CSELAB
Computation Science and Engineering Laboratory
www.cse-lab.ch

ETH zürich

Uncertainty Quantification in Cloud Cavitation Collapse using Multi-Level Monte Carlo

Jonas Šukys

Computational Science and Engineering Laboratory
ETH Zurich, Switzerland

SIAM Conference on Uncertainty Quantification
Lausanne, Switzerland

April 7, 2016

Work in progress in collaboration with

Petros
Koumoutsakos

Ursula
Rasthofer

Panagiotis
Hadjidoukas

Diego
Rossinelli

Fabian
Wermelinger

Cavitation phenomenon

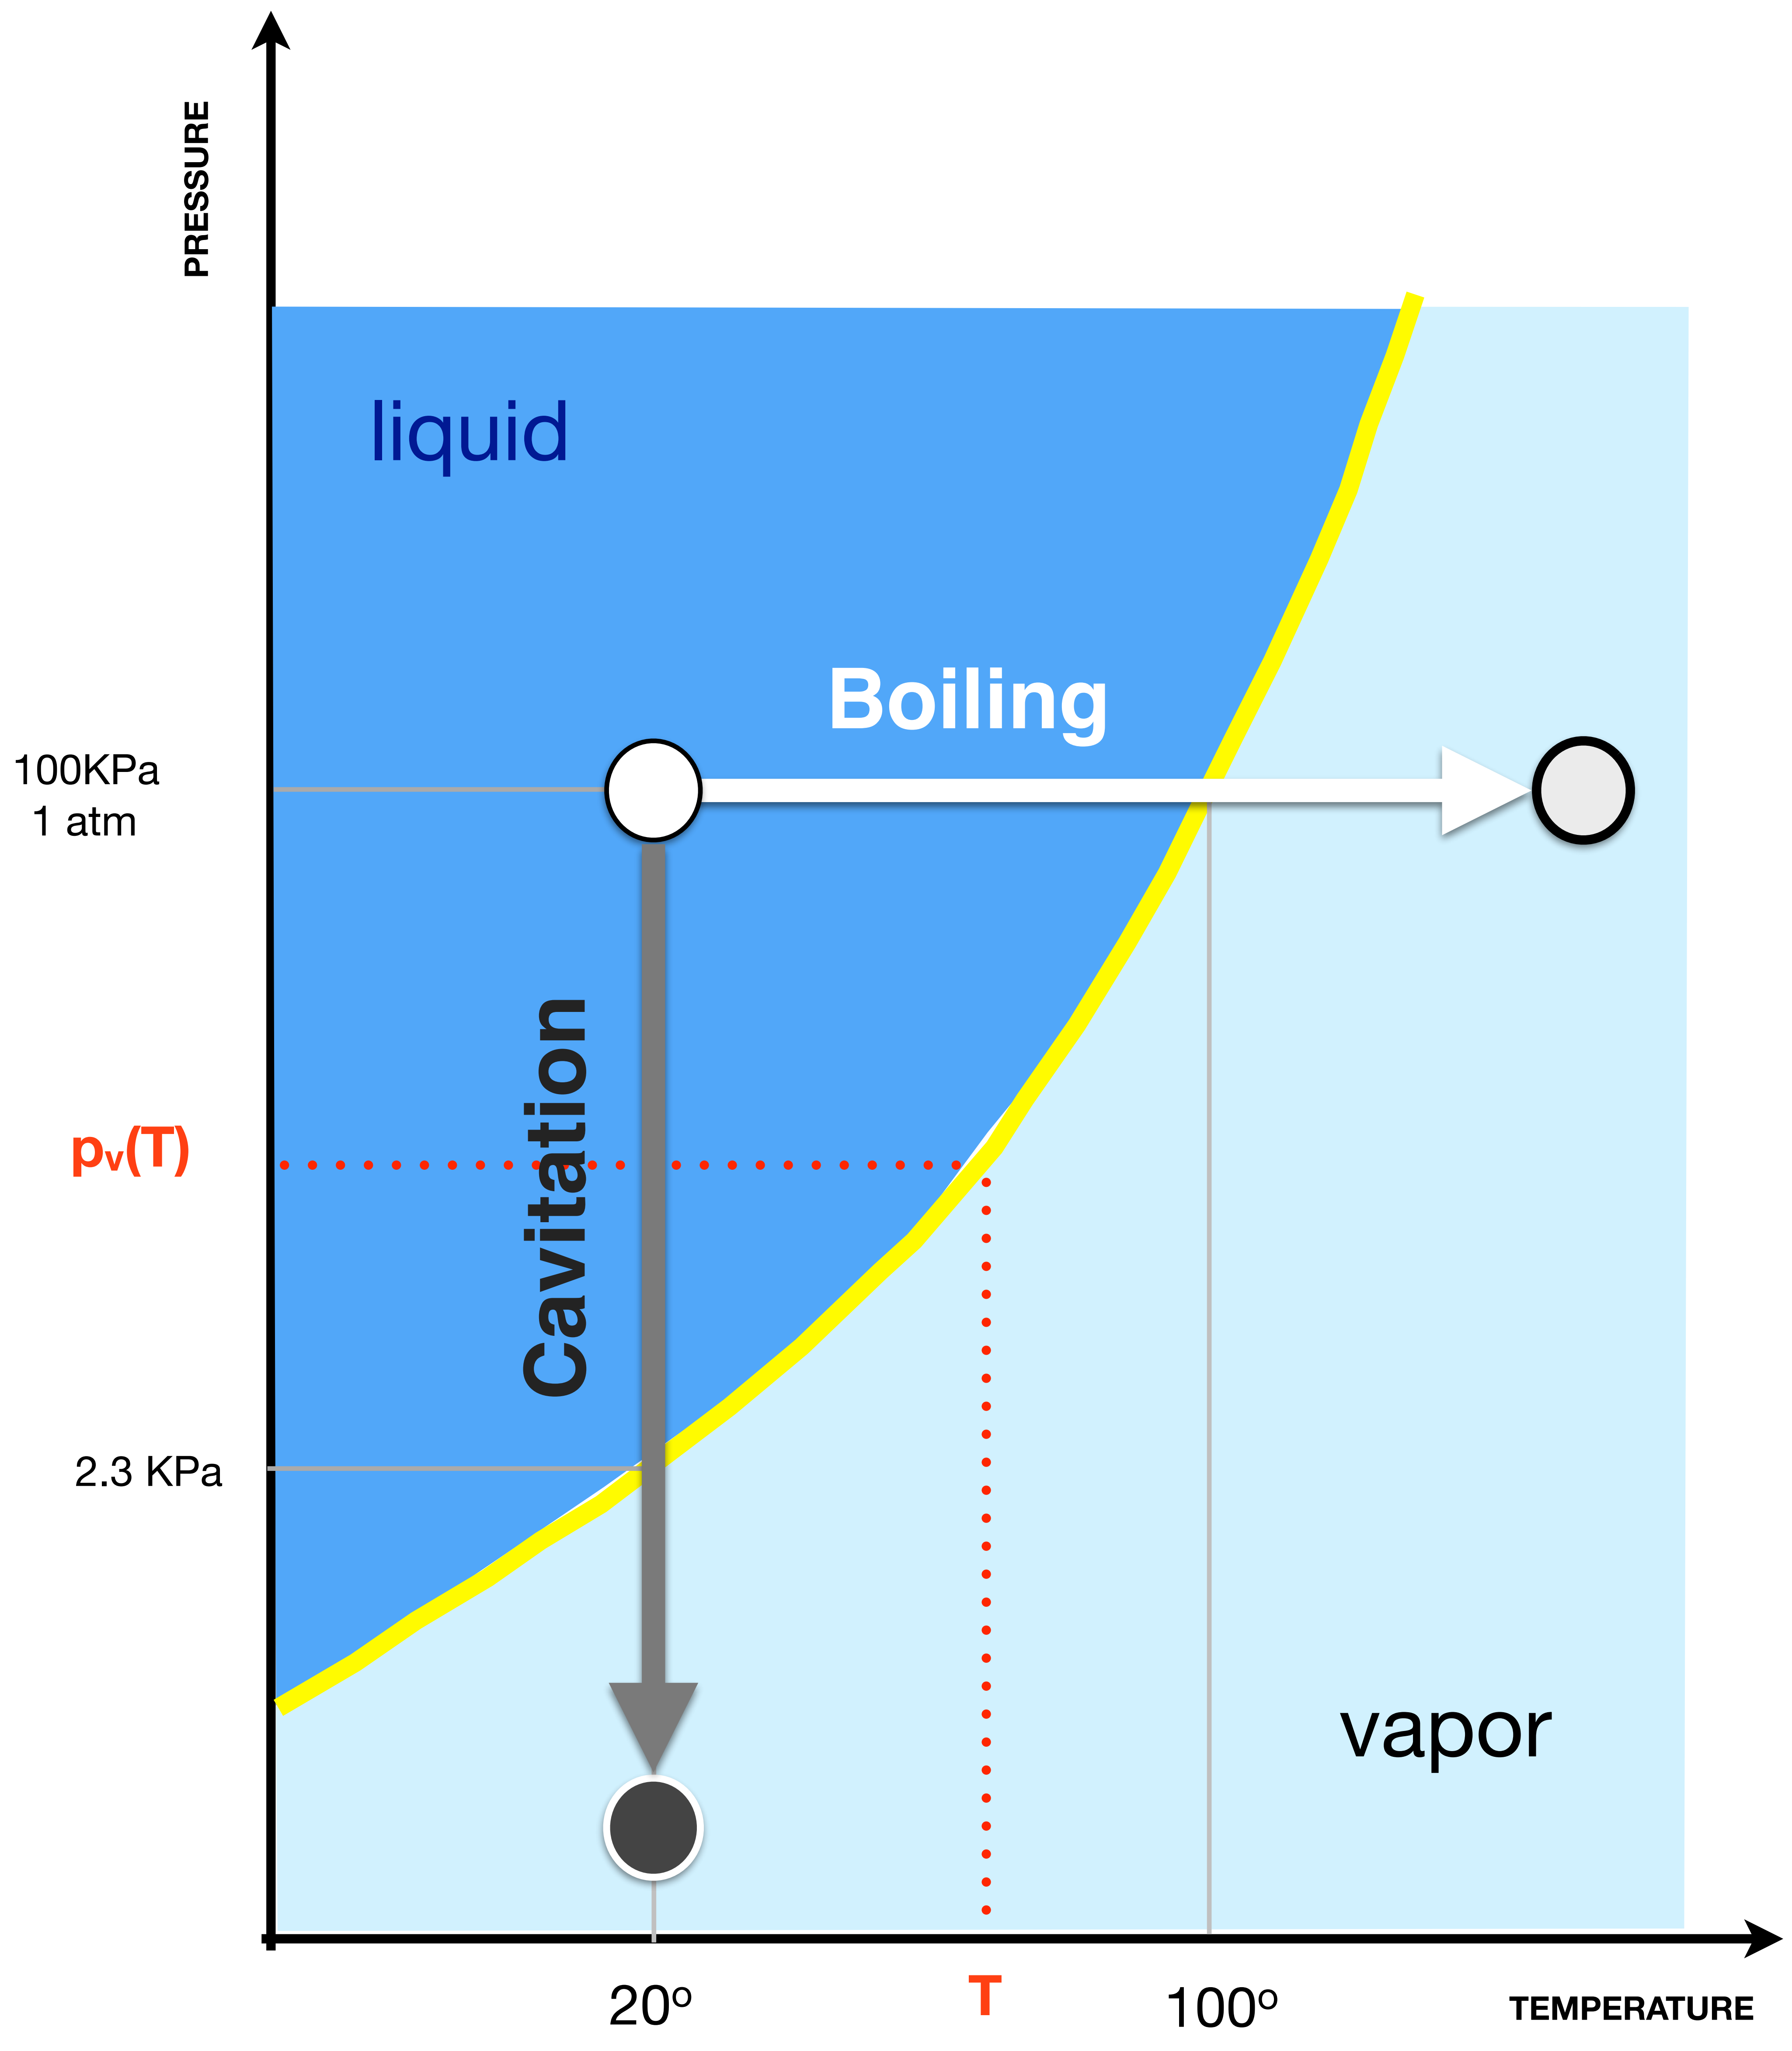

Image courtesy:
C. Koumoutsakos

$$p + \frac{1}{2}\rho u^2 = \text{const.}$$

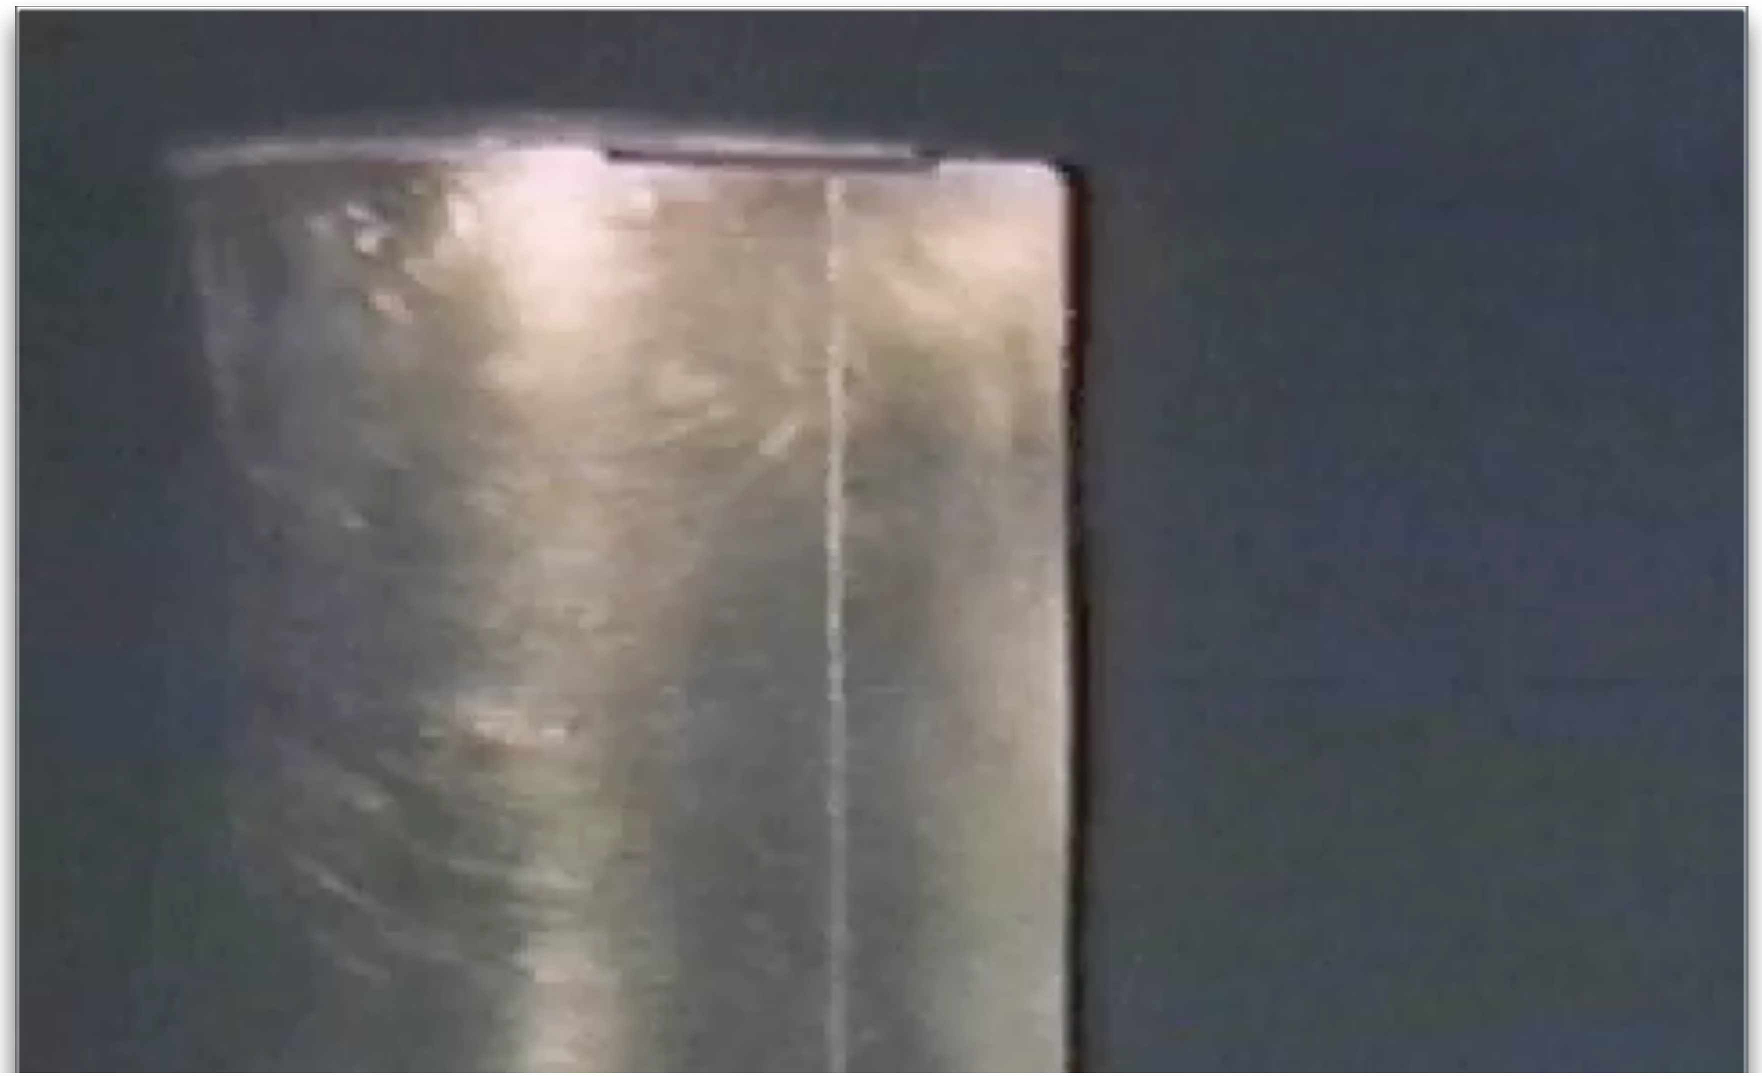

Image courtesy:
C. Brennen

Single cavity bubble collapse

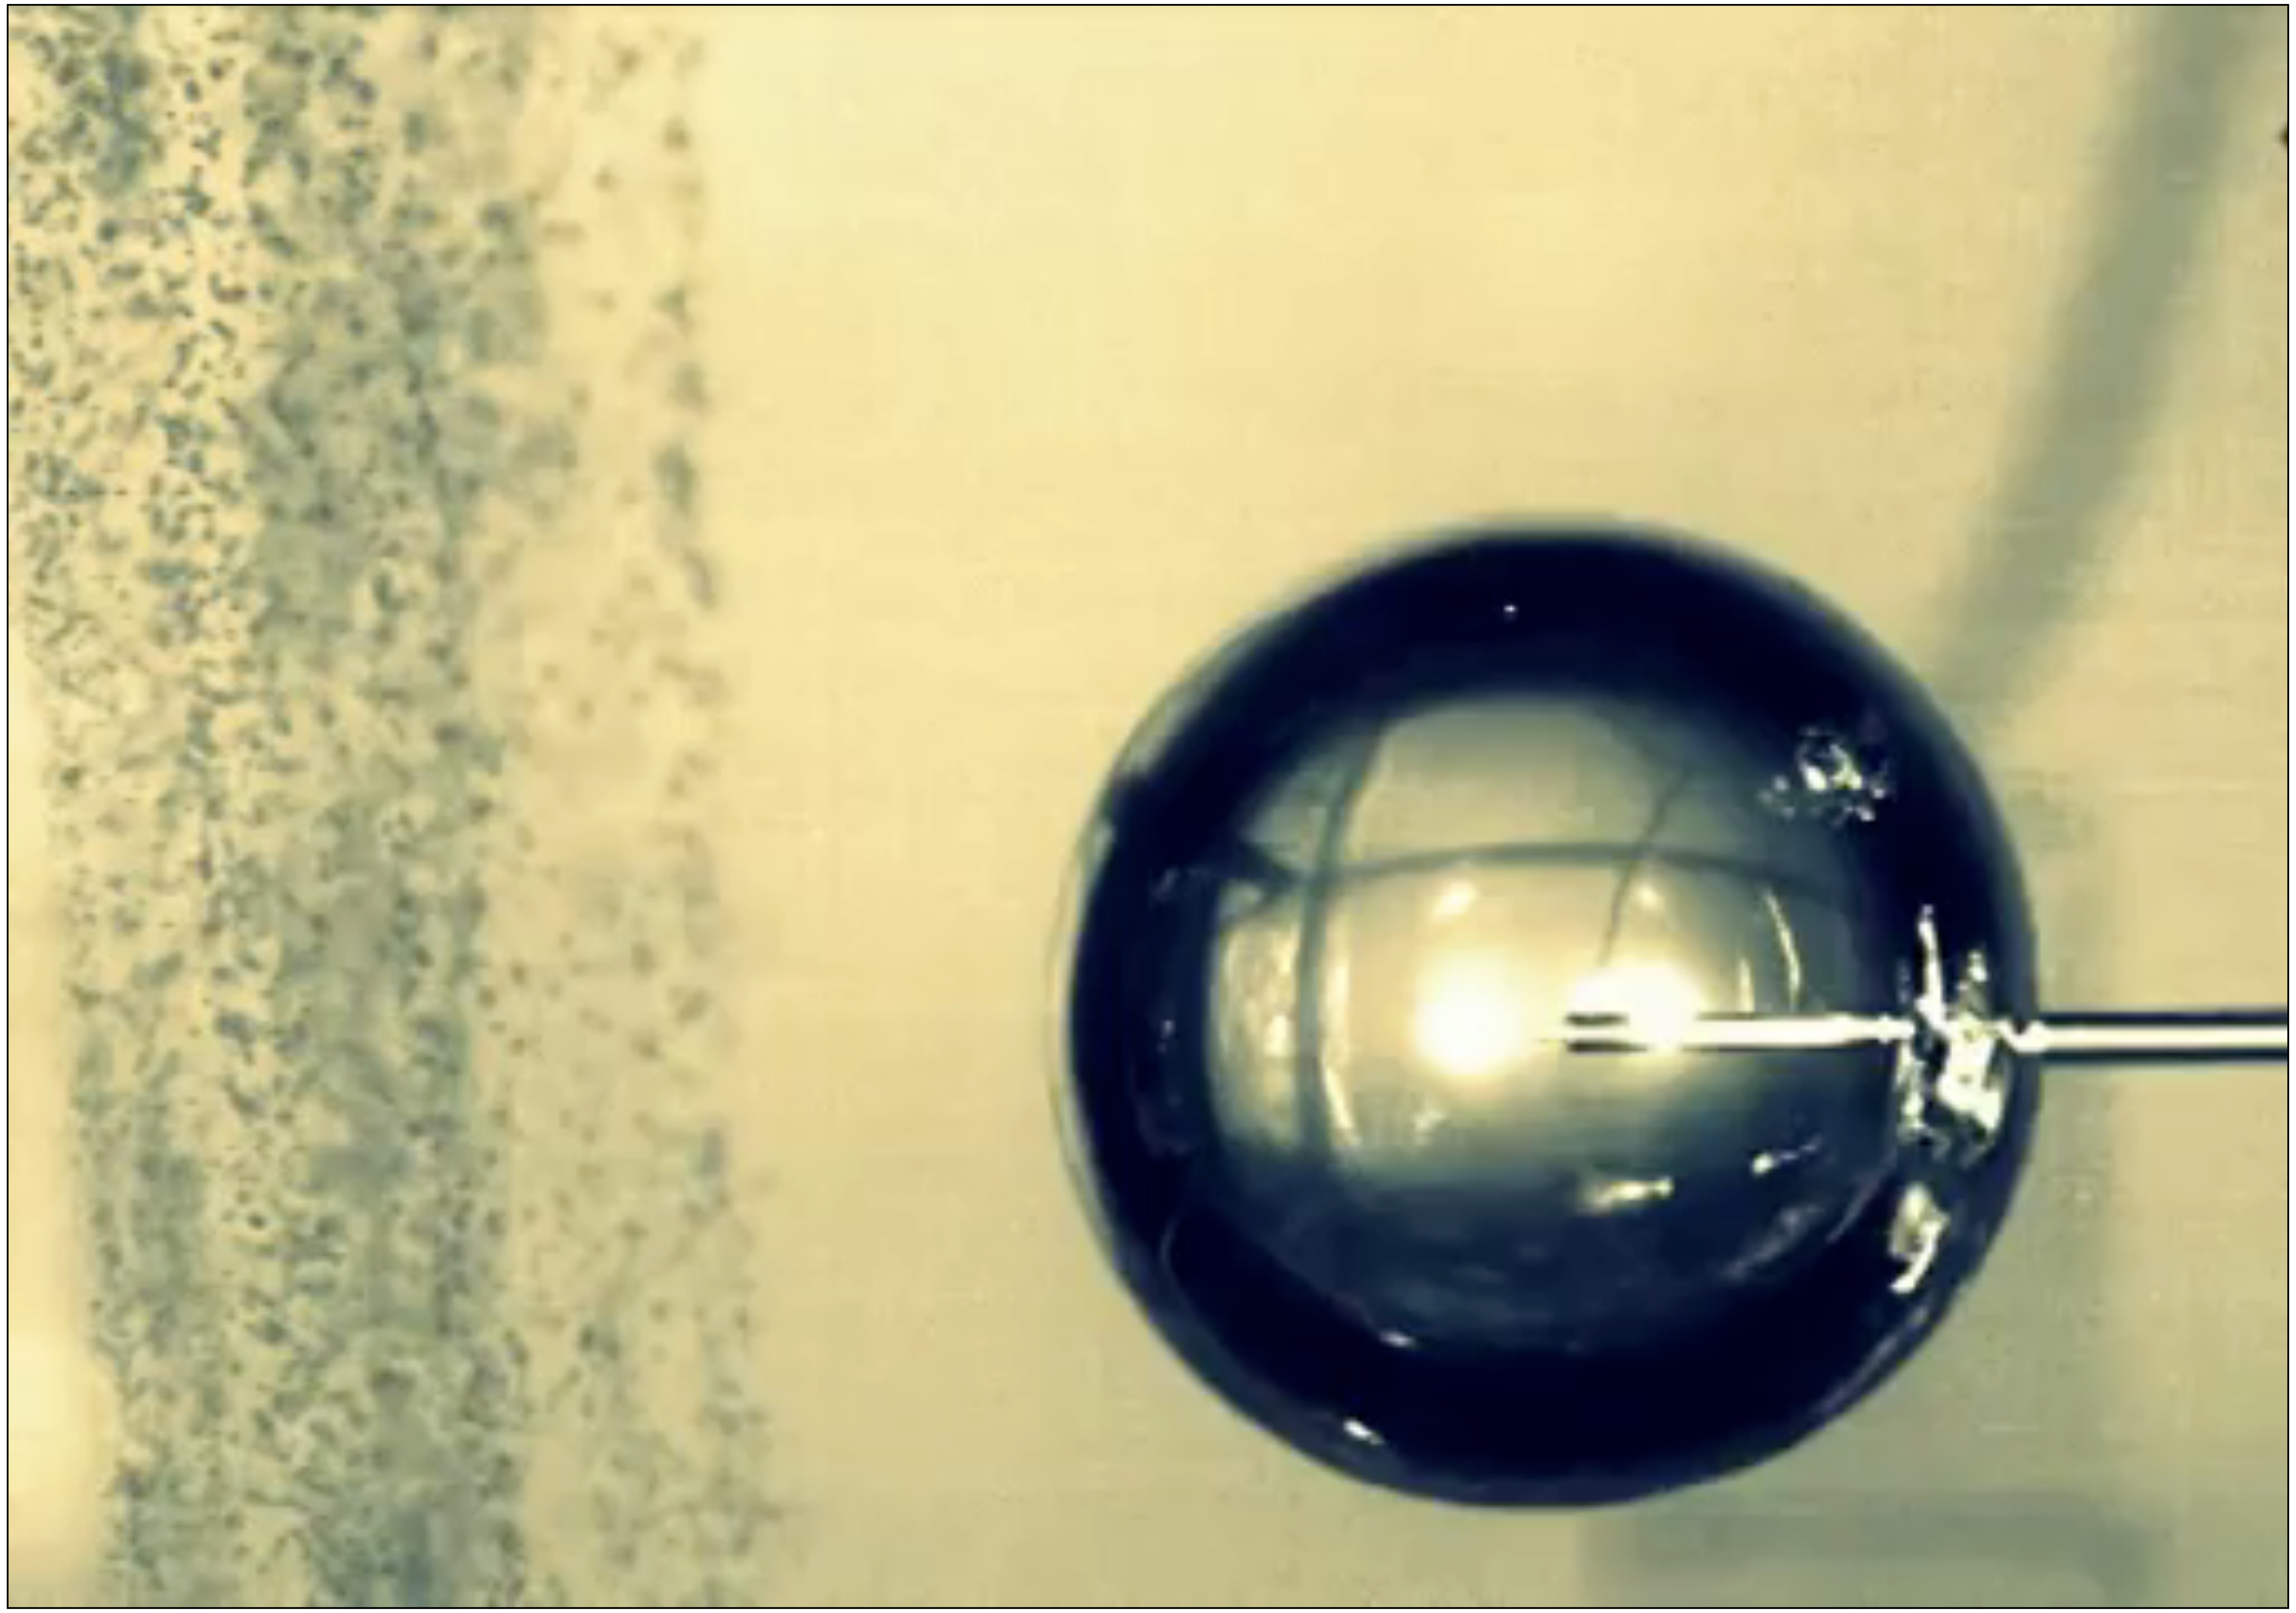

Image courtesy: DynaFlow Inc.

Destructive power of cavitation

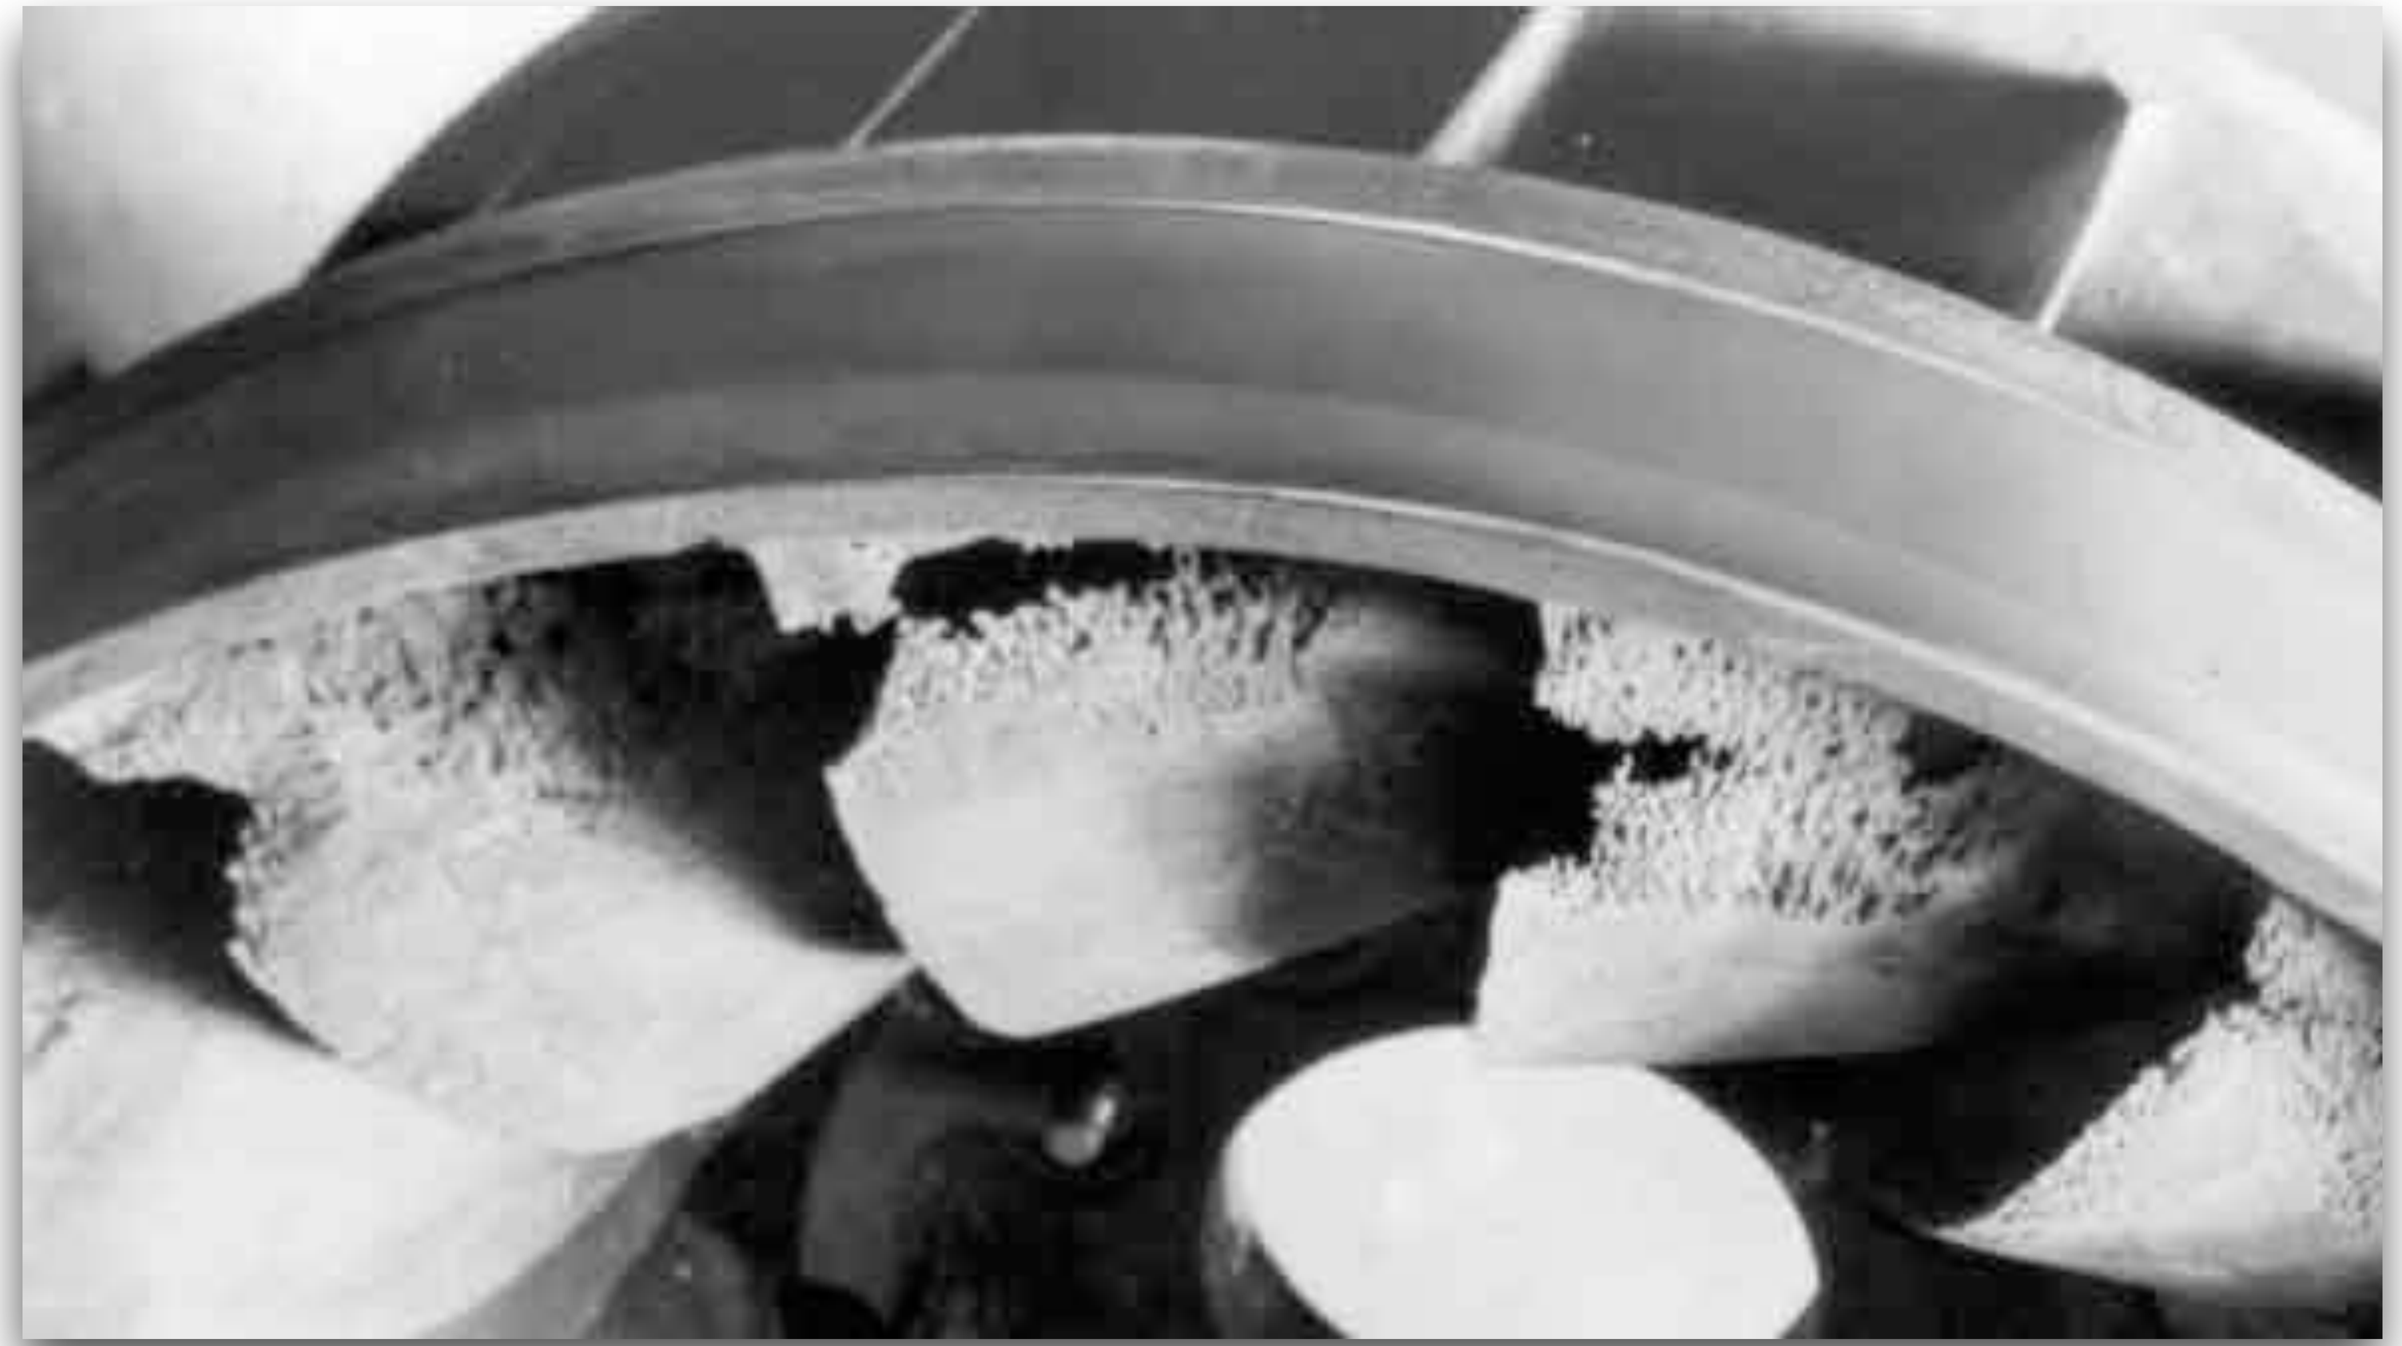

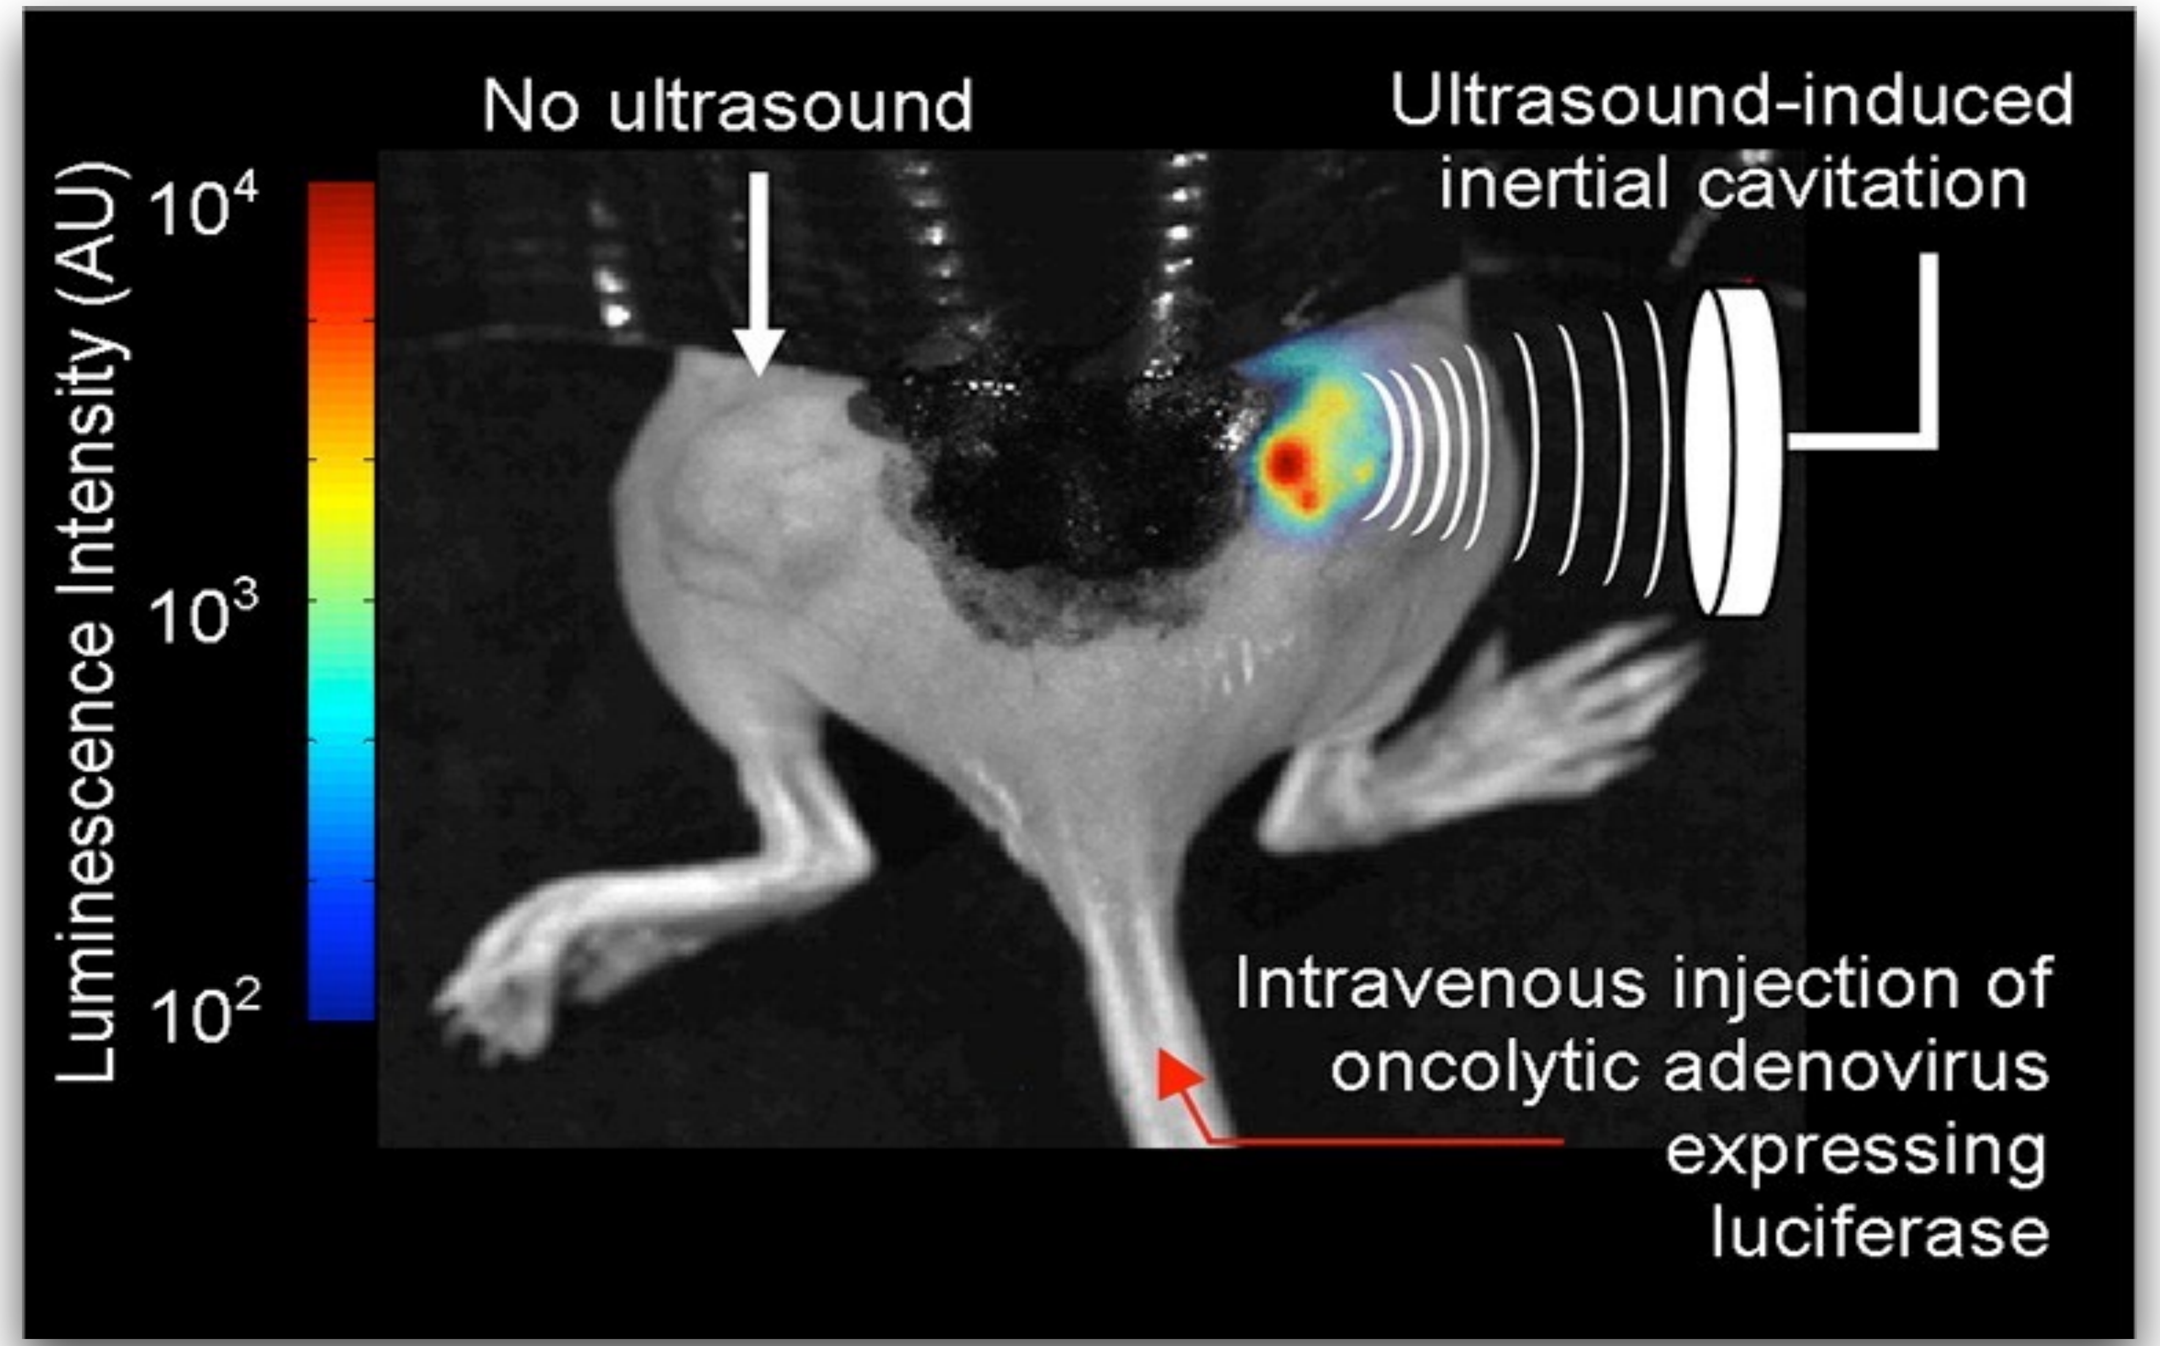

AVOID to maintain performance

- turbines (hydroelectricity, pumps)
- high pressure fuel injectors
- high pressure pipes
- propellers

Image courtesy:
Brennen, "Hydrodynamics of Pumps". Oxford University Press, 1994.

HARNESS for medical treatments

- ultrasonic drug delivery
- kidney shockwave lithotripsy
  - collapse of cavities near stone surface

Image courtesy:
Bazan-Peregrino et al., Cavitation-enhanced delivery of a replicating oncolytic adenovirus to tumors using focused ultrasound.
Journal of Controlled Release Volume 169, Issues 1–2, 2013, pp. 40 - 47.

Governing equations [Kappila] [Masoni] [Allaire]

Multiphase flow equations

$$\begin{cases} (\alpha_1\rho_1)_t + \nabla\cdot(\alpha_1\rho_1\mathbf{u}) = 0, \\ (\alpha_2\rho_2)_t + \nabla\cdot(\alpha_2\rho_2\mathbf{u}) = 0, \\ (\rho\mathbf{u})_t + \nabla\cdot(\rho\mathbf{u}\otimes\mathbf{u} + p\mathbf{I}) = 0, \\ E_t + \nabla\cdot((E+p)\mathbf{u}) = 0. \end{cases}$$

Advection of phase volume fractions

$$(\alpha_2)_t + \mathbf{u}\cdot\nabla\alpha_2 = K(\alpha_{1,2}, \rho_{1,2}, c_{1,2})\nabla\cdot\mathbf{u}.$$

Equation of state (water phase: stiffened)

$$E = \frac{1}{2}\rho\mathbf{u}^2 + \Gamma p + \Pi, \quad \Gamma = \frac{1}{\gamma - 1}, \quad \Pi = \frac{\gamma p_c}{\gamma - 1}.$$

$$p = \frac{(E - \rho\mathbf{u}^2) - (\alpha_1\Pi_1 + \alpha_2\Pi_2)}{\alpha_1\Gamma_1 + \alpha_2\Gamma_2}, \quad \frac{1}{\rho c^2} = \frac{\alpha_1}{\rho_1 c_1^2} + \frac{\alpha_2}{\rho_2 c_2^2}.$$

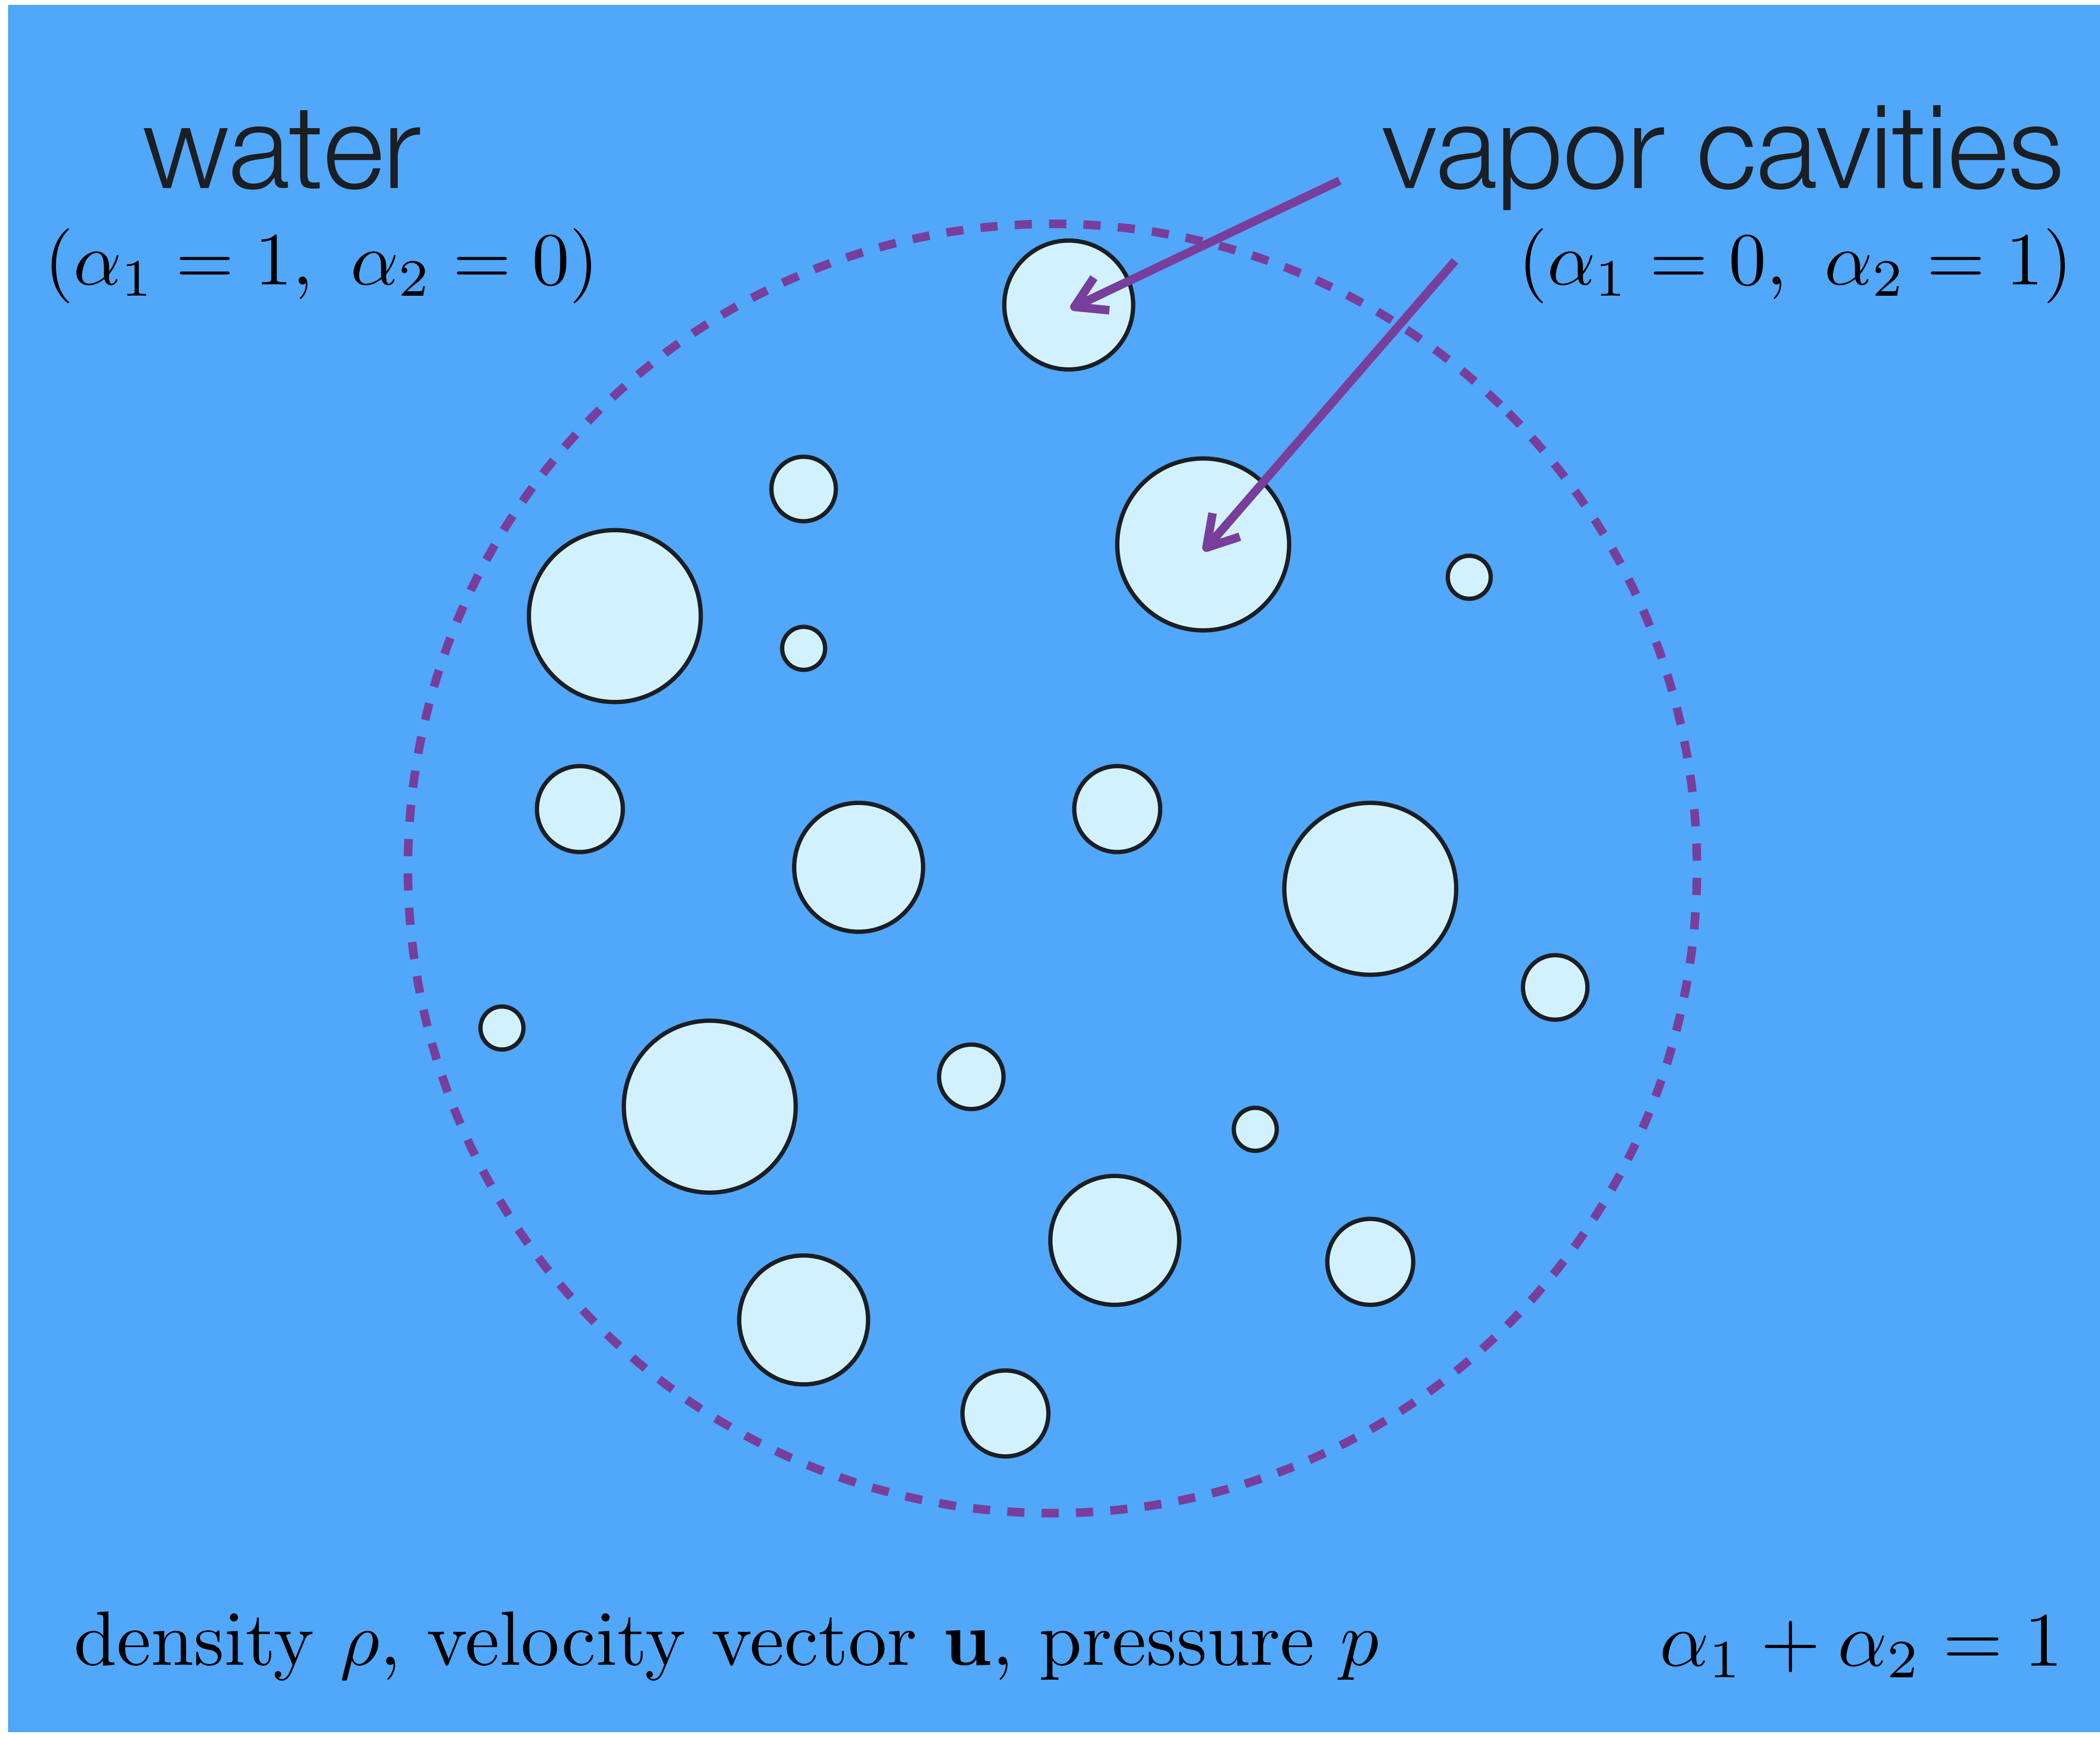

2D slice of a 3D domain

cavity sizes of 50-200 µm (log-Gaussian)

Finite Volume Solver

$$\partial_t \mathbf{U}(\mathbf{x}, t) + \mathrm{div}\, \mathbf{F}(\mathbf{U}, \mathbf{x}) = 0$$

- Cell averages

$$\mathbf{U}_j(t) \approx \frac{1}{|C_j|} \int_{C_j} \mathbf{U}(x,t)dx$$

- Semi-discrete formulation (ODE)

$$\frac{\mathrm{d}}{\mathrm{dt}} \mathbf{U}_j(t) + \frac{1}{\Delta x}\left(\mathbf{F}_{j+\frac{1}{2}} - \mathbf{F}_{j-\frac{1}{2}}\right) = 0$$

- High order reconstruction

WENO3 / WENO5
[Harten, Shu, Osher]

- Approximate Riemann solver HLLC

$$\mathbf{F}_{j+\frac{1}{2}} \approx \mathbf{F}^{\mathrm{HLLC}}_{j+\frac{1}{2}}(\mathbf{U}^+, \mathbf{U}^-)$$

- RK3 time stepping [Gottlieb, Shu, Tadmor]

$$\mathbf{U}_j^n \rightarrow \mathbf{U}_j^{n+1}$$

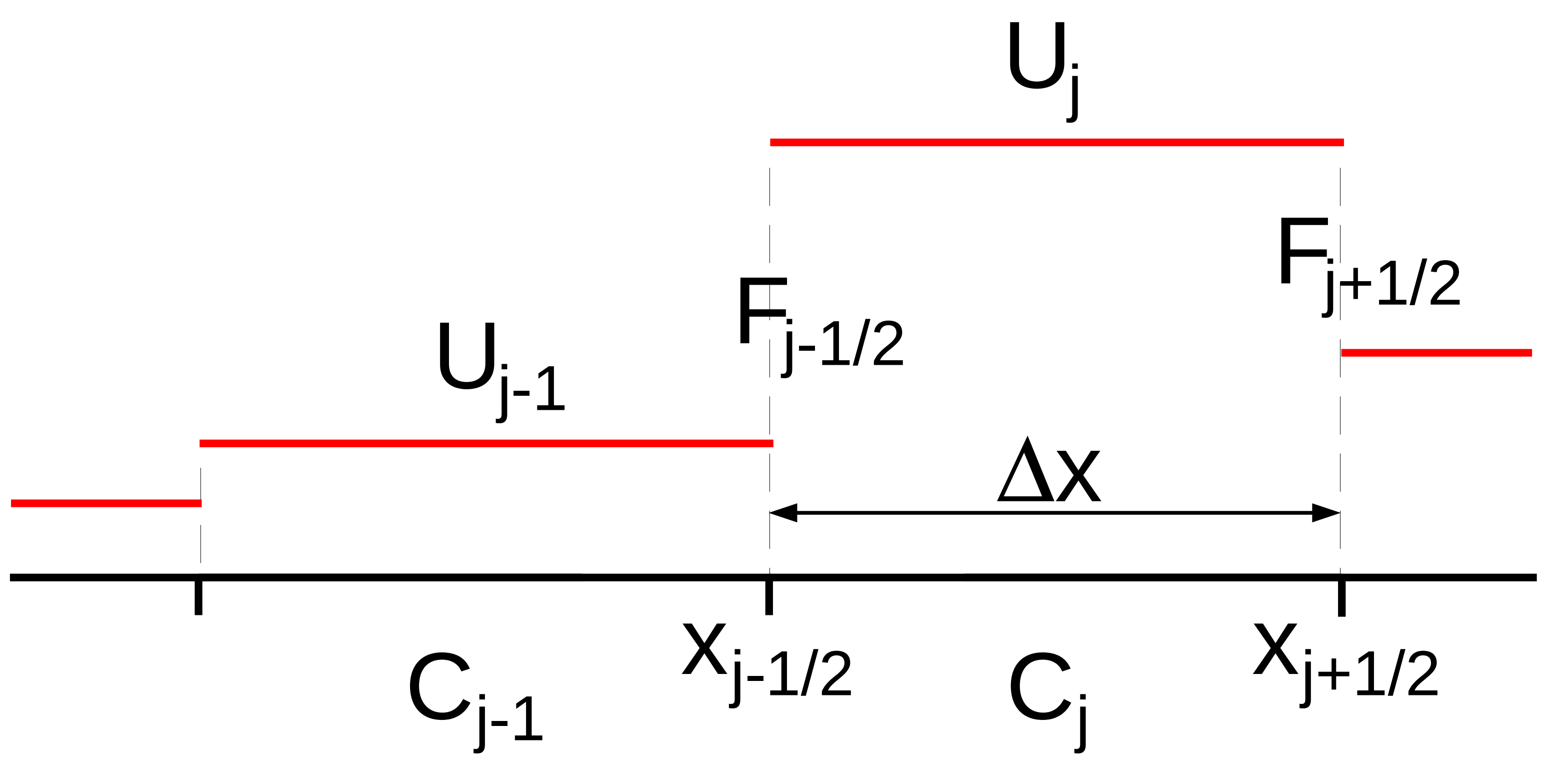

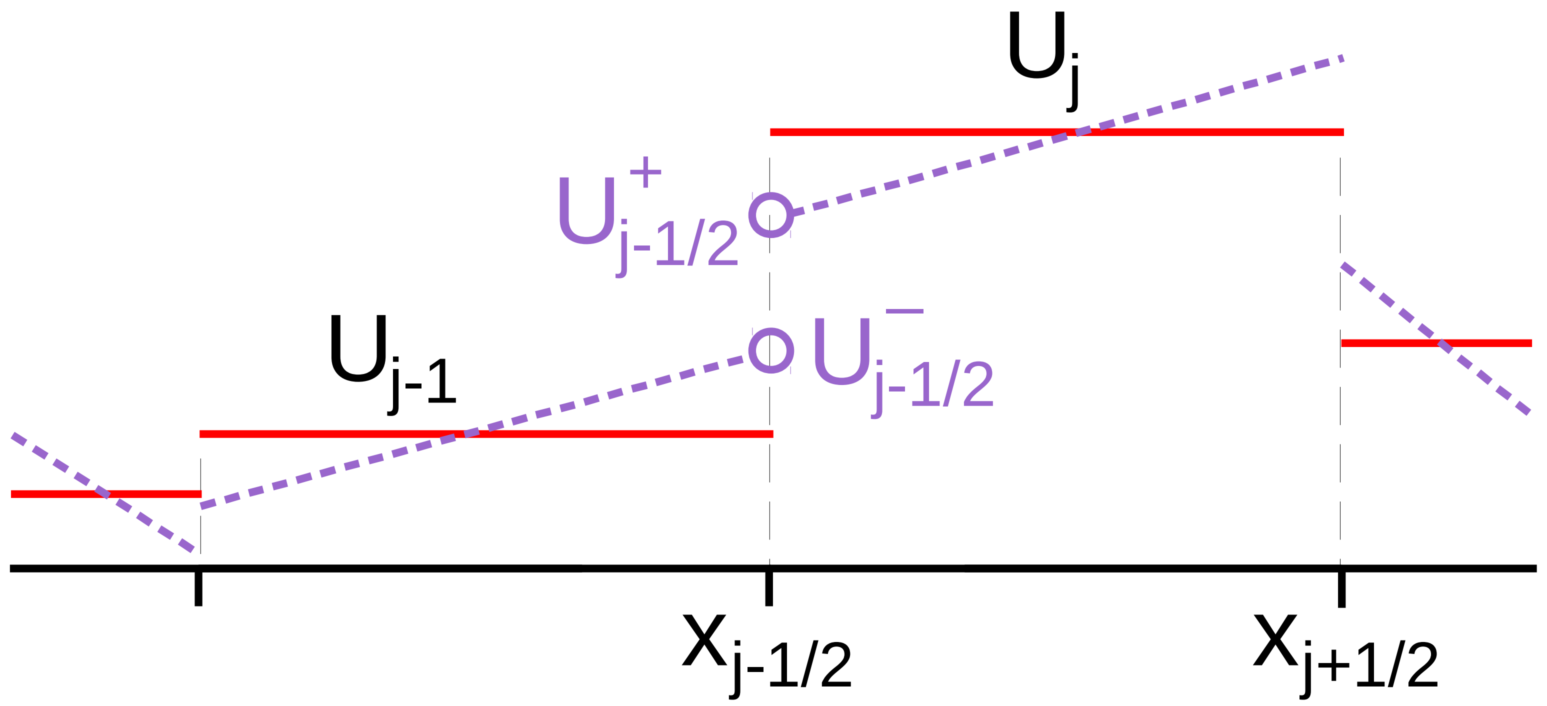

CUBISM-MPCF

Peta-scale Multi-Phase Compressible Flow approximate Riemann solver

[Rossinelli, Hejazialhosseini, Hadjidoukas, Conti, Bergdorf, Wermelinger, Rasthofer, Šukys]

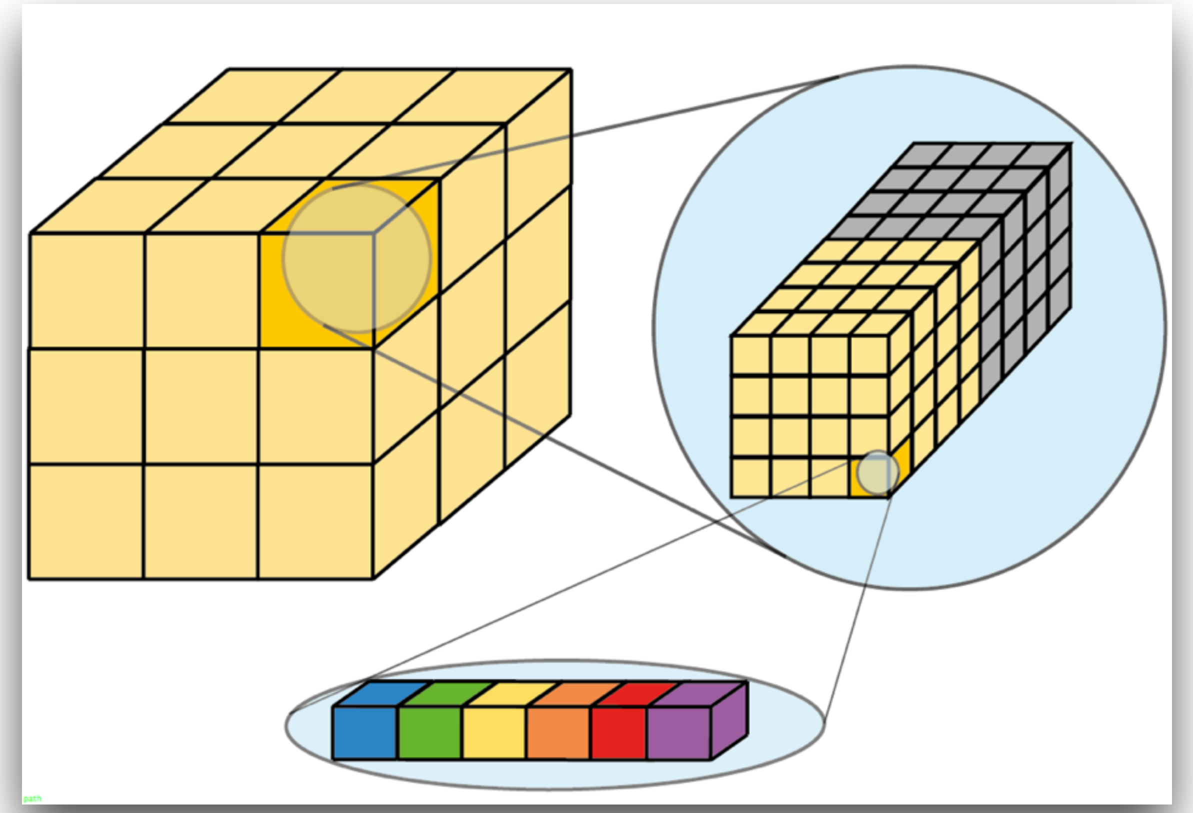

Block-based memory layout

(spatial locality)

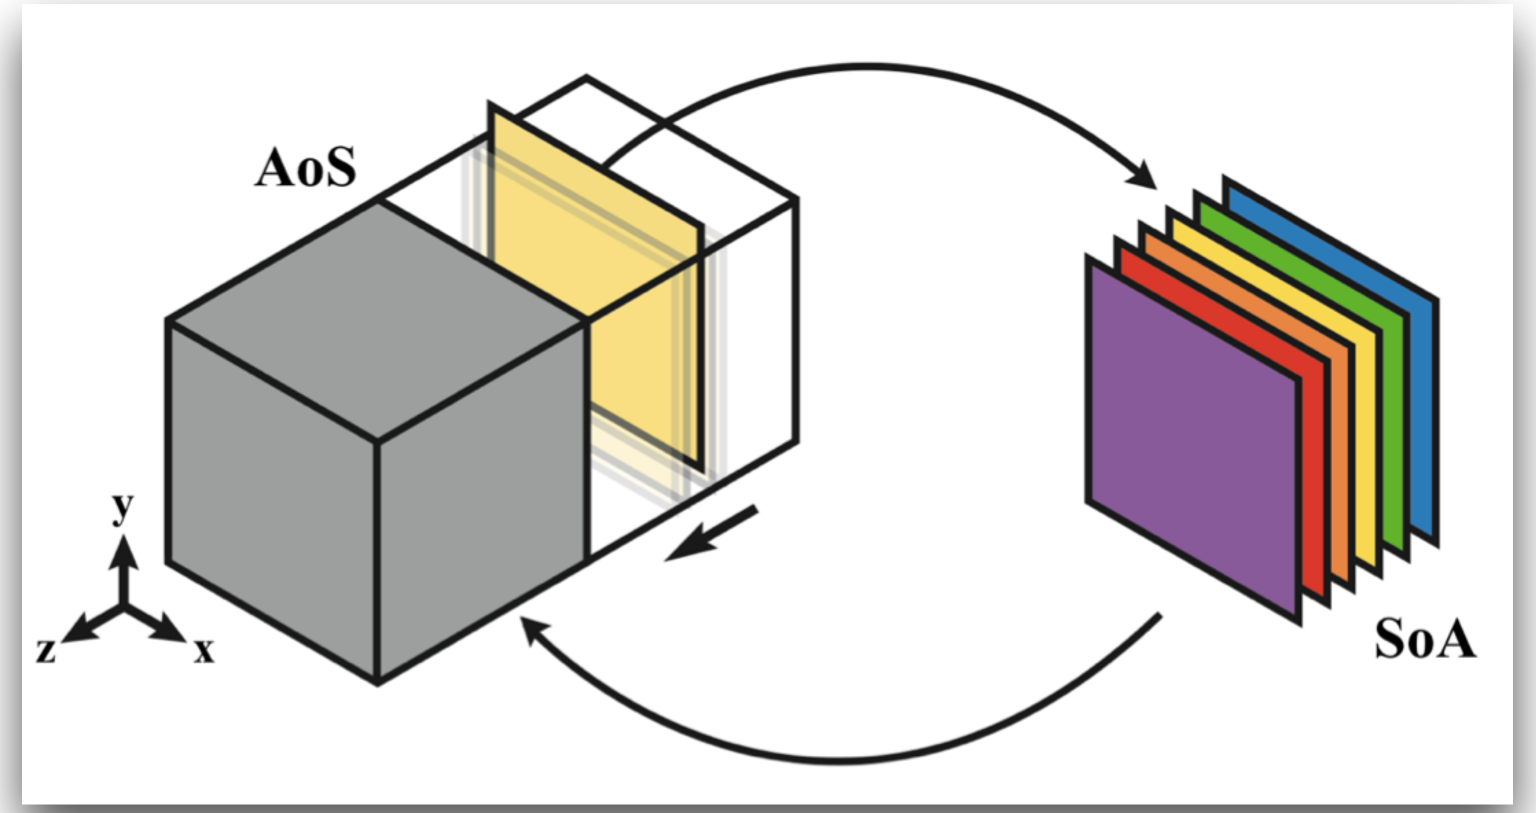

Instruction/data-level parallelism

(**S**tructure **o**f **A**rrays for SSE/QPX vectorization)

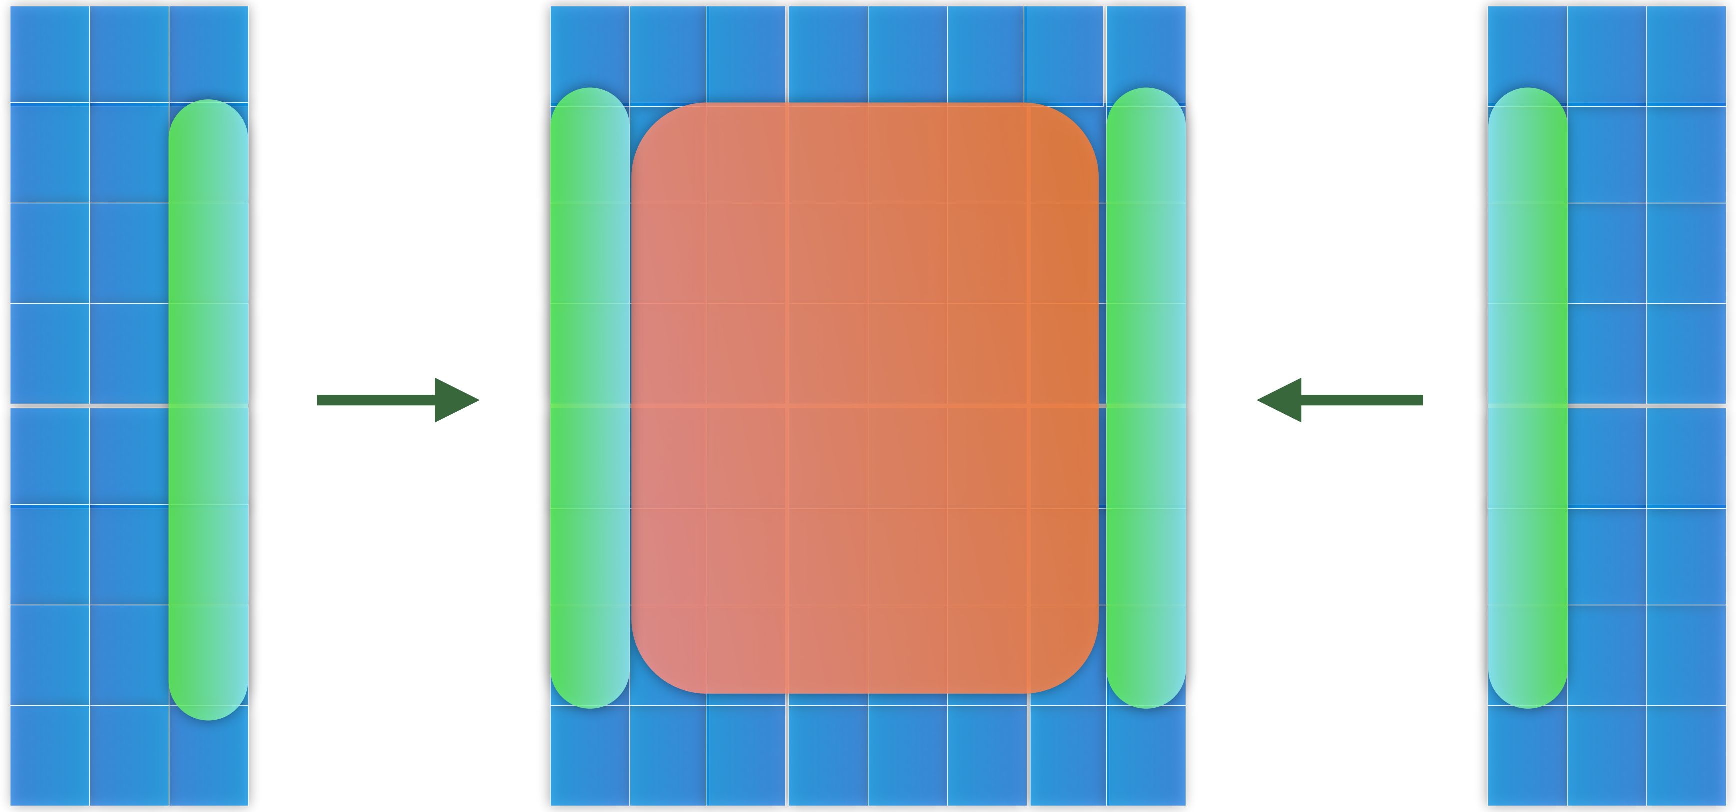

Domain decomposition MPI/OpenMP

(dynamic loop scheduling) (non-blocking P2P communication)

(asynchronous progress for C/T overlap)

- ACM Gordon Bell Prize: **14.4 Pflops (72% peak)** on Sequoia (IBM BlueGene/Q, 1.6M cores)
- Wavelet-based I/O **compression** | ~100x reduction | 1% overhead
- **Fault-tolerance with restart** mechanism | lossless compression ~10x reduction

Petascale simulations of cloud cavitation collapse

50 thousand cavities at 100 bar
0.5 billion mesh elements
25 thousand time steps
25x pressure amplification

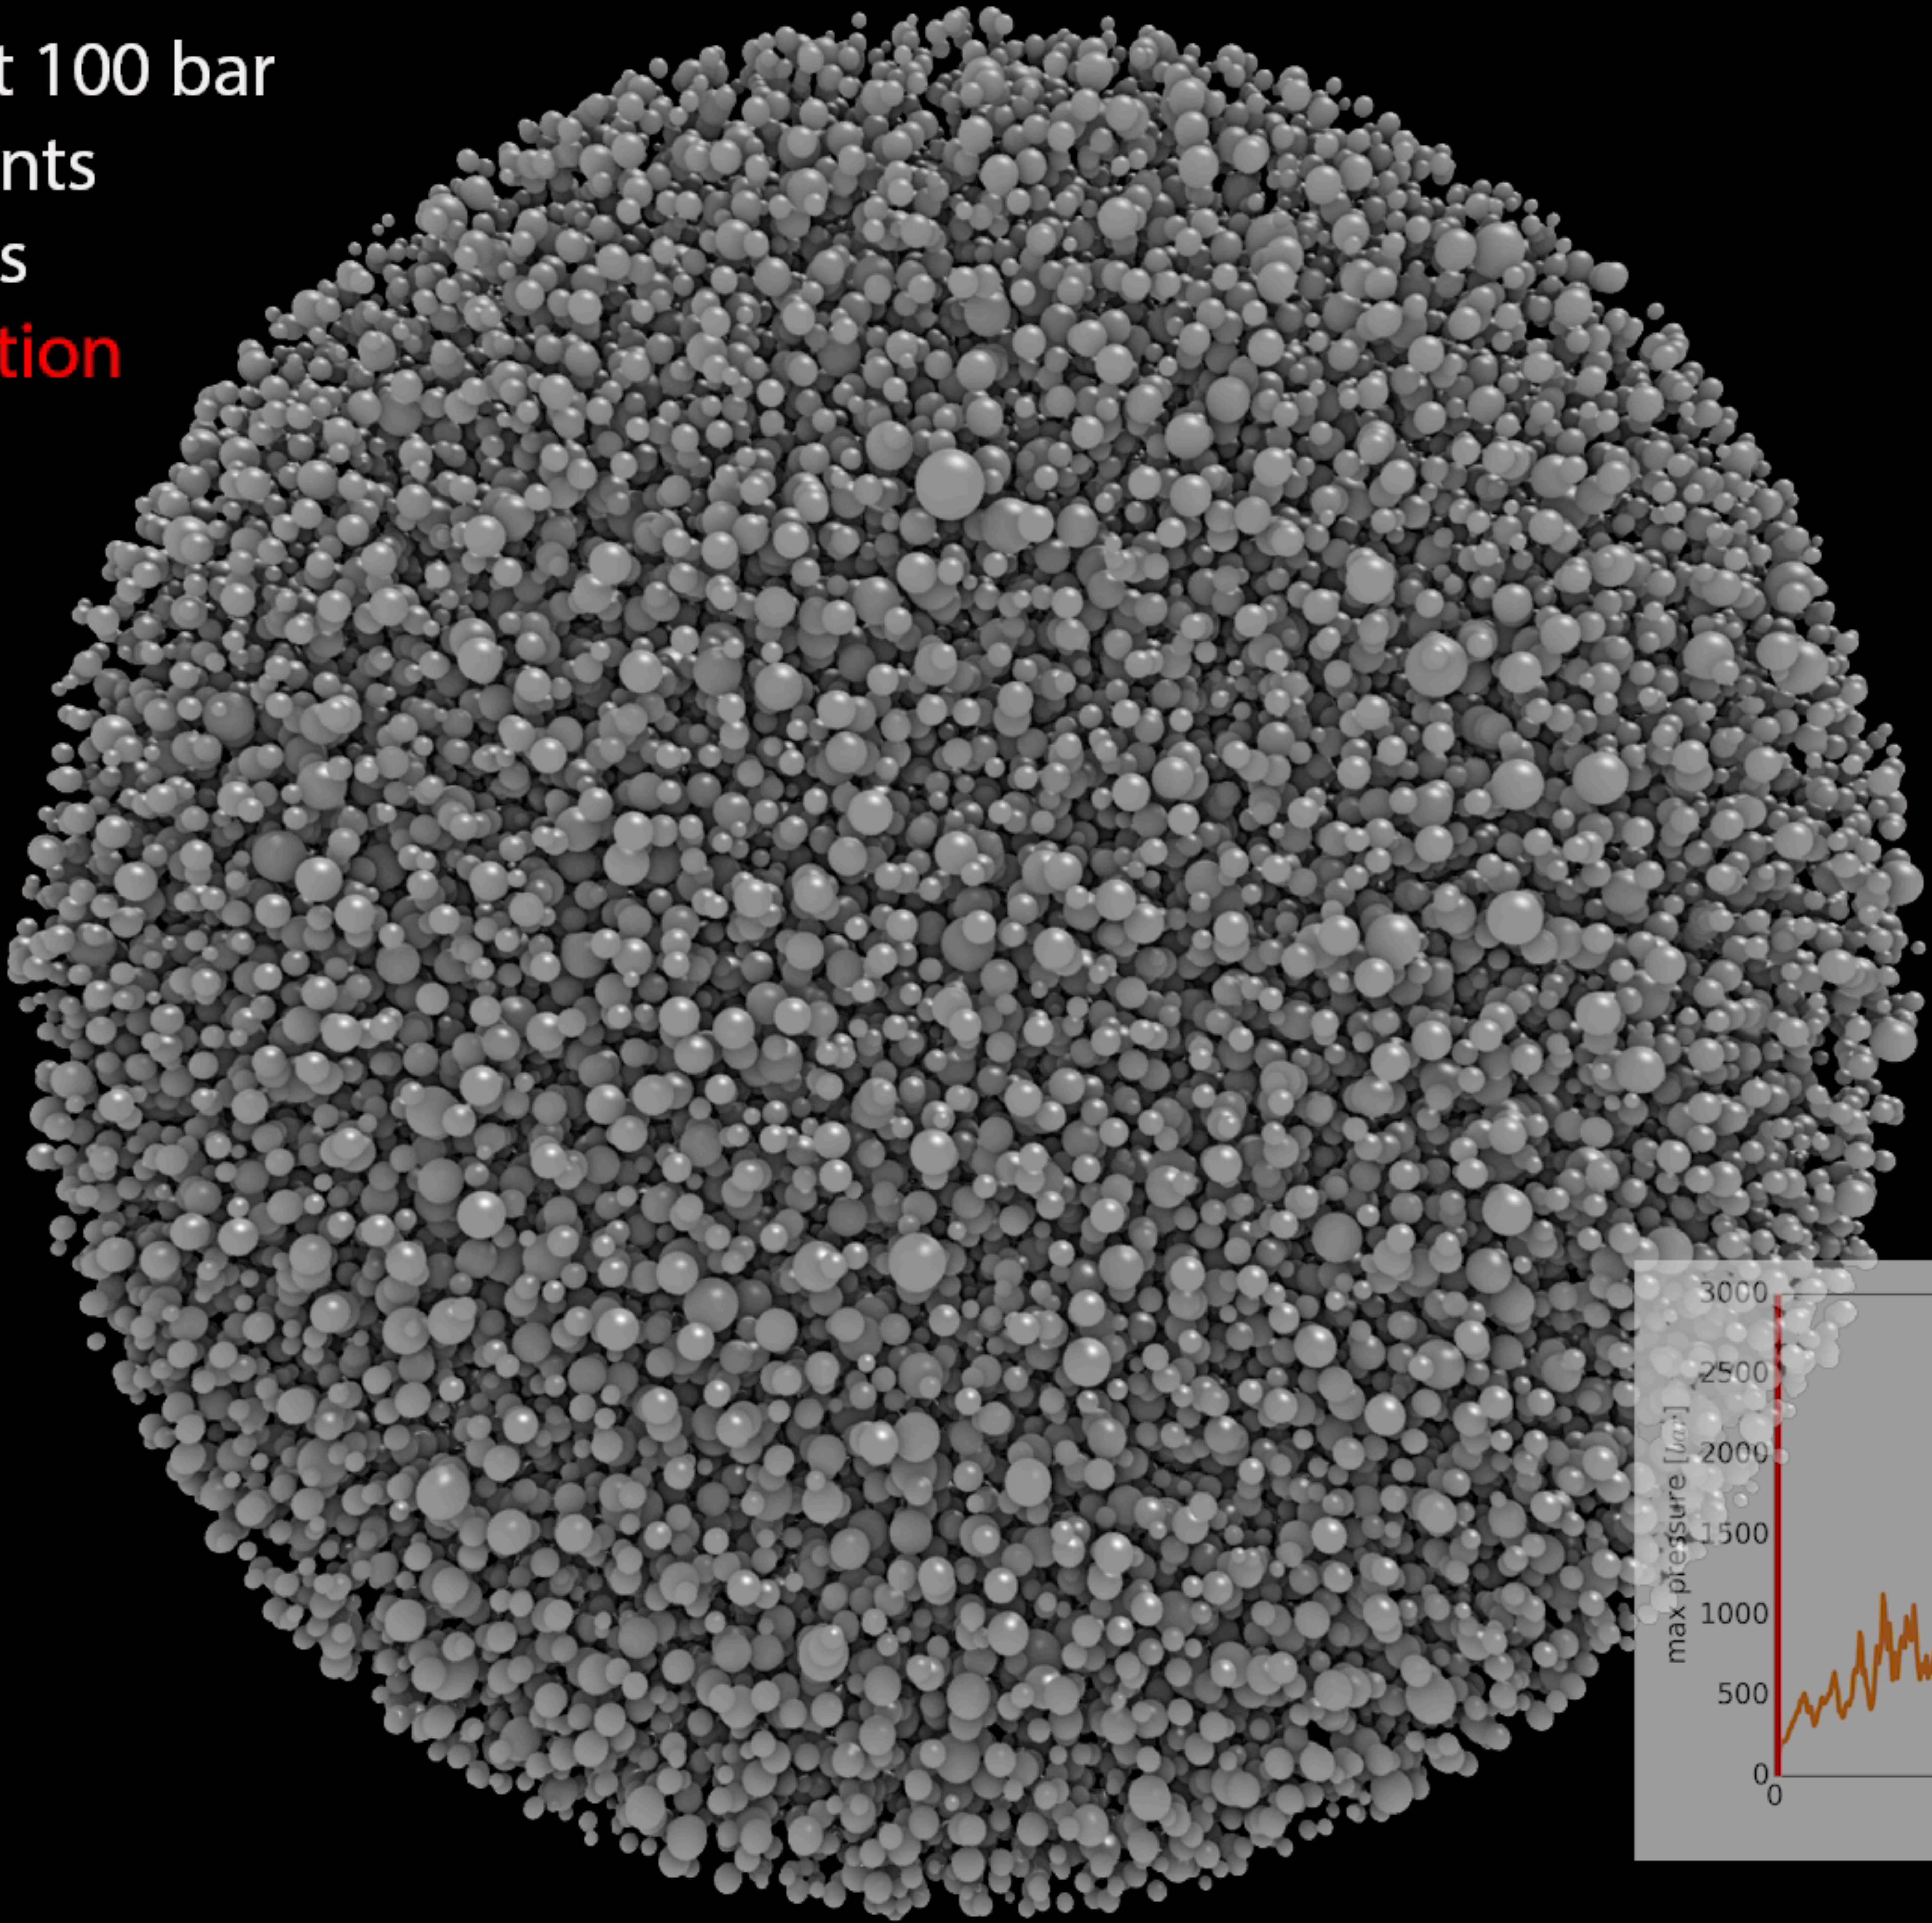

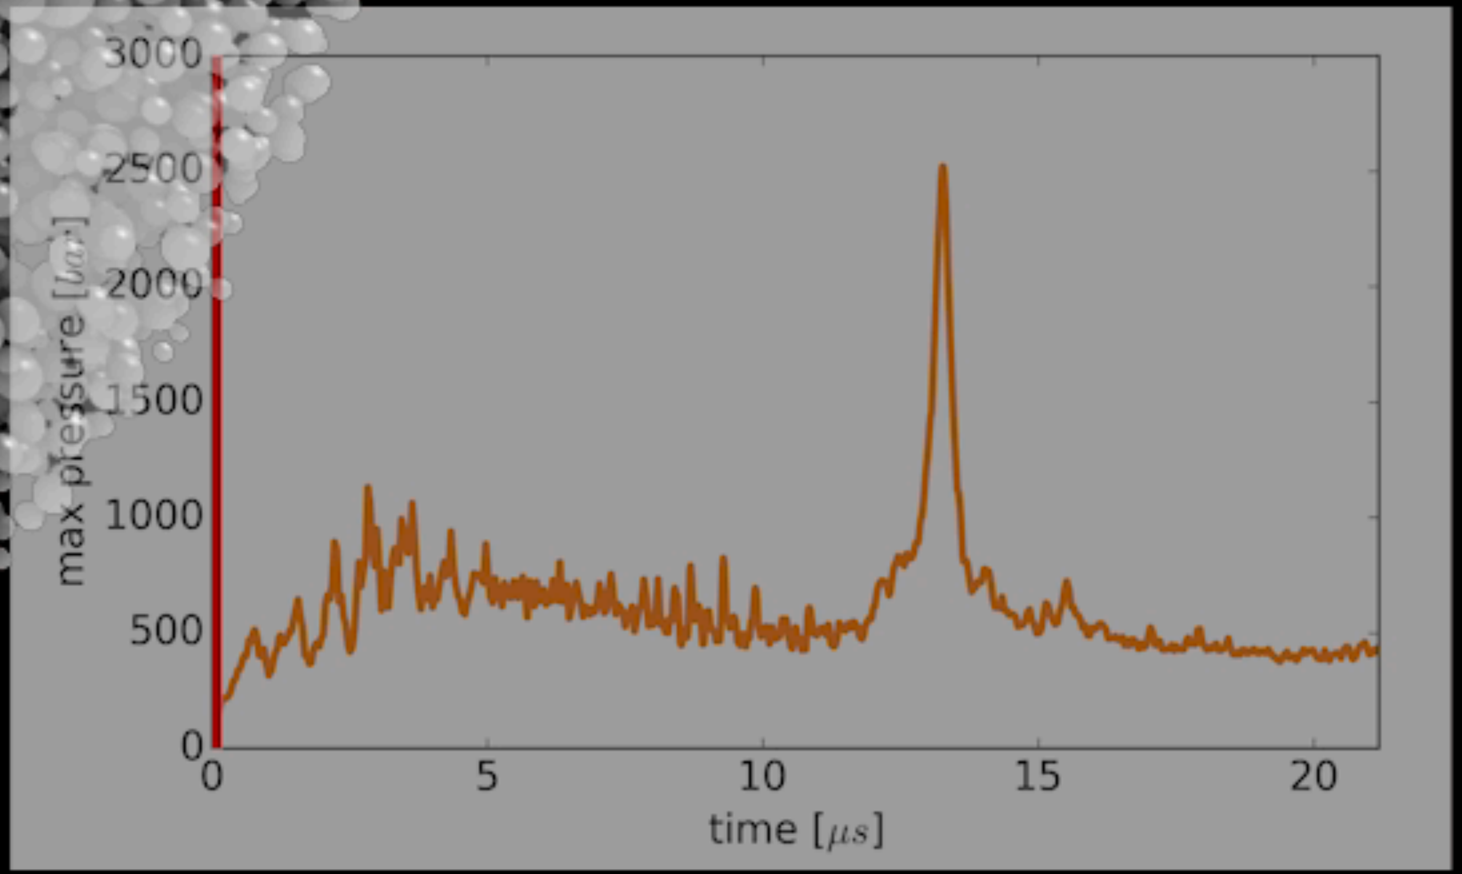

1 / 500 000 X

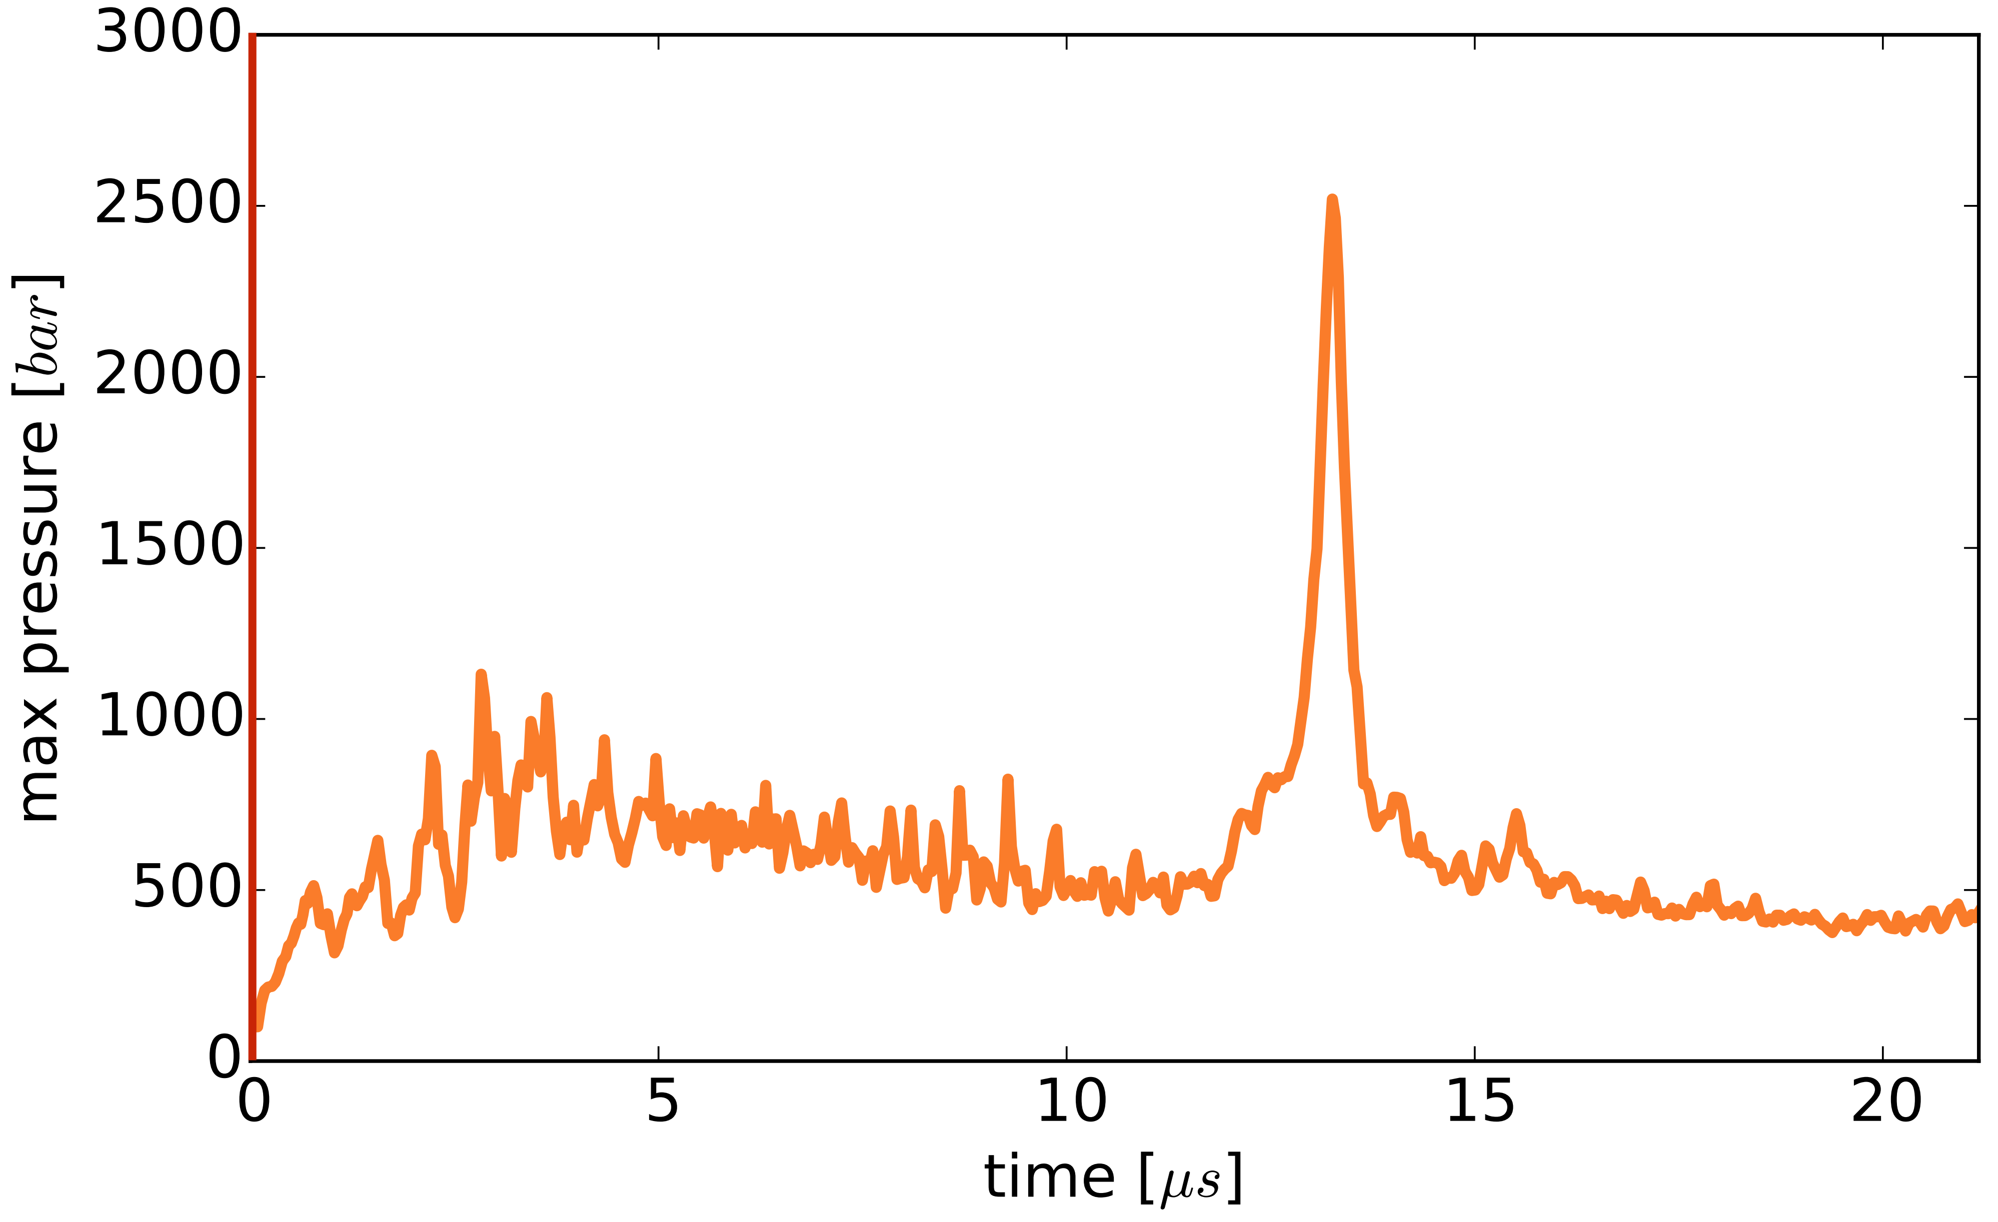

max pressure [bar]
3000
2500
2000
1500
1000
500
0
0
5
10
15
20
time [μs]

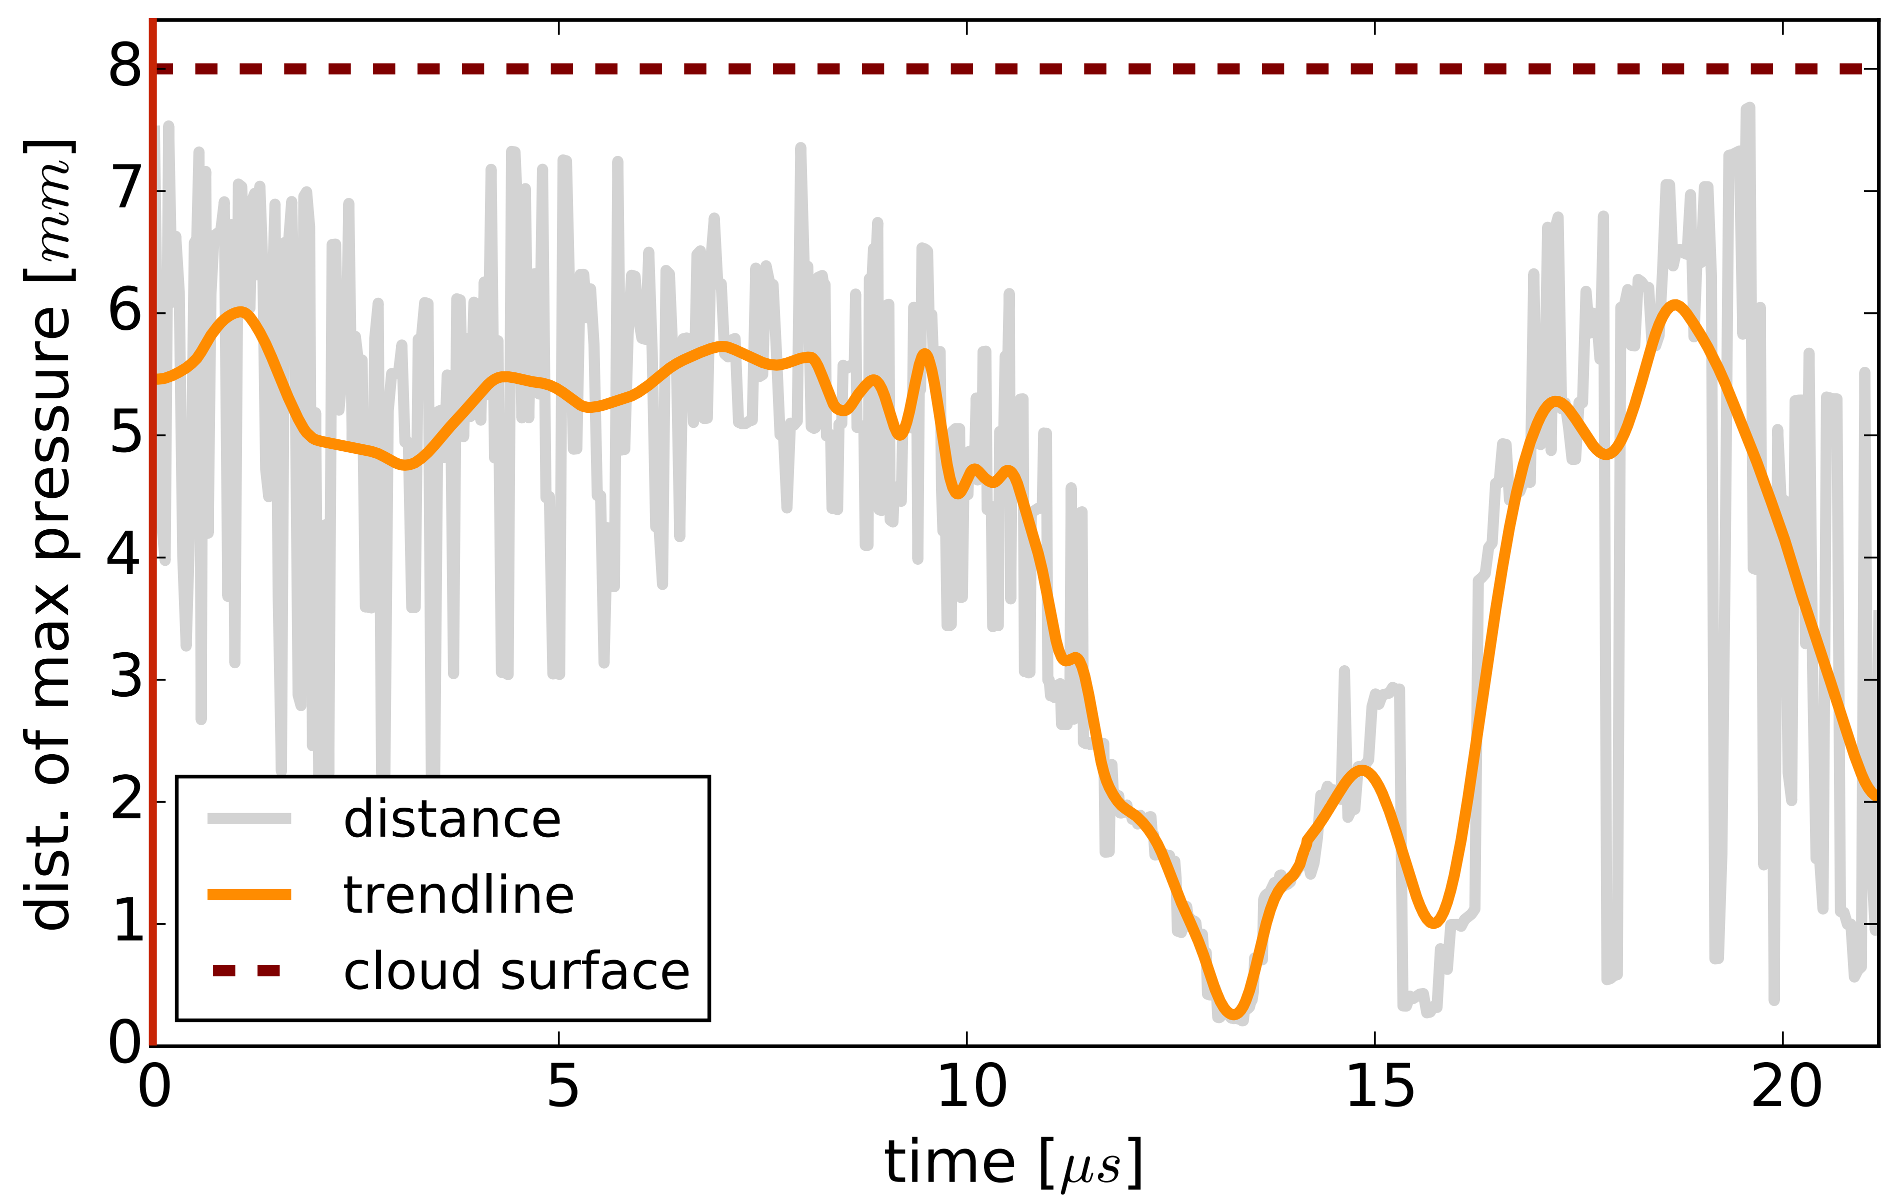

dist. of max pressure [mm]
8
7
6
5
4
3
2
1
0
0
5
10
15
20
time [μs]
distance
trendline
cloud surface

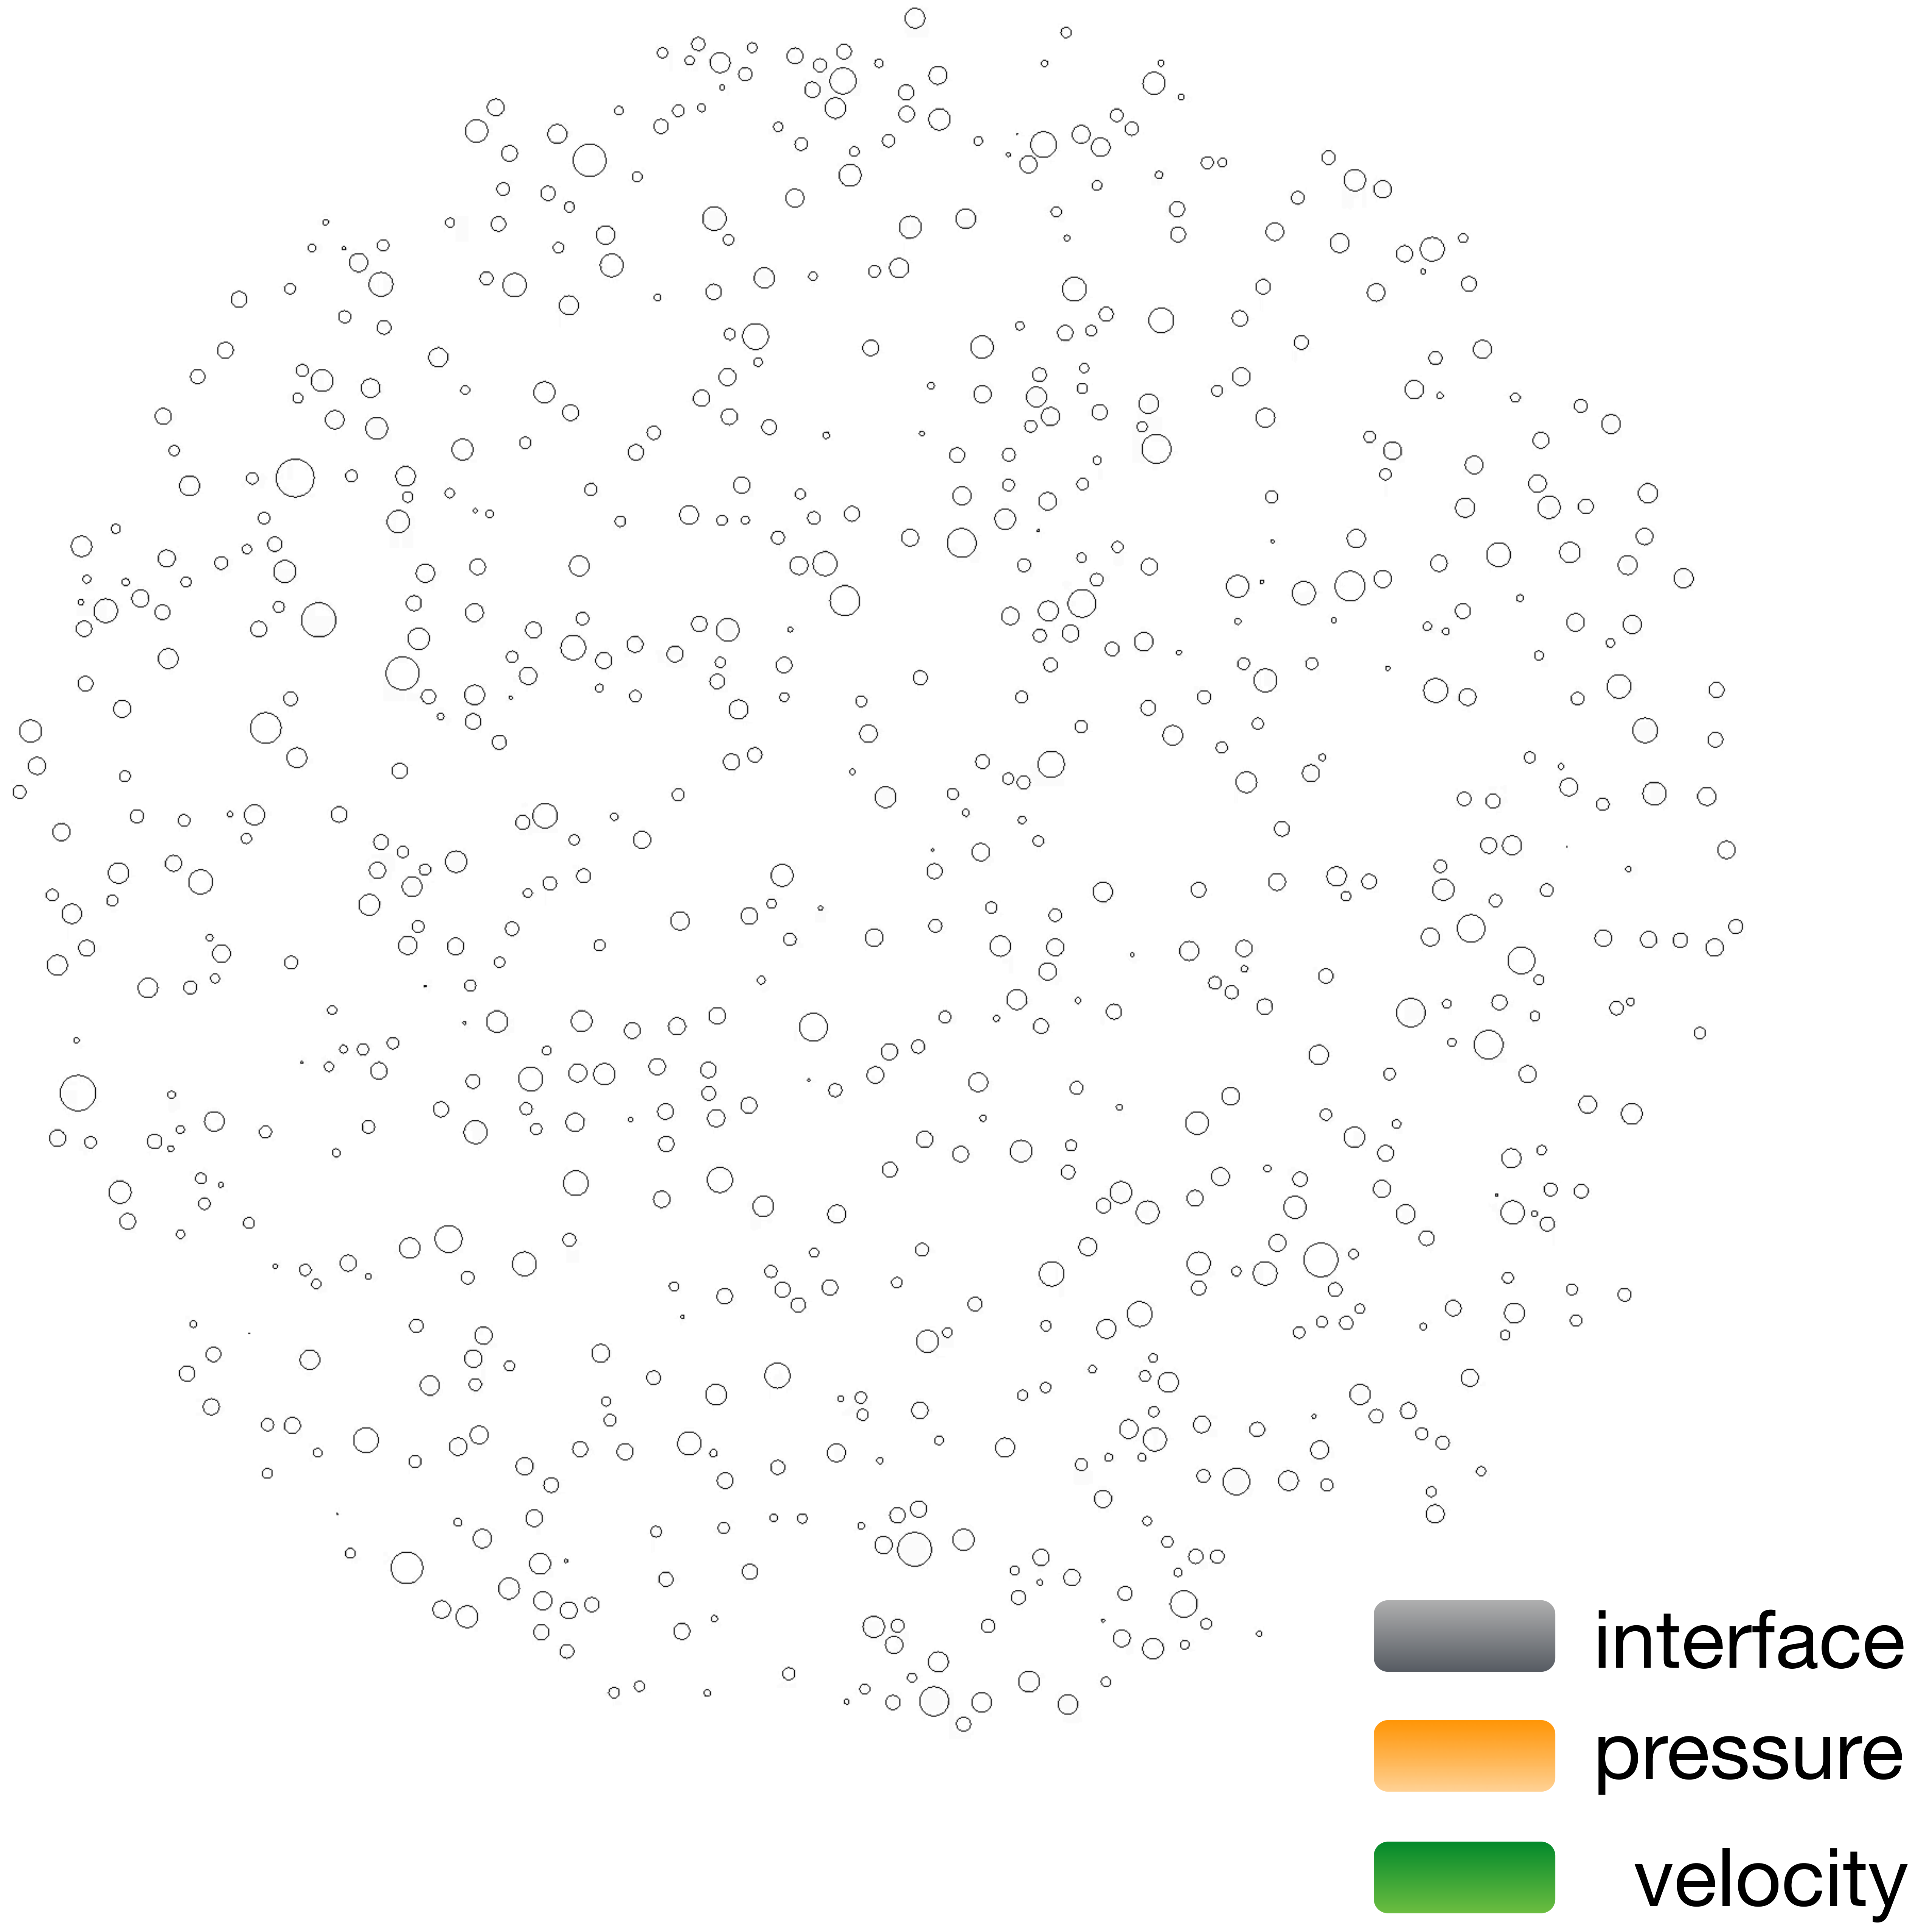

interface
pressure
velocity

Uncertainty quantification in cloud cavitation collapse

Collapse of two random clouds

2 clouds: different statistical realizations **(RNG seeds)** of the initial configuration

Spherical clouds of 100 equally sized (75µm) cavities

Uniformly distributed (random) cavity **positions**

Collapse of two random clouds

2 clouds: different statistical realizations **(RNG seeds)** of the initial configuration

maximum pressure

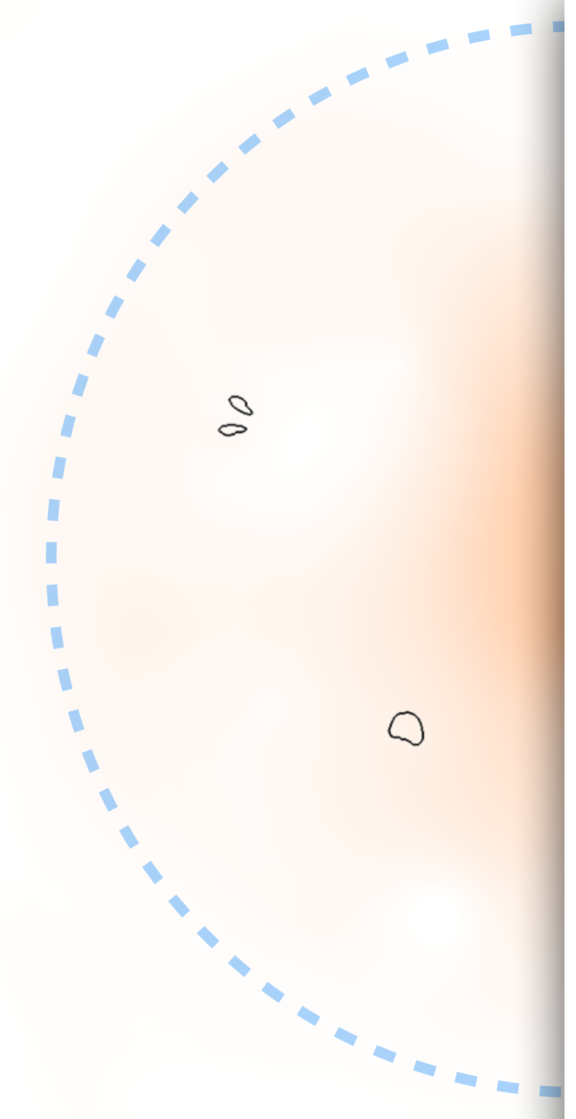

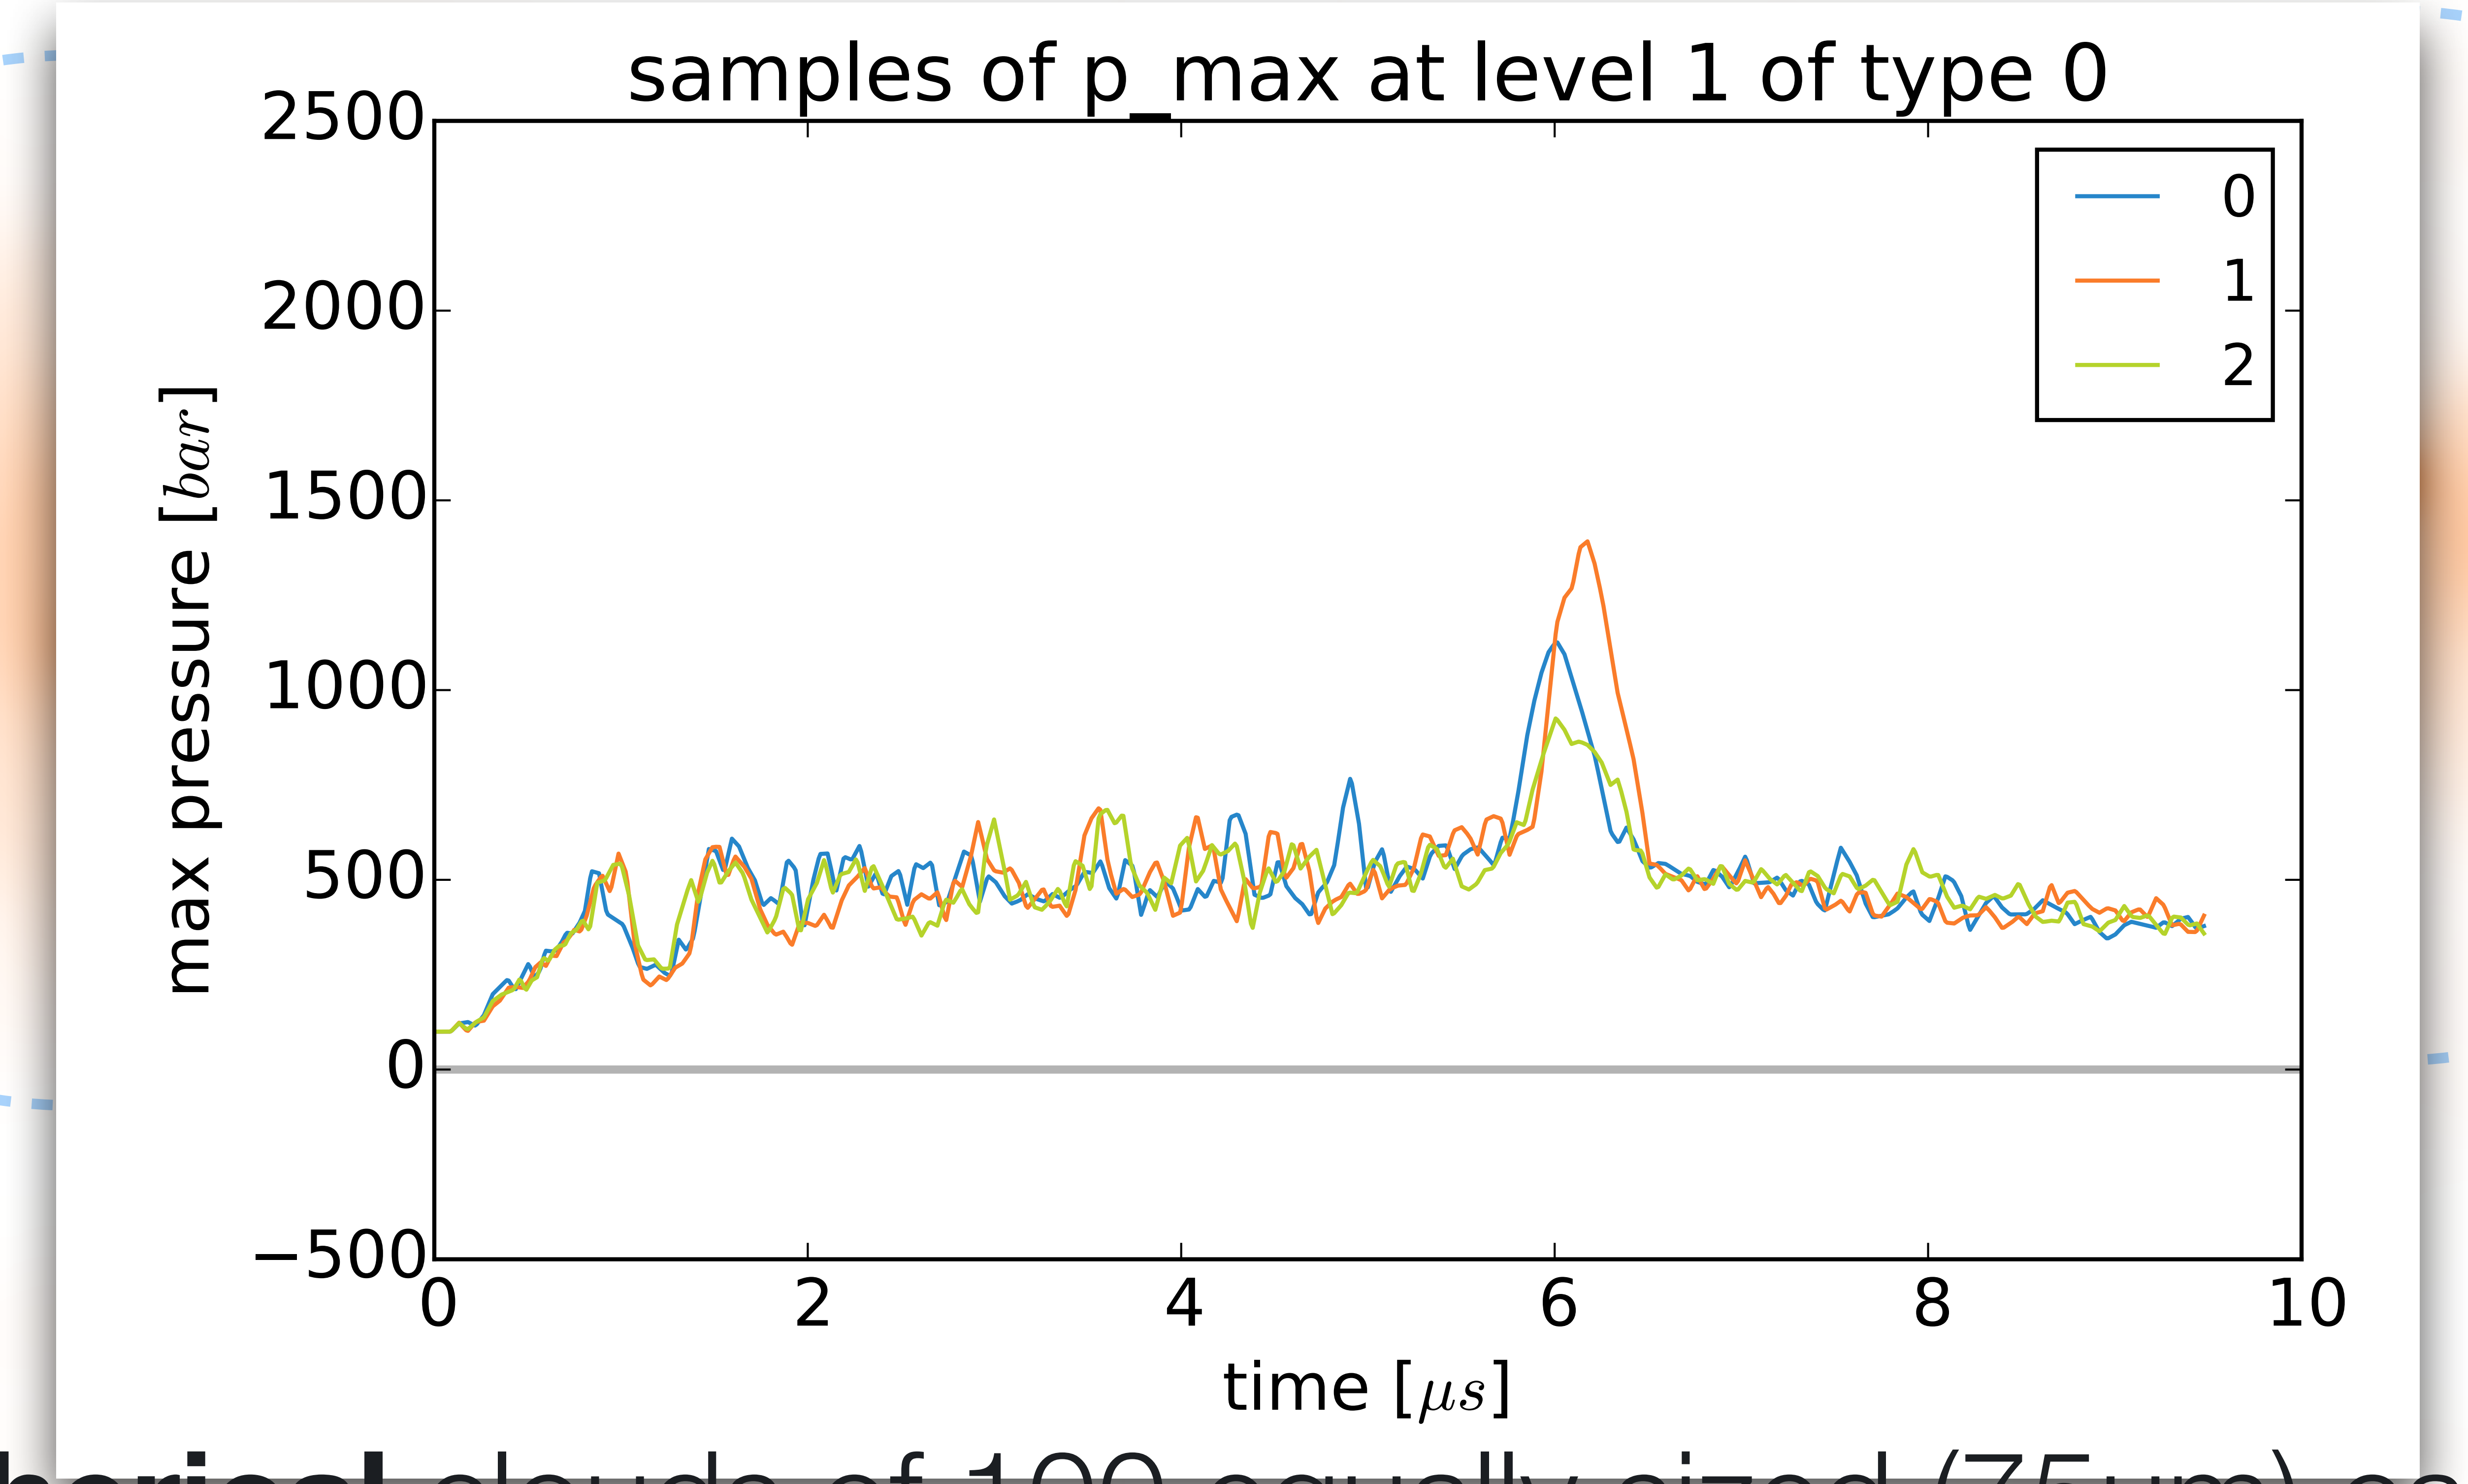

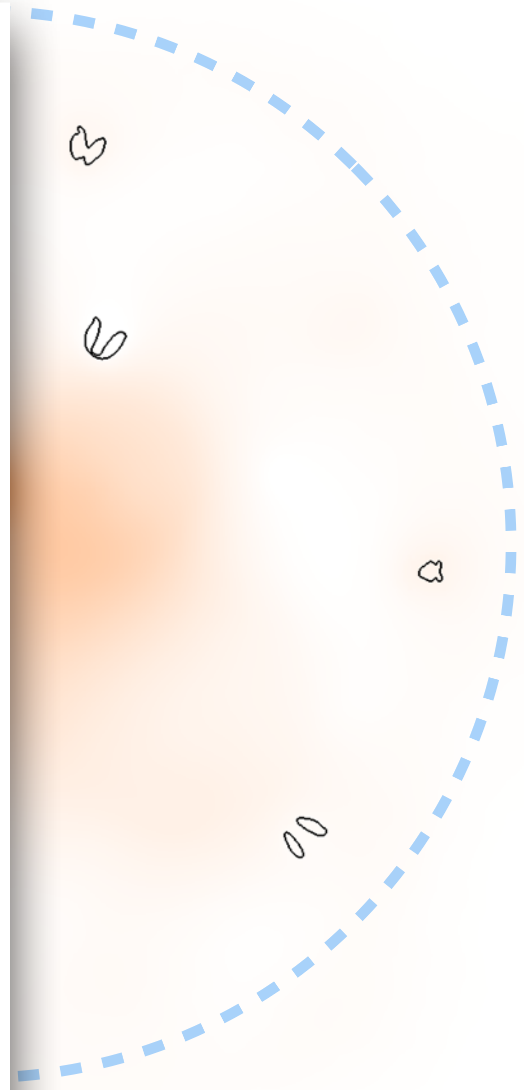

Spherical clouds of 100 equally sized (75μm) cavities

Uniformly distributed (random) cavity **positions**

significant variations

800 - 1500 bar

Previous work and goals

11.6.2010 11:02:04 -4493.9866[ms] 000000000 MotionBLITZ EoSens mini1 Mikrotron GmbH 320x80 @ 15000fps 61μs

PREVIOUS WORK

- Congedo, Goncalves, Rodio
  "About the uncertainty quantification of turbulence and cavitation"
  European Journal of Mechanics B/Fluids 53 (2015) 190–204
- 2D, sDEM [Abgrall, 2015], forward UQ propagation

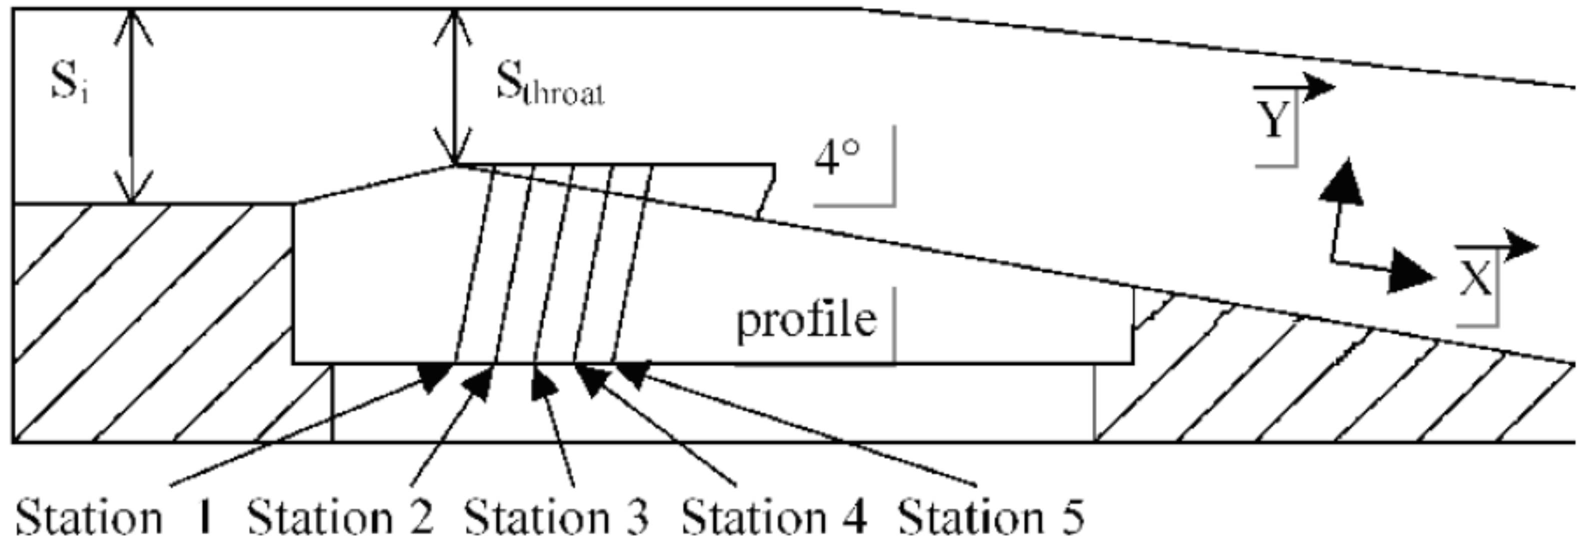

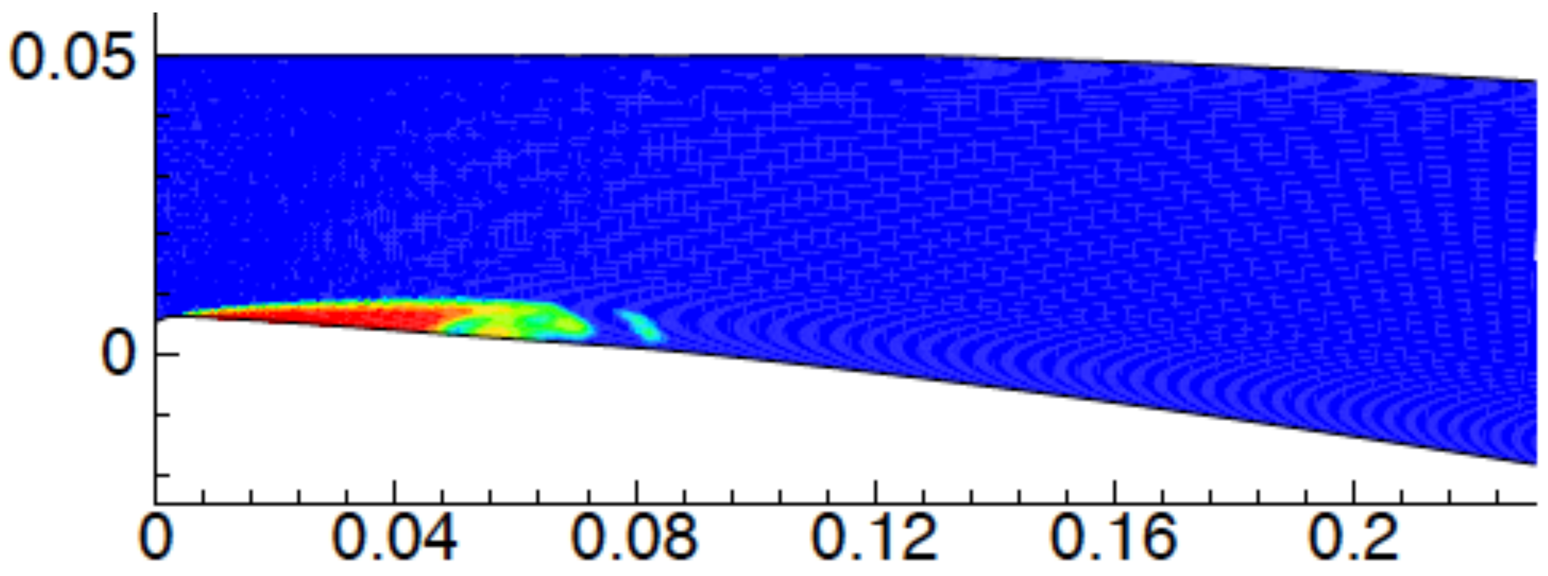

mean of vapor fraction

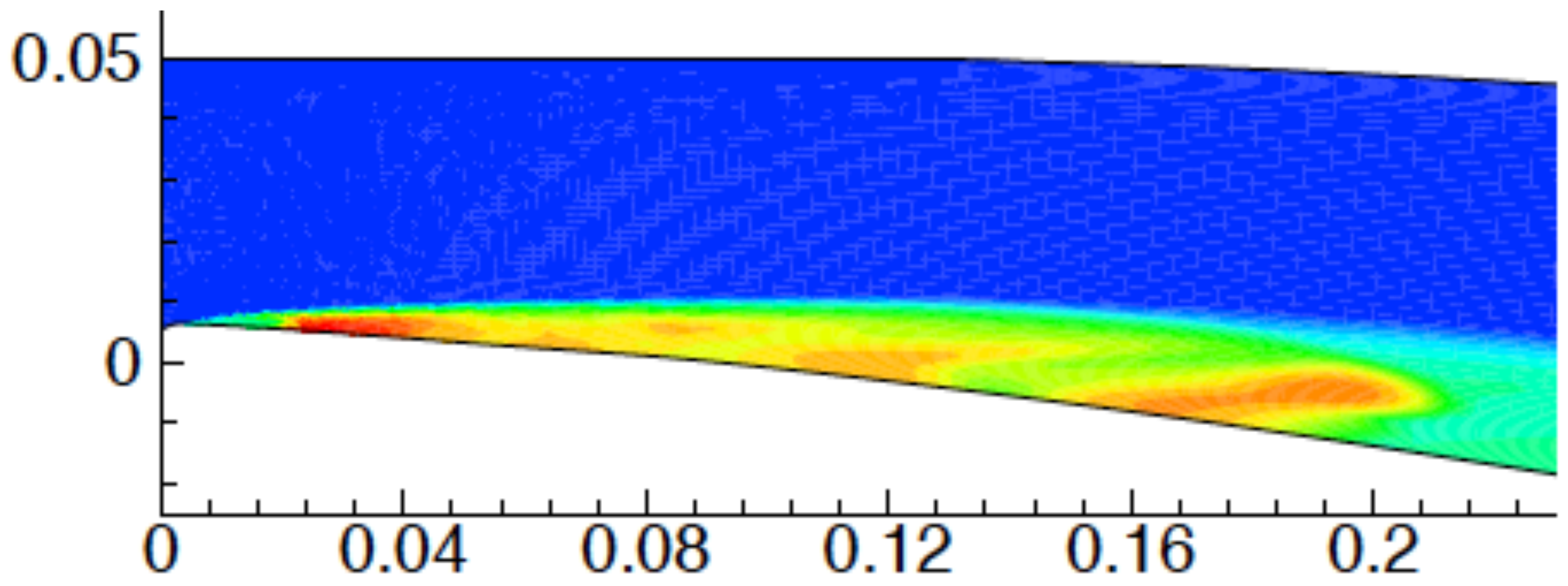

variance of vapor fraction

GOALS

- **Confidence interval** estimation for local integral **quantities of interest**
  - multiple **sensors for pressure**, density, speed of sound, etc.
- Investigation of observed transition from random to focused micro-collapses for spherical clouds
- **Fault tolerance**

Multi-Level Monte Carlo [Heinrich, 1999] [Giles, 2008]

Variance reduction technique using sampling on a hierarchy of mesh resolutions

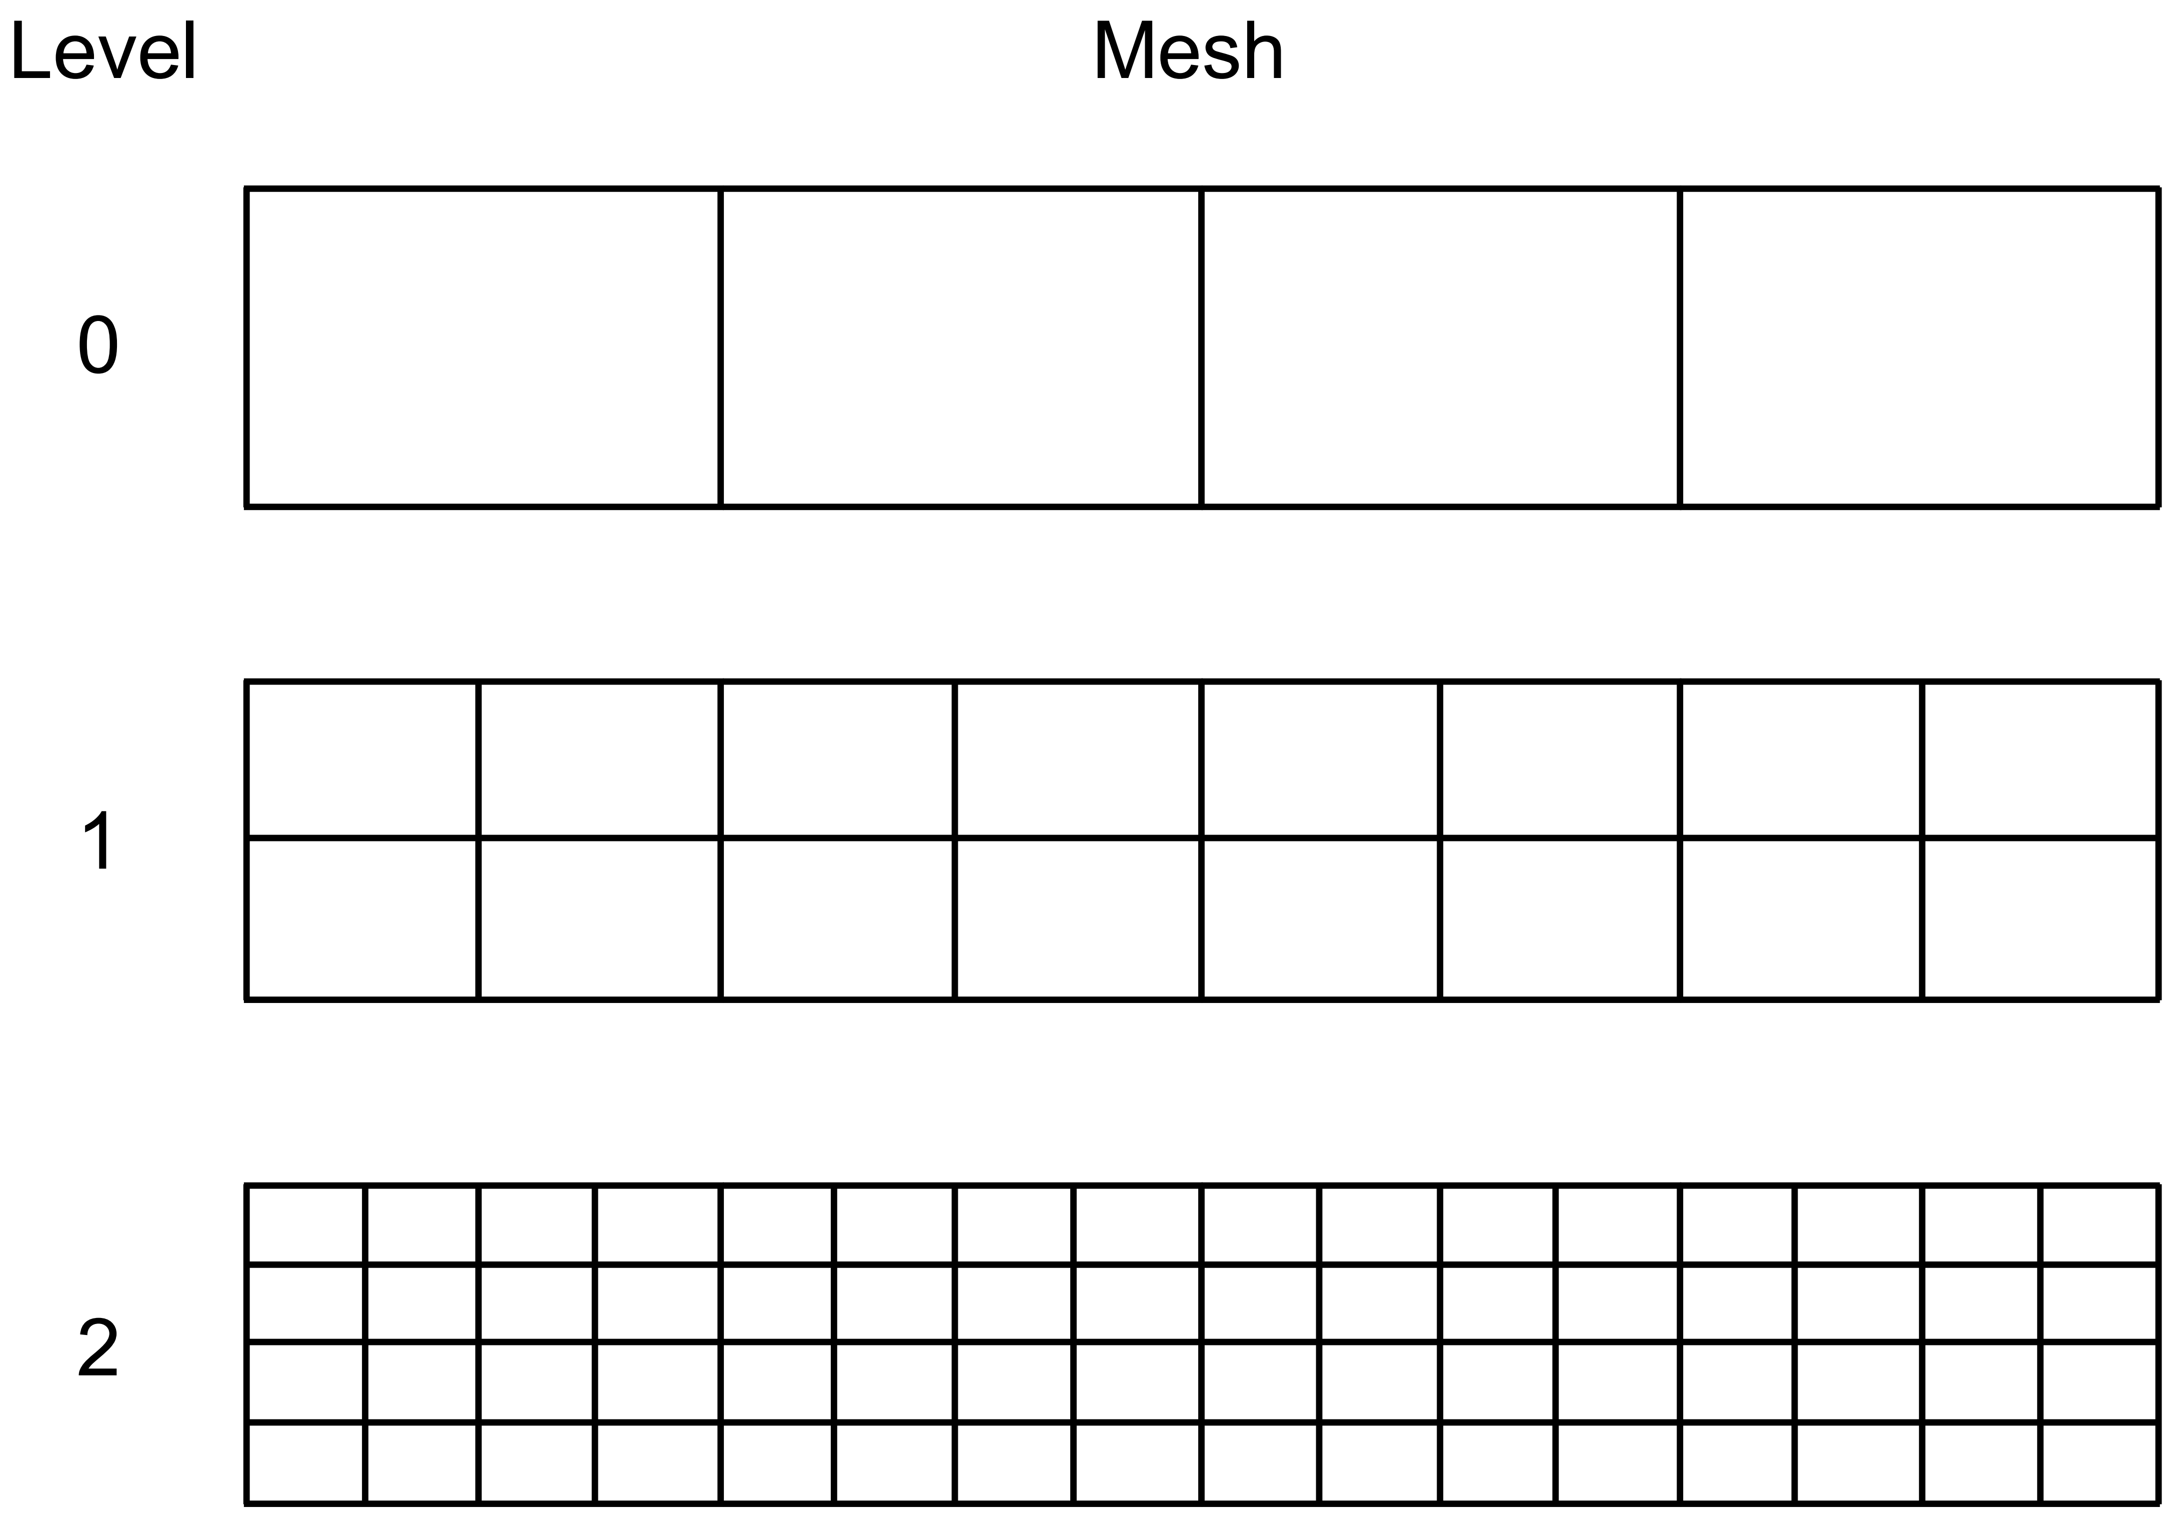

Multi-Level Monte Carlo method

Variance reduction technique using sampling on a hierarchy of mesh resolutions

1. Generate i.i.d. samples of random input quantities for each resolution level 0…L
2. For each level and sample, solve for approximate solutions using Cubism-MPCF
3. Assemble MLMC estimator for statistics of quantities of interest:

$$\mathbb{E}[q_L] = \mathbb{E}[q_0] + \sum_{\ell=1}^{L} \left(\mathbb{E}[q_\ell] - \mathbb{E}[q_{\ell-1}]\right) \approx \frac{1}{M_0}\sum_{i=1}^{M_0} q_0^i + \sum_{\ell=1}^{L} \frac{1}{M_\ell}\sum_{i=1}^{M_\ell}(q_\ell^i - q_{\ell-1}^i).$$

- Sampling **error** of the MLMC estimator is given in terms of **level correlations**:

$$\varepsilon^2 = \frac{\mathbb{V}[q_0]}{M_0} + \sum_{\ell=1}^{L} \frac{\mathbb{V}[q_\ell - q_{\ell-1}]}{M_\ell} \approx \mathbb{V}[q]\left(\frac{1}{M_0} + 2\sum_{\ell=1}^{L}\frac{1-\mathrm{Cor}[q_\ell, q_{\ell-1}]}{M_\ell}\right).$$

Insight to inner workings of MLMC

Multiple samples for each resolution level

semi-transparent lines correspond to the **same sample** (realization) but on a **coarser resolution**

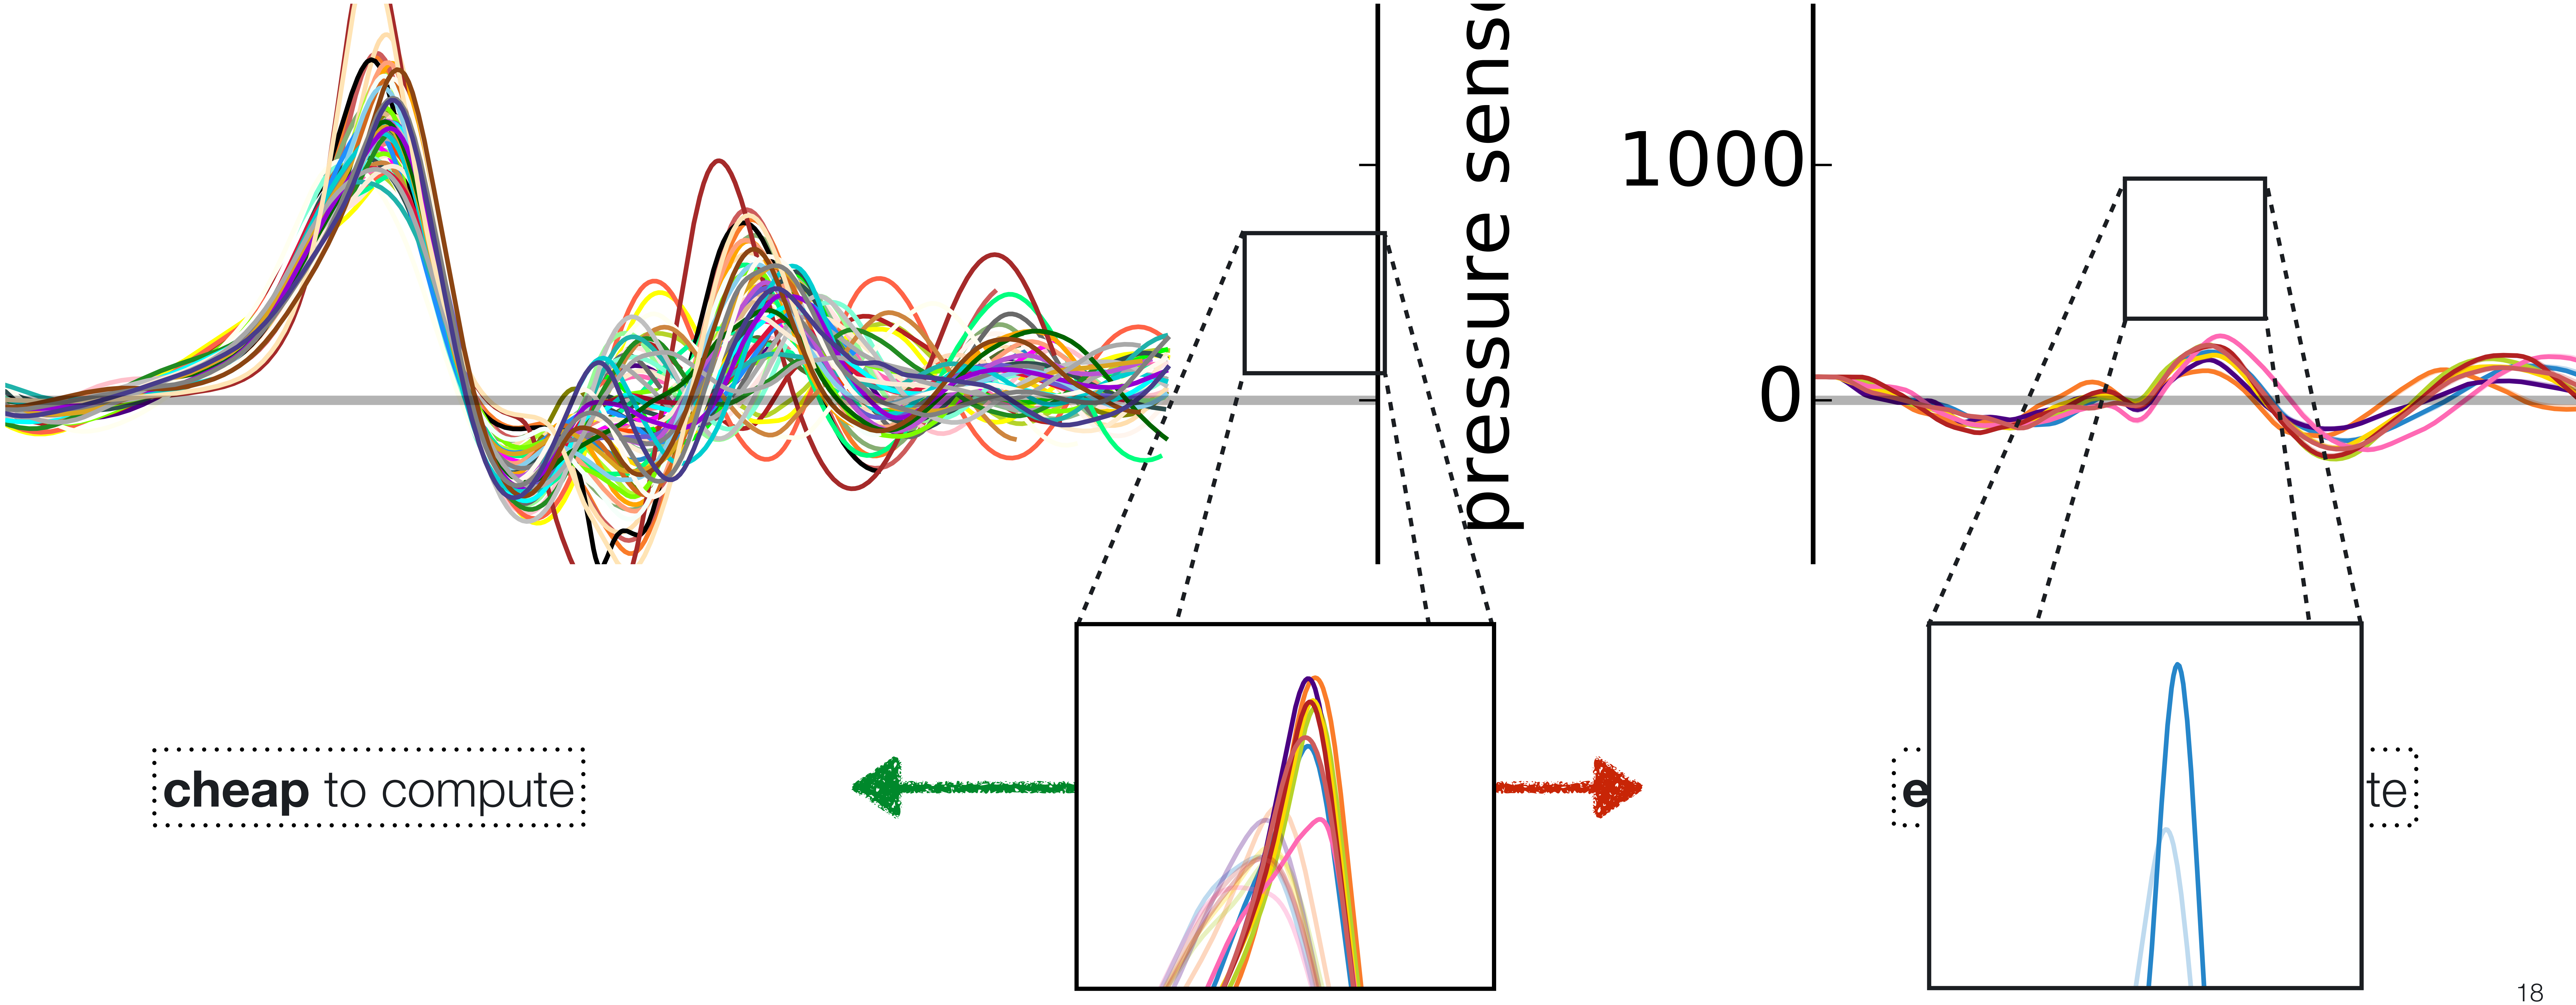

cheap to compute

Insight to inner workings of MLMC

Monte Carlo variance estimates for differences between resolution levels decrease

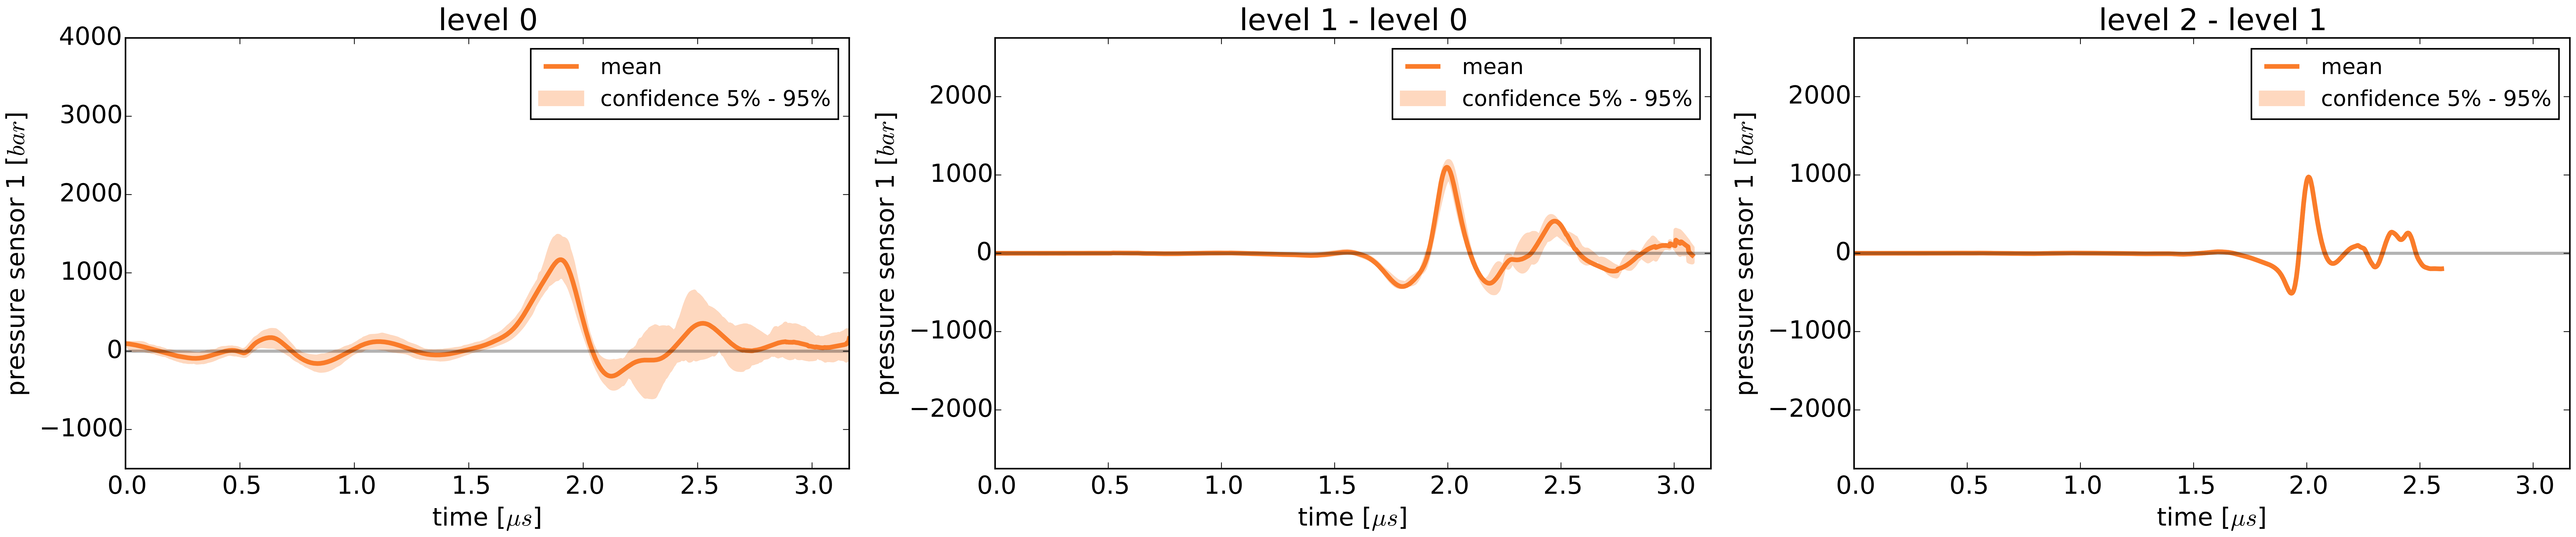

most of sampling is required here (cheap!)

expensive to compute

Insight to inner workings of MLMC

Variance estimates for differences between resolution levels decrease

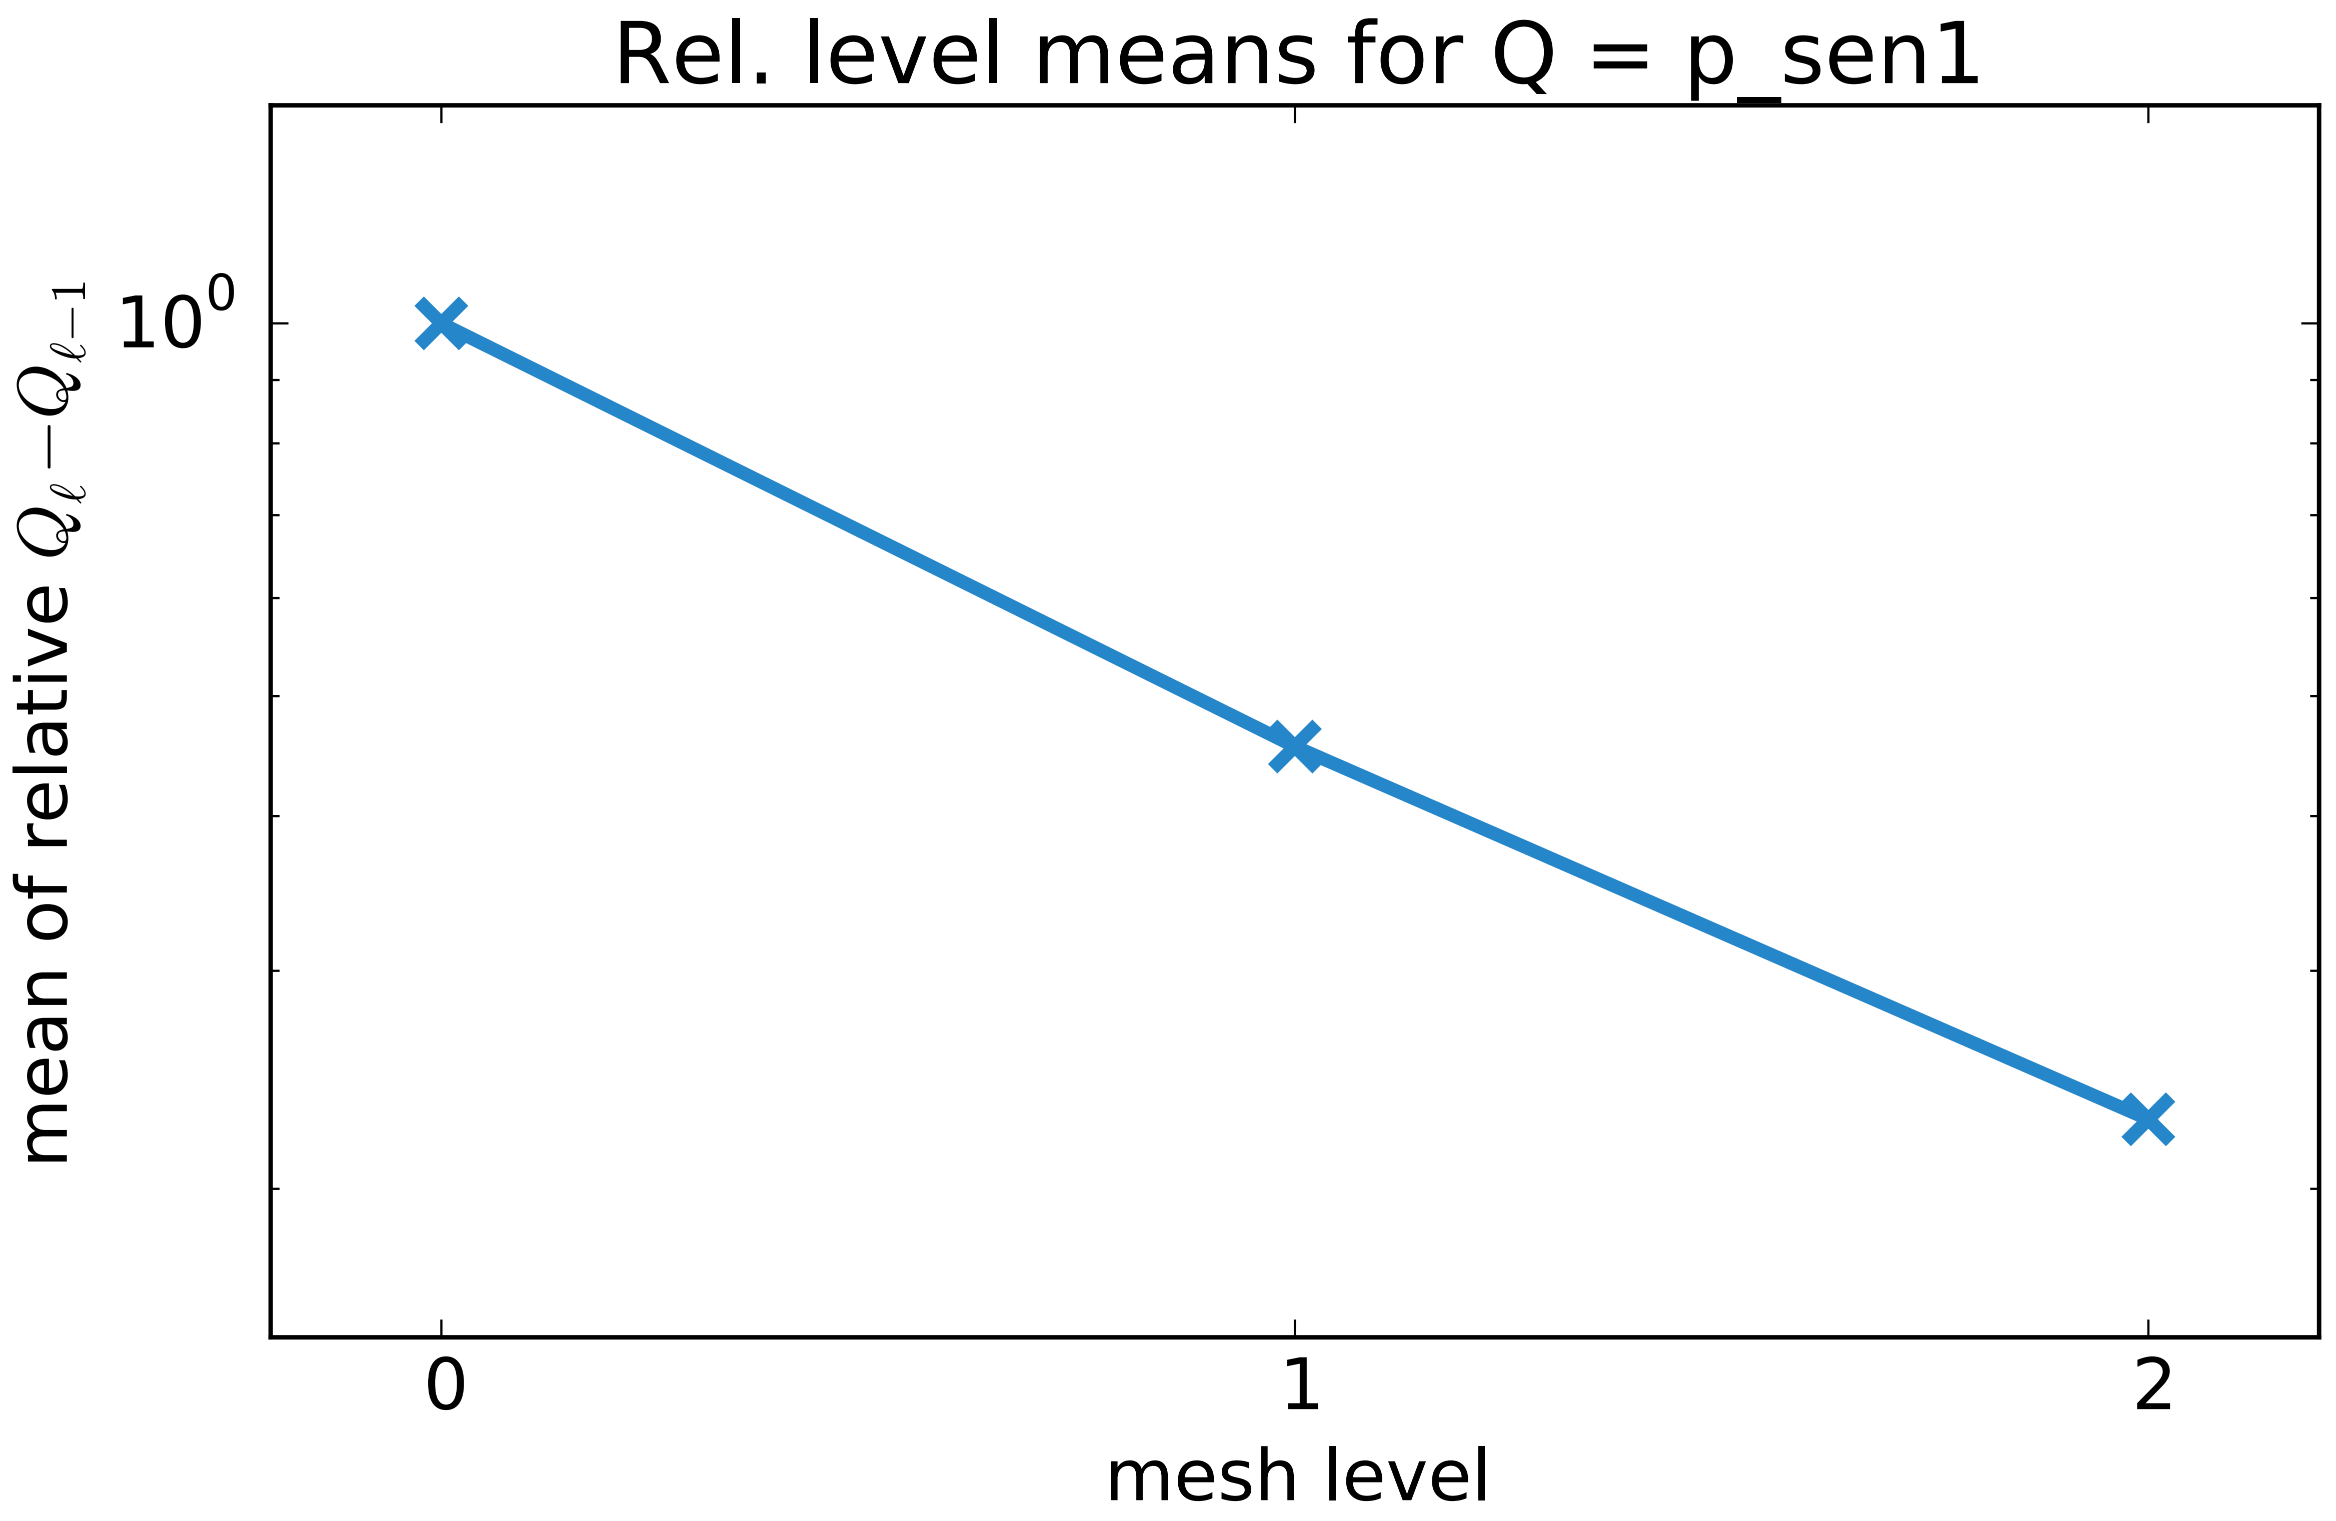

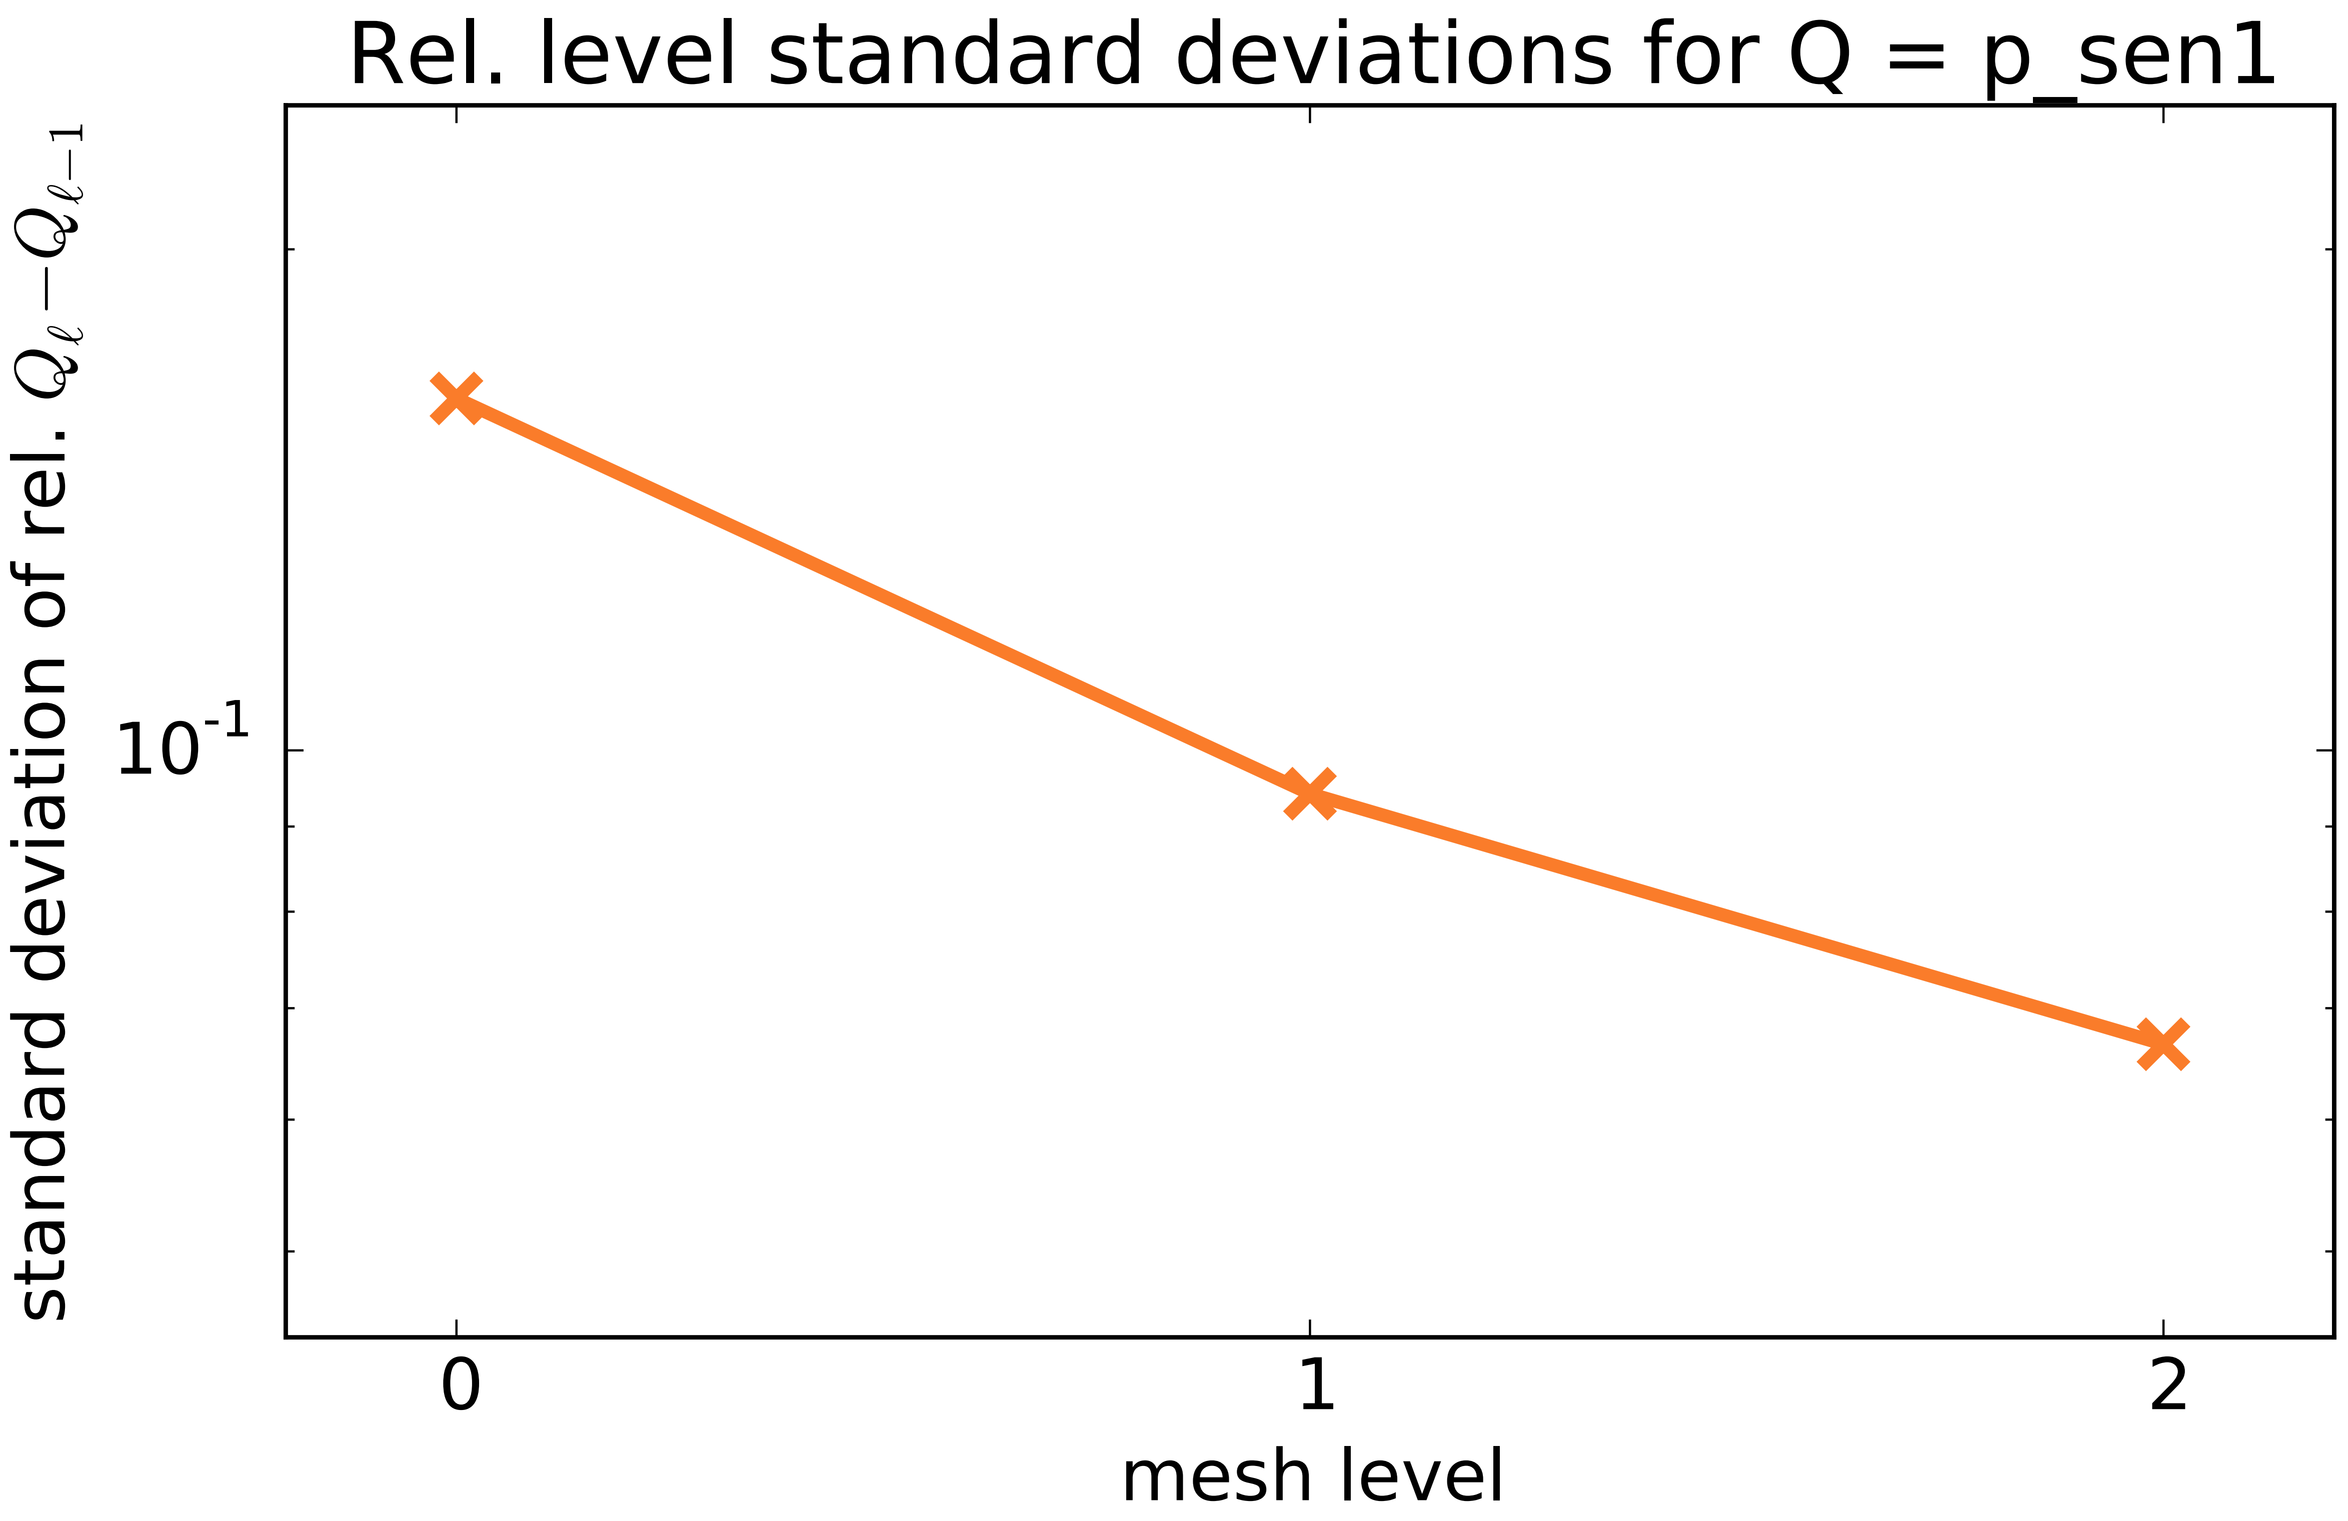

Assembly of the MLMC estimator

Monte Carlo estimates from each resolution level are combined together

$$\mathbb{E}[q_L] = \mathbb{E}[q_0] + \sum_{\ell=1}^{L} \left(\mathbb{E}[q_\ell] - \mathbb{E}[q_{\ell-1}]\right) \approx \frac{1}{M_0}\sum_{i=1}^{M_0} q_0^i + \sum_{\ell=1}^{L} \frac{1}{M_\ell}\sum_{i=1}^{M_\ell}(q_\ell^i - q_{\ell-1}^i).$$

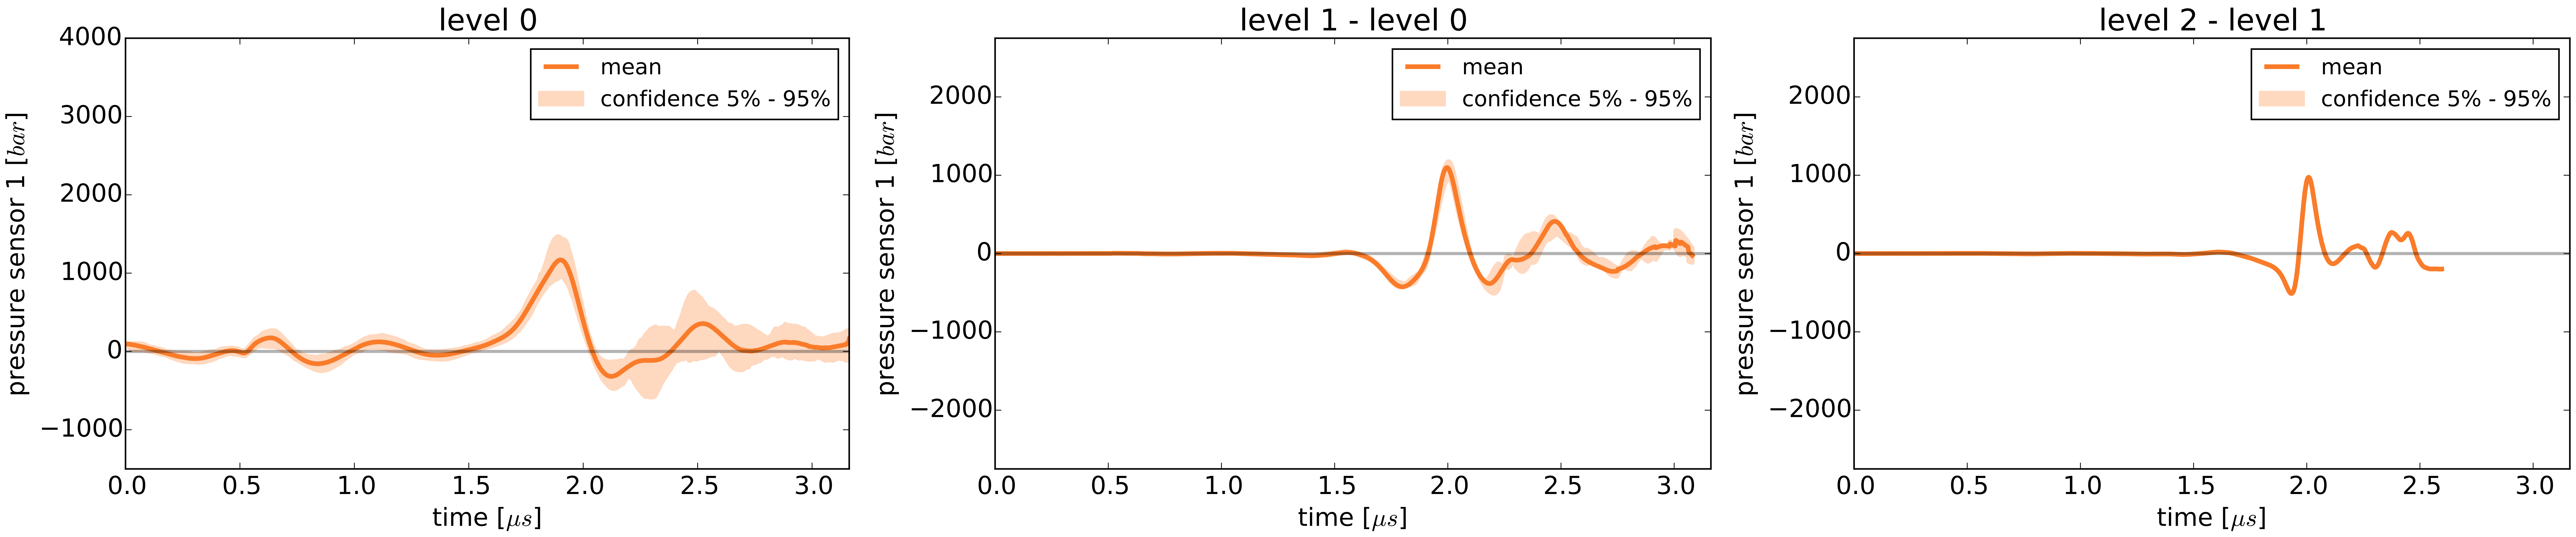

+ +

= MLMC estimate

Results of MLMC

Uncertainty quantification (i.e. mean, confidence intervals) for QoIs

vapor volume

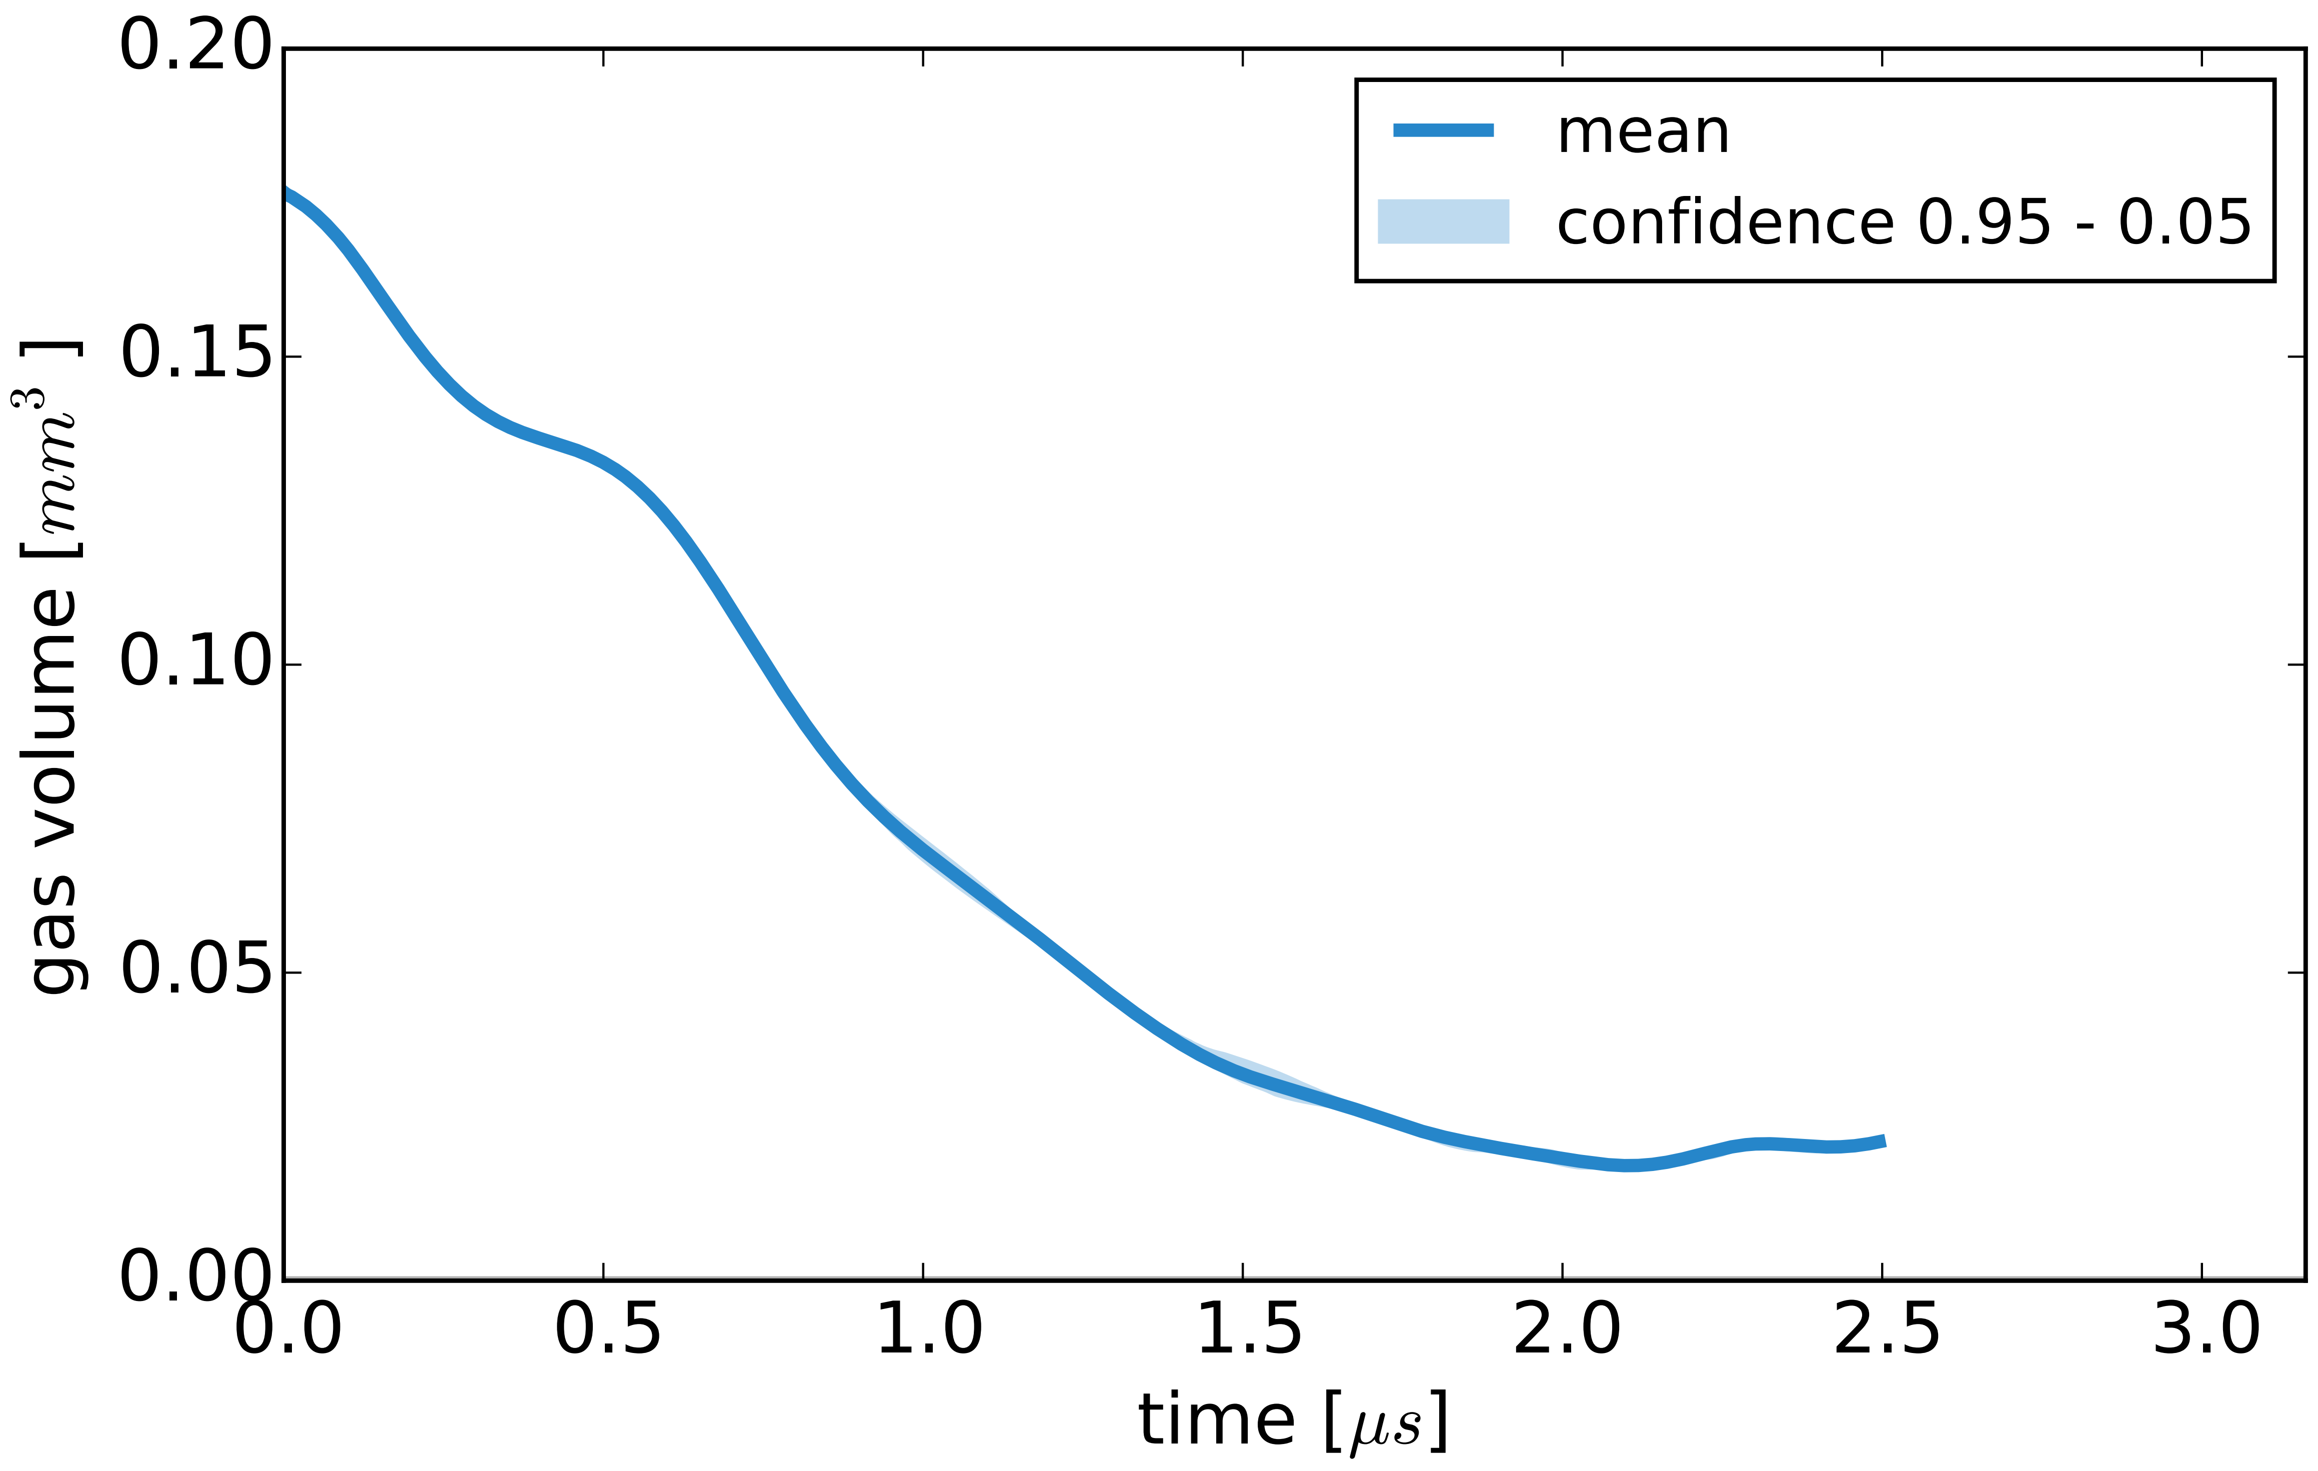

no significant uncertainty

pressure sensor

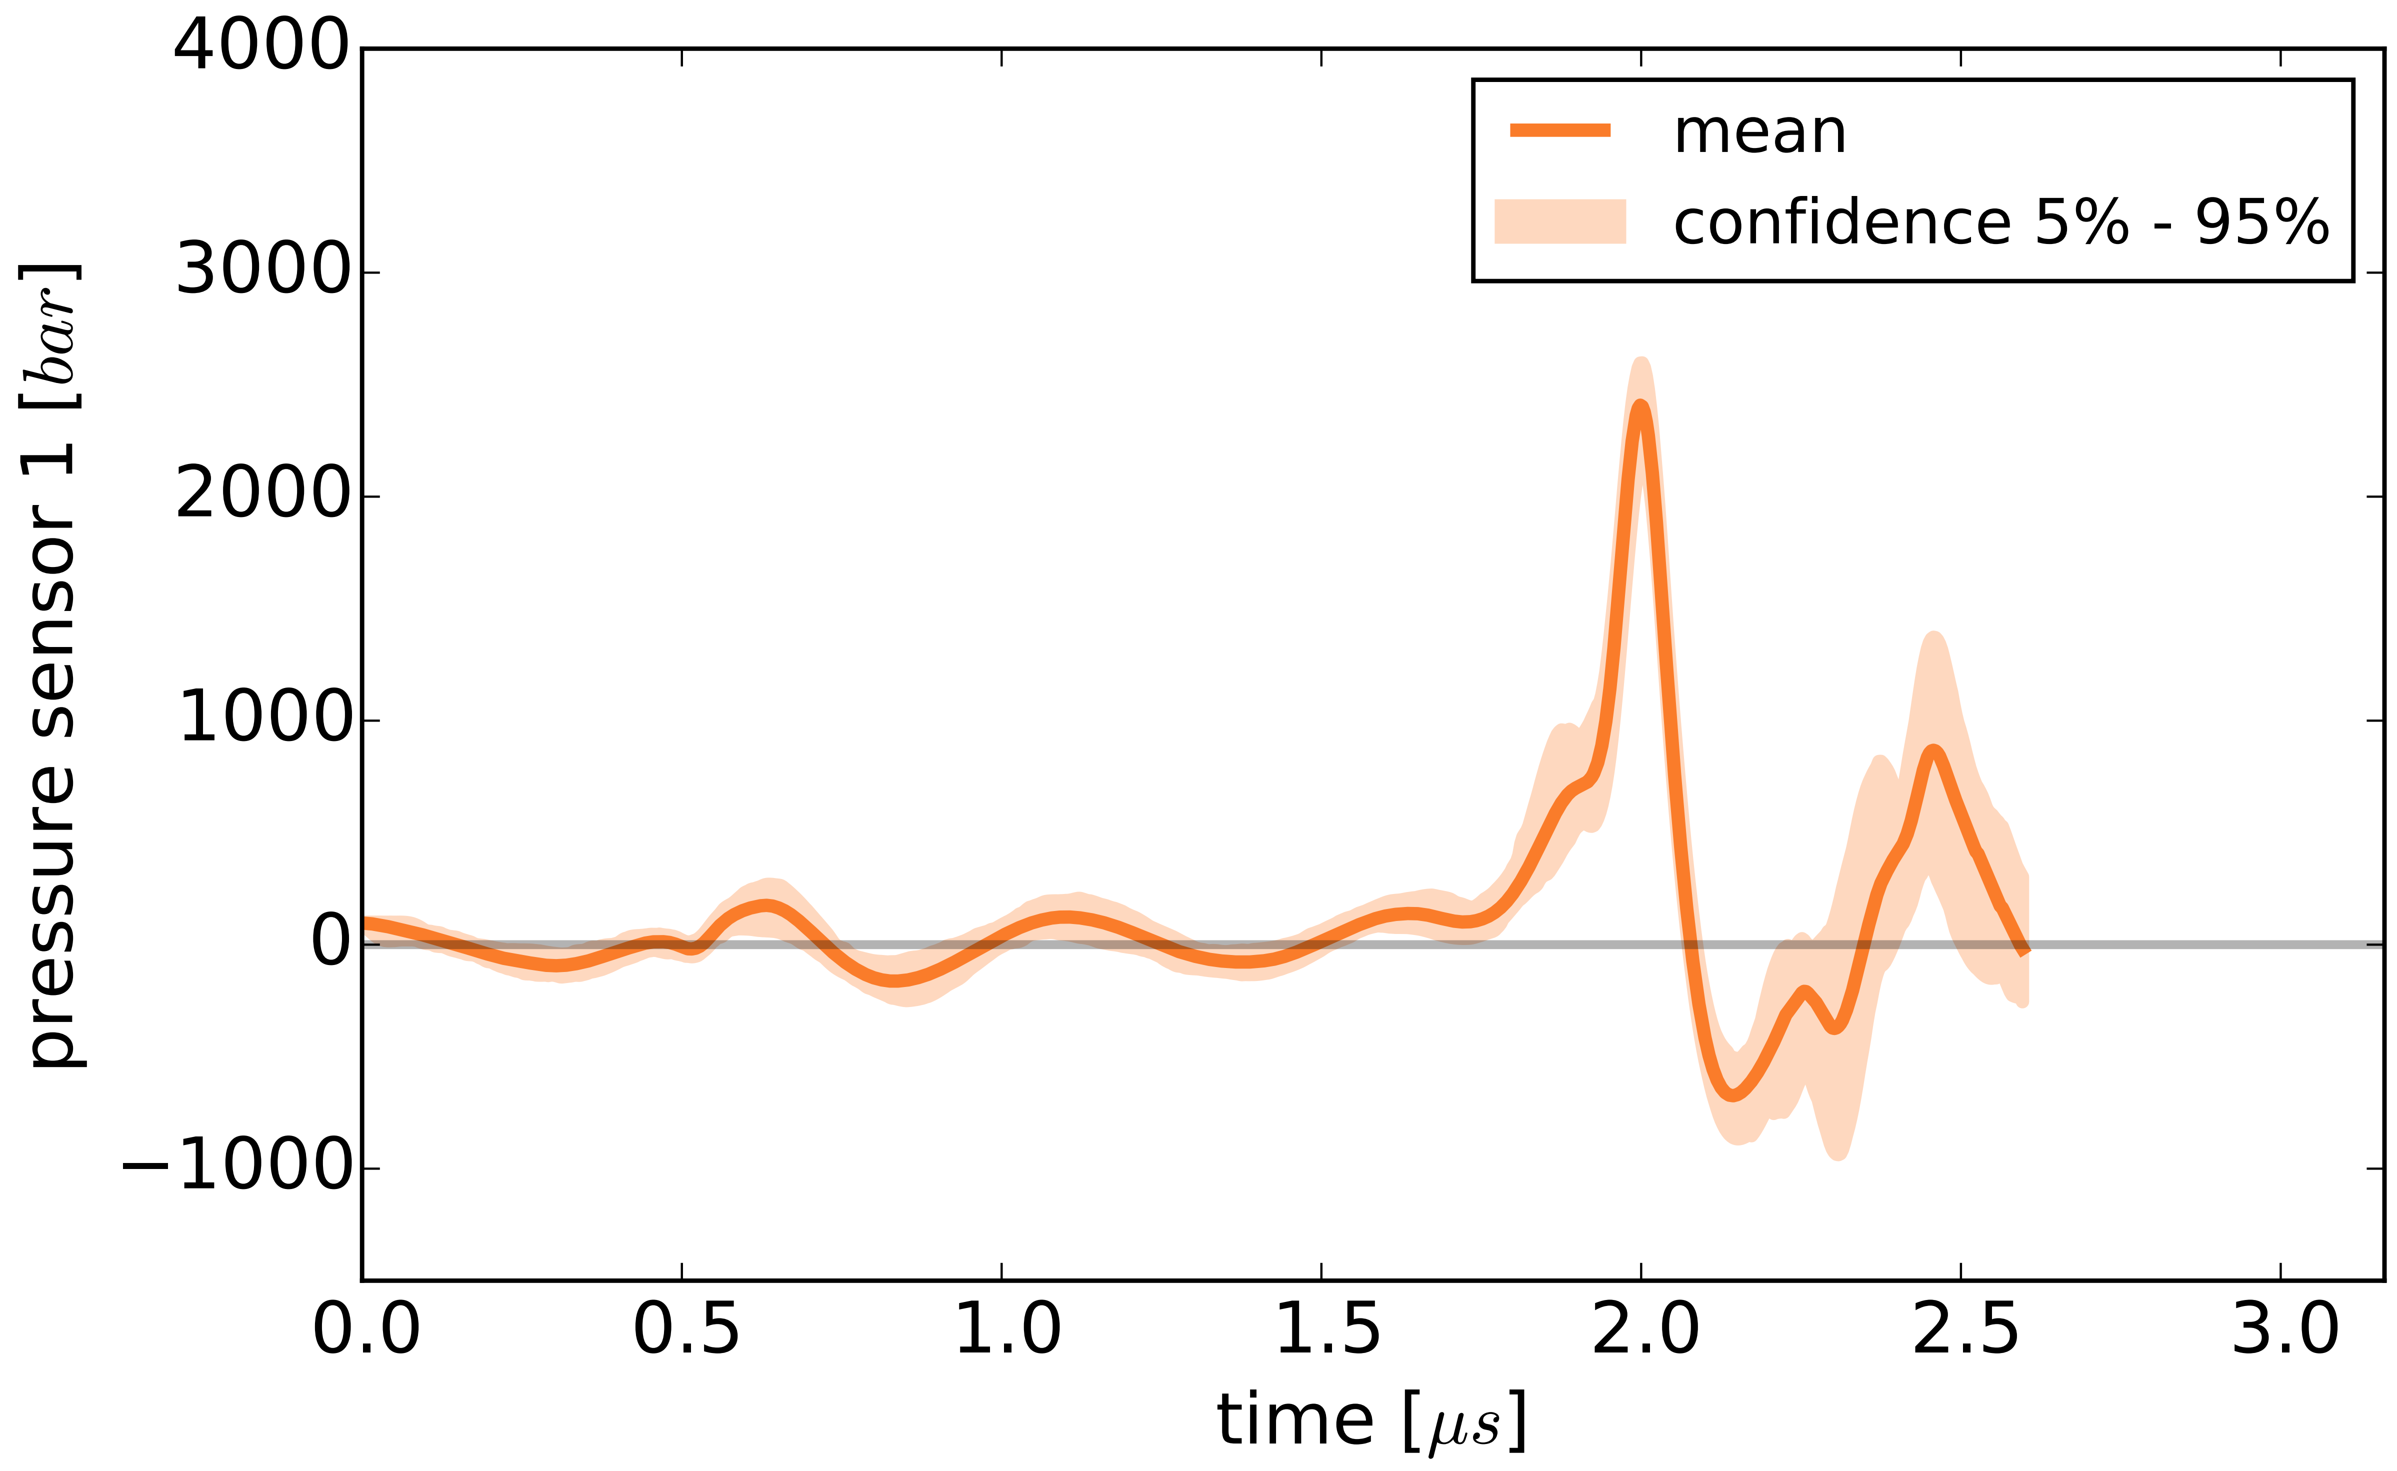

wide 90% confidence interval

2000 bar - 2500 bar

Optimized number of samples

Using empirical estimators for variances and measurements of computations work

- Sampling **error** of the MLMC estimator is given in terms of **level variances**:

$$\varepsilon^2 = \frac{\mathbb{V}[q_0]}{M_0} + \sum_{\ell=1}^{L} \frac{\mathbb{V}[q_\ell - q_{\ell-1}]}{M_\ell} \approx \frac{\sigma_0^2}{M_0} + \sum_{\ell=1}^{L} \frac{\sigma_\ell^2}{M_\ell}.$$

Optimization problem

Given a required tolerance τ and variances σ_ℓ^2 each level, minimize computational work and find optimal number of samples such that tolerance is attained: $\varepsilon \leq \tau$.

Optimized number of samples

Using Lagrange multipliers for derivations, optimized number of samples are given by

$$M_\ell = \left\lceil \frac{1}{\tau^2} \sqrt{\frac{\sigma_\ell^2}{\mathrm{Work}_\ell}} \sum_{k=0}^{L} \sqrt{\sigma_k^2 \mathrm{Work}_k} \right\rceil .$$

Remark: an analogous result is available for a prescribed computational budget (instead of tolerance).

PyMLMC

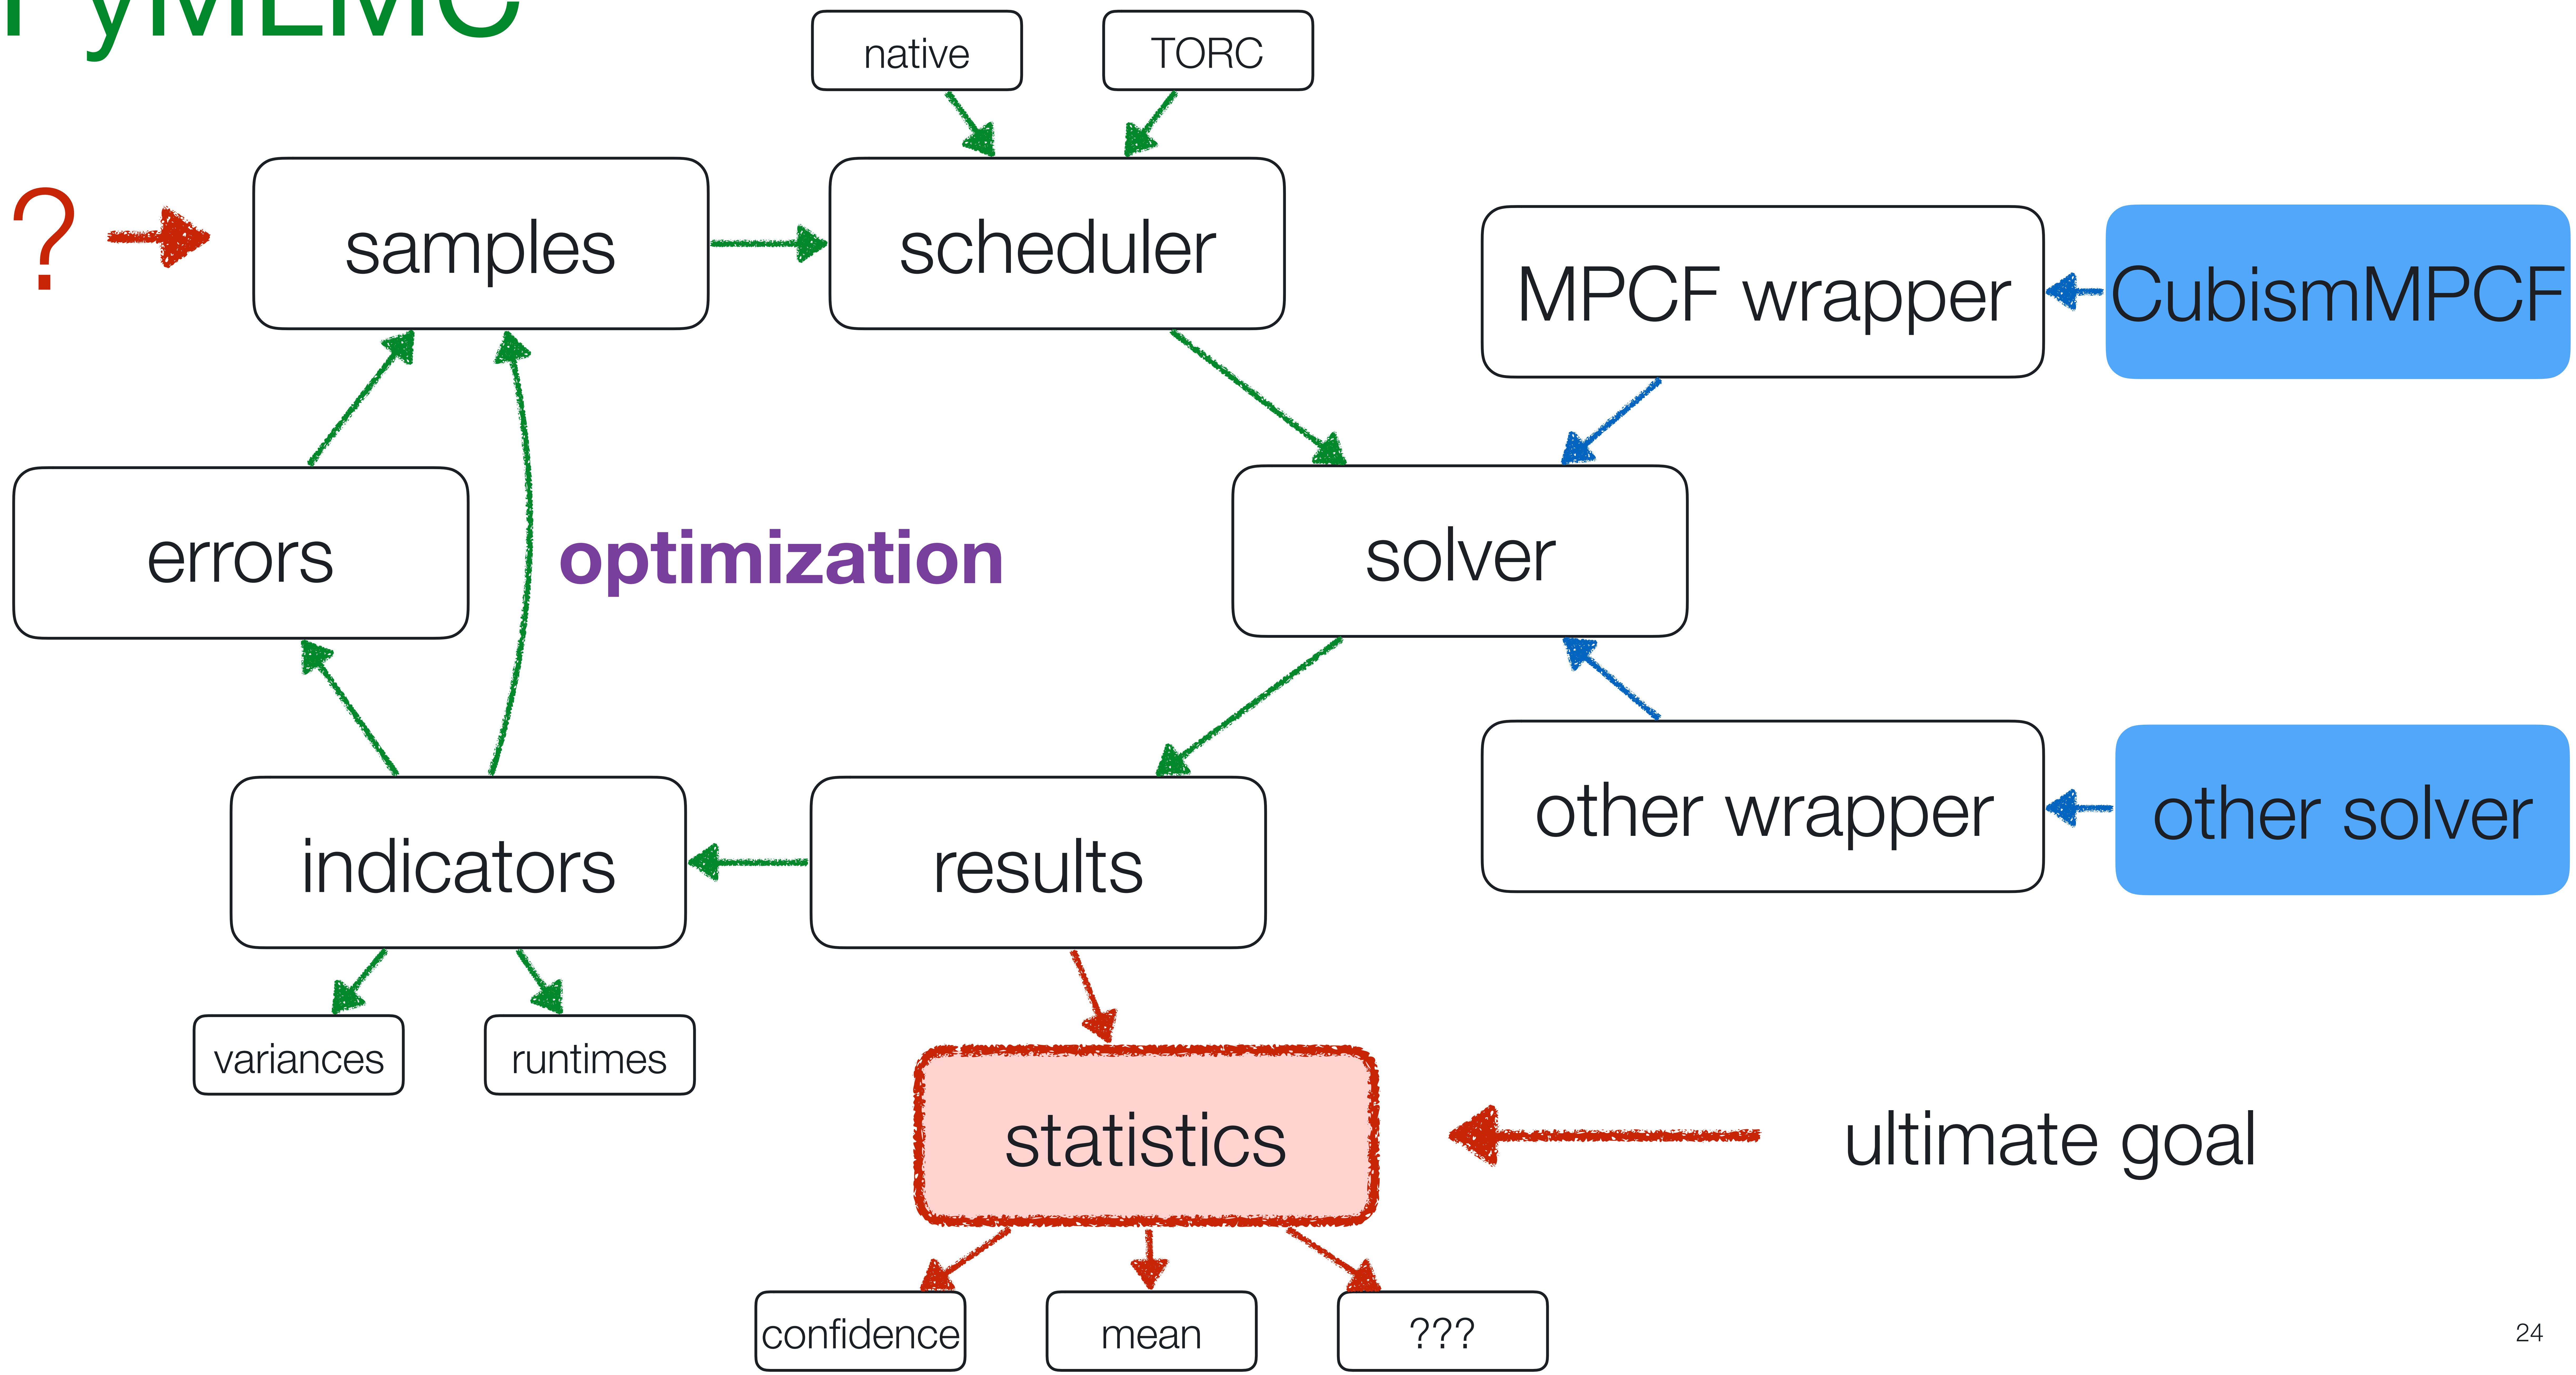

native
TORC
?
samples
scheduler
MPCF wrapper
CubismMPCF
errors
optimization
solver
indicators
results
other wrapper
other solver
variances
runtimes
statistics
ultimate goal
confidence
mean
???

Insight to inner workings of MLMC

Majority of samples computed on lowest levels of resolution - reduced budget

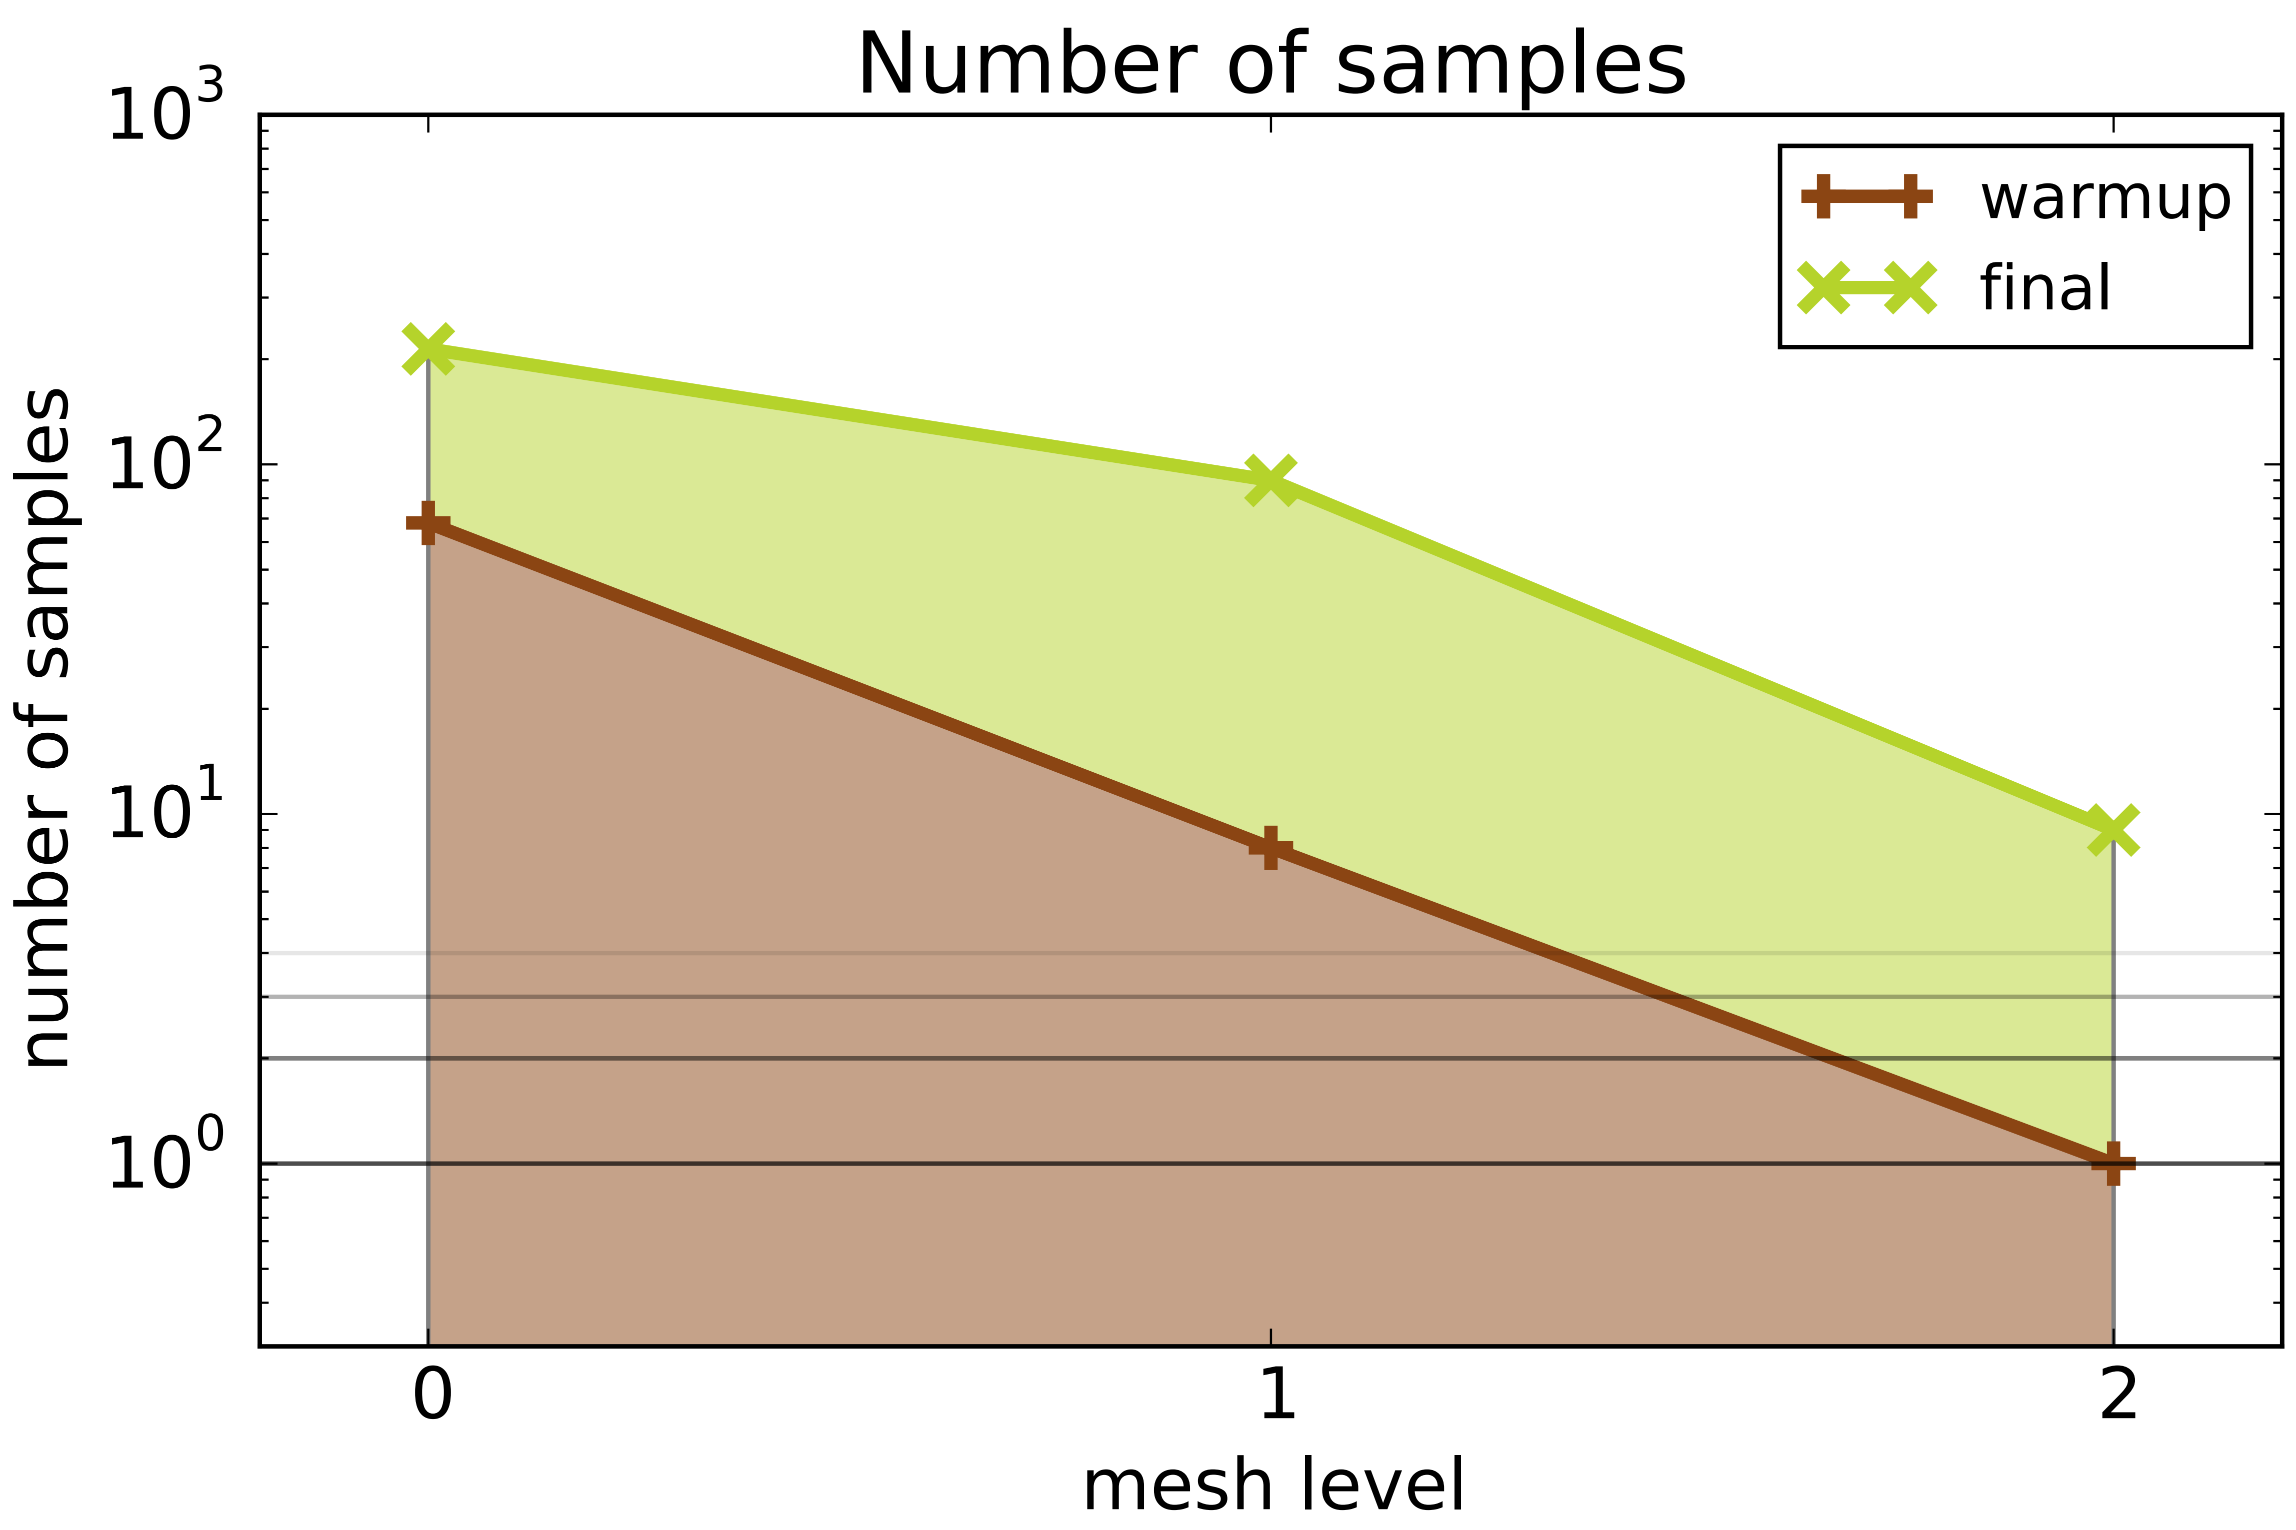

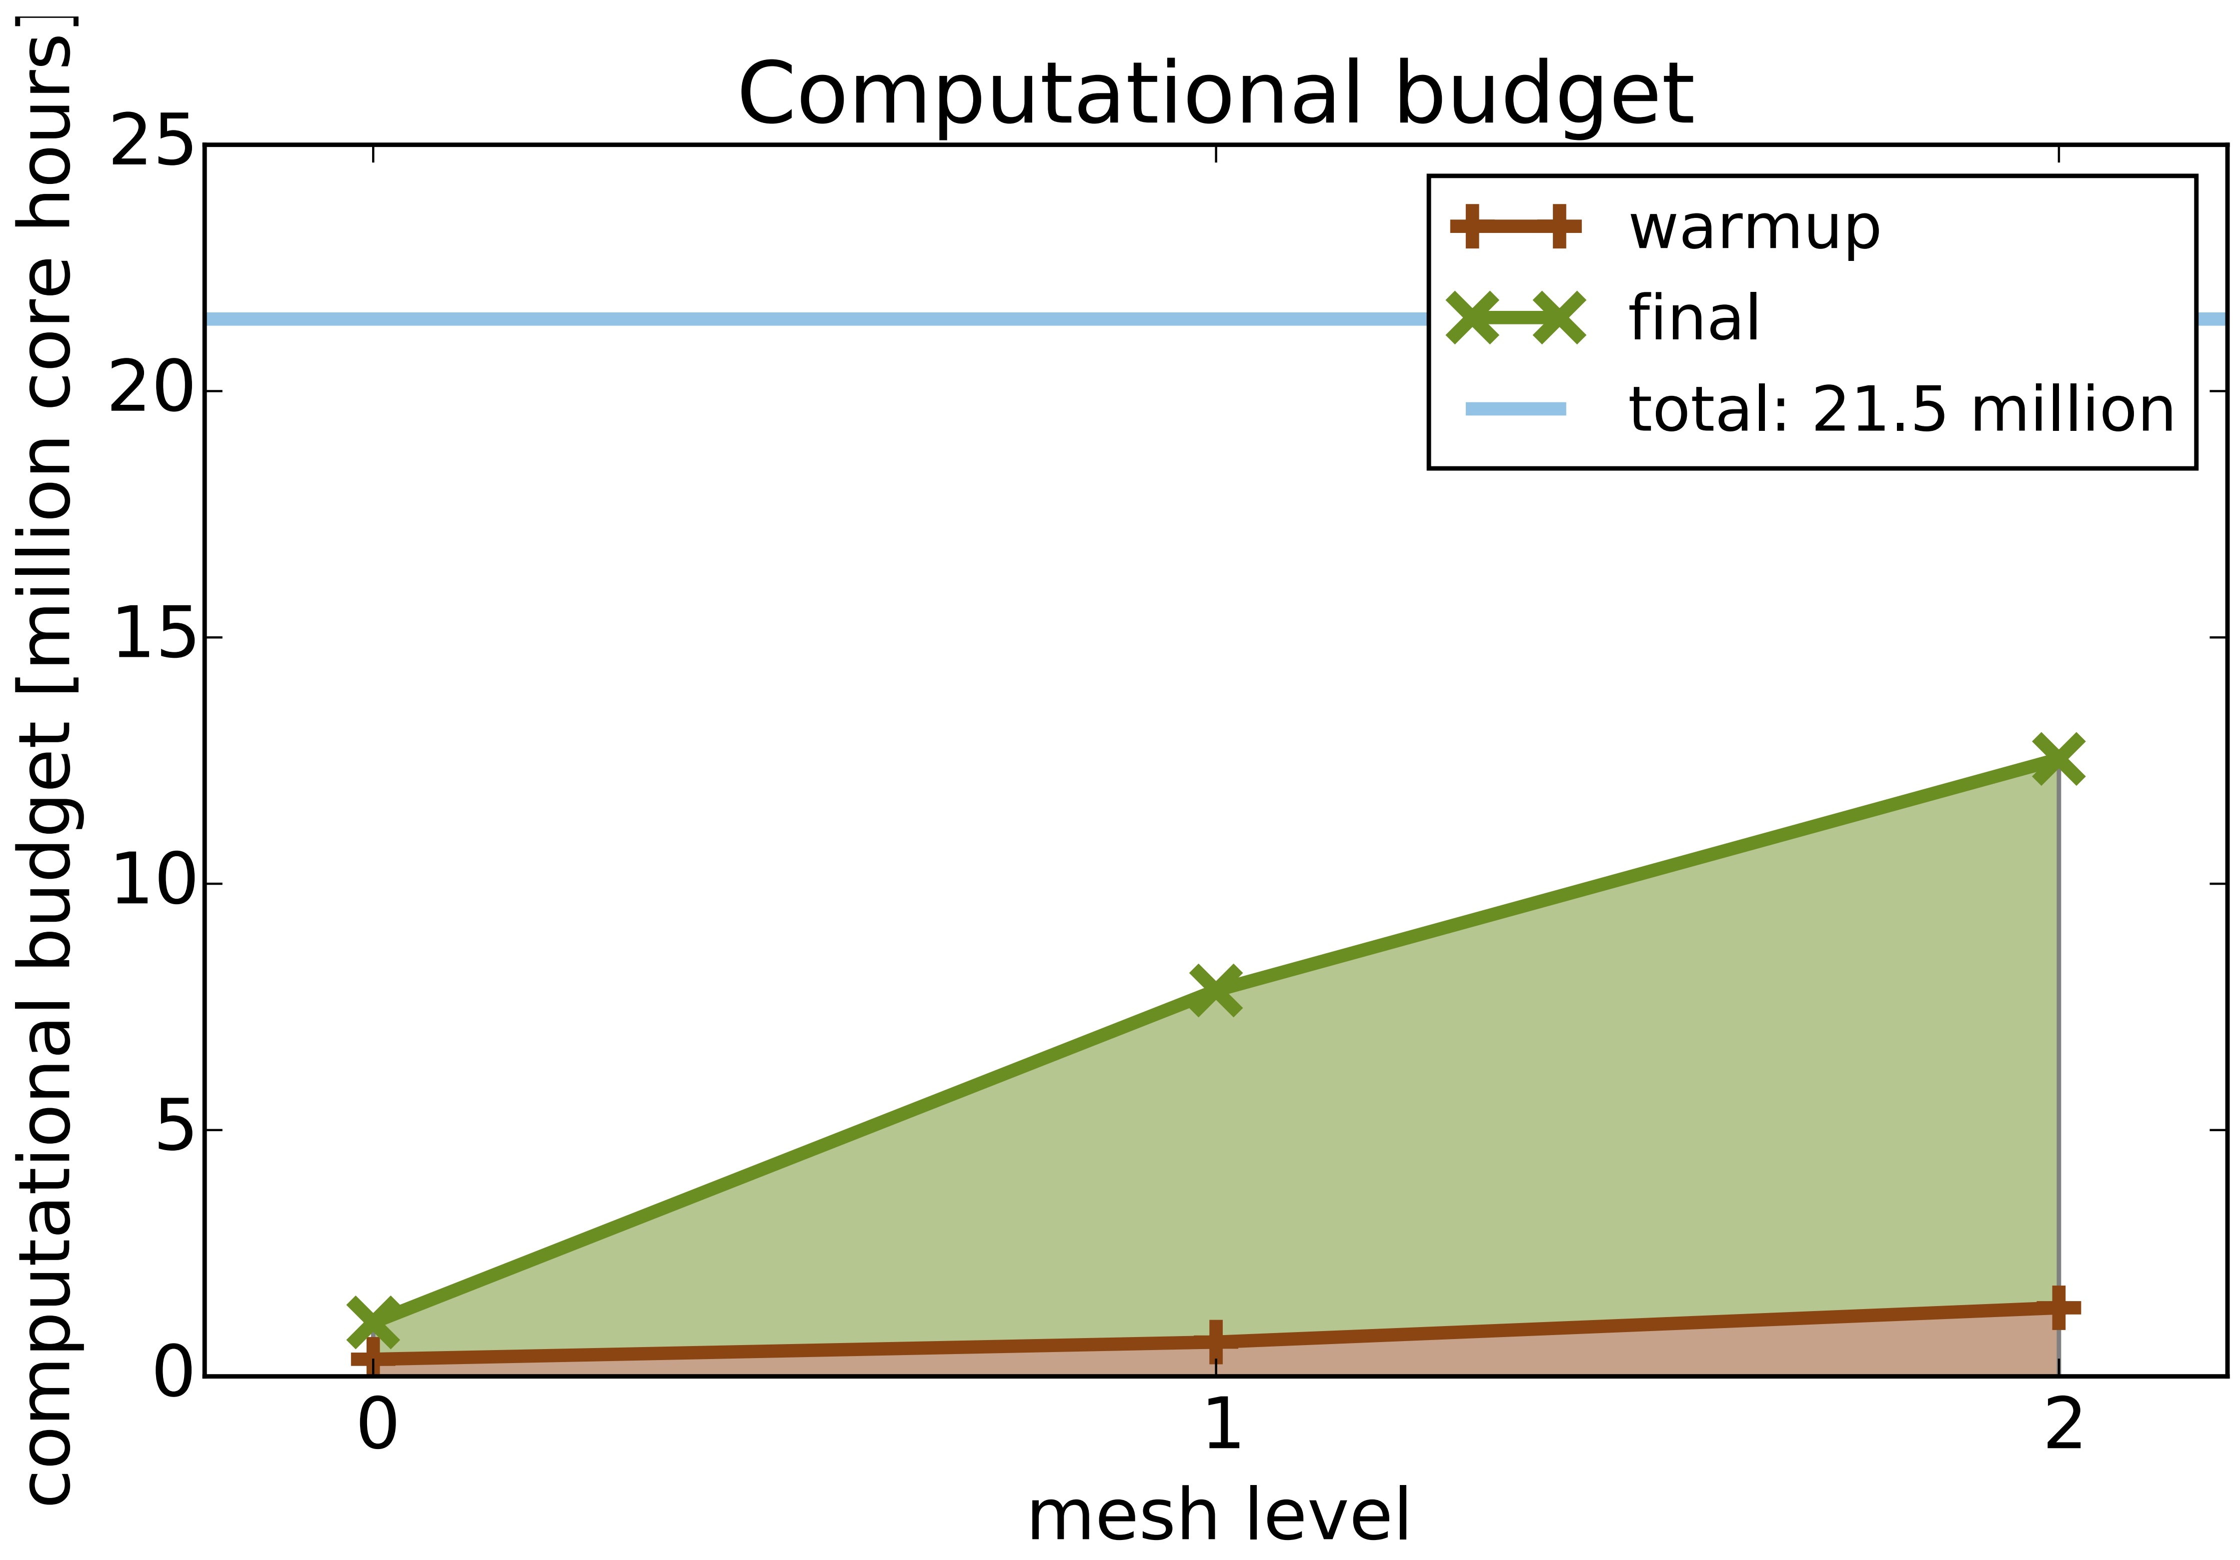

adaptive number of warmup samples

observed speedup: 5.8x

Uncertainty quantification in observed transition from random to focused & synchronous micro-collapses

Uncertainty Quantification

Multi-Level Monte Carlo estimation of mean values and 90% confidence intervals

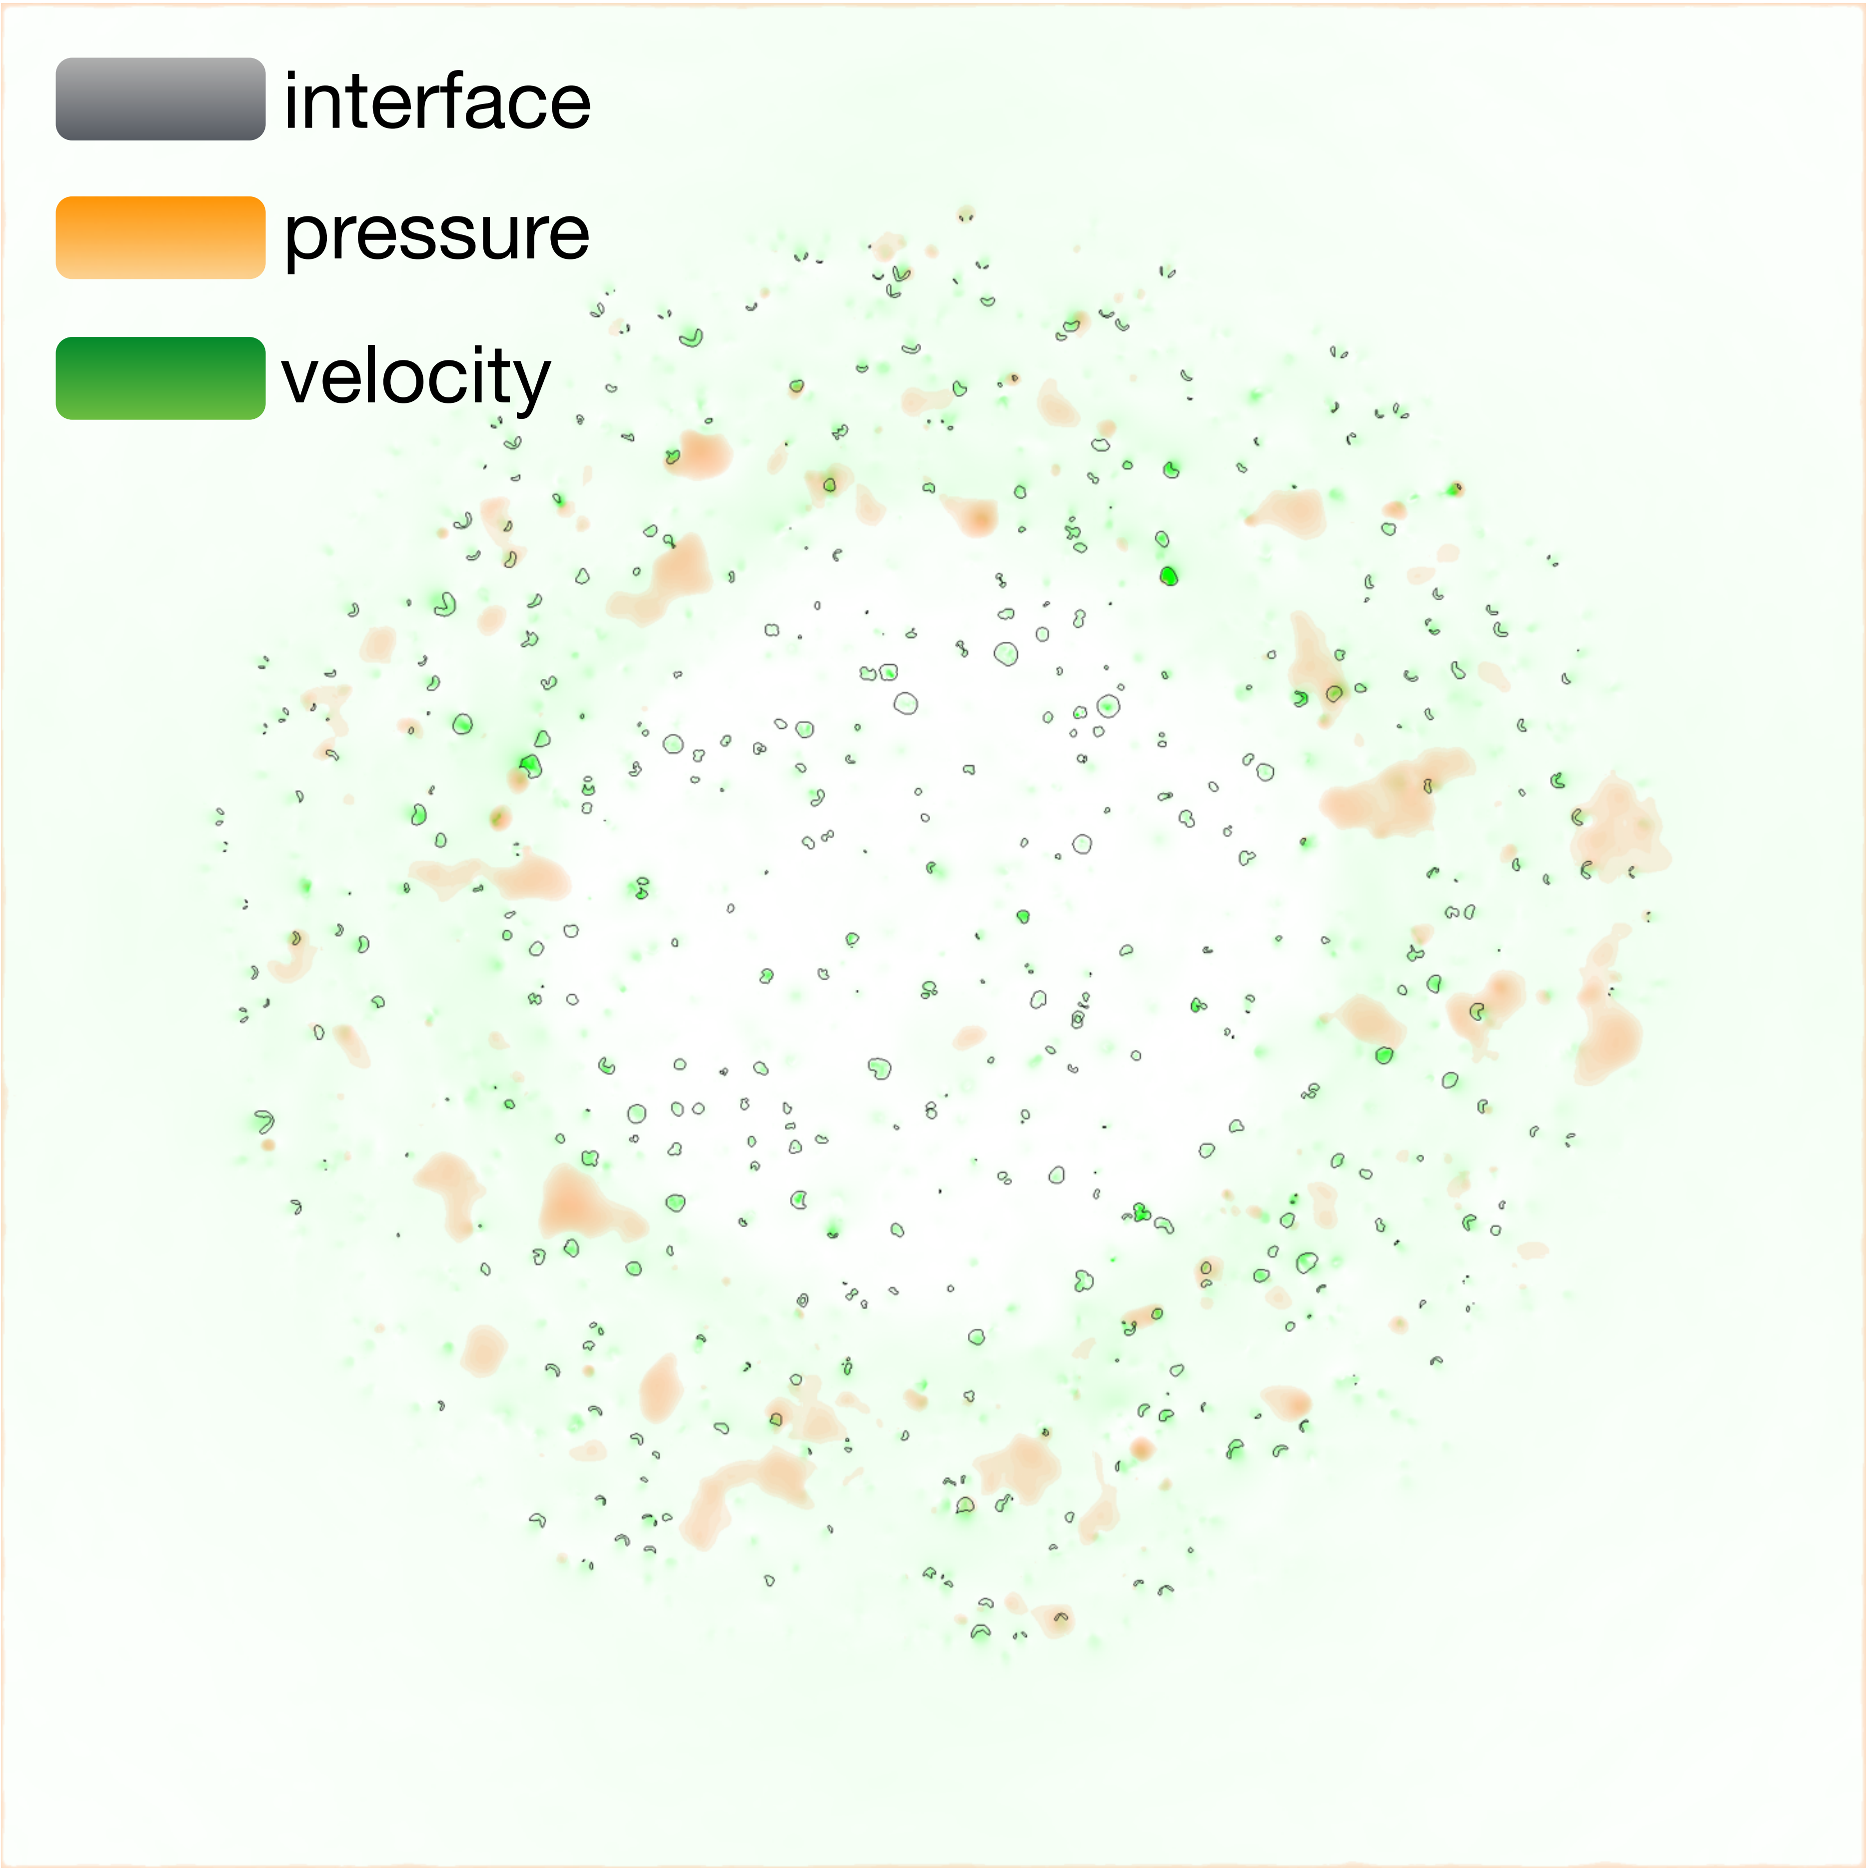

uncertain positions of initial collapses

distance from the cloud center

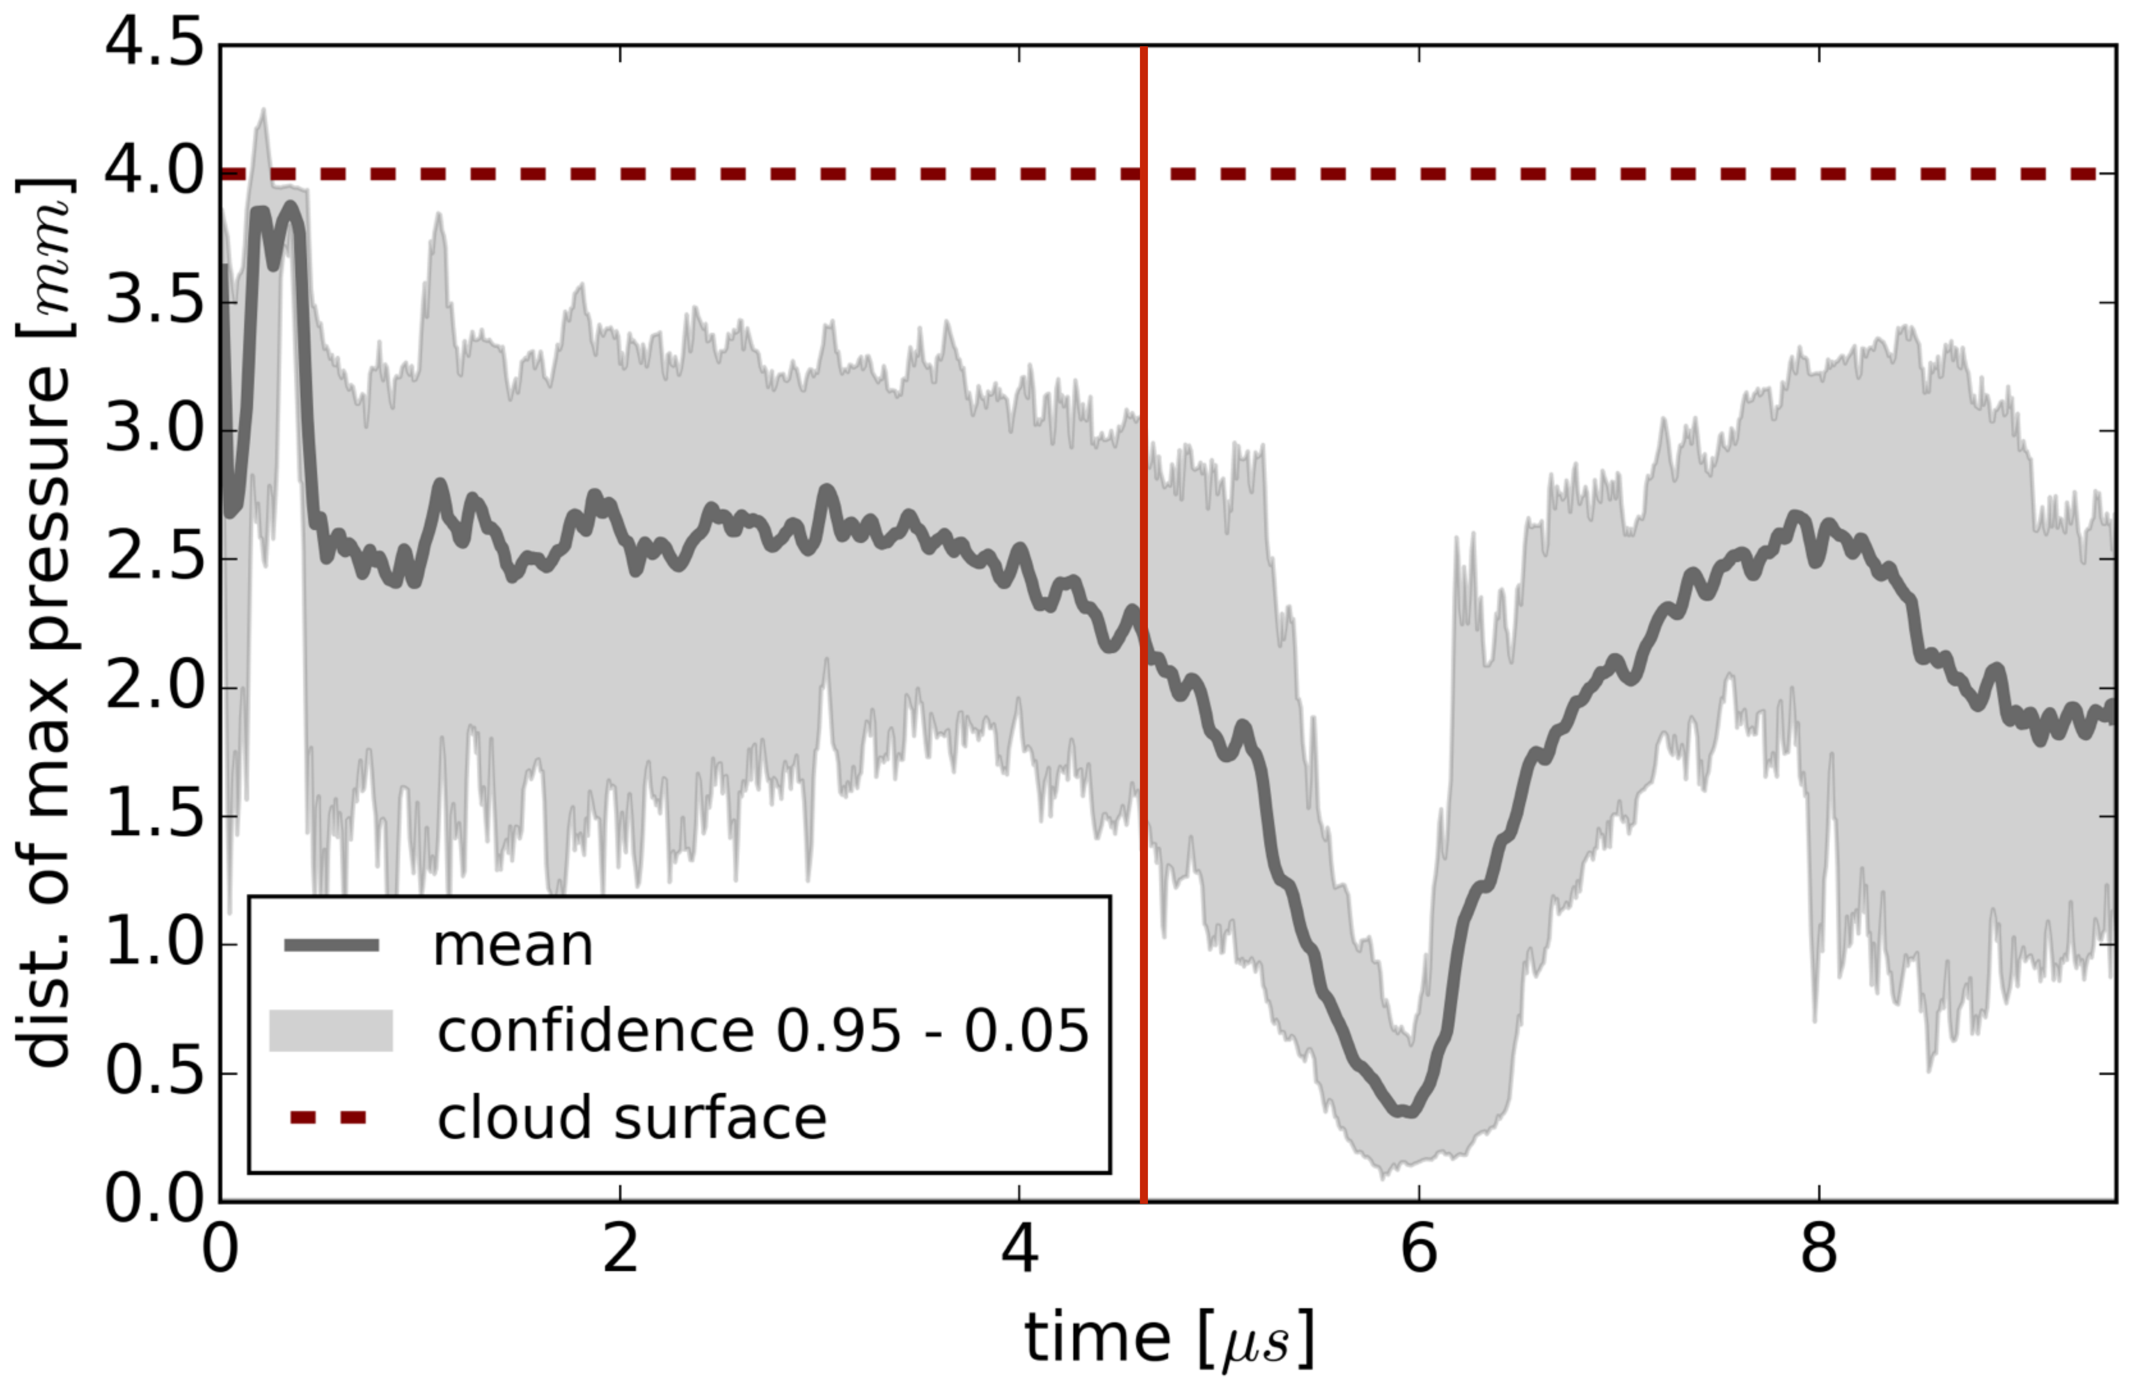

pre-collapse: **wide** confidence interval

collapse: **narrow** confidence interval

Uncertainty Quantification

Multi-Level Monte Carlo estimation of mean values and 90% confidence intervals

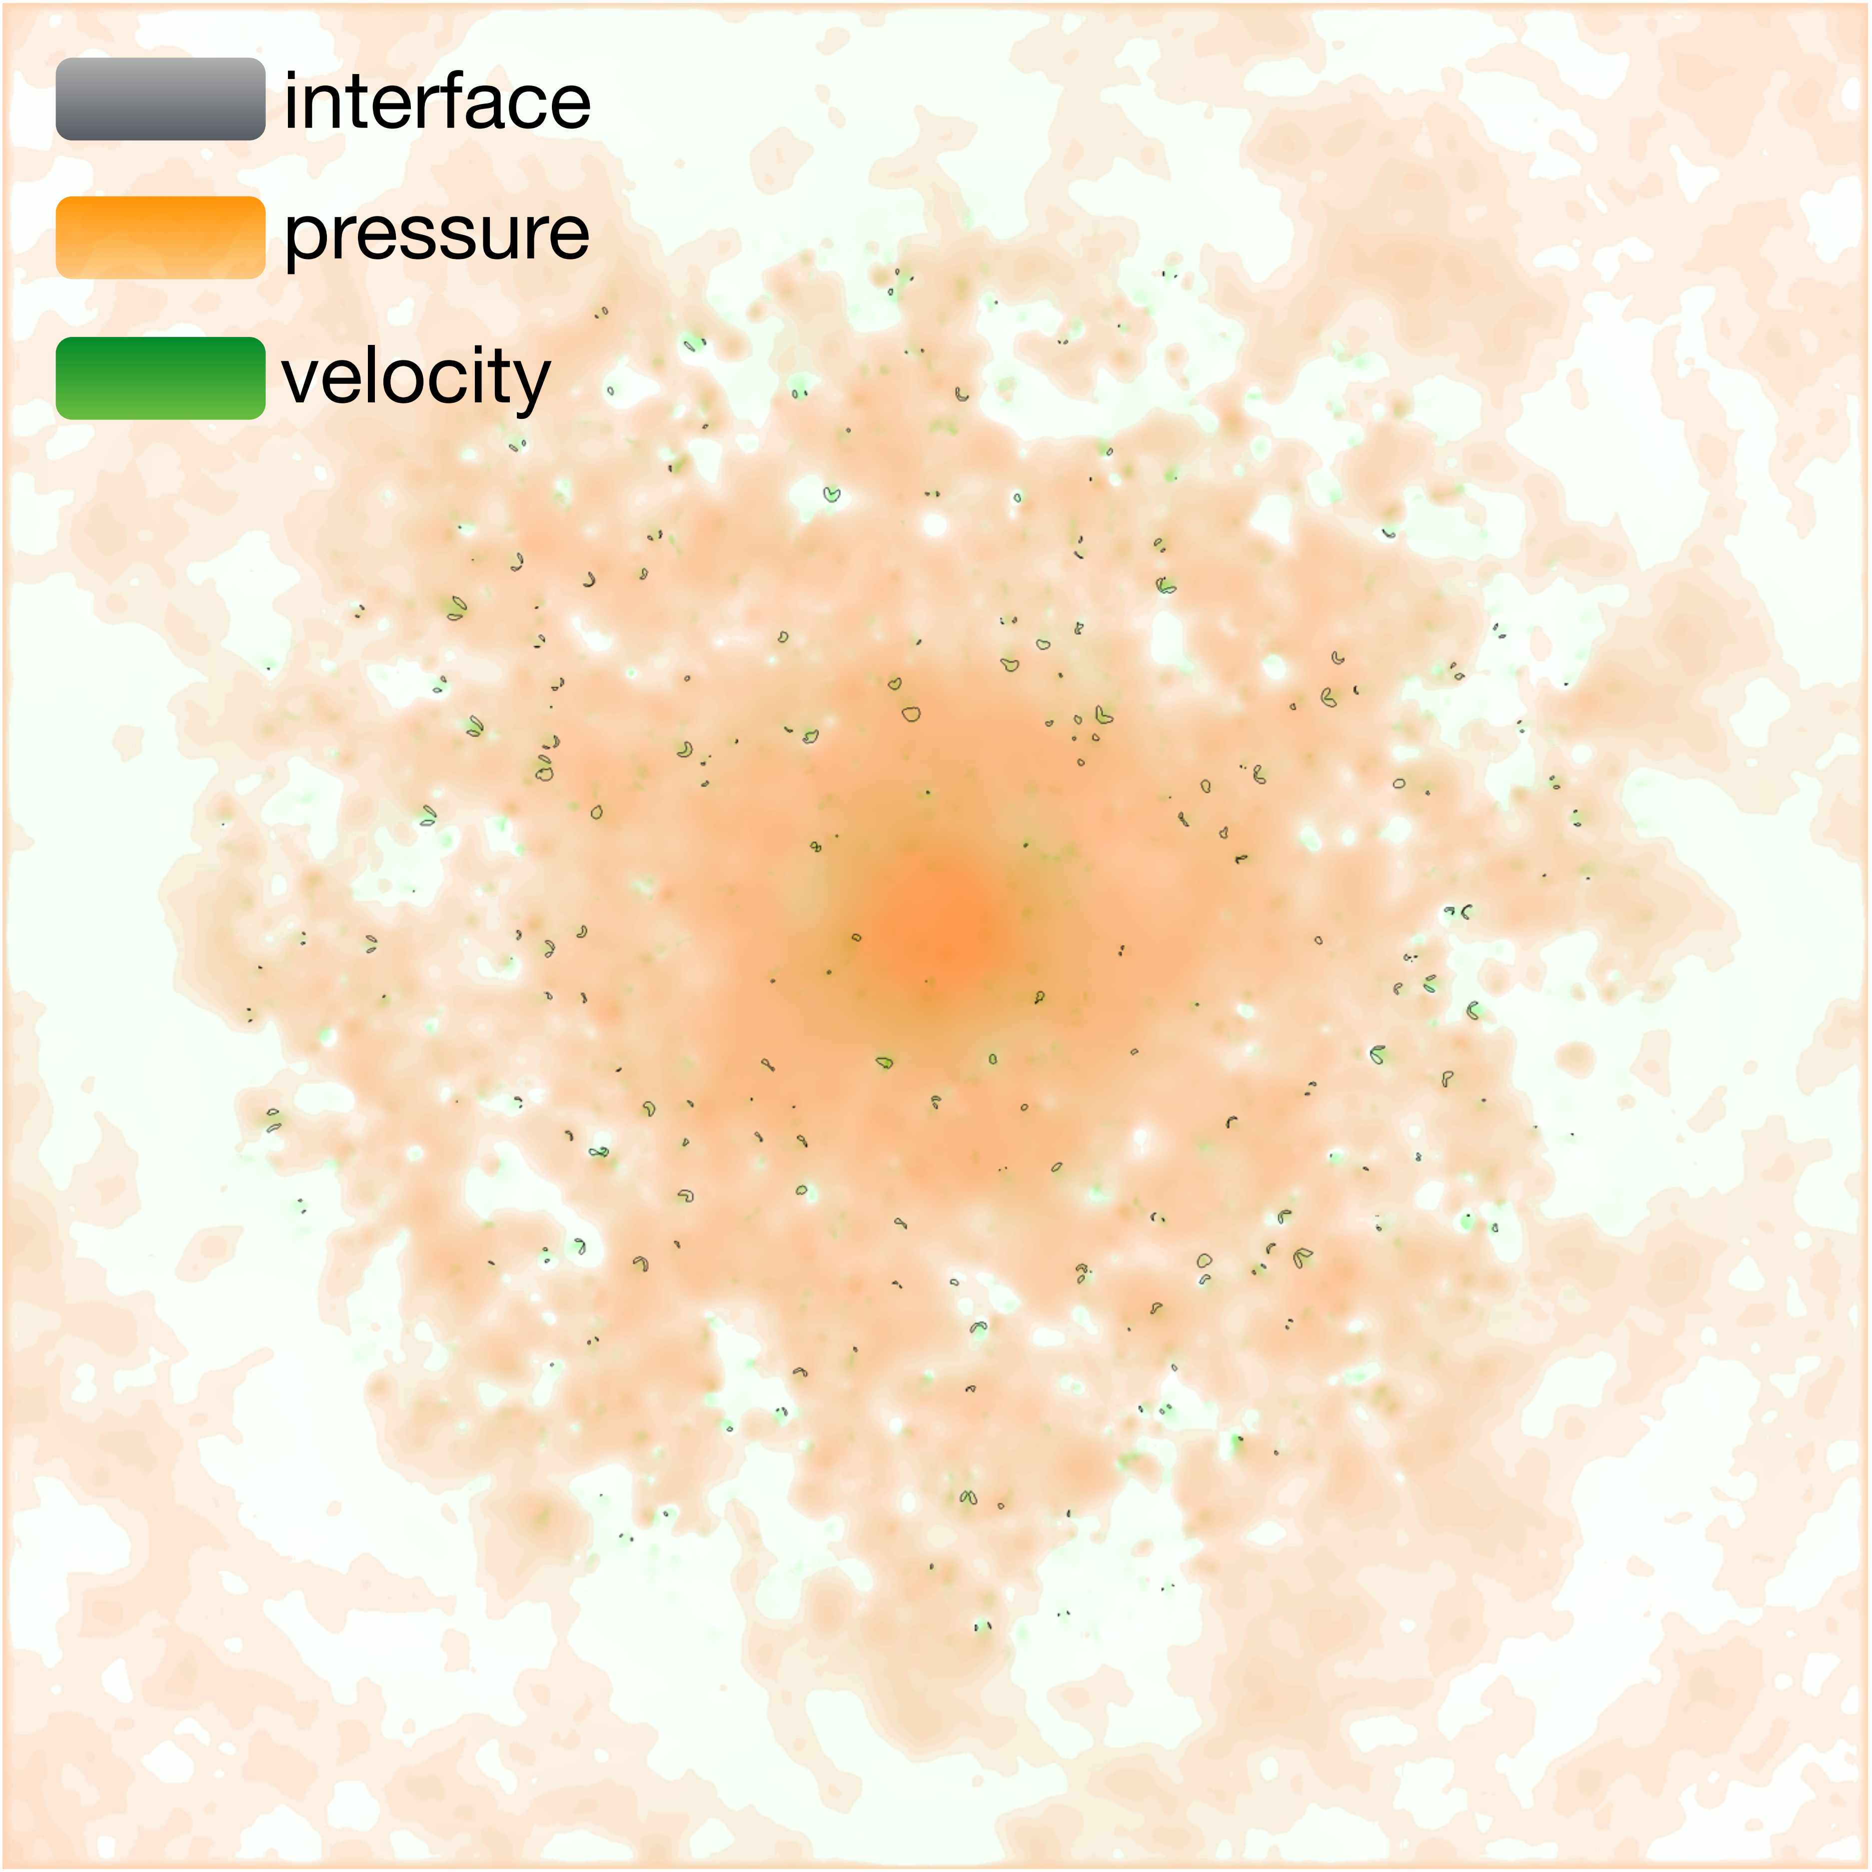

final collapse **certainly** at the cloud center

distance from the cloud center

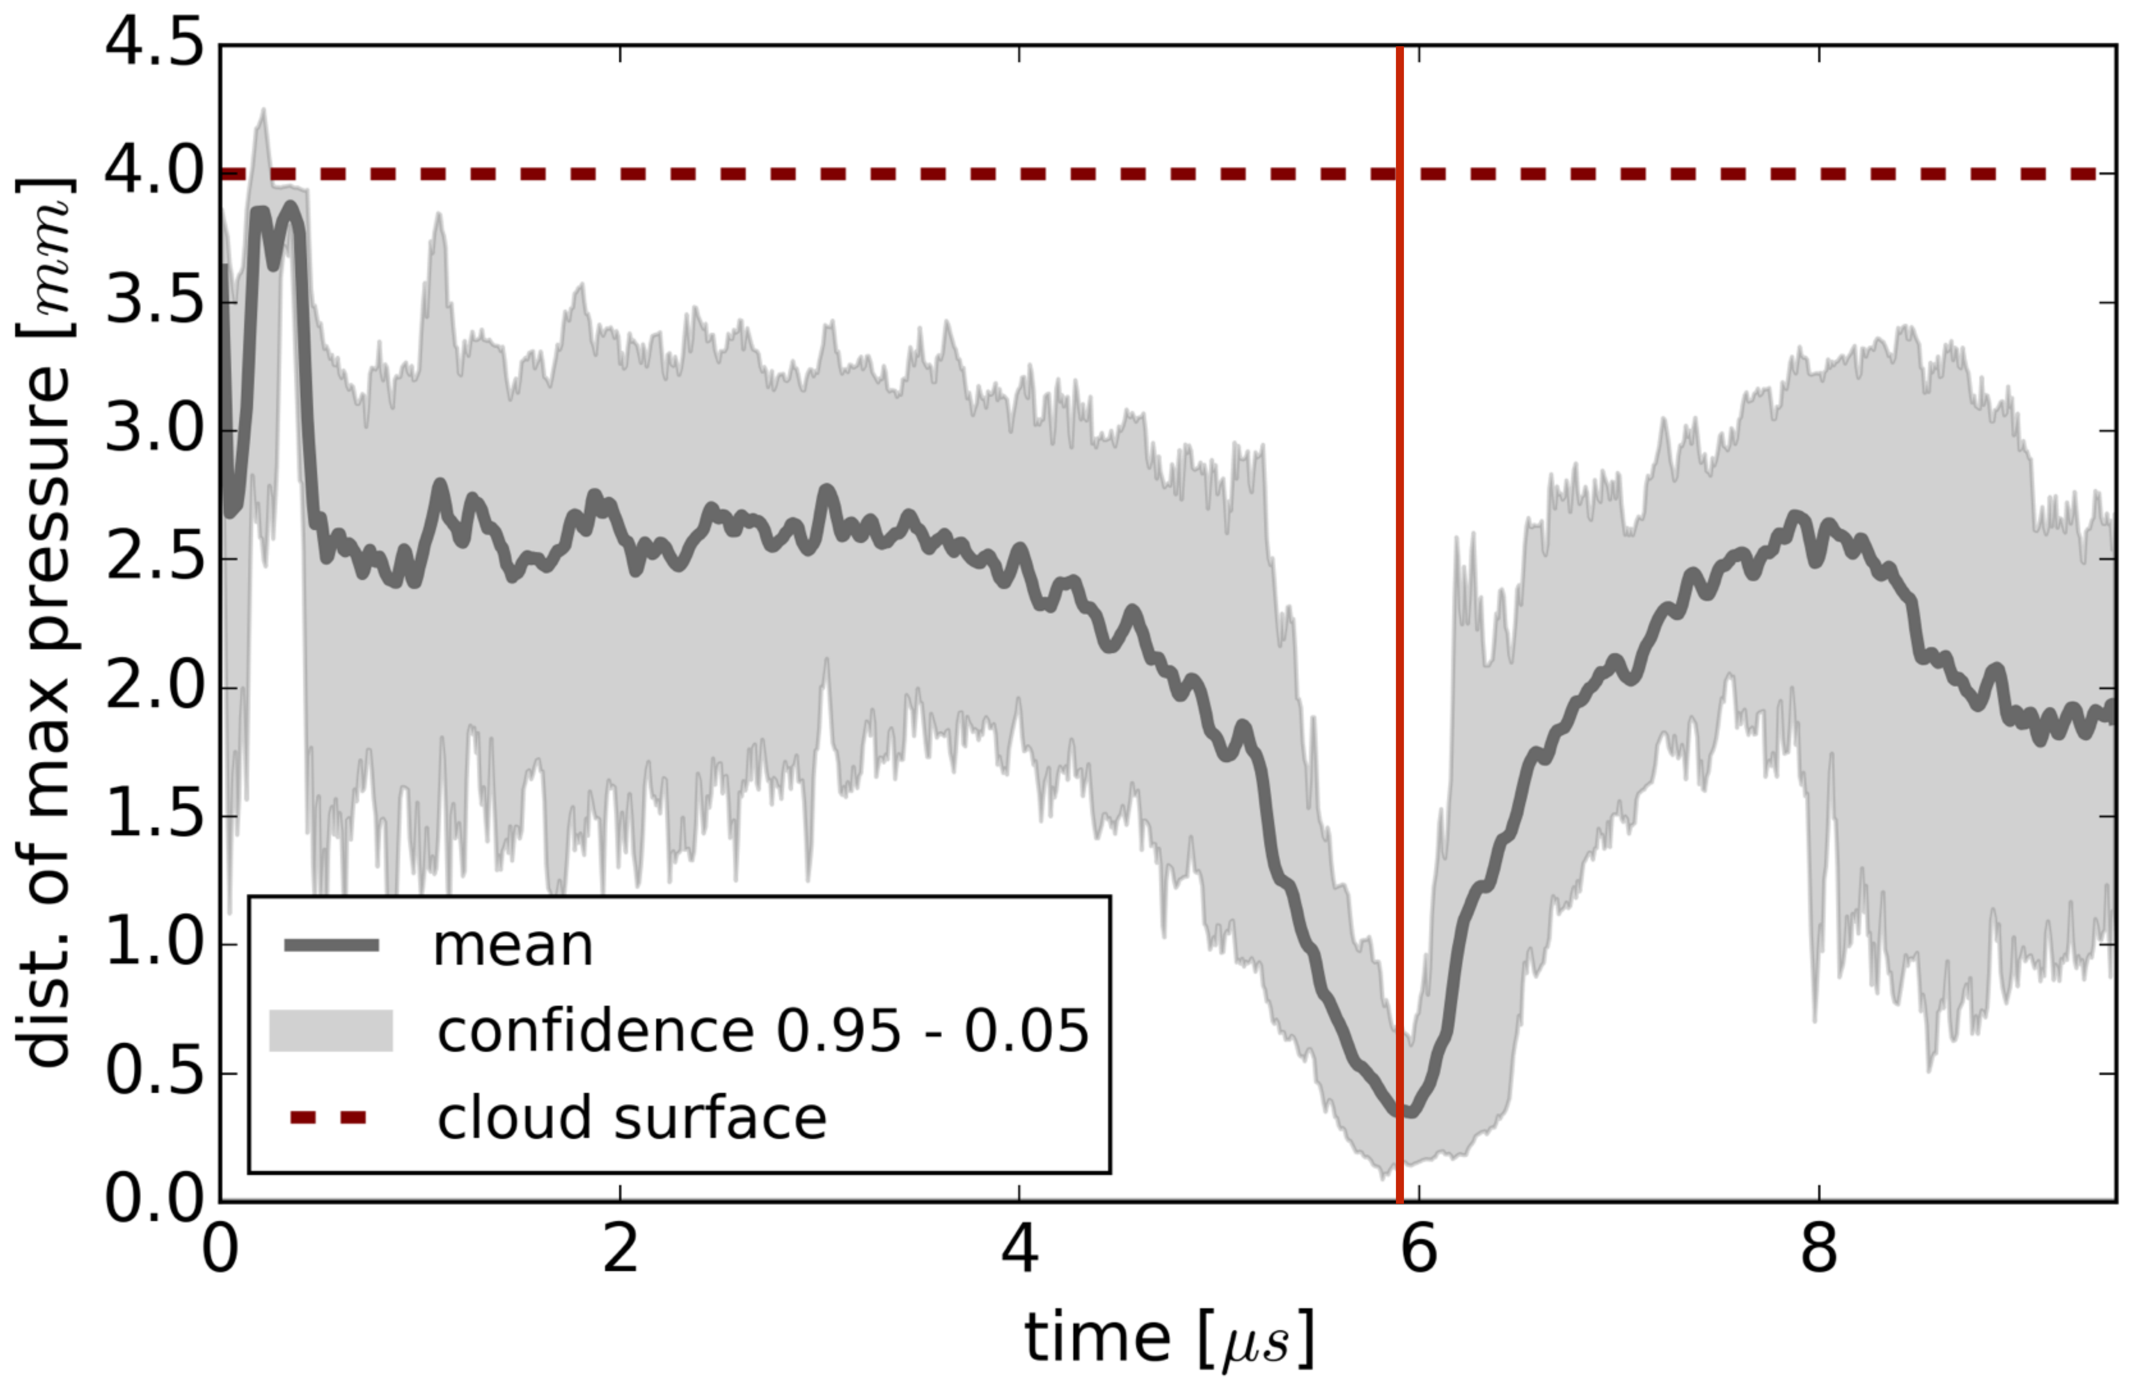

pre-collapse: **wide** confidence interval

collapse: **narrow** confidence interval

Observed **limitations** of MLMC
in cloud cavitation collapse

Observed limitations of MLMC

Accuracy requirements limit the number of coarsening levels

semi-transparent lines correspond to the **same sample** (realization) but on a **coarser resolution**

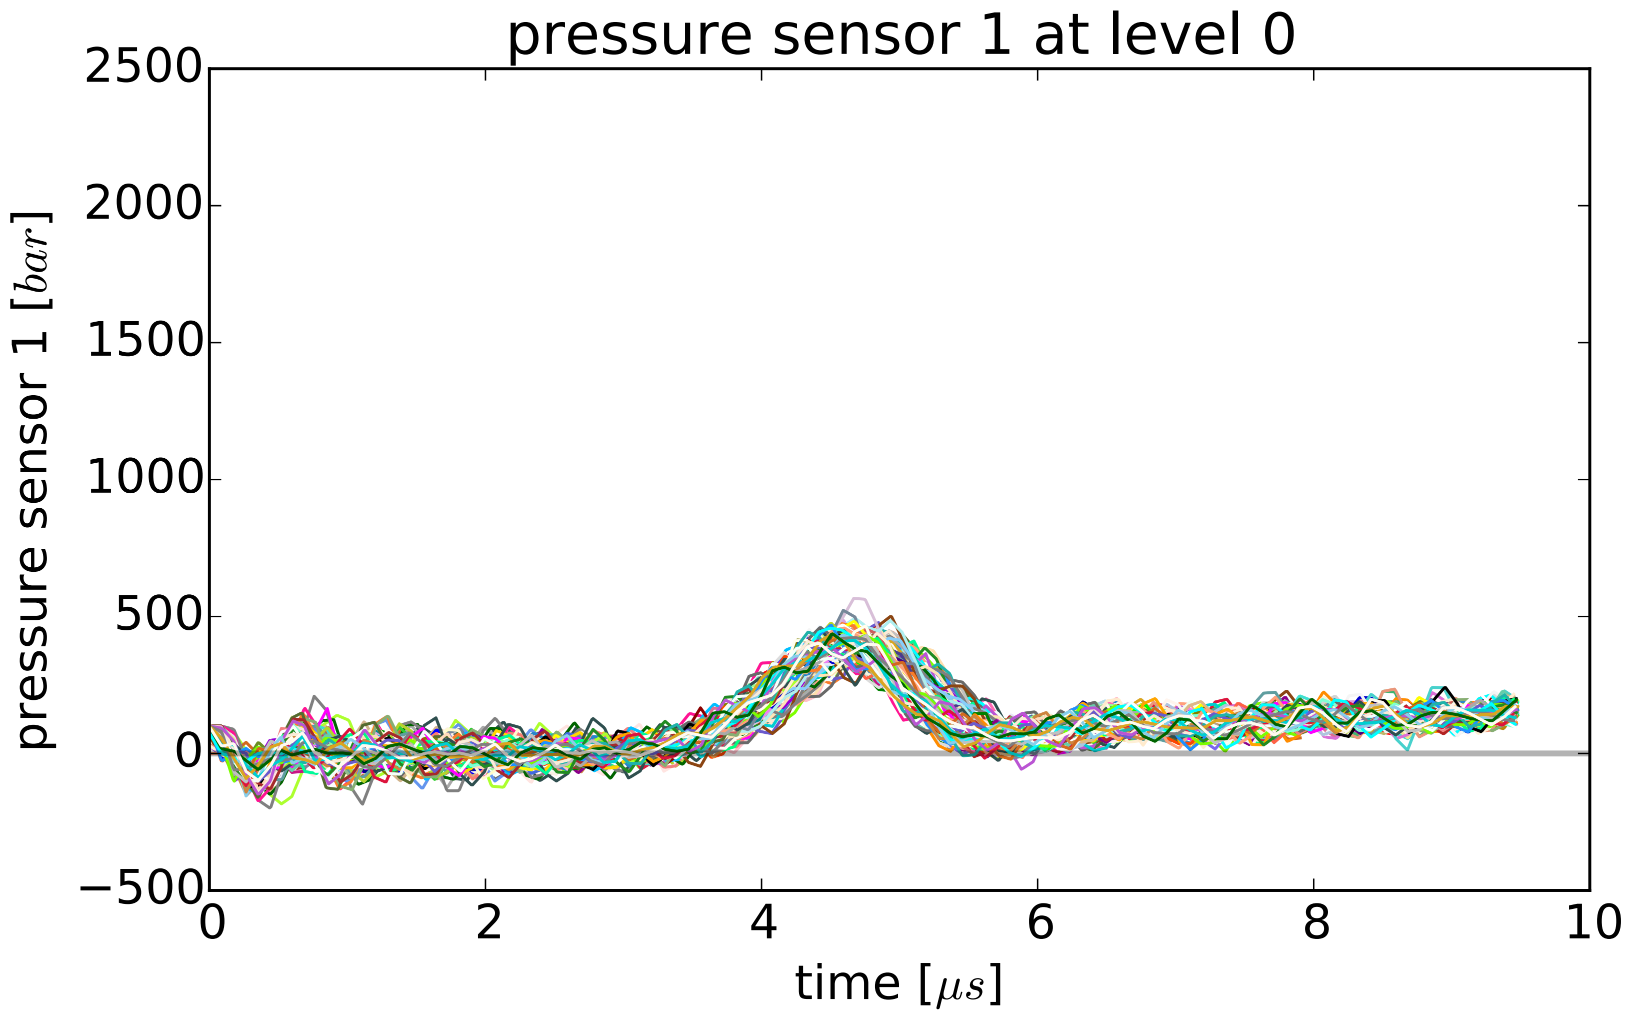

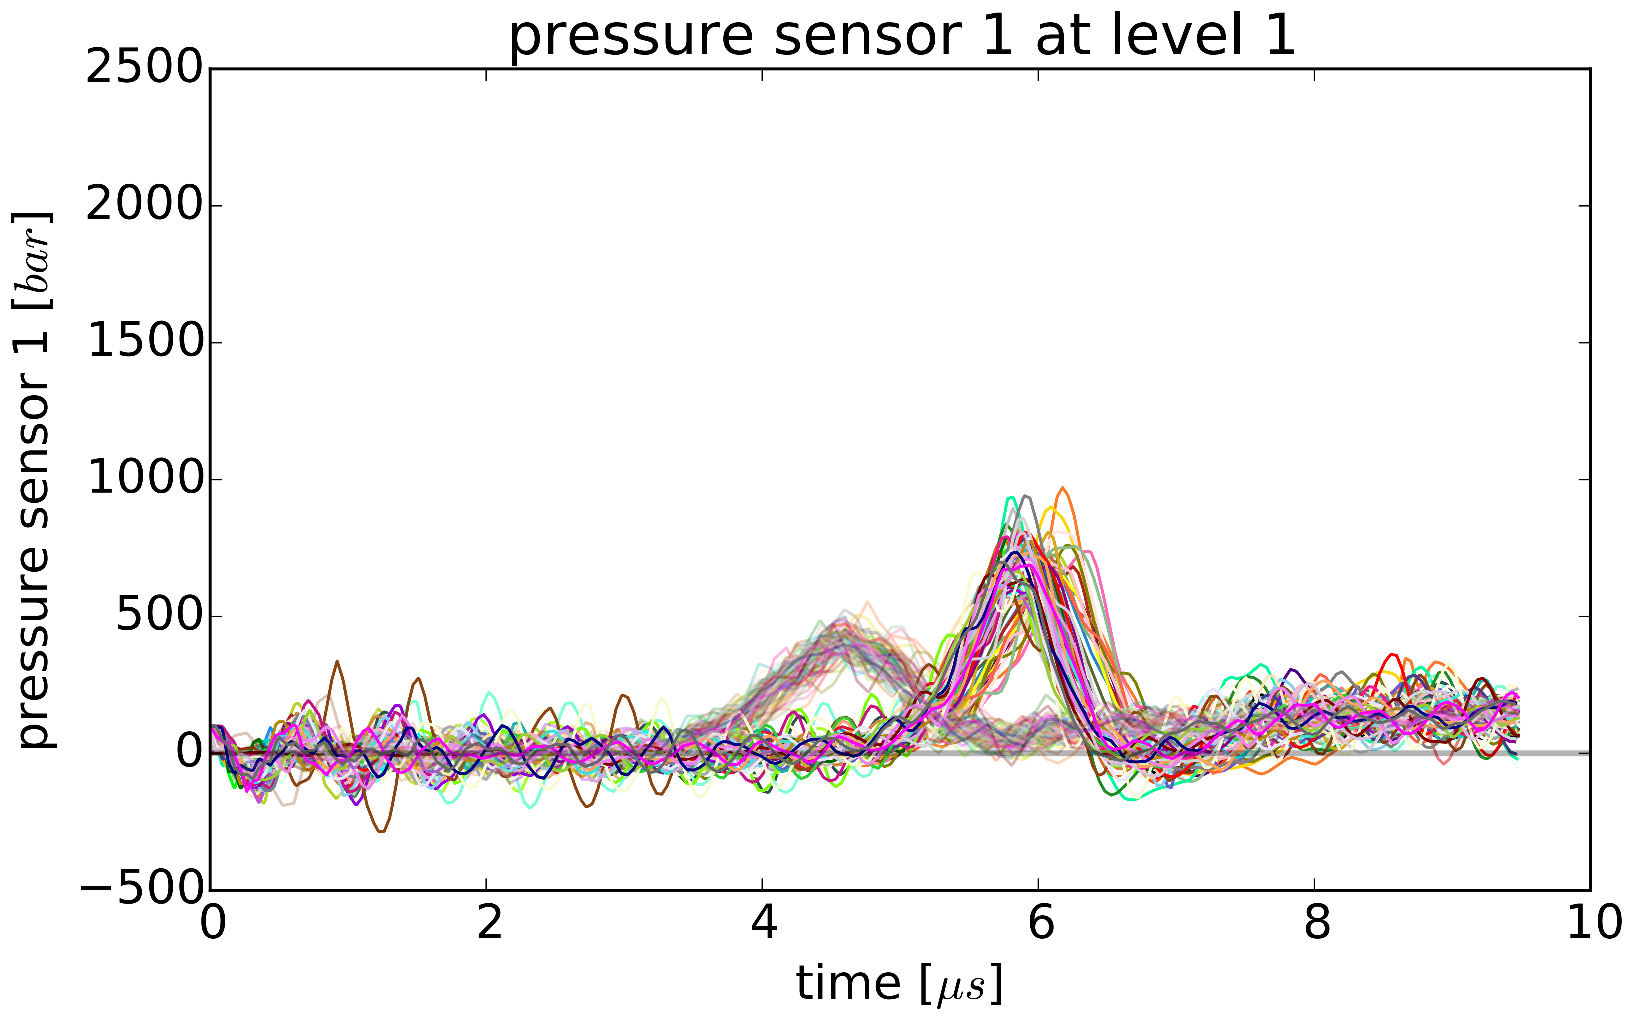

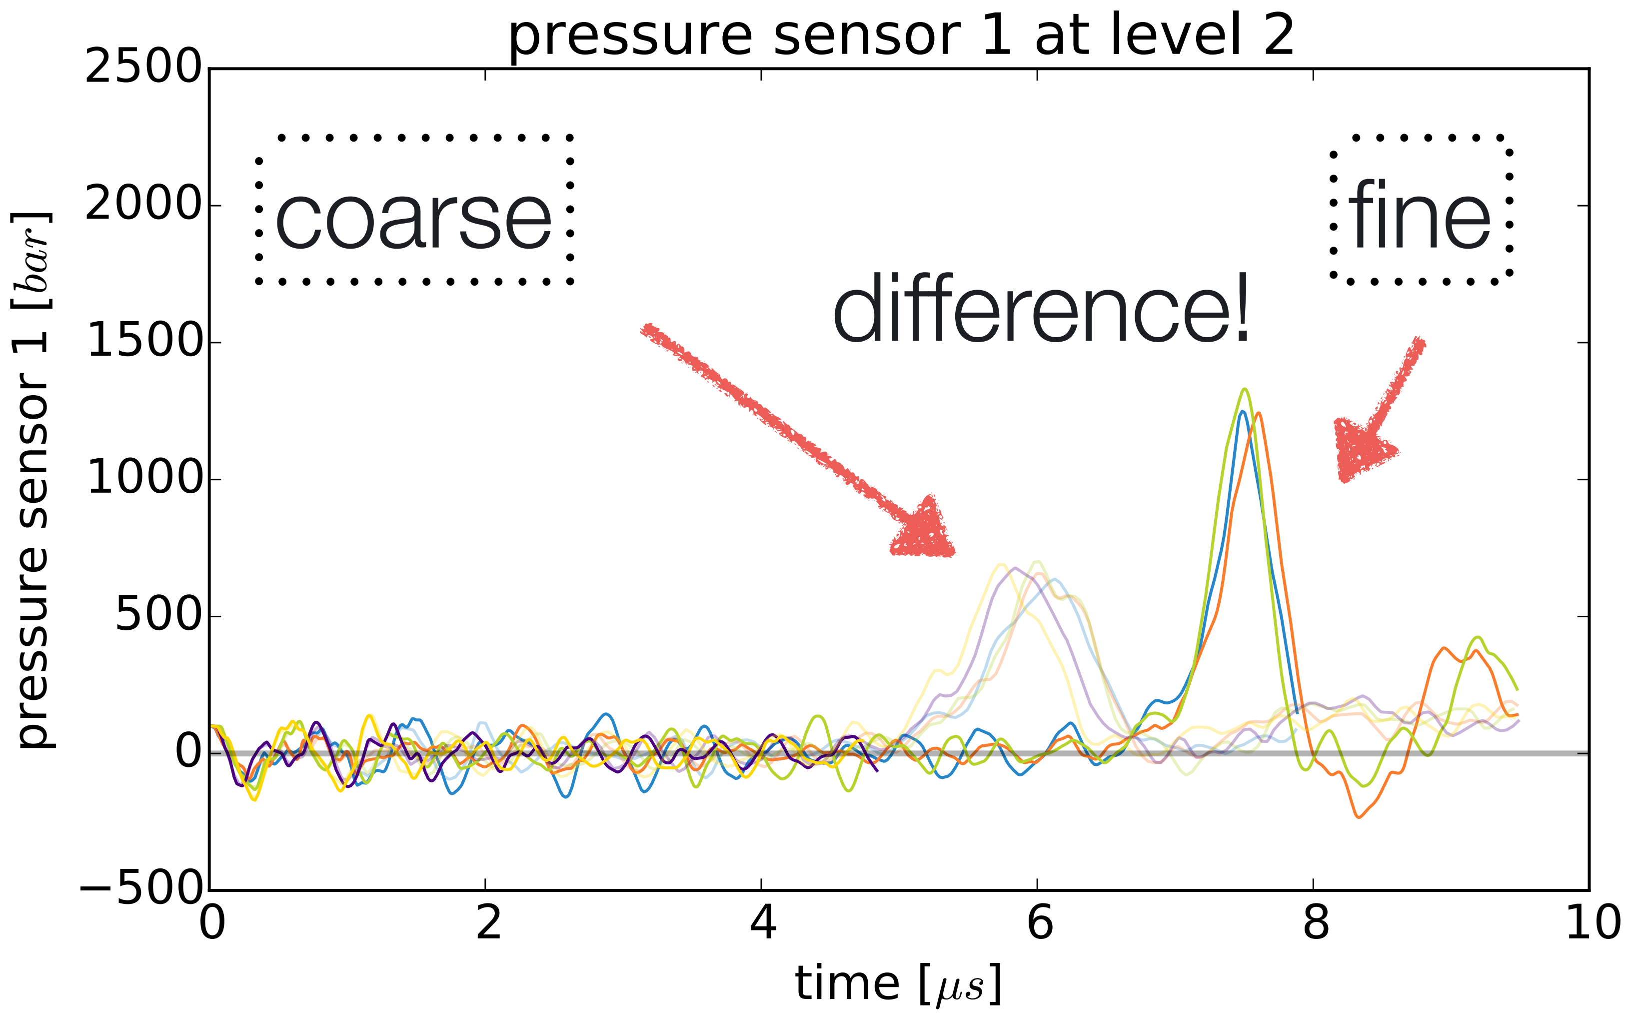

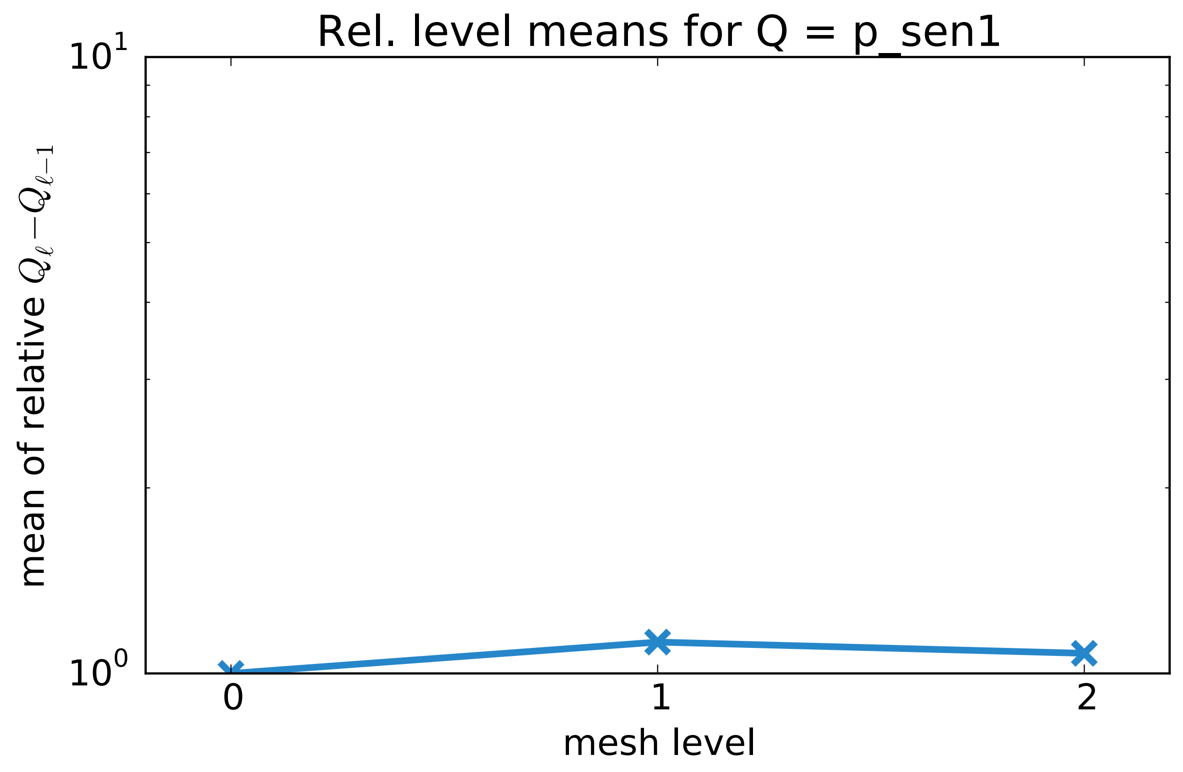

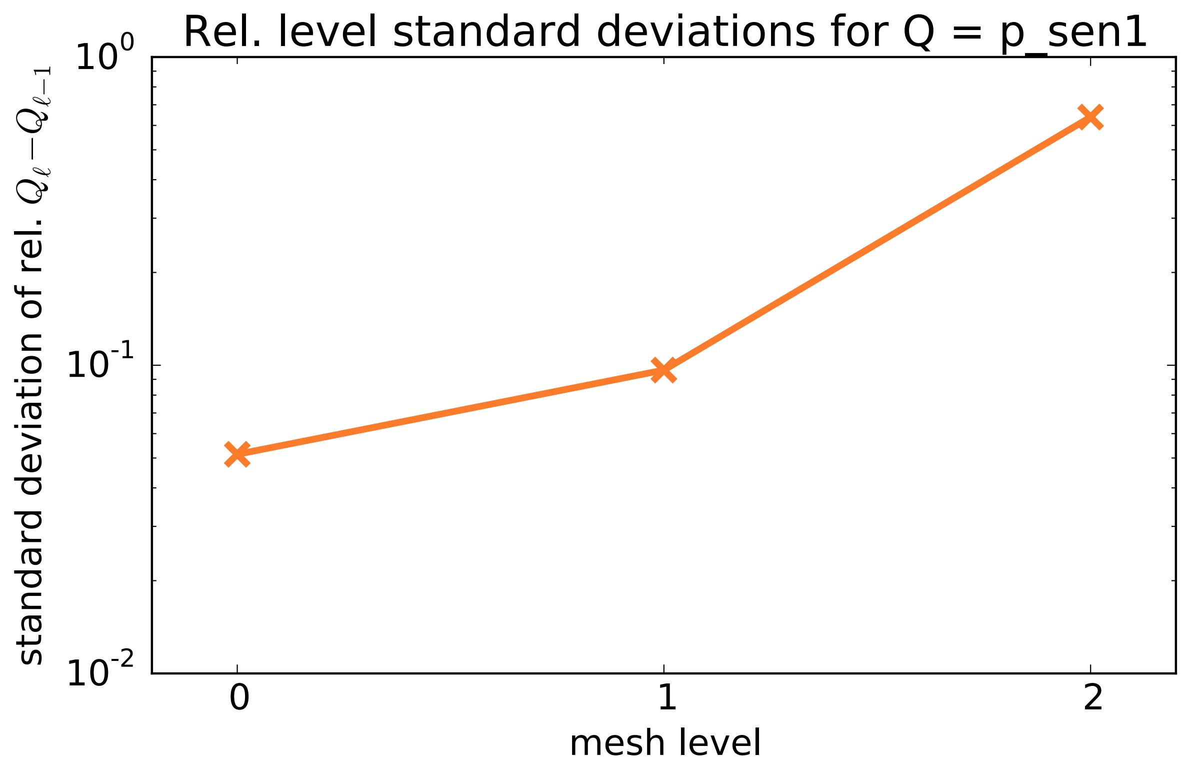

Differences between two consecutive levels do **NOT** decrease

NO proper convergence for increasing resolutions

Summary and outlook

- **Confidence intervals** estimation for local integral **quantities of interest**
  - multiple **sensors for pressure**, density, speed of sound, etc.
  - instead of a single value, full empirical **distribution of peak pressures** could be provided

- Uncertainty quantification on the transition from random to focused micro-collapses

- **Fault tolerance**: in case some samples fail, the rest can still be used to assemble estimators

- Outlook
  - use more levels of resolution to increase the efficiency of MLMC (achieve better speedup)
  - investigate the effect of uncertain and inhomogeneous vapor pressures inside cavities
  - investigate the effect of uncertain cloud geometry (e.g. small surface perturbations in a sphere)

HPC resources

CSCS allocation
Project s500
Piz Daint
Cray XC30
42 176 cores
5 272 GPUs
7.8 PFlops
Switzerland

INCITE allocation
Argonne National Labs
Project “CloudPredict”
MIRA
BlueGene/Q
786 432 cores
10 PFlops
United States

PRACE allocation
Jülich Research Center
Project 091
JUQUEEN
BlueGene/Q
458 752 cores
5.9 PFlops
Germany

PRACE allocation
CINECA
Project 09_2376
FERMI
BlueGene/Q
163 840 cores
2.1 PFlops
Italy

Thanks to

ETH zürich

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

Argonne NATIONAL LABORATORY

FNSNF
FONDS NATIONAL SUISSE
SCHWEIZERISCHER NATIONALFONDS
FONDO NAZIONALE SVIZZERO
SWISS NATIONAL SCIENCE FOUNDATION

U.S. DEPARTMENT OF ENERGY
INCITE
LEADERSHIP COMPUTING

CINECA

IBM

PRACE

JÜLICH
FORSCHUNGSZENTRUM

THANK YOU

Index

Acceleration of UQ simulations using MLMC

Using empirical estimators for variances and measurements of computational work

- Speedup of the MLMC agains plain MC can be estimated as follows

$$\text{speedup} = \boxed{\left(\frac{\mathbb{V}[q_L]}{\varepsilon^2}\text{Work}_L\right)} \Big/ \boxed{\left(M_0 \text{Work}_0 + \sum_{\ell=1}^{L} M_\ell \left(\text{Work}_\ell + \text{Work}_{\ell-1}\right)\right)}.$$

computational work for MC

computational work for MLMC