
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Morev, Ilya; Tekgül, Bulut; Gadalla, Mahmoud; Shahanaghi, Ali; Kannan, Jeevananthan; Karimkashi, Shervin; Kaario, Ossi; Vuorinen, Ville

Fast reactive flow simulations using analytical Jacobian and dynamic load balancing in OpenFOAM

Published in:
Physics of Fluids

DOI:
[10.1063/5.0077437](https://doi.org/10.1063/5.0077437)

Published: 01/02/2022

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Morev, I., Tekgül, B., Gadalla, M., Shahanaghi, A., Kannan, J., Karimkashi, S., Kaario, O., & Vuorinen, V. (2022). Fast reactive flow simulations using analytical Jacobian and dynamic load balancing in OpenFOAM. *Physics of Fluids*, 34(2), Article 021801. <https://doi.org/10.1063/5.0077437>

Fast reactive flow simulations using analytical Jacobian and dynamic load balancing in OpenFOAM

Cite as: Phys. Fluids **34**, 021801 (2022); <https://doi.org/10.1063/5.0077437>

Submitted: 02 November 2021 • Accepted: 11 January 2022 • Published Online: 02 February 2022

 Ilya Morev,  Bulut Tekgül,  Mahmoud Gadalla, et al.



View Online



Export Citation



CrossMark

ARTICLES YOU MAY BE INTERESTED IN

[High-order methods for diffuse-interface models in compressible multi-medium flows: A review](#)

Physics of Fluids **34**, 021301 (2022); <https://doi.org/10.1063/5.0077314>

[Flow and mass transfer characteristics for interacting side-by-side cylinders](#)

Physics of Fluids **34**, 023602 (2022); <https://doi.org/10.1063/5.0078171>

[The rheology of ultra-high molecular weight poly\(ethylene oxide\) dispersed in a low molecular weight carrier](#)

Physics of Fluids **34**, 023304 (2022); <https://doi.org/10.1063/5.0077122>




Author Services

English Language Editing

High-quality assistance from subject specialists

LEARN MORE



Fast reactive flow simulations using analytical Jacobian and dynamic load balancing in OpenFOAM

Cite as: Phys. Fluids **34**, 021801 (2022); doi: [10.1063/5.0077437](https://doi.org/10.1063/5.0077437)

Submitted: 2 November 2021 · Accepted: 11 January 2022 ·

Published Online: 2 February 2022



View Online



Export Citation



CrossMark

Ilya Morev,^{a)} Bulut Tekgül, Mahmoud Gadalla, Ali Shahanaghi, Jeevananthan Kannan, Shervin Karimkashi, Ossi Kaario, and Ville Vuorinen

AFFILIATIONS

Department of Mechanical Engineering, Aalto University School of Engineering, Otakaari 4, 02150 Espoo, Finland

^{a)} Author to whom correspondence should be addressed: ilya.morev@aalto.fi

ABSTRACT

Detailed chemistry-based computational fluid dynamics (CFD) simulations are computationally expensive due to the solution of the underlying chemical kinetics system of ordinary differential equations (ODEs). Here, we introduce a novel open-source library aiming at speeding up such reactive flow simulations using OpenFOAM, an open-source software for CFD. First, our dynamic load balancing model by Tekgül *et al.* ["DLBFoam: An open-source dynamic load balancing model for fast reacting flow simulations in OpenFOAM," *Comput. Phys. Commun.* **267**, 108073 (2021)] is utilized to mitigate the computational imbalance due to chemistry solution in multiprocessor reactive flow simulations. Then, the individual (cell-based) chemistry solutions are optimized by implementing an analytical Jacobian formulation using the open-source library pyJac, and by increasing the efficiency of the ODE solvers by utilizing the standard linear algebra package. We demonstrate the speed-up capabilities of this new library on various combustion problems. These test problems include a two-dimensional (2D) turbulent reacting shear layer and three-dimensional (3D) stratified combustion to highlight the favorable scaling aspects of the library on ignition and flame front initiation setups for dual-fuel combustion. Furthermore, two fundamental 3D demonstrations are provided on non-premixed and partially premixed flames, viz., the Engine Combustion Network Spray A and the Sandia flame D experimental configurations, which were previously considered unfeasible using OpenFOAM. The novel model offers up to two orders of magnitude speed-up for most of the investigated cases. The openly shared code along with the test case setups represent a radically new enabler for reactive flow simulations in the OpenFOAM framework.

© 2022 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0077437>

I. INTRODUCTION

There is an urgent need to develop effective and accurate engineering solutions to combat climate change. Reactive flow simulations play an essential role in academic and industrial research on mitigating harmful emissions. A number of computational fluid dynamics (CFD) packages are available on the market, offering a variety of features together with possible limitations. The major limitations related to software distribution under commercial licenses are their affordability and source code access. The imminent pace of climate change urges us to broaden the use of the latest achievements in CFD beyond the limitations of commercial software to all appropriate applications. Unfortunately, the available open-source reacting flow solvers are currently limited. The most popular open-source general-purpose CFD code—OpenFOAM—poses relatively poor performance for finite rate

chemistry simulations compared to commercial solutions, making it impractical for large-scale industrial applications and thus limiting its use to relatively simple academic cases. However, the importance of open-source solutions was recently recognized at a governmental level. For instance, there is a large ongoing project funded by the European Union on optimizing the OpenFOAM code—exaFoam.¹ In this work, we proceed in a similar track and focus on the search for the bottlenecks in reacting flow simulations within OpenFOAM with the aim of significantly improving computational speed-up.

One of the most fundamental approaches for reactive flow modeling is to use CFD with the direct integration of chemical kinetics, referred to as finite-rate chemistry.² In this approach, each computational cell is treated as an individual chemistry problem with pressure (p), and a thermochemical state vector (Φ) comprising

temperature (T) and species mass fractions (Y_k). The rate of change of the thermochemical state vector, $\partial\Phi/\partial t = f(\Phi, t)$, within a computational cell forms a stiff system of ordinary differential equations (ODEs) that requires a specific class of algorithms for temporal integration.

In reactive CFD applications, chemistry evaluation comprises the most computationally demanding part of the simulation. With the growing complexity of chemical kinetic mechanisms, the cost of solving the chemistry problem often exceeds the cost of fluid flow solution by up to two orders of magnitude.³ Nevertheless, high-fidelity reactive flow simulations can be still feasible through efficient usage of the provided computational resources, in addition to cell-based optimization of the chemistry solution.⁴ In general, considerable efforts have been made in the literature to improve the computational performance of CFD simulations involving turbulence,^{5–9} shock waves,¹⁰ gas rarefaction,¹¹ geometrical parameterization,¹² or chemical reactions^{13–16} through model reduction, machine learning, or efficient parallelization.

In finite-rate chemistry simulations, the high computational cost associated with the chemistry solution originates mainly from two factors. First, the cost of solving the system of ODEs depends on the system stiffness, which is, to some extent, influenced by the local thermochemical state. For example, an oxidation process comprising different radicals at various chemical timescales may constitute a greater computational challenge compared to the local state of burnt mixture. Second, the cost of computing the rate of change of the thermochemical state vector is directly proportional to the chemical mechanism size,^{17,18} which can reach up to thousands of species and even more reactions.^{19,20}

For the chemistry ODE problem, implicit solvers involving Jacobian evaluation are usually preferred over explicit ones for stability and accuracy reasons.²¹ The evaluation of the Jacobian, $\mathcal{J} = \partial f / \partial \Phi$, is a computationally demanding process within an ODE solution routine. ODE solvers often employ finite differencing methods for its evaluation, which may be an expensive operation depending on the chemical mechanism size. Lu *et al.*¹⁷ reported that the computational cost of numerical evaluation of Jacobian scales quadratically with the number of reactions. It has been shown that using an analytically computed Jacobian provides high performance gain when solving the ODE system.^{17,22,23} There are various implementations of analytical Jacobian matrix evaluation in the literature.^{23–25} Recently, Niemeyer *et al.*²⁶ introduced an open-source library pyJac, which generates C subroutines for analytical Jacobian evaluation. They demonstrated its accuracy and out-performance over the available analytical Jacobian generators in the literature.

While the cost of integrating a single chemistry ODE problem poses a performance challenge in reactive CFD simulations, there is a secondary issue that is often left unaddressed. As previously mentioned, the computational cost of solving a chemistry ODE depends on the local stiffness of Φ . The difference in computational complexity throughout the geometrical domain poses a load imbalance issue for multiprocessor applications, where one process becomes the computational bottleneck and causes performance issues. There have been several studies in the recent literature tackling this load imbalance issue via utilizing dynamic load balancing algorithms, often using Message-Passing Interface (MPI) routines. Antonelli *et al.*²⁷ developed an MPI-based model introducing a cell distribution-based load balancing

algorithm. Following this work, Shi *et al.*²⁸ and Kodavasal *et al.*²⁹ both introduced stiffness detection-based load balancing algorithms and employed them in reactive CFD simulations. Zirwes *et al.* developed an MPI-based dynamic load balancing algorithm for chemistry problem distribution for OpenFOAM.³⁰ More recently, Muela *et al.*³¹ presented a dynamic load balancing model, which also utilizes a stiffness detection algorithm that chooses the optimal ODE integration method for each computational cell. These methods introduced computational speed-up around 3–5, depending on the application. We have recently also introduced an open-source dynamic load balancing model called DLBFoam for OpenFOAM, providing speed-up up to a factor of 10.⁴

In this work, we introduce a novel chemistry model that provides speed-up in reactive flow simulations in OpenFOAM by targeting the two major issues in reactive CFD simulations addressed above. We further extend our dynamic load balancing model DLBFoam, and focus on optimizing the cell-wise chemistry solution by introducing two new features. First, we utilize an analytical Jacobian approach using the pyJac library. Second, we make use of the standard linear algebra library (LAPACK)³² to further improve the ODE solution procedure of the chemistry problem, by replacing the LU decomposition and back substitution operations of the Seulex ODE integration algorithm in OpenFOAM with more robust alternatives suitable for dense Jacobian matrices. The effectiveness and robustness of the developed model are demonstrated on two different academic dual-fuel (DF) combustion setups: a two-dimensional (2D) reacting shear layer and a three-dimensional (3D) stratified combustion configuration. After that, the model is further demonstrated in two well-established experimental flame configurations: a non-premixed *n*-dodecane flame [Engine Combustion Network (ECN) Spray A] and a partially premixed methane–air flame (Sandia flame D).

The paper is outlined as follows. First, implementation details of the developed model are presented in Sec. II. Then, in Sec. III we discuss the application of our model to different combustion simulations, represented by the aforementioned test cases. Finally, conclusions are provided in Sec. IV.

II. IMPLEMENTATION

In this section, the implementation details of our model are presented. We note that an earlier in-house version of this implementation was carried out by Kahila *et al.*,³³ details of which can be found in his thesis work.³⁴ This section briefly discusses the DLBFoam model introduced by Tekgül *et al.*⁴ as an improvement to our in-house code, and further describes the efforts toward optimizing the solution of the chemistry ODE problem by utilizing an analytical Jacobian formulation (pyJac) together with efficient linear algebraic routines (LAPACK).

A. Finite rate chemistry

In reactive CFD applications, we numerically solve the governing equations for mass, momentum, energy, and species transport. The total production or consumption rate of chemical species and the heat release rate are represented as source terms in the species transport equations and the energy equation, respectively. Due to the vast difference between the chemistry and flow timescales, commonly, an operator splitting procedure is utilized to separate the calculation of source terms from the flow solution. These source terms define the change of

the local thermochemical state resulting from solving a stiff, nonlinear initial value problem described as follows:

$$\begin{cases} \partial \Phi / \partial t = \left\{ \frac{\partial T}{\partial t}, \frac{\partial Y_1}{\partial t}, \frac{\partial Y_2}{\partial t}, \dots, \frac{\partial Y_{N_{sp}-1}}{\partial t} \right\}^\top = f(\Phi, t), \\ \Phi(t=0) = \Phi_0, \end{cases} \quad (1)$$

where N_{sp} is the number of species in the chemical mechanism. The ODE system in Eq. (1) is integrated in each computational cell over the CFD time step Δt_{CFD} in smaller chemical time steps $\Delta t_{\text{ODE}} \leq \Delta t_{\text{CFD}}$ (see Part C for more information). The source terms are then evaluated and used in species transport and energy equations. Further details on this modeling approach can be found in the work of Imren and Haworth.³⁵

B. Dynamic load balancing: DLBFoam

Next, the previously developed DLBFoam library is briefly introduced here as essential background information. As mentioned earlier, the computational cost of solving the stiff nonlinear ODE may depend on the local thermochemical conditions. In multiprocessor reactive CFD simulations, the process with the highest computational load may take longer to compute than the rest, hence creating a bottleneck within the given CFD time step.

Recently, we introduced an open-source model called DLBFoam⁴ for OpenFOAM, aiming to mitigate this imbalance issue in multiprocessor reactive CFD simulations via dynamic load balancing. DLBFoam uses MPI routines to redistribute the chemistry computational load evenly between processes during the simulation. Figure 1 demonstrates the computational imbalance caused by the direct integration of chemistry, and how it is mitigated by DLBFoam. In addition, DLBFoam introduces a zonal reference cell mapping method, which further lowers the computational cost by mapping a chemistry solution from a reference cell instead of explicitly solving it for ambient regions with low reactivity. In total, we reported around a factor of 10 speed-up in 3D reactive CFD simulations, compared to the standard OpenFOAM chemistry model. Further implementation details on DLBFoam can be found in our previous publication describing it in detail.⁴

C. ODE solver optimization and analytical Jacobian

The new development in the present paper introduces a coupling of (1) the analytical Jacobian formulation introduced in pyJac library, (2) standard linear algebra routines from LAPACK library for the ODE solution, and (3) the DLBFoam library previously introduced by Tekgül *et al.*⁴ Within a CFD time step, chemistry is treated as an independent, stiff system of ODEs. Implicit and semi-implicit algorithms for solving this system of stiff ODEs are usually preferred to ensure the solution stability.²¹

The solution to the nonlinear initial value problem given by Eq. (1) can be obtained using a variant of Newton iteration with the following discretized form:

$$\begin{aligned} \Phi_{n+1} &= \Phi_n + \int_{t_n}^{t_n + \Delta t_{\text{ODE}}} (f_n + \mathcal{J}_n(\Phi_{n+1} - \Phi_n) + \mathcal{O}(\Delta t_{\text{ODE}}^2)) dt \\ &= \Phi_n + f_n \Delta t_{\text{ODE}} + \mathcal{J}_n(\Phi_{n+1} - \Phi_n) \Delta t_{\text{ODE}} + \mathcal{O}(\Delta t_{\text{ODE}}^2), \end{aligned} \quad (2)$$

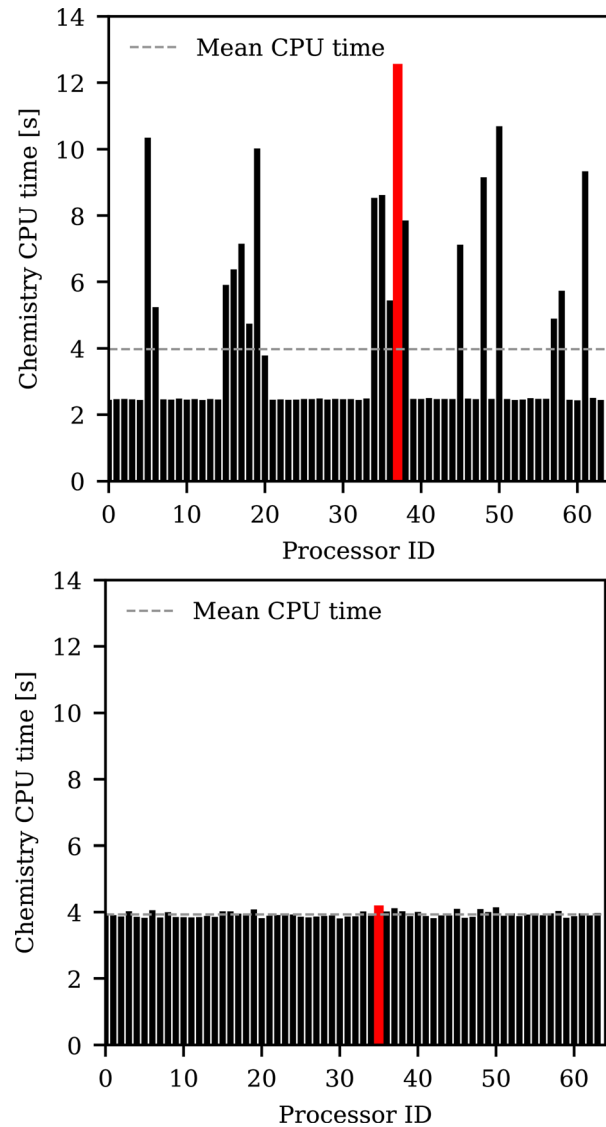


FIG. 1. Top: the process-based computational imbalance in reactive CFD simulations without DLBFoam. The process with the highest computational load (red) creates a bottleneck in the simulation. Bottom: same case with DLBFoam.

which is then linearized by neglecting the higher-order terms. Here, Δt_{ODE} is a subinterval of the full integration interval Δt_{CFD} , which is commonly determined by the underlying ODE solver using a combination of theoretical and heuristic relations. The direct solution of the previous equation for Φ_{n+1} requires the inverse of \mathcal{J} , which is commonly avoided through matrix factorization (i.e., LU decomposition) and back substitution techniques.

Commonly, in reactive CFD codes, the Jacobian of the system of ODEs representing the chemistry is implemented via finite differencing to be used in ODE solver integration. However, using such an approximation introduces a negative effect on the ODE solver accuracy and performance.²⁶ Furthermore, calculating the Jacobian with finite differencing is a computationally expensive procedure.^{17,26}

Instead, utilizing a fully algebraic Jacobian is more efficient and it improves the ODE solver accuracy and performance. Lu *et al.*¹⁷ reported that using an analytical Jacobian reduces its evaluation time from the square of the number of species to a linear dependence.

OpenFOAM features a fully algebraic Jacobian implementation to speed up the solution of chemical kinetics in its recent releases.³⁶ However, we have found that the implemented algebraic Jacobian fails to deliver a fast solution to chemistry ODE integration when high ODE convergence tolerances are utilized. Instead, we utilize the Python-based open-source software pyJac, introduced by Niemeyer *et al.*²⁶ The pyJac package generates C subroutines for the evaluation of an analytical Jacobian for a chemical mechanism. They reported that pyJac performs 3–7.5 times faster than the first-order finite differencing approach. In our proposed model, OpenFOAM's Jacobian calculation subroutines are replaced with the subroutines generated by pyJac.

The fully algebraic Jacobian utilization via pyJac speeds up the Jacobian evaluation time and improves the solution convergence rate. However, the ODE solution method itself is also particularly important for obtaining faster chemistry solution. We observed that the OpenFOAM's native functions used for LU decomposition and back substitution matrix operations of the chemistry ODE system are not suitable for solving dense and small matrices ($<500 \times 500$). Therefore, the LU decomposition and back substitution linear algebraic operations of OpenFOAM are replaced by more robust implementations existing in the open-source library LAPACK.³² The LAPACK routines utilized in our model are very efficient for dense and small matrices, such as the fully dense Jacobian matrix created by the pyJac.

In this study, LAPACK functions were linked to a semi-implicit extrapolation-based Euler method, denoted as Seulex.^{21,35} Seulex requires the specification of relative and absolute tolerances of the solution along with an initial estimate of Δt_{ODE} . As the solution progresses, the solver improves the estimate for Δt_{ODE} at the current thermochemical state using theoretical relations and heuristics. The previous available value is then used as an initial value for the following integration interval. Furthermore, a Δt_{ODE} is recursively split into a number of partitions and each partition is then solved using a low-order method. Finally, the high-order solution over the whole interval Δt_{ODE} is recursively gathered from

solutions of the smaller partitions. If the solution does not satisfy the tolerance criteria, partitioning continues until it is satisfied. A more detailed explanation of the Seulex algorithm can be found in the work by Imren and Haworth.³⁵

As noted earlier, each computational cell has its own independent ODE problem with a thermochemical state vector comprising only intensive mixture properties. Furthermore, the ODE solver guarantees that if the solution has converged, the error will lie within the user-specified tolerance regardless of the width of the integration interval Δt_{CFD} . However, errors related to the operator splitting technique need to be taken care of. The assumption behind operator splitting is that the chemistry timescales are orders of magnitude smaller than the flow timescales or the CFD solver time step Δt_{CFD} . Nevertheless, the time step Δt_{CFD} is to be chosen so that the species and temperature transport effects over a time step are accounted for and the thermochemical state vector is updated correspondingly. We note that in this study we do not consider any chemical source term closures for the filtered or averaged equations. Such turbulence–chemistry interaction (TCI) models could significantly depend on the mesh parameters and the CFD time step size and might require a separate treatment.

The effect of the optimized Jacobian evaluation is demonstrated in Fig. 2 on a zero-dimensional (0D) homogeneous reactor simulation with stoichiometric methane–air mixture at $T = 1200$ K and $p = 13.5$ atm. The GRI-3.0³⁷ chemical kinetic mechanism was used and the absolute and relative ODE solver tolerances were set to 10^{-10} and 10^{-6} , respectively. The integration interval (Δt_{CFD}) was fixed to 10^{-6} s, which corresponds to timescales relevant to reactive CFD applications under engine-relevant conditions. From the figure, it is observed that the time spent on Jacobian and derivative evaluations is reduced by roughly one order of magnitude between Standard and pyJac along the entire simulation time. Furthermore, the exact Jacobian formulation increases the solution accuracy for every Newton iteration, which reduces the amount of iterations required for convergence. Next, with LAPACK we notice a further reduction in the analytical Jacobian retrieval and also less time for all operations only within the sharp gradient zone. The reason is that with LAPACK, the ODE solver uses wider sub-intervals in the stiff zone, hence, fewer function calls for the Jacobian retrieval, and, consequently, fewer linear algebraic operations.

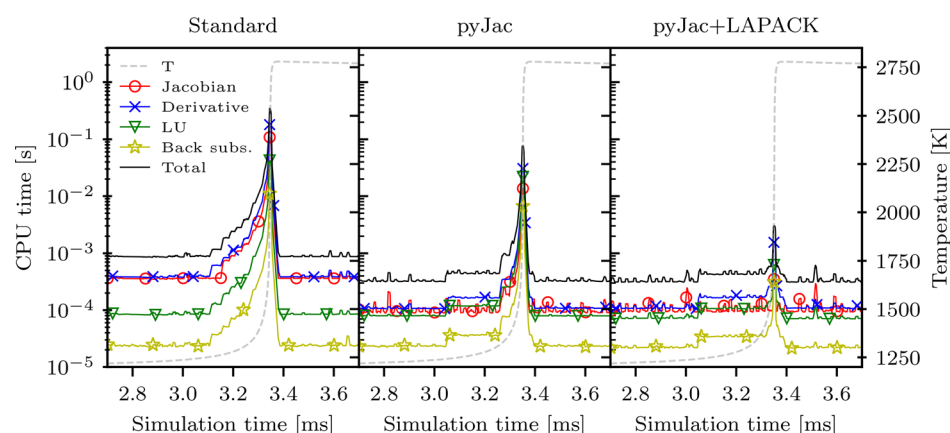


FIG. 2. CPU time spent on different operations of ODE solution algorithm for a single cell 0D problem involving CH_4 chemistry using Standard, pyJac, and pyJac+LAPACK models. Selected time-frame represents the most computationally demanding part of the problem. The GRI-3.0 mechanism (53 species and 325 reactions) is used in the benchmark.

Figure 3 depicts the total execution times for the aforementioned problem with varied absolute tolerances. It can be seen that while the Jacobian retrieval using pyJac reduces the execution time by about one order of magnitude, using LAPACK routines provides another order of magnitude speed-up. The utilization of pyJac together with LAPACK allows using much tighter ODE tolerances.

III. RESULTS

A. An overview of the test cases

After introducing the new features of our model in Sec. II, here the combined effect of the DLBFOam library with pyJac and LAPACK is investigated. Three different models are employed here for performance benchmarking: (1) the standard OpenFOAM chemistry model (referred to as Standard), (2) our original model with dynamic load balancing and zonal reference mapping⁴ (referred to as DLBFOam), and (3) our original model combined with the ODE-related improvements utilizing pyJac and LAPACK routines (referred to as DLBFOam+pyJac). We first demonstrate the computational performance of DLBFOam+pyJac in two academic test cases (2D reacting shear layer and 3D stratified combustion). After that, we demonstrate its performance in two experimental flame configurations (ECN Spray A and Sandia flame D). Schematic diagrams presenting the four test cases are provided in Fig. 4.

For each case, we first use DLBFOam+pyJac to solve a full-scale combustion problem. Then, for each problem we choose a time interval that is considered computationally challenging—e.g., chemistry ODE problem is stiff in parts of the domain due to ignition—and compare the performance of DLBFOam+pyJac against DLBFOam and Standard for a specific number of iterations. A summary of various case-specific benchmark parameters along with the computational speed-up is presented in Table I. Here, we have to limit the number of iterations for benchmarking because of the poor performance of the Standard model we highlighted in Sec. II. In fact, with the chosen strict tolerance values, it is infeasible to carry out a full simulation with




FIG. 3. Dependency of the execution time on the absolute tolerance ($N_{\text{sample}} = 10$) for a single cell 0D problem involving CH_4 chemistry with different ODE convergence tolerances. The GRI-3.0 mechanism (53 species and 325 reactions) is used in the benchmark.

Standard or even benchmark with a larger number of iterations than those stated in the table for all of the studied test cases.

Regarding numerical schemes, a second-order spatial and temporal discretization is utilized for all test cases presented herein. The reacting PISO (Pressure-Implicit with Splitting of Operators) algorithm with two outer-loop correctors is utilized for pressure-velocity coupling,⁴⁰ i.e., the chemistry is solved twice within a single CFD time step. The absolute and relative tolerances of the chemistry ODE solver are set to 10^{-10} and 10^{-6} , respectively. More details about the case setup can be found in test case files, which are openly shared.⁴¹

All simulations and benchmarks were performed on the Mahti supercomputer at CSC—Finnish IT Center for Science. Mahti has 1404 nodes, each with two 64-core 2.6 GHz (boost up to 3.3 GHz) AMD Rome 7H12 CPUs. All nodes are connected to the inter-node communication network with 200 GB/s links.⁴²

B. 2D reacting shear layer

The first test case presents a simple, temporally evolving 2D reacting shear layer that is computationally affordable even on a personal computer. The presented test case is fundamentally related to the ignition and flame initiation of DF combustion in compression ignition engines such as the reactivity controlled compression ignition (RCCI) engine.^{33,43} We demonstrate the computational speed-up and scaling effects of DLBFOam+pyJac with respect to the number of cores.

The schematic of the present numerical setup is illustrated in Fig. 4(a). The setup contains a high reactivity *n*-dodecane jet stream that mixes with the surrounding oxidizer consisting of premixed methane, oxidizer, and exhaust gas recirculation (EGR). The simulation domain is a square box with a side length $D = 1.5$ mm. The number of grid points in both directions is 300, which is based on the pre-estimated laminar flame thickness for *n*-dodecane-methane premixed flame at the corresponding most reactive mixture fraction, $p = 60$ bar, and $T_{\text{reactants}} = 800$ K ($\delta_f \approx 50 \mu\text{m}$) to resolve the flame by 10 grid points. A hyperbolic tangent function is used to generate the shear layer between the *n*-dodecane and oxidizer streams. The *n*-dodecane jet is initially set to 700 K whereas the methane/air mixture is set to 900 K at a constant pressure of 60 bar. The initial conditions corresponding to the ambient premixed mixture are similar to our previous DF spray studies^{33,43–47} as reported in Table II. The *n*-dodecane jet moves initially with a relative velocity of 10 m/s to the methane/air stream and develops a Kelvin–Helmholtz instability. A skeletal chemical mechanism (54 species and 269 reactions) developed by Yao *et al.*³⁸ for *n*-dodecane combustion is used. The performance of this mechanism in DF context has been already demonstrated in our earlier studies.^{33,44,46}

Figure 5 highlights the DF ignition process in the shear layer, where the *n*-dodecane jet ignites the surrounding mixture. It is observed that the first-stage ignition from low-temperature chemistry develops primarily within the *n*-dodecane jet near the mixing layer. Next, the ignition front initiates in the *n*-dodecane region at the most reactive mixture fraction and propagates toward the ambient mixture. Here, the mixture fraction describes the mixing extent of the *n*-dodecane fuel stream and the premixed methane, oxidizer, and EGR. Finally, a premixed flame front is established, which completes the combustion of the ambient methane.

Next, we report the computational speed-up and the parallel efficiency provided by our model for different processor counts, as

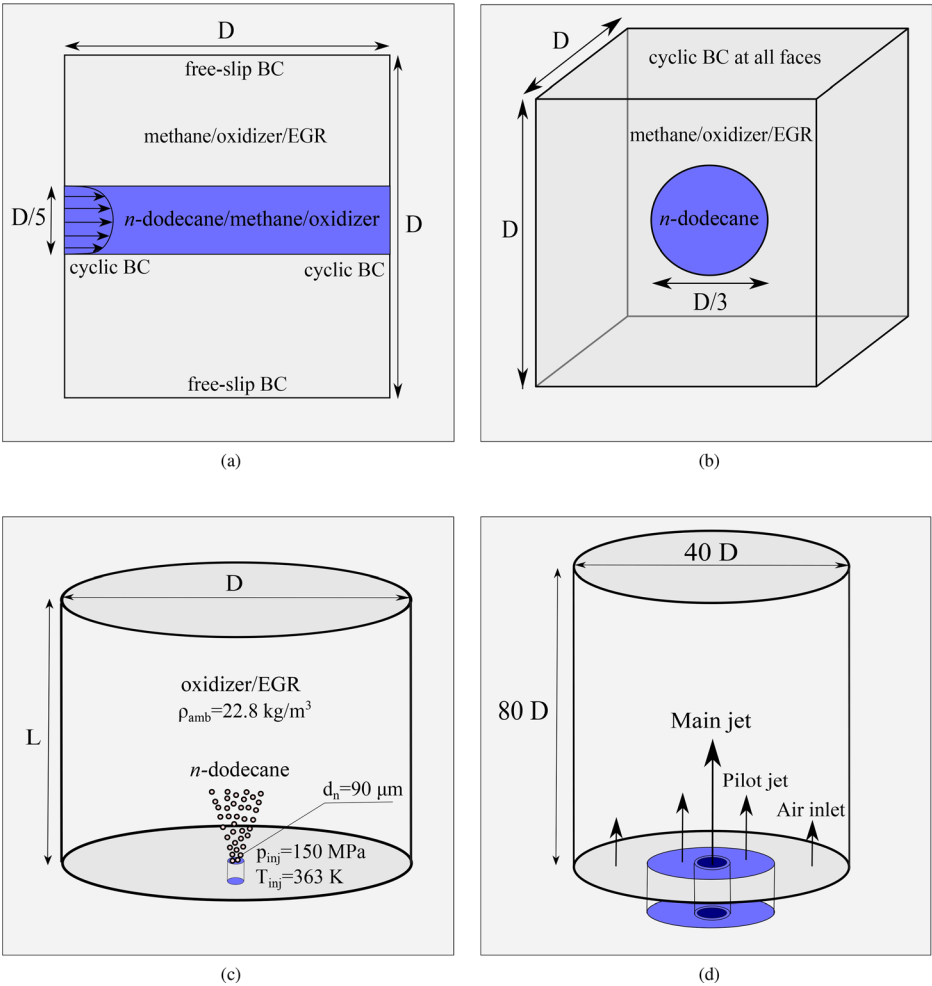


FIG. 4. Schematic diagrams of the demonstrated test cases. (a) 2D reacting shear layer, (b) 3D stratified combustion, (c) ECN Spray A, and (d) Sandia flame D.

TABLE I. List of parameters used in performance benchmarks. Computational time refers to the full-scale simulation using DLBFOam+pyJac. Computational speed-up is reported for DLBFOam+pyJac against Standard. Computational time for ECN Spray A reports the case with $T_{amb} = 900\text{ K}$ and simulation time of 1.5 ms.

	2D shear layer	3D stratified combustion	ECN Spray A	Sandia flame D
Chemical mechanism	Yao ³⁸	Yao	Yao	DRM19 ³⁹
Number of cells	90 000	16.78 M	39 M	1.99 M
Number of processors	8–32	1152	1920	768
Benchmarking samples (CFD iterations)	100	7	65	100
Computational speed-up	24–38	≈400	256	13.5
Computational time (CPU-hr)	65	15 000	128 500	24 500

TABLE II. Initial conditions in molar fractions and corresponding equivalence ratio for premixed ambient mixture in 2D shear layer and 3D stratified combustion cases.

CH ₄ %	O ₂ %	CO ₂ %	H ₂ O %	N ₂ %	ϕ_{CH_4}
3.750	15	5.955	3.460	71.835	0.5

depicted in Fig. 6. The domain is decomposed into 8, 16, and 32 processors to demonstrate the scaling capabilities of the model. All speed-up tests are carried out for 100 constant CFD time steps of 2×10^{-7} s after the DF ignition. The following observations are made from the analysis: (1) the standard model demonstrates a poor parallel scaling efficiency due to the load imbalance. While the strong scaling

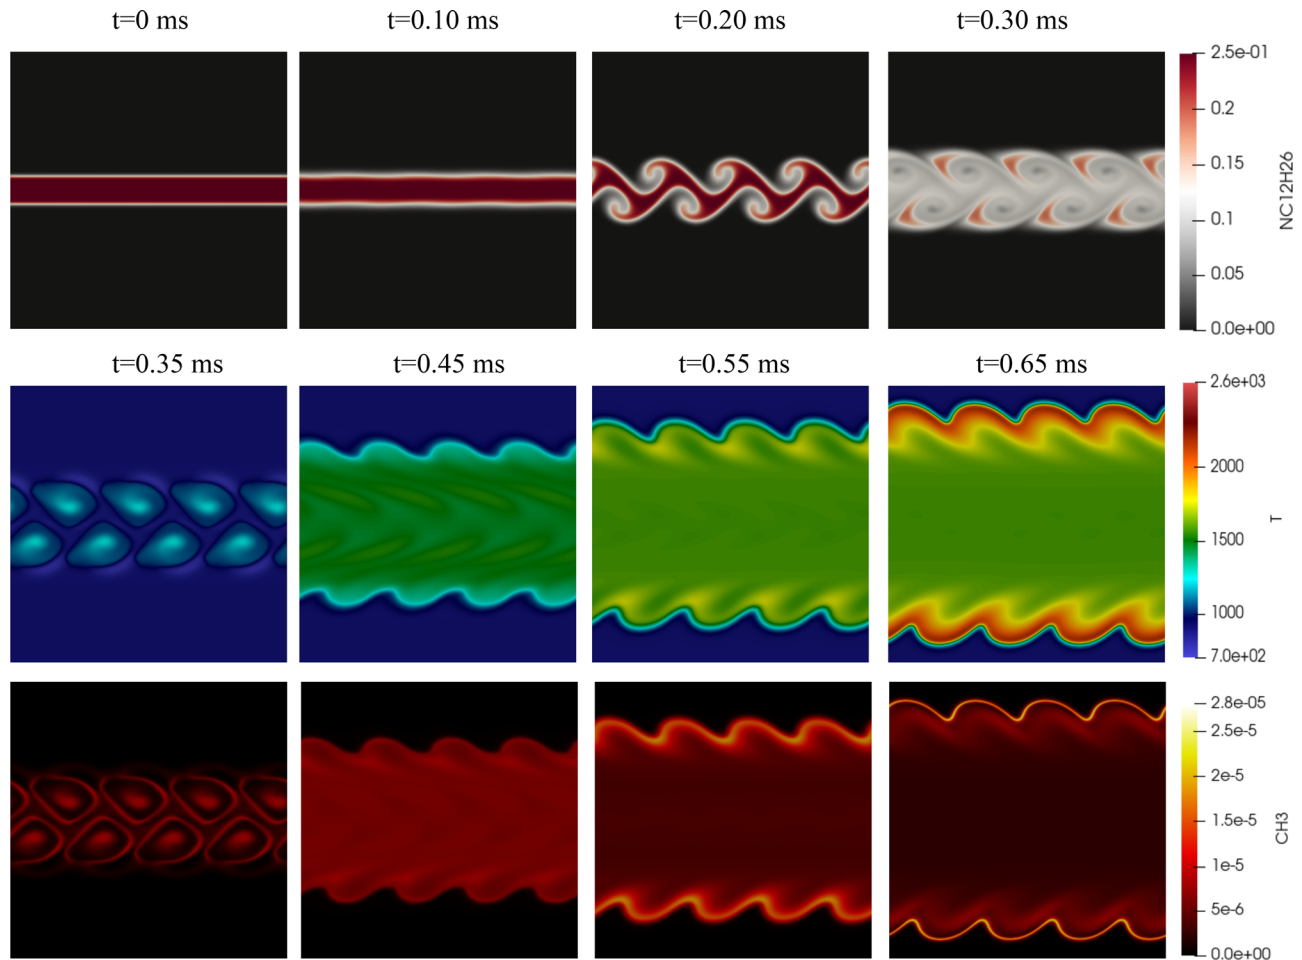


FIG. 5. Temporal evolution of *n*-dodecane (top row), temperature (middle row), and CH_3 (bottom row) fields of the 2D reacting shear layer simulation, where the *n*-dodecane jet initiates the flame and then burns the premixed ambient.

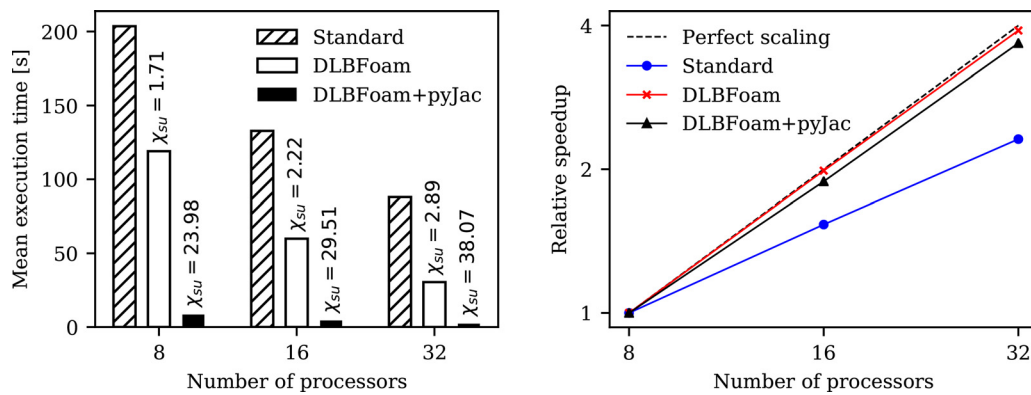


FIG. 6. Left: mean execution times over 100 iterations for the 2D reacting shear layer for different number of processors. Here, χ_{su} is the speed-up factor of DLBFoam or DLBFoam+pyJac compared to Standard. Right: speed-up values for different number of processors relative to the case with eight processors.

efficiency of DLBFoam and DLBFoam+pyJac is almost linear ($>95\%$), the Standard shows a scaling efficiency around 75% for the investigated processor counts. This explains the increased speed-up value for DLBFoam and DLBFoam+pyJac for higher processor counts. (2) The DLBFoam model provides speed-up by a factor of 1.71–2.89 for different processor counts, compared to Standard. (3) The DLBFoam+pyJac model provides speed-up by a factor of 23.98–38.07 compared to Standard.

C. 3D stratified combustion

After providing speed-up benchmarks along with scaling tests in Sec. III B, we next present another academic test case in this section, i.e., 3D stratified combustion (Fig. 7). This case cannot be even tested for ten iterations using Standard within a reasonable computational time. Hence, the purpose here is mainly to show that it is possible to investigate 3D ignition/combustion problems using DLBFoam+pyJac. This DF test case has close relevance to combustion mode design in modern engines such as RCCI.

This test case is conceptually similar to the previous 2D setup by Karimkashi *et al.*⁴⁸ As shown in the schematic of Fig. 4(b), the computational domain is a 3D cube with periodic boundary conditions and a side length $D = 1$ mm. The grid size is $4\text{ }\mu\text{m}$ with 256 grid points along each direction, to capture the flame with at least ten grid points within the laminar flame thickness ($\delta_f \approx 40\text{ }\mu\text{m}$). The initial pressure is set at 60 bar and the initial temperature is homogeneous within the entire computational domain at 800 K, which is relevant to the estimated temperature at the most reactive mixture fraction for DF methane/*n*-dodecane based on 0D simulations.^{44,46,47,49} Here, pure *n*-dodecane is initially constrained in a blob in the middle of the domain with a diameter $d = D/3$. The ambient gas consists of premixed methane, oxidizer, and EGR at $\phi_{\text{CH}_4} = 0.5$, similar to the previous test case, with the composition reported in Table II. As in the previous test case, Yao's mechanism is used.

Turbulence is initialized using the Taylor–Green–Vortex (TGV) structure, which is generated in a separate non-reactive 3D simulation with an initial velocity $U = 50\text{ m/s}$ and reference length $L_{\text{ref}} = D/2\pi$. We let the initialized TGV to evolve in time until the gradient of total




FIG. 7. Temporal evolution of temperature (top row), *n*-dodecane (second row), methane (third row), and OH (bottom row) is shown in cut-planes taken from the 3D stratified combustion simulation around the second-stage ignition time. With gradual temperature rise, ignition kernels are formed at different locations. Consistently, the fuels (*n*-dodecane and methane) are consumed and deflagrative-like fronts are formed and finally merged.

kinetic energy reaches its peak and then, the velocity and pressure fields are mapped to the reactive case as initial fields. The estimated Reynolds number at the start of reacting simulation is $Re \approx 1000$. In this case study, the interactive roles of convection, diffusion, and reaction are investigated. First, the initial *n*-dodecane is mixed with the ambient gas at early simulation time instances, long before the second-stage ignition (τ_2). Here, $\tau_2 \approx 0.28$ ms is defined as the first time instance when $T > 1500$ K. At τ_2 , the stratified mixture of *n*-dodecane and methane/oxidizer is formed before any significant temperature rise in the system ($t = 0.7\tau_2$). Closer to the ignition time, hot spots emerge ($t = 0.9\tau_2$) and deflagrative-like fronts develop at a moderate rate with various wrinkled surfaces ($t = 1.1\tau_2$). Slightly after the ignition time, the fronts evolve ($t = 1.3\tau_2$) as clearly observed from the OH cut-plane. Finally, the fronts merge ($t = 1.8\tau_2$) and combustion is almost complete. We note that at $t = 1.8\tau_2$, only small fractions of *n*-dodecane and methane are left in the system, which are convected from other iso-surfaces to the displayed cut-plane in Fig. 7.

Figure 7 depicts 2D cut-planes for the temporal evolution of temperature (top row), *n*-dodecane (second row), methane (third row), and hydroxyl radical OH (bottom row) wherein time is normalized by τ_2 . It is observed that a stratified mixture of *n*-dodecane and methane/oxidizer is formed before any significant temperature rise in the system ($t = 0.7\tau_2$). Closer to the ignition time, hot spots emerge ($t = 0.9\tau_2$) and deflagrative-like fronts develop at a moderate rate with various wrinkled surfaces ($t = 1.1\tau_2$). Slightly after the ignition time, the fronts evolve ($t = 1.3\tau_2$) as clearly observed from the OH cut-plane. Finally, the fronts merge ($t = 1.8\tau_2$) and combustion is almost complete. We note that at $t = 1.8\tau_2$, only small fractions of *n*-dodecane and methane are left in the system, which are convected from other iso-surfaces to the displayed cut-plane in Fig. 7.

In this case study, a speed-up gain by a factor of 2–3 orders of magnitude is observed with DLBFoam+pyJac compared to Standard. Therefore, it would not be possible to benchmark this problem using Standard for 100 iterations. In particular, we evaluate the speed-up of our model against Standard by restarting the simulation from the first-stage ignition time (0.2 ms) and using the same number of processors (1152 processors) and identical tolerances. While Standard simulation proceeds only 7 iterations in $\sim 76\,000$ s of clock time, the respective time is only about 200 s with DLBFoam+pyJac; i.e., speed-up by a factor of ≈ 400 with DLBFoam+pyJac compared to Standard is achieved. As a final note, the total clock time for simulating this case using DLBFoam+pyJac with 1152 processors for 0.5 ms simulation time is ≈ 13 h. Hence, DLBFoam+pyJac is a key enabler to study such 3D reactive flow configurations with OpenFOAM.

D. ECN Spray A

At this point, it is clear that DLBFoam+pyJac offers great advantages over Standard. In particular, DLBFoam+pyJac was shown to scale almost linearly with the number of processors in an academic case and enable 3D reactive simulations with a relatively large chemical kinetic mechanism and tight solver tolerances. Next, we move on to a more realistic 3D configuration, i.e., the Engine Combustion Network Spray A test case with experimental validation data. Spray A represents an igniting non-premixed diffusion flame under engine-relevant conditions. Hence, it is an optimal test case for reactive CFD large eddy simulation (LES) code validation.

A schematic of the computational setup is demonstrated in Fig. 4(c), where the domain volume corresponds to that of ECN combustion vessel by Sandia. Liquid *n*-dodecane at a temperature of 363 K and pressure of 150 MPa is injected from a $90\text{ }\mu\text{m}$ nozzle (i.e., Spray A) into a hot quiescent ambient gas of density 22.8 kg/m^3 and mixture composition (15% O_2 , 75.15% N_2 , 6.23% CO_2 , and 3.62% H_2O) based on the molar fractions. The simulations are performed for four reacting cases with varied ambient temperature (900, 1000, 1100, and 1200 K) and a single non-reacting case at 900 K. The domain is discretized with a static mesh of $62\text{ }\mu\text{m}$ cell size in the innermost refinement layer to resolve the turbulent mixing and the quasi-steady lifted flame, as shown in Fig. 8, resulting in a total of $25\text{--}39 \times 10^6$ cells depending on the particular case. Moreover, the implicit LES (ILES) approach is employed for turbulent subgrid-scale (SGS) modeling, consistent with our previous spray combustion works,^{27,44–47,50,51} while the no-breakup model approach is used for droplet atomization, which is further discussed by Gadalla *et al.*⁵¹ The volume-rendered spray flame shown in Fig. 8 for ambient $T = 900$ K indicates the mixing and evaporation zone, the low-temperature combustion zone, and the ignition and high-temperature combustion region wherein a fully developed diffusion flame is observed.

The present case demonstrates the capability of our new model to match the well-known ECN Spray A benchmark. First, in Fig. 9(a) liquid and vapor penetrations are validated against experimental data up until 1.5 ms under non-reacting conditions at 900 K. Then, the simulation is repeated in the reacting mode and the pressure rise is monitored in Fig. 9(b) against experiments for the same temporal interval. In Fig. 9(c), the formaldehyde (CH_2O) planar laser-induced fluorescence (PLIF) false color images obtained by Skeen *et al.*⁵² at various snapshots are compared against the corresponding LES data using the new solver. Here, the LES data are circumstantially averaged for each




FIG. 8. Volume rendering of the ECN Spray A diffusion flame of ambient $T = 900$ K after 1.5 ms. The inset in the top left corner presents a 2D projected cut-plane of the OH radical mass fraction.

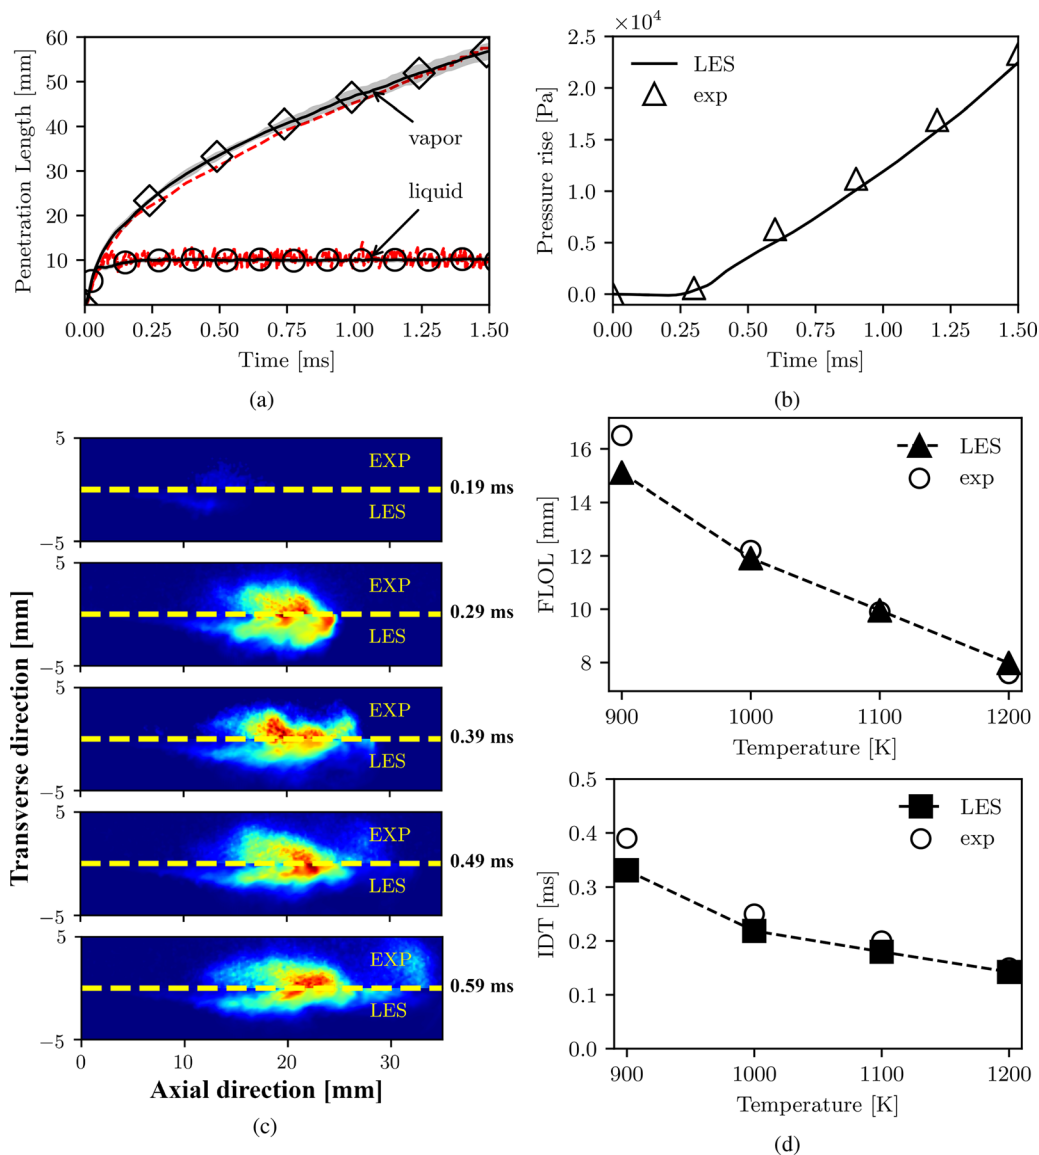


FIG. 9. Validations of the ECN Spray A at ambient $T = 900$ K for (a) liquid and vapor penetrations under non-reactive conditions, (b) pressure rise in the combustion vessel after ignition, (c) formaldehyde fields at various snapshots, and (d) ignition delay time and flame liftoff length (FLOL) at varied ambient temperature levels of the mixture.

sampling point in the axial-radial plane, which are then compared with the experimental data set. This comparison verifies the global trends for Spray A baseline including ignition onset, cool flame, and quasi-steady flame liftoff. Finally, validations of ignition delay times (IDTs) and flame liftoff lengths against experimental data at various temperature levels are presented in Fig. 9(d).

After that, we demonstrate in Fig. 10 the gained speed-up using DLBFoam+pyJac within 100 time steps after 0.3 ms of the simulation and using a constant step size of 2×10^{-7} s. The load imbalance in this test case is rather high. It originates from the large number of ambient cells that are computationally less stiff than those in the ignition region. This imbalance is further increased with reference

mapping so that ambient cells become mostly idle, hence improving the performance gain of dynamic load balancing, which results in a factor of 38 speed-up, compared with Standard. After that, a further speed-up by a factor of 6.7 is attributed to the analytical Jacobian retrieval and the robust linear algebra, i.e., DLBFoam+pyJac, hence a total speed-up by a factor of 256 as compared to Standard.

E. Sandia flame D

The ECN Spray A case demonstrated excellent performance for the LES of non-premixed Spray flame using DLBFoam+pyJac. The last test case is the Sandia flame D⁵³ representing a piloted partially




FIG. 10. Mean execution times over 65 iterations for the ECN Spray A after 0.3 ms of the simulation. χ_{su} is the speed-up factor of DLBFoam or DLBFoam+pyJac compared to Standard.

premixed methane–air flame. Sandia flame D is a well-known configuration from the TNF workshop flame series, which is often studied in the literature, e.g., Refs. 54–57. Therefore, it is considered as a useful benchmark for reactive CFD LES code validation for gas burner flame investigations.

A schematic of the Sandia flame D setup is provided in Fig. 4(d). The main jet is a mixture of CH_4 and air at equivalence ratio $\phi_{\text{CH}_4} = 3.17$. It is injected with a bulk velocity of 49.6 m/s and temperature of 294 K to the domain from a nozzle of diameter $D = 7.2$ mm. The flame is stabilized with a pilot jet, which is a hot (1880 K) lean mixture of C_2H_2 , H_2 , air, CO_2 , and N_2 . The inner and outer diameters of the pilot nozzle are 7.7 and 18.2 mm, respectively, and the flow velocity is 11.4 m/s. Jet inlets are surrounded by air co-flow with an inner diameter of 18.9 mm. A Cartesian mesh is used, and a refinement zone with cell size of $0.75 \times 0.368 \times 0.368$ mm defined from the nozzle exit until $15D$ is realized, resulting in a total of 2×10^6 cell count. Velocity fluctuations in the form proposed by Pitsch and Steiner⁵⁴ are superimposed on the main jet inlet velocity profile to induce turbulent jet behavior. Here, the purpose is to demonstrate that the code produces first-order statistics for this flame configuration. Hence, we do not focus on the turbulence–chemistry interaction modeling but use the ILES approach as a SGS scale model, consistent with the previous test cases. A reduced chemical kinetic mechanism DRM19³⁹ with 21 species and 83 reactions is used. The simulation is performed for 100 ms after which the fields are averaged for another 150 ms, where the flame becomes statistically stable.

Figure 11 depicts time-averaged mixture fraction, temperature and species mass fraction profiles compared to the experimental data⁵⁵ and an instantaneous snapshot of the temperature field. The close agreement between experiments and the predicted mixture fraction and temperature profiles along the center axis indicates that the coupled influence of turbulent mixing and flame heat release is resolved adequately in the core of the jet flame. Minor species profiles, CO and OH, also correspond well with the experiments. The predicted OH mass fractions show correct trends; however, the maximum value is overpredicted compared with the measured data. Overall, mass fraction prediction of short-lived species like OH is a challenging task due

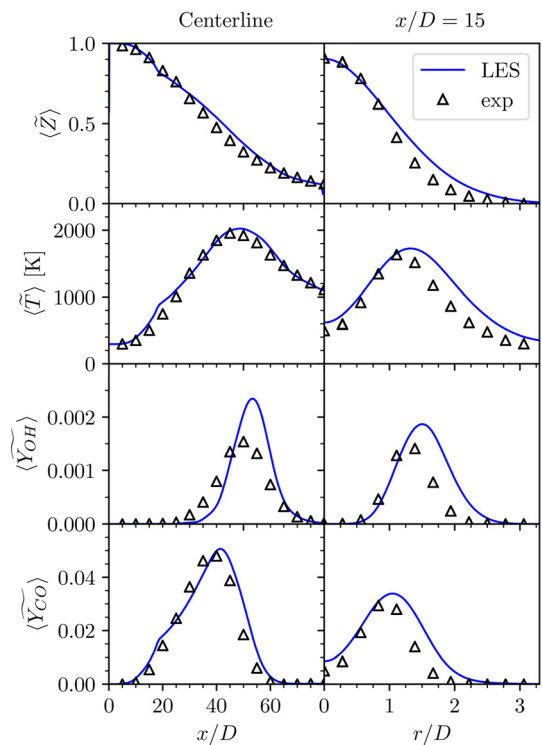
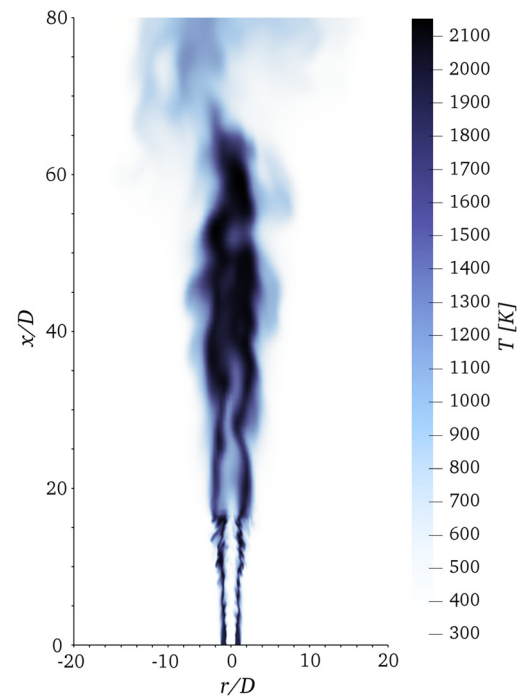


FIG. 11. Top: instantaneous temperature field of Sandia flame D at $t = 0.1$ s. Note that the data attributed to the refined grid are only available to $x/D = 15$ in order to limit the computational cost. Bottom: validations of the Sandia flame D for mean mixture fraction, temperature, and species mass fraction at centerline and $x/D = 15$.

to the nonlinear evolution of the species,^{55–58} and it is beyond the scope of the present study. In the radial direction, the reactions mostly occur in the shear layer between the main jet and co-flow ($r/D = 1$), which is evident in the peaks of the mass fraction profiles. The radial profiles agree well with the experimental data on the rich side. However, mixing is overpredicted in the lean region ($r/D > 1$). Close to the nozzle and away from the central axis, the reaction zone may still be considered relatively laminar.⁵⁴ Therefore, the overpredicted profiles in the lean region could be explained by the utilized unity Lewis number assumption. Moreover, the temperature and species profiles also follow the same trend as the mixture fraction in the flame on the rich and lean sides.

Next, a speed-up test is carried out for a stabilized flame for 100 iterations with a constant time step of 1×10^{-6} s. The results are presented in Fig. 12. The DLBFoam provides speed-up by a factor of 4.52, which is significantly lower than with ECN Spray A. This can be explained by a relatively low number of mapped cells due to the location of the refinement zone in the core of the jet. The optimized ODE solver routines pyJac and LAPACK bring additional speed-up by a factor of 2.99. Here, the lower speed-up is a consequence of the use of the reduced mechanism, which as discussed earlier significantly reduces the computational cost of the chemistry problem. However, the combined effect of DLBFoam+pyJac is significant and brings about 1 order of magnitude speed-up to this case.

IV. CONCLUSIONS

In this work, an open-source dynamic load balancing library for OpenFOAM, namely DLBFoam, with reference mapping feature was re-introduced and improved. In particular, the feature was extended with optimized chemistry ODE solution routines. First, the Jacobian evaluation procedures in OpenFOAM were replaced with an analytical implementation provided by the open-source package pyJac. This step resulted in one order of magnitude speed-up in a 0D homogeneous reactor test case. Then, in order to fully utilize the benefits of the new fully dense Jacobian, LU decomposition and back substitution routines were replaced with more robust ones from the open-source library

LAPACK, resulting in even higher speed-up especially for tighter ODE tolerances.

The improvements were tested in two academic cases and then applied to two experimental flame cases. First, the model was tested in a 2D reacting shear layer problem, demonstrating almost perfectly linear scaling. Speed-up by a factor of ≈ 30 was demonstrated in the 2D case for DLBFoam with pyJac and LAPACK. Second, 3D stratified combustion was studied, which was not even computationally feasible to benchmark with the Standard model due to 2–3 orders of magnitude performance difference. Last, two 3D flames were investigated—ECN Spray A and Sandia flame D—showing speed-up factors of ≈ 256 and ≈ 13.5 , respectively, for DLBFoam with pyJac and LAPACK. We also note that the reported speed-up numbers in comparison to Standard were achieved with fixed ODE tolerances, specific time step sizes, and chemical kinetic mechanisms. It is worth acknowledging that changing those parameters may affect the speed-up estimates. However, the promising results herein indicate that the improvements to DLBFoam offer an avenue to model complex combustion phenomena in OpenFOAM with increased accuracy and computational efficiency. Further investigations in a broader user community could result in more in-depth analysis and development of the present test cases for different combustion models and chemical mechanisms. As future work, more test cases could be established around the proposed reactive CFD approach to account for a broader variety of turbulent combustion conditions. Finally, the source code and the case setup files used in this work are openly available^{41,59}

ACKNOWLEDGMENTS

This study is financially supported by the Academy of Finland (Grant Nos. 318024, 332784, and 332835). Computational resources have been provided by CSC—Finnish IT Center for science and Aalto Science-IT project. We are grateful to Dr. Heikki Kahila (Wärtsilä Finland Oy) and Dr. Petteri Peltonen (VTT Oy) for their precious comments and technical advice throughout the implementation process.

AUTHOR DECLARATIONS

Conflict of Interest

The authors declare that they have no conflicts of interest to disclose.

DATA AVAILABILITY

The data that support the findings of this study are openly available in GitHub repository at <https://github.com/Aalto-CFD/DLBfoam> and <https://github.com/Aalto-CFD/DLBfoam-Advanced-Tutorials> Refs. 59 and 41.

REFERENCES

- ¹See <https://exafoam.eu/overview/> for “information about the exaFoam project” (last accessed October 21, 2021).
- ²T. Poinso and D. Veynante, *Theoretical and Numerical Combustion*, 2nd ed. (RT Edwards, 2005).
- ³N. Peters, *Turbulent Combustion*, *Cambridge Monographs on Mechanics* (Cambridge University Press, 2000).
- ⁴B. Tekgül, P. Peltonen, H. Kahila, O. Kaario, and V. Vuorinen, “DLBFoam: An open-source dynamic load balancing model for fast reacting flow simulations in OpenFOAM,” *Comput. Phys. Commun.* **267**, 108073 (2021).

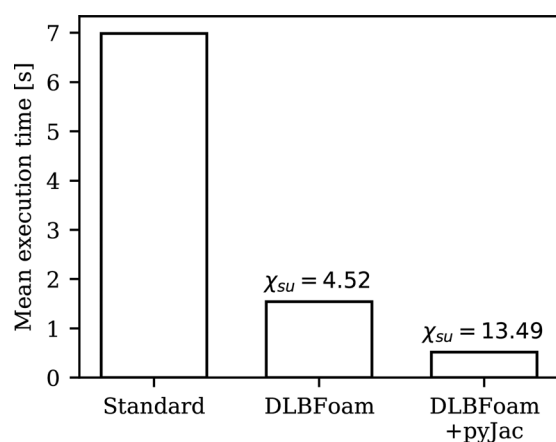


FIG. 12. Mean execution times over 100 iterations for the Sandia flame D after 0.1 s of the simulation. χ_{su} is the speed-up factor of DLBFoam or DLBFoam+pyJac compared to Standard.

- ⁵D. Kochkov, J. A. Smith, A. Alieva, Q. Wang, M. P. Brenner, and S. Hoyer, "Machine learning-accelerated computational fluid dynamics," *Proc. Natl. Acad. Sci. U. S. A.* **118**(21), e2101784118 (2021).
- ⁶M. Gadalla, M. Cianferri, M. Tezzele, G. Stabile, A. Mola, and G. Rozza, "On the comparison of LES data-driven reduced order approaches for hydroacoustic analysis," *Comput. Fluids* **216**, 104819 (2021).
- ⁷S. Hijazi, G. Stabile, A. Mola, and G. Rozza, "Data-driven POD-Galerkin reduced order model for turbulent flows," *J. Comput. Phys.* **416**, 109513 (2020).
- ⁸N. Wandel, M. Weinmann, and R. Klein, "Teaching the incompressible Navier–Stokes equations to fast neural surrogate models in three dimensions," *Phys. Fluids* **33**(4), 047117 (2021).
- ⁹Z. Li, Y. Ju, and C. Zhang, "Parallel large eddy simulations of transitional flow in a compressor cascade with endwalls," *Phys. Fluids* **31**(11), 115104 (2019).
- ¹⁰S. Brahmachary, A. Bhagyarajan, and H. Ogawa, "Fast estimation of internal flowfields in scramjet intakes via reduced-order modeling and machine learning," *Phys. Fluids* **33**(10), 106110 (2021).
- ¹¹S. J. Plimpton, S. G. Moore, A. Borner, A. K. Stagg, T. P. Koehler, J. R. Torczynski, and M. A. Gallis, "Direct simulation Monte Carlo on petaflop supercomputers and beyond," *Phys. Fluids* **31**(8), 086101 (2019).
- ¹²V. Sekar, Q. Jiang, C. Shu, and B. C. Khoo, "Fast flow field prediction over airfoils using deep learning approach," *Phys. Fluids* **31**(5), 057103 (2019).
- ¹³G. Aversano, G. D'Alessio, A. Coussement, F. Contino, and A. Parente, "Combination of polynomial chaos and Kriging for reduced-order model of reacting flow applications," *Results Eng.* **10**, 100223 (2021).
- ¹⁴A. Wehrfritz, O. Kaario, V. Vuorinen, and B. Somers, "Large eddy simulation of *n*-dodecane spray flames using Flamelet Generated Manifolds," *Combust. Flame* **167**, 113–131 (2016).
- ¹⁵V. Jaganath and M. Stoellinger, "Transported and presumed probability density function modeling of the Sandia flames with flamelet generated manifold chemistry," *Phys. Fluids* **33**(4), 045123 (2021).
- ¹⁶F. Contino, H. Jeanmart, T. Lucchini, and G. D'Errico, "Coupling of in situ adaptive tabulation and dynamic adaptive chemistry: An effective method for solving combustion in engine simulations," *Proc. Combust. Inst.* **33**(2), 3057–3064 (2011).
- ¹⁷T. Lu and C. K. Law, "Toward accommodating realistic fuel chemistry in large-scale computations," *Prog. Energy Combust. Sci.* **35**(2), 192–215 (2009).
- ¹⁸C. K. Law, "Combustion at a crossroads: Status and prospects," *Proc. Combust. Inst.* **31**(1), 1–29 (2007).
- ¹⁹Y. Pei, M. Mehl, W. Liu, T. Lu, W. J. Pitz, and S. Som, "A multicomponent blend as a diesel fuel surrogate for compression ignition engine applications," *J. Eng. Gas Turbines Power* **137**(11), 111502 (2015).
- ²⁰S. M. Sarathy, G. Kukkadapu, M. Mehl, T. Javed, A. Ahmed, N. Naser, A. Tekawade, G. Kosiba, M. AlAbbad, E. Singh, S. Park, M. A. Rashidi, S. H. Chung, W. L. Roberts, M. A. Oehlschlaeger, C.-J. Sung, and A. Farooq, "Compositional effects on the ignition of FACE gasolines," *Combust. Flame* **169**, 171–193 (2016).
- ²¹E. Hairer and G. Wanner, *Solving Ordinary Differential Equations II* (Springer Berlin Heidelberg, 1996).
- ²²M. J. McNenly, R. A. Whitesides, and D. L. Flowers, "Faster solvers for large kinetic mechanisms using adaptive preconditioners," *Proc. Combust. Inst.* **35**(1), 581–587 (2015).
- ²³F. Perini, E. Galligani, and R. D. Reitz, "An analytical Jacobian approach to sparse reaction kinetics for computationally efficient combustion modeling with large reaction mechanisms," *Energy Fuels* **26**(8), 4804–4822 (2012).
- ²⁴D. A. Schwer, J. E. Tolsma, W. H. Green, and P. I. Barton, "On upgrading the numerics in combustion chemistry codes," *Combust. Flame* **128**(3), 270–291 (2002).
- ²⁵C. Safta, H. N. Najm, and O. Knio, "TChem—A software toolkit for the analysis of complex kinetic models," Technical Report No. SAND2011-3282, Sandia National Laboratories (2011).
- ²⁶K. E. Niemeyer, N. J. Curtis, and C. J. Sung, "pyJac: Analytical Jacobian generator for chemical kinetics," *Comput. Phys. Commun.* **215**, 188–203 (2017).
- ²⁷L. Antonelli and P. D'Ambra, "Dynamic load balancing for high-performance simulations of combustion in engine applications," in *19th International Euromicro Conference on Parallel, Distributed and Network-Based Processing* (IEEE, 2011), pp. 133–140.
- ²⁸Y. Shi, W. H. Green, H. W. Wong, and O. O. Oluwale, "Accelerating multi-dimensional combustion simulations using GPU and hybrid explicit/implicit ODE integration," *Combust. Flame* **159**(7), 2388–2397 (2012).
- ²⁹J. Kodavasal, K. Harms, P. Srivastava, S. Som, S. Quan, K. Richards, and M. García, "Development of a stiffness-based chemistry load balancing scheme, and optimization of input/output and communication, to enable massively parallel high-fidelity internal combustion engine simulations," *J. Energy Resour. Technol.* **138**(5), 052203 (2016).
- ³⁰T. Zirwes, F. Zhang, P. Habisreuther, J. A. Denev, H. Bockhorn, and D. Trimis, "Optimizing load balancing of reacting flow solvers in OpenFOAM for high performance computing," in *Proceedings of 6th ESI OpenFOAM User Conference* (ESI Group, Hamburg, Germany, 2018), Vol. 6; available at https://cn.esi-group.com/sites/default/files/resource/other/7400/student-abstract_zirwes_karlsruhe-institute-of-technology_optimizing-load-balancing-of-reacting-flow-solvers-in-openfoam-for-high-performance-computing1.pdf.
- ³¹J. Muela, R. Borrell, J. Ventosa-Molina, L. Jofre, O. Lehmkuhl, and C. D. Pérez-Segarra, "A dynamic load balancing method for the evaluation of chemical reaction rates in parallel combustion simulations," *Comput. Fluids* **190**, 308–321 (2019).
- ³²E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. D. Croz, A. Greenbaum, S. Hammerling, A. McKenney, and D. Sorensen, *LAPACK Users' Guide: Third Edition*, 3rd ed. (Society for Industrial and Applied Mathematics, 1999).
- ³³H. Kahila, O. Kaario, Z. Ahmad, M. Ghaderi Masouleh, B. Tekgül, M. Larmi, and V. Vuorinen, "A large-eddy simulation study on the influence of diesel pilot spray quantity on methane-air flame initiation," *Combust. Flame* **206**, 506–521 (2019).
- ³⁴H. Kahila, "Numerical modeling of spray-assisted dual-fuel ignition," Ph.D. thesis (Aalto University, 2019).
- ³⁵A. Imren and D. Haworth, "On the merits of extrapolation-based stiff ODE solvers for combustion CFD," *Combust. Flame* **174**, 1–15 (2016).
- ³⁶See <https://openfoam.org/release/6/> for "OpenFOAM v6 Release Notes" (last accessed February 18, 2021).
- ³⁷G. P. Smith, D. M. Golden, M. Frenklach, N. W. Moriarty, B. Eiteneer, M. Goldenberg, C. T. Bowman, R. K. Hanson, S. Song, W. C. Gardiner, Jr., V. V. Lissianski, and Z. Qin, see <http://combustion.berkeley.edu/gri-mech/> for "GRI-Mech 3.0 (1999)."
- ³⁸T. Yao, Y. Pei, B. J. Zhong, S. Som, T. Lu, and K. H. Luo, "A compact skeletal mechanism for *n*-dodecane with optimized semi-global low-temperature chemistry for diesel engine simulations," *Fuel* **191**, 339–349 (2017).
- ³⁹A. Kazakov and M. Frenklach, see <http://combustion.berkeley.edu/drm/> for "Reduced Reaction Sets Based on GRI-Mech 1.2 (2005)" (last accessed October 21, 2021).
- ⁴⁰R. Issa, B. Ahmadi-Befrui, K. Beshay, and A. Gosman, "Solution of the implicitly discretised reacting flow equations by operator-splitting," *J. Comput. Phys.* **93**(2), 388–410 (1991).
- ⁴¹I. Morev, M. Gadalla, S. Karimkashi, J. Kannan, A. Shahanaghi, B. Tekgül, O. Kaario, and V. Vuorinen, see <https://github.com/Aalto-CFD/DLBFOam-Advanced-Tutorials> for "GitHub Repository: DLBFOam-Advanced-Tutorials" (2021).
- ⁴²See <https://docs.csc.fi/computing/systems-mahti/> for "Mahti Supercomputer Specifications" (last accessed October 21, 2021).
- ⁴³H. Kahila, A. Wehrfritz, O. Kaario, and V. Vuorinen, "Large-eddy simulation of dual-fuel ignition: Diesel spray injection into a lean methane-air mixture," *Combust. Flame* **199**, 131–151 (2019).
- ⁴⁴B. Tekgül, H. Kahila, O. Kaario, and V. Vuorinen, "Large-eddy simulation of dual-fuel spray ignition at different ambient temperatures," *Combust. Flame* **215**, 51–65 (2020).
- ⁴⁵B. Tekgül, H. Kahila, S. Karimkashi, O. Kaario, Z. Ahmad, É. Lendormy, J. Hyvönen, and V. Vuorinen, "Large-eddy simulation of spray assisted dual-fuel ignition under reactivity-controlled dynamic conditions," *Fuel* **293**, 120295 (2021).
- ⁴⁶J. Kannan, M. Gadalla, B. Tekgül, S. Karimkashi, O. Kaario, and V. Vuorinen, "Large eddy simulation of diesel spray-assisted dual-fuel ignition: A comparative study on two *n*-dodecane mechanisms at different ambient temperatures," *Int. J. Engine Res.* **22**, 2521–2532 (2021).

- ⁴⁷J. Kannan, M. Gadalla, B. Tekgül, S. Karimkashi, O. Kaario, and V. Vuorinen, "Large-eddy simulation of tri-fuel ignition: Diesel spray-assisted ignition of lean hydrogen-methane-air mixtures," *Combust. Theory Modell.* **25**, 436–459 (2021).
- ⁴⁸S. Karimkashi, H. Kahila, O. Kaario, M. Larmi, and V. Vuorinen, "A numerical study on combustion mode characterization for locally stratified dual-fuel mixtures," *Combust. Flame* **214**, 121–135 (2020).
- ⁴⁹S. Karimkashi, H. Kahila, O. Kaario, M. Larmi, and V. Vuorinen, "Numerical study on tri-fuel combustion: Ignition properties of hydrogen-enriched methane-diesel and methanol-diesel mixtures," *Int. J. Hydrogen Energy* **45**(7), 4946–4962 (2020).
- ⁵⁰M. Gadalla, J. Kannan, B. Tekgül, S. Karimkashi, O. Kaario, and V. Vuorinen, "Large-eddy simulation of tri-fuel combustion: Diesel spray assisted ignition of methanol-hydrogen blends," *Int. J. Hydrogen Energy* **46**, 21687 (2021).
- ⁵¹M. Gadalla, J. Kannan, B. Tekgül, S. Karimkashi, O. Kaario, and V. Vuorinen, "Large-eddy simulation of ECN spray A: Sensitivity study on modeling assumptions," *Energies* **13**(13), 3360 (2020).
- ⁵²S. A. Skeen, J. Manin, and L. M. Pickett, "Simultaneous formaldehyde PLIF and high-speed Schlieren imaging for ignition visualization in high-pressure spray flames," *Proc. Combust. Inst.* **35**(3), 3167–3174 (2015).
- ⁵³R. S. Barlow and J. H. Frank, "Effects of turbulence on species mass fractions in methane/air jet flames," in *Symposium (International) on Combustion* (Elsevier, 1998), Vol. 27, pp. 1087–1095.
- ⁵⁴H. Pitsch and H. Steiner, "Large-eddy simulation of a turbulent piloted methane/air diffusion flame (Sandia flame D)," *Phys. Fluids* **12**(10), 2541–2554 (2000).
- ⁵⁵A. W. Vreman, B. A. Albrecht, J. A. van Oijen, L. P. de Goeij, and R. J. Bastiaans, "Premixed and nonpremixed generated manifolds in large-eddy simulation of Sandia flame D and F," *Combust. Flame* **153**(3), 394–416 (2008).
- ⁵⁶A. Bertels, B. Kober, A. Rittler, and A. Kempf, "Large-eddy simulation of Sandia flame D with efficient explicit filtering," *Flow, Turbul. Combust.* **102**(4), 887–907 (2019).
- ⁵⁷A. Garmory and E. Mastorakos, "Sensitivity analysis of LES-CMC predictions of piloted jet flames," *Int. J. Heat Fluid Flow* **39**, 53–63 (2013).
- ⁵⁸D. A. Lysenko, I. S. Ertesvåg, and K. E. Rian, "Numerical simulations of the Sandia Flame D using the eddy dissipation concept," *Flow, Turbul. Combust.* **93**(4), 665–687 (2014).
- ⁵⁹B. Tekgül, P. Peltonen, H. Kahila, I. Morev, and M. Gadalla, see <https://github.com/Aalto-CFD/DLBFoam> for "GitHub Repository: DLBFoam (2021)"