CHALMERS

Chalmers Publication Library

Computational optimization of catalyst distributions at the nano-scale

This document has been downloaded from Chalmers Publication Library (CPL). It is the author´s version of a work that was accepted for publication in:

Applied Energy (ISSN: 0306-2619)

Citation for the published paper:
Ström, H. (2015) "Computational optimization of catalyst distributions at the nano-scale". Applied Energy

http://dx.doi.org/10.1016/j.apenergy.2015.10.171

Downloaded from: http://publications.lib.chalmers.se/publication/225120

Notice: Changes introduced as a result of publishing processes such as copy-editing and formatting may not be reflected in this document. For a definitive version of this work, please refer to the published source. Please note that access to the published version might require a subscription.

Chalmers Publication Library (CPL) offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all types of publications: articles, dissertations, licentiate theses, masters theses, conference papers, reports etc. Since 2006 it is the official tool for Chalmers official publication statistics. To ensure that Chalmers research results are disseminated as widely as possible, an Open Access Policy has been adopted.
The CPL service is administrated and maintained by Chalmers Library.

(article starts on next page)

Computational optimization of catalyst distributions at the nano-scale

Henrik Ström[a,b,*]

[a]*Department of Applied Mechanics, Chalmers University of Technology, Gothenburg SE-412 96, Sweden*
[b]*Department of Energy & Environment, Chalmers University of Technology, Gothenburg SE-412 96, Sweden*
[*]*Corresponding author: henrik.strom@chalmers.se; +46317721360*

Abstract

Catalysis is a key phenomenon in a great number of energy processes, including feedstock conversion, tar cracking, emission abatement and optimizations of energy use. Within heterogeneous, catalytic nano-scale systems, the chemical reactions typically proceed at very high rates at a gas-solid interface. However, the statistical uncertainties characteristic of molecular processes pose efficiency problems for computational optimizations of such nano-scale systems. The present work investigates the performance of a Direct Simulation Monte Carlo (DSMC) code with a stochastic optimization heuristic for evaluations of an optimal catalyst distribution. The DSMC code treats molecular motion with homogeneous and heterogeneous chemical reactions in wall-bounded systems and algorithms have been devised that allow optimization of the distribution of a catalytically active material within a three-dimensional duct (e.g. a pore). The objective function is the outlet concentration of computational molecules that have interacted with the catalytically active surface, and the optimization method used is simulated annealing. The application of a stochastic optimization heuristic is shown to be more efficient within the present DSMC framework than using a macroscopic overlay method. Furthermore, it is shown that the performance of the developed method is superior to that of a gradient search method for the current class of problems. Finally, the advantages and disadvantages of different types of objective functions are discussed.

Keywords: Optimization; DSMC; Catalysis; Stochastic optimization; Nanoscale

---

This paper was presented at the 7th International Conference on Applied Energy (ICAE2015), March 28-31, 2015, Abu Dhabi, UAE (Original paper title: “A computational method to optimize the distribution of a catalytically active material inside a nano-scale pore” and Paper No.: 173).

1. Introduction

The most viable route to a reduction of the environmental costs of modern societies is an increase in the efficiencies of processes used in the manufacturing and transportation of products and in the production of energy [1]. Catalysis is a key phenomenon in a great number of relevant industrial processes, including feedstock conversion [2, 3], energy conversion [4], tar cracking [5], emission abatement [6, 7] and optimizations of energy use [8]. At the same time, nanotechnology has emerged as a subject area with a strong potential to enhance energy efficiency in all areas of the energy sector, from energy sources to energy change, distribution, storage and usage [9]. It now seems clear that the future development and optimization of fields such as renewable energy production and emission abatement will depend on the success of research activities related to reactive systems at the nano-scale.

More specifically, the current development of heterogeneous catalysis at the nano-scale is particularly promising [7]. Within heterogeneous, catalytic nano-scale systems, the chemical reactions are allowed to proceed at very high rates at a gas-solid interface (often the surface of a precious metal). The process efficiency is limited by the acceptable cost of the particular material in combination with the surface-to-bulk atom ratio obtainable, since the chemical reaction only occurs at the surface and the interior atoms remain unused. As the gas containing the reactants is typically brought into contact with the catalyst while being forced to flow past it, optimization of the efficiencies of such systems require numerical tools that take into account both the chemistry and the fluid dynamics of the system.

The governing processes in applications where the bounding geometry is of micro- or nanometer size typically span several orders of magnitude in spatial and temporal scales [10-12]. Consequently, there are many inherent difficulties involved in performing non-intrusive, non-destructive experimental investigations of the processes occurring on the smallest scales in such systems. Comprehensive numerical models therefore form an indispensable basis in the research into their behavior.

The multi-scale nature of heterogeneous catalytic systems has led to the development of a number of numerical frameworks dedicated to their study. Multi-scale models for reaction and transport in porous catalysts that are based on the continuum assumption have been proposed by several authors [13, 14]. However, such methods rely on the use of effective transport coefficients in macroscopic balance equations and cannot be used for systems where there is a net convective flow when the mean free path is significant to the bounding geometry (e.g.

when the Knudsen number is larger than 0.015) [15]. For such systems, continuum descriptions are not valid, and the predicted velocity fields are therefore erroneous, leading to inaccurate predictions of momentum, heat and mass transfer. The route to accurate descriptions of molecular flows with chemical reactions is via molecular methods, e.g. by obtaining the solution to the Boltzmann equation rather than the Navier-Stokes equations [15-18]. Solving the Boltzmann equation directly is however very difficult for real-world problems, as it represents a 7D partial differential equation (for the probability distribution of molecular positions and velocities over time). A more efficient approach is then to use a molecular simulation model, such as Direct Simulation Monte Carlo (DSMC) [19]. It has been shown that the DSMC method can be directly related to the Boltzmann equation and that solutions from the two frameworks are consistent [20, 21]. Furthermore, the DSMC method has the additional advantages of allowing treatment of inverse collisions and ternary chemical reactions, which becomes especially problematic in attempts at solving the Boltzmann equation directly [19]. The DSMC method is therefore well suited to describe reactive nano-scale systems [12, 22]. It is, in fact, the most widely used numerical algorithm in kinetic theory [23, 24] and has been experimentally validated for a great number of applications, including nonequilibrium gas flows (e.g. shocks) [24], rarefied gas dynamics (e.g. velocity, temperature and concentration slip) [25], near-vacuum flow of high-temperature gas at supersonic speeds [26], low-pressure deposition processes [27] and temperature-programmed desorption in heterogeneous catalysis [12].

A scientific problem of specific interest for heterogeneous catalytic systems is that of optimizing the catalyst distribution. In such an optimization process, the goal is to come up with a conceptual solution for the optimal design of the catalytic system, within a given design space while respecting a set of design constraints. Fine-tuning of the system with respect to actual real-life performance and manufacturability can then be carried out from an otherwise optimal starting point, resulting in significant reductions of the total development time and cost. However, most state-of-the-art optimization methods developed for reactive fluid flow systems rely on the availability of a system of partial differential equations describing the system in question. Hence, when the system to be optimized is described by a molecular method instead, many well-known optimization methods (such as the adjoint method for aerodynamic shape optimization [28, 29]) cannot be applied directly. Furthermore, as one of the most prominent characteristics of molecular systems is the existence of statistical uncertainties [30, 31], any chosen objective function will always contain some degree of noise. In the optimization of a reactive nano-scale system, it is therefore reasonable to choose an optimization heuristic that can find the approximate global optimum while

handling uncertainties in the objective function and discrete search spaces. One optimization heuristic that fulfils these requirements is the stochastic optimization approach known as simulated annealing [32, 33]. The simulated annealing method is simple to implement, relatively fast and has been found to be more accurate than genetic algorithms and maximum entropy reconstruction techniques in reconstructions of heterogeneous media [34]. Interestingly, the simulated annealing method itself is also classified as a Monte Carlo method, as it is constitutes an adaptation of the Metropolis-Hastings algorithm [35].

There have been previous attempts at using the DSMC method in optimization studies. Recently, Pflug et al. [36] used DSMC to optimize the film thickness uniformity in an industrial physical vapor deposition (PVD) reactor. However, the design of the PVD reactor in question allowed for a single DSMC computation to be used (per sputtering material) to optimize the geometry. For a generic heterogeneous catalytic system, such simplifications are typically not possible. Furthermore, enumeration approaches (repeating simulations for several values of a given design parameter) can be useful for systems that are easily characterized by a small set of design parameters [cf. 37], but are also not generally applicable to problems involving the distribution of a catalytic material over the walls of a fixed system, as the possible designs are too many and cannot easily be grouped together.

Hence, the purpose of the present work is to develop a numerical tool for investigations of reacting, molecular flows that can be applied in optimizations of catalytic systems on the micro- and nano-scales. The main challenge to be faced is thus related to the inevitable fluctuations in any objective function that result from the natural uncertainty pertaining to the molecular regime: when the continuum approximation no longer holds, fluctuations appear in the macroscopic properties derived by averaging over molecular properties [38]. It is shown in the current work that a successful (i.e. robust and computationally efficient) optimization approach for such cases is dependent on a well-balanced combination of the choice of objective function, optimization algorithm and convergence criteria.

2. Modeling

2.1. The Direct Simulation Monte Carlo (DSMC) framework

The route to accurate descriptions of molecular flows with chemical reactions is via molecular methods, e.g. by obtaining the solution to the Boltzmann equation rather than the Navier-

Stokes equations [15-18]. The molecular method chosen in the present work is the well-established Direct Simulation Monte Carlo (DSMC) method [19], which is a probabilistic simulation approach. The fluid flow is modelled using particles that represent a large number of real molecules. The particle motion is updated deterministically using a time step that is smaller than the mean collision time, so that the intermolecular collisions can be decoupled from the molecular motion. The code developed in this work is based on Bird's DSMC method [19] and is able to handle chemistry in the gas phase and to approximate the rate of surface chemistry reactions at gas-solid interfaces [39, 40]. More specifically, the code is designed for performing optimizations of the distribution of an active material over the bounding surfaces of the system. The aim is to make possible a complete optimization during the course of one single simulation run. The code is written in the programming language C.

The DSMC procedure has been described extensively elsewhere [19], and will only be summarized briefly here. With DSMC, the ensemble of molecules is modeled with a reduced number of computational molecules that move in straight lines according to their velocities for a short time step during which no collisions take place. Thereafter, collisions are modelled using random numbers and collision probabilities, which are based on the collision cross section and the relative velocities between pairs of molecules. Here, the hard sphere model is used to determine the outcome of a collision (scattering angles and post-collision velocities) [19], as it is sufficient for the purpose of the present work. If a molecule collides with a wall, an adsorption or wall reaction event can be triggered, and if two molecules collide, a homogeneous chemical reaction can result.

At certain intervals, sampling is performed over the molecules to derive the macroscopic fields of interest, such as the mass-averaged gas velocity, temperature and species or number concentrations. This sampling is performed on a computational mesh that is coarser than the mesh used to calculate collisions, and the two types of cells are typically referred to as cells (or samplings cells) and subcells (or collision cells), respectively. The flow is sampled every fourth time step to obtain samples with only a small degree of correlation. A schematic diagram of the DSMC algorithm within the layout of the complete code is shown in Figure 1.

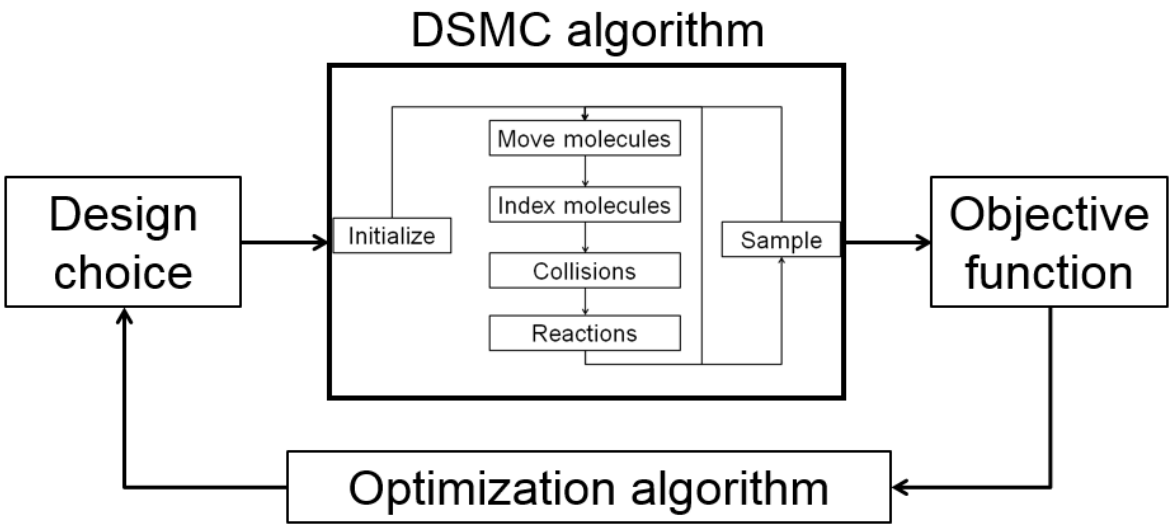

Figure 1. Schematic layout of the DSMC algorithm and the optimization algorithm and their interconnectedness.

2.2. The stochastic optimization heuristic

Macroscopic fields in a DSMC procedure are always deduced by averaging over the computational molecules and therefore have a tendency to contain statistical scatter, as the presence of a statistically significant number of molecules in all sampling cells at all times cannot be guaranteed [30]. (On the contrary, it is precisely this loss of statistically significant averages that causes the continuum approximation to break down and that therefore characterizes the molecular flow regimes). This problem becomes further emphasized with the DSMC method, since the actual number of real molecules is represented by a smaller number of computational molecules, which acts so as to increase the scatter [31]. In addition, the objective function in the optimization process may very well exhibit both global and local optima. Furthermore, it is not possible to test every conceivable geometrical design due to the large computational cost of such an investigation. A suitable compromise is then to use a stochastic optimization method, such as simulated annealing [32, 33, 41, 42]. Simulated annealing cannot be guaranteed to find the global optimum of an objective function, but it can avoid becoming trapped in a local optimum (when there is a better global optimum somewhere else), it prevents premature termination due to scatter in the objective function and it helps avoid the tedious task of investigating every possible design case.

In the current work, the following simulated annealing heuristic is used:

1) Sample the objective function for one initial (randomly chosen) design case.
2) Pick another design case (using an algorithm that has to be specified separately).
3) If the new case is better, move to it. If it is worse, accept it anyway with a certain probability, *P*. This probability is to be a function of the time elapsed in the optimization process and it too has to be specified separately.
4) Repeat steps 2-3 a pre-determined number of rounds or until the objective function reaches a pre-defined threshold value.

At the heart of the simulated annealing algorithm lies the determination of the probability *P*. The original probability function of Kirkpatrick et al. [32] is here modified slightly, so that

$$P = exp\left[\frac{-(f' - f)/f}{T(t)}\right]$$

In this notation, *f* is the value of the objective function, a prime denotes the value for the newer design case, and *T(t)* is the analogue of temperature in a physical annealing process. In this work, the function *T(t)* is defined as

$$T(t) = a\tau(1 - b\, t/t_{max})$$

where *t* is the total time elapsed in the simulation and t_{max} is the time at which the optimization process is stopped. Hence, the tendency to accept a design case that is worse decreases with time. The variable τ represents the convergence criterion for the normalized change in the objective function, and the parameters *a* and *b* thus determine the behavior of *P* in time. These values should be chosen to enable a more global character of the search initially, and to progress towards a local search in the most promising region with time. The optimum values for a generic problem will always be problem-dependent to some extent. In the current work, the values $a = 2$ and $b = 0.95$ were found to produce satisfactory results.

The final component in the optimization routine is the algorithm for picking another design. This component is not prescribed by the simulated annealing algorithm as such, but typically involves a randomized selection of either the step length, the step direction, or both [41]. The following algorithm was applied in the present work:
1) The new design is obtained by moving the catalytically active region of interest a (uniformly distributed) random distance in the interval [$\Delta x/20$, $\Delta x/10$], where Δx is the extent of the domain in coordinate direction *x*.

2) The direction in which to move is by default the direction of increasing value of the objective function, but in 25% of the cases the direction is reversed to introduce a random behavior also to the design picking algorithm.
3) Modifications to the new design choice are made if needed to ensure that the geometric bounds of the system are respected.

A schematic illustration of the implementation of the optimization algorithm into the DSMC framework is shown in Figure 1. For every design that is to be evaluated, the DSMC code needs to run long enough for the samples used to calculate the objective function to converge. It is therefore evident that the choice of objective function and the robustness of the optimization algorithm with regard to fluctuations and sampling errors are of utmost importance in the derivation of the combined procedure.

3. Results and Discussion

The DSMC code is validated in an extensive series of tests, of which only a subset are reported here. Thereafter, the task of finding the optimum position for a catalytically active region inside a three-dimensional pore is used to test the robustness and efficiency of the proposed optimization algorithm.

3.1. Validation of the DSMC code

As the most challenging aspect of the DSMC procedure in relation to molecular motion lies in the modelling of molecular collisions, one fundamental validation test carried out is that of a homogeneous gas at rest in a one-dimensional domain. This test case proves that the code is able to predict the correct solution for a one-dimensional homogeneous gas and that the performance of the random number generator used is acceptable. Indeed, the number of collisions predicted is very close to the theoretical value [19], and the mean collision separation is less than 5% of the cell-width, meaning that collision partners are located within the same subcell.

Next, the performance of the DSMC code is exemplified for a non-isothermal fluid flow test case where the domain is a flow between two planes separated by a distance of 0.5 m. This distance is divided into 40 sampling cells with 10 collision subcells each. The number of

computational molecules is 10^5, and the number density of molecules is 10^{20} m^{-3}. The Knudsen number is approximately 0.03. The lower wall is stationary and maintained at a temperature of 250 K. It has surface properties such that 50% of the incoming molecules are specularly reflected, whereas the rest are diffusively reflected. The upper wall is maintained at 300 K and moves at a velocity of 1 m/s in the plane perpendicular to the gap. At this wall, 80% of the incoming molecules are specularly reflected. For this problem, the performance of the current code is validated by comparing its predictions to benchmark results from one of Graeme Bird's program in the DS suite (DS1V) [43], in line with Bird's recommendation on how to assess the validity of a new DSMC code. The predicted temperature slip is approximately 3.4 times higher at the hotter boundary, which agrees well with the DS1V solution as shown in Figure 2.

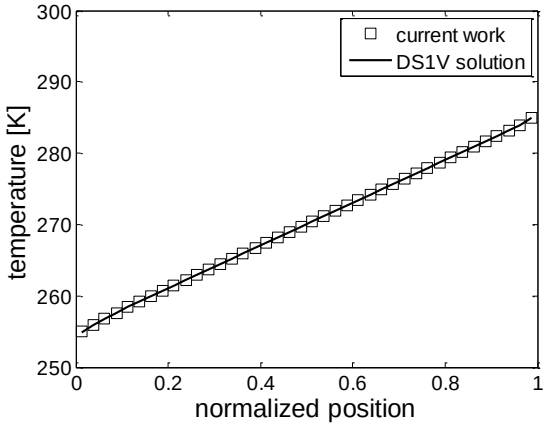

Figure 2. Sampling cell temperature as a function of the normalized position between a stationary and a moving wall of different temperatures. The plane at position 0 is maintained at 250 K and is specularly reflecting to 50% while the plane at position 1 is maintained at 300 K and is specularly reflecting to 80%. The predictions obtained in the current work are in excellent agreement with the DS1V solution.

The ability of the code to handle different molecular species and chemical reactions is validated in a homogeneous chemistry test case. Two stationary walls are separated by a distance of 0.5 m. Both walls have surface properties such that there is 100% specular reflection. The gap between the walls is initially occupied by oxygen (O_2) at a number density of 10^{20} m^{-3} and a temperature of 5000 K. In this test case, two chemical reactions may occur, namely the dissociation and recombination, respectively, of diatomic oxygen and atomic

oxygen: $O_2 \leftrightarrow 2O$. The total duration of the simulation is approximately one second of real time, and the result is shown in Figure 3. Again, the current code is in excellent agreement with the DS1V solution.

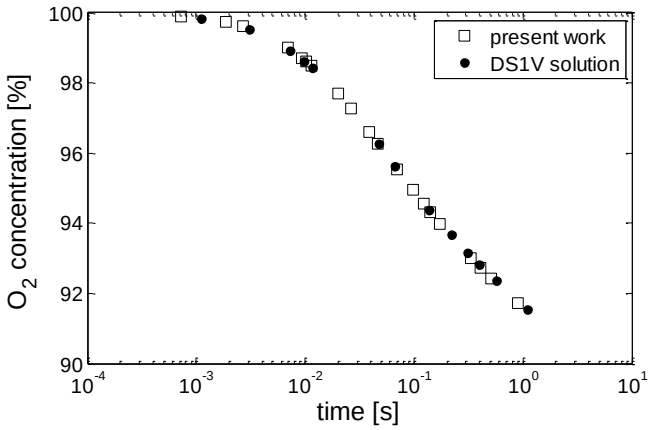

Figure 3. Volume-averaged concentration of diatomic oxygen as a function of time. The gas consists of 100% pure O_2 initially. The predictions obtained in the current work are in excellent agreement with the corresponding DS1V solution.

A comprehensive treatment of surface reactions in the DSMC framework would necessitate detailed modeling of adsorption, desorption, coadsorption, reaction, surface diffusion and the effects of surface defects, which is currently beyond the state-of-the-art for this computational technique, although significant advances are made continuously [12]. Such additional complexities would also add to the computational cost and significantly reduce the efficiency of the optimization. In this work, surface reactions are therefore instead implemented as occurring at a wall with a certain probability [40]. This probability can be tuned to reproduce a physical reaction rate, implying that the main simplification involved is related to the loss of coverage-dependence. This simplification is deemed appropriate in the light of the main goal, which is to combine the DSMC simulation with an inline optimization routine.

As a test case for the surface reaction setup, the wall-catalyzed dissociation of oxygen is simulated at a temperature of 300 K in the same geometry as the homogeneous validation case. Dissociation is prescribed to occur at the walls with a reaction probability of 1%. The temporal evolution of the atomic oxygen concentration profile is depicted in Figure 4. The results presented here agree with what is qualitatively expected for the system under

investigation and thus constitute a verification of the implementation of the wall reaction mechanism into the code.

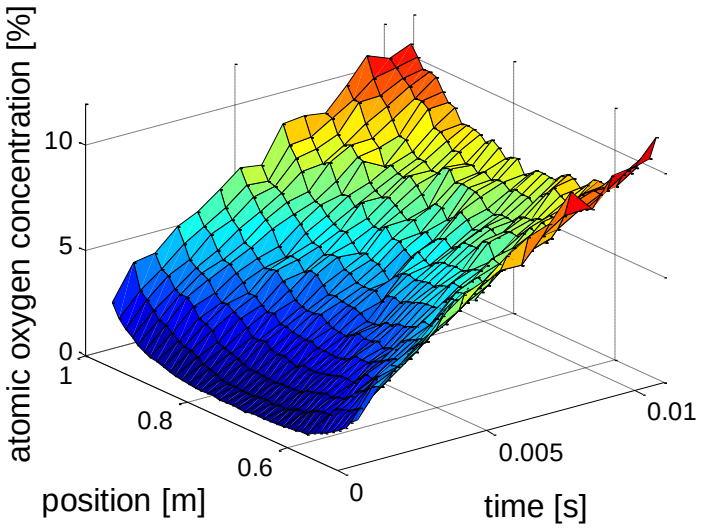

Figure 4. 3D surface plot showing the time-and-space resolved atomic oxygen concentration in the heterogeneous chemistry test case. The gas is initially pure diatomic oxygen (O_2). Upon collision with a wall, there is a 1% probability that a diatomic oxygen molecule dissociates ($O_2 \leftrightarrow 2O$). The code predicts that the concentration of atomic oxygen resulting from this dissociation increases with time and penetrates into the domain.

In conclusion, the observations from these validation tests support the inference that the code can be used for the optimization processes described next.

3.2. Optimization of the position of a catalytically active region

A three-dimensional nano-scale "pore" can be constructed by having two boundaries specified as an inlet and an outlet, and a procedure is implemented by which the pressure difference between these boundaries is maintained throughout the simulation [44]. The remaining four sides of the domain are regular walls, which may be designed with or without protrusions. The aim is to design a numerical framework that can be used to determine the optimal distribution of a limited amount of catalytically active material over this pore wall surface. To simplify the problem setup, the catalytically active material is limited to a single surface

location, but generalization to an arbitrary number of active sites is straightforward. Consequently, in the simulations presented here, the catalytically active region is a 2 nm thin section around the perimeter of a rectangular (100 *x* 100 *x* 300 nm) 3D pore.

3.2.1 Objective function

The most important feature of the objective function is that it converges quickly, which means that the signal should be strong in comparison to the fluctuations present. A comprehensive treatment of a catalytic reaction at the active region would have to account for sticking factors lower than unity [45], temperature-dependence of the chemical reaction rates and changes to the gas phase composition that could potentially affect the molecular flow field. Such effects would however tend to increase both the magnitude of the fluctuations (by making the objective function signal lower because of the lower reaction probability) and to delay the convergence towards a steady state (by introducing changes into the molecular flow field). For the purpose of optimization, the additional value in terms of the accuracy possibly gained from adding such descriptions is small in relation to the computational cost. Hence, in the current work, the chemical reaction is assumed not to influence the molecular flow field. In other words, no actual reaction (wherein molecules change nature) is carried out, but the molecules that have made contact with the catalytically active surface are marked, so that their concentration can be monitored by the code. This approach is equivalent to monitoring the impingement rate on the catalytically active sites, rather than the actual reaction rate [46]. Under the assumption that the flow field does not change significantly with the trace species conversion, this method may thus reduce the computational cost of obtaining converged statistics by several orders of magnitude. Additionally, it could be thought of as a means of probing the mass transfer rate towards the catalytically active sites (e.g., as in CO oxidation experiments over Pt/Al_2O_3 catalysts). It is well known that the reduction of the real number of molecules to a smaller number of computational molecules in DSMC makes the method sensitive to the prediction of rare events, which have low probability and therefore would require a large number of computational molecules to be reproduced correctly. A further advantage with the proposed approach is therefore that it makes use of all computational molecules, rather than the small fraction that reacts with the catalyst surface upon impingement.

The design picking algorithm is implemented to move the catalytically active section around in the domain. This algorithm waits for a steady signal from the outlet sampling of marked

molecules before changing the location of the catalytically active section as proposed by the simulated annealing algorithm. The objective function is judged to have converged when the relative change between two samples is less than 10^{-3}. The convergence history for a typical design with the specified convergence criterion is shown in Figure 5. The simulation for this design is continued from the last state of the simulation for the previous design. It is clear that for the objective function to be useful in finding the optimum location, it need not provide a highly accurate value for the converged number of marked molecules on the outlet. Instead, the accuracy necessary is determined only by the need to be able to tell two different designs apart.

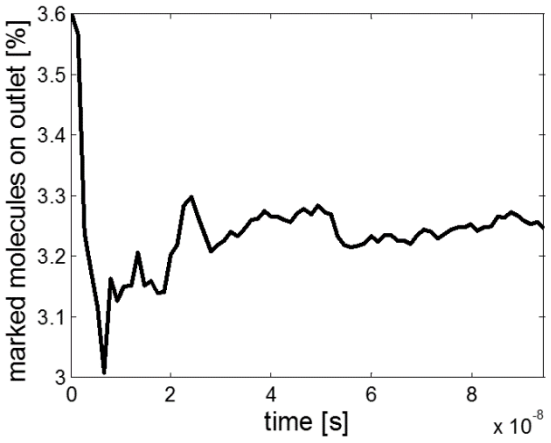

Figure 5. Visualization of the objective function versus time for a given position of the catalytically active slit in the 3D pore.

The effect of the number of computational molecules employed on the fluctuations in the objective function was also investigated, for 10^3, 10^4 and $2{\cdot}10^4$ molecules, respectively. There were no significant adverse effect from employing a smaller number of computational molecules, leading to the decision to use 10^3 molecules in the subsequent optimization runs. For a more complicated pore structure than the current one, the requirements for the number of computational molecules could possibly increase [47].

The computational cost for one DSMC run (i.e. one call to the objective function) depends on the number of computational molecules used, the number of cells used for the discretization of the computational domain and the time step employed. For a small number of computational molecules and a small computational domain (as used here), the time step is the most limiting factor. At atmospheric conditions, the mean collision time is approximately

10^{-10} s, implying that the objective function would converge within 10^4 time steps (the time step is of the order of 10^{-11} s and the total time needed somewhat less than 10^{-7} s (cf. Figure 5)). For the chosen design case, the corresponding run-time for one call on a single CPU is then several hours. However, atmospheric conditions represent an extreme case in the limit of zero Knudsen number, and lower pressures or higher temperatures reduce this time significantly. Similarly, much more complex geometries could significantly increase the number of computational molecules needed, which would increase the computational cost and probably make parallelization of the DSMC algorithms necessary [48].

3.2.2 Sample fluctuations

Fluctuations in the sampled DSMC properties are the main challenge for the optimization algorithm. There are macroscopic overlay methods [49, 50] available for the DSCM framework that are based on the solution of trace species transport equations using the flow field of the other (dominating) species. Such methods thus represent a solution to the problem of treating very rare events in DSMC without having to resort to using an excessive number of computational molecules. However, these approaches are susceptible to numerical errors if the sampling of the macroscopic fields has not yet converged. This is a significant drawback in optimization, and it makes these methods less efficient than the sampling of marked molecules as proposed here.

As an example, consider the sampling of the three velocity components u, v and w in a randomly chosen cell in a three-dimensional pore with a gas at rest at 300 K, as depicted in Figure 6. It is seen that a total time of 0.1 ms is needed to obtain an estimate of the steady-state solution which is correct within approximately 10^{-3} m/s. When the simulation is terminated, hundreds of billions of molecular moves and billions of collisions have been performed, but mass conservation is still only within 0.1% error tolerance. Errors of such magnitude are still too large to be acceptable in the solution of a species transport equation with chemical reaction source terms for a species present in trace amounts.

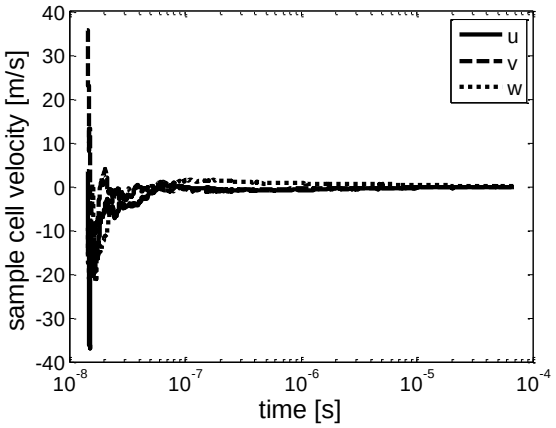

Figure 6. Convergence history for the three sample velocity components in a randomly chosen cell in a 1 *x* 1 *x* 3 μm domain (discretized into 8 *x* 8 *x* 24 sampling cells). All velocity components tend to zero as time increases, which is the result expected for a gas at rest.

3.2.3 Optimization

The optimization algorithm is evaluated for an objective function that contains two local optima (and three extreme points), as depicted in Figure 7. The global optimum is positioned around $z = 0.2$ and is ~2.6%, where z is the normalized position of the catalytically active slit in the streamwise direction. There is also a local maximum (~1.5%) at around $z = 0.7$. Figure 7 represents the converged objective function, but for any sampling from a DSMC simulation there will always be a significant uncertainty due to the presence of noise in the signal. This noise emanates from the molecular uncertainties and is further influenced by the convergence criteria used: in order for computational efficiency not to be lost, the sampling that produces the objective function signal has to be terminated within a realistic time frame, and so the signal will always be somewhat colored by noise. The extreme point that separates the curves leading to the two maxima is located at $z = 0.5$. Hence, for the current objective function – and in the absence of noise – a gradient search optimization process starting from a random location would find the global optimum in 50% of the cases and the other maximum in the remaining 50% of the cases. The aim here is to prove that the simulated annealing algorithm can exhibit superior performance to such an algorithm.

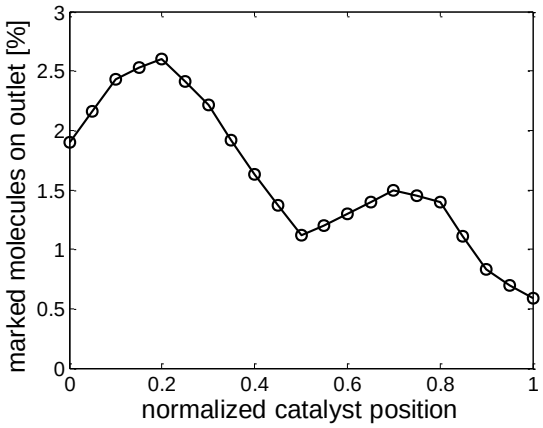

Figure 7. Objective function with two local optima, but only one global, and three extreme points.

The statistics for the results of the simulated annealing algorithm when applied to a system that is described by the aforementioned objective function are shown in Figure 8. The optimization process has been repeated 1000 times for every limit on the number of calls to obtain reliable statistics. The algorithm is typically able to find the correct optimum after 20 function calls. As the total number of calls allowed increases, the number of unsuccessful simulation runs decreases significantly. This behavior is superior to the performance of a gradient search method. A gradient search would only be able to find the local optimum closest to the initial position (and only if allowed a large enough number of calls), and would not in general be able to handle the fluctuations in the objective function. Consequently, a gradient search method could at best produce two ridges of equal height in Figure 8.

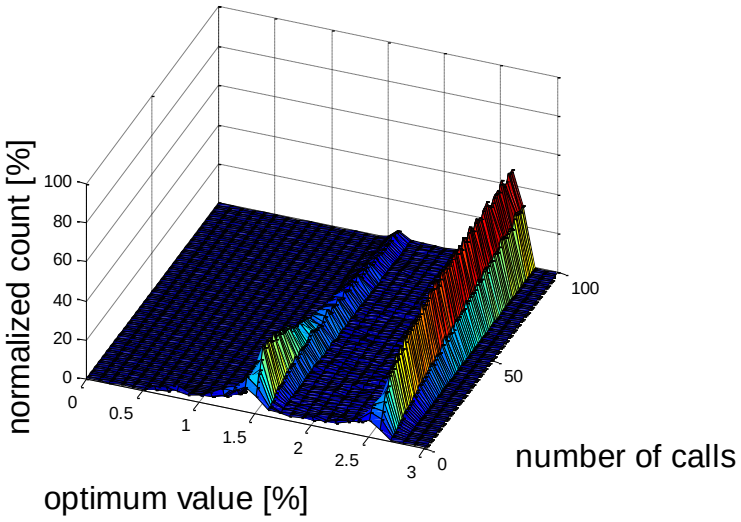

Figure 8. Statistics for the simulated annealing algorithm developed in the present work for the objective function in Figure 7.

Finally, it should be stressed here that the number of calls needed to find the correct optimum is much dependent on the convergence criterion for the objective function (cf. Section 3.2.1). If this convergence criterion is relaxed, each function call will be cheaper but more function calls will be needed. Similarly, if the convergence criteria is tightened, the number of calls needed will decrease further, at the expense of a higher computational cost for each call. The optimum settings for the optimization algorithm itself will therefore depend mostly on the signal-to-noise ratio of the objective function chosen for the system under study. Furthermore, the computational cost for the DSMC simulation is approximately proportional to the number of computational molecules employed, implying that one would like to use as few molecules as possible. At the same time, however, the time needed to obtain statistically converged DSMC results increases with decreasing the number of computational molecules. In practice, the applicability of the DSMC approach is therefore mainly limited by the restriction that the time step must be smaller than the collision time, which implies that weakly rarefied flows (where the number density of molecules is relatively high) are much more computationally expensive than strongly rarefied flows (where the number density is low). In relation to the hierarchy of pores existing in a realistic porous medium, the DSMC technique is therefore most suited to study the behavior in the smaller pores, although it should be stressed that there are no limitations to the validity of the approach for the entire range of pore sizes. For the methodology developed in the current work to be more efficiently applied to a large computational domain spanning a wide pore size distribution, it is likely that some kind of hybrid DSMC/CFD (computational fluid dynamics) method would be most appropriate [51].

4. Conclusions

A Direct Simulation Monte Carlo (DSMC) code has been developed that treats molecular motion in wall-bounded systems with homogeneous and heterogeneous chemical reactions. A simulated annealing optimization algorithm is implemented to allow for optimization of the distribution of a catalytically active material within a 3D pore where the flow field is described by the present code. It is shown that the performance of the simulated annealing method for the current class of problems is superior to that of a gradient search method, in that it enables optimizations also of systems that exhibit global and local optima as well as fluctuations. Furthermore, using a stochastic optimization heuristic to handle the presence of noise in the sampling of the objective function is shown to be more efficient than using a macroscopic overlay method.

To minimize the difficulties involved with handling noisy objective functions, the objective function should be a strong signal, suggesting that probing the mass transfer rate towards the catalytically active sites is more efficient than trying to approximate the actual surface reaction rate, as long as the local coverage does not vary significantly with the position of the catalyst material. The approach described in the present work thus represents a suitable starting-point for addressing a number of important research challenges involving the optimization of reacting nano-scale flows and reacting heterogeneous flows with and without surface diffusion.

Acknowledgements

This work was financed by a pilot project grant from the Chalmers e-Science Centre (CheSC).

References

[1] Dudukovic MP. Reaction engineering: Status and future challenges. Chem Eng Sci 2010;65:3-11.

[2] Lin L, Cunshan Z, Vittayapadung S, Xiangqian S, Mingdong D. Opportunities and challenges for biodiesel fuel. Appl Energy 2011;88:1020-1031.

[3] Salmi T. Chemical reaction engineering of biomass conversion. Adv Chem Eng 2013;42:195-260.
[4] Menon, V, Banerjee, A, Dailly, J, Deutschmann, O. Numerical analysis of mass and heat transport in proton-conducting SOFCs with direct internal reforming, Appl Energy 2015;149:161-175.
[5] Palma CF. Modelling of tar formation and evolution for biomass gasification: A review. Appl Energy 2013;111:129-141.
[6] Liang Z, Ma X, Lin H, Tang Y. The energy consumption and environmental impacts of SCR technology in China. Appl Energy 2011;88:1120-1129.
[7] Chalmers University of Technology, Studies of individual nanoparticles can be the key to future catalysis, https://www.chalmers.se/en/departments/ap/news/Pages/Langhammer-KAW.aspx, 2015 (Accessed on 2015-10-27).
[8] Rahimpour MR, Dehnavi MR, Allahgholipour F, Iranshahi D, Jokar SM. Assessment and comparison of different catalytic coupling exothermic and endothermic reactions: A review. Appl Energy 2012;99:496-512.
[9] Hessen-Nanotech. Application of nanotechnologies in the energy sector, Aktionslinie Hessen-Nanotech 2008;9:1-84.
[10] Nagel DJ, Zaghloul ME. MEMS: micro technology, mega impact. IEEE Circuits Syst Mag 2001;17:14-25.
[11] Konstandopoulos AG, Kostoglou M. Microstructural aspects of soot oxidation in diesel particulate filters. SAE Techn Paper 2004;01-0693.
[12] Pesch GR, Riefler N, Fritsching U, Ciacchi LC, Mädler L. Gas-solid catalytic reactions with an extended DSMC model. AIChE J 2015;61:2092-2103.
[13] Kočí, P, Novák, V, Štěpánek, F, Marek, M, Kubíček, M. Multi-scale modelling of reaction and transport in porous catalysts. Chem Eng Sci 2010;65:412-419.
[14] Maffei, T, Gentile, G, Rebughini, S, Bracconi, M, Manelli, F, Lipp, S, Cuoci, A, Maestri, M. A multiregion operator-splitting CFD approach for coupling microkinetic modeling with internal porous transport in heterogeneous catalytic reactors. Chem Eng J 2016;283:1392-1404.
[15] Cercignani C. Rarefied gas dynamics: From basic concepts to actual calculations. Cambridge University Press; 2000.
[16] Bailey CL, Barber RW, Emerson DR. Is it safe to use Navier-Stokes for gas microflows? European Congress on Computational Methods in Applied Sciences and Engineering, ECCOMAS; 2004.
[17] Gu XJ, Emerson DR. A high-order moment approach for capturing non-equilibrium phenomena in the transition regime. J Fluid Mech 2009;636:177-216.

[18] Xu B, Ju Y. Concentration slip and its impact on heterogeneous combustion in a micro scale chemical reactor. Chem Eng Sci 2005;60:3561-3572.
[19] Bird G. Molecular gas dynamics and the direct simulation of gas flows. Clarendon; 1994.
[20] Bird, GA. Direct Simulation and the Boltzmann equation. Phys Fluids 1970;13:2676-2681.
[21] Wagner, W. A convergence proof for Bird's Direct Simulation Monte Carlo method for the Boltzmann equation, J Stat Phys 1992;66:1011-1044.
[22] Galinsky, M, Sénéchal, U, Breitkopf, C. The impact of microstructure geometry on the mass transport in artificial pores: a numerical approach. Model Sim Eng 2014; Article ID 109036, 7 pages.
[23] Wagner, W. Monte Carlo methods and numerical solutions, Proc 24th Int Symp Rarefied Gas Dynamics, Monopoli (Bari), Italy, 10-16 July, 2004.
[24] Erwin, DA, Pham-Van-Diep, GC, Muntz, EP. Nonequilibrium gas flows. I: A detailed validation of Monte Carlo direct simulation for monatomic gases. Phys Fluids A 1991;3:697-705.
[25] Wadsworth, DC. Slip effects in a confined rarefied gas. I: Temperature slip. Phys Fluids A 1993;5:1831-1839.
[26] Spencer, RL, Taylor, N, Farnsworth, PB. Comparison of calculated and experimental flow velocities upstream from the sampling cone of an inductively coupled plasma mass spectrometer, Spectrochimica Acta B Atom Spectroscopy 2009;64:921-924.
[27] Dorsman, R, Kleijn, CR, Velthuis, JFM, Zijp, JP, van Mol, AMB. Zinc deposition experiments for validation of direct-simulation Monte Carlo calculations of rarefied internal gas flows, J Vac Sci Technol 2007;25:474-479.
[28] Jameson, A. Aerodynamic shape optimization using the adjoint method, VKI Lecture Series on Aerodynamic Drag Prediction and Reduction, von Karman Institute of Fluid Dynamics, Rhode St Genese, 2003.
[29] Giannakoglou, KC, Papadimitriou, DI. Adjoint methods for shape optimization, in: Optimization and Computational Fluid Dynamics (eds. Thévenin, D, Janiga, G), Springer Verlag, 2008.
[30] Plotnikov, MY, Shkarupa, EV. Theoretical and numerical analysis of approaches to evaluation of statistical error of the DSMC method. Comp Fluids 2014;105:251-261.
[31] Oran, ES, Oh, CK, Cybyk, BZ. Direct Simulation Monte Carlo: Recent advances and applications, Ann Rev Fluid Mech 1998;30:403-441.
[32] Kirkpatrick S, Gelatt Jr DA, Vecchi MA. Optimization by simulated annealing. Science 1983;220:671-680.

627 [33] Černý V. Thermodynamical Approach to the Traveling Salesman Problem: An Efficient
628 Solution Algorithm. J Optimiz Theory App 1985;45:41-51.
629 [34] Patelli, E, Schuëller, G. On optimization techniques to reconstruct microstructures of
630 random heterogeneous media, Comp Materials Sci 2009;45:536-549.
631 [35] Metropolis, N, Rosenbluth, AW, Rosenbluth, MN, Teller, AH, Teller, E. Equation of
632 state calculations by fast computing machines. J Chem Phys 1953;21:1087-1092.
633 [36] Pflug, A, Siemers, M, Melzig, T, Rademacher, D, Zickenrott, T, Vergöhl, M. Numerical
634 optimization of baffles for sputtering optical precision filters. Surf Coat Techn 2014;241:45-
635 49.
636 [37] Wei, J-L, Hu, C-D, Xie, Y-L, Liang, L-Z. Optimization of the gas flow in the
637 neutralization region of EAST neutral beam injector, Fusion Eng Des 2013;88:3176-3179.
638 [38] Bird RB, Stewart WE, Lightfoot EN. Transport phenomena, 2nd Edition. John Wiley &
639 Sons, Inc.; 2002.
640 [39] Herdrich G, Fertig M, Petkow D, Steinbeck A, Fasoulas S. Experimental and numerical
641 techniques to assess catalysis. Prog Aerospace Sci 2012;27:48-49.
642 [40] Tomarikawa K, Yonemura S, Tokumasu T, Koido T. Numerical analysis of gas flow in
643 porous media with surface reaction. AIP Conf Proc 2011;1333;796-801.
644 [41] Ekren O, Ekren BY. Size optimization of a PV/wind hybrid energy conversion system
645 with battery storage using simulated annealing. Appl Energy 2010;87:592-598.
646 [42] Robertson JJ, Polly BJ, Collis JM. Reduced-order modeling and simulated annealing
647 optimization for efficient residential building utility bill calibration. Appl Energy
648 2015;148:169-177.
649 [43] Bird GA. The DS2V/3V Program Suite for DSMC Calculations. AIP Conf Proc
650 2005;762:541-546.
651 [44] Wang M, Li Z. Simulations for gas flows in microgeometries using the direct simulation
652 Monte Carlo method. Int J Heat Fluid Flow 2004;25:975-985.
653 [45] Steininger H, Lehwald S, Ibach H. On the adsorption of of CO on Pt(111). Surface Sci
654 1982;123:264-282.
655 [46] Carlsson P-A, Österlund L, Thormälen P, Palmqvist A, Fridell E, Jansson J, Skoglundh
656 M. A transient in situ FTIR and XANES study of CO oxidation over Pt/Al2O3 catalysts. J
657 Catal 2004;226:422-434.
658 [47] Bird, GA. Sophisticated DSMC, DSMC07 Meeting, Santa Fe, 30 September, 2007.
659 [48] Roohi, E, Darbandi, M. Recommendations on performance of parallel DSMC algorithm
660 in solving subsonic nanoflows, Appl Math Model 2012;36:2314-2321.
661 [49] Boyd ID, Candler GV, Levin DA. Dissociation modeling in low density hypersonic
662 flows of air. Phys Fluids 1995;7:1757-1763.

663 [50] Lilley CR, Macrossan MN. A macroscopic chemistry method for the direct simulation of
664 gas flows. Phys Fluids 2004;16:2054-2066.
665 [51] Lofthouse, AJ, Boyd, ID, Wright, MJ. Effects of continuum breakdown on hypersonic
666 aerothermodynamics, AIAA Paper 2006-993.