Towards Petascale LES of Reacting Flow

S. L. Yilmaz*, P. H. Pisciuneri** and P. Givi**
Corresponding author: slyilmaz@pitt.edu

* Center for Simulation and Modeling, University of Pittsburgh, USA.
** Mechanical Engineering and Materials Science, University of Pittsburgh, USA.

Abstract:

The filtered density function (FDF) is considered as one of the most effective means of conducting large eddy simulation (LES) of turbulent reacting flows. The FDF is essentially the counterpart of the probability density function (PDF) methods in Reynolds-averaged Navier-Stokes (RANS) simulations. Alternative formulations of FDF have been introduced in the literature, and these range from FDF closures for a only subset of quantities where the rest is modeled via conventional models, to FDF closures for the joint statistics of all of the relevant physical variables. Common to all FDF models is that the one-point statistics are accounted for in an exact manner; in particular, FDF allows model-free representation of the subgrid effects of highly non-linear chemical kinetics. The scalar FDF is the simplest form of FDF, it considers only the joint statistic of the scalar field, and can be utilized to incorporate any kinetics model in a straightforward manner. This is the most widely used form and is the subject of the current study.

FDF and other methodologies with detailed chemistry computations share a common feature: computational cost. In particular, in its current implementation, the scalar FDF involves concurrent solution of compressible flow dynamics on an Eulerian domain, and the Monte Carlo (MC) simulation of a set of stochastic differential equations (SDEs) over a large number of Lagrangian particles. These particles represent the transient thermo-chemical composition of the fluid. The SDEs have three main components: transport, mixing, and reaction; and these are advanced in fractional steps using splitting schemes. The reaction step involves the computation of detailed kinetics and the solution of a stiff ODE system for each particle; and accounts for 95% of the computation time in a typical simulation.

A common feature of such computations is the computational load imbalance due to the nature of chemical kinetics and inhomogeneity of flow composition throughout the simulation domain and time. At any instant of the simulation, different regions of the flame undergo different stages of combustion. For some regions, the integration of chemistry sub-step can be done very quickly (e.g. in cold air or fuel with no mixing), but for some others implicit integration of a highly stiff ODE system is required (e.g. extinction/reignition regions, regions of high level of fuel/oxidizer mixing, flame fronts, etc.).

We introduce an adaptive domain decomposition strategy which addresses the highly dynamic and transient nature of detailed chemistry computations in the context of a structured Navier

Stokes solver coupled with the Lagrangian MC solver (see Figure). The implementation allows for tremendous improvements in scalability, and is the key enabler of petascale computations. The paper describes the implementation details of the hybrid solver and the domain decomposition strategy, along with presentation of sample petascale LES of realistic flames. We believe that the dynamic parallelization is not only useful, but also is essential for such computations in order to achieve *petascalability*; and that the lessons learned and the strategies presented in this work are applicable to all PDF and FDF methodologies, as well as other simulation techniques that involve detailed chemistry computations.

Keywords: Turbulence, Combustion, Large Eddy Simulation, Filtered Density Function, Petascale Parallelization, Chemical Kinetics.

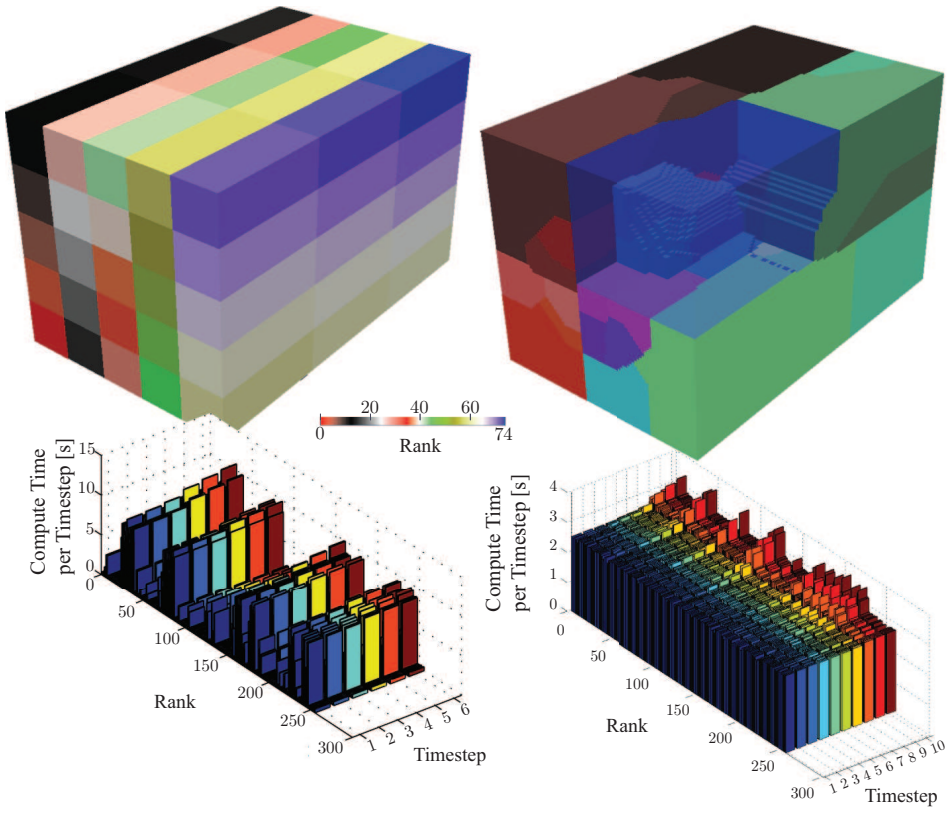

Figure 1: Domain topology with (a) the uniform decomposition, and (b) adaptive decomposition. Non-idle CPU times per time-step for each rank for subsequent time-steps with (c) uniform decomposition, and (d) adaptive irregular decomposition. Petascale computations contain 10,000s of CPU cores. What is shown here for clarity, is a benchmark case with small number of CPUs.