




Contents lists available at ScienceDirect

Digital Chemical Engineering

journal homepage: www.elsevier.com/locate/dche

Original article

A CFD-based digital twin framework for transient MoCl₅ transport in an experimental atomic layer deposition process

Abdulrahman Alghamdi^a, Feiyang Ou^a , Berc Kalanyan^c, James E. Maslar^c ,
Panagiotis D. Christofides^{a,b,*}

^a Department of Chemical and Biomolecular Engineering, University of California, Los Angeles, CA, 90095-1592, USA^b Department of Electrical and Computer Engineering, University of California, Los Angeles, CA 90095-1592, USA^c National Institute of Standards and Technology, Gaithersburg, MD, 20899, USA

ARTICLE INFO

Keywords:

Semiconductor manufacturing
CFD modeling
Atomic layer deposition
Model validation

ABSTRACT

A physics-based digital twin framework is developed for transient molybdenum pentachloride (MoCl₅) transport in an optically-accessible, experimental atomic layer deposition (ALD) reactor at the National Institute of Standards and Technology (NIST). The framework is anchored to high-speed absorption-imaging measurements reported by NIST and is implemented using a three-dimensional, time-resolved computational fluid dynamics (CFD) model that resolves momentum, heat, and species transport during pulsed precursor delivery. Experimentally measured, time-dependent absorbance data are used to inform the inlet precursor waveform, enabling direct comparison of normalized transport behavior between simulation and experiment in the absence of calibrated absolute concentrations. The digital twin of the reactor gas delivery system is validated against multiple experimental observables, including inlet velocity, precursor buildup and decay dynamics, residence time, spatial plume structure, and wafer chuck-level transport behavior, for representative operating conditions. The model reproduces the dominant temporal and spatial trends observed experimentally, including plume symmetry, transport delays, and purge dynamics, while maintaining physically consistent flow scaling near the wafer surface. Image-based comparisons further confirm that the CFD modeling framework captures the transient evolution and symmetry of MoCl₅ plumes under both low- and high-flow conditions. Beyond validation, a parametric study is conducted in which flow rate, chamber pressure, and injection duration are varied using a unified model configuration, including evaluation of an intermediate 300 SCCM condition to assess consistency between baseline cases. The results demonstrate that the model responds smoothly and predictably across the operating space, with wafer chuck-level velocity governed primarily by mass flow rate and precursor exposure transitioning from transport-limited to injection-duration-limited behavior for long pulses. Together, these results establish the present modeling framework as an experiment-informed, reusable digital twin suitable for rapid evaluation of ALD operating conditions and future extension to reaction-coupled modeling and process optimization.

1. Introduction

Relentless device scaling and the transition from planar to three-dimensional transistor architectures have tightened process windows throughout semiconductor manufacturing. As summarized in recent work on multiscale CFD for atomic layer etching (ALE), EUV lithography enhanced FinFET (Wang et al., 2020) scaling from the 22 nm generation to nominal 5 nm nodes has confronted mobility loss and short-channel effects (Jurczak et al., 2009), motivating the gate-all-around (GAA) nanosheet architecture as a successor technology that demands atomic-scale precision in thin-film processing (Loubet et al., 2017; Lee et al., 2020) as entering into 3 nm process era (Chen et al.,

2015; Huang, 2022). In this context, atomic layer deposition (ALD) and its inverse, atomic layer etching (ALE), have become indispensable because their sequential, self-limiting surface reactions enable angstrom-level control of film addition and removal across complex 3D geometries (George, 2010; Kanarik et al., 2015).

Although the chemistry of ALD/ALE is surface-limited by design, reactor-scale transport strongly influences uniformity, throughput, and precursor utilization. The introduction of the ALE chamber optimization study (Yun et al., 2022) catalogs how nonuniform precursor delivery, flow recirculation, and departures from laminar behavior degrade spatial uniformity and can compromise the benefits of self-limitation;

* Corresponding author at: Department of Chemical and Biomolecular Engineering, University of California, Los Angeles, CA, 90095-1592, USA.
E-mail address: pdc@seas.ucla.edu (P.D. Christofides).

<https://doi.org/10.1016/j.dche.2026.100304>

Received 24 February 2026; Received in revised form 20 March 2026; Accepted 20 March 2026

Available online 23 March 2026

2772-5081/© 2026 The Authors. Published by Elsevier Ltd on behalf of Institution of Chemical Engineers (IChemE). This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

it also surveys the emergence of reactor-scale modeling to interrogate such effects. Over roughly the past decade, computational fluid dynamics (CFD) has been applied to ALD (Pan et al., 2014; Shaeri et al., 2015; Peltonen et al., 2018) and related processes to quantify momentum, heat, and mass transport, and also connect reactor design choices with wafer-level outcomes when coupled to microscopic reaction models.

Beyond aiding interpretation, CFD provides the core of a physics-based digital twin for atomic-layer processes: an executable, predictive representation of the reactor that can be interrogated across a wide operating space without risk to hardware or product. Two motivations emphasized in recent multiscale ALD frameworks align directly with digital twin practice: (i) experimental diagnostics are often limited to a small number of locations and are costly to iterate, and (ii) three-dimensional CFD enables systematic sweeps of operating conditions and reactor configurations while providing local state variables (e.g., pressure, temperature, and species concentration) needed to inform microscopic surface-reaction models.

Moreover, reactor chamber configuration constitutes a critical design variable that can be systematically explored and optimized within a digital twin framework. In Yun et al. (2022), a set of four cross-flow reactor geometries was evaluated using multiscale modeling to quantify their impact on transport uniformity, cycle efficiency, and precursor utilization. Among the configurations studied, the inclined-plate design achieved simultaneously improved flow uniformity and reduced half-cycle times. Importantly, this configuration also led to a substantial reduction in annual HF and TMA consumption relative to a conventional reactor, demonstrating that targeted hardware modifications can yield concurrent gains in process performance and manufacturing cost efficiency.

The influence of chamber configuration on reactor efficiency and precursor utilization has also been examined beyond stationary single-wafer systems. Recent work has extended this analysis to spatially moving reactor architectures, where gas-surface interaction dynamics are inherently coupled to substrate motion (Yun et al., 2023). These studies consistently highlight that chamber geometry plays a central role in governing transport efficiency, precursor usage, and overall process economics, underscoring the importance of chamber configuration optimization and digital-twin based reactor design.

High-quality, spatially resolved measurements are needed to calibrate and validate such digital twins. Recent work at NIST (Maslar and Kalanyan, 2025) describes a camera-based absorption imaging system that visualizes molybdenum pentachloride (MoCl_5), a widely applied precursor to deposit molybdenum contained materials (Juppo et al., 1998; Lubnin et al., 2022; Yoshikawa and Kikuchi, 1996), to be transported in an optically accessible vapor-deposition chamber at 93fps with 0.13 mm/pixel sampling. Beyond demonstrating instrumentation details (telecentric imaging and 443 nm LED illumination), they demonstrate that modest flow/pressure conditions produce a gradually advancing, relatively uniform precursor front with short residence time, whereas higher carrier flow and chamber pressure generate a high-velocity jet that impinges on the chuck and creates recirculation zones, lengthening residence time and compromising uniformity. They also note that such 2D image sequences are well suited for training and validating CFD, and that multi-view tomography could enable full 3D precursor reconstructions.

Building on these developments, this work reports on a collaboration with NIST and uses their optically accessible MoCl_5 ALD process as an experimental anchor to construct a physics-based CFD digital twin. Using upstream optical measurements, time-resolved absorption-imaging data and reported operating conditions, a transient three-dimensional model of MoCl_5 delivery is developed and validated against experimentally observed transport behavior. The model is then simulated across multiple operating conditions to examine key transport features governing process performance, including wafer exposure, jet-induced recirculation, residence time, and spatial uniformity. These results establish a validated digital twin foundation suitable for predictive analysis and future coupling with surface-reaction kinetics and process optimization workflows.

2. Chamber geometry and macroscopic CFD model

2.1. Reactor geometry and computational domain

The computational model is based on an optically accessible, perpendicular-flow, single-wafer atomic layer deposition (ALD) process developed at the NIST for time-resolved gas-phase diagnostics (Kimes et al., 2012). The chamber was specifically designed to minimize gas residence time while preserving optical access for in-situ absorption imaging by camera, making it well-suited as an experimental anchor for CFD-based digital-twin development. Similar optical-access reactor configurations enabling visualization of internal flow and reaction phenomena have also been reported in Oikawa et al. (2022) and Sakai et al. (2023).

The chamber consists of three primary components assembled along the vertical flow direction: an expansion cone, a cylindrical chamber body, and an exhaust manifold integrated with a centrally mounted wafer chuck. Precursor and carrier gases are introduced through four symmetrically arranged stainless-steel injection lines located at the top of the expansion cone. Each injection line has an internal diameter of approximately 4.6–4.8 mm (quarter-inch outer diameter tubing), and the four inlets are evenly distributed azimuthally to promote circumferential flow uniformity prior to entering the main chamber volume.

Immediately downstream of the inlets, the expansion cone transitions from a 50 mm inner diameter at the inlet plane to a 102 mm inner diameter over a length of approximately 110 mm. This gradual expansion reduces inlet jet momentum and promotes lateral spreading of the incoming flow before it enters the chamber body. In the experimental system, the expansion cone is actively heated to suppress precursor condensation and maintain thermal uniformity.

The reactor body is an aluminum cuboid with external dimensions of approximately 152 mm \times 152 mm \times 94 mm and a 102 mm-diameter cylindrical bore running axially through its center. Three recessed diagnostic ports are located on a single horizontal plane along the reactor body, providing optical access across the full diameter of the reactor. These ports define the imaging plane used in the NIST absorption measurements and are therefore explicitly included in the computational domain to preserve the correct internal volume and flow path. A fourth port on the same plane serves as a loading access for wafer installation.

Downstream of the reactor body, the flow enters the exhaust manifold, which also houses the wafer chuck. The 76.2 mm-diameter chuck supports a 50.4 mm-diameter wafer positioned coaxially within the reactor. Gas exits the reactor through four radially symmetric exhaust tubes with an internal diameter of 10.2 mm, located in a recessed region surrounding the wafer chuck. This exhaust configuration strongly influences flow symmetry, vortex formation, and gas residence time near the wafer surface.

The computational domain includes the full internal gas volume of the reactor, encompassing the inlet lines, expansion cone, reactor body, wafer region, and exhaust manifold up to the outlet boundaries. External components such as upstream valves, bubblers, and downstream pumping lines are not modeled explicitly; instead, their influence is incorporated through time-dependent inlet conditions and pressure-based outlet boundary conditions consistent with the experimental configuration. This approach preserves geometric fidelity within the reactor while maintaining computational tractability for transient simulations. A schematic of the reactor geometry and the computational mesh used in the CFD model is shown in Fig. 1.

2.2. Governing equations and species transport model

The macroscopic transport of momentum, mass, and chemical species within the ALD reactor was modeled using transient computational

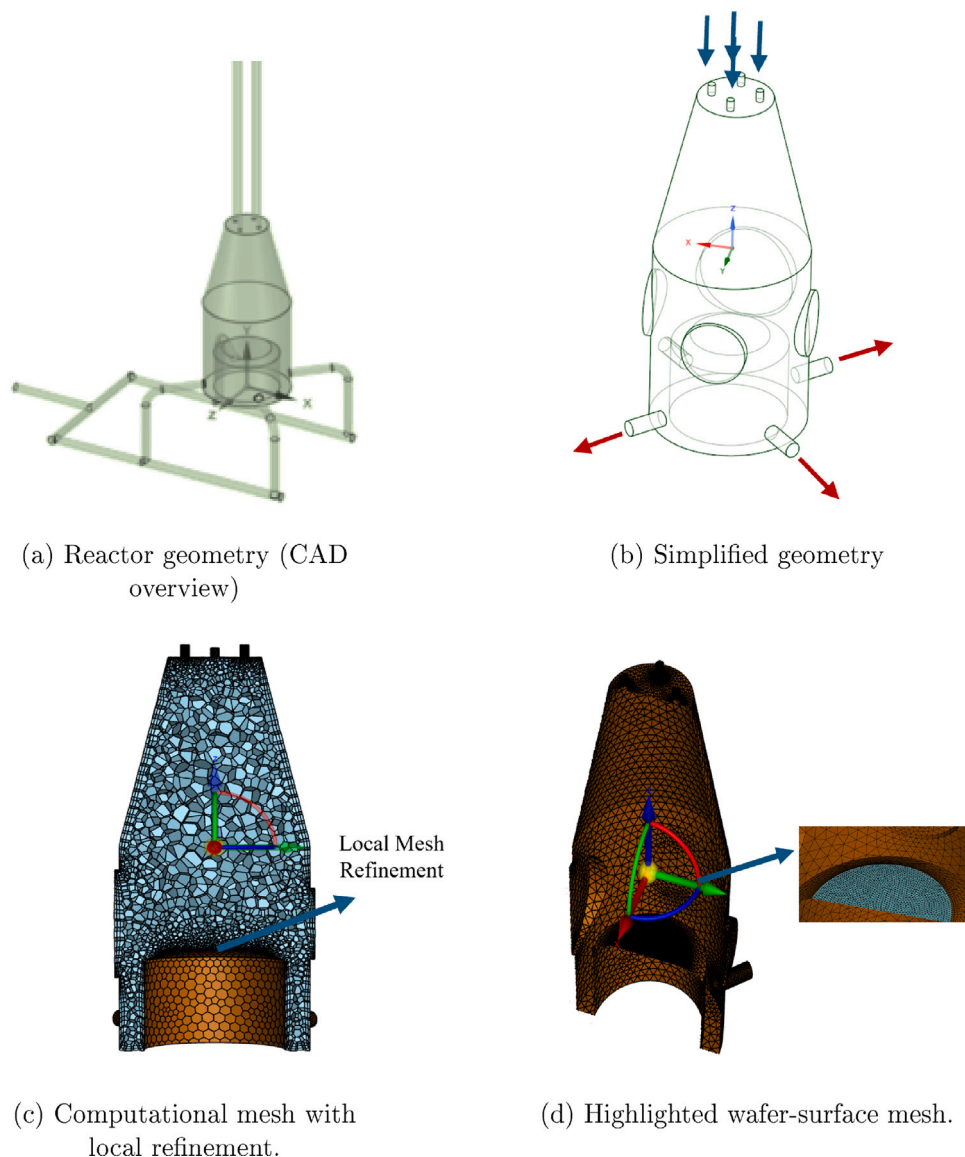


Fig. 1. Reactor geometry and computational mesh used in the CFD model. Top row: geometry views. Bottom row: overall mesh and local refinement near the wafer region.

fluid dynamics (CFD). The gas-phase flow was described by the conservation equations for mass and momentum, coupled with a species transport equation to resolve the spatiotemporal evolution of the precursor concentration. The model focuses exclusively on reactor-scale transport phenomena and intentionally excludes gas-phase chemical reactions, consistent with the separated pulsed injection of precursors and self-limiting feature of atomic layer deposition/etching process.

2.2.1. Mass and momentum conservation

The carrier and precursor gases were treated as a compressible, Newtonian fluid mixture. Conservation of mass is expressed as

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0,$$

where ρ is the gas density and \mathbf{u} is the velocity vector. Momentum transport is governed by the Navier–Stokes equations,

$$\frac{\partial(\rho \mathbf{u})}{\partial t} + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u}) = -\nabla p + \nabla \cdot \boldsymbol{\tau} + \rho \mathbf{g},$$

where p is the static pressure, \mathbf{g} is the gravitational acceleration vector, and $\boldsymbol{\tau}$ is the viscous stress tensor, given for a Newtonian fluid by

$$\boldsymbol{\tau} = \mu [\nabla \mathbf{u} + (\nabla \mathbf{u})^T] - \frac{2}{3} \mu (\nabla \cdot \mathbf{u}) \mathbf{I},$$

with μ denoting the dynamic viscosity and \mathbf{I} the identity tensor. Although the flow velocities are moderate, gas compressibility was retained due to the low operating pressures characteristic of ALD reactors.

The velocity field, \mathbf{u} , obtained from the coupled mass and momentum equations is subsequently post-processed to compute derived observables used in validation, such as wafer-averaged speed and inlet velocity time histories.

2.2.2. Thermodynamic closure and density model

The energy equation was solved with isothermal wall boundary conditions prescribed at the experimental reactor setpoint (393 K). Although the reactor is externally insulated, the reactor walls are actively temperature controlled and thermally massive compared to gas within the reactor in the NIST system; therefore, an isothermal wall condition provides an appropriate representation of the experimental thermal environment.

Gas density was evaluated using the ideal gas equation of state,

$$\rho = \frac{p}{RT},$$

where R is the specific gas constant of the mixture and T is the absolute temperature. This formulation allows density to vary with

Table 1

Mesh-sensitivity setup and quality summary. The 8 mm mesh (baseline) is listed with its reported mesh metrics; across the full mesh sweep (100–4 mm), all meshes maintained acceptable quality, and the total cell count increased by approximately 40% from the coarsest to the finest mesh.

Max cell length (mm)	Growth rate	Total cells (k)	Min orth. quality	Max aspect ratio
100	3.00	240	0.280	24.0
50	3.00	253	0.280	24.0
12	1.25	274	0.219	24.0
8 (baseline)	1.20	279	0.290	25.1
6	1.15	301	0.260	24.5
4	1.10	343	0.240	24.3

both pressure and temperature and is appropriate for the low-pressure (order of 10^2 – 10^3 Pa) operating conditions considered in this work. Because the wall temperature is fixed at the experimental setpoint, density variations arise primarily from transient pressure fluctuations associated with pulsed injection and flow redistribution.

2.2.3. Species transport model

Transport of the precursor species (MoCl_5) was modeled using the transient convection–diffusion equation for the species mass fraction Y_i :

$$\frac{\partial(\rho Y_i)}{\partial t} + \nabla \cdot (\rho \mathbf{u} Y_i) = \nabla \cdot (\rho D_i \nabla Y_i),$$

where D_i is the effective binary diffusivity of MoCl_5 in the argon carrier gas. In ANSYS Fluent, D_i was specified as a constant using the constant–dilute–approx option. The nominal value $D_i = 10^{-4}$ m²/s was selected as an order-of-magnitude estimate consistent with typical gas-phase binary diffusivities for heavy vapor species in inert carrier gases at elevated temperature (Bird et al., 2002). Due to limited validated transport-property data for the MoCl_5/Ar system at 393 K, the diffusivity was treated as an effective parameter rather than being calculated from first-principles correlations.

To assess the importance of diffusion, D_i was varied over two orders of magnitude (10^{-4} – 10^{-2} m²/s). No measurable changes were observed in the reactor-scale transport behavior or in the wafer chuck-level flow/transport trends reported in this work. This confirms that macroscopic transport under the present operating conditions is convection-dominated and that the primary conclusions are not sensitive to the exact value of D_i . Diffusion was retained for physical completeness and to support future extensions that couple gas-phase transport to near-surface concentration gradients.

2.3. Model meshing - sensitivity study

A mesh-sensitivity (grid-independence) study was performed to verify that the predicted flow and species transport results are not influenced by spatial discretization. A set of meshes spanning a wide resolution range was generated, with maximum cell lengths varying from 100 mm (coarsest) to 4 mm (finest). As the mesh is refined, spatial resolution and total cell count increase, while all meshes maintain acceptable orthogonal quality and aspect ratio metrics, as summarized in Table 1. The orthogonal quality and aspect ratio of all candidate meshes in the grid-independence study meet the “good” criteria defined in the ANSYS theoretical guide (ANSYS, 2021).

To assess mesh dependence, three representative quantities were examined: (i) wafer-averaged MoCl_5 mass fraction, (ii) maximum MoCl_5 mass fraction on the wafer, and (iii) inlet velocity. Fig. 2 compares the temporal evolution of these quantities across all mesh resolutions. The results are indistinguishable across the mesh sweep, with only minor deviations confined to the transient peak.

To further verify the adequacy of the near-wall resolution, the non-dimensional wall distance y^+ was evaluated for the baseline mesh. Two wall zones are defined in the model: the wafer surface and the reactor wall, which includes the chamber body. Area-weighted statistics show that the reactor wall exhibits an average y^+ of 0.50 with a maximum value of 1.50, while the wafer surface maintains an average y^+ of

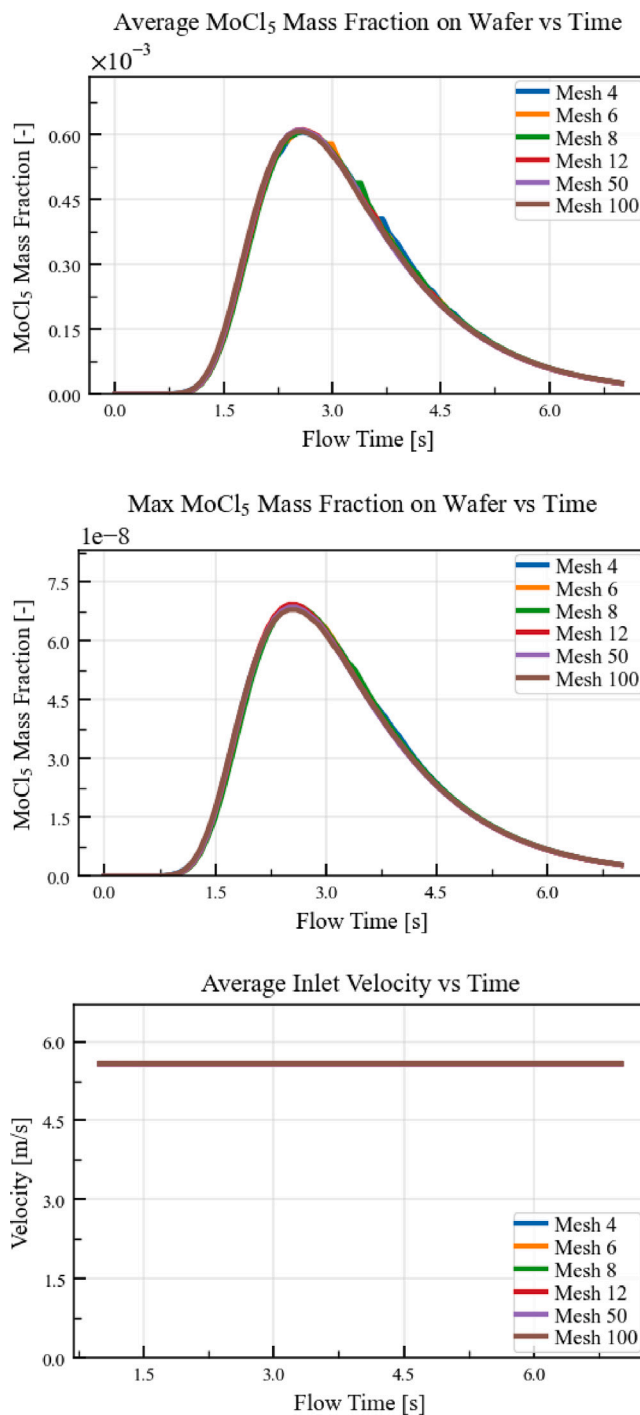


Fig. 2. Mesh-sensitivity results for representative flow and transport observables across the full mesh sweep.

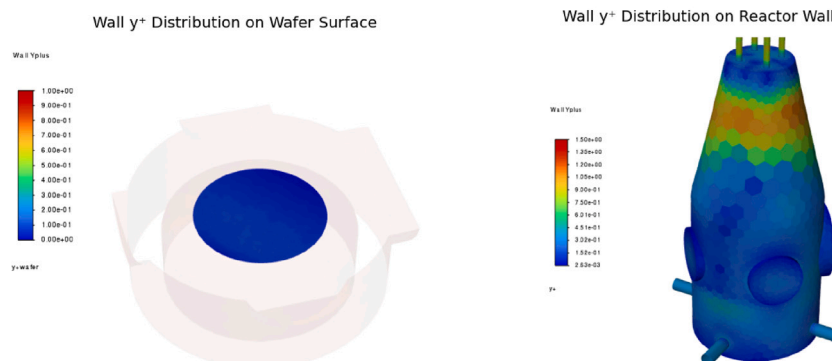
Near-Wall Resolution (y^+) Contours

Fig. 3. Distribution of wall y^+ values across the reactor wall and wafer surface for the baseline mesh. Most reactor surfaces remain below $y^+ \approx 1$, while the wafer surface stays well within the viscous sublayer regime.

0.07 with a maximum of 0.10. These values fall within best-practice guidelines for simulations employing the Scale-Adaptive Simulation (SAS) turbulence model based on the $k-\omega$ SST formulation, where resolving the viscous sublayer typically requires $y^+ \lesssim 1$ on solid boundaries (Menter et al., 2003; ANSYS, 2021). The wafer surface in particular remains deeply within the viscous sublayer regime across the entire domain, ensuring accurate resolution of near-wall shear and species transport. The spatial distribution of y^+ across the reactor walls and wafer surface is shown in Fig. 3, confirming that the baseline mesh provides consistently fine near-wall resolution throughout the reactor.

This limited sensitivity to mesh refinement is expected, as the present model focuses on flow and passive species transport without gas-phase reactions or strongly mesh-dependent source terms. Based on these results, a maximum cell length of 8 mm provides mesh-independent predictions while maintaining adequate near-wall resolution, as confirmed by the y^+ analysis discussed above

2.4. Model settings and assumptions

The governing equations were solved using a pressure-based, transient solver suitable for low-Mach-number compressible flows. Based on characteristic velocities in the domain and argon/ MoCl_5 properties at 393 K, the Mach number remains in the range $M \sim 10^{-5}$ – 10^{-2} across all cases. Species transport was enabled without volumetric chemical reactions, and the precursor species was introduced through time-dependent inlet boundary conditions corresponding to the experimentally prescribed pulsing sequences.

All solid reactor surfaces, including the expansion cone, reactor body, wafer chuck, and exhaust manifold, were modeled as no-slip walls with prescribed constant temperatures at 393 K consistent with the experimental operating conditions, including insulated external and actively controlled wall temperature. Thermal gradients within the gas phase were therefore weak, and buoyancy effects were secondary compared to forced convection induced by the inlet jets and exhaust flow. Gravity was included in the simulations but was not a dominant driver of the flow structure under the examined conditions.

The modeling framework focuses on reactor-scale gas-phase transport within the chamber. Surface adsorption, desorption, and reaction kinetics are intentionally omitted in the present study in order to isolate and validate the macroscopic flow and species transport behavior against experimental measurements. In atomic layer deposition and etching processes, surface reactions are self-limiting and typically consume approximately a monolayer of precursor species per cycle. Consequently, under the convection-dominated operating conditions considered here and the relatively high carrier-gas flow rates used in the experiments, the total amount of precursor removed by surface

reactions represents a negligible fraction of the gas-phase inventory within the reactor volume. As a result, the feedback effect of surface uptake on reactor-scale hydrodynamics, exposure profiles, and gas residence times is expected to be minimal.

This modeling approach is consistent with common multiscale ALD strategies in which reactor-scale CFD is first validated independently before being coupled to microscopic surface-reaction models. The present transport-focused framework therefore provides the macroscopic thermofluidic environment required for future integration of adsorption–reaction kinetics and wafer-scale growth models, which will enable more detailed investigation of precursor utilization and process optimization.

2.5. Flow regime and turbulence modeling

The flow regime within the ALD reactor spans laminar to low-Reynolds-number transitional behavior, depending on the inlet flow rate and operating pressure. Characteristic Reynolds numbers were estimated based on the inlet tube diameter and mean inlet velocity. For the lower-flow cases (e.g., 100 SCCM), the Reynolds number was approximately $\text{Re} \approx 30$ – 40 , while for higher-flow conditions (e.g., 500 SCCM), Re increased to approximately 150–200. These values place the flow in a regime where laminar core flow coexists with localized unsteady vortex dynamics generated by inlet jet impingement and geometric expansion.

To ensure consistent treatment across all operating conditions and to enable automated parametric studies, the Scale-Adaptive Simulation (SAS) turbulence model was employed for all simulations to avoid the potential disadvantages of pure laminar flow simulation (Kim et al., 2003). The SAS approach bridges Reynolds-averaged Navier–Stokes (RANS) and large-eddy simulation (LES) behavior by dynamically adjusting the resolved turbulence length scales based on local flow instabilities and grid resolution. As a result, SAS naturally collapses toward laminar behavior in low-Reynolds-number, steady regions while resolving large-scale unsteady vortical structures where they develop.

For low-flow cases, the SAS model effectively reproduces laminar flow behavior, yielding symmetric precursor plumes and residence-time characteristics consistent with experimental observations. For high-flow cases, the SAS formulation resolves unsteady jet-induced vortices and recirculation zones that strongly influence precursor transport near the wafer and exhaust regions. This capability is critical for accurately capturing plume symmetry, transient concentration buildup and decay, and residence-time distributions under transitional flow conditions.

Standard steady RANS models, such as the $k-\omega$ SST and Transition SST models, were evaluated during model development but were found to excessively damp unsteady vortex dynamics in the high-flow regime,

Table 2

Boundary conditions for baseline validation cases. Common conditions: uniform wall temperature (393 K), no-slip walls, inert wafer (no surface reactions), initial $Y_{\text{MoCl}_5} = 0$ throughout the domain and peak MoCl_5 inlet mass fraction = 3% during pulse.

Flow rate (SCCM)	Outlet pressure P_2 (Pa)	MoCl_5 pulse duration (s)
100	317.6	0.2
500	872.5	0.2

leading to artificial asymmetry and discrepancies with experimental absorbance measurements. By contrast, the SAS model preserves large-scale flow unsteadiness while maintaining numerical robustness near walls, making it well suited for the unified, transient CFD framework adopted in this study.

Based on these considerations, the SAS model was selected as a single turbulence closure for all cases to ensure physical consistency and numerical stability across the full range of flow rates and pulse conditions examined.

2.6. Boundary and initial conditions

Carrier gas and precursor were introduced through the four inlet lines located at the top of the expansion cone. The inlets were modeled using mass-flow boundary conditions derived from the specified volumetric flow rates and inlet line geometry. Pulsed operation was represented by prescribing a time-dependent inlet MoCl_5 mass-fraction profile, constructed from the experimentally measured upstream absorbance trace for each case.

The total inlet mass flow rate was held constant during pulsed operation, consistent with the specified carrier-gas flow conditions. The inlet MoCl_5 mass fraction followed the measured transient profile and was scaled to a peak value of 3% for both baseline cases. Outside the precursor pulse, the inlets supplied pure carrier gas to represent purge conditions. Thus, pulsing was implemented as a time-dependent compositional variation at fixed total mass flow, rather than by adding mass to the system.

The reactor exhaust was modeled using pressure outlet boundary conditions applied at the downstream ends of the exhaust manifold. The outlet pressure was set to the experimentally measured chamber pressure, denoted as P_2 , for each operating condition. The downstream exhaust-line configuration – where the pressure P_3 is measured experimentally – was not included in the primary CFD domain. However, the pressure drop between P_2 and P_3 was quantified separately using dedicated inlet–outlet flow models to confirm that prescribing P_2 as the outlet boundary condition accurately represents the reactor-scale pressure field. Details of the pressure measurement locations and absorbance-based precursor diagnostics are provided in Sections 3.1 and 3.3. To account for potential outlet backflow arising from pressure gradients and turbulent fluctuations, backflow species fractions were prescribed to match the carrier gas composition, thereby promoting numerical stability during transient flow reversals.

Although the outlet boundary condition was prescribed using the measured P_2 , the simulations were fully transient and compressible. Any transient pressure redistribution within the chamber arises naturally from the unsteady jet dynamics and flow recirculation during pulsed operation rather than from an imposed increase in total mass flow.

The initial condition for each simulation corresponded to a fully purged reactor filled with carrier gas at the specified operating pressure and temperature, with zero precursor mass fraction throughout the domain. Precursor injection was initiated after a brief purge period to establish a stable baseline flow field prior to pulsing.

Two representative operating conditions, corresponding to 100 SCCM and 500 SCCM total flow rate, were first simulated to validate the CFD model against experimental absorption measurements. Subsequently, the same boundary-condition framework was extended to a broader set of flow rates, pulse durations, and pressures using a unified,

parametric CFD model. Details of the validation cases and parametric study are presented in Section 4 (see Table 2 for boundary conditions of the baseline validation cases).

3. Results and analysis

3.1. Experimental reference and comparison metrics

CFD results are evaluated against time-resolved experimental measurements from the optically accessible NIST ALD reactor to assess whether the model reproduces the dominant transport behavior governing transient MoCl_5 delivery. The objective of the model is to capture precursor transport dynamics for transport diagnostics and pulse timing analysis. Reactor design comparison is not considered here since the geometry corresponds directly to the experimental reactor configuration, although the framework can be extended to such studies.

Two operating conditions are used as reference cases: 100 SCCM and 500 SCCM, with chamber pressure set to the experimentally measured values. These cases span distinct transport regimes from laminar-dominated flow to laminar-to-transitional behavior and provide a representative basis for model assessment. Validation focuses on comparison of transport behavior and characteristic time scales, including the buildup and decay trends as well as derived quantities such as residence time and inferred inlet pressure and velocity. For the intended diagnostic and pulse scheduling applications, agreement within approximately 20% is considered sufficient to capture the relevant transport dynamics.

The experimental absorbance traces are obtained from optical imaging measurements and therefore include inherent measurement uncertainty. As reported in the experimental study, the baseline noise level of the absorbance signal is approximately $\sigma_{bl} \approx 1 \times 10^{-3}$ absorbance units, while slow baseline drift associated with window deposition during operation is on the order of $(1-2) \times 10^{-4}$ absorbance. In addition, the temporal resolution of the measurements is limited by the camera acquisition rate of approximately 93 fps, corresponding to a timing resolution of $\Delta t \approx 0.011$ s. These uncertainties primarily affect the absolute amplitude and precise timing of the measured signals but do not alter the overall transport behavior observed in the buildup and decay dynamics.

The experimental diagnostics are summarized in Fig. 4 and comprise three complementary measurement approaches: (i) pressure measurements, (ii) upstream in-line absorbance monitoring, and (iii) in-chamber optical absorbance measurements.

Pressure characterization was performed using capacitance diaphragm gauges located at multiple positions within the tool. In particular, P_1 denotes the upstream pressure measured in the inlet line prior to chamber entry, while P_2 corresponds to the chamber pressure measured near the reactor exhaust. An additional gauge (P_3) was located downstream on the exhaust line. These measurements were used to characterize pressure drops through the system and to infer inlet velocity under various operating conditions.

MoCl_5 absorbance was monitored upstream of the chamber using a single-element in-line detector positioned in the inlet line. This measurement captures the transient injection waveform delivered to the reactor and serves as the basis for prescribing the inlet mass-fraction boundary condition in the CFD model.

Within the chamber, two types of optical absorbance measurements were performed. First, time-resolved absorbance signals were extracted at three single-pixel locations along the reactor centerline (Points A, B, and C above the wafer chuck in Fig. 9) as well as over a cylindrical

Experimental Instrumentation Configuration – Pressure and Optical Measurements

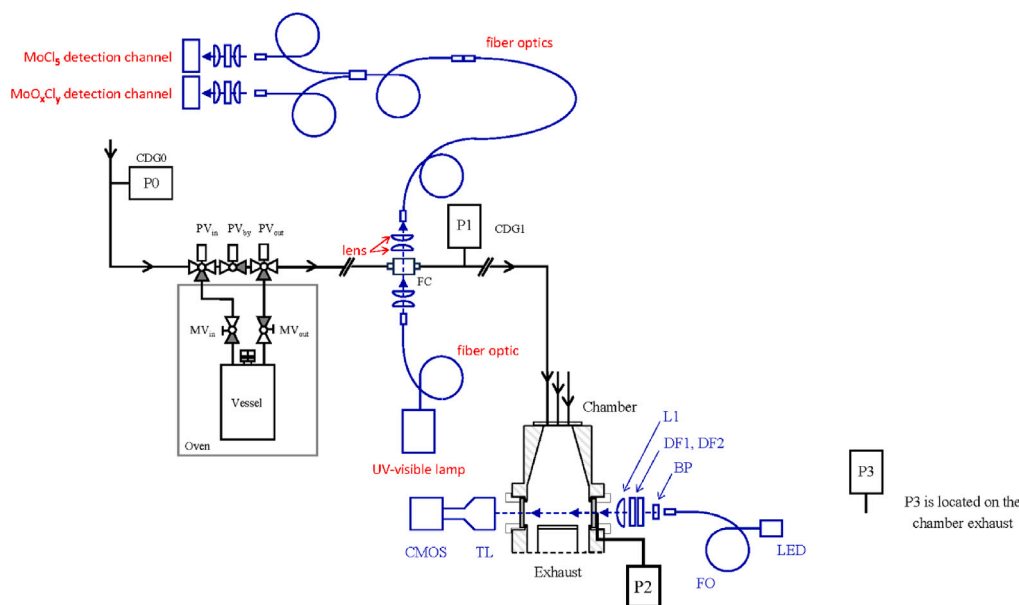


Fig. 4. Integrated experimental measurement configuration used in the NIST MoCl_5 ALD reactor. Capacitance diaphragm gauges measure upstream pressure ($P1$), chamber pressure ($P2$), and exhaust pressure ($P3$). Optical diagnostics include (i) upstream in-line absorbance detection of the injection waveform and (ii) in-chamber absorption imaging through the optical window using LED illumination, bandpass filters, a telecentric lens, and a CMOS camera.

region of interest (20 mm diameter, 106 mm optical path length) to obtain a chamber-averaged signal. Second, full-field absorption imaging was conducted through the optical window using LED illumination, a telecentric lens, bandpass filters, and a CMOS camera, enabling spatially resolved visualization of transient plume evolution.

The absorbance diagnostics provide signals proportional to local MoCl_5 partial pressure but are not calibrated to absolute concentration values. Accordingly, all model-experiment comparisons emphasize relative transport behavior and temporal response rather than absolute concentration magnitude. This includes comparison of precursor arrival timing, buildup and decay dynamics, residence time, spatial distribution, plume symmetry, and wafer-level exposure trends.

Model performance is assessed using the following complementary observables:

- upstream pressure ($P1$), chamber pressure ($P2$), and inferred inlet velocity;
- upstream in-line absorbance waveform used for inlet forcing;
- time-resolved, normalized absorbance signals at Points A, B, C, and the chamber-averaged cylindrical region;
- residence time inferred from decay behavior during purge;
- spatial structure and symmetry of transient precursor plumes from optical imaging;
- wafer chuck-level precursor transport behavior from CFD.

These metrics provide a structured and consistent framework for evaluating the model across operating conditions. After establishing agreement for the 100 SCCM and 500 SCCM reference cases, the same framework is applied without modification to a broader parametric study involving variations in flow rate, pressure, and injection duration, as discussed in Section 4.

3.2. Inlet conditions and flow field behavior

Accurate specification of inlet boundary conditions and the resulting flow field is essential for reliable prediction of precursor transport

within the reactor. Accordingly, inlet conditions and flow behavior are first evaluated for the 100 SCCM and 500 SCCM reference cases, which serve as the baseline for subsequent transport analysis. Subsequently, results for the 300 SCCM case are also presented in Section 4.2.

To faithfully represent the experimentally observed injection dynamics while maintaining numerical tractability, the inlet MoCl_5 mass-fraction profile was implemented in the CFD model as a piecewise-linear function derived from time-resolved upstream absorbance measurements. This approach preserves the key temporal features of the experimental signal – including finite rise time, peak duration, and rapid decay – while providing a stable and computationally efficient boundary condition. The normalized absorbance signal was therefore approximated using a piecewise-linear fit and mapped to a time-dependent inlet mass-fraction profile, with the peak value scaled to 3% for both baseline cases.

Fig. 5 presents a comparison between the measured absorbance signal and the piecewise-linear inlet profile imposed in the simulations for the 100 SCCM and 500 SCCM cases. The agreement demonstrates that the constructed boundary condition captures the dominant temporal characteristics of the experimental injection, thereby ensuring that the subsequent flow and transport simulations are driven by an inlet representation consistent with experimental behavior.

The resulting reactor-scale flow topology is illustrated using pathline visualizations in Fig. 6 for the 100 SCCM case. For both flow rates, the pathlines reveal a consistent global flow structure characterized by jet expansion within the inlet cone, downward transport toward the wafer region, and symmetric recirculation beneath the wafer. Differences between operating conditions arise primarily in velocity magnitude and the degree of local unsteadiness rather than in the underlying flow pattern.

Velocity magnitude contours on a representative vertical mid-plane are shown in Fig. 7. For both the 100 SCCM and 500 SCCM cases, elevated velocities are confined to the immediate vicinity of the inlet exits and rapidly decay within the expansion cone, resulting in relatively low

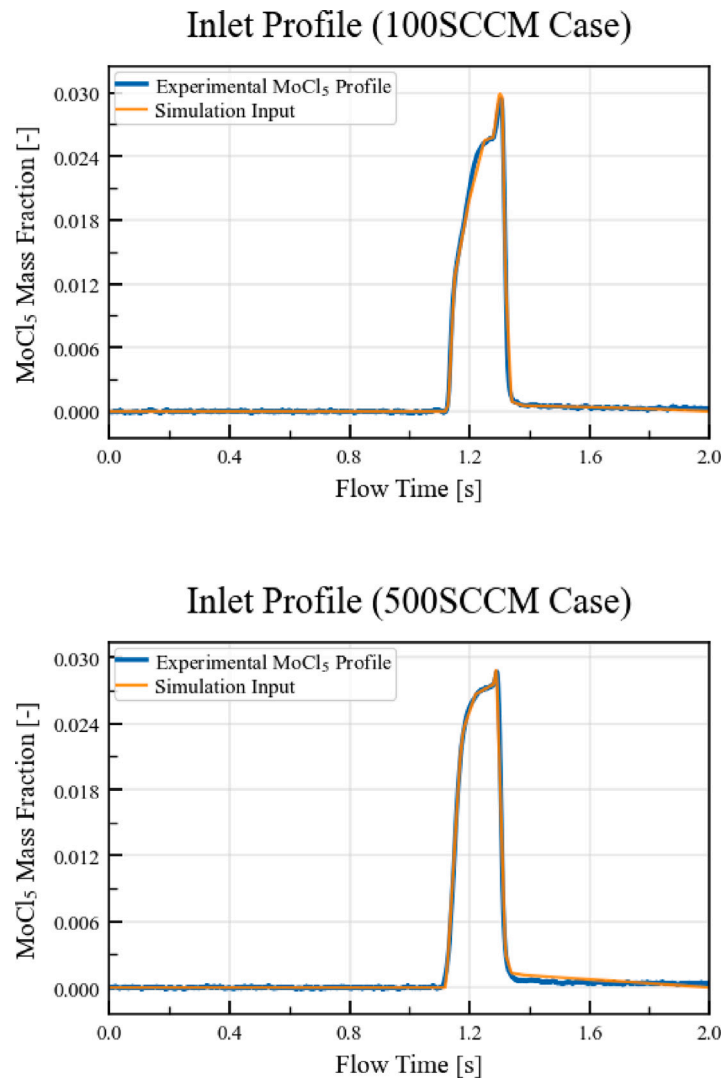


Fig. 5. Experiment-informed inlet forcing for the baseline cases: comparison of upstream absorbance signal and prescribed inlet MoCl₅ mass-fraction profile used as CFD input.

Evolution of Velocity Pathlines Colored by Velocity Magnitude – 100 SCCM Case




Fig. 6. Pathline visualizations of the reactor-scale flow field colored by velocity magnitude for 100 SCCM case. The overall flow topology is similar between the 100 SCCM and 500 SCCM cases.

Velocity Magnitude Contours on the Vertical Mid-Plane




Fig. 7. Velocity magnitude contours on a vertical mid-plane through the reactor for the 100 SCCM and 500 SCCM operating conditions.

velocities throughout the reactor body. Increasing the inlet flow rate primarily amplifies velocity magnitude and enhances local unsteadiness near the inlet region, while preserving the same overall spatial velocity distribution.

For both flow rates, the prescribed inlet boundary conditions produce a high-velocity jet within the injection lines that rapidly decelerates upon entering the expansion cone. The sudden increase in cross-sectional area reduces axial momentum and promotes lateral spreading before the gas enters the main reactor body. This behavior is consistent with the reactor design intent and with experimental observations reported for the NIST system.

At 100 SCCM, the flow field within the reactor body remains steady and symmetric, with smooth streamlines extending from the expansion cone toward the wafer region and exhaust outlets. No large-scale unsteady vortical structures are observed, and the velocity magnitude near the wafer surface remains low, consistent with laminar-dominated transport. Under these conditions, the simulated velocity field agrees with experimentally inferred inlet and wafer chuck-level velocities.

At 500 SCCM, increased inlet momentum leads to the formation of localized unsteady vortical structures downstream of the expansion cone and near the wafer and exhaust regions. These structures arise from interactions between the inlet jet, geometric expansion, and the symmetric exhaust configuration. The widely applied SAS turbulence model (Menter and Egorov, 2006) resolves this unsteady transition region behavior while preserving overall flow symmetry, avoiding the artificial asymmetry encountered with steady RANS models. Despite increased unsteadiness, the spatially-averaged flow field remains symmetric, consistent with experimental imaging.

The transient response of the inlet flow to precursor injection is quantified in Fig. 8, which shows the area-averaged velocity at Inlet 2 as a function of time. For both flow rates, the onset of MoCl_5 injection at $t = 0.5$ s induces a very small, short-lived perturbation in inlet velocity attributable to numerical transients associated with discretization and the abrupt change in boundary conditions, after which the flow rapidly returns to a quasi-steady value.

Quantitative comparison between experimentally inferred and simulated inlet conditions is summarized in Table 3. Simulated inlet velocities agree with experimental estimates to within approximately 1% for both baseline cases, while inlet-line pressure predictions remain within approximately 3%–5%.

3.3. MoCl_5 buildup and decay dynamics

Time-resolved MoCl_5 transport behavior was characterized experimentally using single-point in-line absorbance measurements upstream of the chamber and spatially resolved absorption imaging within the optically accessible NIST ALD reactor. The in-line absorbance provides time-dependent precursor delivery information, while the in-chamber imaging captures spatial and temporal variations in MoCl_5 partial pressure. Absorbance signals are proportional to the local MoCl_5 partial pressure and are therefore used here in normalized form to assess relative transport behavior and temporal dynamics.

Fig. 9 illustrates the experimental sampling locations used for CFD-experiment comparison. Three single-pixel locations were selected along the reactor centerline at approximately 25.5 mm (Point A), 15.7 mm (Point B), and 5.8 mm (Point C) above the wafer chuck surface. Each pixel corresponds to a pathlength-integrated absorbance over an area of approximately $0.131 \text{ mm} \times 0.131 \text{ mm}$. In addition, a chamber-averaged absorbance signal was obtained by integrating over a cylindrical region of interest with a diameter of 20 mm and an optical pathlength of 106 mm, centered at the same vertical location as Point B.

In the CFD model, analogous quantities were extracted to enable direct comparison. Single-point signals were obtained by sampling the local MoCl_5 mass fraction along the reactor centerline using point probes in ANSYS Fluent. The chamber-averaged signal was computed by volume-averaging the MoCl_5 mass fraction over a cylindrical region matching the experimental field of view, defined within the reactor chamber using a custom volume region consistent with the 20 mm diameter optical contour.

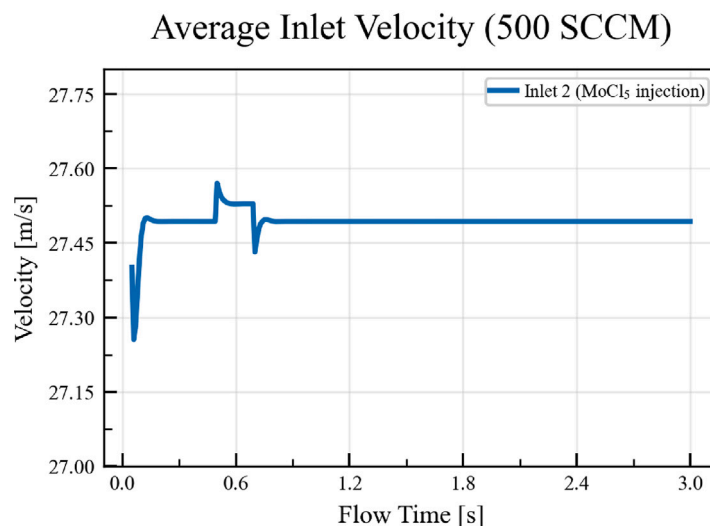
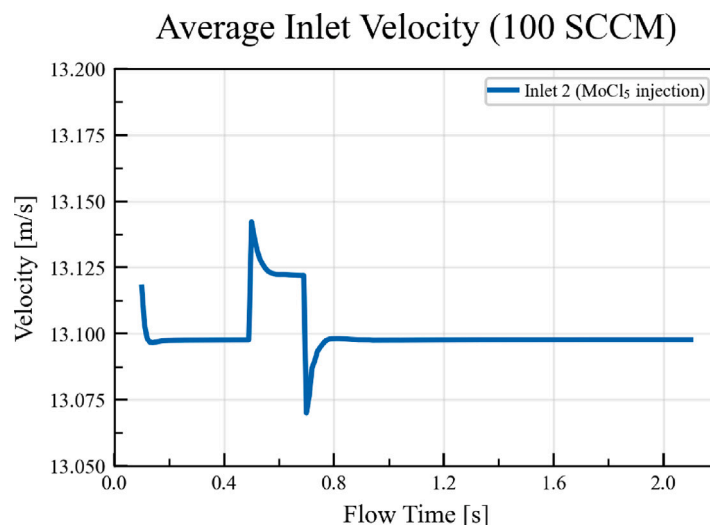


Fig. 8. Time histories of the area-averaged velocity at Inlet 2 for the 100 SCCM and 500 SCCM cases, illustrating the transient response associated with MoCl_5 injection at $t = 0.5$ s.

Table 3

Comparison of experimentally inferred and simulated inlet-line pressure ($P1$) and inlet velocity for the baseline cases. Experimental velocity values are reported at 393 K.

Case	$P1$ Exp. (Pa)	$P1$ Sim. (Pa)	$P1$ Error (%)	Velocity Exp. @393 K (m/s)	Velocity Sim. (m/s)	Velocity Error (%)
0.2s–100 SCCM–317 Pa	998.2	1031.2	+3.3	13.01	13.10	+0.7
0.2s–500 SCCM–872.5 Pa	2371.5	2491.3	+5.1	27.39	27.53	+0.5

Figs. 10 and 11 compare the normalized experimental absorbance signals with the corresponding normalized CFD predictions for the 100 SCCM and 500 SCCM cases, respectively. Results are shown for the chamber-averaged cylindrical region and for the three single-pixel locations (Points A, B, and C).

For both operating conditions, the simulations reproduce the dominant transient behavior observed experimentally, including precursor arrival, peak response, and purge-phase decay. The chamber-averaged signal is captured particularly well for both flow rates, indicating that the model predicts the bulk filling and purge dynamics consistently. Across the point locations, the ordering and decay trends are well represented; however, a modest delay in the experimental rise relative to the simulations is observed, most clearly at the point measurements.

This discrepancy likely reflects a combination of upstream delivery dynamics, near-wall interactions, and geometric details of the physical chamber – such as recesses, seams, and protrusions – not represented in the simplified computational geometry, which can influence local flow development and timing. Overall, the close agreement in buildup and decay behavior across both flow conditions demonstrates that the CFD model captures the primary convection-driven transport time scales within the reactor.

3.4. MoCl_5 residence time

A decay-based residence-time metric was evaluated for the 100 SCCM and 500 SCCM cases. In the NIST imaging study (Maslar and

Absorbance Sampling Locations




Fig. 9. Experimental measurement regions used for absorbance extraction: three single-pixel locations along the reactor centerline (Points A–C) above the wafer chamber for 25.5 mm (A), 15.7 mm (B) and 5.8 mm (C), respectively, as well as a chamber-averaged cylindrical region of interest (20 mm diameter, 106 mm optical pathlength).

Comparison of Normalized Experimental and Simulated MoCl₅ Signals for the 100SCCM Case

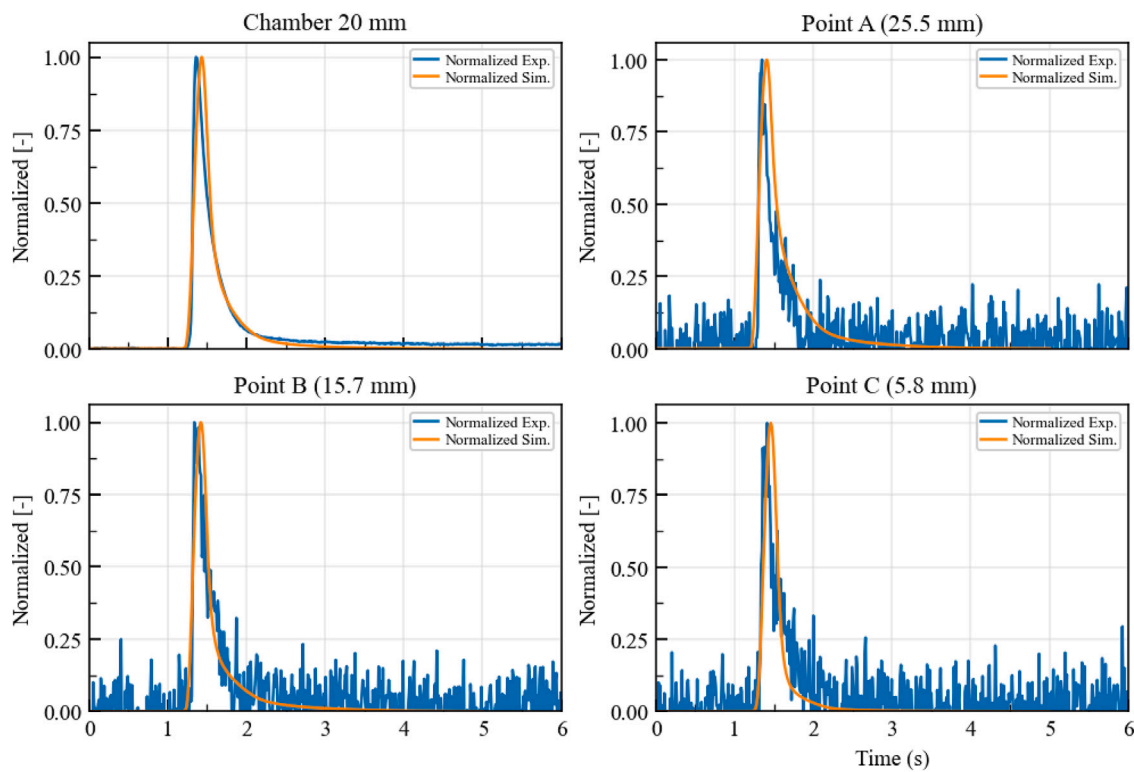


Fig. 10. Comparison between experimentally measured single-pixel data and simulated results for the 100 SCCM case. Despite the higher noise level in the experimental signal, strong agreement is observed with the smoother simulated profile.

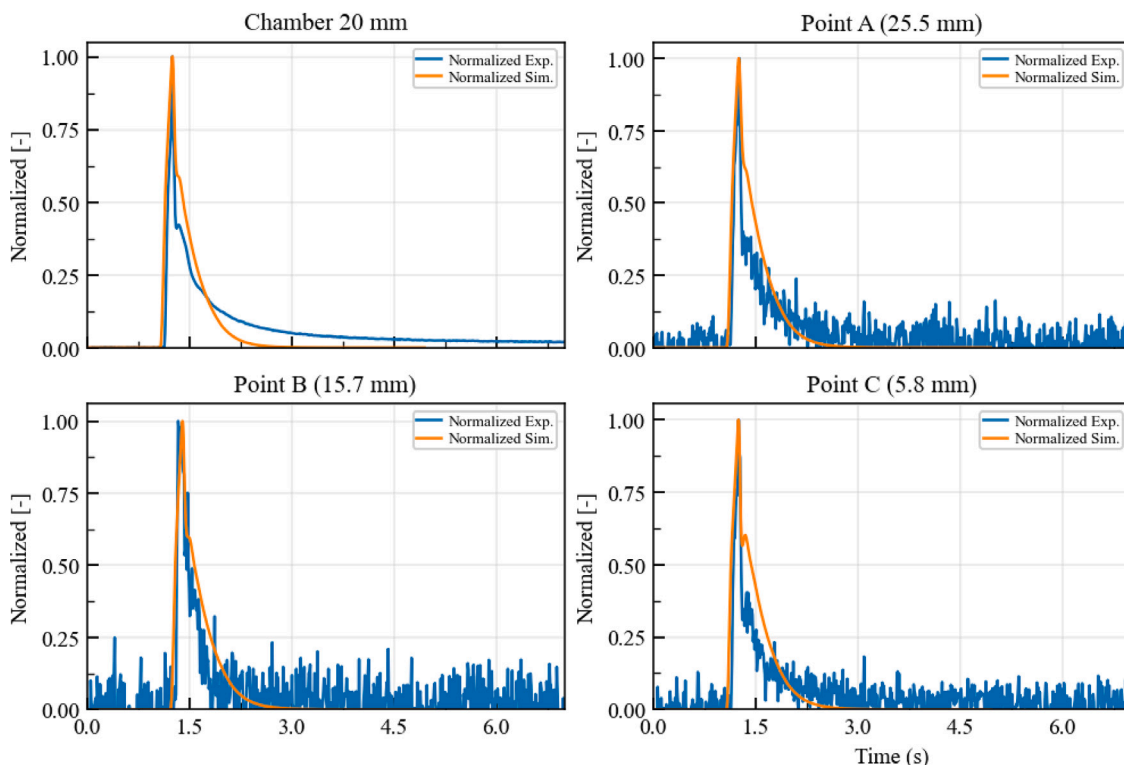
Comparison of Normalized Experimental and Simulated MoCl₅ Signals for the 500SCCM Case

Fig. 11. For the 500 SCCM case, the simulated results exhibit strong agreement with the experimental data during the rising phase. Although minor discrepancies are observed in the decay period, the overall agreement remains close.

Kalanyan, 2025), the reported residence time is inferred from the temporal decay of the absorbance signal observed within the camera field of view following precursor injection, yielding values of approximately 0.8 s for 100 SCCM and 1.4 s for 500 SCCM.

In the CFD model, the closest spatial analogue to this imaged region is a finite volume located above the wafer chuck along the reactor centerline. The MoCl₅ mass fraction was volume-averaged over this region and the residence time was extracted from the simulated decay using a peak-to-1% criterion to provide a consistent and repeatable numerical definition.

Fig. 12 shows the resulting transients. The predicted residence time is 0.75 s for 100 SCCM, in close agreement with the experimental value (difference ~6%). For 500 SCCM, the model yields 1.2 s compared to the experimental 1.4 s (difference ~14%), which remains a reasonable agreement for the high-flow case. Additional LOS-aligned line sampling across the optical depth produced similar values in the range of 1.1–1.2 s, indicating that the discrepancy is not strongly dependent on the post-processing metric. The remaining difference is likely due to the increased sensitivity of the 500 SCCM case to small differences in spatial sampling and local flow structure, particularly under jet-impingement and recirculating conditions, where the late-stage decay is more difficult to match exactly.

3.5. Spatial structure of transient MoCl₅ plumes near the wafer

In addition to time-resolved point and path-integrated measurements, the NIST experiments provide spatially resolved visualization of transient MoCl₅ transport through absorption imaging of the reactor window (Maslar and Kalanyan, 2025). These images reveal the spatial structure, symmetry, and temporal evolution of precursor plumes as they propagate from the inlet region toward the wafer and decay during purge. Such image-based diagnostics complement point measurements

by enabling direct assessment of plume uniformity and large-scale flow symmetry in the wafer vicinity.

To enable comparison with these optical measurements, CFD-predicted MoCl₅ mass-fraction fields were post-processed to generate planar slices that approximate the camera viewing geometry. In the experiments, the telecentric lens and CMOS camera integrate absorbance over the full optical path length through the reactor volume, whereas the CFD post-processing extracts two-dimensional planar fields representing a slice of the three-dimensional domain. Despite this distinction, the extracted planes capture the dominant spatial characteristics of the transient plume – such as symmetry, spreading behavior, and interaction with the wafer region – and are therefore suitable for qualitative and semi-quantitative comparison.

Fig. 13 illustrates the procedure used to extract spatial MoCl₅ fields from the CFD simulations. A set of vertical planes aligned with the optical viewing direction was constructed in the vicinity of the wafer surface to mimic the camera perspective. These planes were further extended upstream toward the inlet region to visualize plume development from injection to wafer impingement. This approach enables consistent spatial comparison between simulated concentration fields and experimentally observed absorbance images, while acknowledging that minor discrepancies may arise from differences between volumetric optical integration and planar CFD sampling.

Using the methodology described above, image-based validation was performed by comparing time-resolved experimental absorbance snapshots with CFD-predicted MoCl₅ fields processed to match the NIST camera view (Maslar and Kalanyan, 2025). For the 100 SCCM case, the comparison is performed directly on a camera-aligned plane with a circular mask corresponding to the reactor window. For the 500 SCCM case, additional vertical planes are included to visualize plume development from the inlet region toward the wafer and to assess plume symmetry under higher-flow conditions.

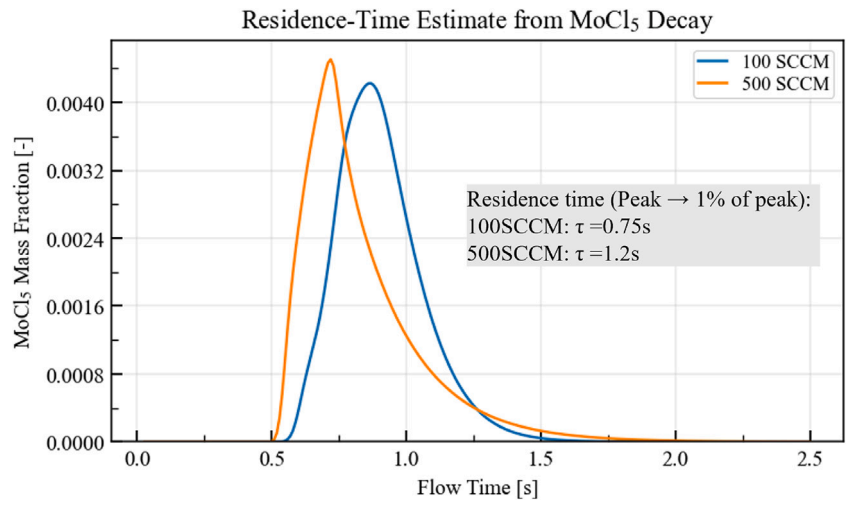


Fig. 12. Simulated MoCl₅ mass-fraction decay in the reactor-center volume above the wafer chuck for 100 SCCM and 500 SCCM. Residence time in the CFD data is defined as the interval from the peak to 1% of the peak value.

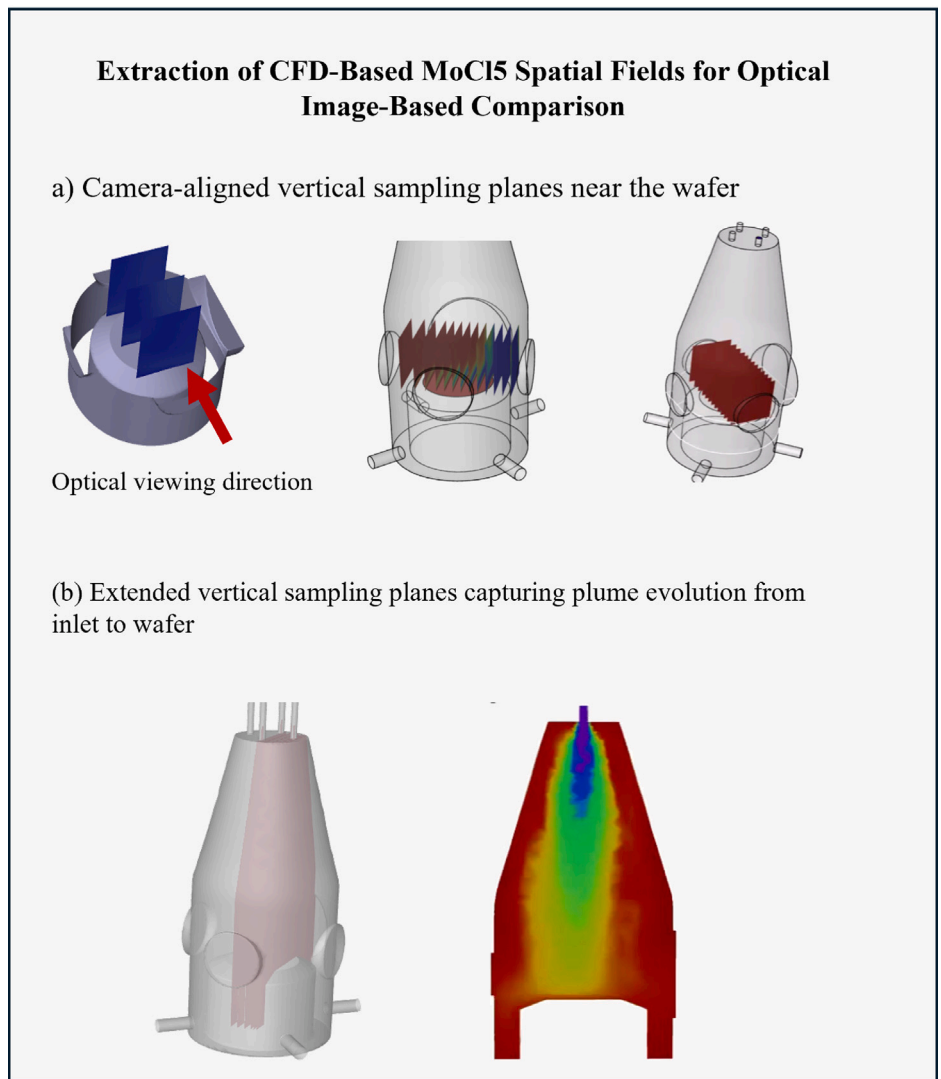
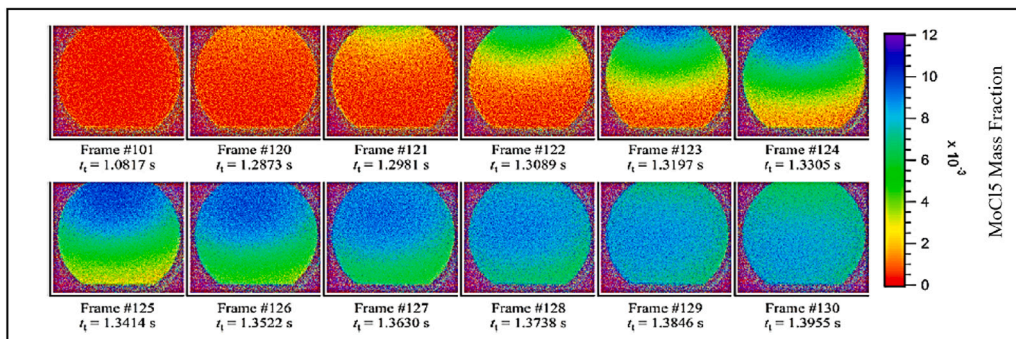


Fig. 13. Extraction of CFD-based MoCl₅ spatial fields for image-based comparison with optical absorbance measurements. Camera-aligned vertical planes near the wafer surface and extended planes toward the inlet are used to capture the dominant plume structure and symmetry observed experimentally.

MoCl₅ Spatial Structure Near the Wafer - 100 SCCM Case

a) Experimental absorbance images (Maslar and Kalanyan, 2025)



b) Simulated MoCl₅ concentration fields near wafer




Fig. 14. Image-based comparison of transient MoCl₅ plume evolution for the 100 SCCM case. **Top:** experimental absorbance snapshots acquired through the reactor window (Maslar and Kalanyan, 2025). **Bottom:** CFD-predicted MoCl₅ fields post-processed on a camera-aligned plane and circularly masked to match the window field of view.

Fig. 14 shows that the 100 SCCM simulation reproduces the main qualitative sequence observed experimentally: initial precursor arrival, development of spatial gradients across the viewing window as the plume spreads toward the wafer region, and subsequent decay during purge with the remaining signal biased toward the lower portion of the field of view. The relative timing of plume evolution and the overall spatial progression are consistent, indicating that the experimentally informed inlet forcing and transport model capture the dominant transient behavior observed in the optical measurements.

For the 500 SCCM condition (Fig. 15), the simulated plume remains globally symmetric about the reactor axis and develops as a coherent core that propagates downstream while diffusing laterally across the vertical planes, consistent with the symmetric structure apparent in the experimental snapshots. The Scale-Adaptive Simulation (SAS) approach preserves this symmetry while resolving the increased unsteadiness associated with higher inlet momentum, providing a physically consistent representation of plume structure under elevated flow conditions.

3.6. Wafer chuck-level velocity and MoCl₅ exposure

In addition to chamber-scale transport behavior, wafer chuck-level conditions play a critical role in determining precursor delivery and, ultimately, film growth characteristics. Two key quantities are examined here: the local gas velocity impinging on the wafer surface and the transient MoCl₅ exposure experienced by the wafer during pulsed injection. Together, these metrics provide complementary insight into flow intensity near the substrate and the effective availability of precursor at the wafer surface.

It is noted that wafer chuck-level precursor exposure depends not only on inlet flow rate, but also on injection duration, chamber pressure, and exhaust configuration. In the present analysis, all comparisons are performed for a fixed pulsing condition corresponding to a 0.2 s MoCl₅ injection, such that observed differences in wafer exposure arise primarily from changes in inlet flow rate.

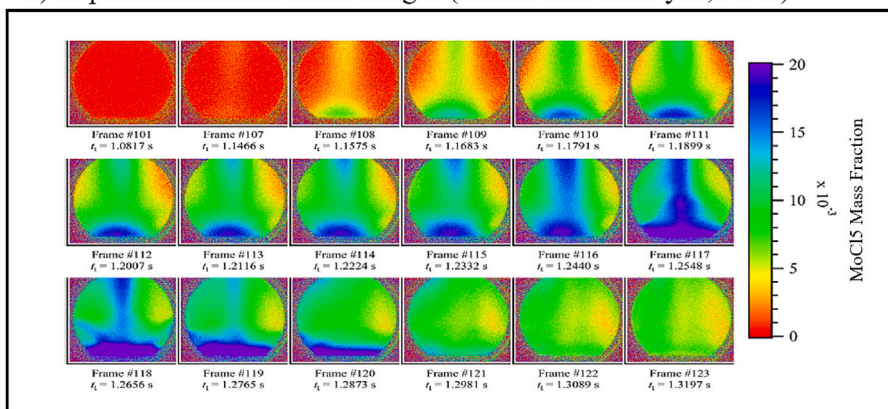
Experimentally, wafer chuck-level velocities were estimated for the 100 SCCM case by NIST based on measured chamber pressures, known inlet geometries, and characteristic transport distances. Using this approach, the wafer velocity was inferred to lie in the range of approximately 0.01–0.1 m/s. In the CFD simulations, wafer chuck-level velocity was quantified by defining a thin control volume immediately above the wafer surface and computing the area-averaged velocity magnitude within this region as a function of time. This approach provides a consistent and physically representative measure of the gas velocity experienced by the wafer.

Fig. 16 compares the simulated wafer chuck-level velocity for the 100 SCCM and 500 SCCM cases. For 100 SCCM, the predicted wafer velocity remains low, on the order of 0.01–0.1 m/s, and falls within the experimentally inferred range, providing additional validation of the near-wafer flow field. In contrast, increasing the inlet flow rate to 500 SCCM results in a substantially higher wafer velocity, reaching approximately 2–3 m/s, corresponding to an increase of roughly one to two orders of magnitude relative to the 100 SCCM case. This increase reflects stronger jet penetration and enhanced convective transport under higher-flow operation.

The corresponding wafer-averaged MoCl₅ exposure is shown in Fig. 17, where the average MoCl₅ mass fraction over the wafer surface is plotted as a function of time. The 100 SCCM case exhibits a sharper

MoCl₅ Spatial Structure Near the Wafer - 500 SCCM Case

a) Experimental absorbance images (Maslar and Kalanyan, 2025)



b) Simulated MoCl₅ concentration fields

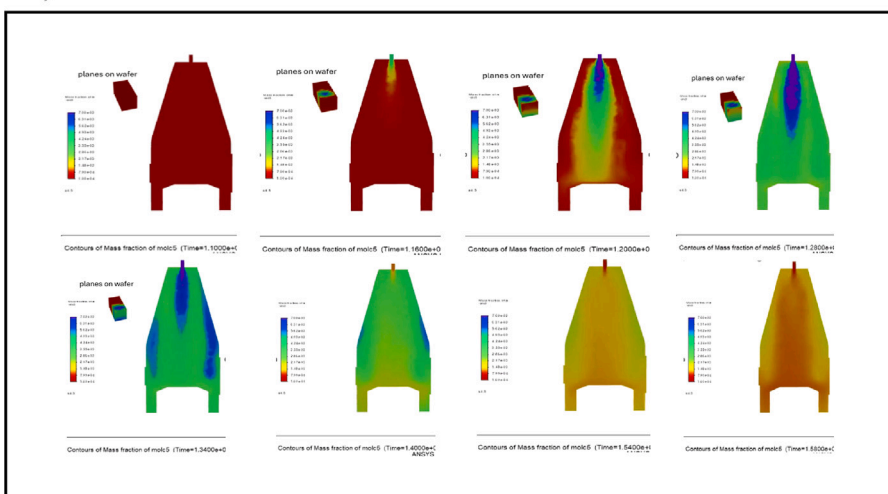


Fig. 15. Image-based comparison of transient MoCl₅ plume evolution for the 500 SCCM case. **Top:** experimental absorbance snapshots (Maslar and Kalanyan, 2025). **Bottom:** CFD-predicted MoCl₅ fields shown on extended vertical planes from the inlet region toward the wafer to visualize plume propagation and symmetry under higher-flow conditions.

buildup and higher peak exposure, followed by rapid decay during purge. At 500 SCCM, the wafer experiences earlier precursor arrival but a lower peak average mass fraction and a more extended decay tail, consistent with enhanced convective transport and more rapid removal of precursor from the wafer vicinity.

Together, these results show that increasing the inlet flow rate significantly alters both the flow intensity at the wafer surface and the temporal characteristics of MoCl₅ exposure. For the fixed 0.2 s injection considered here, higher flow rates lead to earlier arrival, reduced peak wafer-averaged concentration, and stronger convective sweeping of precursor away from the surface. These trends directly influence the spatial distribution of MoCl₅ over the wafer and motivate examination of wafer-scale concentration contours.

Fig. 18 shows the time-resolved spatial distribution of MoCl₅ on the wafer chuck-level surface (the same plane as in Figs. 16 and 17) for the 100 SCCM case, extracted from a wafer-aligned sampling plane in the CFD model. Following the onset of injection, MoCl₅ initially reaches the upstream region of the wafer and spreads rapidly laterally across the surface. As the pulse progresses, concentration gradients diminish and the distribution becomes increasingly smooth and symmetric. During

the decay phase, precursor removal occurs in a spatially coherent manner, with no evidence of localized stagnation or preferential depletion zones.

At intermediate times during the injection pulse, the wafer surface becomes fully enveloped by the MoCl₅ plume, indicating that gas-phase transport is sufficiently rapid to provide complete wafer-scale coverage under these conditions. While this does not imply uniform surface reaction or film thickness, it suggests that macroscopic transport limitations are minimal and that any remaining non-uniformities would be governed primarily by surface kinetics rather than precursor delivery.

4. Parametric study

4.1. Generalized boundary-condition variation

To evaluate robustness of the validated CFD framework across a broader operating envelope, a parametric study was conducted in which only boundary conditions were varied. The objective is not to replicate specific experimental injection waveforms or rank process conditions, but to verify physically consistent model response under variations in pulse duration, flow rate, and chamber pressure,




Fig. 16. Simulated wafer chuck-level velocity as a function of time for the 100 SCCM and 500 SCCM cases. The 100 SCCM case exhibits low velocities consistent with experimentally inferred values, while the 500 SCCM case shows a substantial increase in flow impingement intensity near the wafer.

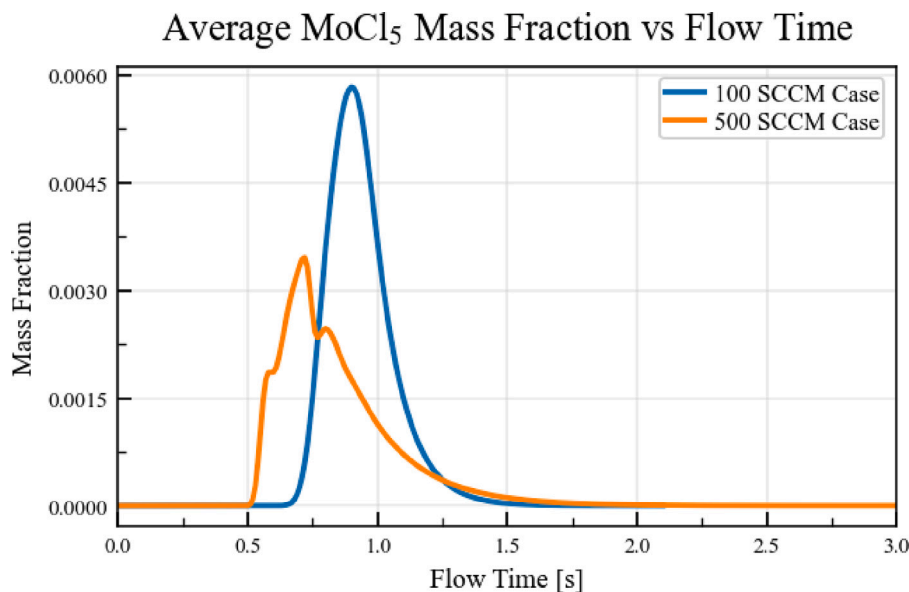


Fig. 17. Time history of the wafer-averaged MoCl_5 mass fraction for the 100 SCCM and 500 SCCM cases. Increasing the inlet flow rate leads to earlier precursor arrival but reduced peak average exposure and a more extended decay tail.

and to demonstrate that the framework generalizes beyond the two experimentally anchored cases.

A unified transient CFD model was used in which geometry, mesh, turbulence model, and numerical settings were held fixed while inlet and outlet boundary conditions were parameterized. In the NIST experiments, short MoCl_5 pulses were primarily used as transport tracers, whereas longer pulses corresponded to deposition-relevant conditions. The pulse-duration sweep therefore spans tracer-like and deposition-like regimes, enabling systematic evaluation of transport scaling without implying process optimization. Chamber pressure for each case was imposed via a pressure-outlet boundary condition consistent with the target value. The operating conditions are summarized in Table 4.

As a consistency check, inlet-line velocity and upstream pressure (P_1) were compared with experimentally inferred values. In the NIST experiments, velocity was estimated by converting the mass flow controller setpoint to volumetric flow rate at 393 K and dividing by inlet-tube area, assuming delivery within the nominal 5% uncertainty. The upstream pressure P_1 was inferred using a Hagen–Poiseuille analysis based on downstream pressure, gas properties, and tubing dimensions. In the CFD model, velocity and pressure were extracted at the same inlet location for direct comparison.

Table 5 shows good agreement across all cases. Remaining discrepancies are modest and attributable to experimental inference uncertainty and simplified inlet/outlet representations used for efficient parametric evaluation.

Temporal Evolution of MoCl₅ Mass Fraction on Wafer surface




Fig. 18. Time-resolved CFD contours of MoCl₅ mass fraction on the wafer surface for the 100 SCCM case, illustrating the spatial evolution of precursor distribution during injection and purge.

Table 4

Parametric operating conditions considered for generalized boundary-condition testing. Pulse duration denotes the MoCl₅ injection time.

#	Case ID	Pulse duration (s)	Flow (SCCM)	Chamber pressure (Pa)
1	0.2s-100SCCM-317Pa	0.2	100	317.6
2	0.2s-500SCCM-872.5Pa	0.2	500	872.5
3	0.5s-100SCCM-2425.9Pa	0.5	100	2425.9
4	0.5s-100SCCM-1367.5Pa	0.5	100	1367.5
5	3s-100SCCM-317Pa	3.0	100	317.6
6	3s-100SCCM-1367.5Pa	3.0	100	1367.5
7	3s-100SCCM-2425.9Pa	3.0	100	2425.9

Table 5

Comparison of experimentally inferred and simulated inlet-line pressure (P_1) and inlet velocity across parametric cases. Experimental velocity values are reported at 393 K.

Case	P_1 Exp. (Pa)	P_1 Sim. (Pa)	P_1 Error (%)	Velocity Exp. @393 K (m/s)	Velocity Sim. (m/s)	Velocity Error (%)
0.2s-100SCCM-317Pa	998.2	1031.2	+3.3%	13.01	13.10	+0.7%
0.2s-500SCCM-872.5Pa	2371.5	2491.3	+5.1%	27.39	27.53	+0.5%
0.5s-100SCCM-2425.9Pa	2598.1	2658.55	+2.3%	5.00	5.57	+11.4%
0.5s-100SCCM-1367.5Pa	1659.6	1747.46	+5.3%	7.83	8.48	+8.3%
3s-100SCCM-317Pa	998.2	1031.2	+3.3%	13.01	13.10	+0.7%
3s-100SCCM-1367.5Pa	1659.6	1742.74	+5.0%	7.83	8.50	+8.6%
3s-100SCCM-2425.9Pa	2598.1	2655.10	+2.2%	5.00	5.58	+11.6%

Time-resolved wafer-averaged MoCl₅ responses were also compared with corresponding experimental absorbance measurements when available. The model reproduced temporal buildup and decay behavior with close agreement in peak timing and qualitative profile shape. Detailed optical comparisons are not repeated here to avoid redundancy; emphasis is placed on generalized transport scaling and internal consistency.

Fig. 19 shows wafer-averaged gas velocity for all cases. After a brief transient, a steady velocity is established. Wafer-level velocity is governed primarily by inlet mass flow rate: the 500 SCCM case yields significantly higher steady velocity than all 100 SCCM cases, while variations in upstream pressure and pulse duration produce minimal differences once flow expands into the reactor. For fixed mass flow rate, inlet pressure does not appreciably alter steady wafer velocity. These trends confirm physically consistent hydrodynamic scaling.

The effect of pulse duration on wafer exposure was evaluated by comparing 0.2 s, 0.5 s, and 3 s injections under otherwise identical conditions (Fig. 20). Increasing pulse duration produces systematic broadening and higher wafer-averaged MoCl₅ magnitude. Short and

medium pulses exhibit transient peaks followed by decay, indicating injection–transport coupling, whereas the 3 s pulse approaches a quasi-steady plateau consistent with near-complete filling of the wafer region. The smooth transition from transport-influenced to duration-limited behavior further confirms consistent model response across tracer-like and deposition-like regimes.

4.2. Intermediate flow case: 300 SCCM

To further evaluate model consistency at intermediate operating conditions, a 300 SCCM case corresponding to a mid-range Reynolds number and chamber pressure was investigated. The experimentally measured downstream pressure $P_2 = 612$ Pa and the corresponding mass flow rate were imposed as the model boundary condition. Under these conditions, the CFD model predicted an inlet pressure of 1978 Pa, corresponding to a relative deviation of approximately 11.6% from the experimental upstream pressure $P_1 = 1773$ Pa. The predicted inlet velocity was 21.74 m/s compared to the experimentally calculated value of 21.17 m/s, resulting in a deviation of approximately 2.7%. The

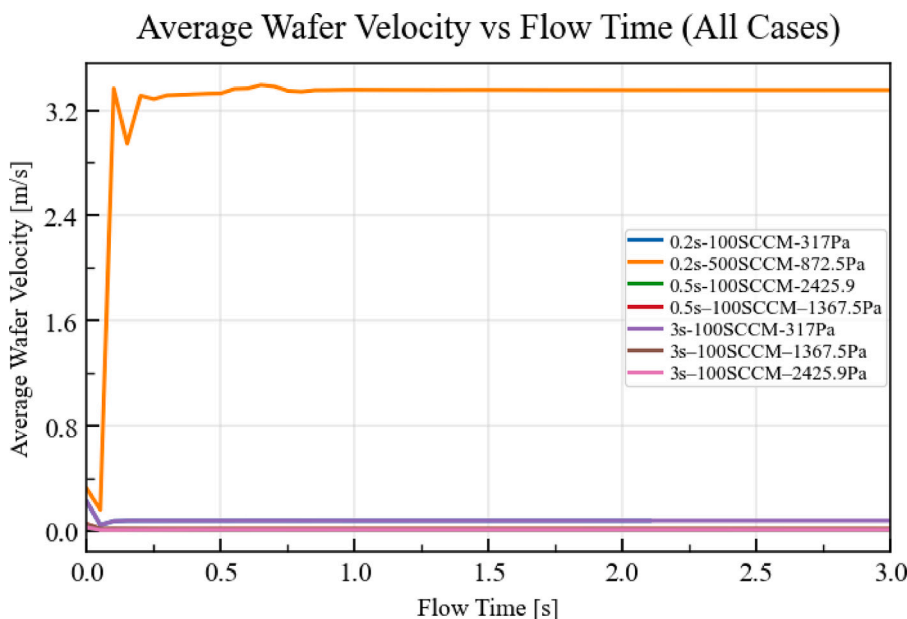


Fig. 19. Time evolution of wafer-averaged gas velocity for all parametric cases. Steady-state wafer velocity is primarily controlled by inlet mass flow rate.

close agreement in velocity confirms consistency in the imposed flow specification, while the pressure deviation remains within acceptable limits considering geometric simplifications and numerical assumptions in the model.

Fig. 21 presents the time-resolved normalized MoCl_5 signals predicted for the 300 SCCM case, compared against the experimental baseline cases at 100 and 500 SCCM to assess intermediate-flow transport behavior. Although no experimental measurements are available for 300 SCCM, the simulated response exhibits a buildup and decay profile that follows the expected trend between the two baseline cases, indicating physically consistent transient dynamics. Minor differences are observed in the decay region, likely reflecting variations in flow rate and associated velocity-driven residence effects.

Using the same peak-to-1% decay criterion and the same reactor-center volume above the wafer chuck defined in Section 3.4, the estimated residence time for the 300 SCCM case is $\tau \approx 1.2$ s. This value is comparable to the predicted 500 SCCM residence time (1.2 s), indicating that the intermediate operating condition produces a decay behavior of similar characteristic duration within the sampled region.

This outcome suggests that the transient purge behavior within the camera-equivalent volume does not scale strictly with volumetric flow rate. In the present configuration, the 300 SCCM case is simulated at a lower operating pressure ($P_2 = 612$ Pa) than the 500 SCCM baseline ($P_2 = 872.5$ Pa), which modifies gas density and consequently influences the local velocity distribution, mixing intensity, and convective removal within the region of interest. As a result, the intermediate-flow case exhibits a decay tail of comparable duration despite the reduced volumetric flow rate (see Fig. 22). Since no experimental residence-time data are available for 300 SCCM, this value should be interpreted as a model-based estimate of localized decay dynamics rather than a validated global chamber residence time.

Overall, the parametric study confirms that the CFD framework exhibits stable and physically consistent behavior across variations in flow rate, chamber pressure, and pulse duration without modification of the underlying numerical configuration and model settings. The predicted hydrodynamic and transient transport trends scale logically with imposed boundary conditions, indicating that the model response is not confined to the two experimentally anchored cases. Instead, the framework demonstrates general applicability within the investigated

operating envelope, supporting its use for intermediate and extended process conditions.

It is important to note that the need for detailed resolution of the convection-dominated flow field in the reactor for a variety of operating conditions necessitates the use of a direct CFD modeling approach as opposed to the use of approximate schemes like multi-compartment models (Shah et al., 2026b,a) which have been proven to work well in the modeling of other applications.

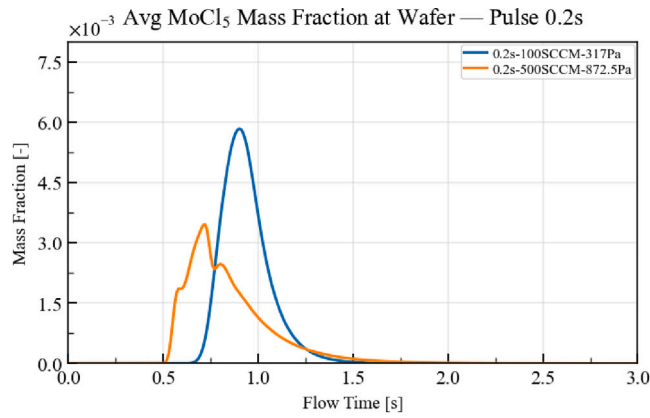
5. Conclusion

A physics-based digital twin for transient MoCl_5 transport in an experimental ALD reactor has been developed and evaluated over a representative operating envelope. The framework integrates time-resolved absorption-imaging and upstream absorbance measurements from the optically accessible NIST reactor with a three-dimensional, transient CFD model to capture the dominant flow and transport phenomena governing pulsed precursor delivery.

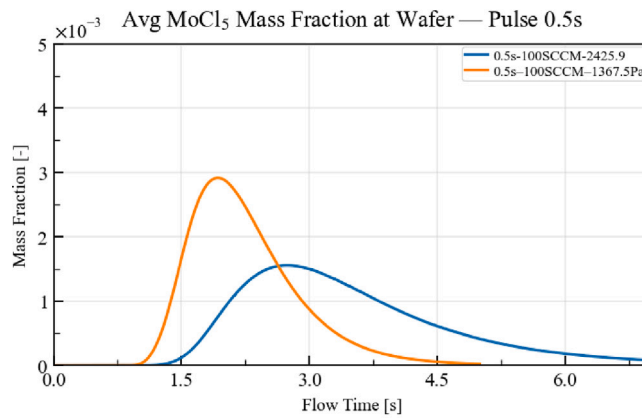
Comparison against experimental measurements demonstrates that the model reproduces key transport observables, including inlet velocity, normalized precursor buildup and decay behavior, residence-time estimates within the camera-equivalent region, plume symmetry and spatial evolution, and wafer chuck-level transport trends across low- and high-flow conditions. Image-based comparisons further confirm that the simulations capture the principal plume structures observed experimentally while maintaining physically consistent near-wafer hydrodynamic scaling.

A unified parametric study demonstrates smooth and predictable response to variations in flow rate, chamber pressure, and injection duration without modifying geometry or numerical settings. The results confirm that wafer chuck-level velocity is governed primarily by mass flow rate and that wafer exposure transitions from transport-limited to injection-duration-limited behavior for long pulses. An additional intermediate 300 SCCM case was evaluated within the same framework, and the predicted transient and residence-time behavior remained physically consistent with the validated baseline conditions.


Overall, the present work establishes a reusable, experiment-informed digital twin for ALD reactor transport analysis, providing a foundation for future coupling to surface reaction kinetics, multi-cycle simulations, and control- or optimization-oriented studies targeting uniformity, throughput, and precursor utilization.



(a) 0.2 s pulse



(b) 0.5 s pulse



(c) 3 s pulse

Fig. 20. Wafer-averaged MoCl₅ mass fraction versus flow time for the pulse-duration sweep.

Simulated 300 SCCM MoCl₅ Signals vs. Experimental Baseline Cases (100 & 500 SCCM)

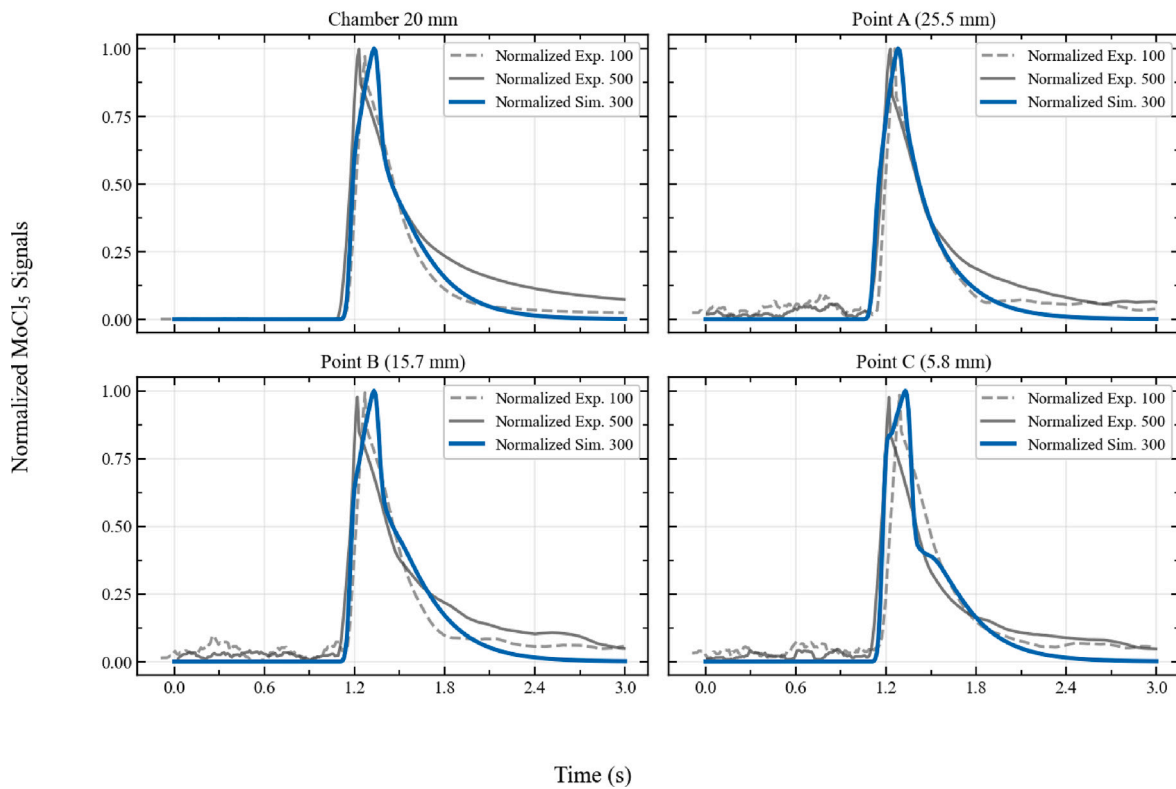


Fig. 21. Simulated 300 SCCM MoCl₅ signals compared against the experimental baseline cases at 100 and 500 SCCM across all measurement locations. No experimental data are available for the 300 SCCM condition; the comparison is shown to illustrate intermediate-flow transport behavior.

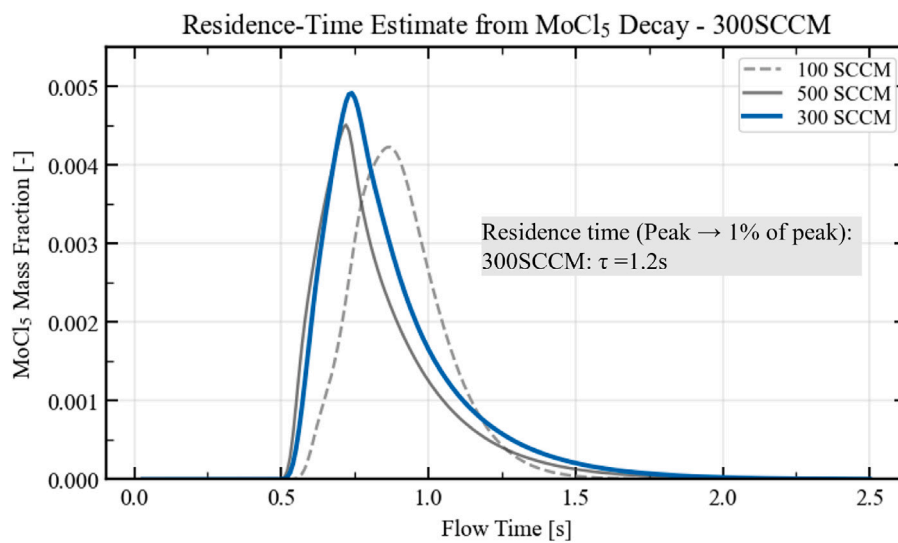


Fig. 22. Simulated MoCl₅ mass-fraction decay in the reactor-center volume above the wafer chuck for the 300 SCCM case compared with the 100 SCCM and 500 SCCM cases. Residence time is defined as the interval from the peak to 1% of the peak value.

CRedit authorship contribution statement

Abdulrahman Alghamdi: Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization. **Feiyang Ou:** Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization. **Berc Kalanyan:** Writing – review & editing, Methodology, Investigation. **James E. Maslar:** Writing – review & editing, Methodology, Investigation. **Panagiotis D. Christofides:** Writing – original draft, Methodology, Investigation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Financial support from the National Science Foundation, United States, CBET-2227241, and the Department of Energy is gratefully acknowledged. This work used computational and storage services associated with the Hoffman2 Shared Cluster provided by the UCLA Office of Advanced Research Computing's Research Technology Group. The authors would like to acknowledge Dr. Vladimir Khromchenko for providing CAD files that helped create Figs. 1(a) and (b).

References

- ANSYS, 2021. Ansys Fluent Theory Guide. ANSYS Inc., Canonsburg, PA.
- Bird, R.B., Stewart, W.E., Lightfoot, E.N., 2002. Transport Phenomena, second ed. John Wiley & Sons.
- Chen, M.-C., Li, K.-S., Li, L.-J., Lu, A.-Y., Li, M.-Y., Chang, Y.-H., Lin, C.-H., Chen, Y.-J., Hou, Y.-F., Chen, C.-C., et al., 2015. TMD FinFET with 4 nm thin body and back gate control for future low power technology. In: Proceedings of IEEE International Electron Devices Meeting. Washington, DC, pp. 32.2.1–32.2.4.
- George, S.M., 2010. Atomic layer deposition: An overview. Chem. Rev. 110, 111–131.
- Huang, J., 2022. Research progresses on suppressing the short-channel effects of field-effect transistor. Highlights Sci. Eng. Technol. 27, 361–367.
- Juppo, M., Vehkamäki, M., Ritala, M., Leskelä, M., 1998. Deposition of molybdenum thin films by an alternate supply of MoCl₅ and Zn. J. Vac. Sci. Technol. A 16, 2845–2850.
- Jurczak, M., Collaert, N., Veloso, A., Hoffmann, T., Biesemans, S., 2009. Review of finFET technology. In: Proceedings of IEEE International SOI Conference. Foster City, CA, pp. 1–4.
- Kanarik, K.J., Lill, T., Hudson, E.A., Sriraman, S., Tan, S., Marks, J., Vahedi, V., Gottscho, R.A., 2015. Overview of atomic layer etching in the semiconductor industry. J. Vac. Sci. Technol. A 33.
- Kim, J., Chakrabarti, K., Lee, J., Oh, K.-Y., Lee, C., 2003. Effects of ozone as an oxygen source on the properties of the Al₂O₃ thin films prepared by atomic layer deposition. Mater. Chem. Phys. 78, 733–738.
- Kimes, W.A., Moore, E., Maslar, J., 2012. Perpendicular-flow, single-wafer atomic layer deposition reactor chamber design for use with in situ diagnostics. Rev. Sci. Instrum. 83, 083106.
- Lee, Y., Kim, G.-H., Choi, B., Yoon, J., Kim, H.-J., Kim, D.H., Kim, D.M., Kang, M.-H., Choi, S.-J., 2020. Design study of the gate-all-around silicon nanosheet MOSFETS. Semicond. Sci. Technol. 35, 03LT01.
- Loubet, N., Hook, T., Montanini, P., Yeung, C.-W., Kanakasabapathy, S., Guillom, M., Yamashita, T., Zhang, J., Miao, X., Wang, J., et al., 2017. Stacked nanosheet gate-all-around transistor to enable scaling beyond finFET. In: Proceedings of 2017 Symposium on VLSI Technology. Kyoto, Japan, pp. T230–T231.
- Lubnin, A., Lebedev, R., Lad'yanov, V., Pushkarev, B., Faizullin, R., Baldaev, L.K., Treshchev, S.Y., 2022. Hydrogen-free CVD deposition of molybdenum coatings. Prot. Met. Phys. Chem. Surfaces 58, 76–83.
- Maslar, J.E., Kalanyan, B., 2025. Visualizing molybdenum pentachloride flow during vapor deposition processes using absorption imaging. Appl. Spectrosc. 79, 1487–1496.
- Menter, F., Egorov, Y., 2006. SAS turbulence modelling of technical flows. In: Direct and Large-Eddy Simulation VI. Springer, pp. 687–694.
- Menter, F., Kuntz, M., Langtry, R., 2003. Ten years of industrial experience with the SST turbulence model. Turbul. Heat Mass Transf. 4, 625–632.
- Oikawa, T., Kuroda, R., Takahashi, K., Shiba, Y., Fujihara, Y., Shike, H., Murata, M., Kuo, C.-C., Da Silva, Y.R.S.C., Goto, T., et al., 2022. A 70-dB SNR high-speed global shutter CMOS image sensor for in situ fluid concentration distribution measurements. IEEE Trans. Electron Devices 69, 2965–2972.
- Pan, D., Li, T., Jen, T.-C., Yuan, C., 2014. Numerical modeling of carrier gas flow in atomic layer deposition vacuum reactor: A comparative study of lattice Boltzmann models. J. Vac. Sci. & Technol. A 32, 01A110.
- Peltonen, P., Vuorinen, V., Marin, G., Karttunen, A.J., Karppinen, M., 2018. Numerical study on the fluid dynamical aspects of atomic layer deposition process. J. Vac. Sci. Technol. A 36, 021516.
- Sakai, Y., Shiba, Y., Inada, T., Goto, T., Suwa, T., Oikawa, T., Hamaya, A., Sutoh, A., Morimoto, T., Shirai, Y., et al., 2023. Visualization and analysis of temporal and steady-state gas concentration in process chamber using 70-dB SNR 1000 fps absorption imaging system. IEEE Trans. Semicond. Manuf. 36, 515–519.
- Shaeri, M.R., Jen, T.-C., Yuan, C.Y., 2015. Reactor scale simulation of an atomic layer deposition process. Chem. Eng. Res. Des. 94, 584–593.
- Shah, P., Nagpal, S., Kwak, D.H., Kim, J.H., Cho, J.H., Kim, J.-W., Roh, K., Park, S.M., Kwon, J.S.-I., 2026a. Comparative analysis of hydrodynamic and reactor design effects on performance in bioreactors. Chem. Eng. J. 527, 171731.
- Shah, P., Nagpal, S., Kwak, D.H., Kim, J.H., Cho, J.H., Park, S.M., Roh, K., Kim, J.-W., Kwon, J.S.-I., 2026b. Development of a compartment modeling framework from axisymmetric CFD models: Extracting flow topology for industrial fermenters. Chem. Eng. Sci. 326, 123502.
- Wang, C.-P., Tsai, Y.-P., Lin, B.-J., Liang, Z.-Y., Chiu, P.-W., Shih, J.-R., Lin, C.-J., King, Y.-C., 2020. On-wafer FinFET-based EUV/eBeam detector arrays for advanced lithography processes. IEEE Trans. Electron Devices 67, 2406–2413.
- Yoshikawa, N., Kikuchi, A., 1996. Growth rate and microstructure of Mo film by chemical-vapour-deposition. Mater. Trans. JIM 37, 283–288.
- Yun, S., Ou, F., Wang, H., Tom, M., Orkoulas, G., Christofides, P.D., 2022. Atomistic-mesoscopic modeling of area-selective thermal atomic layer deposition. Chem. Eng. Res. Des. 188, 271–286.
- Yun, S., Wang, H., Tom, M., Ou, F., Orkoulas, G., Christofides, P.D., 2023. Multiscale CFD modeling of area-selective atomic layer deposition: Application to reactor design and operating condition calculation. Coatings 13, 558.