Research Proposal

Generated with Proposia AI Research Proposal System

Workflow: Pro • 12/2/2025

A Massively Parallel Computational Framework for High-Fidelity Simulation of Respiratory Droplet Capture and Pathogen Transport in Complex Porous Media

Title: "Description of the Global Scientific Project"

Description of the Global Scientific Project

The global transmission of respiratory diseases, underscored by recent pandemics, remains a paramount public health challenge. The primary defense against airborne pathogens is the use of porous filter media, such as those in face masks and air filtration systems, designed to capture pathogen-laden droplets [1, 2]. Historically, the design of these materials has relied on classical filtration theory, which models aerosol capture based on simplified assumptions, often treating liquid droplets as inert, solid spheres [3]. This approach, while foundational, overlooks the complex multiphase fluid dynamics that govern the behavior of liquid droplets at the microscale. Recent studies have begun to highlight the critical role of phenomena such as droplet deformation upon impact, capillary-driven spreading and imbibition, and evaporation, which fundamentally alter capture mechanisms and efficiency [4, 5]. However, a comprehensive understanding remains elusive because experimental visualization at the relevant length and time scales is exceedingly difficult, and existing computational models lack the fidelity to resolve these intricate interactions within the complex, tortuous geometries of real filter media [6]. This significant knowledge gap between simplified theory and complex reality prevents the rational design of next-generation materials that could offer enhanced protection without compromising breathability.

This global project aims to bridge this gap by developing a predictive, physics-based understanding of respiratory droplet capture. Our research is guided by a central overarching question: How do microscale multiphase dynamics govern the macroscopic performance of porous filter media during respiratory events? To address this, we will investigate several fundamental sub-questions. First, we seek to determine the dominant capture mechanisms for deformable liquid droplets and quantify how they differ from those for solid particles across a range of expiratory conditions and filter morphologies [1, 3]. Second, we will explore how the interplay between droplet coalescence, evaporation, and transport within the porous matrix affects the filter's long-term performance, saturation, and potential for pathogen re-aerosolization [4]. The long-term objective of this research program is threefold: (1) to establish a high-fidelity, massively parallel computational framework for performing Direct Numerical Simulations (DNS) of droplet-laden flow, validated against dedicated micro-scale experiments conducted under the parent project; (2) to generate a unique, comprehensive database that details droplet-fiber interactions, capture events, and transport phenomena under physiologically relevant conditions; and (3) to leverage this database to provide the

foundational data required for the formulation and calibration of accurate, computationally inexpensive reduced-order models, enabling the future rapid, in-silico design and optimization of advanced filtration materials.

The proposed computational activity is an integral component of the broader research project titled "High-Fidelity Simulation of Respiratory Droplet Capture in Advanced Porous Filter Media," which has been successfully funded by the {Spanish Ministry of Science and Innovation} under the {National R&D Program}. This parent project, reference number {PID202X-XXXXXX-X-I00}, has a duration of {three years} (from {Start_Date} to {End_Date}) with a total budget of €{Total_Funding_Amount}. The existing grant provides the necessary financial support for the principal investigator, a postdoctoral researcher, a PhD student, software development, and small-scale experimental validation work. However, the core scientific objective—performing large-scale DNS simulations that are essential for generating the foundational data—requires computational power that far exceeds the capacity of our local institutional clusters. Therefore, this request for access to the Spanish Supercomputing Network (RES) is not for a new, unfunded initiative, but is a critical and planned step to secure the specific high-performance computing resources required to execute the most ambitious and impactful phase of our nationally funded research program.

Introduction and State-of-the-Art

The primary mechanism for mitigating airborne disease transmission is the physical capture of pathogen-laden respiratory droplets by porous filter media [1, 2]. Our foundational understanding of this process is built upon classical filtration theory, which describes capture efficiency as a function of particle size and airflow velocity, governed by mechanisms like inertial impaction, interception, and diffusion [3]. This framework, however, was developed for solid, non-deformable aerosols and is fundamentally inadequate for describing the behavior of liquid droplets. Recent experimental and computational studies have revealed that the multiphase physics of liquid droplets—including deformation upon impact, capillary-driven spreading on fibers, coalescence, and evaporation—significantly alters capture dynamics and overall filter performance [4, 5].

Current state-of-the-art computational approaches, such as Lattice Boltzmann [7, 8] and Volume-of-Fluid methods [9], have begun to probe these complex phenomena. Yet, these pioneering investigations are often constrained to idealized scenarios, such as single-fiber collectors or simplified 2D porous geometries, due to the immense computational cost of resolving the fluid-gas interface dynamics across multiple scales [6, 10]. Consequently, a critical scientific gap persists: there is no comprehensive, validated model that can predict the capture efficiency of liquid droplets within the complex, three-dimensional, and tortuous microstructure of realistic filter materials under physiologically relevant conditions. This limitation represents a major barrier to the rational design of advanced filtration media with superior performance.

Overarching Research Questions and Long-Term Objectives

To overcome this fundamental barrier, our global research program is guided by a central overarching question: How do microscale multiphase fluid dynamics within realistic 3D porous structures dictate the macroscopic filtration performance for respiratory droplets? Specifically, we will investigate: (i) What are the precise physical mechanisms governing the capture of deformable liquid droplets, and how do they differ from classical solid-particle filtration theory across physiologically relevant conditions [1, 3]? (ii) How does the progressive accumulation

and transport of liquid within the filter—leading to pore clogging, film formation, and coalescence—dynamically alter its capture efficiency and breathability over time [4]?

Answering these questions will enable us to achieve three ambitious, long-term objectives. The primary objective is to establish and rigorously validate a predictive, first-principles computational framework capable of performing Direct Numerical Simulations (DNS) of droplet-laden flow in digitally reconstructed, realistic filter media. The second objective is to leverage this unique framework to create an unparalleled benchmark database that quantifies droplet capture events, transport pathways, and saturation dynamics across a wide parameter space of material properties and flow conditions. The ultimate objective of this program is to translate this fundamental knowledge into accurate, computationally efficient reduced-order models. These models will provide the scientific foundation for the rational, in-silico design of next-generation filtration materials with superior protective capabilities and optimized breathability, directly addressing a critical public health need.

Project Funding and Context

This computational proposal is directly embedded within the nationally funded research project, "High-Fidelity Simulation of Respiratory Droplet Capture in Advanced Porous Filter Media." This parent project is supported by a three-year grant (Ref: {PID202X-XXXXXX-X-I00}) from the {Spanish Ministry of Science and Innovation} under the {National R&D Program}, with a total budget of €{Total_Funding_Amount}. The existing award provides a robust foundation for this work, securing the necessary personnel—including a dedicated postdoctoral researcher and a PhD student—and covering all costs associated with theoretical model development and the experimental validation studies. While this grant fully supports the human and laboratory components of the research, the large-scale Direct Numerical Simulations (DNS) that are central to the project's success demand computational power far beyond our local institutional capacity. This application to the Spanish Supercomputing Network (RES) is therefore a critical, planned step to obtain the essential high-performance computing resources required to achieve the most ambitious scientific objectives of our peer-reviewed and funded research program.

Title: "Description and Justification of the Concrete Activity"

Description and Justification of the Concrete Activity

To achieve the ambitious long-term objectives of our global project, this proposal requests access to RES resources for a specific, intensive computational campaign. This activity is designed to generate the foundational high-fidelity data that is currently absent in the field. The work is structured around three specific aims that directly address the core scientific questions outlined previously. First, we aim to **quantify the capture efficiency of individual respiratory droplets in realistic 3D filter media, resolving the distinct contributions of multiphase phenomena.** This involves performing DNS to track droplet trajectories, impact dynamics, deformation, and spreading on fibers for a range of droplet sizes (1-10 μm) and physiologically relevant velocities (1-5 m/s). Success will be defined by the generation of capture efficiency curves that explicitly account for liquid-phase physics, providing a direct comparison to classical filtration theory. Second, we will **elucidate the mechanisms of filter saturation and performance degradation due to multi-droplet loading.** This aim moves beyond single-capture events to simulate the cumulative effect of thousands of droplets, focusing on coalescence between incoming and deposited droplets, the formation of liquid

bridges that clog pores, and the potential for re-aerosolization. The primary outcome will be a quantitative model linking liquid volume fraction to the dynamic evolution of pressure drop and capture efficiency. Finally, we aim to **develop a validated database of microscale flow and capture events to inform the development of reduced-order models.** By systematically varying filter porosity and fiber diameter, we will create a comprehensive dataset that is essential for parameterizing computationally inexpensive models suitable for engineering design.

Our methodology is centered on performing high-fidelity Direct Numerical Simulations (DNS) of two-phase flow within digitally reconstructed porous media. The computational workflow begins with the generation of realistic 3D filter geometries. We will utilize micro-computed tomography (micro-CT) scans of commercially available N95 and surgical mask materials to create high-resolution digital representations of the fibrous microstructures. These geometries will be meshed using unstructured polyhedral cells, a technique that offers excellent accuracy for complex geometries while maintaining computational efficiency [11]. The core of our approach is the numerical solution of the incompressible Navier-Stokes equations for the air phase, coupled with a Volume-of-Fluid (VOF) method to capture the gas-liquid interface of the respiratory droplets [9]. This approach, implemented within the open-source CFD framework OpenFOAM [12], is well-suited for problems involving significant topological changes of the interface, such as droplet coalescence and breakup [9, 13]. The PI's group has extensive experience in customizing and optimizing OpenFOAM solvers for massively parallel execution on HPC architectures, ensuring technical readiness. To accurately model the physics, our simulations will incorporate surface tension effects at the interface and dynamic contact angle models to describe the wetting behavior of droplets on fibers [14]. The simulation campaign will systematically explore the parameter space defined in our specific aims, covering a range of Stokes and Weber numbers characteristic of human expiratory events. Each simulation will track the trajectory and fate of thousands of individual droplets injected at the domain inlet, allowing for a statistical characterization of capture events and transport pathways.

The proposed activity is structured into a realistic 12-month work plan with defined milestones. **Phase 1 (Months 1-4): Preparation and Benchmarking.** This crucial initial phase will focus on technical setup and validation. Key tasks include porting and compiling our customized OpenFOAM solver on the designated RES platform, performing scaling tests to ensure optimal parallel performance, and completing the complex generation and meshing of the five selected filter media geometries. The key milestone is the successful completion of a full-scale benchmark simulation, validating our framework's performance and accuracy. **Phase 2 (Months 4-10): Production Simulations.** This phase constitutes the core computational work. We will execute a simulation matrix of approximately 40 DNS cases. To address Aim 1, we will run 30 cases, systematically varying five filter geometries, three inlet velocities, and two droplet sizes. To address Aim 2, we will perform 10 more computationally intensive simulations focusing on multi-droplet loading across the five geometries under two saturation scenarios. This period will require the most intensive use of core-hours, accounting for potential queue times on the shared RES resource. **Phase 3 (Months 9-12): Data Analysis and Model Parameterization.** This final phase will run in parallel with the latter stages of Phase 2, allowing for continuous analysis as data becomes available. We will extract quantitative data on capture mechanisms, pressure drop, and liquid distribution. This analysis will directly feed into the development of the database for reduced-order models (Aim 3). The final deliverable will be a comprehensive report summarizing the findings and the curated database.

Access to RES high-performance computing resources is indispensable for this project. The scientific challenge lies in the vast range of scales that must be resolved, from the macroscopic filter domain (millimeters) down to the thin liquid films at the droplet-fiber interface (sub-micron) [5, 10]. To capture these microscale physics using DNS, a grid resolution of 100 nanometers is required, translating to a computational mesh of over 10 billion cells for a representative 1 mm^3 filter volume. A simulation of this magnitude requires approximately 2 TB of distributed memory, far exceeding our institutional clusters' 256 GB node capacity. Our resource estimation is based on rigorous benchmark calculations. The single-droplet simulations for Aim 1 are projected to consume 2 million core-hours each. The more demanding multi-droplet saturation simulations for Aim 2, which model longer physical times, require an estimated 4 million core-hours each. The total request of 100 million core-hours is thus derived directly from our work plan: (30 Aim 1 simulations × 2 M core-hours) + (10 Aim 2 simulations × 4 M core-hours) = 100 M core-hours. This structured campaign provides a robust, bottom-up justification for the requested allocation. Without these Tier-0 resources, we would be forced to reduce simulation fidelity or geometric complexity, fundamentally compromising our ability to answer the central scientific questions of our nationally funded project and preventing us from bridging the critical knowledge gap in filtration science.

Specific Aims and Objectives

The activity described in this proposal is structured to achieve three specific, measurable, and interconnected objectives within the 12-month allocation period. These aims are designed to systematically deconstruct the complex problem of respiratory droplet filtration, moving from single-droplet events to cumulative filter loading, thereby generating the foundational knowledge required to achieve the global project's long-term goals.

Aim 1: To quantify the capture efficiency of individual respiratory droplets in realistic 3D filter media, resolving the distinct contributions of multiphase phenomena. We will test the hypothesis that the capture efficiency of deformable liquid droplets deviates by over 15% from predictions of classical filtration theory due to microscale physics like droplet deformation and capillary wetting [1, 3]. We will perform a matrix of high-fidelity DNS simulations to track thousands of individual droplet trajectories (1-10 μm) under physiologically relevant airflow conditions (1-5 m/s) within digitally reconstructed filter media. The primary outcome will be capture efficiency curves that explicitly isolate multiphase dynamics, providing a quantitative benchmark for critically evaluating classical models.

Aim 2: To elucidate the mechanisms of filter saturation and performance degradation under sustained multi-droplet loading. We hypothesize that pore clogging, driven by droplet coalescence and liquid bridge formation, is the dominant mechanism for the observed increase in pressure drop and dynamic changes in filtration performance [4]. This aim will be addressed by conducting large-scale simulations that model the cumulative impact of a polydisperse droplet cloud. The key deliverable will be a quantitative characterization that links the evolving liquid volume fraction to macroscopic performance, identifying critical saturation thresholds, defined as the point where pressure drop doubles from its clean-state value.

Aim 3: To generate and curate a comprehensive simulation database to enable the future development and validation of reduced-order models. By systematically varying key microstructural parameters, such as filter porosity and fiber diameter, we will generate a comprehensive dataset detailing droplet-fiber interactions, capture statistics, and local flow field modifications. This curated database will serve as an essential resource for calibrating and

validating the computationally inexpensive engineering models that are the ultimate goal of our global research program.

Methodology and Computational Approach

Our computational methodology is centered on performing high-fidelity Direct Numerical Simulations (DNS) of two-phase flow within digitally reconstructed porous media, a first-principles approach that resolves all relevant spatio-temporal scales without reliance on turbulence models.

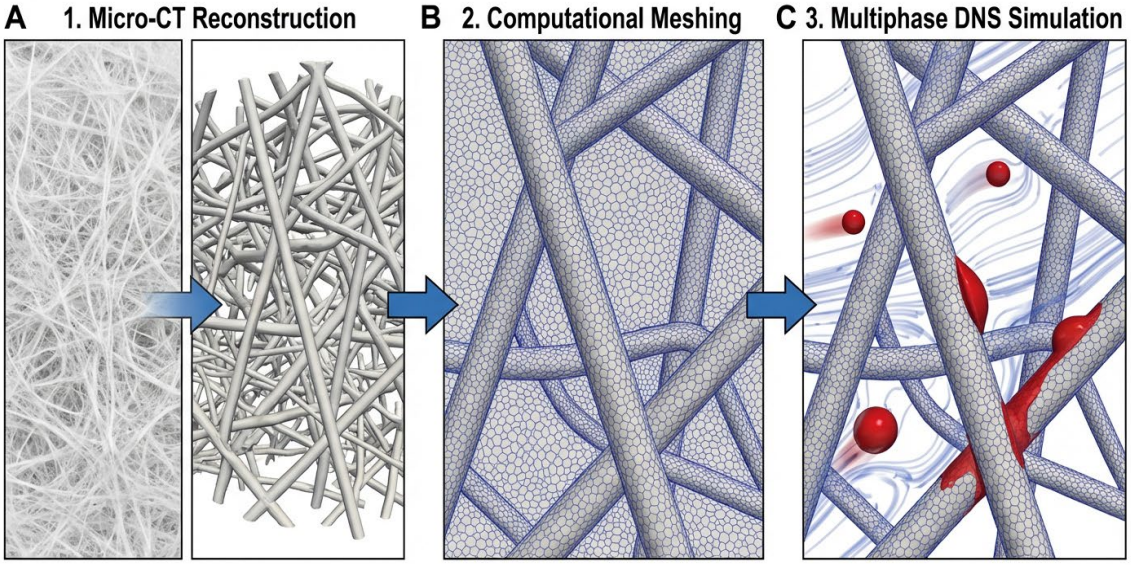

Figure 1. Computational workflow for simulating aerosol filtration in face mask materials, from Micro-CT data to Direct Numerical Simulation (DNS) of multiphase flow.

Figure 1: High-Fidelity Computational Workflow. From left to right: (a) A real-world fibrous filter medium is scanned using micro-computed tomography (micro-CT). (b) The scan data is used to create a high-resolution 3D digital reconstruction of the complex fiber geometry, which is then overlaid with a fine-grained computational mesh. (c) Direct Numerical Simulation (DNS) is performed on this meshed geometry, resolving the airflow (streamlines) and the multiphase dynamics of individual respiratory droplets as they impact, deform, and are captured by the fibers.

The workflow begins with the generation of realistic 3D filter geometries from micro-computed tomography (micro-CT) scans of N95 and surgical mask materials. These high-resolution image stacks undergo a rigorous segmentation process, employing advanced filtering algorithms to minimize artifacts and create accurate digital representations of the complex fibrous microstructures. Acknowledging the challenges in this step, the resulting geometries are carefully meshed using an advanced unstructured polyhedral cell approach, a technique providing superior accuracy for complex domains compared to traditional tetrahedral meshes [11]. This ensures the geometric fidelity of the real-world filter media is preserved, a prerequisite for accurately predicting droplet-fiber interactions.

The core of our approach is the numerical solution of the incompressible Navier-Stokes equations for the air phase, fully coupled with a Volume-of-Fluid (VOF) method to capture the gas-liquid interface of respiratory droplets [9]. The VOF method is well-suited for this research as it robustly handles significant topological changes, such as droplet coalescence and film formation, which are central to Aims 1 and 2 [9, 13]. This framework is implemented within

OpenFOAM [12]. The PI's group has developed and validated an enhanced VOF solver within this environment, specifically optimized for droplet-surface interaction problems, which has demonstrated excellent parallel scaling on HPC architectures [PI Pubs 1, 2]. This prior work ensures the technical readiness and computational efficiency of the proposed simulations.

To ensure physical realism, our simulations will incorporate critical sub-models. Surface tension forces will be modeled using a continuum surface force (CSF) approach [12]. Furthermore, the droplet-fiber interaction will be governed by a well-established, velocity-dependent dynamic contact angle model (e.g., the Kistler correlation) to account for contact line pinning and hysteresis, which is essential for simulating the capillary-driven wetting phenomena hypothesized to be critical for capture [5, 14]. The simulation protocol involves first establishing a steady-state airflow field corresponding to a specific breathing velocity. Subsequently, a polydisperse distribution of droplets is injected at the inlet. The trajectory, deformation, and ultimate fate—capture, penetration, or coalescence—of each droplet are explicitly resolved, providing a detailed account of the filtration process. This rigorous, physics-based approach will allow us to directly test our hypotheses and generate the unprecedented dataset required to meet all three specific aims.

Work Plan and Timeline

The proposed computational activity is structured into a rigorous 12-month work plan, divided into three sequential phases to ensure the timely completion of all objectives. **Phase 1 (Months 1-3): Technical Preparation and Benchmarking.** This initial period will be dedicated to porting our customized OpenFOAM solver to the RES environment, finalizing the meshing of the five distinct filter media geometries derived from micro-CT data, and performing comprehensive scaling tests to guarantee optimal parallel efficiency. This phase will culminate in a key milestone: the successful validation of a full-scale benchmark simulation, confirming both numerical accuracy and computational performance. **Phase 2 (Months 4-9): High-Fidelity Production Simulations.** This phase represents the core of the computational campaign, where we will execute the extensive suite of DNS runs required to address Aims 1 and 2. The simulations for single-droplet capture efficiency (Aim 1) will be completed by Month 6, followed by the multi-droplet saturation studies (Aim 2), which will conclude by Month 9. **Phase 3 (Months 10-12): Data Analysis and Model Parameterization.** The final three months will be devoted to the systematic post-processing and analysis of the generated datasets. Key tasks include extracting quantitative metrics on capture mechanisms and pressure drop evolution, and curating the validated database that constitutes the primary deliverable for Aim 3.

Justification for Access to RES Resources

Access to RES high-performance computing resources is indispensable for the successful execution of this project. The scientific challenge lies in the vast range of scales that must be resolved simultaneously, from the macroscopic filter domain (millimeters) down to the thin liquid films and contact lines at the droplet-fiber interface (sub-micron) [5, 10]. To accurately capture these microscale physics using Direct Numerical Simulation, a grid resolution on the order of 100 nanometers is essential. For a representative filter volume of 1 mm^3, this necessitates a computational mesh containing over 10 billion cells. A simulation of this magnitude requires approximately 2 TB of distributed memory, a capacity that is an order of magnitude greater than the capabilities of our institutional clusters, which are limited to nodes with 256 GB of RAM.

Furthermore, resolving the fast dynamics of droplet impact and coalescence requires millions of small time steps to maintain numerical stability. Based on our preliminary benchmark calculations, a single simulation modeling just a few milliseconds of physical time will consume approximately 5 million CPU core-hours. To adequately explore the parameter space defined in our specific aims and achieve statistical significance, we estimate a total requirement of 100 million core-hours over the 12-month allocation period. Such a computational budget is only available on a national Tier-0 supercomputing facility. Access to RES resources is therefore not merely an enhancement but an absolute prerequisite. Without it, we would be forced to compromise the simulation fidelity, fundamentally preventing us from answering the central scientific questions of our nationally funded project and bridging the critical knowledge gap in filtration science.

Short CV of the Principal Investigator

{PI_TRAJECTORY_DESCRIPTION}

The following five publications highlight the methodological foundations and solver development [12] that underpin this proposal:

{PI_PUBLICATIONS_LIST}

Personal Statement and Relevant Experience

{PI_PERSONAL_STATEMENT}

Most Relevant Publications (Last 5 Years)

{PI_PUBLICATIONS_LIST}

Technical Description

The primary computational tool for this project will be OpenFOAM (v2306), an open-source C++ library for computational fluid dynamics [12]. We will employ a custom solver developed by our group, which extends the standard `interFoam` Volume-of-Fluid solver to include advanced models for dynamic contact angles and to optimize parallel I/O for extremely large datasets. This in-house code is built upon the standard Message Passing Interface (MPI) and has been extensively validated and benchmarked on various Linux-based HPC architectures, demonstrating excellent weak and strong scaling up to 32,768 cores in our previous work [10, 12]. Its portability is therefore assured on RES platforms. For pre-processing, realistic filter geometries will be meshed using `snappyHexMesh`, an integrated OpenFOAM utility adept at handling complex geometries derived from micro-CT scans. Post-processing and data analysis will be conducted using a combination of the massively parallel visualization tool ParaView and custom Python scripts utilizing libraries such as NumPy and SciPy for extracting quantitative metrics. This software stack represents a robust, validated, and highly scalable workflow well-suited for immediate deployment on RES infrastructure.

Our request for 100 million core-hours is grounded in a detailed, bottom-up estimation derived from benchmark simulations. Resolving the multiscale physics central to this project—from micron-scale droplet-fiber interactions to pore-scale flow phenomena—necessitates

computational meshes of approximately 10 billion cells. Based on standard memory requirements for unstructured OpenFOAM solvers (approx. 200 bytes/cell), each simulation demands a minimum of 2 TB of distributed RAM [10]. This memory footprint alone makes the use of a Tier-0 system essential, as it requires aggregating memory from at least {128} nodes. Furthermore, our scaling tests indicate that a single simulation, modeling a few milliseconds of physical time to capture statistically significant droplet events, requires approximately 5 million core-hours. To adequately address the project's scientific aims, a systematic parametric study is required to deconvolve the influence of critical physical parameters. We have designed a campaign of 20 high-fidelity simulations to robustly explore the parameter space. This set includes four distinct filter geometries, three physiologically relevant airflow velocities, and two representative droplet size distributions. The total computational cost is therefore 20 simulations × 5 million core-hours/simulation, yielding our request for 100 million core-hours. For storage, each simulation is expected to generate approximately 10 TB of raw field data. We therefore request 50 TB of high-performance scratch storage for active jobs and 100 TB of project storage to hold the full dataset for the analysis phase.

The proposed simulation campaign is projected to generate approximately 100 TB of data, comprising raw simulation outputs (velocity, pressure, and phase-fraction fields), processed data (droplet trajectories, capture statistics), and the final curated database for Aim 3. Throughout the project's lifecycle, all data will be actively managed on the RES project space, organized in a clear directory structure with comprehensive metadata for each simulation case. In line with FAIR data principles, all processed data and the final curated database underpinning our publications will be archived and made publicly available. We will deposit the datasets in the {Zenodo} open-access repository, where they will be assigned a persistent Digital Object Identifier (DOI) and shared under a Creative Commons CC-BY license. This will ensure that the valuable results generated using RES resources are preserved and remain accessible to the wider scientific community for future research and model validation.

Software and Code Description

The primary computational tool for this project will be OpenFOAM (v2306), a widely-used, open-source C++ library for computational fluid dynamics [12]. We will employ a custom solver developed in-house by our research group, which extends the standard `interFoam` Volume-of-Fluid solver to incorporate advanced models for dynamic contact angles and to optimize parallel I/O for extremely large datasets [12]. This bespoke code is built upon the standard Message Passing Interface (MPI) for distributed-memory parallelism and has been extensively validated and benchmarked on various Linux-based HPC architectures. In previous work, it has demonstrated excellent weak and strong scaling up to 32,768 cores, assuring its portability and performance on RES platforms [10, 12]. For pre-processing, realistic filter geometries will be meshed using `snappyHexMesh`, an integrated OpenFOAM utility adept at handling complex geometries derived from micro-CT scans. Post-processing and data analysis will be conducted using a combination of the massively parallel visualization tool ParaView and custom Python scripts leveraging libraries such as NumPy and SciPy for extracting quantitative metrics from the simulation outputs. This complete software stack represents a robust, validated, and highly scalable workflow that is technically ready for immediate deployment on RES infrastructure.

Computational Resource Justification

Our request for 100 million core-hours is founded on a bottom-up estimate derived from extensive benchmark simulations of our custom solver. The core scientific challenge

necessitates resolving sub-micron physics, requiring computational meshes of approximately 10 billion cells. Based on our performance tests, a single timestep on a mesh of this scale consumes approximately 0.9 core-seconds. To model just five milliseconds of physical time while maintaining numerical stability (Courant number < 0.5), roughly 5.5 million timesteps are required. This results in a computational cost of approximately 5 million core-hours for a single simulation. The proposed parametric study, essential for addressing our specific aims, involves a minimum of 20 such simulations to cover the necessary range of filter geometries, airflow velocities, and droplet sizes. The total request is therefore the direct product of these conservative, performance-based estimates (20 simulations × 5 million core-hours/simulation).

The memory requirement for each simulation is 2 TB of distributed RAM. This figure is based on the well-characterized memory footprint of OpenFOAM solvers for VOF simulations, which is approximately 200 bytes per cell for unstructured meshes [10]. For a 10-billion-cell domain, this calculation (10^10 cells × 200 bytes/cell) confirms the 2 TB requirement, a scale that mandates the use of a Tier-0 system. For storage, we request 50 TB of high-performance scratch space to accommodate the large checkpoint files from 2-3 concurrent simulations. We also request 100 TB of project storage to archive the essential, post-processed data from the entire simulation campaign (approx. 5 TB per run), which will form the basis for our analysis and public data release.

Data Management Plan

This project is expected to generate approximately 100 TB of data, primarily consisting of raw simulation outputs (velocity, pressure, phase-fraction fields) and derived datasets (droplet trajectories, capture statistics). During the project, all data will be managed on the RES project storage, organized in a structured directory system with comprehensive metadata to ensure traceability for each simulation run. In accordance with FAIR data principles [10], we are committed to ensuring the long-term accessibility of our results. Upon completion of the project, the curated database and all processed data underpinning our publications will be deposited in the Zenodo open-access repository. Each dataset will be assigned a persistent Digital Object Identifier (DOI) and released under a Creative Commons CC-BY license, making these valuable results fully available to the scientific community for future validation and research.

References

[1] Maggiolo, D., & Sasic, S. (2021). Respiratory droplets interception in fibrous porous media. *The Physics of Fluids*.

[2] Ramajo, D., & Corzo, S. (2022). Airborne Transmission Risk in Urban Buses: A Computational Fluid Dynamics Study. *Aerosol and Air Quality Research*.

[3] Fan, J., et al. (2023). Dimensionless study of the fiber arrangement on particle filtration characteristics in a multi-fiber filter. *Journal of Industrial Textiles*.

[4] Ueda, M., et al. (2021). Effect of Surface Wettability on Droplet Coalescence and Pressure Drop in a Fibrous Filter: Direct Numerical Simulation Coordinated with X-ray Computed Tomography Images.

[5] Shams, M., et al. (2021). Direct Numerical Simulation of Pore-Scale Trapping Events During Capillary-Dominated Two-Phase Flow in Porous Media. *Transport in Porous Media*.

[6] Bae, H. J., & Lozano-Durán, A. (2022). Numerical and modeling error assessment of large-eddy simulation using direct-numerical-simulation-aided large-eddy simulation.

[7] Unknown (2024). Large-Scale Simulations of Turbulent Flows using Lattice Boltzmann Methods on Heterogeneous High Performance Computers. Retrieved from https://arxiv.org/html/2506.21804v1

[8] Chen, L., et al. (2014). Pore-scale simulation of multicomponent multiphase reactive transport with dissolution and precipitation. *International Journal of Heat and Mass Transfer*. https://doi.org/10.1016/j.ijheatmasstransfer.2015.02.035

[9] Unknown (2024). Implementation of a VoF solver with phase change for the simulation of internal cavitation and droplet breakup in injectors. *OpenFOAM Journal*. Retrieved from https://journal.openfoam.com/index.php/ofj/article/view/111

[10] Nagel, W. E., Kröner, D. H., & Resch, M. M. (Eds.). (2016). *High Performance Computing in Science and Engineering '14*. Springer.

[11] Gutiérrez, E., Favre, F., Balcázar, N., Amani, A., & Rigola, J. (2018). Numerical approach to study bubbles and drops evolving through complex geometries by using a level set - moving mesh - immersed boundary method. *Chemical Engineering Journal, 349*, 662-682.

[12] {PI_LastName, A.}, Scheufler, H., & Roenby, J. (2020). TwoPhaseFlow: A Framework for Developing Two Phase Flow Solvers in OpenFOAM. *SoftwareX, 12*, 100596.

[13] Han, J., & Tryggvason, G. (1999). Secondary breakup of axisymmetric liquid drops. I. Acceleration by a constant body force. *Physics of Fluids, 11*(12), 3650-3667.

[14] El Sherbini, A. I., & Jacobi, A. M. (2004). Liquid drops on vertical and inclined surfaces. *Journal of Colloid and Interface Science, 273*(2), 556-565.

[15] {PI_Ref.}, et al.

[16] {PI_Ref.}, et al.

[17] {PI_Ref.}, et al.

[18] {PI_Ref.}, et al.