Computers and Chemical Engineering 60 (2014) 197–212

Contents lists available at ScienceDirect

Computers and Chemical Engineering

journal homepage: www.elsevier.com/locate/compchemeng

A fully coupled, parallel approach for the post-processing of CFD data through reactor network analysis

A. Stagni, A. Cuoci*, A. Frassoldati, T. Faravelli, E. Ranzi

Department of Chemistry, Materials, and Chemical Engineering, Politecnico di Milano, P.zza Leonardo da Vinci 32, 20133 Milano, Italy

ARTICLE INFO

Article history:
Received 12 March 2013
Received in revised form 1 September 2013
Accepted 4 September 2013
Available online 14 September 2013

Keywords:
Turbulent combustion
NO_x
Non-linear systems
Kinetic post-processing
MPI
Detailed kinetics

ABSTRACT

In this paper we describe and apply a parallel code, named KPPSMOKE, for the prediction of pollutant emissions from combustion devices operating in turbulent conditions. The approach is based on the kinetic post-processing of CFD simulations, which are transformed into equivalent networks of perfectly stirred reactors and solved using a detailed kinetic mechanism (hundreds of species). The numerical algorithm is based on a fully-coupled technique, in which the highly non-linear mass balance equations are solved together, by alternating different resolution methods in order to ensure high accuracy and fast convergence. As a result of KPPSMOKE parallel structure, large reactor networks characterizing industrial devices (10^5–10^6 reactors) can be solved in reasonable times (∼hours).

The accuracy and the reliability of the algorithm was demonstrated on a lab-scale burner and on a full-scale industrial device, i.e. a combustor for aircrafts. The numerical performance was also assessed in terms of parallel efficiency and speedup.

1. Introduction

The emissions of pollutant species like nitrogen oxides (NO_x), CO, unburned hydrocarbons and soot particles from combustion devices are regulated from a legislative standpoint in Europe (2010) as well as in the rest of the world. As a result, the design of devices able to reduce the formation of such compounds is vitally important for the concerned industries, and the ability to predict their formation through computational tools with the required accuracy (which is usually of the order of ppm) is then essential.

Such an accuracy can be ensured only by the combination of two factors: an accurate fluid dynamic description of the equipment and a detailed kinetic mechanism, which includes all the elementary chemical reactions involved in the different reaction paths. A full coupling of them, based on a simultaneous resolution of species, momentum and energy equations, is unfeasible because of their mutual dependence and their high degree of non-linearity. This results in an excessive computational load, which can be barely ensured only by modern supercomputers.

In particular, when dealing with turbulent combustion, several assumptions must be made to face the problem: considering a RANS approach (Poinsot & Veynante, 2005) to combustion phenomena (currently the most used in industrial cases), either the turbulent flow or the kinetic mechanism must be simplified in order to reach a reasonable size of the problem and consequently a reasonable computational time (Fichet, Kanniche, Plion, & Gicquel, 2010).

In the last fifteen years, a new family of approaches was conceived, which keeps accuracy on both CFD and kinetics by splitting the resolution procedure into more steps: named post processing techniques, they allow to ease the resolution of the overall numerical problem by exploiting the peculiarities of the system under investigation.

Ehrhardt et al. (1998) first introduced a method of this kind, based on the following key steps:

(1) First, they used a CFD code with a fine grid to evaluate fluid dynamic fields with global kinetic schemes;
(2) Then, through proper algorithms, the computational domain was split into volume elements, whose size was much larger than those of the original CFD cells, thus avoiding an excessive computational load. Each of these volumes was considered as a reactor, which exchanged fixed flows with the neighbors. In this way a network of reactors was created;
(3) Finally, the created reactor network was solved with the assumption of the volume elements being perfectly stirred reactors.

This strategy was then adopted by Frassoldati, Frigerio, Colombo, Inzoli, & Faravelli (2005) after being refined by Faravelli et al. (2001) and Falcitelli, Pasini, Rossi, & Tognotti (2002), who respectively extended their applicability to liquids and proposed a general algorithm to build a reactor network from the CFD field,

* Corresponding author. Tel.: +39 02 2399 3283; fax: +39 02 7063 8173.
E-mail address: alberto.cuoci@polimi.it (A. Cuoci).

0098-1354/$ – see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.compchemeng.2013.09.002

Nomenclature

Roman symbols

m	reactor mass (kg)
$\dot{m}$	convective flow (kg/s)
J	diffusion flux (kg/(m^2 s))
S	surface (m^2)
V	reactor volume (m^3)
D	diffusion coefficient (m^2/s)
x	mole fraction (–)
Sc	Schmidt number
N_C	number of reactors
N_S	number of species
t	execution time (s)

Greek symbols

ω	mass fraction (–)
μ	dynamic viscosity (kg/(m s))
ρ	density (kg/m^3)
∇x	gradient of x ($1/x$)

Subscripts

i	species
j	reactor
N_S	turbulent property
s	serial quantity
p	parallel quantity

with positive outcomes on the accuracy of the estimation of pollutants, if compared to experimental data.

A similar approach was introduced by Skjøth-Rasmussen et al. (2004), who solved larger cases without simplifying the computational grid, i.e. considering each computational cell of the CFD grid as a perfectly stirred reactor. They decoupled the material balance from the energy balance by importing and keeping fixed the temperature field from the first step (CFD simulation), thus significantly decreasing the computational load.

A major change in this technique was then introduced by Cuoci, Frassoldati, Buzzi-Ferraris, Faravelli, & Ranzi (2007) and Cuoci et al. (2013), who built up a software, called *Kinetic Post Processor* (KPP), able to post-process CFD data and solve the network through an efficient alternation of local and global solution methods. In particular, they fully exploited the physical structure of the problem in the resolution procedure and translated it numerically into a large, sparse non-linear system of equations. Monaghan et al. (2012) successfully applied this approach to study the pathways of formation of NO and NO_2 in a piloted methane–air diffusion flame. To this purpose, starting from a CFD simulation, they created a reactor network made up of ~1100 perfectly stirred reactors, which was post-processed through the KPP, thus obtaining predictions in very good agreement with experimental data. Van Goethem, Risseeuw, Barendregt, & Frassoldati (2010) effectively exploited the same strategy to study the formation of NO_x in industrial furnaces for the production of ethylene, using this tool also in the design phase, in order to optimize the geometries of the desired equipment.

The performances of this tool were then improved through a deeper focus on the numerical strategy taken by the KPP in order to solve the chemical reactor network created downstream of CFD simulation. An evolution of this tool, named KPPSMOKE (as part of OpenSMOKE libraries (Cuoci, Frassoldati, Faravelli, & Ranzi (2011))), was created, thus being able to manage very large reactor networks (up to 10^6 reactors) with very detailed kinetic mechanisms (hundreds of species and thousands of reactions). Indeed, if not treated properly, those systems may have excessive requirements of computational resources.

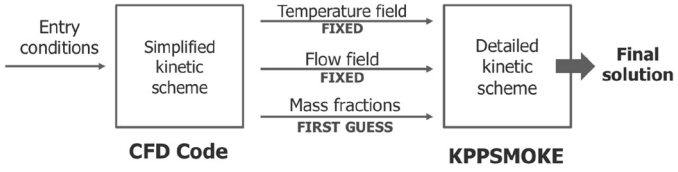

Fig. 1. Kinetic post-processing procedure.

In this paper the development of KPPSMOKE is described in detail by focusing the attention on the two key concepts representing the cornerstones of the newly conceived software: (i) on the one hand, a flexible and efficient strategy to approach the solution was developed, with the help of external libraries for the resolution of the algebraic problems concerned; (ii) on the other, the potentials of the created algorithm were further extended through a parallel distribution of data and tasks over more processes, thus making available a higher memory and computing power.

These two aspects are described in this paper, which is organized as follows: starting from a description of the mathematical model of the reactor network, as well as its constitutive equations (Section 2), the strategy to pursue the solution is explained through an overview of the different techniques adopted (Section 3). Then, Section 4 describes the very core of the work, i.e. the implementation of the numerical model in a parallel code. Its performances are finally benchmarked in Sections 5 and 6, where the comparison of the KPPSMOKE predictions with the experimental data at different scales is presented and its numerical behavior is checked, respectively.

2. Numerical model

The kinetic post-processing approach was developed by relying on one basic assumption: minor species, like radicals and pollutants, do not affect the thermo-fluid dynamic fields (temperature and velocity) in a significant way because of their small amounts (usually ppm). This hypothesis is reasonable as long as they do not have a role in the radiative heat transfer: i.e., while this assumption is largely acceptable for NO_x, CO or polycyclic aromatic hydrocarbons (PAH), it is not for compounds like soot.

Therefore, concerned fluid dynamics can be evaluated upstream through a global, or skeletal, kinetic mechanism, by using one of the available CFD codes for reacting systems (e.g. ANSYS® Fluent (Ansys Inc., 2011), OpenFOAM® (2012), etc.). A fine mesh is usually needed in order to correctly describe the temperature and velocity fields inside the system under investigation. Once this step is complete, the evaluated temperature and velocity fields are exported into the KPPSMOKE and kept fixed during the post-processing analysis (Fig. 1).

The topology of the original computational mesh is imported from the CFD code. Then, a corresponding reactor network model is created in the KPPSMOKE, where each imported cell is assumed as a perfectly stirred reactor, with a fixed temperature and volume. Convective flows among reactors are also obtained in this first step and kept fixed afterwards. These assumptions allow to decouple the mass balances for the species from the energy and momentum balances. Only the former need then to be solved, thus decreasing the overall number of equations and especially their non-linearity (linked to the dependence of kinetic constants from temperature through the Arrhenius' law).

The resulting system is then constituted of $N_C \times N_S$ equations, where N_C and N_S are respectively the number of cells (i.e. reactors) and species, and the unknowns are the related mass fractions. For