

Numerical Simulations of Cavitating Bubbles in Elastic and Viscoelastic Materials for Biomedical Applications

Thesis by
Jean-Sébastien Spratt

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

The logo for the California Institute of Technology (Caltech), featuring the word "Caltech" in a bold, orange, sans-serif font.

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2024
Defended August 25, 2023

© 2024

Jean-Sébastien Spratt
ORCID: 0000-0002-1962-4214

All rights reserved

ACKNOWLEDGEMENTS

The work presented in a Ph.D. thesis represents the significant contributions made by an individual to a given field. However, these contributions are never made in isolation, and it would be unfair to claim sole credit without recognizing the many people who played pivotal roles in making it a reality.

First and foremost, thank you to Dr. Tim Colonius for being an exemplary advisor. Your patience and guidance as I grappled with the projects that made up my Ph.D. have been invaluable. I appreciate you giving me the chance to work with you and learn from you for the past 6 years. I would also like to thank my thesis committee, Dr. Joanna Austin, Dr. John Dabiri, and Dr. Dan Meiron, for useful conversations and feedback that helped guide my project.

I've been lucky to participate in collaborative projects involving researchers across various universities. At Brown, Dr. Mauro Rodriguez and Dr. David Hennan. At the University of Wisconsin, Dr. Selda Buyukozturk, Dr. Christian Frank, and Dr. Jin Yang. At the University of Michigan, Dr. Lauren Mancina and Dr. Eric Johnsen. At Georgia Tech, Anand Radhakrishnan, Henry Le Berre, and Dr. Spencer Bryngelson. Finally, at the Inria Center at the University of Bordeaux, Dr. Kevin Schmidmayer.

I would also like to thank the present and past members of the Colonius group for creating a great work environment: Liam Heidt, Wei Hou, Michael Sleeman, Jose Rodolfo Chreim, Caroline Cardinale, Franz O'Meally, Ethan Eichberger, Hyeoksu Lee, Soto Catsoulis, Rahul Arun, Isabel Scherl and Akhil Nekkanti; Andre da Silva, Phillippe Tosi, Ethan Pickering, Ke Yu, Ben Stevens, Marcus Lee, Omar Kamal, Kevin Schmidmayer, Georgios Rigas, Benedikt Dorschner, Mauro Rodriguez, Lennart Schneiders and Shunxiang Cao.

I want to extend a special thanks to Mauro Rodriguez and Kevin Schmidmayer, for your guidance and mentorship in the early stages of my Ph.D., and to Ethan, Ke, Liam, Jose, and Michael for your friendship. I will also always be grateful to my first-year GALCIT cohort: Kai, Conor, Akshay, Eric, Hiro, Uba, Lucas, Mouâdh, Ben, Harsha, Avinash, Emile, Wesley, Alex, and Kira, without whom I would not have made it to year two

Thanks to my great friend Dr. Greg Roberts for your companionship throughout the Ph.D. journey. Thank you to my parents for never letting me think there was anything I couldn't do, or at least try. Thanks to my brother Tristan for the moral support and encouragement. and thank you to my aunt Caroline and uncle Ken for being my family away from home in the US.

Finally, thank you to my wife Niki, for sticking with me throughout this journey. Over the past 6 years, we moved in together, got engaged, and got married. Without your never-ending support throughout all those stages in our lives while I pursued this degree, I would not have even come close to making it.

ABSTRACT

The interactions of cavitating bubbles with elastic and viscoelastic materials play a central role in many biomedical applications. This thesis makes use of numerical modeling and data-driven approaches to characterize soft biomaterials at high strain rates via observation of bubble dynamics, and to model burst-wave lithotripsy, a focused ultrasound therapy to break kidney stones.

In the first part of the thesis, a data assimilation framework is developed for cavitation rheometry, a technique that uses bubble dynamics to characterize soft, viscoelastic materials at high strain-rates. This framework aims to determine material properties that best fit observed cavitating bubble dynamics. We propose ensemble-based data assimilation methods to solve this inverse problem. This approach is validated with surrogate data generated by adding random noise to simulated bubble radius time histories, and we show that we can confidently and efficiently estimate parameters of interest within 5% given an iterative Kalman smoother approach and an ensemble-based 4D-Var hybrid technique. The developed framework is applied to experimental data in three distinct settings, with varying bubble nucleation methods, cavitation media, and using different material constitutive models. We demonstrate that the mechanical properties of gels used in each experiment can be estimated quickly and accurately despite experimental inconsistencies, model error, and noisy data. The framework is used to further our understanding of the underlying physics and identify limitations of our bubble dynamics model for violent bubble collapse.

In the second part of the thesis, we simulate burst-wave lithotripsy (BWL), a non-invasive treatment for kidney stones that relies on repeated short bursts of focused ultrasound. Numerical approaches to study BWL require simulation of acoustic waves interacting with solid stones as well as bubble clouds which can nucleate ahead of the stone. We implement and validate a hypoelastic material model, which, with the addition of a continuum damage model and calibration of a spherically-focused transducer array, enables us to determine how effective various treatment strategies are with arbitrary stones. We present a preliminary investigation of the bubble dynamics occurring during treatment, and their impact on damage to the stone. Finally, we propose a strategy to reduce shielding by collapsing bubbles ahead of the stone via introduction of a secondary, low-frequency ultrasound pulse during treatment.

PUBLISHED CONTENT AND CONTRIBUTIONS

Radhakrishnan, A., Berre, H. L., Wilfong, B., **Spratt, J.-S.**, Rodriguez, M., Colonius, T., and Bryngelson, S. H. (2023). *Method for portable, scalable, and performant GPU-accelerated simulation of multiphase compressible flow*. Under Review.

Spratt, J.-S. and Colonius, T. (2023). “Numerical Investigation of Bubble Dynamics and Damage in Burst-Wave Lithotripsy.” In Preparation.

Buyukozturk, S., **Spratt, J.-S.**, Henann, D. L., Colonius, T., and Franck, C. (2022). “Particle-assisted laser-induced inertial cavitation for high strain-rate soft material characterization.” In: *Experimental Mechanics* 62.6, pp. 1037–1050. DOI: [10.1007/s11340-022-00861-7](https://doi.org/10.1007/s11340-022-00861-7).

Mancia, L., Yang, J., **Spratt, J.-S.**, Sukovich, J. R., Xu, Z., Colonius, T., Franck, C., and Johnsen, E. (2021). “Acoustic cavitation rheometry.” In: *Soft Matter* 17 (10), pp. 2931–2941. DOI: [10.1039/D0SM02086A](https://doi.org/10.1039/D0SM02086A).

Spratt, J.-S., Rodriguez, M., Bryngelson, S. H., Cao, S., and Colonius, T. (2021). “Eulerian framework for bubble-cloud-kidney stone interaction.” In: *11th International Symposium on Cavitation*. Daejeon, Korea.

Spratt, J.-S., Rodriguez, M., Schmidmayer, K., Bryngelson, S. H., Yang, J., Franck, C., and Colonius, T. (2021). “Characterizing viscoelastic materials via ensemble-based data assimilation of bubble collapse observations.” In: *Journal of the Mechanics and Physics of Solids* 152, p. 104455. ISSN: 0022-5096. DOI: <https://doi.org/10.1016/j.jmps.2021.104455>.

In the second, fifth, and sixth publications, the author implemented methods, performed simulations, and analyzed all results. In the first listed publication, the author participated in the development of the GPU solver which is the subject of the paper, and performed some of the presented simulations. In the third and fourth publications, the author adapted and applied data assimilation methods to experimental data obtained by the first authors of each publication (Dr. Selda Buyukozturk and Dr. Lauren Mancia, respectively), and analyzed and interpreted results of the applied data assimilation methods. At the time of writing, the first publication is under review at the Computer Physics Communications journal, and the second publication is being prepared for submission.

CONTENTS

Acknowledgements	iii
Abstract	v
Published Content and Contributions	vi
Contents	vii
List of Figures	ix
List of Tables	xvii
Nomenclature	xix
Chapter I: Introduction	1
1.1 Motivation	1
1.2 Inertial Microcavitation Rheometry	1
1.3 Data Assimilation	4
1.4 The DA-IMR Approach	8
1.5 Burst-Wave Lithotripsy	10
1.6 Bubble Formation during BWL	12
1.7 Fluid-Solid Solver	13
1.8 Contributions and Outline	14
Chapter II: Data Assimilation for Inertial Microcavitation Rheometry	17
2.1 Bubble Dynamics Model	17
2.2 Data Assimilation Methods	19
2.2.1 The Ensemble Kalman Filter	21
2.2.2 The Iterative Ensemble Kalman Smoother	24
2.2.3 Covariance Inflation	26
2.2.4 A hybrid Ensemble-based 4D-var Method	27
2.3 Framework Validation with Surrogate Data	28
2.3.1 Parameter Estimation Results	28
2.3.2 Uncertainty	32
Chapter III: Application of DA-IMR to Laser and Ultrasound-Induced Cavitation	37
3.1 DA-IMR with Laser-induced Cavitation Data	37
3.1.1 Experimental Setup	38
3.1.2 En4D-Var Estimates for Shear Modulus and Viscosity	39
3.1.3 Accounting for Model Uncertainty	41
3.2 DA-IMR with Laser-induced Cavitation Data in Seeded Gels	46
3.2.1 Experimental Setup and Results	47
3.2.2 DA-IMR Parameter Estimation Results	48
3.2.3 Discussion of Modeling Uncertainties	57
3.3 DA-IMR with Ultrasound-induced Cavitation Data	59
3.3.1 Experiment and Problem Setup	59
3.3.2 En4D-Var Parameter Estimation Results	60

Chapter IV: A Single Framework Fluid–Solid Numerical Solver for Burst-Wave Lithotripsy Simulations	64
4.1 Numerical Methods	64
4.1.1 The 5-equation Model with Hypoelasticity	65
4.1.2 A Continuum Damage Model	67
4.1.3 Spherical Acoustic Wave Generation from a Transducer Source	69
4.1.4 GPU Acceleration	70
4.2 Model Validation	72
4.2.1 Comparison to 1D Analytical Solution	72
4.2.2 2D Bi-Layered Media Simulations	74
4.3 Comparison of BWL Simulations to Experimental Photoelastic Stress Images	77
4.4 Frequency–Dependence of Stresses Generated by BWL	79
Chapter V: Numerical Investigation of Bubble Dynamics and Damage in Burst-Wave Lithotripsy	84
5.1 Problem Setup	85
5.2 Virtual Transducer Array Calibration	87
5.3 Bubble Shielding during Burst-Wave Lithotripsy	89
5.4 Simulations of Bubble Dynamics during BWL	92
5.4.1 1D Keller-Miksis Model	93
5.4.2 3D Plane-Wave–Bubble Simulations	94
5.4.3 Spherically focused Ultrasound–Bubble Cloud Simulations	96
5.5 Dual and Split-Frequency BWL Simulations	99
5.6 Discussion	102
Chapter VI: Conclusions and Future Work	105
6.1 Data Assimilation for Inertial Microcavitation Rheometry	105
6.2 A Numerical Framework for Full Simulations of Burst-Wave Lithotripsy	107
Bibliography	109
Appendix A: Comparing One-dimensional, Reduced-order Spherical Bubble Dynamics Models	126
A.1 Model Descriptions	126
A.1.1 Inertial Microcavitation High Strain-rate Rheometry – Estrada Model)	126
A.1.2 Bubble Dynamics with Phase Change—Preston model	129
A.2 From the Preston Model to the Estrada Model	132
Appendix B: Performance and Scaling of MFC	137
B.1 CPU Scaling and Performance	137
B.2 GPU Acceleration	138
B.2.1 Performance Speedup	138
B.2.2 Scaling Tests	140
B.3 Added Computational Cost of the Hypoelastic and Damage Models	141

LIST OF FIGURES

<i>Number</i>	<i>Page</i>
1.1 Sketch of the IMR experimental setup (A) and time-history of representative captured bubble images (B). Reprinted from Estrada et al. (2018) with permission from Elsevier, © 2018.	4
1.2 Comparison of typical SWL (top) and BWL (bottom) focal pressure waveforms. Reprinted from Maxwell, Cunitz, et al. (2015) with permission from Wolters Kluwer Health, © 2015.	11
1.3 Schematic of the experimental setup for in-vitro BWL used by Maxwell, Cunitz, et al. (2015) and representative of the setup used for most experimental studies of BWL. Reprinted from Maxwell, Cunitz, et al. (2015) with permission from Wolters Kluwer Health, © 2015.	12
2.1 Flowchart of the Ensemble Kalman Filter as applied in the framework of inertial microcavitation rheometry.	24
2.2 Simulated bubble radius and noisy sampled data used to test data assimilation methods (a), alongside simulated bubble-wall velocity, normalized bubble pressure and stress integral (b), plotted over non-dimensional time t^* . Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.	29
2.3 Estimation of shear modulus G and viscosity μ with initial guesses of $G = 3.8$ kPa and $\mu = 0.05$ Pa s (both at 50% error) in cases with surrogate experimental data (simulated radius time histories with added noise). The estimation is plotted over non-dimensional time t^* for the EnKF and IEnKS methods, and over iteration number for the En4D-Var. The dashed black line is the true parameter value, the blue circles are the estimate, and the blue shaded area is the ensemble spread. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.	32
2.4 Histogram of the final estimate for $\log(\text{Ca})$ with the lag 1 IEnKS and fitted normal curve, where n is the number of ensemble members at each value of $\log(\text{Ca})$. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.	33

2.5	Comparison of final ensembles for $\log(\text{Ca})$ (a) and $\log(\text{Re})$ (b) estimates with each tested DA method in the case with 50% initial error in both parameters. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.	33
2.6	En4D–Var estimation results for G (a) and μ (b) , for ten simulated data sets. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.	34
2.7	Histogram for G estimates combining 10 final ensembles for simulated data runs with En4D-Var. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.	35
3.1	Simulated radius vs. time curve given by En4D–Var estimates for G and μ , and experimental measurements for data set 10. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.	40
3.2	En4D–Var estimates for shear modulus (a) and viscosity (b) for 10 experimental data sets. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.	42
3.3	Comparison of final combined ensembles for viscosity estimation in simulated and experimental data. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.	42
3.4	Comparison of online estimation of viscosity in data sets 2 (a) and 3 (b) using the IEnKS–MDA. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.	43
3.5	Comparison of online estimation of the temperature field in the surrounding material for data sets 2 (a) and 3 (b) using the IEnKS–MDA. The x-axis is non-dimensional time, the y-axis represents distance from the bubble wall, and the color of each cell represents temperature.	44
3.6	Bar plot of radius normalized root mean squared errors (NRMSE) for each data set. Also plotted are the previous estimate mean NRMSE (mean of all sets), the mean NRMSE for sets 3 to 5, and the final estimate mean NRMSE (mean of all other sets). Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.	45

3.7	Value of Maximum stretch ratio λ_{\max} vs. laser energy obtained across experiments with varying seed particles. All points represent the mean of 8 experimental data sets. Reproduced from Buyukozturk et al. (2022) with permission from Springer Nature, © 2022.	48
3.8	Results of offline En4D-Var fitting on experimental curves for 1-, 2- and 3-peak fits. Representative examples of multi-peak fits with median NRMSE values (for PA paramagnetic particles at nominal laser energy 117 μJ) are shown for a (a) 1-peak fit (NRMSE = 0.009), (b) 2-peak fit (NRMSE = 0.027), and (c) 3-peak fit (NRMSE = 0.047). NRMSE values with respect to maximum material stretch for (d) 1-peak, (e) 2-peak, and (f) 3-peak fits, and with respect to laser energy for (g) 1-peak, (h) 2-peak, and (i) 3-peak fitting cases. Reproduced from Buyukozturk et al. (2022) with permission from Springer Nature, © 2022.	51
3.9	Offline En4D-Var estimates for Quadratic law Kelvin-Voigt material model parameters, namely viscosity μ , and strain stiffening parameter α . The material parameters estimates are shown for (a) 1-peak, (b) 2-peak, and (c) 3-peak fits. The symbols correspond to sample type and the color corresponds to laser energy, with the associated color bar shown on the far right. Each point is a weighted average of $n = 8$ experiments where error bars are standard deviation. Reproduced from Buyukozturk et al. (2022) with permission from Springer Nature, © 2022.	52
3.10	Critical Mach number, M_{cr} plotted with respect to maximum material stretch, λ_{\max} , Using offline (En4D-Var) estimates given the (a) 1-peak, (b) 2-peak, and (c) 3-peak fitted solutions. Critical Mach numbers were obtained by re-running simulations given these En4D-Var estimates in each case. Reproduced from Buyukozturk et al. (2022) with permission from Springer Nature, © 2022.	54

3.11	Quasi-online IEnKS α and μ estimate variance versus normalized time at nominal laser energies of (a) $\sim 117 \mu\text{J}$ (b) $\sim 254 \mu\text{J}$, (c) $\sim 337 \mu\text{J}$, and (d) $\sim 449 \mu\text{J}$. This is for bubble radius time histories in PA with no beads, where the IEnKS was run using an ensemble size of 48. First bubble collapse occurs at approximately the dashed line for all $n = 8$ samples per experimental condition. Reproduced from Buyukozturk et al. (2022) with permission from Springer Nature, © 2022.	55
3.12	Comparison of results of the quasi-online IEnKS fitting for two extreme cases with normalized radius versus time curves in PA (no beads) at nominal laser energy $449 \mu\text{J}$ and (b) PA (paramagnetic particles) at nominal laser energy $18 \mu\text{J}$. The (c, d) α (e, f) and μ estimates over normalized time for each case are plotted. Reproduced from Buyukozturk et al. (2022) with permission from Springer Nature, © 2022.	57
3.13	Histogram of combined final stress-free radius, shear modulus, stiffening parameter and viscosity ensembles for 0.3% gels with the En4D-var using the QLKV model. Reproduced from Mancina, Yang, et al. (2021) with permission from the Royal Society of Chemistry, © 2021.	62
3.14	Iterative estimation of the stress-free radius, shear modulus, stiffening parameter and viscosity for 0.3% gels with the En4D-var using the QLKV model. Reproduced from Mancina, Yang, et al. (2021) with permission from the Royal Society of Chemistry, © 2021.	63
4.1	Comparing the simulated results (blue dotted line) to the exact solution (black line) for a 1D impact problem with initial conditions described by (4.21) at $t = 64 \mu\text{s}$	73
4.2	Comparing the simulated results (blue dotted line) to the exact solution (black line) for the 1D Riemann problem described by initial conditions (4.22) at $t = 64 \mu\text{s}$	74
4.3	Sketch of the setup for a 2D bi-layered media problem, showing the location of wave source and probe as well as material properties of each layer.	75

4.4	Simulated velocity magnitudes in bi-layered media. A source in the liquid (top half) emits a Ricker wavelet which reflects and propagates through the material interface with the elastic solid (bottom half). Resulting wave patterns are shown at two different times t as labeled.	75
4.5	Velocity in the x (top) and y-direction (bottom) at the probe point in the elastic solid layer, located 1000 m below the source in the y-direction and 2000 m past the source in the x-direction. The probe is 500 m below the water-solid interface.	76
4.6	Setup of the simulation with a submerged epoxy stone on the right and simplified, single-element ultrasound transducer on the left, both submerged in water.	78
4.7	Propagation of pressure in the fluid and maximum principal stress in an epoxy stone submerged in water as a result of a 20-cycle BWL ultrasound pulse from a single-element transducer at 340 kHz, shown for three distinctive times as labeled.	78
4.8	Comparison of computational photoelastic stress images in an epoxy stone obtained in our simulation (top) and experimental images obtained by Sapozhnikov, Maxwell, and Bailey (2020) which are also reported in Maxwell, MacConaghy, et al. (2020) The bottom experimental images are reprinted with permission from Sapozhnikov, Maxwell, and Bailey (2020). © 2020, Acoustical Society of America.	79
4.9	Maximum value of σ_1 in a 5 mm diameter circular stone over the course of a BWL simulation with a single spherical transducer element at three test frequencies.	81
4.10	Maximum value of σ_1 in a 10 mm diameter circular stone over the course of a BWL simulation with a single spherical transducer element at three test frequencies.	82
4.11	Maximum value of σ_1 in a 5 mm length rectangular stone over the course of a BWL simulation with a single spherical transducer element at three test frequencies.	83
5.1	Simplified sketch of the setup used for BWL simulations throughout Chapter 5.	85
5.2	2D example of the 4 edge waves generated by transducer elements, for a transducer with focal length 150 mm and aperture 180 mm with 4 elements. Half of the domain is shown (positive y half-domain). . .	88

5.3	Comparison of probed pressures for the full transducer (black solid line) and the scaled-down transducer used in simulations (dotted colored line) at three locations corresponding to the stone center (a), stone edge (b), and center of the proximal stone surface (c).	88
5.4	Virtual array used in BWL simulations viewed face-on (a), and from the side in the context of the full setup of a BWL simulation (b). . . .	89
5.5	Propagation of the maximum principal stress in the stone and damage caused during BWL at 375kHz when no bubbles are present. Figures (a), (b), (c) show three different time steps, and figure (d) shows maximum value of σ_1 in the top half, and damage field D in the bottom half at the end of the simulation.	90
5.6	Propagation of the maximum principal stress in the stone and damage caused during BWL at 375kHz with a dense bubble cloud ($A = 3 \text{ mm}^2$) present. Bubbles are shown with 2 contour levels for the volume fraction of air: 0.2 (low opacity) and 0.8 (higher opacity). Figures (a), (b), (b) show three different time steps, and figure (d) shows maximum value of σ_1 in the top half, and damage field D in the bottom half at the end of the simulation.	91
5.7	Damage to the stone during BWL with default transducer parameters and a sparser bubble cloud ($A = 2 \text{ mm}^2$) present. Maximum value of σ_1 over time is plotted in the top half of the stone, and damage field D in the bottom half.	92
5.8	Radius time histories for 250 μm bubble under ultrasound at various frequencies with an amplitude of 1 MPa.	94
5.9	Oscillation of a 250 μm bubble under a plane wave at 375 kHz.	94
5.10	Oscillation of a 250 μm bubble under a plane wave at 20 kHz.	95
5.11	Oscillation of a dense bubble cloud ($A = 3 \text{ mm}^2$) under a plane wave at 375 kHz.	95
5.12	Oscillation of a dense bubble cloud ($A = 3 \text{ mm}^2$) under the superposition of plane waves at 375 kHz and 20 kHz.	96
5.13	Oscillation of a dense bubble cloud ($A = 3 \text{ mm}^2$) exposed to BWL transducer ultrasound at 375 kHz.	97
5.14	Oscillation of a dense bubble cloud ($A = 3 \text{ mm}^2$) exposed to BWL transducer ultrasound in the dual frequency configuration.	98
5.15	Oscillation of a dense bubble cloud ($A = 3 \text{ mm}^2$) exposed to BWL transducer ultrasound in the split frequency configuration.	98

5.16	Split frequency transducer configuration: orange elements fire at 20 kHz, blue elements at 375 kHz.	99
5.17	Propagation of the maximum principal stress in the stone and damage caused during BWL with the dual-frequency configuration with a dense bubble cloud ($A = 3 \text{ mm}^2$) present. Bubbles are shown with 2 contour levels for the volume fraction of air: 0.2 (low opacity) and 0.8 (higher opacity). Figures (a), (b), (c) show three different time steps, and figure (d) shows maximum value of σ_1 in the top half, and damage field D in the bottom half at the end of the simulation.	100
5.18	Propagation of the maximum principal stress in the stone and damage caused during BWL with the split-frequency configuration with a dense bubble cloud ($A = 3 \text{ mm}^2$) present. Bubbles are shown with 2 contour levels for the volume fraction of air: 0.2 (low opacity) and 0.8 (higher opacity). Figures (a), (b), (c) show three different time steps, and figure (d) shows maximum value of σ_1 in the top half, and damage field D in the bottom half at the end of the simulation.	101
5.19	Damage to the stone during BWL with split-frequency configuration and a sparse bubble cloud ($A = 2 \text{ mm}^2$) present. Maximum value of σ_1 over time is plotted in the top half of the stone, and damage field D in the bottom half.	102
5.20	Dense bubble cloud ($A = 3 \text{ mm}^2$) deformation during treatment at $t = 30 \mu\text{s}$ for the (a) dual frequency and (b) split frequency cases. Bubbles are shown with 2 contour levels for the volume fraction of air: 0.2 (low opacity) and 0.8 (higher opacity)	103
B.1	Strong scaling (a) of MFC for 500^3 grid points and weak scaling (b) for 50^3 grid points per core, both up to 4096 cores. Scaling test performed on up to 32 compute nodes on the <i>Expansive</i> supercomputer at the San Diego Supercomputer Center.	138
B.2	Strong scaling (a) of MFC for 64M grid points up to 128 GPU with and without CUDA-aware MPI, and weak scaling (b) for 1M grid points per GPU up to $O(10000)$ GPUs. Scaling test performed on the <i>Summit</i> supercomputer at the Oak Ridge Leadership Computing Facility and reproduced from Radhakrishnan et al. (2023).	140

B.3 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 $\alpha = 0.5$	143
--	-----

LIST OF TABLES

<i>Number</i>	<i>Page</i>
2.1 Model parameters as they follow from Estrada et al. (2018). Table reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.	29
2.2 Comparing the accuracy of estimation and run time with 3 different initial guesses for the material parameters G and μ . Runs were performed on a computer with dual 12-core 2.3 GHz processors. Table reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.	30
3.1 Comparing results of estimation with three different initial guesses for the material parameters G and μ . Runs were again performed on a computer with dual 12-core 2.3Ghz processors. Table reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.	39
3.2 Final En4D–Var estimates (discarding three outlier data sets) and standard deviation, along with the average radius normalized root mean squared error. The previous best estimate corresponds to the mean of all 10 data sets, outliers included. Table reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.	45
3.3 Weighted mean and standard deviation of inferred properties using En4D-Var for 0.3% and 1% agarose specimens. Table reproduced from Mancia, Yang, et al. (2021) with permission from the Royal Society of Chemistry, © 2021.	61
4.1 Total time per time-step comparison for a 3D MFC test run with 10 million grid points with two components, with different models activated.	71
4.2 Maximum value of the normalized maximum principal stress over a BWL simulation with a single spherical element at three fixed frequencies, for circular stones with diameter 5 mm and 10 mm; bold values represent the maximum stress values in each case.	81

4.3	Maximum value of the normalized maximum principal stress over a BWL simulation with a single spherical element at three fixed frequencies, for rectangular stones with width 5 mm, and length 5 mm and 10 mm, respectively; bold values represent the maximum stress values in each case.	82
5.1	Comparison of resulting average damage and $\max(\sigma_1)$ values over the stone region for three BWL cases with: no bubbles, a sparse bubble cloud ($A = 2 \text{ mm}^2$) and a dense bubble cloud ($A = 3 \text{ mm}^2$). All cases are with 375 kHz ultrasound.	92
5.2	Comparison of resulting average damage and $\max(\sigma_1)$ values over the stone region between the baseline BWL 375 kHz waveform case and dual frequency case for a sparse bubble cloud ($A = 2 \text{ mm}^2$) and a dense bubble cloud ($A = 3 \text{ mm}^2$).	102
B.1	Speedup of GPU code running on one NVIDIA V100 GPU when compared to one POWER9 CPU core. These speedup numbers are detailed in Radhakrishnan et al. (2023)	139
B.2	Percentage of time-contribution of the WENO Reconstruction, Riemann solver, communication, and all other routines to total compute time.	140
B.3	Total time per time-step comparison for a 3D case with 10 million grid points with two components, with various MFC configurations. Percentages shown correspond to slowdown relative to the base MFC case for each version.	142
B.4	Speedup of GPU code compared to CPU code in 2D and 3D, when running on a full CPU node (128 cores) and full GPU node (4 GPUs) on <i>Expanse</i>	142

NOMENCLATURE

BWL. Burst-Wave Lithotripsy.

CCL. Cavitation Control Lithotripsy.

DA. Data Assimilation.

En4D-Var. Ensemble based 4D-Var hybrid method.

EnKF. Ensemble Kalman Filter.

Gen QLV. General Quadratic Law Kelvin-Voigt model.

IEnKS. Iterative Ensemble Kalman Smoother.

IEnKS-MDA. Iterative Ensemble Kalman Smoother - Multiple Data Assimilation.

IEnKS-SDA. Iterative Ensemble Kalman Smoother - Single Data Assimilation.

IMR. Inertial Microcavitation Rheometry.

LIC. Laser-Induced Cavitation.

MFC. [Multi-component Flow Code](#).

NHKV. Neo-Hookean Kelvin-Voigt model.

NRMSE. Normalized Root-Mean-Squared Error.

PDF. Probability Density Function.

QLKV. Quadratic Law Kelvin-Voigt model.

QS QLV. Quasi-Static Quadratic Law Kelvin-Voigt model.

SWL. Shock-Wave Lithotripsy.

Chapter 1

INTRODUCTION

1.1 Motivation

In biomedical sciences, numerical and modeling approaches to study biological systems often present significant challenges due to their multiscale nature and the interaction of various materials ranging from liquids and soft tissue to elastic solids. Many applications in the medical field involve the—often nonlinear—dynamical behavior of these materials in high strain-rate and/or high stress regimes. One example where such strain-rates and stresses are generated is during cavitation of single or clouds of bubbles. Treatments making use of focused ultrasound, shock waves, or laser surgery can cause nucleation of bubble clouds which collapse rapidly, deforming the surrounding material at high rates and inducing very high pressures at the bubble nucleation site. Understanding the interaction of such cavitating gas bubbles with viscoelastic materials and stiff elastic solids is thus valuable in various applications in the biomedical field.

In this thesis, we present numerical and data-driven methods for the modeling of multiphase physics involving cavitating bubble dynamics in and near elastic and viscoelastic materials. In the first part, we make use of data assimilation with bubbles cavitating in a viscoelastic hydrogel to estimate its mechanical properties. This method to characterize soft materials can aid in the modeling of biomaterials in settings where high strain rates occur. In the second part, we implement a numerical framework for the direct numerical simulation of acoustic wave–bubble–stone interactions. This is used to simulate burst-wave lithotripsy, a treatment to break kidney stones through exposure to repeated bursts of focused ultrasound. In particular, this framework enables a better understanding of the impact and dynamics of bubbles that can form near the stone during treatment.

1.2 Inertial Microcavitation Rheometry

Across various engineering disciplines, particularly in medicine and the biomedical field, knowing the mechanical properties of soft materials such as hydrogels,

polymers, and biomaterials is of great importance (Chaudhuri et al., 2016; Lee and Mooney, 2012; Storrie and Mooney, 2006; Solomon and Jindal, 2007). This includes applications where high strain rates occur in biological tissue, for example during impact and blast exposure (Bar-Kochba et al., 2016; Sarntinoranont et al., 2012; Meaney and Smith, 2011; Nyein et al., 2010; Ramasamy et al., 2011), therapeutic ultrasound (Maxwell, Cain, Duryea, et al., 2009; Xu et al., 2007; Mancina, Vlasisavljevich, Xu, et al., 2017; Mancina, Vlasisavljevich, Yousefi, et al., 2019; Vlasisavljevich, Lin, Warnez, et al., 2015; Bailey, Khokhlova, et al., 2003) or laser surgery (Brujan and Vogel, 2006; Vogel et al., 2008). However, measuring the mechanical properties of soft, viscoelastic materials at these high strain rates (exceeding 10^3 s^{-1}) is challenging, in particular due to the high compliance of these materials (Arora, Narani, and McCulloch, 1999) and the strain rate dependence of their properties (Brujan and Vogel, 2006). At high strain rates, impact (Taylor, 1948; Allen, Rule, and Jones, 1997) or Kolsky bar tests (Chen and Song, 2010) can be used, but neither is adequate for soft materials, as they rely on high loading rates and applied stresses, incompatible with the high compliance of viscoelastic materials of interest (Hu, Zhao, et al., 2010; Hu, You, et al., 2012). Thus, measuring the mechanical properties of soft, viscoelastic materials at high strain rates remains a challenging goal of rheometry.

To this end, methods relying on cavitation in these soft materials have been developed. The bubble dynamics that occur when a bubble collapses in a soft material are sensitive to the material's properties, and thus observation of these dynamics can be exploited to characterize these properties. The first such method, called the Cavitation Rheology Technique (CRT), was introduced by Zimmerlin, Sanabria-DeLong, et al. (2007). They make use of needle-induced cavitation, where a cavity is created in a gel at the tip of a needle by introducing pressurized air. The elastic modulus of the gel is deduced by determining the critical pressure where rapid deformation of the gel occurs. This technique has been studied extensively, and used, for example, to determine the mechanical properties of polymers and biomaterials in the eye and skin (Barney et al., 2020; Zimmerlin, Sanabria-DeLong, et al., 2007; Zimmerlin, McManus, and Crosby, 2010; Cui et al., 2011; Chin et al., 2013; Bentz, Walley, and Savin, 2016). An extension of this method, which does not rely on determining the critical pressure, but instead tracking the volume expansion of the cavity, was introduced by Raayai-Ardakani, Chen, et al. (2019) and Raayai-Ardakani and Cohen (2019) and used to characterize biological tissue (Mijailovic et al., 2021; Van

Sligtenhorst, Cronin, and Brodland, 2006). While effective for smaller strain rates, these methods cannot reach the high strain rate regime of interest to the biomedical applications described above (Chockalingam et al., 2021).

To this end, Estrada et al. (2018) proposed a high-strain rate rheometer to estimate the viscoelastic properties of polyacrylamide gels through observation of the bubble radius time history during a laser-generated cavitation event in a sample of the material. Cavitation occurring in liquids and soft materials on exposure to tensile waves of sufficient amplitude will induce very high strain rates in the material, within the regime of interest. The technique, called Inertial Microcavitation Rheometry (IMR), compares bubble radius time-histories obtained through high-speed imaging of laser-induced cavitation (LIC) in the material of interest, to simulated radius vs. time curves. Using an adequate theoretical cavitation model derived from literature on cavitation in a fluid (Akhatov et al., 2001; Epstein and Keller, 1972; Flynn, 1975; Fujikawa and Akamatsu, 1980; Keller and Kolodner, 1956; Keller and Miksis, 1980; Nigmatulin, Khabeev, and Nagiev, 1981; Prosperetti, 1991; Prosperetti, Crum, and Commander, 1988; Prosperetti and Lezzi, 1986), an array of simulations are run with varying material properties and the difference with experimental data is minimized through least-squares fitting to determine parameters which best match the experiment. In this way, properties of the material can be inferred through observation of these cavitating bubbles. Estrada et al. (2018) demonstrate that they can adequately infer the shear modulus and viscosity of polyacrylamide gels with varying stiffness using this method. This technique can achieve a broad range of strain rates, from $O(10^3)s^{-1}$ to $O(10^8)s^{-1}$. It is the first reliable technique for rheometry of soft materials in this regime. The method is minimally invasive and the setup relatively straightforward, making it broadly applicable to an array of gels or other soft materials with high compliance. The experimental setup, as well as representative bubble images captured by the high-speed camera which are used to determine bubble radius time-histories are shown in figure 1.1.

Results obtained in well-characterized polyacrylamide hydrogels by Estrada et al. (2018) are promising. However, there are a few areas where their method could be improved. First, the computational cost of the least-squares fitting approach used is not easily scalable, particularly to the estimation of additional parameters, as may be necessary given other material models. Only two parameters (shear modulus and viscosity) are estimated by Estrada et al. (2018), which already required N^2 simulations, with N the number of guesses per parameter. Unless a good estimate

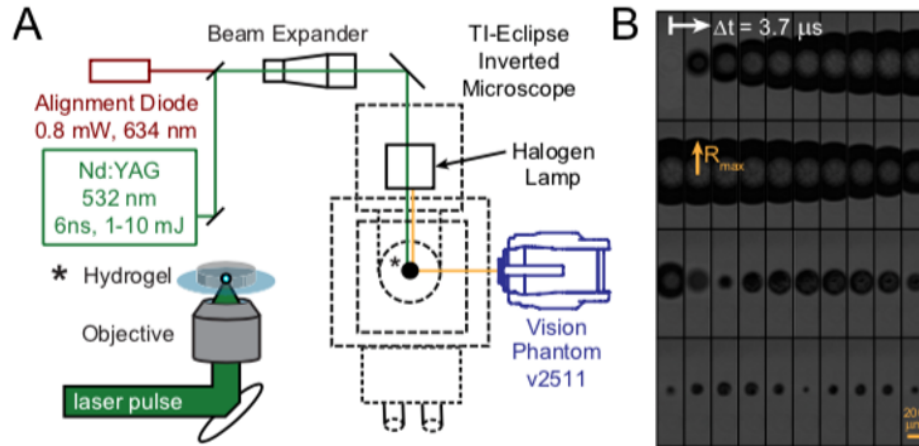


Figure 1.1: Sketch of the IMR experimental setup (A) and time-history of representative captured bubble images (B). Reprinted from Estrada et al. (2018) with permission from Elsevier, © 2018.

for each parameter is known a priori, N must be relatively large to obtain a precise estimate, with its size growing by an order of magnitude with each added significant figure. If material models require a third parameter to estimate, the number of simulations rises to N^3 , and so on for each additional parameter. In the case tested where only two parameters are estimated, in a material that has been extensively studied, this is tractable. However, scaling to other methods may prove difficult.

Second, the reliance of IMR on laser-induced cavitation introduces some uncertainty in the parameter estimation. The physics of laser nucleation are not captured in the model used, particularly in the initial growth phase. Thus, Estrada et al. (2018) fit radius curves starting at the maximum bubble radius. This is a reasonable approach to avoid the complicated growth phase, and some uncertainty cannot be avoided unless a new bubble-dynamics model is used that captures plasma generation by the laser. However, IMR would benefit from a framework that better quantifies this uncertainty during parameter estimation and provides additional information as to the expected error in obtained estimates.

1.3 Data Assimilation

The fundamental idea behind data assimilation (DA) is to combine a mathematical model for a phenomenon with data, or observations, in order to best determine its state and predict its behavior. In many cases, such as those explored in this thesis, we begin with a forward model for the system. That is, we derive a mathematical model

which, given a system at time t , will predict its state at time $t + 1$. This model can be analytical, or a simulation tool providing an approximate solution. At each time when data (experimental or otherwise) is available, it is used to correct and improve the state prediction. This is often done sequentially, where the model forecast and data-informed correction are applied in turn at each time, but can also be used to determine the initial state of a system and estimate particular parameters. Data assimilation can thus be considered a framework to solve inverse problems, where data is combined with a mathematical model to estimate unknown parameters in the initial condition of a dynamical system (Sanz-Alonso, Stuart, and Taeb, 2023; Law, Stuart, and Zygalakis, 2015; Schillings and Stuart, 2017). Most DA methods are designed to account for uncertainty in both the model and data, which makes them versatile and adaptable to contexts with noisy data or using reduced-order models.

Applications of data assimilation are plentiful and span many disciplines. One example is the state estimation for guidance, navigation and control of aircraft, spacecraft, and other vehicles. Here, observation data from various onboard instruments can be combined online with modeled attitude and position for precise determination of the state of the craft or vehicle (Lefferts, Markley, and Shuster, 1982; Kim et al., 2007; Madyastha et al., 2011; Venhovens and Naab, 1999; Wenzel et al., 2006; Antonov, Fehn, and Kugi, 2011). Another is atmospheric and oceanographic modeling, for example for weather forecasting where sparse data points (e.g., pressure, temperature, and wind measurements) are combined with fluid dynamics models to best predict weather and atmospheric conditions. This problem is a great example of how useful DA can be. Indeed, precise direct modeling of weather conditions on an atmospheric scale is not feasible given current (or near-future) computing power. On the other hand, only limited data points can be obtained at weather tracking stations or using weather balloons. Neither modeling tools nor data alone are sufficient to accurately predict future weather conditions. However, combining reduced-order fluid dynamical models with sparse atmospheric measurements in a data assimilation framework is a powerful tool for atmospheric modeling and weather prediction. In fact, many modern DA methods have emerged from atmospheric sciences and weather forecasting, where it is widely used (Houtekamer and Zhang, 2016; Leeuwen, 2009; Lorenc, 1986; Bengtsson, Snyder, and Nychka, 2003; Evensen and Van Leeuwen, 1996). The DA methods used in this thesis are all extensions or variations on perhaps the most well-known and widely used DA

method, which was foundational to the field of data assimilation: the Kalman filter (Kalman, 1960; Kalman and Bucy, 1961).

The Kalman Filter and Kalman Smoother To describe the Kalman filter, we begin by defining the linear dynamics model

$$x_{k+1} = Fx_k + \eta_k \quad (1.1)$$

where $F \in \mathbb{R}^{d \times d}$ is a linear operator for the state transition model, which maps a state $x \in \mathbb{R}^d$ from time k to $k + 1$, and $\eta_k \sim \mathcal{N}(0, \Sigma)$ is the model error, assumed Gaussian with mean 0 and variance Σ . The state is viewed as a random variable $x_k \sim \mathcal{N}(m_k, C_k)$ with mean m_k and covariance C_k . The estimate of the state at a time k is thus the mean of this random variable x_k . The state transition model F is generally based on a mathematical model for the evolution of the dynamical system, and must, in this case, be linear to be represented in matrix form.

We also define a linear observation model

$$y_{k+1} = Hx_{k+1} + v_{k+1}, \quad (1.2)$$

where $H \in \mathbb{R}^{n \times d}$ is the linear observation operator, which maps the state x_{k+1} to measurement space, where $y_{k+1} \in \mathbb{R}^n$ is the observation (data) at time $k + 1$. $v_k \sim \mathcal{N}(0, \Gamma)$ is the measurement noise, again assumed Gaussian with mean 0 and variance Γ . The Kalman filtering process, which estimates the optimal (estimation variance minimizing) state at step $k + 1$ given the dynamical system described by equations (1.1) and (1.2) can be broken down into a forecast step and an analysis step. We track the random Gaussian state variable x_k via its mean m_k and covariance C_k at any given time k . The mean and covariance are forecast and updated using the following. We note that only the final formulas used in the Kalman filter are reported here. A clear and thorough example of their full derivation can be found in Sanz-Alonso, Stuart, and Taeb (2023).

Forecast Step Since the state is represented by a Gaussian random variable (which is independent of the model error noise η , and the forecast operator is linear, we can write

$$\widehat{m}_{k+1} = Fm_k \quad (1.3)$$

$$\widehat{C}_{k+1} = FC_kF^T + \Sigma, \quad (1.4)$$

$$(1.5)$$

where $\hat{\cdot}$ represents forecast variables. Here, we simply propagate the mean and covariance of our state through the forecast operator.

Analysis Step Once the forecast step is done, we introduce our measurement y_{k+1} to correct the forecast state statistics. This yields

$$m_{k+1} = \hat{m}_{k+1} + K_{k+1}(y_{k+1} - H\hat{m}_{k+1}) \quad (1.6)$$

$$C_{k+1} = (I - K_{k+1}H)\hat{C}_{k+1}, \quad (1.7)$$

where

$$K_{k+1} = \hat{C}_{k+1}H^T(H\hat{C}_{k+1}H^T + \Gamma)^{-1} \quad (1.8)$$

is called the *Kalman gain*, and the vector $y_{k+1} - H\hat{m}_{k+1}$ in equation (1.6) is called the *innovation*. The Analysis step can also be re-framed from an optimization perspective, which will be particularly useful when comparing it to other DA methods. We can re-write this analysis step as the minimization of a cost function J_{KF} :

$$m_{k+1} = \operatorname{argmin}_x J_{KF}(x), \quad (1.9)$$

where

$$J_{KF}(x) = \frac{1}{2}\|y_{k+1} - Hx\|_{\Gamma}^2 + \frac{1}{2}\|x - \hat{m}_{k+1}\|_{\hat{C}_{k+1}}^2. \quad (1.10)$$

This cost function helps understand the Kalman filter analysis intuitively. The first term minimizes the difference between the state vector mapped to measurement space and the observation, weighed by the measurement noise covariance. The second term minimizes the difference between the state vector with the projected state vector given the physical model, weighed by the model error covariance.

Another class of DA methods are smoothers. The difference between *filtering* and *smoothing* is that in filtering, data is incorporated into the state estimate sequentially at each time step. In this sense, it is often referred to as an ‘online’ method, as the estimate is adjusted at each time. By contrast, smoothing is done ‘offline,’ incorporating data from multiple time steps to adjust the state estimate. In filtering, only data from time k is used to estimate x_k . In smoothing, data from multiple times (e.g., k to $k+n$) is used to estimate this same state. The Kalman smoother is a simple extension of the Kalman filter, best understood in the optimization framework, where

the cost function is modified to account for data at multiple times, yielding

$$J_{KS}(X) = \frac{1}{2} \|x_0 - m_0\|_{C_0}^2 + \frac{1}{2} \sum_{k=0}^{N-1} \|y_{k+1} - Hx_{k+1}\|_{\Gamma}^2 + \frac{1}{2} \sum_{k=0}^{N-1} \|x_{k+1} - Fx_k\|_{\hat{C}_{k+1}}^2, \quad (1.11)$$

where $X = \{x_0, \dots, x_N\}$ represents the time-series of the state from time 0 to N . This cost function minimizes once again the difference with both data and forecast state, but this time over a series of times 0 to N , and the state at all times is estimated.

1.4 The DA-IMR Approach

We propose a data assimilation approach to inertial microcavitation rheometry. The goal of IMR, namely estimation of the mechanical properties of a viscoelastic material by observation of bubble dynamics, is posed as an inverse problem that is solved using data assimilation. In this framework, we require a bubble dynamics solver for our operator F in the dynamical model given in equation (1.1), and the state x will be comprised of all dependent variables of this solver. The data y is simply the bubble-radius observations, and the observation operator H maps this observed radius to the state space. The key to using data assimilation to estimate material properties is then simply to append these unknown properties to the state vector x , with an initial estimate at initial time for x_0 . As long as the dynamics of the bubble radius evolution are sensitive to changes in these properties, they will be estimated at each step in the filtering. This represents the basis of DA-IMR, our data assimilation approach to inertial microcavitation rheometry.

The goal of DA-IMR is twofold. First, this method can vastly improve the efficiency of IMR. IMR requires least-squares fitting of a large number of curves to experimental data and thus requires N^i bubble collapse simulations, where i is the number of parameters to estimate and N the number of values to test per parameter. This quickly becomes intractable in order to achieve large precision or to estimate many parameters. Second, DA can provide interesting insight into both modeling and experimental error in the estimation process. Tracking covariances associated with the forecast and analysis steps can provide insight into both model uncertainty and experimental error which could be present. Furthermore, the ability of filters to estimate parameters online means that we can obtain time-dependent information about this uncertainty.

However, an issue appears in order to use Kalman filtering (KF) or smoothing (KS) for this application. As described in section 1.3, the operators F and H used in the forecast and analysis steps of the KF and KS must be *linear*. This poses no issue for the observation operator H here, as the radius is a dependent variable of our model, and thus H will simply be a one-to-one map of the radius to itself. To capture bubble dynamics in a viscoelastic material, however, our dynamical model will need to be nonlinear. Indeed, capturing these dynamics with a sufficient level of accuracy usually requires solving a system of at least 2 nonlinear ordinary differential equations (for bubble radius and pressure), and 2 partial differential equations (for temperature and vapor concentration in the bubble). These models will be described in Chapter 2, but the issue in the context of Kalman filtering is their nonlinearity, which means they cannot be represented by an operator F in matrix form. This means a different approach is required to update the Covariance matrix at each step. A simple approach to deal with weakly nonlinear systems is the so-called extended Kalman filter (Jazwinski, 2007; Gelb, 1974; Ghil et al., 1981). This method simply linearizes the dynamics about the mean, using the Jacobian matrix to represent F and calculate the Kalman gain. However, this method is inadequate for highly nonlinear dynamics, as is the case of our bubble dynamics system (Julier and Uhlmann, 2004).

Instead, we turn to a class of methods called ensemble methods, which are widely used in problems with high-dimensional systems and nonlinear dynamics. The driving idea behind ensemble methods is to represent the statistics of the state empirically by randomly sampling q realizations of the state from a given distribution. We thus create an ensemble, the statistics of which can be empirically determined. The mean and covariance of this ensemble are used as a surrogate measure of the state estimate and error covariance. This removes the need to update the covariance with a linear forecast operator, as the covariance is simply calculated empirically from the ensemble, each element of which can be updated with any dynamical model. The particular methods used in DA-IMR, which will be detailed in Chapter 2, are Ensemble Kalman methods, which are ensemble-based extensions of the Kalman filter and smoother described in section 1.3, making use of this empirical sampling (Evensen, 1994; Evensen, 2003; Evensen, 2004; Houtekamer and Derome, 1995; Houtekamer and Mitchell, 1998; Anderson and Anderson, 1999).

1.5 Burst-Wave Lithotripsy

The second part of this thesis, in Chapters 4 and 5, focuses on a different biomedical application: the treatment of kidney stone disease with lithotripsy. Renal calculi occur in millions of patients around the world every year, with the number of individuals affected in the United States rising above 10% of the population in the past decade (Scales et al., 2012; Hill et al., 2022). Various treatments exist to break and remove kidney and urinary stones. While surgical removal is effective, it presents complication risks, is expensive, and can be difficult depending on the stone location. Instead, an effective and safe non-invasive treatment can create better outcomes for patients. Extracorporeal Shock-wave lithotripsy (SWL) is a common, non-invasive treatment for kidney stone disease. In SWL stones are subject to multiple rounds of focused shock waves to break them into smaller pieces that can then be passed naturally (Chaussy, Brendel, and Schmiedt, 1980; Chaussy, Schmiedt, et al., 1982; Chaussy and Fuchs, 1989; Sackmann et al., 1988). SWL is applied non-invasively under local or general anesthesia to avoid pain for the patient, and $O(1000)$ shocks are delivered at a rate of approximately 1 Hz. Various studies have examined treatment parameters such as shock wave rate, lithotripter technology, and shock wave amplitude variation (Lingeman et al., 2009; Jendeborg et al., 2017).

While SWL avoids the need for surgery and is generally effective, the shock waves generated have very high pressure magnitudes, which can cause damage to surrounding tissue, internal bleeding, and other adverse effects in patients (Evan et al., 1998; McAteer and Evan, 2008; Bailey, Pishchalnikov, et al., 2005; Janetschek et al., 1997). Burst-wave lithotripsy (BWL) was developed as an alternative, to improve comminution and mitigate risks associated with SWL. Instead of shock waves, BWL uses repeated short bursts of ultrasound (Maxwell, Cunitz, et al., 2015). The frequency is $O(100)$ kHz, with pulses of 10 to 20 cycles of ultrasound occurring at a pulse repetition frequency (PRF) of 1 to 10 Hz. In BWL, peak positive pressures are greatly reduced compared to SWL, where pressure magnitudes can reach up to 100 MPa. During treatment, this peak pressure is generally limited so that negative pressure magnitudes need not exceed 7 MPa to minimize the risk of cavitation-induced injury. An added advantage is that BWL does not require anesthesia, as patients tolerate the procedure well (Harper, Metzler, et al., 2021). Figure 1.2 compares the SWL and BWL waveforms, highlighting the difference in pressure amplitude between the treatments.

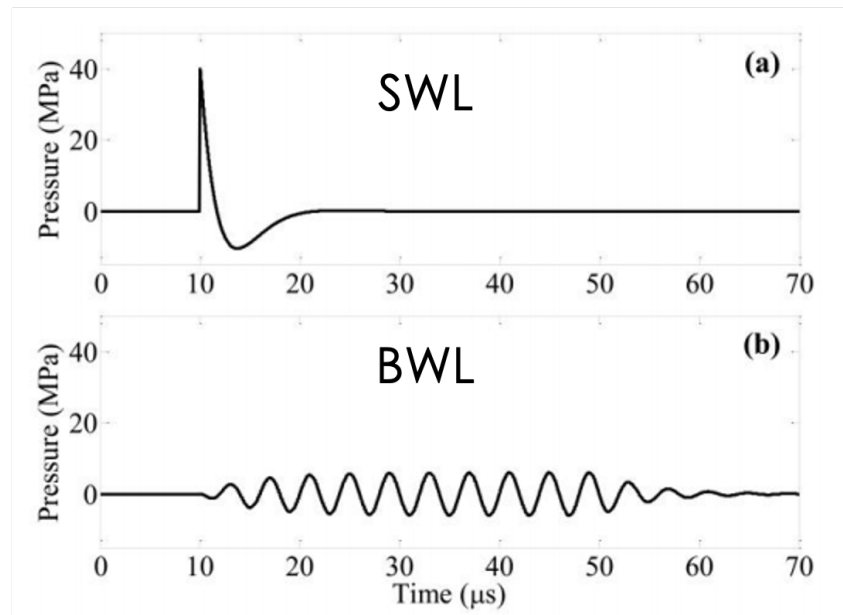


Figure 1.2: Comparison of typical SWL (top) and BWL (bottom) focal pressure waveforms. Reprinted from Maxwell, Cunitz, et al. (2015) with permission from Wolters Kluwer Health, © 2015.

Experimental studies of BWL, both in-vitro (Maxwell, Cunitz, et al., 2015) and in-vivo in a pig model (Maxwell, Wang, et al., 2019), have shown that it can break stones effectively. Maxwell, Cunitz, et al. (2015) fractured stones of various compositions, including stones known to be difficult to break with SWL, and compared treatment efficacy at various frequencies. Notably, they observed the comminution of stones into smaller fragments at higher frequencies. Compared to SWL, BWL thus leads to a more uniform fragment size that can be controlled by ultrasound frequency. A schematic of the experimental setup used by Maxwell, Cunitz, et al. (2015), taken from their paper, is shown in figure 1.3. Sapozhnikov, Maxwell, and Bailey (2021) and Bailey, Maxwell, et al. (2022) have shown that higher frequencies of ultrasound can enhance stone comminution for smaller calculi, further pointing to the fact that this frequency-dependence may be exploited to adapt BWL to stone size. More recently, clinical trials have shown promise in treating human patients (Harper, Metzler, et al., 2021). In a trial on 19 patients, Harper, Lingeman, et al. (2022) successfully fragmented 91 % of treated stones, achieving a median stone comminution of 90 % of stone volume.

These results show that BWL can be very effective. Studies have looked at the structure of elastic waves in stones during treatment experimentally (Maxwell, MacConaghy, et al., 2020), and developed models to predict maximum stress in




Figure 1.3: Schematic of the experimental setup for in-vitro BWL used by Maxwell, Cunitz, et al. (2015) and representative of the setup used for most experimental studies of BWL. Reprinted from Maxwell, Cunitz, et al. (2015) with permission from Wolters Kluwer Health, © 2015.

the stone in idealized settings (Sapozhnikov, Maxwell, and Bailey, 2021). However, as a novel treatment, many questions remain as to the mechanisms of stone fragmentation during BWL. A complicating factor is that high pressure magnitudes near the proximal surface of the stone can induce cavitation, resulting in bubble clouds forming between the transducer and the stone.

1.6 Bubble Formation during BWL

As mentioned, a remaining challenge of BWL is the propensity for cavitation to occur, particularly near the stone's proximal surface. In SWL it is generally accepted that cavitation, while undesirable from the point of view of injury, is important for complete stone comminution. While stress waves generated by the focused shock alone tend to result in large fragments, bubble-collapse-induced erosion produces much smaller ones. Inter-pulse times of 0.5 to 1s in SWL ensure that cavitation bubbles induced in each pulse are largely dissolved by the following one

Some approaches have tried to exploit this stone-erosion potential of bubble cloud collapse by designing a two-frequency waveform aimed at nucleating and controlling the collapse of bubble clouds on kidney stones (Ikeda et al., 2006; Yoshizawa et al., 2009; Matsumoto and Yoshizawa, 2005; Brujan and Matsumoto, 2012). This method, sometimes called cavitation control lithotripsy (CCL), employs an initial focused ultrasound pulse with very high frequency (usually 1 to 5 MHz) to create

a localized bubble cloud. A second pulse at or around 545 kHz trails the high frequency pulse, and forces the bubble cloud to collapse near the stone surface, delivering high pressures to erode the stone in the process. This approach exploits cloud cavitation to erode the stone. While the pressure magnitudes used for CCL are marginally higher than those of BWL — Yoshizawa et al. (2009) report peak positive focal pressure of 16 MPa — it is an efficient method to deliver very high localized pressure on the stone surface, which can reach $O(1 \text{ GPa})$ as the bubble cloud collapses (Shimada, Kobayashi, and Matsumoto, 1999).

By contrast to SWL and CCL, the resonating stress waves associated with the ultrasound wavelength appear to be sufficient to produce small enough fragments in burst-wave lithotripsy without relying on cavitation-induced damage (Maxwell, Cunitz, et al., 2015; Maxwell, MacConaghy, et al., 2020; Chen, Samson, et al., 2020; Raskolnikov, Bailey, and Harper, 2022). In developing treatment protocols, it is therefore generally sought to avoid cavitation through pausing treatment to allow bubbles to dissolve, or to utilize interleaved pulses of lower-amplitude ultrasound to induce bubble coalescence and dissolution (Duryea, Roberts, et al., 2014; Duryea, Cain, et al., 2015; Duryea, Tamaddoni, et al., 2015; Tamaddoni, Roberts, Duryea, et al., 2016; Tamaddoni, Roberts, and Hall, 2019). Apart from minimizing the risk of injury, minimizing cavitation is important in BWL because the short inter-pulse times do not allow the full dissolution of bubbles from pulse to pulse. The continual presence of bubbles on the proximal surface of the stone can then lead to a substantial shielding of the stone from the incident acoustic energy (Maeda, Maxwell, et al., 2018).

To better understand the mechanisms of stone comminution during BWL, a tool capable of simulating the response to BWL of an arbitrary stone, while modeling the dynamics of bubbles that may shield the stone during treatment, would be valuable.

1.7 Fluid-Solid Solver

To accurately simulate ultrasound–bubble cloud interactions, and cavitation of bubbles which may be present during treatment, a high-order accurate multi-phase compressible flow solver is required. To this end, we use an interface capturing method with the so-called 5-equation multiphase flow model of Kapila et al. (2001), solved with a WENO shock-capturing finite volume scheme. This is paired with

an HLL-type Riemann solver and total-variation diminishing time-stepping. The solver used, now open source, is the Multi-component Flow Code (MFC). Also implemented is a model to propagate spherical acoustic waves, as generated by the transducer during BWL. Previous numerical studies with this solver have examined the energy shielding effects of bubbles by using a subgrid bubble model and calculating the scattering of acoustic energy (Maeda, Maxwell, et al., 2018; Maeda and Colonius, 2018; Maeda and Colonius, 2019). However, these did not include a solid mechanics model or any way to determine the elastic response of the stone to the ultrasound.

Various approaches exist to pair a solid-mechanical model with a fluid dynamics solver. A difficulty in this case is that modeling the response of a solid generally requires tracking material deformation. However, the Eulerian framework used in MFC and other structured-grid CFD solvers does not track strains, and thus the stresses in the solids cannot be calculated using typical stress-strain relations. To address this problem and maintain a single-framework, monolithic solver, we adopt a hypoelastic material model following Rodriguez and Johnsen (2019). This approach is capable of modeling solid mechanics in a fully Eulerian framework by taking appropriate derivatives of constitutive stress-strain relations. To determine elastic stresses in this framework, we only require strain derivatives, which are straightforward to calculate in an Eulerian framework where velocity is tracked at each point in the domain. Thus, such a hypoelastic solver can naturally be implemented in a CFD solver such as MFC, without the need to switch the frame of reference used. Details on the hypoelastic approach and its implementation are discussed in Chapter 4, in section 4.1.1. There, we also detail the addition of a continuum damage model following Cao et al. (2019). This model makes use of the stresses obtained in the stone via the hypoelastic model to calculate a damage field, which tracks the accumulation of high stresses over time in the stone. Thus, the full interaction of ultrasound waves with gas bubbles and a submerged stone can be modeled to better understand mechanisms key to improving the efficacy of BWL.

1.8 Contributions and Outline

In this thesis, we first present the development of a data assimilation framework for the characterization of soft, viscoelastic materials via observation of bubble collapse (DA-IMR). This includes the implementation of a hybrid DA method

which is particularly effective for parameter estimation in this context. We show its successful application to data from various experiments using different constitutive models, demonstrating the accuracy and versatility of material property estimation with DA-IMR. Beyond parameter estimation, we gain insight into the limitations of the physical models used when violent bubble collapse occurs, particularly with high energy laser-induced cavitation. We then present the implementation of a comprehensive numerical framework for high-fidelity direct numerical simulations of acoustic wave–bubble–stone interactions. This framework is used to investigate bubble dynamics and stone damage in burst-wave lithotripsy. The reduction of stress and damage in stones during BWL when bubbles are present is quantified, and strategies to mitigate bubble shielding and improve treatment efficacy by modulating the frequency of individual transducer elements are proposed.

Chapter 2 details the development of the DA-IMR framework. A reduced-order spherical bubble dynamics model, and various ensemble-based data assimilation methods used in conjunction with this model are presented. Surrogate bubble radius time history data is generated via numerical simulations of inertial bubble collapse, and used to compare the performance of each data assimilation method. We conclude on the most accurate and efficient methods in the context of soft material characterization, and validate the DA-IMR framework with these DA methods.

In Chapter 3, this framework is tested on experimental cavitation data in three subsequent studies. We begin by applying DA-IMR on the laser-induced cavitation data of Estrada et al. (2018), which serves as a benchmark for our method with experimental data. Then, we apply DA-IMR to results with laser-induced cavitation (LIC) in gels with seed particles, which expands the stretch-ratio regime where DA-IMR can be used. The DA methods not only estimate the mechanical properties of the gel, but show the limitations of the physical model used for violently cavitating bubbles. Finally, the method is tested with ultrasound-induced cavitation data, which we show avoids some of the issues observed with bubble nucleation physics in LIC.

Next, in Chapter 4, we shift to direct numerical simulations of acoustic wave–bubble–stone interactions. This chapter describes the developed framework in detail, including all the implementations to make BWL simulations possible. We begin by describing the multiphase flow solver used for the interactions of acoustic waves with liquids and gas bubbles. Then, we present the hypoelastic model implemented into

this solver to simulate the elastic response of solids, and validate it against analytical and experimental results. Additionally, we present the addition of a continuum damage model, the calibration of a virtual transducer array to mimic experimental conditions with a much smaller simulated acoustic source, and the acceleration of our code using GPUs to enable simulations at high enough resolutions to capture all scales required in BWL simulations. Example simulations are performed to study the ultrasound frequency–dependence of maximum stress in stones of various shapes and sizes.

In Chapter 5, this numerical framework is used to study bubble dynamics and damage in the stone during BWL. We begin by quantifying the damage mitigation in the stone due to bubble shielding, showing the significant impact that bubble clouds can have on treatment efficacy. We then study the dynamics of characteristic bubble clouds when exposed to spherically focused ultrasound acoustic fields generated by a virtual BWL transducer array. We show that modulating the ultrasound frequency of individual transducer elements by introducing a secondary low-frequency wave can cause bubbles to collapse ahead of the stone. Applying this to full BWL simulations with bubble clouds present in front of a model kidney stone, we explore strategies to maximize the efficiency and efficacy of single BWL pulses in the presence of bubbles.

Finally, in Chapter 6, we summarize the conclusions of this thesis and provide guidance on potential future efforts to expand on the work presented.

Chapter 2

DATA ASSIMILATION FOR INERTIAL MICROCAVITATION RHEOMETRY

Part of this chapter is adapted from sections 1-4 of Spratt, Rodriguez, Schmidmayer, et al. (2021). We begin by describing the material–bubble-dynamic model used for inertial microcavitation rheometry throughout Chapter 2 and Chapter 3, with slight modifications to the viscoelastic material model where needed. Then, we present three data assimilation methods and details of their implementation for our problem. Finally, we compare each method using synthetic data, generated by adding random noise to results from our spherical bubble dynamics model. Thus, we can gauge the relative performance of each method in the absence of modeling uncertainties. We show that two of the tested methods can predict material parameters accurately, and are thus chosen as candidates to use with experimental data in Chapter 3.

2.1 Bubble Dynamics Model

A physical model for the dynamics of collapsing bubbles is required to characterize the viscoelastic properties of surrounding materials. Many spherical bubble dynamics models exist. Of particular relevance here are those for cavitation in soft materials (Gaudron, Warnez, and Johnsen, 2015; Yang and Church, 2005) and specific numerical methods for solving them (Warnez and Johnsen, 2015; Barajas and Johnsen, 2017). We use the model of Estrada et al. (2018), which adopts approximations validated in previous spherical-bubble models (Prosperetti and Lezzi, 1986; Prosperetti, Crum, and Commander, 1988; Akhatov et al., 2001; Epstein and Keller, 1972; Keller and Miksis, 1980; Preston, Colonius, and Brennen, 2007). Key assumptions of this model are that the motion of the bubble and its contents are spherically symmetric, the bubble pressure is spatially uniform (homobaricity), the temperature of the surrounding material is constant, and that there is no mass transfer of the non-condensable gas across the bubble wall. While the effects of vapor mass transfer may be negligible in the regime of interest to this chapter (Barajas and Johnsen, 2017), we opt to include these in the present model formulation as they are important in other regimes where this method may be applied (Preston, Colonius, and Brennen, 2007).

The Keller–Miksis equation models the radius evolution (Keller and Miksis, 1980),

$$\begin{aligned} \left(1 - \frac{\dot{R}}{c}\right) R \ddot{R} + \frac{3}{2} \left(1 - \frac{\dot{R}}{3c}\right) \dot{R}^2 = \frac{1}{\rho} \left(1 + \frac{\dot{R}}{c}\right) \left(p_b - \frac{2s}{R} + S - p_\infty\right) \\ + \frac{1}{\rho} \frac{R}{c} \overline{\left(p_b - \frac{2s}{R} + S\right)}, \end{aligned} \quad (2.1)$$

where R is the bubble radius, c the material speed of sound, ρ the material density, p_b the bubble internal pressure, s the bubble-wall surface tension, S the stress integral (see (2.7)), and p_∞ the far-field pressure. The Rayleigh-Plesset equation (Rayleigh, 1917; Plesset, 1948; Plesset, 1949; Plesset and Prosperetti, 1977) is often used for spherical bubble dynamics. However, it assumes incompressibility of the surrounding material, whereas the Keller-Miksis equation relaxes the incompressibility assumption near the bubble wall, taking into account the speed of sound of the surrounding material. These effects prove to be important in the context of IMR, though values for the Mach number are not expected to exceed $\dot{R}/c \approx 0.3$ (Akhatov et al., 2001), and thus the Keller-Miksis equation is used.

Under the model assumptions, no mass or energy conservation equations are needed outside the bubble. Furthermore, the conservation of momentum simplifies to an ordinary differential equation for the bubble pressure

$$\dot{p}_b = \frac{3}{R} \left[-\kappa p_b \dot{R} + (\kappa - 1) K(T(R)) \frac{\partial T}{\partial r} \Big|_{r=R} + \kappa p_b \frac{C_{p,v}}{C_p(k_v(R))} \frac{D}{1 - k_v(R)} \frac{\partial k_v}{\partial r} \Big|_{r=R} \right], \quad (2.2)$$

where κ is the specific heat ratio of vapor, K the thermal conductivity, T the gas temperature, C_p the specific heat, D the binary diffusion coefficient, and k_v the vapor mass fraction. Subscripts g , v , and m refer to gas, vapor, and mixture properties, respectively. Applying conservation of energy in the bubble interior yields an equation for the bubble temperature:

$$\rho_m C_p \left(\frac{\partial T}{\partial t} + v_m \frac{\partial T}{\partial r} \right) = \dot{p}_b + \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 K \frac{\partial T}{\partial r} \right) + \rho_m (C_{p,v} - C_{p,g}) D \frac{\partial k_v}{\partial r} \frac{\partial T}{\partial r}, \quad (2.3)$$

where v_m is the radial mixture velocity and ρ_m the mixture density. The boundary condition $T(R) = T_\infty$ follows from the model assumptions. The radial mixture velocities are computed as

$$v_m(r, t) = \frac{1}{\kappa p_b} \left[(\kappa - 1) K \frac{\partial T}{\partial R} - \frac{1}{3} r \dot{p}_b \right] + \frac{C_{p,v} - C_{p,g}}{C_{p,m}} D \frac{\partial k}{\partial r}, \quad (2.4)$$

with associated kinematic boundary condition

$$v_m(R) = \dot{R} - \frac{D}{1 - k_v(R)} \left. \frac{\partial k_v}{\partial r} \right|_{r=R}. \quad (2.5)$$

Fick's law describes the mass diffusion process in the bubble. Casting the conservation of mass inside the bubble in terms of the mixture density, the vapor mass fraction inside the bubble is

$$\frac{\partial k_v}{\partial t} + v_m \frac{\partial k_v}{\partial r} = \frac{1}{\rho_m} \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \rho_m D \frac{\partial k_v}{\partial r} \right). \quad (2.6)$$

Under the assumption of equilibrium phase change at the bubble wall, the associated boundary condition at the wall is $p_{v,sat}(T(R)) = \mathcal{R}_v k(R) \rho_m(R) T(R)$, where $p_{v,sat}$ is the saturation pressure of the vapor and \mathcal{R}_v is the gas constant of the vapor.

Equations (2.1), (2.2), (2.3), and (2.6) form a system of equations, which is evolved in time with an implicit Runge–Kutta algorithm that uses the trapezoidal rule and backwards differentiation at each step (TR–BDF2) (Hosea and Shampine, 1996; Shampine and Reichelt, 1997; Shampine, Reichelt, and Kierzenka, 1999). The partial differential equations for temperature and vapor mass fraction are discretized in space via a uniform grid and computed using second-order-accurate central finite differences. Estrada et al. (2018) showed that the finite-deformation neo-Hookean Kelvin–Voigt model can adequately represent the material response at high strain rates in problems of interest, and that a more complicated model did not improve the goodness of fit of simulated radius curves with experimental data. In this framework, the material is modeled with a parallel spring (neo-Hookean elastic response with shear modulus G) and dashpot (linear viscous response with viscosity μ). The stress integral in (2.1) is

$$S = -\frac{G}{2} \left[5 - \left(\frac{R_0}{R} \right)^4 - 4 \frac{R_0}{R} \right] - \frac{4\mu\dot{R}}{R}, \quad (2.7)$$

where R_0 is the equilibrium bubble radius (Gaudron, Warnez, and Johnsen, 2015). Details on the derivation of this spherical bubble dynamics model, and comparison with the model of Preston, Colonius, and Brennen (2007), are shown in Appendix A.

2.2 Data Assimilation Methods

Two difficulties that drive the choice of data assimilation method are the nonlinearity of the dynamics and large state vector required to discretize the partial differential

equations adequately. The former rules out the standard linearized Kalman filter (EKF) (Kalman, 1960) and the latter renders its direct nonlinear extensions (e.g., the unscented Kalman filter, UKF) computationally prohibitive. Instead, ensemble-based methods (Evensen, 1994) are considered. They combine computational efficiency with nonlinear dynamics by approximating the state covariance via statistics of a finite (and typically small) ensemble. We consider three specific ensemble methods: an ensemble Kalman filter (EnKF), an iterative ensemble Kalman smoother (IEnKS), and a hybrid ensemble-based 4D–Var method.

The discretized equations of section 2.1 are rewritten as a nonlinear operator F , and we define the linear observation function H that maps the state \mathbf{x} to measurement space. This yields the discrete-time dynamical system

$$\mathbf{x}_{k+1} = F(\mathbf{x}_k) + \boldsymbol{\eta}_k, \quad (2.8)$$

$$\mathbf{y}_k = H(\mathbf{x}_k) + \mathbf{v}_k, \quad (2.9)$$

where

$$\mathbf{x}_k \in \mathbb{R}^d, \mathbf{y}_k \in \mathbb{R}^n,$$

$$\boldsymbol{\eta}_k \sim \mathcal{N}(0, \Sigma), \mathbf{v}_k \sim \mathcal{N}(0, \Gamma),$$

$$F : \mathbb{R}^d \rightarrow \mathbb{R}^d, H : \mathbb{R}^d \rightarrow \mathbb{R}^n.$$

\mathbf{x}_k is the d -dimensional state at time k comprised of all the dependent variables plus the unknown parameters

$$\mathbf{x} = \{R, \dot{R}, p_b, S, \mathbf{T}, \mathbf{C}, \log(\text{Ca}), \log(\text{Re})\}, \quad (2.10)$$

which are the bubble-wall radius, velocity, bubble pressure, stress integral, the discretized temperature and vapor concentration fields inside the bubble, and the log-Cauchy and log-Reynolds numbers, respectively. The Cauchy and Reynolds numbers are here defined as

$$\text{Ca} \equiv \frac{p_\infty}{G}, \quad \text{Re} \equiv \frac{\sqrt{\rho p_\infty} R_{\max}}{\mu}. \quad (2.11)$$

These quantities appear in the nondimensionalized model equations of section 2.1 for the surrounding material and the shear modulus G and viscosity μ , our material properties of interest, can be computed via (2.11). When the state vector is stepped forward in time with equation (2.8), the forecast operator F simply maps $\log(\text{Ca})$ and $\log(\text{Re})$ to themselves, as they are constant in the physical model. Appending

them to the state vector enables the estimation of these parameters within the data assimilation framework in a straightforward way. We note that the logarithm of these parameters is taken to avoid negative (and thus non-physical) values of the Cauchy and Reynolds numbers (Ca and Re) occurring during the analysis step of the assimilation algorithms (described in sections 2.2.1, 2.2.2, and 2.2.4).

The variable \mathbf{y}_k is the n -dimensional observation (data) at time k . $\boldsymbol{\eta}_k$ is the unknown process noise (or model error) added to $H(\mathbf{x}_k)$ to retrieve \mathbf{y}_k . It is assumed to be Gaussian with zero mean and standard deviation Σ . Similarly, \mathbf{v}_k is the assumed Gaussian measurement noise added to $F(\mathbf{x}_k)$ to obtain \mathbf{x}_{k+1} , with zero mean and unknown standard deviation Γ . Throughout this study, the only available measurement is the bubble radius. This means that \mathbf{y}_k is comprised of a single element: the bubble radius. The observation operator H is then the linear map from the state vector to its first element R . In the following, the linear operator H is sometimes represented as the matrix \mathbf{H} for clarity ($\mathbf{H}\mathbf{x} = H(\mathbf{x})$).

The following methods estimate the full state vector \mathbf{x} (including parameters of interest $\log(\text{Ca})$ and $\log(\text{Re})$) based on observations of \mathbf{y} . The EnKF and IEnKS are online (or quasi-online) methods—they optimize the value of \mathbf{x} at each time through the simulation. The IEnKS is deemed quasi-online because it uses data from a few future times as well. The estimation trails the simulation time by a fixed number of time steps called the lag. Alternatively, the En4D–Var is an offline method, which only optimizes the initial condition for \mathbf{x} , taking into account data from the entire time domain.

2.2.1 The Ensemble Kalman Filter

The ensemble Kalman filter (Evensen, 1994) represents the probability density function (PDF) for the state of the dynamics through the statistics of an ensemble of q state vectors. It does not require an adjoint, or deriving a tangent linear operator to the physical model (Evensen, 2003; Evensen, 2009a; Evensen, 2009b). Starting with suitably randomized initial conditions, each ensemble member is propagated through the physical model, and the predictions are then corrected using the ensemble statistics. The ensemble is initialized with a guess for the initial condition \mathbf{x}_0 as the mean, and a given covariance corresponding to the expected error covariance. In practice, each ensemble member is independently sampled from a normal distribution with mean \mathbf{x}_0 and the assumed covariance matrix. Several

initialization strategies exist depending on the system and its dynamics. In the present case, the nonlinear dynamics render a systematic approach difficult. Instead, the ensemble is initialized with a covariance corresponding to our best estimate based on simulations, and adjusted through trial and error to optimize results. An ensemble size of $q = 48$ is used for the tests. This has been shown to give accurate results while keeping computational costs modest. Indeed, in testing with the current framework, larger ensemble sizes led to negligible improvement in parameter estimation but significantly increased the cost of tested DA methods. At any given time, the estimated value for the state vector is then taken to be the ensemble average.

$$\bar{\mathbf{x}}_k = \frac{1}{q} \sum_{j=1}^q \mathbf{x}_k^{(j)}. \quad (2.12)$$

The filter is broken down into a forecast and an analysis step. In the forecast step, the physical model is used to step the state forward in time with (2.8). Each representation of the state vector $\mathbf{x}_k^{(j)}$ in the ensemble at time k is propagated through F with $\hat{\mathbf{x}}_{k+1}^{(j)} = F(\mathbf{x}_k^{(j)})$. Next, in the analysis step, if an experimental measurement \mathbf{y}_{k+1} is available at the current time step $k+1$, then it is used to correct the forecast. As described in the dynamical system equations, each ensemble member is mapped to measurement space $H(\mathbf{x}_{k+1})$. The analysis proceeds by minimizing a cost function involving the difference between $H(\mathbf{x})$ and the data point \mathbf{y} , while accounting for measurement noise and model error. This cost function is similar to that of the Kalman filter and is given by

$$J(\mathbf{x}) = \frac{1}{2} \|\mathbf{y}_k - H(\mathbf{x})\|_{\mathbf{R}}^2 + \frac{1}{2} \|\mathbf{x} - \hat{\mathbf{x}}_k\|_{\mathbf{C}_k}^2 \quad (2.13)$$

$$= \frac{1}{2} [\mathbf{y}_k - H(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y}_k - H(\mathbf{x})] + \frac{1}{2} [\mathbf{x} - \hat{\mathbf{x}}_k]^T \mathbf{C}_k^{-1} [\mathbf{x} - \hat{\mathbf{x}}_k], \quad (2.14)$$

where \mathbf{R} is the measurement noise covariance matrix, which is an input to the algorithm, and \mathbf{C}_k is the ensemble covariance at time step k . A key difference with the Kalman filter is that this covariance is calculated empirically in this case. It is defined as

$$\mathbf{C}_k = \mathbf{A}_k (\mathbf{A}_k)^T, \quad (2.15)$$

where \mathbf{A}_k is the state perturbation matrix

$$\mathbf{A}_k = \frac{1}{\sqrt{q-1}} \left[\mathbf{x}_k^{(1)} - \bar{\mathbf{x}}_k, \dots, \mathbf{x}_k^{(q)} - \bar{\mathbf{x}}_k \right]. \quad (2.16)$$

In fact, the minimization does not make use of the covariance matrix directly, but instead uses the state perturbation matrix and scaled output perturbation matrix $\mathbf{H}\mathbf{A}_k$ defined as

$$\mathbf{H}\mathbf{A}_k = \frac{1}{\sqrt{q-1}} \left[\mathbf{y}_k^{(1)} - \bar{\mathbf{y}}_k, \dots, \mathbf{y}_k^{(q)} - \bar{\mathbf{y}}_k \right]. \quad (2.17)$$

The optimization is carried out by finding the minimizer \mathbf{x}_k satisfying

$$\mathbf{x}_k = \hat{\mathbf{x}}_k + \mathbf{A}_k \cdot \mathbf{w}_k, \quad (2.18)$$

with \mathbf{w}_k a correction coefficient. This restricts the solution to the subspace spanned by the scaled perturbation matrix around the prior estimate $\hat{\mathbf{x}}_k$. The optimization can be restated as

$$\mathbf{w}_k = \underset{\mathbf{w} \in \mathbb{R}^q}{\operatorname{argmin}} J(\mathbf{w}), \quad (2.19)$$

where

$$J(\mathbf{w}) = \frac{1}{2} \|\mathbf{w}\|^2 + \frac{1}{2} \|\mathbf{y}_k - H(\hat{\mathbf{x}}_k) - \mathbf{H}\mathbf{A}_k(\mathbf{w})\|_{\mathbf{R}}^2. \quad (2.20)$$

The solution is unique, and using the Woodbury matrix identity to write the inversion in measurement space, can be written as

$$\mathbf{w}_k = (\mathbf{H}\mathbf{A}_k)^{\top} [\mathbf{R} + (\mathbf{H}\mathbf{A}_k)(\mathbf{H}\mathbf{A}_k)^{\top}]^{-1} (\mathbf{y}_k - H(\hat{\mathbf{x}}_k)). \quad (2.21)$$

Performing this inversion in the measurement space is in most cases more computationally efficient. Here this is clear, as the measurement space is comprised of only one variable (the bubble radius). Once the minimizer is found and the analysis step complete, covariance inflation is applied to the ensemble to correct for the (typical) underestimation of the variance with finite (typically small) ensembles (see section 2.2.3 for details on covariance inflation). Finally, the forecast step can be repeated. Figure 2.1 shows a flow chart of the EnKF method for a visual representation of the steps described above. This flow chart will remain the same for the IEnKS method described in section 2.2.2, with changes to the analysis step only.

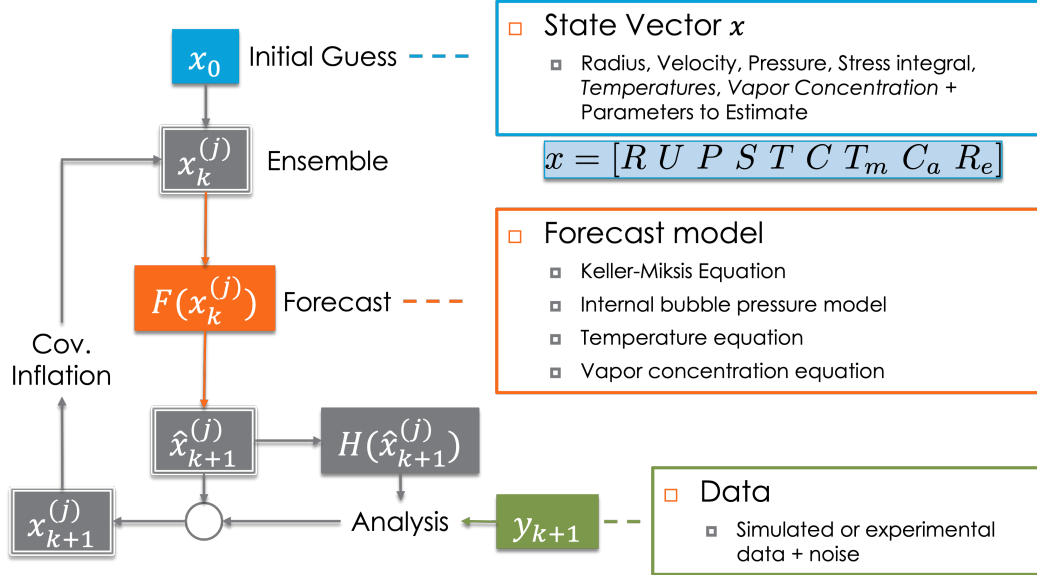


Figure 2.1: Flowchart of the Ensemble Kalman Filter as applied in the framework of inertial microcavitation rheometry.

2.2.2 The Iterative Ensemble Kalman Smoother

Minimizing deviation from data not just at the current estimation step, but at future times, can help to smooth out estimation and focus on longer-term trends. The IEnKS uses information from one or multiple future time steps in its assimilation, and can thus be an effective tool. While the ensemble initialization and forecast step are the same as that of the EnKF, the difference in the analysis step is twofold. First, the cost function is modified to minimize difference with data at a single or multiple future times (Evensen and Leeuwen, 2000). The assimilation thus trails the simulation by a number of time steps (called the lag). Second, it is no longer minimized analytically but iteratively using a Gauss–Newton algorithm.

The IEnKS method used here is from Bocquet and Sakov (2013a) and Sakov, Oliver, and Bertino (2012). Bocquet and Sakov (2013b) have shown it to be effective for state and parameter estimation problems with highly nonlinear dynamics. Its cost function can take two forms referred to as ‘single data assimilation’ (SDA) or ‘multiple data assimilation’ (MDA) (Bocquet and Sakov, 2013a). The IEnKS–SDA cost function penalizes difference with measurements at a single time step $k + L$, where L corresponds to the lag of the smoother. It is given by

$$J(\mathbf{x}) = \frac{1}{2} \|\mathbf{y}_{k+L} - H \circ F_{k \rightarrow (k+L)}(\mathbf{x})\|_{\mathbf{R}}^2 + \frac{1}{2} \|\mathbf{x} - \hat{\mathbf{x}}_k\|_{\mathbf{C}_k}^2, \quad (2.22)$$

where $H \circ F$ denoted the composition of H and F .

On the other hand, the IEnKS–MDA cost function minimizes this difference over a data assimilation window (DAW) from $k + 1$ to $k + L$, and is expressed as

$$J(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^L \beta_i \| \mathbf{y}_{k+i} - H \circ F_{k \rightarrow (k+i)}(\mathbf{x}) \|_{\mathbf{R}}^2 + \frac{1}{2} \| \mathbf{x} - \hat{\mathbf{x}}_k \|_{\mathbf{C}_k}^2, \quad (2.23)$$

where β_i are weights attributed to given time steps with $\sum \beta_i = 1$. Again, a solution of the form $\mathbf{x} = \hat{\mathbf{x}}_k + \mathbf{A}_k \cdot \mathbf{w}$ is sought, but a Gauss–Newton method is used (Bocquet and Sakov, 2013a). The minimizer \mathbf{w} is found by iterating following

$$\mathbf{w}_{(i+1)} = \mathbf{w}_{(i)} - \mathcal{H}_{(i)}^{-1} \Delta J_{(i)}(\mathbf{w}_{(i)}), \quad (2.24)$$

where i is the iteration number, and \mathcal{H} is the approximate Hessian

$$\mathcal{H}_{(j)} = (q - 1)\mathbf{I} + \mathbf{H}\mathbf{A}_{(j)}^{\mathbf{T}}\mathbf{R}^{-1}\mathbf{H}\mathbf{A}_{(j)}, \quad (2.25)$$

where \mathbf{I} is the $q \times q$ identity matrix. The gradient is given by

$$\Delta J_{(j)} = -\mathbf{H}\mathbf{A}_{(j)}^{\mathbf{T}}\mathbf{R}^{-1}[\mathbf{y}_{k+L} - H \circ F_{k+L \leftarrow k}(\mathbf{x}_k)] + (q - 1)\mathbf{w}_{(j)}. \quad (2.26)$$

For the smoother, $\mathbf{H}\mathbf{A}$ is more complicated than it is for the filter as it involves differences with measurements at future time steps. This quantity is akin to a tangent linear operator from ensemble to measurement space and has to be estimated. Following Bocquet and Sakov (2013a), a finite-difference estimate is used:

$$\mathbf{H}\mathbf{A}_{(j)} \approx \frac{1}{\alpha} H \circ F_{k+L \leftarrow k}(\mathbf{x}_k^{(j)} \mathbf{1}^{\mathbf{T}} + \alpha \mathbf{A}_k) \left(\mathbf{I} - \frac{\mathbf{1} \cdot \mathbf{1}^{\mathbf{T}}}{q} \right), \quad (2.27)$$

with scaling factor $\alpha \ll 1$ and $\mathbf{1} = (1 \ \dots \ 1)^{\mathbf{T}}$ a vector of length q . The iteration is repeated until a threshold $\mathbf{w}_{(i+1)} - \mathbf{w}_{(i)} < \epsilon$, or a fixed number of iterations is reached. Once the optimal value \mathbf{w}_{opt} is obtained, $\mathbf{x}_{\text{opt}} = \hat{\mathbf{x}} + \mathbf{A}_k \cdot \mathbf{w}_{\text{opt}}$ is calculated and a new ensemble \mathbf{E}_k is sampled at time step k with

$$\mathbf{E}_k = \mathbf{x}_{\text{opt}} \mathbf{1}^{\mathbf{T}} + \sqrt{q - 1} \mathbf{A}_k \mathcal{H}_{\text{opt}}^{-1/2} \mathbf{I}. \quad (2.28)$$

This completes the analysis step. When using the MDA variant, the Hessian and gradient of J are found with

$$\mathcal{H}_{(j)} = (q - 1)\mathbf{I} + \sum_{i=1}^L \mathbf{H}\mathbf{A}_i^{\mathbf{T}} \beta_i \mathbf{R}^{-1} \mathbf{H}\mathbf{A}_i \quad (2.29)$$

$$\Delta J_{(j)} = - \sum_{i=1}^L \mathbf{H}\mathbf{A}_i^{\mathbf{T}} \beta_i \mathbf{R}^{-1} [\mathbf{y}_{k+i} - H \circ F_{k+i \leftarrow k}(\mathbf{x}_i)] + (q - 1)\mathbf{w}_{(j)}. \quad (2.30)$$

2.2.3 Covariance Inflation

While the EnKF and IEnKS may converge, ensemble methods are subject to intrinsic sampling error (Bocquet, 2011; Luo and Hoteit, 2011). This sampling error results from the finite ensemble size q used to represent the statistics of a system of often much higher dimension. As Leeuwen (1999) explains, the EnKF tends to underestimate error variances, particularly for small ensemble sizes. There exist different ways to address this sampling error, but a simple approach is covariance inflation (Whitaker and Hamill, 2012), where we correct

$$\mathbf{x}^{(j)} = \bar{\mathbf{x}} + \alpha(\mathbf{x}^{(j)} - \bar{\mathbf{x}}) + \boldsymbol{\lambda}^{(j)}. \quad (2.31)$$

Here, $\bar{\mathbf{x}}$ denotes the ensemble average, as defined in (2.12), after the analysis step. Parameters α and $\boldsymbol{\lambda}$ correspond to multiplicative and additive inflation parameters, respectively.

There exist many schemes for multiplicative inflation, the most simple of which is picking a scalar α (usually $1.005 \leq \alpha \leq 1.05$). This can work well but requires extensive tuning to optimize the value for each run or data set. Instead, Whitaker and Hamill (2012) propose a scheme they call ‘Relaxation to Prior Spread’ (RTPS). Here, the value for α is found at each time step using

$$\alpha_i = 1 + \theta \left(\frac{\sigma_i^b - \sigma_i^a}{\sigma_i^a} \right), \quad (2.32)$$

where σ_i^a and σ_i^b are the prior and posterior ensemble standard deviation for the i^{th} element of the state vector (α is a vector here), and θ is a scalar (usually $0.5 \leq \theta \leq 0.95$). As this expression for α shows, this scheme inflates the covariance more in regions where the analysis led to a large correction. Whitaker and Hamill (2012) test this method and compare it to other approaches, showing that it performs well by adequately preventing underestimation of error variances and better converging on accurate estimates. We tested this covariance inflation scheme, and similarly found that it performs as well or better than a simple scalar α for data assimilation in our context. After some tuning across data sets, this RTPS model was used with $\theta = 0.7$. Additive covariance inflation was not found to significantly affect results and introduced some stability issues with larger magnitudes of $\boldsymbol{\lambda}$. Therefore, $\boldsymbol{\lambda} = \mathbf{0}$ is used.

2.2.4 A hybrid Ensemble-based 4D-var Method

As with ensemble Kalman methods, ensembles can be used with 4D-Var to estimate covariance empirically, thus reducing computational cost (Gustafsson and Bojarova, 2014; Liu, Xiao, and Wang, 2008; Caya, Sun, and Snyder, 2005; Trémolet, 2007). The present method (En4D-Var) is a fully offline extension of the IEnKS-MDA method. Again, the ensemble is initialized in the same way as EnKF, but the cost function is here

$$J(\mathbf{x}) = \frac{1}{2} \sum_k \beta_k \|y_k - H \circ F_{k \leftarrow 0}(\mathbf{x})\|_{\mathbf{R}}^2 + \frac{1}{2} \|\mathbf{x} - \hat{\mathbf{x}}_k\|_{\mathbf{C}_0}^2. \quad (2.33)$$

The difference with the IEnKS-MDA cost function is the data assimilation window size. Rather than minimizing over a few time steps forward and then stepping through time, the minimization is done over the entire time domain and only the initial state vector is corrected. Each new iteration is initialized with the corrected initial state (including parameters to estimate). The same minimization procedure as described in section 2.2.2 is used. When the minimization has converged, a final simulation is run with the forecast model only. In cases where only a few iterations are necessary, this method reduces computational cost as compared to the IEnKS-MDA. The time dimension is still fully included, but each point in time is only assimilated once per iteration. Furthermore, this retains the advantage of ensemble methods. As opposed to classical 4D-Var, there is no need to linearize the state function and find the tangent linear adjoint operator. This novel adaptation of the IEnKS method is well suited to the present problems given that our interest is the estimation of material properties which are, at the outset, assumed to be constant.

Whether online or offline, the structure of all the presented methods provides a significant computational advantage, as compared to the least squares fitting used by Estrada et al. (2018). For simple least squares fitting, a large array of cases with varying G and μ need to be simulated. This requires $N_G \times N_\mu$ simulations, where N_G and N_μ are the number of shear moduli and viscosities to test. To obtain precise estimates (to several decimal points), these must be large. On the other hand, only an initial guess for these parameters is needed in the data assimilation methods presented. Here, the cost will increase with increasing ensemble size q (at different rates depending on the method), as each ensemble member is stepped through time independently. However, a small value of $q = 48$ is shown to suffice for the present estimation, resulting in low computational time while providing good parameter estimates for most methods, as shown in sections 2.3.1 and 3.1.2. This advantage

becomes even more apparent if the method is scaled to estimate more parameters. For least squares fitting, the number of simulations to run would be multiplied by the number of values N to test for each new parameter, quickly becoming unfeasible for precise estimation. On the other hand, these data assimilation methods would incur a minimal added computational cost associated with the added state variable. In this chapter, we only estimate two parameters given our constitutive material model, namely G and μ , found through estimation of $\log(\text{Ca})$ and $\log(\text{Re})$. However, we will see examples in Chapter 3 where we estimate up to 4 material parameters given a different viscoelastic model. There, the number of simulations to run with simple least-squares fitting would be $\prod_{i=1}^4 N_i$, with each N_i the number of guesses for each parameter i to estimate. In contrast, the added computational cost for any of the data assimilation methods presented is negligible, as we are simply adding a single additional element to the state vector which has $O(100)$ elements.

2.3 Framework Validation with Surrogate Data

2.3.1 Parameter Estimation Results

Synthetic data where the true shear modulus and viscosity are known is generated from the model (section 2.1) and used to test the data assimilation methods in a setting where there is no modeling error. Bubble radius time-history data from the simulation is sampled every $3.7 \mu\text{s}$ to match the 270,000 frames per second image capture rate in available experiments. Random Gaussian noise is added to these samples to mimic experimental data. The standard deviation of this noise is set at $\sigma = 0.02$, which is greater than the estimated noise in the experiments. Two polyacrylamide gels were examined with nominal values of shear modulus and viscosity determined by Estrada et al. (2018). For the stiff gel: $G_{\text{stiff}} = 7.69 \text{ kPa}$, $\mu_{\text{stiff}} = 0.101 \text{ Pa s}$, and for the soft gel: $G_{\text{soft}} = 2.12 \text{ kPa}$, $\mu_{\text{soft}} = 0.118 \text{ Pa s}$. Since similar estimation accuracy was achieved in both cases, we report results for the stiff gel only. The other material properties used are taken from Estrada et al. (2018) and given in table 2.1. No uncertainty is added to these parameters in the present study to match their conditions and focus on estimating G and μ .

An example simulated radius curve and sampled surrogate measurements (with noise added) with these parameters is shown in figure 2.2a, plotted against non-dimensional time

$$t^* = \frac{t}{R_{\text{max}}} \sqrt{\frac{p_{\infty}}{\rho}}. \quad (2.34)$$

Parameter	Value	Parameter	Value
ρ	1060 kg/m ³	c	1430 m/s
p_∞	101.3 kPa	s	5.6×10^{-2} N/m
D	24.2×10^{-6} m ² /s	κ	1.4
$C_{p,g}$	1.62 kJ/kg K	$C_{p,v}$	1.00 kJ/kg K
A	5.3×10^{-5} W/m K ²	B	1.17×10^{-2} W/m K ²
p_{ref}	1.17×10^8 kPa	T_{ref}	5200 K
T_∞	298.15 K		

Table 2.1: Model parameters as they follow from Estrada et al. (2018). Table reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.

With the simulated data, the evolution over time of all variables in the state vector is known. For example, bubble-wall velocity, bubble pressure and stress integral are plotted in figure 2.2b.




Figure 2.2: Simulated bubble radius and noisy sampled data used to test data assimilation methods (a), alongside simulated bubble-wall velocity, normalized bubble pressure and stress integral (b), plotted over non-dimensional time t^* . Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.

A set of initial guesses for the shear modulus and viscosity, ranging from 10% to 100% error from the true values, were used to test each method. Table 2.2 summarizes results for a subset of these cases, representing 10, 50, and 100% initial error in G and μ . In each case, ensembles were initialized as Gaussian with these erroneous material properties as the mean, and standard deviation increasing with increased error. That is, the spread of the initial ensemble was made wider for cases with more error, to account for the increased uncertainty in the initial guess.

To match the experimental data (as used in section 3.1 in Chapter 3), simulated data is limited to the first three peaks of the bubble collapse. This corresponds to approximately 35 points given the initial conditions and frame rate. Estrada et al. (2018) found that limiting the data to this region led to better parameter estimation. Similarly, we find that the model fails to fit the radius measurements after this time. Reasons for this reduced accuracy at later times are discussed in results with experimental data in Chapter 3.

Method	G estimate (%error) [kPa]	μ estimate (%error) [Pa s]	Run time [s]
Guess 1	8.50 (+10%)	0.09 (-10%)	–
EnKF	7.234 (5.93%)	0.098 (2.61%)	428
IEnKS–SDA (lag 1)	7.364 (4.24%)	0.110 (8.58%)	852
IEnKS–MDA (lag 3)	6.682 (13.11%)	0.100 (0.92%)	8076
En4D–Var	7.150 (7.03%)	0.099 (1.80%)	679
Guess 2	3.80 (-50%)	0.05 (-50%)	–
EnKF	3.988 (48.1%)	0.057 (43.1%)	375
IEnKS–SDA (lag 1)	4.203 (45.4%)	0.080 (20.9%)	904
IEnKS–MDA (lag 3)	7.390 (3.90%)	0.086 (15.2%)	9755
En4D–Var	7.396 (3.82%)	0.100 (0.52%)	690
Guess 3	15.0 (+100%)	0.20 (+100%)	–
EnKF	13.649 (77.5%)	0.175 (73.1%)	495
IEnKS–SDA (lag 1)	10.272 (33.6%)	0.142 (40.7%)	800
IEnKS–MDA (lag 3)	10.078 (31.1%)	0.121 (19.9%)	9802
En4D–Var	10.210 (32.7%)	0.118 (16.6%)	611

Table 2.2: Comparing the accuracy of estimation and run time with 3 different initial guesses for the material parameters G and μ . Runs were performed on a computer with dual 12-core 2.3 GHz processors. Table reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.

Table 2.2 shows that with a relatively good initial guess with 10% error, the assimilation methods perform adequately. For example, the EnKF tracks the correct values for shear modulus and viscosity within 6% and 3%, respectively. With a moderate initial error of 50%, however, the EnKF loses accuracy and barely improves on the initial guess. In some cases, the EnKF was observed to be unstable, and the initial ensemble covariance had to be limited to prevent divergence. This limited the ability of the filter to estimate the parameters of interest, and thus despite its computational efficiency, the EnKF is eliminated from further consideration for this application.

The IEnKS and En4D–Var performed better than the EnKF for the 50% error case. The estimation was stable while varying initial conditions and covariance. Still, the lag 1 IEnKS–SDA only resulted in marginal improvements in the parameter values. The lag 3 IEnKS–MDA, on the other hand, resulted in further improvement, but at a high computational cost. This cost is associated with the calculation of the Hessian (see equation (2.29)) and gradient of the cost function (see equation (2.30)), which now involves three future time steps. The En4D–Var performs best in this test case, achieving good estimation with a comparably fast computational time. We note that while the En4D–Var was run for fifteen iterations in each case, the material property estimation converged by the fifth iteration. Thus, results and run-time after five iterations are reported.

Estimation results in the case with 50% error are presented in figure 2.3. Figure 2.3d shows the suitability of the En4D–Var: both parameters converge to accurate estimates within a few iterations. Overall, figure 2.3 also highlights the value of looking over a time horizon. While the EnKF and lag 1 IEnKS appear to disbelieve the data too much throughout the run, taking into account multiple times enables the lag 3 IEnKS–MDA to adjust to new information well, notably around collapse. Additionally, by nature, the IEnKS–MDA assimilates over the same data points multiple times. Data at any given time-step k , including near collapse point, will be used to minimize the cost functions at three separate times $k - 4$ to $k - 1$ for a lag 3 IEnKS–MDA. This redundancy appears important to parameter estimation in this context and given this data capture rate. Indeed, the IEnKS–MDA significantly corrects the viscosity estimate around each collapse, and the shear modulus estimate during the second collapse. Assimilating data from single time-steps appears to be insufficient given the short time scales of bubble cavitation and limited data. Smoothing over multiple times far improves performance around collapse points, which, given the IEnKS–MDA results, appear to hold the most pertinent data to make the necessary corrections.

In the case with 100% error in the initial guess, the relative performances of each method are similar to the 50% error case, but the three smoothers stagnate at 20 to 40% errors for μ and G . Weighted by computational expense, the En4D–Var performs best, but the IEnKS–MDA should not be discarded. Indeed, the time-varying estimation provides additional information about potentially time-dependent modeling uncertainties. While the physical model used assumes a constant shear modulus

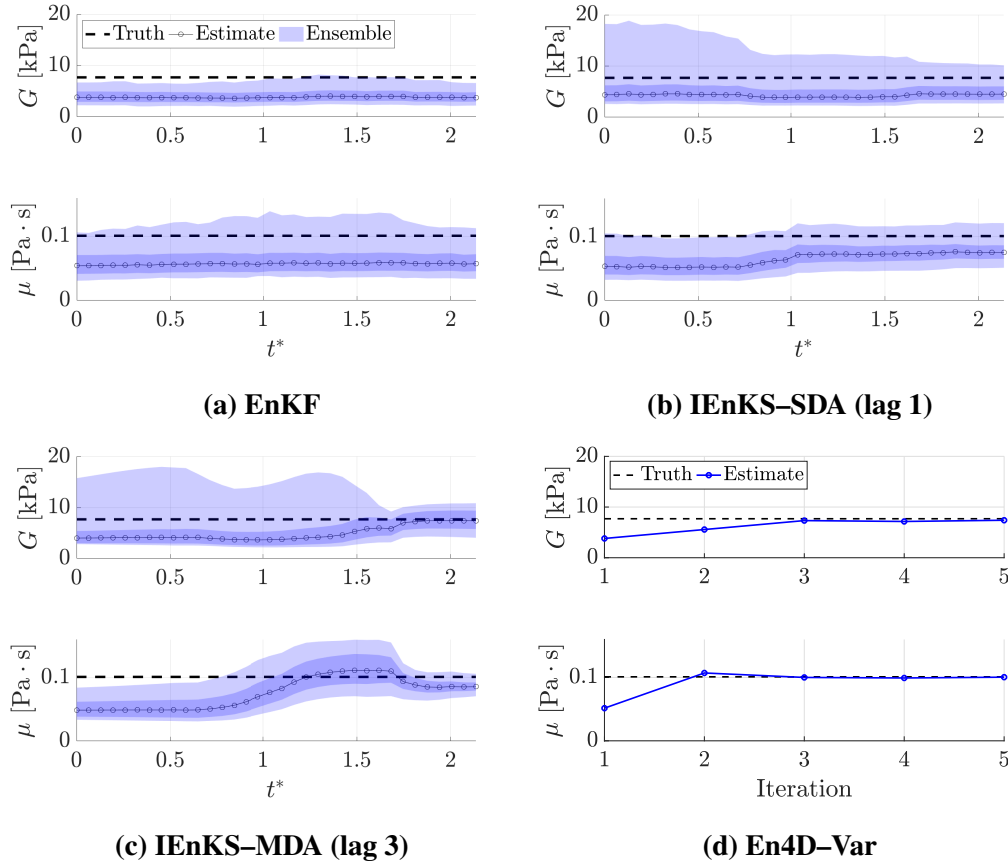


Figure 2.3: Estimation of shear modulus G and viscosity μ with initial guesses of $G = 3.8$ kPa and $\mu = 0.05$ Pa s (both at 50% error) in cases with surrogate experimental data (simulated radius time histories with added noise). The estimation is plotted over non-dimensional time t^* for the EnKF and IEnKS methods, and over iteration number for the En4D-Var. The dashed black line is the true parameter value, the blue circles are the estimate, and the blue shaded area is the ensemble spread. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.

and viscosity, the quasi-online IEnKS-MDA can uncover potential limitations of this assumption. This issue is further examined in section 3.1.

2.3.2 Uncertainty

Ensemble methods carry information about error statistics of the estimated parameters in the final ensemble. One way to visualize ensembles is through a histogram, an example of which is shown in figure 2.4 for the logarithm of the Cauchy number with the lag 1 IEnKS-SDA estimator. Despite the nonlinearity of the model, the tested methods track only the first two statistical moments of an assumed Gaussian

filtering or smoothing PDF. Previous works (e.g., Evensen and Leeuwen (2000), Yang, Kalnay, and Hunt (2012), and Katzfuss, Stroud, and Wikle (2016)) have discussed that adequate results can still be achieved with a nonlinear model where this assumption must break down to some degree. Our results for the IEnKS and En4D-Var results above confirm that this is the case in this example.

Figure 2.5 shows a comparison of the fitted histograms for the methods for the case with 50% initial error in both parameters. Despite imperfect estimation, the En4D-Var converges significantly better than other methods given the limited data. The IEnKS-MDA curve displays the least variance of the Kalman methods, as expected.




Figure 2.4: Histogram of the final estimate for $\log(\text{Ca})$ with the lag 1 IEnKS and fitted normal curve, where n is the number of ensemble members at each value of $\log(\text{Ca})$. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.

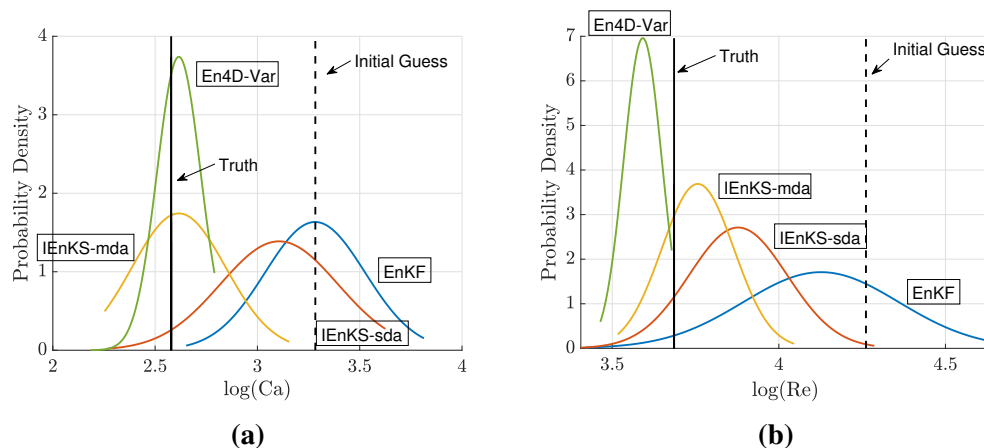


Figure 2.5: Comparison of final ensembles for $\log(\text{Ca})$ (a) and $\log(\text{Re})$ (b) estimates with each tested DA method in the case with 50% initial error in both parameters. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.

Anticipating experimental results where 10 data sets are available, the En4D–Var was run 10 times with the same ground truth but different (random) noise. Results are shown in figure 2.6 for shear modulus and viscosity estimates over the data sets. The dashed black lines correspond to the truth, and the blue line to the mean estimate over the 10 runs. Results across these 10 data sets are fairly uniform (standard deviation of 0.98 kPa for G , 0.009 Pa s for μ), confirming that reliable estimates are obtained despite noisy measurements across data sets. Figure 2.7 shows a histogram combining final ensembles for shear modulus to visualize overall results. As each of the 10 ensembles should be approximately normal, a Gaussian curve is expected when combining them. Figure 2.7 indeed shows an approximately normal distribution, as does the equivalent histogram for viscosity (as shown in section 3.1.3 in figure 3.3a).

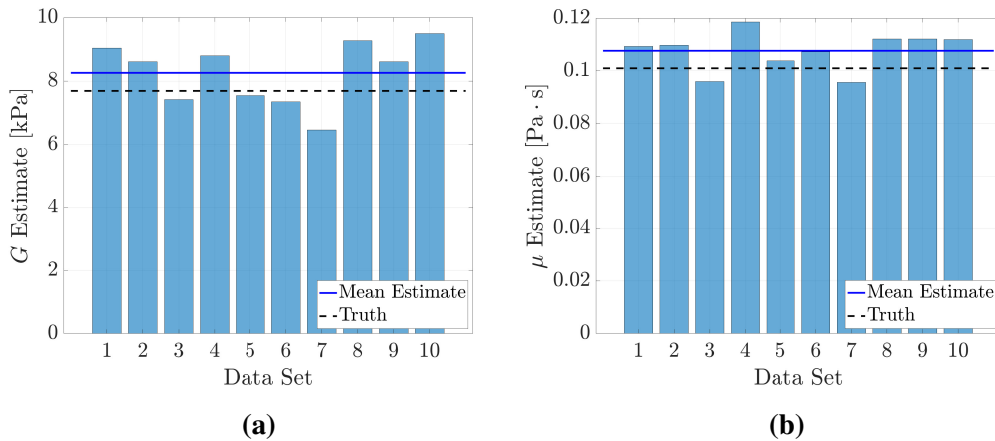


Figure 2.6: En4D–Var estimation results for G (a) and μ (b), for ten simulated data sets. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.

Based on these results with simulated data, given reasonable initial guesses as to the shear modulus and viscosity, we can confidently expect to estimate both parameters to within 5% using the 10 available data sets. Multiple initial guesses can be tested and their fits with experimental radius histories compared. In practice, an iterative process can be used to formulate a good initial parameter guess. That is, the final estimates of a first data assimilation run can be used as an initial guess for the next, and so on until results converge or an adequate radius fit is achieved. As shown in table 2.2, even with a large error, both IEnKS methods and the En4D–Var significantly improve on the initial guess, thus only a few iterations suffice to obtain a good estimate. The En4D–Var is particularly well suited for this, as it

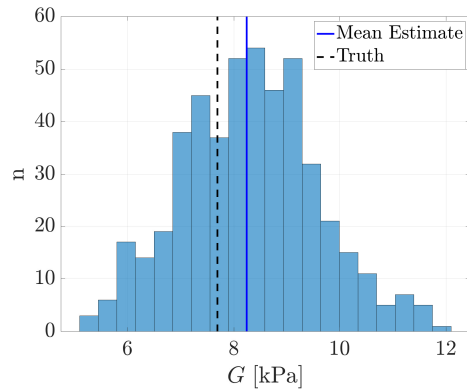


Figure 2.7: Histogram for G estimates combining 10 final ensembles for simulated data runs with En4D-Var. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.

remains stable even with large initial error and a wide spread in the ensemble, as long as all parameters retain physical values across the ensemble. Therefore, it is straightforward to formulate an initial guess with less than 50% error, and thus obtain results comparable to the second test case with simulated data. When applying this to experimental data, the En4D-Var serves as the baseline given its performance. IEnKS-MDA can also be tested for its ability to estimate parameters quasi-online.

We note that in all presented results, values for the β_i coefficients in the IEnKS-MDA and En4D-Var methods, which attribute weights to different time steps (see equation (2.23)), were held constant across all time-steps. Different values were tested for these coefficients to observe their impact on estimation results. For the IEnKS-MDA, we attempted to increase or decrease the coefficients as we stepped forward over the data assimilation window (linearly with varying slope). Neither approach, across all test cases, had a positive impact on the estimation. Thus we kept constant values for the coefficients, with $\beta_i = \frac{1}{3}$ for each $i \in [0, 3]$ for the data assimilation window of 3 time steps used.

For the En4D-Var methods, given the large amount of information gained near the collapse points during estimation, we attempted to increase the β_i coefficients near collapse points. That is, we posited that by increasing the weight attributed to data near collapse, and increasing the sensitivity of our estimates to collapse points, we may improve the accuracy of results. However, as with the IEnKS-MDA, no improvement in parameter estimation was observed given this strategy. Furthermore, for consistency, modifying coefficients in this way would require

additional parameter tuning for each data set, as collapse occurs at different times in each. Hence, as with the IEnKS–MDA method, we pick constant coefficients of $\beta_i = \frac{1}{N}$ for the En4D–var, with N the total number of time-steps. This yields good results, without any added tuning required for β_i coefficients.

Based on the results presented in this chapter, DA-IMR—in particular the IEnKS–MDA and En4D–Var methods—appears to be a powerful framework for material parameter estimation, expanding on the capabilities of IMR and maintaining high accuracy. However, there are limitations to working with synthetic data. While noise with a high variance is added to the surrogate data, this is white noise added to data generated with the same model used in our DA forecast model. This means that the model used throughout this section matches exactly with the data (except for the added noise), which is a significant simplification of real-world cases, where the model carries uncertainty associated with its assumptions (as listed in section 2.1). We have shown that if the model used is perfect, we can accurately estimate material parameters even with high levels of noise, but this does not account for systematic modeling errors that could be present in experimental data. In the next chapter, we address this by applying DA-IMR to experimental data to further test its performance. Given the results from this chapter, we expect DA-IMR to perform well, as long as the physical model we use accurately captures the physics of bubble collapse in the hydrogels used.

*Chapter 3*APPLICATION OF DA-IMR TO LASER AND
ULTRASOUND-INDUCED CAVITATION

Parts of this chapter are adapted from Section 5 of Spratt, Rodriguez, Schmidmayer, et al. (2021), Chapters 4-5 of Buyukozturk et al. (2022) and Chapters 2-4 of Mancina, Yang, et al. (2021). This chapter presents applications of our developed DA-IMR framework in three separate cases with varying cavitation methods, material properties, experimental setups, and using different material models. This establishes the accuracy and versatility of our method, and provides concrete applications of the framework. We begin by applying the data assimilation methods found to be optimal in Chapter 2 to experimental laser-induced cavitation data, where bubble radius time-histories are obtained through high-speed imaging of cavitating bubbles in a gel subject to a high amplitude laser pulse. Here, we show that our method accurately retrieves material parameters with experimental data which the surrogate data used in Chapter 2 was designed to mimic. We discuss these experimental results in detail and compare them to previous results obtained with the same data sets. Then, we present an application with laser-induced cavitation in gels with nucleation seed particles. These particles enable cavitation with lower laser energy, to access new finite deformation and material stretch regimes. A new, more robust material model is also used here requiring estimation of new model parameters. We show the advantage of parameter estimation at lower energy, and use the DA methods to pinpoint uncertainties in our physical model. Finally, we present results of our DA framework in ultrasound-induced cavitation data. Using ultrasound reduces modeling uncertainty seen with laser-induced cavitation, as no plasma is formed in the material during cavitation. This highlights a new setting where DA-IMR can be used.

3.1 DA-IMR with Laser-induced Cavitation Data

In Chapter 2, all presented results were obtained with simulated data. There, the model used in our DA framework exactly matches that of the data, as the same forecast function in our DA methods is used to generate this data. This, of course, represents an idealization of experimental scenarios, as the model used may

contain intrinsic error. Thankfully, these DA methods are well suited to address both experimental and model error, and we will show that we can predict material properties with a high degree of accuracy. In the absence of known ‘true’ material parameters, we will use the discrepancy between simulated radius vs. time curves and experimental data as the main metric to determine the adequacy of our results.

3.1.1 Experimental Setup

We begin by briefly describing the experimental setup for the first test case examined in this chapter, the data from which is used throughout section 3.1. This data was presented in Estrada et al. (2018), where a more detailed description of the setup for data collection can be found. Here, a polyacrylamide is used as the medium for cavitation, the properties of which will be estimated. After the polyacrylamide gel is prepared, each cavitation event is induced with a 6ns pulse of a “user-adjustable 1–50 mJ, frequency-doubled Q-switched 532 nm Nd:YAG laser.” These cavitation events are triggered at different locations in the same large batch of polyacrylamide to maximize uniformity of material properties across experiments. Bubble radius is captured approximately every 3.7 μs , processing 270 000 fps high-speed camera output by subtracting a reference image from each frame and fitting a circle. A few sources of error may be present. Nonuniformity of the polyacrylamide gel or discrepancies across data sets could cause the bubble to lose spherical symmetry. However, each was triggered at least 5 maximum bubble radii away from a previous location or edge of the gel, to prevent any boundary effects which could introduce non-sphericity, for example through the formation of microjets. Estrada et al. (2018) report that the maximum \dot{R}/c was approximately 0.4, and thus below a regime where significant non-spherical effects may be introduced during the initial collapse (Brujan, Nahen, et al., 2001; Sagar and el Moctar, 2020). Laser pulses may also vary slightly across runs, affecting the energy deposited in the system and thus initial growth conditions. In practice, a difference in maximum bubble radii was observed, with $R_{\text{max}} = 388 \pm 35 \mu\text{m}$ across experiments with the stiff gel, and $R_{\text{max}} = 430 \pm 17 \mu\text{m}$ with the soft gel. Ten experimental data sets in the stiff gel were used for the following results, in part to address this potential lack of uniformity in experimental conditions. As with the test case with surrogate data in Chapter 2, we obtain similar results for the stiff and soft gels, and report only results with the stiff gel here

3.1.2 En4D–Var Estimates for Shear Modulus and Viscosity

The noise magnitude in the experimental data is smaller than what was used in the simulated data with the same data rate. Therefore, if the model is adequate and the noise accurately represented as Gaussian, the IEnKS and En4D–Var should yield comparable or better estimation results with the experimental data. As with the simulated data, the assimilation window is limited to the initial collapse and two subsequent rebounds to match the setup used by Estrada et al. (2018), who estimated in the stiff-gel: $G_{\text{stiff}} = 7.69 \pm 1.12$ kPa and $\mu_{\text{stiff}} = 0.101 \pm 0.023$ Pa s. The results are compared to theirs. Our estimation is initialized with three different initial guesses, detailed in table 3.1. Similarly to the surrogate truth data from section 2.3.1, initial guesses with 10%, 50%, and 100% difference from the Estrada et al. (2018) estimates are chosen.

Method	G estimate $\pm \sigma$ [kPa]	μ estimate $\pm \sigma$ [Pa s]	Run time [s]
Guess 1	8.50 (+10% diff)	0.09 (-10% diff)	–
IEnKS–SDA (lag 1)	7.93 ± 1.68	0.096 ± 0.012	2751
IEnKS–MDA (lag 3)	7.51 ± 1.50	0.089 ± 0.016	9536
En4D–Var	7.41 ± 1.63	0.093 ± 0.014	609
Guess 2	3.80 (-50% diff)	0.05 (-50% diff)	–
IEnKS–SDA (lag 1)	4.32 ± 0.46	0.085 ± 0.013	2832
IEnKS–MDA (lag 3)	6.67 ± 1.43	0.083 ± 0.016	10052
En4D–Var	6.53 ± 1.58	0.090 ± 0.014	585
Guess 3	15.0 (+100% diff)	0.20 (+100% diff)	–
IEnKS–SDA (lag 1)	9.46 ± 2.76	0.114 ± 0.014	2871
IEnKS–MDA (lag 3)	8.57 ± 1.52	0.103 ± 0.015	9222
En4D–Var	8.24 ± 1.58	0.098 ± 0.016	535

Table 3.1: Comparing results of estimation with three different initial guesses for the material parameters G and μ . Runs were again performed on a computer with dual 12-core 2.3Ghz processors. Table reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.

Table 3.1 summarizes the results from the three different initial material parameter guesses for the three methods. These estimates correspond to the mean estimate over all 10 experimental data sets. The standard deviation σ of the results is also reported.

The En4D–Var estimates for shear modulus and viscosity all fall within the error bounds provided by Estrada et al. (2018). Results are close to theirs in the test

cases considered. The estimates are uniform, with only 8.5% difference between viscosity estimates, and 23% difference in the shear modulus results across the data sets. This larger difference in the shear modulus estimates and in the associated standard deviations is expected, as we have found the radius curves to be relatively more sensitive to μ than G . Finally, the average normalized root mean squared error (NRMSE) for bubble radius is low at 2.16×10^{-2} for guess 1, indicating that a good fit was achieved with this method. This NRMSE is obtained by re-running the simulation with the forecast model and final parameter estimates, and comparing the results to experimental data using equation (3.1). Given that the guess-1 results lead to the smallest radius error, our initial shear and viscosity modulus estimates are $G = 7.41 \pm 1.63$ kPa and $\mu = 0.093 \pm 0.014$ Pa s, respectively. An example bubble radius curve is shown in figure 3.1, for one of the experimental data sets (data set 10).




Figure 3.1: Simulated radius vs. time curve given by En4D-Var estimates for G and μ , and experimental measurements for data set 10. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.

The standard deviation across the 10 IEnKS-SDA runs was comparable to the En4D-Var. However, the estimates varied significantly based on the initial guess: from 4.32 kPa to 9.46 kPa for shear modulus, and from 0.085 Pa s to 0.114 Pa s for viscosity. Except for estimates from guess 1, the shear modulus estimates are also outside the bounds given by Estrada et al. (2018), and the radius fit is significantly worse than that of the En4D-Var, with an average NRMSE of 9.10×10^{-2} .

On the other hand, while still worse than that of the En4D-Var, the IEnKS-MDA estimates are within the Estrada et al. (2018) margin, and the radius fit is better

than that of the IEnKS–SDA (NRMSE = 6.79×10^{-2}). The IEnKS–MDA thus represents the best tested quasi-online method, as expected from the simulated data results of section 2.3.1. It is important to note here that the bubble radius fits were all obtained by re-running simulations with final shear modulus and viscosity estimates and comparing to experimental measurements. This is a fair way to compare the ability of each method to estimate these parameters. However, the radius fit obtained online during assimilation with the IEnKS methods is better (and comparable with the En4D–Var estimates), given that the radius is also being directly corrected at each time-step as part of the state vector. For parameter estimation, the En4D–Var is the best tested method, but the IEnKS–MDA is a good quasi-online estimator. This is particularly useful for the discussion in section 3.1.3, where we make use of this time-varying estimation.

3.1.3 Accounting for Model Uncertainty

The estimates obtained for the shear modulus and viscosity from the previous section show that ensemble data assimilation methods can be effectively used for estimation of viscoelastic material properties. A further look at the results, though, provides more information than simply this estimate. Examining estimates for each variable across the 10 tested experimental data sets, as shown in figure 3.2, there appears to be a discrepancy between data sets 3, 4, 5 and the rest for the viscosity. While the shear modulus estimation shows no discernible trend (despite the previously mentioned larger spread in results), the viscosity data appears to be split between two estimates. The red line in figure 3.2b shows the mean estimate of data sets 3 to 5 ($\mu = 0.074$ Pa s) and the green line that of the rest ($\mu = 0.102$ Pa s).

Figure 3.3 compares the histogram obtained when collating the $10 \times q$ final ensemble members for viscosity from 10 runs with different simulated data but the same ground truth (Figure 3.3a) and the 10 runs done with experimental data (Figure 3.3b). As discussed in section 2.3.2, we expect to approximately retrieve a Gaussian distribution around the estimate, as is the case for the simulated run in figure 3.3a. However, figure 3.3b shows an apparent bimodal distribution. The lower viscosity peak corresponds to the mean estimate of data sets 3 to 5, and the higher peak to that of the rest.

To understand what may be causing this discrepancy in results, it is useful to consider the IEnKS–MDA and its quasi-online estimation of viscosity. Figure 3.4 shows a




Figure 3.2: En4D-Var estimates for shear modulus (a) and viscosity (b) for 10 experimental data sets. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.

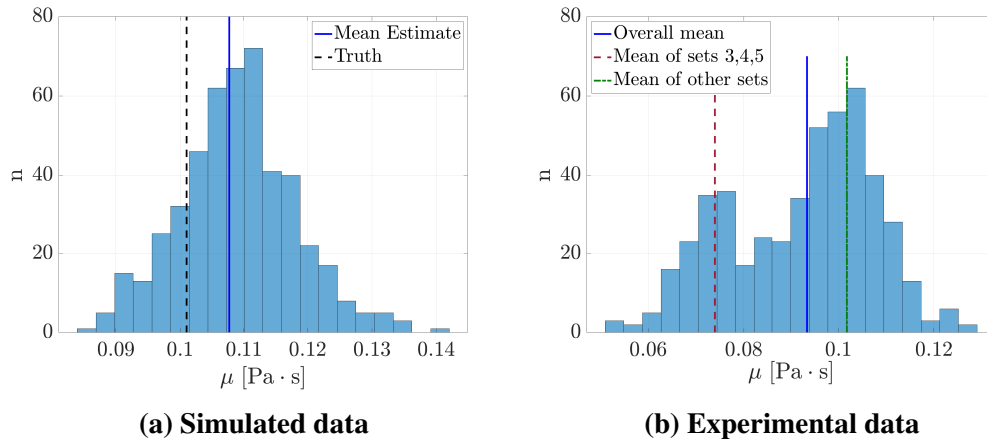


Figure 3.3: Comparison of final combined ensembles for viscosity estimation in simulated and experimental data. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.

comparison between viscosity estimation for data set 2 (figure 3.4a) and data set 3 (figure 3.4b). These are representative of data sets with a high and low viscosity estimate, respectively. They result in estimates of $\mu = 0.098$ Pa s for data set 2, and $\mu = 0.077$ Pa s for data set 3. Comparing the two data sets, it appears that the assimilation begins similarly, correcting to a higher viscosity estimate during the first collapse. However, there is a divergence between the behavior of the smoother after each collapse point, particularly the second one (around $t = 65$ μ s). Figure 3.4a shows a slow decrease and convergence towards a higher viscosity value, with negligible change at the second collapse point. However, this estimate drops

sharply after these collapse points in figure 3.4b. In fact, the data around each collapse causes the viscosity estimate to sharply drop in data set 3, which does not occur in data set 2. This behavior is representative of what is seen in data sets 3 through 5, but does not occur in the rest of the runs.




Figure 3.4: Comparison of online estimation of viscosity in data sets 2 (a) and 3 (b) using the IEnKS–MDA. Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.

Given the physical model used, the viscosity should be constant and such drops in the parameter are not expected. The model alone thus cannot adequately capture the behavior of the gel seen by the IEnKS in these data sets. We can posit that a violent collapse in these data sets is causing inelastic behavior in the material, and thereupon this perceived change in material properties (Yang, Cramer, and Franck, 2020). More work will be needed to determine the exact cause, but this could perhaps result from fracture, damage to the polymer network in the gel, or combustion in the gas phase (Movahed et al., 2016; Kundu and Crosby, 2009; Raayai-Ardakani, Earl, and Cohen, 2019).

To further examine this, we perform two simulations where we slightly modify our spherical bubble dynamics model. Here assumption of constant temperature in the surrounding material is relaxed, and we apply conservation of energy to obtain a partial differential equation for this temperature field. Appendix A describes the changes to the model used to calculate material temperature, based on the model of Preston, Colonius, and Brennen (2007). Specifically equation (A.10) is used for temperature in the surrounding medium. The material temperature field T_m is added to the state vector x , and estimated along with the temperature field inside the bubble and all other variables. While it has predominantly been used to estimate constant model parameters G and μ so far, IEnKS provides quasi-online estimates of all state

variables, which now include the temperature field in the surrounding gels. Figure 3.5 shows a comparison of heat-maps of this estimated temperature field for data sets 2 and 3, where the x-axis represents non-dimensional time $t^* = \frac{t}{R_{max}} \sqrt{\frac{p_{\infty}}{\rho}}$, the y-axis the distance from the bubble wall, and the color of each cell is non-dimensional temperature.

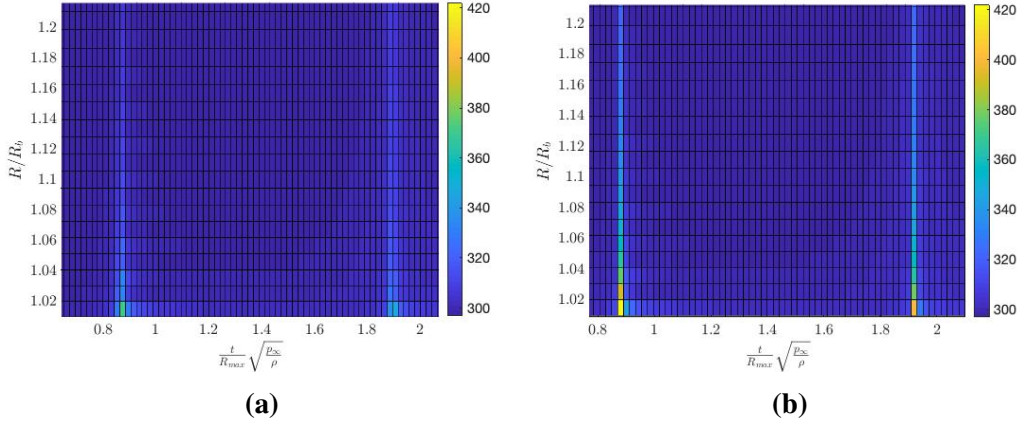


Figure 3.5: Comparison of online estimation of the temperature field in the surrounding material for data sets 2 (a) and 3 (b) using the IEnKS-MDA. The x-axis is non-dimensional time, the y-axis represents distance from the bubble wall, and the color of each cell represents temperature.

In both cases, temperature peaks are visible at two distinct times, which correspond to the first and second collapse of the bubble. However, the magnitude of these temperature peaks is larger for data set 3, the same data set which saw the sharp drop in viscosity estimate at collapse time. This could indicate that the violence of collapse in this data set results in high temperatures at the bubble wall, which in turn causes damage to the material. Regardless of physical cause, this time-dependent behavior is not accounted for in the model, but is captured by the IEnKS-MDA as a drop in the perceived viscosity. While the exact bounds of the physical model used are not known and would require more data to determine, this shows that in these particular data sets, model accuracy is far reduced after the first collapse.

Figure 3.6 compares the normalized root mean squared error (NRMSE) across all data sets, given by

$$\text{NRMSE} = \sqrt{\frac{1}{N} \sum_i \frac{(R_{\text{sim}}(t_i) - R_{\text{exp}}(t_i))^2}{R_{\text{exp}}(t_i)^2}} \quad (3.1)$$

where R_{exp} is the experimental bubble radius time history, and R_{sim} is the simulated time history given the estimated material properties (at the corresponding times).




Figure 3.6: Bar plot of radius normalized root mean squared errors (NRMSE) for each data set. Also plotted are the previous estimate mean NRMSE (mean of all sets), the mean NRMSE for sets 3 to 5, and the final estimate mean NRMSE (mean of all other sets). Reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.

Figure 3.6 shows a higher error in the estimated bubble radius curves fit for data sets 3 through 5, which is expected given the heightened model uncertainty in these data sets. Because of this uncertainty and higher error, we discard these three sets as outliers, which yields the final IEnKS–MDA-informed En4D–Var estimate reported in table 3.2. Notable is the drop in standard deviation for viscosity as compared to the previous En4D–Var estimate and the reduced NRMSE.

Estimate	$G \pm \sigma$ [kPa]	$\mu \pm \sigma$ [Pa s]	NRMSE
Estrada et al. (2018)	7.69 ± 1.12	0.101 ± 0.023	
Previous	7.41 ± 1.63	0.093 ± 0.014	2.16×10^{-2}
Final	7.81 ± 1.80	0.102 ± 0.006	1.95×10^{-2}

Table 3.2: Final En4D–Var estimates (discarding three outlier data sets) and standard deviation, along with the average radius normalized root mean squared error. The previous best estimate corresponds to the mean of all 10 data sets, outliers included. Table reprinted from Spratt, Rodriguez, Schmidmayer, et al. (2021) with permission from Elsevier, © 2021.

Given the similarity to the values obtained by Estrada et al. (2018) and the low NRMSE of the resulting radius time history simulations, these results show that ensemble-based data assimilation can successfully estimate the mechanical properties of hydrogels at these high strain rates. The En4D–Var method efficiently provides an accurate estimate of material properties. In addition, the IEnKS–MDA

method can elucidate time-dependent behavior, even when not fully accounted for by the forecast model used, by estimating parameters quasi-online. These methods account for both model and experimental error, dealing with noisy measurements or inaccuracies in the model well. It is hypothesized here that the bubble collapses are damaging the polyacrylamide gel in certain test cases, leading to a reduced estimated viscosity after each subsequent collapse, a physical effect not accounted for in the model. Finally, added benefits of these algorithms include adaptability to different numerical or viscoelastic models, and scalability to further parameter estimation, with negligible computational cost for additional parameters. The following chapters will demonstrate these benefits, as we apply these methods to new material models and estimate additional parameters.

3.2 DA-IMR with Laser-induced Cavitation Data in Seeded Gels

In the results of section 3.1, we applied our DA framework to experimental data sets within the same deformation regime. The material stretch ratio, defined (assuming a spherical bubble) as the maximum value of the hoop stretch over the bubble collapse

$$\lambda_{\max} = R_{\max}/R_0, \quad (3.2)$$

where R_0 is the long-time bubble equilibrium radius, varies little across experiments. This introduces some limitations to the scope of the DA-IMR method. Namely, across various materials and applications, maximum deformation can vary broadly. The breakdown process undergone during laser-induced cavitation is highly dependent on the material's properties, and can lead to drastically different stretch ratios in different materials. Thus, testing our characterization framework in different stretch ratio regimes would inform how broadly it can be applied. To address this, Buyukozturk et al. (2022) introduce a novel experimental approach, whereby the stretch ratio can be controlled within the same material by introducing seed particles and varying laser energy. The seed particles are used as a nucleation site to cavitate bubbles, reducing the dependence of bubble dynamics on properties of the material itself. Indeed, the particles make it far easier for cavitation to occur, and lower laser energy can be used. Varying both laser energy and the nature of the seed particles, a wide range of stretch ratios can be achieved for a single material.

Furthermore, it is known that the bubble dynamics physical model used, described in section 2.1 in Chapter 2, can be limited in the case of extreme loading rates (Gent and Wang, 1991; Barney et al., 2020; Hashemnejad and Kundu, 2015; Hutchens,

Fakhouri, and Crosby, 2016). As discussed in section 3.1.3, the model does not account for material fracture or other failure mechanisms, and assumes a consistent viscoelastic response. By using our data assimilation approach, we can continue to quantify modeling uncertainty and identify regimes or critical points in the bubble collapse where our models may be failing to adequately capture the physics. Applying DA-IMR to the experimental method of Buyukozturk et al. (2022), which extends stretch ratio regimes where cavitation can occur, will much more broadly assess the adequacy of our DA-IMR method in determining material properties of viscoelastic materials.

3.2.1 Experimental Setup and Results

We briefly describe the experimental setup used and results obtained by Buyukozturk et al. (2022), to which we apply our DA-IMR framework in section 3.2.2. Overall, the setup is similar to that of Estrada et al. (2018), used to generate the data analyzed in section 3.1. The hydrogel used is again polyacrylamide, with a quasi-static shear modulus of $G_{\infty} = 461 \pm 4$ Pa. Four samples are prepared. In three of them, seed particles are distributed throughout the gels, respectively microspheres of glass, stainless steel, and paramagnetic coated polyethylene. For details on the gel preparation and laser setup used to nucleate bubbles, see section 2 of Buyukozturk et al. (2022). The key differences with the setup used to generate the data of section 3.1 and shown in figure 1.1 are that the laser will be focused on seed particles (in gels containing them) to initiate cavitation, and the use of neutral density (ND) filters to attenuate the laser energy, thus modulating the energy between 18.4 ± 1.0 μ J and 449.2 ± 5.0 μ J. Finally, the high-speed camera used to image bubbles is capable of far greater frame rates than in the previous setup. In the data sets used in the following, bubble radius is captured every 1 μ s, at 1 million frames per second (compared to 270,000 previously). We saw in section 3.1.3 when using the En4D-Var method that the estimator was sensitive to data around the first bubble collapse. However, the time-resolution of those experiments meant that only a few data points were available near this time, given the violence of the dynamics. Increasing the time-resolution of our data by almost a factor of 4 means we will be able to resolve the collapse much better, and improve material parameter estimation.

As expected, the use of seed particles enables cavitation across various stretch ratio regimes. Both particle type and laser energy have an impact on the maximum bubble

radius. For cases with no seed particles, the stretch ratio hovered around $\lambda_{\max} = 9$, with values averaging $\lambda_{\max} = 9.2$ for the highest laser energy (449 μJ), and reducing to $\lambda_{\max} = 8.8$ at the lowest laser energy where cavitation can still occur with no seed particles present, around 118 μJ . In contrast, the stretch ratio can be reduced all the way to a mean of $\lambda_{\max} = 4.5$ in cases with the paramagnetic beads and laser energy 18.4 μJ . While the reduced laser energy does affect stretch ratio, the nature of seed particles has a large effect as well. For example, with a laser energy of 118 μJ , the stretch ratio is $\lambda_{\max} = 8.8$ in the gel without particles, $\lambda_{\max} = 5$ for paramagnetic particles, $\lambda_{\max} = 5.6$ for glass particles, and $\lambda_{\max} = 7.2$ for steel particles. Figure 3.7 summarizes the stretch ratios across all experimental data sets. We note that with seed particles present, laser energy was capped at 254 μJ to avoid cases where loss of sphericity was observed.




Figure 3.7: Value of Maximum stretch ratio λ_{\max} vs. laser energy obtained across experiments with varying seed particles. All points represent the mean of 8 experimental data sets. Reproduced from Buyukozturk et al. (2022) with permission from Springer Nature, © 2022.

3.2.2 DA-IMR Parameter Estimation Results

We implement ensemble-based data assimilation techniques described in Chapter 2 and used in section 3.1 to fit the hydrogel's material properties and account for uncertainties in the physics model. In particular, the IEnKS-MDA and En4D-Var are used, with some minor modifications described here. First, some elements of

the bubble dynamics model described in section 2.1 are improved. The Keller-Miksis equation (equation (2.1) (Keller and Miksis, 1980)) is still used, with the physics of a two-phase mixture of condensable water vapor and non-condensable gas in the bubble tracked, as shown to be appropriate in the literature (Estrada et al., 2018; Yang, Cramer, and Franck, 2020; Akhatov et al., 2001; Nigmatulin, Khabeev, and Nagiev, 1981; Barajas and Johnsen, 2017; Vincent et al., 2014). The same assumptions described in section 2.1 are maintained, however, the viscoelastic material model used is changed. Indeed, Yang, Cramer, and Franck (2020) showed that the quadratic law Kelvin-Voigt model, an adaptation of the Fung model (Fung, 2013) where higher order strain stiffening terms are used (Movahed et al., 2016), is best for modeling very high strain rate behavior of gels such as the presently used polyacrylamide. The use of this more general model may help estimation as we expand DA-IMR to new material stretch regimes, as compared to results from section 3.1.

Rather than estimating a shear modulus G as done in section 3.1, we fix a quasi-static shear modulus G_∞ and estimate a strain-stiffening parameter α . The stress integral used with this model is then defined as

$$S_{Quad} = \frac{(3\alpha - 1)G_\infty}{2} \left[5 - \left(\frac{R_o}{R}\right)^4 - \frac{4R_o}{R} \right] - \frac{4\mu\dot{R}}{R} + 2\alpha G_\infty \left[\frac{27}{40} + \frac{1}{8} \left(\frac{R_o}{R}\right)^8 + \frac{1}{5} \left(\frac{R_o}{R}\right)^5 + \left(\frac{R_o}{R}\right)^2 - \frac{2R}{R_o} \right]. \quad (3.3)$$

To account for these changes, the state vector used (previously equation (2.10)) becomes

$$\mathbf{x} = \{R, \dot{R}, p_b, S, \mathbf{T}, \mathbf{C}, \alpha, \mu\}. \quad (3.4)$$

Otherwise, the IEnKS and En4D-Var implementation of sections 2.2.2 and 2.2.4 remain unchanged. A lag of 3 is used with the IEnKS.

Now, in section 3.2.2.1, we perform bubble radius fits using the offline (En4D-Var) method to determine the Quadratic law Kelvin-Voigt material parameters up through the first (1-peak), second (2-peak), and third (3-peak) collapse. Material parameters are then used to calculate the critical Mach numbers for each initial bubble collapse using the IMR framework. For the online method, we assess the time-varying fitted material properties of the quasi-online (IEnKS) method to identify deviations from the model during the bubble time evolution in section 3.2.2.2. We note that

the analysis of results of the following sections 3.2.2.1 and 3.2.2.2 was done in collaboration with Dr. Selda Buyukozturk and the other co-authors of Buyukozturk et al. (2022).

3.2.2.1 Multi-peak Fitting with the Offline (En4D-Var) Method

Numerical results presented in this section use the offline (En4D-Var) method for determining the viscoelastic material parameters in PA. Initial guesses into the solver for viscosity, μ , and strain stiffening parameter, α , are 0.05 Pa·s and 0.5, respectively. These are derived from a preliminary run with the former version of IMR using a least squares fitting scheme for fitting material properties Yang, Cramer, and Franck, 2020.

As a measure of the $R(t)$ goodness of fit given the En4D-Var parameter estimates, as with section 3.1, we use the normalized radius root mean squared error (NRMSE). Representative $R(t)$ fits with median NRMSE values for each fitted peak case are shown in Fig. 3.8. For a cavitation bubble in PA (paramagnetic seed particle), induced at nominal laser energy 117 μ J, the 1-peak case has a median NRMSE value of 0.009 (Fig. 3.8a), 2-peak case has a median NRMSE value of 0.027 (Fig. 3.8b), and the 3-peak a median NRMSE value of 0.047 (Fig. 3.8c).

We compare $R(t)$ curves obtained with the En4D-Var parameter estimates to experimental data to calculate the NRMSE. Figs. 3.8(d-f) compare λ_{max} to the NRMSE for each peak fit. As the number of fitted peaks increases, the NRMSE also increases. 1-peak fits (Fig. 3.8d) exhibit the lowest NRMSE, with most errors within 0.02. However, the error for the 2-peak fits (Fig. 3.8e) are mostly within 0.05, with some large scatter exhibited by the steel LIC cases. The 3-peak fit case (Fig. 3.8f) NRMSE values are mostly within 0.15, with a few outliers. Given the number of fitted peaks, stretch ratio does not have a noticeable effect on the NRMSE. However, with increasing number of peaks fitted, the NRMSE from the 1- to 3-peak fits increases by almost an order of magnitude. Thus, subsequent oscillatory bubble dynamics beyond the first peak are not as accurately described by the current theoretical IMR framework.

Next, we examine the effect of laser energy on NRMSE. In the 1-peak fit case (Fig. 3.8g), LIC bubbles in PA with no particles have the lowest NRMSE compared to the cases with PA with seed particles. While the general trend indicates increasing

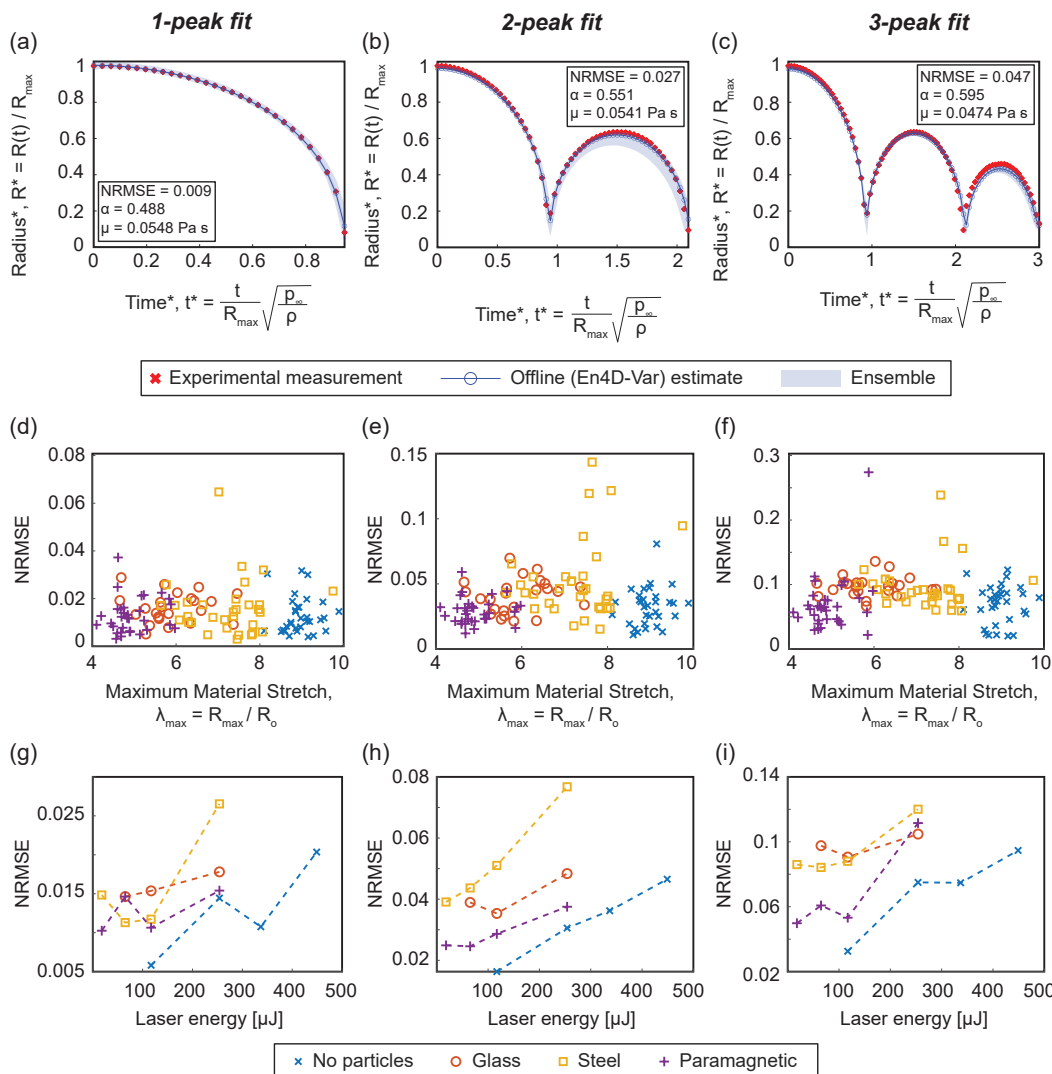


Figure 3.8: Results of offline En4D-Var fitting on experimental curves for 1-, 2- and 3-peak fits. Representative examples of multi-peak fits with median NRMSE values (for PA paramagnetic particles at nominal laser energy 117 μJ) are shown for a (a) 1-peak fit (NRMSE = 0.009), (b) 2-peak fit (NRMSE = 0.027), and (c) 3-peak fit (NRMSE = 0.047). NRMSE values with respect to maximum material stretch for (d) 1-peak, (e) 2-peak, and (f) 3-peak fits, and with respect to laser energy for (g) 1-peak, (h) 2-peak, and (i) 3-peak fitting cases. Reproduced from Buyukozturk et al. (2022) with permission from Springer Nature, © 2022.

NRMSE with increasing laser energy, the bubbles cavitating in PA with seed particles tend to have similar errors at the lowest energies and diverge more with increasing energy. This trend continues in the 2- and 3-peak fit cases, as shown in Figs. 3.8(h, i). Overall, increasing the laser energy increases the fit error.

For the range of samples and laser energies presented, the strain stiffening parameter, α , and viscosity, μ , were estimated using the En4D-Var with the IMR framework and Quadratic law Kelvin-Voigt stress integral (equation (3.3)). A summary of fitted material parameters for all cases is shown in Fig. 3.9, where the weighted-mean viscosity, μ , and strain stiffening parameter, α , values are plotted against each other with standard deviation error bars for each experimental condition. (The type of sample is denoted by symbol and the energy level is denoted by color, in accordance with the color bar shown.) The 1-peak fit case in Fig. 3.9a contains a large cluster of material parameters located within the range of $\mu \approx 0.04 - 0.08$ and $\alpha \approx 0.4 - 0.6$ with standard deviation values of $\sigma_\mu = 0.008$ and $\sigma_\alpha = 0.08$. A few higher energy cases fall outside of the cluster. However, as the number of fitted peaks increases to 2-peak (Fig. 3.9b) and 3-peak (Fig. 3.9c) fits, an increase in both μ and α values is observed, though the viscosity term is more affected. Another noticeable trend is that the standard deviation of fitted material parameters for a given experimental condition increases with respect to the laser energy. Thus, as energy increases, so does the spread in fitted material parameters, particularly for the 2- and 3-peak fitted cases. Notably though, parameter estimates for lower energy cases remain more clustered, and do not increase as much as more peaks are fitted. This may indicate that the IMR framework remains suitable over multiple peaks for lower laser energies.

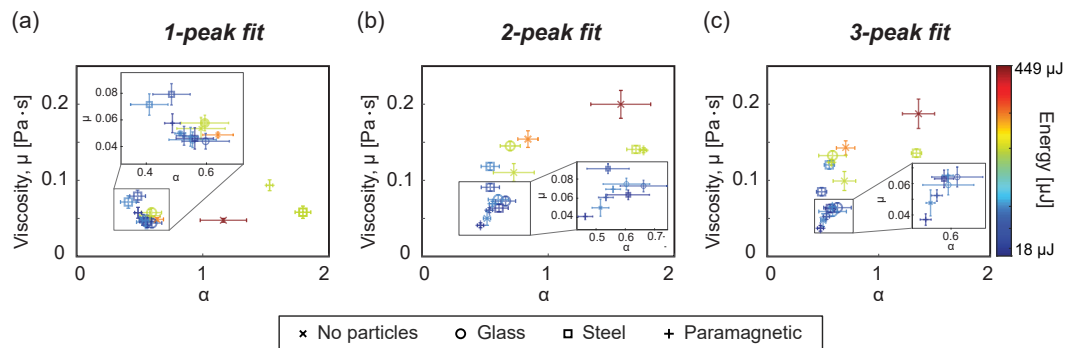


Figure 3.9: Offline En4D-Var estimates for Quadratic law Kelvin-Voigt material model parameters, namely viscosity μ , and strain stiffening parameter α . The material parameters estimates are shown for (a) 1-peak, (b) 2-peak, and (c) 3-peak fits. The symbols correspond to sample type and the color corresponds to laser energy, with the associated color bar shown on the far right. Each point is a weighted average of $n = 8$ experiments where error bars are standard deviation. Reproduced from Buyukozturk et al. (2022) with permission from Springer Nature, © 2022.

Upon collapse, inertial cavitation bubbles emit acoustic waves that propagate through the hydrogel matrix. The Mach number, M , over time is a function of the velocity of the bubble wall and the finite wave speed of the material, given by

$$M = \dot{R}/c. \quad (3.5)$$

The critical Mach number, M_{cr} , is the maximum Mach number of an inertial bubble, typically occurring at the first bubble collapse. Previous work suggests that M_{cr} values exceeding 0.08 define the condition for a violent collapse (Yang, Cramer, and Franck, 2020), which may introduce additional physical phenomena such as damage and inelastic material behavior. For more details on the IMR theoretical framework, as well as the derivation of non-dimensionalized expressions required for computations, see Estrada et al. (2018).

Given the estimated viscoelastic material parameters, the time evolution of bubble radii is simulated with the physical model equations, and the critical Mach number for each experiment is calculated using equation (3.5). In Fig. 3.10a, the critical Mach number, M_{cr} , with respect to each corresponding maximum stretch ratio, λ_{max} , is plotted for the 1-peak case and it is found that as the stretch ratio increases, so does the critical Mach number. The 2-peak (Fig. 3.10b) and 3-peak (Fig. 3.10c) cases follow the same trend, though there is an observable increase in scatter, particularly for the steel seed particle case. Thus, there seems to be a dependency of critical Mach number on the stretch ratio. Comparing the critical Mach number to laser energy, we refer to Fig. 3.7, where increasing energy only slightly increases the maximum stretch ratio. Given this, we can conclude that the critical Mach number is similarly minimally affected by laser energy.

The offline (En4D-Var) method applies the IMR framework to the experimental kinematic bubble radii using the Quadratic law Kelvin-Voigt material model to assess the goodness of the bubble radius history fit and viscoelastic material properties for up to 3-peak fits. The radius fit error, quantified as NRMSE, is generally lower for all cases at lower energy levels, regardless of the number of peaks fitted. However, the radius fit is worse, increasing by almost an order of magnitude, as we increase the DA window from 1- up through 3-peaks with respect to both maximum material stretch and laser energy. While the calculated critical Mach number did not have a noticeable change with respect to input laser energy, it was found that the critical Mach number increases with maximum material stretch ratio, a trend consistent regardless of the number of bubble radius peaks fitted. Lastly, the fitted viscoelastic

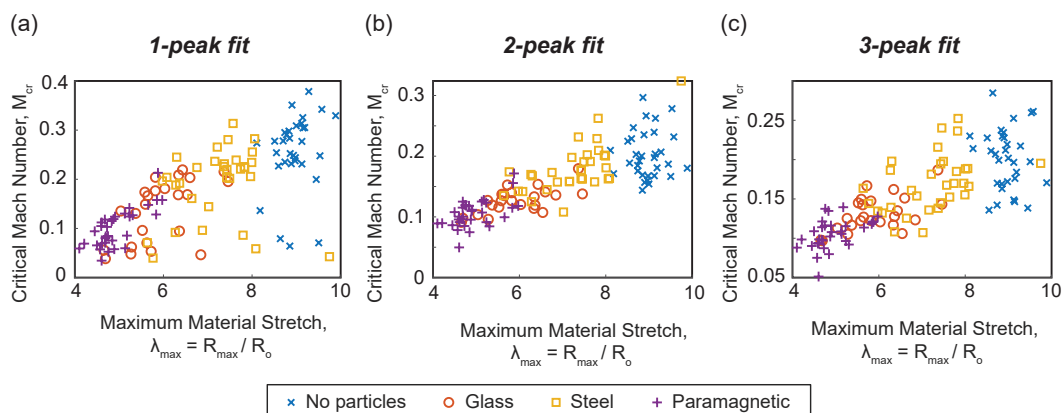


Figure 3.10: Critical Mach number, M_{cr} plotted with respect to maximum material stretch, λ_{max} , Using offline (En4D-Var) estimates given the (a) 1-peak, (b) 2-peak, and (c) 3-peak fitted solutions. Critical Mach numbers were obtained by re-running simulations given these En4D-Var estimates in each case. Reproduced from Buyukozturk et al. (2022) with permission from Springer Nature, © 2022.

material properties tend to cluster with a 1-peak fit, but the spread in mean values, as well as the standard deviation of values at a given experimental condition, increases with respect to higher laser energy and number of peaks fitted. The latter particularly affects the fitted viscosity parameter. Critical Mach number increases with respect to material stretch as before, leading to violent collapses at higher material stretches ($M \geq 0.08$).

3.2.2.2 Quasi-online (IEnKS) Results

Next, we implement the quasi-online (IEnKS) method, where the estimation trails the simulation by a fixed number of time steps. This is particularly useful because the time-varying estimation provides additional information on IMR model uncertainties. In this section, we will focus on two examples that indicate deviations from the theoretical framework, while comparing cavitation events nucleated under different experimental conditions.

For the first example, we focus on bubbles nucleated in PA (no seed particles) at all four laser energies, ranging from nominal values of 117 μJ to 449 μJ . In Fig. 3.11, we plot the ensemble's variance in these estimates over time. The time evolution of curves spans from the first bubble radius peak for all cases, where initial bubble collapse for each example occurs approximately at the normalized time indicated

by the dashed vertical lines. For the nominal laser energy of $117 \mu\text{J}$ case shown in Fig. 3.11a, the time-varying α parameter gradually increases then dips after the first bubble collapse. In Figs. 3.11(b,c), as the laser energy is increased, the α and μ variance have an increased peak. This is even more apparent at the highest laser energy case shown in Fig. 3.11d, where at a nominal laser energy of $449 \mu\text{J}$, both peak α and μ variances increase almost five times the peak values of the lowest energy case. Thus, as laser energy increases, the quasi-online (IEnKS) method produces more uncertain fitting results. This is most likely due to missing explicit mathematical descriptions accounting for additional physical phenomena that, at present, are not represented in the current IMR theoretical framework.

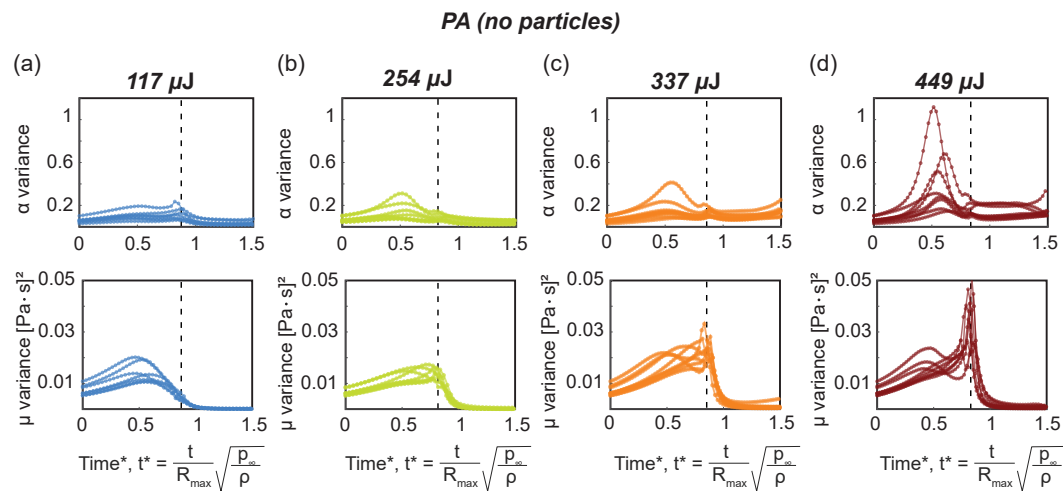


Figure 3.11: Quasi-online IEnKS α and μ estimate variance versus normalized time at nominal laser energies of (a) $\sim 117 \mu\text{J}$ (b) $\sim 254 \mu\text{J}$, (c) $\sim 337 \mu\text{J}$, and (d) $\sim 449 \mu\text{J}$. This is for bubble radius time histories in PA with no beads, where the IEnKS was run using an ensemble size of 48. First bubble collapse occurs at approximately the dashed line for all $n = 8$ samples per experimental condition. Reproduced from Buyukozturk et al. (2022) with permission from Springer Nature, © 2022.

The second representative example we present with the quasi-online (IEnKS) method compares a bubble in a PA hydrogel (no seed particles) at the highest energy to a bubble in a PA hydrogel (paramagnetic seed particles) at the lowest energy. The non-dimensionalized time versus radius curves for each case and online estimation are shown in Figs. 3.12(a,b). These cases have λ_{max} values of 9.5 and 4.7, with low NRMSE values. The data assimilation starts at maximum normalized radius, $R^* = 1$, and normalized time, $t^* = 0$. In Fig. 3.12c, we consider a time-varying α of the PA hydrogel (no seed particles) at highest energy first ($\lambda_{max} = 9.5$). There

is a gradual increase followed by a decrease in α as the bubble collapses. However, at initial collapse, α experiences a significant increase. Upon the second peak collapse, we observe another increase in the estimation. The time-varying viscosity (Fig. 3.12e) follows a similar trend, where it maintains a steady value until the first collapse, when a large jump increase is observed. Upon second collapse, the estimator adjusts to a higher viscosity estimate again.

We compare this case to a bubble nucleated within PA on a paramagnetic particle at the lowest laser energy ($\lambda_{max} = 4.7$). In this case, α has no discernible jump at collapse, but instead displays a gradual drop-off in the estimation value over time (Fig. 3.12d). Likewise, for the viscosity estimate, we observe relatively constant parameter values in both parameters at collapse, and a slight dip upon first collapse (Fig. 3.12f). It should be noted that while the parameter estimates in subsequent radius peaks are constant, this does not imply a better fit to the model during that portion of the bubble dynamics. Instead, despite using covariance inflation in the quasi-online method, the ensemble eventually converges when the best estimates are found. Thus, subsequent peaks are not a better fit to the model, but rather, the estimator can no longer improve on its estimates for μ and α .

Overall, the low laser energy (paramagnetic seed particle) case exhibits a smoother estimation around first collapse, which is likely due to the smaller exhibited material stretch. This is in stark contrast to the high energy PA (no seed particles) case, with large material stretch, where there is a discernible increase in both parameter estimates and uncertainties at the collapse points. Given the current viscoelastic IMR theoretical framework, the viscosity and strain-stiffening parameters of the surrounding PA hydrogel should be estimated to be the same under both conditions, yet there are observable differences between these cases. These stark contrasts in time-varying material behavior are consistent with previous observations of violent bubble collapses in PA gel. This notion is also reflected by a much higher critical Mach number in the case of no seed particles, compared to that of the paramagnetic particle case (see Fig. 3.10).

The selected examples described in this section illustrate the time-varying material parameter estimates, as well as the model uncertainties with respect to material stretch and laser energy. We see the greatest uncertainty in time-varying material parameters at the bubble collapse points, especially at the first inertial collapse. However, we see a convergence in material parameter variation with decreasing energy

for a given case, as well as at lower material maximum stretch regimes, indicating that for lower material stretches or nucleation energies, the DA-IMR framework is appropriate for accurately describing the laser-induced cavitation dynamics.

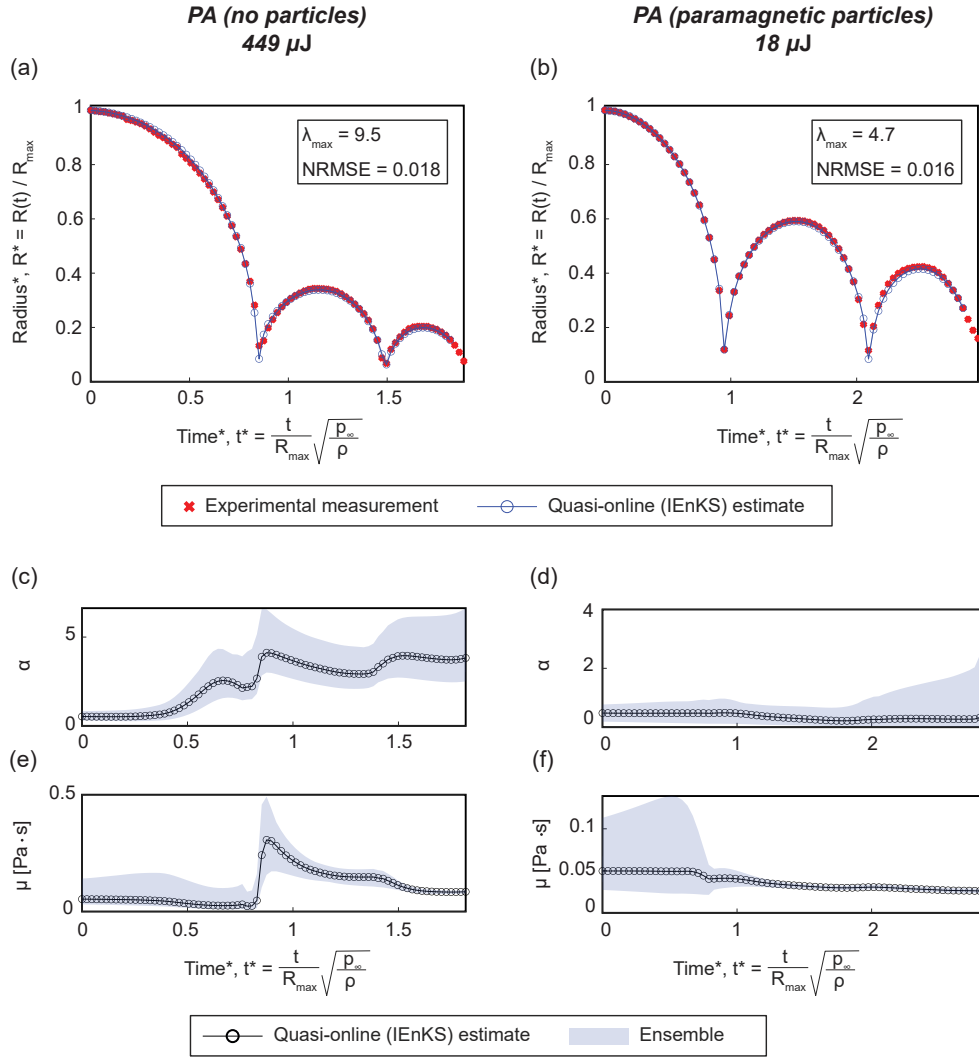


Figure 3.12: Comparison of results of the quasi-online IEnKS fitting for two extreme cases with normalized radius versus time curves in PA (no beads) at nominal laser energy 449 μJ and (b) PA (paramagnetic particles) at nominal laser energy 18 μJ . The (c, d) α (e, f) and μ estimates over normalized time for each case are plotted. Reproduced from Buyukozturk et al. (2022) with permission from Springer Nature, © 2022.

3.2.3 Discussion of Modeling Uncertainties

Overall in section 3.2, we showed that the IEnKS and En4D-Var were able to fit the parameters of a quadratic law Kelvin-Voigt constitutive model to the broad

range of data sets, while quantifying associated errors and critical Mach numbers at initial collapse. Important trends were observed given these DA results. With the high temporal resolution available in the experimental data, fitting bubble radius to the first peak only consistently led to a decrease in the NRMSE. While data from previous studies meant that using a long time-scale was vital to obtaining good parameter estimates, it appears that with this image capture rate, limiting to the first peak yields best results. The decrease in result reliability past the first collapse may be a result of the limit on compressibility effects captured in the model. Indeed the Keller-Miksis equation only accounts for compressibility effects near the bubble wall region (Keller and Miksis, 1980; Estrada et al., 2018). While this often suffices, the increase in results NRMSE past the first collapse point indicates that this assumption may not hold, particularly in cases where the critical Mach number exceeds $M_{cr} = 0.08$ (Anderson, 2009). As stretch ratio, and thus Mach number, increases, the physical model used may be limited in its ability to fully capture the bubble dynamics.

It was seen that this effect was more pronounced when increasing the laser energy. While for single-peak fitting, an increase in the stretch ratio did not correlate with an increased fitting error, a higher laser energy consistently increased the NRMSE of results, as seen in figures 3.8, 3.9, and in the increased variance in parameters estimates with the IEnKS with higher energies seen in figure 3.11. This suggests that at higher laser energies, the physics of plasma formation optical breakdown of the polyacrylamide at cavitation inception have an impact beyond our modeling capabilities. The current approach of beginning our data assimilation window at peak bubble radius no longer suffices to avoid capturing this initial breakdown, which affects the physics beyond the initial bubble growth, as assumed by our approach. Lowering energy, for example by introducing seed particles, is thus an effective way to carry out DA-IMR to avoid these effects. However, another way to address this may be to modify the cavitation inception method used, and avoid difficulties caused by the laser nucleation physics altogether. In the next section, we describe our application of the DA-IMR framework to ultrasound-induced cavitation data. There, high intensity focused ultrasound is used to induce cavitation in hydrogels instead of using laser nucleation.

3.3 DA-IMR with Ultrasound-induced Cavitation Data

3.3.1 Experiment and Problem Setup

Mancia, Yang, et al. (2021) introduce 'Acoustic Cavitation Rheometry' as an alternative to the laser-induced cavitation (LIC) rheometry used in sections 3.1 and 3.2. Here, cavitation is generated via focused ultrasound, the negative pressure peaks of which can cause bubble nucleation about existing material defects or microbubbles trapped in a material. This can avoid some of the issues observed in LIC and described in section 3.2.3, while still causing bubbles to grow explosively in media of interest, and attain similar strain rates as with LIC (Wilson et al., 2019). The reader is referred to section 2.1 of Mancia, Yang, et al. (2021) for details on the experimental setup. Key differences with the setup of experiments in section 3.1 and 3.2 are the following. First, agarose gel is used as the cavitation medium rather than polyacrylamide. The 0.3% and 1% gels are prepared following Vlaisavljevich, Lin, Maxwell, et al. (2015), and have quasi-static shear moduli of approximately 380 Pa and 720 Pa, respectively. Next, to induce cavitation, a 16-element transducer array fires 1.5 cycles of focused ultrasound (with a single negative half-period) at 1 MHz with a peak negative amplitude of -24 MPa into the gel. Each bubble, generated by acoustic forcing at the transducer focal point alone (Maxwell, Cain, Hall, et al., 2013), is at least 5 mm away from the previous nucleation site. Images were captured every $2.5 \mu\text{s}$ (at 400,000 frames per second), giving a better temporal resolution than that at which the DA-IMR methods are validated in Chapter 2.

Given the lack of plasma formation, we now use data from the initial growth of the bubble as well, and are no longer restricted to begin at the first bubble peak. Our data assimilation window thus includes the full first growth and collapse, so while the temporal resolution is worse than that of section 3.2, the number of data points remains large enough to converge on well-fitted material parameters. We continue to use the same bubble dynamics model described in section 2.1, and three distinct material models are tested. First, the Neo-Hookean Kelvin-Voigt model (Gaudron, Warnez, and Johnsen, 2015) used in section 3.1, where the stress integral is given by equation (2.7). Then, two variants of the quadratic law Kelvin-Voigt model (Fung, 2013; Yang, Cramer, and Franck, 2020), with the stress integral given by equation (3.3). The first variant is identical to that used previously in section 3.2, where the shear modulus is fixed at the quasi-static value G_∞ . The second generalizes this so that both G and the strain stiffening parameter α can vary. Additionally, instead of

using the long-time equilibrium bubble radius to estimate R_0 as done previously, this parameter is here estimated along with the material model parameters, with an initial estimate taken to be the initial bubble radius. Finally, a time-shift parameter t_s is added to our estimation, which is used to initialize bubble growth. This yields a new, expanded state vector for our En4D-Var estimator with up to 5 constant parameters to estimate appended with the generalized QLKV model:

$$\mathbf{x} = \{R, \dot{R}, p_b, S, \mathbf{T}, \mathbf{C}, G, \mu, \alpha, R_0, t_s\}. \quad (3.6)$$

For the neo-Hookean Kelvin-Voigt model, $\alpha = 0$ is fixed, for the quasi-static QLKV model we set $G = G_\infty$, and for the generalized QLKV model all parameters are allowed to vary. The main goal of this study is material parameter estimation, so we report results with the En4D-Var method only (as detailed in section 2.2.4). An ensemble size of $q = 48$ is used in all cases, which is found to be sufficient while keeping computational costs low. The initial ensemble is sampled from a Gaussian distribution centered around an initial state vector. The method is implemented with both gels, and for each material model.

3.3.2 En4D-Var Parameter Estimation Results

We now analyze the results obtained with En4D-Var parameter estimation. Initial guesses as to each material property are determined based on the estimates obtained with IMR in table 1 of Mancina, Yang, et al. (2021). For example, in the 0.3% gel case, the initial guesses with all models are $R_0 = 0.5 \mu\text{m}$ and $\mu = 0.1 \text{ Pa}\cdot\text{s}$. For the Neo-Hookean model, the shear modulus initial guess is $G = 10 \text{ kPa}$. For both QLKV models, the initial guess is $\alpha = 0.03$, and while the shear modulus is fixed at $G = 0.38 \text{ kPa}$ in the quasi-static case, the initial guess is $G = 0.5 \text{ kPa}$ in the general QLKV case.

Table 3.3 summarizes the results with En4D-Var for both the 0.3% and 1% agarose specimens. These results are comparable with those obtained with simple least-squares fitting done by Mancina, Yang, et al. (2021) in table 1 for these material parameters. A key difference, however, appears in the comparatively much smaller stress-free radius and stiffening parameter standard deviations. In fact, the estimates for these parameters are close to the initial guess in all cases. This indicates that given these initial guesses, the En4D-Var is unable to improve on these results and converges quickly to the original value. This behavior is further discussed below.

Model	G [kPa]	α [10^{-2}]	μ [Pa s]	R_0 [μm]
0.3% gel				
NH	9.66 ± 0.55	0	0.086 ± 0.028	0.51 ± 0.07
QS QLKV	0.38 ± 0.16	3.0 ± 0.1	0.097 ± 0.026	0.50 ± 0.03
Gen QLKV	0.50 ± 0.05	3.0 ± 0.3	0.094 ± 0.026	0.51 ± 0.06
1% gel				
NH	36 ± 3.57	0	0.14 ± 0.031	1.03 ± 0.04
QS QLKV	7.2 ± 0.33	2.4 ± 0.15	0.16 ± 0.038	1.29 ± 0.06
Gen QLKV	7.7 ± 0.89	2.5 ± 0.12	0.16 ± 0.036	1.29 ± 0.05

Table 3.3: Weighted mean and standard deviation of inferred properties using En4D-Var for 0.3% and 1% agarose specimens. Table reproduced from Mancina, Yang, et al. (2021) with permission from the Royal Society of Chemistry, © 2021.

The radius-normalized RMS errors obtained with the En4D-Var are also similar to those obtained with IMR. For example, in the 0.3% gel case, the normalized RMS errors range from 0.02 to 0.09, with a mean of 0.04 with the Neo-Hookean model. For the quasi-static QLKV model, they range from 0.02 to 0.07 with a mean of 0.05. Finally, for the general QLKV model, the range is 0.02 to 0.06 with a mean of 0.04. The increase in model complexity thus appears to better capture material behavior, reducing the error. Details on these parameter estimates can be found in sections 3 and 4.1-4.3 of Mancina, Yang, et al. (2021), but we here focus on the additional information that our En4D-Var implementation provides.

Indeed, the En4D-Var results can further inform the uncertainties associated with the material property estimates obtained, e.g., due to variations in samples.

Figure 3.13 summarizes the En4D results with the QLKV model for each material property. These histograms combine all final ensemble members across the 19 data sets (with an ensemble size of 48, the total number of ensemble members is thus $19 \times 48 = 912$). An approximately Gaussian distribution is obtained for all four quantities, the mean of which are our estimates for each quantity, thus confirming that the En4D-Var estimates are uniform across all data sets.

Figure 3.14 shows the iterative estimation of each parameter for all 19 data sets, again with the QLKV model. It appears that the viscosity estimates are relatively scattered, which points to a relatively high sensitivity to changes in viscosity in our estimator. By contrast, the spread in stress-free radius, shear modulus and stiffening parameter is quite narrow around the initial guess, indicating that in this regime

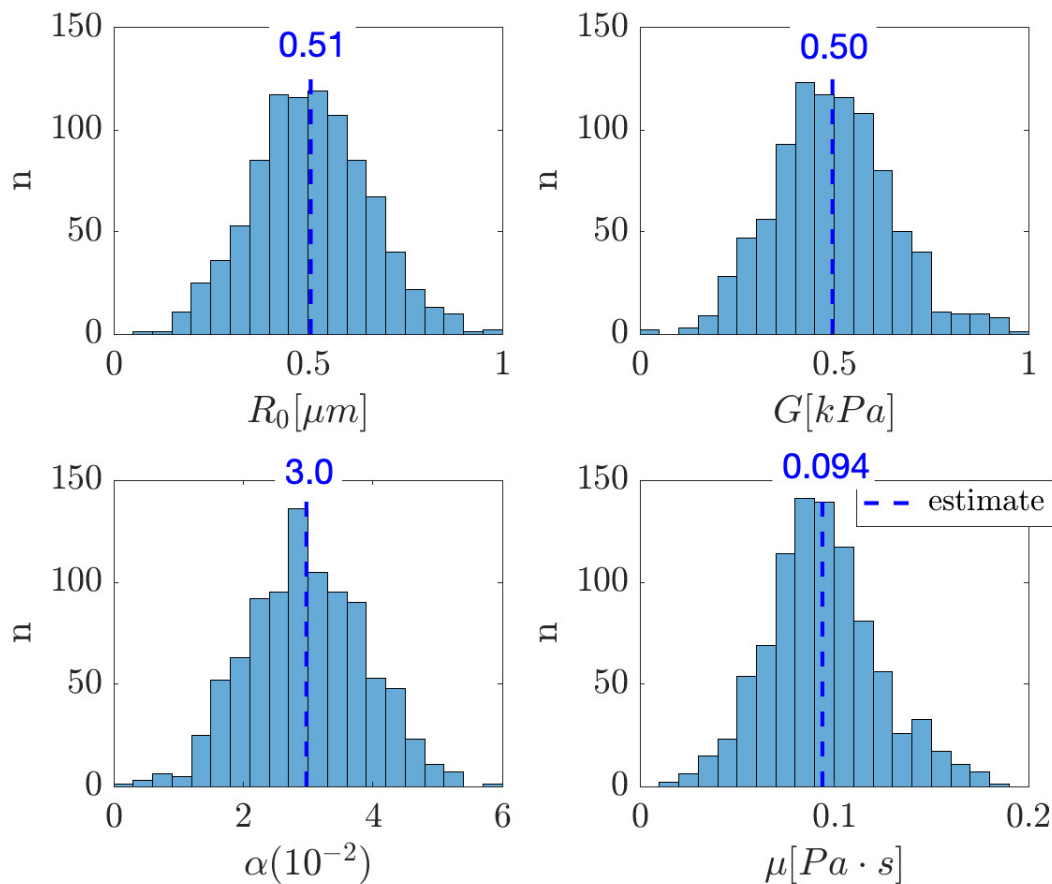


Figure 3.13: Histogram of combined final stress-free radius, shear modulus, stiffening parameter and viscosity ensembles for 0.3% gels with the En4D-var using the QLV model. Reproduced from Mancía, Yang, et al. (2021) with permission from the Royal Society of Chemistry, © 2021.

and given this guess the estimator could not improve the fit significantly. Overall, this uncertainty assessment demonstrates that the acoustic cavitation extension of IMR applied with data assimilation methods provides robust parameter estimates comparable to those obtained with the traditional IMR optimization approach.

Overall, the En4D-Var method was successfully applied to the ultrasound-cavitation framework for material estimation, or acoustic cavitation rheometry. As compared to the traditional least-squares fitting method, we obtained comparable material parameters in a manner far more scalable and computationally efficient, while additionally providing insight into modeling uncertainties. It has confirmed that acoustic cavitation rheometry is a powerful framework for viscoelastic material property estimation at high strain rates, which avoids some of the issues associated with laser-induced cavitation discussed in section 3.2. Applying En4D to this data

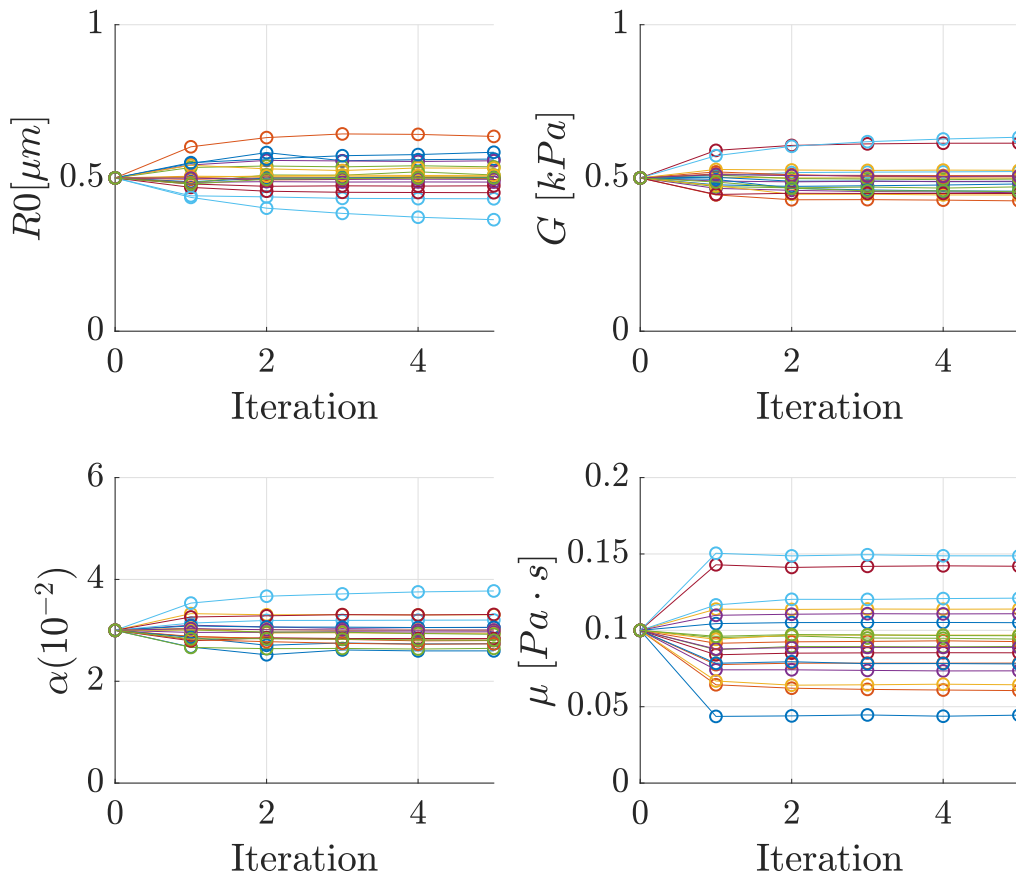


Figure 3.14: Iterative estimation of the stress-free radius, shear modulus, stiffening parameter and viscosity for 0.3% gels with the En4D-var using the QLKV model. Reproduced from Mancía, Yang, et al. (2021) with permission from the Royal Society of Chemistry, © 2021.

has shown that DA-IMR is a versatile framework, which in all tested cases has functioned well and provided additional physical insight into the bubble dynamics of the problem.

This concludes the first part of the thesis, focused on data assimilation for viscoelastic material characterization. While the interaction of ultrasound and bubble dynamics, as examined in section 3.3, play a central role in the following chapters and particularly in Chapter 5, we will now shift to a new application of interest.

*Chapter 4***A SINGLE FRAMEWORK FLUID–SOLID NUMERICAL SOLVER FOR BURST-WAVE LITHOTRIPSY SIMULATIONS**

Parts of this chapter are adapted from Spratt, Rodriguez, Bryngelson, et al. (2021), Spratt and Colonius (2023) and sections 2.2 and 6.2 of Radhakrishnan et al. (2023). In this chapter, we present our numerical solver implemented for full simulations of burst-wave lithotripsy. In section 4.1, this modeling framework is described, including the 5-equation multiphase model with hypoelasticity, a continuum damage model, a one-way acoustic wave source for the transducer, and the GPU implementation of the code. In section 4.2, our hypoelastic model implementation is validated through comparison to analytical results. In section 4.3, we present first simulations of BWL, and compare simulated stress patterns in stones to experimental photoelastic imaging. Finally, we use the solver in section 4.4 to study the frequency-dependence of stress produced in stones of various sizes and shapes during BWL.

4.1 Numerical Methods

Simulating burst-wave lithotripsy (BWL) requires modeling of various interactions. In the most simple case, we require a solver that can model acoustic (ultrasound) waves traveling through the surrounding fluid and impinging on a solid stone. The elastic response of the kidney stones must be fully represented to determine treatment efficacy and ensure correct interactions with the acoustic field and surrounding fluid. This requires modeling the elastic response of the material in the solid region. Furthermore, the BWL process frequently causes bubbles to cavitate near the proximal surface of the stone. Accounting for this requires a high-order accurate multi-phase compressible flow solver, capable of simulating acoustic wave–bubble interactions, which could include rapid bubble collapse. In all, BWL simulations thus require a multiphase fluid mechanics solver that can additionally model the elastic response of a solid material.

To meet these requirements, we use an interface-capturing method with the 5-equation multiphase flow model (Kapila et al., 2001; Allaire, Clerc, and Kokh,

2002). This is complemented with a hypoelastic model to capture the behavior of solid regions while remaining in an Eulerian framework (Rodriguez and Johnsen, 2019), and a continuum damage model to better understand damage accumulation occurring as a result of the treatment (Cao et al., 2019). Specifically, these models are added to the open-source Multi-component Flow Code (MFC) to carry out our simulations, which uses a diffuse interface method with HLL-type Riemann solvers and supports the 5-equation model (Bryngelson et al., 2021). Spherical acoustic waves generated by the BWL therapy transducer elements are simulated using the volumetric source-term approach of Maeda and Colonius (2017), which can simulate one-way waves off of any smooth surface. Finally, to fully resolve bubbles on the order $O(100 \mu\text{m})$ in our simulations while capturing the entire stone volume, we require high resolution. This is facilitated by the port of our code to GPUs, without which computational costs for the simulations presented in this paper would have been far greater. We describe each of these models and, in turn, the GPU implementation in the remainder of this section.

4.1.1 The 5-equation Model with Hypoelasticity

We use the hypoelastic material model of Rodriguez and Johnsen (2019), which is based on a small-strain approximation (Eringen, 1962) that facilitates the incorporation of elasticity in an Eulerian framework. This enables adding this model to our existing multiphase flow solver, without sacrificing any of its capabilities. While the cases shown in the present paper only include elastic materials (stones), our implementation uses the Kelvin-Voigt material model, making it generalizable to viscoelastic materials. Details on the Kelvin-Voigt material can be found in Chapters 2 and 3, where it is used to represent the viscoelastic hydrogels characterized using DA-IMR.

The MFC code (Bryngelson et al., 2021) uses the so-called 5-equation multiphase flow model (Kapila et al., 2001; Allaire, Clerc, and Kokh, 2002). *Without* hypoelasticity, it is given by (for 2 components with superscripts ⁽¹⁾ and ⁽²⁾):

$$\frac{\partial}{\partial t} \begin{bmatrix} \alpha^{(1)} \\ \alpha^{(1)}\rho^{(1)} \\ \alpha^{(2)}\rho^{(2)} \\ \rho u_i \\ \rho E \end{bmatrix} + \frac{\partial}{\partial x_j} \begin{bmatrix} \alpha^{(1)}u_j \\ \alpha^{(1)}\rho^{(1)}u_j \\ \alpha^{(2)}\rho^{(2)}u_j \\ \rho u_i u_j + p\delta_{ij} - \tau_{ij}^v \\ (\rho E + p)u_j - u_i\tau_{ij}^v \end{bmatrix} + \begin{bmatrix} -\alpha^{(1)} - K \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \frac{\partial u_j}{\partial x_j} = \mathbf{0}, \quad (4.1)$$

with $\alpha^{(k)}$ the volume fraction of the k -th component, and ρ , u , p , E the mixture density, velocity, pressure, and total energy, respectively. τ^v is the viscous stress tensor, and K is the interface compressibility term, which is important in cavitation problems to accurately represent the sound speed in mixture regions (Schmidmayer, Bryngelson, and Colonius, 2020)

$$K = \frac{\rho^{(2)}c^{(2)2} - \rho^{(1)}c^{(1)2}}{\frac{\rho^{(2)}c^{(2)2}}{\alpha^{(2)}} + \frac{\rho^{(1)}c^{(1)2}}{\alpha^{(1)}}}. \quad (4.2)$$

We now describe modifications of this model to model the elastic response of solids, using a hypoelastic material model (Rodriguez and Johnsen, 2019).

This model relies on using a Lie objective temporal derivative (which preserves thermodynamic consistency) of the stress-strain relation to transform strains to strain-rates (Altmeyer, Rouhaud, et al., 2015; Altmeyer, Panicaud, et al., 2016). Strain rates are computed as gradients of velocity $\dot{\epsilon}_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right)$, consistent with fluids. The 5-equation model given by equation (4.1) is modified by first adding an elastic shear stress term $\tau_{ij}^{(e)}$ to the Cauchy stress tensor

$$\sigma_{ij} = -p\delta_{ij} + \tau_{ij}^{(v)} + \tau_{ij}^{(e)}. \quad (4.3)$$

This elastic shear stress will appear in the momentum and energy equations. The total energy is then also modified with an elastic contribution $e^{(e)} = \frac{\tau_{ij}^{(e)}\tau_{ij}^{(e)}}{4\rho G}$, yielding

$$E = e + \frac{\|\mathbf{u}\|^2}{2} + e^{(e)}. \quad (4.4)$$

Here, e is the internal energy, and G is the medium shear modulus. Now, the Lie derivative of the elastic stresses is given by

$$\dot{\tau}_{ij}^{(e)} = \frac{\partial \tau_{ij}^{(e)}}{\partial t} + u_k \frac{\partial \tau_{ij}^{(e)}}{\partial x_k} - \tau_{kj}^{(e)} \frac{\partial u_i}{\partial x_k} - \tau_{ik}^{(e)} \frac{\partial u_j}{\partial x_k} - \tau_{ij}^{(e)} \frac{\partial u_k}{\partial x_k}. \quad (4.5)$$

Combining this, and the fact that for an isotropic Kelvin-Voigt material, $\dot{\tau}_{ij}^{(e)} = 2G\dot{\epsilon}_{ij}^{(d)}$ (see section 3.1.1 of Rodriguez and Johnsen (2019)), an evolution equation for the elastic stresses is derived and appended to our modified 5-equation model to

yield a final system of equations, given here for two materials:

$$\frac{\partial}{\partial t} \begin{bmatrix} \alpha^{(1)} \\ \alpha^{(1)}\rho^{(1)} \\ \alpha^{(2)}\rho^{(2)} \\ \rho u_i \\ \rho E \\ \rho \tau_{il}^e \end{bmatrix} + \frac{\partial}{\partial x_j} \begin{bmatrix} \alpha^{(1)}u_j \\ \alpha^{(1)}\rho^{(1)}u_j \\ \alpha^{(2)}\rho^{(2)}u_j \\ \rho u_i u_j + p\delta_{ij} - \tau_{ij}^v - \tau_{ij}^e \\ (\rho E + p - \tau_{ij}^e)u_j - u_i \tau_{ij}^v \\ \rho \tau_{il}^e u_j \end{bmatrix} + \begin{bmatrix} -\alpha^{(1)} - K \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \frac{\partial u_j}{\partial x_j} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ S_{il}^e \end{bmatrix}, \quad (4.6)$$

with

$$S_{il}^e = \rho \left(\tau_{kj}^{(e)} \frac{\partial u_i}{\partial x_k} + \tau_{ik}^{(e)} \frac{\partial u_j}{\partial x_k} - \tau_{ij}^{(e)} \frac{\partial u_k}{\partial x_k} + 2G\dot{\epsilon}_{ij}^{(d)} \right). \quad (4.7)$$

As with the original 5-equation model of MFC, the stiffened-gas equation of state (Le Métayer, Massoni, and Saurel, 2005) closes the system, where for the k^{th} component

$$p_k = (\gamma_k - 1)\rho_k e_k - \gamma_k \pi_{\infty, k}, \quad (4.8)$$

where γ is the specific heat ratio and p_i is the liquid stiffness, which are fitted to produce correct wave speeds in each material (Le Métayer, Massoni, and Saurel, 2004; Le Métayer, Massoni, and Saurel, 2005).

4.1.2 A Continuum Damage Model

We will use two metrics to determine the efficacy of treatment across our simulations. As brittle materials are expected to fail in tension, we measure the maximum tensile principal stress σ_1 in the stone over the course of a BWL pulse. In 2D, this is given by

$$\sigma_1 = \frac{\sigma_{11} + \sigma_{22}}{2} + \sqrt{\left(\frac{\sigma_{11} - \sigma_{22}}{2}\right)^2 + \tau_{12}^2}, \quad (4.9)$$

where each diagonal stress term σ_{ii} is given by $\sigma_{ii} = -p + \tau_{ii}^{(v)} + \tau_{ii}^{(e)}$.

In 3D, the maximum principal stress is obtained by finding the largest root of the cubic equation:

$$\sigma^3 - I_1\sigma^2 + I_2\sigma - I_3 = 0, \quad (4.10)$$

where the I_i terms are the stress invariants:

$$I_1 = \sigma_{11} + \sigma_{22} + \sigma_{33} \quad (4.11)$$

$$I_2 = \sigma_{11}\sigma_{22} + \sigma_{22}\sigma_{33} + \sigma_{11}\sigma_{33} - \tau_{12}^2 - \tau_{23}^2 - \tau_{13}^2 \quad (4.12)$$

$$I_3 = \sigma_{11}\sigma_{22}\sigma_{33} + 2\tau_{12}\tau_{23}\tau_{13} - \sigma_{11}\tau_{23}^2 - \sigma_{22}\tau_{13}^2 - \sigma_{33}\tau_{12}^2. \quad (4.13)$$

The maximum principal stress will indicate where and in which simulations large peak tensile stresses occur in the stone. This alone is a good predictor of stone damage.

Second, we implement a continuum damage model following Cao et al. (2019). Here, a damage state variable $D(X, t) \in [0, 1)$ is introduced, which is calculated as the integral over time

$$D(X, t) = \int_0^t (\bar{\alpha} \max(\sigma_1(X, \tau) - \sigma^*, 0))^s d\tau, \quad (4.14)$$

where σ_1 is the max. principal stress and σ^* , s , $\bar{\alpha}$ are constant model parameters determined empirically. $D = 0$ indicates a fully intact material, whereas a point in the material where $D = 1$ indicates it is fully damaged. This damage is reflected in the material properties, where in an isotropic linear elastic material, D modifies the local Young's modulus \mathcal{E} as

$$\mathcal{E}(X, t) = \mathcal{E}_0(1 - D(X, t)). \quad (4.15)$$

In our case, the material stiffness is parametrized by the shear modulus G , and given that $G = \frac{\mathcal{E}}{2(1+\nu)}$, this yields the following equation for the local time dependent shear modulus

$$G(X, t) = G_0(1 - D(X, t)), \quad (4.16)$$

where G_0 is the initial material shear modulus. Thus, if D reaches a value of $D = 1$ at a given position in the solid, the material is considered fully damaged, and the shear modulus reduced locally to $G = 0$.

We note that in the results presented in Chapter 5, this damage model is modified and implemented as a one-way model only, meaning that we do not modify local shear modulus as described in equation (4.16). While Cao et al. (2019) looked at shock-wave lithotripsy, our simulations only span a single BWL pulse. Relative to that caused by a shock wave in SWL, the stresses generated by a single BWL pulse with an originally intact stone are small. Indeed, full stone comminution usually requires at least $O(1000)$ ultrasound pulses during BWL. Because of this, the feedback incorporated in the model whereby the damage state modifies local material properties is not necessary for a single BWL pulse, and we only calculate the accumulation of damage D . Thus, the damage model can be thought of as a qualitative way of tracking the accumulation of stress. In Cao et al. (2019), a damage value of $D = 1$ meant that the stone was completely broken, but that is not the case here where damage values are used as a qualitative comparison tool.

The stones used in our simulations in Chapter 5 have identical properties to those of Cao et al. (2019). We thus utilize the same empirical model parameters in equation (4.14), except for σ^* which is reduced to 4 MPa to account for the time-scale of these simulations of BWL as compared to SWL. Given these modifications, the damage state $D(X, t) \in [0, 1)$ serves here as a measure and visualization tool of the accumulation of large maximum principal stress σ_1 over the course of the BWL pulse. Plots of this damage over a cross-section of the stone are used to compare the efficacy of BWL across simulations in Chapter 5.

4.1.3 Spherical Acoustic Wave Generation from a Transducer Source

To model the therapy transducer used in BWL, an acoustic source model capable of generating spherical ultrasound waves is necessary. To this end, we make use of the volumetric source term approach of Maeda and Colonius (2017), which can generate one-way acoustic waves from any smooth surface. In this case, we insert individual spherically focused transducer elements with a given aperture and focal length, each of which serves as a source for a spherically focused sine wave. Following section 4.2 of Maeda and Colonius (2017), source terms are added to the mass, momentum (in the longitudinal and radial direction), and energy equations in our model, given by

$$\Omega(\xi, t) = \begin{bmatrix} f(t) \\ g_z(\xi, t) \\ g_r(\xi, t) \\ \frac{c_0^2 f(t)}{\gamma-1} \end{bmatrix}, \quad (4.17)$$

where

$$f(t) = \frac{p_a}{c_0} \sin(\omega(t - t_0)) + \frac{p_a}{A} \left(\frac{1}{\omega} \cos(\omega^2(t - t_0)) - \frac{1}{\omega} \right) \quad (4.18)$$

$$g_z(\xi, t) = -p_a \sin(\omega(t - t_0)) \cos(\xi) \quad (4.19)$$

$$g_r(\xi, t) = -p_a \sin(\omega(t - t_0)) \sin(\xi). \quad (4.20)$$

Here, f is the ultrasound frequency, p_a is the pressure pulse amplitude, c_0 the speed of sound, $\omega = 2\pi f$ the angular frequency, and A the element aperture. ξ is the polar angle parametrizing the transducer element as $(x, y) = [r_0 \cos(\xi), r_0 \sin(\xi)]$. The full model derivation and examples of its use are given in Maeda and Colonius (2017).

For the simulations of Chapter 5, we implement a custom transducer surface with 18 spherically focused elements placed in two concentric rings (of 6 and 12 elements). The transducer element pressure amplitudes are calibrated to match experimental waveforms. Details are provided in the description of the problem setup in section 5.1, and in section 5.2 where a virtual array is created to simulate a transducer within a reduced domain size.

4.1.4 GPU Acceleration

For the past two decades, CPU clock speed increase has slowed and come to a plateau. Instead, much of the increase in computational power has come from increased parallelization, particularly through the use of GPUs. Modern supercomputers make effective use of GPU acceleration, which has become central to the continued scaling of computational capabilities. Given CFL number limitations that govern simulations in computational fluid dynamics, the minimization of wall time for each time step is critical. To this end, the MFC code, including all models described thus far in section 4.1, has been ported to GPUs. Details of this implementation can be found in Radhakrishnan et al. (2023), and are summarized here.

The GPU implementation is done using OpenACC (Wienke et al., 2012), which is a directive-based programming language. OpenACC directives are read at compile-time, only when compiling with GPU-enabled compilers. Thus, GPU kernels are only created when necessary hardware is available, and OpenACC directives are otherwise ignored. This enables the use of a single codebase for CPU and GPU execution, while retaining high performance comparable to using GPU-specific programming languages such as CUDA (Khalilov and Timoveev, 2021). Furthermore, OpenACC is not GPU-vendor dependent, and has increased support across various computing ecosystems, meaning a flexible and future-proof implementation as GPU architecture evolves (Jarmusch et al., 2022). Despite its relative ease of implementation, it enables custom parallelization strategies and data management between CPU host and GPU devices, serving the needs of a relatively large Fortran code such as MFC.

The results are a performant and highly scalable GPU-accelerated code. Speedup numbers vary depending on the hardware and models used. In one representative test case reported in Radhakrishnan et al. (2023), we see a 300× speedup when comparing a single NVIDIA V100 GPU to a POWER9 CPU core. In another test

case comparing a full GPU node (4 NVIDIA V100 GPUs) to a CPU node (128 cores AMD EPYC 7742) on the Expanse supercomputer at the San Diego Supercomputer Center, a $5.8\times$ speedup is achieved with the hypoelastic and damage models turned on in a node-to-node comparison (see table 4.1 and Appendix B for full results of this example). We note that this speedup would be even more significant if using NVIDIA A100 GPUs, which are approximately 1.7 times faster than NVIDIA V100s.

The code sees near-ideal weak scaling (where the problem size scales with the number of processors) up to $O(10000)$ GPUs. That is, running a test 3D two-component case with 1 million points/GPU is only 3% slower on 13824 GPUs than on a single GPU. Strong scaling (where the problem size is fixed for each run) sees MFC retain 84% of ideal performance when increasing the number of GPUs 8-fold, from 8 to 64 GPUs. This decreased performance is attributed to the increased MPI communication required as the number of GPUs increases. We note that for all GPU simulations presented in Chapter 5, the number of GPUs did not exceed 32, where scaling performance remains near-ideal (see figure 2 in Radhakrishnan et al. (2023)). Overall, the GPU implementation of the code enables runs at much higher resolutions than previously possible with the CPU code alone.

When the hypoelastic model is activated, we have seen in section 4.1.1 that up to 6 additional equations must be solved. That is, there are up to 6 distinct elastic stresses σ_{ij} in 3D, as this is a symmetric 3×3 tensor. Also, calculating the damage state following the method described in section 4.1.2 requires calculating the maximum principal stress at each point in the domain. These models thus both increase the computational cost of simulations. Table 4.1 compares run time per time-step for a 3D test case with the base version of MFC, with the hypoelastic model turned on, and with both the hypoelastic and damage models. This shows that the GPU

Version	Base MFC	Hypoelastic	Hypoelastic and Damage
CPU node (128 cores)	8.8 s	19.4 s	22.5 s
GPU node (4 GPUs)	2.0 s	3.7 s	3.9 s

Table 4.1: Total time per time-step comparison for a 3D MFC test run with 10 million grid points with two components, with different models activated.

code reduces the additional burden of the implemented models. On one CPU node, the added cost of turning on both models is 155%. By contrast, this added cost is reduced to 98% on one GPU node. The added cost of including just the hypoelastic

model is similarly lower (84% vs 119%) on GPUs. This is likely because the performance of the GPU code is more limited by routines with intense memory transfer requirements. Calculations required for both the hypoelastic and damage models are highly parallelizable, and thus constitute a smaller added cost relative to the entire time step operations on GPUs, as detailed in Appendix B.

Having described our numerical framework including implemented hypoelastic and damage models, acoustic source model, and the GPU acceleration of the code, we now turn to validation cases, particularly for the implemented hypoelastic model.

4.2 Model Validation

4.2.1 Comparison to 1D Analytical Solution

In 1D, following Gavrilyuk, Favrie, and Saurel (2008), we can obtain analytical solutions to Riemann problems in elastic media as long as ρG is kept constant. We consider two such 1D Riemann problems in elastic media taken from Gavrilyuk, Favrie, and Saurel (2008) and Rodriguez and Johnsen (2019), one of which was also reported in Spratt, Rodriguez, Bryngelson, et al. (2021). The first is an impact problem in an elastic material with shear modulus $G = 500$ MPa. There is a velocity discontinuity at the interface, and the full initial conditions are given by

$$(\rho, u, p, \tau_e) = \begin{cases} (1000, 10, 1 \times 10^5, 0) & \text{for } x \in [0, 0.5], \\ (1000, -10, 1 \times 10^5, 0) & \text{for } x \in [0.5, 1]. \end{cases} \quad (4.21)$$

The resulting density, velocity, pressure, and normal stress at $t = 64 \mu\text{s}$ are shown in figure 4.1. The simulation is run on a 1D grid with 400 points. We see good agreement between the simulation results and analytical solution.




Figure 4.1: Comparing the simulated results (blue dotted line) to the exact solution (black line) for a 1D impact problem with initial conditions described by (4.21) at $t = 64 \mu\text{s}$.

The next validation case is akin to a Sod shock tube with a 1000 : 1 pressure ratio, but with an added velocity discontinuity in the y -direction at the interface in an elastic material with high shear modulus $G = 10 \text{ GPa}$. The initial conditions are

$$(\rho, u, v, p, \tau_e) = \begin{cases} (10^3, 0, 100, 10^8, 0) & \text{for } x \in [0, 0.5), \\ (10^3, 0, -100, 10^5, 0) & \text{for } x \in [0.5, 1]. \end{cases} \quad (4.22)$$

To simulate this case with MFC, we run a 'quasi-1D' simulation with the 2D code. By this, we mean that initial conditions are uniform along the y -direction (which only has a few grid points). This way, we achieve the y -velocity discontinuity required in this example. We plot the resulting state (along the centerline) at time $t = 64 \mu\text{s}$ in figure 4.2. Density, velocity in the x -direction, pressure, normal stress, velocity in the y -direction, and elastic shear stress τ_{12} are shown for $x \in [0, 1]$. Note




Figure 4.2: Comparing the simulated results (blue dotted line) to the exact solution (black line) for the 1D Riemann problem described by initial conditions (4.22) at $t = 64 \mu\text{s}$.

that these values are plotted along the $y = 0$ axis, but are equal for any y as the initial condition is uniform in the y -direction. Five waves are observed: a shock propagating to the right, a rarefaction propagating to the left, a contact discontinuity in the center, and shear waves propagating in both directions. Simulation results show good agreement with the exact solution, notably demonstrating correct p -wave and s -wave speeds in this elastic material.

4.2.2 2D Bi-Layered Media Simulations

Next, we consider a two-dimensional simulation of bi-layered media to examine transmitted s - and p -waves through a water–elastic-medium interface. This wave transmission is vital to the understanding of the elastic response of stones to BWL. Here, an acoustic source placed in the top fluid layer, 500 m above the interface, generates one cycle of a 10 Hz Ricker wave. This waveform, which is a derivative of a Gaussian pulse, is frequently used as a surrogate for seismograph data (Ryan, 1994). The elastic material has density $\rho = 2500 \text{ kg/m}^3$ and theoretical wave-speeds $c_L = 3400 \text{ m/s}$ and $c_T = 1963 \text{ m/s}$. A probe is placed 500 m below the material interface, and 2000 m from the source in the x -direction. The full setup is shown in figure 4.3.




Figure 4.3: Sketch of the setup for a 2D bi-layered media problem, showing the location of wave source and probe as well as material properties of each layer.

This example is taken from Komatitsch, Barnes, and Tromp (2000), and thus we can compare the resulting wave pattern to that observed in their spectral-element code simulations. Figure 4.4 shows the velocity magnitudes in the domain at two distinct times, where the reflected and transmitted waves can be observed.




Figure 4.4: Simulated velocity magnitudes in bi-layered media. A source in the liquid (top half) emits a Ricker wavelet which reflects and propagates through the material interface with the elastic solid (bottom half). Resulting wave patterns are shown at two different times t as labeled.

The simulation displays all expected waves. These are labeled as follows: The initial wave emitted from the source (a), a reflected (b) and transmitted (c) p-wave, a transmitted s-wave (d), and two refracted waves (e,f). Calculated p and s-wave speeds in the elastic medium match the expected theoretical values within 0.1%, and comparing to figure 1 of Komatitsch, Barnes, and Tromp (2000) shows great alignment.

We also compare our simulation results at the probed location (shown as the blue diamond in figure 4.3) to an analytical solution. For this comparison, we use analytical results obtained with the Gar6more2D solver of Diaz and Ezziani (2010). They compute the analytical solution to wave propagation in 2D bi-layered media using the Cagniard-De Hoop method (Cagniard et al., 1963; De Hoop, 1960). Details of their implementation can be found in Diaz and Ezziani (2010). Specifically, we plot the velocity in the x and y-direction for the first 1.5 s from wave emission in figure 4.5. This shows good agreement between the probed velocities from our

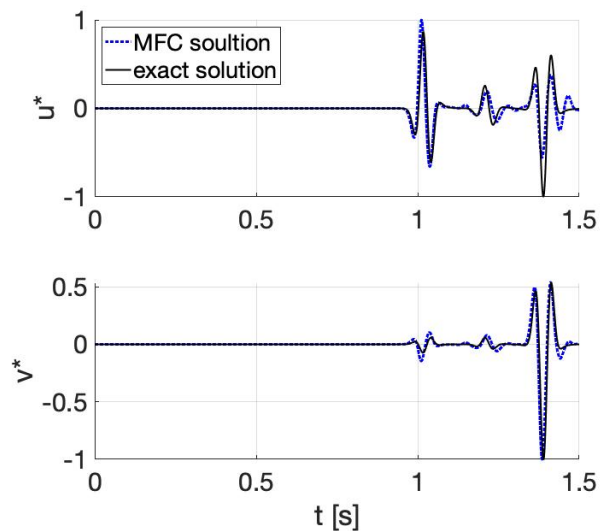


Figure 4.5: Velocity in the x (top) and y-direction (bottom) at the probe point in the elastic solid layer, located 1000 m below the source in the y-direction and 2000 m past the source in the x-direction. The probe is 500 m below the water-solid interface.

simulation and the analytical solution, showing accurate propagation of elastic waves through the interface and solid material. Three waves are visible in this plot. The first, at $t \approx 1$ s, corresponds to the transmitted p-wave labeled (c) in figure 4.4. The second, at $t \approx 1.2$ s is the refracted wave labeled (e), and the final wave at $t \approx 1.4$ s is the transmitted s-wave labeled (d).

The three validation cases presented in section 4.2 have shown the accuracy of our simulation framework, as compared to analytical results. Section 4.3 presents a validation case for a simulation of BWL, where results are compared to experimental stress measurements.

4.3 Comparison of BWL Simulations to Experimental Photoelastic Stress Images

Experimentally, it can be difficult to determine the optimal transducer setup for stone comminution. Surrogate kidney stones can be used, and time to comminution recorded in various settings, but this does not help understand elastic wave propagation patterns leading to stone damage. To address this, Sapozhnikov, Maxwell, and Bailey (2020) used an ingenious method to visualize stresses occurring in stones in real-time during experiments. For this, they make use of photoelastic stress imaging. This method works by shining a polarized light through a stone (which must be transparent, and optically isotropic) undergoing BWL. The local stresses in the stone modifies the polarization of the light shining through it, and thus the projected light displays patterns of stress inside the stone. However, this cannot be done in retrieved human kidney stones, or other typical kidney stone phantoms used experimentally, as the light beam cannot traverse these. Instead, Sapozhnikov, Maxwell, and Bailey (2020) use transparent, optically isotropic epoxy, which is prepared to match kidney stone sound speeds and density relatively well. Another advantage of these experiments, is that they are a good point of comparison for BWL simulations.

In this section, we show simulations of BWL using the numerical framework described in section 4.1, and compare them to the photoelastic images obtained by Sapozhnikov, Maxwell, and Bailey (2020) and Maxwell, MacConaghy, et al. (2020). We replicate their experimental setup and use cylindrical stones with length 20.8 mm, diameter 6.3 mm. The stone material properties are identical to the epoxy stones used: $\rho = 1100 \text{ kg/m}^3$, $c_L = 2440 \text{ m/s}$, $c_T = 1295 \text{ m/s}$. Figure 4.6 shows the setup of this simulation, where the transducer is simplified as a single spherically focused element with aperture 90 mm and focal length 70 mm, and a cylindrical stone is submerged in water. The transducer fires 20 cycles of ultrasound at 340 kHz.

In figure 4.7, we show the pressure in the surrounding liquid, and maximum principal stress σ_1 in the stone during the BWL pulse at three distinct times.




Figure 4.6: Setup of the simulation with a submerged epoxy stone on the right and simplified, single-element ultrasound transducer on the left, both submerged in water.

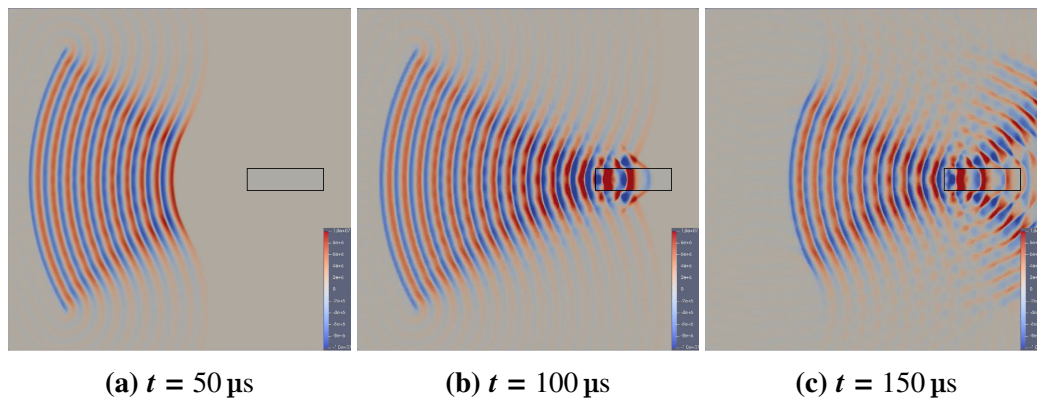


Figure 4.7: Propagation of pressure in the fluid and maximum principal stress in an epoxy stone submerged in water as a result of a 20-cycle BWL ultrasound pulse from a single-element transducer at 340 kHz, shown for three distinctive times as labeled.

We compare this to experimental photoelastic images. Using the method outlined by Sapozhnikov, Maxwell, and Bailey (2020), computational photoelastic images can be obtained from the stresses calculated in our simulations. Section 2 of Sapozhnikov, Maxwell, and Bailey (2020) derives light propagation equations dependent on local stresses. These equations are integrated over the stone cross-section to obtain a 2D distribution of light intensity, matching what would be seen in experimental photoelastic images. Details on this method can be found in sections 2 and 3 of Sapozhnikov, Maxwell, and Bailey (2020). In figure 4.8, we compare the computa-

tional photoelastic images obtained from the MFC simulation (top) to experimental images of Maxwell, MacConaghy, et al. (2020) and Sapozhnikov, Maxwell, and Bailey (2020) (bottom).




Figure 4.8: Comparison of computational photoelastic stress images in an epoxy stone obtained in our simulation (top) and experimental images obtained by Sapozhnikov, Maxwell, and Bailey (2020) which are also reported in Maxwell, MacConaghy, et al. (2020) The bottom experimental images are reprinted with permission from Sapozhnikov, Maxwell, and Bailey (2020). © 2020, Acoustical Society of America.

While quantitative comparisons with photoelastic images are challenging given the dependence of computational photoelastic modeling on empirical parameters (Sapozhnikov, Maxwell, and Bailey, 2020), good qualitative agreement between our simulation results and the experimental data is observed. This shows that our simulation framework can adequately model the behavior of elastic waves in a stone exposed to a BWL pulse. We now use the hypoelastic model to study the frequency dependence of stress patterns in stones during BWL.

4.4 Frequency–Dependence of Stresses Generated by BWL

Having validated our numerical framework, including with a BWL example, we apply it to simulate BWL at various frequencies, and in stones with varying shapes. The goal is here to determine a frequency which optimally damages stones of typical composition and size seen in human kidney stones. In Chapter 5, the problem is complicated by introducing bubbles to shield the stone, but we begin here by applying our framework in cases without bubbles present. Thus, we aim to determine a *baseline* frequency to be used going forward.

Experimental studies have explored BWL treatment frequency (Maxwell, Cunitz, et al., 2015; Maxwell, MacConaghy, et al., 2020). These have led to the generally accepted range of 350 kHz to 400 kHz, which is used in clinical trials (Harper, Metzler, et al., 2021; Harper, Lingeman, et al., 2022). By deriving an analytical model for mechanical stress in a spherical stone during BWL, Sapozhnikov, Maxwell, and Bailey (2021) explore the stone-size-dependence of maximum stress at a given frequency (or conversely, the frequency-dependence for a fixed stone size). Their study focuses on smaller stones, and shows that higher frequencies of ultrasound plane-waves increase stress in small stones. With plane waves, they show that stresses in circular stones remain approximately equal to the pressure wave magnitude when the wavelength is much larger than the stone diameter. However, as the frequency increases and the wavelength drops below twice the stone diameter, stress in the stone can increase to many times this value. Here, we calculate stresses generated in circular and rectangular stones at various ultrasound frequencies guided by their findings, but using our numerical framework with a spherical transducer element, and stone composition matching what will be used for full BWL simulations in Chapter 5.

Given the large beam-width (compared to the stone diameter) in BWL, we expect our results with the spherical waves to closely match the observations of Sapozhnikov, Maxwell, and Bailey (2021). They find that for calcium oxalate monohydrate (COM) stones with $\rho = 1823 \text{ kg/m}^3$ $c_L = 4476 \text{ m/s}$ $c_T = 2247 \text{ m/s}$, large stresses are generated for ultrasound frequencies around 350 kHz for 5mm diameter stones, and 175 kHz for 10mm stones. We note that higher stress peaks are observed at higher frequencies, but with narrow frequency bands. In practice, using the listed frequencies where high stresses are sustained over a larger frequency range is preferable. In a setup matching that of figure 4.6, we expose circular stones with a 5mm and 10mm diameter to three ultrasound frequencies ranging from 175 kHz to 700 kHz. Table 4.2 reports values for $\max(\sigma_1)/p_0$, the maximum principal stress over the simulation normalized by the focal pressure, which is calibrated to $p_0 = 6.5 \text{ MPa}$ for all tested frequencies.

	frequency [kHz]	$\max(\sigma_1)/p_0$
5 mm stone	175	2.18
	350	3.22
	700	2.17
10 mm stone	175	3.18
	350	1.97
	700	2.80

Table 4.2: Maximum value of the normalized maximum principal stress over a BWL simulation with a single spherical element at three fixed frequencies, for circular stones with diameter 5 mm and 10 mm; bold values represent the maximum stress values in each case.

The trend observed by Sapozhnikov, Maxwell, and Bailey (2021) is also seen in our results, with the largest stresses observed at 350 kHz for the smaller stone, and 175 kHz for the larger stone. $\max(\sigma_1)$ fields in the stone for each of these cases are shown in figure 4.9 for the 5 mm-diameter stone, and figure 4.10 for the 10 mm-diameter stone.




Figure 4.9: Maximum value of σ_1 in a 5 mm diameter circular stone over the course of a BWL simulation with a single spherical transducer element at three test frequencies.

Large areas of high tensile stress are observed at 350 kHz for the 5 mm stone, and at 175 kHz for the 10 mm stone, as expected. We note the interesting pattern observed for the 10 mm stone at 700 kHz in figure 4.10c. This case, which table 4.2 also showed has a relatively large $\max(\sigma_1)$ value, corresponds to one of the narrower stress peaks discussed in Sapozhnikov, Maxwell, and Bailey (2021). At this frequency, resonance in the stone is causing localized stress peaks of high magnitude. However, in practice, this regime is difficult to reach unless stone size




Figure 4.10: Maximum value of σ_1 in a 10 mm diameter circular stone over the course of a BWL simulation with a single spherical transducer element at three test frequencies.

is precisely known a priori. Furthermore, while the peak stress values in this case are high, they are far more localized.

To examine the shape-dependence of these results, we calculated the equivalent $\max(\sigma_1)$ values in rectangular stones, which are shown in table 4.3

	frequency [kHz]	$\max(\sigma_1)/p_0$
5 × 5 mm	175	2.35
	350	2.38
	700	1.77
10 × 5 mm	175	4.47
	350	1.85
	700	1.54

Table 4.3: Maximum value of the normalized maximum principal stress over a BWL simulation with a single spherical element at three fixed frequencies, for rectangular stones with width 5 mm, and length 5 mm and 10 mm, respectively; bold values represent the maximum stress values in each case.

The same trends are observed, with highest values of $\max(\sigma_1)$ at 350 kHz in the small stone, and 175 kHz in the large stone. Representative stress patterns are shown for the 5 mm length stone in figure 4.11.

While the $\max(\sigma_1)/p_0$ values reported for this rectangular stone in table 4.3 were similar for 175 kHz and 350 kHz, figure 4.11 shows a preferable stress pattern with the higher frequency of 350 kHz, where the peak stresses occur across a large region of the stone, centered in the middle of the stone.

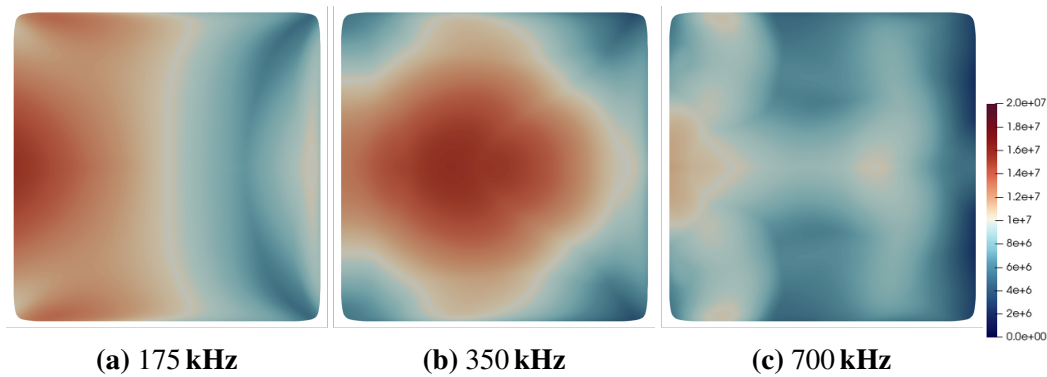


Figure 4.11: Maximum value of σ_1 in a 5 mm length rectangular stone over the course of a BWL simulation with a single spherical transducer element at three test frequencies.

Overall, these results have shown that previous experimental results, as well as theoretical results with plane waves, appear to hold for BWL with spherically focused ultrasound waves. Moreover, these results hold for different stone shapes, with length-scale seemingly the most important factor to predict frequencies causing larger peak stresses. While we see high peak stresses at higher resonant frequencies in certain cases, as seen with the 10mm diameter stone at 700 kHz, a frequency near 350 kHz is most reliable, as Sapozhnikov, Maxwell, and Bailey (2021) showed that high stresses are generated in a broader band of frequencies around this value. In subsequent simulations in Chapter 5, we use 375 kHz as our baseline frequency, which is near this 350 kHz value, but appeared in testing to yield slightly higher stresses in the material used, which has higher wave-speeds. Given that higher frequencies yield higher stresses in smaller stones, testing at a slightly higher frequency is also advantageous as it will continue to deliver high stress to the stone as it fragments and its length scale is reduced.

*Chapter 5*NUMERICAL INVESTIGATION OF BUBBLE DYNAMICS AND
DAMAGE IN BURST-WAVE LITHOTRIPSY

Parts of this chapter are adapted from Spratt and Colonius (2023). In Chapter 4, we presented BWL simulation results involving the interaction of focused ultrasound and stones. However, as discussed in section 1.6 of Chapter 1, it has been observed that bubbles can nucleate during lithotripsy and impact treatment efficacy (Maeda, Kreider, et al., 2015; Maeda, Maxwell, et al., 2018; Maeda, Colonius, et al., 2018; Hunter et al., 2018). In particular, Maeda, Maxwell, et al. (2018) looked at energy shielding of bubble clouds over single pulses and saw that they can reduce delivery of acoustic energy to the stone by as much as 90%. In this chapter, using the numerical setup described in detail in Chapter 4, we quantify the impact of shielding on wave propagation and the resulting damage within the stone. Additionally, we investigate further strategies to minimize shielding by modulating the ultrasound frequency of individual transducer elements, introducing a secondary low-frequency ultrasound wave that causes bubbles to collapse ahead of the stone. We examine different configurations of this strategy, aiming to maximize the efficiency and efficacy of treatment in the presence of bubbles.

As an outline of the chapter, we begin by describing the problem setup in section 5.1. Then, the creation of a virtual transducer array to replicate waves generated from a transducer with a much smaller acoustic source to reduce computational costs is detailed in section 5.2. In section 5.3, we use this framework to show the impact that the presence of bubbles near the proximal part of the stone has on the stresses and damage generated in the stones. Section 5.4 examines bubble dynamics when modulating transducer element frequencies according to two distinct strategies, to attempt to encourage bubble collapse. In section 5.5, we run full (wave–bubble–stone) BWL simulations once again with these shielding mitigation strategies to test their efficacy. Results are discussed in section 5.6.

5.1 Problem Setup

A simple sketch of a side-view of the experimental setup for BWL which our simulations replicate is shown in figure 5.1. We use a spherically focused transducer with focal length $F = 150$ mm and total aperture is $A = 180$ mm. The transducer is focused on the center of a cylindrical stone with length 6 mm and diameter 5 mm. The stone (and transducer array) are submerged in water. The sketch is not drawn to scale, as the transducer aperture and focal length are much larger than the stone length scale. In section 5.2, we provide more details on the transducer, and show how a virtual array is created to replicate it with a much smaller aperture and focal length. This way, we minimize computational costs by reducing the required domain size.




Figure 5.1: Simplified sketch of the setup used for BWL simulations throughout Chapter 5.

The kidney stones in our simulations are modeled using the material parameters of a BegoStone (Liu and Zhong, 2002), an artificial surrogate for kidney stones, which has been used in various experimental studies of SWL and BWL (Maxwell, Cunitz, et al., 2015; Zwaschka et al., 2018). The properties of BegoStone fall within the range of density and sound speeds seen in human kidney stones. In our simulations, we use a stone density of $\rho = 1995 \text{ kg/m}^3$, and longitudinal and transverse wave speeds $c_L = 4159 \text{ m/s}$, $c_T = 2319 \text{ m/s}$ to match that of Cao et al. (2019) for a BegoStone prepared as detailed by Esch et al. (2010). The cylindrical stone has a length of 6 mm and a diameter of 5 mm, and is submerged in water.

Finally, to model a bubble cloud in front of the stone, we randomly distribute air bubbles with radii following a log-normal distribution with a mean of $250 \mu\text{m}$. The standard deviation is determined through observation of experimental and simulated bubble clouds in Maeda, Maxwell, et al. (2018). Our setup is agnostic as to how bubbles formed, and we begin simulations with air bubbles already present. Given the small time-scale of simulations, bubbles will only collapse to a certain point. While bubbles nucleated during BWL may contain some water vapor, the dynamics of air bubbles will adequately capture these dynamics up to final collapse.

Maeda, Maxwell, et al. (2018) performed a number of simulations with varied initial bubble cloud thickness h and void fraction β_0 . Using a subgrid bubble model, they excite bubble nuclei with ultrasound analogous to that used in BWL, and track the evolution of these bubble clouds via their total area A when projected on a 2D plane. They observe that the bubble cloud projected area plateaus near a value dependent on the initial thickness and void fraction of the bubble cloud. Furthermore, this projected area value correlates to shielding factor, with larger values of A corresponding to higher shielding. Given the use of resolved bubbles in our simulations, we do not simulate the growth phase of bubbles from the original nuclei, and instead initialize our bubble clouds with a given projected area corresponding to these plateau values. The projected area value $A = 2 \text{ mm}^2$ used in the following sections corresponds to a case with initial void fraction of $\beta_0 = 8 \times 10^{-5}$. A denser bubble cloud with projected area $A = 3 \text{ mm}^2$ is also considered. Given the results of Maeda, Maxwell, et al. (2018), both bubble clouds are expected to shield the stone effectively.

We set the mean bubble radius of $250 \mu\text{m}$ to match the upper bound for radii values seen in experiments and simulations of BWL with bubbles present (Maeda, Maxwell, et al., 2018; Maeda, Colonius, et al., 2018), while remaining large enough to be resolved in our simulations. This enables a resolution of ~ 50 grid points per bubble diameter. Furthermore, while small bubbles readily oscillate when exposed to ultrasound at treatment frequencies, the dynamics of these larger bubbles are less sensitive to these waves, as will be shown in simulations in section 5.4. The persistence of larger bubbles over multiple ultrasound pulses is a likely cause for reduced treatment efficacy, which our setup can examine. Standoff distance of the bubble cloud is chosen to match observed values in Maeda, Maxwell, et al. (2018). A representative setup for a full simulation with a bubble cloud present is shown in the next section, in figure 5.4b.

5.2 Virtual Transducer Array Calibration

To accurately capture the focused ultrasound waves used experimentally (Maeda and Colonius, 2017; Maeda, Maxwell, et al., 2018), we model an 18-element transducer, which consists of an inner ring of 6 elements, and an outer ring of 12 elements. In order to maximize resolution and capture bubble dynamics accurately while reducing computational expense, our simulations use a virtual transducer array, which is a spherical projection of the full transducer array with a much reduced aperture and focal length of $A/12$ and $F/12$. The virtual array, still composed of 18 elements, retains the correct incident wave angles to the stone and is calibrated to account (in an approximate way) for diffraction effects (edge waves) associated with each transducer element.

To calibrate this virtual array, we first ran a large simulation with a full 18-element transducer with focal length $F = 150$ mm and aperture $A = 180$ mm firing at 375 kHz. This simulation was calibrated to achieve a peak negative focal pressure of -7 MPa. Thinking of the transducer as two concentric rings, there are four edge waves to consider: one traveling inward and one traveling outward radially from each concentric ring. Figure 5.2 shows this for a 2D example, with each of the edge waves labeled 1 through 4. Wave 4 propagates out of the domain, and thus needn't be considered. What remains is to understand the effects of the first 3 waves. We hypothesized that waves 1 and 3, as they propagate inward, would create a focusing effect and increase the pressure near the centerline between transducer and stone center. Conversely, as wave 2 propagates away radially, we conjectured that this would deflect waves from the outer ring of the transducer away from the stone. With a smaller virtual array, these edge waves have far less time to develop before the ultrasound reaches the stone, so we reconstruct these effects through calibration of the input pressure of each transducer element.

We split the virtual transducer into its inner and outer ring, and varied the pressure amplitude for each ring. We gradually increased the amplitude for the six inner elements, and decreased that of the twelve outer elements, while maintaining a peak focal pressure of -7 MPa through trial and error. At a pressure amplitude ratio of approximately 1.7 : 1 from outer-ring to inner-ring elements, we reached good agreement between probed pressures across the stone region. Figure 5.3 shows the agreement between the probed pressures at three locations corresponding to the




Figure 5.2: 2D example of the 4 edge waves generated by transducer elements, for a transducer with focal length 150 mm and aperture 180 mm with 4 elements. Half of the domain is shown (positive y half-domain).

stone center, edge, and center of the proximal surface for the full and reduced-size transducer for 5 cycles of a 375kHz BWL pulse.




Figure 5.3: Comparison of probed pressures for the full transducer (black solid line) and the scaled-down transducer used in simulations (dotted colored line) at three locations corresponding to the stone center (a), stone edge (b), and center of the proximal stone surface (c).

We note, for clarity, that these simulations are run without a stone, but use the coordinates of these locations on the stone (when present) to probe the pressure. Good agreement is observed between the waveforms over the 5 cycles. While this represents a small difference over the 20-cycle ultrasound pulse used, we do see that the positive focal pressure peaks for the leading and trailing cycles are overestimated by the scaled-down transducer. The resulting scaled-down 18-element virtual array used in simulations is shown in figure 5.4, including in the context of a full BWL simulation in figure 5.4b.

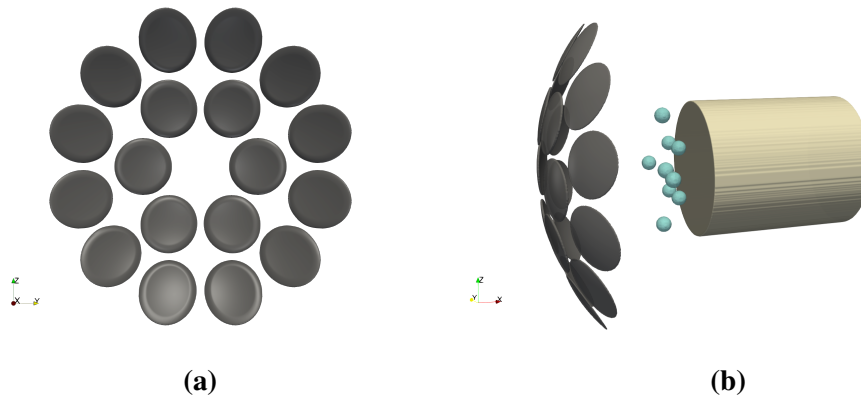


Figure 5.4: Virtual array used in BWL simulations viewed face-on (a), and from the side in the context of the full setup of a BWL simulation (b).

5.3 Bubble Shielding during Burst-Wave Lithotripsy

In this section, we utilize a 20-cycle pulse at 375 kHz and compare the results with and without bubbles, in order to corroborate past evidence on shielding while quantifying the effect of bubbles on wave propagation (and damage) within the stone. The simulations shown in this section are 3D, quarter-domain simulations with symmetric boundary conditions in the y and z-directions. As a measure of the performance of our GPU-accelerated numerical framework described in section 4.1.4 of Chapter 4, running on 64 (NVIDIA A40) GPUs on the Delta supercomputer at the National Center for Supercomputing, such a simulation takes approximately 20 hours to run.

Figure 5.5 shows the propagation of waves in the fluid and stone at various times during the simulation. Pressure is plotted in the liquid along a slice of the domain in the y-normal direction, whereas the maximum principal stress σ_1 and continuum damage field D are shown in the top and bottom halves of the stone, respectively. The final damage state in the stone can be seen in figure 5.5d, with the maximum value over time of the maximum principal tensile stress $\text{Max}(\sigma_1)$ in the top half of the stone, and the final damage field D in the bottom half.

We see that in the absence of bubbles, high stress magnitudes are seen across large areas of the stone, and we see two regions of high damage near the stone centerline.




Figure 5.5: Propagation of the maximum principal stress in the stone and damage caused during BWL at 375kHz when no bubbles are present. Figures (a), (b), (c) show three different time steps, and figure (d) shows maximum value of σ_1 in the top half, and damage field D in the bottom half at the end of the simulation.

We now introduce a bubble cloud with a projected area of 3 mm^2 in front of the stone (which we will refer to as the ‘dense’ bubble cloud), with the closest bubbles at a standoff distance $250 \text{ }\mu\text{m}$, and bubbles randomly distributed along the entire radius of the stone. The resulting maximum principal stress and damage in the stone, as well as bubble cloud oscillations, are shown in figure 5.6. Figure 5.6d shows that the presence of a bubble cloud influences the damage magnitude and pattern in the stone, as acoustic energy can no longer be effectively delivered to the center of the stone. The magnitude of maximum principal stress and damage to the stone is reduced by 16 % and 83 %, respectively (when averaged over the stone volume). We observe slight deformation in the bubbles ahead of the stone over the course of the pulse, but they remain mostly intact.




Figure 5.6: Propagation of the maximum principal stress in the stone and damage caused during BWL at 375kHz with a dense bubble cloud ($A = 3 \text{ mm}^2$) present. Bubbles are shown with 2 contour levels for the volume fraction of air: 0.2 (low opacity) and 0.8 (higher opacity). Figures (a), (b), (b) show three different time steps, and figure (d) shows maximum value of σ_1 in the top half, and damage field D in the bottom half at the end of the simulation.

To determine the effect that bubble cloud void fraction has on the stress and damage to the stone, we simulated a second test case with a sparser bubble cloud of projected area 2 mm^2 . Figure 5.7 shows the final damage state in this case. The pattern is similar to that seen in the previous case in figure 5.6d, though as expected the stress magnitude and final damage are larger.

To quantify the shielding effects of these bubble clouds, we compare the mean damage in the stone $\frac{1}{V} \iiint_V D$ and mean maximum value of σ_1 over the course of the simulation: $\frac{1}{V} \iiint_V \max(\sigma_1)$, where V is the stone volume. Values are reported in table 5.1. We can see the expected trend, with far higher damage occurring in




Figure 5.7: Damage to the stone during BWL with default transducer parameters and a sparser bubble cloud ($A = 2 \text{ mm}^2$) present. Maximum value of σ_1 over time is plotted in the top half of the stone, and damage field D in the bottom half.

the case with no bubbles present, and an increase in shielding as bubble cloud void fraction increases.

Case	$\frac{1}{V} \iiint_V D$	$\frac{1}{V} \iiint_V \max(\sigma_1)$
No Bubbles	0.052	7.45 MPa
Sparse bubble cloud	0.032	7.18 MPa
Dense bubble cloud	0.009	6.28 MPa

Table 5.1: Comparison of resulting average damage and $\max(\sigma_1)$ values over the stone region for three BWL cases with: no bubbles, a sparse bubble cloud ($A = 2 \text{ mm}^2$) and a dense bubble cloud ($A = 3 \text{ mm}^2$). All cases are with 375 kHz ultrasound.

5.4 Simulations of Bubble Dynamics during BWL

In this section, we aim to understand frequency-typical bubble dynamics in BWL. We begin with two simplified cases to better understand expected behavior of bubbles during treatment. In the first, a simple 1D spherical bubble dynamics model is used, similar to the bubble dynamics model used with DA-IMR in Chapters 2 and 3, to get an idea of the frequency-dependence of bubble dynamics for bubbles of sizes typical of BWL. Next, we will use full 3D simulations of plane ultrasound waves interacting with bubbles and bubble clouds to test hypotheses given by the 1D model. Finally, we conclude this section by using the virtual transducer array

described in section 5.2 to model the interactions of spherically focused waves from a multi-element transducer with bubble clouds typical of BWL.

5.4.1 1D Keller-Miksis Model

In the cases simulated in section 5.3, we saw only small oscillations in the bubbles as the ultrasound waves traveled through the cloud. Here, we explore different modifications to the BWL waveform that affect the bubble dynamics more significantly, hoping to collapse or shrink the bubbles, and thus increase damage to the stone by reducing the void fraction of the bubble cloud. Ideally, sufficiently violent collapse of bubbles can be induced, leading to breakup into smaller bubbles and ultimately their dissolution. Given the size of the bubbles ($\approx 250\mu\text{m}$), a lower ultrasound frequency will be needed to observe large oscillations in the bubble radius over the time scale of a BWL pulse. To remain within the capabilities of ultrasound transducers, however, we cannot reduce the frequency below $O(10\text{ Hz})$. As a first approximation of expected bubble dynamics, we implement a solver for spherical bubble dynamics using the Keller-Miksis equation (Keller and Miksis, 1980; Brennen, 1995; Spratt, Rodriguez, Schmidmayer, et al., 2021), which accounts for an incident plane wave with amplitude P and angular frequency ω :

$$\begin{aligned} \ddot{R} \left(\frac{4\mu}{\rho} - R(\dot{R} - c) \right) = & \frac{1}{2} \dot{R}^3 + \dot{R} \Delta(R) - c \left(\frac{3}{2} \dot{R}^2 + \frac{4\mu \dot{R}}{\rho R} + \frac{2s}{\rho R} - \Delta(R) \right) \\ & + R \dot{R} \Delta'(R) + \left(1 + \frac{\dot{R}}{c} \right) \frac{Pc}{\rho} \sin(\omega) \left(t + \frac{R}{c} \right), \end{aligned} \quad (5.1)$$

where R is the bubble radius, c the sound speed and s the surface tension. Δ is defined as $\Delta(R) = \frac{1}{\rho}(p_b(R) - p_\infty)$, with p_b the bubble pressure assumed uniform, and p_∞ the freestream pressure. We use this model to simulate the response of a $250\mu\text{m}$ bubble to several ultrasound frequencies, ranging from 375 kHz to 20 kHz. The bubble radius time history over $30\mu\text{s}$ for representative frequencies is shown in figure 5.8.

Here, we see that while the therapy frequency of 375 kHz only causes very small oscillations of the bubble, lower frequencies induce a more violent collapse. A value of 20 kHz is found to be effective at collapsing bubbles around $250\mu\text{m}$ over the course of a BWL pulse, while remaining high enough to be within the range of existing ultrasound transducers. Thus, this frequency is chosen for our low-frequency 'bubble collapsing' wave. In section 5.4.2, we turn to fully 3D plane-




Figure 5.8: Radius time histories for 250 μm bubble under ultrasound at various frequencies with an amplitude of 1 MPa.

wave–bubble simulations to confirm whether this result obtained under a spherical bubble assumption holds for a fully resolved 3D bubble.

5.4.2 3D Plane-Wave–Bubble Simulations

We first compare the oscillations of a single bubble under a plane wave at 375 kHz and 20 kHz with a magnitude of 3 MPa. This is shown in figures 5.9 and 5.10. Different bubble dynamics are observed. As expected, the low frequency wave collapses the bubble early in the simulation, which rebounds but remains small throughout. By contrast, the bubble radius oscillates only slightly when exposed to the 375 kHz ultrasound. This matches what was expected based on the Keller-Miksis model simulations and the radius curves shown in figure 5.8.

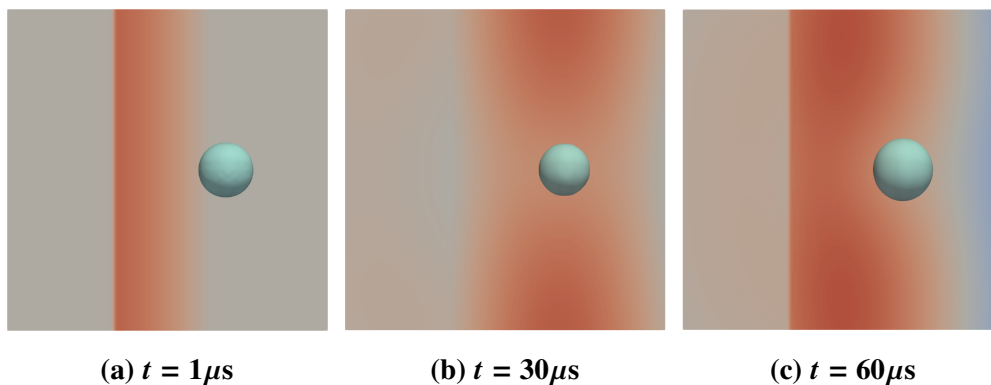


Figure 5.9: Oscillation of a 250 μm bubble under a plane wave at 375 kHz.

We also simulate plane waves interacting with the dense bubble cloud used in section 5.3. The dynamics of the bubble cloud are such that there is some deformation of




Figure 5.10: Oscillation of a 250 μm bubble under a plane wave at 20 kHz.

the bubbles in both cases, but the overall results remain that the low frequency wave more effectively collapses and deforms bubbles in the cloud. Figure 5.11 shows the behavior of a bubble cloud under 375 kHz ultrasound. In Figure 5.12, a 20 kHz pulse of equal amplitude is superposed on the 375 kHz pulse. This causes large deformation in the bubble cloud, with many bubbles largely collapsed by the end of the simulation in the dual-frequency case, with a contour level of $\alpha^{(air)} = 0.9$




Figure 5.11: Oscillation of a dense bubble cloud ($A = 3 \text{ mm}^2$) under a plane wave at 375 kHz.

This shows that adding a low frequency 'bubble collapsing pulse' to the usual 375 kHz 'therapy pulse' may help reduce shielding by collapsing bubbles and thus reduce the void fraction of the cloud. While this appears to be effective, further study of cloud bubble dynamics under ultrasound pulses could help formulate other strategies. Given the high dependence of bubble dynamics on bubble radius and characteristics of bubble clouds, determining a systematically optimal frequency




Figure 5.12: Oscillation of a dense bubble cloud ($A = 3 \text{ mm}^2$) under the superposition of plane waves at 375 kHz and 20 kHz.

may prove difficult, as these characteristics vary case-by-case during BWL. By choosing this low frequency, and setting the phase of the ultrasound so that the pressure pulse begins as positive, we can ensure the fast collapse of bubbles across a range of sizes and for different bubble cloud configurations. Thus, we reduce the need for fine-tuning, and opt for a strategy that should be broadly applicable given various bubble cloud configurations.

With a full 18-element therapy transducer, there are multiple ways this low frequency pulse can be added. A first approach is to simply superpose the 20 kHz sine wave signal to each transducer element. Alternatively, individual elements can be controlled to fire at a single frequency. Both approaches are examined in section 5.4.3 for transducer–bubble–cloud simulations.

5.4.3 Spherically focused Ultrasound–Bubble Cloud Simulations

To examine the dynamics of a bubble cloud in an acoustic field akin to that of BWL, the simulations presented in this section make use of the virtual array described in section 5.2. In all these simulations, the denser bubble cloud with projected area 3 mm^2 is used, where the interaction between bubbles in the cloud has a larger effect on its dynamics. Three cases are shown. First, a reference simulation with only the 375 kHz ultrasound is shown in figure 5.13.




Figure 5.13: Oscillation of a dense bubble cloud ($A = 3 \text{ mm}^2$) exposed to BWL transducer ultrasound at 375 kHz.

Then, a case where a 20 kHz ultrasound with a focal amplitude of 1 MPa is superposed on each transducer element, which we will call the ‘dual frequency configuration,’ is shown in figure 5.14. Finally, a simulation where the 6 inner ring elements fire the 20 kHz pulse, while the outer ring fires at the usual 375 kHz, referred to as the ‘split frequency configuration,’ is shown in figure 5.15. Figure 5.16 shows this split between inner ring and outer ring, with elements colored in orange at 20 kHz and those in blue at 375 kHz. In all tested cases, focal pressure was calibrated not to exceed a peak of 7 MPa to remain within a safe treatment regime.

The first thing to note is the similarity between the bubble cloud dynamics in the plane-wave and transducer cases, apparent when comparing figures 5.13 and 5.14 to the results of section 5.4.2. The notable difference being that the cloud collapses faster under the incident spherical wave, mostly due to higher pressures near the focus of the transducer. The deformation of the bubbles in the dual and split-frequency cases is similar, though the final bubble cloud shape is slightly different. In the dual-frequency case, the cloud is slightly more compact and concave to the right, while it is more spread out and concave to the left in the split-frequency case. While this appears to be a minor difference, we will see in section 5.6 that it has an impact on the shielding of the cloud.




Figure 5.14: Oscillation of a dense bubble cloud ($A = 3 \text{ mm}^2$) exposed to BWL transducer ultrasound in the dual frequency configuration.

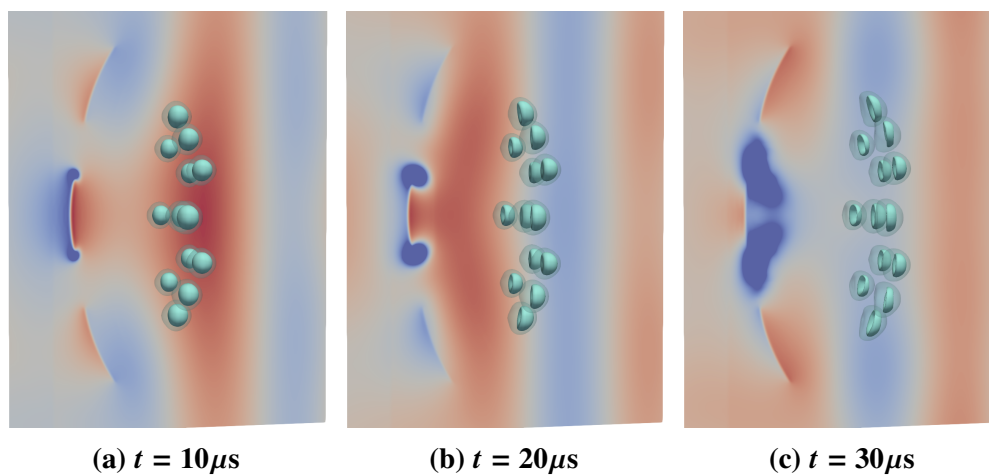


Figure 5.15: Oscillation of a dense bubble cloud ($A = 3 \text{ mm}^2$) exposed to BWL transducer ultrasound in the split frequency configuration.

Overall, both the split-frequency and dual-frequency configurations are effective at collapsing the bubble clouds quickly during the course of the ultrasound pulse. From these transducer–bubble cloud simulations alone, it appears that both strategies could be viable to increase the damage done to the stone in BWL in the presence of a bubble cloud. In the following section 5.5, we examine both strategies in the context of a full transducer–bubble–cloud–stone simulation, and compare the damage to the stone with each configuration.




Figure 5.16: Split frequency transducer configuration: orange elements fire at 20 kHz, blue elements at 375 kHz.

5.5 Dual and Split-Frequency BWL Simulations

Adding the stone back into the simulations, we now carry out full 3D simulations of BWL with the two strategies described in the previous section. We begin by looking at the dual-frequency configuration. While there may be similarities with the bubble cloud behavior observed in section 5.4.3, the reflection of acoustic waves off of the surface of the stone impacts bubble cloud dynamics. Figure 5.17 shows the propagation of σ_1 and damage state D throughout the simulation, as well as the evolution of the bubble cloud in the dual-frequency case. We see that as with the simulation shown in figure 5.14, bubbles are significantly deformed, though bubbles deform and collapse faster in this case due to the higher pressure magnitudes experienced in the bubble cloud due to reflected waves off the stone.

However, as seen in figure 5.17d which shows the final damage state, this configuration does not improve damage to the stone compared to the case with only the 375 kHz wave shown in figure 5.6d. With this transducer configuration, the bubble cloud continues to effectively shield the stone, and it is thus not an effective method to reduce shielding.

In contrast, we now turn to the split-frequency configuration. Figure 5.18 shows a time history of this simulation with a dense bubble cloud. We see that there is again a significant deformation of bubbles and a reduction in void fraction. In this case, though, the bubbles are stretched and pushed out radially, away from the center of the proximal stone surface. This difference in bubble cloud dynamics appears to




Figure 5.17: Propagation of the maximum principal stress in the stone and damage caused during BWL with the dual-frequency configuration with a dense bubble cloud ($A = 3 \text{ mm}^2$) present. Bubbles are shown with 2 contour levels for the volume fraction of air: 0.2 (low opacity) and 0.8 (higher opacity). Figures (a), (b), (c) show three different time steps, and figure (d) shows maximum value of σ_1 in the top half, and damage field D in the bottom half at the end of the simulation.

have an important effect on stone shielding, as the damage to the stone is much more significant in this case.

Indeed, figure 5.18d shows the high damage and large maximum magnitude of σ_1 observed in this case. We also simulated this configuration with the sparse bubble cloud, with similar results apart from minor differences in the damage patterns. The split-frequency configuration is far more effective than the dual-frequency case, causing high damage and stress magnitudes in the stone.

Table 5.2 compares average damage field value D and maximum σ_1 in the stone over the course of the simulation in the case with only the 375 kHz ultrasound and




Figure 5.18: Propagation of the maximum principal stress in the stone and damage caused during BWL with the split-frequency configuration with a dense bubble cloud ($A = 3 \text{ mm}^2$) present. Bubbles are shown with 2 contour levels for the volume fraction of air: 0.2 (low opacity) and 0.8 (higher opacity). Figures (a), (b), (c) show three different time steps, and figure (d) shows maximum value of σ_1 in the top half, and damage field D in the bottom half at the end of the simulation.

the split-frequency case. In both cases, the split frequency significantly improves damage to the stone. This is particularly effective with the dense bubble cloud, where we observe a nine-fold increase in the mean damage value in the stone, and a 25% increase in the mean value of $\max(\sigma_1)$. This method thus appears to be highly effective at reducing bubble cloud shielding in both tested cases.




Figure 5.19: Damage to the stone during BWL with split-frequency configuration and a sparse bubble cloud ($A = 2 \text{ mm}^2$) present. Maximum value of σ_1 over time is plotted in the top half of the stone, and damage field D in the bottom half.

Case	wave form	$\frac{1}{V} \iiint_V D$	$\frac{1}{V} \iiint_V \max(\sigma_1)$
Sparse bubble cloud	BWL at 375 kHz	0.032	7.18 MPa
	split-frequency	0.080	7.95 MPa
Dense bubble cloud	BWL at 375 kHz	0.009	6.28 MPa
	split-frequency	0.082	7.84 MPa

Table 5.2: Comparison of resulting average damage and $\max(\sigma_1)$ values over the stone region between the baseline BWL 375 kHz waveform case and dual frequency case for a sparse bubble cloud ($A = 2 \text{ mm}^2$) and a dense bubble cloud ($A = 3 \text{ mm}^2$).

5.6 Discussion

For both the dual-frequency and split-frequency configurations, bubble cloud void fraction reduction occurs on the same time scale, as shown in sections 5.4.3 and 5.5. There is a difference, however, in the dynamics and deformation of the bubbles. It appears that the split frequency configuration causes favorable conditions for the propagation of acoustic waves through the bubble cloud. In figure 5.20, we show a side view of the bubble clouds from the dual-frequency case and split-frequency cases after $30 \mu\text{s}$ (from the same simulations as figures 5.17 and 5.18).

Here we see that the shapes of the bubble clouds are quite different. Figure 5.20a shows a more compact bubble cloud which is slightly concave to the right, while the bubble cloud in figure 5.20b is concave to the left and more spread out in the y-z




Figure 5.20: Dense bubble cloud ($A = 3 \text{ mm}^2$) deformation during treatment at $t = 30 \mu\text{s}$ for the (a) dual frequency and (b) split frequency cases. Bubbles are shown with 2 contour levels for the volume fraction of air: 0.2 (low opacity) and 0.8 (higher opacity)

plane, with the edges of the cloud extending beyond the stone edge. This difference in shape can be explained by the fact that the bubbles are being pushed out by the low frequency waves. In the dual frequency case, these low-frequency waves are coming from the entire transducer, pushing the outer bubbles inward towards the center. By contrast, the narrower beam of low-frequency waves in the split-frequency case pushes bubbles at the center of the cloud towards the stone, and as these bubbles interact with the stone surface they are pushed out radially, causing the cloud to spread out.

In fact, this effect was visible in simulations from section 5.4.3, in figures 5.14c and 5.15c, even without the stone present. The bubble cloud in the former is more compact, and in the latter spreads out radially. Given these examples, it appears that over a single ultrasound pulse, a strategy that spreads a bubble cloud out may be more effective than one aimed simply at collapsing bubbles or reducing their void fraction.

We note here that these results are limited to the time-scale of a single BWL pulse. Even with a fast pulse repetition frequency, the time between pulses is three orders of magnitude larger than the time-span of a BWL pulse (0.1 s vs. $60 \mu\text{s}$). While this method of spreading out the bubble cloud is effective at causing high damage to the stone over a single pulse, the resulting bubble dynamics between pulses are not yet well understood. Intuitively, spreading bubbles away from the stone as is

done with the split-frequency should only improve the energy delivery to the stone in the subsequent pulse, and so the effects of this strategy could stack up with each subsequent pulse. However, this remains to be tested, and is beyond the capabilities of our numerical methods.

Finally, in a clinical setting, the focal point accuracy required for the transducer may be difficult to achieve. The schism between damage caused in the split-frequency and dual-frequency cases may not be as obvious in a case where the transducer focal point is not exactly in the center of the stone. With low frequency wave in the center of the beamwidth, however, the tendency should remain for bubbles to spread out and away from the focal point of the transducer, and thus retain the benefits observed above even with imperfect focusing and a stone with a more complex topology.

Overall, the presented strategy in the split-frequency configuration appears to be a promising approach to reduce the shielding of bubble clouds by collapsing and dispersing bubbles forming ahead of the stone. Both the dense and sparse bubble clouds presented were effective in shielding the stone and reducing damage. However, in the tested cases, the split-frequency transducer configuration improved damage to the stone. It is a promising way to improve the efficacy and efficiency of BWL in cases with bubbles present ahead of the stone.

*Chapter 6***CONCLUSIONS AND FUTURE WORK**

The goal of this thesis was to implement and apply numerical models and data-driven methods to study biological systems involving the interaction of cavitating gas bubbles with elastic and viscoelastic materials. We successfully developed a data assimilation framework (DA-IMR) to aid in the modeling of biomaterials in low strain, high strain-rate settings, and applied it to various experimental data sets with distinct materials, bubble nucleation methods, and viscoelastic material models. DA-IMR was thoroughly tested for soft material characterization, and has improved both practicality (from O(days) to O(minutes)) and accuracy of the existing framework. We have shown its usage in a broad array of cases, and its ability to adapt to varied material models or additional parameters to estimate, with negligible added computational cost. Additionally, its use helped further our understanding of the limitations of current physical models used in inertial microcavitation rheometry. We have also implemented a numerical framework for the direct numerical simulation of acoustic-wave–bubble–elastic solid interactions, enabling the full simulation of burst wave lithotripsy. Through these simulations, we have investigated bubble dynamics and damage occurring in stones during BWL, and proposed strategies to mitigate bubble shielding, thus improving damage delivered to kidney stones, and accelerating their comminution. Key conclusions from each chapter of this thesis and suggestions for future work are summarized below

6.1 Data Assimilation for Inertial Microcavitation Rheometry

In Chapter 2, we presented a data-assimilation approach for the characterization of viscoelastic materials via observation of bubble collapse. This extension of inertial microcavitation rheometry, which we call DA-IMR, accurately estimated the mechanical properties of soft materials using ensemble-based data assimilation methods. With surrogate data obtained by adding random noise to simulated bubble radius vs. time curves (at frame rates replicating available high-speed imaging equipment), we validated this approach and demonstrated its accuracy and reliability in the context of laser-induced cavitation in hydrogels. The iterative ensemble Kalman smoother (IEnKS) and a hybrid ensemble-based 4D-Var method (En4D-Var)

provided the best results. The En4D-Var reduced computational cost by up to two orders of magnitude, obtaining accurate parameter estimates efficiently, accurately, and in a scalable framework. The IEnKS, while less efficient, provided useful information about time-dependent modeling uncertainty.

In Chapter 3, the DA-IMR framework was applied to experimental data across three test cases. In the first, it was used with experimental data similar to that used in the validation case in Chapter 2. Here, it was shown to provide accurate estimates for shear modulus and viscosity of a polyacrylamide gel with experimental laser-induced cavitation data, despite experimental inconsistency, model error, and noisy data. Additionally, we gained insight into data sets where the spherical bubble dynamics model used may have inadequately captured physics during violent bubble collapse. Next, we applied DA-IMR to experimental data for laser-nucleated cavitation in hydrogels with seeded particles. Here, a different stretch ratio regime was reached by nucleating bubbles on the surface of micro-particles in the gels. In this regime, and using a more general viscoelastic material model, we showed the adaptability and versatility of DA-IMR. At the same time, we established shortcomings in the physical model near the bubble collapse point in cases with high laser energy, as revealed by the quasi-online IEnKS estimation. This highlighted model uncertainties linked to plasma formation and the optical breakdown process during laser-induced cavitation. Finally, DA-IMR was used with ultrasound-induced cavitation, which addresses the issues seen previously with laser-induced cavitation. Here, the En4D-Var again accurately estimated material properties, using three different material models and estimating two additional model parameters.

Possible Improvements and Future Work The modeling uncertainties seen with laser-induced cavitation, particularly those analyzed in section 3.2, point to the need to capture additional physics in the bubble-dynamics models used. Some of the concerns around modeling uncertainties due to violent bubble collapse were addressed by switching to ultrasound-nucleated cavitation. However, with a better understanding of the plasma physics and rapid collapse occurring following laser-induced cavitation, this could become a more accurate and dependable framework for material characterization. Extending DA-IMR to more complete models, including those which could account for non-sphericity of bubbles, could also expand the capabilities of this framework. While such models are likely to be far more computationally expensive, the gained efficiency and parallelization capability of

DA-IMR suggest that such an approach would be computationally feasible. Such an adaptation could help gain insight into the non-spherical collapse of bubbles in viscoelastic materials, beyond the scope of material characterization.

While a significant advantage of DA-IMR is its efficiency and scalability due to the parallelizable nature of ensemble-based methods, this efficiency can be reduced given the need to tune algorithm parameters. In these initial runs, parameters such as initial modeling and measurement error covariance, number of ensemble members, and covariance inflation (both additive and multiplicative) needed to be tested and set through trial and error to obtain optimal results. An approach to tune these parameters based on the statistics of available data could further improve the efficiency of the method.

6.2 A Numerical Framework for Full Simulations of Burst-Wave Lithotripsy

In Chapter 4, we presented our implementation of a framework to model the full interaction of ultrasound waves from a spherically focused transducer array with gas bubbles and a submerged stone. This model was validated, and shown to match experimental stress imaging results during BWL well. Furthermore, the GPU implementation of the code was shown to provide significant acceleration and near-ideal scaling, enabling high resolution simulations. Calculating the maximum principal stresses generated in stones of various shapes and sizes, we confirmed that a frequency range of 350 kHz to 375 kHz appears effective in breaking kidney stones of typical size. That is, it maximizes the stresses and damage generated in stones of various sizes typical of human renal calculi.

In Chapter 5, we used the framework presented in Chapter 4 to investigate bubble dynamics and damage in burst-wave lithotripsy. We quantified the shielding of bubble clouds on waves and damage in the stone. We determined that introducing a low-frequency 20 kHz wave in addition to the usual BWL frequency, in this case 375 kHz, can effectively collapse bubbles ahead of a kidney stone during treatment. When splitting the therapy transducer between a low-frequency inner ring and high-frequency outer ring, shielding due to the tested bubble clouds was effectively reduced. We saw that the manner in which the low-frequency wave is introduced is important to maximize damage, and that bubble dispersion plays a role in this shielding reduction. We have presented a transducer configuration that enabled high damage in the tested cases, with a low-frequency pulse in the center of the transducer

to disperse bubbles, and high-frequency pulse from the outer ring to deliver damage to the stone.

Perspectives on Future Applications Only a few strategies to minimize bubble shielding have been tested here. Many other approaches exist, be it with the 18-element transducer used here or other transducer and acoustic source configurations. There are likely other ways to reduce shielding, and other approaches focused on pushing bubbles away from the stone, may yield better results. Furthermore, given the vast difference in time scales between BWL pulses, and inter-pulse time, only single pulses can be simulated with this method. Experimental trials or reduced order modeling methods that can cover the time-scale of multiple successive pulses could shed light on how well this method would translate to a therapy setting.

While multi-pulse simulations remain a challenge, the far reduced computational cost afforded by the GPU implementation broadens how this method can be applied. Indeed, iterative optimization approaches are possible with relatively high resolutions, even in full 3D simulations, provided the domain is small enough. Given the possibility of simulating damage in any arbitrary stones, one could envisage a treatment protocol whereby a stone is first imaged (using ultrasound or other imaging techniques), and an optimization is performed given this stone's size and shape using MFC. In simulations without bubbles present, the necessary resolution to obtain accurate stress and damage outputs is low, and such optimizations could thus readily be performed.

BIBLIOGRAPHY

- Akhatov, I., Lindau, O., Topolnikov, A., Mettin, R., Vakhitova, N., and Lauterborn, W. (2001). “Collapse and rebound of a laser-induced cavitation bubble.” In: *Physics of Fluids* 13.10, pp. 2805–2819. DOI: [10.1063/1.1401810](https://doi.org/10.1063/1.1401810).
- Allaire, G., Clerc, S., and Kokh, S. (2002). “A five-equation model for the simulation of interfaces between compressible fluids.” In: *Journal of Computational Physics* 181.2, pp. 577–616. DOI: [10.1006/jcph.2002.7143](https://doi.org/10.1006/jcph.2002.7143).
- Allen, D., Rule, W., and Jones, S. (1997). “Optimizing material strength constants numerically extracted from taylor impact data.” In: *Experimental Mechanics* 37.3, pp. 333–338.
- Altmeyer, G., Panicaud, B., Rouhaud, E., Wang, M., Roos, A., and Kerner, R. (2016). “Viscoelasticity behavior for finite deformations, using a consistent hypoelastic model based on rivlin materials.” In: *Continuum Mechanics and Thermodynamics* 28.6, pp. 1741–1758. DOI: [10.1007/s00161-016-0507-0](https://doi.org/10.1007/s00161-016-0507-0).
- Altmeyer, G., Rouhaud, E., Panicaud, B., Roos, A., Kerner, R., and Wang, M. (2015). “Viscoelastic models with consistent hypoelasticity for fluids undergoing finite deformations.” In: *Mechanics of Time-Dependent Materials* 19.3, pp. 375–395. DOI: [10.1007/s11043-015-9269-5](https://doi.org/10.1007/s11043-015-9269-5).
- Anderson, J. (2009). *Fundamentals of Aerodynamics*. McGraw-Hill Education. ISBN: 9780073398105.
- Anderson, J. L. and Anderson, S. L. (1999). “A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts.” In: *Monthly Weather Review* 127.12, pp. 2741–2758.
- Antonov, S., Fehn, A., and Kugi, A. (2011). “Unscented Kalman filter for vehicle state estimation.” In: *Vehicle System Dynamics* 49.9, pp. 1497–1520.
- Arora, P. D., Narani, N., and McCulloch, C. A. (1999). “The compliance of collagen gels regulates transforming growth factor- β induction of α -smooth muscle actin in fibroblasts.” In: *The American Journal of Pathology* 154.3, pp. 871–882.
- Bailey, M., Khokhlova, V., Sapozhnikov, O., Kargl, S., and Crum, L. (2003). “Physical mechanisms of the therapeutic effect of ultrasound (a review).” In: *Acoustical Physics* 49, pp. 369–388. DOI: [10.1134/1.1591291](https://doi.org/10.1134/1.1591291).
- Bailey, M. R., Maxwell, A. D., Cao, S., Ramesh, S., Liu, Z., Williams, J. C., Thiel, J., Dunmire, B., Colonius, T., Kuznetsova, E., Kreider, W., Sorensen, M. D., Lingeman, J. E., and Sapozhnikov, O. A. (2022). “Improving burst wave lithotripsy effectiveness for small stones and fragments by increasing frequency: theoretical

- modeling and ex vivo study.” eng. In: *Journal of Endourology* 36.7, pp. 996–1003. ISSN: 1557-900X (Electronic); 0892-7790 (Print); 0892-7790 (Linking). DOI: [10.1089/end.2021.0714](https://doi.org/10.1089/end.2021.0714).
- Bailey, M. R., Pishchalnikov, Y. A., Sapozhnikov, O. A., Cleveland, R. O., McAteer, J. A., Miller, N. A., Pishchalnikova, I. V., Connors, B. A., Crum, L. A., and Evan, A. P. (2005). “Cavitation detection during shock-wave lithotripsy.” In: *Ultrasound in Medicine & Biology* 31.9, pp. 1245–1256. DOI: [10.1016/j.ultrasmedbio.2005.02.017](https://doi.org/10.1016/j.ultrasmedbio.2005.02.017).
- Bar-Kochba, E., Scimone, M. T., Estrada, J. B., and Franck, C. (2016). “Strain and rate-dependent neuronal injury in a 3d in vitro compression model of traumatic brain injury.” In: *Scientific Reports* 6.1, p. 30550. DOI: [10.1038/srep30550](https://doi.org/10.1038/srep30550).
- Barajas, C. and Johnsen, E. (2017). “The effects of heat and mass diffusion on freely oscillating bubbles in a viscoelastic, tissue-like medium.” In: *The Journal of the Acoustical Society of America* 141.2, pp. 908–918. ISSN: 0001-4966. DOI: [10.1121/1.4976081](https://doi.org/10.1121/1.4976081).
- Barney, C. W., Dougan, C. E., McLeod, K. R., Kazemi-Moridani, A., Zheng, Y., Ye, Z., Tiwari, S., Sacligil, I., Riggelman, R. A., Cai, S., et al. (2020). “Cavitation in soft matter.” In: *Proceedings of the National Academy of Sciences* 117.17, pp. 9157–9165.
- Bengtsson, T., Snyder, C., and Nychka, D. (2003). “Toward a nonlinear ensemble filter for high-dimensional systems.” In: *Journal of Geophysical Research: Atmospheres* 108.D24.
- Bentz, K. C., Walley, S. E., and Savin, D. A. (2016). “Solvent effects on modulus of poly (propylene oxide)-based organogels as measured by cavitation rheology.” In: *Soft Matter* 12.22, pp. 4991–5001.
- Bocquet, M. (2011). “Ensemble Kalman filtering without the intrinsic need for inflation.” In: *Nonlinear Processes in Geophysics* 18.5, pp. 735–750. DOI: [10.5194/npg-18-735-2011](https://doi.org/10.5194/npg-18-735-2011).
- Bocquet, M. and Sakov, P. (2013a). “An iterative ensemble Kalman smoother.” In: *Quarterly Journal of the Royal Meteorological Society* 140.682, pp. 1521–1535. DOI: [10.1002/qj.2236](https://doi.org/10.1002/qj.2236).
- Bocquet, M. and Sakov, P. (2013b). “Joint state and parameter estimation with an iterative ensemble Kalman smoother.” In: *Nonlinear Processes in Geophysics* 20.5, pp. 803–818. DOI: [10.5194/npg-20-803-2013](https://doi.org/10.5194/npg-20-803-2013).
- Brennen, C. E. (1995). *Cavitation and Bubble Dynamics*. Oxford University Press.

- Brujan, E.-A. and Matsumoto, Y. (2012). “Collapse of micrometer-sized cavitation bubbles near a rigid boundary.” In: *Microfluidics and Nanofluidics* 13.6, pp. 957–966. DOI: [10.1007/s10404-012-1015-6](https://doi.org/10.1007/s10404-012-1015-6).
- Brujan, E.-A., Nahen, K., Schmidt, P., and Vogel, A. (2001). “Dynamics of laser-induced cavitation bubbles near an elastic boundary.” In: *Journal of Fluid Mechanics* 433, pp. 251–281. DOI: [10.1017/S0022112000003347](https://doi.org/10.1017/S0022112000003347).
- Brujan, E.-A. and Vogel, A. (2006). “Stress wave emission and cavitation bubble dynamics by nanosecond optical breakdown in a tissue phantom.” In: *Journal of Fluid Mechanics* 558, pp. 281–308.
- Bryngelson, S. H., Schmidmayer, K., Coralic, V., Meng, J. C., Maeda, K., and Colonius, T. (2021). “MFC: an open-source high-order multi-component, multi-phase, and multi-scale compressible flow solver.” In: *Computer Physics Communications* 266, p. 107396. ISSN: 0010-4655. DOI: <https://doi.org/10.1016/j.cpc.2020.107396>.
- Buyukozturk, S., Spratt, J.-S., Henann, D. L., Colonius, T., and Franck, C. (2022). “Particle-assisted laser-induced inertial cavitation for high strain-rate soft material characterization.” In: *Experimental Mechanics* 62.6, pp. 1037–1050. DOI: [10.1007/s11340-022-00861-7](https://doi.org/10.1007/s11340-022-00861-7).
- Cagniard, L., Flinn, E. A., Hewitt Dix, C., and Mayer, W. G. (1963). “Reflection and Refraction of Progressive Seismic Waves.” In: *Physics Today* 16.2, pp. 64–64. ISSN: 0031-9228. DOI: [10.1063/1.3050759](https://doi.org/10.1063/1.3050759).
- Cao, S., Zhang, Y., Liao, D., Zhong, P., and Wang, K. G. (2019). “Shock-induced damage and dynamic fracture in cylindrical bodies submerged in liquid.” In: *International Journal of Solids and Structures* 169, pp. 55–71. ISSN: 0020-7683. DOI: <https://doi.org/10.1016/j.ijsolstr.2019.04.002>.
- Caya, A., Sun, J., and Snyder, C. (2005). “A comparison between the 4DVAR and the ensemble Kalman filter techniques for radar data assimilation.” In: *Monthly Weather Review* 133.11, pp. 3081–3094. DOI: [10.1175/MWR3021.1](https://doi.org/10.1175/MWR3021.1).
- Chaudhuri, O., Gu, L., Klumpers, D., Darnell, M., Bencherif, S. A., Weaver, J. C., Huebsch, N., Lee, H.-p., Lippens, E., Duda, G. N., et al. (2016). “Hydrogels with tunable stress relaxation regulate stem cell fate and activity.” In: *Nature Materials* 15.3, pp. 326–334.
- Chaussy, C., Brendel, W., and Schmiedt, E. (1980). “Extracorporeally induced destruction of kidney stones by shock waves.” In: *The Lancet* 316.8207. Originally published as Volume 2, Issue 8207, pp. 1265–1268. ISSN: 0140-6736. DOI: [https://doi.org/10.1016/S0140-6736\(80\)92335-1](https://doi.org/10.1016/S0140-6736(80)92335-1).
- Chaussy, C., Schmiedt, E., Jocham, B., Brendel, W., Forssmann, B., and Walther, V. (1982). “First clinical experience with extracorporeally induced destruction of

- kidney stones by shock waves.” In: *The Journal of Urology* 127.3, pp. 417–420. ISSN: 0022-5347. DOI: [https://doi.org/10.1016/S0022-5347\(17\)53841-0](https://doi.org/10.1016/S0022-5347(17)53841-0).
- Chaussy, C. G. and Fuchs, G. J. (1989). “Current state and future developments of noninvasive treatment of human urinary stones with extracorporeal shock wave lithotripsy.” In: *The Journal of Urology* 141.3, Part 2, pp. 782–789. ISSN: 0022-5347. DOI: [https://doi.org/10.1016/S0022-5347\(17\)41010-X](https://doi.org/10.1016/S0022-5347(17)41010-X).
- Chen, T. T., Samson, P. C., Sorensen, M. D., and Bailey, M. R. (2020). “Burst wave lithotripsy and acoustic manipulation of stones.” eng. In: *Current Opinion in Urology* 30.2, pp. 149–156. ISSN: 1473-6586 (Electronic); 0963-0643 (Print); 0963-0643 (Linking). DOI: [10.1097/MOU.0000000000000727](https://doi.org/10.1097/MOU.0000000000000727).
- Chen, W. W. and Song, B. (2010). *Split Hopkinson (Kolsky) bar: design, testing and applications*. Springer Science & Business Media.
- Chin, M. S., Freniere, B. B., Fakhouri, S., Harris, J. E., Lalikos, J. F., and Crosby, A. J. (2013). “Cavitation rheology as a potential method for in vivo assessment of skin biomechanics.” In: *Plastic and Reconstructive Surgery* 131.2, 303e.
- Chockalingam, S., Roth, C., Henzel, T., and Cohen, T. (2021). “Probing local non-linear viscoelastic properties in soft materials.” In: *Journal of the Mechanics and Physics of Solids* 146, p. 104172. ISSN: 0022-5096. DOI: <https://doi.org/10.1016/j.jmps.2020.104172>.
- Cui, J., Lee, C. H., Delbos, A., McManus, J. J., and Crosby, A. J. (2011). “Cavitation rheology of the eye lens.” In: *Soft Matter* 7.17, pp. 7827–7831.
- De Hoop, A. (1960). “A modification of cagniard’s method for solving seismic pulse problems.” In: *Applied Scientific Research, Section B* 8.1, pp. 349–356.
- Diaz, J. and Ezziani, A. (2010). “Analytical solution for waves propagation in heterogeneous acoustic/porous media. part i: the 2d case.” In: *Communications in Computational Physics* 7.1, p. 171.
- Duryea, A. P., Cain, C. A., Roberts, W. W., and Hall, T. L. (2015). “Removal of residual cavitation nuclei to enhance histotripsy fractionation of soft tissue.” In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 62.12, pp. 2068–2078.
- Duryea, A. P., Roberts, W. W., Cain, C. A., Tamaddoni, H. A., and Hall, T. L. (2014). “Acoustic bubble removal to enhance swl efficacy at high shock rate: an in vitro study.” In: *Journal of Endourology* 28.1. PMID: 23957846, pp. 90–95. DOI: [10.1089/end.2013.0313](https://doi.org/10.1089/end.2013.0313).
- Duryea, A. P., Tamaddoni, H. A., Cain, C. A., Roberts, W. W., and Hall, T. L. (2015). “Removal of residual nuclei following a cavitation event: a parametric study.” In:

- IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 62.9, pp. 1605–1614. DOI: [10.1109/TUFFC.2014.006601](https://doi.org/10.1109/TUFFC.2014.006601).
- Epstein, D. and Keller, J. B. (1972). “Expansion and contraction of planar, cylindrical, and spherical underwater gas bubbles.” In: *The Journal of the Acoustical Society of America* 52.3B, pp. 975–980. DOI: [10.1121/1.1913203](https://doi.org/10.1121/1.1913203).
- Eringen, A. C. (1962). *Nonlinear theory of continuous media*. McGraw-Hill.
- Esch, E., Simmons, W. N., Sankin, G., Cocks, H. F., Preminger, G. M., and Zhong, P. (2010). “A simple method for fabricating artificial kidney stones of different physical properties.” In: *Urological Research* 38.4, pp. 315–319. DOI: [10.1007/s00240-010-0298-x](https://doi.org/10.1007/s00240-010-0298-x).
- Estrada, J. B., Barajas, C., Henann, D. L., Johnsen, E., and Franck, C. (2018). “High strain-rate soft material characterization via inertial cavitation.” In: *Journal of the Mechanics and Physics of Solids* 112, pp. 291–317. ISSN: 0022-5096. DOI: <https://doi.org/10.1016/j.jmps.2017.12.006>.
- Evan, A. P., Willis, L. R., Lingeman, J. E., and McAteer, J. A. (1998). “Renal trauma and the risk of long-term complications in shock wave lithotripsy.” In: *Nephron* 78.1, pp. 1–8.
- Evensen, G. (1994). “Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics.” In: *Journal of Geophysical Research: Oceans* 99.C5, pp. 10143–10162. DOI: [10.1029/94JC00572](https://doi.org/10.1029/94JC00572).
- Evensen, G. (2003). “The ensemble Kalman filter: Theoretical formulation and practical implementation.” In: *Ocean Dynamics* 53.4, pp. 343–367. DOI: [10.1007/s10236-003-0036-9](https://doi.org/10.1007/s10236-003-0036-9).
- Evensen, G. (2004). “Sampling strategies and square root analysis schemes for the enf.” In: *Ocean Dynamics* 54.6, pp. 539–560. DOI: [10.1007/s10236-004-0099-2](https://doi.org/10.1007/s10236-004-0099-2).
- Evensen, G. (2009a). *Data Assimilation: The Ensemble Kalman Filter*. Springer.
- Evensen, G. (2009b). “The ensemble Kalman filter for combined state and parameter estimation.” In: *IEEE Control Systems Magazine* 29.3, pp. 83–104. DOI: [10.1109/MCS.2009.932223](https://doi.org/10.1109/MCS.2009.932223).
- Evensen, G. and Leeuwen, P. J. van (2000). “An ensemble Kalman smoother for nonlinear dynamics.” In: *Monthly Weather Review* 128, pp. 1852–1867.
- Evensen, G. and Van Leeuwen, P. J. (1996). “Assimilation of geosat altimeter data for the agulhas current using the ensemble Kalman filter with a quasigeostrophic model.” In: *Monthly Weather Review* 124.1, pp. 85–96.

- Flynn, H. G. (1975). “Cavitation dynamics. I. A mathematical formulation.” In: *The Journal of the Acoustical Society of America* 57.6, pp. 1379–1396. ISSN: 0001-4966. DOI: [10.1121/1.380624](https://doi.org/10.1121/1.380624).
- Fujikawa, S. and Akamatsu, T. (1980). “Effects of the non-equilibrium condensation of vapour on the pressure wave produced by the collapse of a bubble in a liquid.” In: *Journal of Fluid Mechanics* 97.3, pp. 481–512.
- Fung, Y.-c. (2013). *Biomechanics: mechanical properties of living tissues*. Springer Science & Business Media.
- Gaudron, R., Warnez, M. T., and Johnsen, E. (2015). “Bubble dynamics in a viscoelastic medium with nonlinear elasticity.” In: *Journal of Fluid Mechanics* 766, pp. 54–75. ISSN: 1469-7645. DOI: [10.1017/jfm.2015.7](https://doi.org/10.1017/jfm.2015.7).
- Gavrilyuk, S., Favrie, N., and Saurel, R. (2008). “Modelling wave dynamics of compressible elastic materials.” In: *Journal of Computational Physics* 227.5, pp. 2941–2969. DOI: [10.1016/j.jcp.2007.11.030](https://doi.org/10.1016/j.jcp.2007.11.030).
- Gelb, A. (1974). *Applied Optimal Estimation*. IT Pro. MIT Press.
- Gent, A. and Wang, C. (1991). “Fracture mechanics and cavitation in rubber-like solids.” In: *Journal of Materials Science* 26.12, pp. 3392–3395.
- Ghil, M., Cohn, S., Tavantzis, J., Bube, K., and Isaacson, E. (1981). “Applications of estimation theory to numerical weather prediction.” In: *Dynamic Meteorology: Data Assimilation Methods*. Ed. by L. Bengtsson, M. Ghil, and E. Källén. New York, NY: Springer New York, pp. 139–224. DOI: [10.1007/978-1-4612-5970-1_5](https://doi.org/10.1007/978-1-4612-5970-1_5).
- Gustafsson, N. and Bojarova, J. (2014). “Four-dimensional ensemble variational (4d-en-var) data assimilation for the high resolution limited area model (hirlam).” In: *Nonlinear Processes in Geophysics* 21.4, pp. 745–762.
- Harper, J. D., Lingeman, J. E., Sweet, R. M., Metzler, I. S., Sunaryo, P. L., Williams Jr, J. C., Maxwell, A. D., Thiel, J., Cunitz, B. W., Dunmire, B., et al. (2022). “Fragmentation of stones by burst wave lithotripsy in the first 19 humans.” In: *The Journal of Urology* 207.5, pp. 1067–1076.
- Harper, J. D., Metzler, I., Hall, M. K., Chen, T. T., Maxwell, A. D., Cunitz, B. W., Dunmire, B., Thiel, J., Williams, J. C., Bailey, M. R., et al. (2021). “First in-human burst wave lithotripsy for kidney stone comminution: initial two case studies.” In: *Journal of Endourology* 35.4, pp. 506–511.
- Hashemnejad, S. M. and Kundu, S. (2015). “Nonlinear elasticity and cavitation of a triblock copolymer gel.” In: *Soft Matter* 11.21, pp. 4315–4325.

- Hill, A. J., Basourakos, S. P., Lewicki, P., Wu, X., Arenas-Gallo, C., Chuang, D., Bodner, D., Jaeger, I., Nevo, A., Zell, M., et al. (2022). "Incidence of kidney stones in the united states: the continuous national health and nutrition examination survey." In: *The Journal of Urology* 207.4, pp. 851–856.
- Hosea, M. and Shampine, L. (1996). "Analysis and implementation of TR–BDF2." In: *Applied Numerical Mathematics* 20.1, pp. 21–37. ISSN: 0168-9274. DOI: [https://doi.org/10.1016/0168-9274\(95\)00115-8](https://doi.org/10.1016/0168-9274(95)00115-8).
- Houtekamer, P. L. and Derome, J. (1995). "Methods for ensemble prediction." In: *Monthly Weather Review* 123.7, pp. 2181–2196. DOI: [https://doi.org/10.1175/1520-0493\(1995\)123<2181:MFEP>2.0.CO;2](https://doi.org/10.1175/1520-0493(1995)123<2181:MFEP>2.0.CO;2).
- Houtekamer, P. L. and Mitchell, H. L. (1998). "Data assimilation using an ensemble Kalman filter technique." In: *Monthly Weather Review* 126.3, pp. 796–811. DOI: [https://doi.org/10.1175/1520-0493\(1998\)126<0796:DAUAEK>2.0.CO;2](https://doi.org/10.1175/1520-0493(1998)126<0796:DAUAEK>2.0.CO;2).
- Houtekamer, P. L. and Zhang, F. (2016). "Review of the ensemble Kalman filter for atmospheric data assimilation." In: *Monthly Weather Review* 144.12, pp. 4489–4532.
- Hu, Y., You, J.-O., Auguste, D. T., Suo, Z., and Vlassak, J. J. (2012). "Indentation: a simple, nondestructive method for characterizing the mechanical and transport properties of ph-sensitive hydrogels." In: *Journal of Materials Research* 27.1, pp. 152–160.
- Hu, Y., Zhao, X., Vlassak, J. J., and Suo, Z. (2010). "Using indentation to characterize the poroelasticity of gels." In: *Applied Physics Letters* 96.12, p. 121904.
- Hunter, C., Maxwell, A. D., Cunitz, B., Dunmire, B., Sorensen, M. D., Williams James C., J., Randad, A., Bailey, M., and Kreider, W. (2018). "Impact of stone type on cavitation in burst wave lithotripsy." In: *Proceedings of Meetings on Acoustics* 35.1, p. 020005. ISSN: 1939-800X. DOI: [10.1121/2.0000950](https://doi.org/10.1121/2.0000950).
- Hutchens, S. B., Fakhouri, S., and Crosby, A. J. (2016). "Elastic cavitation and fracture via injection." In: *Soft Matter* 12.9, pp. 2557–2566.
- Ikeda, T., Yoshizawa, S., Tosaki, M., Allen, J. S., Takagi, S., Ohta, N., Kitamura, T., and Matsumoto, Y. (2006). "Cloud cavitation control for lithotripsy using high intensity focused ultrasound." eng. In: *Ultrasound in Medicine & Biology* 32.9, pp. 1383–1397. ISSN: 0301-5629 (Print); 0301-5629 (Linking). DOI: [10.1016/j.ultrasmedbio.2006.05.010](https://doi.org/10.1016/j.ultrasmedbio.2006.05.010).
- Janetschek, G., Frauscher, F., Knapp, R., Hofle, G., Peschel, R., and Bartsch, G. (1997). "New onset hypertension after extracorporeal shock wave lithotripsy: age related incidence and prediction by intrarenal resistive index." In: *The Journal of Urology* 158.2, pp. 346–351.

- Jarmusch, A., Liu, A., Munley, C., Horta, D., Ravichandran, V., Denny, J., Friedline, K., and Chandrasekaran, S. (2022). “Analysis of validating and verifying openacc compilers 3.0 and above.” In: *2022 Workshop on Accelerator Programming Using Directives (WACCPD)*, pp. 1–10. DOI: [10.1109/WACCPD56842.2022.000006](https://doi.org/10.1109/WACCPD56842.2022.000006).
- Jazwinski, A. (2007). *Stochastic Processes and Filtering Theory*. Dover Books on Electrical Engineering Series. Dover Publications.
- Jendeberg, J., Geijer, H., Alshamari, M., Cierzniak, B., and Lidén, M. (2017). “Size matters: the width and location of a ureteral stone accurately predict the chance of spontaneous passage.” In: *European Radiology* 27.11, pp. 4775–4785. DOI: [10.1007/s00330-017-4852-6](https://doi.org/10.1007/s00330-017-4852-6).
- Julier, S. and Uhlmann, J. (2004). “Unscented filtering and nonlinear estimation.” In: *Proceedings of the IEEE* 92.3, pp. 401–422. DOI: [10.1109/JPROC.2003.823141](https://doi.org/10.1109/JPROC.2003.823141).
- Kalman, R. E. (1960). “A new approach to linear filtering and prediction problems.” In: *Journal of Basic Engineering* 82.1, pp. 35–45.
- Kalman, R. E. and Bucy, R. S. (1961). “New results in linear filtering and prediction theory.” In: *Journal of Basic Engineering* 83.1, pp. 95–108. ISSN: 0021-9223. DOI: [10.1115/1.3658902](https://doi.org/10.1115/1.3658902).
- Kapila, A. K., Menikoff, R., Bdzil, J. B., Son, S. F., and Stewart, D. S. (2001). “Two-phase modeling of deflagration-to-detonation transition in granular materials: reduced equations.” In: *Physics of Fluids* 13.10, pp. 3002–3024. DOI: [10.1063/1.1398042](https://doi.org/10.1063/1.1398042).
- Katzfuss, M., Stroud, J. R., and Wikle, C. K. (2016). “Understanding the ensemble Kalman filter.” In: *The American Statistician* 70.4, pp. 350–357. DOI: [10.1080/00031305.2016.1141709](https://doi.org/10.1080/00031305.2016.1141709).
- Keller, J. B. and Kolodner, I. I. (1956). “Damping of underwater explosion bubble oscillations.” In: *Journal of Applied Physics* 27.10, pp. 1152–1161. ISSN: 0021-8979. DOI: [10.1063/1.1722221](https://doi.org/10.1063/1.1722221).
- Keller, J. B. and Miksis, M. (1980). “Bubble oscillations of large amplitude.” In: *The Journal of the Acoustical Society of America* 68.2, pp. 628–633. DOI: [10.1121/1.384720](https://doi.org/10.1121/1.384720).
- Khalilov, M. and Timoveev, A. (2021). “Performance analysis of cuda, openacc and openmp programming models on tesla v100 gpu.” In: *Journal of Physics: Conference Series* 1740.1, p. 012056. DOI: [10.1088/1742-6596/1740/1/012056](https://doi.org/10.1088/1742-6596/1740/1/012056).
- Kim, S.-G., Crassidis, J. L., Cheng, Y., Fosbury, A. M., and Junkins, J. L. (2007). “Kalman filtering for relative spacecraft attitude and position estimation.” In:

- Journal of Guidance, Control, and Dynamics* 30.1, pp. 133–143. DOI: [10.2514/1.22377](https://doi.org/10.2514/1.22377).
- Komatitsch, D., Barnes, C., and Tromp, J. (2000). “Wave propagation near a fluid-solid interface: a spectral-element approach.” In: *Geophysics* 65.2, pp. 623–631. DOI: [10.1190/1.1444758](https://doi.org/10.1190/1.1444758).
- Kundu, S. and Crosby, A. J. (2009). “Cavitation and fracture behavior of polyacrylamide hydrogels.” In: *Soft Matter* 5 (20), pp. 3963–3968. DOI: [10.1039/B909237D](https://doi.org/10.1039/B909237D).
- Law, K., Stuart, A., and Zygalakis, K. (2015). “Data assimilation.” In: *Cham, Switzerland: Springer* 214, p. 52.
- Le Métayer, O., Massoni, J., and Saurel, R. (2004). “Élaboration des lois d’état d’un liquide et de sa vapeur pour les modèles d’écoulements diphasiques.” In: *International Journal of Thermal Sciences* 43.3, pp. 265–276. ISSN: 1290-0729. DOI: <https://doi.org/10.1016/j.ijthermalsci.2003.09.002>.
- Le Métayer, O., Massoni, J., and Saurel, R. (2005). “Modelling evaporation fronts with reactive riemann solvers.” In: *Journal of Computational Physics* 205.2, pp. 567–610. ISSN: 0021-9991. DOI: <https://doi.org/10.1016/j.jcp.2004.11.021>.
- Lee, K. Y. and Mooney, D. J. (2012). “Alginate: properties and biomedical applications.” In: *Progress in Polymer Science* 37.1, pp. 106–126.
- Leeuwen, P. J. van (1999). “Comment on “Data assimilation using an ensemble Kalman filter technique”.” In: *Monthly Weather Review* 127.6, pp. 1374–1377. DOI: [10.1175/1520-0493\(1999\)127<1374:CODAUA>2.0.CO;2](https://doi.org/10.1175/1520-0493(1999)127<1374:CODAUA>2.0.CO;2).
- Leeuwen, P. J. van (2009). “Particle filtering in geophysical systems.” In: *Monthly Weather Review* 137.12, pp. 4089–4114.
- Lefferts, E., Markley, F., and Shuster, M. (1982). “Kalman filtering for spacecraft attitude estimation.” In: *Journal of Guidance, Control, and Dynamics* 5.5, pp. 417–429. DOI: [10.2514/3.56190](https://doi.org/10.2514/3.56190).
- Lingeman, J. E., McAteer, J. A., Gnessin, E., and Evan, A. P. (2009). “Shock wave lithotripsy: advances in technology and technique.” In: *Nature Reviews Urology* 6.12, pp. 660–670. DOI: [10.1038/nrurol.2009.216](https://doi.org/10.1038/nrurol.2009.216).
- Liu, C., Xiao, Q., and Wang, B. (2008). “An ensemble-based four-dimensional variational data assimilation scheme. part i: technical formulation and preliminary test.” In: *Monthly Weather Review* 136.9, pp. 3363–3373.

- Liu, Y. and Zhong, P. (2002). “BegoStone—a new stone phantom for shock wave lithotripsy research (L).” In: *The Journal of the Acoustical Society of America* 112.4, pp. 1265–1268. ISSN: 0001-4966. DOI: [10.1121/1.1501905](https://doi.org/10.1121/1.1501905).
- Lorenc, A. C. (1986). “Analysis methods for numerical weather prediction.” In: *Quarterly Journal of the Royal Meteorological Society* 112.474, pp. 1177–1194.
- Luo, X. and Hoteit, I. (2011). “Robust ensemble filtering and its relation to covariance inflation in the ensemble Kalman filter.” In: *Monthly Weather Review* 139.12, pp. 3938–3953. DOI: [10.1175/MWR-D-10-05068.1](https://doi.org/10.1175/MWR-D-10-05068.1).
- Madyastha, V., Ravindra, V., Mallikarjunan, S., and Goyal, A. (2011). “Extended kalman filter vs. error state kalman filter for aircraft attitude estimation.” In: *AIAA Guidance, Navigation, and Control Conference*, p. 6615.
- Maeda, K. and Colonius, T. (2017). “A source term approach for generation of one-way acoustic waves in the Euler and Navier–Stokes equations.” In: *Wave Motion* 75, pp. 36–49. ISSN: 0165-2125. DOI: <https://doi.org/10.1016/j.wavemoti.2017.08.004>.
- Maeda, K. and Colonius, T. (2018). “Eulerian–lagrangian method for simulation of cloud cavitation.” In: *Journal of Computational Physics* 371, pp. 994–1017. ISSN: 0021-9991. DOI: <https://doi.org/10.1016/j.jcp.2018.05.029>.
- Maeda, K. and Colonius, T. (2019). “Bubble cloud dynamics in an ultrasound field.” In: *Journal of Fluid Mechanics* 862, pp. 1105–1134. DOI: [10.1017/jfm.2018.968](https://doi.org/10.1017/jfm.2018.968).
- Maeda, K., Colonius, T., Maxwell, A., Kreider, W., and Bailey, M. (2018). “Modeling and numerical simulation of the bubble cloud dynamics in an ultrasound field for burst wave lithotripsy.” In: *Proceedings of Meetings on Acoustics* 35.1, p. 020006. ISSN: 1939-800X. DOI: [10.1121/2.0000946](https://doi.org/10.1121/2.0000946).
- Maeda, K., Kreider, W., Maxwell, A., Cunitz, B., Colonius, T., and Bailey, M. (2015). “Modeling and experimental analysis of acoustic cavitation bubbles for burst wave lithotripsy.” In: *Journal of Physics: Conference Series* 656, p. 012027. DOI: [10.1088/1742-6596/656/1/012027](https://doi.org/10.1088/1742-6596/656/1/012027).
- Maeda, K., Maxwell, A. D., Colonius, T., Kreider, W., and Bailey, M. R. (2018). “Energy shielding by cavitation bubble clouds in burst wave lithotripsy.” In: *The Journal of the Acoustical Society of America* 144.5, pp. 2952–2961. DOI: [10.1121/1.5079641](https://doi.org/10.1121/1.5079641).
- Mancia, L., Vlaisavljevich, E., Xu, Z., and Johnsen, E. (2017). “Predicting tissue susceptibility to mechanical cavitation damage in therapeutic ultrasound.” In: *Ultrasound in Medicine & Biology* 43.7, pp. 1421–1440. ISSN: 0301-5629. DOI: <https://doi.org/10.1016/j.ultrasmedbio.2017.02.020>.

- Mancia, L., Vlaisavljevich, E., Yousefi, N., Rodriguez, M., Ziemlewicz, T. J., Lee, F. T., Henann, D., Franck, C., Xu, Z., and Johnsen, E. (2019). “Modeling tissue-selective cavitation damage.” In: *Physics in Medicine & Biology* 64.22, p. 225001.
- Mancia, L., Yang, J., Spratt, J.-S., Sukovich, J. R., Xu, Z., Colonius, T., Franck, C., and Johnsen, E. (2021). “Acoustic cavitation rheometry.” In: *Soft Matter* 17.10, pp. 2931–2941. DOI: [10.1039/d0sm02086a](https://doi.org/10.1039/d0sm02086a).
- Matsumoto, Y. and Yoshizawa, S. (2005). “Behaviour of a bubble cluster in an ultrasound field.” In: *International Journal for Numerical Methods in Fluids* 47.6-7, pp. 591–601. DOI: [10.1002/flid.833](https://doi.org/10.1002/flid.833).
- Maxwell, A. D., Cain, C. A., Hall, T. L., Fowlkes, J. B., and Xu, Z. (2013). “Probability of cavitation for single ultrasound pulses applied to tissues and tissue-mimicking materials.” In: *Ultrasound in Medicine & Biology* 39.3, pp. 449–465.
- Maxwell, A. D., Wang, Y.-N., Kreider, W., Cunitz, B. W., Starr, F., Lee, D., Nazari, Y., Williams Jr, J. C., Bailey, M. R., and Sorensen, M. D. (2019). “Evaluation of renal stone comminution and injury by burst wave lithotripsy in a pig model.” In: *Journal of Endourology* 33.10, pp. 787–792.
- Maxwell, A. D., Cain, C. A., Duryea, A. P., Yuan, L., Gurm, H. S., and Xu, Z. (2009). “Noninvasive thrombolysis using pulsed ultrasound cavitation therapy – histotripsy.” In: *Ultrasound in Medicine & Biology* 35.12, pp. 1982–1994. ISSN: 0301-5629. DOI: <https://doi.org/10.1016/j.ultrasmedbio.2009.07.001>.
- Maxwell, A. D., Cunitz, B. W., Kreider, W., Sapozhnikov, O. A., Hsi, R. S., Harper, J. D., Bailey, M. R., and Sorensen, M. D. (2015). “Fragmentation of urinary calculi in vitro by burst wave lithotripsy.” In: *Journal of Urology* 193.1, pp. 338–344. DOI: [10.1016/j.juro.2014.08.009](https://doi.org/10.1016/j.juro.2014.08.009).
- Maxwell, A. D., MacConaghy, B., Bailey, M. R., and Sapozhnikov, O. A. (2020). “An investigation of elastic waves producing stone fracture in burst wave lithotripsy.” In: *The Journal of the Acoustical Society of America* 147.3, pp. 1607–1622. DOI: [10.1121/10.0000847](https://doi.org/10.1121/10.0000847).
- McAteer, J. A. and Evan, A. P. (2008). “The acute and long-term adverse effects of shock wave lithotripsy.” In: *Seminars in Nephrology* 28.2, pp. 200–213. DOI: [10.1016/j.semnephrol.2008.01.003](https://doi.org/10.1016/j.semnephrol.2008.01.003).
- Meaney, D. F. and Smith, D. H. (2011). “Biomechanics of concussion.” In: *Clinics in Sports Medicine* 30.1. Concussion in Sports, pp. 19–31. ISSN: 0278-5919. DOI: <https://doi.org/10.1016/j.csm.2010.08.009>.
- Mijailovic, A. S., Galarza, S., Raayai-Ardakani, S., Birch, N. P., Schiffman, J. D., Crosby, A. J., Cohen, T., Peyton, S. R., and Van Vliet, K. J. (2021). “Localized characterization of brain tissue mechanical properties by needle induced cavitation

- rheology and volume controlled cavity expansion.” In: *Journal of the Mechanical Behavior of Biomedical Materials* 114, p. 104168.
- Movahed, P., Kreider, W., Maxwell, A. D., Hutchens, S. B., and Freund, J. B. (2016). “Cavitation-induced damage of soft materials by focused ultrasound bursts: a fracture-based bubble dynamics model.” In: *The Journal of the Acoustical Society of America* 140.2, pp. 1374–1386. DOI: [10.1121/1.4961364](https://doi.org/10.1121/1.4961364).
- Nigmatulin, R., Khabeev, N., and Nagiev, F. (1981). “Dynamics, heat and mass transfer of vapour-gas bubbles in a liquid.” In: *International Journal of Heat and Mass Transfer* 24.6, pp. 1033–1044.
- Nyein, M. K., Jason, A. M., Yu, L., Pita, C. M., Joannopoulos, J. D., Moore, D. F., and Radovitzky, R. A. (2010). “In silico investigation of intracranial blast mitigation with relevance to military traumatic brain injury.” In: *Proceedings of the National Academy of Sciences* 107.48, pp. 20703–20708. ISSN: 0027-8424. DOI: [10.1073/pnas.1014786107](https://doi.org/10.1073/pnas.1014786107).
- Plesset, M. S. and Prosperetti, A. (1977). “Bubble dynamics and cavitation.” In: *Annual Review of Fluid Mechanics* 9.1, pp. 145–185. DOI: [10.1146/annurev.fl.09.010177.001045](https://doi.org/10.1146/annurev.fl.09.010177.001045).
- Plesset, M. (1948). “Dynamics of cavitation bubbles.” In: *Journal of Applied Mechanics* 16, pp. 228–231.
- Plesset, M. (1949). “The dynamics of cavitation bubbles.” In: *Journal of Applied Mechanics* 16.1, pp. 277–282.
- Preston, A. T., Colonius, T., and Brennen, C. E. (2007). “A reduced-order model of diffusive effects on the dynamics of bubbles.” In: *Physics of Fluids* 19.12, p. 123302. DOI: [10.1063/1.2825018](https://doi.org/10.1063/1.2825018).
- Preston, A. T. (2004). “Modeling heat and mass transfer in bubbly cavitating flows and shock waves in cavitating nozzles.” PhD thesis. California Institute of Technology.
- Prosperetti, A. and Lezzi, A. (1986). “Bubble dynamics in a compressible liquid. part 1. first-order theory.” In: *Journal of Fluid Mechanics* 168, pp. 457–478. DOI: [10.1017/S0022112086000460](https://doi.org/10.1017/S0022112086000460).
- Prosperetti, A. (1991). “The thermal behaviour of oscillating gas bubbles.” In: *Journal of Fluid Mechanics* 222, pp. 587–616.
- Prosperetti, A., Crum, L. A., and Commander, K. W. (1988). “Nonlinear bubble dynamics.” In: *The Journal of the Acoustical Society of America* 83.2, pp. 502–514. DOI: [10.1121/1.396145](https://doi.org/10.1121/1.396145).

- Raayai-Ardakani, S., Chen, Z., Earl, D. R., and Cohen, T. (2019). “Volume-controlled cavity expansion for probing of local elastic properties in soft materials.” In: *Soft Matter* 15.3, pp. 381–392.
- Raayai-Ardakani, S. and Cohen, T. (2019). “Capturing strain stiffening using volume controlled cavity expansion.” In: *Extreme Mechanics Letters* 31, p. 100536.
- Raayai-Ardakani, S., Earl, D. R., and Cohen, T. (2019). “The intimate relationship between cavitation and fracture.” In: *Soft Matter* 15 (25), pp. 4999–5005. DOI: [10.1039/C9SM00570F](https://doi.org/10.1039/C9SM00570F).
- Radhakrishnan, A., Berre, H. L., Wilfong, B., Spratt, J.-S., Rodriguez, M., Colonius, T., and Bryngelson, S. H. (2023). *Method for portable, scalable, and performant GPU-accelerated simulation of multiphase compressible flow*. Under Review.
- Ramasamy, A., Hill, A. M., Masouros, S., Gibb, I., Bull, A. M., and Clasper, J. C. (2011). “Blast-related fracture patterns: a forensic biomechanical approach.” In: *Journal of The Royal Society Interface* 8.58, pp. 689–698.
- Raskolnikov, D., Bailey, M. R., and Harper, J. D. (2022). “Recent advances in the science of burst wave lithotripsy and ultrasonic propulsion.” In: *BME Frontiers* 2022. DOI: [10.34133/2022/9847952](https://doi.org/10.34133/2022/9847952).
- Rayleigh, L. (1917). “Viii. On the pressure developed in a liquid during the collapse of a spherical cavity.” In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 34.200, pp. 94–98. ISSN: 1941-5990. DOI: [10.1080/14786440808635681](https://doi.org/10.1080/14786440808635681).
- Rodriguez, M. and Johnsen, E. (2019). “A high-order accurate five-equations compressible multiphase approach for viscoelastic fluids and solids with relaxation and elasticity.” In: *Journal of Computational Physics* 379, pp. 70–90. DOI: [10.1016/j.jcp.2018.10.035](https://doi.org/10.1016/j.jcp.2018.10.035).
- Ryan, H. (1994). “Ricker, ormsby, klander, butterworth – a choice of wavelets.” In: *Canadian Society of Exploration Geophysicists Recorder* 19.07.
- Sackmann, M., Delius, M., Sauerbruch, T., Holl, J., Weber, W., Ippisch, E., Hage-lauer, U., Wess, O., Hepp, W., Brendel, W., et al. (1988). “Shock-wave lithotripsy of gallbladder stones.” In: *New England Journal of Medicine* 318.7, pp. 393–397.
- Sagar, H. J. and el Moctar, O. (2020). “Dynamics of a cavitation bubble near a solid surface and the induced damage.” In: *Journal of Fluids and Structures* 92, p. 102799. ISSN: 0889-9746. DOI: <https://doi.org/10.1016/j.jfluidstructs.2019.102799>.
- Sakov, P., Oliver, D. S., and Bertino, L. (2012). “An iterative EnKF for strongly nonlinear systems.” In: *Monthly Weather Review* 140.6, pp. 1988–2004. DOI: [10.1175/mwr-d-11-00176.1](https://doi.org/10.1175/mwr-d-11-00176.1).

- Sanz-Alonso, D., Stuart, A., and Taeb, A. (2023). *Inverse Problems and Data Assimilation*. London Mathematical Society Student Texts. Cambridge University Press.
- Sapozhnikov, O. A., Maxwell, A. D., and Bailey, M. R. (2020). “Modeling of photoelastic imaging of mechanical stresses in transparent solids mimicking kidney stones.” eng. In: *The Journal of the Acoustical Society of America* 147.6, p. 3819. ISSN: 1520-8524 (Electronic); 0001-4966 (Print); 0001-4966 (Linking). DOI: [10.1121/10.0001386](https://doi.org/10.1121/10.0001386).
- Sapozhnikov, O. A., Maxwell, A. D., and Bailey, M. R. (2021). “Maximizing mechanical stress in small urinary stones during burst wave lithotripsy.” In: *The Journal of the Acoustical Society of America* 150.6, pp. 4203–4212. DOI: [10.1121/10.0008902](https://doi.org/10.1121/10.0008902).
- Sarntinoranont, M., Lee, S. J., Hong, Y., King, M. A., Subhash, G., Kwon, J., and Moore, D. F. (2012). “High-strain-rate brain injury model using submerged acute rat brain tissue slices.” In: *Journal of Neurotrauma* 29.2. PMID: 21970544, pp. 418–429. DOI: [10.1089/neu.2011.1772](https://doi.org/10.1089/neu.2011.1772).
- Scales, C. D. J., Smith, A. C., Hanley, J. M., and Saigal, C. S. (2012). “Prevalence of kidney stones in the United States.” eng. In: *European Radiology* 62.1, pp. 160–165. ISSN: 1873-7560 (Electronic); 0302-2838 (Print); 0302-2838 (Linking). DOI: [10.1016/j.eururo.2012.03.052](https://doi.org/10.1016/j.eururo.2012.03.052).
- Schillings, C. and Stuart, A. M. (2017). “Analysis of the ensemble Kalman filter for inverse problems.” In: *SIAM Journal on Numerical Analysis* 55.3, pp. 1264–1290.
- Schmidmayer, K., Bryngelson, S. H., and Colonius, T. (2020). “An assessment of multicomponent flow models and interface capturing schemes for spherical bubble dynamics.” In: *Journal of Computational Physics* 402, p. 109080. ISSN: 0021-9991. DOI: <https://doi.org/10.1016/j.jcp.2019.109080>.
- Shampine, L. F. and Reichelt, M. W. (1997). “The matlab ODE suite.” In: *SIAM Journal on Scientific Computing* 18.1, pp. 1–22.
- Shampine, L. F., Reichelt, M. W., and Kierzenka, J. A. (1999). “Solving index-1 daes in matlab and simulink.” In: *SIAM Review* 41.3, pp. 538–552.
- Shimada, M., Kobayashi, T., and Matsumoto, Y. (1999). “Dynamics of cloud cavitation and cavitation erosion.” In: *Proceedings of the 1999 3rd ASME/JSME Joint Fluids Engineering Conference, FEDSM'99, San Francisco, California, USA, 18-23 July 1999 (CD-ROM)*. American Society of Mechanical Engineers.
- Solomon, W. and Jindal, V. (2007). “Modeling changes in rheological properties of potatoes during storage under constant and variable conditions.” In: *LWT-Food Science and Technology* 40.1, pp. 170–178.

- Spratt, J.-S., Rodriguez, M., Bryngelson, S. H., Cao, S., and Colonius, T. (2021). “Eulerian framework for bubble-cloud-kidney stone interaction.” In: *11th International Symposium on Cavitation*. Daejeon, Korea.
- Spratt, J.-S. and Colonius, T. (2023). “Numerical Investigation of Bubble Dynamics and Damage in Burst-Wave Lithotripsy.” In Preparation.
- Spratt, J.-S., Rodriguez, M., Schmidmayer, K., Bryngelson, S. H., Yang, J., Franck, C., and Colonius, T. (2021). “Characterizing viscoelastic materials via ensemble-based data assimilation of bubble collapse observations.” In: *Journal of the Mechanics and Physics of Solids* 152, p. 104455. ISSN: 0022-5096. DOI: <https://doi.org/10.1016/j.jmps.2021.104455>.
- Storrie, H. and Mooney, D. J. (2006). “Sustained delivery of plasmid dna from polymeric scaffolds for tissue engineering.” In: *Advanced Drug Delivery Reviews* 58.4, pp. 500–514.
- Tamaddoni, H. A., Roberts, W. W., Duryea, A. P., Cain, C. A., and Hall, T. L. (2016). “Enhanced high-rate shockwave lithotripsy stone comminution in an in vivo porcine model using acoustic bubble coalescence.” In: *Journal of Endourology* 30.12. PMID: 27762629, pp. 1321–1325. DOI: [10.1089/end.2016.0407](https://doi.org/10.1089/end.2016.0407).
- Tamaddoni, H. A., Roberts, W. W., and Hall, T. L. (2019). “Enhanced shockwave lithotripsy with active cavitation mitigation.” In: *The Journal of the Acoustical Society of America* 146.5, pp. 3275–3282. ISSN: 0001-4966. DOI: [10.1121/1.5131649](https://doi.org/10.1121/1.5131649).
- Taylor, G. I. (1948). “The use of flat-ended projectiles for determining dynamic yield stress i. Theoretical considerations.” In: *Proceedings of the Royal Society A* 194.1038, pp. 289–299.
- Trémolet, Y. (2007). “Incremental 4D–Var convergence study.” In: *Tellus A: Dynamic Meteorology and Oceanography* 59.5, pp. 706–718. DOI: [10.1111/j.1600-0870.2007.00271.x](https://doi.org/10.1111/j.1600-0870.2007.00271.x).
- Van Slijtenhorst, C., Cronin, D. S., and Brodland, G. W. (2006). “High strain rate compressive properties of bovine muscle tissue determined using a split Hopkinson bar apparatus.” In: *Journal of Biomechanics* 39.10, pp. 1852–1858.
- Venhovens, P. J. T. and Naab, K. (1999). “Vehicle dynamics estimation using Kalman filters.” In: *Vehicle System Dynamics* 32.2-3, pp. 171–184.
- Vincent, O., Marmottant, P., Gonzalez-Avila, S. R., Ando, K., and Ohl, C.-D. (2014). “The fast dynamics of cavitation bubbles within water confined in elastic solids.” In: *Soft Matter* 10.10, pp. 1455–1461.
- Vlaisavljevich, E., Lin, K.-W., Maxwell, A., Warnez, M. T., Mancina, L., Singh, R., Putnam, A. J., Fowlkes, B., Johnsen, E., Cain, C., et al. (2015). “Effects of

- ultrasound frequency and tissue stiffness on the histotripsy intrinsic threshold for cavitation.” In: *Ultrasound in Medicine & Biology* 41.6, pp. 1651–1667.
- Vlaisavljevich, E., Lin, K.-W., Warnez, M. T., Singh, R., Mancina, L., Putnam, A. J., Johnsen, E., Cain, C., and Xu, Z. (2015). “Effects of tissue stiffness, ultrasound frequency, and pressure on histotripsy-induced cavitation bubble behavior.” In: *Physics in Medicine & Biology* 60.6, pp. 2271–2292. ISSN: 1361-6560 (Electronic); 0031-9155 (Print); 0031-9155 (Linking). DOI: [10.1088/0031-9155/60/6/2271](https://doi.org/10.1088/0031-9155/60/6/2271).
- Vogel, A., Linz, N., Freidank, S., and Paltauf, G. (2008). “Femtosecond-laser-induced nanocavitation in water: implications for optical breakdown threshold and cell surgery.” In: *Physical Review Letters* 100.3, p. 038102.
- Warnez, M. T. and Johnsen, E. (2015). “Numerical modeling of bubble dynamics in viscoelastic media with relaxation.” In: *Physics of Fluids* 27.6, p. 063103. ISSN: 1089-7666. DOI: [10.1063/1.4922598](https://doi.org/10.1063/1.4922598).
- Wenzel, T. A., Burnham, K., Blundell, M., and Williams, R. (2006). “Dual extended kalman filter for vehicle state and parameter estimation.” In: *Vehicle System Dynamics* 44.2, pp. 153–171.
- Whitaker, J. S. and Hamill, T. M. (2012). “Evaluating methods to account for system errors in ensemble data assimilation.” In: *Monthly Weather Review* 140.9, pp. 3078–3089. DOI: [10.1175/MWR-D-11-00276.1](https://doi.org/10.1175/MWR-D-11-00276.1).
- Wienke, S., Springer, P., Terboven, C., and Mey, D. an (2012). “OpenACC — first experiences with real-world applications.” In: *Euro-Par 2012 Parallel Processing*. Ed. by C. Kaklamanis, T. Papatheodorou, and P. G. Spirakis. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 859–870. ISBN: 978-3-642-32820-6.
- Wilson, C. T., Hall, T. L., Johnsen, E., Mancina, L., Rodriguez, M., Lundt, J. E., Colonius, T., Henann, D. L., Franck, C., Xu, Z., et al. (2019). “Comparative study of the dynamics of laser and acoustically generated bubbles in viscoelastic media.” In: *Physical Review E* 99.4, p. 043103.
- Xu, Z., Raghavan, M., Hall, T. L., Chang, C., Mycek, M., Fowlkes, J. B., and Cain, C. A. (2007). “High speed imaging of bubble clouds generated in pulsed ultrasound cavitation therapy – histotripsy.” In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 54.10, pp. 2091–2101.
- Yang, J., Cramer, H. C., and Franck, C. (2020). “Extracting non-linear viscoelastic material properties from violently-collapsing cavitation bubbles.” In: *Extreme Mechanics Letters* 39, p. 100839. DOI: [10.1016/j.eml.2020.100839](https://doi.org/10.1016/j.eml.2020.100839).
- Yang, S.-C., Kalnay, E., and Hunt, B. (2012). “Handling nonlinearity in an ensemble Kalman filter: Experiments with the tree-variable Lorenz model.” In: *Monthly Weather Review* 140.8, pp. 2628–2646.

- Yang, X. and Church, C. C. (2005). “A model for the dynamics of gas bubbles in soft tissue.” In: *The Journal of the Acoustical Society of America* 118.6, pp. 3595–3606. ISSN: 0001-4966. DOI: [10.1121/1.2118307](https://doi.org/10.1121/1.2118307).
- Yoshizawa, S., Ikeda, T., Ito, A., Ota, R., Takagi, S., and Matsumoto, Y. (2009). “High intensity focused ultrasound lithotripsy with cavitating microbubbles.” In: *Medical & Biological Engineering & Computing* 47.8, pp. 851–860. DOI: [10.1007/s11517-009-0471-y](https://doi.org/10.1007/s11517-009-0471-y).
- Zimberlin, J. A., McManus, J. J., and Crosby, A. J. (2010). “Cavitation rheology of the vitreous: mechanical properties of biological tissue.” In: *Soft Matter* 6.15, pp. 3632–3635.
- Zimberlin, J. A., Sanabria-DeLong, N., Tew, G. N., and Crosby, A. J. (2007). “Cavitation rheology for soft materials.” In: *Soft Matter* 3.6, pp. 763–767.
- Zwaschka, T. A., Ahn, J. S., Cunitz, B. W., Bailey, M. R., Dunmire, B., Sorensen, M. D., Harper, J. D., and Maxwell, A. D. (2018). “Combined burst wave lithotripsy and ultrasonic propulsion for improved urinary stone fragmentation.” In: *Journal of Endourology* 32.4, pp. 344–349. ISSN: 1557-900X (Electronic); 0892-7790 (Print); 0892-7790 (Linking). DOI: [10.1089/end.2017.0675](https://doi.org/10.1089/end.2017.0675).

Appendix A

COMPARING ONE-DIMENSIONAL, REDUCED-ORDER SPHERICAL BUBBLE DYNAMICS MODELS

In the following, we detail two recent models for one-dimensional modeling of spherical bubble dynamics. The first is the model developed by Estrada et al. (2018) and used in IMR and DA-IMR. The second is the model derived by Preston, Colonius, and Brennen (2007). We highlight key assumptions made in the derivation of each set of equations, comparing them across models. In the following, for clarity, we refer to these models as the *Estrada Model* and *Preston Model*. These notes complement the presentation of the spherical bubble dynamics model in Chapter 2, section 2.1, adding detail as to the model derivation, and comparing it to another similar reduced-order 1D bubble dynamics model.

A.1 Model Descriptions

A.1.1 Inertial Microcavitation High Strain-rate Rheometry – Estrada Model

The spherical bubble dynamics model derived in Estrada et al. (2018) and used in IMR involves the evolution of four quantities:

1. The bubble radius evolved with the Keller-Miksis equation (A.1) with stress integral given by equation (A.2).
2. The vapor mass fraction field inside the bubble, evolved with equation (A.6).
3. The temperature field inside the bubble, following equation (A.7).
4. The bubble pressure (assumed uniform), according to equation (A.8).

To come to this system of only four equations (two ordinary differential equations for radius and pressure, and two partial differential equations for the vapor mass fraction and temperature fields), a number of assumptions are made by Estrada et al. (2018).

Assumptions First, it is assumed that the motion of the bubble and its contents are spherically symmetric, thus reducing this to a 1-dimensional system. We also assume no water flux with respect to the surrounding material. That is, the surroundings remain undrained. As long as the surrounding material has high water content and simulations remain on a small time-scale (as is the case in all cases presented in this thesis), this assumption holds. Next, it is assumed that surrounding material is nearly incompressible, with compressible effects only considered in the region near the bubble wall, following the approach of Keller and Miksis (1980). A low Mach number approximation (Yang and Church, 2005) is used in the bubble, and we assume that the bubble contents are a mixture of 2 components, namely water vapor and a non-condensable gas. The pressure in the bubble is assumed spatially uniform, and we track a single bubble pressure value. We assume that the temperatures of the water vapor and non-condensable gas are equal, and track a single mixture temperature. Finally, we assume that the temperature of the surrounding material is constant. The Estrada model equations, as given in section 2.1, are as follows:

Radius The Keller-Miksis equation is used for the time-evolution of the bubble radius.

$$\begin{aligned} \left(1 - \frac{\dot{R}}{c}\right) R \ddot{R} + \frac{3}{2} \left(1 - \frac{\dot{R}}{3c}\right) \dot{R}^2 = \frac{1}{\rho} \left(1 + \frac{\dot{R}}{c}\right) \left(p_b - \frac{2\gamma}{R} + S - p_\infty\right) \\ + \frac{1}{\rho} \frac{R}{c} \overline{\left(p_b - \frac{2\gamma}{R} + S\right)}. \end{aligned} \quad (\text{A.1})$$

Stress Integral Given a neo-Hookean Kelvin-Voigt model for the surrounding material, the stress integral takes the form

$$S = -\frac{G}{2} \left[5 - \left(\frac{R_0}{R}\right)^4 - 4\frac{R_0}{R} \right] - \frac{4\mu\dot{R}}{R}. \quad (\text{A.2})$$

Vapor Mass Fraction It is assumed the bubble contents are a mixture of water-vapor and non-condensable gas (subscripts v and g , respectively). A mixture density $\rho_m = \rho_v + \rho_g$ is defined, as well as a mixture radial velocity $v_m = k_v v_v + k_g v_g$, where k_v and k_g are the vapor and gas mass fractions, respectively. Balance of mass

equations for each species are given by

$$\frac{\partial \rho_v}{\partial t} + \frac{1}{\partial r^2} \frac{\partial}{\partial r} (r^2 (\rho_v v_m + j_v)) = 0, \quad (\text{A.3})$$

$$\frac{\partial \rho_g}{\partial t} + \frac{1}{\partial r^2} \frac{\partial}{\partial r} (r^2 (\rho_g v_m + j_g)) = 0, \quad (\text{A.4})$$

where j_v and j_g represent radial mass flux for each species relative to the mixture. Combining these two equations yields a balance of mass equation for the mixture:

$$\frac{\partial \rho_m}{\partial t} + \frac{1}{r^2} \frac{\partial}{\partial r} (r^2 \rho_m v_m) = 0, \quad (\text{A.5})$$

which can be expressed in terms of k to obtain the final expression

$$\frac{\partial k_v}{\partial t} + v_m \frac{\partial k_v}{\partial r} = \frac{1}{\rho_m} \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \rho_m D \frac{\partial k_v}{\partial r} \right). \quad (\text{A.6})$$

(This requires using $\rho_v = k \rho_m$ and equation (A.5), as well as Fick's law).

Temperature

$$\rho_m C_p \left(\frac{\partial T}{\partial t} + v_m \frac{\partial T}{\partial r} \right) = \dot{p}_b + \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 K \frac{\partial T}{\partial r} \right) + \rho_m (C_{p,v} - C_{p,g}) D \frac{\partial k_v}{\partial r} \frac{\partial T}{\partial r}. \quad (\text{A.7})$$

Pressure

$$\begin{aligned} \dot{p}_b = \frac{3}{R} \left[-\kappa p_b \dot{R} + (\kappa - 1) K(T(R)) \frac{\partial T}{\partial r} \Big|_{r=R} \right. \\ \left. + \kappa p_b \frac{C_{p,v}}{C_p(k_v(R))} \frac{D}{1 - k_v(R)} \frac{\partial k_v}{\partial r} \Big|_{r=R} \right]. \end{aligned} \quad (\text{A.8})$$

We note that if we relax the last assumption listed of constant temperature and lack of water flux in the surrounding material, the temperature and vapor mass fraction equations are split into two separate equations: one for bubble contents and one for the surrounding material. In the following, subscripts 1 and 2 denote mixture quantities inside and outside the bubble, respectively.

Temperature (relaxed model assumptions)

$$\begin{aligned} \rho_1 C_{p,1} \left(\frac{\partial T_1}{\partial t} + v_1 \frac{\partial T_1}{\partial r} \right) = \dot{p}_b + \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 K_1 \frac{\partial T_1}{\partial r} \right) \\ + \rho_1 C_{p,1} \frac{C_{p,v} - C_{p,g}}{C_{p,1}} D_1 \frac{\partial k_{1,v}}{\partial r} \frac{\partial T_1}{\partial r}, \end{aligned} \quad (\text{A.9})$$

$$\rho_2 C_{p,2} \left(\frac{\partial T_2}{\partial t} + v_2 \frac{\partial T_2}{\partial r} \right) = \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 K_2 \frac{\partial T_2}{\partial r} \right) + 12 \mu_2 \left(\frac{v_2}{r} \right)^2, \quad (\text{A.10})$$

Vapor Mass Fraction (relaxed model assumptions)

$$\frac{\partial k_{1,v}}{\partial t} + v_1 \frac{\partial k_{1,v}}{\partial r} = \frac{1}{\rho_1} \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \rho_1 D_1 \frac{\partial k_{1,v}}{\partial r} \right) \quad (\text{A.11})$$

$$\frac{\partial k_{2,g}}{\partial t} + v_2 \frac{\partial k_{2,g}}{\partial r} = \frac{1}{\rho_2} \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \rho_2 D_2 \frac{\partial k_{2,g}}{\partial r} \right). \quad (\text{A.12})$$

A.1.2 Bubble Dynamics with Phase Change—Preston model

We now summarize the equations used by Preston, Colonius, and Brennen (2007) to solve for the bubble dynamics in the (i) liquid phase, (ii) gaseous phase, and (iii) bubble wall interface. The equations are:

- Liquid phase
 1. Bubble radius evolution
 2. Conservation of mass of the dissolved air/gas in the liquid
 3. Conservation of energy in the energy in the liquid (temperature)
- Gaseous phase (mixture of air and water vapor)
 1. Conservation of mass of the mixture
 2. Conservation of mass concentration of water vapor
 3. Conservation of momentum of the mixture
 4. Conservation of energy for the mixture
- Interface conditions
 1. Conservation of (air) mass
 2. Conservation of energy at the interface
 3. Interface conditions

Liquid phase

1. Integration of the liquid momentum equation from the bubble wall to infinity yields

$$\tilde{p}|_{r=R} - \tilde{p}_\infty(t) = R\dot{V} + \frac{3}{2}V^2 + \frac{4}{Re} \frac{V}{R} + \frac{2}{WeR} + \frac{4}{3} \frac{\mu}{Re} \left[\frac{\partial u}{\partial r} - \frac{u}{r} \right]_{r=R}, \quad (\text{A.13})$$

where r is the radial direction, $\tilde{p}|_{r=R}$ is the bubble wall pressure, $\tilde{p}|_{\infty}(t)$ specified far-field liquid pressure, R bubble radius, $V \equiv \dot{R}$, and μ and u the gas viscosity and velocity, respectively.

2. Conservation of energy in the liquid

$$\frac{\partial \tilde{T}_l}{\partial t} + u_l \frac{\partial \tilde{T}_l}{\partial r} = \frac{1}{Pr} \frac{1}{Re} \frac{\partial}{\partial r} \left(r^2 \frac{\partial \tilde{T}_l}{\partial r} \right) + \frac{12}{Re} \left(\frac{u_l}{r} \right)^2, \quad (\text{A.14})$$

where \tilde{T}_l is the liquid temperature.

3. Conservation of mass of the dissolved non-condensable gas in the liquid

$$\frac{\partial C_a}{\partial t} + u_l \frac{\partial C_a}{\partial r} = \frac{1}{Sc} \frac{1}{Re} \frac{\partial}{\partial r} \left(r^2 \frac{\partial C_a}{\partial r} \right), \quad (\text{A.15})$$

where $C_a \equiv \rho_a / \rho_l$ is the mass concentration of the dissolved gas in the liquid.

The natural frequency,

$$\omega_o = \sqrt{\frac{3(p'_{\infty o} - p'_v) + 4S'/R'_o}{\rho'_l R_o'^2}}, \quad (\text{A.16})$$

where is used to non-dimensionalize the equations and yield the following definitions

$$Re = \frac{\rho'_l R_o'^2 \omega_o'}{\mu'_l}, \quad We = \frac{\rho'_l R_o'^3 \omega_o'^2}{S'}, \quad Sc = \frac{\mu'_l}{\rho'_l D'_l}, \quad Pr = \frac{\mu'_l c'_{pl}}{\kappa'_l}.$$

Gaseous phase

1. Conservation of the mass for the mixture

$$\frac{\partial \rho}{\partial t} + \frac{1}{r^2} \frac{\partial}{\partial r} (r^2 \rho u) = 0, \quad (\text{A.17})$$

where ρ and u are the mixture density and velocity, respectively.

2. Conservation of mass vapor in the mixture yields

$$\frac{\partial \rho C}{\partial t} + \frac{1}{r^2} \frac{\partial}{\partial r} (r^2 \rho C u) = \frac{1}{Sc} \frac{1}{Re} \frac{\partial}{\partial r} \left(r^2 \rho D \frac{\partial C}{\partial r} \right), \quad (\text{A.18})$$

where $C \equiv \rho_v / \rho$ is the mass concentration of vapor. Conservation of momentum for the mixture becomes

$$\begin{aligned} \frac{\partial \rho u}{\partial t} + \frac{1}{r^2} \frac{\partial}{\partial r} (r^2 \rho u^2) = \\ - \frac{\partial \tilde{p}}{\partial r} + \frac{1}{Re} \frac{4}{3} \left(\frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \mu \left[\frac{\partial u}{\partial r} - \frac{u}{r} \right] \right) + \frac{\mu}{r} \left[\frac{\partial u}{\partial r} - \frac{u}{r} \right] \right), \end{aligned} \quad (\text{A.19})$$

where pressure is $\tilde{p} = \rho \mathcal{R} \tilde{T}$ and the gas temperature is \tilde{T} . The effective gas constant is determined by a weighted average $\mathcal{R} = C \mathcal{R}_v + (1 - C) \mathcal{R}_a$.

3. Conservation of energy for the mixture

$$c_v \left(\frac{\partial \rho \tilde{T}}{\partial t} + \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \rho u \tilde{T} \right) \right) = \frac{1}{Re} \frac{1}{Pr} \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 k \frac{\partial \tilde{T}}{\partial r} \right) - \tilde{p} \frac{1}{r^2} \frac{\partial}{\partial r} (r^2 u) + \frac{4}{3} \frac{\mu}{Re} \left[\frac{\partial u}{\partial r} - \frac{u}{r} \right]^2, \quad (\text{A.20})$$

where $c_v = C c_{v_v} + (1 - C) c_{v_a}$.

Interface conditions

1. Conservation of mass of air

$$\dot{m}''_a = \frac{-1}{Sc} \frac{\rho D}{Re} \frac{\partial C}{\partial r} - \rho(1 - C)(u - V) = \frac{1}{Sc} \frac{\partial C_a l}{Re \partial r}, \quad (\text{A.21})$$

where $\dot{m}''_a = \dot{m}''_a' / \rho_l R'_o \omega'_o$.

2. Conservation of energy can be written as

$$L \dot{m}''_v = \frac{1}{Pr} \frac{1}{Re} \left[\frac{\partial \tilde{T}}{\partial r} - k \frac{\partial \tilde{T}}{\partial r} \right] + (u - V) \left[\rho \left(c_v \tilde{T} + \frac{1}{2} u^2 \right) + \tilde{p} - \frac{4}{3} \frac{\mu}{Re} \left(\frac{\partial u}{\partial r} - \frac{u}{r} \right) \right], \quad (\text{A.22})$$

where $L = L' / (R'_o \omega'_o)^2$ and $\dot{m}''_v = \dot{m}''_v' / \rho_l R'_o \omega'_o$.

3. Additional interface conditions are

$$\tilde{T} = \tilde{T}_l = \tilde{T}_w, \quad (\text{A.23})$$

and

$$C_a = H \tilde{p}_a = H \rho (1 - C) \mathcal{R}_a \tilde{T}, \quad (\text{A.24})$$

4. Mass flux of vapor out of the bubble

$$\dot{m}''_v = \alpha \frac{\tilde{p}_{v_{\text{sat}}}(\tilde{T}) - \tilde{p}_v}{\sqrt{2\pi \mathcal{R}_v \tilde{T}}}, \quad (\text{A.25})$$

which is equated to the mass flux of vapor due to bulk motion and reciprocal diffusion on the gas side of the interface,

$$\dot{m}''_v = \frac{-1}{Sc} \frac{\rho D}{Re} \frac{\partial C}{\partial r} - \rho(1 - C)(u - V), \quad (\text{A.26})$$

which is the same as equation (A.21) but for the vapor. Finally using the perfect gas law, we can re-write the vapor pressure at the bubble wall as

$$\tilde{p}_v = \rho C \mathcal{R}_v \tilde{T}. \quad (\text{A.27})$$

A.2 From the Preston Model to the Estrada Model

In this section, we aim to recover the Estrada model described in section from the more general the Preston model. Four assumptions must be made for this, which we describe in detail here, before applying them to the equations of section A.1.2 to recover the set of four equations of section A.1.1. While these are mentioned in section A.1.1, their limitations and regimes of validity are further explained here. They are as follows:

1. Insoluble gas

Given the fast timescales of growth and collapse of the bubble (and low solubility of air in water), the effect of transfer of mass of the non-condensable gas may be negligible. When looking over many repetitions of the cycle, the mass of gas in the bubble will increase (and with it the equilibrium radius), but if looking at a few bubble collapse and growth cycles only, this effect can be neglected and we can thus remove equation (A.15).

Additionally, interface conditions (A.21) and (A.24) can be removed, and we can rewrite (A.26) as

$$\dot{m}_v'' = \frac{D}{Sc} \frac{\rho_w}{Re} \frac{\partial C}{\partial r} \Big|_{r=R}. \quad (\text{A.28})$$

2. Cold Liquid

If the liquid temperature is low enough, temperature changes in the liquid can be neglected because the slope of the vapor saturation curve is small at low temperatures, and temperature in the liquid remains relatively constant in these cases. Thus we can remove equation (A.14) and replace the interface condition with:

$$\tilde{T}_w = \tilde{T}_{l_\infty} = \text{constant}. \quad (\text{A.29})$$

3. Equilibrium Phase change

equation (A.25) restricts the rate of phase change at the interface. If the bubble

motion is slow enough, though, this rate is not reached and instead the finite rate of diffusion of vapor in the non-condensable gas near the wall limits this rate. We can then assume the vapor pressure at the wall is in equilibrium with the bubble wall temperature, enabling the substitution of $\tilde{p}_v = \tilde{p}_{v,\text{sat}}(\tilde{T})$ into equation (A.27) to yield the new interface condition

$$\rho_w C_w = \frac{\tilde{p}_{v,\text{sat}}(\tilde{T}_w)}{\mathcal{R}_v \tilde{T}_w}. \quad (\text{A.30})$$

4. Homobaricity

Assuming uniform pressure in the bubble yields an ODE for the internal bubble pressure (by integrating equation (A.19) and combining with equation (A.17)):

$$\frac{d\tilde{p}}{dt} = \frac{-3\gamma}{R} \left[\tilde{p}V - \frac{\gamma-1}{\gamma} \frac{k_w}{PrRe} \frac{\partial \tilde{T}}{\partial r} \Big|_{r=R} - \mathcal{R}_v \tilde{T}_w \dot{m}_v'' \right]. \quad (\text{A.31})$$

Furthermore, equation (A.20) can be re-written as

$$\frac{\partial \tilde{T}}{\partial t} + U \frac{\partial \tilde{T}}{\partial r} = \frac{1}{PrRe\rho c_p} \frac{1}{r^2} \frac{\partial}{\partial r} \left(kr^2 \frac{\partial \tilde{T}}{\partial r} \right) + \frac{\gamma-1}{\gamma} \tilde{T} \frac{\dot{\tilde{p}}}{\tilde{p}}. \quad (\text{A.32})$$

Given these four assumptions, the Preston model simplifies to equations (A.13), (A.31), (A.32) and (A.18). The only remaining necessary boundary conditions are equations (A.29) and (A.30).

We compare the model obtained with these four assumptions to the Estrada model:

Liquid Phase: Momentum Equation Preston (2004) integrates the momentum equation for the liquid phase assuming constant liquid viscosity and constant liquid density. This yields the Rayleigh-Plesset equation, except for a relaxed assumption of polytropic behavior of the internal bubble pressure

$$\tilde{p}|_{r=R} - \tilde{p}_\infty(t) = R\dot{V} + \frac{3}{2}V^2 + \frac{4}{Re} \frac{V}{R} + \frac{2}{WeR} + \frac{4}{3} \frac{\mu}{Re} \left[\frac{\partial u}{\partial r} - \frac{u}{r} \right]_{r=R}. \quad (\text{A.33})$$

Estrada et al. (2018), however, use a Keller-Miksis formulation for this (relaxing the incompressibility assumption of Preston (2004) and Preston, Colonius, and Brennen

(2007)), given by

$$\begin{aligned} \left(1 - \frac{\dot{R}}{c}\right) R \ddot{R} + \frac{3}{2} \left(1 - \frac{\dot{R}}{3c}\right) \dot{R}^2 = \frac{1}{\rho} \left(1 + \frac{\dot{R}}{c}\right) \left(p_b - \frac{2\gamma}{R} + S - p_\infty\right) \\ + \frac{1}{\rho} \frac{R}{c} \left(p_b - \frac{2\gamma}{R} + S\right). \end{aligned} \quad (\text{A.34})$$

Both of these equations are commonly used to describe spherical bubble dynamics, with the Keller-Miksis a slightly more general formulation, whereby the surrounding material is no longer assumed to be fully incompressible.

Given the insoluble gas and cold liquid assumptions detailed above, there is no longer a need for mass and energy conservation equations in the liquid phase.

Gas Phase: Mass Conservation Equation (A.18) is used for mass conservation in the gas phase in the Preston model:

$$\frac{\partial \rho C}{\partial t} + \frac{1}{r^2} \frac{\partial}{\partial r} (r^2 \rho C u) = \frac{1}{Sc} \frac{1}{Re} \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \rho D \frac{\partial C}{\partial r} \right). \quad (\text{A.35})$$

This can be rewritten (expanding partial derivatives on LHS) as:

$$\left(C \frac{\partial \rho}{\partial t} + \rho \frac{\partial C}{\partial t} \right) + \left(\frac{r^2 \rho u}{r^2} \frac{\partial C}{\partial r} + \frac{C}{r^2} \frac{\partial (r^2 \rho u)}{\partial r} \right) = \frac{1}{Sc} \frac{1}{Re} \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \rho D \frac{\partial C}{\partial r} \right) \quad (\text{A.36})$$

$$\iff C \left(\frac{\partial \rho}{\partial r} + \frac{1}{r^2} \frac{\partial (r^2 \rho u)}{\partial r} \right) + \rho \left(\frac{\partial C}{\partial t} + u \frac{\partial C}{\partial r} \right) = \frac{1}{Sc} \frac{1}{Re} \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \rho D \frac{\partial C}{\partial r} \right), \quad (\text{A.37})$$

but by conservation of mass for the mixture,

$$\frac{\partial \rho}{\partial r} + \frac{1}{r^2} \frac{\partial (r^2 \rho u)}{\partial r} = 0 \quad (\text{A.38})$$

And thus dividing by ρ we obtain (a non-dimensional form of) Estrada's equation:

$$\frac{\partial C}{\partial t} + u \frac{\partial C}{\partial r} = \frac{1}{\rho} \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \rho D \frac{\partial C}{\partial r} \right). \quad (\text{A.39})$$

Gas Phase: Momentum Equation Given the homobaricity assumption, equation (A.31) is used in the Preston model:

$$\frac{d\tilde{p}}{dt} = \frac{-3\gamma}{R} \left[\tilde{p}V - \frac{\gamma-1}{\gamma} \frac{k_w}{PrRe} \frac{\partial \tilde{T}}{\partial r} \Big|_{r=R} - \mathcal{R}_v \tilde{T}_w \dot{m}_v'' \right]. \quad (\text{A.40})$$

We can re-write this by multiplying through by γ and substituting in our expression for \dot{m}_v'' in (A.28):

$$\frac{d\tilde{p}}{dt} = \frac{3}{R} \left[-\gamma \tilde{p}V + (\gamma-1) \frac{k_w}{PrRe} \frac{\partial \tilde{T}}{\partial r} \Big|_{r=R} + \gamma \mathcal{R}_v \tilde{T}_w \rho_w \frac{D}{1-C_w} \frac{\partial C}{\partial r} \Big|_{r=R} \right], \quad (\text{A.41})$$

and using (A.30), we can write

$$\mathcal{R}_v \tilde{T}_w \rho_w = \tilde{p} \frac{\mathcal{R}_v \tilde{T}_w \rho_w}{\tilde{p}} \quad (\text{A.42})$$

$$= \tilde{p} \frac{C_{p,v}}{C_p}, \quad (\text{A.43})$$

which yields the Estrada estimate

$$\frac{d\tilde{p}}{dt} = \frac{3}{R} \left[-\gamma \tilde{p}V + (\gamma-1) \frac{k_w}{PrRe} \frac{\partial \tilde{T}}{\partial r} \Big|_{r=R} + \gamma \tilde{p} \frac{C_{p,v}}{C_p} \frac{D}{1-C_w} \frac{\partial C}{\partial r} \Big|_{r=R} \right]. \quad (\text{A.44})$$

Gas Phase: Energy Equation Preston (2004) show that equation (A.32), given the homobaricity assumption, can be written as

$$\frac{\partial \tilde{T}}{\partial t} + U \frac{\partial \tilde{T}}{\partial r} = \frac{1}{PrRe\rho c_p} \frac{1}{r^2} \frac{\partial}{\partial r} \left(kr^2 \frac{\partial \tilde{T}}{\partial r} \right) + \frac{\gamma-1}{\gamma} \tilde{T} \frac{\dot{\tilde{p}}}{\tilde{p}}, \quad (\text{A.45})$$

while the Estrada model uses:

$$\rho_m C_p \left(\frac{\partial T}{\partial t} + v_m \frac{\partial T}{\partial r} \right) = \dot{p}_b + \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 K \frac{\partial T}{\partial r} \right) + \rho_m (C_{p,v} - C_{p,g}) D \frac{\partial k_v}{\partial r} \frac{\partial T}{\partial r}. \quad (\text{A.46})$$

Expanding U in equation (A.45) as described in section 2.4.4 of Preston (2004) and with some algebra, the Estrada model equation can be recovered given this assumption.

Boundary Conditions Finally, the same boundary conditions are used in both models given the added assumptions. The temperature at the bubble wall is

$$\tilde{T}_w = \tilde{T}_{l_\infty} = \text{constant}, \quad (\text{A.47})$$

and the vapor mass fraction equilibrium is given by

$$\rho_w C_w = \frac{\tilde{p}_{v,\text{sat}}(\tilde{T}_w)}{\mathcal{R}_v \tilde{T}_w}. \quad (\text{A.48})$$

This completes the comparison of the two models. Given the four key assumptions listed at the beginning of this section, the Preston model simplifies to the Estrada model, with the exception of the bubble radius evolution equation, which is given by the Rayleigh-Plesset equation in Preston, Colonius, and Brennen (2007), and by the Keller-Miksis equation in Estrada et al. (2018). We note that this is a particularly important distinction in the context of bubble dynamics in a viscoelastic material, as the stress integral S in the Keller-Miksis equation is defined by the material model used. The use of the Keller-Miksis equation with this stress integral term is central to the IMR and DA-IMR methods described in Chapters 2 and 3, where terms in the material-model-specific expression of S are estimated to characterize soft materials.

Appendix B

PERFORMANCE AND SCALING OF MFC

The majority of simulation results shown in Chapters 4 and 5 of this thesis were performed using the Multi-component Flow Code (MFC) (Bryngelson et al., 2021; Radhakrishnan et al., 2023), with the additional models described in section 4.1. MFC, with the aforementioned models now implemented, is open source and can be accessed at <https://mflowcode.github.io/>. In this appendix, we report on performance and scaling results of MFC on both CPUs and GPUs, and provide details on the performance impact of the implemented hypoelastic and damage models.

B.1 CPU Scaling and Performance

MFC employs an interface-capturing approach, high-order accurate finite-volume weighted essentially non-oscillatory (WENO) reconstructions of the primitive flow variables at the cell-boundaries and Harten-Lax-van Leer-type (HLL and HLLC) approximate Riemann solvers for upwinding. Explicit total-variation diminishing (TVD) Runge-Kutta schemes are used to march the solution in time. To solve the physical model, the five-equation numerical model described in Chapter 4, section 4.1 is used. Inter-processor communication is handled by MPI (message-passing interface). The algorithm used by MFC is well-suited for parallel applications, due to the large number of operations required per grid point to evolve the solution. To ascertain the scaling of MFC, we performed strong and weak scaling tests up to 4096 CPU cores. A weak scaling test keeps the number of grid points *per processor* fixed. Ideal weak scaling means that the wall time remains the same as we increase the number of processors with the problem size. In a strong scaling test, on the other hand, the problem size is kept fixed, and we report the decrease in wall time as the number of processors increases. the results of both scaling tests are shown in figure B.1.

This demonstrates ideal weak scaling (figure B.1b) and near-ideal strong scaling (figure B.1a) of MFC. The slight deviation from ideal scaling in the strong scaling test as problem size increases to $O(1000)$ cores is expected. This results from




Figure B.1: Strong scaling (a) of MFC for 500^3 grid points and weak scaling (b) for 50^3 grid points per core, both up to 4096 cores. Scaling test performed on up to 32 compute nodes on the *Expansive* supercomputer at the San Diego Supercomputer Center.

the increase in the number of halo regions, where the edge of the domain of each block is shared with the adjacent block, as the number of processors increases. These scaling tests were performed on the *Expansive* supercomputer at the San Diego Supercomputer Center, with up to 32 nodes in 3D with third-order accurate WENO spatial reconstruction and TVD Runge-Kutta time marching schemes. The average grid time per flow variable was found to be $3.7483 \mu\text{s}$. We now describe the implementation of the code on GPUs to accelerate its performance.

B.2 GPU Acceleration

B.2.1 Performance Speedup

A benchmark of the CPU implementation of the code revealed that the majority of the computational cost came for only two routines, namely the WENO-reconstruction and Riemann solver, which combined account for more than 90% of the compute time per time step. The first row of table B.2 shows the distribution of computational cost across routines on CPUs for a representative 2-component, 3D test case.

This highlighted the potential for straightforward acceleration of the code by making use of GPU offloading for these two subroutines, in a heterogeneous (combining CPU and GPU) approach. The GPU programming framework used is OpenACC (Wienke et al., 2012), which is a directive-based programming language readily implemented within a Fortran code such as MFC. OpenACC renders a CPU/GPU

heterogeneous approach easy to implement, as OpenACC directives can be added to specific subroutines to create GPU kernels executed in parallel on GPU accelerators.

The GPU offloading of these two subroutines proved effective, with $O(100)\times$ speedup for each individual routine when comparing 1 GPU to 1 CPU core. In particular, the WENO reconstruction is nearly $500\times$ faster on an NVIDIA V100 GPU when compared to a single POWER9 CPU core. This heterogeneous CPU/GPU implementation was overall approximately 50 times faster than the base CPU code, when comparing total wall time per time step. Benchmarking this implementation, however, revealed that there was potential for much greater speedup, as up to 80% of the cost now came from data transfer, up from under 5% with the CPU code. This is attributed to back-and-forth data transfers which are now required between CPU and GPU at each time step.

While the GPU offloading of other routines with low operation counts may not yield significant acceleration in terms of their individual compute times, a significant advantage can be gained by running all computing on GPUs. That is, we would avoid any extraneous data transfer back and forth from host (CPU) to device (GPU), which is a significant downside of the heterogeneous CPU/GPU approach. We thus offloaded all operations to run on GPUs, as detailed in Radhakrishnan et al. (2023), which significantly increased the speedup to more than $300\times$ (from $50\times$ with the heterogeneous code). Speedup numbers for key routines and overall code are reported in table B.1

Routine	Speedup
WENO reconstruction	486
Riemann Solver	189
Communication	125
Total	305

Table B.1: Speedup of GPU code running on one NVIDIA V100 GPU when compared to one POWER9 CPU core. These speedup numbers are detailed in Radhakrishnan et al. (2023)

This table shows that while speedup of other routines did not match that of the WENO reconstruction, their speedup was sufficient to achieve a high overall speedup of the code, approximately 6 times faster than the heterogeneous implementation due to the decreased data transfer and communication required.

Table B.2 compares benchmarking of key routines in the code for the final GPU implementation and CPU code. This reveals that while WENO reconstruction remains the most expensive routine, its relative cost has been far diminished due to its high speedup. Other routines and communications now represent a higher contribution to total compute time, given their lower GPU parallelization potential. Given that WENO reconstruction represented near three quarters of the compute time on CPU, the very high speedup of that routine drives the overall speedup of the code, despite the smaller acceleration seen in other routines.

Case	WENO	Riemann Solver	Communication	Rest of Code
CPU code	71.7 %	18.7 %	3.8 %	5.8 %
GPU code	43.3 %	33.3 %	5.8 %	17.6 %

Table B.2: Percentage of time-contribution of the WENO Reconstruction, Riemann solver, communication, and all other routines to total compute time.

B.2.2 Scaling Tests

Scaling tests for the GPU implementation of MFC were performed and reported in Radhakrishnan et al. (2023). Weak and strong scaling results are reproduced here in figure B.2.

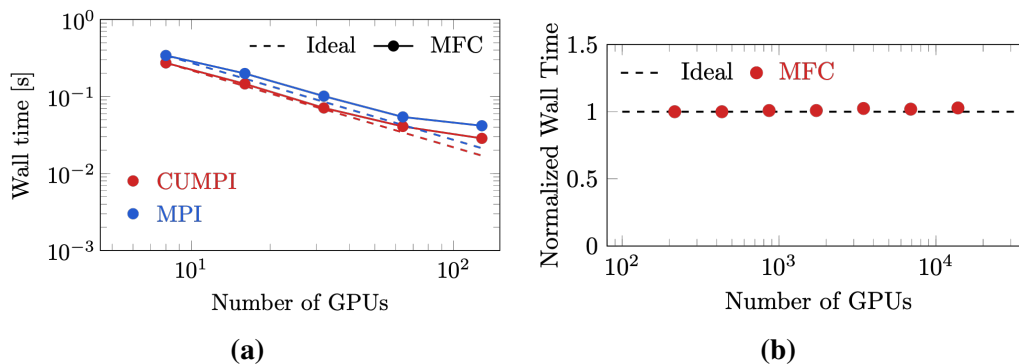


Figure B.2: Strong scaling (a) of MFC for 64M grid points up to 128 GPU with and without CUDA-aware MPI, and weak scaling (b) for 1M grid points per GPU up to $O(10000)$ GPUs. Scaling test performed on the *Summit* supercomputer at the Oak Ridge Leadership Computing Facility and reproduced from Radhakrishnan et al. (2023).

As with the CPU code, we see ideal weak scaling results. We highlight that this weak scaling test was run on up to $O(10000)$ GPUs, which given the speedup detailed in section B.2.1 corresponds to ideal weak scaling up to the performance-equivalent

of $O(1M)$ CPU cores. Strong scaling results also shows near-ideal scaling up to 64 GPUs, but begins to deviate for larger numbers of GPUs, where MPI communication becomes dominant in compute time. We note that in all simulations shown in this thesis, number of GPUs used never exceeded 64, which was sufficient to run even the largest cases shown in under 24 hours. To conclude this appendix on the performance and scaling of MFC, we discuss the slowdown caused by the additional models added to the code, described in Chapter 4, section 4.1.

B.3 Added Computational Cost of the Hypoelastic and Damage Models

The added computational cost incurred by adding the hypoelastic and damage models are discussed in the last part of section 4.1.4. As mentioned, when the hypoelastic model is turned on, up to 6 additional equations are solved in 3D, one for each distinct element of the 3×3 elastic stress tensor. Furthermore, the damage model requires calculating an additional damage field D dependent on the maximum principal stress, which must be calculated at all points in the solid domain. This has been shown to add up to 155% computational cost on CPUs, or 98% computational cost on GPUs for 3D simulations, as reported in section 4.1.4. However, this incurs no change to the scaling performance of the code given that the parallel infrastructure remains the same. For completeness, we report the full slowdown numbers in 1D, 2D and 3D for both CPU and GPU, in table B.3. Note that the 3D numbers reported are the same as those from table 4.1. Percentage slowdown relative to Base MFC (without the added hypoelastic and damage models) are reported for each added model. Numbers are for a simulation with 10M grid points, with a solid immersed in a high pressure fluid, run on a CPU node with 128 cores, or a GPU node with 4 GPUs on the *Expanse* supercomputer at the San Diego Supercomputer Center.

We see that for 1D and 2D cases, the slowdown on GPUs is marginally better than that of CPUs for corresponding cases. However, in 3D there is a significant difference in the added computational cost. The hypoelastic model alone is 84% slower as compared to 119% slower on CPUs, and when adding both models, the slowdown is 98% on GPUs and 155% on CPUs. This is because the six added equations in 3D represent a significant added cost to the WENO reconstruction, which is much faster on GPUs. Thus, this translates to a more minor slowdown in that version of the code.

Version	Dimension	Base MFC	Hypoelastic	Hypoelastic and Damage
CPU	1D	1.76 s	2.18 s (24%)	2.50 s (42%)
	2D	3.70 s	5.46 s (48%)	5.83 s (58%)
	3D	8.83 s	19.4 s (119%)	22.5 s (155%)
GPU	1D	0.40 s	0.48 s (20%)	0.54 s (34%)
	2D	0.87 s	1.25 s (45%)	1.36 s (56%)
	3D	1.99 s	3.66 s (84%)	3.93 s (98%)

Table B.3: Total time per time-step comparison for a 3D case with 10 million grid points with two components, with various MFC configurations. Percentages shown correspond to slowdown relative to the base MFC case for each version.

We also report the speedup numbers for the GPU implementation with both the hypoelastic and damage models turned on in 2D and 3D in table B.4. These are node-to-node comparisons, where we compare a full GPU node with 4 NVIDIA V100 GPUs to a full CPU node with 128 CPU cores, again on *Expanse*.

Version	CPU node time	GPU node time	Speedup
2D	0.58 s	0.14 s	4.28×
3D	2.25 s	0.39 s	5.73×

Table B.4: Speedup of GPU code compared to CPU code in 2D and 3D, when running on a full CPU node (128 cores) and full GPU node (4 GPUs) on *Expanse*

We see approximately a 4.3× and 5.7× speedup in 2D and 3D, respectively. This means that in a node-to-node comparison, 1 GPU-hour on that system with the hypoelastic and damage models turned on is approximately equal to 136 CPU core-hours for a 2D simulation and 183 CPU core-hours for a 3D simulation. Where available, running the GPU version of the code is thus also advantageous with the hypoelastic and damage models, enabling faster and scalable computations.

Example Simulation To conclude this section on the performance of MFC on GPUs with hypoelasticity, we show the results of a simulation of shock-induced collapse of air bubbles near a model kidney stone, reported in Radhakrishnan et al. (2023). Here, we consider the 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 properties of the stone are those of a BegoStone (Liu and Zhong, 2002), a kidney stone phantom used in lithotripsy research, which has been used throughout

this thesis. It has a density of $\rho = 1995 \text{ kg/m}^3$ and longitudinal and transverse wavespeeds $c_L = 4159 \text{ m s}^{-1}$ and $c_T = 2319 \text{ m s}^{-1}$.

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 the *Summit* supercomputer at the Oak Ridge Leadership Computing Facility for 25×10^3 time steps, corresponding to a wall-time of 30 minutes.




Figure B.3: 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 $\alpha = 0.5$.

Figure B.3 shows maximum principal stresses in the stone at non-dimensional time $\tau \equiv tc/R_0 = 2.03$ for air sound speed c and initial bubble radius R_0 . These stresses follow from the shock, and the subsequent bubble collapses. We observe larger stresses near the center of the bubble cloud and stone cross-section.