



Institut für Numerische Simulation

Rheinische Friedrich-Wilhelms-Universität Bonn

Wegelerstraße 6 • 53115 Bonn • Germany
phone +49 228 73-3427 • fax +49 228 73-7527
www.ins.uni-bonn.de

A. Rüttgers, M. Griebel, L. Pastrik, H. Schmied,
D. Wittmann, A. Scherrieble, A. Dinkelmann,
T. Stegmaier

**Simulation of the oil storage process in the scopa
of specialized bees**

INS Preprint No. 1404

June 2015

Simulation of the oil storage process in the scopa of specialized bees

A. Rüttgers^{a,*}, M. Griebel^{a,b}, L. Pastrik^c, H. Schmied^c, D. Wittmann^c, A. Scherrieble^d, A. Dinkelmann^d, T. Stegmaier^d

^aInstitute for Numerical Simulation (INS), University of Bonn, Wegelerstr. 6, 53115 Bonn, Germany

^bFraunhofer Institute for Algorithms and Scientific Computing (SCAI), Schloss Birlinghoven, 53754 Sankt Augustin, Germany

^cInstitute of Crop Science and Resource Conservation (INRES), University of Bonn, Melbweg 42, 53115 Bonn, Germany

^dInstitute of Textile Technology and Process Engineering Denkendorf (ITV Denkendorf), Körschtalstr. 26, 73770 Denkendorf, Germany

Abstract

Several species of specialized bees possess special structures to store and transport floral oils. By using closely spaced hairs at their back legs, the so called scopa, these bees can absorb and release oil droplets without loss. The high efficiency of this process is a matter of ongoing research. Based on recent x-ray microtomography scans from a bee's scopa, we build a three-dimensional geometric computer model. Then, using *NaSt3DGPF*, a two-phase flow solver developed at the Institute for Numerical Simulation, we compute the micro flow in the scopa model. Our calculations reveal the laminar to turbulent air flow in the scopa during flight. Furthermore, we simulate the deformation of an oil droplet in the scopa due to surface tension effects on a microscopic length scale. Our results are in good agreement with measurements for an oil-wetted scopa at steady state which are obtained from x-ray scans. Both simulations are relevant for the understanding of the process of oil absorption and transportation in the real scopa of a bee. Due to the large computational complexity of the problem, massively parallel computations are essential for our simulations.

Keywords: bee morphology, two-phase flow, high-performance computing, large eddy simulation, simulation on x-ray computed tomography data

1. Introduction

Several species of specialized bees such as Centridini, Tetrapediini, Ctenoplectridae, see Fig. 1(a), have developed structures to store and transport fatty floral oils instead of nectar (cf. Dressler [1], Buchmann [2]). The floral oil is mixed with pollen and fed to the bee's larvae. The advantage of collecting floral oils instead of nectar stems from the oil's much (8 times) higher energy content. The oil is also used for coating breeding cells to prevent water intrusion due to its hydrophobic behavior.

The oil is collected from the flowers by bristles at the front legs of the bees. It is then transferred from the front legs to closely spaced hairs at the bee's back legs, the so called scopa (cf. Fig. 1(b)), to allow for transportation of the oil to the larvae. This transport structure can be used repeatedly and shows a high efficiency of transported oil mass compared to the scopa's weight.

The scopa has a multiscale structure as its size ranges from millimeters to micrometers. It consists of several types of main hairs (cf. Fig. 1(c)) with a length of 2-3 mm and a diameter of approximately 40 μm . About 20 lateral hairs with lengths of 400 μm and diameters of 5-12 μm branch out from each main hair. In total, the scopa

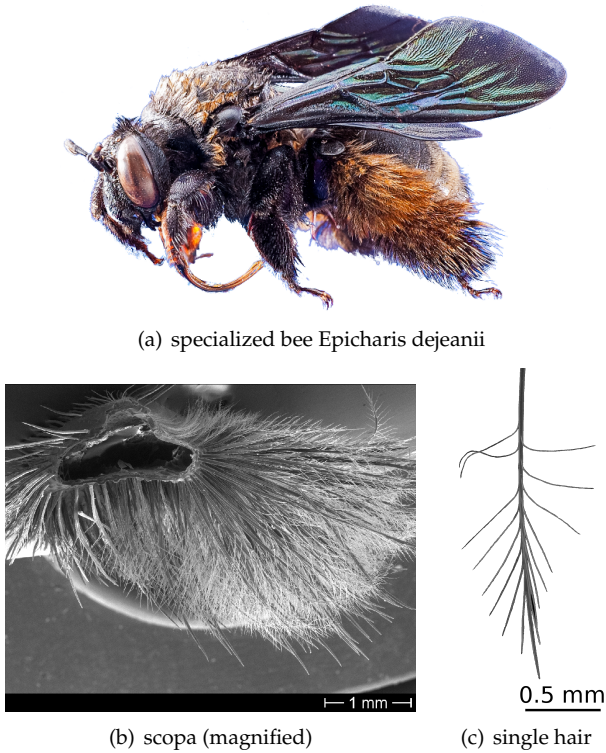
contains about 900 hairs. They form a three-dimensional grid structure in which the oil is stored. Furthermore, the scopa has an oleophilic surface which, beside the geometry, further facilitates oil collection and storage.

The properties of the multiscale microstructure of the scopa are of great practical interest. Man-made structures with an oil storage efficiency comparable to that of the biological role model are not known yet. One application is the oil separation of washing solutions. These solutions occur, for instance, in the process of car engine cleaning. The washing solution can only be given to the wastewater after the oil has been separated out. In practice, expensive steel coalescer are used for the separation process. Structures that increase the efficiency of the process and which are cheaper to produce are intensively searched for. The scopa of bees might be a role model for such future structures. Therefore, its precise properties need to be better understood. As the effects in the scopa take place on a microscopic length scale and are difficult to measure, simulations of the biological system are necessary. But to the authors' knowledge, there have been no simulations of the oil storage process in a bee's scopa so far. Instead of that, numerical simulations in literature primarily focus on the insect's wing design, see, e.g., Young et al. [3]. In the following, we present the first 3D two-phase flow simulations of an oil droplet within the microscopic scopa.

The remainder of this article is organized as follows: First, in Section 2, we describe the micro-CT measurement

*Corresponding author. Tel.: +49 228 733170; fax: +49 228 737527

Email addresses: ruetters@ins.uni-bonn.de (A. Rüttgers),
griebel@ins.uni-bonn.de (M. Griebel)



(a) specialized bee *Epicharis dejeanii*

(b) scopa (magnified)

(c) single hair

Figure 1: Photos were taken by the Institute of Crop Science and Resource Conservation (INRES).

of the scopa and the post processing of the measured data necessary to derive a geometric model of the scopa. We then consider the mathematical model for flow dynamics within the scopa. To this end, we state the governing equations and their spatial and temporal discretization used for the two-phase flow solver *NaSt3DGPF* in Section 3. Special focus is given to the boundary conditions on a microscopic length scale. We also explain our domain decomposition approach for parallel computations and discuss the parallel speed-up behavior for the considered problem. In Section 4 we present the simulation results for two specific situations: First, the laminar to turbulent air flow around the scopa's hairs is computed in Section 4.1. Second, in Section 4.2, the deformation of an oil droplet in the scopa geometry is simulated and the outcome is compared to experimental measurements of an oil filled scopa. Finally, in Section 5, we evaluate our findings and give some conclusions.

2. Processing of micro-CT data

The geometric complexity of the scopa is illustrated in Fig. 1(b). In principle, two different approaches are possible to allow for simulations. The structure can either be modeled in a tedious process by hand with a common computer-aided design (CAD) program or it can directly be measured in an experiment from a real scopa by, e.g., a CT scan. The advantage of a CAD modeled geometry is that the surface of the geometry is properly defined and

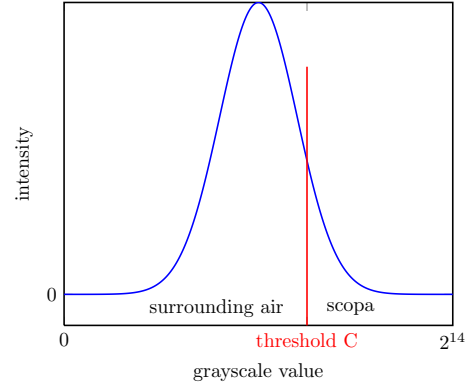


Figure 2: Histogram with typical distribution of the density values.

inaccuracies and artifacts due to measurement errors are avoided. On the other hand, a modeled geometry might not capture all relevant details of the scopa's structure. Nowak, Kakade and Annapragada [4], for instance, simulate the airflow and the particle deposition in two different geometries of the human lung. While the first geometry is an idealized model, the second geometry bases on a CT-scan and achieves more accurate results.

For this reason, we derive a geometric model from a CT microtomography measurement of the bee species *Epicharis dejeanii* that lives in Southern America. The measurement was performed at the Institute of Textile Technology and Process Engineering Denkendorf (ITV Denkendorf). To this end, the micro-CT *nanotom m* build by GE Sensing & Inspection Technologies GmbH was used. *Nanotom m* has been funded by the Federal Ministry of Education & Research (BMBF) in two projects for the analysis of fiber-reinforced composites and of surface coatings.

The micro-CT consists of an x-ray tube and a digital detector for image capturing. The x-ray tube allows long-time measurements for up to 8 hours. Focused on a small target of micrometer size, the x-rays have a power of up to 15 watt. The DXR detector has 3072×2400 pixels with a pixel size of $100 \mu\text{m}$. It allows for a scanning resolution up to an accuracy of $0.8 \mu\text{m}$. The sensor has a 14 bit color depth so that 16384 different gray-scale values are used to distinguish small density variations. A CT scan results in a large number of voxel cells which are represented and processed with the software VG-Studio MAX 2.2 by Volume Graphics.

A typical distribution from a micro-CT density measurement is shown in the histogram plot in Fig. 2. The density in each resolved grid cell can have values between 1 and $16384 = 2^{14}$. Thus, there is no natural distinction between solid cells from the scopa cells and the surrounding air cells which renders direct numerical simulations impossible. In the area of medical image segmentation, one viable approach to allow for this distinction is thresholding, see Pham, Xu and Prince [5] and the references cited therein. We here give the necessary informa-

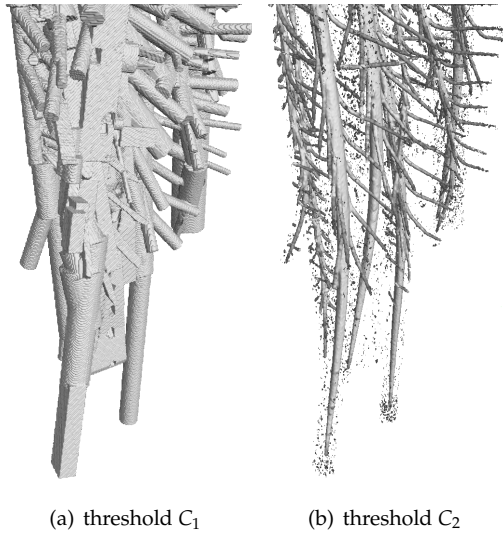


Figure 3: Scopa visualization for different threshold values.

tion for its application in a CFD solver. First, a threshold value C has to be chosen such that all voxel cells with a gray-scale value smaller than C are set to air cells. Consequently, cells with a gray-scale value larger than C are identified as a part of the scopa. Fig. 3 compares two different threshold values for the scopa data set. The threshold value C_1 in Fig. 3(a) was deliberately chosen too small so that a larger part of the surrounding volume was identified as a part of the scopa. The same data set is shown in Fig. 3(b) for a different threshold value $C_2 > C_1$. The corresponding geometry occupies less volume and agrees better with the biological data. Since most volume cells of the data set correspond to air, the threshold value is in practice chosen relatively large as indicated by the red vertical line in Fig. 2.

The CT measurement and a first post-processing of the scopa from the bee *Epicharis dejeanii* at the ITV Denkendorf required about ten days of work. In the end, the scopa was derived as shown in Fig. 4(a). Here, the physical dimensions are $7.1 \times 7.1 \times 5.5 \text{ mm}^3$ with a resolution of $3072 \times 3072 \times 2400$ voxel cells. These voxel data can not directly be used as input for our flow solver *NaSt3DGPf* due to its huge size and complexity. We thus restrict our simulations to a subdomain that contains 8 single hairs (cf. Fig. 4(b)). There, the grid resolution still contains $439 \times 633 \times 1289$ cells with a mesh width of $\Delta x = 2.3 \mu\text{m}$. Consequently, the diameter of a main hair is resolved with about 17 voxel cells. 4-5 voxel cells are used to model the basis of the lateral hairs which reduces to 2 voxel cells in the tips. The threshold value that is associated with the geometries in Fig. 4(a) and Fig. 4(b) is $C = 12400$. With this threshold, the geometry is fully captured but contains minor artifacts close to the surface. This artificial noise is illustrated in Fig. 5(a). The artifacts have to be removed before numerical simulations are feasible. To this end, we have used the following approach:

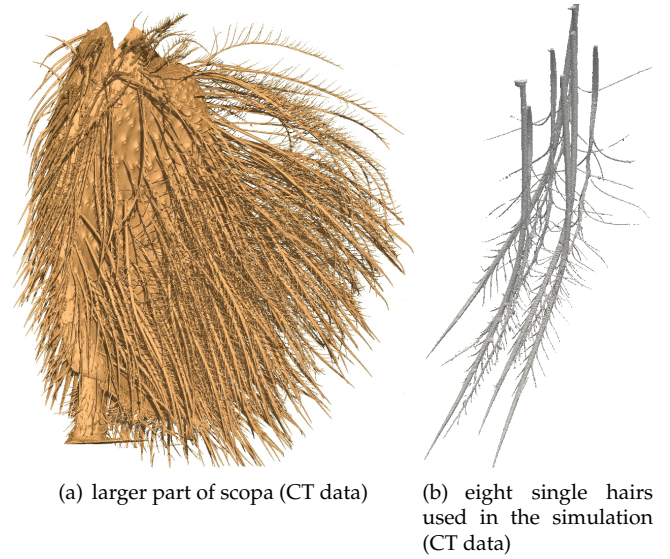


Figure 4: Scopa data for threshold $C = 12400$ from micro-CT measurement at ITV Denkendorf.

First, we assign all voxel cells that have been identified as a part of the geometry, i.e. with a density value larger than $C = 12400$, to their associated connected components. Then, if a connected component does not consist of at least two geometry cells in each coordinate direction, the complete connected component is flagged as liquid. Here, the minimum number of two grid cells is necessary to ensure that boundary conditions can be set correctly in our flow solver later on. This is a standard approach for flow solvers with a finite difference based Cartesian grid. This way, artifacts such as those colored blue in Fig. 6 are removed. Furthermore, we demand that the total number of cells in a component is larger than 11. This approach removes, for instance, the green structure in Fig. 6 if its extension to the third coordinate direction (not shown here) is smaller than three grid cells. We note that the stated number of 11 cells is not derived from theory but from a heuristic approach which delivers good results for this specific geometry. The resulting geometry is depicted in Fig. 5(b). Here, all artifacts have been removed but the actual shape of the scopa geometry is still conserved.

In Fig. 5 (b) we note that the hairs have a bumpy surface when the magnification is largest. This is an effect of the image segmentation with thresholding. A typical approach to reduce the bumpy surface is to apply a volume and topology preserving smoothing on the geometry, see, for instance, the book of Weickert [6] on anisotropic diffusion in image processing. The smoothed geometry can be used as basis for elaborated numerical methods to describe boundaries such as the immersed boundary method (IB) by Peskin [7], see also Mittal and Iaccarino [8] for a review of IB and cut-cell methods. On the other hand, a disadvantage of smoothing is that fine-grain details in the CT data might get lost.

In this work, we avoid a smoothing of the geometry




Figure 5: Smoothing of CT data is necessary to allow for numerical simulations.

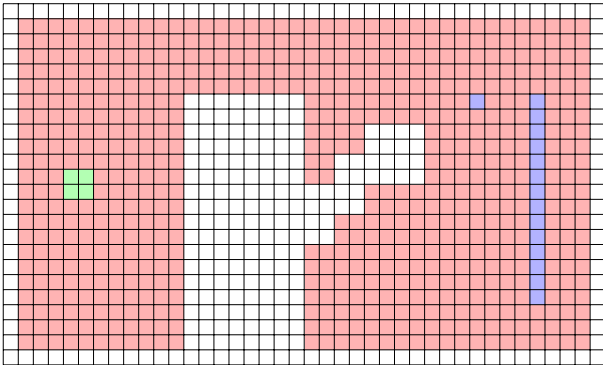


Figure 6: Illustration of flag field approach and artifacts, colored in blue and green, that have to be removed.

and directly perform simulations on the mesh obtained from thresholding. *NaSt3DGPF* uses a Cartesian grid, the so called flag field, to resolve geometries (cf. Fig. 6). In the present work, the flag field coincides with the CT generated voxel mesh. In general, a flag field approach delivers only a coarse geometry approximation. In this application, however, the CT dataset after thresholding does not contain any additional geometric information as those resolved with a Cartesian mesh. Consequently, the flag field approach used in *NaSt3DGPF* does not limit the ac-

curacy even though the biological role model involves a complicated multiscale geometry.

Although the scopa model in Fig. 5(b) shows high agreement to the biological role model, apart from the limitations due to thresholding, it cannot yet be used for numerical simulations with *NaSt3DGPF*. This becomes clear from a cut through the scopa as in Fig. 7(a). The scopa hairs are hollow. In nature, this decreases the weight of the scopa and leads to a better ratio of oil storage capability compared to the structure’s own weight. However, flow solvers with a flag field based approach as *NaSt3DGPF* do not distinguish between air cells inside the hollow hairs and from the surrounding. For this reason, the airflow within the scopa hairs would also be computed in the hollow inside of a hair which is physically meaningless but may cause numerical problems since, in this case, the matrix of the linear system of equations that has to be solved (cf. Section 3.4) can contain rows with zero entries only. To prevent this, we just fill up the interior part of each hair. This does not affect the simulation results but results in stability of the flow solver.

A simple approach for filling the interior part of a Cartesian grid geometry is the following: Let us assume that the eight scopa hairs are oriented in vertical direction as shown in Fig. 4(b). Then, air cells within a scopa hair are surrounded by geometry cells in all perpendicular directions. Consequently, a scopa cell lies within a neighborhood of m cells in these directions. We convert such cells into geometry cells. In practice, m is chosen comparatively small, i.e. $m = 3-5$, and the algorithm is applied several times. If m would be chosen too large, the space between two scopa hairs might accidentally be interpreted as an interior cell as well. This approach is sufficient for the subsequent applications. It is, to the best of our knowledge, not yet described in the literature. The resulting geometry is shown in Fig. 7(b). It will now be used for our numerical simulations.

3. Governing equations and numerical discretization

The diameters of the smallest hairs of the bee’s scopa have a size in the order of about $5 \mu\text{m}$. In this article, we still rely on the continuum assumption for the oil and air flow. The continuum assumption can fail when the mean free path of a molecule is in the order of the smallest significant length scales. This ratio is expressed by the dimensionless Knudsen number. According to Karniadakis, Beskok and Aluru [9], the length scales in our case lie somewhere between the continuous flow regime and the slip flow regime. In this regime, the Navier-Stokes equations are valid. However, fluid slip has to be taken into account which will be discussed in Section 3.3.

We first state the mathematical flow model. More details can be found for one-phase flows in Griebel, Dornseifer and Neunhoeffer [10] and for two-phase flows in Groß and Reusken [11]. Section 3.4 gives details on the




Figure 7: Comparison of original and filled scopa hair.

numerical discretization in *NaSt3DGPf*. The parallelization approach and speedup measurements are presented in Section 3.5.

3.1. Mathematical model for the flow system

We consider a rectangular domain $\Omega \subset \mathbb{R}^3$ that contains two immiscible incompressible fluid phases $\Omega_1(t)$, $\Omega_2(t)$ for which $\bar{\Omega} = \bar{\Omega}_1 \cup \bar{\Omega}_2$ and $\Omega_1 \cap \Omega_2 = \emptyset$. The

interface $\Gamma(t) = \partial\Omega_1(t) \cap \partial\Omega_2(t)$ separates both phases. The conservation of mass and momentum for the pressure $p = p(x, t)$ and the velocity field $\mathbf{u} = \mathbf{u}(x, t)$ in both subdomains Ω_i for $i = 1, 2$ are written as

$$\rho_i \left(\frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla \right) \mathbf{u} - \operatorname{div} \boldsymbol{\sigma}_i = \rho_i \mathbf{g} \quad \text{in } \Omega_i \times [0, T], \quad (1)$$

$$\operatorname{div} \mathbf{u} = 0 \quad \text{in } \Omega_i \times [0, T], \quad (2)$$

$$[\boldsymbol{\sigma} \mathbf{n}]_{\Gamma} = \tau \kappa \mathbf{n}, \quad [\mathbf{u}]_{\Gamma} = 0 \quad \text{at } \Gamma \times [0, T] \quad (3)$$

with appropriate initial and boundary conditions for \mathbf{u} and p . In equation (1), $\boldsymbol{\sigma}_i = -p\mathbf{I} + \mu_i \mathbf{D}(\mathbf{u})$ denotes the stress tensor and $\mathbf{D}(\mathbf{u}) = \nabla \mathbf{u} + (\nabla \mathbf{u})^T$ represents the deformation tensor. The viscosities μ_i , the corresponding densities ρ_i and the gravitation force \mathbf{g} are assumed to be constant. Furthermore, (3) describes the coupling at the interface Γ . Here, the surface tension leads to a jump of the normal stresses whereas \mathbf{u} is continuous over the interface.

The interface $\Gamma(t)$ is implicitly given as the zero contour of a signed distance function φ , the so called level-set function, i.e. $\Gamma(t) = \{x \in \Omega \mid \varphi(x, t) = 0\}$. The level-set function is used to indicate the phase-dependent material parameters, see Osher and Sethian [12] and Sussman, Smereka and Osher [13]. φ is known at the beginning and is then transported with the fluid field \mathbf{u} according to

$$\frac{\partial \varphi}{\partial t} + \mathbf{u} \cdot \nabla \varphi = 0 \quad \text{in } \Omega \times [0, T]. \quad (4)$$

The transport equation (4) does not necessary guarantee that $|\nabla \varphi| = 1$, i.e. that φ remains a distance function. In practice, (4) is augmented with a Hamilton-Jacobi type reinitialization equation to ensure $|\nabla \varphi| = 1$ close to the interface, for details see Sussman, Smereka and Osher [13].

Note that the equations (1) and (2) also cover one phase flows which will be performed in Section 4.1. In this case, we set $\rho = \rho_1 = \rho_2$ and $\mu = \mu_1 = \mu_2$ in Ω and $\Gamma = \emptyset$.

3.2. Turbulence model

For turbulence modeling of the one phase flow system in Section 4.1, we employ a Large Eddy Simulation (LES). There, in contrast to $\kappa\epsilon$ -models and $\kappa\omega$ models, the number of parameters is strongly reduced. As these parameters are in general not exactly known, the LES model is less prone to calibration errors. A comprehensive overview of Large Eddy models can be found in Sagaut [14].

The basic idea of LES models is to resolve the large scales of the flow field whereas the small scales are modeled. LES represents a very efficient technique for high Reynolds number flow. It has been applied, for instance, to simulate the turbulent flow at $Re = 2.1 \cdot 10^6$ around an airfoil, see Mary and Sagaut [15].

The LES model was first introduced by Smagorinsky [16] in 1963 and was later analyzed by Deardorff [17]. Smagorinsky assumed that the flow is in constant equilibrium with no accumulation of energy over time. In this approach, a space averaging filter is applied to the equations for conservation of momentum (1) and conservation of mass (2). This leads in the case of one-phase flows to the filtered equations

$$\rho \left(\frac{\partial}{\partial t} + \bar{\mathbf{u}} \cdot \nabla \right) \bar{\mathbf{u}} - \operatorname{div} \bar{\boldsymbol{\sigma}} + \operatorname{div} \boldsymbol{\tau} = \rho \mathbf{g} \quad \text{in } \Omega \times [0, T], \quad (5)$$

$$\operatorname{div} \bar{\mathbf{u}} = 0 \quad \text{in } \Omega \times [0, T] \quad (6)$$

with $\bar{\mathbf{u}}$ and $\bar{\boldsymbol{\sigma}} = -\bar{p}\mathbf{I} + \mu\bar{\mathbf{D}}(\bar{\mathbf{u}})$ as spatially filtered quantities. Compared with (1), the equation for the conservation of momentum (5) now contains an additional stress tensor $\boldsymbol{\tau}$. This tensor models stress on a sub-grid scale. Smagorinsky's approach [16] relates the sub-grid stresses $\boldsymbol{\tau}$ to the coarse scale deformation tensor $\bar{\mathbf{D}}(\bar{\mathbf{u}}) = \nabla \bar{\mathbf{u}} + (\nabla \bar{\mathbf{u}})^T$ according to

$$\boldsymbol{\tau} = -\nu_t \bar{\mathbf{D}}(\bar{\mathbf{u}}).$$

Here, the eddy viscosity ν_t , a factor of proportionality, is defined as

$$\nu_t = (C_s \Delta)^2 (2\bar{\mathbf{D}} : \bar{\mathbf{D}})^{1/2} \quad (7)$$

where C_s is the Smagorinsky constant and Δ is the local filter size. Subsequently, the filter size Δ used for a grid with constant mesh widths Δx , Δy and Δz in the three coordinate directions is

$$\Delta = \sqrt{\Delta x^2 + \Delta y^2 + \Delta z^2}.$$

However, other filter sizes would be possible as well.

In Section 4.1 we use a Smagorinsky constant $C_s = 0.0825$. Although C_s can be estimated from ab initio assumptions, in practice it often represents a degree of freedom to adjust the numerical results. To this end, different values of C_s have been proposed in the literature. For instance, Johnson [18] (pp. 14-88) states that values for C_s^2 in literature vary from approximately 0.005 to 0.05, i.e. C_s is in the range of 0.07 - 0.22. The lower values of C_s are normally chosen for channel flows while the larger values of C_s are used in situations with isotropic turbulence. An advantage of the Smagorinsky model stems from the fact that it dissipates the exact amount of energy provided that C_s is chosen correctly.

3.3. Boundary conditions

The adequate choice of the fluid boundary conditions at the solid-liquid interface is an open problem since the 19th century. Bernoulli first assumed that the fluid's movement at the interface is equivalent to the movement of the interface. These boundary conditions are known as no slip boundary conditions. No slip boundary conditions are normally prescribed in macroscopic flow fields.




Figure 8: Illustration of generalized Navier boundary conditions.

They have been validated in various experiments, for instance, for water flowing through glass pipes.

The situation is however different for microscopic flows. Here, several experiments have shown fluid slip at the contact line. An overview of the so-called moving contact line problem is given by Qian, Wang and Sheng [19]. Slip for liquids such as water often occurs at hydrophobic surfaces, and is also observed for the flow of non-Newtonian polymer solutions.

A problematic feature of the no slip model which cannot be ignored on a microscopic length scale is that the predicted shear rate at the interface is unlimited. This leads to a singularity at the interface. The occurrence of singularities gives one explanation for the poor predictions of the no slip model on this length scale.

The generalized Navier condition, proposed by Navier [20] in 1824, removes this singularity by weakening the no slip assumption. The model has been made popular by Huh and Scriven [21] in 1971. The quantity of slip is measured by the *slip length* l_s . This parameter can be derived by extending the velocity profile into the solid domain. The slip length then specifies the distance from the boundary to the zero value of the extrapolated velocity profile. This principle is illustrated in Fig. 8.

Following Shikhmurzaev [22], the fluid velocity \mathbf{u} close to the interface can be written as

$$\mathbf{u} = \beta \mathbf{n}^T \cdot \boldsymbol{\tau}_s \cdot (\mathbf{Id} - \mathbf{n} \otimes \mathbf{n}) \quad (8)$$

with \mathbf{n} the unit normal, β the coefficient of sliding friction or slip coefficient and $\boldsymbol{\tau}_s = -p\mathbf{Id} + \mu(\nabla \mathbf{u} + (\nabla \mathbf{u})^T)$ the stress tensor for a viscous Newtonian fluid. The Navier slip boundary condition is then derived as the tangential projection of (8) using $(\mathbf{Id} - \mathbf{n} \otimes \mathbf{n})$ where $(\mathbf{Id} - \mathbf{n} \otimes \mathbf{n})^2 = (\mathbf{Id} - \mathbf{n} \otimes \mathbf{n})$ and

$$\mathbf{u} (\mathbf{Id} - \mathbf{n} \otimes \mathbf{n}) = \beta \mathbf{n}^T \cdot \boldsymbol{\tau}_s \cdot (\mathbf{Id} - \mathbf{n} \otimes \mathbf{n}). \quad (9)$$

Then, we define $\mathbf{u}_{\text{slip}} = \mathbf{u} (\mathbf{Id} - \mathbf{n} \otimes \mathbf{n})$ as slip velocity. In the literature, equation (9) is often written in a simplified form in which \mathbf{n} is aligned with one of the coordinate directions. For instance, at the wall $y = 0$ the unit normal is $\mathbf{n} = (0, 1, 0)^T$ and for a generalized slip flow $\mathbf{u} = (u, 0, w)$ on this wall, as illustrated in Fig. 8, equation (9) becomes

$$\mathbf{u}_{\text{slip}} = \beta \mu \left(\frac{\partial u}{\partial y}, 0, \frac{\partial w}{\partial y} \right)^T. \quad (10)$$

The term $\beta\mu$ in (10) has a length unit and is exactly the slip length l_s that is shown in Fig. 8, i.e. $l_s = \beta\mu$.

As illustrated in Fig. 8, different boundary conditions correspond to different slip lengths l_s , i.e.

$$l_s \begin{cases} = 0, & \text{for the no slip BC,} \\ > 0, & \text{for the generalized slip BC,} \\ \rightarrow \infty, & \text{for the free slip BC.} \end{cases} \quad (11)$$

On a microscopic length scale, the slip length l_s is difficult to measure. In general, l_s is close to zero on a completely wetting surface and therefore the no slip assumption is valid. On the other hand, for non-wetting surfaces, such as for the flow of water in hydrophobic micro channels, slip lengths of several micrometers have been measured experimentally. Maali and Bhushan [23] give a review on techniques that are used to determine l_s on hydrophobic surfaces. In certain pressure-drop experiments slip lengths larger than $20\ \mu\text{m}$ were measured.

The exact slip length l_s for an oil flow in the scopa, which has an oleophilic surface, is not yet known. Due to the oleophilic surface of the scopa, we expect l_s to be much smaller than in the cases considered by Maali and Bhushan [23]. In Section 4.2 we compare the effect of different slip lengths on the simulation result. Furthermore, in Section 3.4 we discuss the discretization of the Navier-Slip boundary condition and show that the slip length is then linked with the mesh width h . Nevertheless, the approach allows us to get rid of the stress singularity and to obtain a realistic droplet behavior on the microscopic length scale.

For the boundary conditions of the level set function φ a suitable approach is required. Dirichlet boundary conditions could be used to enforce a static contact angle at the contact line but this might lead to unphysical results since the contact angle varies over time. Therefore, a dynamic contact angle model is required. This is still an area of active research, see Shikhmurzaev [22]. For our two-phase flow simulations we use, for reasons of simplicity, a more straight forward approach: Homogeneous Neumann boundary conditions are employed for φ whereas the wetting of the scopa is modeled with the slip length l_s . Then, the angle at the contact line is always 90° but the apparent macroscopic contact angle computed over several cells can still vary over time. Using this approach, an oleophilic behavior of the surface can be approximated quite well. Of course, this does not reproduce all features of a dynamic contact angle model but gives sufficiently good results.

3.4. Numerical discretization

NaSt3DGPF employs a finite difference (FD) scheme on a staggered grid for the spatial discretization of the flow equations (1)–(6) for one-phase turbulent flows and for two-phase flows with surface tension. Therefore, we subdivide Ω into rectangular cells and evaluate the unknowns p , φ in the cell centers and the velocity field \mathbf{u}

on the cell faces. This ensures a strong coupling between pressure and velocity field and avoids spurious modes. Apart from the convective terms in (1), (4) and (5), we discretize all spatial derivatives with second-order central differences. For the discretization of the convective terms, we use a 5th-order WENO scheme due to Jiang and Shu [24] to avoid oscillatory solutions. This high-order scheme requires the grid values from adjacent cells, which complicates the parallelization described in Section 3.5. For modeling surface tension effects, the continuum surface force method by Brackbill, Kothe and Zemach [25] is used. More details on *NaSt3DGPF*'s spatial discretization can be found in Croce, Griebel and Schweitzer [26] and in Griebel et al. [10].

The temporal discretization in *NaSt3DGPF* differs for single flows and multiphase flows. The filtered equations (5) and (6) are treated with a semi-implicit pressure correction scheme proposed by Bell, Colella and Glaz [27] which is of second-order accuracy in time. Here, the pressure field p and the diffusive velocity terms are treated implicitly. This leads to a diffusion equation for p with flow phase dependent coefficients and to three modified Helmholtz equations for an intermediate velocity \mathbf{u}^* in every time step. Here, the equation for p is solved with an AMG-preconditioned BiCGStab solver and a SSOR-preconditioned CG method is used for the Helmholtz system. Details on the implementation in *NaSt3DGPF* can be found in Verleye et al. [28]. Furthermore, details on the employed Algebraic Multigrid Method (AMG) are given in Metsch [29].

The two-phase flow system in (1)–(4) is solved implicitly in the pressure and explicitly in the velocity field. For this purpose, Chorin's projection method [30] is used to decouple \mathbf{u} and p . Again, the pressure diffusion equation is solved with an AMG-preconditioned BiCGStab solver. The velocity field is treated with an explicit 2nd-order Runge-Kutta method. In time, the transport equation for the level set function (4) is treated with a 3rd-order Runge-Kutta scheme. Due to the splitting of the velocity and the pressure field, the ansatz is of first-order accuracy in time.

The explicit treatment of the velocity terms leads to a Courant-Friedrichs-Lewy (CFL) condition to ensure numerical stability. This condition limits the size of a time step. It includes velocity contributions due to convective and diffusive velocity terms and due to volume and surfaces forces. An efficient parallelization strategy is therefore important to obtain results within a reasonable amount of time.

The numerical simulation of two-phase flows with the level set technique often suffers from numerical diffusion in the reinitialization step. One approach to circumvent volume loss are so-called volume correction methods. For all two-phase simulations in Section 4.2, we employ a global volume correction scheme that is described in Croce et al. [26] and in Croce [31]. Here, in each time step a new zero contour line of the level set function φ is

sought such that the resulting mass of both phases coincides with the initial mass. For this purpose, we define a Picard iteration which adds a constant c to φ and consequently shifts the zero contour line of φ . Let $l = 1, \dots$ denote the steps of our iterative procedure in each time step, ϵ a given error tolerance, V_0 the initial fluid volume and $V(\varphi_l)$ the fluid volume of φ at iteration l . We then perform for $l = 1, \dots$ the iteration

$$\begin{aligned} c &\leftarrow \omega (V_0 - V(\varphi_l)) \\ \varphi_{l+1} &\leftarrow \varphi_l + c \\ &\text{until } \frac{|V(\varphi_{l+1}) - V_0|}{V_0} < \epsilon. \end{aligned}$$

The parameter ω is used to convert the volumetric quantities into a distance quantity.

In Section 3.3 we discussed a generalized Navier boundary condition. In the following, we describe its discretization on a staggered grid and show how the boundary condition is linked with the mesh width. As an example, we consider the boundary condition (10) for the first velocity component at the wall $y = 0$. This can be straightforwardly generalized to the discretization of (8).

For simplicity, we assume a discretization with equidistant mesh width $h = \Delta x = \Delta y = \Delta z$. In an FD scheme on a staggered grid, the x -component of the velocity field is only computed at the grid points $u_{i,j,k} = u(ih, (j - 0.5)h, (k - 0.5)h)$ with $i = 1, \dots, i_{\max}$, $j = 1, \dots, j_{\max}$ and $k = 1, \dots, k_{\max}$. This is extended with a boundary strip that has the index positions $i = 0$ and $i = i_{\max} + 1$ and analogously for j and k to allow for the setting of boundary conditions. Due to the staggered grid, the x -component of the slip velocity has to be averaged from the nearest grid point $u_{i,1,k}$ in the flow domain and from the nearest grid point $u_{i,0,k}$ in the boundary strip (ghost cells). Using central differences, equation (10) is discretized with second-order accuracy in h as

$$\frac{u_{i,0,k} + u_{i,1,k}}{2} = l_s \frac{u_{i,1,k} - u_{i,0,k}}{h}. \quad (12)$$

Equation (12) is solved for $u_{i,0,k}$ and set as boundary condition in the ghost cells. We note from Fig. 8 that the magnitude of the slip velocity is between zero (no slip case) and the fluid velocity close to the boundary $\mathbf{u}_{\text{fluid}}$ (free slip case), i.e.

$$\|\mathbf{u}_{\text{slip}}\| = r \|\mathbf{u}_{\text{fluid}}\| \text{ with } r \in [0, 1]. \quad (13)$$

In this example, we have $\mathbf{u}_{\text{fluid}} \approx (u_{i,1,k}, v_{i,1,k}, w_{i,1,k})$ and by combining equation (12) and (13) we obtain a relation between l_s and the factor r for the percentage of slip velocity magnitude to fluid velocity magnitude as

$$l_s = \frac{r h}{2 - 2r}. \quad (14)$$

In Section 4.2 we consider, beside others, the case $r = 0.5$ in which the slip velocity is half the size of the fluid velocities. In this case, l_s has to be chosen as half of the mesh width h in *NaSt3DGPF*.




Figure 9: Illustration of the domain decomposition approach for parallelization in 2D.

3.5. Parallelization

The large amount of data obtained from the micro CT measurement and the underlying problem requirements demand an efficient parallelization strategy. *NaSt3DGPF* is parallelized by a domain decomposition approach. For this purpose, the flow domain is subdivided into subdomains, and each subdomain is assigned to a different processor. Processors with a common boundary face have to exchange data in every time step, which is illustrated in Fig. 9. Data exchange is performed with the Message Passing Interface (MPI) on parallel CPUs.

The subsequent results have been computed on the parallel CPU cluster *Atacama* at the University of Bonn. The cluster consists of 78 Dell PowerEdge M620 compute nodes with 1248 Intel Xeon CPU E5-2650 2.60GHz CPUs in total. Each compute node contains 16 CPUs and we always employ all processors per node in the simulations. The system has a main memory of 4992 GB. The MPI communication is conducted with 56 Gb/sec (4X FDR) Infiniband. Altogether, the system shows a Linpack performance of 20630 GFlop/s with a parallel efficiency of 80%. The cluster is operated by the Institute for Numerical Simulation and the Sonderforschungsbereich 1060 at University of Bonn.

Before we discuss our simulation results in Sections 4.1 and 4.2, we focus on the weak and strong scaling of *NaSt3DGPF*. The scaling behavior of a *NaSt3DGPF* porting to NVIDIA's CUDA framework to allow for GPU computing of two-phase flows has been investigated by Zaspel and Griebel [32]. Furthermore, the parallel scaling behavior of *NaSt3DGPF* for non-Newtonian one-phase flows was studied by Griebel and Rüttgers [33]. Since a geometry with a comparable complexity was not used for parallel scaling measurements of *NaSt3DGPF* before, we present our recent results in the following.

We investigate the weak scaling for an one-phase test problem. For this purpose, we consider three different subdomains l_1, l_2, l_3 of the eight single hairs shown in Fig. 4(b). These subdomains have grid resolutions of $219 \times 165 \times 219$ cells for l_1 , $439 \times 165 \times 417$ cells for l_2 and $439 \times 330 \times 833$ cells for l_3 . Consequently, the number of grid cells increases by a number of four from subdomain to subdomain. We visualize all three domains in Fig. 10.




Figure 10: Comparison of simulation domains for the scale-up measurement.

Table 1: Weak scaling/ scale-up results with respect to l_1 on the cluster *Atacama*.

subdomain	# cores	absolute time [s]	efficiency [%]
l_1	16	757	100
l_2	64	731	104
l_3	256	865	88

The cuboid colored in blue represents subdomain l_1 and the red and black cuboid show the simulation domains l_2 and l_3 .

As the size of the simulation domain l_i increases with the index i , the ratio of geometry to surrounding fluid cells differs in each case. The ratio of geometry cells to the total number of grid cells is roughly 0.016 for l_1 ($7.9 \cdot 10^6$ grid points containing 127.732 obstacle cells), 0.01 for l_2 ($3 \cdot 10^7$ grid points containing 296.240 obstacle cells) and 0.007 for l_3 ($1.2 \cdot 10^8$ grid points containing 840.349 obstacle cells). Therefore, only a small amount of grid cells in Fig. 10 is actually used to represent parts of the scopa geometry.

In Table 1 we list the scale-up results for our turbulent flow simulations on l_1 , l_2 and l_3 . The simulation itself will be discussed in detail in Section 4.1. Here, we concentrate on the aspect of parallelization. Table 1 lists the required physical time for the simulation of 40 time steps of the semi-implicit scheme described in Section 3.4. Ideally, the additional number of processors compensates the increase in the grid cell number.

We employ 16, 64 and 256 CPU cores for the simulations on the three subdomains. Against our expectations,

Table 2: Strong scaling/ speed-up results with respect to 64 CPUs on the cluster *Atacama* for subdomain l_3 .

#cores	absolute timing [s]	speedup	efficiency [%]
64	3391	1.0	100
128	1702	2.0	100
256	865	3.9	98
512	478	7.1	89

the absolute time decreases from l_1 to l_2 so that the theoretical scaling is 104%. This behavior can be explained by the different ratio of geometry cells to the total number of grid cells. Subdomain l_1 has the largest ratio of 0.016, which has two different effects: On the one hand, each grid cell that resolves the geometry and is not used to compute the air flow decreases the number of unknowns in the linear system of equations. Theoretically, this decreases the problem size. On the other hand, l_1 contains many more surface cells on which boundary conditions have to be set. This deteriorates the condition number of the corresponding matrix so that more iterations of the iterative solver are necessary. In practice, the second effect seems to be dominant which leads to the lower computing time for subdomain l_2 . For l_3 we observe a decrease to 88% parallel scale-up which is still comparatively high. In this case, the decrease in parallel efficiency results from the communication overhead. Note that the high efficiency which is observed for l_3 might slightly decrease, if the lower geometry cell ratio for l_3 is taken into account. Nevertheless, we are able to reproduce the good weak scaling results from Zaspel and Griebel [32] and from Griebel and Rüttgers [33] for the considered complex geometry.

Next, we investigate the strong scaling of *NaSt3D-GPF* for the eight hair geometry. We list the results in Table 2. In this case, we compare the absolute time for the simulation of 40 time steps of a turbulent air flow described in Section 4.1 with a number of processors which ranges from 64 to 512. Note that the minimum number of processors used for the strong scaling performance measurement is 64 since the large number of grid cells requires a minimum amount of compute nodes due to memory requirements. In each case, we consider subdomain l_3 shown in Fig. 10 with a mesh that consists of $439 \times 330 \times 833$ cells. The speed-up results are nearly optimal with an efficiency of about 90% or better. In contrast to the weak scaling results, the simulation domain is now identical for the different number of processors so that the complex geometry does not affect the strong scaling results. We conclude that *NaSt3DGPF* is perfectly adapted to perform massively parallel computations even in a geometry as complex and irregular as the bee's scopa.

4. Numerical results

Now, we discuss the simulation results computed with *NaSt3DGPF*. In nature, the whole process of oil storage and transport is very complex due to the involved scales. The process of oil storage in the scopa geometry takes fractions of seconds up to seconds depending on whether the complete scopa is filled with oil or only a small part of it. The process of oil transport requires time in the order of minutes as several flowers are searched for floral oil, and the oil has to be transported to the breeding cells. We therefore subdivide the simulation of this process into two parts:

- First, we discuss simulation results of a laminar to turbulent air flow around the scopa hairs that occurs during the flight process. In this case, we restrict ourselves to the simulation of a semi turbulent one-phase flow. This is due to the fact that the involved time scale is relatively large and higher flow velocities occur. We therefore require a semi-implicit scheme that allows for larger time steps. This is available in *NaSt3DGPF* for the one-phase flow solver. Furthermore, the matrix has a better condition in the one-phase flow case which further reduces the required amount of CPU time. The simulation will reveal regions of strong vortex activity at which the floral oil is most likely to get lost. For this reason, the simulation is of high practical relevance for the considered problem.
- Second, we simulate the oil storage process within the scopa. For this purpose, the two-phase flow solver with surface tension is used. As this problem possesses a much larger mathematical complexity, the mesh resolution is lower in this case compared to the one-phase problem. Here, we will compare our simulation results with that of a CT measurement of the oil filled scopa.

4.1. Simulation of turbulent air flow around the scopa hairs

The first application of *NaSt3DGPF* for the CT measured geometry is the simulation of a turbulent air flow around the scopa hairs. The corresponding situation in nature is the flight process of the bee species *Epicharis dejeanii* whose typical flying speed is used here. Due to the complexity of the eight scopa hairs for a resolution with $439 \times 633 \times 1289$ cells, see Fig. 4(b), we concentrate ourselves here to the central part of the geometry. First, we investigate in a model problem if such a restriction has influence on the simulation results, i.e. if such a simplification can be applied without changing the outcome too much. In a second step, we present the simulation results that are adapted to the biological prototype. In contrast to the simulations in Section 4.2, experimental comparisons are not available here. However, we refer to Griebel

and Rüttgers [33] for a validation of *NaSt3DGPF* for a one-phase flow problem for which an analytical solution exists.

In the first test case, we analyze three different but overlapping subdomains of the scopa hair geometry. To this end, we reuse the simulation domains that were employed in the weak-scaling analysis of *NaSt3DGPF* in Section 3.5, see Fig. 10. The three nested subdomains have a resolution of $l_1 = 219 \times 165 \times 219$ cells, $l_2 = 439 \times 165 \times 417$ cells and $l_3 = 439 \times 330 \times 833$ cells. The physical size of the subdomains is $0.5 \text{ mm} \times 0.39 \text{ mm} \times 0.5 \text{ mm}$ for l_1 , $1.0 \text{ mm} \times 0.39 \text{ mm} \times 0.97 \text{ mm}$ for l_2 and $1.0 \text{ mm} \times 0.77 \text{ mm} \times 1.9 \text{ mm}$ for l_3 . The basic idea is to perform a turbulent flow simulation for each subdomain with the same inflow condition and to compare the result on the common domain of all three meshes. All three meshes contain the part of the flow domain that is resolved on l_1 . In Fig. 10 the common domain of all three simulations is indicated in the center of the light blue cuboid together with a slice on which we depict our three different solutions.


The typical flying speed of a bee is $\tilde{u} \approx 24 \text{ km/h} = 6\frac{2}{3} \text{ m/s}$. We prescribe an air flow with $\mathbf{u} = (\tilde{u}, 0, 0)$ on the inflow face in x -direction. As the channel length in x -direction is $l_m \approx 1 \text{ mm}$ for l_2 and l_3 , we obtain a Reynolds number $Re = (l_m \tilde{u} \rho) / \mu \approx 500$ for air with density $\rho = 1.29 \text{ kg/m}^3$ and viscosity $\mu = 1.72 \cdot 10^{-5} \text{ Pa}\cdot\text{s}$. This Reynolds number lies in the transition regime between laminar and turbulent flow behavior.

Next, the nondimensionalization of all units is performed by a scaling with reference units. To this end, we employ the channel length in x -direction $l_m \approx 1 \text{ mm}$ as a characteristic length scale and use $u_m = 1 \text{ m/s}$ as reference velocity. Then, we obtain $t_m = l_m / u_m = 1 \text{ ms}$ as characteristic time scale and $6\frac{2}{3}$ as inflow velocity.


For the simulation of turbulent flow, we use our Large Eddy turbulence model with $C_s = 0.0825$. This value lies in the range of 0.07 - 0.22 as stated in Section 3.2. A comparatively low value for C_s was chosen since the turbulent viscosity is oriented in x -direction. On the outflow face in x -direction, we prescribe homogeneous Neumann boundary conditions while we set free slip boundary conditions on the lateral four faces. Furthermore, we employ no slip boundary conditions on the surface of the hairs.

The three simulations were conducted up to a physical time of $t/t_m = 1$ for all subdomains l_1, l_2, l_3 . In Fig. 11 we compare the outcome of the turbulence model on a slice with normal $\mathbf{n} = (0, 0, 1)$ at an l_m -normalized height of 0.72 in z -direction. This slice goes through the center of the light blue cuboid in Fig. 10. To allow for a better comparison in Fig. 11, we restrict the x/y -planes for l_2 and l_3 so that they coincide with l_1 . Consequently, the x/y -cut for l_1 in Fig. 11(a) has 219×165 cells, Fig. 11(b) contains additional cells in x -direction and the mesh of Fig. 11(c) is extended in the directions x and y .


In general, the results of the three simulations for the turbulent viscosity ν_t as defined in (7) agree very well.



(a) $l_1 = 219 \times 165 \times 219$ cells



(b) $l_2 = 439 \times 165 \times 417$ cells (shown is an x/y -extract with 219×165 cells)



(c) $l_3 = 439 \times 330 \times 833$ cells (shown is an x/y -extract with 219×165 cells)

Figure 11: Comparison of the turbulent viscosity v_t at $z/l_m = 0.72$ for three simulations in which different parts of the scopa have been resolved. The complete simulation domains and the position of the slice are shown in Fig. 10.

The turbulent viscosity is largest at the side and behind the scopa hairs in the flow direction. Note that the scopa in Fig. 11 is colored white. An exact analysis, in which the absolute difference of the results is measured, shows that the peak values of the turbulent viscosity differ by


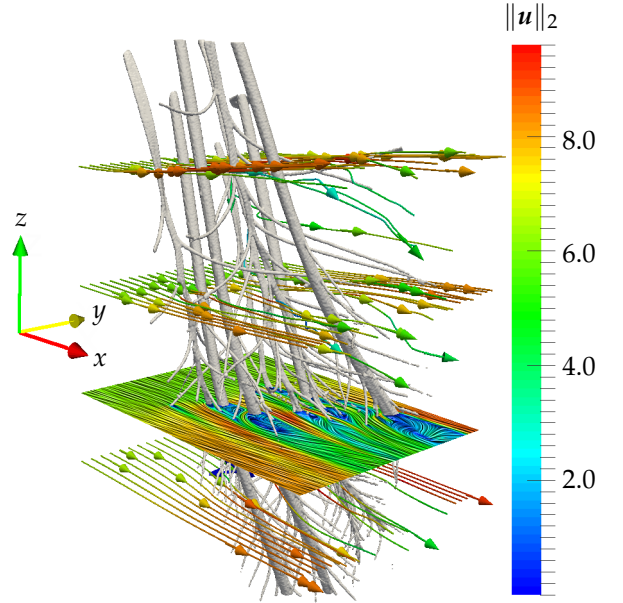


Figure 12: LIC streamline visualization of the airflow around the scopa hairs. In the top figure, the flow field is shown on a slice at $z/l_m = 0.74$ and at different streamlines with initial heights $z/l_m = 0.5$, $z/l_m = 1.1$ and $z/l_m = 1.47$. In the bottom figure, the velocity field is shown on the same slice from a different perspective.

at most 10%. Another difference is visible around the lowest scopa hair in y -direction in Fig. 11(c). The simulations l_1 and l_2 predict a much lower turbulent viscosity at this single lateral hair. This is due to the free slip boundary at the bottom side in these two simulations so that the velocity gradient is very low there. As the free slip boundary for l_3 is further away and is not shown here, higher flow variations occur at the lateral hair and lead to a larger turbulent viscosity. In conclusion, a subdomain of the scopa shows flow structures similar to the full geometry as long as boundary interference is avoided. We can therefore consider interesting subdomains of the full geometry shown in Fig. 4(b). This, however, still leads to an enormous computational effort.

We now focus on the largest domain l_3 with a resolution of $439 \times 330 \times 833$ cells for a detailed analysis of the airflow within the bee's scopa. The simulation was

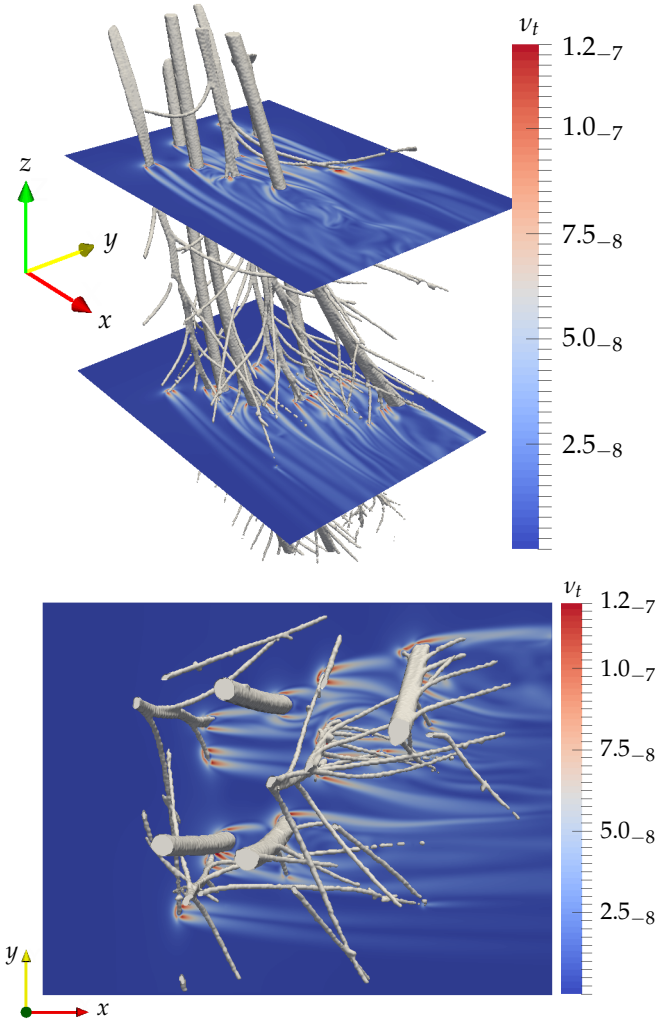


Figure 13: Turbulent viscosity v_t visualized on planes through the scopula hairs with height $z/l_m = 0.74$ and $z/l_m = 1.47$. The figure at the bottom shows the slice at $z/l_m = 0.74$ from a different perspective.

conducted up to a physical time of $t/t_m = 1$. The computation of the described problem ran for 160 hours on 256 CPUs of the HPC cluster *Atacama*. The simulation required 25.000 time steps with an average time step width of $\Delta t/t_m = 4 \cdot 10^{-5}$.

Fig. 12 shows the laminar to turbulent flow around the hairs of the scopula. At the top, we display several streamlines and visualize them on a slice at $z/l_m = 0.74$ in z -direction with a Line Integral Convolution (LIC) filter. At the bottom of Fig. 12, we show the flow field on the same slice from a different perspective. Several minor vortices form directly behind the hairs. The vortices partially overlap due to the proximity of the various side hairs. However, sufficiently far away from the hairs there is laminar flow as expected for $Re \approx 500$. For comparison, Fig. 13 shows the turbulent viscosity v_t defined in (7) on slices at positions $z/l_m = 0.74$ and $z/l_m = 1.47$. Again, the turbulent viscosity is largest in the vortex regions directly behind the scopula in the flow direction. Because of

the accurate resolution of the hairs from the CT measurement, both figures now show small scale structures that would not have been resolved on a coarser mesh. Note that the CT model allows for a resolution of the hair's geometry which indeed captures all relevant flow phenomena for this problem.

4.2. Simulation of the oil storage process within the scopula

The second problem which we consider is the process of oil absorption within the scopula. In the following, we discuss our simulation results for a microscopic oil droplet that spreads between the branched hairs. In a first step, the mesh independence of the simulation is investigated, i.e. we show that our CT-adapted mesh is sufficiently fine to resolve all quantities in the flow domain. We then investigate the effect of the slip length l_s that has been introduced in Section 3.3. Finally, we discuss the results of our simulations of the physical process and compare them with a direct CT-measurement of an oil-filled scopula. Note that *NaSt3DGPf* has been recently used to establish a three-dimensional two-phase benchmark configuration. In this benchmark, the outcome of *NaSt3DGPf* has been compared with the results of two other flow solvers *DROPS* and *openFOAM*, see [34]. This indicates the suitability of *NaSt3DGPf* for this problem.

Again, our mesh widths are prescribed by the micro CT scan resolution $\Delta x = 2.3 \mu\text{m}$ which depicts the finest grid resolutions in our numerical simulation. Nevertheless, we still have to ensure that the resolution of the flow domain is sufficient to simulate the oil absorption process. For this purpose, we employ the following approach: We coarsen the scopula mesh two times which results in meshes with a resolution of $\Delta x = 3.1 \mu\text{m}$ and $\Delta x = 4.6 \mu\text{m}$. Then, we investigate the oil droplet deformation on these three different meshes and show that the steady state of the droplet is similar in each case. This shows that the employed finest mesh is indeed sufficient.

The scopula subdomain that we use for this grid convergence study is shown on the right hand side of Fig. 14. In this case, the grid has a resolution of $160 \times 160 \times 350$ cells and a physical size of $0.37 \text{ mm} \times 0.37 \text{ mm} \times 0.82 \text{ mm}$. The coarsening of the scopula geometry, i.e. of the three dimensional flag field, is performed with a cubic interpolation to the reduced mesh. In the flag field description a value of zero is used to denote a fluid cell and a value of one is used to denote a cell from the scopula. As the interpolated flag field contains intermediate floating point values, flag cells with a value below 0.5 are rounded to fluid cells. Otherwise, they are rounded to geometry cells, i.e. they are rounded up to 1.0. This coarsening procedure can lead to an insufficient number of geometry cells for the description of the hairs on a coarser mesh as already discussed in Section 2 and as illustrated in Fig. 6. Here, lateral hairs of the scopula are erased that do not consist of sufficiently many cells. Consequently, the side hairs vanish on coarser grids. This is avoided by converting

Table 3: Material properties for the oil storage simulations.

ρ_{oil}	μ_{oil}	ρ_{air}	μ_{air}	τ
910	0.1	1.2	$1.8 \cdot 10^{-5}$	$2.8 \cdot 10^{-2}$
[kg/m ³]	[kg/m s]	[kg/m ³]	[kg/m s]	[kg/s ²]

air cells into geometry cells if such a situation is encountered. On the other hand, this leads to an increase in size of the lateral hairs compared to the main hairs. The coarsened meshes with resolutions of $120 \times 120 \times 262$ cells and $80 \times 80 \times 175$ cells are shown on the left hand side and in the middle of Fig. 14. The effect of the increased lateral hair size can be seen in mesh l_1 on the left hand side of Fig. 14.

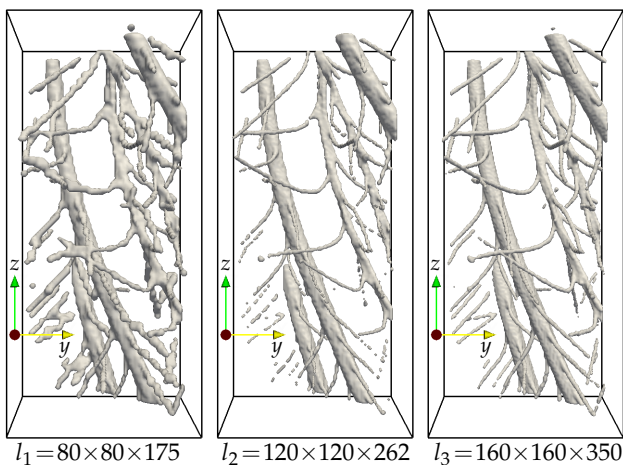


Figure 14: Comparison of the meshes for the convergence study. The CT measured mesh is shown on the right hand side. The results of the two coarsening steps are displayed on the left hand side and in the middle.

For all three meshes, we simulate the deformation of a microscopic oil droplet with largest extension of 0.56 mm in z -direction. The material properties for the two-phase system of oil and surrounding air are listed in Table 3. Note that the simulations in this section, in contrast to those in Section 4.1, now employ dimensional units. We furthermore prescribe the gravitational force $\mathbf{g} = (0, 0, -9.81) \text{ m/s}^2$ in z -direction. We use this parameters for all simulations in this section. At the beginning of the simulation, the droplet has an ellipsoidal shape and a small initial velocity $(0, 0, -0.5) \text{ m/s}$ in vertical direction. In all three simulations the droplets have the same surface area but, as the scopa geometry that is surrounded by the droplet is slightly different due to the coarsening procedure, the droplets' masses differ slightly. The initial position of the oil droplet and its steady state at about $t = 0.3 \text{ ms}$ are shown for all three simulations in Fig. 15. Since the initial velocity was deliberately chosen very small, only a slight movement in vertical direction occurred. We observe that the droplet evolves along the scopa hairs. As shown in Fig. 15 the steady state result is similar in all simulations. This indicates that the employed mesh is sufficiently fine.

Next, we investigate the influence of the slip length

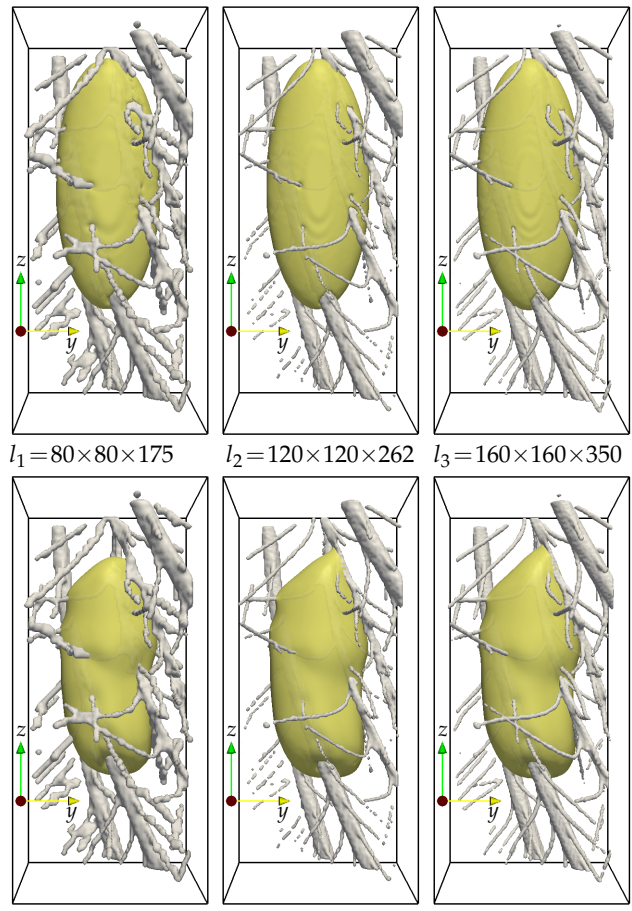


Figure 15: Comparison of the initial and the steady state at $t = 0.3 \text{ ms}$ of an oil droplet for different mesh accuracies.

l_s on the contact line, see Section 3.3. For this purpose, we perform four simulations of a deforming oil droplet that coincide in all parameters but for the slip length l_s . As before, the material parameters are listed in Table 3. Furthermore, we reuse the same scopa subdomain with a physical size of $0.37 \text{ mm} \times 0.37 \text{ mm} \times 0.82 \text{ mm}$ and a mesh resolution of $160 \times 160 \times 350$ grid cells. In this case, a smaller droplet with an extension of 0.26 mm in z -direction and an initial velocity $(0, 0, -0.5) \text{ m/s}$ has been simulated since only the contact line at one lateral hair is investigated. This subdomain, however, is representative for the other parts of the scopa. Using relation (14) the slip length l_s is chosen such that u_{slip} is zero for the no slip boundary condition and it is 10%, 50% and 90% of the fluid velocity \mathbf{u} close to the boundary. The corresponding slip lengths for the generalized slip boundary conditions are $0.13 \mu\text{m}$, $1.15 \mu\text{m}$ and $10.35 \mu\text{m}$. We avoid free slip boundary conditions as they do not take interactions between fluid and scopa into account.

We show the droplet's position at steady state in Fig. 16. In the first row, the droplet is shown for the no slip boundary condition and for a slip velocity condition $u_{\text{slip}} = 0.9 \mathbf{u}$ with the surrounding geometry. Then, both results are compared on a zoomed extract that primarily

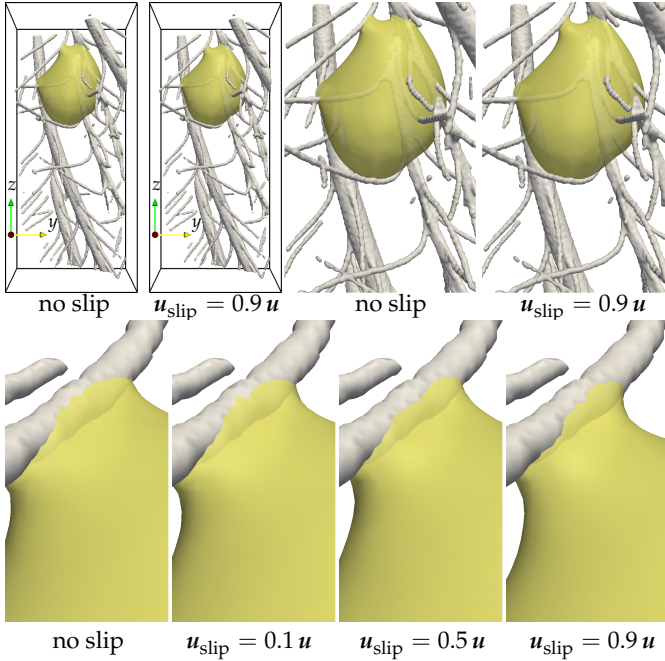


Figure 16: Influence of the slip length l_s on the fluid/geometry interaction.

focuses on the oil droplet. On this length scale both results closely agree with each other. The influence of l_s becomes only visible at the contact line that is shown in the second row of Fig. 16. Depending on the slip velocity, the wetting of the scopula is different. It is largest for the no slip boundary condition. The more the slip velocity is increased, the less of the structures is wetted. Note that the wetting of the lateral scopula hair disappears completely if free slip boundary conditions are employed. Consequently, the largest difference in the wetting behavior occurs between $u_{\text{slip}} = 0.9 u$ and the free slip boundary condition.

It is known from experimental measurements that the scopula has an oleophilic surface so that large parts of the scopula are wetted with oil. As mentioned before, fluid slip is observed in experiments on a microscopic length scale, see Karniadakis, Beskok and Aluru [9]. In Section 3.3 we roughly estimated the slip length l_s as lower or equal to $1 \mu\text{m}$. A slip length $l_s = 1 \mu\text{m}$ approximately corresponds to a slip velocity $u_{\text{slip}} = 0.5 u$ and we use this boundary condition for the following simulations. We note that there is no profound verification for this choice. As a result, the subsequent simulation has an uncertainty with respect to the wetting of the scopula at the contact line. This uncertainty is shown in the second row of Fig. 16. Since the influence of r from (13) in the range $[0, 0.9]$ is hardly noticeable from a more macroscopic perspective, see first row of Fig. 16, we will discuss the influence of l_s on the following results only when the contact line is taken into account.

Using the parameters obtained from the previous test studies, we simulate the deformation of an oil droplet

on the scopula's hairs and compare the outcome with experimental measurements. The simulation has a physical size of $0.46 \text{ mm} \times 0.77 \text{ mm} \times 1.1 \text{ mm}$ and a resolution of $200 \times 330 \times 389$ grid cells. This subdomain is smaller than the simulation domain with a size of $1.0 \text{ mm} \times 0.77 \text{ mm} \times 1.9 \text{ mm}$ that was considered in Section 4.1. Nevertheless, the current simulation is the most demanding with respect to computing time. The simulation required about 4 weeks on 256 CPUs on *Atacama*, see Section 3.5, to reach its steady state at about $t = 1.0 \text{ ms}$. The number of time steps was 56.000 and the typical time step width was $\Delta t \approx 1.8 \cdot 10^{-5} \text{ ms}$. This is a much more severe restriction than for the one phase flow simulations in Section 4.1 since the involved fluid velocities are lower by a factor of 10.

As before, we place an ellipsoidal shaped droplet in the upper part of the simulation domain and investigate its deformation over time. The droplet's axes have lengths of 0.28 mm , 0.51 mm and 0.26 mm in the x , y and z coordinate directions, respectively. We list the material parameters in Table 3. Furthermore, we prescribe a small initial velocity of the droplet of $(0, 0, -0.4) \text{ m/s}$ in z -direction. In reality, a much larger percentage of the scopula is wetted. In such a configuration, however, it would be more difficult to recognize the oil distribution along the scopula hairs which is our primary quantity of interest. The oil absorption takes place due to capillary motion that is caused by a complex interplay of surface tension and adhesive forces. Beside these effects, another driving force of the wetting process is the initial velocity of the oil droplet.

The importance of surface forces on this length scale is emphasized by the Bond number Bo . This dimensional unit measures the importance of surface tension forces compared to body forces and is defined as

$$Bo = \frac{\rho_{\text{oil}} g l^2}{\tau}$$

with g as gravitational acceleration and l as a characteristic length scale; the other material parameters are listed in Table 3. We employ the maximum droplet length of 0.51 mm in y -direction as characteristic unit and obtain $Bo \approx 0.08$ as Bond number. In general, scenarios with Bo smaller than 1 are considered as dominated by surface tension.

Fig. 17(a) shows the initial state of the simulation with the oil droplet in the upper part of the simulation domain. The droplet then moves downwards and deforms over time in such a way that its surface area is minimized. Fig. 17(b) and Fig. 17(c) show intermediate states of the simulation at $t = 0.25 \text{ ms}$ and $t = 0.5 \text{ ms}$, respectively. We display the droplet close to its steady state at $t = 1.0 \text{ ms}$ in Fig. 17(d).

In Fig. 18 we visualize the magnitude $\|\mathbf{u}\|_2$ of the velocity field and the zero contour of the level set function φ (indicated with a solid white line) on a slice with unit normal $(1, 0, 0)$ through half of the flow domain in

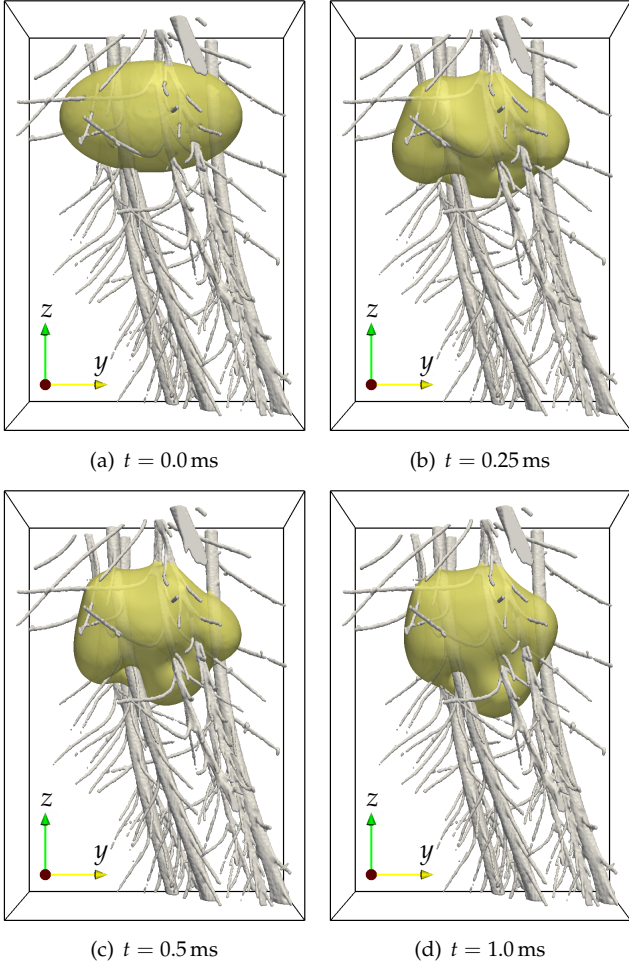


Figure 17: Oil droplet within the scopula at four different positions in time.

x -direction at $x = 0.23$ mm. In the initial state of the system the air phase is at rest and the oil phase possesses an initial velocity of $(0, 0, -0.4)$ m/s. The simulation shows that the upper part of the droplet reaches its steady state quite soon at $t \approx 0.4$ ms. The bottom side of the droplet deforms for a longer period of time until it reaches its steady state at about $t = 1.0$ ms. At this point in time, the velocity magnitude $\|\mathbf{u}\|_2$ is below 0.1 m/s on the lower droplet side and close to zero nearly everywhere. Therefore, only minor changes of the droplet occur afterwards.

The droplet's average velocity component w over time is shown in Fig. 19. Here, using $\Omega_1(t)$ to denote the oil phase at time t , we calculate the medium velocity $\bar{\mathbf{u}} = (u, v, w)$ of the droplet as

$$\bar{\mathbf{u}}(t) = |\Omega_1(t)|^{-1} \int_{\Omega_1(t)} \mathbf{u} \, dx.$$

At about $t \approx 0.12$ ms, its velocity component w has decreased by a factor of 0.5 compared to the initial velocity. In the final stage of the simulation, only a small part of the droplet deforms while most of it has already reached a steady state. Then, the average velocity component w is

close to zero at about $t \approx 1.0$ ms.

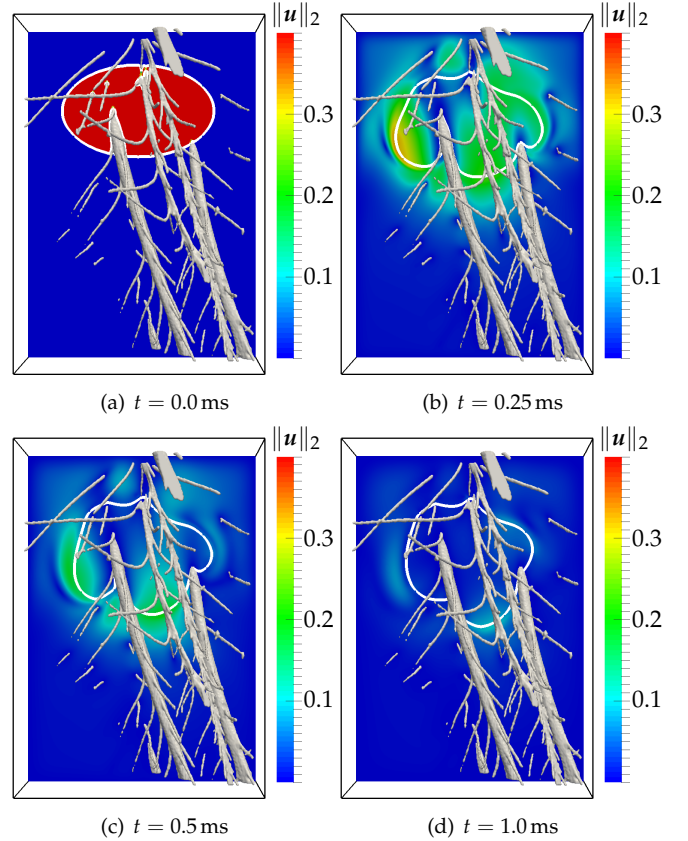


Figure 18: Magnitude $\|\mathbf{u}\|_2$ of the velocity field and free surface, indicated with a solid white line, on a central slice at four different positions in time.




Figure 19: Microscopic oil droplet's velocity component w over time.

We investigate a further quantity of interest, the sphericity $\Psi \in [0, 1]$ introduced by Wadell [35], which measures how much the shape of the droplet differs from a sphere. This quantity is defined as

$$\Psi(t) = |\Gamma(t)|^{-1} \pi^{1/3} (6|\Omega_1(t)|)^{2/3}.$$

Here, the surface area of a sphere with volume $|\Omega_1(t)|$ is divided by the surface area $|\Gamma(t)| := \text{meas}_2 \Gamma(t) = \int_{\Gamma(t)} 1 \, ds$ of the actual droplet. A perfect spherical shape corresponds to $\Psi = 1$ whereas Ψ becomes smaller the

more the droplet is deformed. We plot Ψ over time in Fig. 20. At the beginning of the simulation, Ψ is about 0.73 due to the ellipsoidal initial shape of the droplet. Since the Bond number Bo is comparatively large, the surface tension force leads to an increase of Ψ over time. This increase, however, is not monotone. From $t \approx 0.1$ ms to $t \approx 0.3$ ms, Ψ slightly decreases from 0.84 to 0.825 until it increases again to a final value of about 0.9. The decrease is partially caused by the movement of the droplet in vertical direction. Due to this movement, the droplet is flattened and extends in the perpendicular x/y -plane. This effect can be seen by comparing Fig. 17(b) and Fig. 17(c). At $t = 0.25$ ms the droplet is more stretched in the directions x and y than later on at $t = 0.5$ ms. In the end however, the droplet's movement comes to an end and the surface tension enforces a spherical shape only disturbed by the scopa geometry.




Figure 20: Plot of the sphericity Ψ over time.

The primary quantity of interest is the oil distribution along the scopa hairs. For this purpose, we examine the upper part of the droplet in detail. Fig. 21 shows a zoomed view from two different perspectives. The oil is spread along the lateral scopa hairs. Furthermore, the oil droplet has a concave shape in the spacing between two scopa hairs. In Fig. 16 we have analyzed the influence of the slip length l_s on a test simulation. In contrast to the results in Fig. 17 to Fig. 20, the choice of l_s will affect the wetting that is shown here in detail. From the previous results we expect a slightly larger wetting of the lateral hairs with oil if l_s is decreased.

Next, we compare the simulated oil distribution shown in Fig. 21 with the process as it takes place in nature. In Fig. 22 a scopa wetted with fatty floral oil is shown in a microscopic photography. The oil encloses a large part of the structure. Unfortunately, a direct comparison of Fig. 22 with the experiment is not possible as the magnification is not sufficient to closely analyze the contact line. For this reason, a CT data measurement of an oil wetted scopa from the species *Epicharis dejeanii* was performed at ITV Denkendorf. First, a large oil droplet was placed in the center of the scopa. The oil then spread in all directions along the scopa. After reaching its steady state, a measurement was taken with the micro-CT *nanotom m*. Since the measurement was primarily performed to allow for comparisons with the simula-



Figure 21: Two different perspectives on the upper part of the oil droplet at steady state.

tions, the scanning resolution was lower than for the CT model used for the simulations. The scanning resolution used here was $\Delta x \approx 5.7 \mu\text{m}$, the number of voxel cells was $1169 \times 1283 \times 2242$ and the resulting size of the measurement was $6.75 \times 7.41 \times 12.95 \text{ mm}^3$. Consequently, the mesh obtained in this case is coarser than the mesh described in Section 2 which is the basis for the numerical simulations.

Table 4: Wetting of lateral hairs with floral oil in the CT scan and in the simulation.

	CT data	simulation
wetting height	7 – 20 μm	5 – 12 μm
number of investigated hairs	10	7


The full CT-measurement data of the oil filled scopa is shown in Fig. 23(a). Here, the oil is colored in yellow analogously to the coloring used for the simulations. Due to




Figure 22: Floral oil within bee scopa. Photo was taken at the Institute of Crop Science and Resource Conservation (INRES).

the setup for the micro-CT, the arrangement in Fig. 23(a) is upside down. Fig. 23(b) and Fig. 23(c) zoom into the boundary layer of the spread oil. In Fig. 23(c) a zoomed detail is shown that is comparable in size to the simulation domain in Fig. 21. In both cases, a similar oil distribution along the scopa's hairs and a concave shape of the oil's surface can be observed. This qualitative agreement shows that our simulation is able to capture the relevant features of the system at steady state.


For a quantitative comparison of the steady state results, we have measured the oil distribution along the small lateral hairs above the surrounding oil surface as shown in Fig. 23(c) and Fig. 21. The results are listed in Table 4. Since the wetting of the lateral hairs strongly varies, the table gives a range that covers all considered hairs. For the measurement of the CT data we used the software VG-Studio MAX 2.2. The simulation results were processed with the program ParaView 4.1.0. The second row in Table 4 states the number of lateral hairs on which the oil coverage was investigated. The wetting occurs along 7 lateral hairs in the simulation. The CT scan contains hundreds of lateral hairs wetted with oil so that only a subset could be used to measure the stated range of $7 - 20 \mu\text{m}$ for the lateral hair wetting. This subset consists of 10 randomly chosen hairs. These hairs were measured on zoomed extracts as in Fig. 23(c). A larger subset would lead to more precise results for the CT data but, on the other hand, this would not increase the accuracy of the simulation results. Nevertheless, the results in Table 4 indicate that the simulation underestimates the wetting along the lateral scopa hairs. One reason for this could be the simplified numerical treatment of the dynamic contact angle as discussed in Section 3.3. Furthermore, the wetting height would slightly increase for simulations with slip lengths smaller than $1 \mu\text{m}$. Moreover, there is a large overlap of the data ranges which shows that the general trend can indeed be reproduced in the



(a) Full scopa geometry (CT data).



(b) Zoom into system of branched hairs (CT data).



(c) Oil distribution in one single hair (CT data).

Figure 23: Oil distribution within the bee's scopa. The model results from a micro CT measurement at ITV Denkendorf.

simulations.

As the CT measurement data gives details on the steady state only, a comparison of the dynamics of the oil filling with the simulation is not possible. For the correct description of the dynamics of the system, the influence of the dynamic contact angle becomes more important than in the steady state case. For this reason, the error in the dynamic case, shown in Figure 17 to Figure 20, is presumably larger than for steady state, see Figure 21 and Table 4. We note that capturing of the dynamics of such systems on a micrometer length scale is difficult and still an area of active research.

5. Conclusions

We reported on our numerical flow simulations in the scopa of a specialized bee which are, at least to our knowledge, the first ones in this area. First, we discussed the postprocessing of a CT measured geometry and gave specific details on noise reduction and geometry optimization to allow for simulations with a Cartesian grid based flow solver. Our approach to fill the interior part of the hollow hairs is, as far as we know, not reported in literature yet. Then, we simulated the airflow through a CT measured geometry of the scopa. Furthermore, we reproduced the deformation of an oil droplet in a complex hair geometry on a microscopic length scale.

This work provides the following new insights into the process of oil transport in the microscopic scopa:

- As the vortex intensity in the simulation in Section 4.1 was largest directly behind the hairs of the scopa, the forces that act on a partially filled scopa are largest there as well. This is the region in which floral oil is most likely to get lost during transport.
- The influence of the slip length l_s on the wetting of the lateral hairs was investigated. For slip lengths below $1 \mu\text{m}$ differences in the scopa wetting only occurred at the contact line but the general oil distribution was identical.
- The oil distribution is substantially influenced by the side hairs of the scopa. Table 4 gives a qualitative estimate of the lateral hair wetting. Although, this effect was slightly underestimated in the simulation, the general shape of the oil wetting along the lateral hairs could be reproduced.

Our results clearly indicate the potential of simulations of complex biological systems for which direct experimental measurements are not possible.

Future simulations might remove the restrictions that were necessary to obtain results for such an enormously complex system within a reasonable amount of time:

- The simulations of the one and two-phase systems were restricted to resolutions with $439 \times 330 \times 833$

and with $200 \times 330 \times 389$ cells, respectively. Therefore, the simulations could only deal with eight scopa hairs. A long-term goal is the simulation of the process in the full scopa. For the performed CT measurements, this would lead to a simulation with $3072 \times 3072 \times 2400$ cells in total.

- The full process in nature takes place on different time scales. This was the reason to separately consider the process of oil absorption in Section 4.2 and the process of transportation at the bee's flight in Section 4.1. Ideally, this process would have to be simulated on its complete time scale.
- The simulation of the full problem requires the solution of the involved dynamic contact angle problem. We employed here a model which is based on the slip length l_s . Improved models exist, see [22], but their numerical verification needs still to be done.

Apart from the third problem, the restrictions can be overcome with massively parallel simulations on larger parallel systems. This emphasizes the importance of adequately scaling multi-CPU and multi-GPU flow solvers for the simulation of complex biological systems in the future.

Let us finally mention that the simulation with one specific scopa is not sufficient for a general statement on the oil storage properties of the bee's scopa. Bazilevs et al. [36], for instance, perform fluid-structure simulations on four patient-specific models of cerebral aneurysm that vary in shape and size. Therefore, many different scopae need to be measured and simulations need to be carried out in the spirit of uncertainty quantification, see, e.g., Le Maître and Knio [37] or Xiu [38].

Acknowledgments

This research was carried out under the project *Speicherung und Transport von Öl in Borsten spezialisierter Bienen und in faserbasierten Strukturen: Strukturprinzipien, Wirkmechanismen, Numerische Simulation und technologische Umsetzung* by the Deutsche Forschungsgemeinschaft (DFG).

References

- [1] R. Dressler, Biology of orchid bees (euglossini), *Ann. Rev. Ecol. Syst.* 13 (1982) 373–394.
- [2] S. Buchmann, The ecology of oil flowers and their bees, *Ann. Rev. Ecol. Syst.* (1987) 343–369.
- [3] J. Young, S. Walker, R. Bomphrey, G. Taylor, A. Thomas, Details of insect wing design and deformation enhance aerodynamic function and flight efficiency, *Sci.* 325 (2009) 1549–1552.
- [4] N. Nowak, P. Kakade, A. Annapragada, Computational fluid dynamics simulation of airflow and aerosol deposition in human lungs, *Ann. Biomed. Eng.* 31 (2003) 374–390.
- [5] D. Pham, C. Xu, J. Prince, Current methods in medical image segmentation, *Annu. Rev. Biomed. Eng.* 2 (2000) 315–337.

- [6] J. Weickert, *Anisotropic Diffusion in Image Processing*, volume 1, Teubner, Stuttgart, 1998.
- [7] C. Peskin, The immersed boundary method, *Acta Numer.* 11 (2002) 479–517.
- [8] R. Mittal, G. Iaccarino, *Immersed boundary methods*, *Annu. Rev. Fluid Mech.* 37 (2005) 239–261.
- [9] G. Karniadakis, A. Beskok, N. Aluru, *Microflows and Nanoflows: Fundamentals and Simulation*, volume 29 of *Interdisciplinary Applied Mathematics*, Springer, 2005.
- [10] M. Griebel, T. Dornseifer, T. Neunhoffer, *Numerical Simulation in Fluid Dynamics, a Practical Introduction*, SIAM, Philadelphia, 1998.
- [11] S. Groß, A. Reusken, *Numerical Methods for Two-Phase Incompressible Flows*, volume 40 of *Springer Series in Computational Mathematics*, Springer, 2011.
- [12] S. Osher, J. Sethian, Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulations, *J. Comput. Phys.* 79 (1988) 12–49.
- [13] M. Sussman, P. Smereka, S. Osher, A level set approach for computing solutions to incompressible two-phase flow, *J. Comput. Phys.* 114 (1994) 146–159.
- [14] P. Sagaut, *Large Eddy Simulation for Incompressible Flows*, Springer, Berlin, 2002.
- [15] I. Mary, P. Sagaut, Large eddy simulation of flow around an airfoil near stall, *AIAA J.* 40 (2002) 1139–1145.
- [16] J. Smagorinsky, General circulation experiments with the primitive equations: I. the basic experiment, *Mon. Weather Rev.* 91 (1963) 99–164.
- [17] J. Deardorff, A numerical study of three-dimensional turbulent channel flow at large Reynolds numbers, *J. Fluid Mech.* 41 (1970) 453–480.
- [18] R. Johnson, *The Handbook of Fluid Dynamics*, CRC Press, 1998.
- [19] T. Qian, X. Wang, P. Sheng, Molecular hydrodynamics of the moving contact line in two-phase immiscible flows, *Commun. Comput. Phys.* 1 (2006) 1–52.
- [20] C. Navier, Mémoire sur les lois du mouvement des fluides, *Mémoires de l’Académie Royale des Sciences de l’Institut de France* 6 (1823) 389–440.
- [21] C. Huh, L. Scriven, Hydrodynamic model of steady movement of a solid/liquid/fluid contact line, *J. Colloid Interface Sci.* 35 (1971) 85–101.
- [22] Y. Shikhmurzaev, *Capillary Flows with Forming Interfaces*, CRC Press, 2007.
- [23] A. Maali, B. Bhushan, Measurement of slip length on superhydrophobic surfaces, *Phil. Trans. R. Soc. A* 370 (2012) 2304–2320.
- [24] G. Jiang, C. Shu, Efficient implementation of weighted ENO schemes, *J. Comput. Phys.* 126 (1996) 202–228.
- [25] J. Brackbill, D. Kothe, C. Zemach, A continuum method for modeling surface tension, *J. Comput. Phys.* 100 (1992) 335–354.
- [26] R. Croce, M. Griebel, M. Schweitzer, Numerical simulation of bubble and droplet-deformation by a level set approach with surface tension in three dimensions, *Int. J. Numer. Methods Fluids* 62 (2009) 963–993.
- [27] J. Bell, P. Colella, H. Glaz, A second-order projection method for the incompressible Navier-Stokes equations, *J. Comput. Phys.* 85 (1989) 257–283.
- [28] B. Verleye, M. Klitz, R. Croce, D. Roose, S. Lomov, I. Verpoest, Computation of the permeability of textiles with experimental validation for monofilament and non crimp fabrics, in: *Computational Textile*, volume 55 of *Studies in Computational Intelligence*, Springer Berlin / Heidelberg, 2007, pp. 93–110.
- [29] B. Metsch, *Algebraic Multigrid (AMG) for Saddle Point Systems*, Dissertation, Institut für Numerische Simulation, Universität Bonn, 2013.
- [30] A. Chorin, Numerical solution of the Navier-Stokes equations, *Math. of Comput.* 22 (1968) 745–762.
- [31] R. Croce, *Numerische Simulation der Interaktion von inkompressiblen Zweiphasenströmungen mit Starrkörpern in drei Raumdimensionen*, Dissertation, Institut für Numerische Simulation, Universität Bonn, Bonn, Germany, 2010.
- [32] P. Zaspel, M. Griebel, Solving incompressible two-phase flows on multi-GPU clusters, *Comput. & Fluids* 80 (2013) 356 – 364. Selected contributions of the 23rd International Conference on Parallel Fluid Dynamics ParCFD2011.
- [33] M. Griebel, A. Rüttgers, Multiscale simulations of three-dimensional viscoelastic flows in a square-square contraction, *J. Non-Newtonian Fluid Mech.* 205 (2014) 41–63.
- [34] J. Adelsberger, P. Esser, M. Griebel, S. Groß, M. Klitz, A. Rüttgers, 3D incompressible two-phase flow benchmark computations for rising droplets, *Proceedings of the 11th World Congress on Computational Mechanics (WCCM XI)*, Barcelona.
- [35] H. Wadell, Volume, shape and roundness of quartz particles, *J. Geol.* 43 (1935) 250–280.
- [36] Y. Bazilevs, M. Hsu, Y. Zhang, W. Wang, X. Liang, T. Kvamsdal, R. Brekken, J. Isaksen, A fully-coupled fluid-structure interaction simulation of cerebral aneurysms, *Comput. Mech.* 46 (2010) 3–16.
- [37] O. Le Maître, O. Knio, *Spectral Methods for Uncertainty Quantification: With Applications to Computational Fluid Dynamics*, Springer, 2010.
- [38] D. Xiu, *Numerical Methods for Stochastic Computations: A Spectral Method Approach*, Princeton University Press, 2010.