A study of direct and Krylov iterative sparse solver
 techniques to approach linear scaling of the integration
 of Chemical Kinetics with detailed combustion
 mechanisms
 Federico Perini^{a,*}, Emanuele Galligani^b, Rolf D. Reitz^a,

^aUniversity of Wisconsin–Madison Engine Research Center, 1500 Engineering Drive, Madison, WI 53706, U.S.A.

^bDipartimento di Ingegneria Enzo Ferrari, Università di Modena e Reggio Emilia, strada Vignolese 905/A, 41125 Modena, Italy

10 Abstract

The integration of the stiff ODE systems associated with chemical ki-11 netics is the most computationally demanding task in most practical com-12 bustion simulations. The introduction of detailed reaction mechanisms in 13 multi-dimensional simulations is limited by unfavorable scaling of the stiff 14 ODE solution methods with the mechanism's size. In this paper, we com-15 pare the efficiency and the appropriateness of direct and Krylov subspace 16 sparse iterative solvers to speed-up the integration of combustion chemistry 17 ODEs, with focus on their incorporation into multi-dimensional CFD codes 18 through operator splitting. A suitable preconditioner formulation was ad-19 dressed by using a general-purpose incomplete LU factorization method for 20 the chemistry Jacobians, and optimizing its parameters using ignition delay 21 simulations for practical fuels. All the calculations were run using a same ef-22 ficient framework: SpeedCHEM, a recently developed library for gas-mixture 23 kinetics that incorporates a sparse analytical approach for the ODE system 24 functions. The solution was integrated through direct and Krylov subspace 25 iteration implementations with different backward differentiation formula in-26 tegrators for stiff ODE systems: LSODE, VODE, DASSL. Both ignition de-27 lay calculations, involving reaction mechanisms that ranged from 29 to 7171 28 species, and multi-dimensional internal combustion engine simulations with 29 the KIVA code were used as test cases. All solvers showed similar robustness, 30 and no integration failures were observed when using ILUT-preconditioned 31

Preprinter withombustion and Flame

Email addresses: perini@wisc.edu (Federico Perini),

November 16, 2013

6

8

9

emanuele.galligani@unimore.it (Emanuele Galligani), reitz@engr.wisc.edu (Rolf D. Reitz)

Krylov enabled integrators. We found that both solver approaches, coupled 32 with efficient function evaluation numerics, were capable of scaling computa-33 tional time requirements approximately linearly with the number of species. 34 This allows up to three orders of magnitude speed-ups in comparison with 35 the traditional dense solution approach. The direct solvers outperformed 36 Krylov subspace solvers at mechanism sizes smaller than about 1000 species, 37 while the Krylov approach allowed more than 40% speed-up over the direct 38 solver when using the largest reaction mechanism with 7171 species. 39

Keywords: chemical kinetics, sparse analytical Jacobian, Krylov iterative
 methods, preconditioner, ILUT, stiff ODE, combustion, SpeedCHEM

42 **1. Introduction**

The steady increase in computational resources is fostering research of 43 cleaner and more efficient combustion processes, whose target is to reduce 44 dependence on fossil fuels and to increase environmental sustainability of the 45 economies of industrialized countries [1]. Realistic simulations of combus-46 tion phenomena have been made possible by thorough characterization and 47 modelling of both the fluid-dynamic and thermal processes that define local 48 transport and mixing in high temperature environments, and of the chemical 49 reaction pathways that lead to fuel oxidation and pollutant formation [2]. De-50 tailed chemical kinetics modelling of complex and multi-component fuels has 51 been achieved through extensive understanding of fundamental hydrocarbon 52 chemistry [3] and through the development of semi-automated tools [4] that 53 identify possible elementary reactions based on modelling of how molecules 54 interact. These detailed reaction models can add thousands of species and 55 However, often simple phenomenological models are preferred reactions. 56 to even skeletal reaction mechanisms that feature a few tens/hundreds of 57 species in "real-world" multidimensional combustion simulations, due to the 58 expensive computational requirements that their integration requires, even 59 on parallel architectures, when using conventional numerics. The stiffness of 60 chemical kinetics ODE systems, caused by the simultaneous co-existence of 61 broad temporal scales, is a further complexity factor, as the time integrators 62 have to advance the solution using extremely small time steps to guarantee 63 reasonable accuracy, or to compute expensive implicit iterations while solv-64 ing the nonlinear systems of equations [5]. 65

66

Many extremely effective approaches have been developed in the last few 67 decades to alleviate the computational cost associated with the integration of 68 chemical kinetics ODE systems for combustion applications, targeting: the 69 stiffness of the reactive system; the number of species and reactions; the 70 overall number of calculations; or the main ordinary differential equations 71 (ODE) system dynamics, by identifying main low-dimensional manifolds [6– 72 16]. A thorough review of reduction methods is addressed in [2]. From the 73 point of view of integration numerics, some recent approaches have shown 74 the possibility to reduce computational times by orders of magnitude in com-75 parison with the standard, dense ODE system integration approach, without 76 introducing simplifications into the system of ODEs, but instead improving 77 the computational performance associated with: 1) efficient evaluation of 78 computationally expensive functions, 2) improved numerical ODE solution 79 strategy, 3) architecture-tailored coding. 80

As thermodynamics-related and reaction kinetics functions involve expen-81 sive mathematical functions such as exponentials, logarithms and powers, ap-82 proaches that make use of data caching, including in-situ adaptive tabulation 83 (ISAT) [17] and optimal-degree interpolation [18], or equation formulations 84 that increase sparsity [19], or approximate mathematical libraries [20], can 85 significantly reduce the computational time for evaluating both the ODE sys-86 tem function and its Jacobian matrix. As far as the numerics internal to the 87 ODE solution strategy, the adoption of sparse matrix algebra in treating the 88 Jacobian matrix associated with the chemistry ODE system is acknowledged 89 to reduce the computational cost of the integration from the order of the 90 cube number of species, $O(n_s^3)$, down to a linear increase, $O(n_s)$ [21–23]. 91

Implicit, backward-differentiation formula methods (BDF) are among 92 the most effective methods for integrating large chemical kinetics problems 93 [24, 25] and packages such as VODE [26], LSODE [27], DASSL [28] are widely 94 adopted. Recent interest is being focused on applying different methods that 95 make use of Krylov subspace approximations to the solution of the linear sys-96 tem involving the problem's Jacobian, such as Krylov-enabled BDF methods 97 [29], Rosenbrock-Krylov methods [30], exponential methods [31]. As far as 98 architecture-tailored coding is concerned, some studies [32, 33] have shown 99 that solving chemistry ODE systems on graphical processing units (GPUs) 100 can significantly reduce the dollar-per-time cost of the computation. How-101 ever, the potential achieveable per GPU processor is still far from full opti-102 mization, due to the extremely sparse memory access that chemical kinetics 103 have in matrix-vector operations, and due to small problem sizes in compari-104

son to the available number of processing units. Furthermore, approaches for
chemical kinetics that make use of standardized multi-architecture languages
such as OpenCL are still not openly available.

108

Our approach, available in the SpeedCHEM code, a recently developed re-109 search library written in modern Fortran language, addresses both issues 110 through the development of new methods for the time integration of the 111 chemical kinetics of reactive gaseous mixtures [18]. Efficient numerics are 112 coupled with detailed or skeletal reaction mechanisms of arbitrarily large 113 size, e.g., $n_s \geq 10^4$, to achieve significant speed-ups compared to traditional 114 methods especially for the small-medium mechanism sizes, $100 \le n_s \le 500$, 115 which are of major interest to multi-dimensional simulations. In the code, 116 high computational efficiency is achieved by evaluating the functions associ-117 ated with the chemistry ODE system throughout the adoption of optimal-118 degree interpolation of expensive thermodynamic functions, internal sparse 119 algebra management of mechanism-related quantities, and sparse analytical 120 formulation of the Jacobian matrix. This approach allows a reduction in 121 CPU times by almost two orders of magnitude in ignition delay calculations 122 using a reaction mechanism with about three thousand species [34], and was 123 capable of reducing the total CPU time of practical internal combustion en-124 gine simulations with skeletal reaction mechanisms by almost one order of 125 magnitude in comparison with a traditional, dense-algebra-based reference 126 approach [35, 36]. 127

In this paper, we describe the numerics internal to the ODE system solution 128 procedure. The flexibility of the software framework allows investigations of 129 the optimal solution strategies, by applying different numerical algorithms 130 to a numerically consistent, and computationally equally efficient, problem 131 formulation. Some known effective and robust ODE solvers, as LSODE [27], 132 VODE [26], DASSL [28], used for the combustion chemistry integration [18], 133 are based on backward differentiation formulae (BDF) that include implicit 134 integration steps, requiring repeated linear system solutions associated with 135 the chemistry Jacobian matrix, as part of an iterative Newton procedure. 136 Based on the same computational framework and BDF integration proce-137 dure, the effective efficiency of the sparse direct and preconditioned iterative 138 Krylov subspace solvers reported in Table 1 was studied. To consider the 139 effects of reaction mechanism size, a common matrix of reaction mechanisms 140 for typical hydrocarbon fuels and fuel surrogates, spanning from $n_s = 29$ up 141 to $n_s = 7171$ species was used, whose details are reported in Table 1. 142

The paper is structured as follows. A description of the modelling ap-143 proach adopted for the simulations is reported, including the major steps of 144 the BDF integration procedure. Then, a robust preconditioner for the chem-145 istry ODE Jacobian matrix is defined, as the result of the optimization of a 146 general-purpose preconditioner-based incomplete LU factorization [37, 38]. 147 Integration of ignition delay calculations at conditions relevant to practi-148 cal combustion systems is compared for the range of reaction mechanisms 149 chosen, at both solver techniques and with different BDF-based stiff ODE 150 integrators. Finally, the direct and Krylov subspace solvers are compared for 151 modeling practical internal combustion engine simulations. 152

Mechanism	ns	nr	fuel	composition	ref.
ERC n-heptane	29	52	n-heptane	$[nC_7H_{16}:1.0]$	[39]
ERC PRF	47	142	PRF25	$[nC_7H_{16}:0.75,iC_8H_{18}:0.25]$	[40]
Wang PRF	86	392	PRF25	$[nC_7H_{16}:0.75,iC_8H_{18}:0.25]$	[41]
ERC multiChem	128	503	PRF25	$[nC_7H_{16}:0.75,iC_8H_{18}:0.25]$	[42]
LLNL n-heptane (red.)	160	1540	n-heptane	$[nC_7H_{16}:1.0]$	[43]
LLNL n-heptane (det.)	654	2827	n-heptane	$[nC_7H_{16}:1.0]$	[44]
LLNL PRF	1034	4236	PRF25	$[nC_7H_{16}:0.75,iC_8H_{18}:0.25]$	[44]
LLNL mehtyl-decanoate	2878	8555	methyl-decanoate	[md:1.0]	[34]
LLNL n-alkanes	7171	31669	diesel surrogate	$\left[nC_7H_{16}:0.4, nC_{16}H_{34}:0.1, nC_{14}H_{30}:0.5\right]$	[45]

×
÷ď
Ē
÷
U
s,
Ξ.
÷
-
.Ħ
50
20
.Ħ
÷
ð
÷
Ч
0
Ŧ
Ч
é.
5
-
ŝ
В
5
٠Ħ
H
La
-
ă
ā
H
d
5
· 🗄
G
ā
ē
Ч
e)
Ч
÷
Ļ
0
⊳
5
-Ĕ
2
E
5
Ć
\cup
••
d)
Ĭ
9
-a
H

Description	ODE integration method	order	Linear system solution solution	Preconditioning
LSODE, direct	fixed-coefficient BDF	variable	direct, dense	no
LSODES, direct sparse	fixed-coefficient BDF	variable	direct, sparse	no
LSODKR, Krylov	fixed-coefficient BDF	variable	scaled GMRES	ILUT
VODE, direct sparse	variable-coefficient BDF	variable	direct, sparse	no
VODPK, Krylov	variable-coefficient BDF	variable	scaled GMRES	ILUT
DASSL, direct sparse	DAE form, variable-coefficient BDF	variable	direct, sparse	no
DASKR, Krylov	DAE form, variable-coefficient BDF	variable	scaled GMRES	ILUT
Table 2: Ove	rview of the chemical kinetics ODE solution :	strategies co	ompared in this study.	
		0	· Constant second war and no di	

-	study.
•	this
	Ξ
-	compared
	strategies
•	solution
	JUD
	netics UDE
	kinetics UDE
	emical kinetics UDE
	e chemical kinetics UDE
	the chemical kinetics UDE
	of the chemical kinetics UDE
	verview of the chemical kinetics UDE
	Uverview of the chemical kinetics UDE
	3.2: Uverview of the chemical kinetics UDE

¹⁵³ 2. Description of the numerical methodology

¹⁵⁴ 2.1. A sparse, analytical Jacobian framework for chemical kinetics

Significant advancements in combustion chemistry codes have occurred 155 in the last few years, but comparison with earlier implementations [46] is not 156 a straightforward task. Often access to the source codes is unavailable, and 157 insufficient description of the numerical algorithms or to the architecture-158 tailored optimizations is provided. Furthermore, the size of the chemical ki-159 netics problem is typically for small-medium systems $(n_s \leq 10^4)$ [47], which 160 makes the adoption of advanced parallel algorithms quite inefficient. Thus, 161 meaningful comparisons of different numerical methods should be done adopt-162 ing a same code framework, with the same level of optimization, on the same 163 computing architecture. 164

The SpeedCHEM chemistry code [18] is a computationally efficient chemical 165 kinetics library developed by the authors. The library, written in Modern 166 Fortran standard [48], and compiled with the portable Fortran frontend of 167 the GNU suite of compilers [49], exploits a sparse, analytical formulation for 168 all the chemical and thermodynamic functions involved in the description 160 of the reacting system. The library includes thermal property calculation, 170 mixture states, and reaction kinetics, for the analytical calculation of the 171 Jacobian matrix associated with the chemical kinetic ODE system. Inter-172 nal BLAS routines have been specifically developed in order not to rely on 173 architecture-dependent or copyright-protected libraries, while exploiting the 174 vectorization features of the language for most low-level code tuning, and 175 leaving most of architecture-related optimization to the compiler. 176 177

Still a number of low-level optimizations are needed to enable efficient evalua-178 tion of the functions that describe the evolution of the species in the reacting 179 environment, as well as their derivatives. Polymorphic procedure handling 180 has been implemented to enable compliance with a number of different im-181 plicit stiff ODE integrators, including VODE [26], LSODE [27], DASSL [28], 182 RADAU5 [5], RODAS [5]. Where not originally available, capability for a 183 direct, sparse solution of the linear system associated with the implicit ODE 184 integration step has been implemented using the code's internal sparse alge-185 bra library. 186

Chemical Kinetics problem. We refer to the *Chemical Kinetics problem* as the initial value problem (IVP) where the laws that define temporal evolution of

mass fractions of a set of species in a reactive environment are described as the result of the reactions involving them. Considering a set of n_s species, $i = 1, ..., n_s$ and n_r reactions, $j = 1, ..., n_r$, every generic reaction is defined as the interaction between a set of reactant species and a set of products:

$$\sum_{i=1}^{n_s} \nu'_{j,i} M_i \rightleftharpoons \sum_{k=1}^{n_s} \nu''_{j,k} M_k, \ j = 1, ..., n_r;$$
(1)

where M_i represents species names, and matrices ν' and ν'' contain the stoichiometric coefficients of the species involved in the reactions. As the number of species involved in every reaction is typically small, both these matrices can be extremely sparse. Every reaction features a reaction progress variable rate, q_j , that defines, for the given thermodynamic state of the reactive system, its tendency towards (forward, q_f) producing new products or (backward, q_b) producing reactants:

$$q_j = q_{f,j} - q_{b,j} = k_{f,j} \prod_{i=1}^{n_s} \left(\frac{\rho Y_i}{W_i}\right)^{\nu'_{j,i}} - k_{b,j} \prod_{i=1}^{n_s} \left(\frac{\rho Y_i}{W_i}\right)^{\nu''_{j,i}},\tag{2}$$

where $k_{f,j}$ and $k_{b,j}$ are reaction rate constants, of the type either k = f(T), k = f(T, p) or $k = f(T, p, \mathbf{Y})$, depending on the complexity of the reaction formulation. The products that involve species concentrations C_i are calculated using mass fractions Y_i , molecular weights W_i , and reactor density ρ . The rate of change of every species' mass fraction is given by:

$$\frac{dY_i}{dt} = \dot{Y}_i = \frac{W_i}{\rho} \sum_{j=1}^{n_r} \left[\left(\nu''_{j,i} - \nu'_{j,i} \right) q_j \right].$$
(3)

¹⁹⁹ Mass conservation, $\sum_{i} \dot{Y}_{i} = 0$ or $\sum_{i} Y_{i} = 1$, is guaranteed if atom balance ²⁰⁰ in every reaction is verified, i.e., $\sum_{i} \nu'_{j,i} E_{k,i} = \sum_{i} \nu''_{j,i} E_{k,i} \forall j = 1, ..., n_{r}, k =$ ²⁰¹ 1, ..., n_{e} , where the matrix $E_{k,i}$ contains the number of atoms of element k in ²⁰² species i. While the chemical kinetics equations can be incorporated in any ²⁰³ reactive environment formulation, a constant-volume reactor is used, as this ²⁰⁴ is the assumption that is typically adopted when doing operator splitting in ²⁰⁵ multi-dimensional combustion codes [50–53], or for ignition delay calculations ²⁰⁶ in shock tubes:

$$\frac{dT}{dt} = \dot{T} = -\frac{1}{\bar{c}_v} \sum_{i=1}^{n_s} \frac{U_i \dot{Y}_i}{W_i}.$$
(4)

Here, \bar{c}_v represents the mass-based mixture averaged constant volume specific heat, and $U_i = U_i(T)$ are the molar internal energies of the species at temperature T. Note however that the approach is also applicable to the sparser, constant-pressure reactor formulation [19].

Sparse matrix algebra. Elementary reactions according to the theory of ac-211 tivated complex represent the probability that collisions among molecules 212 result in changes in the chemical composition of the colliding molecules [54]. 213 Practical combustion mechanisms can incorporate third bodies, but most re-214 actions typically follow the common initiation, propagation, branching, ter-215 mination pathway structure. This means that, while most species interact 216 with the hydrogen and oxygen radicals of the hydrogen combustion mecha-217 nism, large hydrocarbon classes have typically just few interactions among 218 themselves. Looking at the Jacobian matrix \mathbf{J} of different reaction mecha-219 nisms, as represented in Figure 1, it is clear that the basic hydrogen combus-220 tion mechanism, that typically involves the 10-20 smaller species depending 221 on the formulation [55], is almost dense, and can represent a significant part 222 of the overall reaction mechanism especially when it is very skeletal. How-223 ever, as more hydrocarbon classes are included in the mechanism, the overall 224 sparsity increases significantly, and can be more than 99.5% for the large 225 reaction mechanisms. 226

227

Accordingly, a sparse matrix algebra library was coded to model the stoi-228 chiometric index matrices $\nu', \nu''(n_r \times n_s)$, the element matrix $\mathbf{E}(n_s \times n_e)$, the 229 ODE system Jacobian matrix $\mathbf{J}(n_s + 1 \times n_s + 1)$, as well as all the inter-230 nal calculation matrices involving analytical derivatives such as $\partial \mathbf{q} / \partial \mathbf{Y}$, for 231 instance. The Compressed Sparse Row (CSR) format [56, Chapt. 2] was 232 adopted as the standard to model matrix-vector calculations, matrix-matrix 233 calculations, as well as other overloaded operator functions including sum, 234 scalar-matrix multiplication, comparison and assignment. Even with dynam-235 ically allocated matrix structures, the sparsity structures of the matrices are 236 only calculated once at the reaction mechanism initialisation, as they do not 237 change. It was chosen not to consider any dynamic reaction mechanism re-238 duction, to avoid any additional error in the solution procedure. Finally, full 239 encapsulation of the overloaded methods allows their easy implementation in 240 other architectures such as GPUs. 241

242

As most ODE integrators for solving stiff problems perform implicit inte-



Figure 1: Sparsity of the Jacobian matrix versus reaction mechanism dimension.

gration steps that require evaluation of the Jacobian matrix and the solution 244 of a linear system, or its composition/scaling with some diagonal matrix (see, 245 e.g., $[5, \S8.8]$, routines that perform direct solution of a sparse linear system 246 were implemented using a standard four-step solution method: 1) optimal 247 matrix ordering using the Cuthill-McKee algorithm [57, 58, §4.3.4], 2) sym-248 bolic sparse LU factorization, 3) numerical LU factorization, 4) solution by 249 forward and backward substitution. The first two steps of the procedure are 250 only evaluated at the problem initialization, as they only involve the matrix 251 sparsity pattern. 252

Figure 2 shows the effect of optimal matrix reordering to reduce fill-in 253 during matrix factorization of a large reaction mechanism. Even if the origi-254 nal sparsity of the Jacobian matrix is extremely high, the LU decomposition 255 destroys almost all of it when rows and columns are not correctly ordered. 256 The fill-in phenomenon, i.e. the number of elements in a matrix that switch 257 from having a zero value to a non-zero value after application of an algorithm 258 [59], cannot be avoided when using direct solution methods, and tends to be-259 come significant for large reaction mechanisms. However, if the matrix is 260 properly ordered, the number of non-zero elements in the LU decomposition 261 can be not bigger than 2-3 times the number of non-zero elements in the non-262 decomposed matrix. After the Jacobian matrix has been decomposed into a 263

lower and an upper triangular matrix, there is no need for matrix inversion
as the linear system can be eventually solved in two steps by simple forward
and backward substitution:

$$\mathbf{A}x = b \longrightarrow \mathbf{L}\mathbf{U}x = b \longrightarrow y = \mathbf{L}^{-1}b \longrightarrow x = \mathbf{U}^{-1}y.$$
(5)

Maximizing the use of sparse algebra, and consistently using the same 267 CSR format for all calculations throughout the code, including for building 268 the analytical Jacobian matrix, allows efficient memory usage and minimiza-269 tion of cache misses, that have an impact on the total CPU time. This 270 provides a universal formulation that does not require reaction-mechanism-271 dependent tuning, and avoids the use of symbolic differentiation tools or 272 reaction-mechanism-dependent routines that would require recompilation of 273 the code for every different reaction mechanism used. 274

Optimal-degree interpolation. Optimal-degree interpolation was implemented 275 as a procedure to speed-up the evaluation of thermodynamic functions that 276 involve both species and reactions. Most thermodynamic- and chemistry-277 related are temperature-dependent functions, and include powers, exponen-278 tials, fractions, that require many CPU clock cycles to be completed [2]. 279 Functions describing reaction behavior span from simple reaction rate formu-280 lations such as the universally adopted extended Arrhenius form, to equilib-281 rium constant formulations that are required to compute backward reaction 282 rates, to more complex reaction parameters such as the fall-off centering fac-283 tor in Troe pressure-dependent reactions [60]. Functions describing species 284 and mixture thermodynamics, such as the JANAF format laws adopted in 285 the present code, typically feature polynomials of temperature powers and 286 logarithms. 287

As memory availability is not a restriction in modern computing architectures, caching values of these functions at discretized temperature points, and interpolating using simple interpolation functions, proved to be extremely efficient.

Also the derivatives with respect to temperature were evaluated analytically and tabulated at initialization at equally-spaced temperature intervals of $\Delta T = 10K$. This tabulation allows different degrees of the interpolation, to be calculated only on the basis of the relative position of the current temperature value within the closest neighboring points [18]:

$$f(T) = (T_{n+1} - T) / (T_{n+1} - T_n) = (T_{n+1} - T) / \Delta T.$$
(6)



Figure 2: Sparsity patterns of the Jacobian matrix of the LLNL PRF mechanism [44], non-permuted (\mathbf{J}) and permuted (\mathbf{J}') using the Cuthill-McKee algorithm, and of their sparse LU decompositions.

At the initialization of every function's table, an optimal degree of interpolation is cached, defined as the minimum interpolation degree that allows the interpolation procedure to predict the function's values within a required relative accuracy, set as $\varepsilon_{tab} \leq 10^{-6}$ in this study, corresponding to 6 exact significant digits. The optimal-degree interpolation procedure also prevents Runge's oscillatory phenomenon by imposing an accuracy check on the interpolated function.

Performance comparisons between interpolated and analytically evaluated 304 thermodynamic functions are shown in Figure 3 using the reaction mecha-305 nisms in Table 1. The speed-up factor achieved in CPU time when calling dif-306 ferent thermodynamics-related functions using the tabulation-interpolation 307 procedure rather than their analytical formulation is plotted against the num-308 ber of species in the reaction mechanism, n_s . The plot shows that significant 309 speed-ups of up to about two orders of magnitude are achieved for the most 310 complex reaction properties, even for equilibrium-based reversible reactions 311 and Troe fall-off reactions, whose number in a reaction mechanism can be 312 much smaller than the total number of reactions. The speed-up achieved 313 in evaluating species thermodynamic functions, whose nature is already of a 314 polynomial formulation in the JANAF standard, is smaller but quickly adds 315 up to about 2 times for practical reaction mechanisms. 316

Overall code performance. Computational scaling of the most expensive func-317 tions associated with the integration of the chemical kinetics problem is re-318 ported in Figure 4. The results, taken on a desktop PC running on an Intel 319 Core i5 processor with 2.40 GHz clock speed, and with 667 MHz DDR3 RAM, 320 show less than linear computational performance scaling with the problem di-321 mensions of both the ODE vector function r.h.s. and of the Jacobian matrix, 322 and approximately linear scaling, $O(n_s)$, for the complete sparse numerical 323 factorization and solution of the linear system associated with the Jacobian 324 matrix itself. The significant computational efficiency achieved through scal-325 ability and low-level optimization of the sparse analytical Jacobian approach 326 is also evident as the CPU time required by Jacobian evaluation is always 327 only between 5.4 and 12.0 times greater than the time needed for evaluat-328 ing the ODE system r.h.s. throughout all the reaction mechanisms tested. 329 without any changes to the compiled code. With a standard finite difference 330 approach computational cost is estimated to be between 29 and 7200 times 331 greater than the r.h.s. evaluation for the smallest and the largest reaction 332 mechanisms in the group of Table 1. Note that the number of reactions 333



Figure 3: Performance scaling of optimal-degree interpolation vs. analytical evaluation of some temperature-dependent functions.

plays a direct role in the CPU time requirements of the ODE system func-334 tion and the analytical Jacobian, due to presence of reaction-based functions 335 and sparse matrix-vector (spMxV) multiplications involving $(n_r \times n_s)$ matri-336 ces. However, Figure 4 shows that the cost of solving the linear system with 337 the Jacobian, which is species-based, is usually dominant. Furthermore, in 338 order to have a number of reactions large enough to break the linear scaling 339 of the sparse algorithms on the Jacobian, n_r would have to scale at least 340 with $O(n_s^2)$, which would unrealistically correspond to having about $\approx 10^6$ 341 reactions per $\approx 10^3$ species. 342

The code results were previously validated against both homogeneous reactor simulations and multi-dimensional RANS simulation of engine combustion [18, 35], where it was shown that no differences could be observed between the SpeedCHEM solution and the solution obtained using a reference research chemistry code, when using the same integration method and tolerances.

Profiling of ignition delay calculations showed how much of the total CPU time was spent among the most important tasks, as represented in Figure 5. The linear system solution was still the most computationally demanding task across all mechanism sizes, requiring from roughly 40% up to more than 80% of the total CPU time as the problem dimensions increase. All the other computationally demanding tasks, such as thermodynamic property evalu-



Figure 4: Computational time cost of the functions associated with the ODE system for different reaction mechanisms: r.h.s, Jacobian matrix, direct solution of the linear system (numerical factorization, forward and backward substitution).

ations, sparse matrix algebra, law of mass action, that make up the most
expensive parts of the r.h.s and Jacobian function calls, demanded less than
20% of the total CPU time, with the rest of the time devoted to internal
ODE solver operation, as well as dynamic memory allocation.

Finally, RAM memory requirements were found to scale slightly less-than-358 linearly with mechanism dimensions, as an effect of increased Jacobian spar-359 sity for larger hydrocarbon dimensions. Tabulation of temperature-dependent 360 functions appears to be the most storage-demanding feature, whose require-361 ments linearly increase with the number of species (for all the mechanisms 362 tested, the numbers of reactions always followed Lu and Law's rule-of thumb, 363 $n_r \approx 5n_s$, [2]). However, only about 86MB of RAM were needed by the 364 SpeedCHEM solver at the largest mechanism. Considering that tabulated 365 data requirements undergo read-only access during the simulations and are 366 shared when parallel multi-dimensional calculations are run, these require-367 ments appear to be negligible compared to current memory availability, even 368 on personal desktop machines. 360



Figure 5: Subdivision of total CPU time among internal procedures during time integration of different reaction mechanisms.



Figure 6: RAM storage requirements of the SpeedCHEM code with and without tabulated data for optimal-degree interpolation. LSODE solver storage featured.

370 3. Solution of the chemistry ODE Jacobian using Krylov subspace 371 methods

Krylov subspace methods represent a class of iterative methods for solving 372 linear systems that feature repeated matrix-vector multiplications instead of 373 using matrix decomposition approaches of the stationary iterative methods 374 (see, e.g., [47, 61]). These solvers find extensive applications in problems 375 where building the system solution matrix is expensive, or the matrix it-376 self is even not known, or where sufficient conditions [62, Chapt. 7] for the 377 convergence of the iteration matrix of the stationary iterative methods are 378 not known. Their practical applicability is however constrained to special 379 matrix properties, such as symmetry, or to the need to find appropriate pre-380 conditioning of the linear system matrix, to guarantee convergence of the 381 method [61, Chapt. 2,3]. Iterative methods such as GMRES [63] are guar-382 anteed convergence to the exact solution in n iterations, n being the max-383 imum Krylov subspace dimension, equal to the linear system dimensions. 384 Thus, if no strategy to accelerate their convergence rate towards the solu-385 tion within a reasonable tolerance is used, their computational performance 386 can be bad. Furthermore, round-off and truncation errors intrinsic to float-387 ing point arithmetic can make the iterations numerically unstable, especially 388 when the problem is ill-conditioned, i.e., its solution is extremely sensitive to 389 perturbations. These problems have limited their implementation in widely 390 available ODE integration packages. 391

392

393 3.1. ODE integration procedure

Backward differentiation formulae (BDF) methods [64] are amongst the 394 most reliable stiff ODE integrators currently available, due to their ability to 395 allow large integration time steps, even in presence of steep transients and the 396 simultaneous co-existence of multiple time scales. The methods exploit the 397 past behavior of the solution for improving the prediction of its future values. 398 BDF methods are incorporated in many software packages for the integration 399 of systems of ordinary differential equations (ODE) and differential-algebraic 400 equations (DAE), including the LSODE package, the base integrator in this 401 study. Given a general initial value problem (IVP) for a vector ODE function 402 \mathbf{f} in the unknown \mathbf{y} of the form: 403

$$\dot{\mathbf{y}} = \mathbf{f}(t, \mathbf{y}), \, \mathbf{y}(t=0) = \mathbf{y}_{\mathbf{0}},\tag{7}$$

BDF methods build up the solution at successive steps in time, $t_n, n = 1, 2, ...,$ based on the solution at q previous steps $t_{n-j}, j = 1, 2, ..., q$, with the following scheme:

$$\mathbf{y}_{\mathbf{n}} = \mathbf{y}(t_n) = \sum_{j=1}^{q} \alpha_j \mathbf{y}_{\mathbf{n}-\mathbf{j}} + h\beta_0 \dot{\mathbf{y}}_{\mathbf{n}}, \tag{8}$$

where the number of previous steps, q, also defines the order of accuracy of the method, α_j and β_0 are method-specific coefficients that depend on the order of the integration, and $h = \Delta t = t_n - t_{n-1}$ is the current integration time-step. These methods are implicit as the evaluation of the ODE system function at the current timestep, $\dot{\mathbf{y}}_{\mathbf{n}} = \mathbf{f}(t_n, \mathbf{y}_{\mathbf{n}})$, is needed to derive the solution. Thus, a nonlinear system of equations has to be solved, which accounts for most of the total CPU time during the integration:

$$\mathbf{F}_{\mathbf{n}}(\mathbf{y}_{\mathbf{n}}) = \mathbf{y}_{\mathbf{n}} - h\beta_0 \mathbf{f}(t_n, \mathbf{y}_{\mathbf{n}}) - \mathbf{a}_{\mathbf{n}} = 0, \qquad (9)$$

where $\mathbf{a_n}$ is the sum of terms involving the previous time steps, $\mathbf{a_n} = \sum_{j=1}^q \alpha_j \mathbf{y_{n-j}}$. VODE and DASSL incorporate variable α_j and β_0 coefficients whose values are updated based on the previous step sizes to allow for more effective reusage of the previous values of the solution. However, the fixed coefficient form is more desirable in this study. LSODE handles the solution in Equation (9) through a variable substitution, $\mathbf{x_n} = h\dot{\mathbf{y_n}} = (\mathbf{y_n} - \mathbf{a_n})/\beta_0$, as follows:

$$\mathbf{F}_{\mathbf{n}}(\mathbf{x}_{\mathbf{n}}) = \mathbf{x}_{\mathbf{n}} - h\mathbf{f}(t_n, \mathbf{a}_{\mathbf{n}} + \beta_0 \mathbf{x}_{\mathbf{n}}) = 0.$$
(10)

This nonlinear system has to be solved using some version of the Newton-Raphson iterative procedure:

$$\widetilde{\mathbf{x}}_{\mathbf{n}}^{(1)} = \text{initial guess;}$$
(11)

 $\forall m = 1, 2, ..., (\text{until convergence}):$

$$\mathbf{F}_{\mathbf{n}}'(\tilde{\mathbf{x}}_{\mathbf{n}}^{(\mathbf{m})}) = \mathbf{I} - h\beta_0 \mathbf{J}(t_n, \mathbf{a}_n + \beta_0 \tilde{\mathbf{x}}_{\mathbf{n}}^{(\mathbf{m})}), \qquad (12)$$

solve:
$$\mathbf{F}'_{\mathbf{n}}(\tilde{\mathbf{x}}_{\mathbf{n}}^{(\mathbf{m})})\mathbf{r}_{\mathbf{n}} = -\mathbf{F}_{\mathbf{n}}(\tilde{\mathbf{x}}_{\mathbf{n}}^{(\mathbf{m})}),$$
 (13)
 $\tilde{\mathbf{x}}_{\mathbf{n}}^{(\mathbf{m}+1)} = \tilde{\mathbf{x}}_{\mathbf{n}}^{(\mathbf{m})} + \mathbf{r}_{\mathbf{n}}^{(\mathbf{m})};$

where the procedure is typically simplified by assuming that the Jacobian matrix does not change across the iterations (simplified-Newton method), thus $\mathbf{F}'(\mathbf{\tilde{x}_n^{(m)}}) \approx \mathbf{F}'(\mathbf{\tilde{x}_n^{(1)}})$, and Equation (13) provides the actual linear system that has to be solved throughout the whole time integration process. The first guess $\mathbf{\tilde{x}_n^{(1)}}$ for the solution is estimated applying an explicit integration step (predictor), which is then progressively modified towards the correct solution by the Newton procedure (corrector).

429 3.2. Iterative linear system solution

The introduction of preconditioned Krylov subspace linear system solution methods into a so-called Newton-Krylov iteration, enters in Equation (13) without modifying the structure of the iterative Newton procedure, as the solution found by the Krylov subspace method is some approximation of the correct solution, and satisfies:

$$\mathbf{F}_{\mathbf{n}}'(\tilde{\mathbf{x}}_{\mathbf{n}}^{(\mathbf{m})})\mathbf{r}_{\mathbf{n}} = -\mathbf{F}_{\mathbf{n}}(\tilde{\mathbf{x}}_{\mathbf{n}}^{(\mathbf{m})}) + \boldsymbol{\varepsilon}_{\boldsymbol{n}}, \qquad (14)$$

where the residual $\varepsilon_{\mathbf{n}}$ represents the amount by which the approximate solution $\mathbf{r}_{\mathbf{n}}$ found by the Krylov method fails to match the Newton iteration equation. It has been demostrated [65] that a simple condition, such as the imposition of a tolerance constraint, $\varepsilon_{\mathbf{n}} < \delta$, is enough to obtain any desired degree of accuracy in $\tilde{\mathbf{x}}_{\mathbf{n}}^{(\mathbf{m})}$ at the end of the Newton procedure. The solution $\mathbf{r}_{\mathbf{n}}$ of Equation (14) can be seen as a step of the Inexact Newton method [66] if it satisfies

$$\|\boldsymbol{\varepsilon}_{\boldsymbol{n}}\| \le \eta_n \left\| \mathbf{F}_{\mathbf{n}}(\tilde{\mathbf{x}}_n^{(m)}) \right\|$$
(15)

for a vector norm ||.||. The parameter η_n is called a *forcing term* and it plays a crucial role to avoid the oversolving phenomenon, i.e., the computation of unnecessary iterations of the inner iterative solver (such as the Krylov method) [67]. The convergence of Inexact Newton methods is guaranteed under standard assumptions for **F** (see, e.g., [61, pp. 68,90],[68, p. 70]).

The Krylov subspace method to iteratively solve the linear system of Eq. 448 (14) was a scaled version of the GMRES solver by Saad [63] in LSODKR. The 449 algorithm solves general linear systems in the form $A\mathbf{x} = \mathbf{b}$, where A is a large 450 and sparse $n \times n$ matrix, and **x** and **b** are *n*-vectors. The idea underlying 451 the Krylov subspace methodology is that of refining the solution array by 452 starting with an initial guess \mathbf{x}_0 , and updating it at every iteration towards 453 a more accurate solution. The linear system can be written in residual form, 454 i.e., the exact solution is a deviation $\mathbf{x} = \mathbf{x}_0 + \mathbf{z}$ from the initial estimate 455

(typically, $\mathbf{x_0} = 0$), and given the initial residual $\mathbf{r_0} = \mathbf{b} - A\mathbf{x_0}$ we have $A\mathbf{z} = \mathbf{r_0}$. If the Krylov subspace $\kappa_k(A, \mathbf{r_0})$ generated at the k-th iteration by the matrix and the residual r.h.s. vector is the one defined by the images of $\mathbf{r_0}$ under the first k powers of A,

$$\kappa_k(A, \mathbf{r_0}) = \operatorname{span}\left\{\mathbf{r_0}, A\mathbf{r_0}, A^2\mathbf{r_0}, \dots, A^{k-1}\mathbf{r_0}\right\},\tag{16}$$

it can be demonstrated that the solution to the linear system lies in the Krylov subspace of the problem's dimension, $\kappa_n(A, \mathbf{r_0})$. Thus, GMRES looks for the closest vector in this space to the solution, as the vector that satisfies the least squares problem:

2

$$\min_{\mathbf{z}\in\kappa_k(A,\mathbf{r_0})} \|\mathbf{r_0} - A\mathbf{z}\|.$$
(17)

The least squares problem of Equation (17) is solved using Arnoldi's method 464 [69], an algorithm that incrementally builds an orthonormal basis for $\kappa_k(A, \mathbf{r_0})$ 465 at every iteration, i.e., a set of orthogonal vectors starting from a set of 466 linearly independent vectors. The least squares problem can be solved in 467 this equivalent, well-defined orthonormal space rather than on the undefined 468 Krylov subspace. Overall, the number of operations per Krylov subspace iter-469 ation is expensive, as they involve a minimization process where the number 470 of unknowns, n, is greater than the current dimensions of the process, k. 471 However, this procedure can be effective if convergence to a desired tolerance 472 can be reached in a number of Krylov iterations that is much smaller than 473 the overall dimensions of the linear system. In practice, the convergence rate 474 can only be accelerated to practical numbers of iterations if some sort of 475 preconditioning is used (see, e.g., [70, 71]). 476

477 3.3. Preconditioning of the linear system

The practical applicability of preconditioned Krylov subspace methods typically depends much more on the quality of the preconditioner than on the Krylov subspace solver [63]. Preconditioning in the iterative solution implies computing the linear solution on a different, equivalent formulation, i.e., by left-multiplication of boths sides of the equation by a (left) preconditioner matrix:

$$P^{-1}A\mathbf{x} = P^{-1}\mathbf{b}, or \tag{18}$$
$$\bar{A}\mathbf{x} = \bar{\mathbf{b}}.$$

The role of applying a preconditioner matrix is thus that of reducing the span of the eigenvalues of A, and substituting it with a new matrix that is possibly closer to identity, $P^{-1} \approx A^{-1}$, $\bar{A} \approx I$, for which convergence would be in a single step. However, every step of the Newton procedure of Equation (14) will take the following form:

$$[P_n^{-1}\mathbf{F}'_{\mathbf{n}}(\tilde{\mathbf{x}}_{\mathbf{n}}^{(\mathbf{m})})]\mathbf{r}_{\mathbf{n}} = [P_n^{-1}(-\mathbf{F}_{\mathbf{n}}(\tilde{\mathbf{x}}_{\mathbf{n}}^{(\mathbf{m})}))] + \boldsymbol{\varepsilon}_{\boldsymbol{n}}.$$
(19)

Thus, an additional direct linear system solution needs to be solved at every 489 Newton iteration in the following form: $\mathbf{c} = P_n^{-1}(-\mathbf{F}_n(\tilde{\mathbf{x}}_n^{(m)}))$. From Equa-490 tion (19) it is evident that it is extremely important that the appropriate 491 preconditioner has a much simpler structure than the original matrix, so 492 that its factorization and solution process for the r.h.s. can be kept as simple 493 as possible. But, it has to approximate the inverse matrix sufficiently well 494 so that the convergence of the Krylov procedure is improved in comparison 495 to the non-preconditioned algorithm. 496

497

Ultimately, if the problem's eigenvalues span many orders of magnitude or 498 the preconditioner is too rough, the number of iterations needed for the so-499 lution may make the method more expensive than a standard, direct solver, 500 and the solver may fail, even if a solution exists. On the other hand, if the 501 preconditioner provides a too accurate representation of the inverse of the 502 linear system matrix, i.e., $P \approx A, P^{-1}A \approx I$, the dimensions of the Krylov 503 subspace, and similarly the method's convergence rate, will tend to reduce. 504 and the cost of building the preconditioned matrix and solving the r.h.s will 505 typically be unviable. 506

507

It has been shown that simple preconditioners such as Jacobi diagonal 508 preconditioning are not accurate enough for chemical kinetics problems [72]. 509 Thus, we have chosen the Incomplete LU preconditioner with dual Trunca-510 tion strategy (ILUT) by Saad [38, 47]. This preconditioner approximates the 511 original sparse matrix A by an incomplete version of its LU factorization. 512 The factorization is incomplete in the sense that not all of the elements that 513 should be included in the L and U triangular matrices are kept in the in-514 complete representation. This can significantly reduce the problem of fill-in, 515 both making the factorization much sparser and faster to compute, while 516 still maintaining the LU representation of matrix A within a reasonable de-517 gree of approximation. Elements are 'dropped off' the decomposition in case: 518

⁵¹⁹ 1) their absolute value is smaller than the corresponding diagonal element ⁵²⁰ magnitude, by a certain *drop-off* tolerance δ ; 2) only the l_{fil} elements with ⁵²¹ greatest absolute value are kept in every row and every column of the fac-⁵²² torization, thus defining a maximum level of fill-in.

It is interesting to notice that preconditioners for chemical kinetics problems can also be based on chemistry-based model reduction techniques such as the Directed Relation Graph (DRG) [29]. In this respect, usage of a full sparse analytical Jacobian formulation in companion with a purely algebraic preconditioner is promising as:

• the chemistry Jacobian matrix already physically models mutual interactions among species. Its matrix's sparsity structure is exactly the adjacency representation of a graph, where the path lengths among species (i.e., the intensity of their links) are the Jacobian matrix elements, so it already contains the pieces of information that are modeled using techniques such as the DRG method [13];

in fact, dropping elements from the Jacobian matrix decomposition
 can be interpreted as the effort of removing the smallest contributions
 where species-to-species interactions are smaller;

application of element drop-off makes the matrix factorization faster, and its definition does not rely on chemistry-based considerations, such as reaction rate thresholds, whose validity would not be universal; furthermore, it retains its validity also in case the matrix to be factorized is scaled (for example when diagonal terms are added to the Jacobian matrix, such as in Equation (12)).

Figure 7 compares complete and incomplete LU factorizations of the Jaco-543 bian matrix for a reaction mechanism with more than 1000 species [44], at 544 a reactor temperature $T_0 = 1500K$, pressure $p_0 = 50bar$, and with uniform 545 species mass fractions $Y_i = 1/n_s, i = 1, ..., n_s$, so that all the elements in the 546 Jacobian matrix were active. The ILUT decompositions have been obtained 547 by keeping a constant maximum level of fill-in, $l_{fil} = n_s$, and using different 548 drop-off tolerance values, $\delta = \{0.1, 10^{-3}, 10^{-6}, 10^{-10}\}$. As the Figure shows, 549 the sparse complete LU decomposition leads to about 95% fill-in rate, fea-550 turing 40539 non-zero elements instead of the Jacobian's 20765 non-zeroes. 551 Note that the last line and the last column of the Jacobian matrix, as well 552

as that of the *LU* decomposition, are dense, as the energy closure equation does uniformly depend on all of the species.

The ILUT factorization shows instead a significant reduction in the total 555 number of non-zero elements, and at the drop-off value recommended by Saad 556 [38], $\delta = 0.001$, only 22.7% of the elements of the complete LU factorization 557 survive the truncation strategy. The final number of elements monotonically 558 increases when reducing the drop-off tolerance value; even at the smallest 559 value of the threshold, $\delta = 10^{-10}$, where only the values that are at least 560 ten orders of magnitude smaller than their diagonal element are cancelled. 561 Thus, a significant 28.4% of elements are dropped from the complete LU, 562 conveying the extreme stiffness of the chemical kinetics Jacobian. 563

564 4. Results

The robustness and the computational accuracy of Krylov subspace meth-565 ods incorporated in stiff ODE solution procedures are compared with direct 566 solution methods first for constant-volume, homogeneous reactor ignition cal-567 culations. The LSODES solver [27] is used for the direct, complete sparse 568 solution of the linear system. An interface subroutine has been implemented 569 to provide the Jacobian matrix, analytically calculated in sparse form using 570 the chemistry library, columnwise as requested by the solver. Even if dense 571 columns are requested by the solver interface, including zeroes, we provide the 572 solver with the sparsity structure of the optimally permuted Jacobian matrix 573 at every call. The direct linear system solution is then computed by LSODES 574 at every step by applying complete sparse numerical LU factorization, and 575 then forward and backward substitution. Symbolic LU factorization, i.e., 576 calculation of the sparsity structure of the L and U matrices, is computed 577 by the solver once per call. The sparsity structure of the analytical Jaco-578 bian matrix provided by the SpeedCHEM library never changes throughout 579 the integration, as zero values that may arise in case some species are not 580 currently present, are replaced by the smallest positive number allowed by 581 floating-point accuracy (2.2×10^{-308}) , in double precision). 582

The LSODKR version of the LSODE package is used to deal with a preconditioned Newton-Krylov implicit iteration. The LSODKR solver uses SPIGMR, a modified version of the preconditioned GMRES algorithm [63], that also scales the different variables so that the accuracy for every component of the solution array responds to the tolerance formulation of LSODE's integration algorithm:



Figure 7: Comparison between sparsity pattern of complete and various incomplete LU decompositions for the Jacobian matrix of a reaction mechanism with $n_s = 1034$ [44].

$$|\tilde{y}_i - y_i| \le RTOL_i |y_i| + ATOL_i.$$
⁽²⁰⁾

Here, $RTOL_i$ and $ATOL_i$ are user-defineable relative and absolute values, 589 that have been kept constant in this study at 10^{-8} and 10^{-20} , respectively. 590 The LSODKR solver interfaces to the chemistry library in two steps. First, 591 the Jacobian calculation and preprocessing phase is called; Saad's ILUT pre-592 conditioner has been coupled to the SpeedCHEM library so that precondition-593 ing could be performed at the desired drop-off and fill-in parameter values. 594 and eventually a sparse matrix object containing the L and U parts could be 595 generated. As a second step, a subroutine for directly solving linear systems 596 involving the stored copy of the preconditioner as the linear system matrix 597 has been developed, using the chemistry code's internal sparse algebra library 598 [18]. Thus, LSODKR computes the Jacobian matrix only when the available 590 copy is too old and unaccurate, and then performes multiple system solutions 600 as requested by the r.h.s. of the preconditioned Krylov solution procedure of 601 Eq. (19). 602

603

As the critical factor that determines the total CPU time when integrat-604 ing chemical kinetics is the reaction mechanism size [2], we have investigated 605 a range of different reaction mechanisms, from $n_s = 29$ (skeletal n-heptane 606 ignition for HCCI simulations [39]), to $n_s = 7171$ (detailed mechanism for 607 2-methyl alkanes up to C20 and n-alkanes up to C16 [45]), as summarized in 608 Table 1. Skeletal or semi-skeletal mechanisms with 20-200 species of interest 609 in multi-dimensional combustion simulations were included. A representa-610 tive multi-component fuel composition was used to activate more reaction 611 pathways and keep most elements in the Jacobian matrix active. A stan-612 dardized matrix of 18 ignition delay simulations was established. The initial 613 conditions mimic validation operating points for reaction mechanisms tai-614 lored for internal combustion engine simulations and span different pressures 615 $p_0 = [20.0, 50.0] bar$, temperatures $T_0 = [650, 800, 1000] K$ and fuel-air equiv-616 alence ratios, $\phi_0 = [0.5, 1.0, 2.0]$. The temporal evolution of all of the 18 617 initial value problems has been integrated through a same total time interval 618 of $\Delta t = 0.1s$. 619

620 4.1. Study of optimal preconditioner settings

⁶²¹ A parameter study involving the ILUT preconditioner drop-off tolerance ⁶²² δ and maximum fill-in parameter l_{fil} was conducted to determine the optimal tradeoff between computational complexity of the incomplete factorization, and the convergence rate of the Krylov solver. The matrix of ignition delay calculations was run serially on a single CPU for every reaction mechanism, spanning drop-off tolerances $\delta = [10^{-3}, 10^{-6}, 10^{-8}, 10^{-10}]$ and maximum fillin parameters, expressed in terms of fraction of the Jacobian matrix size, $f_{fil} = [0.25, 0.50, 0.75, 1.0]$, where $f_{fil} = l_{fil}/(n_s + 1)$, corresponding to a total of 288 ignition delay simulations per solver.

In Figure 8, the computational performance of the simulations as a function
of the two ILUT preconditioner settings is reported. CPU time values are
normalized versus the corresponding simulation times requested by running
the same ignition delay calculations using LSODES with the direct solver.
Finally, response surfaces were reconstructed using ordinary kriging.

The shape of the response surfaces was consistent for most reaction mecha-635 nisms. In particular, dependency on the drop-off tolerance order appeared to 636 be much stronger than the maximum fill-in parameter. Only for the small-637 est and more dense reaction mechanisms, some slight effect of the maximum 638 fill-in parameter was seen when the value was as small as about 25% of the 639 row length, and the total computational time was increased due to the slower 640 convergence rate of the Krylov iterative solver. As for the actual speed-up 641 achievable using the Krylov subspace solver, it appears that no improvement 642 could be obtained for reaction mechanisms smaller than $n_s \approx 650$. The 643 total CPU time required by LSODKR was consistently faster than the di-644 rect solver for the largest and sparsest mechanisms. A significant speed-up 645 of about two times was observed only for the largest reaction mechanism, 646 with $n_s = 7171$. More important, coupling of the ILUT preconditioner with 647 the Krylov subspace solver proved to be an extremely robust and efficient 648 methodology, as no solver convergence failures were observed throughout the 649 whole set of simulations, even with tight integration tolerances. Ultimately, 650 optimal settings for the ILUT preconditioner of $\delta = 10^{-8}$ and $f_{fil} = 0.8$ were 651 selected based on the following considerations: 652

- 653 654
- 655
- 656 657
- local minima were consistently observed at drop-off tolerance values not larger than $\delta = 10^{-6}$. Coarser tolerances, such as of the order of $\delta = 10^{-3}$, as suggested by Saad [38], have too slow convergence rates for chemical kinetics problems, and the total simulation time was as much as 3-4 times bigger than with direct system solution;
- 658 659

• minima appeared at different fill-in parameters f_{fil} across different reaction mechanisms. The effect of this parameter on the overall-speed



Figure 8: Influence of ILUT preconditioner settings on relative CPU time performance of the preconditioned Krylov solver-enabled integration for different reaction mechanisms. Contours represent normalized CPU times against CPU time using LSODES and direct sparse solver. (x marks) actual measurement points, (lines) response surface contours using kriging, (O marks) interpolated local minima.

- up was much less relevant, but in some cases allowing complete rows, 660 i.e., $f_{fil} = 1.0$, was too expensive; 661
- extreme drop-off tolerances of $\delta = 10^{-10}$ never yielded faster solutions, 662 meaning that a more accurate and expensive LU decomposition was 663 not compensated enough by the faster convergence rate of the Krylov 664 solver.

665

In Figure 9, cumulative integration metrics are reported for medium and large 666 size mechanisms. Total number of time steps taken, ODE system function 667 and Jacobian matrix evaluations, and linear system solutions are plotted as 668

a function of the drop-off tolerance value, δ , at the optimal fill-in parameter 669 value $f_{fil} = 0.8$, in comparison with the direct solver metrics. The plots show 670 consistent behavior between the two mechanisms, despite the large difference 671 in size and sparsity. Relevantly, the coarsest drop-off tolerance, $\delta = 10^{-3}$, was 672 extremely unefficient in both CPU time performance and integration metrics, 673 as the roughness of the ILUT decomposition requires the ODE integrator to 674 take an extremely high number of integration steps to achieve the desired 675 solution accuracy. Monotonic reduction of all integration metric parameters 676 was seen as the accuracy of the LU decomposition improved. Interestingly, 677 the LSODKR solver took a smaller number of timesteps to completion than 678 LSODES. However, despite similar numbers of Jacobian matrix evaluations 679 (both solvers share the same Jacobian reusage strategy), the overall number 680 of function evaluations and linear system solutions was significantly higher 681 with the Newton-Krylov iteration. In particular, the average number of 682 Krylov subspace iterations per timestep (estimated equal to the number of 683 calls to the linear system solution subroutine), was between 4.47 and 4.98 for 684 the preconditioner tolerances smaller than 10^{-4} . 685

The reaction mechanism size did impact the main solution metrics, but it did affect the total CPU time. The small mechanism showed an optimal trade-off point at $\delta = 10^{-8}$, while the large mechanism reached a plateau in CPU time at $\delta \leq 10^{-8}$. The knee when using the large reaction mechanism was smaller, but the CPU time using the optimal tolerance was still smaller than when using the smallest δ , by about 4.1%.

692 4.2. Overall computational efficiency

693 4.2.1. Ignition delay calculations

The computational efficiency of the preconditioned BDF-Krylov solver 694 LSODKR was compared to both LSODE, that uses direct methods and dense 695 matrix algebra, with sparsely-evaluated analytical Jacobian converted into 696 dense form, and LSODES, with the previously described configuration. Fig-697 ure 10 reports the average CPU time averaged over the 18 ignition delay cases 698 versus reaction mechanism size. Interestingly, despite the fact that Krylov 699 solver techniques have been historically developed for large-scale applica-700 tions, the performance of the Krylov solver was consistently more efficient 701 than the dense solver also for small mechanism sizes, even if the Jacobian 702 matrix was computed in sparse analytical form in that case too. The di-703 rect solver outperformed the Krylov solver for all mechanism sizes smaller 704 than about one thousand with greater efficiency at the smallest mechanisms, 705



Figure 9: Cumulative integration metrics for 18 ignition delay simulations using the LSODKR solver with ILUT preconditioning, versus ILUT drop-off tolerance δ . (top) ERC multiChem mechanism [42], (bottom) LLNL detailed methyl-decanoate mechanism [34]. Marks represent corresponding metrics for the LSODES solver.



Figure 10: Average CPU times per integration over 18 ignition delay simulation cases, using different linear system solution strategies and the fixed-coefficient LSODE BDF method.

where the Krylov solver required up to about 34% more CPU time. The benefit of the preconditioned Krylov solver was observed only for the largest reaction mechanism size, with $n_s = 7171$, where it outperformed the direct sparse solver by 45.4%.

Integration metrics, plotted against reaction mechanism size in Figure 11, 710 show that the structure, including the sparsity, does not affect the solver 711 behavior in terms of the integration parameters; similarly, the differences 712 in behavior observed between the Krylov and the direct sparse solver are 713 seen throughout all the mechanism sizes. One only significant difference is 714 seen when looking at the number of Jacobian evaluations taken by the direct 715 solver for the largest mechanism, when its exact formulation is calculated to 716 bring the Newton procedure in the implicit step down to convergence. 717 718

A closer look at the BDF method is reported in Figure 12, where ignition of a stoichiometric PRF25 fuel-air mixture was simulated using the ERC multiChem mechanism [42], at initial conditions $p_0 = 50bar, T_0 = 800K$, using the LSODE solver. This mechanism was the stiffest due to the extremely lumped reaction pathways for hydrocarbon classes up to C_{16} (cf. Figure 11). Here, the accuracy of all the solutions was not affected by how differently



Figure 11: Integration metrics for 18 ignition cases versus reaction mechanism size. (blue) LSODES with direct solver, (red) LSODKR Krylov solver with $ILUT(\delta = 10^{-8}, f_{fil} = 0.8)$ preconditioning.

the linear system solution was addressed internally within the implicit step 725 iterations. However, slight differences are seen when looking at the time 726 stepping, i.e., at how large time steps the ODE solvers took to compute the 727 solution, as an effect of solution algorithm and round-off and truncation er-728 rors arising from the different procedures. 1) During low-temperature ignition 729 (0.3ms < t < 0.6ms), the solvers (dense direct, sparse direct and Krylov iter-730 ative) show surprisingly similar timestepping, including spikes towards small 731 time step values occurring when the solvers failed to converge due to severe 732 stiffness. 2) during the no-heat-release, preparatory phase when chain initia-733 tions lead to generation of extremely small species mass fractions, the Krylov 734 solver kept integrating with larger timesteps while both direct solvers had 735 sudden time step reductions. This suggests that suppression of the smallest 736 elements in the LU decomposition aids the solution convergence when many 737 species have very small mass fraction values. 3) after ignition, when the reac-738 tions are near equilibrium and the system's stiffness is at its maximum, again 739 the Krylov solver outperformed both other solvers in terms of both taking 740 larger time steps and of achieving smoother integration (fewer restarts). 741 Finally, integration performance was studied in comparison with other stiff 742 ODE solvers that adopt preconditioned Krylov subspace methods: VODPK 743 [73] and DASKR [74], the Krylov enabled versions of VODE and DASSL, 744 respectively. Simulations with the VODE and DASSL ODE solvers using 745 direct linear system solution were run too, where the internal Jacobian com-746 munication and linear system solution algebra parts were replaced by using 747 SpeedCHEM's sparse matrix library. Figure 13 summarizes the computational 748 performance when the optimal ILUT preconditioner parameters were used. 749 All three solvers (LSODKR, VODPK, DASKR) show consistent, efficient be-750 havior and overall scaling of the order of $O(n_s)$ was achieved, both when using 751 the direct sparse linear system solver, or the preconditioned Krylov method. 752 LSODKR and VODPK in particular showed very similar trends, though 753 VODPK was slightly more efficient for large reaction mechanisms. The im-754 plicit formulation in DASKR, that makes it suitable for differential-algebraic 755 (DAE) systems, appeared less efficient when using internal Newton-Krylov 756 iteration, but the version that used SpeedCHEM's sparse matrix algebra im-757 plementation, DASSL, was the fastest solver in the group. The variable coef-758 ficient formulation in VODE made its behavior different than LSODE when 759 switching from direct sparse algebra to the preconditioned Krylov method: 760 the latter being consistently faster than the direct approach throughout most 761 of the mechanisms tested. LSODE with sparse direct algebra was signif-762



Figure 12: Linear system solver comparison for ignition of a PRF25 fuel using the ERC multiChem mechanism and LSODE method, $\phi = 1.0, p_0 = 50bar, T_0 = 800K$: (top) predicted temperature profile, (bottom) integration time steps.



Figure 13: Average CPU times per integration over 18 ignition delay simulation cases, using (solid lines) ILUT preconditioned Krylov subspace solvers LSODKR, VODPK, DASKR; or the corresponding direct solver versions (dashed lines) with custom sparse matrix algebra: LSODES, VODE, DASSL.

icantly faster than its Krylov version for all the reaction mechanism sizes $n_s \leq 654$.

765 4.2.2. Internal Combustion Engine simulations

As pointed out in a previous study [18], the computational efficiency 766 of stiff ODE solvers is different when incorporating them within operator-767 splitting calculation frameworks, such as for the integration of chemical ki-768 netics in multidimensional combustion CFD simulations. The KIVA codes 769 [50, 52, 53] model internal combustion engine flows with sprays and chemical 770 reactions. The code makes use of operator-splitting for spray and chemistry 771 terms. The chemistry evolves in every cell of the computational grid inde-772 pendently as in separate, constant-volume adiabatic reactors. The chemistry 773 integration time requested for every cell is thus the fluid flow solver time-774 step, and due to accuracy constraints of the Arbitrary Lagrangian-Eulerian 775 methodology, this can be as small as $\Delta t = 10^{-7}$ to $10^{-6}s$. This made coupling 776 with the KIVA CFD code a challenging benchmark for the present ODE solu-777

tion methods. In fact, typical integration time steps of stiff chemistry ODE 778 solvers accidentally fall in this range, as shown in Figure 12. Thus, their 779 overall computational performance when coupled to a CFD code relies not 780 only in their capability to perform large time steps, but more importantly, in 781 being able to perform very few Newton iterations within the same step, and 782 achieving high order accuracy, as Jacobian matrix and solution reusage for 783 large time-steps are scarcely possible. Techniques are available to reduce the 784 computational burden by reducing the CFD domain size to a limited number 785 of CFD cell 'clusters' or 'zones' [75] from which the solution can be mapped 786 back to the individual cells based on a proximity criterion [17], but at the 787 cost of introducing some error. However, as the aim of this study was to 788 achieve a solution without approximations, high-dimensional cell clustering 789 features were not used. 790

791

In order to compare the computational performance of direct and Krylov 792 subspace solvers for internal combustion engine simulations, closed-valve sim-793 ulations of a light duty optical research engine [76] were chosen. The sim-794 ulations model a partially premixed combustion strategy, operated through 795 a single, early injection of PRF25 fuel, at a boosted, 55% EGR, 3barIMEP796 low-load operating condition. For computational efficiency, a 1/7th sector 797 mesh with 18399 cells at bottom dead center was used. A customized ver-798 sion of the KIVA-3v code that incorporates the ERC models for spray injec-799 tion, atomization, multi-component evaporation, droplet collision and gas-800 phase interaction, as well as improved turbulence modelling was used. A 801 full description of the code's submodels and of the numerical setup is re-802 ported in [77]. Five different operating conditions were studied, including 803 different swirl ratios $R_s = [1.5, 2.2, 3.5]$, and different fuel injection pressure 804 values, namely $p_{ini} = [500, 860, 1220]$ bar. Chemistry was solved using the 805 SpeedCHEM library, and the relatively small, skeletal ERC PRF mechanism, 806 which has $n_s = 47$ and $n_r = 142$ [40]. The maximum time-step allowed by 807 the fluid flow solver was $3 \times 10^{-6} s$. All the simulations were run in serial 808 mode on one processor, on a computing cluster featuring AMD Opteron pro-809 cessors with 4 cores per CPU and a clock speed of 2700 MHz. 810

811

The CPU time comparison of the simulations for the three solvers LSODE, VODE, DASSL in both Direct and preconditioned Krylov configurations is reported in Figure 14. Specific chemistry times are separated from the rest of the code operation, that includes the fluid flow and spray solution. Clearly,



Figure 14: Solver wall-time comparison for 5 KIVA engine case simulations.

both the VODE and the LSODE solvers showed similar overall performance 816 when using the direct solver. Usage of the Krylov method in LSODE provided 817 a significant speed-up of the solution (an average of -23%). For VODE, the 818 advantage of the direct solver seen for ignition simulations was not seen, but 819 similar performance was observed between the two methods. The DASSL 820 solver showed significant benefits in adopting the Krylov iteration method as 821 LSODE, but its overall computational times were larger than for the other 822 solvers. Noticeably, all solvers performed consistently across all cases, and 823 no Krylov method failures were observed with the optimal preconditioner 824 settings despite the wide variety of reactor conditions. Also, no noticeable 825 differences could be observed in the predicted engine pressure and tempera-826 ture traces, as shown for example in Figure 15. The overall agreements of all 827 the simulations were within a relative error of 0.63% for the pressure traces. 828 In order to quantify the computational efficiency of the solvers during the 829 simulation, a *grid-index* value was used. This index is defined as the average 830 CPU time spent on every cell of the grid at every time step of the KIVA 831 solver. Taking the simulation with $R_s = 1.5$ and $p_{inj} = 860 bar$ as represen-832 tative, details of the grid index behavior during the simulations are reported 833



Figure 15: Solution comparison for the $R_s = 1.5$, $p_{inj} = 860bar$ case. (solid) in-cylinder pressure trace, (dashed) apparent rate of heat release. Chemistry integration using direct vs. preconditioned Krylov LSODE solver.

in Figure 16. Confirming the results of the ignition delay integrations, it is 834 seen that the Krylov method is much more suitable for the integration of the 835 low-temperature region than the direct solvers. Grid index values achieved 836 by the Krylov solvers at crank angles smaller than about $-30^{\circ}ATDC$ out-837 performed the corresponding direct solvers by up to more than three times. 838 On the contrary, during high temperature ignition, the Krylov solvers were 839 always outperformed by the direct solvers, except for LSODE, where both 840 approaches had similar computational time requirements. Finally, later after 841 ignition, when the mixture reaches near-equilibrium conditions, and strong, 842 but almost equal, forward and backward reaction contributions compete, the 843 Krylov subspace methods performed similarly, or faster, than the correspond-844 ing direct ones. 845

846

In conclusion, the present engine simulation study confirmed the results ob-847 served for the ignition delay calculations. However, the integrators were 848 forced to small time steps by the CFD solver and required repeated Ja-849 cobian evaluations/solutions due to the operator-splitting procedure. Tis 850 enhances differences in performance among the solvers. In particular, the 851 Krylov method significantly outperformed the direct solver in the region of 852 the compression stroke where no fuel is present within the mixture, and the 853 temperature is low, and again late after ignition, where the kinetics are near 854 equilibrium. The direct solvers were instead consistently faster than their 855 Krylov-enabled counterparts throughout the whole fuel injection and com-856 bustion duration. Amongst the solvers, both versions of LSODE showed the 857 best performance. 858

5. Concluding remarks

The computational performance of direct and preconditioned iterative 860 Krylov subspace solvers for the solution of the stiff ODE systems associ-861 ated with combustion chemical kinetics has been studied. The chemical 862 kinetics of reactive gaseous mixtures was been modeled using the SpeedCHEM 863 package, an efficient Fortran library for combustion chemistry that features 864 sparse analytical formulations for the most important functions, optimal-865 degree interpolation of expensive thermodynamic properties, as well as a full 866 library for internal sparse algebra that handles both matrix and vector com-867 putations and interfaces to stiff ODE solvers. The LSODES and LSODKR 868 packages were compared, as they implement the same implicit BDF method 869



Figure 16: Grid index comparison for the $R_s = 1.5$, $p_{inj} = 860bar$ case, Direct vs. preconditioned Krylov ODE solver operation: (top) VODE, (center) LSODE, (bottom) DASSL.

and the same Jacobian re-usage strategy. Saad's sparse incomplete LU fac-870 torization with dual truncation strategy (ILUT) was used as the problem 871 preconditioner, in combination with the SPIGMR solver when the Krylov 872 subspace solution technique was used. Other BDF solvers, such as DASSL 873 and VODE, including their Krylov-enabled versions, DASKR and VODPK, 874 were also compared to the LSODE solver. Comparisons were also made with 875 the SpeedCHEM code coupled to the KIVA-ERC CFD code to perform internal 876 combustion engine simulations of a light duty diesel engine operating with a 877 partially premixed combustion strategy. 878

879

The study showed that Krylov subspace methods, with a robust, generalpurpose preconditioner, are extremely efficient for simulating combustion chemistry problems. The direct solver approach however provided greater computational efficiency for smaller reaction mechanisms that are of practical interest for multidimensional combustion simulations, where $n_s \leq 10^3$. In particular:

• the sparse analytical Jacobian formulation made it possible to find a 886 suitable preconditioner for the integration. The ILUT's strategy that 887 drops the least important elements in the decomposition is similar to 888 removing interactions among species whose mutual links are weak. In 889 this sense, the approach has the same effects as on-the-fly mechanism 890 reduction approaches such as the DRG. The simplification is achieved 891 without using expensive species relation techniques that would make 892 the construction of the reduced Jacobian matrix much more demanding, 893 since the graph links between the species, as well as its structure, would 894 have to be re-calculated at every timestep; 895

- optimal preconditioner settings of $\delta = 10^{-8}$ and $f_{fil} = 0.8$ were found after an optimization study. These settings suggest that accurate factorization is needed. Despite this accuracy constraint, the final number of non-zeros can be as low as about half that of the complete factorization;
- no solver failures due to bad convergence of the Krylov subspace method
 were observed when using the ILUT preconditioner;
- integration metrics at the optimal preconditioner settings were close to the metrics of the direct solver method, confining that the ILUT

is a robust, general-purpose preconditioner, and the optimal settings
 chosen were appropriate for this class of problems. A smaller number
 of integration time steps was needed for the Krylov solver;

- the overall performance of the direct solver was consistently faster than
 the Krylov subspace solver for reaction mechanisms smaller than about
 1000 species;
- the dimensions of the chemistry ODE system were evidently not large enough to fully exploit the potential of Krylov solvers;

when incorporated into multidimensional engine CFD simulations with
a skeletal reaction mechanism, the Krylov method was faster than
the direct method during the low-temperature, pre-ignition phase, and
had similar performance after ignition, at near-equilibrium conditions.
During ignition, the direct solvers always outperformed their Krylovenabled versions;

a computational speed-up of nearly 25% on chemistry could be achieved by adopting the LSODKR solver in comparison with LSODES, that has already been shown to be the fastest among the direct solvers within CFD simulations. Greater speed-ups could be achieved by adopting a hybrid solution strategy where the most appropriate solver method is chosen based on the local mixture conditions.

Finally, an important finding is that both methods achieve scaling of 925 the computational cost of the solution from $O(n_s^3)$ down to about $O(n_s)$ for 926 the very sparse combustion chemistry models that describe high molecular 927 weight hydrocarbon fuels. This allows about two orders of magnitude CPU 928 time savings for multi-dimensional CFD simulations, and up to three orders 929 of magnitude savings for the largest hydrocarbon fuel oxidation mechanisms 930 currently available in the literature. This also sets the path for further stud-931 As the computational cost due to the solution of the linear system ies. 932 associated with the chemistry Jacobian is still the most computationally de-933 manding task, provided that advanced numerics are adopted to compute 934 both the Jacobian and the r.h.s. of the system of equations (sparse algebra, 935 analytical formulation, optimal-degree interpolation of expensive functions), 936 it is possible that further improvements could be achieved by improving the 937 analytical representation of the Jacobian's LU factorization, or, when cou-938 pled with CFD codes, by removing the operator splitting simplification and 930

solve for both transport and chemistry using the same method. The issue of reaction mechanism size, that is large for dense algebra, but yet too small for extended parallelism, must be addressed to allow the integration of chemical kinetics problems on massively parallel architectures such as GPUs to be really efficient. Possibly an approach similar to CFD is needed, where every reactor is solved onto one single graphical unit, and shared memory provides a common pool for tabulated thermodynamic functions.

947 6. Acknowledgements

The authors gratefully acknowledge funding by the Sandia National Laboratories and the Princeton Combustion Energy Frontier Research Center (CEFRC).

951

952 7. Notes

⁹⁵³ The SpeedCHEM library is available upon request to the authors.

- [1] R. D. Reitz. Directions in internal combustion engine re-954 search, Combustion and Flame 160(1)(2013)1 8. 955 doi:http://dx.doi.org/10.1016/j.combustflame.2012.11.002. 956
- 957 URL http://www.sciencedirect.com/science/article/pii/S001021801200315X
- [2] T. Lu, C. K. Law, Toward accommodating realistic fuel chemistry in large-scale
 computations, Progress in Energy and Combustion Science 35 (2) (2009) 192 215.
 doi:http://dx.doi.org/10.1016/j.pecs.2008.10.002.
- 961 URL http://www.sciencedirect.com/science/article/pii/S036012850800066X
- [3] W. J. Pitz, C. J. Mueller, Recent progress in the development of diesel surrogate fuels, Progress in Energy and Combustion Science 37 (3) (2011) 330 350.
 doi:http://dx.doi.org/10.1016/j.pecs.2010.06.004.
- 965 URL http://www.sciencedirect.com/science/article/pii/S0360128510000535
- [4] W. H. Green, Predictive Kinetics: A New Approach for the 21st Century, in: G. B.
 Marin (Ed.), Chemical Engineering Kinetics, Vol. 32 of Advances in Chemical Engineering, Academic Press, 2007, pp. 1 313. doi:http://dx.doi.org/10.1016/S0065-2377(07)32001-2.

- [5] E. Hairer, S. P. Norsett, G. Wanner, Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems, 2nd Edition, Springer, Berlin, 1991.
- [6] S. Lam, D. Coussis, Understanding complex chemical kinetics with computational singular perturbation, Symposium (International) on Combustion 22 (1) (1989) 931
 975 941. doi:http://dx.doi.org/10.1016/S0082-0784(89)80102-X.
- 976 URL http://www.sciencedirect.com/science/article/pii/S008207848980102X
- 977 [7] S. Vajda, P. Valko, T. Turanyi, Principal component analysis of kinetic
 978 models, International Journal of Chemical Kinetics 17 (1) (1985) 55–81.
 979 doi:10.1002/kin.550170107.
- 980 URL http://dx.doi.org/10.1002/kin.550170107
- [8] N. J. Brown, G. Li, M. L. Koszykowski, Mechanism reduction via principal component
 analysis, International Journal of Chemical Kinetics 29 (6) (1997) 393–414.
- [9] M. Valorani, D. A. Goussis, Explicit Time-Scale Splitting Algorithm for Stiff Problems: Auto-ignition of Gaseous Mixtures behind a Steady Shock, Journal of Computational Physics 169 (1) (2001) 44 79. doi:http://dx.doi.org/10.1006/jcph.2001.6709.
 URL http://www.sciencedirect.com/science/article/pii/S0021999101967099
- [10] U. Maas, S. Pope, Simplifying chemical kinetics: Intrinsic low-dimensional manifolds in composition space, Combustion and Flame 88 (34) (1992) 239 - 264. doi:http://dx.doi.org/10.1016/0010-2180(92)90034-M.
- 990 URL http://www.sciencedirect.com/science/article/pii/001021809290034M

⁹⁷⁰ URL http://www.sciencedirect.com/science/article/pii/S0065237707320012

- [11] J. C. Keck, Rate-controlled constrained-equilibrium theory of chemical reactions in
 complex systems, Progress in Energy and Combustion Science 16 (2) (1990) 125 –
 154. doi:http://dx.doi.org/10.1016/0360-1285(90)90046-6.
- 994 URL http://www.sciencedirect.com/science/article/pii/0360128590900466
- [12] G. Li, H. Rabitz, A general analysis of exact lumping in chemical kinetics, Chemical Engineering Science 44 (6) (1989) 1413 1430. doi:http://dx.doi.org/10.1016/0009-2509(89)85014-6.
- URL http://www.sciencedirect.com/science/article/pii/0009250989850146
- ⁹⁹⁹ [13] T. Lu, C. K. Law, A directed relation graph method for mechanism reduction, Proceedings of the Combustion Institute 30 (1) (2005) 1333 1341.
 ¹⁰⁰¹ doi:http://dx.doi.org/10.1016/j.proci.2004.08.145.
- 1002 URL http://www.sciencedirect.com/science/article/pii/S0082078404001973
- [14] P. Pepiot-Desjardins, H. Pitsch, An efficient error-propagation-based reduction method for large chemical kinetic mechanisms, Combustion and Flame 154 (12) (2008) 67 - 81. doi:http://dx.doi.org/10.1016/j.combustflame.2007.10.020.
 URL http://www.sciencedirect.com/science/article/pii/S0010218007003264
- [15] L. Elliott, D. B. Ingham, A. G. Kyne, N. S. Mera, M. Pourkashanian, C. W.
 Wilson, Reaction Mechanism Reduction and Optimization Using Genetic Algorithms, Industrial and Engineering Chemistry Research 44 (4) (2005) 658–667.
 arXiv:http://pubs.acs.org/doi/pdf/10.1021/ie049409d, doi:10.1021/ie049409d.
 URL http://pubs.acs.org/doi/abs/10.1021/ie049409d
- [16] F. Perini, J. L. Brakora, R. D. Reitz, G. Cantore, Development of reduced and optimized reaction mechanisms based on genetic algorithms and element flux analysis, Combustion and Flame 159 (1) (2012) 103 – 119. doi:http://dx.doi.org/10.1016/j.combustflame.2011.06.012.
- 1016 URL http://www.sciencedirect.com/science/article/pii/S0010218011001921
- 1017[17] S. Pope, Computationally efficient implementation of combustion chemistry us-1018ing in situ adaptive tabulation, Combustion Theory and Modelling 1 (1)1019(1997) 41-63. arXiv:http://www.tandfonline.com/doi/pdf/10.1080/713665229,1020doi:10.1080/713665229.
- 1021 URL http://www.tandfonline.com/doi/abs/10.1080/713665229
- [18] F. Perini, E. Galligani, R. D. Reitz, An Analytical Jacobian Approach to
 Sparse Reaction Kinetics for Computationally Efficient Combustion Modeling
 with Large Reaction Mechanisms, Energy & Fuels 26 (8) (2012) 4804–4822.
 arXiv:http://pubs.acs.org/doi/pdf/10.1021/ef300747n, doi:10.1021/ef300747n.
- 1026 URL http://pubs.acs.org/doi/abs/10.1021/ef300747n
- [19] D. A. Schwer, J. E. Tolsma, W. H. Green, P. I. Barton, On upgrading the numerics in combustion chemistry codes, Combustion and Flame 128 (3) (2002) 270 - 291. doi:http://dx.doi.org/10.1016/S0010-2180(01)00352-2.
- 1030 URL http://www.sciencedirect.com/science/article/pii/S0010218001003522

- [20] R. A. Whitesides, M. J. McNenly, D. L. Flowers, Optimizing Time Integration of Chemical-Kinetic Networks for Speed and Accuracy, in: 8th US National Combustion Meeting, 2013.
- [21] K. V. Puduppakkam, L. Liang, C. V. Naik, E. Meeks, S. L. Kokjohn, R. D. Reitz, Use of Detailed Kinetics and Advanced Chemistry-Solution Techniques in CFD to Investigate Dual-Fuel Engine Concepts, SAE Int. J. Engines 4 (2011) 1127–1149. doi:10.4271/2011-01-0895.
- 1038 URL http://dx.doi.org/10.4271/2011-01-0895
- [22] G. B. Ferraris, D. Manca, BzzOde: a new C++ class for the solution of stiff and non-stiff ordinary differential equation systems, Computers & Chemical Engineering 22 (11) (1998) 1595 - 1621. doi:http://dx.doi.org/10.1016/S0098-1354(98)00233-6.
- 1042 URL http://www.sciencedirect.com/science/article/pii/S0098135498002336
- [23] V. Damian, A. Sandu, M. Damian, F. Potra, G. R. Carmichael, The kinetic preprocessor KPP-a software environment for solving chemical kinetics, Computers and Chemical Engineering 26 (11) (2002) 1567 1579. doi:http://dx.doi.org/10.1016/S0098-1354(02)00128-X.
- 1047 URL http://www.sciencedirect.com/science/article/pii/S009813540200128X
- [24] A. Sandu, J. Verwer, M. V. Loon, G. Carmichael, F. Potra, D. Dabdub, J. Seinfeld, Benchmarking stiff ode solvers for atmospheric chemistry problems-I. implicit vs explicit, Atmospheric Environment 31 (19) (1997) 3151 – 3166. doi:http://dx.doi.org/10.1016/S1352-2310(97)00059-9.
- 1052 URL http://www.sciencedirect.com/science/article/pii/S1352231097000599
- [25] A. Sandu, J. Verwer, J. Blom, E. Spee, G. Carmichael, F. Potra, Benchmarking stiff
 ode solvers for atmospheric chemistry problems II: Rosenbrock solvers, Atmospheric
 Environment 31 (20) (1997) 3459 3472. doi:http://dx.doi.org/10.1016/S13522310(97)83212-8.
- 1057 URL http://www.sciencedirect.com/science/article/pii/S1352231097832128
- [26] P. Brown, G. Byrne, A. Hindmarsh, VODE: A Variable-Coefficient ODE Solver,
 SIAM Journal on Scientific and Statistical Computing 10 (5) (1989) 1038–1051.
 arXiv:http://epubs.siam.org/doi/pdf/10.1137/0910062, doi:10.1137/0910062.
 URL http://epubs.siam.org/doi/abs/10.1137/0910062
- 1001 0101 100p1, , opublibiam.org, a01, abb, 1011101, 0010002
- 1062 [27] A. C. Hindmarsh, LSODE and LSODI, two new initial value ordinary
 1063 differential equation solvers, SIGNUM Newsletter 15 (4) (1980) 10–11.
 1064 doi:10.1145/1218052.1218054.
- 1065 URL http://doi.acm.org/10.1145/1218052.1218054
- [28] L. R. Petzold, Description of DASSL: a differential/algebraic system solver, Tech.
 rep., Sandia National Laboratories SAND-82-8637 (1982).

- [29] M. J. McNenly, R. A. Whitesides, D. L. Flowers, Adaptive Preconditioning Strategies
 for Integrating Large Kinetic Mechanisms, in: 8th US National Combustion Meeting,
 2013.
- [30] P. Tranquilli, A. Sandu, Rosenbrock-Krylov Methods for Large Systems of Differential
 Equations, ArXiv 2013arXiv:1305.5481.
- of[31] F. 1073 Bisetti, Integration large chemical kinetic mechanisms via exponential methods with Krylov approximations to Jacobian matrix 1074 Theory and Modelling (3)(2012)387 - 418.functions, Combustion 161075 arXiv:http://www.tandfonline.com/doi/pdf/10.1080/13647830.2011.631032, 1076 doi:10.1080/13647830.2011.631032. 1077
- 1078 URL http://www.tandfonline.com/doi/abs/10.1080/13647830.2011.631032
- [32] Y. Shi, W. H. Green, H.-W. Wong, O. O. Oluwole, Redesigning combustion modeling algorithms for the Graphics Processing Unit (GPU): Chemical kinetic rate evaluation and ordinary differential equation integration, Combustion and Flame 158 (5) (2011)
 836 - 847. doi:http://dx.doi.org/10.1016/j.combustflame.2011.01.024.
- 1083 URL http://www.sciencedirect.com/science/article/pii/S0010218011000393
- [33] J. C. Linford, J. Michalakes, M. Vachharajani, A. Sandu, Multi-core acceleration of chemical kinetics for simulation and prediction, in: Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis, SC '09, ACM, New York, NY, USA, 2009, pp. 7:1–7:11. doi:10.1145/1654059.1654067.
- 1088 URL http://doi.acm.org/10.1145/1654059.1654067
- [34] O. Herbinet, W. J. Pitz, C. K. Westbrook, Detailed chemical kinetic oxidation mechanism for a biodiesel surrogate, Combustion and Flame 154 (3) (2008) 507 – 528. doi:http://dx.doi.org/10.1016/j.combustflame.2008.03.003.
- 1092 URL http://www.sciencedirect.com/science/article/pii/S0010218008000771
- [35] F. Perini, E. Galligani, G. Cantore, R. Reitz, Validation of a Sparse Analytical Jacobian Chemistry Solver for Heavy-Duty Diesel Engine Simulations with Comprehensive Reaction Mechanisms, SAE Technical Paper 2012-01-1974doi:10.4271/2012-01-1974.
 URL http://dx.doi.org/10.4271/2012-01-1974
- [36] F. Perini, A. Krishnasamy, Y. Ra, R. D. Reitz, Computationally Efficient Simulation
 of Multi-component Fuel Combustion using a Sparse Analytical Jacobian Chemistry
 Solver and High-dimensional Clustering, in: Proceedings of the ASME Internal Com bustion Engine Division's 2013 Fall Technical Conference ICEF2013-19039, 2013.
- [37] R. S. Varga, Factorization and normalized iterative methods, in: R. E. Langer (Ed.),
 Boundary Problems in Differential Equations, University of Wisconsin Press, 1960,
 pp. 121–142.
- [38] Y. Saad, ILUT: A dual threshold incomplete LU factorization, Numerical Linear
 Algebra with Applications 1 (4) (1994) 387-402. doi:10.1002/nla.1680010405.
 URL http://dx.doi.org/10.1002/nla.1680010405

- [39] A. Patel, S.-C. Kong, R. D. Reitz, Development and Validation of a Reduced Reaction Mechanism for HCCI Engine Simulations, SAE Technical Paper 2004-010558doi:10.4271/2004-01-0558.
- 1110 URL http://dx.doi.org/10.4271/2004-01-0558
- [40] Y. Ra, R. D. Reitz, A reduced chemical kinetic model for IC engine combustion
 simulations with primary reference fuels, Combustion and Flame 155 (4) (2008) 713
 738. doi:http://dx.doi.org/10.1016/j.combustflame.2008.05.002.
- 1114 URL http://www.sciencedirect.com/science/article/pii/S0010218008001351
- 1115 [41] H. Wang, PRF reaction mechanism, personal communication (2013).
- [42] Y. Ra, R. D. Reitz, A combustion model for IC engine combustion simulations
 with multi-component fuels, Combustion and Flame 158 (1) (2011) 69 90.
 doi:http://dx.doi.org/10.1016/j.combustflame.2010.07.019.
- 1119 URL http://www.sciencedirect.com/science/article/pii/S0010218010002075
- [43] R. Seiser, H. Pitsch, K. Seshadri, W. Pitz, H. Curran, Extinction and autoignition
 of n-heptane in counterflow configuration, Proceedings of the Combustion Institute
 28 (2) (2000) 2029 2037. doi:http://dx.doi.org/10.1016/S0082-0784(00)80610-4.
- 1123 URL http://www.sciencedirect.com/science/article/pii/S0082078400806104
- [44] H. Curran, P. Gaffuri, W. Pitz, C. Westbrook, A comprehensive modeling
 study of iso-octane oxidation, Combustion and Flame 129 (3) (2002) 253 280.
 doi:http://dx.doi.org/10.1016/S0010-2180(01)00373-X.
- 1127 URL http://www.sciencedirect.com/science/article/pii/S001021800100373X
- [45] C. K. Westbrook, W. J. Pitz, O. Herbinet, H. J. Curran, E. J. Silke, A comprehensive detailed chemical kinetic reaction mechanism for combustion of n-alkane hydrocarbons from n-octane to n-hexadecane, Combustion and Flame 156 (1) (2009) 181 – 199. doi:http://dx.doi.org/10.1016/j.combustflame.2008.07.014.
- 1132 URL http://www.sciencedirect.com/science/article/pii/S0010218008002125
- [46] R. Kee, F. M. Rupley, A. Miller, Chemkin-II: A Fortran chemical kinetics package for
 the analysis of gas-phase chemical kinetics, Tech. rep., Sandia National Laboratories
 SAND-89-8009 (1989).
- [47] Y. Saad, Iterative Methods for Sparse Linear Systems, EngineeringPro collection,
 Society for Industrial and Applied Mathematics (SIAM, 3600 Market Street, Floor
 6, Philadelphia, PA 19104), 2003.
- 1139 URL http://books.google.com/books?id=h9nwszYPblEC
- [48] J. Reid, The new features of Fortran 2008, SIGPLAN Fortran Forum 27 (2) (2008)
 8-21. doi:10.1145/1408643.1408645.
- 1142 URL http://doi.acm.org/10.1145/1408643.1408645
- [49] R. M. Stallman, et al., Gnu compiler collection (November 1987).
 URL http://gcc.gnu.org

- [50] A. Amsden, P. Orourke, T. Butler, KIVA-2: A computer program for chemically reactive flows with sprays, NASA STI/Recon Technical Report N 89 (1989) 27975.
- [51] M. J. Holst, Notes on the kiva-ii software and chemically reactive fluid mechanics,
 Tech. rep., Lawrence Livermore National Laboratory (1992).
- [52] A. A. Amsden, KIVA-3V: A block-structured KIVA program for engines with vertical
 or canted valves, Tech. rep., Los Alamos National Lab., NM (United States) (1997).
- [53] D. J. Torres, M. F. Trujillo, KIVA-4: An unstructured ALE code for compressible
 gas flow with sprays, Journal of Computational Physics 219 (2) (2006) 943 975.
 doi:http://dx.doi.org/10.1016/j.jcp.2006.07.006.
- 1154 URL http://www.sciencedirect.com/science/article/pii/S002199910600338X
- [54] H. Eyring, The Activated Complex in Chemical Reactions, The Journal of Chemical Physics 3 (2) (1935) 107–115. doi:10.1063/1.1749604.
- 1157 URL http://link.aip.org/link/?JCP/3/107/1
- [55] A. A. Konnov, Remaining uncertainties in the kinetic mechanism of hydrogen combustion, Combustion and Flame 152 (4) (2008) 507 528.
 doi:http://dx.doi.org/10.1016/j.combustflame.2007.10.024.
- 1161 URL http://www.sciencedirect.com/science/article/pii/S0010218007003318
- [56] I. Duff, A. Erisman, J. Reid, Direct Methods for Sparse Matrices, Monographs on numerical analysis, Oxford University Press, Incorporated, 1989.
- E. Cuthill, J. McKee, Reducing the bandwidth of sparse symmetric matrices, in:
 Proceedings of the 1969 24th national conference, ACM '69, ACM, New York, NY, USA, 1969, pp. 157–172. doi:10.1145/800195.805928.
- ¹¹⁶⁷ URL http://doi.acm.org/10.1145/800195.805928
- [58] A. George, J. W. Liu, Computer Solution of Large Sparse Positive Definite Systems,
 Prentice Hall Professional Technical Reference, 1981.
- [59] R. Tewarson, Sparse Matrices, Mathematics in Science and Engineering, Academic
 Press, 1973.
- [60] J. Troe, Fall-off Curves of Unimolecular Reactions, Berichte der Bunsengesellschaft
 fur physikalische Chemie 78 (5) (1974) 478–488. doi:10.1002/bbpc.19740780510.
 URL http://dx.doi.org/10.1002/bbpc.19740780510
- [61] C. Kelley, Iterative Methods for Linear and Nonlinear Equations, Frontiers in Applied Mathematics, Society for Industrial and Applied Mathematics, 1995.
- 1177 URL http://books.google.com/books?id=c3m7Ro7EhOAC
- [62] J. Ortega, Numerical Analysis: A Second Course, Classics in Applied Mathematics,
 Society for Industrial and Applied Mathematics, 1990.
- 1180 URL http://books.google.com/books?id=DboVs3T2UnoC

- [63] Y. Saad, M. Schultz, GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems, SIAM Journal on Scientific and Statistical Computing 7 (3) (1986) 856–869. arXiv:http://epubs.siam.org/doi/pdf/10.1137/0907058, doi:10.1137/0907058.
- 1185 URL http://epubs.siam.org/doi/abs/10.1137/0907058
- [64] C. Gear, Numerical initial value problems in ordinary differential equations, Prentice Hall series in automatic computation, Prentice-Hall, 1971.
- 1188 URL http://books.google.com/books?id=e9QQAQAAIAAJ
- [65] P. N. Brown, A. C. Hindmarsh, Matrix-free methods for stiff systems of ODE's, SIAM
 J. Numer. Anal. 23 (3) (1986) 610–638. doi:10.1137/0723039.
 URL http://dx.doi.org/10.1137/0723039
- [66]R. Dembo. S. Eisenstat. Τ. Steihaug, Inexact Newton Meth-1192 Numerical 19(1982)ods, SIAM Journal on Analysis (2)400 - 408.1193 arXiv:http://epubs.siam.org/doi/pdf/10.1137/0719025, doi:10.1137/0719025. 1194 URL http://epubs.siam.org/doi/abs/10.1137/0719025 1195
- [67] S. C. Eisenstat, H. F. Walker, Choosing the Forcing Terms in an Inexact Newton
 Method, SIAM J. Sci. Comput 17 (1994) 16–32.
- [68] W. Rheinboldt, Methods for Solving Systems of Nonlinear Equations, CBMS-NSF
 Regional Conference Series in Applied Mathematics, Society for Industrial and Applied Mathematics, 1998.
- 1201 URL http://books.google.com/books?id=lnJznmVnJ2wC
- ¹²⁰² [69] W. E. Arnoldi, The principle of minimized iterations in the solution of the matrix ¹²⁰³ eigenvalue problem, Quart. Appl. Math 9 (1) (1951) 17–29.
- [70] I. C. Ipsen, C. D. Meyer, The Idea Behind Krylov Methods, The American Mathe matical Monthly 105 (1998) 889–899.
- [71] A. Greenbaum, Iterative Methods for Solving Linear Systems, Frontiers in Applied Mathematics, Society for Industrial and Applied Mathematics, 1997.
 URL http://books.google.com/books?id=WMDNLxruosC
- 1208 URL http://books.google.com/books?id=WwMDNLxrwocC
- [72] S.-L. Kim, J.-Y. Choi, I.-S. Jeung, Y.-H. Park, Application of approximate chemical Jacobians for constant volume reaction and shock-induced combustion, Applied Numerical Mathematics 39 (1) (2001) 87 – 104. doi:http://dx.doi.org/10.1016/S0168-9274(01)00054-X.
- 1213 URL http://www.sciencedirect.com/science/article/pii/S016892740100054X
- [73] G. D. Byrne, Pragmatic experiments with Krylov methods in the stiff ODE settings,
 in: Computational Ordinary Differential Equations (J.R. Cash and I. Gladwell, eds.),
 Clarendon Press, Oxford, UK, 1992.

- [74] P. Brown, A. Hindmarsh, L. Petzold, Using Krylov Methods in the Solution of Large-Scale Differential-Algebraic Systems, SIAM Journal on Scientific Computing 15 (6) (1994) 1467–1488. arXiv:http://epubs.siam.org/doi/pdf/10.1137/0915088, doi:10.1137/0915088.
- 1221 URL http://epubs.siam.org/doi/abs/10.1137/0915088
- [75] F. Perini, High-dimensional, unsupervised cell clustering for computationally efficient
 engine simulations with detailed combustion chemistry, Fuel 106 (2013) 344 356.
 doi:http://dx.doi.org/10.1016/j.fuel.2012.11.015.
- 1225 URL http://www.sciencedirect.com/science/article/pii/S0016236112008915
- [76] I. W. Ekoto, W. F. Colban, P. C. Miles, S. W. Park, D. E. Fos-1226 R. D. Reitz, U. Aronsson, O. Andersson, UHC and CO Emister, 1227 sions Sources from a Light-Duty Diesel Engine Undergoing Dilution-Controlled 1228 Low-Temperature Combustion, SAE International Journal of Engines 2 (2) 1229 (2010) 411–430. arXiv:http://saeeng.saejournals.org/content/2/2/411.full.pdf+html. 1230 doi:10.4271/2009-24-0043. 1231
- 1232 URL http://saeeng.saejournals.org/content/2/2/411.abstract
- [77] F. Perini, A. Dempsey, R. Reitz, D. Sahoo, B. Petersen, P. Miles, A Computational
 Investigation of the Effects of Swirl Ratio and Injection Pressure on Mixture Preparation and Wall Heat Transfer in a Light-Duty Diesel Engine, SAE Technical Paper
 2013-01-1105doi:10.4271/2013-01-1105.
- ¹²³⁷ URL http://dx.doi.org/10.4271/2013-01-1105