaths in gray. Conductive subpaths are successions of straight lines between δ -sphere centers (see Sec. 4.2). As internal air is perfectly mixed, there is no explicit heat migration within the fluid and convective subpaths are only sketched as straight lines from the surface location where convection is started to the surface location where it ends.

to boundary conditions. In supplementary material (Sec. 1), we show a typical engineering study which compares the responses of the internally/externally-insulated buildings when changing the




Fig. 6. Temporal evolution (in s) of the lower room temperature (in K) of a farm being insulated by inside (blue) or outside (orange). The heater being set to 290K (circle) or 300K (triangle), the temperature is similar for both insulations at stationary states. However, it increases much faster when the insulation is inside, resulting in a better thermal comfort. One single Monte Carlo simulation is performed at the steady state with 10000 samples. The stored information then allows to very rapidly replay the temperature estimation (see Fig. 1), now in transient regime at any time.




Fig. 7. Validation of the propagator implementation (triangles) against a Monte Carlo simulation (circles) on the computation of the lower-room probe temperature and its uncertainty (as in Fig. 6), which are similar in both cases.

altitude and the weather condition from a foggy atmosphere to a clear sky. After one Monte Carlo run has been made, predicting temperature under either of these conditions is instantaneous. Such predictions can be made using only four propagator values (four scalars to be multiplied by their associated boundary temperature values). They represent the four imposed temperatures in our scene: the yearly average temperature at 3 meters deep into the ground (related to the altitude), the heater temperature, the outside air temperature (related to the altitude) and the radiance temperature in directions coming from the upper atmosphere (related to weather condition).

Transient states of thermal comfort. In the same spirit as for the heat exchanger, the propagator can also be used to replay the whole dynamics of heating the room. As an illustration, we computed the two propagators (internal / external insulations) for the temperature of the room at stationary state. From these, we can estimate this temperature any sooner in time (Fig. 6). Here again, with only one single Monte Carlo run for each insulation mode, we capture the information required for any transient dynamic. In our example, with the same heating temperature, the external insulation leads to a lower stationary temperature of about 2 K, but more strikingly,




Fig. 8. Typical of a Monte Carlo resolution, the computation time (blue) of our coupled method is linear wrt the number of samples N and the normalized variance (i.e. $\bar{var} = var / (mean - 280)$ in red) wrt $1/\sqrt{N}$. Considering a farm insulated by the interior, two probes are set (left) inside the lower room (triangle) and at the external wall (circle). When the probe is outside the farm, the relative variance is far bigger because the source of heat is hard to reach (due to the insulation) and the computation time is shorter because most paths stop directly in the known radiative environment.

the time needed to reach it can be up to one month, while the internal insulation brings thermal comfort within one hour after the heating is turned on. The retarding effect of exterior insulation depends of course on the properties of the insulation materials and the walls and can be adjusted, for instance to phase the retarded effect of solar sources with the house occupation times. The walls of the considered farm are made of very thick and dense rocks that explain the extremely long duration of transient regimes, but modern buildings have walls that are thermally less capacitive.

5.4 Computation time, variance and memory

We have shown above the immediate applicability of our simple physical model of heat-transfer. Here, we give an overview of the performance characteristics of our method, typical of Monte Carlo approaches. In Fig. 8, the computation time and variance are provided as functions of the number of samples, when computing the temperature of the lower room of the farm. As mentioned in Section 5.2, the computation time depends on the underlying physical phenomena, especially the ability of the paths to quickly reach the source. In our case, conductive paths are longer to compute than radiative and convective ones, because of the number of walk steps required to reach the boundary. Also, adverse probabilities at interfaces (in Eq. 23) can trap trajectories within a geometry and negatively impact the computation time as described in the next Section 5.5.

As for classical rendering, point sources are difficult to reach, resulting in a higher variance. Indeed, we compare in Figure 8 the variance estimates when the probe is put inside vs outside the farm. The variance for the outside probe is larger because less than 2.5% of the paths actually reach the heater due to the wall insulation which prevents access to the hot source (more details in the supplementary material (Sec. 1)).

As the construction of the paths is time-consuming, their replay with different boundary or initial conditions as well as observation time is of high interest. This storage is linear wrt the image resolution and the number of samples per pixel to be generic: for spatio-temporal replays and error bar recomputation. A path-replay




Fig. 9. Path entrapment due to high conductivity ratios. Sampling heat transfer paths in electronics is a well representative example of how computation times may be dominated by entrapment. Electric connections conduct heat efficiently and tend to be quasi-isothermal. In path sampling terms, this is translated by very long paths inside the connectors because each time the path reaches the boundary of a connector the probability is high that it is reinjected inside the connector. Here the connectors are spheres with only a very small contact area with other solids of similar conductivity levels. In the displayed example, one single conductive path (in blue, from entry point in red to exit point in white) very longly visits the sphere before it finds an exit. Mesh courtesy of DeepConcept.

to change any uniform boundary condition in a scene only requires to store the source identifier for each sample. When these conditions are non uniform, the spatio-temporal impact of the path on the source is needed. To replay the simulation for a previous observation time, we need to additionally store the lifetime of each path if we assume a homogeneous initial temperature. Our prototype currently stores 116 bytes per path, reusing a *Stardis* existing structure which was originally neither designed to replay images nor observation times. This implementation, which could be optimized, is already practical. However, considering heterogeneous initial temperatures would drastically change the storage cost because time passage between medium would have to be saved.

5.5 A limit case with path-entrapment

We now detail a practical limitation we have identified, namely a case with path-entrapment. In the example presented in Fig. 9, a highly conductive solid which is embedded inside another less conductive solid, when a conductive path reaches the highly conductive solid, it tends to stay inside it. When a path reaches the interface, the probability that it is re-injected inside the very conductive part is high. So it revisits this part over and over until it finally exits. But at this exit, it is injected in the other solid, still in the close vicinity of the interface and has therefore a strong probability to be trapped again, etc, see Figure 9. Of course, this is indeed physically meaningful: a highly conductive solid is quasi-isothermal and this is here translated by the fact that wherever the path starts, it visits the whole solid several times and loses the information about its initial position. The expectation of the Monte Carlo weight is the same whatever the starting location, i.e. the temperature is the same. Although trivial, this feature leads to computation times that can be disastrous. Of course we can think of changing the model, stipulating the solid is isothermal from the start, as we did for perfectly mixed fluid cells, but this is not an easy task and we should rather retain that, at present stage, even such a very simple limit is enough to damage our proposal.

6 PERSPECTIVES

We have shown how heat transfer rendering can be addressed by using one single path statistics to make the infrared image of an object with no need to compute its temperature. For this, we have developed the theoretical framework to couple the three modes of heat transfer, starting from a model suitable for a large number of scenes, still with some constraints. The latter are our starting points for drawing perspectives:

Constraints:

- (1) Solids are opaque and fluid are transparent
- (2) Fluid cells are perfectly mixed
- (3) Properties are uniform and constant in each sub-volume
- (4) All parameters fields are prescribed
- (5) Linearization assumptions have been made (e.g. radiative exchange coefficient h_R assumed independent of temperature)

What should we do to extend the application range to more advanced heat transfer or to other physics? With our framework, a key observation is that the problem is expressed as a set of kinetic equations. Solutions of such PDEs are left with two kinds of parameters: sources and (kinetic) coefficients, as detailed in Sec. 6.1. Hence, the coupling can pass either by a *source* or by a *kinetic coefficient* (see Figure 10).

By posing the constraints of Sec. 3.4, the coupling can pass only by the sources. Moreover, in our radiative transfer model, the sources were linear functions of solid temperature and, for the conduction and convection model, the convective, conductive and radiative fluxes appeared linearly in the boundary conditions. Then, we have essentially designed our algorithms based on the well-known double randomization technique: when a Monte Carlo algorithm exists for one phenomenon, and when the corresponding Monte Carlo weight is a linear function of the solution of a Monte Carlo algorithm for another phenomenon, then we can branch the two algorithms to form a new Monte Carlo algorithm for the coupled problem, with paths switching from one phenomenon to the other.

Hence, our question becomes more specifically: what should we do to extend the application range, still keeping the benefits of double randomization, whether the coupling is via a source or via a kinetic coefficient? For coupling via more complex sources, the answer is straightforward: the proposal of Sec. 4 is a self-consistent framework that leads to the chaining of sub-paths and paves the way to easily incorporating additional phenomena to our heat transfer problem or addressing fully distinct physics (Sec. 6.2). By contrast, coupling by kinetic coefficient will call for additional algorithmic design, starting from existing preliminary work (Sec. 6.3).

6.1 Set of coupled kinetic equations

With the simplified model exposed in Sec. 3.4, an essential point is that, in the present paper, we illustrated the *linear coupling of linear kinetics models*.

In a kinetic model, the addressed quantity η is transported at a given velocity \vec{c} (by a spatio-temporal transport operator $\mathcal{T}[\eta] = \frac{\partial \eta}{\partial t} + \vec{c} \cdot \nabla \eta$, typically like in transient rendering) and reacts with an external material that pulls it back to a known kinetic source η^* with a characteristic time $1/\kappa$, where κ is named here the kinetic




Fig. 10. Illustrations of the different cases considered in Sec. 6.

In Sect. 6.1 The problem is expressed as a set of coupled kinetic equations, and we define η^* as *the source* and κ as *the kinetic coefficient*. Coupling can pass by the sources or by the kinetic coefficient.

(a) Sect. 6.2 Coupling via the sources. A path is sampled as if estimating η alone (e.g. a radiative path, in red). At the end of the path the Monte Carlo weight should be η^* (e.g. the temperature at the absorption location). But η^* is only known as the solution of another physics (e.g. conduction). A new path (in blue) is then sampled as if estimating this other physics. At the end of this second path the Monte Carlo weight should be the local value of the source for this other physics. Again this source may be only known as the solution of another physics, etc. This is the chaining of sub-paths proposed in Sec. 4.

In Sect. 6.3 Situation 1: uncoupled prescribed heterogeneous κ -fields are considered, e.g. volume radiation in a heterogeneous cloud, to introduce null-collisions techniques. Here, locations of null-collision are represented by empty dots along the path. In a standard null-collision algorithm, collision locations are exponentially sampled using an upper bound $\hat{\kappa}$ of the kinetic coefficient and the acceptance probability is $\kappa/\hat{\kappa}$.

(b) Sect. 6.3 Situation 2: coupling with an independent physics via the kinetic coefficient. When κ is only known as the solution of another physics, at each collision along the main path (red dots), the collision type cannot be decided before launching a new path (in black), as if estimating κ itself. This new path produces a Monte Carlo weight w_κ . The main path is then continued with this gathered information: the collision is accepted with probability $w_\kappa/\hat{\kappa}$. As w_κ is random, this acceptance probability is random.

(c) Sect. 6.3 Situation 3: nonlinear coupling via the kinetic coefficient. The path launched for κ may itself deal with unknown kinetic coefficients and therefore require the further launching of new paths before it is continued, etc, leading to branching paths (three recursive branching levels are displayed, in red, black and brown). When κ depends itself upon the first level quantity (e.g. h_R depends on the temperature), branching paths open the way to tackle non linear kinetics. In the particular case of one single nonlinear kinetic equation (e.g. the original Boltzmann equation), all paths would address the same quantity (and would be represented with the same color here).

coefficient:

$$\mathcal{T}[\eta] = -\kappa\eta + \kappa\eta^*. \quad (36)$$

Diffusive or advecto-diffusive models are not of the kinetic type, but using the walk-on- δ -sphere approximation, each step is modeled using a kinetic equation (see Eq. 32 where $\mathcal{T}[\eta] \equiv \frac{\partial \theta_S}{\partial t}$ is only temporal, as for transport in a uniform field, spatialization being recovered because $\eta^* \equiv \theta_S^*$ is an integral over the sphere). Note that when using a convective exchange coefficient to simplify the convective model, this model also gets a kinetic form. So, in practice, we were dealing with the coupling of kinetic equations.

One major feature of kinetic equations is that their solutions are always accessible as sums of sources (in the sense of Green theory, i.e. comprising boundary fluxes and initial conditions) multiplied by their propagators, each propagator involving negative exponentials of the integral of the kinetic coefficient over time or space. Formally, the solution of each kinetic equation, considered independently, writes

$$\eta(t) = \sum \mathcal{L} \left\{ \exp \left(- \int \kappa(t) dt \right) \right\} \eta^*. \quad (37)$$

The coupling passes via a source when η^* depends on the solution of another kinetic equation (Sec. 6.2). The coupling passes via a kinetic coefficient when $\kappa(t)$ depends on the solution of another kinetic equation (Sec. 6.3). In some cases, both can be present.

6.2 Coupling via the sources

Linearly coupling via the sources is already sufficient to alleviate some of the constraints listed above. We explain in supplementary material (Sec. 4) how the algorithm of Section 4 can actually be

extended to recover most of the full initial heat transfer model of Section 3. For instance, had we considered volume radiation, we would still have been dealing with kinetic equations because the Radiative Transfer Equation is kinetic by nature. This alleviates constraint (1), provided that we still leave aside the nonlinearity of radiation (otherwise, see below). True advection add neither specific difficulty as far as coupling is concerned if we still assume that the velocity field is prescribed and independent of temperature (otherwise, see below). This alleviates constraint (2).

Coupling via the sources can also be used to incorporate additional physics: as sketched in Fig. 10(a), it suffices that one of the sources be a linear function of this additional physics to justify that when the first path ends at this source, it is chained with a sub-path for the new physics. Let us take the example of a three-dimensional electric flow inside a solid, and that heats the solid because of the local resistance to the electric flow (Joule effect). Assuming that we have a Monte Carlo algorithm for this flow and its associated power density, can we couple it to a Monte Carlo heat transfer algorithm such as those we illustrated in Section 4? The answer is yes because this additional coupling with electrical dissipation is via a volume source in the physics of solid heat-conduction: along the walk-on-sphere algorithm, when the volume source is needed, an electric sub-path would be initiated to fetch the missing information. All such enrichments of the phenomenology are straightforward provided that the coupling is linear and passes via the sources and from this point of view, the proposal of Sec. 4, although depicted through a practical example, may be considered as a closed formal framework.

6.3 Coupling via the kinetic coefficient

As an introduction to the additional algorithmic work required for tackling coupling via the kinetic coefficient, let us consider first Eq. 37 with regards to the third constraint (3).

If we consider subvolume properties as uniform and constant by parts, then κ is uniform along a subpath within a subvolume. In this case, the reexpression of Eq. 37 to eliminate the time integral is immediate. This is how we constructed Eq. 29 and 32. If now the field of κ is continuously heterogeneous and time dependent, the Monte Carlo algorithm solving η must account for the fact that, in Equation 37, κ is varying along the path when addressing the integral $\int \kappa(t)dt$ over the propagation path.

We can here face three situations:

- (1) the heterogeneous field and its time dependence are prescribed (uncoupled situation)
- (2) the field has to be estimated from a model which does not depend upon the heat transfer (constraint (4), coupling with an independent physics),
- (3) the field has to be estimated from a model in which the solution itself depends upon the heat transfer (constraint (5), nonlinear coupling)

In order to tackle heterogeneous time-dependent kinetic coefficient, practical solutions are already available based on the concept of *null-collisions*. This concept is indeed originally introduced as a way to stick rigorously to the Monte Carlo framework when progressing through a heterogeneous field. To take an example in radiative transfer, let us consider a volume with a heterogeneous distribution of colliders (κ field) so that the exponential of the integral can not be known analytically. Let us consider null colliders in addition to the true colliders. When a photon encounters a null collider, it continues its path as if no collision had occurred so that the transfer of radiation is unchanged.

6.3.1 Situation 1: uncoupled prescribed heterogeneous κ -fields. In standard practice, if the field is prescribed, the amount of null colliders can be adjusted locally so that, in total (true colliders plus null colliders), the new apparent coefficient field becomes homogeneous over the subvolume, and the difficult question of integrating along the path within the exponential just vanishes. When the field of collisions is known, Monte Carlo practice is then: set an upper bound value from the known field, and at each collision location, deciding whether the collision is true or null is made using only the knowledge of the local value of the true extinction coefficient versus the upper bound value. That would allow to alleviate constraint (3). Null collisions are sketched in Fig. 10(a), inspired of multiple scattering in clouds for instance, by plotting collisions as dots along the path, some of them leading to the path being continued as if no collision had occurred.

6.3.2 Situation 2: coupling with an independent physics. In this case, κ depends on the solution of another physics, which can be solved on its own.

Even if \mathcal{L} is linear in Equation 37, the dependence is inside the exponential. For discussion, let us assume that in Eq. 36 transport is only in time ($\nabla\eta$ is null and $\mathcal{T}[\eta] = \frac{\partial\eta}{\partial t}$). But now the kinetic

coefficient is itself modeled as the solution of a second kinetic equation. For this second equation, all we posit is that a Monte Carlo algorithm is available that estimates $\kappa(t)$ as the expectation of a weight $W_{\kappa,t}$:

$$\begin{aligned} \frac{\partial\eta}{\partial t} &= -\kappa\eta + \kappa\eta^* \\ \kappa(t) &= \mathbb{E}(W_{\kappa,t}) \end{aligned} \quad (38)$$

Using the same derivation as for Equations 29 and 34, the formal solution of this coupled model would have the form of Equation 37, the expectation of W_{κ} appearing in place of κ inside the temporal integral and inside the exponential.

Yet, as this integral is inside the exponential, and the exponential is itself inside an expectation over the random sampling of the path, the problem does not write anymore as the expectation of an expectation: double randomization does not apply directly.

Basically, what plays the role of the true extinction coefficient κ becomes unknown: it is the solution of a second kinetic equation to be solved together with the first one. This question has been investigated in [Galtier et al. 2016] and [Terrée et al. 2022] and although the presentations were very much domain-specific, a common formalism can be extracted from these works.

They both start with a standard null-collision algorithm. Let us assume that there exists a known upper bound \hat{w}_{κ} to $W_{\kappa,t}$, meaning that for each realization $w_{\kappa,t}$ of $W_{\kappa,t}$ at any time, $w_{\kappa,t} < \hat{w}_{\kappa}$. This upper bound is therefore also an upper bound for $\kappa(t)$. We can use it to define the null-collision amount:

$$\frac{\partial\eta}{\partial t} = -\hat{w}_{\kappa}\eta + \hat{w}_{\kappa} \left(\frac{\kappa}{\hat{w}_{\kappa}}\eta^* + \left(1 - \frac{\kappa}{\hat{w}_{\kappa}}\right)\eta \right) \quad (39)$$

At this stage, all we have is more extinction (the kinetic coefficient is replaced with its upper bound) and the surplus in extinction is compensated by the fact that the new source (inside the parenthesis) is the sum of the old source and the sought value of η (a recursion), weighted by the true-collision probability $\frac{\kappa}{\hat{w}_{\kappa}}$ and its complementary null-collision probability.

After this first-step of introducing null-collisions, [Galtier et al. 2016] and [Terrée et al. 2022] propose that κ may be replaced by its expectation inside the new source and that true-collision and null-collision probabilities can be made random themselves. An additional sampling is then required to decide these probabilities. Once the probabilities have been sampled, they are used to sample the Bernoulli random variable:

$$\begin{aligned} \frac{\kappa}{\hat{w}_{\kappa}}\eta^*(t) + \left(1 - \frac{\kappa}{\hat{w}_{\kappa}}\right)\eta(t) &= \frac{\mathbb{E}(W_{\kappa,t})}{\hat{w}_{\kappa}}\eta^*(t) + \left(1 - \frac{\mathbb{E}(W_{\kappa,t})}{\hat{w}_{\kappa}}\right)\eta(t) \\ &= \mathbb{E} \left[\frac{W_{\kappa,t}}{\hat{w}_{\kappa}}\eta^*(t) + \left(1 - \frac{W_{\kappa,t}}{\hat{w}_{\kappa}}\right)\eta(t) \right] \\ &= \mathbb{E} \left[\mathcal{H}_{true}(\mathcal{P})\eta^*(t) + \mathcal{H}_{null}(\mathcal{P})\eta(t) \right] \end{aligned} \quad (40)$$

with $\mathcal{P} = \frac{W_{\kappa,t}}{\hat{w}_{\kappa}}$ and $\mathcal{H}_{null}(\mathcal{P}) = 1 - \mathcal{H}_{true}(\mathcal{P})$. The test $\mathcal{H}_{true}(\mathcal{P})$ is a Bernoulli random variable of probability \mathcal{P} and the fact that this probability is itself a random variable simply means that the Bernoulli test is made after \mathcal{P} is realized. In summary, true-collision and null-collision probabilities are random: before deciding the collision type, a path is sampled corresponding to the physics of κ , as if estimating κ alone, the resulting Monte Carlo weight is used

in place of κ to compute the collision-type probabilities and this random probability is used to continue the algorithm exactly the same way it would do in a standard null-collision algorithm.

Null collisions offer us a practical way to set up a single path statistics for physics that are coupled via a kinetic coefficient: the path is sampled through a field where the local values of the kinetic coefficient are unknown, and, at each collision, a κ -subpath is launched to fetch the missing information. This alleviates constraint (4). Figure 10(b) illustrate this first level of path-branching with a main path (for η) and secondary paths (for κ) launched at each collision location. In [Galtier et al. 2016], κ was a gaseous absorption cross-section and κ -subpaths were randomly visiting molecular spectroscopic databases to gather information about millions of absorption lines.

As a limit to this approach, we need to assume that there exists an upper bound to the Monte Carlo weight. This is not always possible. For instance, when there is a volume source inside a solid, the Monte Carlo weight will not only carry the temperature at the end of the path but will also include a contribution of the volume sources accumulated along all the visited location. Since here the visiting time does not have an upper bound, nor does the Monte Carlo weight.

6.3.3 Situation 3: nonlinear kinetics. In [Terrée et al. 2022], the addressed question is fluid mechanics in rarefied gases, η is the distribution function (the equivalence of radiance for molecules instead of photons), κ is the molecular collision cross-section and the approach is essentially the same.

But, from the present point of view, this contribution opens a new perspective: in short, molecules collide with molecules and κ is η itself. This has two implications:

- the kinetic model for η is nonlinear (the original Boltzmann equation);
- when a κ -subpath is launched because of missing information about the local value of κ , as κ is η , the new path is practically an η -path that may itself require the launching of a κ -subpath, etc.

A nonlinear kinetic model is therefore translated into a branching-path statistics. Such statistics remain widely unknown [Dimov and Gurov 2000; Rasulov et al. 2019; Rioux-Lavoie et al. 2022], but there are already cases with effective application potentials where they can be deployed successfully.

In the present context, such branching paths techniques open avenues to alleviate constraint (5), so as to consider the effect of wider temperature ranges upon radiative transfer, i.e. to take into account situations where the approximation considering h_R as independent of the temperature does not hold. Such branching paths are sketched in Fig. 10: along the main path (for η), sub-paths are launched at each collision location (for κ), and these sub-paths themselves initiate sub-paths for unknown kinetic coefficients, etc. Three recursive branching levels are represented. The particular case of one single nonlinear kinetic equation is when each launched sub-path is of the same type as the primary path.

7 FUTURE WORKS

In the introduction and related works, we highlighted that there were, up to our knowledge, only few reported attempts to address the coupling of several physics in a single path space Monte Carlo algorithm. Figure 4 illustrates very well the expected benefits of our coupling proposal, and specifically the better scalability of Monte Carlo methods over deterministic approaches to solve complex problems.

Beyond the range of cases that can be readily handled on the basis of the work presented above, we have identified two limitations that would deserve future work: the risk of path-entrapment and the existence of an upper bound when using null-collision. In addition, we believe that several improvements should also be tackled as fundamental research topics.

Brownian motion is among the leading requirements. Pragmatically, we used a walk-on- δ -sphere algorithm to model diffusion, in association with heterogeneous advection and heterogeneous sources. By construction, walk-on- δ -sphere loses general guaranties about convergence, while walk-on-sphere ensures a logarithmic convergence under homogeneous hypothesis (thanks to the tangent sphere approach). Keeping such convergence guaranties in heterogeneous situations is a wide research question in the walk-on-sphere literature, but also in an active literature about walks on other shapes than spheres [Deaconu and Lejay 2006].

A very similar comment holds for *branching path statistics*. If the reported practical successes are convincing, there is still a lack of a general framework, and we believe that making a connection with the active literature opened by the work of Dimov [2008] should help for nonlinear coupling and more specifically nonlinear kinetics.

In our proposal, we did not consider *importance sampling* nor *next event estimation* but we pointed-out convergence difficulties associated with images of a well isolated house, where paths starting from an outside camera and reaching the heating system were scarce. If similar questions are already well studied in the rendering community, they remain open questions for thick diffusive media, typically for radiative transfer in systems with huge optical thicknesses. The difficulty has been well identified in the nuclear shielding literature since the early ages of Monte Carlo simulations [Hammersley 2013], but it remains open. Several links should be explored in this still-active literature.

Sensitivity modeling is another promising topic evoked in Section 4. We computed propagators for replay, and propagators are sensitivities to the sources. We also insisted on the information carried by the paths themselves: even simply plotting them is already useful as far as analysis is concerned. Obviously, this is true for other parameters than sources (conductivities, capacities, geometry, etc). Following the spirit of differential light transport, we believe that this path information could be translated into the computation of sensitivities to these other parameters. This would be highly useful in practical engineering terms, e.g., for inverse problem solving.

Along this paper, all concepts were introduced and discussed from the heat-transfer perspective. This is how the question was introduced in Villefranque et al. [2022], thinking of the energy transition and climate-change communities. But even in these communities, the research questions for which computer graphics Monte Carlo

simulations could be useful widely extend heat transfer. Fluid mechanics, electrical engineering, biology modeling are only few of the phenomena addressed by a broad scientific community that has the same flexibility and scalability requirements as those addressed by computer graphics. We hope that the coupling of Monte Carlo path sampling algorithms will pave the way for new fruitful connections.

ACKNOWLEDGMENTS

This work was supported by the ANR under grand No.: ANR-18-CE46-0012 (<https://anr.fr/Project-ANR-18-CE46-0012>).

REFERENCES

Next Limit 2016. *Maxwell Multilight*. Next Limit. <https://maxwellrender.com>

ANSYS, Inc. 2017. *ANSYS Fluent*. ANSYS, Inc. <https://www.ansys.com/products/fluids/ansys-fluent>

Meso-Star 2018. *Stardis*. Meso-Star. <https://www.meso-star.com/projects/stardis/stardis.html>

Comsol 2020. *Comsol Multiphysics®*. Comsol. <https://www.comsol.com/>

EDF 2021. *SYRTHES*. EDF. <https://www.edf.fr/en/the-edf-group/inventing-the-future-of-energy/r-d-global-expertise/our-offers/simulation-software/syrthes>

Marco Ament, Christoph Bergmann, and Daniel Weiskopf. 2014. Refractive Radiative Transfer Equation. *ACM Trans. Graph.* 33, 2, Article 17 (apr 2014), 22 pages. <https://doi.org/10.1145/2557605>

Benedikt Bitterli and Eugene d'Eon. 2022. A Position-Free Path Integral for Homogeneous Slabs and Multiple Scattering on Smith Microfacets. *Computer Graphics Forum* 41, 4 (2022), 93–104. <https://doi.org/10.1111/cgf.14589> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14589>

Subrahmanyan Chandrasekar. 1960. *Radiative Transfer*. Dover Publications.

Xavier Chermain, Basile Sauvage, Jean-Michel Dischler, and Carsten Dachsbacher. 2020. Procedural Physically-based BRDF for Real-Time Rendering of Glints. *Comput. Graph. Forum (Proc. Pacific Graphics)* 39, 7 (2020), 243–253.

W. A. Coleman. 1968. Mathematical Verification of a Certain Monte Carlo Sampling Technique and Applications of the Technique to Radiation Transport Problems. *Nuclear Science and Engineering* 32, 1 (1968), 76–81. <https://doi.org/10.13182/NSE68-1> arXiv:<https://doi.org/10.13182/NSE68-1>

Jérémi Dauchet, Jean-Jacques Beziau, Stéphane Blanco, Cyril Caliot, Julien Charon, Christophe Coustet, Mouna El Hafi, Vincent Eymet, Olivier Farges, Vincent Forest, et al. 2018. Addressing nonlinearities in monte carlo. *Scientific reports* 8, 1 (2018), 1–11.

Jérémi Dauchet, Stéphane Blanco, Jean-François Cornet, Mouna El Hafi, Vincent Eymet, and Richard Fournier. 2013. The practice of recent radiative transfer Monte Carlo advances and its contribution to the field of microorganisms cultivation in photobioreactors. *Journal of Quantitative Spectroscopy and Radiative Transfer* 128 (2013), 52–59.

Madalina Deaconu, Samuel Herrmann, and Sylvain Maire. 2017. The walk on moving spheres: a new tool for simulating Brownian motion's exit time from a domain. *Mathematics and Computers in Simulation* 135 (2017), 28–38.

Madalina Deaconu and Antoine Lejay. 2006. A random walk on rectangles algorithm. *Methodology and Computing in Applied Probability* 8, 1 (2006), 135–151. <https://doi.org/10.1007/s11009-006-7292-3>

Madalina Deaconu and Antoine Lejay. 2007. Simulation of exit times and positions for Brownian motions and diffusions. In *PAMM: Proceedings in Applied Mathematics and Mechanics*, Vol. 7. Wiley Online Library, 1081401–1081402.

Ivan Dimov and Todor Gurov. 2000. Monte Carlo algorithm for solving integral equations with polynomial non-linearity. Parallel implementation. *Pliska Studia Mathematica Bulgarica* 13, 1 (2000), 117p–132p.

Ivan T Dimov. 2008. *Monte Carlo methods for applied scientists*. World Scientific.

Richard Phillips Feynman and Laurie M Brown. 2005. *Feynman's thesis: a new approach to quantum theory*. World Scientific.

M. Galtier, S. Blanco, C. Caliot, C. Coustet, J. Dauchet, M. El Hafi, V. Eymet, R. Fournier, J. Gautrais, A. Khuong, B. Piaud, and G. Terrée. 2013. Integral formulation of null-collision Monte Carlo algorithms. *Journal of Quantitative Spectroscopy and Radiative Transfer* 125 (2013), 57–68. <https://doi.org/10.1016/j.jqsrt.2013.04.001>

Mathieu Galtier, Stéphane Blanco, Jérémi Dauchet, Mouna El Hafi, Vincent Eymet, Richard Fournier, Maxime Roger, Christophe Spiesser, and Guillaume Terrée. 2016. Radiative transfer and spectroscopic databases: A line-sampling Monte Carlo approach. *Journal of Quantitative Spectroscopy and Radiative Transfer* 172 (2016), 83–97.

Victor Gattepaille, Jérémi Dauchet, Fabrice Gros, Matthieu Roudet, Caroline Supplis, and Jean-François Cornet. 2018. Integral formulations of multi-scale models for the optimization of solar photo-catalytic processes. In *International Congress/3rd days of GdR Solar Fuels: Artificial Photosynthesis and Solar Fuels*.

Ibón Guillén, Julio Marco, Diego Gutierrez, Wenzel Jakob, and Adrien Jarabo. 2020. A General Framework for Pearlescent Materials. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 39, 6 (Nov. 2020). <https://doi.org/10.1145/3414685.3417782>

Yu Guo, Miloš Hašan, and Shuang Zhao. 2018. Position-Free Monte Carlo Simulation for Arbitrary Layered BSDFs. *ACM Trans. Graph.* 37, 6, Article 279 (dec 2018), 14 pages. <https://doi.org/10.1145/3272127.3275053>

John Hammersley. 2013. *Monte carlo methods*. Springer Science & Business Media.

Eric Heitz, Johannes Hanika, Eugene d'Eon, and Carsten Dachsbacher. 2016. Multiple-scattering microfacet BSDFs with the Smith model. *ACM Trans. Graph.* 35, 4 (July 2016), 58:1–58:14. <https://doi.org/10.1145/2897824.2925943>

Nicolas Holzschuch and Romain Pacanowski. 2017. A Two-Scale Microfacet Reflectance Model Combining Reflection and Diffraction. *ACM Trans. Graph.* 36, 4, Article 66 (jul 2017), 12 pages. <https://doi.org/10.1145/3072959.3073621>

John R Howell and Kyle J Daun. 2021. The past and future of the Monte Carlo method in thermal radiation transfer. *Journal of Heat Transfer* 143, 10 (2021).

Chi-Ok Hwang and Michael Mascagni. 2003. Analysis and comparison of Green's function first-passage algorithms with "Walk on Spheres" algorithms. *Mathematics and computers in simulation* 63, 6 (2003), 605–613.

Kiyosi Itô, P Henry Jr, et al. 1996. *Diffusion processes and their sample paths: Reprint of the 1974 edition*. Springer Science & Business Media.

Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. 2022. *Mitsuba 3 renderer*. <https://mitsuba-renderer.org>.

Adrian Jarabo, Julio Marco, Adolfo Muñoz, Raul Buisan, Wojciech Jarosz, and Diego Gutierrez. 2014. A Framework for Transient Rendering. *ACM Trans. Graph.* 33, 6, Article 177 (nov 2014), 10 pages. <https://doi.org/10.1145/2661229.2661251>

Mark Kac. 1947. Random walk and the theory of Brownian motion. *The American Mathematical Monthly* 54, 7P1 (1947), 369–391.

Mark Kac et al. 1951. On some connections between probability theory and differential and integral equations. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability*. The Regents of the University of California.

James T. Kajiya. 1986. The Rendering Equation. *SIGGRAPH Comput. Graph.* 20, 4 (aug 1986), 143–150. <https://doi.org/10.1145/15886.15902>

Sabelfeld Karl. 2019. First passage Monte Carlo algorithms for solving coupled systems of diffusion-reaction equations. *Applied Mathematics Letters* 88 (2019), 141–148.

Edward A Kopylov and Kirill A Dmitriev. 2000. Light propagation visualization as a tool for 3D scene analysis in lighting design. *Computers & Graphics* 24, 1 (2000), 31–39. [https://doi.org/10.1016/S0097-8493\(99\)00135-1](https://doi.org/10.1016/S0097-8493(99)00135-1)

Peter Kutz, Ralf Habel, Yining Karl Li, and Jan Novák. 2017. Spectral and Decomposition Tracking for Rendering Heterogeneous Volumes. *ACM Trans. Graph.* 36, 4, Article 111 (jul 2017), 16 pages. <https://doi.org/10.1145/3072959.3073665>

P Lapeyre, Stéphane Blanco, Cyril Caliot, J Dauchet, Mouna El Hafi, Richard Fournier, Olivier Farges, Jacques Gautrais, and Maxime Roger. 2020. Monte-Carlo and sensitivity transport models for domain deformation. *Journal of Quantitative Spectroscopy and Radiative Transfer* 251 (2020), 107022.

Antoine Lejay and Sylvain Maire. 2010. Simulating diffusions with piecewise constant coefficients using a kinetic approximation. *Computer Methods in Applied Mechanics and Engineering* 199, 29–32 (2010), 2014–2023.

Antoine Lejay and Sylvain Maire. 2013. New Monte Carlo schemes for simulating diffusions in discontinuous media. *J. Comput. Appl. Math.* 245, 1 (2013), 97–116. <https://doi.org/10.1016/j.cam.2012.12.013>

Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. 2018. Differentiable Monte Carlo Ray Tracing through Edge Sampling. *ACM Trans. Graph.* 37, 6, Article 222 (dec 2018), 11 pages. <https://doi.org/10.1145/3272127.3275109>

Sylvain Maire and Giang Nguyen. 2016. Stochastic finite differences for elliptic diffusion equations in stratified domains. *Mathematics and Computers in Simulation* 121 (2016), 146–165.

Julio Marco, Ibón Guillén, Wojciech Jarosz, Diego Gutierrez, and Adrian Jarabo. 2019. Progressive Transient Photon Beams. *Computer Graphics Forum* 38, 6 (2019), 19–30. <https://doi.org/10.1111/cgf.13600> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13600>

Johannes Meng, Marios Papas, Ralf Habel, Carsten Dachsbacher, Steve Marschner, Markus Gross, and Wojciech Jarosz. 2015. Multi-Scale Modeling and Rendering of Granular Materials. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 34, 4 (July 2015). <https://doi.org/10/gfzndr>

Mervin E. Muller. 1956. Some Continuous Monte Carlo Methods for the Dirichlet Problem. *The Annals of Mathematical Statistics* 27, 3 (1956), 569 – 589. <https://doi.org/10.1214/aoms/1177728169>

Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically Based Rendering: From Theory to Implementation (3rd ed.)* (3rd ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. 1266 pages.

Yang Qi, Dario Seyb, Benedikt Bitterli, and Wojciech Jarosz. 2022. A bidirectional formulation for Walk on Spheres. *Computer Graphics Forum (Proceedings of EGSR)* 41, 4 (July 2022). <https://doi.org/10.1111/cgf.14586>

Matthias Raab, Daniel Seibert, and Alexander Keller. 2008. Unbiased Global Illumination with Participating Media. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*,

Alexander Keller, Stefan Heinrich, and Harald Niederreiter (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 591–605.

Joël Randrianandrasana, Patrick Callet, and Laurent Lucas. 2021. Transfer Matrix Based Layered Materials Rendering. *ACM Trans. Graph.* 40, 4, Article 177 (jul 2021), 16 pages. <https://doi.org/10.1145/3450626.3459859>

Abdujabbor Rasulov, Gulnora Raimova, and Matyokub Bakoev. 2019. Monte Carlo Solution of Dirichlet Problem for Semi-linear Equation. In *Finite Difference Methods. Theory and Applications: 7th International Conference, FDM 2018, Lozenetz, Bulgaria, June 11–16, 2018, Revised Selected Papers 7*. Springer, 443–451.

Damien Rioux-Lavoie, Ryusuke Sugimoto, Tümay Özdemiir, Naoharu H Shimada, Christopher Batty, Derek Nowrouzezahrai, and Toshiya Hachisuka. 2022. A Monte Carlo Method for Fluid Simulation. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–16.

Karl K Sabelfeld. 2019. A global random walk on spheres algorithm for transient heat equation and some extensions. *Monte Carlo Methods and Applications* 25, 1 (2019), 85–96.

Karl K Sabelfeld and Nikita Popov. 2020. Monte Carlo tracking drift-diffusion trajectories algorithm for solving narrow escape problems. *Monte Carlo Methods and Applications* 26, 3 (2020), 177–191.

Karl K Sabelfeld and Dmitrii Smirnov. 2021. A global random walk on grid algorithm for second order elliptic equations. *Monte Carlo Methods and Applications* 27, 3 (2021), 211–225.

Rohan Sawhney and Keenan Crane. 2020. Monte Carlo Geometry Processing: A Grid-Free Approach to PDE-Based Methods on Volumetric Domains. *ACM Trans. Graph.* 39, 4, Article 123 (jul 2020), 18 pages. <https://doi.org/10.1145/3386569.3392374>

Rohan Sawhney, Bailey Miller, Ioannis Gkioulekas, and Keenan Crane. 2023. Walk on Stars: A Grid-Free Monte Carlo Method for PDEs with Neumann Boundary Conditions. arXiv:2302.11815 [cs.GR]

Rohan Sawhney, Dario Seyb, Wojciech Jarosz, and Keenan Crane. 2022. Grid-free Monte Carlo for PDEs with spatially varying coefficients. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–17.

Evan Shellshear and Robin Ytterlid. 2014. Fast distance queries for triangles, lines, and points using SSE instructions. *Journal of Computer Graphics Techniques Vol 3*, 4 (2014).

Shlomi Steinberg, Pradeep Sen, and Ling-Qi Yan. 2022. Towards Practical Physical-Optics Rendering. *ACM Trans. Graph.* 41, 4, Article 132 (jul 2022), 24 pages. <https://doi.org/10.1145/3528223.3530119>

Guillaume Terrée, Mouna El Hafi, Stéphane Blanco, Richard Fournier, Jérémie Dauchet, and Jacques Gautrais. 2022. Addressing the gas kinetics Boltzmann equation with branching-path statistics. *Physical Review E* 105, 2 (2022), 025305.

Jean Marc Tregan, Jean Luc Amestoy, Megane Bati, Jean-Jacques Beziau, Stéphane Blanco, Laurent Brunel, Cyril Caliot, Julien Charon, Jean-Francois Cornet, Christophe Coustet, Louis d’Alençon, Jeremi Dauchet, Sebastien Dutour, Simon Eibner, Mouna El Hafi, Vincent Eymet, Olivier Farges, Vincent Forest, Richard Fournier, Mathieu Galtier, Victor Gattepaille, Jacques Gautrais, Zili He, Frédéric Hourdin, Loris Ibarrart, Jean-Louis Joly, Paule Lapeyre, Pascal Lavielle, Marie-Helene Lecureux, Jacques Lluc, Marc Miscovic, Nada Mourtaaday, Yaniss Nyffenegger-Péré, Lionel Pelissier, Lea Penazzi, Benjamin Piaud, Clément Rodrigues-Viguiier, Gisele Roques, Maxime Roger, Thomas Saez, Guillaume Terrée, Najda Villefranque, Thomas Vourc’h, and Daniel Yaacoub. 2023. Coupling radiative, conductive and convective heat-transfers in a single Monte Carlo algorithm: A general theoretical framework for linear situations. *PLOS ONE* 18, 4 (04 2023), 1–54. <https://doi.org/10.1371/journal.pone.0283681>

Eric Veach. 1998. *Robust Monte Carlo methods for light transport simulation*. Stanford University.

Delio Vicini, Sébastien Speierer, and Wenzel Jakob. 2021. Path Replay Backpropagation: Differentiating Light Paths using Constant Memory and Linear Time. *Transactions on Graphics (Proceedings of SIGGRAPH)* 40, 4 (Aug. 2021), 108:1–108:14. <https://doi.org/10.1145/3450626.3459804>

Najda Villefranque, Richard Fournier, Fleur Couvreur, Stéphane Blanco, Céline Cornet, Vincent Eymet, Vincent Forest, and Jean-Marc Tregan. 2019. A Path-Tracing Monte Carlo Library for 3-D Radiative Transfer in Highly Resolved Cloudy Atmospheres. *Journal of Advances in Modeling Earth Systems* 11, 8 (2019), 2449–2473. <https://doi.org/10.1029/2018MS001602>

Najda Villefranque, Frédéric Hourdin, Louis d’Alençon, Stéphane Blanco, Olivier Boucher, Cyril Caliot, Christophe Coustet, Jérémie Dauchet, Mouna El Hafi, Vincent Eymet, Olivier Farges, Vincent Forest, Richard Fournier, Jacques Gautrais, Valéry Masson, Benjamin Piaud, and Robert Schoetter. 2022. The “teapot in a city”: A paradigm shift in urban climate modeling. *Science Advances* 8, 27 (2022), eabp8934. <https://doi.org/10.1126/sciadv.abp8934> arXiv:https://www.science.org/doi/pdf/10.1126/sciadv.abp8934

Ingo Wald, Sven Woop, Carsten Benthin, Gregory S Johnson, and Manfred Ernst. 2014. Embree: a kernel framework for efficient CPU ray tracing. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–8.

Beibei Wang, Wenhua Jin, Jiahui Fan, Jian Yang, Nicolas Holzschuch, and Ling-Qi Yan. 2022. Position-Free Multiple-Bounce Computations for Smith Microfacet BSDFs.

ACM Trans. Graph. 41, 4, Article 134 (jul 2022), 14 pages. <https://doi.org/10.1145/3528223.3530112>

A. Wilkie, S. Nawaz, M. Droske, A. Weidlich, and J. Hanika. 2014. Hero Wavelength Spectral Sampling. *Computer Graphics Forum* 33, 4 (2014), 123–131. <https://doi.org/10.1111/cgf.12419> arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12419

Cheng Zhang, Lifan Wu, Changxi Zheng, Ioannis Gkioulekas, Ravi Ramamoorthi, and Shuang Zhao. 2019. A Differential Theory of Radiative Transfer. *ACM Trans. Graph.* 38, 6, Article 227 (nov 2019), 16 pages. <https://doi.org/10.1145/3355089.3356522>

Shuang Zhao, Wenzel Jakob, and Tzu-Mao Li. 2020. Physics-based differentiable rendering: from theory to implementation. In *ACM siggraph 2020 courses*. 1–30.

Tobias Zirr, Marco Ament, and Carsten Dachsbacher. 2015. Visualization of Coherent Structures of Light Transport. *Computer Graphics Forum* 34, 3 (2015), 491–500. <https://doi.org/10.1111/cgf.12661> arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12661

A BASICS OF HEAT TRANSFER

A.1 Conduction and convection models for solid and fluid volumes

We detail how Equations 8 and 12 are derived from the basic laws of thermodynamics. For continuous matter with a uniform mass density, the first principle of thermodynamics (the law of energy conservation) relates the internal energy density $u \equiv u(\vec{x}, t)$ to the energy flux density vector $\vec{j} \equiv \vec{j}(\vec{x}, t)$ as

$$\frac{\partial u}{\partial t} + \vec{\nabla} \cdot \vec{j} = 0 \quad (41)$$

with

$$\frac{\partial u}{\partial t} = \rho c \frac{\partial \theta}{\partial t} \quad (42)$$

where ρ is the mass density and c the specific heat capacity. For an opaque solid, \vec{j} reduces to the conduction of heat (no advection, no radiation) and is given by Fourier’s law:

$$\vec{j} = -k \vec{\nabla} \theta \quad (43)$$

where k is the thermal conductivity. Eq. 8 is obtained by replacing Eq. 42-43 into Eq. 41.

We now consider a cavity inside a solid. The cavity is filled with a fluid. At any location \vec{x} inside the cavity, the fluid temperature is $\theta_F \equiv \theta_F(\vec{x}, t)$. At any location \vec{y} at the surface of the solid interface, the solid temperature is $\theta_S \equiv \theta_S(\vec{y}, t)$. Newton’s law is used to replace the detailed description of the boundary layer of the fluid along the surface: $\theta_F(\vec{y}, t)$ is assumed to be the temperature of the fluid outside the boundary layer ; therefore, $\theta_F(\vec{y}, t) \neq \theta_S(\vec{y}, t)$ (the boundary layer is replaced by a temperature discontinuity) ; then, the flux density at \vec{y} is :

$$\varphi(\vec{y}, t) = h(\vec{y})(\theta_F(\vec{y}, t) - \theta_S(\vec{y}, t)). \quad (44)$$

Choosing the value of the proportionality coefficient h requires either experimental or numerical fluid mechanics (characterizing the boundary layer). We further assume that outside the boundary layer, the cavity is perfectly uniform and mixed which means that θ_F is independent of \vec{y} . We assume a constant pressure. The first principle for an open system is then the same as Eq. 41 with enthalpy instead of internal energy is

$$\frac{dH}{dt} = \int_S \varphi(\vec{y}) d\vec{y} \quad (45)$$

where \mathcal{S} is the fluid boundary surface. As for Eq. 42, the enthalpy relates to the isobaric specific heat capacity c_p :

$$\frac{dH}{dt} = \int_{\mathcal{V}} \rho c_p \frac{\partial \theta_F}{\partial t} d\vec{x} = \mathcal{V} \rho c_p \frac{d\theta_F}{dt} \quad (46)$$

where \mathcal{V} denotes the fluid volume. Eq. 12 is obtained by replacing Eq. 44 and 46 into Eq. 45.

A.2 Finite difference on the Laplace operator

We detail the derivation of Eq. 32 from Eq. 9. Each term of the Laplace operator $\Delta \theta_S := \frac{\partial^2 \theta_S}{\partial x^2} + \frac{\partial^2 \theta_S}{\partial y^2} + \frac{\partial^2 \theta_S}{\partial z^2}$ is approximated by two finite differences as in any cartesian coordinate system, e.g.

$$\begin{aligned} \frac{d\theta^2}{dx^2} &= \frac{d}{dx} \left(\frac{d\theta}{dx} \right) \simeq \frac{1}{\delta} \left(\left. \frac{d\theta}{dx} \right|_{x+\frac{\delta}{2}} - \left. \frac{d\theta}{dx} \right|_{x-\frac{\delta}{2}} \right) \\ &\simeq \frac{1}{\delta} \left(\frac{\theta(x+\delta) - \theta(x)}{\delta} - \frac{\theta(x) - \theta(x-\delta)}{\delta} \right) \\ &= \frac{\theta(x+\delta) - 2\theta(x) + \theta(x-\delta)}{\delta^2}. \end{aligned}$$

Summing these terms, $\Delta \theta_S = \frac{1}{\delta^2} (-6\theta_S + 6\bar{\theta}_S)$ where the source term is defined as:

$$\begin{aligned} \bar{\theta}_S := \frac{1}{6} [&\theta(x+\delta, y, z) + \theta(x-\delta, y, z) \\ &+ \theta(x, y+\delta, z) + \theta(x, y-\delta, z) \\ &+ \theta(x, y, z+\delta) + \theta(x, y, z-\delta)]. \end{aligned}$$

As the Laplace operator is independent of the coordinate system, we can apply any rotation and then average the source on the δ -sphere which we denote θ_S^* .

B ALGORITHMS

We provide the pseudo-codes for the method presented in Sec. 4. The sub-paths construction for each physical mode are given in separate algorithms, respectively Algo. 1 for convection, 2 for radiative transfer and 3 for conduction. The selection of a physical mode occurs at boundaries, as shown in Algo. 4.

Algorithm 1: Convective sub-path

Input: A time t

Output: A temperature θ

$$\tau_{conv} \leftarrow \frac{\rho c \mathcal{V}}{h \mathcal{S}};$$

$\tau_{b,conv} \leftarrow$ exponential sampling with a mean τ_{conv} ;

$t \leftarrow t - \tau_{b,conv}$; /* Time rewind */

if $t \leq t_I$ **then**

return $\theta_F(t_I)$; /* Initial temperature */

else

$\vec{x}_{\mathcal{S}} \leftarrow$ sampling of $\vec{X}_{\mathcal{S}}$ proportionally to h ;

return $\theta_B(\vec{x}_{\mathcal{S}})$; /* Algo. 4 */

end

Algorithm 2: Radiative sub-path

Input: A position \vec{x} and a direction $\vec{\omega}$

Output: A temperature θ

Loop

$s \leftarrow$ Trace a ray from \vec{x} along $\vec{\omega}$;

if $s = \infty$ **then**

return $\theta_{R,ambient}$;

else

$\vec{x} \leftarrow \vec{x} + s\vec{\omega}$;

$r \leftarrow$ uniform sampling in $[0, 1)$;

if $r < \epsilon(\vec{x})$ **then**

return $\theta_B(\vec{x})$; /* Algo. 4 */

else

$\vec{\omega} \leftarrow$ Sample a direction wrt p_r at \vec{x} ;

end

end

EndLoop

Algorithm 3: Conductive sub-path

Input: A position \vec{x} and a time t

Output: A temperature θ

Loop

$\vec{\omega} \leftarrow$ isotropic sampling of a direction;

$s^+ \leftarrow$ Trace a ray from \vec{x} along $+\vec{\omega}$;

$s^- \leftarrow$ Trace a ray from \vec{x} along $-\vec{\omega}$;

$\tilde{\delta} \leftarrow \min(\min(\delta, s^+), \min(\delta, s^-))$;

$$\tau_{cond} \leftarrow \frac{\rho c \tilde{\delta}^2}{6k};$$

$\tau_{b,cond} \leftarrow$ exponential sampling with a mean τ_{cond} ;

$t \leftarrow t - \tau_{b,cond}$; /* Time rewind */

if $t \leq t_I$ **then**

return $\theta_S(t_I)$; /* Initial temperature */

else

$\vec{x} \leftarrow \vec{x} + \tilde{\delta}\vec{\omega}$;

if \vec{x} is on boundary **then**

return $\theta_B(\vec{x})$; /* Algo. 4 */

end

end

EndLoop
