



Leading High-Performance Computational Software and Multi-domain Architecture for Multi-Physics Modeling Techniques

Candidate Interview for Staff Computational Modeling
Software Developer, Lam Research

Mohammed (Moe) Kamel

April 22nd, 2025

LAM Research Waltham MA

INTRODUCTION

- ✓ Senior R&D Engineer at Ansys – Fluent HPC & Parallel Team
- ✓ 15+ years in high-performance simulation software, specializing in CFD, HPC, and multiphysics
- ✓ PhD in Aerospace Engineering with deep expertise in modeling & simulation, numerical methods and distributed solvers
- ✓ Proven impact in solver architecture, domain decomposition, Heterogenous Parallel CPU/GPU enablement
 - Passion for building performant, robust simulation platforms that scale across multi-domain, multi-physics semiconductor use cases
 - Strategic focus on exploring new frontiers of integrating AI/ML with multiphysics modeling and simulation to have AI/ML-augmented solvers



CORE PROFICIENCIES AND RESEARCH SCOPE

Domain	Highlights
Programming	C/C++, FORTRAN, MPI, CUDA, OpenMP, Scheme, Python
HPC Infrastructure	Domain decomposition, Solver Parallelization, Distributed Memory Design, Heterogenous CPU/GPU computation
Multiphysics Co-Simulation	Turbulent High-speed Flows; Conjugate Heat Transfer; Aero-Optics; Chemically Reacting Flows; Fluid/Structure Interaction; Particle Tracking
Computational Geometry	Mesh morphing, smoothing & dynamic adaption; Geometrical-aware Partitioning; Many-to-many mapping
Numerical Algorithms	AMG & Krylov Linear Solver; Preconditioning; Explicit/Implicit hybrid approaches; Ray tracing; etc. ...
Software Engineering	End-to-end contributor: infrastructure, APIs, integration, CI/CD testing with QA and customer support in industrial-scale CAE software

AGENDA

Meshing & Domain Decomposition

Model-Weighted Partitioning

Geometric-Aware Partitioning

Parallel Mesh Morphing


Multidomain Architecture

Particle Tracking Using DPM

Conjugate Heat Transfer

Aero-Optics

Reduced-Data Postprocessing






MODEL-BASED DOMAIN DECOMPOSITION


DOMAIN DECOMPOSITION

MODEL-WEIGHTED PARTITIONING

- **ParMeTiS** is used in the domain to achieve minimal deviation in cell and face counts across all partitions and to minimize edge-based inter-partition communication.
- Multi-constraint partitions incorporate user-defined, physics-based controls to achieve optimal load balancing when simulating multi-physics flows.
- This approach is applied in multiphysics computations such as Discrete Phase Model (DPM), Conjugate Heat Transfer (CHT), Volume of Fluid (VOF) In-Situ Adaptive Tabulation (ISAT), and Finite Rate Chemistry.
- Model-based weight may need to be filtered for better quality
- Infrastructure of data exchangers is maintain & scaled-up for inter-partitions communications





A dual graph with vertex weight vectors of size two is constructed from the particle-in-cells mesh (Ref. 1)



Using flame sensor for stiff-chemistry weight partitioning







GEOMETRICAL-AWARE PARTITIONING

- Cells sharing stretched faces are encapsulated to enhance AMG local coarsening and enhance scalability.
- Encapsulation are also applied to prevent partition interfaces to cross through poor-quality mesh elements
- Encapsulation are also applied to prevent partition interfaces to cross through poor-quality mesh elements





COMPUTATIONAL GEOMETRY & MESHING

ACCOMPLISHMENT IN PARALLEL MESH MORPHING



Speedup compared to Fluent 2019R1 default settings



Impeller

Speedup compared to Fluent 2019R1 default settings




MULTIDOMAIN ARCHITECTURE FOR MULTIPHYSICS CO-SIMULATIONS

MULTIDOMAIN ARCHITECTURE (MDA)


DISCRETE PARTICLE METHOD (DPM) DOMAIN

- In a single domain approach: Both flow and DPM calculations share the same CFD computational domain
- In multidomain approach: DPM calculations are performed in a distinct second domain (DPM 2nd Domain)
- **In DPM 2nd Domain:**
 - Cells containing particles are evenly distributed among machines and stored in the **primary ranks**
 - Ensuring that each machine receives a fair share of the particle tracking load
 - Maximize the performance boost using MPI/OpenMP hybrid approach
 - Concentrating particles in one partition per machine
 - Avoid any overhead of small threads
 - It increases memory usage by approximately 50% due to additional data exchange between two domains



DISCRETE PARTICLE METHOD (DPM) DOMAIN

LM Combustor, Single injector rig, Experiments by Joshi et. al, Particle evaporation, 246K cells, 1M particles





Hybrid MPI/OpenMP approach is 4X faster than using MPI only
DPM second domain adds 2X speedup


MULTIDOMAIN ARCHITECTURE (MDA)

GENERAL DATA STRUCTURE


- Handles multiple domains with different meshes, models, settings and parameters of an intrinsic co-simulation
- Provides generic API's for developers to construct domain couplings and exchangers between multiple domains within customized co-simulation
- Each Domain is encapsulated with all relevant characteristics to have robust handling of different meshes
- Has a swapping mechanism to switch between different meshes for computation, post-processing and User Interface purposes



MULTIDOMAIN ARCHITECTURE (MDA) CONJUGATE HEAT TRANSFER (CHT)



Fully thermally coupled at interface



Conjugate Heat Transfer

Fully Coupled CHT

Fluid and solid in a single system

Coupled system of equations
(Continuity of heat flux and
temperature treated implicitly)

Convergence more efficient due to
“coupled” system of equations at
the interface.

Transient CHT:
characteristic time scale of fluid
many orders smaller than solid as
such for solid many more time
steps are needed due to small
time step size dictated by fluid

Loosely Coupled CHT

Decouple fluid and solid systems

Fluid and solid equations solved
separately and coupled at
interface via interpolation of BC


Convergence depends on each
solver and data transfer
convergence

Transient CHT:
Solid and fluid domains can have
their own time steps hence
coupled methods more efficient
than conjugate methods.


APPROACHES OF LOOSELY COUPLED CHT

A - Instantaneous-Implicit Loosely Coupled CHT Using Multidomain Architecture


A1 - Fully thermally coupled at interface
At every fluid time step



A2 Loosely Coupling
(No exchange, **most of the simulation time**)




A3 - Fully/Implicit thermally coupled at interface
At coupling frequency




B - Time-Averaged Explicit Loosely Coupled CHT Using Multidomain Architecture


B1 - Fully thermally coupled at interface
At every fluid time step



B2 Loosely Coupling
(No exchange, **most of the simulation time**)




B3 - Explicit thermally coupled at interface
At coupling frequency




APPROACHES OF LOOSELY COUPLED CHT

- **Efficient and Fully Parallelized Direct data exchange** through maps across the solid/fluid Interface
 - Supports **Non-Conformal & Conformal** Solid-Fluid Interfaces
 - **No need to re-create the exchange mapping** across the solid-fluid interface from scratch
 - **Use the already created interface data** for local mapping between fluid and solid cells



Mapping across face and its shadow
of Conformal Interface




Mapping across face and its shadow of **intersected faces**
of Non-conformal Interface

MULTIDOMAIN ARCHITECTURE (MDA) PERFORMANCE ENHANCEMENT USING LOOSELY COUPLED CHT


Electric motor Case (7.6 Million cells)

- Solid cell count/Fluid cell count = 0.7
- 15% total speedup in overall computation time
- 1.7X speedup in energy equation computation



Generic Motor (50 Million cells)

- Solid cell count/Fluid cell count = 1
- 10% total speedup in overall computation time
- 2X speedup in energy equation computation





MULTIDOMAIN ARCHITECTURE (MDA)

REDUCED-DATA POSTPROCESSING


- Each domain is partitioned and properly distributed within processes
- MDA Exchanger performs flow data interpolation from the dense CFD mesh to octree-based post-processing mesh
 - Different meshes / Different partition topologies Data Exchanger
- Cell-locator mapping is used to identify **donor elements (cells and faces)** of the CFD domain for each of the **receptor elements** in the post-processing domain
- Data Interpolation is gradient-based 2nd accurate

$$\Phi_{receptor} = \Phi_{donor} + \nabla \Phi_{donor} \cdot (\vec{x}_{receptor} - \vec{x}_{donor})$$



MULTIDOMAIN ARCHITECTURE (MDA)

AERO-OPTICS



Flow direction



y
x

Kamel et al. AIAA Paper 2016-1462

Aberrated wavefront

Instantaneous density field

Planar wavefront

- The **aero-optical** phenomenon is concerned with the aberration effects of **compressible turbulent flows** induced by solid surfaces near the projection or viewing optical aperture


AERO-OPTICAL PHENOMENON (ASSUMPTIONS)

The propagation of electromagnetic waves through turbulent medium of heterogenous index of refraction is discussed theoretically by [Monin and Yaglom](#)^[1]:

- No charges and No electric current in the medium
- The distance propagation $\ll$ distance of turbulent disturbance propagation with light speed $L \ll \frac{2\pi}{\omega_{turb}} c$
 - The turbulent flow is **frozen** with respect to electromagnetic waves propagations
 - Amplitudes is **not a function of time** (at a given instant of flowfield)
 - It can be treated as monochromatic waves of single frequency
- Kolmogorov turbulence length scale is much larger than optical wavelength $\rightarrow$ **no depolarization effect**
- The fluctuations of index of refraction is small enough **to distort only the phases** of electromagnetic waves not their amplitudes
- The optical beam can be represented by **collimated planar wavefront** in the **near-aperture region**.

MULTIDOMAIN ARCHITECTURE (MDA) BUILDING THE OPTICAL GRID

- Each surface is limited for a single beam of arbitrary propagation vector
- Flat and curved window surfaces are supported for a beam creation
- Each face elements is considered as a root of a ray gridline
- Donor CFD cells are marked to create beam ray points to match the resolution



MULTIDOMAIN ARCHITECTURE (MDA)

CO-CURRENT AERO-OPTICAL SIMULATION

- The instantaneous density values are interpolated from the flow solver mesh onto beam grid(s)

$$\rho_{optical} = \rho_{cv} + (\mathbf{x}_{optical} - \mathbf{x}_{cv}) \cdot \nabla \rho|_{cv}$$

- Inside the beam grid, the index of refraction is calculated using Gladstone Dale relation

$$n(\hat{x}, \hat{y}, \hat{z}, t) = 1 + K_{GD} \rho(\hat{x}, \hat{y}, \hat{z}, t)$$

- The local phase the wavefront in terms of OPL is calculated by integration along the near turbulent field:

$$OPL(\hat{x}, L_{\hat{y}}, \hat{z}, t) = \int_0^{L_{\hat{y}}} n(\hat{x}, \hat{y}, \hat{z}, t) d\hat{y}$$

- Steady Lensing Term (OPD_{mean}) $\overline{OPD} = \overline{OPL} - \langle OPL \rangle$




- Steady-Lensing-Removed OPD (OPD') $OPD' = OPL - \langle OPL \rangle$

- High-Order Terms of OPD (OPD_{HO}) Using Least Square Method to minimize G

$$G = \iint_{A_p} [OPD'(x', z') - C_0(t) - C_x(t)x' - C_z(t)z']^2 dx' dz'$$



$$OPD_{LO}(x', z') = C_0(t) - C_x(t)x' - C_z(t)z'$$

$$OPD_{HO}(x', z') = OPD'(x', z') - OPD_{LO}(x', z')$$




MULTIDOMAIN ARCHITECTURE (MDA)


VALIDATION OF AERO-OPTICAL SIMULATION



Experiment Conditions	
R	2"
M_{in}	0.5
Re_{θ}	8.5×10^3
Re_R	6.3×10^5
θ	120°
δ_{in}	$0.1378 R$



Cylindrical Turret (Gordeyev et al. 2010)



Data Source	$\overline{OPD_{rms}}$	Mesh Size	Error%
	$(\rho_{\infty}/\rho_{SL})DM_{\infty}^2$		
Experimental measurement [1]	1.50×10^{-6}	—	—
Fluent Simulation (Mesh A)	3.86×10^{-6}	3×10^6	157%
Fluent Simulation (Mesh B)	1.98×10^{-6}	25×10^6	32%
Previous Numerical Simulation (Fine Mesh) [2]	1.12×10^{-6}	63×10^6	25%

[1] Gordeyev et al. (AIAA Paper 2010-4492)

[2] Kamel et al. (AIAA Paper 2016-1462), Kamel, M., Univ. of Notre Dame PhD. Diss. 2017



VISION & SUMMARY

FORWARD VISION HYBRID HPC+AI/ML FOR NEXT GENERATION PROCESSING MODELING

- Augment traditional solvers with:
 - Distributed Model Simulation over Heterogenous Parallel Computing
 - Surrogate models for sub-grid closure or fast approximations
 - ML-based preconditioner tuning and solver convergence prediction
 - Reduced-order modeling for parametric studies
- Strategic direction: **data-enhanced physics**, not replacement

SUMMARY

Summary – Readiness to Contribute at Lam

- Experience in large-scale solver development and physics modeling
- Proven architectural leadership in HPC and multiphysics infrastructure
- Strong contributor across domains — engineering rigor + systems thinking
- Ready to translate expertise from aerospace, automotive & energy sectors to semiconductor challenges
- Committed to innovation, collaboration, and real-world impact in advanced manufacturing



THANK YOU

Mohammed (Moe) Kamel

mkamel@alumni.nd.edu

[LinkedIN: mohammedkamel](#)