

Computational Physics



Method for scalable and performant GPU-accelerated simulation of multiphase compressible flow [☆]

Anand Radhakrishnan ^a, Henry Le Berre ^a, Benjamin Wilfong ^a, Jean-Sebastien Spratt ^c, Mauro Rodriguez Jr. ^d, Tim Colonius ^c, Spencer H. Bryngelson ^{a,b,*}

^a School of Computational Science & Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

^b Daniel Guggenheim School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

^c Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA 91125, USA

^d School of Engineering, Brown University, Providence, RI 02912, USA

ARTICLE INFO

Keywords:

Computational fluid dynamics
Heterogeneous computing
Multiphase flows

ABSTRACT

Multiphase compressible flows are often characterized by a broad range of space and time scales, entailing large grids and small time steps. Simulations of these flows on CPU-based clusters can thus take several wall-clock days. Offloading the compute kernels to GPUs appears attractive but is memory-bound for many finite-volume and -difference methods, damping speedups. Even when realized, GPU-based kernels lead to more intrusive communication and I/O times owing to lower computation costs. We present a strategy for GPU acceleration of multiphase compressible flow solvers that addresses these challenges and obtains large speedups at scale. We use OpenACC for directive-based offloading of all compute kernels while maintaining low-level control when needed. An established Fortran preprocessor and metaprogramming tool, Fypp, enables otherwise hidden compile-time optimizations. This strategy exposes compile-time optimizations and high memory reuse while retaining readable, maintainable, and compact code. Remote direct memory access realized via CUDA-aware MPI and GPUDirect reduces halo-exchange communication time. We implement this approach in the open-source solver MFC [1]. Metaprogramming results in an 8-times speedup of the most expensive kernels compared to a statically compiled program, reaching 46% of peak FLOPs on modern NVIDIA GPUs and high arithmetic intensity (about 10 FLOPs/byte). In representative simulations, a single NVIDIA A100 GPU is 7-times faster compared to an Intel Xeon Cascade Lake (6248) CPU die, or about 300-times faster compared to a single such CPU core. At the same time, near-ideal (97%) weak scaling is observed for at least 13824 GPUs on OLCF Summit. A strong scaling efficiency of 84% is retained for an 8-times increase in GPU count. Collective I/O, implemented via MPI3, helps ensure the negligible contribution of data transfers (< 1% of the wall time for a typical, large simulation). Large many-GPU simulations of compressible (solid-)liquid-gas flows demonstrate the practical utility of this strategy.

1. Introduction¹

Multiphase compressible flows are ubiquitous, with examples such as the atomization of liquid droplets [2], bubble cavitation [3], and shock-wave attenuation of nuclear blasts [4]. Collapsing bubble clouds in cavitating flows can result in shock waves that lead to large pressures. This has applications in a wide variety of engineering problems, such as the design of mechanical heart valves [5], burst-wave lithotripsy [6], and minimizing blast-induced trauma [7]. These high pressures also

cause erosion of industrial equipment in flow around hydrofoils [8], pumps [9], and propellers [10]. Simulation of multiphase phenomena is thus critical to enable engineering design and minimize equipment damage.

Courant–Friedrichs–Lewy (CFL) constraints for compressible flow restrict permissible time step sizes. Thus, many time steps are needed to simulate the relevant physical phenomena, making the minimization of wall time for each time step critical. Since 2004, CPU clock speeds have plateaued, ending Dennard scaling. Thus, the compute capabilities of

[☆] The review of this paper was arranged by Prof. Andrew Hazel.

* Corresponding author at: School of Computational Science & Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA.

E-mail address: shb@gatech.edu (S.H. Bryngelson).

¹ All code available at <https://github.com/MFlowCode/MFC>.

modern supercomputers stem primarily from GPU accelerators. Leveraging GPUs is thus essential to extracting meaningful speedups from state-of-the-art supercomputers.

Multiphase compressible flow algorithms consist mostly of level 1 BLAS (vector) stencil operations, so most are memory bound [11]. The low arithmetic intensity of these kernels prevents efficient use of the GPU's compute capabilities. Faster GPU kernels can also lead to prominent MPI communication and I/O transfer times. This can exacerbate strong scaling behavior for smaller problem sizes with large numbers of GPUs. We present strategies that address the concerns above to obtain satisfactory performance on accelerators. The portability of our acceleration strategies is ensured by conducting tests on various architectures. Large multi-GPU simulations are included to emphasize the benefits of our strategies in pertinent applications.

GPU speedups, scaling tests, and example simulations are conducted with the open-source solver MFC [1]. Interface capturing methods [12], particularly the 5- and 6- equation models [13–15], are used to represent the multi-component flow. These equations are discretized and solved using a shock-capturing finite volume scheme that uses high-order accurate WENO reconstructions [16,17]. The Riemann problem is then solved using an HLLC approximate Riemann solver [18], and the solution is evolved with a total-variation-diminishing (TVD) Runge–Kutta time stepper [19]. The GPU acceleration strategies outlined in this work take advantage of the increased arithmetic intensity of high-order accurate methods typically used for multicomponent flow simulations.

Large-scale compressible flow simulations have been conducted on CPUs for a long time and still gather substantial attention [20]. However, attempts at optimizing such solvers for GPU acceleration are less unified as the GPU hardware landscape evolves. We cite solvers like STREAMS (version 1 and 2) [21,22] and ZEFR [23] as just a couple of demonstrative examples of these efforts. STREAMS simulates compressible turbulent flow via a flux vector splitting method and achieves 250-times speedups on a single NVIDIA V100 GPU over an Intel Skylake CPU core. They observe 97% weak scaling efficiency for up to 1024 V100 GPUs along with 90% strong scaling efficiency for an 8-times increase in GPU count. ZEFR employs similar strategies and observes similar accelerations, simulating single-phase compressible flows and retaining 70% strong scaling efficiency for an 8-times increase in GPU count. These solvers do not address the challenges associated with multicomponent flows, which is part of our focus.

The flexibility associated with the GPU programming model is a concern of increasing importance. CUDA offers reliable performance on NVIDIA GPUs and has been a GPU-programming mainstay since its inception in 2007. However, other vendors like AMD and Intel are deploying competitive GPU accelerators in the most capable new supercomputers, like OLCF Frontier, CSC LUMI, and ALCF Aurora [24–26]. Vendor-specific programming models like CUDA are insufficient to take advantage of the capability these new supercomputers bring. Deploying performant and vendor-agnostic fluid flow solvers on new computers requires adopting more flexible programming models. Here, we use OpenACC [27], a performance-competitive directive-based model with established support for NVIDIA GPUs [28] and increasing support for AMD and Intel hardware [29]. The FluTAS solver [30] also adopted OpenACC, solving the incompressible multiphase flow problem via a finite difference scheme. However, speedups are limited in the incompressible flow case due to communication times associated with pressure-Poisson solves and Fourier transforms. In their study, FluTAS displayed linear weak scaling and a 40% retention in strong scaling efficiency for an 8-times increase in GPU count on the MeluXina supercomputer [31]. URANOS [32] uses OpenACC similarly for turbulent compressible flows, demonstrating the approach's efficacy for compressible CFD applications. However, the algorithms employed degrade performance by 20% when weak scaling up to 300 GPUs, and the speedups are limited compared to the ones we present here.

We describe the computational models used to formulate the governing equations in section 2. The numerical method that solves the

discrete conservation laws is outlined in section 3. Section 4 describes optimal GPU acceleration and MPI communication strategies. Results for model validation, GPU speedups, and scaling tests are presented in section 5. The benefits of GPU acceleration for large multiphase problems are illustrated via example simulations in section 6. Section 7 highlights the relevant conclusions from this work.

2. Computational model

We briefly describe the multicomponent models used in GPU acceleration tests. These models are reduced from the non-equilibrium Baer–Nunziato model [33].

2.1. 5-equation models

The so-called Kapila 5-equation model [13] is obtained from the non-equilibrium Baer–Nunziato model [33] under the assumptions of velocity and pressure equilibrium between the phases. The equations for 2 components are

$$\begin{aligned}\frac{\partial \alpha_1 \rho_1}{\partial t} + \nabla \cdot (\alpha_1 \rho_1 \mathbf{u}) &= 0, \\ \frac{\partial \alpha_2 \rho_2}{\partial t} + \nabla \cdot (\alpha_2 \rho_2 \mathbf{u}) &= 0, \\ \frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u} + p \mathbf{I}) &= 0, \\ \frac{\partial \rho E}{\partial t} + \nabla \cdot [(\rho E + p) \mathbf{u}] &= 0, \\ \frac{\partial \alpha_1}{\partial t} + \mathbf{u} \cdot \nabla \alpha_1 &= K \nabla \cdot \mathbf{u},\end{aligned}$$

where ρ , \mathbf{u} , p , and E are the mixture density, velocity, pressure, and energy, and α_i are the volume fractions of component i . The system of equations is closed using an equation of state (EOS). Here, we use the stiffened gas EOS, which can faithfully model many liquids and gases [34]:

$$\rho E = \frac{1}{\gamma - 1} p + \frac{\gamma \pi_\infty}{\gamma - 1}, \quad (1)$$

though other relations can be used as appropriate. For a 2-component problem, we have

$$K = \frac{\rho_2 c_2^2 - \rho_1 c_1^2}{\frac{\rho_2 c_2^2}{\alpha_2} + \frac{\rho_1 c_1^2}{\alpha_1}}, \quad (2)$$

and $K \nabla \cdot \mathbf{u}$ represents the expansion and compression of each phase in mixture regions and ensures thermodynamic consistency via the conservation of phase entropy. This admits a consistent representation of the sound speed in the mixture region, though it can lead to numerical instabilities due to the non-conservative source term in the volume fraction advection equation [16].

The $K \nabla \cdot \mathbf{u}$ term can be ignored in cases where mixture compression effects are unimportant, though it is unclear how to determine this a priori. For example, a case where they are known to be important is spherical bubble dynamics [35]. If one can ignore this term, the equations degenerate to the Allaire model [14]. Though the Allaire model is conservative, it does not strictly obey the second law of thermodynamics.

2.2. A 5-equation model with hypoelasticity

A hypoelastic material model represents the elastic response of solids [36]. The model is obtained by modifying the 5-equation model. An elastic shear stress term $\tau_{ij}^{(e)}$ modifies the Cauchy stress tensor as

$$\sigma_{ij} = -p \delta_{ij} + \tau_{ij}^{(v)} + \tau_{ij}^{(e)}, \quad (3)$$

where $\boldsymbol{\tau}^{(v)}$ is the viscous stresses. An elastic contribution $e^{(e)}$ contributes to the total energy E as

$$E = e + \frac{\|\mathbf{u}\|^2}{2} + e^{(e)} \quad \text{where} \quad e^{(e)} = \frac{(\boldsymbol{\tau}_{ij}^{(e)})^2}{4\rho G}. \quad (4)$$

Additional equations are required to track the elastic stresses. In 3D, this is 6 additional equations, one for each stress term $\tau_{ij}^{(e)}$ where $i, j \in \{1, 2, 3\}$ and $\boldsymbol{\tau}^{(e)}$ symmetric. With elastic contributions and additional equations, the hypoelastic 5-equation model for 2 materials is

$$\begin{aligned} \frac{\partial \alpha_1 \rho_1}{\partial t} + \nabla \cdot (\alpha_1 \rho_1 \mathbf{u}) &= 0, \\ \frac{\partial \alpha_2 \rho_2}{\partial t} + \nabla \cdot (\alpha_2 \rho_2 \mathbf{u}) &= 0, \\ \frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u} + p \mathbf{I}) + \nabla \cdot (\boldsymbol{\tau}^{(e)} + \boldsymbol{\tau}^{(v)}) &= 0, \\ \frac{\partial \rho E}{\partial t} + \nabla \cdot [(\rho E + p)\mathbf{u}] - \nabla \cdot [(\boldsymbol{\tau}^{(e)} + \boldsymbol{\tau}^{(v)})\mathbf{u}] &= 0, \\ \frac{\partial \alpha_1}{\partial t} + \mathbf{u} \cdot \nabla \alpha_1 &= K \nabla \cdot \mathbf{u}, \\ \frac{\partial \tau_{il}^{(e)}}{\partial t} + \nabla \cdot (\tau_{il}^{(e)} \mathbf{u}) &= S_{il}^{(e)}, \end{aligned}$$

where

$$S_{il}^{(e)} = \rho \left(\tau_{kj}^{(e)} \frac{\partial u_i}{\partial x_k} + \tau_{ik}^{(e)} \frac{\partial u_j}{\partial x_k} - \tau_{ij}^{(e)} \frac{\partial u_k}{\partial x_k} + 2G \dot{\epsilon}_{ij}^{(d)} \right). \quad (5)$$

2.3. 6-equation model with p -relaxation

The numerical instabilities introduced by the Kapila model can be alleviated via the pressure disequilibrium model of Saurel et al. [15]. The system of equations is first evolved without the source terms, followed by a pressure relaxation step under the assumption of infinite stiffness for the pressure relaxation coefficient as discussed in Schmidmayer et al. [35].

3. Numerical method

Herein, we describe the numerical method that evaluates the governing 5/6-equation models of section 2.

3.1. Finite volume method (FVM)

A finite volume numerical scheme that follows Coralic and Colonna [16] is used to discretize the governing equations

$$\frac{\partial \mathbf{q}}{\partial t} + \frac{\partial \mathbf{F}^x(\mathbf{q})}{\partial x} + \frac{\partial \mathbf{F}^y(\mathbf{q})}{\partial y} + \frac{\partial \mathbf{F}^z(\mathbf{q})}{\partial z} = \mathbf{s}(\mathbf{q}) - \mathbf{h}(\mathbf{q}) \nabla \cdot \mathbf{u}, \quad (6)$$

where \mathbf{q} and \mathbf{F} represent the conservative variables and fluxes in the governing equations. The finite volume method represents the conservative variables $\mathbf{q}_{i,j,k}$ centered at the location (x_i, y_j, z_k) . The dimensions of the cell are

$$I_{i,j,k} = [x_{i-1/2}, x_{i+1/2}] \times [y_{j-1/2}, y_{j+1/2}] \times [z_{k-1/2}, z_{k+1/2}], \quad (7)$$

with grid spacing

$$\Delta x_i = x_{i+1/2} - x_{i-1/2}, \quad \Delta y_j = y_{j+1/2} - y_{j-1/2}, \quad \Delta z_k = z_{k+1/2} - z_{k-1/2}. \quad (8)$$

The PDE (6) is integrated in space across each cell center as

$$\begin{aligned} \frac{\partial \mathbf{q}_{i,j,k}}{\partial t} &= \frac{1}{\Delta x_i} (\mathbf{F}_{i-1/2,j,k}^x - \mathbf{F}_{i+1/2,j,k}^x) + \frac{1}{\Delta y_j} (\mathbf{F}_{i,j-1/2,k}^y - \mathbf{F}_{i,j+1/2,k}^y) \\ &+ \frac{1}{\Delta z_k} (\mathbf{F}_{i,j,k-1/2}^z - \mathbf{F}_{i,j,k+1/2}^z) + \mathbf{s}(\mathbf{q}_{i,j,k}) - \mathbf{h}(\mathbf{q}_{i,j,k}) (\nabla \cdot \mathbf{u})_{i,j,k}, \end{aligned} \quad (9)$$

where $\mathbf{q}_{i,j,k}$ are the volume averaged conservative variables, and $\mathbf{s}(\mathbf{q}_{i,j,k})$ and $\mathbf{h}(\mathbf{q}_{i,j,k})$ are the volume averaged source terms. The flux term $\mathbf{F}^x(\mathbf{q}_{i+1/2,j,k})$ is obtained by averaging over the finite volume cell face $A_{i+1/2,j,k}$ centered at $(x_{i+1/2}, y_j, z_k)$. The other flux terms in (9) are obtained using a similar procedure. The divergence term is computed as

$$\begin{aligned} (\nabla \cdot \mathbf{u})_{i,j,k} &= \frac{1}{\Delta x_i} (u_{i+1/2,j,k}^x - u_{i-1/2,j,k}^x) + \frac{1}{\Delta y_j} (u_{i,j+1/2,k}^y - u_{i,j-1/2,k}^y) \\ &+ \frac{1}{\Delta z_k} (u_{i,j,k+1/2}^z - u_{i,j,k-1/2}^z), \end{aligned} \quad (10)$$

where $u_{i+1/2,j,k}^x$ is the x direction of the velocity averaged across the cell face at grid point $A_{i+1/2,j,k}$.

3.2. Shock capturing via reconstruction

The numerical scheme in section 3.1 requires the reconstruction of the fluxes \mathbf{F} and the velocity \mathbf{u} at the cell faces. The state variable \mathbf{q} at a cell face $A_{i+1/2,j,k}$ is reconstructed from the face's left and right sides, resulting in a discontinuity. The resulting flux and velocity at the cell face are obtained by solving a Riemann problem at the interface [37]:

$$\begin{aligned} \mathbf{F}_{i+1/2,j,k}^x &= \mathbf{F}^x(\mathbf{q}_{i+1/2,j,k}^L, \mathbf{q}_{i+1/2,j,k}^R) \\ \mathbf{u}_{i+1/2,j,k}^x &= \mathbf{u}^x(\mathbf{q}_{i+1/2,j,k}^L, \mathbf{q}_{i+1/2,j,k}^R). \end{aligned}$$

The superscripts L and R denote the reconstructed state variable at the left and right cell faces. A first-order, total variation diminishing approximation of the state variables at the interface $A_{i+1/2,j,k}$ follows as

$$\mathbf{q}_{i+1/2,j,k}^L = \mathbf{q}_{i,j,k} \quad \text{and} \quad \mathbf{q}_{i+1/2,j,k}^R = \mathbf{q}_{i+1,j,k}. \quad (11)$$

This scheme suppresses spurious oscillations at interfaces, but can lose accuracy via their smearing. For this, high-order accurate reconstructions at the interface help keep interfaces sharp for a given grid resolution. Here, we will show results for a fifth-order-accurate WENO reconstruction, though the method is also performant for the third-order accurate variant.

WENO reconstruction follows from a convex combination of interpolating polynomials on candidate stencils. A $(2k - 1)$ th-order WENO reconstructed state variable $f_{i+1/2,j,k}$ is obtained by a weighted sum of k candidate polynomials as

$$f_{i+1/2,j,k} = \sum_{r=0}^k \omega_{i+1/2,j,k}^r f_{i+1/2,j,k}^r. \quad (12)$$

The weights ω^r are obtained from ideal weights w^r using smoothness indicators β^r . Additional implementation details for WENO3 and WENO5 are in Shu [38]. High-order WENO reconstructions are not total variation diminishing (TVD). This makes it susceptible to spurious oscillations at material interfaces. To suppress these oscillations, the conservative variables \mathbf{q} are converted to the primitive ones before reconstruction [16].

3.2.1. Approximate Riemann solver

The Riemann problem in section 3.2 is solved using an approximate HLLC (Harten-Lax-van Leer contact) Riemann solver [39]. The HLLC Riemann solver admits three discontinuities in the solution with wave speeds s_L , s_* , and s_R . The resulting state at the cell interface $A_{i+1/2,j,k}$ is given as

$$\mathbf{q}_{i+1/2,j,k} = \begin{cases} \mathbf{q}_{i+1/2,j,k}^L & 0 \leq s_L \\ \mathbf{q}_{i+1/2,j,k}^{L*} & s_L \leq 0 \leq s_* \\ \mathbf{q}_{i+1/2,j,k}^{R*} & s_* \leq 0 \leq s_R \\ \mathbf{q}_{i+1/2,j,k}^R & 0 \geq s_R \end{cases}. \quad (13)$$

The wave speeds s_L and s_R are estimated using the state variables $q_{i+1/2,j,k}^L$ and $q_{i+1/2,j,k}^R$. In order to calculate the intermediate states $q_{i+1/2,j,k}^{L*}$ and $q_{i+1/2,j,k}^{R*}$, continuity of normal velocity ($u^{L*} = u^{R*} = u^*$) and pressure ($p^{L*} = p^{R*} = p^*$) is imposed across the contact discontinuity with speed $s_* = u^*$. The flux at the cell interface is then calculated as

$$F(q_{i+1/2,j,k}) = \begin{cases} F(q_{i+1/2,j,k}^L) & 0 \leq s_L \\ F(q_{i+1/2,j,k}^L) + s_L(q_{i+1/2,j,k}^{L*} - q_{i+1/2,j,k}^L) & s_L \leq 0 \leq s_* \\ F(q_{i+1/2,j,k}^R) + s_R(q_{i+1/2,j,k}^{R*} - q_{i+1/2,j,k}^R) & s_* \leq 0 \leq s_R \\ F(q_{i+1/2,j,k}^R) & 0 \geq s_R \end{cases} \quad (14)$$

3.3. Boundary conditions

Boundary conditions are implemented by allocating buffers at the domain edges. Buffer sizes are determined based on the order of WENO reconstruction. Boundary conditions for time-dependent hyperbolic systems require knowledge of solutions exterior to the computational domain. We use characteristic-based boundary conditions for this purpose [40].

3.4. Time stepping

Once the right-hand side of (9) is determined using spatial reconstruction coupled with a Riemann solver, the solution is evolved in time by discretizing the time derivative. For this purpose, a high-order total variation diminishing Runge–Kutta time stepper [19] achieves temporal accuracy while suppressing spurious oscillations.

4. Implementation strategy

4.1. Domain decomposition and I/O

Distributed computing techniques are essential to improve performance at large scales. This is achieved by decomposing the domain into 3D blocks across multiple processors. The block dimensions for each processor are kept uniform across all dimensions instead of splitting across a single dimension (slabs). This is done to minimize the data communicated at the processor boundaries.

We utilize a structured mesh with non-uniform spacing to discretize the domain. Cartesian and Cylindrical geometries are supported, with the axisymmetric option available for cylindrical grids. For 3D cylindrical geometries, fast Fourier transforms maintain finite spacing along the circumferential direction near the center. This is implemented with the FFTW package on CPUs and cuFFT on NVIDIA GPUs [41]. Grid stretching through a hyperbolic tangent function is done to locally refine the grids near locations of interest [42].

High-order reconstruction at processor boundaries requires knowledge of the state variables outside the local domain. Communication of the boundary data among adjacent processors, called the halo exchange, is thus required across all dimensions at each time step. Blocking send and receive MPI calls transfer data in the halo region [43]. Data in the halo regions are packed into 1D buffers for compatibility with the MPI subroutines and then unpacked into the requisite data structures. Ideal weak and strong scaling is observed on multiple CPU cores [1].

Unsteady compressible flow simulations require I/O data saves at a fixed number of time steps to obtain the temporal variation of the solution. We use a parallel I/O framework for the file systems, which enables collective MPI read and write functions [44]. I/O data dumps can thus be performed in a distributed manner across all processors, ensuring scalability of the code on modern supercomputers. Silo files are exported from the data for analysis and visualization [45].

```
#:for NORM_DIR, dir in [(1, 'x'), (2, 'y'), (3, 'z')]
if (norm_dir == ${NORM_DIR}) then
!$acc parallel loop gang vector collapse (3) private(dvd,
poly, beta, alpha, omega)
do l = is3%beg, is3%end !third coordinate direction
do k = is2%beg, is2%end !second coordinate direction
do j = is1%beg, is1%end !first coordinate direction
!$acc loop seq
do i = 1, sys_size
!$acc loop seq
do q = -2, 1 !compute divided differences
dvd(q) = v_${dir}$(j + q + 1, k, l, i) &
- v_${dir}$(j + q, k, l, i)
end do
!$acc loop seq
do q = 0, 2
poly(q) = v_${dir}$(j, k, l, i) &
+ poly_coef_${dir}$(j, q, 0)*dvd(q+1) &
+ poly_coef_${dir}$(j, q, 1)*dvd(q)
beta(q) = beta_coef_${dir}$(j, q,
0)*dvd(1-q)*dvd(1-q) &
+ beta_coef_${dir}$(j, q,
1)*dvd(1-q)*dvd(-q) &
+ beta_coef_${dir}$(j, q, 2)*dvd(
-q)*dvd(-q)
end do
alpha = d_${dir}$(:, j)/(beta**2)
omega = alpha/sum(alpha)
vL_${dir}$(j, k, l, i) = sum(omega*poly)
end do
end do
end do
end if
#:endfor
```

Listing 1: A Fortran90+Fypc code snippet of the WENO5 reconstruction kernel.

4.2. GPU offloading

All computation within a time step is offloaded to GPUs via OpenACC [27]. After applying the initial condition in the pre-processing step, the state variables are copied to the GPUs. The state variables are transferred back to the CPUs only during I/O data saves, typically occurring once every one thousand steps for large problems. Directive-based offloading in OpenACC only requires the specification of the location of independent loops and the level of parallelization. This allows the compiler to decide the optimum kernel parameters for maximum speedup. Another advantage of directive-based offloading is maintaining a common codebase for the CPU and GPU versions and enabling accelerators via a compiler flag. cuTENSOR and cuFFT libraries are used for optimized tensor reshapes and Fourier transforms on GPUs.

Our OpenACC implementation uses well-established directives, making it portable. The code is compatible with NVHPC, GNU, and Cray (CCE) compilers on NVIDIA GPUs and AMD GPU support for CCE and GNU. Here, we conduct simulations on NVIDIA V100 (OLCF Summit) and A100 GPUs (OLCF Wombat) using the NVHPC compiler. CPU simulations are also supported on various architectures using NVHPC and GCC compilers with appropriate optimization flags. We test CPU simulations on Intel Xeon Cascade Lake CPUs (PSC's Bridges2), IBM Power9 CPUs (OLCF Summit), and Arm CPUs (OLCF Wombat).

A code snippet example of an OpenACC kernel used in MFC is given in Listing 1. The kernel uses 5th-order-accurate WENO to reconstruct the state variables at the cell faces. WENO reconstruction constitutes about 40% of the total time step on GPUs, making it the most expensive OpenACC kernel. Various optimization techniques used in Listing 1 to improve kernel performance are detailed in section 4.3.

High-order WENO reconstruction of state variables entails larger stencils. Larger stencils are commensurate with larger halo regions at the processor boundaries and longer communication times. Halo exchange times on GPUs can be significant owing to faster kernel times

over CPUs. Moreover, the proportion of data on the processor boundaries increases with decreasing problem size. Minimizing MPI communication time is thus essential to retain GPU speedups for smaller problem sizes. For this purpose, we use CUDA-aware MPI and GPUDirect RDMA to accelerate communication by using fast GPU interconnects [46]. We observe 4-times performance improvement for halo exchanges using CUDA-aware MPI over conventional MPI communication with Spectrum MPI 10.4 on OLCF Summit.

4.3. Optimization

OpenACC kernels use gangs, workers, and vectors to distribute the workload, which maps to blocks, warps, and vectors in CUDA notation. The `parallel loop` constructs in OpenACC split the loop iterations across gangs by default. This usually leads to each block using a single OpenACC vector, resulting in a huge waste of resources. The `parallel loop` construct in Listing 1 is thus augmented with a `gang` vector clause that enables the splitting of the loop iterations across multiple gangs of fixed vector length [47]. The compiler chooses the optimum number of gangs and vectors for efficient resource allocation [47]. Furthermore, the three loops across multiple dimensions in Listing 1 are combined into a single loop using the `collapse(3)` clause. This allows the compiler to choose optimal gang and vector sizes for a specific architecture based on the total problem size, thus eliminating the effect of any skewness in the grid dimensions. The fourth loop in Listing 1 benefits from serialization due to its smaller loop bounds, with `sys_size` typically between 5 and 9 for multiphase problems.

Compressible flow algorithms are primarily vector operations, typically leading to low arithmetic intensity. Potential for improved GPU performance was observed in the computationally intensive WENO reconstruction kernel, outlined in Listing 1, through a high degree of memory reuse. Flattened multidimensional arrays are preferred over user-defined data types, allowing for aggressive compiler optimizations in GPU kernels. A 6-times performance improvement was observed using multidimensional arrays for the kernel in Listing 1 over user-defined data types for a 3D problem with 1M (10^6) grid points.

Memory coalescence is needed to saturate the GPU's global memory bandwidth. This can be achieved by ensuring that successive threads in a parallelized loop access consecutive memory locations. Previous serialized implementations benefited from switching the inner and outer loops in Listing 1 corresponding to iteration variables `j` and `k`. This allowed for reusability of the coefficient variables (`poly_coef`, `beta_coef`) across the iterations of the inner loop. However, reordering the two loops, as in Listing 1, ensured thread coalescence across inner iterations and proved beneficial for GPU implementation.

Dimensional splitting in the finite volume method requires independent reconstructions and flux additions across all dimensions. This allows for temporarily reshaping the state variables to achieve coalesced memory access. The additional cost of reshaping state variables is mitigated by the high degree of reuse of the reshaped variables in the computationally intensive GPU kernels. The state variables (`v`, `vL`) in Listing 1 are reshaped so that the fastest-changing index `j` matches the direction of reconstruction. Reshaping the state variables for the kernel in Listing 1 results in a 10-times speedup for a 3D problem with 1M grid points per NVIDIA V100 GPU. The cuTENSOR library reshapes these arrays on NVIDIA GPUs, though we observe similar performance when reshaping manually via array loops.

Metaprogramming techniques, in this case, enabled by the Fypp preprocessor [48], further improves GPU kernel performance. It allows for user inputs to be passed as compile-time constants, which are used to allocate fixed-size thread-local arrays in GPUs. Their availability as fixed-size arrays allows the compiler to introduce additional optimizations in the kernel and thus primarily (fast) register memory accesses. Fixed parameters are also used to eliminate large conditional blocks in kernels, leading to improved times via judicious use of available GPU registers. Kernel duplication across multiple dimensions in Listing 1 is also

eliminated through the use of Fypp macros (`#:for NORM_DIR`). This strategy results in 8- and 2-times speedups of the two most expensive kernels associated with WENO and the Riemann solver, respectively. OpenACC does not automatically inline serial subroutines within GPU kernels. This can cause a 10-times slowdown of kernels with sequential subroutine calls. Fypp enables manual inlining of these subroutines to retain compiler optimizations and speedups. Fypp is only a metaprogramming tool, so it does not generate any code that a programmer could not write themselves. However, it does write code a programmer is *unlikely* to write themselves out of the sheer work and repeated optimizations required. In the case of MFC, fypp also enables a reduced line count by about a factor of 10.

The same core kernel code is used for CPUs and GPUs. After optimization, we found a minor speedup in the CPU simulations before GPU porting. Thus, no meaningful performance difference is observed on CPUs by optimizing for GPUs via OpenACC offloading.

4.4. Validation

MFC has been validated on test cases via comparisons to experimental results for various physical problems such as shock-bubble interaction, shock-droplet, spherical bubble collapse, and Taylor–Green vortices [1]. The results are verified to be high-order accurate by monitoring error convergence. Solutions are also free of any spurious oscillations at material interfaces.

5. Performance results

The 5-equation model in section 2.1 is used to obtain performance metrics for both CPUs and GPUs. Similar GPU speedups and scaling are observed for the 5-equation hypoelastic model in section 2.2 and the 6-equation model in section 2.3.

5.1. Scaling

The scalability of the code is probed by examining wall times for a 3D 2-component problem with air and water on OLCF Summit. The hardware details of Summit are available in Vergara Larrea et al. [49]. In brief, Summit is an IBM system of AC922 POWER nodes. It contains 4600 compute nodes, each with 2 IBM POWER9 processors and 6 NVIDIA V100 GPUs. Each POWER9 chip is connected via NVLINK with a bandwidth of 25 GB/s. Each node has 96 GB of HBM2 GPU memory. The wall times are averaged across 10 time steps during a simulation, which is sufficient for variability to be sufficiently small that it is visually indistinguishable on subsequent plots. The processor with maximum wall time is used to examine results.

5.1.1. Weak scaling

The problem size per GPU or CPU is 1M grid points in 3D with 1 MPI rank per processor. Here, the problem size increases proportionately to the number of processors to test weak scaling. Fig. 1 shows the weak scaling performance using NVIDIA V100 GPUs and POWER9 CPUs on OLCF Summit. The wall times are normalized using the base case with 216 GPUs and 112 CPU cores, respectively. A near-ideal efficiency of 97% is observed for at least 13824 GPUs. Efficiency is within 1% of ideal performance for at least 14336 CPU cores. A higher weak scaling efficiency on CPUs over GPUs can be attributed to the negligible contribution of halo exchanges.

5.1.2. Strong scaling

We use the same 3D 2-component problem as the weak scaling test on Summit to observe strong scaling performance. This is achieved by fixing the total problem size and increasing the number of processors. The base case uses 8 GPUs or CPU cores, with 1 MPI rank per core. For GPU cases, this corresponds to one core binding to each GPU. We use problem sizes of 64 and 16 M grid points with the base case using 8 and




Fig. 1. Weak scaling for a 3D 2-phase problem on (a) NVIDIA V100 GPUs and (b) POWER9 CPUs.




Fig. 2. Strong scaling results on POWER9 CPUs for a 3D 2-component problem with sized as labeled.

Table 1

Strong scaling simulation times for the 64M grid point case of Fig. 3 (b).

GPUs	CUMPI		MPI	
	Actual [s]	Ideal [s]	Actual [s]	Ideal [s]
8	0.273	0.273	0.341	0.341
16	0.146	0.137	0.199	0.171
32	0.0710	0.0683	0.101	0.0853
64	0.0408	0.0341	0.0542	0.0426
128	0.0286	0.0171	0.0417	0.0213

2 M grid points per GPU or CPU, respectively. For simulations with GPU quantities that are not multiple of 6, the extra GPUs are spilled onto the next node.

Fig. 2 shows strong scaling performance on POWER9 CPUs on OLCF Summit. We observe ideal scaling varying from one to six nodes. This indicates that communication is negligible for CPU simulations and problems of appreciable size, which is discussed further in section 5.2. This contrasts against GPU simulations, for which computational time is markedly shorter but network communication times are not. Fig. 3 shows strong scaling performance using NVIDIA V100 GPUs with and without GPU-aware MPI on OLCF Summit. An increase in processor count leads to decreased problem size per GPU. Deviation from ideal performance is expected due to the increasing share of MPI communication.

CUDA-aware MPI and the GPUDirect RDMA use fast GPU interconnects to achieve up to 4-times faster halo exchanges. This can be observed for the largest problem size (64M) in Fig. 3 and in more detail in Table 1, retaining 84% of ideal performance for a factor of 8 increase in processor count. A larger deviation from ideal performance is observed without CUDA-aware MPI. The problem size is chosen to be close to the memory limit of an NVIDIA V100-16 GB SMX2 GPU. However, deviations from ideal behavior are observed for smaller problem sizes (16M) at larger numbers of GPUs (64 and 128). This can be attributed to MPI communication dominating the simulation time ($> 50\%$ of wall

Table 2

Percentage contributions of the most expensive subroutines per time step. The GPUs are NVIDIA V100s, and the CPUs are IBM POWER9s. The problem size per CPU or GPU as labeled (1M and 8M).

Subroutine	GPU(1M)%	GPU(8M)%	CPU(1M)%	CPU(8M)%
Reconstruction	37.0	43.3	72.9	71.7
Riemann Solve	24.7	33.3	16.4	18.7
Communication	17.6	5.81	3.64	3.83
Other	20.7	17.6	7.06	5.77

time) as the compute required per GPU diminishes more quickly but the communication does not. Of course, one expects a full plateau and nearly 100% of the time spent doing communication for a sufficiently large number of GPUs.

5.2. Profiles and I/O

We monitor the contribution of various subroutines in the code to the total simulation time for GPUs and CPUs. This is done by measuring the wall times of the most expensive routines for a single time step. The wall times are averaged across 10-time steps, and the processor with maximum wall time is used for runs with multiple processors. Table 2 shows the percentage contribution of most expensive subroutines for a 3D 2-component problem using 4 NVIDIA V100 GPUs on OLCF Summit. CPU simulations use GNU compilers with `-Ofast` optimization flag. We test a relatively small problem with 4M grid points (1M per processor or accelerator) and a larger problem with 32M grid points (8M per processor or accelerator). The larger problem size is chosen to be close to the hardware limit of an NVIDIA V100-16 GB SMX2 GPU (16 GB).

WENO reconstruction is observed to be the most expensive kernel in MFC, taking around 40% and 70% of the time step for GPUs and CPUs. A relatively smaller contribution of the WENO kernel on GPUs indicates large speedups for this kernel. MPI communication time on GPUs has a decreasing contribution to the overall simulation time with increasing problem size. The contribution of MPI communication on CPUs is relatively insignificant ($\approx 1\%$) and invariant to problem size. It should be noted that I/O data saves, while insignificant on CPUs, require the equivalent of about 10 time step units on GPUs. However, CFL-restricted time step sizes mean data saves are typically conducted once every 1000 steps. Hence, simulation times remain largely unaffected by I/O transfers on all architectures.

5.3. Kernel performance

The analysis of system profiles in section 5.2 indicates that WENO and Riemann solver kernels account for over 75% of the simulation time on GPUs for large problem sizes (8 M per GPU). Individual speedups of the most expensive subroutines using GPUs over CPUs for the larger simulation in section 5.2 are given in Table 3. Large speedups (480-times) for the most expensive WENO kernel on GPUs result from coa-




Fig. 3. Strong scaling analysis on NVIDIA V100 GPUs with and without CUDA-aware MPI (CUMPI) for a 3D 2-component problem with problem sizes (number of grid points) as labeled.

Table 3

Speedups (times faster) for the most expensive subroutines and the overall time step using V100 GPUs over a POWER9 CPU for a 3D 2-component problem with 8M grid points on OLCF Summit. This problem uses 4 GPUs and 4 CPU cores.

Subroutine	Speedup
Reconstruction	486
Riemann Solver	189
Communication	125
Total	305

Table 4

Comparison of wall times per time step on various architectures. The Intel Xeon Gold chips are the Cascade Lake architecture, and Ampere indicates the Ampere Altra Q80-30 chip. The A/V100 GPU simulations use the NVHPC v22.1 compiler, and the CPUs use GNU v11.1 (the fastest among tested compilers).

	# Cores	Time [s]	Slowdown
A100	—	0.28	Ref.
V100	—	0.50	1.7
Xeon	40	2.1	7.3
Ampere	40	2.7	9.2
Power9	42	3.5	12




Fig. 4. Roofline analysis for the two most expensive kernels, WENO and Riemann solvers. The kernels and rooflines are associated with double precision (DP) computation and fused multiply-add (FMA) operations.

lesced memory access and a high degree of reuse. The Riemann solvers kernel achieves lower speedups on GPUs due to a large number of thread-local variables and a low degree of memory reuse. Speedups in the MPI subroutines are observed due to faster buffer packing and unpacking on GPUs and CUDA-aware MPI. A 300-times speedup is observed for the time step using a single V100 GPU over a POWER9 CPU core. On a single Summit compute node, this translates to a 40-times speedup using the GPUs over only the CPUs.

A roofline analysis of the two most expensive kernels, WENO reconstruction and Riemann solvers, on NVIDIA A100 and V100 GPUs and double precision rooflines using NVIDIA Nsight Compute [50] are shown in Fig. 4. A high arithmetic intensity of 10 FLOP/byte and over 40% of the peak double precision FLOPs are observed for the WENO kernel on both GPUs. The kernel achieves a bandwidth of 305 GB/s on the V100 GPU, with the peak bandwidth being 900 GB/s. Kernel optimizations and high memory reuse are largely responsible for efficiently utilizing compute resources. In contrast, the Riemann solvers have a

lower arithmetic intensity (2 FLOP/byte) and only achieve about 10% of the peak double precision FLOPs on both GPUs. This can be attributed to low memory reuse and an abundance of conditional statements in this kernel. However, efforts to alleviate performance issues by employing single-precision arithmetic in the Riemann solvers kernel are currently being pursued.

5.4. Architecture comparisons

The portability of our acceleration strategies is verified by testing performance on various hardware architectures. We tested performance on several CPUs: Ampere Altra Q80-30 (located on OLCF Wombat), Intel Xeon Gold Cascade Lake (SKU 6248 m, PSC Bridges2), and IBM Power9 (OLCF Summit). NVHPC and GCC11 compilers were tested with `-fast` and `-ofast` compiler optimization flags, respectively. GPU performance was analyzed for the NVIDIA V100 (OLCF Summit) and A100 (OLCF Wombat) using the NVHPC 22.1 compiler with the `-ofast` flag. All computations are double precision. A 3D two-component problem with 16 M grid points on 2 MPI ranks was used for testing. The wall time was averaged over 10 time steps.

Table 4 shows average wall times and relative performance metrics for the different hardware. The “Time” column has little absolute meaning, with the relative performance being the most meaningful (also shown in the last column). In Table 4, the performance of a single GPU is compared to that of the CPU chip, with the CPU wall times normalized by the number of CPU cores per chip.

The results show that the A100 GPU is 1.72-times faster than the V100 on OLCF Summit, faster than even the peak double-precision performance would anticipate between the two cards (a factor of 1.24). This can be attributed to higher memory bandwidth (1.7-times) and faster GPU interconnects (2-times) on an A100 over a V100. A single A100 also gives a 7.3-times speedup over the Intel Xeon Cascade Lake, the fastest CPU chip. Increasing the number of GPU ranks has a negligible effect ($\approx 15\%$) on performance due to the relatively small

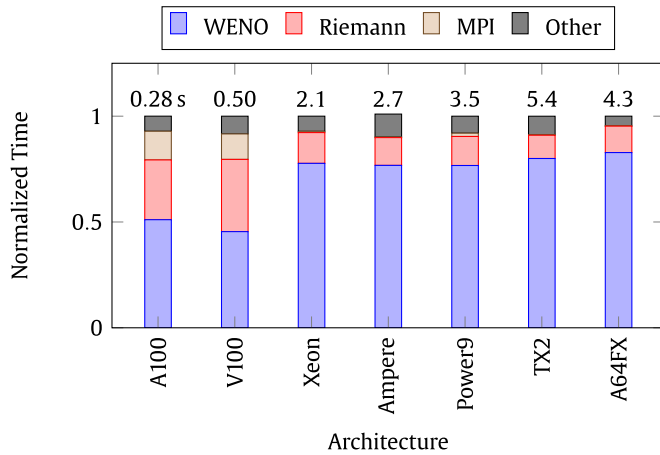


Fig. 5. Cost breakdown of different MFC subroutines on various architectures for a 3D 2-component problem with 16 M grid points. Cases V100 and A100 have all compute kernels on the respective GPUs, so the associated CPU architecture is not meaningful. Numbers above the bars indicate the absolute wall time in seconds, also shown in Table 4, which are the fastest times amongst tested compilers (NVHPC and GCC). The CPU cases use GCC v11.1 compilers except for the A64FX case, which used NVHPC 22.1 compilers. The GPU cases use the NVHPC 22.1 compiler.

contribution of the MPI subroutine section 5.3. The GNU11 compilers give shorter wall times than the NVHPC compilers on all CPU architectures. The Ampere Altra CPUs are 1.4-times faster when compared to the POWER9s and 1.2-times slower than the Intel Xeons. The Arm-based CPUs (Ampere Altra) are more energy- and cost-efficient than their x86 counterparts (Intel Xeon) under their reduced instruction set (RISC). While performance on x86 CPUs is superior, future trends indicate a growing popularity for Arm chips in high-performance computing applications [51]. The performance for different Arm chips, including the TX2 and A64FX, and their compilers are discussed for various applications, including MFC, in Elwasif et al. [52].

Fig. 5 shows a timestep normalized breakdown of the duration of the most expensive MFC routines. The two left columns indicate kernel times on GPUs; the rest are CPU-only. All computation is offloaded to accelerators, with CPUs being used only for I/O operations and halo exchanges on systems with no support for GPU-aware MPI. MPI communication constitutes a meaningful proportion of the total time on GPUs while being negligible on CPUs. This can be attributed to much larger speedups in computation on GPUs over CPUs as compared to MPI communication. The relative proportion of various routines remains nearly constant across CPU and GPU architectures. In addition to the architectures of Table 4, Fig. 5 includes results for older Arm architectures like the ThunderX2 (TX2) and Fujitsu A64FX. These processors are modestly slower than more modern x86 or Arm chips, though the profiles are qualitatively similar to the other CPUs.

6. Example simulations

We demonstrate the capability and flexibility of the method described via a large multiphase simulation of different test configurations. GPU speedups are consistent with those of section 5.

6.1. Cavitating bubble cloud

The first example case simulates the collapse of 50 air bubbles in water at ambient pressure (1 atm) near a wall when subject to a 10 atm pressure wave. The density of the fluids follows from those at standard temperature and pressure. The stiffened-gas equation of state represents both fluids [34]. This example simulation uses 216 M grid points and 60 grid points across each initial bubble diameter. We simulate for 3×10^5

time steps, corresponding to a dimensionless time $\tau \equiv tc/R_0 = 42$ for air sound speed c and equilibrium bubble radii R_0 .

Fig. 6 shows the wall pressure contours and streamlines of the collapsing bubble cloud and the isocontours of air volume fraction. Here, we use 216 GPUs (36 nodes) on Summit, or 1 M points per GPU, which amounts to 3 hours of wall-time for this simulation. Using POWER9 CPUs on the same compute nodes would require more than 2 days for the same problem. We save the simulation state to disk every 1000 time steps. Each data export requires about ten steps worth of wall time, so this cost has a negligible impact on performance.

6.2. Shock–bubble–cloud–stone interaction

The approach of this manuscript can also simulate the shock-induced collapse of air bubbles near a model kidney stone. This, in part, demonstrates the efficiency and capabilities of the method. The configuration of interest is similar in spirit to shock- and burst-wave lithotripsy [6,53].

Here, we consider a dispersion of 17 air bubbles initially near a model stone and submerged in water. The impinging shock has a Mach number $M_s = 7.92$. The BegoStone material represents the stone [54], a kidney stone phantom in lithotripsy research trials [55]. BegoStone has a density of $\rho = 1995 \text{ kg/m}^3$ and longitudinal and transverse wavespeeds $c_L = 4159 \text{ m/s}$ and $c_T = 2319 \text{ m/s}$.

The hypoelastic model described in section 2.2 represents elastic stresses in the stone and carries out this simulation. The simulation domain is $2.67D$ in the mean flow direction and $1D$ in the transverse and spanwise directions, where D is the stone’s diameter. A structured $1600 \times 600 \times 600$ Cartesian grid (576M grid points) discretizes the domain. This simulation was conducted on 576 GPUs (96 nodes) on OLCF Summit for 25×10^3 time steps, corresponding to a wall-time of 30 minutes.

Fig. 7 shows maximum principal stresses (defined in the usual way) in the stone at $\tau \equiv tc/R_0 = 2.03$ for air sound speed c and initial bubble radius R_0 . These stresses follow from the shock, and the later bubble collapses. As expected, we observe larger stresses near the center of the bubble cloud and stone cross-section.

6.3. Atomizing droplet

The third example simulation shows the atomization of a 3D water droplet in air impinged by a Mach 1.46 shock wave. The domain is $20D$ in the mean flow direction and $10D$ in the transverse and spanwise directions, where D is the initial droplet diameter. A $2000 \times 1000 \times 1000$ Cartesian grid (2B grid points) discretizes the domain. Stretching locally refines and stretches the grid near the droplet as

$$x_{\text{stretch}} = x + \frac{x}{a_x} \left[\log \left(\cosh \left(\frac{a_x(x - x_a)}{L} \right) \right) + \log \left(\cosh \left(\frac{a_x(x - x_b)}{L} \right) \right) - 2 \log \left(\cosh \left(\frac{a_x(x_b - x_a)}{2L} \right) \right) \right] \quad (15)$$

where a_x is the stretch magnitude, L is the domain length, and x_a and x_b control where stretching occurs. The droplet is centered at the origin, with $x_a = -1.2D$ and $x_b = 1.2D$ and stretching factor $a_x = 4$ across all coordinate directions. This flow is simulated for 2×10^5 time steps with a dimensionless timestep $\Delta \tau = 9.6 \times 10^{-6}$, and the simulation state is saved every 1000 time steps. The simulation is conducted on OLCF Summit using 960 GPUs (160 nodes), which amounts to 4 hours of wall-time.

Fig. 8 shows the vorticity (ω) magnitude isosurface $\|\omega\| = 10^5$ at dimensionless time $\tau = 1.43$, where $\tau \equiv tu_g/D\sqrt{\rho_g/\rho_l}$, for shock velocity u_g and post-shock gas and liquid densities ρ_g and ρ_l . The shock wave leads to two counter-rotating vorticity streams in the wake of the




Fig. 6. Illustrative MFC simulation of a collapsing bubble cloud near a wall. Results show the bubbles as contours of $\alpha_1 = 0.5$ and pressures and streamlines as labeled. In (a), red colors are larger pressures, and blue colors are lower. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)




Fig. 7. Illustrative simulation of a kidney stone near a collapsing bubble cloud. Reds indicate higher stresses, and blues indicate lower stresses. The bubble and stone isosurfaces are shown for volume fraction $\alpha = 0.5$.




Fig. 8. Vorticity isocontours $\|\omega\| = 10^5$ around a shedding water droplet.

droplet, forming a recirculation region as shown in Fig. 8. This simulation agrees with previous studies of droplet aerobreakup [56], though it requires shorter wall times and makes judicious use of the latest accelerators.

7. Conclusions

We present a GPU-accelerated multiphase compressible flow solver, MFC, capable of simulating various physical phenomena. Optimization techniques are used to facilitate efficient memory use in computationally expensive kernels. This enables higher arithmetic intensity over

typical CFD algorithms for compressible flow. Further performance improvements are observed through metaprogramming techniques via Fypp. Portability of the implementation is ensured through directive-based offloading via OpenACC, and performance is tested on NVIDIA GPUs as well as Intel Xeon, IBM POWER9, and ARM CPUs. A 40-times speedup is observed on a single Summit node using NVIDIA V100 GPUs over POWER9 CPUs.

Multi-GPU performance is examined by conducting weak and strong scaling tests. Near ideal weak scaling (with 3%) is observed for up to 13824 GPUs on Summit. CUDA-aware MPI coupled with the GPUDirect RDMA further reduces communication overhead on GPUs, reducing

simulation time for all cases tested. Level-3 MPI parallel I/O ensures that the cost of data dumps is negligible for large simulations. Scaling and speedup results show significant improvement over previous multiphase implementations [30] and are comparable to previous single-phase implementations [21,23].

The ability to conduct large multiphase simulations was tested by conducting large multi-GPU simulations across many compute nodes. The example simulations performed $\mathcal{O}(10^5)$ time steps within a few hours on NVIDIA GPUs instead of a few days on current x86 and ARM multicore CPU processors. This strategy is poised to enable the efficient use of existing and upcoming exascale systems like OLCF Frontier and LLNL El Capitan for state-of-the-art compressible multiphase flow simulations.

CRediT authorship contribution statement

Anand Radhakrishnan: Writing – review & editing, Writing – original draft, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Henry Le Berre:** Writing – original draft, Software, Methodology, Conceptualization. **Benjamin Wilfong:** Software, Methodology. **Jean-Sebastien Spratt:** Software. **Mauro Rodriguez:** Software. **Tim Colonius:** Supervision, Funding acquisition. **Spencer H. Bryngelson:** Writing – review & editing, Writing – original draft, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Spencer Bryngelson reports financial support was provided by Office of Naval Research.

Data availability

I have shared all code and the link is in the footnote on page 1.

Acknowledgements

We acknowledge useful discussions of this work from Brent Leback and Mat Colgrove (NVIDIA), Stéphan Ethier (Princeton), Nicholson Koukpaizan (Oak Ridge National Lab), Pedro Costa (TU Delft), and Luca Brandt (KTH, Sweden). SHB acknowledges the support of this work via the US Office of Naval Research under grant number N00014-22-1-2519 (PM Julie Young), hardware gifts from the NVIDIA Corporation, and use of OLCF Summit and Wombat under allocation CFD154. TC acknowledges support via the US Office of Naval Research under grant number N00014-22-1-2518 (PM Julie Young). This work used Bridges2 at the Pittsburgh Supercomputing Center through allocation PHY210084 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296.

References

- [1] S.H. Bryngelson, K. Schmidmayer, V. Coralic, J.C. Meng, K. Maeda, T. Colonius, MFC: an open-source high-order multi-component, multi-phase, and multi-scale compressible flow solver, *Comput. Phys. Commun.* 266 (2021) 107396.
- [2] J.C. Meng, Numerical Simulations of Droplet Aerobreakup, Ph.D. thesis, California Institute of Technology, 2016.
- [3] C. Brennen, Cavitation in medicine, *Interface Focus* 5 (2015) 20150022.
- [4] A. Chauvin, G. Jourdan, E. Daniel, L. Houas, R. Tosello, Experimental investigation of the propagation of a planar shock wave through a two-phase gas-liquid medium, *Phys. Fluids* 23 (2011) 113301.
- [5] P. Johansen, Mechanical heart valve cavitation, *Expert Rev. Med. Devices* 1 (2004) 95–104.
- [6] K. Maeda, A.D. Maxwell, W. Kreider, T. Colonius, M.R. Bailey, Investigation of the energy shielding of kidney stones by cavitation bubble clouds during burst wave lithotripsy, in: J. Katz (Ed.), *Proceedings of the 10th International Symposium on Cavitation (CAV2018)*, ASME Press, 2018, pp. 626–630.
- [7] P. Movahed, W. Kreider, A.D. Maxwell, S.B. Hutchens, J.B. Freund, Cavitation-induced damage of soft materials by focused ultrasound bursts: a fracture-based bubble dynamics model, *J. Acoust. Soc. Am.* 140 (2016) 1374–1386.
- [8] J.H. Seo, S.K. Lele, Numerical investigation of cloud cavitation and cavitation noise on a hydrofoil section, in: *CAV2009*, 2009, p. 15.
- [9] L. d'Agostino, M.V. Salvetti (Eds.), *Cavitation Instabilities and Rotordynamic Effects in Turbopumps and Hydroturbines: Turbopump and Inducer Cavitation, Experiments and Design*, CISM International Centre for Mechanical Sciences, vol. 575, Springer International Publishing, Cham, 2017.
- [10] L. Jofre, J. Urzay, Transcritical diffuse-interface hydrodynamics of propellants in high-pressure combustors of chemical propulsion systems, *Prog. Energy Combust. Sci.* 82 (2021) 100877.
- [11] K. Datta, S. Kamil, S. Williams, L. Oliker, J. Shalf, K. Yelick, Optimization and performance modeling of stencil computations on modern microprocessors, *SIAM Rev.* 51 (2009) 129–159.
- [12] R. Saurel, C. Pantano, Diffuse-interface capturing methods for compressible two-phase flows, *Annu. Rev. Fluid Mech.* 50 (2018) 105–130.
- [13] A.K. Kapila, R. Menikoff, J.B. Dzil, S.F. Son, D.S. Stewart, Two-phase modeling of deflagration-to-detonation transition in granular materials: reduced equations, *Phys. Fluids* 13 (2001) 3002–3024.
- [14] G. Allaire, S. Clerc, S. Kokh, A five-equation model for the simulation of interfaces between compressible fluids, *J. Comput. Phys.* 181 (2002) 577–616.
- [15] R. Saurel, F. Petitpas, R.A. Berry, Simple and efficient relaxation methods for interfaces separating compressible fluids, cavitating flows and shocks in multiphase mixtures, *J. Comput. Phys.* 228 (2009) 1678–1712.
- [16] V. Coralic, T. Colonius, Finite-volume WENO scheme for viscous compressible multicomponent flows, *J. Comput. Phys.* 274 (2014) 95–121.
- [17] G.-S. Jiang, C.-W. Shu, Efficient implementation of weighted ENO schemes, *J. Comput. Phys.* 126 (1996) 202–228.
- [18] E.F. Toro, *Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction*, 3 ed., Springer, Dordrecht/New York, 2009.
- [19] S. Gottlieb, C.-W. Shu, Total variation diminishing Runge–Kutta schemes, *Math. Comput.* 67 (1998) 73–85.
- [20] U. Rasthofer, F. Wermelinger, P. Karnakov, J. Šukys, P. Koumoutsakos, Computational study of the collapse of a cloud with 12500 gas bubbles in a liquid, *Phys. Rev. Fluids* 4 (2019) 063602.
- [21] M. Bernardini, D. Modesti, F. Salvatore, S. Pirozzoli, STREAmS: a high-fidelity accelerated solver for direct numerical simulation of compressible turbulent flows, *Comput. Phys. Commun.* 263 (2021) 107906.
- [22] M. Bernardini, D. Modesti, F. Salvatore, S. Sathyanarayana, G. Della Posta, S. Pirozzoli, STREAmS-2.0: supersonic turbulent accelerated Navier–Stokes solver version 2.0, *Comput. Phys. Commun.* 285 (2023) 108644.
- [23] J. Romero, J. Crabill, J. Watkins, F. Witherden, A. Jameson, ZEFR: a GPU-accelerated high-order solver for compressible viscous flows using the flux reconstruction method, *Comput. Phys. Commun.* 250 (2020) 107169.
- [24] V. Melesse Vergara, R. Budiardja, M. Davis, M. Ezell, J. Hanley, C. Zimmer, M. Brim, W. Elwasif, D. Dietz, Approaching the Final Frontier: Lessons Learned from the Deployment of HPE/Cray EX Spock and Crusher supercomputers, Technical Report, Oak Ridge National Lab (ORNL), Oak, Ridge, TN (United States), 2022.
- [25] T. Zwinger, J. Heikonen, P. Manninen, LUMI supercomputer for European researchers, Technical Report, Copernicus Meetings, 2023.
- [26] H. Jiang, Intel's Ponte Vecchio GPU: Architecture, Systems & Software, in: 2022 IEEE Hot Chips 34 Symposium (HCS), IEEE Computer Society, 2022, pp. 1–29.
- [27] S. Wienke, P. Springer, C. Terboven, D. an Mey, OpenACC—first experiences with real-world applications, in: *European Conference on Parallel Processing*, Springer, 2012, pp. 859–870.
- [28] M. Khalilov, A. Timoveev, Performance Analysis of CUDA, OpenACC and OpenMP Programming Models on TESLA V100 GPU, *Journal of Physics: Conference Series*, vol. 1740, IOP Publishing, 2021, p. 012056.
- [29] A. Jarmusch, A. Liu, C. Munley, D. Horta, V. Ravichandran, J. Denny, K. Friedline, S. Chandrasekaran, Analysis of validating and verifying OpenACC compilers 3.0 and above, in: 2022 Workshop on Accelerator Programming Using Directives (WACCPD), IEEE, 2022, pp. 1–10.
- [30] M. Crialesi-Esposito, N. Scapin, A.D. Demou, M.E. Rosti, P. Costa, F. Spiga, L. Brandt, FluTAS: a GPU-accelerated finite difference code for multiphase flows, *Comput. Phys. Commun.* 284 (2023) 108602.
- [31] S. Varette, Uni. Lu HPC Annual Report 2020, Technical Report, University of Luxembourg, 2021.
- [32] F. De Vanna, F. Avanzi, M. Cogo, S. Sandrin, M. Bettencourt, F. Picano, E. Benini, URANOS: a GPU accelerated Navier–Stokes solver for compressible wall-bounded flows, *Comput. Phys. Commun.* 287 (2023) 108717.
- [33] N. Andrianov, G. Warnecke, The Riemann problem for the Baer–Nunziato two-phase flow model, *J. Comput. Phys.* 195 (2004) 434–464.
- [34] O. Le Métayer, R. Saurel, The Noble–Abel stiffened-gas equation of state, *Phys. Fluids* 28 (2016) 046102.

- [35] K. Schmidmayer, S.H. Bryngelson, T. Colonius, An assessment of multicomponent flow models and interface capturing schemes for spherical bubble dynamics, *J. Comput. Phys.* 402 (2020) 109080.
- [36] M. Rodriguez, E. Johnsen, A high-order accurate five-equations compressible multiphase approach for viscoelastic fluids and solids with relaxation and elasticity, *J. Comput. Phys.* 379 (2019) 70–90.
- [37] R. Menikoff, B.J. Plohr, The Riemann problem for fluid flow of real materials, *Rev. Mod. Phys.* 61 (1989) 75.
- [38] C.-W. Shu, *Numerical Methods for Hyperbolic Conservation Laws (AM257)*, Lecture Notes, 2006.
- [39] E.F. Toro, The HLLC Riemann solver, *Shock Waves* 29 (2019) 1065–1082.
- [40] K.W. Thompson, Time-dependent boundary conditions for hyperbolic systems, II, *J. Comput. Phys.* 89 (1990) 439–461.
- [41] M. Frigo, S.G. Johnson, FFTW: An Adaptive Software Architecture for the FFT, *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP'98 (Cat. No. 98CH36181)*, vol. 3, IEEE, 1998, pp. 1381–1384.
- [42] M. Vinokur, On one-dimensional stretching functions for finite-difference calculations, *J. Comput. Phys.* 50 (1983) 215–234.
- [43] W.D. Gropp, E. Lusk, A. Skjellum, E.E. Lusk, *Using MPI: Portable Parallel Programming with the Message-Passing Interface*, vol. 1, MIT Press, 1999.
- [44] R. Thakur, W. Gropp, E. Lusk, Data sieving and collective I/O in ROMIO, in: *Proceedings. Frontiers' 99. Seventh Symposium on the Frontiers of Massively Parallel Computation*, IEEE, 1999, pp. 182–189.
- [45] M. Collette, M. Miller, Silo & HDF5 I/O Scaling Improvements on BG/P Systems, *Technical Report, Lawrence Livermore National Lab. (LLNL)*, Livermore, CA (United States), 2010.
- [46] H. Wang, S. Potluri, D. Bureddy, C. Rosales, D.K. Panda, GPU-aware MPI on RDMA-enabled clusters: design, implementation and evaluation, *IEEE Trans. Parallel Distrib. Syst.* 25 (2013) 2595–2605.
- [47] S. Chandrasekaran, G. Juckeland, *OpenACC for Programmers: Concepts and Strategies*, Addison-Wesley Professional, 2017.
- [48] B. Aradi, Fypp: python-powered fortran metaprogramming, <https://github.com/aradi/fypp>, 2021, *GitHub Repository*.
- [49] V.G. Vergara Larrea, W. Joubert, M.J. Brim, R.D. Budiardja, D. Maxwell, M. Ezell, C. Zimmer, S. Boehm, W. Elwasif, S. Oral, et al., Scaling the summit: deploying the world's fastest supercomputer, in: *High Performance Computing: ISC High Performance 2019 International Workshops, Frankfurt, Germany, June 16–20, 2019*, Springer, 2019, pp. 330–351.
- [50] *Nsight Compute Documentation*, <https://docs.nvidia.com/nsight-compute/NsightCompute/index.html>, 2023, *User Manual*.
- [51] D. Yokoyama, B. Schulze, F. Borges, G. Mc Evoy, The survey on ARM processors for HPC, *J. Supercomput.* 75 (2019) 7003–7036.
- [52] W. Elwasif, S. Bastrakov, S.H. Bryngelson, M. Bussmann, S. Chandrasekaran, F. Giorba, M.A. Clark, A. Debus, W. Godoy, N. Hagerty, J. Hammond, D. Hardy, J.A. Harris, O. Hernandez, B. Joo, S. Keller, P. Kent, H. Le Berre, D. Lebrun-Grandie, E. MacCarthy, V.G.M. Vergara, B. Messer, R. Miller, S. Oral, J.-G. Piccinali, A. Radhakrishnan, O. Simsek, F. Spiga, K. Steiniger, J. Stephan, J.E. Stone, C. Trott, R. Widera, J. Young, Early application experiences on a modern GPU-accelerated Arm-based HPC platform, in: *HPC Asia '23, International Workshop on Arm-Based HPC: Practice and Experience (IWAHPCE)*, Singapore, 2023.
- [53] M. Tanguay, *Computation of bubbly cavitating flow in shock wave lithotripsy*, Ph.D. thesis, California Institute of Technology, 2004.
- [54] Y. Liu, P. Zhong, BegoStone-A new stone phantom for shock wave lithotripsy research, *J. Acoust. Soc. Am.* 112 (2002) 1265–1268.
- [55] T.A. Zwaschka, J.S. Ahn, B.W. Cunitz, M.R. Bailey, B. Dunmire, M.D. Sorensen, J.D. Harper, A.D. Maxwell, Combined burst wave lithotripsy and ultrasonic propulsion for improved urinary stone fragmentation, *J. Endourol.* 32 (2018) 344–349.
- [56] J.C. Meng, T. Colonius, Numerical simulation of the aerobreakup of a water droplet, *J. Fluid Mech.* 835 (2018) 1108–1135.