Cambridge Centre for Computational Chemical Engineering

University of Cambridge

Department of Chemical Engineering

Preprint ISSN 1473 – 4273

Real-time evaluation of a detailed chemistry HCCI engine model using a tabulation technique

Sebastian Mosbach[1], Ali M. Aldawood[1], Markus Kraft[1]

released: 28 June 2007

[1] University of Cambridge
Department of Chemical Engineering
New Museums Site, Pembroke Street
Cambridge CB2 3RA
United Kingdom

Preprint No. 48

c4e

Key words and phrases: HCCI engine modelling, detailed chemistry, storage/retrieval, transient simulation

Edited by
Cambridge Centre for Computational Chemical Engineering
Department of Chemical Engineering
University of Cambridge
Cambridge CB2 3RA
United Kingdom.

Fax: + 44 (0)1223 334796
E-Mail: c4e@cheng.cam.ac.uk
World Wide Web: http://www.cheng.cam.ac.uk/c4e/

Abstract

A storage/retrieval scheme has been implemented for a Stochastic Reactor Model (SRM) for Homogeneous Charge Compression Ignition (HCCI) engines which enables fast evaluation in transient multi-cycle simulations. The SRM models combustion, turbulent mixing, and convective heat transfer during the closed-volume part of the engine cycle employing detailed chemical kinetics. In contrast to previously developed storage/retrieval techniques which tabulate chemistry only, our method stores, retrieves, and interpolates output quantities of the entire internal combustion engine model, i.e. the SRM. These quantities include ignition timing, cumulative heat release, maximum pressure rise rate, and emissions of CO, CO_2, unburnt hydrocarbons, and NO_x, as functions of equivalence ratio, octane number, and inlet temperature for instance. The new tool is intended to be used for performing a variety of otherwise exceedingly expensive computational tasks such as multi-cycle multi-cylinder simulations, transient operation and control, optimization of engine operating parameters, design of experiments, and identification of parameters for achieving stable HCCI operation over a wide range of conditions. Using transient control as an example, we show that, when coupled to a commercial 1D CFD engine modelling package, the tabulation scheme makes such simulations feasible and convenient.

Contents

1 Introduction **3**

2 The Stochastic Reactor Model (SRM) **5**

2.1 Brief outline 5

2.2 Validation 5

3 The tabulation scheme **6**

4 Tabulating the SRM **7**

5 HCCI transient control simulations **12**

6 Conclusions **14**

1 Introduction

Homogeneous Charge Compression Ignition (HCCI) engines promise high thermal efficiency combined with low levels of nitric oxide and particulate matter emissions. However, due to the absence of an immediate means of triggering ignition, stable operation over a wide range of conditions and transient control have proven most challenging and have so far prevented commercialization. Numerous potential strategies for mastering these difficulties have been investigated in the past, including for instance (multiple) direct injection, exhaust gas recirculation (EGR), variable valve timing (VVT), variable compression ratio (VCR), octane number (ON) variation, and fast thermal management.

Modelling and simulation can contribute significantly to understanding the fundamental processes involved in HCCI and speed up engine development and design processes, in addition to reducing their cost. Various modelling approaches have been pursued, ranging from single- and multi-zone models (e.g. [1]), via models based on Computational Fluid Dynamics (CFD, e.g. [2]) and Large Eddy Simulation (LES, e.g. [3]), to approaches based on Probability Density Function (PDF) transport methods (e.g. [4, 5]). From a user's point of view, the trade-off between the computational expense and the quality or detail of the predictions plays a major role in the decision which method to employ. One of the key difficulties in HCCI simulation is the fact that HCCI is determined predominantly by chemical kinetics, which is time-consuming to simulate accurately, especially when coupled to turbulent flow as found inside the engine cylinder.

Amongst a variety of attempts to accelerate combustion simulation in general, in the past, most techniques have focussed on speeding up the evaluation of chemical kinetics. We do not aim for an exhaustive review of the extensive literature in this field here, but instead restrict ourselves to storage/retrieval methods, because the method presented in this paper belongs to the same category (even though it is not restricted to chemistry only). The basic idea is to store the results of evaluations of a certain function, and when subsequent evaluations close to the already evaluated points are requested, to reuse some of the stored information. This usually involves some form of interpolation (typically with polynomials) between available data points. Frequently, significant gains in computational efficiency are reported, in some cases of up to four orders of magnitude.

One of the most commonly used storage/retrieval techniques for chemical kinetics is *In Situ* Adaptive Tabulation (ISAT) [6], in which the accessed region of the chemical composition space is tabulated by calculating the state space evolution on demand during the simulation, rather than beforehand in a preprocessing phase. Linear interpolation is used in order to obtain missing data points during retrieval. ISAT was successfully employed in engine simulations, for example for predicting NO_x with detailed kinetics coupled to a hybrid CFD/PDF method [7]. In [2], an improvement of ISAT, called Database for On-Line Function Approximation (DOLFA), was developed. Coupled to the CFD-code Star CD, significant speed-up was achieved in applications to port and direct injected HCCI.

Along a similar line of thought, referred to as repro-modelling, (orthogonal) polynomials of up to eighth order [8, 9] have been used to parametrize the evolution of chemical kinetic systems in composition space.

Response Surface Methodology (RSM) is another technique based on similar ideas,

i.e. fitting a model, usually a simple algebraic expression such as a low-order polynomial, to data. It is most frequently employed, though, to fit empirical models to experimental measurements using regression. See [10] for a comprehensive introduction and overview.

In [11], as part of a technique called solution mapping, the RSM was applied to the optimization of kinetic parameters in a reaction mechanism describing methane oxidation. Piecewise Reusable Implementation of Solution Mapping (PRISM) [12, 13] is an extension of that technique employing quadratic response surfaces on hypercubical subdivisions of the composition space.

As an example outside the area of chemical kinetics, approximating experimental measurements by a response surface has also proven extremely useful for design and simulation problems in aerospace and automotive applications [14]. Focussing on specific aircraft or vehicle components other than a combustion engine, such as a throttle, the RSM has been shown to enable rapid controller design and to be suitable for inclusion into real-time simulations.

Criteria by which the above-mentioned methods can be distinguished include (1) the type of function (e.g. order of polynomial) which is used to represent model responses, (2) whether the representation is on a single domain or piecewise on multiple domains or on a mesh, and (3) whether the representation reproduces all known data points exactly or in general only approximately (as in a least-squares fit for example).

In [15], the one-dimensional CFD engine simulation package GT-Power® was coupled to Simulink® in order to control ignition timing by adjusting the residual gas fraction. As user-defined combustion model, they used an empirical expression (fitted to experimental data) which can be evaluated at minimal computational expense. This highlights the necessity of fast evaluation for control and multi-cylinder studies.

The purpose of this article is to demonstrate the feasibility of a storage/retrieval technique developed not only for chemical kinetics, but for an entire HCCI engine model. The model we consider is the Stochastic Reactor Model (SRM) employed previously (e.g. [4, 5]), which simulates combustion, turbulent mixing, and convective heat transfer during the closed-volume part of the engine cycle using detailed chemical kinetics. Predicted quantities include ignition timing, cumulative heat release, maximum pressure rise rate, and emissions of CO, CO_2, unburnt hydrocarbons, and NO_x. The newly developed tool enables fast evaluation of the SRM for applications involving a large number of cycles. We believe it will be useful not only for identifying parameters for achieving stable HCCI operation over a wide range of conditions, but also for other computationally intensive tasks like control, optimization, multi-cycle multi-cylinder simulations, etc. The value of providing a tool for fast model evaluation lies mainly in putting the solution of such problems within reach – which would otherwise be impractical. Here, we focus on explaining fundamentals and properties of our tabulation method and on showing its practicability rather than on a particular application, which is deferred to a subsequent publication.

This paper is structured as follows. In section 2, we briefly outline the main features of the Stochastic Reactor Model (SRM) for HCCI engines. In section 3 we describe the implemented tabulation method in general terms and apply it in section 4 to our engine model. Subsequently, we demonstrate the usefulness of the developed software tool in

section 5 using an HCCI transient control example in which we control auto-ignition timing over a range of equivalence ratios by varying the octane number of the fuel. In the final section we summarize our findings and indicate directions of future work.

2 The Stochastic Reactor Model (SRM)

2.1 Brief outline

The Stochastic Reactor Model (SRM) has been explained in detail in previous publications (e.g. [4]), which shall not be repeated here. Besides, as far as the present work is concerned, in-depth understanding of the model is not required. However, its main features which should be kept in mind are the following. The SRM is based on Probability Density Function (PDF) transport [16] methods, uses detailed chemical kinetics, contains a turbulent mixing model, and accounts for convective heat transfer between the charge and the cylinder wall. The detailed chemical kinetic scheme we employ is designed for Primary Reference Fuels (PRFs), covers NO_x chemistry, and contains 157 species and 1552 reactions. Output quantities of the model include distributions of composition and temperature, and in particular emissions of CO, CO_2, unburnt hydrocarbons (UHC), and NO_x. We note that accurate representation of in-cylinder inhomogeneities is crucial for predicting quantities like emissions and maximum pressure rise rates. The SRM has also been successfully applied to simulate partial charge stratification resulting from direct injection [5, 17, 18].

One of the key parameters of the model is the number N_{par} of stochastic particles. The statistical fluctuations (inherent to any PDF method) decrease with increasing particle number, or more precisely, the statistical error is proportional to $N_{\text{par}}^{-1/2}$. At the same time, the computational effort is proportional to N_{par}. As with every computational model, all physical predictions should be asymptotically independent of all numerical parameters. We find that for the number of particles, this regime is reached at about $N_{\text{par}} = 100$. Compared to other modelling approaches, the SRM has shown an excellent ratio between computational expense and quality of predictions (notably of emissions).

2.2 Validation

For the validation of our model, we used measurements carried out at Sandia National Laboratories on a Cummins B-series medium-duty 6-cylinder Diesel engine modified for single cylinder HCCI operation, as reported in [1]. Some basic information about the engine is given in Table 1(a). The engine was run port-injected on PRF80, a primary reference fuel consisting of 80% iso-octane and 20% n-heptane by volume, which has octane number 80 by definition. We used a crevice volume of 2.1% of the clearance volume as suggested in [1]. We fixed parameters of our model by matching pressure profiles for motored and for fired operation (see Fig. 1(b)) at steady state for a single operating point at an equivalence ratio of 0.4.

Bore	102 mm
Stroke	120 mm
Displaced volume	981 cm^3
Connecting rod length	192 mm
Compression ratio (CR)	14
Intake valve opening (IVO)	3° BTDC
Intake valve closing (IVC)	155° BTDC
Exhaust valve opening (EVO)	120° ATDC
Exhaust valve closing (EVC)	8° ATDC
Speed	1200 RPM
Intake pressure	1 bar
Intake temperature	341 K

(a) *Engine specification and operating condition.*

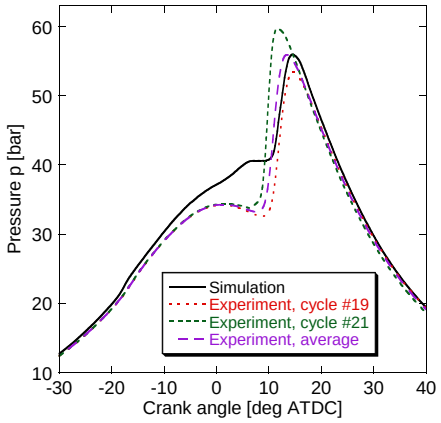

(b) *Comparison of simulated with experimental in-cylinder pressure profiles.*

Figure 1: *Validation of the simulation against experimental data.*

3 The tabulation scheme

In a number of relevant applications, as mentioned in the introduction, it is necessary or at least very useful to evaluate a function at many points within a certain region of the parameter space. If a single evaluation is costly, it appears natural to reuse previous evaluations if possible. In general, however, when an evaluation is requested, the query point does not coincide with one of the already evaluated points. Therefore, the model response needs to be interpolated making use of information stored in the table. Here, we use cubic natural splines for this purpose, which are piecewise third order polynomials whose coefficients are uniquely determined by the requirement that they are twice continuously differentiable through the points where they are fitted together, in conjunction with the 'natural' boundary condition that the second derivatives vanish on the boundary. There exist standard algorithms for implementing cubic natural spline interpolation [19]. We note that, since the tabulated surfaces are differentiable, derivatives of quantities can be used in optimization algorithms.

The underlying point set in parameter space where the tabulated function is evaluated exactly is in principle arbitrary. For the sake of simplicity, we restrict ourselves to Cartesian lattices on a hypercubical region with possibly irregular grid spacing.

Figure 2 illustrates the interpolation generically for a non-linear function f of two parameters x_1 and x_2, tabulated on a 4×4 regularly spaced Cartesian lattice. In order to obtain an approximation of the function f at the query point $(x_1^{\mathrm{q}}, x_2^{\mathrm{q}})$, given the exact function values on the vertices of the lattice, the interpolation proceeds in two stages. Firstly, for each of the lattice coordinates on the x_1-axis, a one-dimensional interpolation along the x_2-direction is carried out, leading to the points indicated by squares. Subsequently, another one-dimensional interpolation is carried out, but this time for $x_2 = x_2^{\mathrm{q}}$ along the x_1-direction, yielding the desired approximation of f at the query point $(x_1^{\mathrm{q}}, x_2^{\mathrm{q}})$, indicated by a hexagon in the figure.

Under favorable conditions, up to several thousand quantities can be tabulated without

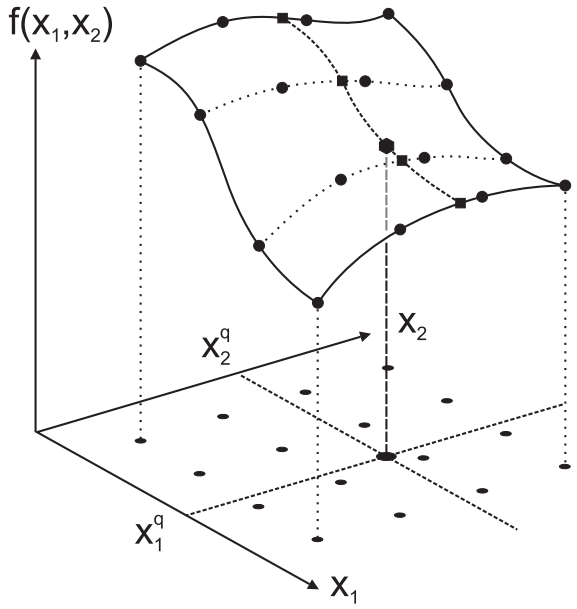

Figure 2: *Sketch of a non-linear function f of two parameters x_1 and x_2 tabulated on a 4×4 regularly spaced Cartesian lattice. Interpolation is required for retrieval at the query point (x_1^q, x_2^q).*

significant difficulties. In contrast to that, the number of parameters, which the tabulated quantities can depend on, is absolutely critical. As pointed out in [6], generally the tabulation of any function in high dimensions is not feasible. This is essentially a consequence of the fact that the number of points which need to be evaluated scales very poorly with the dimension, i.e. the number of parameters. Depending on the non-linearities of the tabulated quantities and available computing facilities, our tabulation method can be useful in up to four dimensions. Non-linearities are a main cause of tabulation errors or interpolation errors [19]. On the one hand, if a quantity is almost linear, i.e. if it exhibits only small curvature, it can be easily tabulated with a sparse grid. On the other hand, severely non-linear quantities, for example quasi-discontinuous functions, possess locally large curvatures and therefore require a relatively dense grid in those regions.

In summary, recalling the criteria mentioned in the introduction, our tabulation uses cubic natural splines to interpolate between points on a hypercubical, irregularly spaced Cartesian lattice, and reproduces each of the evaluated points exactly.

4 Tabulating the SRM

From an abstract point of view, our engine model, the SRM, is simply a function which can be tabulated like any other. The SRM assigns to a set of parameters (say, equivalence ratio, octane number, and temperature) at inlet valve closure a set of values at exhaust valve opening (crank angle at 50% heat release, maximum pressure rise rate, emissions, etc.). Symbolically:

$$\text{SRM} : (\Phi, \text{ON}, T_{\text{IVC}}, \dots) \mapsto (\text{CA50}, \text{PRR}_{\text{max}}, X_{\text{CO}}, X_{\text{NO}_x}, \dots)$$

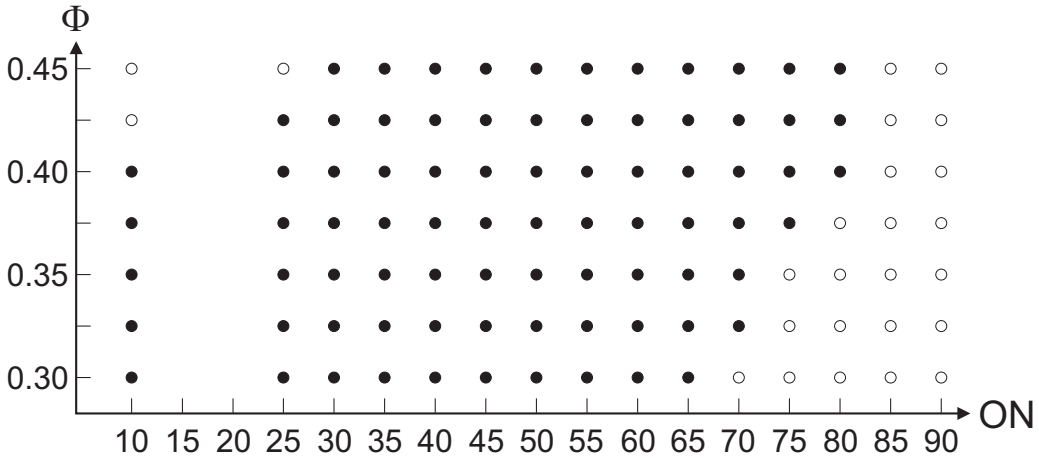

Figure 3: *Lattice in parameter space used for the transient control example in section 5. The empty circles represent the vertices removed during the masking procedure.*

On a conventional desktop computer, it takes between one and two hours to evaluate this function at a single point in parameter space. It should be pointed out that, even though the tabulation technique described here can be applied in principle to any engine model, indeed to any reasonably well-behaved function, if a single model evaluation is too costly (such as CFD with detailed chemistry as in [2] or [7]), it becomes impractical.

In this paper, since the focus is on feasibility, we consider for simplicity only the two-dimensional parameter space spanned by the parameters equivalence ratio Φ and octane number ON. As ranges we have chosen 0.3-0.45 for Φ and 10-90 for ON, on a 7×15 grid which is shown in Fig. 3. All remaining engine operating parameters, such as speed, inlet temperature, pressure, etc. are kept constant. The most important of the tabulated quantities include the crank angle at 50% heat release (CA50), cumulative heat release (CHR), maximum pressure rise rate ($\mathrm{PRR}_{\mathrm{max}}$), and emissions of CO, CO_2, unburnt hydrocarbons (UHC), and NO_x. In total, 900 quantities have been tabulated, which include in-cylinder pressures and temperatures at 443 different crank angles each throughout the cycle.

Frequently, the region of interest within the parameter space is not a hypercube. But in order to expend minimal amounts of computation time, ideally, the function to be tabulated should be evaluated only in regions which are *a priori* known to be of interest. For example, one would like to tabulate an engine model only in regions where combustion takes place and not in large regions of misfire. While such knowledge is beforehand often unavailable, here, where this applies as well, there exists a way to at least partially avoid this difficulty. As mentioned in section 2, the convergence properties for the quantities of interest are known. In practice, it turns out that for many important quantities such as CA50 and CHR, results obtained with 10 particles are close to the 'converged' results at $N_{\mathrm{par}} = 100$. Therefore we employ the following strategy, which we refer to as 'vertex masking'. We calculate a preliminary table with $N_{\mathrm{par}} = 10$, discard 'irrelevant' vertices, and then evaluate the remaining ones with full accuracy, i.e. 100 particles. For CA50, we chose as valid range -5 to 10 CAD ATDC and for the cumulative heat release a threshold of 600 J, values below that indicating misfire. Why we chose such a wide interval of ignition timings, ranging even to a few degrees before TDC, will become clear in the

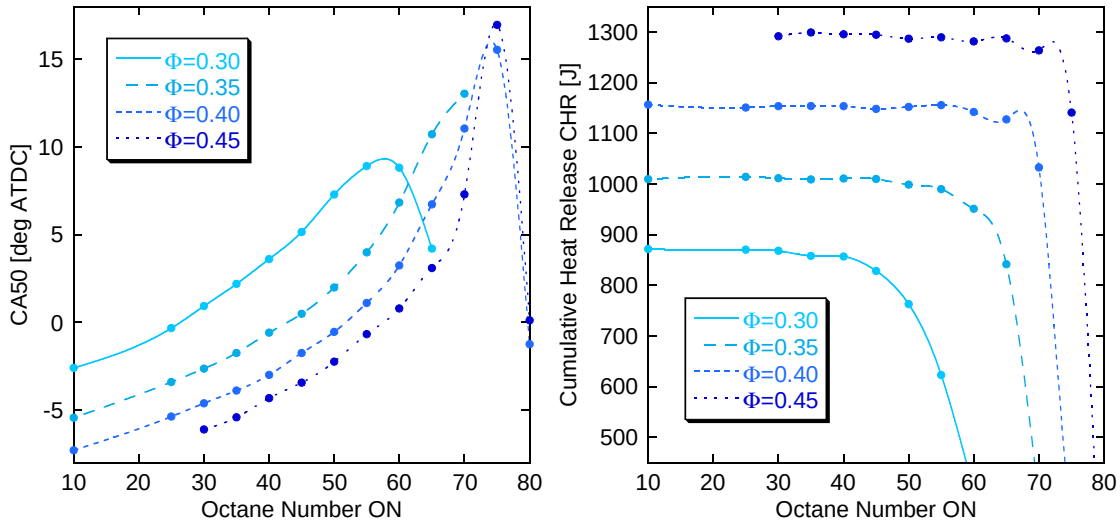

(a) *Crank angle at 50% heat release (CA50).* **(b)** *Cumulative Heat Release (CHR).*

Figure 4: *Tabulated CA50 and cumulative heat release as functions of octane number for different equivalence ratios. The dots correspond to the actual evaluations of the SRM and are therefore exact, whereas the lines show the interpolation.*

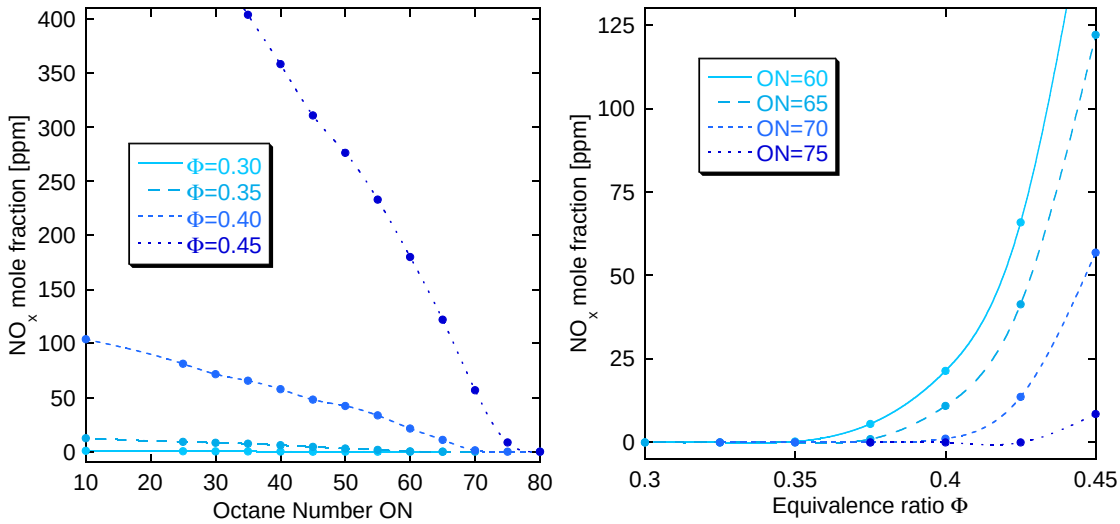

(a) *NO_x as function of octane number for different equivalence ratios.* **(b)** *NO_x as function of equivalence ratio for different octane numbers.*

Figure 5: *Tabulated NO_x emissions. As before, the dots are exact evaluations, and the lines show the interpolation between them.*

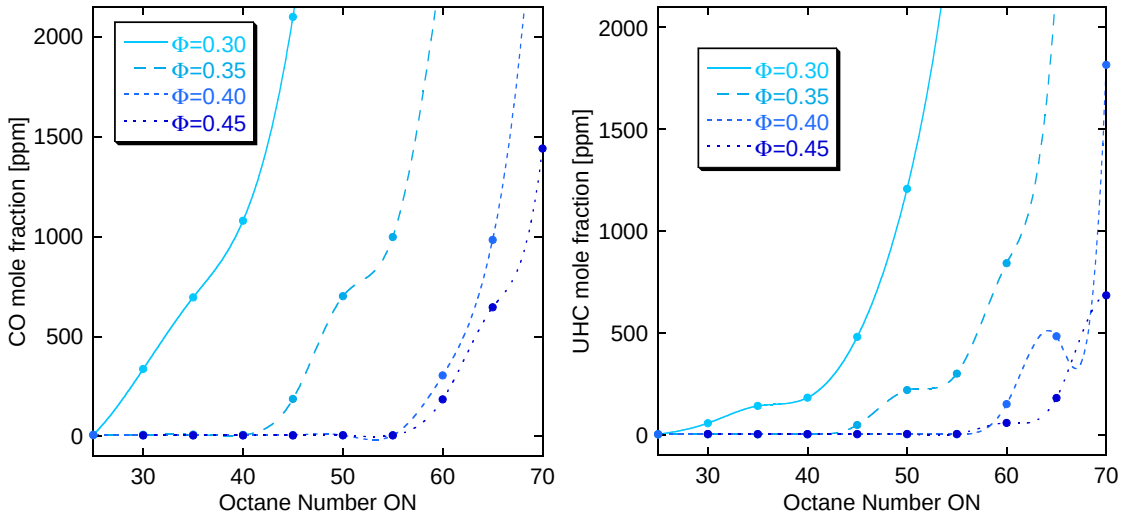

(a) *CO as function of octane number for different equivalence ratios.*

(b) *Unburnt hydrocarbons as function of octane number for different equivalence ratios.*

Figure 6: *Tabulated CO and UHC emissions. As before, the dots are exact evaluations, and the lines show the interpolation between them.*

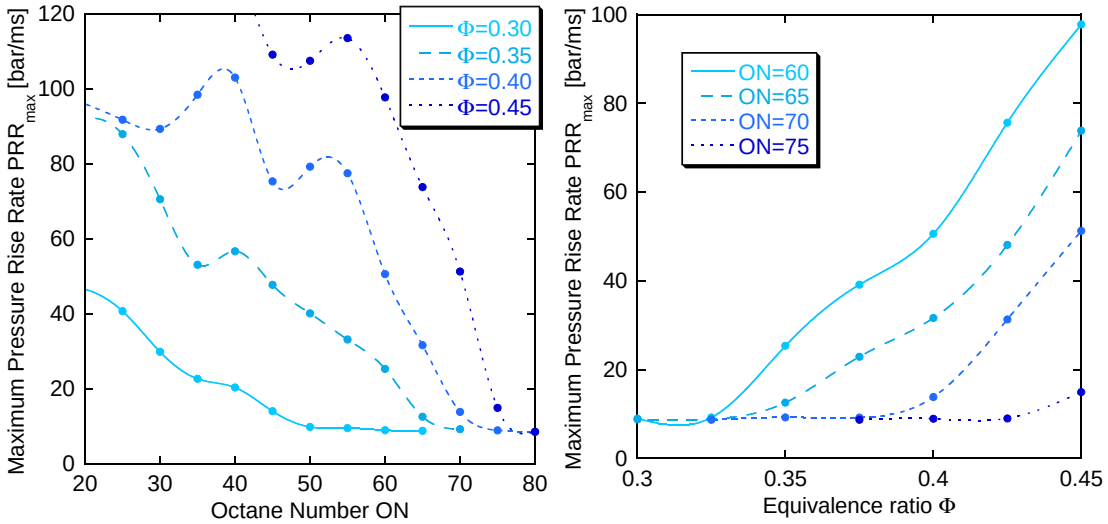

(a) *Maximum pressure rise rate as function of octane number for different equivalence ratios.*

(b) *Maximum pressure rise rate as function of equivalence ratio for different octane numbers.*

Figure 7: *Tabulated maximum pressure rise rates. As before, the dots are exact evaluations, and the lines show the interpolation between them. Fluctuations are more pronounced, as pressure rise rates depend sensitively on higher moments of the temperature distribution.*

following section. After applying this mask, 80 out of 105 vertices (about 76%) remained 'active' (see Fig. 3). Note that a vertex can be omitted only if it and *all* its immediate lattice neighbors are out of range. This introduces a number of subtleties into the interpolation routines, and care needs to be taken when implementing them. We emphasize that for tabulating the SRM, at present, CPU-time is the limiting factor, neither hard disk space nor random access memory.

Naturally, the SRM regarded as a function is a high-dimensional object containing substantial amounts of information. This implies that in order to be visualized, it is necessary to focus on suitable subsets, for which we choose sections, i.e. single quantities as functions of one parameter, while all others are held constant. These projections into two-dimensional space can be easily plotted.

In Fig. 4, several examples of such sections can be seen. Figure 4(a) shows CA50 and Fig. 4(b) the cumulative heat release as functions of octane number for different equivalence ratios. In these and the following figures, the dots represent actual evaluations of the SRM and the lines show the interpolation between them. The vertices which appear to be missing have been removed during the masking procedure. The drop in CA50 for large octane numbers is an immediate consequence of the fact that these cases misfire as can be seen in Fig. 4(b). Even though there is no main combustion event in these cases, the low-temperature heat release is still taking place, hence the CA50-values. In Fig. 4(b), we also observe tabulation errors arising from numerical artefacts of the interpolation, which typically occur in regions of large curvature, for example for $\Phi = 0.4$ and ON between 60 and 70. Inserting additional vertices would eliminate these oscillations. It should be emphasized, though, that in most cases the model errors far outweigh the tabulation errors. The most significant source of errors are the shortcomings in the chemical kinetic scheme, or physical/chemical effects not modelled at all, e.g. turbulence-chemistry interaction, deactivation of radicals at the cylinder wall, influence of lubrication oils, etc.

Figure 5(a) shows NO_x emissions as function of octane number for different equivalence ratios. Figure 5(b) shows the orthogonal section, i.e. the dependence of NO_x on equivalence ratio for different octane numbers. The increase for smaller octane numbers and higher equivalence ratios is not an intrinsic fuel property, but is primarily due to the rise in peak temperatures, which in turn is a consequence of the advance in ignition timing.

Figure 6 shows emissions of CO (Fig. 6(a)) and unburnt hydrocarbons (UHC, Fig. 6(b)) as function of octane number for different equivalence ratios. Again, the rise with increasing octane number is not an intrinsic fuel property, but simply due to the fact that ignition timing retards. CO and UHC emissions are examples for quantities exhibiting strong nonlinearities: In the critical regime close to misfire, these emissions increase very rapidly, which represents almost discontinuous behavior, making them more difficult to tabulate since a denser grid would be required as explained above.

Figure 7 shows the maximum pressure rise rate PRR_{max} as function of both octane number (Fig. 7(a)) and equivalence ratio (Fig. 7(b)). Since PRR_{max} depends relatively sensitively on the *variance* of the temperature distribution rather than on the mean, fluctuations play a more significant role. Once again, the observed trend is in part due to the advance in ignition timing with decreasing octane number and increasing equivalence ratio.

We note that all the responses shown within the domain of interest are sufficiently complicated so that fitting *global* algebraic expressions such as polynomials to them, as it is frequently done within the framework of RSM, is not straightforward.

5 HCCI transient control simulations

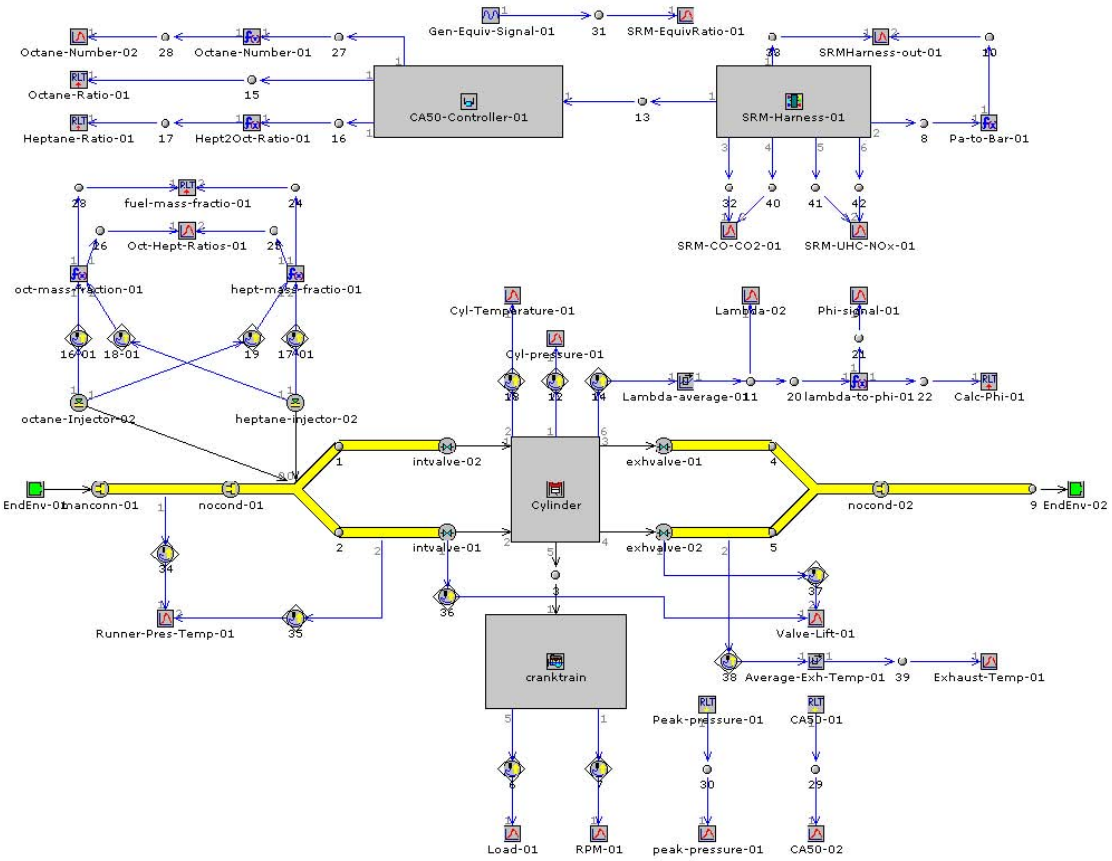

Figure 8: *Representation of the modelled engine including monitors and controllers in a GT-Power® 6.2 engine map.*

The SRM as well as its tabulated version cover only the closed-volume part of the engine cycle. Therefore, in order to conduct full-cycle and multi-cycle simulations, engine breathing needs to be included. This is achieved here through coupling to GT-Power®, a commercial engine simulation package based on one-dimensional CFD. A map representing the Sandia engine as described in section 2 is shown in Fig. 8.

As a generic example for how the tabulation can be useful for engine simulations involving many cycles, we chose a simple control setup. Imitating a varying load situation at constant speed, we impose an arbitrary equivalence ratio profile, and the Octane Number (ON) of the fuel is altered by a controller such that ignition timing is held at a specified target value. We use a PID controller and sensors as provided by GT-Power®, which can also be seen in the engine map in Fig. 8. The values for the Proportional, Integral, and

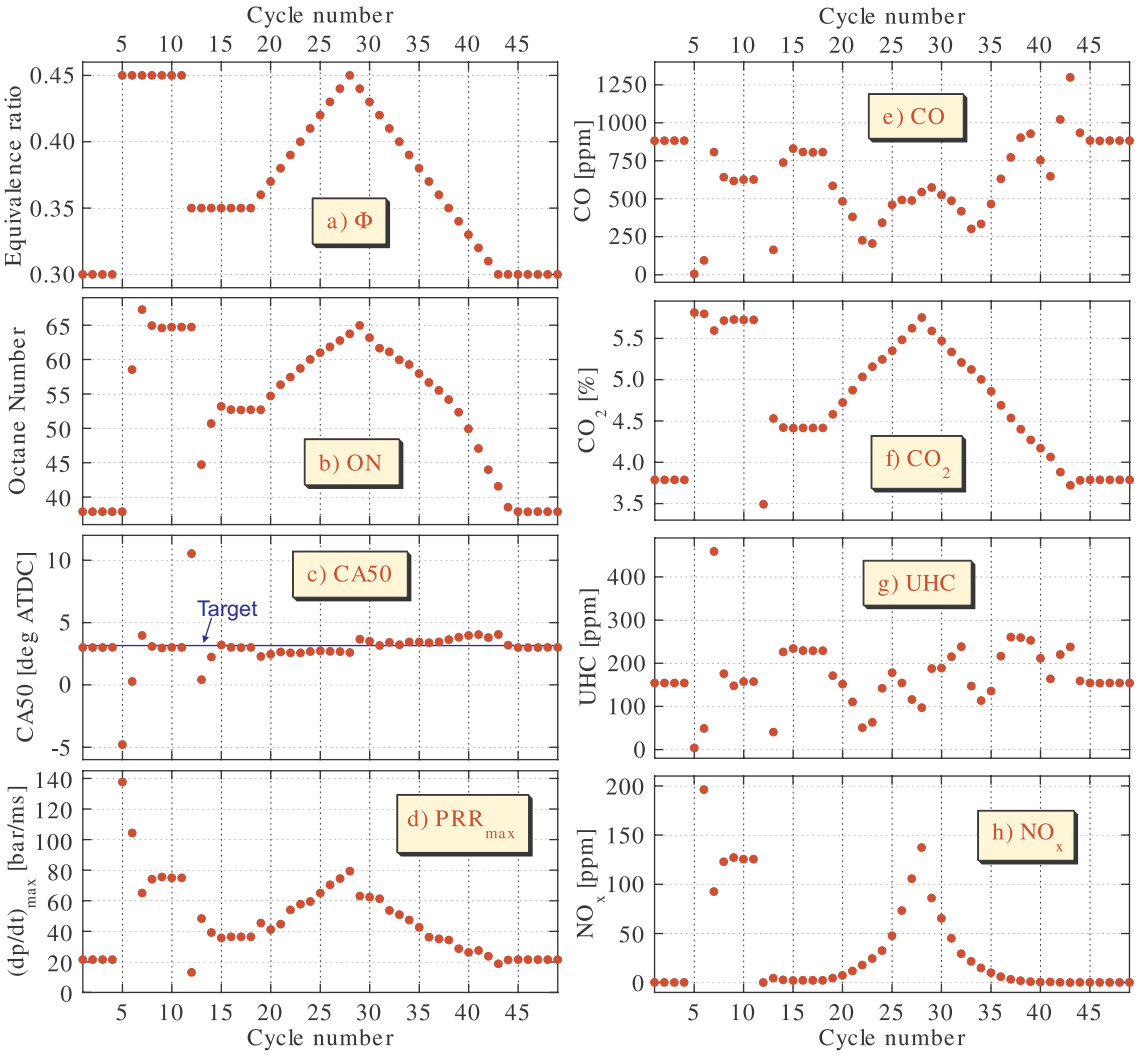

Figure 9: *Cycle-to-cycle evolution of various engine quantities. An arbitrary equivalence ratio profile is imposed (a). A PID-controller adjusts octane number (b) such that CA50 assumes its target value of 3 CAD ATDC (c). The resulting evolution of, for instance, the maximum pressure rise rate (d), CO (e), CO_2 (f), unburnt hydrocarbons (g), and NO_x (h).*

Derivative Gains are 0.01, 0.3, and 0.01 respectively. Example results showing the cycle-to-cycle evolution of several quantities in a typical multi-cycle simulation are given in Fig. 9. The first few cycles prior to reaching steady state are not shown. At this particular setting, the controller manages to re-adjust CA50 to its target value of 3 CAD ATDC in response to the step changes in equivalence ratio within four cycles. The variation in equivalence ratio and CA50 affects a number of other quantities such as the maximum pressure rise rate $\mathrm{PRR}_{\mathrm{max}}$ and emissions as seen in the figure. The NO_x value for the 5^{th} cycle is 378 ppm (not shown). This large value is mostly due to the fact that ignition in this case happens well before TDC (CA50 is about -5 CAD ATDC), leading to a very large peak temperature. Both $\mathrm{PRR}_{\mathrm{max}}$ and NO_x correlate strongly with the equivalence ratio (or ON) but partly also with CA50, as ignition closer to TDC favors rapid pressure rise and high peak temperatures. CO and UHC emissions on the other hand show very little correlation because they are predominantly caused by heat losses, imperfect mixing, and crevice volumes. We note that the experimentally measured values of CO and UHC emissions at steady state for $\Phi = 0.44$ are about 470 ppm and 630 ppm respectively, which compares favorably with the simulated values. For the 12^{th} cycle in the simulation, due to the drop in equivalence ratio CA50 retards to about 10.5 CAD ATDC, implying that combustion is incomplete to a significant extent. As a consequence, CO and UHC emissions reach the very high values of 4750 ppm and 2030 ppm respectively, both of which are not shown in the figure.

We emphasize that the total number of simulations run for creating the table is small compared to the number of cycles run using the table in a typical application, which shows that the tabulation is not merely a redistribution of CPU-time. In practice, during the process of tuning the controller, changing imposed profiles, testing the response, etc., the number of tabulated cycles significantly exceeds the number of direct calculations. In GT-Power®, the total CPU-time for a 50-cycle simulation is of the order of minutes. It is such time-scales which make for example tuning the parameters of a controller feasible. Furthermore, we note that retrieval from the table requires only milliseconds which would be sufficiently fast even for real-time evaluations.

6 Conclusions

We have implemented a storage/retrieval technique for a Stochastic Reactor Model (SRM) for HCCI engines which enables fast evaluation in transient multi-cycle simulations. The SRM uses detailed chemical kinetics, accounts for turbulent mixing and convective heat transfer, and predicts ignition timing, cumulative heat release, maximum pressure rise rates, and emissions of CO, CO_2, unburnt hydrocarbons, and NO_x.

As an example, we have shown that, when coupled to a commercial 1D CFD engine modelling package, the tabulation scheme enables convenient simulation of transient control, using a simple table on a two-dimensional parameter space spanned by equivalence ratio and octane number.

We believe that the developed computational tool will be useful in identifying parameters for achieving stable operation and control of HCCI engines over a wide range of conditions. Furthermore, our tabulation tool enables multi-cycle and multi-cylinder

simulations, and thereby allows to study conveniently phenomena like cycle-to-cycle and cylinder-to-cylinder variations. In particular, simulations of transient operation and control, design of experiments, and optimization of engine operating parameters become feasible.

Acknowledgments

This work has been partially funded by Aramco Overseas Company B.V., contract number 6600014846, under the title "Modeling of Petrol Fuel Combustion in HCCI", and by the EPSRC, grant number EP/D068703/1. The authors express their gratitude to Magnus G. Sjöberg for kindly providing engine data.

References

[1] M. Sjöberg and J. E. Dec. Smoothing HCCI heat-release rates using partial fuel stratification with two-stage ignition fuels. SAE Paper No. 2006-01-0629, 2006.

[2] Y. Zhang, R. Rawat, G. Schmidt, D. C. Haworth, I. Veljkovic, and P. Plassmann. Acceleration of detailed chemistry calculation in multidimensional engine modeling using DOLFA. International Multidimensional Engine Modeling User's Group Meeting 2005, Detroit, MI, 2005.

[3] R. X. Yu, X. S. Bai, A. Vressner, A. Hultqvist, B. Johansson, J. Olofsson, H. Seyfried, J. Sjöholm, M. Richter, and M. Aldén. Effect of turbulence on HCCI combustion. SAE Paper No. 2007-01-0183, 2007.

[4] S. Mosbach, M. Kraft, A. Bhave, F. Mauss, J. H. Mack, and R. W. Dibble. Simulating a homogenous charge compression ignition engine fuelled with a DEE/EtOH blend. SAE Paper No. 2006-01-1362, 2006.

[5] S. Mosbach, H. Su, M. Kraft, A. Bhave, F. Mauss, Z. Wang, and J.-X. Wang. Dual injection HCCI engine simulation using a stochastic reactor model. *Int. J. Engine Research*, 8(1):41–50, 2007. doi:10.1243/14680874JER01806.

[6] S. B. Pope. Computationally efficient implementation of combustion chemistry using in situ adaptive tabulation. *Combust. Theory Modelling*, 1(1):41–63, 1997. doi:10.1080/713665229.

[7] E. H. Kung, S. Priyadarshi, B. C. Nese, and D. C. Haworth. A CFD investigation of emissions formation in HCCI engines, including detailed NOx chemistry. International Multidimensional Engine Modeling User's Group Meeting 2006, Detroit, MI, 2006.

[8] T. Turányi. Application of repro-modeling for the reduction of combustion mechanisms. *Proc. Combust. Inst.*, 25:949–955, 1994.

[9] T. Turányi. Parameterization of reaction mechanisms using orthonormal polynomials. *Comput. Chem.*, 18(1):45–54, 1994. doi:10.1016/0097-8485(94)80022-7.

[10] R. H. Myers and D. C. Montgomery. *Response Surface Methodology: Process and Product Optimization Using Designed Experiments*. Wiley Series in Probability and Statistics. Wiley-Interscience, 2nd edition, 2002.

[11] M. Frenklach, H. Wang, and M. J. Rabinowitz. Optimization and analysis of large chemical kinetic mechanisms using the solution mapping method — combustion of methane. *Prog. Energy Combust. Sci.*, 18:47–73, 1992. doi:10.1016/0360-1285(92)90032-V.

[12] S. R. Tonse, N. W. Moriarty, N. J. Brown, and M. Frenklach. PRISM: Piecewise reusable implementation of solution mapping. an economical strategy for chemical kinetics. *Israel Journal of Chemistry*, 39:97–106, 1999.

[13] S. R. Tonse, N. W. Moriarty, M. Frenklach, and N. J. Brown. Computational economy improvements in PRISM. *Int. J. Chem. Kin.*, 35(9):438–452, 2003. doi:10.1002/kin.10140.

[14] P. Stewart, P. J. Fleming, and S. A. MacKenzie. Real-time simulation and control systems design by the response surface methodology and designed experiments. *International Journal of Systems Science*, 34(1):837–850, 2003. doi:10.1080/00207720310001640287.

[15] K. Chang, G. A. Lavoie, A. Babajimopoulos, Z. S. Filipi, and D. N. Assanis. Control of a multi-cylinder HCCI engine during transient operation by modulating residual gas fraction to compensate for wall temperature effects. SAE Paper No. 2007-01-0204, 2007.

[16] S. B. Pope. PDF methods for turbulent reactive flows. *Prog. Energy Combust. Sci.*, 11:119–192, 1985. doi:10.1016/0360-1285(85)90002-4.

[17] H. Su, A. Vikhansky, S. Mosbach, M. Kraft, A. Bhave, K.-O. Kim, T. Kobayashi, and F. Mauss. A computational study of an HCCI engine with direct injection during gas exchange. *Combust. Flame*, 147(1-2):118–132, 2006. doi:10.1016/j.combustflame.2006.07.005.

[18] H. Su, S. Mosbach, M. Kraft, A. Bhave, S. Kook, and C. Bae. Two-stage fuel direct injection in a Diesel fuelled HCCI engine. SAE Paper No. 2007-01-1880, 2007.

[19] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C++ The Art of Scientific Computing*. Cambridge University Press, Cambridge, 2nd edition, 2002.

List of abbreviations

ATDC	After Top Dead Centre
BTDC	Before Top Dead Centre
CA50	Crank Angle at 50% heat release
CAD	Crank Angle Degree
CFD	Computational Fluid Dynamics
CHR	Cumulative Heat Release
CR	Compression Ratio
DOLFA	Database for On-Line Function Approximation
EGR	Exhaust Gas Recirculation
EVC	Exhaust Valve Closure
EVO	Exhaust Valve Opening
GT	Gamma Technologies®
HCCI	Homogeneous Charge Compression Ignition
ISAT	*In-Situ* Adaptive Tabulation
IVC	Inlet Valve Closure
IVO	Inlet Valve Opening
LES	Large Eddy Simulation
ON	Octane Number
PDF	Probability Density Function
PID	Proportional Integral Differential
PRF	Primary Reference Fuel
PRISM	Piecewise Reusable Implementation of Solution Mapping
PRR_{max}	Maximum Pressure Rise Rate
RPM	Revolutions Per Minute
RSM	Response Surface Methodology
SRM	Stochastic Reactor Model
TDC	Top Dead Centre
UHC	Unburnt Hydrocarbons
VCR	Variable Compression Ratio