

# Implementing detailed chemistry and in-cylinder stratification into 0/1-D IC engine cycle simulation tools

Aaron R. Coble<sup>1</sup>, Andrew J. Smallbone<sup>1</sup>, Amit N. Bhave<sup>1</sup>, Sebastian Mosbach<sup>2</sup>, Markus Kraft<sup>2</sup>, Patrick Niven<sup>4</sup>, Steven Amphlett<sup>3</sup>

released: 1 January 2011

<sup>1</sup> cmcl innovations  
Sheraton House  
Castle Park  
Cambridge, CB3 0AX  
United Kingdom  
E-mail: [anbhave@cmclinnovations.com](mailto:anbhave@cmclinnovations.com)

<sup>2</sup> Department of Chemical Engineering  
and Biotechnology  
University of Cambridge  
New Museums Site  
Pembroke Street  
Cambridge, CB2 3RA  
United Kingdom  
E-mail: [mk306@cam.ac.uk](mailto:mk306@cam.ac.uk)

<sup>3</sup> Ricardo UK Ltd  
Shoreham Technical Centre  
Shore-by-Sea  
West Sussex, BN43 5FG  
United Kingdom

<sup>4</sup> Ricardo Inc.  
4000 Ricardo Drive  
Van Buren Twp  
MI 48111  
USA

Preprint No. 115



**Edited by**

Computational Modelling Group  
Department of Chemical Engineering and Biotechnology  
University of Cambridge  
New Museums Site  
Pembroke Street  
Cambridge CB2 3RA  
United Kingdom

**Fax:** + 44 (0)1223 334796

**E-Mail:** [c4e@cam.ac.uk](mailto:c4e@cam.ac.uk)

**World Wide Web:** <http://como.cheng.cam.ac.uk/>



## Abstract

Employing detailed chemistry into modern engine simulation technologies has potential to enhance the robustness and predictive power of such tools. Specifically this means significant advancements in the ability to compute the onset of ignition, low and high temperature heat release, local extinction, knocking, exhaust gas emissions formation etc. resulting in a set of tools which can be employed to carry out virtual engineering studies and add additional insight into common IC engine development activities such as computing IMEP, identifying safe/feasible operating ranges, minimizing exhaust gas emissions and optimizing operating strategy. However the adoption of detailed chemistry comes at a greater computational cost, this paper investigates the means to retain computational robustness and ease of use whilst reducing computational timescales.

This paper focuses upon a PDF (Probability Density Function) based model based on the Stochastic Reactor Model (SRM), which has gained increasing attention from academics and industry for its capabilities to account for in-cylinder processes such as chemical kinetics, fuel injection, turbulent mixing, heat transfer etc. whilst retaining in-cylinder stratification of mixture composition (i.e. fuel equivalence ratio) and temperature. Among the techniques considered here are: a standard KIVA 3V simulation, down-sampling from 3D CFD composition-space to stochastic particles using sequential coupling of KIVA 3V and SRM, the use of detailed chemical kinetics within SRM, chemical mechanism reduction, down-sampling of a chemical mixture space within the SRM, and parallelization of chemistry solution within SRM. The experimental engine setup studied is that used by Cao et. al. [6], employing Premixed Charge Compression Ignition (PCCI), which is a Low Temperature Combustion (LTC) strategy for diesel engines. This paper demonstrates how equivalent results can be achieved with a reduction in computational time from 28 days to 10 minutes. In order to enable modern engineers to exploit detailed chemistry in and IC engine context, the SRM Suite was coupled with a 1D engine cycle simulation tool. An example is presented which demonstrates how the adoption of SRM Suite can be implemented into standard 0/1-D toolkits with corresponding predictive capability.

# Contents

<b>1</b>	<b>Introduction</b>	<b>4</b>
1.1	Exploiting the potential of these models for IC engine applications . . . .	4
<b>2</b>	<b>Experimental setup</b>	<b>6</b>
2.1	Engine Configuration . . . . .	6
2.2	High performance computing facility . . . . .	6
<b>3</b>	<b>Description of modelling techniques</b>	<b>7</b>
3.0.1	SRM SUITE . . . . .	7
3.0.2	Sequential coupling with CFD . . . . .	8
3.0.3	Fuel oxidation and emission formation models . . . . .	8
3.0.4	Degree of mixing space down-sampling . . . . .	9
3.0.5	Parallelization of chemistry solution . . . . .	9
3.0.6	Coupling with 0D/1D engine cycle simulation software . . . . .	10
<b>4</b>	<b>Results</b>	<b>11</b>
4.1	Comparison predictive capability . . . . .	11
4.2	Model convergence relative to detail of chemistry . . . . .	11
4.3	Model convergence relative to degree of down-sampling . . . . .	13
4.4	Computational cost relative to degree of down-sampling . . . . .	14
4.5	The impact of parallelization on simulation timescales . . . . .	15
4.6	Coupling SRM with 0D/1D engine cycle simulation software . . . . .	17
<b>5</b>	<b>Summary</b>	<b>20</b>
	<b>References</b>	<b>21</b>
	<b>Citation Index</b>	<b>24</b>

## List of Figures

1	Advantages of the adoption of detailed chemistry in an engine context . .	5
2	Caterpillar MCTE Experimental Setup: Normalized Injection Rate Profile	7

3	Execution flow for 0D/1D-coupled parallelized SRM simulation with 8 stochastic particles distributed across 4 cores . . . . .	10
4	Comparison of pressure profiles for experimental data, CFD, and SRM with detailed and reduced chemistry . . . . .	12
5	Comparison of NOx emissions for experimental data, CFD, and SRM with detailed and reduced chemistry . . . . .	12
6	Comparison of percent error between model and experiment for peak pressure and NOx emissions with respect to computational time for detailed and reduced chemistry . . . . .	13
7	Comparison of percent error between model and experiment for peak pressure and NOx emissions with respect to the number of stochastic particles for detailed and reduced chemistry . . . . .	14
8	Comparison of computational time with respect to the number of stochastic particles for detailed chemistry . . . . .	15
9	Comparison of computational time with respect to the number of stochastic particles for reduced chemistry . . . . .	16
10	Comparison of computational time with respect to computational power for detailed chemistry. . . . .	17
11	Comparison of computational time with respect to computational power for reduced chemistry. . . . .	18
12	Overview of methodology for implementing detailed chemistry 0D/1D engine cycle simulation software . . . . .	18
13	Engine map in 0D/1D engine cycle simulation software using SRM SUITE as an external cylinder . . . . .	19
14	Setting SRM SUITE input parameters within 0D/1D engine cycle simulation software . . . . .	19
15	SRM SSUITE simulation, using detailed chemistry, running automatically within a coupled 0D/1D engine cycle simulation . . . . .	20

## List of Tables

1	Basic Experimental Configuration of Caterpillar MCTE . . . . .	6
2	HPC single node specification . . . . .	7
3	Description of chemical models . . . . .	9

# 1 Introduction

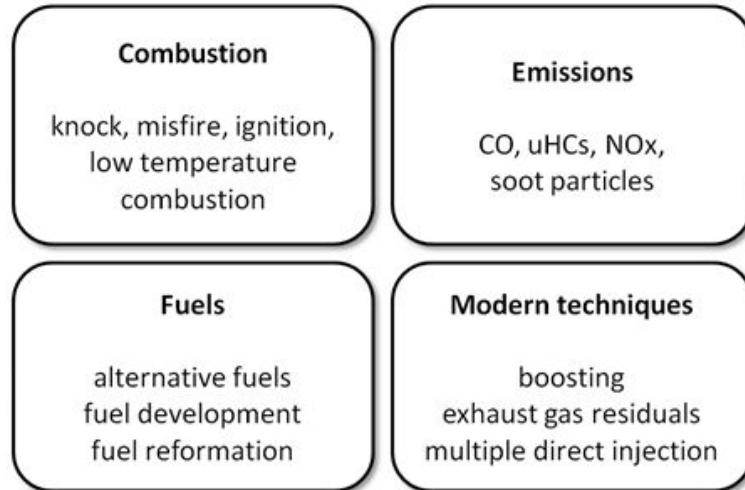
In the last decade, advancements in our understanding of combustion chemistry have yielded increasingly advanced models which describe the oxidation of hydrocarbon fuels in air at engine-like pressures and temperatures [24]. Coupling advanced chemical sub-models with already well established fluid dynamics codes has yielded a new generation of advanced simulation tools, beginning the era of predictive combustion science [14].

Over recent years, hydrocarbon fuel-oxidation models have extended in their size and scope from those initially describing the oxidation of the lower molecular weight fuels, C1 to C4 [27] to include molecules of size relevant to modern IC engines such as iso-octane, n-heptane, toluene, n-decane, n-do-decane, cycloalkanes etc. [9, 21]. These have extended to form surrogate fuels of gasoline [5, 22, 27] and diesel fuels [20] which mimic properties of commercial fuels such as octane rating, composition, H/C ratio, heating value etc.. At each step, these have systematically been validated against many fundamental experiments involving rapid compression machines, shock tubes, laminar flames, perfectly stirred reactors etc. yielding fundamental measurements of in-situ chemical species etc. [9, 24, 27]. One major advantage of these models is that they can properly describe the kinetics associated with ignition, auto-ignition, flame propagation, emissions formation etc. over the full range of pressures, temperatures and mixture compositions reported in modern IC engines.

## 1.1 Exploiting the potential of these models for IC engine applications

The applications and potential application of detailed chemistry to tackle automotive technology challenges is summarized in Figure 1, with improved results for (a) combustion i.e. knock, misfire, ignition, low temperature combustion etc. (b) Fuels i.e. adoption of conventional fuels or advanced fuels, gasoline, diesel, bio-fuels, hydrogen natural gas etc., (c) Emissions regulated gaseous (CO, uHCs, NO<sub>x</sub>) or particulates such as soot, and (d) modern technologies, since these models are "physics-based" simulations are robust and well validated, they remain robust in terms of boosting, high EGR and multiple direct injection strategies.

Whilst improved predictive capability is a major advantage of detailed chemistry, the main challenge is identifying the means to exploit these models and implement them into standard engine development codes. The computational timescales associated with implementing detailed chemistry into an engine context has historically limited the scope of the adoption of such models. When applied to 3D CFD applications a simulation of a single cycle involving a detailed chemical kinetic mechanism (comprised of 872 reactions and 188 species) can take up to 28 days [6] this means that alternative approaches for chemical modeling must be employed such as reduction strategies [15], or the adoption of alternative combustion simulation models such as PDF (Probability Density Function) approaches [3, 4, 16–19, 25], multi-zone (hear defined as fewer than ten zones) or single zone solutions [1, 2, 10, 12, 13]. However since mixture preparation in modern IC engines



**Figure 1:** *Advantages of the adoption of detailed chemistry in an engine context.*

is now often controlled using a direct injection event, in-cylinder composition is highly stratified making the adoption of multi-or single-zone solutions insufficient for resolving the full range of mixture composition required to account for ignition, flame propagation and emissions formation. In particular, particulates which have proven to be generated in regions comprised of less than 3 % of the overall mixture [19].

Over the last decade, PDF approaches such as the SRM are gaining increasing attention from academia and industry due to their ability to account for in-cylinder stratification in terms of mixture composition and temperature in amenable timescales [3, 4, 16–19, 25]. These then enable the adoption of detailed chemistry into their solution with previous studies taking between 90 seconds and 90 minutes depending on the detail required in the chemical model [3, 4, 16–19, 23, 25]. In the past, this model has also been employed to carry out large parametric sweeps of engine control variables such as ignition-timing, EGR fraction etc. results were then stored in look-up tables and coupled with 0/1D engine tools.

Though available for decades, the cost of hardware for parallelized computing has reduced dramatically in recent years, corresponding with a widespread increase in availability. As a result, multi-core desktop machines have become commonplace for consumers and High Performance Computing (HPC) facilities, such as the one used in this study, are now within reach of academic researchers and small enterprise. The SRM approach can take advantage of the rising power of parallelization to reduce simulation timescales by distributing the solution of chemistry; this improvement in timescale is particularly significant when employing detailed chemistry, as is demonstrated in this study.

This paper has two main purposes:

- (1) To demonstrate the computational speed up associated with the parallelization of the SRM across up to 88 cores - results are then compared for emissions with 3D CFD for cases with a reduced and detailed chemical fuel oxidation and emissions formation model.
- (2) The SRM is then implemented into a standard 0/1D cycle simulation code for full

engine breathing and combustion simulation and for ease of operation within standard engine simulation toolkit.

## 2 Experimental setup

### 2.1 Engine Configuration

The experimental data used in this study is from a Caterpillar multi-cylinder test engine (MCTE) and is the same as the data used by Cao et. al. [6]. The basic engine configuration and the normalized injection rate profile are given in Table 1 and Figure 2 respectively.

**Table 1:** *Basic Experimental Configuration of Caterpillar MCTE.*

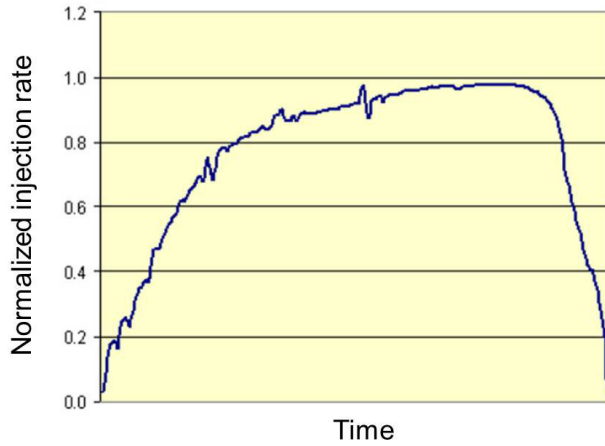
Parameter	Value
Bore (mm)	137.1
Stroke (mm)	171.5
Connecting Rod Length (mm)	270.0
Displacement (6 cylinders)(litres)	15.2
Compression Ratio ( - )	10:1
Engine Speed (rpm)	1200
BMEP (kPa)	850
Nozzle Hole Size (mm)	0.129
Injection System	MEUI
Injection Pressure (MPa)	170 - 190
Nominal AVL Swirl ( - )	0.3
Start of Injection(aTDC)	-38.0
End of Injection (aTDC)	-23.0
Inject Fuel Mass (mg)	126.0

Engine-out  $\text{NO}_x$  emissions data were collected from the experimental engine during steady-state operation. Both external Exhaust Gas Recycling (EGR) and Intake Valve Actuation (IVA) were used for phasing of combustion; cylinder-to-cylinder variation of combustion characteristics was minimized through the use of IVA, which is supported by alignment in heat release rate across the cylinders.

### 2.2 High performance computing facility

The combustion simulations performed in this study were executed using a High Performance Computing (HPC) Cluster in order to maintain a consistent platform for benchmarking all of the simulations, including parallelized simulations. The HPC facility consists of an SGI Altix Cluster with 168 cores, 176 GB of RAM, and an end-to-end MPI latency typically less than 2 microseconds. The 160 cores are spread across 20 dual





**Figure 2:** *Caterpillar MCTE Experimental Setup: Normalized Injection Rate Profile.*

socket XE320 1U nodes; each node constitutes a single symmetric multiprocessing unit (SMP) with 8GB of RAM and two quad-core 3.00 GHz Intel Harpertown processors with a 12MB cache and 1.6GHz front-side bus. The simulations were run on a 64-bit installation of the Fedora Core (Release 12: Constantine) Linux distribution, using the MPICH2 (release 1.2.1p1) [11] implementation of the Message Passing Interface (MPI) for parallelized simulations. The specification for a single node of the HPC is outlined in Table 2.

**Table 2:** *HPC single node specification.*

Number of Cores	8
Number of Processors	2
Clock Speed (GHz)	3.00
RAM (GB)	8
Cache (MB)	12
Bus Speed (GHz)	1.6
Architecture	64-Bit
Linux Distribution	Fedora Core (Release 12)
MPI Implementation	MPICH2 release 1.2.1p1

## 3 Description of modelling techniques

### 3.0.1 SRM SUITE

The primary technology used throughout this study is CMCL Innovations' SRM SUITE, which is a Stochastic Reactor Model (SRM) for in-cylinder combustion simulation. The SRM SUITE is a PDF-based model derived from the joint composition PDF transport

equation for scalars, assuming statistical homogeneity. The SRM is zero-dimensional, which means that quantities are independent of space. Note, however, that the model does not assume spatial homogeneity, but rather statistical homogeneity, i.e., the statistics of turbulence -in other words the PDF - is the same everywhere. The difference is crucial, since in-homogeneities are the key to predicting combustion emissions. These "stochastic particles" are not to be confused with physical atoms or molecules, fluid parcels, or zones in a multi-zone approach. They carry no geometric or spatial information whatsoever, but as an ensemble constitute a statistical representation of the PDF. Each particle carries with it only the mass fractions and the temperature. In this approach, the composition-space in the engine is down-sampled to a given number of "stochastic particles", which are assumed to be homogenous, following the distribution of in-cylinder conditions. The SRM SUITE accounts for detailed chemical kinetics, turbulent micro-mixing and convective heat-transfer and supports study of combustion, emissions, and fuels in 0/1-D timescales; unlike multi-zone models it also accounts for fluctuations in quantity. The PDF transport equation was solved using a Monte Carlo technique with second-order operator splitting and the Euclidean Minimal Spanning Tree (EMST) model [26] describes the turbulent mixing, based on the proximity of particles in composition space. The SRM has previously been applied to HCCI/PCCI engines [3, 4, 6, 16–19, 25] an overview of the SRM methodology can be found in reference [17].

### 3.0.2 Sequential coupling with CFD

This study does not directly make use of sequential coupling of 3D CFD and SRM; however, comparison can be made between the results achieved in this study and those achieved by Cao et. al. [6] using a segregated sequential-coupling approach, since the experimental and simulation configuration considered are the same. This allows for further benchmarking of predictive capability and computational cost.

Sequentially coupled approach using 3D CFD with chemical kinetics and SRM with detailed chemistry is that it allows insight into the fuel and temperature distribution, particularly during Low Temperature Chemistry (LTC), but without running an end-to-end 3D CFD simulation with chemical kinetics, which is cost prohibitive and suffers from uncertainties in the reaction rate closure. The results below show that comparable results can be achieved using the standalone SRM SUITE with detailed kinetics, avoiding the expense of the initial 3D CFD calculation, albeit without explicitly considering the role of engine geometry.

### 3.0.3 Fuel oxidation and emission formation models

The two fuel oxidation and emissions formation models are summarized in Table 3. The level of detail employed by a fuel oxidation and emissions formation model is characterized by its number of reactions and chemical species; an increased number of reactions/species results in a more robust solution, but at an increased computational cost. The development of the detailed model including its advanced soot formation chemistry has been detailed in [19]. The general applications of this particular reaction mechanism have been in the solution of combustion and emissions for SI, HCCI, PPCI, diesel, etc.

[7].

The reduced model has fewer reactions but has been extended to account for  $\text{NO}_x$  formation chemistry. However, to retain a faster computational solution, no chemistry for soot formation processes was included. During development, the objective of this model was to characterize diesel fuelled combustion in reasonable timescales such that it could be applied with SRM as a plug-in to standard 0/1D engine development computational toolkits.

**Table 3:** *Description of chemical models*

Name	Reduced Model	Detailed Model
Description	cmcl diesel surrogate with $\text{NO}_x$ V1.2	n-heptanes/ iso-octane with soot and $\text{NO}_x$
Number of Species	38	188
Number of Reactions	50	872

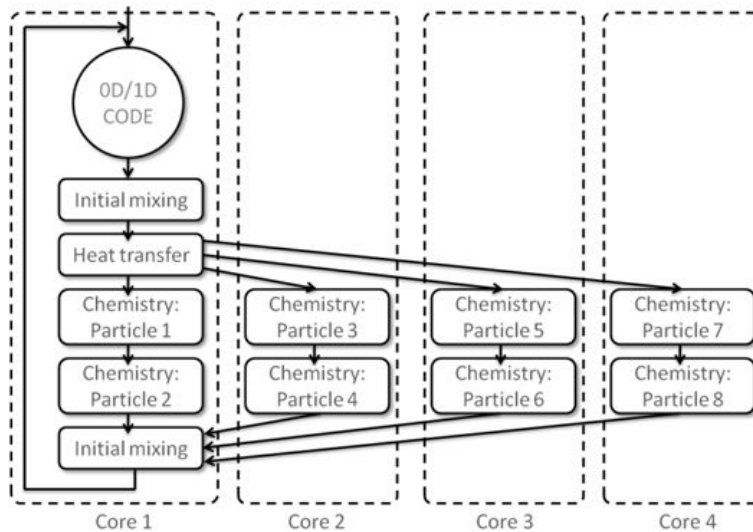
### 3.0.4 Degree of mixing space down-sampling

In the SRM SUITE, the user can select the number of "stochastic particles" onto which the in-cylinder composition space is down-sampled: the larger the number of particles, the finer the granularity with which the particles capture the compositional characteristics within the cylinder. Generally, more particles will result in greater predictive capability, because the composition is more accurately represented. Given that the chemical kinetics equations are solved per particle, a larger number of particles also results in a greater computational cost, as we will see in the results below. The number of stochastic particles required for convergence of results depends on the chemical mechanism employed; a reduced-chemistry mechanism may require 10 or fewer particles, while detailed-chemistry mechanisms can require from 100-500 particles. Fewer particles are required when using reduced chemistry, since the compositional space to be represented has already been simplified through the reduction of chemistry. The relationship between degree of down-sampling, predictive capability and computational cost is discussed below.

### 3.0.5 Parallelization of chemistry solution

The solution of chemistry equations in the SRM, using an Ordinary Differential Equation (ODE) Solver, can be very computationally expensive, particularly when using a detailed chemical mechanism. Fortunately, the process is relatively independent from one stochastic particle to another, so it easily lends itself to parallelization. Below we will see the relationship between the degree of parallelization, the level of detail in the chemical mechanism, the number of particles, and computational time.

In the parallelized version of the SRM SUITE, the Message Passing Interface (MPI) is used to spawn sub-processes for each core available, each of which is each used to compute the chemistry for an individual particle (or multiple particles in sequence if the number of particles exceeds the number of core) in a time-step of the SRM, and to manage the communication between these distributed sub-processes. A time-step in both the parallelized and standard version of the SRM follows the same profile apart from the distribution of the chemistry computation: some initial mixing is performed, heat transfer is computed, chemistry is solved either on a single core or distributed using MPI, and some final mixing is performed. This process is illustrated in Figure 3 for a 0D/1D-coupled SRM simulation involving 8 stochastic particles and distributed across 4 cores.



**Figure 3:** Execution flow for 0D/1D-coupled parallelized SRM simulation with 8 stochastic particles distributed across 4 cores

Calculating the turbulent mixing is another computationally expensive process in the execution of a simulation time-step. Unfortunately, this process interacts heavily with all of the stochastic particles, so it cannot be parallelized easily or without incurring a substantial synchronization overhead. Methods for reducing computational costs in the turbulent mixing portion of the simulation, without parallelization and whilst maintaining predictive capability, present an interesting area for future research.

### 3.0.6 Coupling with 0D/1D engine cycle simulation software

Coupling of the SRM SUITE with industry standard 0D/1D engine cycle simulation software contributes directly to achieving the goals motivating this work. The methodologies examined in this study aim to deliver in-cylinder combustion simulations, with enhanced predictive capabilities, in a timescale comparable to 0D/1D tools. The SRM SUITE has already been coupled with most major 0D/1D engine cycle simulation software [7, 8] and, with these timescales achieved, can be used seamlessly in an end-to-end powertrain simulation, providing detailed results for fuels, combustion, and emissions. Typically the SRM is coupled as an external cylinder available within the 0D/1D tool, which is called each

simulation timestep during the valve-closed portion of the engine cycle, so it is possible to use any of the methodologies discussed here, such as parallelization or mechanism reduction, within the coupled codes.

## 4 Results

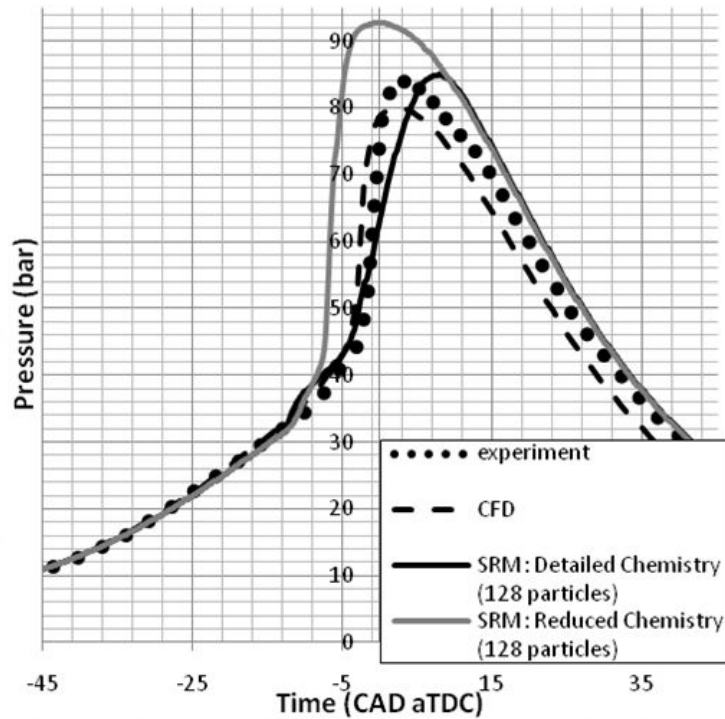
### 4.1 Comparison predictive capability

Before examining performance with respect to computational cost, it is most important to examine the predictive capability in terms of combustion and emissions for each of the simulation technologies considered. This task has been carried out previously for many applications [3, 4, 16–19, 25], Figure 4 compares the pressure profiles obtained from the experimental engine, 3D CFD, and the SRM with both detailed and reduced chemistry employed; the CFD data is from Cao [6]. The fit to the model of the CFD simulation and the SRM with detailed chemistry are comparable, while the SRM with reduced chemistry performs slightly less favorably. It must be noted that all models employed here are sensitive to uncertainty, it is considered that the initial temperature of the simulation at inlet valve closure is not measured exactly and can be considered to be known within 25K. In PPCI combustion, which is kinetically driven compared to say diesel combustion, sensitivity to initial temperature is significant and would certainly impact on the ignition timing and comfortably cover the range of cases observed. As such it is very difficult to justify that any of these simulation techniques are superior based upon the results of the single cycle. Only by running equivalent simulations and comparing model results over a range of operating points could such a statement be made.


$\text{NO}_x$  emissions are compared between the experimental data, CFD simulation, and the SRM (with both detailed and reduced chemistry) in Figure 5; again, the CFD data is from Cao [6]. These results confirm Cao's earlier result [6] showing an outstanding improvement of the SRM with detailed chemistry over the CFD simulation, in terms of  $\text{NO}_x$  emissions. The SRM with reduced chemistry also outperforms the CFD simulation, but without achieving the quality of emissions prediction produced when employing the detailed chemical mechanism. This loss of predictive emissions capability is probably a result of the reduction in the number of chemical species and reactions; this is consistent with the observations of [23] which indicated that the proposed mechanism is adequate for trend based analysis.

### 4.2 Model convergence relative to detail of chemistry

The SRM has different degrees and rates of convergence for combustion and emissions, depending on whether detailed or reduced chemistry is used; these results, for computation parallelized across 16 cores, are illustrated in Figure 6. When employing reduced chemistry, the model converges very rapidly, in approximately 5 seconds for peak pressure and 1 minute for  $\text{NO}_x$  emissions, but does not converge to the same degree as when using detailed chemistry. As was explained in the previous section, the 10 % error in



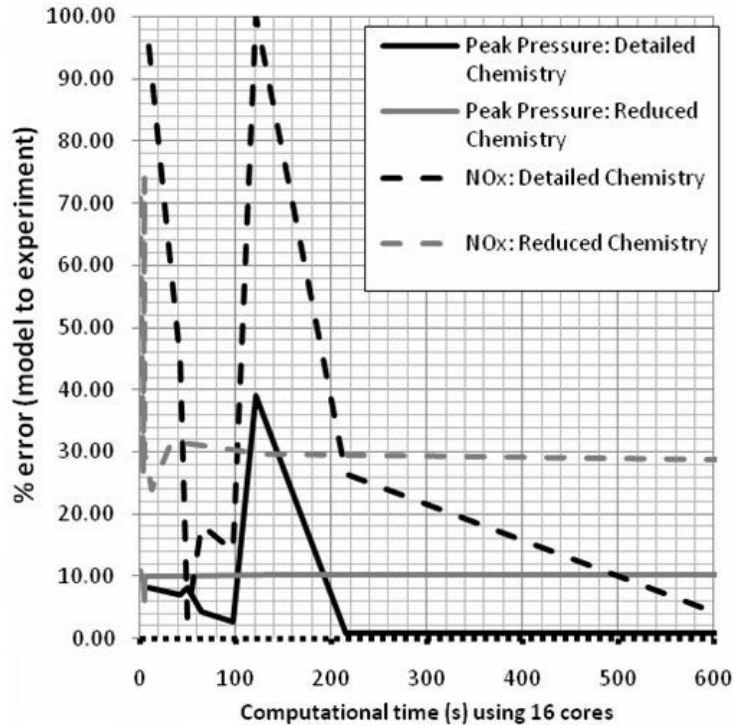
**Figure 4:** Comparison of pressure profiles for experimental data, CFD, and SRM with detailed and reduced chemistry



**Figure 5:** Comparison of NOx emissions for experimental data, CFD, and SRM with detailed and reduced chemistry

peak pressure could be reduced by a nominal adjustment of initial temperature in order to tune the model. A simulation completing in approximately 4 minutes (using 64 stochastic particles), is sufficient to achieve convergence for peak pressure, when using detailed chemistry. However, a simulation taking approximately 10 minutes (using 128 stochastic

particles) is required to achieve convergence for both peak pressure and NO<sub>x</sub> emissions. The model converges more slowly when using detailed chemistry, but achieves an excellent correspondence with the experimental data.




**Figure 6:** Comparison of percent error between model and experiment for peak pressure and NO<sub>x</sub> emissions with respect to computational time for detailed and reduced chemistry

These results suggest that the SRM should be employed with detailed chemistry, whenever very accurate combustion and emissions predictions are required, whilst achieving moderate timescales. In contrast, reduced chemistry can be used to achieve drastically reduce timescales and reasonably accurate combustion predictions, but delivers only rough estimates of NO<sub>x</sub> emissions.

### 4.3 Model convergence relative to degree of down-sampling

In addition to the size of the chemical mechanism used, the degree to which the compositional space is down-sampled into "stochastic particles" affects the degree of convergence of the model (as well as the computational time). Figure 7 demonstrates that rate at which down-sampling of the compositional space impacts model convergence for both the detailed and reduced chemical mechanism.

Notice that the reduced chemical mechanism converges much more quickly (at approximately 8 particles for peak pressure and 64 particles for NO<sub>x</sub> emissions) with respect to the number of stochastic particles, when compared to the detailed chemical mechanism. This is because the amount of stratification possible in the mixture space is much



**Figure 7:** Comparison of percent error between model and experiment for peak pressure and NOx emissions with respect to the number of stochastic particles for detailed and reduced chemistry

less when using the reduced mechanism, since there are fewer chemical species, so fewer stochastic particles and a larger amount of granularity is sufficient to capture the composition of the mixture space. In contrast, the detailed mechanism has more species and has the potential for greater stratification of the mixture; therefore finer granularity is required when down-sampling in order to sufficiently characterize the composition space.


These results are consistent with previous studies where around one hundred stochastic particles have proven sufficient for characterizing heat release and exhaust gas emissions and to reduce this would potentially impact upon model predictive capability [3, 4, 16–19, 25].

#### 4.4 Computational cost relative to degree of down-sampling

This section examines the rate at which the computational cost of a simulation grows relative to an increasing number of stochastic particles, corresponding to less down-sampling and improved characterization of the mixture space. The growth in computational cost is dependent on the level detail in the chemical mechanism employed, as can be seen by comparing Figures 8 and 9. The plot in Figure 8 shows that the computational time of the model, employing detailed chemistry, grows linearly with the number of stochastic particles, whenever it is run on a single core; this is a positive result in terms of scalability. The reason for this rate of growth is that the chemistry, which tends to dominate



computational cost, is solved sequentially for each particle in turn, whatever the number of particles. If the same model is parallelized, the rate of growth in computational time increases over time, but remains bounded by the linear rate of growth for a single core. This non-linear growth rate occurs because the parallelization of the chemistry solution fixes the computational cost of the chemistry at a constant rate, for a given number of cores and increasing number of particles, up until the number of particles reaches the next even multiple of the number of cores, at which point the number of particles solved in sequence on a particle core increases; note that this constant rate of growth between points of increase is not strictly observed, because the computational cost of the turbulent mixing code still increases with the number of particles and overtime dominates.

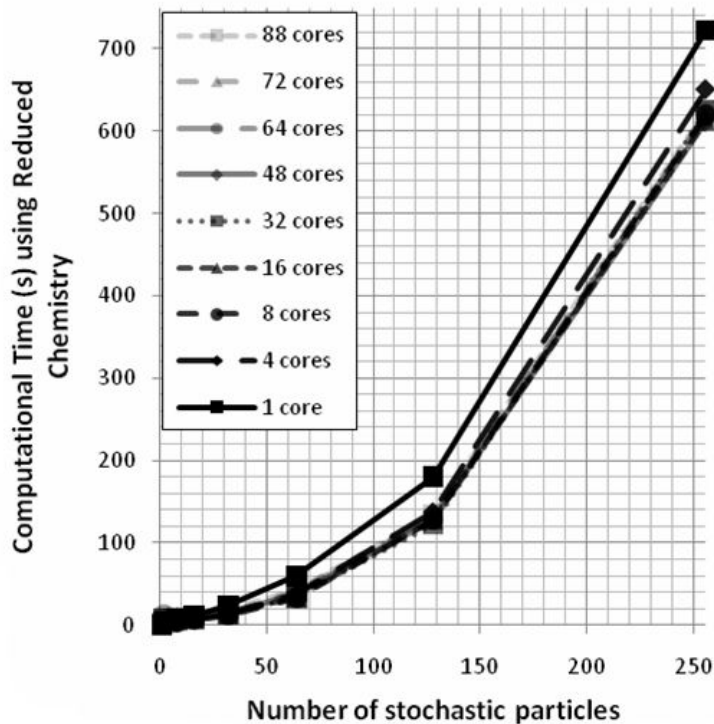


**Figure 8:** Comparison of computational time with respect to the number of stochastic particles for detailed chemistry

Figure 9 shows that the computational time of the SRM, employing reduced chemistry, grows exponentially with respect to the number of stochastic particles; this effect increases with parallelization. The reason for this growth rate is that the cost of computing the chemistry is drastically reduced, so the cost of computing turbulent mixing, which grows exponentially, dominates from the outset.

#### 4.5 The impact of parallelization on simulation timescales

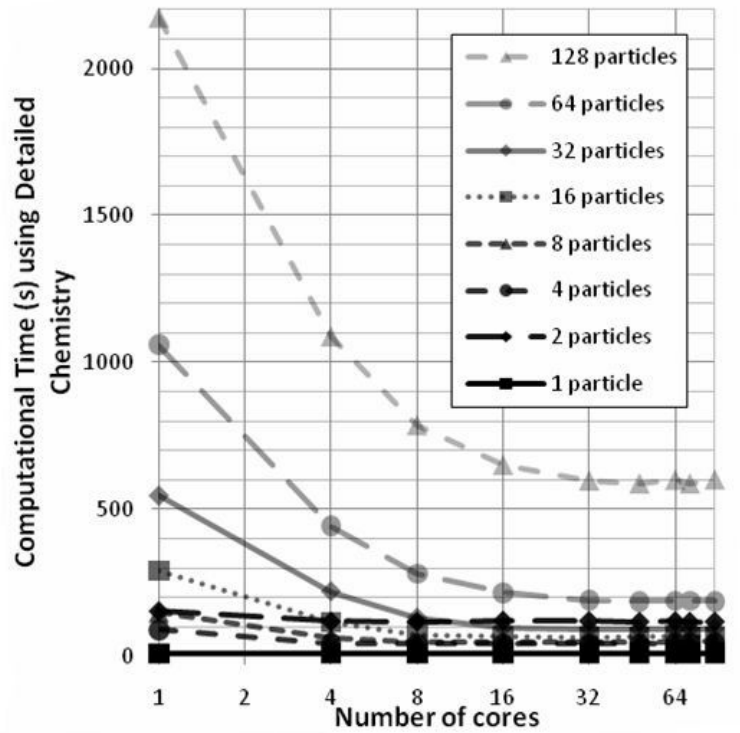
Parallelization of chemistry solution in the SRM presents an opportunity for dramatic reduction of simulation timescales. Figures 10 and 11 provide an overview of the impact of parallelization on computational time when employing detailed and reduced chemistry



**Figure 9:** Comparison of computational time with respect to the number of stochastic particles for reduced chemistry

respectively; a logarithmic scale has been used for the X-axis so that the trends are easier to view. The plots in Figure 10 show that computational time decreases exponentially with increasing parallelization and that this reduction is bounded by the point where there is one core per 2 particles. As a result, simulations involving a greater number of particles benefit more from parallelization, whereas simulations with few particles quickly reach their saturation point. Furthermore, beyond 8-16 cores, a point of diminishing returns is reached, where most of the benefit in reduction of timescale has been gained and the increasing costs for computing power probably will not be justified for most users. The positive aspect of this is that a relatively affordable 8 core desktop machine is sufficient to reap substantial benefits in simulation timescale through parallelization, particularly when using the moderate number of particles (approximately 128) required to achieve the degree of combustion and emissions convergence that are the benefits of employing the detailed chemical mechanism. For those with greater requirements and needing to use more particles, a 16 core facility may still be within reach and the point of diminishing returns is delayed slightly with the increased number of particles.

The computational time when employing reduced chemistry already represents a significant improvement compared with the detailed mechanism, but unfortunately does not benefit significantly from parallelization. The reason for this lack of improvement with parallelization, shown in Figure 11, is that the computational time devoted to solving chemistry is already sufficiently low such that the time to compute turbulent mixing dominates from the outset. A reduction in timescale is only observed with a large number of

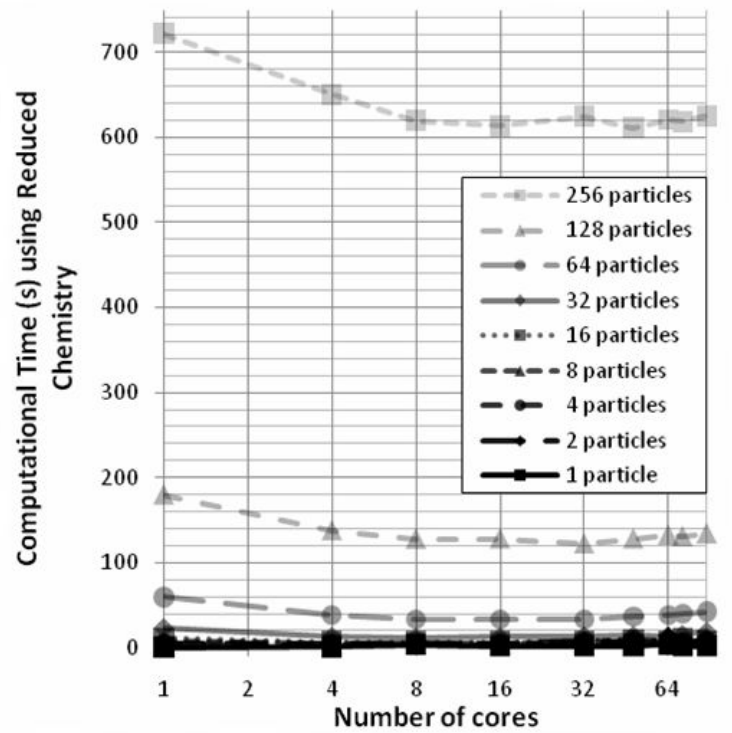


**Figure 10:** Comparison of computational time with respect to computational power for detailed chemistry.

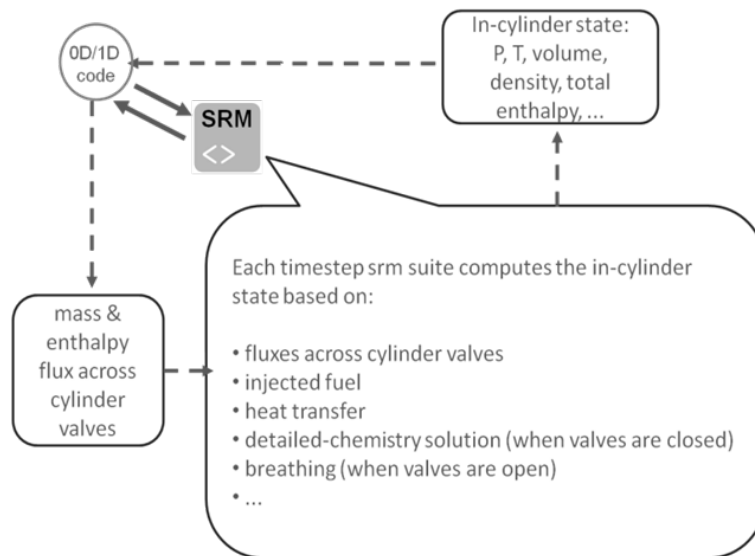
particles, but using such a large number of particles has previously been shown of little benefit (Figure 7). The simulation timescale of the SRM, using reduced chemistry and run on a single core, is still shorter than employing detailed chemistry and parallelizing across 8 or 16 nodes. This suggests that the model with reduced chemistry can be gainfully employed on a single-core machine and with a dramatically shortened timescale, whenever the emissions convergence offered by detailed chemistry is not required.

#### 4.6 Coupling SRM with 0D/1D engine cycle simulation software

Despite the advantages demonstrated, wide-spread adoption of detailed chemistry in an engine simulation context will not occur unless these tools are integrated into industry standard engine cycle simulation software in a smooth and easy-to-use workflow. The SRM SUITE described above has been coupled to most major 0D/1D engine cycle simulation tools, providing seamless integration of detailed chemistry for fuels, combustion, and emissions. Typically, this coupling is implemented as an external cylinder available within the 0D/1D software. The coupled SRM is called by the 0D/1D code each simulation timestep during the valve-closed portion of the engine cycle, in order to simulate in-cylinder conditions and events. Fluxes, in terms of mass and energy, are passed to the external cylinder by the 0D/1D code and the SRM returns the in-cylinder state to the 0D/1D code based on its calculations. The interaction of the coupled SRM with the 0D/1D code is summarized in Figure 12.



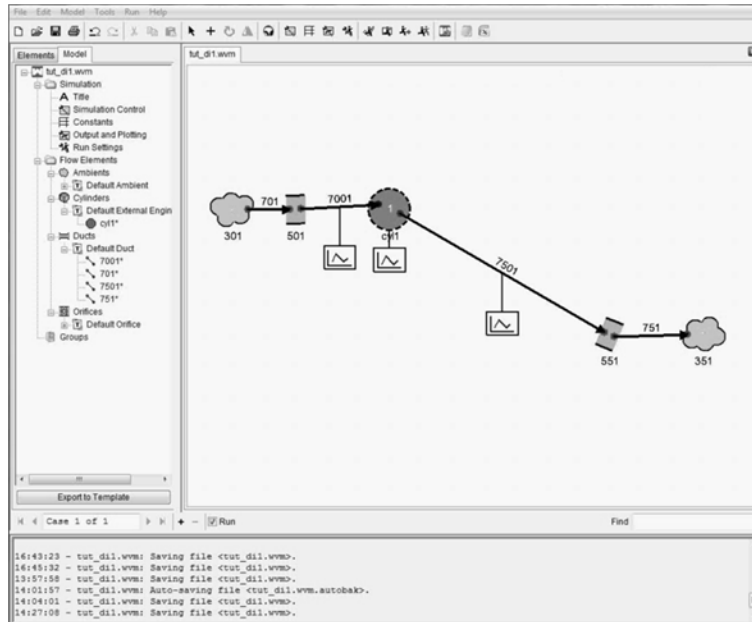
**Figure 11:** Comparison of computational time with respect to computational power for reduced chemistry.



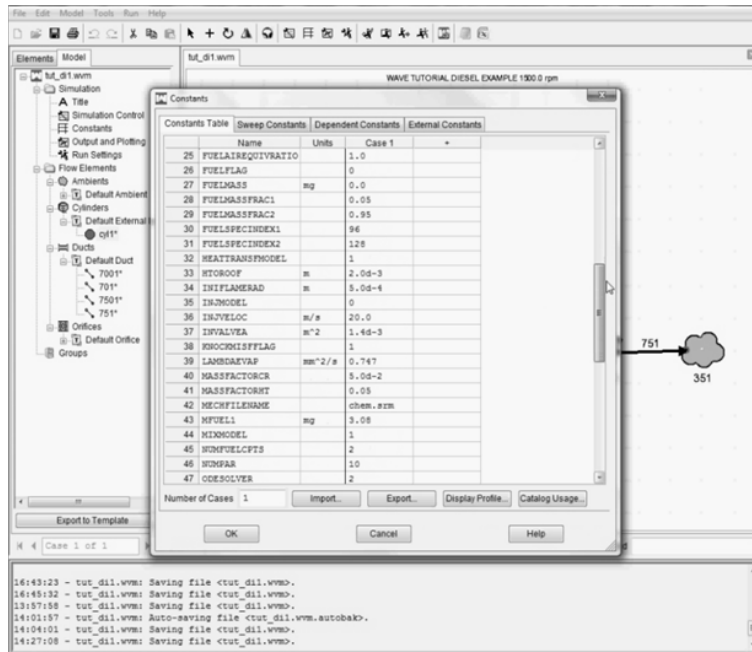
**Figure 12:** Overview of methodology for implementing detailed chemistry 0D/1D engine cycle simulation software

This approach makes use of detailed chemistry accessible, since the model parameters can be set directly in the 0D/1D tool, the in-cylinder simulation is automatically run within the end-to-end engine simulation, and the final results can be viewed within the engine cycle

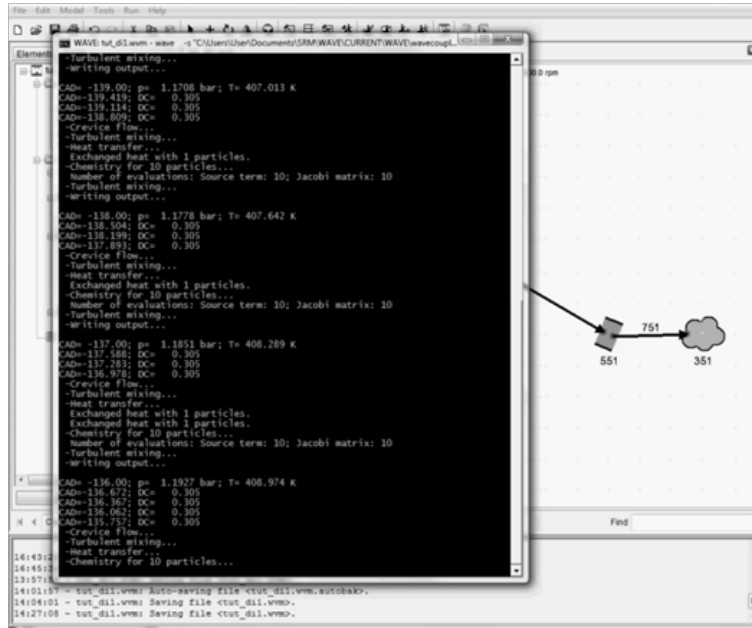
simulation software’s native visualization tools. Figures 13, 14, 15 show screenshots of the SRM SUITE being used with detailed chemistry within a leading industrial engine cycle simulation tool.



**Figure 13:** Engine map in OD/1D engine cycle simulation software using SRM SUITE as an external cylinder



**Figure 14:** Setting SRM SUITE input parameters within OD/1D engine cycle simulation software



**Figure 15:** SRM SUITE simulation, using detailed chemistry, running automatically within a coupled 0D/1D engine cycle simulation

## 5 Summary

- Detailed chemistry in the SRM provides excellent agreement for combustion and emissions in a fraction of the time for 3D CFD. Detailed chemistry requires less down-sampling to a larger number of particles in order to achieve convergence for combustion and emissions results.
- Modest parallelization, requiring only 8-16 cores, can substantially speed up simulations using detailed. Using 8 cores is 3.3 times faster than using a single core (for 128 particles), bringing the computational time from 36 mins to 10mins.
- Reduced chemistry dramatically cuts the computational time, but only provides trends for emissions results. Major down-sampling of the composition-space to a very small number of particles further reduces timescales and still achieves convergence of results.
- Parallelization is not really beneficial with reduced chemistry, since the chemistry can already be solved very quickly on a single-core PC due to the small number of species, reactions, and particles required.
- Since 0D/1D compatible timescales are achievable, the technologies identified can be integrated into full engine-cycle simulations, allowing for a seamless workflow whilst delivering detailed results in terms of fuels, combustion, and emissions.
- SRM, employing both detailed and reduced chemistry, has been coupled with most major 0D/1D engine cycle simulation tools, providing easy-to-use access to this technology directly through industry standard tools.

## References

- [1] *Ricardo Software WAVE 8.2 User's Manual*. 2009.
- [2] E. Abdi Aghdam, A. A. Burluka, T. Hattrell, K. Liu, C. G. W. Sheppard, J. Neumeister, and N. Crundwell. Study of cyclic variation in an si engine using quasi-dimensional combustion model. SAE Paper No. 2007-01-0939, 2007.
- [3] A. M. Aldawood, S. Mosbach, and M. Kraft. HCCI combustion phasing transient control by hydrogen-rich gas: Investigation using a fast detailed-chemistry full-cycle model. SAE Paper No. 2009-01-1134, 2009.
- [4] A. Bhave, M. Balthasar, M. Kraft, and F. Mauss. Analysis of a natural gas fuelled Homogeneous Charge Compression Ignition engine with exhaust gas recirculation using a stochastic reactor model. *Int. J. Engine Research*, 5(1):93–104, 2004.
- [5] A. Burluka, K. Liu, C. G. W. Sheppard, A. J. Smallbone, and R. Woolley. The influence of simulated residual and NO concentrations on knock onset for PRFs and gasolines. SAE Paper No. 2004-01-2998, 2004. <http://dx.doi.org/10.4271/2004-01-2998>.
- [6] L. Cao, H. Su, S. Mosbach, M. Kraft, and A. Bhave. Studying the influence of direct injection on PCCI combustion and emissions at engine idle condition using two dimensional CFD and stochastic reactor model. SAE Paper No. 2008-01-0021, 2008.
- [7] J. Etheridge, A. Bhave, A. J. Smallbone, A. Coble, S. Mosbach, and M. Kraft. Optimisation of injection strategy, combustion phasing and exhaust gas emissions using advanced simulation technologies, 2011.
- [8] J. E. Etheridge, S. Mosbach, M. Kraft, H. Wu, and N. Collings. A detailed chemistry multi-cycle simulation of a gasoline fueled HCCI engine operated with NVO. SAE Paper No. 2009-01-0130, 2009.
- [9] J. T. Farrell, N. P. Cernansky, F. L. Dryer, and C. K. Law. Development of an experimental database and kinetic models for surrogate diesel fuels. SAE Paper No. 2007-01-0201, 2007.
- [10] J. W. Fox, W. K. Cheng, and J. B. Heywood. A model for predicting residual gas fraction in spark-ignition engines. SAE Technical Paper No. 931025, 1993.
- [11] W. Gropp, E. Lusk, D. Ashton, P. Balaji, D. Buntinas, R. Butler, A. Chan, D. Goodell, J. Krishna, G. Mercier, R. Ross, R. Thakur, and B. Toonen. *MPICH2 User's Guide: Version 1.2.1*. 2009.
- [12] T. Hattrell, C. G. W. Sheppard, A. A. Burluka, J. Neumeister, and A. Cairns. Burn rate implications of alternative knock reduction strategies for turbocharged SI engines. SAE Paper No. 2006-01-1110, 2006.

- [13] J. Heywood. *Internal Combustion Engine Fundamentals*. McGraw Hill, 1988.
- [14] C. K. Law. Combustion at a crossroads: Status and prospects. *Proceedings of the Combustion Institute*, 2007.
- [15] T. Lu and C. K. Law. Toward accommodating realistic fuel chemistry in large-scale computations, 2009.
- [16] S. Mosbach, M. Kraft, A. Bhave, F. Mauss, J. H. Mack, and R. W. Dibble. Simulating a homogenous charge compression ignition engine fuelled with a DEE/EtOH blend. SAE Paper No. 2006-01-1362, 2006.
- [17] S. Mosbach, H. Su, M. Kraft, A. Bhave, F. Mauss, Z. Wang, and J.-X. Wang. Dual injection HCCI engine simulation using a stochastic reactor model. *International Journal of Engine Research*, 8(1):41–50, 2007.
- [18] S. Mosbach, A. M. Aldawood, and M. Kraft. Real-time evaluation of a detailed chemistry HCCI engine model using a tabulation technique. *Combustion Science and Technology*, 180(7):1263–1277, 2008.
- [19] S. Mosbach, M. Celnik, R. Abijeet, M. Kraft, H. Zhang, S. Kubo, and K. Kim. Towards a detailed soot model for internal combustion engines. *Combustion and Flame*, 156:1156 – 1165, 2009.
- [20] W. J. Pitz and C. J. Mueller. Recent progress in the development of diesel surrogate fuels. *Progress in Energy and Combustion Science*, 2010. <http://dx.doi.org/10.1016/j.pecs.2010.06.004>.
- [21] K. V. Puduppakkam, C. V. Naik, L. Liang, E. Meeks, and B. G. Bunting. Combustion and emissions modeling of a gasoline HCCI engine using model fuels. SAE Paper No. 2009-01-0669, 2009.
- [22] A. J. Smallbone, N. Morgan, A. Bhave, M. Kraft, R. F. Cracknell, and G. Kalghatgi. Simulating combustion of practical fuels and blends for modern engine applications using detailed chemical kinetics. SAE Paper No. 2010-01-0572, 2010. <http://dx.doi.org/10.4271/2010-01-0572>.
- [23] A. J. Smallbone, A. N. Bhave, A. R. Coble, S. Mosbach, M. Kraft, and R. M. McDavid. Identifying optimal operating points in terms of engineering constraints and regulated emissions in modern diesel engines. SAE Paper No. 2011-01-1388, 2011. <http://dx.doi.org/10.4271/2011-01-1388>.
- [24] G. P. Smith, D. M. Golden, M. Frenklach, N. W. Moriarty, B. Eiteneer, M. Goldenberg, C. Thomas Bowman, R. K. Hanson, S. Song, W. C. Gardiner, V. V. Lissianski, and G. Z. Qin. [http://www.me.berkeley.edu/gri\\_mech/](http://www.me.berkeley.edu/gri_mech/).
- [25] H. Su, A. Vikhansky, S. Mosbach, M. Kraft, A. Bhave, K.-O. Kim, T. Kobayashi, and F. Mauss. A computational study of an HCCI engine with direct injection during gas exchange. <http://dx.doi.org/10.1016/j.combustflame.2006.07.005>.



- [26] S. Subramaniam and S. B. Pope. A mixing model for turbulent reactive flows based on Euclidean minimum spanning trees. *Combustion and Flame*, 115:487–514, 1998. doi:10.1016/S0010-2180(98)00023-6.
- [27] H. Wang, X. You, A. V. Joshi, S. G. Davis, A. Laskin, F. Egolfopoulos, and C. K. Law. USC mech version ii. high-temperature combustion reaction model of h<sub>2</sub>/co/c<sub>1</sub>-c<sub>4</sub> compounds, May 2007. [http://ignis.usc.edu/USC\\_Mech\\_II.htm](http://ignis.usc.edu/USC_Mech_II.htm).

## Citation index

Abdi Aghdam et al. [2], 4  
Aldawood et al. [3], 4, 5, 8, 11, 14  
Bhave et al. [4], 4, 5, 8, 11, 14  
Burluka et al. [5], 4  
Cao et al. [6], 1, 4, 6, 8, 11  
Etheridge et al. [7], 9, 10  
Etheridge et al. [8], 10  
Farrell et al. [9], 4  
Fox et al. [10], 4  
Gropp et al. [11], 7  
Hattrell et al. [12], 4  
Heywood [13], 4  
Law [14], 4  
Lu and Law [15], 4  
Mosbach et al. [16], 4, 5, 8, 11, 14  
Mosbach et al. [17], 4, 5, 8, 11, 14  
Mosbach et al. [18], 4, 5, 8, 11, 14  
Mosbach et al. [19], 4, 5, 8, 11, 14  
Pitz and Mueller [20], 4  
Puduppakkam et al. [21], 4  
Ric [1], 4  
Smallbone et al. [22], 4  
Smallbone et al. [23], 5, 11  
Smith et al. [24], 4  
Su et al. [25], 4, 5, 8, 11, 14  
Subramaniam and Pope [26], 8  
Wang et al. [27], 4