

HPC as a key enabler for high-fidelity simulation of liquid atomization in aerospace applications

Xiaoyi Li and Marios C. Soteriou

Thermal Fluid Science Department, United Technologies Research Center,
East Hartford, CT 06108

Email: xiaoyi.li@utrc.utc.com

Abstract. High-fidelity simulations of the atomization of liquid fuel jets play a critical role in optimizing the combustor performance in aerospace engines. The complex processes of liquid atomization in industrially relevant configurations occur in spatial and temporal scales that span several orders of magnitude. Direct numerical simulations that are able to capture details in all scales rely on sufficiently high grid resolution and require the use of high-performance computing (HPC) to obtain results within reasonable period of execution time. The aim of the current work is to improve the scalability of a continuum multiphase flow solver and explore the feasibility of performing a high-resolution validation test by using DOE-sponsored HPC facilities. Inefficient parallel communication patterns that become bottlenecks of the solver in massively parallel computations are identified and removed. As a result, the weak scaling of the solver is improved from below hundreds to above thousands of processors. Several numerical techniques that effectively reduce the grid count of the simulation are examined with respect to their impact on load balancing and scaling to large problem size. Simulation of a high-resolution validation case of liquid atomization using thousands of processors is under way at the Oak Ridge Leadership Computing Facility on Jaguar.

1. Introduction

Atomization of fuel jets to generate micron-sized droplets is critical to the performance of combustors encountered in industrial and military aerospace applications, such as gas turbines, augmentors, scramjets and ramjets, and rockets. Increased fuel area-to-volume ratio due to atomization significantly enhances the fuel evaporation rate, which contributes to better fuel-air mixing and subsequent combustion of fuels. While the magnitude of fuel-air ratio has an impact on engine efficiency and emissions, the spatiotemporal distribution of fuel vapor determines combustion dynamics that critically affects engine stability when coupled with acoustics.

From the breakup of centimeter-sized liquid jet column to pinch-off of micron-sized ligaments to form droplets, liquid atomization manifests itself as a complex multiphysics multiscale process. Different dynamic forces due to gas flow, liquid inertia, and surface tension compete with each other, controlling the various

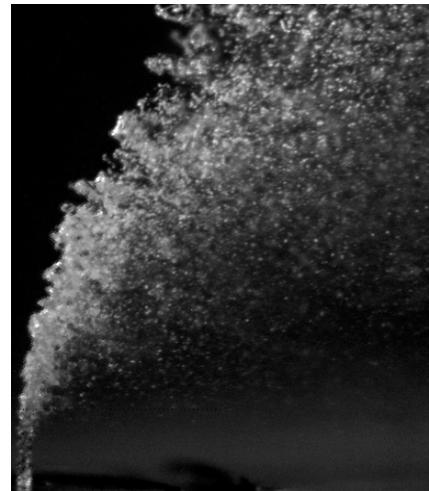


Figure 1. Experimental snapshot of liquid atomization in cross-flowing gas.

instabilities in multiphase flow. Near liquid-gas interface, either Rayleigh-Taylor instability [1,2] due to density difference or Kelvin-Helmholtz instability [3,4] due to aerodynamic shear, or the combination of both, drives the large-scale growth of surface waves and liquid breakup. On the other hand, capillary Plateau-Rayleigh instability [5] due to surface tension plays a dominant role in the pinch-off of highly stretched thin ligament when its size drops below certain threshold. In the experimental snapshot of liquid jet in cross flow in Fig. 1, various multiphase instabilities occur simultaneously at different spatial locations of liquid-gas interface. As a result, complicated multiscale liquid structure forms, making it difficult to quantitatively understand the process. Moreover, since liquid properties (such as density, viscosity, and surface tension) depend on operating conditions (such as temperature and pressure), the liquid in practical conditions can break up and atomize, following a drastically different path from the liquid in ambient conditions. The extremely hostile high-temperature, high-pressure environment in such applications, however, makes it prohibitively challenging to conduct experimental study.

Despite the complexity of liquid atomization process, it is known that the multiphase flow is well described by Navier-Stokes equations with source term due to surface tension acting at the interface. A direct simulation approach based on such first principles has been developed at UTRC with the aim to provide trustworthy data at industrially relevant operating conditions. The approach is based on the coupled level-set and volume-of-fluid (CLSVOF) method developed by Sussman and co-workers [6] to track/capture liquid-gas interface and an Eulerian-Lagrangian droplet transformation approach developed by Herrmann [7] to capture small-scale physics using point particles and alleviate the needs of resolving flow inside droplets. The solver is also enhanced by a robust multigrid preconditioned conjugate gradient (MGPCG) approach for pressure, an adaptive mesh refinement (AMR) technique for liquid-gas interface to reduce grid count [8, 9]. Note that other methods of similar flavor such as a refined level set grid method [10] and a MARS method [11] were also applied to simulate liquid breakup and atomization process, although no method has been truly validated by experimental data.

Validation and execution of all such direct simulations has encountered a huge computational challenge, i.e., the difficulty of resolving within a single simulation a broad range of spatial and temporal scales that emerge in the flow. For instance, a liquid ligament can measure only a few microns in diameter just before pinch-off, whereas the scale of injectors or combustors can be 10^4 orders of magnitude larger. On a fixed (Eulerian) grid, the pinch-off of thin ligaments and formation of droplets require a non-negligible grid density for processes to be correctly resolved. Therefore sufficiently high grid resolution is required to capture the smallest-scale liquid structures. With the help of the technical expertise and computational resources from DOE's National Energy Research Scientific Computing center (NERSC) and the Oak Ridge Leadership Computing Facility (OLCF), the scalability of the CLSVOF solver at UTRC has been significantly improved and the impact of each element of the solver on its parallel performance has been better understood.

In this paper, we first briefly describe the CLSVOF numerical algorithms and various cost-reduction elements in Section 2. A coarse grid example of liquid jet atomization in cross-flowing gas is simulated using 64 processors. The physical results are demonstrated in Section 3 and the needs of higher resolution simulation using HPC are revealed. In Section 4, the removal of some communication bottlenecks in the solver is described and the improvement of the scalability of uniform grid calculation to thousands of processors is

demonstrated. The impact of various grid-reduction techniques on parallel performance is also discussed. Conclusions of current work are provided in Section 5.

2. Numerical algorithms

The flow solver is built upon the CLSVOF approach to track the evolution of liquid-gas interface. The detailed description of CLSVOF can be found elsewhere [6], together with several validation studies. Briefly, the Navier-Stokes equations for incompressible flow of two immiscible fluids (such as liquid and gas) are written in terms of a smooth level set function ϕ , whose zero level represents the time-evolving interface. In addition to the evolution equation for ϕ , the transport equation for the cell liquid volume fraction (the volume-of-fluid function, F) is solved. The normals used in the VOF reconstruction step are determined from the level set function. The volume fractions are then used with the normals to construct a volume preserving distance function ϕ . In this way, volume is preserved by implementing a “local” mass fix at every iteration. Second-order accurate curvature is calculated from F by the method of height function. The method takes advantages of both the accurate geometric interface representation in level set methods and the volume-preserving properties in volume of fluid methods. For the multiphase flow, a single-fluid approach is adopted, that is, properties of density and dynamic viscosity are function of ϕ . Velocity extrapolation based on ϕ from the liquid phase is used to approach the solution of the corresponding one-phase method in the limit of uniform vapor pressure at large liquid-to-gas density ratios. An accurate and stable variable density pressure projection solver based on MGPCG method is used to solve the sparse matrix system that results from discretizing an elliptic equation with discontinuous coefficients and source terms.

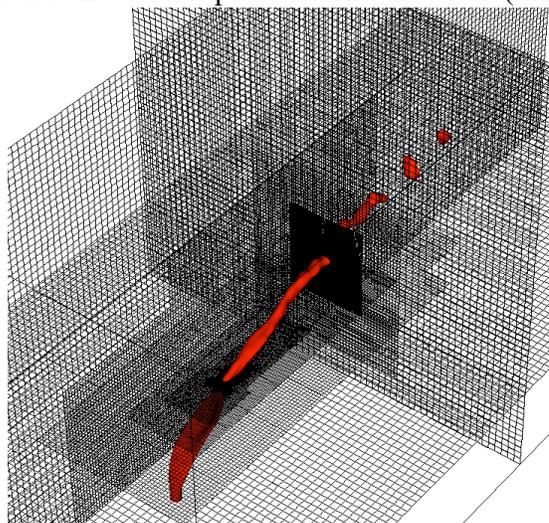


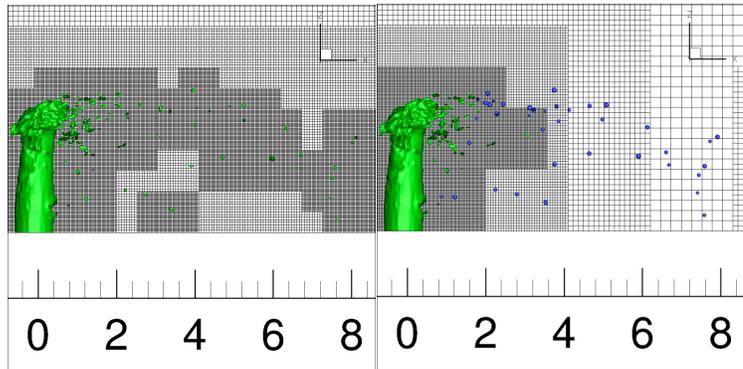
Figure 2. Example of the multilevel grid used in block-structured adaptive mesh refinement.

In order to reduce grid-count in multiphase simulations, block-structured adaptive mesh refinement (AMR) techniques have been used as the underlying framework of the solver. When AMR is active, cells that are crossed by the liquid-gas interface are tagged for refinement. Starting from the base level, boxes (with a minimum size of, say, 32^3 cells) are combined to cover all the tagged cells within assigned coverage efficiency (Fig. 2). This set of blocks with the same grid spacing forms level 1. This level is in turn tagged for refinement at the interface, and the process is repeated until the required grid resolution is achieved. During the simulation, the data on the fine level are either copied from a previous time step or, when the grid structure has changed locally, conservatively interpolated from the underlying coarse level. The interface, however, is always embedded in the finest grid level to avoid gross interpolation errors. In a time step, the calculation is carried out on all levels, and the updated data on a fine level are averaged to the underlying coarser one. Assuming the optimal coverage closely follow the interface, it is reasonable to estimate that doubling the

grid density in three dimensions corresponds to increasing the storage and the computational time by a 2^3 factor instead of a 2^4 factor (three dimensions plus time) for a grid without AMR. This 50% saving in resources is repeated at every new refinement level.

Even though AMR restricts dense grid to a compact zone across liquid-gas interface, a grand numerical challenge still exists because of the presence of a mist of atomized droplets occupying a large portion of the domain. AMR treatment essentially leads to uniformly refined grid in such regions. To further reduce grid-count, one should realize that the small-scale physics of drag and evaporation of individual droplet (neglecting droplet internal flow) can be reasonably well resolved with established models. Therefore a numerical transformation has been developed to change the way the liquid phase is described from Eulerian level-set/volume-of-fluid

representation to Lagrangian particle representation in dilute regions where droplets are well equilibrated with the gas flow and unlikely to collide with each other. Essentially, the droplets are transformed into Lagrangian particles tracked by models such as standard drag law. The



(a)

(b)

Figure 3. Snapshot of LJIC simulation (a) without Eulerian-Lagrangian transformation and (b) with the transformation.

transformation approach, when combining with block-structured AMR technique, significantly reduces the cost of simulation.

When droplets are removed from the Eulerian description, in the region interested by the transformation, the hierarchy of refinement levels quickly reverts to the underlying base level. The cost of tracking Lagrangian droplets is negligible compared with the flow solver cost. Thus, the overall simulation becomes faster because of a relatively coarse grid far from injection.

The cost benefit of the Eulerian-Lagrangian transformation is demonstrated in Figures 3 and 4. Two startup simulations of liquid jet in crossflow (LJIC) without 2(a) and with 2(b) transformation are performed. With transformation activated, the refined grid for droplets is reverted to the base grid, and the grid count is significantly reduced (Fig. 3).

In this transient stage, the cost of the simulation without transformation continues to increase due to newly generated droplets while the cost of simulation with transformation decreases to a constant value. Without transforming the droplets into particles, AMR boxes would eventually cover a substantial portion of the domain, canceling the computational advantage of adaptive mesh refinement.

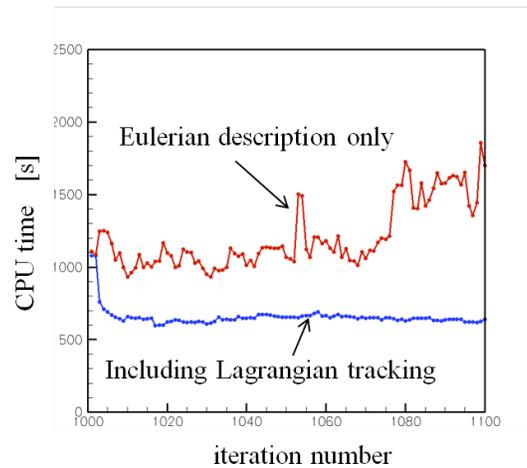


Figure 4. Cost per step as a function of time step for liquid jet in crossflow simulation with/without droplet transformation.

The parallelization of the CLSVOF solver is implemented by spreading the computations on different AMR boxes onto different processors. The AMR boxes having the same grid resolution form a union to cover each AMR level, and the union of base-level boxes covers the whole computational domain. In the presence of AMR, the communications between different processors can be complex, including neighbor box communications within one level and overlapping box communications across different levels. The impact of such communication patterns on the solver’s parallel performance is discussed in Section 4.

3. Coarse grid demonstration

The example chosen for demonstrating the approach comes from a previous validation against the liquid jet in crossflow experiment carried out by Leong and Hautman [12], where downstream droplet size, velocity, and volumetric flux data were acquired using a Phase Doppler Particle Analyzer (PDPA). The measured properties of the test liquid (Jet-A fuel) were density $\rho_l = 780 \text{ kg/m}^3$, surface tension $\sigma = 0.024 \text{ N/m}$, and dynamic viscosity $\mu_l = 0.0013 \text{ kg/m}\cdot\text{s}$. Air was at standard conditions, with inlet velocity of 69 m/s. The orifice size was $d_0 = 0.762 \text{ mm}$, with nominal mass flow rate $\dot{m} = 15.3 \text{ kg/h}$. Droplet sampling was located at 33 orifice diameters downstream of the injector.

In the simulation shown in Fig. 5, the coordinate system has the x -axis in the crossflow direction and the z -axis in the spanwise direction. The origin is located at the center of the jet orifice. The computational domain is a box of 5.12 cm in the x -direction (1.28 cm upstream of the orifice) and of 2.56 cm in the y and z directions. The base resolution of $256 \times 128 \times 128$ nodes and 3 levels of refinement bring the minimum grid size to $\Delta x = 25 \mu\text{m}$.

The no-slip boundary condition is imposed at the $z = 0$ plane, except at the jet orifice. Inlet boundary is imposed at the $x = -1.28 \text{ mm}$ plane, and out-flow boundary conditions are imposed on the remaining boundary planes. A plug flow profile is used for both the liquid and gas inlet. While trivial, this setting can be thought of as a preliminary study where the effect of inlet turbulence is excluded, leaving only to fluid mechanic instabilities the task of atomizing the jet.

The simulation starts from the instant when the jet is injected into a well-developed crossflow gas ($t = 0$). The liquid jet penetrates into the gas flow; and once the jet has reached full penetration, the computational methodology described in the previous section yields a CPU

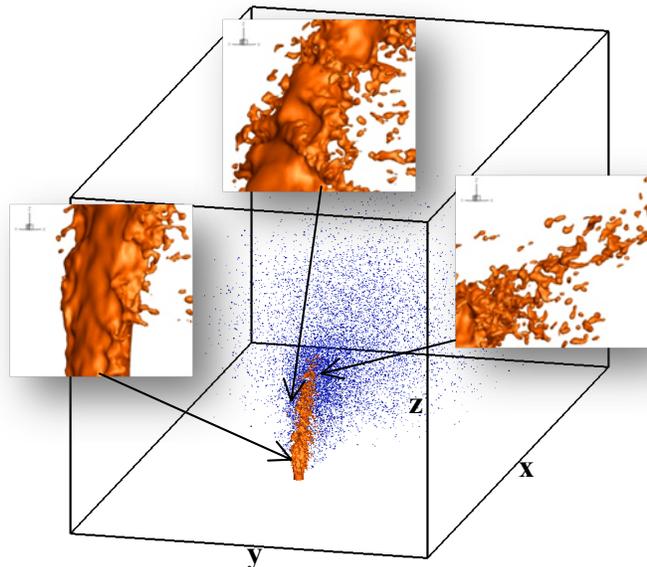


Figure 5. LJC simulation at $t = 1.686 \text{ ms}$. The jet column is represented by 0-isosurface (orange color) of level set on the Eulerian grid. The spray (Lagrangian droplets) is shown as a scatter of blue spheres whose diameter is properly scaled. Different segments of the iso-surface, from near injection to breakup, are also displayed in the three inserts.

time of ~ 50 seconds per time step ($\Delta t = 0.17 \mu\text{s}$) on eight 8-core, 32 GB, 2.66 MHz nodes with InfiniBand switch.

The primary break-up of a liquid jet in crossflow proceeds through the fragmentation of the liquid column into ligaments/large drops and through the stripping of small droplets from the column surface, as shown in the inserts of Fig. 5. Different modes of surface waves due to turbulence, capillary, and aerodynamic forces play a critical role in the formation of ligaments and drops. Reviews of jet breakup studies can be found in [13].

Figure 6 demonstrates the effects of grid refinement on the formation of droplets and their subsequent transport. In frame (a), a coarse base grid of $128 \times 64 \times 64$ with three levels of refinement is used, compared with a fine base grid of $256 \times 128 \times 128$ with the same number of levels in frame (b). Although the large-scale jet penetration is

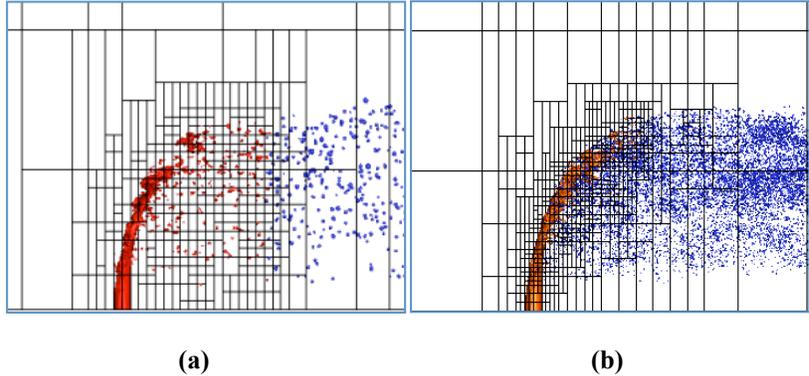


Figure 6. LJIC simulation at different resolutions: (a) coarse grid and (b) fine grid.

little affected by grid resolution, the small-scale breakup details are quite different under different resolutions. In particular, the size of resulting droplets decreases with decreasing grid size. Since the volume loading is the same, the fine-grid simulation generates a larger number of droplets. The result suggests that simulation of the breakup processes has not yet reached grid convergence.

A quantitative comparison is displayed in Fig. 7, where the droplet size distribution from PDPA measurements at a plane 33 orifice diameters downstream of injection is compared with the frequency distribution from the calculation using the $256 \times 128 \times 128$ base grid. A stationary spray distribution is assumed to occur at the sampling plane after $t = 1.686$ ms (the time of the snapshot in Fig. 5). Particles are sampled until $t = 1.884$ ms, for a total of 3470 droplets. Both distributions in Fig. 7 are normalized to unitary area. The data are concentrated between 10 and 50 μm , whereas the calculated droplets spread in a range between 30 and 120 μm . This is not completely unexpected, since the minimum Δx of the simulation is only 25 μm . One, or possibly two, additional levels of refinement should be used to reach the smallest scale in the measurements. The current results underscore the challenge posed to the modeling community by liquid atomization,

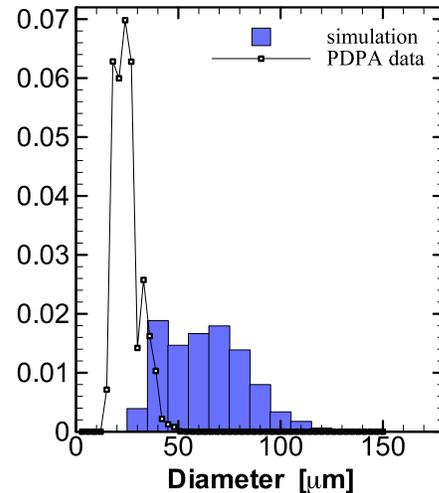


Figure 7. Droplet size distribution from PDPA (symbols) and from simulation.

particularly in the high-shear regime. The results clearly indicate that current grid resolution is not sufficient to completely reproduce the physical problem and HPC has to be exploited to perform higher-resolution, higher-cost simulations.

4. Scalability improvements and performance results

Although the CLSVOF solver provides accurate solution for some cases at low resolution, the parallel performance, specifically the scalability, of the solver needs to be verified and improved before a large simulation can be executed.

A preliminary scaling test was first performed on NERSC Hopper, and the behavior of the solver is depicted in Fig. 8. The test case is the startup LJIC simulation using two uniform grids of $256 \times 128 \times 128$ (total grid count is 4.2 million) and $512 \times 256 \times 256$ (total grid count is 33.6 million), respectively. The strong scaling (red dashdot line) for the 4.2M problem is reasonably good up to ~ 100 processors. The cost per step decreases almost linearly with increasing number of processors (denoted by N below).

The slope in log-log scale is very close to the ideal strong scaling slope, especially when N is small. The sharp leveling off of the strong scaling curves for the 4.2M problem beyond $N=128$ is due to insufficient number of boxes (box of 32^3 with grid count 33K) to be allocated evenly to all processors. Some processors are essentially idling, and the cost starts to increase with further increase of processors because of the additional communication among more processors. This is a typical situation of insufficient computation on some processors and load imbalance. Despite the fact that the solver behaves reasonably well for the 4.2M problem, serious deviation from this behavior appears when the problem size increases by 8 times to 33.6M. The strong scaling levels off even when N is still small, and the slope significantly deviates from the ideal scaling slope. As shown by the weak scaling curve (blue dashdot line), as the problem size and N increase proportionally, the cost does not remain constant, as would be expected in the ideal case, and significantly increases in the large N range. In many solvers, communication cost will sooner or later build up as the problem size increases. However, a solver with good performance should delay this buildup to very large value of $N=N_{cr}$ and allow scalable computation at intermediate size $N < N_{cr}$. The desired value of N_{cr} is determined by the desired problem size as well as acceptable cost per step.

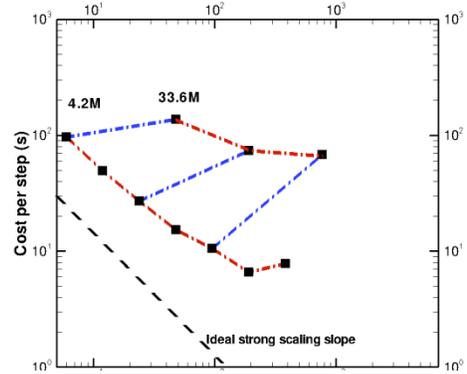


Figure 8. Strong and weak scaling of non-optimized solver.

A careful analysis sheds some light on the scaling behavior shown in Fig. 8. Since the scaling is reasonable when N is small, the increased communication overhead must be related to the increased problem size. For the uniform-grid CLSVOF solver, as the problem size increases, the domain is essentially decomposed into more boxes of the same size (say, 32^3). If the amount of communication relies on number of boxes, and the number of boxes increases with problem size, the weak scaling will be hurt significantly. With this line of reasoning, a number of MPI operations that depend on the total number of boxes have been identified with the help of NERSC staff. Such MPI operations when called too frequently freeze the code at the communication stage. We also found these operations can be reduced by grouping individual communication calls. The modification has improved the code scalability, as demonstrated in Fig. 9. We were able to verify the scalability (strong and weak) of the solver up to $\sim 10^3$ cores on Hopper. The strong scaling is close to ideal for both the 4.2M problem and the 33.6M problem. For the 33.6M problem, the optimized solver gives much better scaling than does the nonoptimized solver. The scaling is also reasonably good for the 268.8M problem (with uniform grid of $1024 \times 512 \times 512$) when $N < 2000$. The weak scaling, represented by the blue solid curves, indicates good parallel behavior of the code up to $\sim 10^3$ processors and up to $\sim 250M$ problem size with an acceptable cost of $\sim 50s$ per step.

Although some numerical techniques described in Section 2 can effectively reduce the grid count of the simulation, they may not be designed to have the optimum parallel efficiency. In the presence of AMR, for example, boxes on each refinement level are distributed to designated processors based on which subdomain they belong to. Since some subdomains may not have a refined grid, some processors may not have refined boxes allocated to them, and they often have to idle until other processors finish their refined-level computations. The situation is even worse when the Eulerian-Lagrangian transformation is introduced. Fine-level boxes used to resolve small droplets are discarded, and only coarse-level boxes are present in significant portion of the domain. In such cases, the computational load is highly imbalanced, and the solver will not scale with increasing number of refinement levels. A plausible alternative is to let each refinement level work independently to distribute their boxes evenly among processors. However, the cross-level data transfer (interpolation and averaging operations between overlapping boxes) will become too expensive when performed through network between different processors.

The impact of AMR and Eulerian-Lagrangian transformation on the solver scalability is demonstrated in Fig. 10. Since the AMR refinement and grid count keep changing as time proceeds in the LJIC problem, we select the instant of $N_{\text{step}}=1500$ as the sample point in a $256 \times 128 \times 128$ grid with three levels of refinement case. The grid

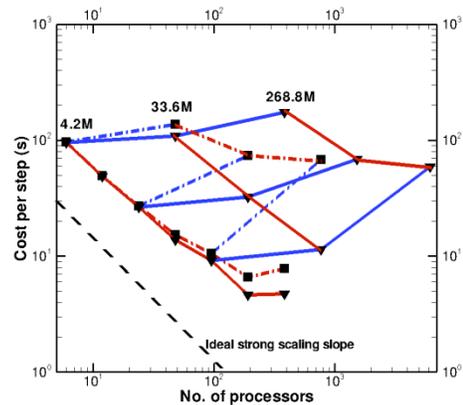


Figure 9. Strong and weak scaling of non-optimized solver (dashdot) and optimized solver (solid).

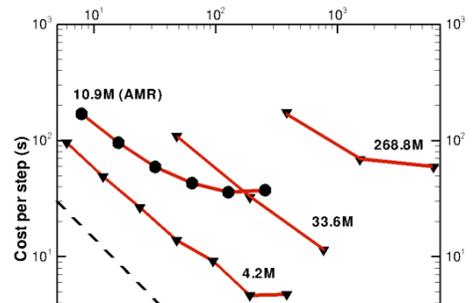


Figure 10. Strong scaling of the optimized solver using uniform grid (delta) and AMR (delta).

count at that point is 10.9M. In Fig. 10, the scaling of AMR case deviates from ideal scaling much more significantly compared with the scaling without AMR. Poor weak scaling is not unexpected for the case using AMR. On the other hand, AMR does reduce the grid count significantly and therefore reduces the actual processing time on each node. From parallel efficiency point of view, then, in order to achieve good balance between node processing and network communication, there is an optimum choice of grid-based resolution and refinement levels for certain desired minimum grid resolution. The choice also depends on the physical problem in terms of how much space needs to be resolved by fine grid.

5. Conclusions

One important contribution of the present work is to show the critical demand of using HPC to perform high-resolution, high-cost liquid atomization simulations that have huge impact on combustor design in aerospace applications. With the aid of NERSC and OLCF, UTRC has made significant progress toward developing a highly scalable, high-fidelity, multiphase solver to be used for atomization simulations. It has been shown that the solver can scale up to $\sim 10^3$ processors and up to $\sim 250M$ problem size with an acceptable cost of ~ 50 s per step. The study also suggests a careful selection of number of refinement levels when AMR is used in order to achieve good balance between node processing and network communication. With this developed knowledge and capability, a high-resolution validation simulation of liquid jet in crossflow is under way at OLCF on Jaguar.

Acknowledgments

The authors acknowledge the funding support from United Technologies Corp. They also acknowledge the computer resource and staff support from National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, and from the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory, which is supported by the Office of Science of the Department of Energy under Contract DE-AC05-00OR2272. In particular, the expertise and time commitment of Dr. Viraj Paropkari and Dr. David Skinner from the SciDAC Outreach Center are sincerely acknowledged. The technical help from Professor Mark Sussman at Florida State University is gratefully acknowledged.

References

1. Rayleigh, "Investigation of the character of the equilibrium of an incompressible heavy fluid of variable density," in *Proceedings of the London Mathematical Society*, 14, 170–177, 1883.
2. Taylor, G. I., "The instability of liquid surfaces when accelerated in a direction perpendicular to their planes," in *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 201 (1065), 192–196, 1950.
3. Kelvin, "Hydrokinetic solutions and observations," *Philosophical Magazine*, 42, 362–377, 1871

4. Helmholtz, H., "On the discontinuous movements of fluids," *Monthly Reports of the Royal Prussian Academy of Philosophy in Berlin*, 23, 215–228, 1868.
5. Rayleigh, "On the instability of jets," *Proc. London Math. Soc.*, s1-10(1), 4-13, 1878
6. Sussman, M., Smith, K. M., Hussaini, M.Y., Ohta, M., and Zhi-Wei, R., "A sharp interface method for incompressible two-phase flows," *Journal of Computational Physics*, 221 469-505, 2007.
7. Herrmann, M., "A parallel Eulerian interface tracking/Lagrangian point particle multi-scale coupling procedure," *Journal of Computational Physics*, (229), 745-759, 2010.
8. Li, X., Arienti, M., Soteriou, M., "Towards an efficient, high fidelity methodology for liquid jet atomization computations," paper AIAA-210, *48th AIAA Aerospace Sciences Meeting*, Orlando, FL, 2010.
9. Li, X., Soteriou, M., Arienti, M., and Sussman, M., "High-fidelity simulation of atomization and evaporation in a liquid jet in cross-flow," AIAA-2011-0099, *49th AIAA Aerospace Sciences Meeting*, Orlando, FL, 2011.
10. Herrmann, M., "A balanced force refined level set grid method for two-phase flows on unstructured flow solver grids," *Journal of Computational Physics*, 227(4) 2674-2706, 2008.
11. Inoue, C., Watanabe, T., and Himeno, T., "Liquid sheet dynamics and primary breakup characteristics at impingement type injector," *45th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, Denver, CO, AIAA 2009-5041, 2009.
12. Leong, M. Y., Hautman, D. J., ILASS Americas, *15th Annual Conference on Liquid Atomization and Spray Systems*, Madison, WI, May 2002.
13. Wu, P. K., Kirkendall, K. A., Fuller, R. P., and Nejad, A. S., "Break-up processes of liquid jets in subsonic crossflows," *Journal of Propulsion and Power*, 13 64-73, 1997.