

SYSTEMATIC STUDY OF THE PERFORMANCE IMPROVEMENT OF AN IN SITU ADAPTIVE TABULATION TECHNIQUE

Andrea Cristina Carvalho dos Anjos, andreacca@puc-rio.br

Luís Fernando Figueira da Silva, luisfer@esp.puc-rio.br

Department of Mechanical Engineering, Pontifícia Universidade Católica do Rio de Janeiro
Rua Marquês de São Vicente, 225, Gávea - Rio de Janeiro, RJ - Brasil - 22451-900

Abstract. *This work is part of a study whose objective is to minimize the evaluation costs of an automatic technique of chemical kinetics mechanisms tabulation. Hence, this paper presents a state of the art discussion of the technique named in situ adaptive tabulation. Modifications in the tabulation of the existing algorithm designed to improve the storage capacity and simulation time are discussed. A systematic study of the overall accuracy and costs of using this technique is presented for the reactive flow CO/O₂ using a pairwise-mixing stirred reactor (PMSR). The obtained results allow to characterize the performance savings, in terms of memory usage, accuracy and CPU time of the proposed improvements.*

Keywords: *tabulation, performance, combustion*

1. INTRODUCTION

Currently, the numerical simulation of combustion processes is a critical task involved in the design of novel combustion apparatus, such as low pollutant emissions combustors. However, a combustion model which faithfully describes the physical-chemical processes presents, as the main difficulty, the vast computational effort required for solution of the governing nonlinear equations, involving a large number of variables and space-time scales, which describes mass, momentum and energy transport. Most often, the largest part of this computational burden is associated to the Arrhenius laws embedded in the chemical sources terms.

Since the computational cost associated to the chemical kinetics mechanisms is inherently expensive, due to the Arrhenius dependence of the reaction rates, the purpose of this work is to reduce such cost by proposal and evaluation of techniques that can be coupled to existing codes, without loss of accuracy. This technique is expected to become a building block of transported probability density function large eddy simulations (Andrade et al., 2011 CST).

To this end, the technique called in situ adaptive tabulation - ISAT, (Cunha, 2010 and Pope, 1997), is modified in order to reduce the storage costs, which were found to be the main drawback of the implemented technique.

2. LITERATURE REVIEW

The in situ adaptive tabulation (ISAT) is a technique developed by Pope (1997) for cost reduction of high dimensionality problems. Given the initial states of a function, the final states are approximated by a linear extrapolation from a neighboring solution using a sensitivity matrix, the error of these approximations can be controlled.

This technique proved to be 1000 times faster than the direct numerical integration of the governing equations (mass, moment, energy), when used in solution of problems that employ the method of the probability density function (PDF) transport. The most important aspect is that there are no restrictions on the physical system to which the technique can be applied; therefore, its use is independent of assumptions about the structure of the flame in combustion simulations. A detailed description of the original ISAT technique may be found elsewhere in Cunha (2010) and Pope (1997). Herein, we review the most recent developments related to such a technique.

Lu and Pope (2009) present a new approach called ISAT-5, which is also based on a tabulation which construction is concurrent simulation. Only the composition space region accessed during the simulation is tabulated. The ISAT tabulates, in a binary tree, the function $\mathbf{f}(\mathbf{x})$. In case of interest, where the goal is solving the turbulent combustion products, \mathbf{x} is the initial thermochemical state and $\mathbf{f}(\mathbf{x})$ is the result of an ordinary differential equations (ODE) solution, given as the composition at the final interval time. In this technique, the entries in the table are referred to as leaves of a binary tree. The n th leaf includes: its location $\mathbf{x}^{(n)}$ (initial state); the value of the function $\mathbf{f}^{(n)} = \mathbf{f}(\mathbf{x}^{(n)})$ (final state); the ellipsoid of accuracy and the gradient matrix $\mathbf{A}^{(n)}$, also called sensitivity matrix. Given an error tolerance, ε_{tol} , the linear approximation, based on neighboring solutions, is returned with error, ε , smaller than ε_{tol} . The region where this occurs is called region of accuracy, which is approximated by an ellipsoid of accuracy (EOA). The costs of searching and testing whether a point is covered by the EOA can be reduced by introducing an affine space of dimension smaller than the composition vector dimension, which contains the orthogonal projections of the query point \mathbf{x} and the ellipsoid E, whose geometry is recalculated every time the affine space is reset.

According to Lu and Pope (2009), the ISAT consists of three fundamentals steps: retrieve, growth and addition. Given a query mixture composition, the retrieve attempt objective is to identify an ellipsoid of accuracy that covers that

point. Four search types can be used in an attempt to identify the EOA. If this is found, a linear approximation based on that leaf is returned. The lower cost of searching is achieved by this algorithm when all four search types are used. The growth attempt is used when the retrieve attempt is not successful. In this case, the leaf's ellipsoid of accuracy can only be grown if the linear approximation is accurate. One or more leaves near to the query point are evaluated and selected for the trial of growth. For each leaf, the linear approximation error (ϵ) is evaluated and, if smaller than the given error tolerance (ϵ_{tol}), the EOA increases to cover the query point. It is possible to identify leaves that are candidates for growth, which leads to a reduction of the computational time required for simulation. The addition is performed when the recovery and growth attempt are not successful; this step can be done by calculating the ellipsoids, followed by inclusion of a new leaf on the table. If the table is already full, the linear approximation is returned without altering the table or replacing the last leaf on a most recently used leaves list or replacing the last leaf on a most frequently used leaves list. A new development is the error checking and correction (ECC), which is particularly effective in reducing the error incurred, with low computational cost. This consists in occasionally checking the incurred error, and, if this error exceeds ϵ_{tol} , the EOA decreases. Recoveries are randomly selected for ECC treatment, with the probability chosen so that, on average, the frequency of ECC events are fixed. Tests made by the authors show that the reliable control of errors, ECC, halves the CPU time required for simulation.

Abrol (2009) presents a modified ISAT formulation, called mISAT, which aims to improve the recovery rate of the technique. The application of this version shows an improved performance when compared to the original algorithm, in terms of the number of recoveries made in a search on the binary tree. Earlier versions of ISAT involve the use of a single binary search tree to build the database. However, the disadvantage of this unique binary tree is that the record closest to the desired composition is not always recovered. To overcome this deficiency, some search techniques involving multiple binary trees are suggested to increase the probability of finding the closest record. This work shows that various structures of binary trees can be used to distribute, on the database, the records that are stored as leaves at the end of the branching binary tree structure. These various tree structures allow more accurate approximations of the sensitivity matrix. Such distribution techniques are based on uniform distribution and k-means clustering. The use of centered k-means clustering results in a more uniform spatial distribution of the leaves records. An alternate framework for building a database with the ISAT balanced tree, known as AVL tree, is suggested to reduce the search time.

Pope and Ren (2009) focus on issues concerning the efficient implementation of detailed chemical kinetics mechanisms in computational combustion. It presents a description of the computationally efficient operator splitting scheme to alleviate the large computational requirements in reactive flows. By using operator splitting, part of the governing equations which contains the chemical reaction term is separated from other parts such as transport terms. The scheme based on Strang splitting requires, for each time step Δt , two $\Delta t/2$ reactions sub-steps and one Δt transport sub-step. When applied to a one-dimensional flame computation of Methane/air, using detailed chemical kinetics, an overall acceleration factor of about 7.5 is reached after the ISAT table is built and filled. Tests made by the authors indicate that the proposed splitting schemes achieved a second order accuracy in time.

Liu and Pope (2005) present a study of the characterization of the error inherent in the use of ISAT algorithm, determining the accuracy of local and global errors associated with storage and recovery. A turbulent flame using a skeletal mechanism for Methane with 16 species and 41 reactions is modeled. A discussion about the three different strategies for EOA growth is presented, and the use of the simplest growth mode is recommended. The authors discuss the cumulative distribution function of the local error and characterize the global error for the particular configuration studied. Also investigated are the possible reasons for inaccuracy due to the ellipsoid of accuracy growth. The EOA growth increases the number of recoveries in ISAT with a specified table size; however, this is also the major cause of large local errors. There are three possible reasons for inaccuracies when the EOA grows. First, it is possible that over the region of accuracy (ROA) the tabulated function is significantly nonlinear. Therefore, the analysis of Taylor series would be inaccurate, leading to the possibility of non-monotonic behavior of the error, ϵ . If this is the case then, given a query point, there may be regions where ϵ is greater than ϵ_{tol} , whereas the query point itself has ϵ smaller than ϵ_{tol} . In such a situation, the EOA increases to cover these inaccurate regions, introducing the possibility of subsequent inaccurate recoveries. This problem, if it exists, tends to disappear when ϵ_{tol} becomes sufficiently small, because the ROA shrinks and the accuracy of the Taylor series increases in this ROA. Second, the ROA may not be convex, with a hyperbolic nature. In this case, all growth strategies may result in EOAs which include inaccurate regions after growing. Thirdly, even if the ROA is convex (eg, ellipsoidal), the EOA growth strategies can lead to the inclusion of imprecise regions. The authors demonstrate that large local errors are not caused by non-monotonic behavior of ϵ , but they are associated with non-convex regions of accuracy. This work also characterizes the global error in the calculations using the transported PDF, resulting from local errors incurred in the ISAT algorithm. This is accomplished by considering the error of the average mass fractions of chemical species and temperature. For the lowest ϵ_{tol} , the calculated results are found to be improved, indicating that the ISAT errors are small when compared to statistical errors for $\epsilon_{tol} \leq 10^{-4}$; and the global errors vary linearly with ϵ_{tol} .

3. IN SITU ADAPTIVE TABULATION

The above mentioned works represent the state of the art of in situ adaptive tabulation techniques applied to the simulation of combustion problems. A common feature, among all, is the attempt of improving the basic ISAT algorithm (Pope, 1997) by different strategies. The present work is also motivated by such an improvement attempt, and thus extends the work of Cunha (2010) by a simple performance improvement strategy.

In what follows, the ISAT technique applied to a classical partially stirred reactor model is briefly described. Then, the proposed modifications are presented. The obtained results are analyzed in terms of the technique performance, memory and CPU usage.

Under the hypothesis of homogeneous flow within a control volume, typical of the so-called Partially Stirred Reactor PaSR (Sabel'nikov and Figueira da Silva, 2002), a reactive flow thermodynamic state in a reactor may be given by the composition vector:

$$\varphi \equiv (h, p, Y_1, \dots, Y_{n_s}), \quad (1)$$

where h is the enthalpy, p is the pressure and Y_i is the mass fraction of each of the n_s chemical species. A PaSR is a continuous flow reactor, characterized by a residence time of the burning mixture, a mixing time of the reactants, a chemical time, and the composition of the fresh gases mixture.

Considering \mathbf{F} the rate of change associated to micro mixing transport and \mathbf{S} the rate of change associated to the chemical reactions which are given by Arrhenius laws, this composition evolves in time according to the equation:

$$\frac{d\varphi}{dt} = \mathbf{F}(t) + \mathbf{S}(\varphi, t). \quad (2)$$

Classically, the reactor is modeled by an ensemble of stochastic particles which describe the state of the burning gases. In order to describe the micro mixing process, the pairwise-mixing stirred reactor (PMSR) model, proposed by Pope (1997), which is further characterized by a pairwise time, τ_p . At each time step, the particles are paired, and in each pair, the micro mixing model is applied. All the particles exchange partners based on a fixed pairwise time and relax to the mean composition of the partners, based on a prescribed mixing time scale, τ_m . Further details may be found in Cunha and Figueira da Silva (2010).

The ISAT algorithm works as follows: given the initial state, $\varphi_0 \equiv (h, p, Y_1, \dots, Y_{n_s})$, the final states are approximated by a linear extrapolation, $\mathbf{R}^l(\varphi, t) \equiv \mathbf{R}(\varphi_0, t) + \mathbf{A}(\varphi_0, t)\delta\varphi$, where $\delta\varphi = \varphi - \varphi_0$; this extrapolation occurs from a neighboring solution using a sensitivity matrix, \mathbf{A} . The technique attempts to control the error of these approaches, $\varepsilon \equiv \|\mathbf{R}(\varphi, t) - \mathbf{R}^l(\varphi, t)\|$, by defining an ellipsoid of accuracy (EOA) around the initial state, which is adjusted according to the given error tolerance.

The tabulation on the tree is made while the simulation is proceeds as performed and begins without any leaf at the tree. The ISAT records, which are referred to as leaves of a binary tree, contain the information about the initial states, the reaction mapping, the mapping gradient matrix and the ellipsoid of accuracy matrix. Figure 1 allows verifying, through a simplified diagram of the ISAT method, that the technique consists of three basics steps: an attempt to retrieve, an attempt to grow and an addition.

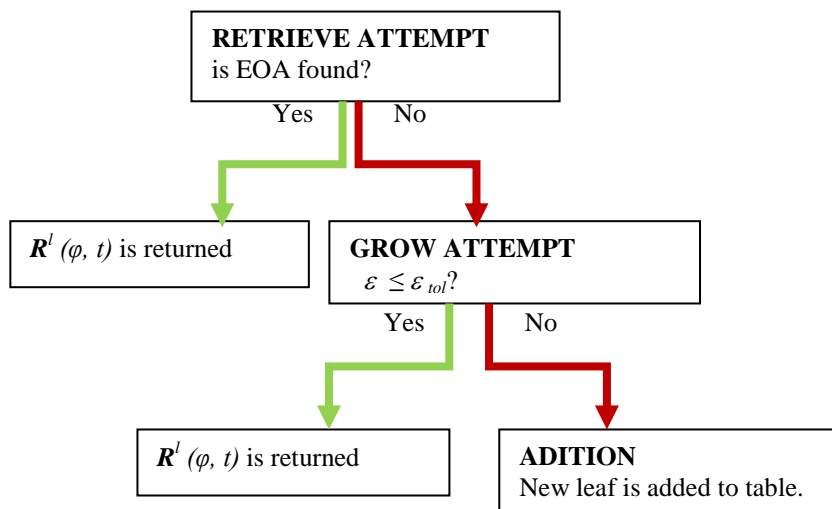


Figure 1. Simple diagram of the ISAT steps.

In the original version of the technique implemented by Cunha (2010), the ISAT tabulation begins at $t=0$, i.e., at the beginning of the simulation. As consequence, the initial query states that are passed to the ISAT algorithm may include mixture compositions that are not representative of the steady state behavior of the combustion process simulated. Indeed, such is the case, for instance, when the initial enthalpy and composition prevailing within the reactor are far from an equilibrium state or from the state of the fresh mixture. In case where the mixing process is slow when compared to the residence time, the initial reactor transient is long, thus, during this initial transient, there is a large probability that the leaves being added or grown in the binary tree are not pertinent to the steady state reactor operation.

As consequence, the resulting binary tree would be far larger than actually needed, possible also leading to relatively inaccurate recoveries. This would be aggravated in situations where the maximum number of the tree entries is fixed, which is most often the case. The modifications introduced in the present work address such issues by initiating the ISAT algorithm after an initial transient, where only direct integrations are performed. The results of parametrical analysis are presented where is examined the influence of the initial tabulation time, t_i , on the various performance parameters of the ISAT technique. The chosen values of t_i lie in range $0 \leq t_i \leq 8 \tau_r$. The examined parameters are:

- Number of stochastic particles;
- Initial tabulation time, t_i ;
- Total simulation time, τ^* ;
- Residence time, τ_r .

For each simulated case, the influence of these four parameters is demonstrated on the following performance results:

- Memory computer usage;
- CPU time;
- Global error;
- Mean error;
- Binary tree height;

Considering ψ a property, the mean relative error of that property, over an interval $\Delta\tau$, is defined as:

$$\langle \varepsilon_{r,\psi} \rangle_M \equiv \frac{1}{\Delta\tau} \int_t^{t+\Delta\tau} \varepsilon_{r,\psi}(t') dt', \quad (3)$$

and the global error is defined as:

$$\varepsilon_g \equiv \frac{1}{\Delta\tau} \int_t^{t+\Delta\tau} \frac{\| \langle \varphi \rangle(t')_{DI} - \langle \varphi \rangle(t')_{ISAT} \|}{\| \langle \varphi \rangle(t')_{DI} \|} dt', \quad (4)$$

where $\langle \varphi \rangle$ denotes the average composition vector, and DI denotes the calculation using direct integration of the governing equation and ISAT denotes the use of the original algorithm.

Each of these results, r , are compared with the previous version of the technique and the savings, s , are given in terms of percentage calculated as:

$$s = \frac{r_{old} - r_{new}}{r_{old}} \times 100\% . \quad (5)$$

More detailed information about the implemented technique and the parameters definitions can be found in Cunha and Figueira da Silva, 2010.

4. RESULTS AND DISCUSSION

In order to allow for comparisons involving total simulation times, much larger than the residence time, a CO/O₂ mixture, which is one of the simplest chemical kinetics mechanisms available, was modeled using the Gardiner (2000) mechanism, with 4 species and 3 reactions. Even if this chemical mechanism is simple, indeed, it allows for characterizing the proposed modifications. Table 1 presents the initial conditions for simulation; these parameters are used in all simulated cases.

Table 1. Initial conditions for simulation.

ISAT error tolerance	1.0×10^{-3}
equivalence ratio	0.7
inlet gases temperature (K)	300
inlet gases pressure (atm)	1
initial system temperature (K)	2950
initial system pressure (atm)	1
mixing time (μs) - τ_m	100
pairwise time (μs) - τ_p	100

The parametrical study performed considered four different cases, for which the variable parameters are shown in Tab.2. In these cases the binary tree is not saturated by the additions, i.e., the maximum number of allowed entries was not reached. The table below presents the number of stochastic particles inside the reactor, the number of residence times to start the tabulation, the total number of residence times simulated and the duration of each residence time.

Table 2. Variable parameters of analysis.

	CASE 1	CASE 2	CASE 3	CASE 4
number of particles	100	256	1024	1024
initial tabulation time - t_r	$2\tau_r$	$2\tau_r$	$2\tau_r$	$2\tau_r$
total simulation time - τ^*	$12\tau_r$	$12\tau_r$	$12\tau_r$	$12\tau_r$
residence time (μs) - τ_r	200	200	200	1000

Note that in cases 1, 2 and 3 the number of stochastic particles within the reactor is progressively increased. Case 4 has the same particle number as case 3, but the residence time is five times greater; in this configuration, where the pairwise and mixing time scales are small when compared with the residence time, the reactor behaves as if the pairing and mixing process occurred instantly. This situation slowly resembles a pairwise stirred reactor.

Previous results have shown that in situations where the PaSR is close to extinction, such as the present one, the value of the statistical seed may drastically influence the computed results. Choosing a single seed for all computations does introduce bias, but avoids this anomalous behavior of the chosen mixed model. The cases were simulated using the same seed for all statistical processes, i.e., inflow, outflow and mixing.

The first analysis presented is related to the number of ISAT steps, as illustrated in Fig. 2 and 3. That figures present the case 1 results with the ISAT characteristics outputs, which are the number of additions, growths, retrieves and direct evaluations and also show the tree height. Figure 2 refers to the new strategy with the modifications of the algorithm and Fig. 3 refers to the previous strategy. The comparisons between these figures allow verifying, for instance, the drastic difference in the trees size. Indeed, the new strategy presents a tree with a more uniform distribution, as its height is 23 times less than the original strategy tree. Another advantage revealed by the improvements is the fact that there are more recoveries than additions. Since the necessary calculations to add a new leaf in a binary tree is four orders of magnitude more time consuming than that to make a recovery, it can be conclude that the new strategy is more efficient than its predecessor, with regard to processing time. Although no direct evaluations were performed in previous formulation, this simulation is longer to be performed, because more additions than retrieves were made.




Figure 2. Characteristics outputs for the new strategy.




Figure 3. Characteristics outputs for the original formulation.

The trends indicated in Fig. 2 and 3 are systematically reported for the four cases studied in Tab. 3, where the performance parameters, defined in section 3, are given in terms of relative savings. A first examination of this table allows verifying that the simple strategy proposed here, leads to substantial savings of all outputs parameters. In particular, it is observed that as the number of particles is increased, both the CPU time and tree height savings increase reaching 28% and 99%, respectively, for case 3. This indicates that the resulting binary tree height is considerably smaller, suggesting greater homogeneity and less time searching, which leads to a greater efficiency of the technique. All the simulations show that the tree size tends rapidly to a steady state.

This steady state behavior is characteristic of a search table that adequately tabulates the accessed composition space. The CPU time savings is also significant and is expected to increase with the particle number even further, which could be of particular interest to large eddy simulations.

Table 3. Relative savings for each evaluated case.

	CASE 1	CASE 2	CASE 3	CASE 4
memory %	22.8	20.63	17.84	10.79
CPU time %	18.32	18.45	28.27	18.09
global error %	40.56	83.17	65.25	63.31
mean error %	41.38	65.62	26.54	50.74
height %	95.75	97.83	99.12	99.79

The mean and global errors exhibit similar saving trends with respect to the influence of the particle number. The obtained gains are due to the compositions that are calculated during the transient regime and that are not all tabulated. Also, the number of growth is not greater than the number of addition, the EOA grows to provide a better estimate of accuracy in the region, which is not observed in the previous formulation.

Both these errors seem to possess a non-monotonic behavior which leads to the presence of a maximum for 256 particles. These could be related to the actual value of the computed mean and global errors, which decrease as the number of particles within the reactor increase. Thus a relative savings would measure the difference between quantities that are very small.

In all cases above, the tests show that the number of recoveries exceeds the number of additions in new formulation. This indicates that the region accessed by the calculations has more compositions and it needs less direct evaluations, i.e., the tree is richer in the new ISAT formulation. This behavior was not observed in the former formulation, where the number of addition equals the number of direct evaluation.

Finally, one of the major drawbacks found in the original ISAT was the memory demand. The new strategy decreases the memory costs by 18%, in average, compared to the original one.

The cases studied here show, properly, the advantages of the new formulation. Considering all tests, we observe that the average reduction in the tree height is 98.12%; the CPU time was reduced about 20.78%; the global and mean errors were decreased more than 63% and 46%, respectively.

This underscores the benefits of the proposed strategy. Future work will examine the savings incurred in more complex chemical systems, such as Methane/air mixtures.

5. ACKNOWLEDGEMENTS

This work has been performed while Luis Fernando Figueira da Silva was on leave from Institute Pprime, *Centre National de la Recherche Scientifique*. Support from Petrobras is gratefully acknowledged.

6. REFERENCES

- Abrol, Sidharth, 2009, "Advanced Tabulation Techniques for Faster Dynamic Simulation", Ph.D. Dissertation, The University of Texas, Austin.
- Andrade, F.O., Figueira da Silva, L.F., Mura, A., 2011, "Large Eddy Simulation of turbulent premixed combustion at moderate Damköhler numbers stabilized in a high speed flow." *Combustion Science and Technology*.
- Cunha, A.B., 2010, "Reduction of Complexity in Combustion Thermochemistry", M.Sc. Dissertation, Pontifícia Universidade Católica do Rio de Janeiro, Rio de Janeiro.
- Cunha, A.B.; Figueira da Silva, L.F., 2010, "Characterization of an Adaptive Technique to Reduce Combustion Thermochemistry", *Proceedings of ENCIT 2010*.
- Gardiner, W.C., 2000, "Gas Phase Combustion Chemistry", Springer, Ney York. Liu, B.J.D. e Pope, S.B., 2005, "The Performance of In Situ Adaptive Tabulation in Computations of Turbulent Flames", *Combustion, Theory and Modelling*, Vol. 4, pp. 549–568.
- Lu, L., Pope, S.B., 2009, "An Improved Algorithm for In Situ Adaptive Tabulation", *Journal of Computational Physics*, Vol. 228, pp. 361–386.

- Sabel'nikov, V., Figueira da Silva, L.F., 2002 “Partially stirred reactor: study of the sensitivity of the Monte-Carlo simulation to the number of stochastic particles with the use of a semi-analytic, steady-state, solution to the pdf equation.” *Combustion and Flame*, Vol. 129, n. 1–2, pp. 164–178.
- Pope, S.B., 1997, “Computationally efficient implementation of combustion chemistry using in situ adaptive tabulation”, *Combustion, Theory and Modelling*, Vol. 1, No. 1, pp. 41–63.
- Pope, S.B., Ren, Z., 2009, “Efficient Implementation of Chemistry in Computational Combustion ”, *Flow Turbulence Combustion*, Vol. 82, pp.437–453