Hybrid Message-Passing Interface-Open Multiprocessing Accelerated Euler–Lagrange Simulations of Microbubble Enhanced HIFU for Tumor Ablation

Jingsen Ma[1]
Dynaflow, Inc.,
10621-J Iron Bridge Road,
Jessup, MD 20724
e-mail: jingsen@dynaflow-inc.com

Xiaolong Deng
Dynaflow, Inc.,
10621-J Iron Bridge Road,
Jessup, MD 20724

Chao-Tsung Hsiao
Dynaflow, Inc.,
10621-J Iron Bridge Road,
Jessup, MD 20724

Georges L. Chahine
Dynaflow, Inc.,
10621-J Iron Bridge Road,
Jessup, MD 20724

Microbubble enhanced high intensity focused ultrasound (HIFU) is of great interest to tissue ablation for solid tumor treatments such as in liver and brain cancers, in which contrast agents/microbubbles are injected into the targeted region to promote heating and reduce prefocal tissue damage. A compressible Euler–Lagrange coupled model has been developed to accurately characterize the acoustic and thermal fields during this process. This employs a compressible Navier–Stokes solver for the ultrasound acoustic field and a discrete singularities model for bubble dynamics. To address the demanding computational cost relevant to practical medical applications, a multilevel hybrid message-passing interface (MPI)-open multiprocessing (OpenMP) parallelization scheme is developed to take advantage of both scalability of MPI and load balancing of OpenMP. At the first level, the Eulerian computational domain is divided into multiple subdomains and the bubbles are subdivided into groups based on which subdomain they fall into. At the next level, in each subdomain containing bubbles, multiple OpenMP threads are activated to speed up the computations of the bubble dynamics. For improved throughput, the OpenMP threads are more heavily distributed to subdomains where the bubbles are clustered. By doing this, MPI load imbalance issue due to uneven bubble distribution is mitigated by OpenMP speedup locally for those subdomains hosting more bubbles than others. The hybrid MPI-OpenMP Euler–Lagrange solver is used to conduct simulations and physical studies of bubble-enhanced HIFU problems containing a large number of microbubbles. The phenomenon of acoustic shadowing caused by the bubble cloud is then analyzed and discussed. Efficiency tests on two different machines with 48 processors are conducted and indicate 2–3 times speedup with the same hardware by introducing an OpenMP parallelization in combination with the MPI parallelization.
[DOI: 10.1115/1.4057050]

Introduction

High intensity focused ultrasound (HIFU) [1,2] uses focused ultrasound waves to treat some diseases with the intense heat from its thermal effects. It is noninvasive and has been applied in many therapeutic and surgical medical interventions such as for tissue ablation in cancer treatment, particularly, of deep-seated tumors such as in the brain, liver, and kidney.

In recent years, the inclusion of bubbles in the form of contrast agents to HIFU has been investigated as means to enhance heating. The increase in temperature levels has been observed in some in vitro [3,4] and in vivo experiments [5,6]. However, the way microbubbles contribute to the heat increase is not well understood. Recently, we use multiphase numerical modeling to simulate this process and study the effects of microbubbles in the HIFU treatments but encounter computational efficiency challenge when a problem involves a large number of microbubbles, which is the main issue to be addressed through a hybrid open multiprocessing (OpenMP)-message-passing interface (MPI) parallelization strategy in this work.

In microbubbles enhanced HIFU with, the propagation and focusing of ultrasonic waves are influenced by the bubbles themselves. Within the low void fraction regime, an increase in void fraction usually results in stronger enhancement effects. However, the physical interaction gets more complicated as the bubble number and local void fraction become high enough for the ultrasonic waves to be strongly absorbed and/or reflected. The underlying driver is a phenomenon known as acoustic shadowing, which leads to an increase in the pressure in the pure liquid at the front edge of the cloud and a reduction of the energy in the focal region, thus reducing the effectiveness of the treatment. The displacement of the energy focus may also result in tissue damage ahead of the bubble region. Studying the effects of acoustic shadowing is a critical topic in promoting applications of HIFU with microbubbles.

For this purpose, we have used our Eulerian–Lagrangian approach, which couples the acoustic field with bubble dynamics [7,8], to study the phenomenon. Similarly, as its success in other traditional multiphase flow applications such as in Refs. [9–11], in HIFU study this kind of Eulerian–Lagrangian method has shown intrinsic advantages owing to its capability to include the complex physics of nonlinear interactions between the oscillating bubbles and the ultrasonic field in the computation of heat deposition. However, owing to high computational cost in real applications, efficient parallelization schemes for significant speedup of computation are critical and highly desirable.

The two major parallelization methods adopted on CPU-based high performance computing techniques are MPI and OpenMP.

[1]Corresponding author.

Manuscript received November 18, 2022; final manuscript received February 27, 2023; published online March 28, 2023. Assoc. Editor: Tony Jun Huang.

MPI is an application programmer interface (API) designed for multinodes, distributed-memory parallel computer architectures, while the latter OpenMP, is also an API to explicitly use direct multithreaded shared memory parallelism.

In our recent work in Ref. [12], an MPI shared-memory parallelization scheme was implemented and indicated good scalability is possible. Both the computational domain and the bubbles, based on their locations, are decomposed into several subdomains or subgroups, with each handled by one MPI process. The difficulty encountered in Euler–Lagrange mapping, i.e., to smoothly spread the Lagrangian bubble effects into the fixed Eulerian grids, which could belong to different subdomains, was overcome by utilization of ghost cells surrounding each fluid subdomain with minimal overhead. However, this method becomes inadequate when the bubble number is so large that the MPI processes handling subdomains with bubbles inside them proceed much slower than the others containing fewer or no bubbles. The load imbalance due to localized bubble presence became a bottleneck for the previous MPI-only parallelization method [12]. To address this issue, under the upper level of MPI parallelization, a second level of shared-memory OpenMP parallelization [13] is introduced, where the numbers of OpenMP threads (*nTh*) under each MPI process is varied according to the numbers of microbubbles present in different subdomains. By doing so, this multilevel MPI-OpenMP parallelization strategy takes advantage of both MPI's scalability due to the distributed memory feature, and OpenMP's load balancing capability due to its shared memory nature.

The paper is structured as follows. First, the main framework of the two-way coupled, compressible Euler–Lagrange model is briefly introduced. Then, the multilevel MPI-OpenMP parallelization scheme is described. After an introduction of the simulation problem conditions and numerical setup, efficiency tests and analysis of the parallelization scheme are conducted. Particularly, efficiency comparison between different combinations of number of MPI processes (*nPr*) and number of OpenMP threads (*nTh*) is evaluated and discussed. Using this hybrid parallelization model, simulations involving a large number of bubbles and acoustic shadowing are conducted and the physics of the shadowing is analyzed. This is followed by a developed discussion of the major findings.

Eulerian–Lagrangian Model

The three-dimensional Eulerian–Lagrangian compressible model employed in this study has been documented and validated in our previous studies [7,8,12]. The primary goal of this paper is to accelerate, through a multilevel MPI-OpenMP parallelization, the applications of this model to the HIFU problem where a large number of bubbles is involved. As no additional model development was performed, below we only give an abbreviated summary of the main features of the model to avoid unnecessary repeats. Readers may refer to the references provided above for model details.

Eulerian Compressible Mixture Flow Solver. The governing equations for unsteady compressible flow considering acoustic wave propagation through a two-phase medium (specifically, in this work, tissue/phantom medium and the gas inside the bubbles) are as follows:

$$\begin{aligned}\frac{\partial \rho_m}{\partial t} + \nabla\cdot(\rho_m \mathbf{u}) &= 0\\ \frac{\partial \rho_m \mathbf{u}}{\partial t} + \nabla\cdot(\rho_m \mathbf{u}\mathbf{u} + \boldsymbol{p}_m \mathbf{I} + \boldsymbol{\sigma}) &= 0\\ \frac{\partial \rho_m E_m}{\partial t} + \nabla\cdot((\rho_m E_m + p_m)\mathbf{u}) + \boldsymbol{\sigma}\mathbf{u} + q) &= 0\end{aligned} \tag{1}$$

where t is time, ρ_m is the mixture density, E_m its total energy, $\mathbf{u}$ its velocity, p_m its pressure, σ is the shear stress tensor, $\boldsymbol{I}$ is the unit tensor, and q is a heat source term. The mixture density can be related to the volume fraction α_i of each component by

$$\rho_m = \sum_i \rho_i \alpha_i, \quad \text{where} \quad \sum_i \alpha_i = 1 \tag{2}$$

The total energy is given by $E_m = e_m + 0.5\mathbf{u}^2$ where e_m is the mixture of internal energy with $\rho_m e_m = \sum_i \rho_i e_i \alpha_i$. Also, p_m and e_m are related to ρ_m with a mixture of stiffened equation of state [14] having the following form:

$$\begin{aligned}&p_m = (\Gamma - 1)\rho_m e_m - \Gamma\Pi\\ &\Gamma = 1 + \frac{1}{\sum_i \frac{\alpha_i}{\gamma_i - 1}}, \quad \Pi = \frac{\Gamma - 1}{\Gamma}\sum_i \alpha_i \frac{\gamma_i \pi_i}{\gamma_i - 1}\end{aligned} \tag{3}$$

Here γ_i is the ratio of specific heats and π_i is a constant for each material in the stiffened equation of state. The sound speed is calculated as

$$c_m = \sqrt{\frac{1}{\rho_m}\sum_i \frac{\alpha_i}{\gamma_i(p_m + \pi_i)}} \tag{4}$$

The numerical scheme uses a fully conservative finite volume method utilizing a higher order scheme, MUSCL-based method along with an approximate Riemann solver [15,16].

Lagrangian Modeling of Dispersed Phase. A Lagrangian bubble tracking method is adopted for the bubbles dispersed in the liquid. The variations of the bubble radius, R, and its time differentiation (denoted by dots) are obtained by the following modified Rayleigh–Plesset–Keller–Herring equation [17]:

$$\begin{aligned}&\left(1 - \frac{\dot{R}}{c_m}\right)R\ddot{R} + \frac{3}{2}\left(1 - \frac{\dot{R}}{3c_m}\right)\dot{R}^2 = \frac{(\mathbf{u} - \mathbf{u}_b)^2}{4}\\ &+ \frac{1}{\rho_m}\left(1 + \frac{\dot{R}}{c_m} + \frac{R}{c_m}\frac{d}{dt}\right)\begin{bmatrix} p_v + p_g - p_m - \dfrac{2\sigma}{R} - \\ 4\mu_m \dfrac{\dot{R}}{R} - \dfrac{4}{3}G\left[1 - \left(\dfrac{R_0}{R}\right)^3\right]\end{bmatrix}\end{aligned} \tag{5}$$

where μ_m is the mixture dynamic viscosity, c_m is the local mixture sound speed, $\mathbf{u}_b$ is the bubble translation velocity, p_g is the gas pressure inside the bubble, and p_v is the liquid vapor pressure. σ is the surface tension, and G is the shear elasticity of the medium surrounding the bubble.

The bubble trajectory follows the motion equation of Johnson and Hsieh [18]:

$$\begin{aligned}\frac{d\mathbf{u}_b}{dt} = &-\frac{3}{\rho_m}\nabla p + \frac{3}{4}\frac{C_D}{R}(\mathbf{u} - \mathbf{u_b})|\mathbf{u} - \mathbf{u_b}|\\ &+ \frac{3}{2\pi}\frac{C_L}{R}\sqrt{\frac{\mu_m}{\rho}}\frac{(\mathbf{u} - \mathbf{u_b}) \times \boldsymbol{\omega}}{\sqrt{|\boldsymbol{\omega}|}} + \frac{3}{R}(\mathbf{u} - \mathbf{u_b})\dot{R}\end{aligned} \tag{6}$$

where the first-term on the right-hand side is pressure gradient force, the second and the third are drag and lift forces with C_D and C_L as the drag and lift coefficient, respectively, and $\boldsymbol{\omega}$ is the local vorticity. The last term accounts for bubble acceleration due to its size variation.

Eulerian–Lagrangian Two-Way Coupling. The two-way coupling between the Eulerian continuum-based model and the Lagrangian discrete bubble model is realized as follows [19–21]:

- The mixture flow field has an evolving mixture density, determined by bubble presence and varies in space and time, which satisfies mass, momentum, and energy conservation.

- The bubble dynamics and motion of the individual bubbles in the flow field are controlled by the two-phase medium's local properties and gradients.
- The local properties of the mixture are determined by the bubble size and distribution by providing α for Eq. (2).

The key of this coupling procedure is to map to the fixed Eulerian grids the Lagrangian bubbles solutions, specifically, to derive the void fraction α, i.e., bubble volume fraction in each computational cell, from the instantaneous bubble sizes and locations, for which a Gaussian distribution scheme [19–21] is employed to smoothly "spread" each bubble volume effect over all neighboring cells

Modeling of Heat Deposition. Once the multiphase flow field solution is obtained from the above-described approach, heat source terms, $q_{\mathrm{US,AC}}$, due to the ultrasound absorption and acoustic emission can be derived as below:

$$q_{\mathrm{US,AC}} = \mu_b \varepsilon_{kk}^2 + 2\mu\left(\varepsilon_{ij} - \frac{1}{3}\varepsilon_{kk}\delta_{ij}\right)^2 \tag{7}$$

where μ and μ_b are the shear and bulk viscosities of the fluid and ε is the strain rate tensor [8].

Very importantly, the heat addition due to the viscous damping of bubbles can be derived by summing up that due to a single bubble [22]

$$q_{\mathrm{vis}}^b = (4\pi R^2)\cdot\left(4\mu\frac{\dot{R}^2}{R}\right) \tag{8}$$

Then q_{VIS} in a given control volume is computed through a procedure very similar to that of void fraction computation [13], which sums up the smoothly spread contributions from each bubble within the "influence range":

$$q_{\mathrm{VIS}} = \sum_{j=1}^{N_i} \frac{q_{\mathrm{vis}}^b V_j^b}{\sum_{k=1}^{N\mathrm{cells}} v_{k,j} V_k^{\mathrm{cell}}} \tag{9}$$

A heat transfer equation is then solved to compute the heat deposition effect on the temperature due to the heat sources $q_{\mathrm{US,AC}}$ and q_{VIS}

$$\rho C_p \frac{\partial T}{\partial \tau} = K\nabla^2 T + q_{\mathrm{US,AC}} + q_{\mathrm{VIS}} \tag{10}$$

where ρ, C_p, and K are the density, specific heat, and thermal conductivity of the medium, T is the temperature and τ is the time.

Parallelization Algorithm and Implementation

Algorithm and Overall Procedure. As illustrated in Fig. 1, under MPI parallelization (level I), the Eulerian computational domain is subdivided into several subdomains and both the Eulerian grid and the Lagrangian Bubbles are assigned to different subdomains based on their locations. During each computation time-step, each involved MPI processor handles one subdomain and all processors operate simultaneously to:

(a) Solve Eqs. (1) and (2) for the fluid computation,
(b) Solve Eqs. (3) and (4) for the bubbles,
(c) Map the Lagrangian bubble solutions into the Eulerian grids of each subdomain.

At every time-step, the interprocessor communications are realized through "ghost" cells surrounding each subdomain (see Fig. 2) which are used to receive the neighboring cells' information from the neighboring subdomains. Note

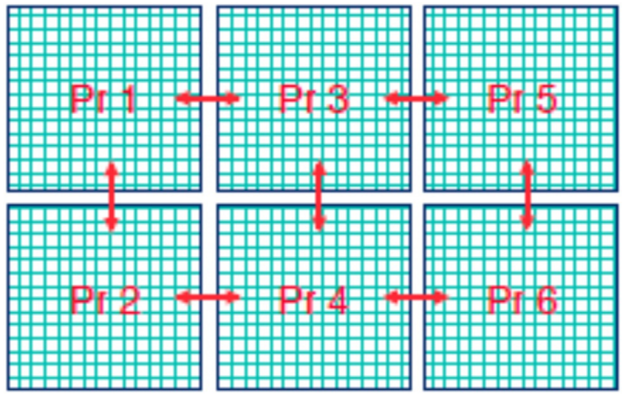

Fig. 1 Illustration of the computation domain decomposition for MPI parallelization (level I) [12]

conventionally ghost cells do not participate in computation and are merely used to create a transfer zone to load data from the "real" cells along the edge of a neighboring subdomain which is the case in step (a). However, to enable mapping in the Eulerian–Lagrangian scheme in step (c) without the need of sending bubble information across neighboring subdomains, a novel utilization of these ghost cells has been explored in our previous work [12]. There the ghost cells first participate in the computations by allowing bubbles near the edges of a subdomain to contribute their effects into ghost cells and then copy their values to the real cells along the edge of its neighboring subdomain. Therefore Lagrangian bubbles are able to smoothly spread, with minimum overhead, their effects into all surrounding fluid cells independent of whether they are within the same subdomain or not [12].

Note that substeps (b) and (c) above are only executed for those MPI processes corresponding to subdomains that include Lagrangian bubbles while being skipped for other subdomains with no bubbles. This may not be an issue when the simulation does not involve too many bubbles such that the computation cost is dominated by substep (a) as in our previous work in Ref. [12].

However, the significant overhead due to load imbalance, i.e., some subdomains host many bubbles while the other may have much fewer or no bubbles may occur when a case to be simulated, such as the one being studied in this paper, involves a large

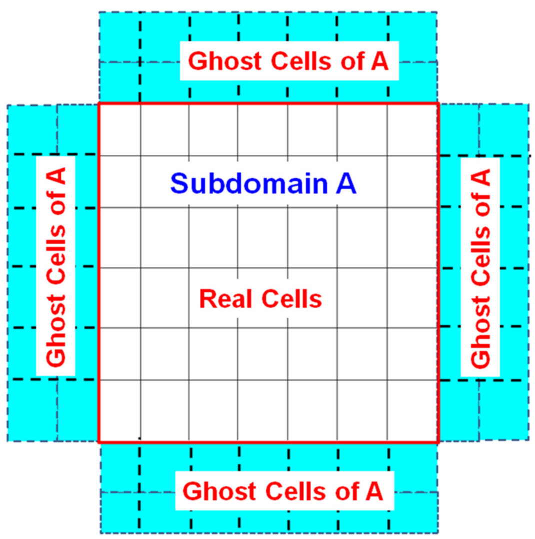

Fig. 2 Definition of N layers of "ghost cells" surrounding a subdomain (here $N = 2$). The red lines indicate the border of subdomain. The white cells inside subdomain are real cells. Those in blue along each border line are ghost cells.

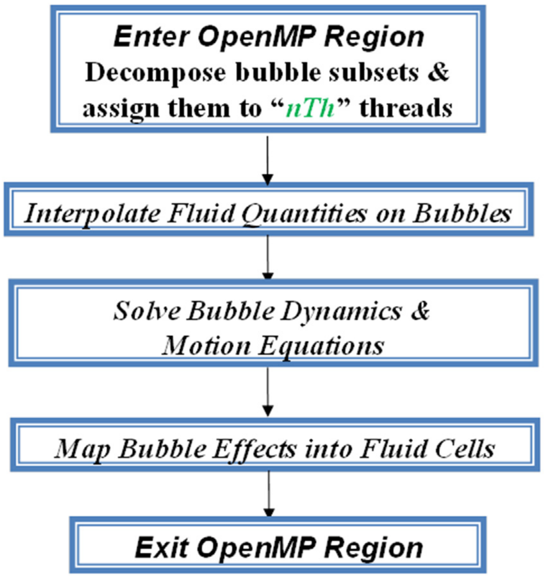

Fig. 3 Illustration of OpenMP parallelization (level II) for bubble computations

number of bubbles. Thus, a second-level parallelization, using OpenMP, to locally speed up bubble computation is employed.

Open Multiprocessing Parallelization (Level II) Under Subdomain. As illustrated in Fig. 3, the bubbles are divided into a number of subsets, which equals the number of threads, nTh, based on their indices with each subset handled independently by one thread [13]. As a result, rather than having the code solve Eqs. (4) and (5) for each bubble one by one in sequence for all the bubbles in a subdomain, this is now shared by nTh threads which work simultaneously. So the computational time for bubbles is reduced to nearly $1/nTh$. Each thread can efficiently access a contiguous global memory space. Therefore, the OpenMP implementation is straightforward. In terms of coding, this is done by enclosing bubble loops with OpenMP directives: "!$omp parallel do" … "!$omp end parallel." Note that when $nTh = 1$, the procedure automatically becomes MPI alone and OpenMP parallelization is turned off. Note this is executed at every time-step during the bubble dynamics computation thus no extra modification to the fluid computation is needed.

Simulation and Results

Problem Setup and Simulation Conditions. In previous work [7,8], we presented our model of microbubble-enhanced HIFU and tested the effects of different parameters, including void fraction, microbubble cloud size and location, , etc., but did not study acoustic shadowing. Here, to test the code with MPI-OpenMP hybrid parallelization, we consider the setup in the experiments of Ref. [3]. A transducer of radius 20 mm and focal length 43 mm is used to insonate a phantom tissue made of polyacrylamide gel (see the sketch in Fig. 4). The sonication frequency, i.e., the frequency of acoustic excitation wave is 2.2 MHz and the pressure amplitude of the generated acoustic waves is 0.1 MPa. For this simulation, the lower boundary condition is designed to change pressure periodically to apply the pressure wave and the energy on the boundary is then updated accordingly. Microbubbles are located inside a polyacrylamide gel cylindrical volume of radius 5 mm and height 10 mm, around the geometric focus of the transducer. The bottom, or the front, of the bubble region is the flat face of the cylindrical volume closest to the transducer. In the present simulations, all the bubbles have a 1.3 μm uniform initial diameter. The density of the gel is 1060 kg/m^3, its viscosity is 0.01 Pa·s, its shear elasticity is 0.1 MPa, its speed of sound is 1540 m/s, its specific heat capacity is 5100 $J/kg \cdot K$, and its thermal conductivity is 0.7 W/mK [22]. Acoustic wave propagation calculations are carried out for 0.0352 ms and the heat source terms are averaged over 0.005 ms ($\sim$ 10 cycles). The heat equation is then solved for 60 s using the derived source terms.

In this work, we use an axisymmetric setup for the simulations, with 608 grids in the r direction and 1414 grids in the z direction for the full computational domain of dimensions (0 cm, 2.95 cm) $\times$ (0 cm, 6.0 cm). The whole domain illustrated in Fig. 4 is split into 30 subdomains for MPI parallelization. Bubbles are located in the small confined $(r \times z)$ domain as (0 cm, 0.005 m) $\times$ (0.035 m, 0.045 m). Due to the axisymmetry, it corresponds to a cylindrical region centered at $z = 0.04$ m with the centerline along with the z-axis. For this axisymmetric study, the bubbles are randomly distributed in a wedge region with an angle of 0.01 radian. In this way, to simulate an initial void fraction of 10^{-5}, a total number of 189,524 bubbles is distributed in the small, confined domain. In this study, air bubbles with standard properties are used.

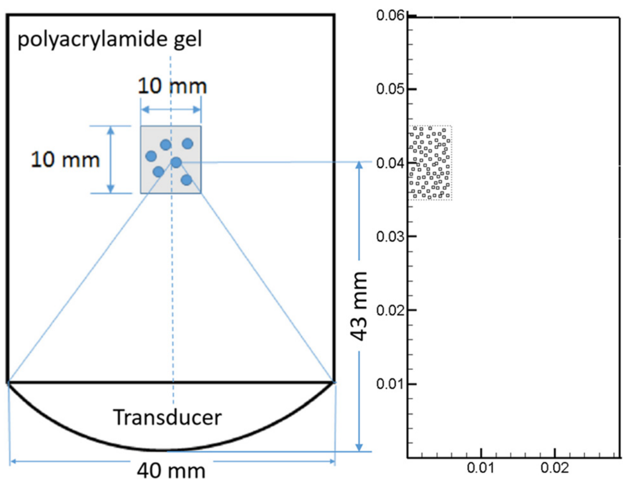

Fig. 4 Illustration of the numerical simulation setup, similar to the experimental setup in Kajiyama et al. [3] (The left drawing is a schematic of the setup. The right drawing shows the actual axisymmetric computational domain with a bubble cloud inside where the coordinate unit is "meter.").

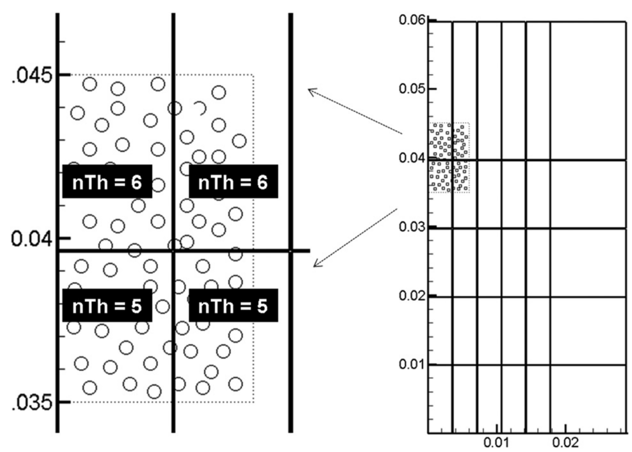

Fig. 5 Illustration of MPI domain decomposition and OpenMP threads distribution. The left is a zoomed view of the subdomains with bubbles inside and OpenMP thread numbers (*nTh*) used by each. The right shows the full computational domain divided into 30 subdomains (*nPr* = 30).

Parallelization Efficiency. To examine the improvements due to the inclusion of second level of OpenMP parallelization, efficiency comparison is first conducted on our in-house, shared memory, multicore multiprocessing workstation, DYNA-OMP, which has four twelve-core AMD™ Opteron 2.80 GHz CPUs with 128 GB total shared memory. To utilize all 48 cores, the full computational domain is decomposed into 30 subdomains (see the sketch in Fig. 5) with each handled by one MPI processor ($n\mathrm{Pr} = 30$). Four of these subdomains contain microbubbles and 5 or 6 OpenMP threads (nTh) are activated in these four subdomains. Hence, all 48 cores are used such that 26 processors handle MPI for the 26 subdomains with no bubbles, 22 processors ($2 \times 6 + 2 \times 5$) handle the OpenMP computations with four of these also handling MPI for the remaining four MPI subdomains. To compare the results, a pure MPI run was also made by decomposing the full domain into 48 subdomains handled by 48 individual MPI processes ($n\mathrm{Pr} = 48$, $nTh = 1$). After verifying in Fig. 6 that both MPI and Hybrid MPI-OpenMP parallelization methods produce the same results physically, the efficiency comparison is conducted in Fig. 7, which indicates the averaged wall time per step is reduced by $\sim 55\%$, from 4.9 s to 2.2 s.

Similar tests were also conducted on DYNAFLOW's other multinode, multiprocessor Linux cluster DYNA-HIVE, which consists of ten nodes, with each equipped with twelve 3G Xeon CPUs and 27 GB memory. Here, the average wall time per step reduced from 3.1 s to 1.1 s, a speedup of $\sim 280\%$. These tests indicate a two to three times of speedup with the same hardware by introducing an OpenMP parallelization in combination with MPI parallelization.

Simulation Results and Analysis. Next, we discuss the results obtained from simulations using the above selected hybrid MPI-OpenMP parallelization threads allocations in Fig. 5. At the arbitrary end of the simulations, at time 0.0352 ms, the field pressure contours are shown in Fig. 8. In the left plot, the acoustic waves travel through the gel field without any bubble interaction and the peak is at the focus, at about $z = 0.04$ m. In the right plot, due to the reflection and absorption of the acoustic wave by the bubble cloud, the pressure amplitude drops in the bubble cloud region *t* and the peak pressure is moved from the liquid-only focal point to a point ahead of the cloud ($z = 0.0358$ m) reflecting the acoustic shadowing effect. This location is more clearly delineated in the following Fig. 9, which shows the temperature increase due to the ultrasound excitation at the end of the simulation. The figure clearly indicates that, in presence of the bubble cloud, the peak pressure moves from $z \sim 0.04$ m on the left plot to $z \sim 0.0358$ m on the right plot. Also, we can observe from the higher values of the temperature distribution in the presence of the bubbles, increased heating resulting from the acoustic shadowing.

Figure 10 shows the sound distribution in the computational domain. The sound speed distribution in the pure liquid region is uniform and corresponds to the sound speed in the gel, 1540 m/s. In the bubble region far from the axis, where the acoustic excitation is unable to penetrate, the sound speed is $\sim$1420 m/s,

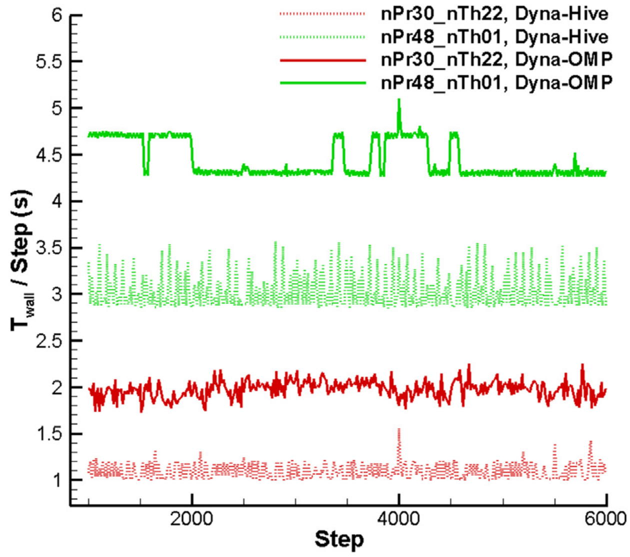

Fig. 7 Comparison of temporal variations of total bubble volumes, $\sum V_b - \sum V_{b,0}$, between using MPI alone and hybrid MPI-OpenMP parallelization methods

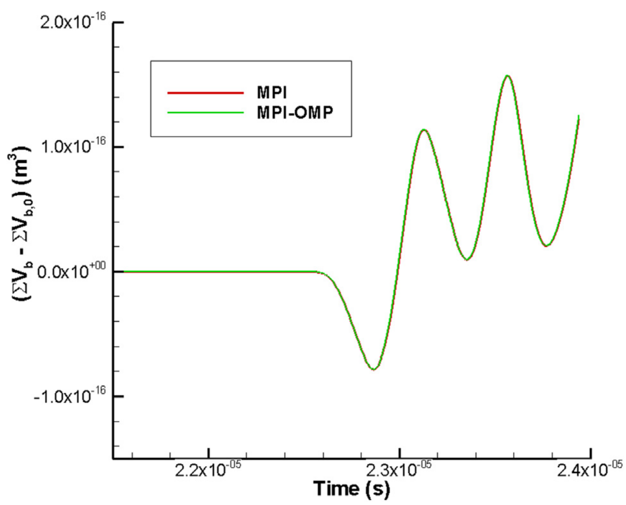

Fig. 6 Comparison of wall time histories per step using MPI alone or hybrid MPI-OpenMP on two different machines, DYNA-OMP and DYNA-HIVE. Herein, *nPr* and *nTh* denote number of MPI processes and number of OpenMP threads, respectively.

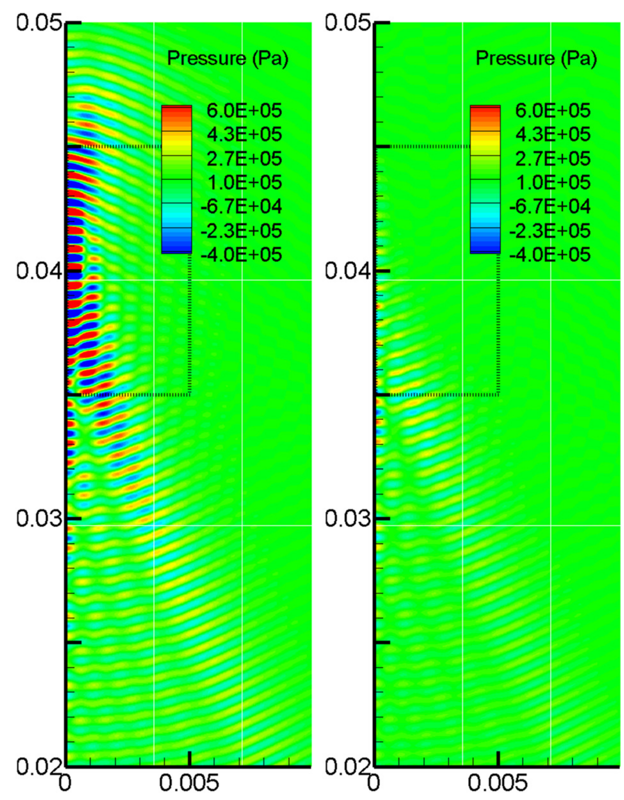

Fig. 8 Pressure contours of the case without bubbles (left) and with bubbles and void fraction 10^{-5} (right), at time 0.0352 ms. The dotted rectangles show the region where bubbles were located on the right.

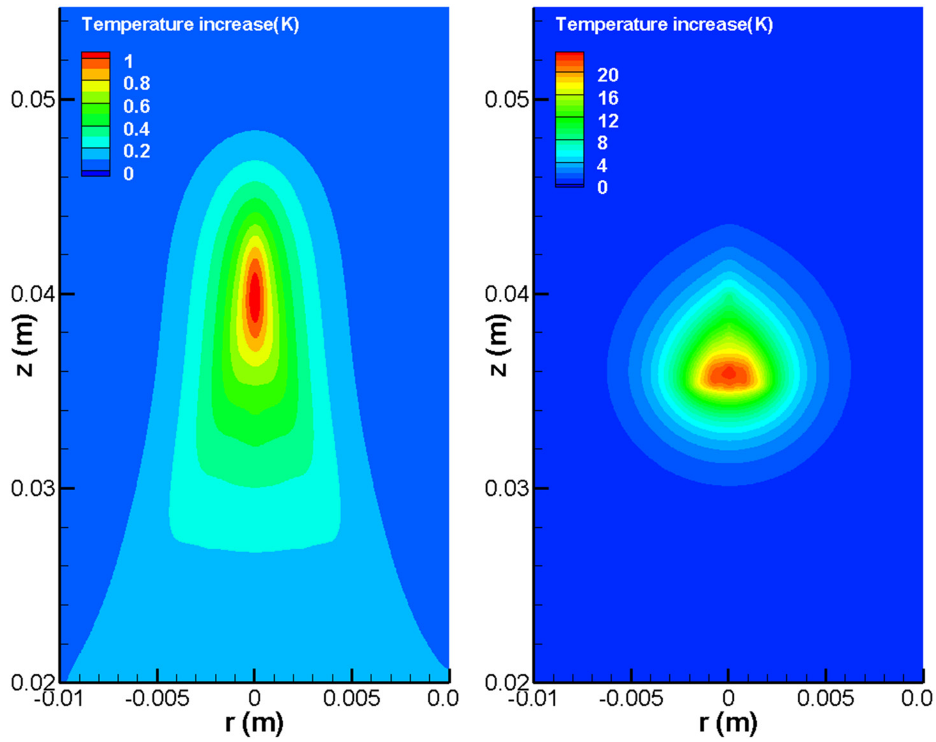

Fig. 9 Contours of temperature increase of the case without bubbles (left) and in presence of bubbles of initial void fraction 10^{-5} (right) at the end of the computation, $t = 0.0352$ ms (Note the existence of bubbles enhance the temperature rise significantly, so the contour scales are different in the two plots in order to show the temperature increase distribution more clearly)

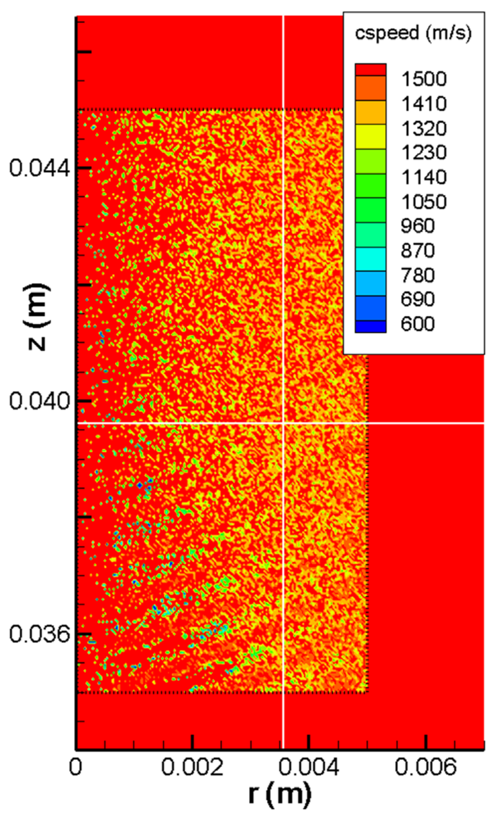

Fig. 10 Contour of sound speed of the flow field in the case with initial void fraction 10^{-5}, at time 0.0352 ms. The dotted rectangles show the region where bubbles were located.

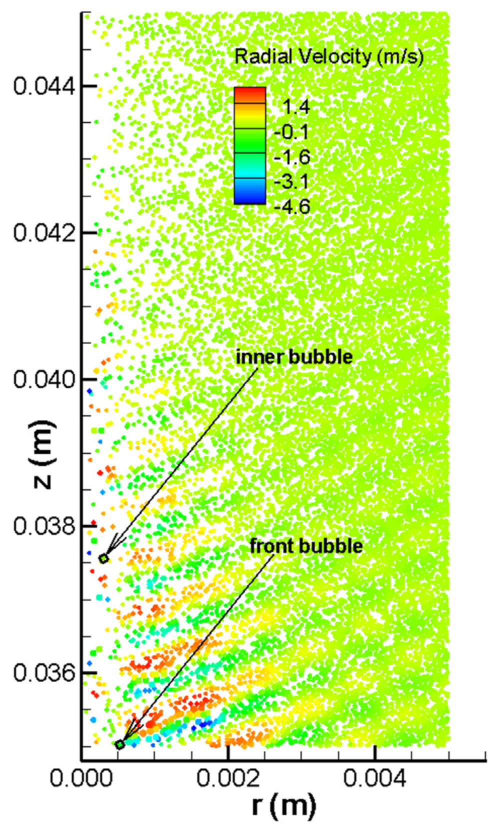

Fig. 11 Bubble distribution (with sizes increased by a factor of 100 to be visible) and radial velocities of the bubbles in the at end of the computation $t = 0.0352$ ms. Two representative bubbles are indicated for more detailed analysis.

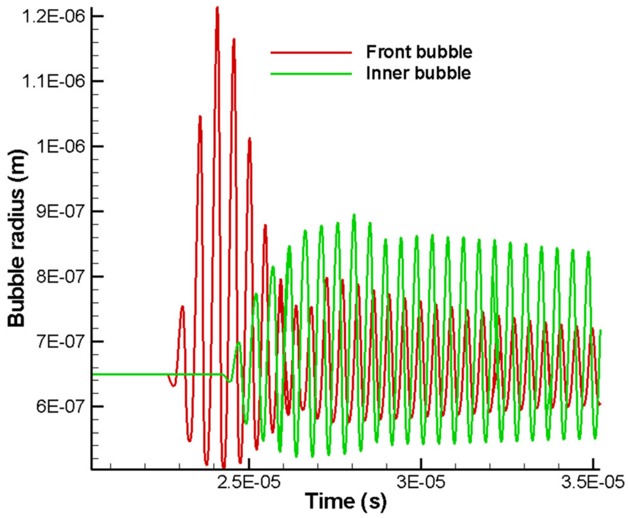

Fig. 12 Time evolution of the radii of two selected bubbles: one near the front of bubble cloud (front bubble, red curve) and one 2.5 mm from the front of bubble cloud (inner bubble, green curve). The two bubbles are indicated in Fig. 11.

a little lower than in the pure gel and corresponding to the 10^{-5}. The change of sound speed in each cell is due to the change of void fraction. Closer to the axis and near the cloud region facing the transducer, the sound speed distribution is richer with regions with sound speeds much lower than in the gel and also other locations with higher sound speed. Acoustic shadowing also results from these strong dynamics and the emission of pressure waves from the excited bubbles that block the majority of energy from penetrating the inside of the bubble cloud.

Figure 11 illustrates the radial velocities of the bubbles ($\dot{R}$ in Eq. (5)) at the end of the computation. The size of the bubbles is reflected by the diameter of each dot, which is enhanced by a factor of 100 for visualization purposes. It can be observed that some of the bubbles near the cloud bottom become much larger than those far away from the bottom because of the high-pressure fluctuations occurring in that region. The radial velocities are

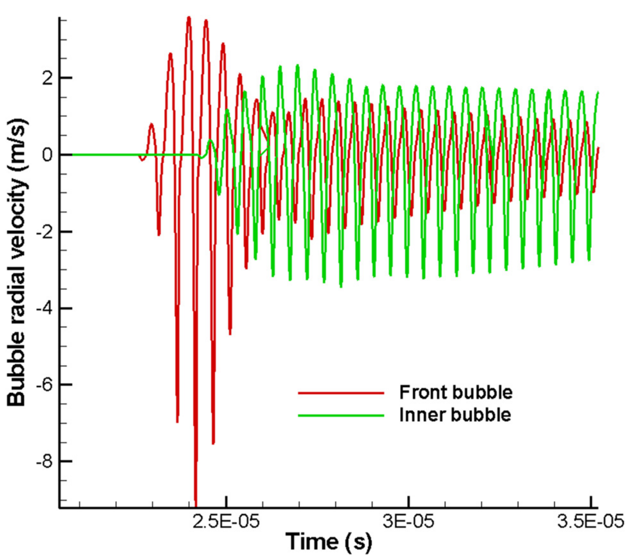

Fig. 13 Time evolution of the radial velocities of two selected bubbles: one near the front of bubble cloud (front bubble, red curve) and one 2.5 mm from the front of bubble cloud (inner bubble, green curve). The two bubbles are indicated in Fig. 11.

illustrated using the colors of the dots that represent the bubbles. The change of bubble size leads to the change of local void fraction, hence, directly driving a change in the sound speed as shown in Fig. 10.

To better visualize the bubble dynamics coupling, the radii versus time and the radial velocities of two selected bubbles are shown in Figs. 12 and 13, respectively. The front bubble is near the bottom edge of the cloud and the inner bubble is 2.5 mm inside the cloud. The plots show that when the acoustic wave reaches a bubble, it starts to oscillate. Bubble size oscillations change the local void fraction and so the local sound speed and acoustic impedance change accordingly, which will again influence the bubble dynamics. In this case, the front bubble oscillates more strongly in the early time when the front bubble region starts to be activated by the acoustic wave and has a big impedance difference to the inside. Later when the acoustic wave penetrates into the inner bubble region and inner bubbles are activated, a new balance starts to be established so the impedance difference from the front region to the inside is smaller then the front bubble oscillation is weaker and the inner bubble oscillates more strongly. Here again, the shadowing effect is illustrated by the occurrence of strong oscillations for the bubbles near the bottom edge of the cloud, while moderate oscillations occur ~2.5 mm inside the cloud. At the end of this simulation, the inner bubble is near steady-state but the front bubble still needs a longer time to reach steady-state. In future study on HIFU physics, longer-time simulations will be needed to get a steady-state for the temperature calculation.

Figure 14 compares the temperature rise, which results in tumor ablation and thus is used as an indicator of treatment efficacy, along the centerline in the presence and absence of a bubble cloud. Again, the presence of the bubbles enhanced the heating effects so that the peak temperature is much higher in the case of bubbles. Second, we can see the peak pressure shifting from 0.04 m to the

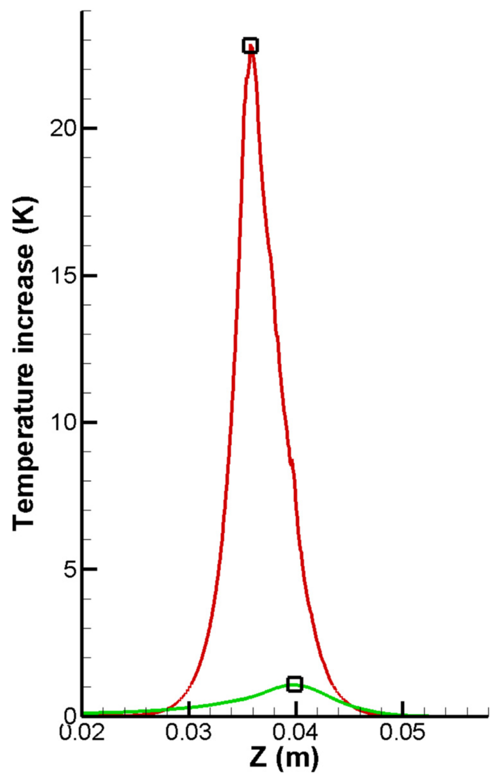

Fig. 14 Temperature increase curves along the centerline of microbubble cloud. With microbubbles (initial void fraction 10^{-5}), the temperature increase (red curve) is much higher than that without microbubbles (green curve) in the bubble cloud region, and the peak point shifted closer to the front.

lower edge of the bubble cloud indicating the cloud shadowing effect.

Summary and Discussion

Microbubble enhanced HIFU is of great interest in biomedical applications such as in tumor ablations, where contrast agents/microbubbles are injected into the targeted region to promote heating and reduce prefocal tissue damage. A multilevel MPI-OpenMP parallelized three-dimensional Eulerian–Lagrangian model was developed and applied to such a problem where bubbles are clustered near the focal regions.

This parallelization scheme adopts, at the upper level, domain decomposition for both the continuum Eulerian domain and the Lagrangian discrete bubbles, with both computations inside each subdomain handled by one MPI process. A second level of OpenMP parallelization is included to speed up the bubble computations inside subdomains, which contain bubbles. With this approach, the load imbalance due to localized bubble presence in some subdomains and not others is mitigated. This new scheme successfully takes advantage of both the scalability of MPI and the load balancing capability of OpenMP. This new procedure was shown to result in about 2 and 3 times speedup on the tested machines versus pure MPI parallelization. Note there is still space for further improvements by developing an OpenMP thread distribution methodology which may automatically assign proper numbers of different threads according to the numbers of bubbles that exist in different subdomains.

Using this hybrid MPI-OpenMP accelerated Eulerian–Lagrangian model, microbubble enhanced HIFU problem with a large number of bubbles per unit volume was simulated. The acoustic shadowing phenomenon of the bubble cloud was observed and discussed. When the bubble number density is high enough and the input pressure is strong enough, the bubble volume oscillations generate local pressure fluctuations that oppose the acoustic transducer excitation input energy and prevent it from penetrating the bubble cloud. Under these conditions, the highest pressures and temperatures are displaced from the transducer focal point to a location in front of the bubble cloud. This not only reduces the expected effectiveness of the treatment but may also result in unwanted damage to the surrounding tissue. This effort appears to indicate a path to tool useful for procedure optimization and planning before treatment.

In performing this research, we advanced the methodology and indicated application. Future efforts for the tool could explore improved OpenMP distribution methodologies. Although the present effort did display significant gains, the is some user intervention. Additionally, coupling the methods to observed tumors of complex shape would be an additional improvement to the method.

Funding Data

- NIH under SBIR Phase II (Grant No. 2R44CA213866-02; Funder ID: 10.13039/100000002).

References

[1] Kennedy, J. E., Wu, F., ter Haar, G. R., Gleeson, F. V., Philips, R. P., Middleton, M. R., and Cranston, D., 2004, "High-Intensity Focused Ultrasound for the Treatment of Liver Tumours," Ultrasonics, **42**(1–9), pp. 931–935.

[2] Hynynen, K., and Clement, G., 2009, "Clinical Applications of Focused Ultrasound—The Brain," Int. J. Hyperthermia, **23**(2), pp. 193–202.

[3] Kajiyama, K., Yoshinaka, S., Takagi, Y., and Matsumoto, Y., 2010, "Micro-Bubble Enhanced HIFU," Phys. Procedia, **3**, pp. 305–314.

[4] Razansky, D., Einziger, P. D., and Adam, D. R., 2006, "Enhanced Heat Deposition Using Ultrasound Contrast Agent—Modeling and Experimental Observations," IEEE Trans. Ultrason. Ferroelectr. Freq. Control, **53**, pp. 137–147.

[5] Chung, D. J., Cho, S. H., Lee, J. M., and Hahn, S. T., 2012, "Effect of Microbubble Contrast Agent During High Intensity Focused Ultrasound Ablation on Rabbit Liver In Vivo," Eur. J. Radiol., **81**(4), pp. e519–e523.

[6] Kaneko, Y., Maruyama, T., Takegami, K., Watanabe, T., and Mitsui, H., 2005, "Use of a Microbubble Agent to Increase the Effects of High Intensity Focused Ultrasound on Liver Tissue," Eur. Radiol., **15**, pp. 1415–1420.

[7] Gnanaskandan, A., Hsiao, C.-T., and Chahine, G. L., 2018, "Numerical Simulation of High Intensity Focused Ultrasound (HIFU) Using a Fully Compressible Multiscale Model," Proceedings of the 10th International Symposium on Cavitation (CAV2018), Baltimore, MD, May 14–16.

[8] Gnanaskandan, A., Hsiao, C.-T., and Chahine, G. L., 2019, "Modeling of Microbubble-Enhanced High-Intensity Focused Ultrasound," Ultrasound Med. Biol., **115**, pp. 25–45.

[9] Dinsmore, C., Aminfar, A., and Princevac, M., 2017, "Dissipative Effects of Bubbles and Particles in Shear Flows," ASME J. Fluids Eng., **139**(6), p. 061302.

[10] Kinzel, M. P., Lindau, J. W., and Kunz, R. F., 2019, "An Assessment of Computational Fluid Dynamics Cavitation Models Using Bubble Growth Theory and Bubble Transport Modeling," ASME J. Fluids Eng., **141**(4), p. 041301.

[11] Khojasteh-Manesh, M., and Mahdi, M., 2019, "Evaluation of Cavitation Erosion Intensity in a Microscale Nozzle Using Eulerian–Lagrangian Bubble Dynamic Simulation," ASME J. Fluids Eng., **141**(6), p. 061303.

[12] Ma, J., Gnanaskandan, A., Hsiao, C. T., and Chahine, G. L., 2021, "Message Passing Interface Parallelization for Two-Way Coupled Euler-Lagrange Simulation of Microbubble Enhanced HIFU," ASME J. Fluids Eng., **143**(8), p. 081105.

[13] Ma, J., Hsiao, C. T., and Chahine, G. L., 2015, "Shared-Memory Parallelization for Two-Way Coupled Euler-Lagrange Modeling of Cavitating Bubbly Flows," ASME J. Fluids Eng., **137**(12), p. 121106.

[14] Pelanti, M., and Shyue, K.-M., 2014, "A Mixture-Energy-Consistent Six-Equation Two-Phase Numerical Model for Fluids With Interfaces, Cavitation and Evaporation Waves," J. Comput. Phys., **259**, pp. 331–357.

[15] Colella, P., 1985, "A Direct Eulerian MUSCL Scheme for Gas Dynamics," SIAM J. Sci. Stat. Comput., **6**(1), pp. 104–117.

[16] Kapahi, A., Hsiao, C., and Chahine, G., 2015, "A Multi-Material Flow Solver for High Speed Compressible Flows," Comput. Fluids, **2**, pp. 163–179.

[17] Plesset, M. S., and Prosperetti, A., 1977, "Bubble Dynamics and Cavitation," Annu. Rev. Fluid Mech., **9**, pp. 145–185.

[18] Johnson, V. E., and Hsieh, T., 1966, "The Influence of the Trajectories of Gas Nuclei on Cavitation Inception," Sixth Symposium on Naval Hydrodynamics, Washington, DC, Sept. 27, pp. 163–179.

[19] Ma, J., Hsiao, C.-T., and Chahine, G. L., 2015, "Spherical Bubble Dynamics in a Bubbly Medium Using an Euler–Lagrange Model," Chem. Eng. Sci., **128**, pp. 64–81.

[20] Ma, J., Hsiao, C.-T., and Chahine, G. L., 2015, "Euler-Lagrange Simulations of Bubble Cloud Dynamics Near a Wall," ASME J. Fluids Eng., **137**(4), pp. 041301–041310.

[21] Ma, J., Hsiao, C.-T., and Chahine, G. L., January 2018, "Numerical Study of Acoustically Driven Bubble Cloud Dynamics Near a Rigid Wall," Ultrason. Sonochem., **40**, pp. 944–954.

[22] Okita, K., Sugiyama, K., Takagi, S., and Matsumto, Y., 2013, "Microbubble Behavior in an Ultrasound Field for High Intensity Focused Ultrasound Therapy Enhancement," J. Acoust. Soc. Am., **134**, pp. 1576–1585.