Oil & Natural Gas Technology

DOE Award No.: DE-FC26-03NT15405

Final Report

Experimental Investigation and High Resolution Simulation of In-Situ Combustion Processes

Submitted by:
Prof. Margot Gerritsen and Prof. Tony Kovscek
Department of Energy Resources Engineering
367 Panama Street
Stanford, CA 94305-2220

Prepared for:
United States Department of Energy
National Energy Technology Laboratory

Office of Fossil Energy

Disclaimer

Acknowledgements

Financial support for this was provided by the U. S. Department of Energy under Award No. DE-FC26-03NT15405.

Additional contributors to this work include Mr. O.G. Awoleke, Dr. L. M. Castanier, Mr M. Cinar, Dr. Morten Kristensen, Mr. J. Cristofari, Mr. B. He, Mr. Q. Chen, Mr. J. Kozdon, Dr. J.V. Lambers, Dr. B. T. Mallison, Dr. J. Nilsson, Mr. R. Younis.

Executive Summary

This final technical report describes work performed for the project "Experimental Investigation and High Resolution Numerical Simulator of In-Situ Combustion Processes," DE-FC26-03NT15405. In summary, this work improved our understanding of in-situ combustion (ISC) process physics and oil recovery. This understanding was translated into improved conceptual models and a suite of software algorithms that extended predictive capabilities. We pursued experimental, theoretical, and numerical tasks during the performance period. The specific project objectives were (i) identification, experimentally, of chemical additives/injectants that improve combustion performance and delineation of the physics of improved performance, (ii) establishment of a benchmark one-dimensional, experimental data set for verification of in-situ combustion dynamics computed by simulators, (iii) develop improved numerical methods that can be used to describe in-situ combustion more accurately, and (iv) to lay the underpinnings of a highly efficient, 3D, in-situ combustion simulator using adaptive mesh refinement techniques and parallelization. We believe that project goals were met and exceeded as discussed next.

Chemical additives to improve ISC process efficiency and reliability were pursued using ramped temperature oxidation experiments complemented by scanning electron microscopy to examine the solid matrix prior to and following combustion. A suite of crude oils, water-soluble metallic salt additives, clay, and silica solids were chosen and various combinations screened for reactivity. These experiments elucidated the role that metallic additives play in the cracking, low temperature oxidation and pyrolysis reactions leading to fuel formation during in-situ combustion. Key findings reaffirmed the important role of surface area provided by clay, in contrast to silica, for heterogeneous oxidation reactions. Results indicate that water soluble metallic additives exchange with cations in clay. Thus, metallic additives can be deployed relatively easily in field operations. The effectiveness of various additives was found to rank with hydrated cation size according to the series: $Sn^{2+} > Fe^{3+} > Al^{3+} > Zn^{2+} > Mg^{2+} > Cu^{2+} > Cd^{2+}$. Combustion tube runs were also conducted to illustrate the effectiveness of metallic additives in aiding the propagation of fire front.

A new recovery concept was proposed to achieve a greater degree of in-situ upgrading during ISC operations. The concept was tested via ramped temperature oxidation runs and combustion tube experiments. In the proposed recovery process, a solvent is injected to dilute the oil as well as to induce asphaltene precipitation. The removal of asphaltenes increases gravity of the oil. The solvent reduces oil viscosity in-situ and extracts the lighter crude-oil fractions. Combustion cleans the near-well region and stimulates thermally the oil production. Cycles of solvent injection and combustion are repeated. Both solvent injection and in-situ combustion are technically effective, but this combination appears to never been tried before. Light hydrocarbons, such as pentane, proved to be effective at inducing

asphaltene precipitation and diluting the oil. Further, pentane was easily separated from the produced oil. The crude oils samples employed for the study were Hamaca (9-11 °API) from the Orinoco belt of Venezuela and West Sak (20-22 °API) from the North Slope of Alaska. The precipitates from Hamaca resulting from solvent injection were sufficient for combustion whereas the process was challenged for West Sak.

Analysis of ramped temperature oxidation kinetics tests with the conventional approach illustrated many short comings. In the conventional approach, the many oxidation reactions are lumped into two or three main reactions and the apparent activation energy of the lumped reaction is estimated from the oxygen consumption during the test. The standard procedure is subject to interpretation, errors in the assumed lumped model, as well as errors from lack of redundancy. Conventional methods of kinetic analysis for crude oil assume, a priori, a two-step or three-step oxidation reaction. Any in-consistency between the assumed model and the real system is reflected as error in model parameters. Isoconversional techniques, as applied here, provide model-free methods for estimating activation energy and naturally deconvolve multi-step reactions. The iso-conversional approach obtains the reaction kinetics for a given rock/oil sample at identical reaction extent from multiple experiments with different heating rates. The technique has never been applied to ISC kinetics, to our knowledge.

Combustion in three-dimensional, heterogeneous systems was examined from an experimental perspective. An extensive literature exists on the nature of the reactions that occur when air comes in contact with crude oil during in-situ combustion (ISC); however the effect of reservoir heterogeneity on the ISC process has not been experimentally addressed in prior studies. We probed the effect of pressure, temperature, injection flow-rate, and matrix properties on the combustion of a 19.7 °API oil in a tight-formation sand. The experimental perspective includes both oil oxidation kinetics and combustion-front dynamics in a 1-m long combustion tube. Most importantly, geological effects on the ISC process are investigated by incorporating various degrees of porous medium heterogeneity in the combustion tube runs. From the experimental data, pre and post burn CT scans of the combustion tube, and postmortem analyses, we deduce the effect of both small and large-scale heterogeneities on the spatial and temporal propagation of the combustion front.

The foundations of an advanced, in-situ combustion numerical simulator were established. We developed an operator splitting method for ISC processes based on fractional stepping (Younis2006, Kristensen2008b, Chapters 8 and 11). We designed a specialized, high order integration method for efficient integration of kinetics. This method contains a phase change detection algorithm that significantly improves the robustness of the time-stepping algorithm. We implemented this algorithm in conjunction with the operator splitting method and tested it with various chemical reaction models. (Kristensen2007, Chapter 7).We implemented a 1D ISC simulation with spatial grid adaptivity, that has been shown to improve computational efficiency. It is based on a

patched gridding approach. In 3D we instead elected to use cell-based refinements as discussed below. (Younis2006, Chapter 10). We formulated a new asynchronous time stepping algorithm so that temporal adaptivity can be applied as well as spatial adaptivity. This time-stepping method can be used in conjunction with the operator splitting method. It will lead to considerable savings for ISC processes as small time step sizes can be restricted to critical areas such as combustion fronts. (Mallison2008, Chapter 11). We used our 1D code to rigorously assess parameter and model sensitivity. We used 1D combustion tube experiments as well as simulation results reported in the literature for comparison (Kristensen2007, Kristensen2008, Chapter 9).

Adaptive mesh refinement (AMR) in 3D was explored as a possible technique to speed up ISC simulations. In particular, we developed a cell-based, 3D adaptive spatial gridding strategy, which we label CCAR (Cartesian, Cell-based Anisotropically Refined) grids (Nilsson2004, Nilsson 2005, Chapter 10).The hanging nodes in the adaptive grid strategy are effectively dealt with using a new flow solver. The resulting matrix-vector system is solved using an efficient parallel algorithm based on algebraic multi-grid (Nilsson2004, Chapter 10).We developed a new multi-D transport solver for use on adaptive grids that reduces sensitivity to grid orientation for adverse mobility ratio flows, such as those encountered in gas injection and ISC processes. (Kozdon2008). We designed two new transmissibility upscaling methods for cell-based adaptive grids, for use with two-point and multi-point flux approximations, respectively. These upscaling methods lead to highly accurate flow solutions even in case of full tensor anisotropy in the underlying permeability field. (Gerritsen2008, Lambers2008, Lambers2009, Chapter 10).

The simulator functions were validated for field-scale calculations with realistic geology and the resolution of combustion fronts were examined, as well as the computational efficiency of AMR under realistic conditions. In particular,w e extended all gridding and upscaling methods, mentioned in task 7, to handle complex heterogeneities. (Chapter 10. We developed an effective algorithm to generate adapted grids based on reservoir heterogeneity (permeability contrasts). (Nilsson2004, Nilsson2005, Chapter 10). After the development of our new gridding strategies, flow solvers, transport solvers, upscaling methods, time-stepping methods and specialized kinetics solvers, we now have all the ingredients needed for a full field simulator. All pieces are integrated in our combustion tube simulator (VCT). Multi-dimensional tests have been conducted on the underlying algorithms, and a new field simulator is forthcoming as well the pieces are now in place.

Publications and Technical Transfer Activities

Technical transfer and dissemination of project results was achieved by publication and presentations at events attended by industry, as detailed next.

Awoleke, O.G., 2007. In-Situ Combustion in Heterogeneous Media, SPE 113022 presented at the 2007 SPE Annual Technical Conference and Exhibition, Anaheim, CA, 11-14 November.

Awoleke, O. G., 2007. An Experimental Investigation of In-Situ Combustion in Heterogeneous Media, M.S. Thesis Stanford University. Online at http://ekofisk.stanford.edu/pereports/web /default.htm

Castanier, L. M. and A. R. Kovscek, " Heavy oil upgrading in–situ via solvent injection and combustion: A "new" method" Proceedings of the EAGE 67th Conference & Exhibition — Madrid, Spain, 13 - 16 June 2005.

Chen, Q., 2005. Microscopic Investigation of Fluid Flow through High Porosity Porous Media and In-Situ Combustion with Aqueous Metallic Additives. ," M.S. Thesis Stanford University. Online at http://ekofisk.stanford.edu/pereports/web /default.htm

Cinar, M, L. M. Castanier, and A. R. Kovscek, "Improved Analysis of the Kinetics of Crude-Oil In-Situ Combustion," SPE 113948, Proceedings of the SPE Western Regional Meeting, Bakersfield, CA. 31 Mar - 2 Apr 2008.

Cristofari, J., L. M. Castanier, and A. R. Kovscek, "Laboratory Investigation of the Effect of Solvent Injection on In-Situ Combustion" SPE Journal, to appear 2008.

Cristofari, J, 2006. Laboratory Investigation of the Effect of Solvent Injection on In-Situ Combustion for Viscous Oils, M.S. Thesis Stanford University. online: http://ekofisk.stanford.edu/pereports/web /default.htm

Cristofari, J., L. M. Castanier, and A. R. Kovscek, "Laboratory Investigation of the Effect of Solvent Injection on In-Situ Combustion" SPE 99752, Proceedings of the SPE Symposium on Improved Oil Recovery, Tulsa, OK Apr 22-26, 2006.

Gerritsen, M., A. R. Kovscek, L. Castanier, J. Nilsson, R. Younis, and B. He, "Experimental Investigation and High Resolution Simulator of In-Situ Combustion: 1. Simulator Design and Improved Combustion With Metallic Additives," SPE 86962, Proceedings of the SPE International Thermal Operations and Heavy Oil Symposium and Western Regional Meeting, Bakersfield, CA, Mar 16-18, 2004.

M. G. Gerritsen and J. Lambers, Integration of Local-Global Upscaling and Grid Adaptivity for Simulation of Subsurface Flow in Heterogeneous Formations, Computational Geosciences 12, 193-218, 2008

He, B., 2004. The Effect of Metallic Salt Additives on In-Situ Combustion Performance," M.S. Report Stanford University. Online at http://ekofisk.stanford.edu/pereports/web /default.htm

He, B., Q. Chen, L. M. Castanier, and A. R. Kovscek, "Improved In-Situ Combustion Performance with Metallic Salt Additives," SPE 93901, Proceedings of the SPE Western Regional Meeting, Irvine, CA, Mar 30-Apr 1, 2005.

Kovscek, A. R., “Heavy-Oil and Thermal Recovery Research Activities at Stanford University." BP Exploration, Anchorage, AK, March 21, 2005.

Kovscek. A. R., “Upgrading and Modification of Crude Oil via In-Situ Combustion,” Shell International Science Symposium, The Hague, November 2, 2006.

Kozdon, J., M. G. Gerritsen and M. Christie, The Grid Orientation Revisited: Near-well, Early-time Effects and Solution Coupling Methods, Transport in Porous Media, 73(3), 255-277, 2008(a)

Kozdon, J., B. Mallison and M. Gerritsen, Robust Multi-D Transport Schemes with Reduced Grid Orientation Effects, Transport in Porous Media, 2008(b)

Kristensen, M.R., M.G. Gerritsen, P.G. Thomsen, M.L. Michelsen, and E.H. Stenby, Impact of Phase Behavior Modeling on In-Situ Combustion Process Performance,SPE 113947, SPE/DOE Symposium on Improved Oil Recovery, 20-23 April 2008, Tulsa, Oklahoma, USA, 2008(a)

Kristensen, M.R ,M. G. Gerritsen, P. G. Thomsen, M. L. Michelsen, E. H. Stenby, An Equation-of-State Compositional In-Situ Combustion Model: A Study of Phase Behavior Sensitivity, Transport in Porous Media, 2008(b)

Kristensen, M.R, M. G. Gerritsen, P. G. Thomsen, M. L. Michelsen, E. H. Stenby, Efficient Integration of Stiff Kinetics with Phase Change Detection for Reactive Reservoir Processes, Transport in Porous Media, Volume 69, Number 3, 383-409, 2007(a)

Kristensen, M.R, M. G. Gerritsen,, P. G. Thomsen, M. L. Michelsen and E. H. Stenby, Coupling Chemical Kinetics and Flashes in Reactive, Thermal and Compositional Reservoir Simulation, SPE 106218, SPE Reservoir Simulation Symposium, 26-28 February 2007, Houston, Texas, U.S.A., 2007(b)

Lambers, J.V. and M. G. Gerritsen, "Spatially-varying Compact Multi-point Flux Approximations for 3-D Adapted Grids with Guaranteed Monotonicity", Proceedings of the 11th European Conference on the Mathematics of Oil Recovery, Bergen, Norway, 2008(a)

Lambers, J.V., M. G. Gerritsen and B. Mallison, Accurate Local Upscaling with Compact Multi-Point Transmissibility Calculations, Computational Geosciences, Special Issue on Multiscale Methods for Flow and Transport in Heterogeneous Porous Media, 2008(b)

Lambers, J.V. and M. G. Gerritsen, An Integration of Multi-Level Local-Global Upscaling and Adaptivity, SPE 97250, SPE Annual Technical Conference and Exhibition, 9-12 October 2005, Dallas, Texas, 2005

Mallison, B.T., M. G. Gerritsen, G. Kreiss, Asynchronous Time Integration of Flux-conservative Transport,Proceedings of the 11th European Conference on the Mathematics of Oil Recovery, Bergen, Norway, 2008

Nilsson, J., M. G. Gerritsen, R. Younis, A novel Adaptive Anisotropic Grid Framework for Efficient Reservoir Simulation, SPE 93243, SPE Reservoir Simulation Symposium, 31 January-2 Feburary 2005, The Woodlands, Texas

Nilsson, J., M. Gerritsen and R. Younis, Parallel Anisotropic Cartesian Grid Adaptation for In-Situ Combustion Simulations, 9th European Conference on the Mathematics of Oil Recovery, Cannes, 2004

Younis. R. and M. G. Gerritsen, Multi-Scale Process Coupling by Adaptive Fractional Stepping: An In-Situ Combustion Model, SPE 93458, SPE/DOE Symposium on Improved Oil Recovery, 22-26 April 2006, Tulsa, Oklahoma, USA (submitted to SPE J)

Table of Contents

1. Project Introduction

The world continues to rely substantially on petroleum fossil fuels as a primary energy source. While the number of new discoveries of petroleum reservoirs decreases, the need to produce the known reservoirs more effectively increases. So far, only the easiest to produce petroleum accumulations have been utilized, and advanced technology is needed to access more difficult to produce reservoirs and to address in a cost-effective manner the volume of oil that remains after primary and secondary operations. By 2007, proved world oil reserves were estimated at 1.3 trillion barrels. A proved reserve is oil that with reasonable certainty can be recovered in future years, assuming existing economic and operating conditions (IEA, 2007). It is further estimated that Canada and Venezuela alone have heavy oil resources, which far exceed the world's conventional reserves, but which cannot yet be economically recovered.

Thermal recovery processes, which rely on a viscosity reduction of the oil through heat injected (steam or hot water) or in-situ generated, are well suited to unlock effectively these heavy oil resources, and conventional residul oil, in an environmentally sound manner. Among the thermal recovery processes, ISC has been a source of interest and effort for several decades. ISC is the process of injecting air (or air enriched with oxygen) into oil reservoirs to oxidize a portion of the crude oil and enhance recovery through the heat and pressure produced. Contrary to other thermal recovery processes the main part of the energy required to displace the oil in ISC is generated inside the reservoir from the heat released by chemical reactions between oxygen and fractions of the crude oil.

ISC also possesses the potential as a low carbon dioxide emissions thermal recovery process. As carbon oxides (carbon dioxide and carbon monoxide) are generated during combustion and produced, the gas phase is easily separated from produced oil and water, compressed, and mixed with enriched air to create an injection gas with appropriate amounts of oxygen and inert gases. In short, carbon dioxide is recycled through the reservoir.

Thus, ISC is technically and economically an attractive process, particularly because the portion of the crude burned is likely to be the heaviest and least valuable. Whereas it is generally classified as a technique that is applicable for heavy oils because of the dramatic reduction in oil viscosity with temperature, ISC also promotes production through thermal expansion and gas drive caused by combustion gases. The process has proven economical in recovering heavy oil from shallow reservoirs and lighter oil from deep reservoirs, where other processes, such as steam injection, are unattractive (Sarathi,1999). When applied to deep light oil reservoirs, the process is often referred to as *high pressure air injection* (HPAI), whereas the term *in-situ combustion* traditionally has been used for heavy oil reservoirs. Throughout this report we refer to air injection into both heavy and light oil reservoirs as in-situ combustion, but important differences between the two settings are highlighted.

1.1. In-Situ Combustion Processes

The ISC process is generally described based on the respective directions of front propagation and air flow. In *forward combustion* oil is ignited at the injection well and the combustion front is propagated towards the production well by continuous injection of air. In *reverse combustion*, on the other hand, the oil in the production well is ignited and the combustion front propagates against the air flow. Compared to forward combustion the temperature in reverse combustion remains high between the combustion zone and the production well, which has the effect of promoting the flow of oil. The disadvantage of the method is that a residue of solid fuel may block the pores and slow down displacement. Moreover, spontaneous ignition of oil may occur near the injection well consuming most of the oxygen and causing a breakdown in the supply of oxygen to the combustion front (Prats,1986). For these reasons, forward combustion is solely practiced in the field.

As air is injected and the combustion front progresses into the reservoir, several zones are between the injector and the producer. A common misconception of the combustion front propagation is that of an underground fire. Instead, the front is more of a glow passing slowly through the reservoir, similar to the hot zone of a burning cigarette. Figure 1.1 gives a schematic representation of characteristic temperature and saturation zones in ISC. Starting from the injection well the burned zone is the volume already swept by the combustion zone. The burned zone contains the injected air and possibly a residue of burned fuel. The combustion zone has the greatest temperature and this is where most of the energy is generated. Injected oxygen reacts with residual hydrocarbons generating carbon oxides and water. Hydrocarbons contacted by the leading edge of the high temperature zone undergo thermal cracking and vaporization. Mobilized light components are transported downstream where they mix with the original crude. The heavy residue, which is normally referred to as coke, is deposited on the core matrix and is the main fuel source for the combustion process. Downstream of the vaporization zone is the steam plateau which is formed from water of combustion and vaporization of formation water. Further downstream the steam condenses into a hot water bank when the temperature drops below the steam saturation temperature. The leading edge of the hot water bank is the primary area of oil mobilization where the oil is banked by the hot water. A more detailed description may be found in (Prats,1986) and (Brigham,2004).

A recently developed operation technique involving horizontal wells is Toe-to-Heel Air Injection (THAI) (Greaves,2005), in which air is injected in a vertical well and oil is produced from a long horizontal well. As the combustion front is propagated, the heated oil is drained into the producer, a technique similar to Steam Assisted Gravity Drainage (SAGD) (Butler,1981). While the THAI process has some features that are attractive, it relies, essentially, on plugging of the production section of the well with coke to avoid combustion in the well itself. Hence, application of THAI must be approached with a degree of caution.

When only air is injected into the reservoir the process is referred to as *dry* combustion. A disadvantage of this process is that a large part of the energy generated remains behind the combustion front as heat absorbed in the reservoir rock. Because of the low heat capacity of air, the heat in the burned zone is not transported efficiently by the injected air. To overcome this inefficiency water can be injected either simultaneously or alternately which is referred to as *wet* combustion. The water-air ratio determines whether water is transported through the combustion zone in a vapor or liquid state. At sufficiently high water-air ratio liquid water will enter or even pass all the way through the combustion zone leading to quenched or partially quenched combustion.

The actual mechanisms responsible for oil displacement in ISC vary with the type of oil. For heavy oils the increase in oil mobility with elevated temperatures is the primary mechanism assisted by gas flood effects and hot and cold water drive. In the case of lighter oils, the flue gas mixture resulting from combustion provides the primary mobilizing force for the oil downstream of the combustion zone (Moore,2002b). The gas-oil mixture may be immiscible, or partly or completely miscible.

Many economically successful ISC projects exist (Moore1995,Moore2002b). Application of the process is, however, not widespread for a variety of reasons including the relatively large expense required to evaluate prospects. It is often regarded as a high-risk process due to the current lack of reliable numerical and experimental tools for accurate and efficient prediction of field performance.

1.1.1. Characteristic Behavior of ISC Processes

In this section we discuss qualitatively the characteristic behavior observed when burning different kinds of oil under varying conditions in laboratory scale tube tests. Only little is known about ISC processes at the field scale. However, laboratory combustion tube tests provide a good testbed for assessing burning characteristics and studying effects of varying process conditions (e.g. pressure, initial oil and water saturations).

Figure 1.2 shows characteristic temperature profiles observed in ISC laboratory tube tests. Profiles are shown for dry and wet combustion of a typical heavy oil and for dry combustion of a light oil. For dry combustion with little or no initial water saturation the temperature profile shows a single front. No significant steam plateau is formed. The temperature remains fairly high behind the combustion front due to the limited ability of the injected gas to carry the heat forward. The major difference between heavy and light oil combustion is the temperature regime in which the process operates. Heavy oils typically burn at peak temperatures around $600^{\circ}\mathrm{C}$, whereas light oils burn at much lower temperatures around $300^{\circ}\mathrm{C}$. When water is injected in the wet combustion process, the heat absorbed in the rock in the burned region is effectively swept forward giving a sharper temperature peak. Moreover, vaporization of the water forms a steam bank ahead of the combustion front.

A phenomenon, that has been observed at laboratory scale and theorized to have a stabilizing effect on combustion fronts at the field scale, is the process of blocking and unblocking. Especially in highly oil saturated reservoirs this will be important. Due to the efficient oil mobilization in ISC, the oil forms a bank ahead of the combustion front and the gas saturation may drop below the critical saturation leading to vanishing gas permeability. With no gas permeability the system is blocked. For one-dimensional combustion tube tests this can lead to a break-down of the process. In a heterogeneous field, however, blocking of certain paths in the reservoir may have a stabilizing effect on the combustion front. The mobilized oil will go to high permeability paths. Blocking of these paths will redirect flow to less permeable regions, thereby preventing fingering phenomena. The benefits of the blocking process at the field scale are still rather hypothetical, and no simulation results have been published capturing this effect. Accurate prediction of blocking and unblocking is likely to be very sensitive to the quality of the relative permeability data. Three phase relative permeability data is, in general, difficult to obtain, and very little is known about the effects of temperature and composition changes on relative permeabilities for processes like ISC.

1.1.2. Chemical reactions

A typical crude oil is a complex mixture of several hundreds of different chemical species. Hence, the chemical reactions that happen in the presence of oxygen are equally complex and numerous. Deriving detailed reaction pathways for crude oil oxidation in porous media is far beyond current capabilities. Determining reactions and reaction rates in ISC is approached in an empirical way.

Today it is generally accepted that three classes of reactions dominate in ISC: (1) low temperature oxidation reactions (LTO), (2) medium temperature reactions and (3) high temperature oxidation reactions (HTO) (Brigham,2004). The temperature ranges associated with each group of reactions are roughly 150 to $300^{\circ}C$ for LTO, 300 to $450^{\circ}C$ for medium temperature reactions and above $450^{\circ}C$ for HTO.

The nature of the reactions taking place in each regime depends on the type of oil. For heavy oils the LTO reactions are oxygen addition reactions producing partially oxygenated compounds (e.g. alcohols, ketones, aldehydes) and only few carbon oxides. Lighter oils undergo full H/C bond breaking combustion reactions in the LTO regime but may undergo oxygen addition reactions at lower temperatures. Oxygen addition increases oil viscosity, and thus operating in this regime is inefficient. In a forward combustion process LTO reactions may take place when oxygen is available downstream of the combustion front due to either channeling around the front or insufficient consumption in the front (Fassihi,1984a). The effects of LTO on ISC performance have been studied by several researchers [e.g][Dabbous1974,Adegbesan1987,Freitag2005). Although generally recognized as being important to prediction of ISC performance, no reliable LTO reaction models exist today. Freitag2005 observed a shift in reaction order for oxygen partial pressure with temperature and concluded that at least two different reaction mechanisms govern LTO. The negative effects of LTO are mainly an increase in oil

viscosity and enhanced trapping of oil due to gas phase shrinkage caused by the oxygen removed from the gas phase in oxygen addition reactions (Moore2002a).

At intermediate temperatures after the LTO reactions a series of cracking or pyrolysis reactions take place. These reactions are the primary source of coke formation for the HTO reactions. Coke formation dictates the quantity of fuel available for combustion. Excessive fuel deposition retards the rate of advance of the combustion front whereas insufficient fuel deposition may not provide enough heat for self-sustained combustion. Finally, in the HTO reactions the coke is burned, generating carbon oxides and water.

Both LTO and HTO reactions are believed to be heterogeneous gas/liquid or gas/solid reactions, in which oxygen from the gas stream diffuses to the surface of the oil/coke, adsorbs on the surface, reacts with a hydrocarbon component, and finally combustion products desorb and diffuse back into the gas stream. It is widely accepted that the overall reaction is kinetically controlled.

1.1.3. Physical Complexity of ISC Processes

ISC is a physically complex enhanced oil recovery process. Driven by complex chemical reactions, the oil mobility is increased with the elevated temperatures, and the fluids are displaced by a combination of steam, water, and gas drive. The spatial as well as temporal scales in ISC vary over many orders of magnitude. The bulk of the chemical reactions take place in the narrow reaction zone that may be less than a meter in thickness compared to reservoir scales of hundreds or thousands of meters. Moreover, combustion reactions often occur in fractions of a second, whereas the temporal scales associated with convective transport may be running to days or years. Accurate prediction of field performance in such a multiscale process is an immense challenge requiring a hierarchical approach, in which both spatial and temporal resolution is adapted in order to capture the crucial input from all levels of activity.

The overall performance of an ISC process is governed in a complex way by reservoir heterogeneity, well configurations, injection rates and composition, initial oil saturation and distribution, and both thermodynamic and chemical properties of the rock and fluids. Reliable prediction of field performance requires a fully integrated approach in which the important contributions from all levels are taken into account. ISC is indeed a multiphysics process bringing together multiphase porous media flow, chemical kinetics, and phase equilibria.

The spatial scales affecting ISC span from large geological features such as faults of the size of the reservoir to the very small scale at which chemical reactions happen in the combustion zone. Faults, fractures, and the placement of wells determine global flow patterns, but local displacement efficiency is governed by small scale heterogeneity in porosity and permeability of the reservoir, and by the chemical, kinetic, and thermodynamic behavior of the fluids. Permeability fields are often obtained from high-resolution geocellular models having gridblock sizes on the order of a meter. Reservoir simulations, however, are carried out using gridblocks that are 1--2 orders of magnitude larger due to computational constraints. Upscaling of the permeability or transmissibility field, in which local

flow behavior is taken into account, is routinely done. As mentioned above, the bulk of the chemical reactions happen in a narrow combustion zone being less than a meter in thickness compared to standard gridblock sizes of, say, 50 meters. Thus, the spatial scale for chemical reactions is smaller than the smallest scale normally resolved in reservoir simulations. Consequently, the temperature profile on the simulation grid will be too smooth, and important phenomena such as ignition/extinction or quenching may not be predicted correctly.

Relating to temporal scale, a number of different processes may be identified in ISC, each having its own characteristic scale. Most of the existing ISC models include convective mass transport, heat transport by convection and conduction, kinetically controlled chemical reactions, and fluid phases in thermodynamic equilibrium (Crookston1979,Grabowski1979,Coats1980). The phase equilibrium assumption implicitly states that the timescales for the interphase mass transfer processes occurring when phases come to an equilibrium state, are much faster than all other timescales. Of the remaining processes, the chemical reactions are likely to occur on timescales that are again much faster than the scales for mass and heat transport.

Although being multiscale in nature, the question remains whether all the processes in ISC represent essential physics that need to be resolved in a simulation. The goal of ISC simulation is to provide reliable predictions of performance, typically in terms of production, for a given ISC project. The production certainly depends on large-scale features such as well placement, but the small scale behavior, spatial as well as temporal, in the combustion zone may be equally important. ISC processes are driven by chemical reactions, and chemical kinetics depends strongly on temperature. Thus failing to capture temperature peaks and, in general, smoothing out temperature profiles on too coarse a grid leads to inaccurate prediction of reaction, that in turn affects the amount of heat released and combustion gases evolved, ultimately resulting in wrong predictions of oil displacement. Ahead of the combustion front (see Figure 1.1) the oil is mobilized by a combination of steam, water, and gas drive. Lighter oil components vaporize easily and are transported downstream. The compositional behavior in this region will determine the amount and composition of the oil left behind as fuel for the combustion. Therefore, accurate prediction of phase behavior as well as flow is likely to impact overall performance. Errors at this small scale feed into overall production calculations, thereby rendering the results unreliable. Hence, the important processes in ISC are indeed multiscale with strong nonlinear interactions between different scales and efficient computational methods must be developed that handle this multiscale nature.

1.1.4. A Review of ISC Models

The majority of mathematical models for ISC emerged in the literature around 1980. Among the first researchers to develop a realistic model for ISC was Farouq1977. He presented a two-dimensional model with three phases (oil, water, and gas) and four fixed components (water, oil, oxygen, and inert gas). The oil component was

assumed non-volatile, and the combustion process was modeled simply as direct oxidation of the oil component. Capillary effects were included. Heat was assumed to be transported by convection and conduction in the reservoir and by conduction to the underlying and overlying formations.

Crookston (1979) and Grabowski (1979) reported on the development of ISCOM, a general purpose thermal simulator which was the predecessor of the commercial code STARS (Steam, Thermal and Advanced Processes Reservoir Simulator) from Computer Modelling Group, Inc. The model operated in three dimensions and included a variable number of components distributed among four phases (water, oil, gas, and solid). The number and type of chemical reactions were variable and both capillary and gravity forces were included. Heat transfer was by convection and conduction in the reservoir and by conduction in the underlying and overlying formations. Mass transfer was allowed between all phases under equilibrium conditions.

Among the most comprehensive models reported so far is the model by Coats (1980). It included a variable number of components and reactions and allowed the components to be distributed in any or all of the four phases (water, oil, gas, and coke). Gravity and capillary forces were incorporated. The permeability field was allowed to change with changing concentration of coke. Heat transfer was by convection, conduction, and radiation within the reservoir and by conduction in the cap and base rock. All models discussed so far assumed kinetically controlled chemical reactions and modeled phase equilibrium using K-value correlations.

A fundamentally different approach was recently presented by Akkutlu (2003). Inspired by the filtration combustion literature, they proposed a one-dimensional model and devised an analytical solution. They used a transport limiting approach for reactions and treated the combustion zone as a discontinuous front and derived appropriate jump conditions relating the change in variables across the front. The problem then reduced to modeling the combustion front dynamics and the regions on either side of the front. The analytical framework was developed for a single reaction in homogeneous media, but in subsequent papers the authors considered extensions to layered systems (Akkutlu2002) and two sequential reactions (Akkutlu2004).

The industry standard today for ISC simulations is STARS (CMG, 2004). It is a multi-component, general purpose thermal simulator with capabilities for handling both equilibrium and kinetically controlled chemical reactions. The simulator has extensive features including coupling with a geomechanics module. It uses an ideal fluid description with K-value based phase equilibrium and ideal phase properties (e.g. densities and enthalpies). With respect to numerical solution techniques, the simulator implements a standard control-volume finite difference scheme with a fully implicit (FIM) time discretization and an option for use of the adaptive implicit method (AIM), where gridblock variables are treated either explicitly or implicitly based on local CFL (Courant-Friedrichs-Lewy) criteria. Recently the simulator was extended with dynamic gridding capabilities (Christensen2004) for tracking of fronts (e.g. combustion or steam). In its current version, functionality is not available for adapting to multiscale behavior in time by tailored timestepping for

individual physical subprocesses such as convective transport and chemical reactions.

1.2. Project Goals

ISC has already been proven as an effective recovery process, and the interest in the process is likely to increase dramatically over the next decades, both as a primary recovery process for the world's vast heavy oil reserves and as a tertiary process in waterflooded, light-oil reservoirs. Predictive mathematical models and robust and efficient simulators are needed to improve our understanding of the ISC process and to enable safe and cost-effective design of ISC projects.

In this project we designed a suite of numerical strategies to improve the accuracy, reliaility and efficiency of ISC simulators. We also conducted an extensive suite of laboratory experiments to further our understanding of the physical behavior of these processes and investigate ways to improve their performance in the field.

1.2.1. Project Goals

Broadly speaking, ISC processes are challenging to describe physically and computationally because they involve a large number of subprocesses with strongly varying spatial and temporal scales. Of particular concern is kinetics: the ISC process is largely driven by chemical reactions that take place on very short spatial scales and very fast temporal scales.

In our research we pursue

1. Improved measurement of the kinetics of LTO, MTO, and HTO reactions through establishing a test suite of oil combustion kinectics and the application of so-called ramped temperature oxidation to crude oil oxidation.
2. Modification of ISC kinetics through the introduction of water-soluble metallic additives. Such additives ultimately can be injected in the aqueous phase prior to combustion.
3. Exploration of the interplay of porous medium heterogeneity and combustion via combustion tube experiments.
4. Maximum in-situ upgrading through a process of cyclic injection of solvent to precipiate asphaltenes thereby increasing the API gravity of crude oil and air injection to remove crude oil precipitates through combustion.
5. Adaptive mesh refinement in space to allow reasonable resolution of temperature, combustion and propagation fronts
6. Adaptive grid refinement in time to prevent local fast processes from affecting computational efficiency in regions with slower processes.
7. Splitting methods to decouple physical subprocesses
8. Specialized solvers for each physical subprocess
9. Improved upscaling methods for processes at small spatial scales.

In this project, we have made substantial progress in all areas. We have laid the foundations for area 7. However, successful splitting requires a detailed analysis of the significant scales in a ISC process. Although such an analysis can be performed for specific cases, no general procedure is yet available. We hope to continue in this area in subsequent projects.

Area 9 we identify as an area in need of considerable future funding. This can not be performed without very close interactions with an experimental program. We have made a start in this area as you can see from our report on the experimental program.

1.3. Physical and Mathematical Modeling of ISC Processes

1.3.1 Common assumptions

A component in a reservoir can be transported by molecular diffusion and convection. The relative importance of these processes is measured by the Peclet number defined as the ratio of a characteristic time for diffusion to a characteristic time for convection. Peclet numbers for typical reservoir flows of interest are often large enough that the effects of diffusion can be neglected. Similarly, heat can be transported in a reservoir by conduction and convection. As for mass transport, the relative importance of heat conduction and convection is measured by the heat Peclet number. Using typical reservoir parameters for flow rate, interwell distance, and heat diffusivity, it can be shown that the heat Peclet number is 1--3 orders of magnitude lower than the equivalent mass Peclet number. Thus, the effect of heat conduction is normally included in thermal simulation models. Especially in the start-up of ISC processes, when flow rates are small, heat conduction is important.

A condition generally assumed to prevail in ISC, as well as in all other reservoir processes, is that the fluids are in thermodynamic equilibrium at every point within the reservoir (Prats1986). This equilibrium assumption implicitly states that the timescales for the interphase mass transfer processes occurring when phases come to an equilibrium state, are much faster than all other timescales. Consider the simple example of an oil phase and a gas phase contained in a pore of diameter $[1]\mu m$. Using a typical value for a diffusion coefficient of $[10^{-9}]m^2 s$, the average time required for a molecule to diffuse across the pore would be on the order of $[10^{-3}]s$. This time is likely to be much smaller than the time required for flow to change the compositions significantly within the pore. Thus, in a non-reacting system it is reasonable to assume thermodynamic equilibrium. The validity of the assumption for reactive systems is unclear, but in lack of better alternatives we shall assume thermodynamic equilibrium conditions.

In addition to thermodynamic equilibrium it is commonly assumed that the fluids and reservoir rock are in thermal equilibrium.

1.3.2 System specifications

We model a porous media system containing n_c chemical (pseudo) components distributed among three mobile fluid phases (oil, water, and gas) and one immobile solid phase. Components existing in the solid phase are assumed to do so

exclusively. The mobile components are assumed to exist in at most two phases, with the water component partitioning into the water and gas phases, and the remaining components partitioning into the oil and gas phases. The chemical components are indexed as

$$i \in \mathsf{I} = \mathsf{I}^{ol} \cup \mathsf{I}^{w} \cup \mathsf{I}^{s}, \tag{1.1}$$

where

$$\begin{aligned} \mathsf{I}^{ol} &= \{1, \ldots, n_{ol}\}, \\ \mathsf{I}^{w} &= \{n_{ol}+1\}, \\ \mathsf{I}^{s} &= \{n_c - n_s + 1, \ldots, n_c\}. \end{aligned} \tag{1.2}$$

Thus, the total number of components is $n_c = n_{ol} + n_s + 1$.

We represent the porous medium by Ω with boundary $\mathsf{S} = \partial\Omega$ and outward pointing normal vector $\mathbf{n}$. Points in the domain are denoted $\mathbf{x} = [x\ y\ z]^T$. Intrinsic to the porous medium are the total bulk volume denoted, V_T, the void pore volume, V_p, and the void porosity, $\phi_v = V_p V_T$. The pore space is entirely occupied by the fluid and solid phases. We assume that the reservoir rock is incompressible and chemically inert. It is convenient to define the fluid porosity, ϕ_f, relating the bulk volume to the volume open to flow. If $V_f = V^o + V^w + V^g$ denotes the total fluid volume, then the fluid porosity is given by

$$\phi_f \equiv \frac{V_f}{V_T} = \frac{V_p - V^s}{V_T} = \phi_v - \sum_{i \in \mathsf{I}^s} \frac{C_i}{\xi_i^s}, \tag{1.3}$$

where C_i is the concentration per unit bulk volume and ξ_i^s is the molar density of solid component i. The volumetric saturation of phase j is defined as

$$S^j = \frac{V^j}{V_f} = \frac{V^j}{\phi_f V_T}, \qquad j \in \{o, w, g\}. \tag{1.4}$$

1.3.3 Governing equations

The governing equations for thermal, compositional and reactive porous media flow can be derived from n_c component conservation equations, an energy conservation equation, and a volume conservation equation. The i-th component conservation equation can be written in integral form as

$$\frac{d}{dt}\int_\Omega C_i d\Omega + \int_{\mathsf{S}} \mathbf{q}_i^m \cdot \mathbf{n} d\mathsf{S} = \int_\Omega \left(Q_i^{m,reac} + Q_i^{m,well}\right) d\Omega, \tag{1.5}$$

where $i \in \mathsf{I}$, $\mathbf{q}_i^m$ is the mass flux, and $Q_i^{m,reac}$ and $Q_i^{m,well}$ are the mass source densities due to chemical reactions and wells, respectively. We can express the phase and total component concentrations as

$$c_i^j = \phi_f x_i^j \xi^j S^j, \qquad i \in \mathsf{I}^{ol} \cup \mathsf{I}^{w} \tag{1.6}$$

$$C_i = \sum_j c_i^j, \tag{1.7}$$

$$c_i^s = \left(\phi_v - \phi_f\right) x_i^s \xi^s, \qquad i \in \mathsf{I}^s, \tag{1.8}$$

$$C_i = c_i^s, \tag{1.9}$$

where $j \in \{o, g\}$ for $i \in \mathsf{I}^{ol}$, and $j \in \{w, g\}$ for $i \in \mathsf{I}^w$. x_i^j is the mole fraction of component i in phase j, and ξ^j is the molar density of phase j. In (1.5) we note that

$$\mathbf{q}_i^m = \mathbf{0}, \qquad i \in \mathsf{I}^s. \tag{1.10}$$

Neglecting capillary forces, we can express the component flux term in (1.5) using the standard multi-phase extension of Darcy's law

$$\mathbf{u}^j = -\frac{k_r^j}{\mu^j}\mathbf{k}\left(\nabla P - \rho^j g \nabla D\right) \tag{1.11}$$

where $\mathbf{u}^j$ is the phase flow velocity, $\mathbf{k}$ is the absolute permeability, k_r^j is the relative permeability, μ^j is the viscosity, ρ^j is the mass density, g is the gravitational constant, and D is the reservoir depth. The total component flux is the sum of the fluxes in the individual phases, given by

$$\mathbf{q}_i^m = \sum_j x_i^j \xi^j \mathbf{u}^j, \qquad i \in \mathsf{I}^{ol} \cup \mathsf{I}^w. \tag{1.12}$$

The energy conservation equation is

$$\frac{d}{dt}\int_\Omega U d\Omega + \int_\mathsf{S}\left(\mathbf{q}^{h,adv} + \mathbf{q}^{h,cond}\right)\mathbf{n} d\mathsf{S} = \int_\Omega \left(Q^{h,reac} + Q^{h,well}\right) d\Omega, \tag{1.13}$$

where U is the total internal energy of the system, $\mathbf{q}^{h,adv}$ is the heat transport due to advection, $\mathbf{q}^{h,cond}$ is the heat transport due to conduction, and $Q^{h,reac}$ and $Q^{h,well}$ are the heat source densities due to chemical reactions and wells, respectively. The total internal energy can be expressed as

$$U = \left(1 - \phi_v\right) U^r + \left(\phi_v - \phi_f\right) \xi^s U^s + \phi_f \sum_j \xi^j S^j U^j, \tag{1.14}$$

where U^r is the volumetric internal energy of the reservoir rock and U^j is the molar internal energy of phase j. The advective heat transport in (1.13) is the sum of the heat carried by each mobile phase

$$\mathbf{q}^{h,adv} = \sum_j h^j \xi^j \mathbf{u}^j, \tag{1.15}$$

where h^j is the molar enthalpy. Fourier's law is used to express the transport due to conduction

$$\mathbf{q}^{h,cond} = -\mathbf{k}_c \nabla T, \tag{1.16}$$

where the tensorial quantity $\mathbf{k}_c$ is the effective thermal conductivity of the saturated medium.

Finally, the principle of volume conservation expresses that the fluid and solid phases must fill the pore space exactly

$$V_p = \sum_j V^j, \tag{1.17}$$

where $j \in \{o, w, g, s\}$.

In addition to the conservation requirements, we assume that the system is in thermodynamic equilibrium as expressed by the equality between chemical potentials, or equivalently fugacities, of a component in all phases. For components partitioning between the oil and gas phases we have

$$f_i^o = f_i^g, \qquad i \in \mathsf{I}^{ol}, \tag{1.18}$$

where f_i^j is the fugacity of component i in phase j.

1.3.4 Boundary conditions

We assume no flow of mass or heat across the reservoir boundaries corresponding to Neumann conditions

$$\mathbf{q}_i^m \cdot \mathbf{n} = 0, \qquad \mathbf{x} \in \partial\Omega, \tag{1.19}$$

$$\left(\mathbf{q}^{h,adv} + \mathbf{q}^{h,cond}\right)\mathbf{n} = 0, \qquad \mathbf{x} \in \partial\Omega. \tag{1.20}$$

We note that thermal reservoir models usually account for heat exchanges with the surroundings, typically through heat conduction to the reservoir cap and base rock (see e.g., Coats (1980),Crookston (1979)). Since in this thesis we focus primarily on one-dimensional problems, we model heat exchanges as a source/sink term in the equation.

We consider wells completed in a single gridblock and two types of well operation: fixed single-phase injection rate for injectors and fixed back pressure for producers. Typically, a fixed gas rate (air or oxygen enriched air) is specified, or both water and gas rates are specified as for wet combustion processes. The wells are indexed as $w \in \mathsf{I}_w$, which consists of injectors and producers $\mathsf{I}_w = \mathsf{I}_w^{inj} \cup \mathsf{I}_w^{prod}$.

The component source density due to wells can be written as

$$Q_i^{m,well} = -\sum_j \frac{x_i^j \xi^j k_r^j}{\mu^j} (WI)_w (P - P_w), \qquad w \in \mathsf{I}_w^{prod} \tag{1.21}$$

for producers, where $(WI)_w$ is the well index, P is the block pressure, and P_w is the well bottom-hole pressure. For injectors, $Q_i^{m,well}$ can be computed directly from the specified injection rates. Again, since we are concerned primarily with one-dimensional combustion tube problems, we omit the well equation for injectors and include the injection rates directly in the equation. For standard block-centered grids the well index in (1.21) is

$$(WI)_w = \frac{2\pi kh}{\ln r_0 r_w + s}, \tag{1.22}$$

where k is the effective permeability, h is the formation thickness, r_w is the well radius, r_0 is the pressure equivalent radius, and s is a skin factor. Expressions for r_0 can be found in Aziz (1979) for general cartesian grids.

The heat source density due to wells can be written as

$$Q^{h,well} = -\sum_j \frac{h^j \xi^j k_r^j}{\mu^j} (WI)_w (P - P_w), \qquad w \in \mathsf{I}_w^{prod}. \tag{1.23}$$

Heat exchange with the surroundings is modeled using an external proportional heat controller

$$Q^{h,htr} = -u_a (T - T^{ext}) \tag{1.24}$$

where u_a is a heat transfer coefficient (or controller gain) and T^{ext} is the controller set point.

1.3.5 Constitutive models

We use a dual strategy for phase behavior and fluid property modeling, using both a cubic EoS and an ideal mixture approach, where, in the latter, phase behavior is modeled from pressure-temperature K-value correlations.

In the ideal mixture approach the K-values are calculated either from interpolation in tables generated from EoS flashes of the original oil or from the Wilson correlation

$$\ln K_i = \frac{P_i^{crit}}{P} + 5.373(1 + \omega_i)\left(1 - \frac{T_i^{crit}}{T}\right), \tag{1.25}$$

where P_i^{crit}, T_i^{crit}, and ω_i are the critical pressure and temperature and accentric factor, respectively, for component i.

In the EoS approach, the fugacities in (1.18) are obtained from a cubic EoS of the form

$$(Z^j)^3 + c_2 (Z^j)^2 + c_1 (Z^j) + c_0 = 0, \tag{1.26}$$

where Z^j is the phase compressibility factor. The coefficients c_2, c_1 and c_0 depend on pressure, temperature and phase composition. A common choice of EoS in reservoir simulation is the Peng-Robinson equation (Peng1976), which is used in this work. Details on the EoS and computation of fugacities can be found in e.g. Michelsen (2004) or Whitson (2000).

All chemical reactions are assumed to be kinetically driven. The reactions are modeled using standard Arrhenius rate relations, where the rates depend on component concentrations in a specific phase. Both homogeneous and heterogeneous reactions are allowed. n_r chemical reactions are assumed to occur, which are indexed as $\gamma \in \mathsf{I}^{rx}$. Let $\mathbf{A} \in \mathsf{R}^{n_c \times n_r}$ denote the stoichiometry matrix, where element $A_{i\gamma}$ is the stoichiometry coefficient for component i in reaction γ (negative for reactants, positive for products). To conserve mass, the constraint

$$\mathbf{M}^T \mathbf{A} = \mathbf{0} \tag{1.27}$$

must hold, where $\mathbf{M} = [M_1 \cdots M_{n_c}]^T$ is the vector of component molecular weights.

Generally, two kinds of reactions dominate in ISC processes: cracking/pyrolysis reactions and oxidation (full/partial) reactions. For oxidation

reactions the partial pressure of oxygen is used in the rate expressions. If the set of triples $(i, j, m) \in \mathrm{J}(\gamma)$ index the involved reactants, phases and reaction orders for reaction γ, then the reaction rate may be expressed as

$$R_\gamma = K_\gamma \left(P x_{O_2}^g \right)^{n_\gamma} \prod_{(i,j,m)\in \mathrm{J}(\gamma)} \left(c_i^j \right)^m, \tag{1.28}$$

where n_γ is the reaction order in oxygen partial pressure and c_i^j is the concentration of component i in phase j according to the definitions (1.6) and (1.8). The rate constants are temperature dependent according to Arrhenius's law

$$K_\gamma = \alpha_\gamma \exp\left(-\frac{E_a^\gamma}{R_g T} \right), \tag{1.29}$$

where α_γ is the frequency factor, E_a^γ is the activation energy and R_g is the universal gas constant.

From the reaction rates we can express the mass and heat source densities in (1.5) and (1.13). The net production of component i in chemical reactions is

$$Q_i^{m,reac} = \sum_\gamma A_{i\gamma} R_\gamma. \tag{1.30}$$

Similarly, the net heat generation is

$$Q^{h,reac} = \sum_\gamma (-\Delta H)_\gamma^r R_\gamma, \tag{1.31}$$

where $-\Delta H^r$ is the heat of reaction.

Relative permeability

Three-phase relative permeabilities are, in general, obtained from two-phase experiments, where the relative permeability of the ``middle'' phase is interpolated from two-phase data. We use the Stone II model (Stone1973) to calculate oil phase relative permeability from water-oil and gas-oil curves, which can be entered either as tables or correlations. Temperature effects and possible wettability changes due to chemical reactions (e.g. deposition of coke) have not been considered in this work.

Density

Phase densities are evaluated from the EoS by solving Eqn. (1.26) for Z^j. The molar and mass phase densities are

$$\xi^j = \frac{P}{R_g T Z^j}, \qquad \rho^j = \xi^j \sum_i x_i^j M_i. \tag{1.32}$$

In the ideal case we assume $Z^g = 1$ and obtain pure component liquid densities from

$$\xi_i^o = \xi_{i,ref}^o \exp\left(c_i^p (P - P_{ref}) - c_i^t (T - T_{ref}) \right) \tag{1.33}$$

where $\xi^o_{i,ref}$ is the density at the reference consitions, and c_i^p and c_i^t are the compressibility and thermal expansion coefficients, respectively. Linear mixing of molar volumes is assumed

$$\frac{1}{\xi^o} = \sum_i \frac{x_i^o}{\xi_i^o}. \tag{1.34}$$

Water density is calculated according to a correlation similar to (1.33).

Viscosity

Accurate modeling of oil phase viscosity is important for ISC processes, since the reduction in oil viscosity with temperature is the primary mechanism for enhanced recovery. Pure component viscosities can be entered either through tables or obtained from the correlation

$$\mu_i^o = a_i^o \exp\left(\frac{b_i^o}{T}\right), \tag{1.35}$$

where a_i^o and b_i^o are correlation constants. The oil phase viscosity is then obtained by logarithmic mixing of the pure component viscosities

$$\ln \mu^o = \sum_i x_i^o \ln \mu_i^o. \tag{1.36}$$

Accuracy in gas phase viscosity is generally less important. The correlation used here is

$$\mu_i^g = a_i^g T^{b_i^g}, \tag{1.37}$$

and the mixing rule used is

$$\mu^g = \sum_i \frac{x_i^g \sqrt{M_i}\, \mu_i^g}{x_i^g \sqrt{M_i}}. \tag{1.38}$$

Water viscosity is obtained from a correlation similar to (1.35).

Enthalpy

Phase enthalpies are computed from ideal gas heat capacity correlations and vaporization enthalpy correlations. For a pure component gas we have

$$C_{p,i}(T) = C_{p,i}^1 + C_{p,i}^2 T + C_{p,i}^3 T^2 + C_{p,i}^4 T^3, \tag{1.39}$$

$$h_i^{g,id} = \int_{T_{ref}}^{T} C_{p,i}(\overline{T})\, d\overline{T}, \tag{1.40}$$

where T_{ref} is the temperature of the reference state. In the ideal case the vaporization enthalpy is obtained from the correlation (CMG, 2004)

$$h_i^{vap} = hv_i^1 \left(T_i^{crit} - T\right)^{hv_i^2}, \tag{1.41}$$

where hv_i^1 and hv_i^2 are correlation constants. Thus, for the ideal case we have

$$h^g = \sum_i x_i^g h_i^{g,id} \tag{1.42}$$

for the gas phase and

$$h^o = \sum_i x_i^o \left(h_i^{g,id} - h_i^{vap} \right) \tag{1.43}$$

for the oil phase. In the non-ideal case the residual enthalpies are obtained from the EoS leading to

$$h^g = \sum_i x_i^g h_i^{g,id} + h^{g,res} \tag{1.44}$$

$$h^o = \sum_i x_i^o h_i^{g,id} + h^{o,res}, \tag{1.45}$$

where $h^{j,res}$ is the residual enthalpy of phase j. The water phase enthalpy is obtained from correlations similar to (1.41).

Internal energy

The volumetric internal energy of the reservoir rock is modeled as

$$C_v(T) = C_v^1 + C_v^2 T, \tag{1.46}$$

$$U^r = \int_{T_{ref}}^{T} C_v(\overline{T}) d\overline{T}. \tag{1.47}$$

The solid phase heat capacity is assumed constant, and the molar internal energy of the fluid phases is obtained from the enthalpy by subtraction of the mechanical work

$$U^j = h^j - \frac{P}{\xi^j}, \qquad j \in \{o, w, g\}. \tag{1.48}$$

Thermal conductivity

The mobile and immobile phases have associated with them a scalar thermal conductivity k_c^j. Moreover, the rock has an associated tensorial thermal conductivity, denoted $\mathbf{k}_c^r$. The effective thermal conductivity of the saturated porous medium is computed as a weighted volume average

$$\mathbf{k}_c = (1 - \phi_v)\mathbf{k}_c^r + \left[(\phi_v - \phi_f) k_c^s + \phi_f \sum_j S^j k_c^j \right] \mathbf{I}, \tag{1.49}$$

where $\mathbf{I}$ denotes the identity matrix. In this thesis we assume the rock conductivity to be isotropic.

1.3.6 Primary equations and variables

According to Gibbs' phase rule, to fix the intensive state of a system with n_c components and n_p phases, we must specify

$$f = n_c - n_p + 2 \tag{1.50}$$

intensive variables. Since the bulk volume, V_T, is a known quantity, we need an additional $n_p - 1$ phase volume fractions (saturations) to fix the amounts of each phase and thereby the extensive state of the system. Consequently, the number of degress of freedom is

$$n_c - n_p + 2 + (n_p - 1) = n_c + 1. \tag{1.51}$$

In a thermal-compositional reservoir model the state of the system is therefore uniquely determined by $n_c + 1$ independent variables. Thus, we must select $n_c + 1$ primary variables and equations when solving the model.

The natural choice of primary equations are the n_c mass balances (1.5) and the energy balance (1.13). For the choice of variables there are, however, several possibilities. Cao2002 review compositional models and group them according to choice of variables. Generally, there are two classes of variable sets: *natural variables* and *overall quantity variables*. Natural variables are based on saturations and mole fractions, whereas overall quantity variables are based on total amounts of each component. The argument, often made in favor of natural variables, is that they appear *naturally* in the conservation equations, and that derivatives (needed in implicit formulations) therefore are straightforward and inexpensive to compute. When working with overall quantities (such as total moles of each component) as variables, some derivatives will be more expensive to evaluate, which increases the cost of Jacobian computations. The main drawback of natural variables is the need to switch variables when the phase state in a gridblock changes. Since natural variables are intensive phase properties, disappearance or reappearance of a phase requires a switch of primary variables, because the phase properties are only defined when the phase is present. Variable switching is administratively difficult and may affect convergence of the nonlinear iterations in implicit formulations.

In reactive systems, a reformulation of the governing equations is often used based on the *extent of reaction* (Aris1989). For single phase, batch reactor problems this is straightforward and leads to a system of equations in n_r unknowns instead of n_c, where n_r is the number of independent chemical reactions. The extent of reaction concept has been generalized to open systems involving flow (Friedly1991,Friedly1992). For certain reactive porous media transport problems in water resources modeling this generalized extent of reaction formulation leads to a reduced and simpler set of equations to be solved (Friedly1992,Krautle2005). For multiphase ISC problems, however, where the transport operator is nonlinear and components move between phases due to phase behavior, the reformulation is not directly applicable.

In this thesis we choose as primary variables the n_c overall component concentrations, C_i, and temperature and pressure. With overall quantities as variables, we avoid the complicated variable switching logic when phases change. Since the volume constraint (1.17) is nonlinear in the primary variables, we include it as an additional primary equation and work with the full set of component

concentrations. Thus, we solve simultaneously a set of $n_c + 2$ equations per gridblock, where temperature is aligned with the energy balance, pressure is aligned with the volume constraint, and the component concentrations are aligned with the respective component balances. We note that the volume constraint as well as the component balances for all immobile solid components can be eliminated from the primary equation set at the Jacobian level, since both types of equations are local to each gridblock.

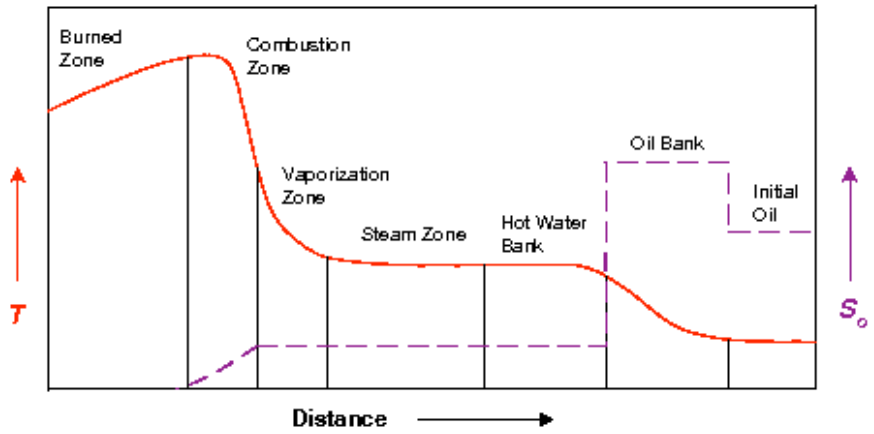

Figure 1.1: Schematic representation of characteristic temperature and saturation distributions in forward in-situ combustion [adapted from][not to scale]Prats1986.

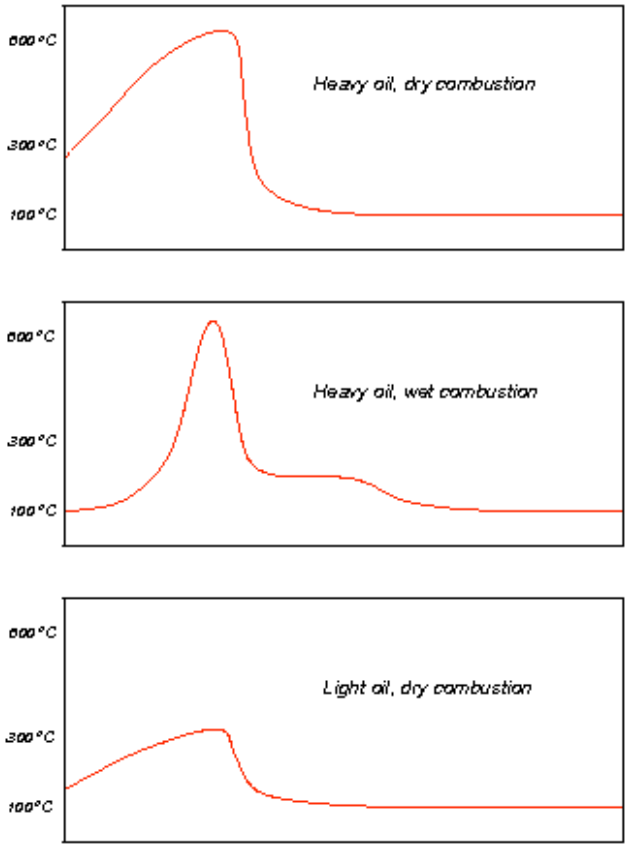

Figure 1.2: Characteristic temperature profiles in ISC for heavy/light oil and dry/wet combustion.

2. Experimental

The experimental aspects of this project make use of a common apparatus for measurement of the kinetics of in-situ combustion as well as our ability to propagate a combustion front within porous media. This apparatus is described shortly. Additionally, portions of our work make use of scanning electron microscopy (SEM) and so our SEM facilities are also described. The question of heat losses is important to the design of the combustion tube and so is also addressed in this section. The specialized procedures for the various experimental studies are discussed later.

2.1 Kinetics Cell and Combustion Tube

The kinetics cell, or alternately the combustion tube, is coupled to a gas analyzer, traveling thermocouple, and data logging system. The combustion tube measures the ability of a combustion front to propagate through a given combination of crude oil, water, and porous medium. The kinetics cell is used to characterize activation energies of combustion as well as oxygen consumption. The system is diagramed schematically in Fig 2.1. Equipment and manufacturer are summarized in Table 2.1 Briefly, the gas analyzer collects effluent gas composition. Oxygen and nitrogen are provided by gas cylinder and metered using a mass flow controller. Temperatures are recorded every centimeter along the length of the sandpack during a combustion tube run, while in the kinetics cell experiment, temperatures are measured in the center of the cell.

The combustion tube is made of stainless steel (316) with a wall thickness of 0.14 cm. It is 1.0 meter long with an inner diameter of 7.5 cm. It is packed with a multiphase mixture of sand, clay, oil, and water as detailed later. A 240V, 1000W igniter coil surrounds the tube between 10 to 15 cm from the top of the tube. Prior to ignition, the combustion tube is placed vertically in an insulator jacket and the annular space between the tube and the wall of the jacket is filled with insulating vermiculite or FiberFrax™ insulation(Unifrax Corp., Niagara Falls, NY) . A Matheson mass flow controller (model 8240, range 0.1-5 SLPM) is connected upstream of the tube to control the rate of air injection. At the outlet of the combustion tube, liquid is separated from the gaseous products. Centrifuge vials (50ml volume) are used to collect the produced liquid for later analysis.

The gas leaving the separator is cooled by a condenser and then dehumidified using drierite. A backpressure regulator elevates the system backpressure to 100 psig (690 kPa), or greater. A part of the effluent stream then flows through two acid scrubbers containing potassium permanganate. All the other effluent is vented. The scrubbed effluent is then directed to a Xentra gas analyzer (model 4200, 0.1% error). This is a compact gas analyzer, measuring oxygen, carbon dioxide, carbon monoxide, and methane concentration. Measurements are made once per minute. The outlet streams from the analyzer are vented.

A thermal well (3.2 mm O.D., 316 stainless steel) is set in the center of the tube and spans from top to bottom. During the combustion process, a traveling thermocouple measures the temperature as a function of time and distance from the

top. A measurement interval of 1-5 cm is generally used depending on the temperature profile and front velocity. The interval is chosen so that a complete traverse of the tube occurs with negligible movement of the combustion front.

The kinetics cell is a thick-walled stainless steel (316) cylinder that is 13.3 cm long and 4.82 cm (O.D) in diameter, Fig. 2.2 . The cylinder is sealed by means of knives and annealed copper gaskets. The 0.76 mm (0.030 in) knives were machined into the end face of the cylinder and the endcaps. With pressure at the two ends, the knives on the end of cylinder and the endcaps cut into the gaskets providing a contact seal. Air is injected from the bottom of the cell. Prior to injection, air flows through coiled 3.2 mm (1/8 in.) tubing and is preheated to the same temperature as the kinetics cell. On the top of the cell air exits and the temperature is measured.

The kinetics cell is packed with a mixture of oil, water, and sand. Two thin-walled stainless steel cups are put inside the kinetics cell. The upper cup is 7.1 cm long with I.D of 2.7 cm. This cup is filled with the oil and sand mixture. The lower cup is filled with dry sand. It serves as a preheater and disperses the inlet air to obtain one-dimensional air flow. Both cups are perforated on the bottom to allow air flow. Sand migration is prevented by 200-mesh stainless steel screens placed across the bottom of each cup.

2.2 Heat Losses

Minimization of heat losses is of substantial concern during a combustion tube experiment. Excessive heat losses slow front progression and lead to an extinguished burn . Heat losses are controlled here by insulating extensively around the combustion tube and using moderately large air injection rates that lead to significant heat generation and relatively rapid progression of the combustion front. In the design of the experiments, we attempt to keep the ratio of the heat lost upon the heat generated during combustion at 10% or less. In this way, heat losses have little effect on the dynamics of front progression.

Heat losses are analyzed in a fashion similar to the heat losses from an insulated tubing (Prats, 1986). The dominant resistance to heat loss is the FiberFrax insulation packed in the annular space between the combustion tube and the containment vessel. The overall heat transfer coefficient is written

$$U = \left[\frac{1}{2\pi\lambda_{ins}} ln\left(\frac{r_{ins}}{r_t}\right) \right]^{-1} \tag{2.1}$$

where r refers to radius, λ is the thermal conductivity, the subscript ins refers to insulated, and the subscript t refers to the tube. For calculations, we take λ_{ins} as 0.04 W/m-K (manufacturer's measurement), r_{ins} as 6.10 cm, and r_t as 3.81 cm. Hence, U is found to be 0.53 W/m^2-K.

The heat lost per unit time, Q_l is calculated as

$$Q_l = UA\Delta T \tag{2.2}$$

For the surface area, we take a length of 0.5 m because this is the practical length of the tube combusted during an experiment. In a typical experiment with heavy oil, 275°C is roughly the average tube temperature within the combustion. The quantity ΔT is taken as 275 - 20 °C. The average rate of energy loss is 16 W.

The energy generated from combustion is estimated using Fig. 5.12 of Burger et al. (1985) For complete combustion (i.e., no carbon monoxide production) and an atomic hydrogen to carbon ratio of 1 for the fuel, it is estimated that combustion releases 18,000 kJ per m^3 of O_2. At the injection rate of 3 SLPM of air, the flux of O_2 is 1.05×10^{-5} m^3/s. Hence, the rate of heat released by combustion is 190 W.

Our estimate of the ratio of heat lost to heat generated is 16/190 or about 8 %. The fraction of heat lost during a combustion tube experiment is minor in comparison to the heat generated and heat losses have little effect on front propagation.

2.3 Scanning Electron Microscope

Microscopic imaging and elemental analysis were employed to evaluate any textural or compositional changes in the sand-clay mixture upon addition of metallic additive. Samples for SEM analysis were prepared as grain mounts. Sand-clay mixtures were placed on double sided adhesive carbon tabs and mounted on cylindrical aluminum SEM mounts. Grain mounts were coated with a monolayer of gold or a gold-palladium mixture to provide a conductive coating on the surface of the sample. This reduces possible problems with surface charging that is common during SEM analysis.

A JEOL JSM-500LV SEM system was used to collect digital images. The system captures secondary electrons as well as back-scattered electrons. Additionally, an EDAX Sapphire Si(Li) energy dispersive system (EDS) with GENESIS software allowed elemental analysis and mineral identification.

Table 2.1 – Equipment used in the experimental apparatus.

Mass flow controller:	Brooks Instruments, model 5850E
Gas flow meter:	Omega, model FMA-A23
Gas analyzer:	Servomex Xentra analyzer, model 4200C.
Temperature controller:	Omega, model CN 8201.
Data logger:	Measurement Computing, model USB-TC
Digital thermometer:	Omega, model 2176A.
Back pressure regulator:	Tescom, model 26 – 1727.

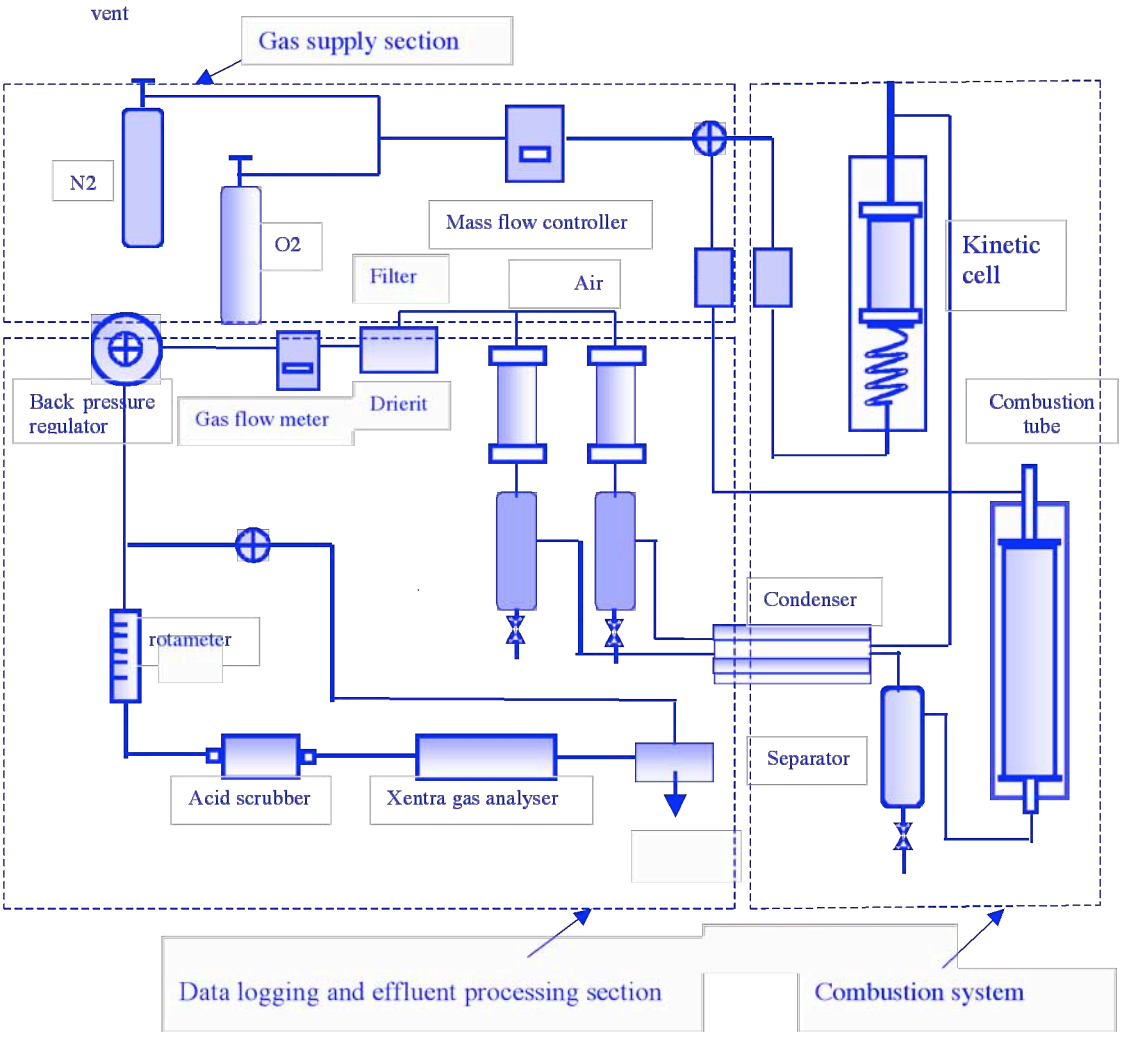

Figure 2.1. Schematic of experimental apparatus for in-situ combustion kinetics and combustion tube tests.

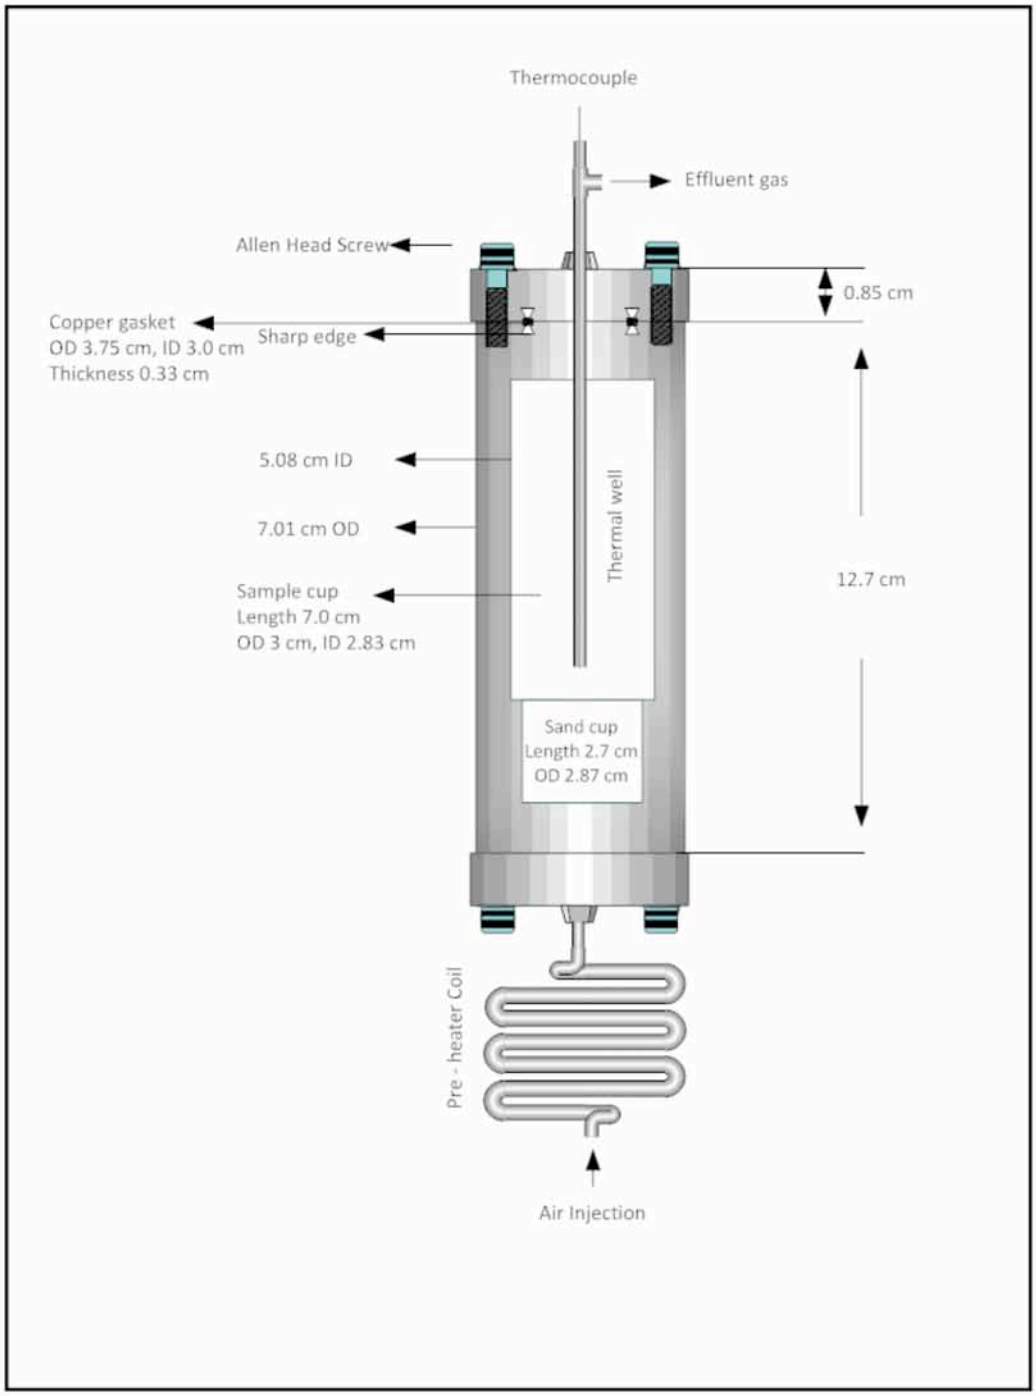

Figure 2.2 – Kinetics cell.

3. Water Soluble Additives to Improve ISC

A key mechanism for successful ISC is effective oxidation of a small fraction of the oil in place to generate heat and pressure. Common, water-soluble metallic salts are shown to play an important role as catalysts for some combinations of crude oil/brine/solid matrix (Shallcross et al., 1991; Castanier and Brigham, 2003). Metallic additives enhance oxidation and cracking of hydrocarbons and thereby affect the nature and the amount of fuel formed. The mechanism of the catalytic effect is, as of yet, unknown and is the focus of this experimental study. Tube runs that gauge combustion performance and ramped temperature oxidation tests that measure the kinetics of combustion are used in parallel. We propose cation exchange of metallic salts with clay as a mechanism to create activated sites that enhance combustion reactions between oil and oxygen. Sand and clay surfaces are examined with scanning electron microscopy for evidence of cation exchange and alteration of surface properties by metallic salts.

In this study, metallic additives for in situ combustion are investigated both quantitatively and qualitatively. In particular, the catalytic effects of metallic additive with surface-active clay or surface-inert materials are compared. Clay, in the presence of metallic additive, enhanced the extent of oxidation reactions. On the other hand, the effect of metallic additive is not obvious when clay is substituted with a silica powder that has significant surface area. An hypothesis is that the metal cations substitute into the clay providing activated sites that improve oxidation. This section proceeds by presenting the procedures and experimental design. Results from combustion tube runs, kinetics cell experiments, and scanning electron microscopy (SEM) follow. The various kinetic data obtained are then discussed.

3.1 Procedures

To observe and confirm combustion performance without and with metallic additive, two tube runs using Cymric light crude oil (34 °API, Kern, Co., CA) were completed. One run had metallic additive on the top half of the tube and the second run had additive across the whole tube. Additional runs are summarized in the M.S. report of He (2004).

To investigate the effect of additive on activation energy and oxidation, a variety of kinetic runs were made with and without metallic additive. The Cymric light crude as well as Cymric heavy crude were used. Also the effect of clay was investigated, by comparing runs in which kaolinite was used with runs employing surface-inert silica powder (silicon dioxide floated powder, about 240 mesh, Fisher Chemical). The silica powder has roughly the same surface area and size as the kaolinite, but is chemically inert. The grain-size distribution of the sand is given in Table 3.1. The sand is fired at 800 °C to deactivate any minerals prior to use. To fire the sand, it is put in a stainless steel container. The temperature is increased in stages. The sand is held at a particular temperature for about an hour. This process is continued until the temperature reaches 800 °C. The sand is held at 800 °C for at least 8 hours.

In general, experiments are prepared and conducted as follows: sand preparation, temperature and data recording, liquid collection, and flue gas analysis. The combustion tube is pressure tested at 100 psi (690 kPa) prior to packing with the oil-water-sand mixture. Then the packing mixture is prepared. Sand and clay or other solid powders are weighed and put in a basin. The sand is stirred until a homogeneous mixture is obtained. Water and then oil are added to the sand mixture and stirred. If metallic additives are used, they are first dissolved in water. The main additive examined here is the ion Fe^{3+} provided in the form of ferric nitrate ($Fe(NO_3)_3 \bullet 9H2O$, >99% purity, Baker & Adamson Chemicals). The thermal well is put in the center along the tube before packing. The mixture is then packed into the tube at a constant rate and with equal stress to give a relatively even distribution of porosity and permeability along the tube. About 5 cm^3 of linseed oil saturated sand are added to the top of the pack to bring about quick and uniform ignition. On the very top of sand mixture about 5 cm of clean sand is added to prevent heat damage to the top flange. The tube is then bolted together, put in the jacket, and assembled. The inlet connects to the gas supply line; the outlet is connected to a separator. After line connection, a pressure test of 100 psi (690 kPa) is again carried out using nitrogen. Then the igniter is turned on and nitrogen injection begins at roughly 1 SLPM. When temperature reaches 550 °C, the igniter is turned off and gas is switched from nitrogen to air. This is the start time of data logging. Air injection is held constant and inlet pressure stabilizes at about 100 psi (690 kPa) throughout the experiment.

When the combustion front moves to a position of 90 cm from the top or an obvious rise in oxygen concentration of the effluent is observed and maintained, the combustion is quenched by switching from air to nitrogen After cooling off, the equipment is dissembled. The tube is unpacked and sand samples collected as a function of distance from the top.

Preparation of the apparatus for a kinetics run begins by mixing sand and other powders in a 100ml beaker. Measured amounts of water and oil are again mixed until a homogeneous mixture is obtained. Metallic additive, if any, is dissolved in water. The sand mixture is tamped into the cup using a stainless steel measuring spoon.

After the cell is full, the end plug is put on and tightened. The cell is pressure tested at 140 psi (965 kPa) with nitrogen. Then the kinetics cell is placed into the furnace, Figs. 2.1and 2.2, and connected. The furnace is turned on and temperature increased at about 60 °C/hr. Simultaneously, air is injected at 0.25 SLPM. Backpressure is maintained at 45 psi (310 kPa). The gas analyzer measures the component concentration in flue gas.

3.2 Results–Tube Runs

Two tube runs were performed with Cymric light (34 °API) crude oil to explore the benefits of metallic additive and whether additive transported within the combustion tube after packing. Previous research showed that this oil did not sustain combustion without additive (Castanier et al., 1992). The properties of the oil and sand pack are listed in Table 3.2.

In the first tube run, nitrogen was injected at 3 SLPM. Once the temperature reached 400 °C, the igniter was switched off and air injection begun. The air injection rate was also 3 SLPM. The combustion lasted for 9 hrs.
Metallic additive packing is only present in the first half of the tube. The additive concentration is at 0.5% by weight. The front position and temperature versus time are plotted in Fig. 3.1. The front moves at a constant velocity of about 0.14 cm/sec although some deviation is observed. The front temperature varied between 400-450 °C until 370 mins where the front was 56 cm from the top. This is slightly beyond the additive/no additive interface at 52 cm. After that position, the front temperature dropped dramatically and front movement slowed significantly.
From the temperature profiles versus time in Fig. 3.2, it is observed that between 0 and 60 cm, the temperature of the combustion front remains around 400-450°C. After passing the additive/no additive interface, the next peak temperature measured was only 270 °C and declined thereafter. It is concluded that combustion efficiency declined at 370 mins and proceeded to die out. From Fig. 3.3, it is observed that at 370 mins the corresponding CO_2 and CO concentrations dropped as expected. At about 450 min, oxygen approached 20%, indicating thorough die out of this run. It is observed that the combustion performance declines markedly after the front passed the metallic additive interface. Additionally, it does not appear that the additive transported any significant distance once placed within the pack indicating that it is immobile and bound to some portion of the solid packing material.

The facilities were disassembled and checked after combustion. The top 45 cm of the sand pack is quite clean, without obvious oil residue. From 45 cm on, some coke formation is observed. That is unburned fuel. Large coke lumps were observed at about 60 cm, Further along the sandpack was still highly saturated with oil consistent with combustion ceasing.

Another combustion tube run was carried out with similar experimental conditions except that metallic additive was placed across the whole tube at 0.5% weight. Also, the sand used in this run was somewhat more coarse and uniform. Table 3.2 documents packing for this run. Nitrogen was injected at 3 SLPM. The heater was turned off when temperature rose to 400 °C and injection was shifted to air at 3 SLPM.

The combustion was stopped at about 6.5 hours after ignition. Very good combustion was observed throughout the experiment. The temperature at the bottom of the tube exceeded 100 °C after 6.5 hours and so the experiment was quenched to prevent flange damage.

The front position and temperature versus time are plotted in Fig. 3.4. The front moves at a constant velocity of about 0.14 cm/sec, almost without oscillation. The temperature profile is plotted in Fig. 3.5. The front temperature remained at about 400 °C, although at the beginning the temperature was as high as 500 °C due to combustion of the linseed oil and the igniter. In the last temperature profile measurement, the front temperature was as high as 428 °C. Oil deposited in the bottom of the tube due to gravity drainage and liquid flood increasing the fuel for combustion and thereby elevating temperature. Figure 3.6 shows the effluent gas concentration versus time. Although the oxygen concentration was around 8% during most of the experiment, at the beginning and the end of the experiment,

oxygen was as low as 3%. Correspondingly, CO_2 is high initially and late in the experiment and kept constant at 7% during most of the run. In summary, combustion is relatively constant and efficient throughout the experiment.
The experimental facilities were disassembled and checked after combustion. The sand was divided into four sequences. From the top to 45 cm is very clean sand without obvious oil residue. It shows excellent combustion. From 45 cm on, some coke is observed. Large coke residues are observed at about 70 cm consistent with shutdown of the experiment.

3.3 Results–Kinetics Experiments

In addition to the Cymric light crude oil, kinetics runs were carried on Cymric heavy oil (12 °API). This provided a contrast in catalytic effect between heavy and light oil. For these tests, 42 g of sand, 3 g of silica powder or kaolinite, 4 cm^3 of water, and 4 cm^3 of oil are mixed together. If the test includes metallic additive, 0.5 g are added to the water. Most of the kinetic runs were duplicated to check repeatability and sensitivity. The sensitivity analysis teaches that the results are more sensitive to the measurement of O_2 than those of CO and CO_2. The average sensitivity of O_2 measurement is 5%, while that of CO and CO_2 are all under 2%.

Cymric light crude oil. Four runs were performed using Cymric light crude oil. Two runs used kaolinite with Fe^{3+} and two used no additive. The effluent gas curves are plotted in Figs. 3.7 and 3.8. Another two runs were made with silica powder. It has significant, but chemically inert, surface area. Effluent gas curves are plotted in Figs. 3.9 and 3.10. Kinetics cell data is interpreted using accepted procedures (Mamora, 1995; Fassihi et al 1984; Fassahi and Brigham, 1984; Shallcross et al., 1991). Activation energies for LTO and HTO for the four runs are summarized in Table 3.3.

In the two runs with silica powder substituted for kaolinite, it is observed that oxygen consumption increased by 10% in the presence of Fe^{3+}. Additionally, the activation energies are decreased by 8% and 5% for LTO and HTO, respectively. Another two runs are conducted with kaolinite. The catalytic effect of Fe^{3+} is strikingly enhanced when kaolinite is present. Comparison of Figs. 3.7 and 3.8 shows that the LTO reaction is catalyzed significantly. The peak oxygen consumption during LTO for the run with Fe^{3+} is 5.1%, while that of the run without Fe^{3+} is just 3%, an increase of 67%. The oxygen consumption peak for HTO increased by 35% for the run with $Fe3^{+}$. Compare the pair using kaolinite with the pair using silica powder; it is obvious that the kaolinite greatly enhanced the catalytic effect of the metallic additive. Kaolinte itself has some catalytic effect. By comparison of the run with kaolinite and silica powder, both without inclusion of Fe^{3+}, it is observed that the role of kaolinte is enlarged when of Fe^{3+} is included.

Cymric heavy crude oil. Cymric heavy crude oil runs are conducted in a fashion similar to the light oil. The catalytic effect of Fe^{3+} is not obvious on the system using silica powder in place of kaolinite as found by comparing the oxygen consumption, in Figs. 3.11 and 3.12. In the runs with kaolinite, the catalytic effect of Fe^{3+} is encouraging, especially modification of the performance of the HTO reaction.

Oxygen consumption increased by 51% from 12.2% for the run without Fe^{3+} to 18.4% for the run with Fe^{3+}. A similar comparison for LTO reveals an increase of only 8.7%, refer to Figs. 3.13 and 3.14. The same trend is observed with respect to activation energy, Table 3.4. Apparently, the light oil catalytic affects are mainly associated with LTO, while that of heavy oil affected the HTO.

3.4 Results–SEM Imaging

Imaging and quantitative microscopic analysis of the various mixtures of sand, clay, silica powder powder, and metallic additive was undertaken to further our understanding of the mechanism of enhancement to combustion. It is our working hypothesis that the water-soluble cationic metals ion exchange with ions residing in the clay to create an activated site that acts in a manner analogous to a heterogeneous catalyst. Hence, the clay structure with additive and without additive was probed to gauge any gross morphological changes in clay structure and to verify that the metal catalyst had ion exchanged into the clay matrix.

The metallic salts tin chloride ($SnCl_2$) and iron nitrate ($Fe(NO_3)_3$) were employed for this phase of the investigation. The mixes are prepared as discussed for the kinetics cell tests above and dried under house vacuum. After SEM mount preparation, all samples are sputtered with a gold-palladium mixture to improve surface conductivity thereby enhancing image quality.

This phase of the investigation is still underway and is subject to further interpretation. The following observations have, nonetheless, emerged. First, the introduction of metallic catalyst ions has not changed remarkably the morphology of the clay within mixtures of sand and clay. Figure 3.15 contains a representative result that contrasts the clay structure in the absence, Fig. 3.15(a), and presence, Fig. 3.15(b), of the ion Fe^{3+}. Images are at 1000x magnification and note the horizontal bar illustrating a dimension of 10 μm. Figure 3.16 presents views of the same systems at 4000x. Now, note the horizontal bar indicating a size of 5 μm. Figure 3.16(b) may illustrate a greater fraction of clay agglomerations that are roughly 1 μm is size. Nevertheless, at these two scales, clay morphology is remarkably similar in the presence or absence of aqueous metallic additives.

The SEM imaging also provided information about the location of metal ions on the solid substrate. At first glance, attempting to determine the location of metal ions on the solid substrate is akin to searching for a needle in a haystack. Various probes available for the SEM have aided in the search. In particular, back-scattered electron mode produces images that are proportional to atomic number. Figure 3.17 illustrates a localized deposit of iron (Fe^{3+}) on what appears to be a quartz sand grain at 1400x. The inset shows the iron deposit at 6000x. The identification of iron was verified by elemental analysis (EDS). The remainder of material in this view is the silica powder. Figure 3.18 presents results from the system sand-kaolinite-tin chloride at 6000x. The composition of the region in the SEM image marked by a square is given in the spectrum below the image. The size, shape, and elemental composition (Al, Si, O) of the solid material in the figure indicate that is clay. The remarkable aspect is the relatively high abundance of tin (Sn) within the square. Consistent with the hypothesis, tin has participated in cation exchange and is

emplaced within the clay. Similar results, not shown, are obtained for the mixture sand-kaolinite-iron nitrate.

3.5 Discussion

Metallic additives affect combustion performance significantly. Here, fuel deposition is a decisive factor altered by the additives. If oil is too light, not enough fuel is deposited to sustain the combustion. Whereas if the oil is heavy, excessive fuel deposition leads to a larger requirement for air and significant consumption of the oil. For Cymric light crude oil, the catalytic effect is obvious in the LTO reaction where fuel deposition is increased to sustain combustion. During LTO, partially oxygenated compounds are produced, such as carboxylic acids, aldehydes, ketones, and alcohols (Brigham and Castanier, 2008). Clays generally have a great affinity for such polar, oxygenated compounds. The organic molecules likely wedge directly into the clay interlayer structure. Removal of the partially oxygenated products from the reaction disrupts the equilibrium of the LTO reaction and shifts the oxidation reaction towards the products. In this way, clay with metallic additive has a positive effect on the LTO reaction.

Based on the work above, and former research (Shallcross et al, 1991, Castanier et al., 1992) employing metallic additives and Huntington Beach oil, a preliminary reactivity series is obtained by a series of standards. These standard include the increase in the extent of oxidation reactions, oxygen utilization efficiency, peak temperature and temperature profile, and heat of combustion. The series reads:

$$Sn^{2+} > Fe^{3+} > Al^{3+} > Zn^{2+} > Mg^{2+} > Cu^{2+} > Cd^{2+}$$

This ranking is developed from experiments with specific crude oils and clays. Further work is needed to verify that the series is general.

An accepted method to compare the relative ease of replacement of one cation by another on clay surfaces is to compute the ratio of the cation valance upon the hydrated radius of the ion (e..g, Grim, 1968). Thus, both electrical charge as well as the hydrated ion size are taken into account. The greater the ratio, the greater is the charge density. Consequentlly, cations are more tightly held on negatively charged portions of the clay. Table 3.5 summarizes valence, ionic radius, and hydrated ionic radius for the reactivity series above. No values of hydrated radius for Sn^{2+} could be found in the literature. The remaining cations follow the trend of decreasing charge to hydrated radius ratio with decreasing combustion reactivity. Apparently, the more tightly bound the cation, the greater is its beneficial effect on combustion.

Clay, even without additive, also has an effect on crude-oil combustion (Burger et al., 1985). One effect is via physical perspective, and the other is chemical. Physically, clay has large external and internal surface area that improves the deposition of oil on the sand matrix and constitutes a strong carrier for various reactions in different stages of the combustion. Kaolinite, and other clays, have found use in petroleum refining as catalysts in the cracking of large organic

molecules in crude oil. In fact, the low-temperature reactions are more biased to include kaolinite as catalyst (Burger et al., 1985). Similarly, the sand matrix affects combustion performance via surface area. Fine sand with a wide grain-size distribution appears to develop significant oil deposition that provides substantial fuel for combustion.

3.6 Summary

A general point from this study is that the conventional method for the interpretation of ramped-temperature oxidation studies does not give a full indication of the role of metallic additives in enhancing combustion. For this particular study, the following specific conclusions arise.

1. Iron, introduced as water soluble Fe^{3+}, enhanced fuel deposition for 34 °API crude oil from Cymric. Whereas for 12 °API crude oil from Cymric, the addition of Fe^{3+} enhanced the HTO reactions producing more complete combustion.
2. There was no evidence of the transport of Fe^{3+} during combustion tube runs even though water is clearly mobile during these experiments.
3. Analytical microscopic analysis via SEM teaches that the Fe^{3+} partitions from the aqueous phase onto the solid substrate (sand and clay) used for packing the combustion tube and kinetics cell. The SEM analysis indicates that the water soluble metallic additives participate in cation exchange and come to reside within the kaolinite.
4. The data collected is consistent with a mechanism whereby water soluble metallic salt ions exchange into the clay matrix. This creates activated sites that enhance the kinetics of combustion favorably.

Table 3.1 Cumulative distribution of sand grain sizes.

meshsize	cumu %
20	0.267016
45	0.354857
60	0.807155
80	0.911867
100	0.948517
120	0.982257
140	0.984584
170	0.994183
200	0.997673
220	1

Table 3.2 Description of combustion tube runs.

length (cm)	oil (cm^3)	water (cm^3)	sand (cm^3)	kaoli-nite (g)	$Fe(NO_3)_3$ (g)
run 1: first half of tube contains additive					
48	350	150	1800	140	9
47	350	150	1800	140	0
run 2: entire tube contains additive					
96	740	320	3850	300	19

Table 3.3 Kinetics results for Cymric light crude oil.

kaolinite or silica	additive	Peak T, LTO, °C	Peak T, HTO, °C	E/R, LTO	E/R, HTO	Max. O_2 LTO, %	Max. O_2 HTO, %
kaolinite	none	275	385	9647	10669	3	2.8
kaolinite	Fe^{3+}	275	355	8535	11522	5.1	3.8
silica powder	none	287	410	8331	10681	2.4	2.6
silica powder	Fe^{3+}	280	370	7886	9775	2.9	2.7

Table 3.4. Kinetic run and analysis results for Cymric heavy crude oil.

kaolinite or silica	additive	Peak T, LTO, °C	Peak T, HTO, °C	E/R, LTO	E/R, HTO	Max. O_2 LTO, %	Max. O_2 HTO, %
kaolinite	none	270	395	9090	12319	12.6	12.2
kaolinite	Fe^{3+}	280	355	9085	9427	13.7	18.7
silica powder	none	265	410	7293	10489	6.3	18.4
silica powder	Fe^{3+}	260	380	7098	10640	7	18.4

Table 3.5. Summary of cation radii and charge density. Hydrated ionic radii from Nightingale (1959) and Baes and Mesmer (1977).

	valence	ionic radius[25] (Å)	hydrated ionic radius (Å)	CHIRR* ($Å^{-1}$)
Sn^{2+}	2	1.19	-	-
Fe^{3+}	3	0.55	4.28	0.701
Al^{3+}	3	0.54	4.75	0.632
Zn^{2+}	2	0.74	4.3	0.465
Mg^{2+}	2	0.72	4.28	0.467
Cu^{2+}	2	0.73	4.19	0.477
Cd^{2+}	2	0.95	4.26	0.469

* CHIRR = valence: hydrated ionic radius

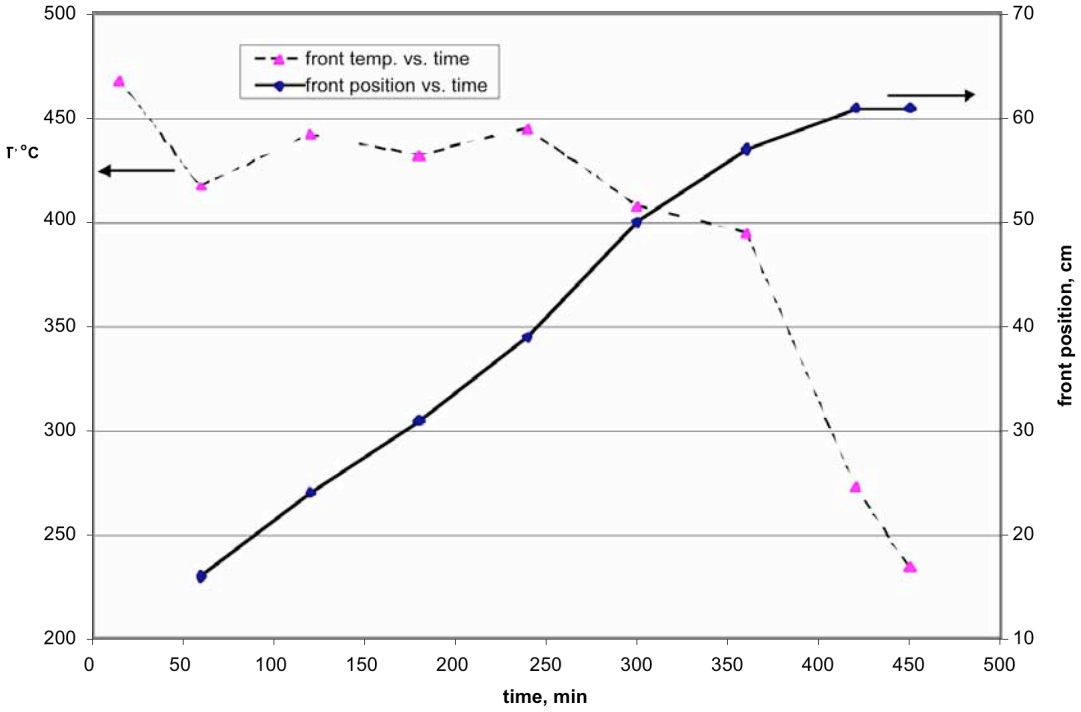

Figure 3.1. Front position and temperature versus time, first tube run.

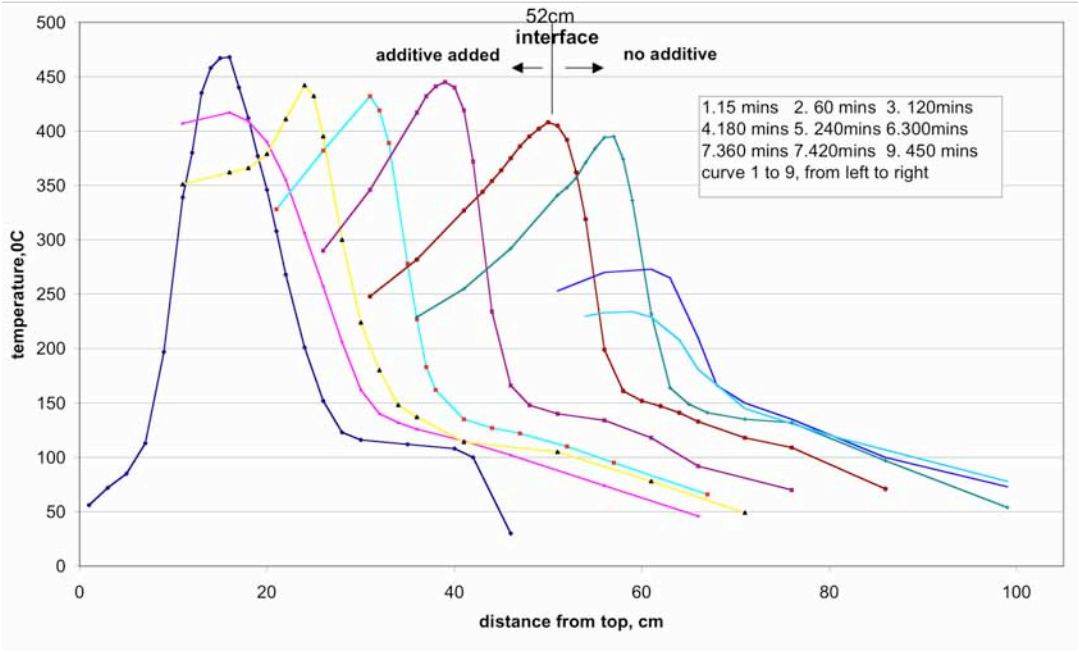

Figure 3.2. Temperature profiles versus distance, first tube run.

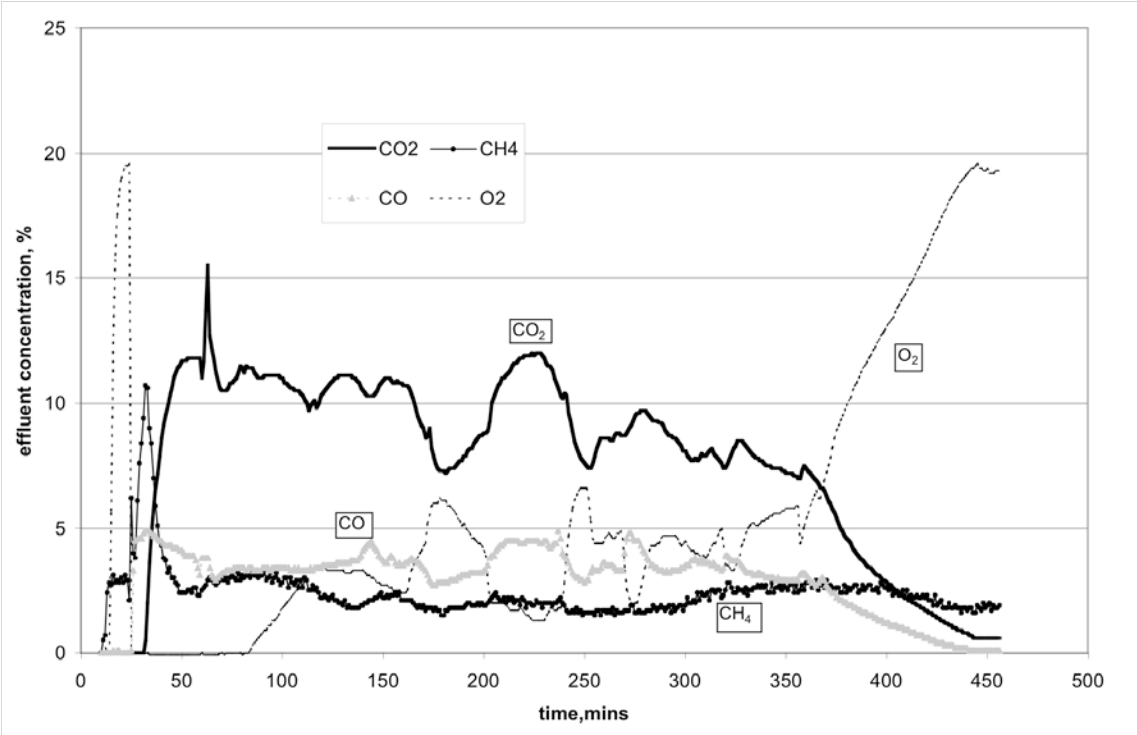

Figure 3.3. Effluent gas concentrations versus time, first tube run.

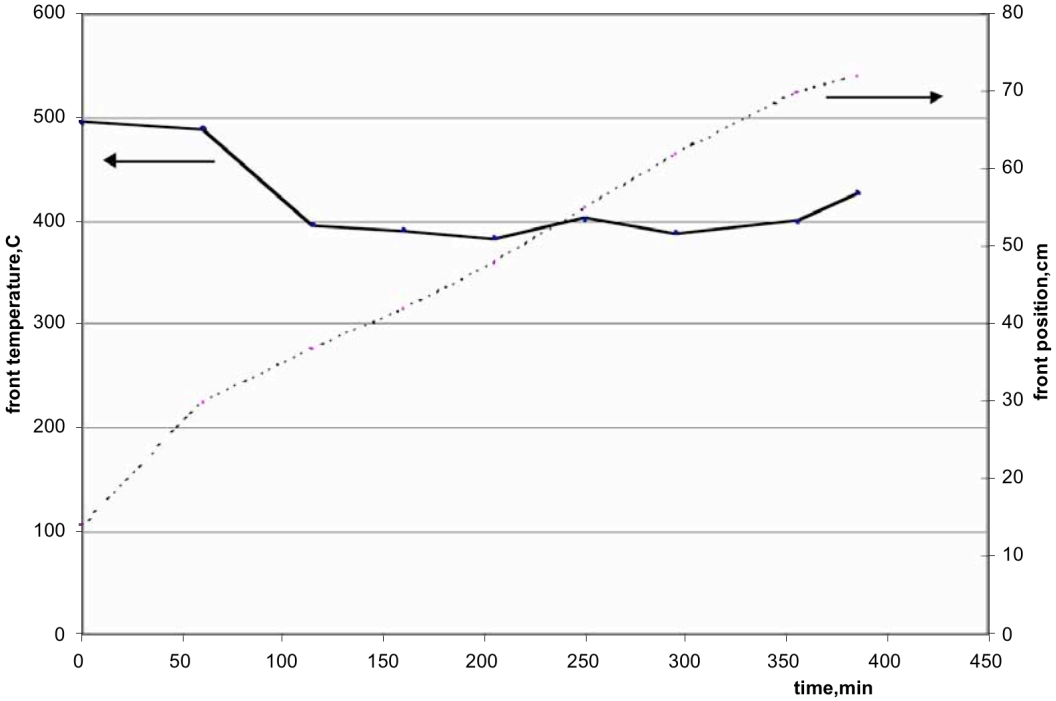

Figure 3.4. Front position and temperature versus time, second tube run.

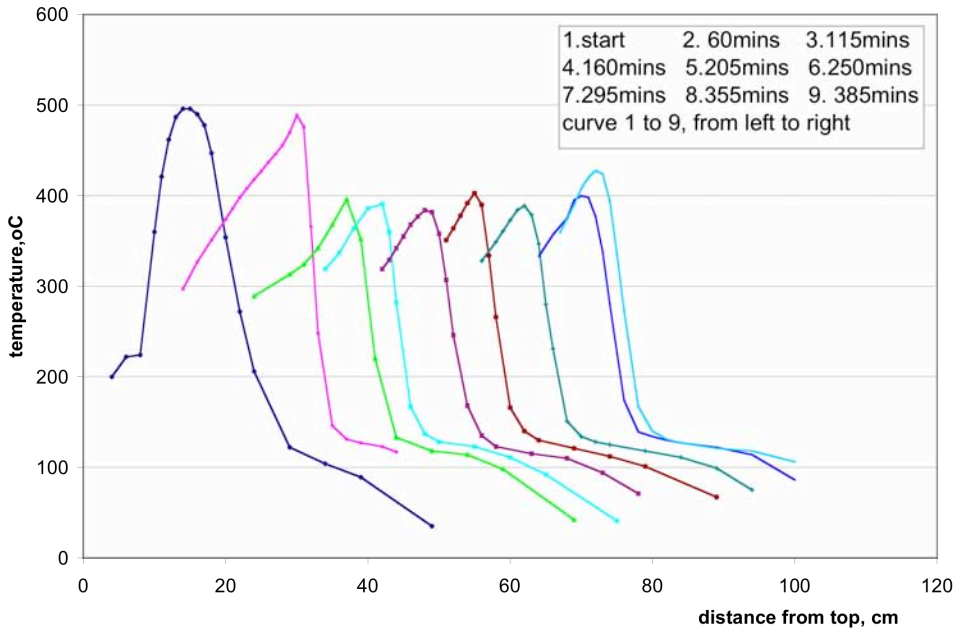

Fig. 3.5 Temperature profiles versus distance, second tube run.

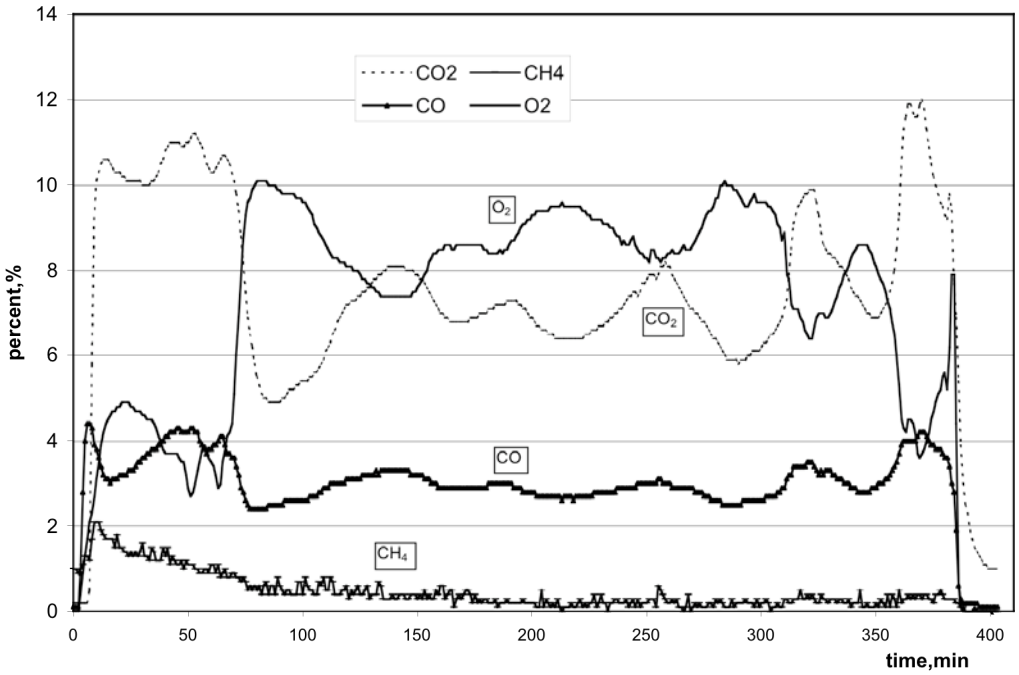

Figure 3.6. Effluent gas concentrations versus time, second tube run.

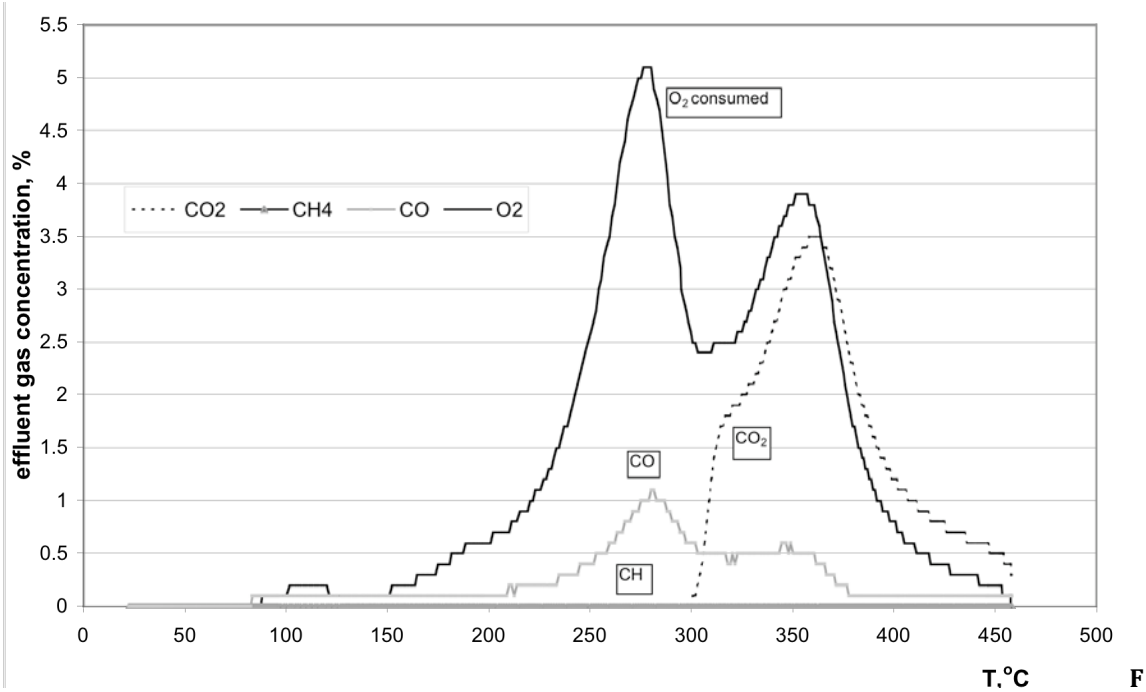

igure 3.7. Kinetic run of Cymric light crude oil with kaolinte and Fe^{3+}.

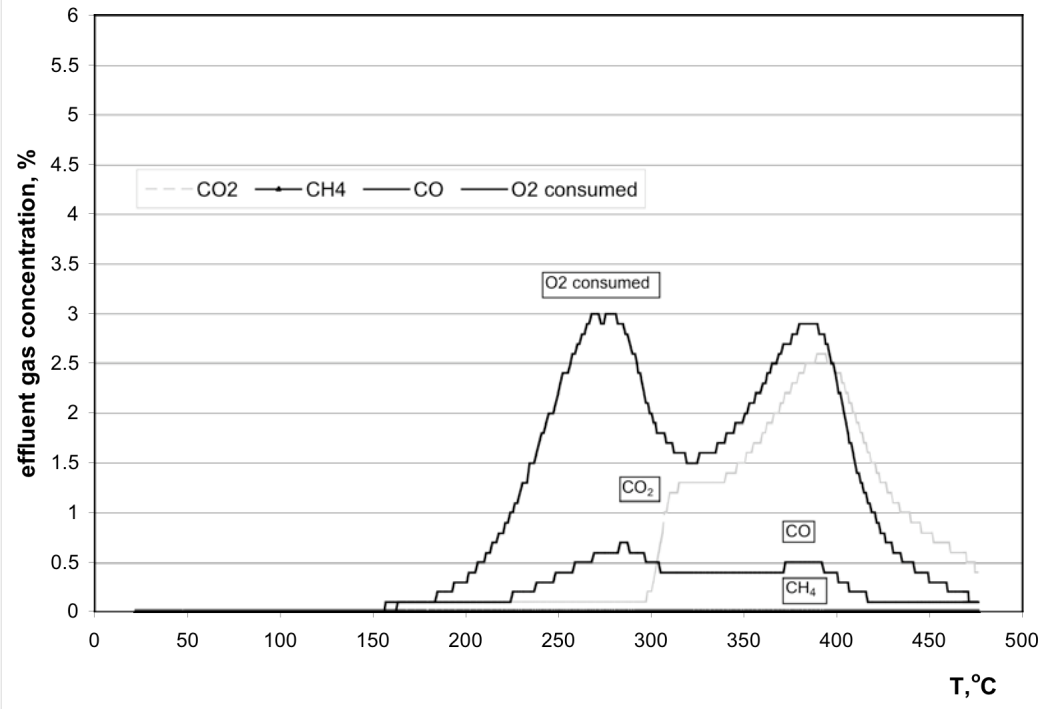

Figure 3.8. Kinetic run of Cymric light crude oil with kaolinte and no Fe^{3+}.

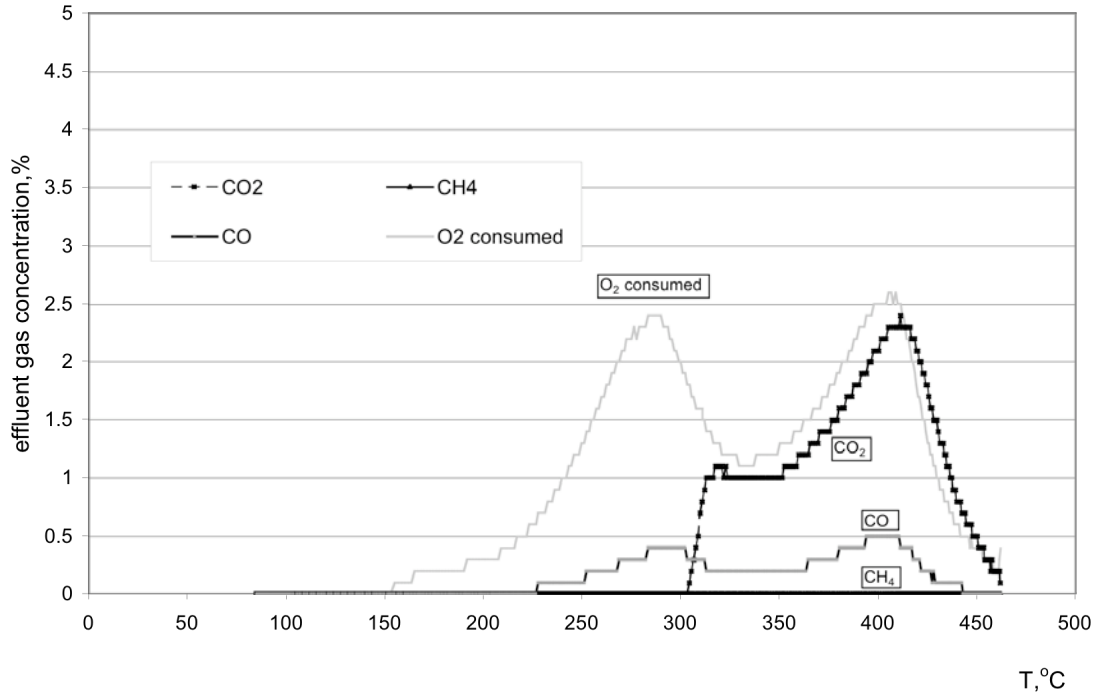

Figure 3.9. Kinetic run with Cymric light crude oil with silica and no Fe^{3+}.

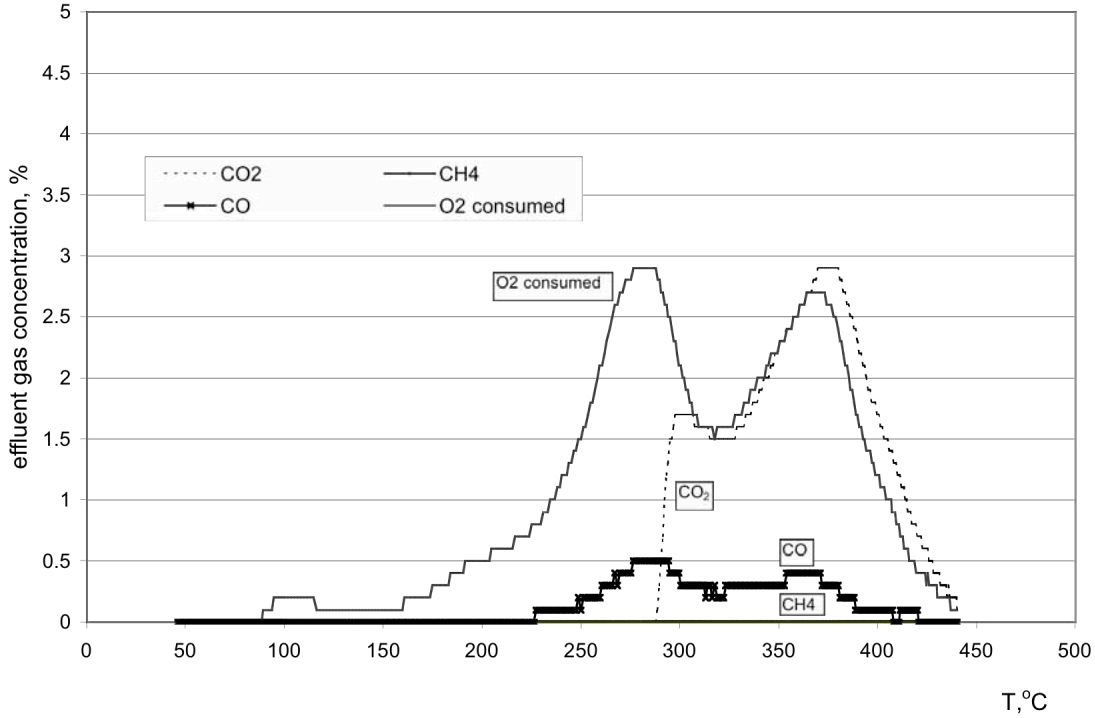

Figure 3.10. Kinetic run of Cymric light crude oil with sillica and Fe^{3+}.

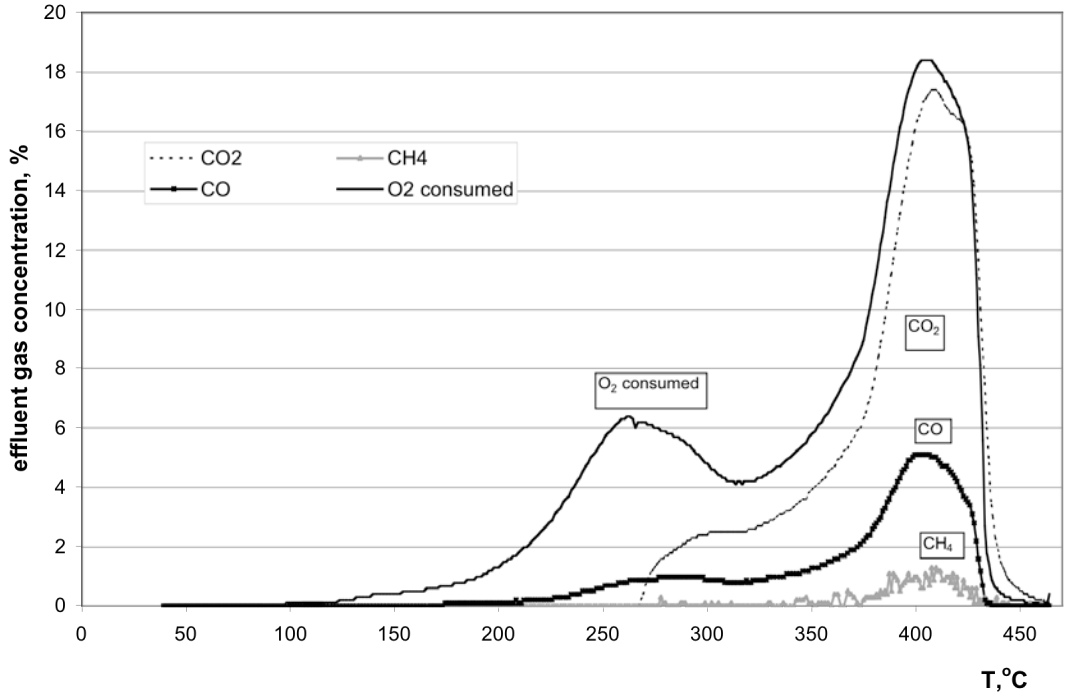

Figure 3.11 Kinetic run of Cymric heavy crude oil with silica and no Fe^{3+}.

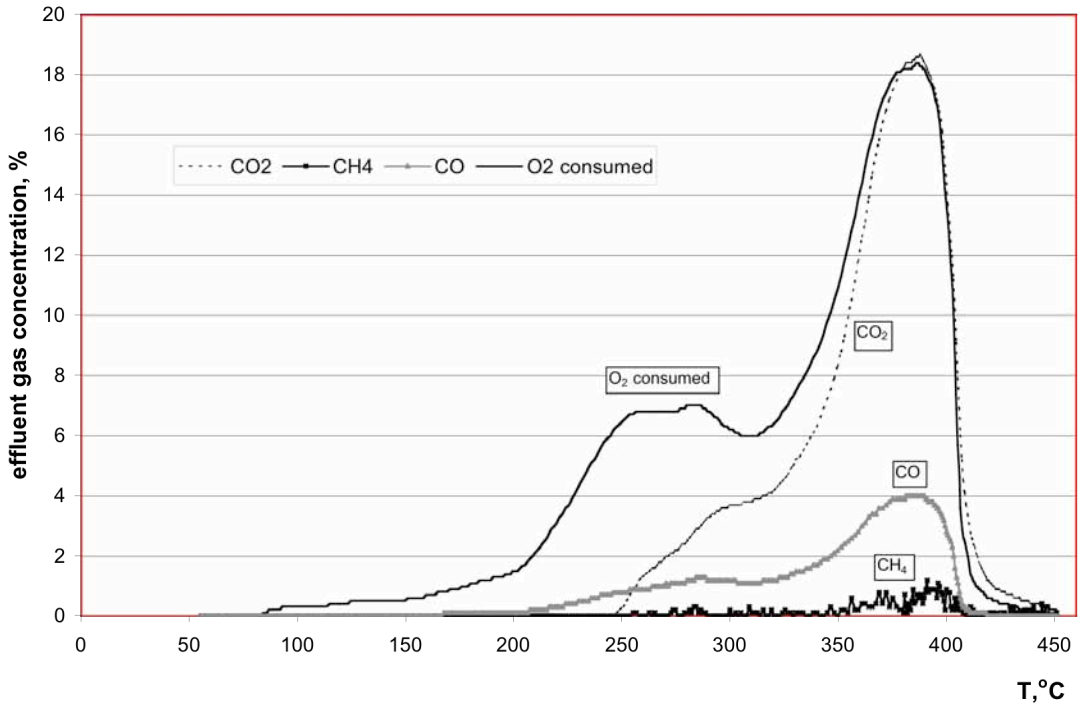

Figure 3.12 Kinetic run of Cymric heavy crude oil with silica and Fe^{3+}.

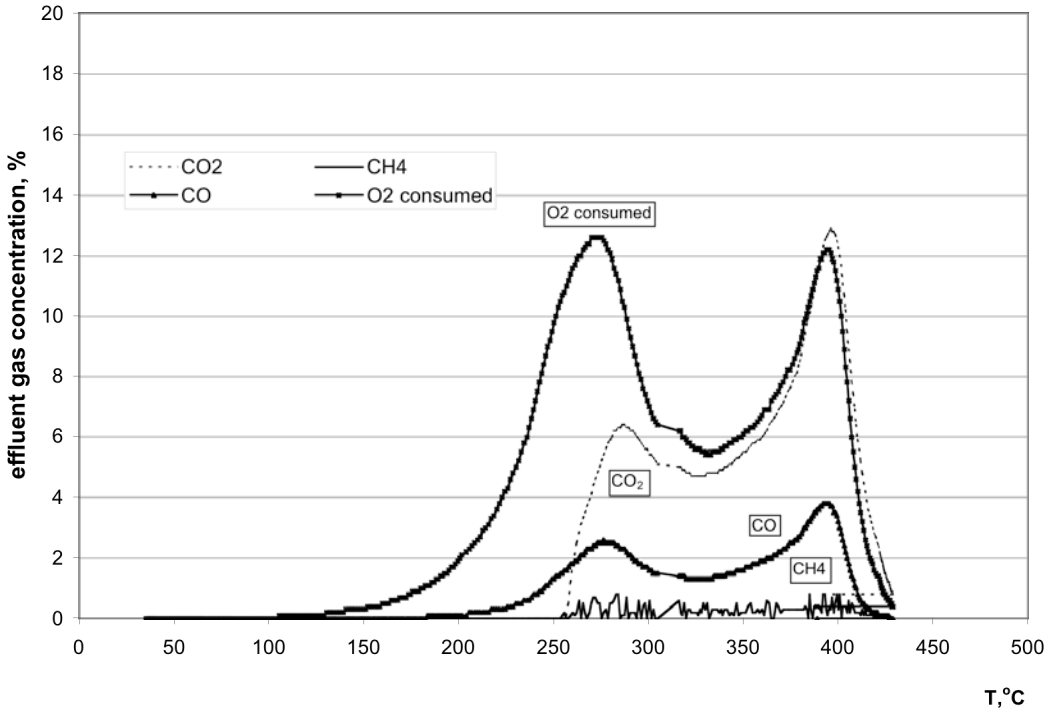

Figure 3.13. Kinetic run of Cymric heavy crude oil with kaolinte and no Fe^{3+}.

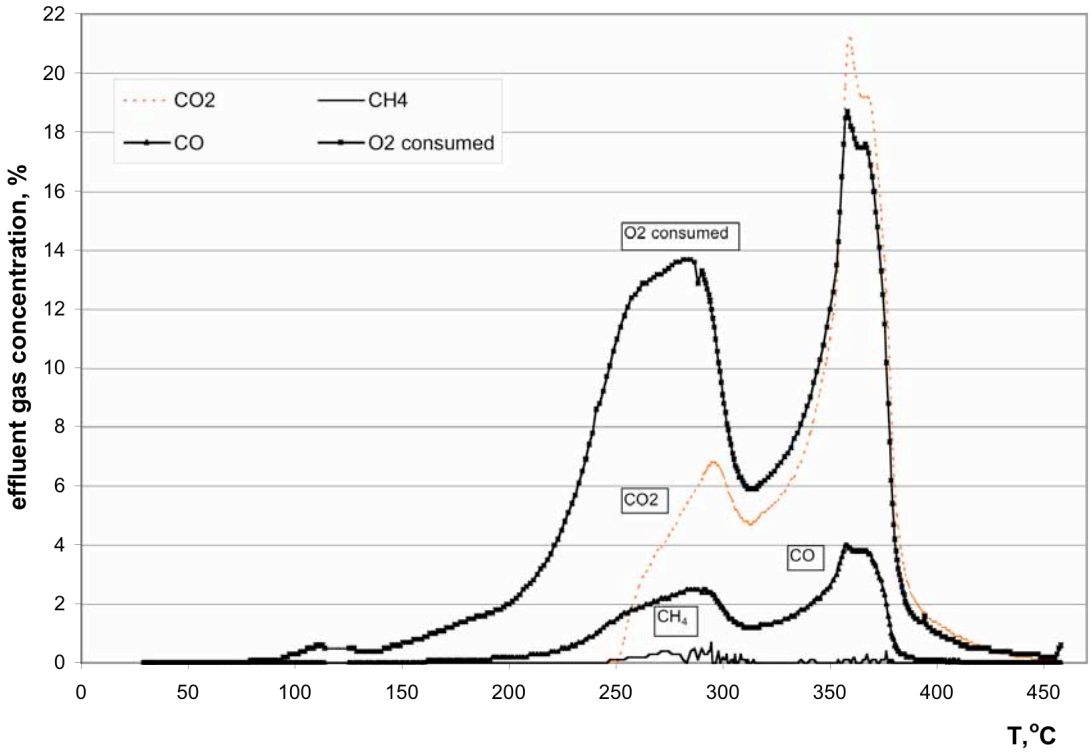

Figure 3.14. Kinetic run of Cymric heavy crude oil with kaolinite and Fe^{3+}.

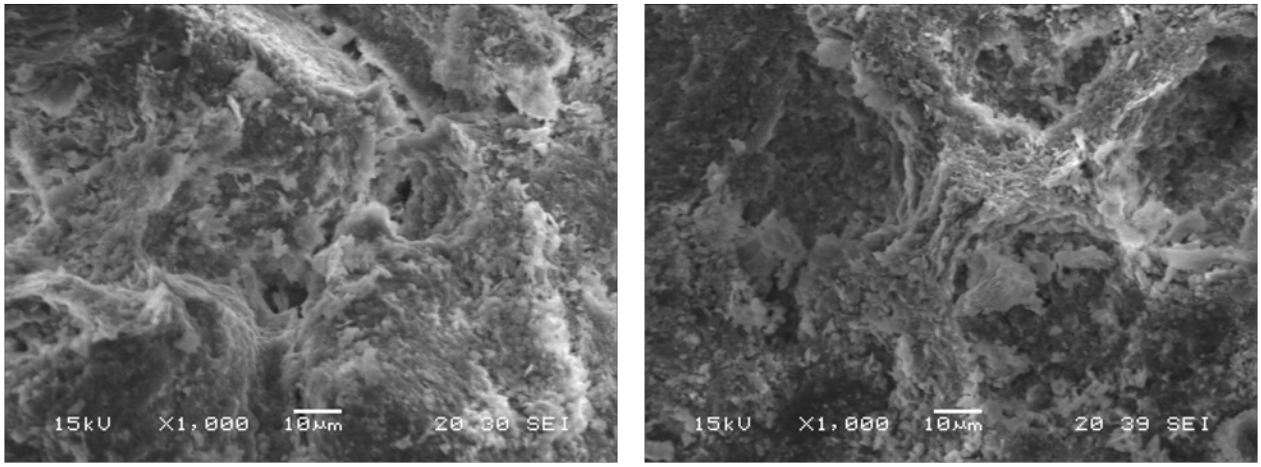

(a) (b)

Figure 3.15. (a) Mixture of sand and kaolinite and (b) mixture of sand, kaolinite, and iron (Fe^{3+}), 1000x.

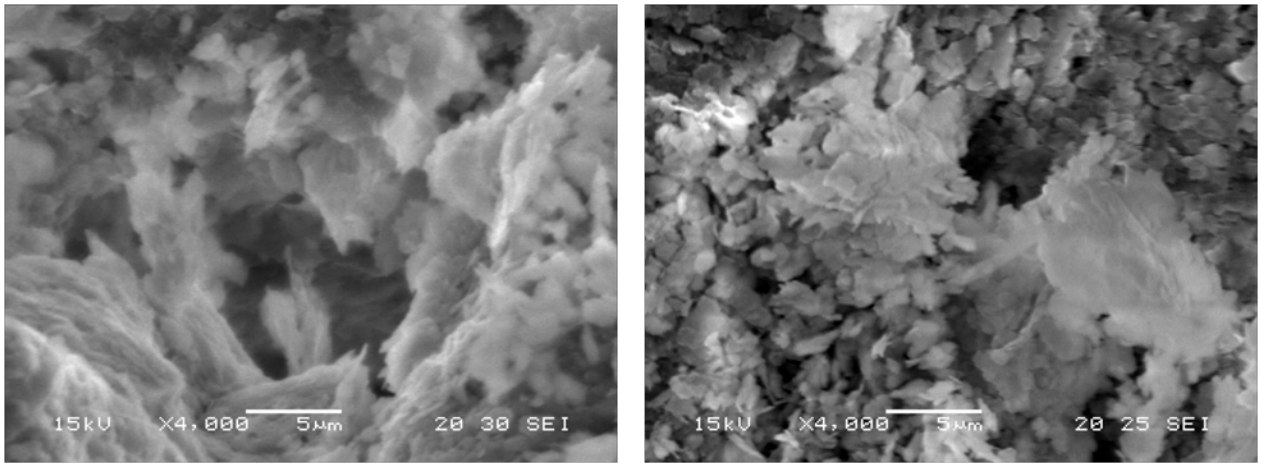

(a) (b)

Figure 3.16. (a) Mixture of sand and kaolinite and (b) mixture of sand, kaolinite, and iron (Fe^{3+}), 4000x.

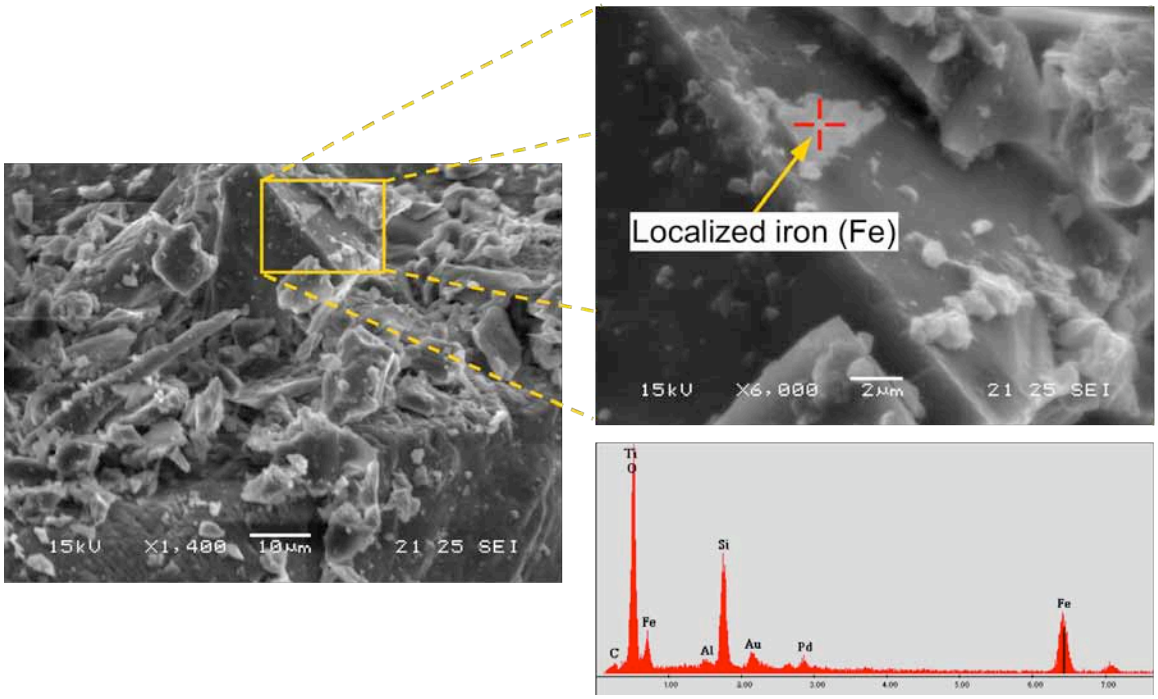

Fig. 3.17. Mixture of sand, amorphous silica powder, and iron (Fe^{3+}).EDS spectrum is from cross on upper right image. Left image is at 1400x and right image is at 6000x.

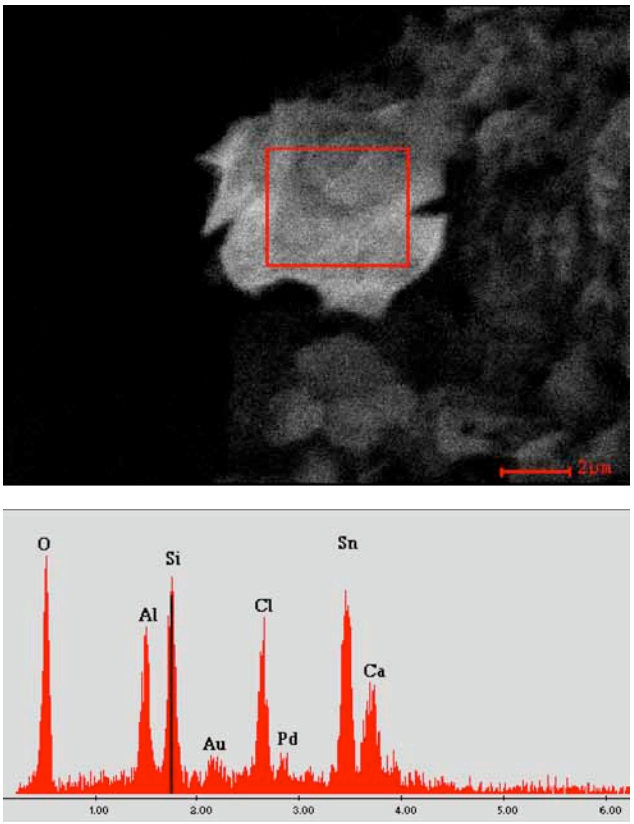

Figure 3.18. Localized tin in mixture of sand, kaolinite, and tin (Sn^{2+}) at 6000x. EDS spectrum is from square.

4. Improved Analysis of the Kinetics of Crude-Oil In-Situ Combustion

One factor that has limited application of ISC is lack of predictability at both the laboratory and field scale. Proper experimental kinetic analysis (using, for instance, ramped temperature oxidation) provides critical parameters, such as activation energy, for modeling efforts. Conventional methods of kinetic analysis for crude oil assume, a priori, a two-step or three-step oxidation reaction. Any inconsistency between the assumed model and the real system is reflected as error in model parameters. We have explored iso-conversional techniques that provide model-free methods for estimating activation energy and naturally deconvolve multi-step reactions. The iso-conversional approach obtains the reaction kinetics for a given rock/oil sample at identical reaction extent from multiple experiments with different heating rates. The technique has never been applied to ISC kinetics, to our knowledge. Our work in this study falls into two major sections. First, the applicability of the iso-conversional method to multiple step oxidation of oil in porous media is established via synthetic examples. Second, and on the practical side, the combustion kinetics are reported for two quite different crude oil samples thereby adding to the knowledge base of crude-oil and rock-matrix characteristics that make for successful ISC. The iso-conversional analysis of ramped temperature oxidation data appears to provide new insight into systems where ISC is/is not successful.

4.1 Introduction

Modeling of complex in-situ combustion (ISC) processes requires an understanding the behavior of different physical phenomena such as phase change, chemical reactivity, heat and mass transfer. In this study, the focus is the kinetics of in-situ combustion, their improved measurement and accurate representation of crude oil reactivity. Reactions during combustion of oil in porous media are generally grouped into three functional groups based upon the relative temperature at which reactions occur: low temperature oxidation (LTO), fuel deposition (MTO) and fuel combustion (HTO) (Fassihi, 1981).

As summarized by Sarathi (1999) the study of in-situ combustion kinetics is undertaken for the following reasons:

- to characterize the reactivity of the oil,
- to determine the conditions required to achieve ignition and or to determine if self ignition will take place in the reservoir upon air injection,
- to gain insight into the nature of fuel formed and its impact on combustion, and
- to establish parameter values for the kinetic (reaction rate) models used in the numerical simulation of ISC processes.

Crude oil is a complex material formed by hundreds of components. In order to model the oxidation reactions that occur in a combustion process an extended

compositional analysis and a large number of kinetic expressions are required. It is not easy, however to describe all of the reactions associated with oxidation of all components. Only for the simplest hydrocarbon molecules (e.g. CH_4) detailed models are available (Law, 2006). Even if such models were available, it would be inefficient to implement them in computer simulators as rigorous models would require many computational steps. Instead simple models based on the pseudo-unimolecular assumption appear in the literature. A detailed summary of the literature on kinetics of oxidation reactions is given by Fassihi (1981) and Sarathi (1999).

In this study, we begin with the basic concepts of the kinetics of combustion in porous media. Then we introduce various isoconversional methods with their mathematical background. After, we illustrate a synthetic example and test the applicability of iso-conversional methods to oil combustion in porous media. Later, we introduce experimental procedures and experimental results. In conclusion, we discuss the results and further work.

4.2 Kinetics of combustion in porous media

Reactions during oil combustion involve a complex set of sequential and parallel reactions that are extremely difficult to characterize. Earlier attempts of characterizing kinetics of oil combustion described reactions in two main categories: low temperature oxidation (LTO) and high temperature oxidation (HTO) (Tadema, 1959). Later studies proved an existence of a third group, so called middle temperature oxidation (Fassihi, 1981).

Conventional analysis of oil combustion uses lumped species such as fuel and describes the kinetics using a pseudo-unimolecular reaction. In the petroleum engineering literature, the rate of reaction is usually expressed in the following form,

$$-\frac{dC_c}{dt} = Ae^{-\frac{E}{RT}} C_c^a C_{O_2}^b \tag{4.1}$$

where a and b represents exponents for carbon and oxygen concentration, respectively, E is the activation energy (J/mol), T is absolute temperature (K), R is the gas constant (J/gmol K) and A is a pre-exponential factor (Arrhenius constant, 1/s). In this case, we use any measured equivalent to concentration. The only difference is the pre-exponential factor, all other variables are unchanged. Partial pressure is such an example. Replacing concentration of oxygen with partial pressure we have,

$$-\frac{dC_c}{dt} = A'e^{-\frac{E}{RT}} C_c^a p_{O_2}^b \tag{4.2}$$

Similar forms are applied by various authors (Bousaid and Ramey, 1968; Burger and Sahuqet, 1972; Fassihi et al, 1984). Eq. (4.2) is valid under the assumptions of homogenous, elementary reactions. Validity of Eq. (4.2) is assumed in this study.

In kinetic analysis, there are three unknowns that are of particular interest: activation energy, pre-exponential factor and reaction model. Proper experimental kinetic analysis provides these parameters along with error estimates for modeling efforts.

Bousaid and Ramey (1968) used a method based on Eq. (4.2) for the analysis of isothermal data. They assumed that the reaction order with respect to oxygen is unity. If we take the logarithm of Eq. (4.2) we have,

$$\log\left(\frac{-\frac{dC}{dt}}{p_{O_2}}\right) = \log(k) + a\log C_c \tag{4.3}$$

A plot of the left hand side Eq. (4.3) versus carbon concentration, yields a line with slope a (the reaction order with respect to carbon concentration) and the intercept gives the logarithm of the rate constant. The logarithm of the Arrhenius Equation for rate constant yields,

$$\log(k) = \log(A') - \frac{E}{RT} \tag{4.4}$$

From experiments at different constant temperatures, rate constant values are estimated and plotted against one over temperature. The slope gives the activation energy and the intercept gives the logarithm of the Arrhenius constant.

Although the method of Bousaid and Ramey (1968) is straight forward, it requires isothermal conditions and assumes first order reaction with respect to oxygen. Isothermal conditions are not easily achieved in kinetic experiments as the reactions occurring in combustion of oil are mostly exothermic. Also multiple reactions take place in combustion at different temperatures. It is more convenient to analyze the data under non-isothermal conditions.

Later Fassihi (1981) proposed a method for the analysis of non-isothermal oxidation. In his analysis, the relation between oxygen consumption rate and carbon consumption is used to derive the change in oxygen concentration. That is, a quantity we measure directly. This method requires that the order of reaction with respect to carbon concentration be known a priori. A trial and error procedure is needed to estimate the reaction orders.

In his study on fuel formation, Abu-Khamsim (1984) used non-linear regression techniques to refine the kinetic parameter estimates. He suggested a recursive process in which at each cycle regression is repeated in order to find the appropriate model and associated parameters.

Later, In 1993 Mamora proposed a new oxidation reaction model that includes two partially overlapping reactions and assesses the effects of surface area (Mamora, 1993) This method needs the kinetic model to be known or assumed a priori.

In the petroleum engineering literature, the common way to calculate kinetic parameters is based on Eq. (4.2) and lumped reactions/species. Pre-exponential

factors, activation energies and, exponents a and b are calculated for different groups of reactions by fitting appropriate portions of data with a model. Different approaches have different assumptions but all of them merge at the elementary reaction assumption. The question remains regarding how valid are the assumptions of the various methods.

Conventional model fitting or straight line fitting methods require the kinetic model to be known or assumed a priori. By fitting a model to the experimental data, activation energy and pre-exponential factor are estimated. In fact, reaction models take various forms. Any inconsistency between the assumed model and the real system is reflected as errors in model parameters as well as model predictions. In short, all the conventional analysis methods make assumptions about the reaction model. At least a form of a reaction model is assumed. On the other hand, iso-conversional methods provide reaction model-free methods to estimate activation energy (Vyazovkin, 1997). Extensive reviews of the applicability and usefulness are given in references (Vyazovkin and Wight, 1999; Burnham and Dinh, 2007). These methods require a series of experiments to be conducted at different heating rates on an identical sample mixture of a crude oil, water and solid matrix.

4.3 Iso-Conversional methods

Experiments indicate that the rate of reaction is a function of both the concentration and the temperature of the material (Levenspiel, 1999). Thus, a single-step kinetic equation for a homogenous reaction is written,

$$-\frac{dC}{dt} = \vartheta(T,C) = k(T)f(C) \tag{4.5}$$

where *k(T)* is the rate constant (units depends on reaction model) and *f(C)* is the reaction model. The rate constant is expressed as a function of temperature given by the Eq. (4.4) Combining Eq. (4.5) with the antilog of Eq. (4.4) we have,

$$-\frac{dC_c}{dt} = Ae^{-\frac{E}{RT}}f(C) \tag{4.6}$$

In terms of fractional conversion, α, Eq. (4.6) is represented as,

$$\frac{d\alpha}{dt} = Ae^{-\frac{E}{RT}}f(\alpha) \tag{4.7}$$

The iso-conversional principle states that at constant extent of conversion, the reaction rate is only a function of the temperature. Eq. (4.7) is the basis of all iso-conversional methods. Taking the logarithm of Eq. (4.7) gives,

$$\ln\left(\frac{d\alpha}{dt}\right) = \ln(A) + \ln\left[f(a)\right] - \frac{E}{RT} \tag{4.8}$$

At constant values of conversion *f(α)* is assumed to be constant. This is analogous to assuming that the chemistry of the process is independent of temperature and dependent only on the level of conversion [14]. Thus, the reaction model is independent of heating rate. As a consequence, for different temperatures at the same conversion levels or iso-conversional values, *f(α)* is equal. Then,

$$\ln\left(\frac{d\alpha}{dt}\right) = m - \frac{E}{RT_\alpha} \tag{4.9}$$

where

$$m = \ln(A) + \ln\left[f(a)\right] \tag{4.10}$$

Here t_α and T_α are time and temperature, respectively, for any particular experiment at conversion α. The symbol E_α is the activation energy at conversion α. For different temperatures, we plot the left hand side (LHS) of Eq. (4.10) with respect to $-1/T_\alpha$. The slope of the graph gives E/R and the intercept is m. This method is first introduced by Friedman (1964) and is referred to in the literature as the method of Friedman or the differential iso-conversional method.

This method needs rate values to be calculated and the rate values are usually calculated by numerical differentiation. It is a known fact that numerical differentiation amplifies error in experimental data. Therefore results are sensitive to experimental noise (Vyazovkin, 1999).

In order to avoid numerical differentiation, integral iso-conversional methods are developed. A linear ramped temperature heating rate is expressed as,

$$T = T_0 + \beta t \tag{4.11}$$

Here, β is heating rate (T/time), T_0 is base temperature (T). By combining Eq. (4.11) with Eq. (4.8) and with some algebraic manipulation, we have,

$$\ln(\beta)_\alpha = \ln\int_x^\infty x^{-2}\exp(-x)dx - \ln\left[\frac{R}{AE}g(\alpha)\right] \tag{4.12}$$

where,

$$x = \frac{E}{RT} \tag{4.13}$$

and

$$g(\alpha) = \int_0^\alpha \frac{d\alpha}{f(\alpha)} \tag{4.14}$$

The integral on the right of Eq. (4.12) has no general analytical solution. For hyperbolic and parabolic heating programs this integral has analytical solutions.

The integral iso-conversional methods differ in their approximation to the solution of this integral.

The Kissenger, Akahira, Sunose (KAS) (Kissenger, 1957; Akahira and Sunose, 1971) method uses the approximation,

$$\int_{x}^{\infty} x^{-2}\exp(-x)dx = x^{-2}\exp(-x) \tag{4.15}$$

This approximation is assumed to be valid for values of x between 20 and 50. Substituting Eq. (4.15) in Eq. (4.12) we have,

$$\ln\left(\frac{\beta}{T}\right)_{\alpha} = -\frac{E}{R}\left(\frac{1}{T_{\alpha}}\right) - \ln\left[\frac{R}{AE}g(\alpha)\right] \tag{4.16}$$

The Ozawa-Flynn-Wall (OFW) (Ozawa, 1965; Flynn and Wall, 1966) method uses another approximation.

$$\int_{x}^{\infty} x^{-2}\exp(-x)dx = \exp(-1.052x - 5.33) \tag{4.17}$$

Approximation is assumed to be valid for values of x between 20 and 60. Substituting Eq. (4.17) in Eq. (4.13) we have,

$$\ln(\beta)_{\alpha} = -1.052\frac{E}{R}\left(\frac{1}{T_{\alpha}}\right) - 5.33 - \ln\left[\frac{R}{AE}g(\alpha)\right] \tag{4.18}$$

The procedure to estimate activation energy from the above methods requires that a series of experiments be conducted with different heating rates. In general, the KAS approximation gives better results compared to the OFW method. The major disadvantage of these integral methods is their approximations.

The heating program of the experiments is described to control the experiment. Note that the KAS and OFW methods require linear heating given by Eq. (4.12). By the virtue of thermal effects of reactions (endothermic or exothermic behavior), the temperature of the system deviates from the prescribed heating program. A computational method that accounts for any variation of temperature is proposed by Vyazovkin(1997).

As shown by Vyazovkin (1997), those deviations of temperature yield erroneous values of activation energy, when such a data set is used with a model that assumes linear heating. Also, all integral iso-conversional techniques have internal errors related with the approximation assumed. In order to overcome those issues, a nonlinear iso-conversional method is proposed by Vyazovkin (1997). Integrating Eq. (4.8)

$$g(\alpha) = \int_{0}^{\alpha}\frac{d\alpha}{f(\alpha)} = A\int_{0}^{t_{\alpha}} e^{\frac{-E}{RT(\alpha)}}dt \tag{4.19}$$

and

$$J\left[E,T(t)\right]=\int_{o}^{t_\alpha} e^{-\frac{E}{RT_\alpha}}dt \tag{4.20}$$

It is assumed that the reaction model is independent of temperature, for a given conversion and experiments conducted with different heating histories. One can then write,

$$A_\alpha J\left[E_\alpha,T_1(t_\alpha)\right]=A_\alpha J\left[E_\alpha,T_2(t_\alpha)\right]=\cdots=A_\alpha J\left[E_\alpha,T_n(t_\alpha)\right] \tag{4.21}$$

A_α values cancel out and we treat this problem as a minimization of the following function.

$$\Phi(E_\alpha)=\sum_{i}^{n}\sum_{j}^{n}\frac{J\left[E_\alpha,T_i(t_\alpha)\right]}{J\left[E_\alpha,T_j(t_\alpha)\right]}=\min\ i\neq j \tag{4.22}$$

The expected value of Eq. (4.22) is the number of combinations $\binom{n}{2}$ where n is the number of experiments.

$$\text{Expected Value}\left[\Phi(E_\alpha)\right]=\frac{2n!}{2!(n-2)}=\frac{n!}{(n-2)!} \tag{4.23}$$

In theory a reaction occurs at any temperature greater than absolute zero. Then, we should reorganize the integral equation given by Eq. (4.20),

$$J\left[E,T(t)\right]=J_o+\int_{t_o}^{t_\alpha} e^{-\frac{E}{RT(\alpha)}}dt \tag{4.24}$$

here,

$$J_0=\int_{0}^{t_0} e^{-\frac{E}{RT(t)}}dt \tag{4.25}$$

While time goes from 0 to t_0, temperature goes from absolute zero to T_0. Defining a quasi-linear (β^*) heating rate which is the slope of temperature against time. Then,

$$t_0=\frac{T_0}{\beta^*} \tag{4.26}$$

While taking experimental measurements, we assume that the starting time is 0. Then we should add t_0 to our time measurements and J_0 can be calculated by the approximation given by Senum and Yang (1979).

$$J_o \approx (B^*)^{-1} \frac{E}{R} \frac{\exp(-x)}{x} \frac{\left(x^2 + 10X + 18\right)}{\left(x^3 + 12x^2 + 36x + 24\right)} \tag{4.27}$$

Details of implementation follow. Required iso-conversional values of temperature and time are estimated by polynomial interpolation (Press et al, 1986). In order to minimize Eq. (4.22), the Brent algorithm [20] is used. This algorithm requires, a bracketing triplet of abscissas Φ_1, Φ_2, Φ_3 such that Φ_2 is between Φ_1 and Φ_3, and $\Phi_2(E_2)$ is less than both $\Phi_1(E_1)$ and $\Phi_3(E_3)$ is given as input. Choosing,

$$E_2 = \frac{E_1 + E_3}{2} \tag{4.28}$$

We can further improve our initial guess (E_2) by assuming that $\Phi(E_\alpha)$ (Eq. (4.22)) is approximated by a quadratic parabola.

$$\Phi(E_\alpha) = rE_\alpha^2 + sE_\alpha + t \tag{4.29}$$

Then, the minimum of a quadratic parabola is equal to,

$$E_0 = -\frac{s}{2r} \tag{4.30}$$

As an aside,

$$-\frac{s}{2r} = -\frac{E_3^2(\Phi_1 - \Phi_2) + E_1^2(\Phi_1 - \Phi_3) + E_2^2(\Phi_3 - \Phi_1)}{2\left[E_3(\Phi_2 - \Phi_1) + E_2(\Phi_1 - \Phi_3) + E_1(\Phi_3 - \Phi_2)\right]} \tag{4.31}$$

E_1 ,E_3 and E_0 are given as input. Simpson's rule is used in calculations of integrals (Press et al. 1986)

The above mentioned integral iso-conversional methods are derived by assuming constant activation energies. If the process involves a variation of activation energy with respect to conversion, the estimated activation energies become averaged over the region 0-α. And the resulting activation energies undergo smoothing over the region (Vyazovkin, 1997). To overcome this issue, an iso-conversional method that accounts for variation in the activation energy is proposed by Vyazovkin (2001).

By modifying Eq. (4.21) as given below, the resulting activation energies become average over the small interval $\Delta\alpha$.

$$J\left[E, T(t)\right] = \int_{t_{\alpha-\Delta\alpha}}^{t_\alpha} e^{-\frac{E}{RT(t)}} dt \tag{4.32}$$

4.4 Synthetic example

In order to test the applicability of iso-conversional methods to multistep parallel and series reactions that are characteristic of ISC, a synthetic example is generated. In order to mimic reactions in kinetic runs, we consider a two-step, irreversible, series-parallel reaction.

$$\begin{aligned} A + B &\xrightarrow{k_1} R \\ R + B &\xrightarrow{k_2} S \end{aligned} \tag{4.33}$$

Here the reaction is in parallel with respect to B (analog of oxygen) and in series with respect to A,R,S. The rate expressions are,

$$\begin{aligned} \frac{dC_A}{dt} &= -k_1 C_A \\ \frac{dC_B}{dt} &= -k_1 C_A - k_2 C_R \\ \frac{dC_R}{dt} &= k_1 C_A - k_2 C_R \\ \frac{dC_S}{dt} &= k_1 C_R \end{aligned} \tag{4.34}$$

In this example B is analogous to oxygen in kinetic runs. In kinetic runs the rate of reaction does not depend on oxygen concentration as oxygen concentration is constant in the system. The parameters assumed are given in Table 4.1. The Arrhenius relationship, Eq. (4.4), describes the rate constants k_1 and k_2 that have activation energies of E_1 and E_2, respectively.
The initial conditions for the rate equations are,

$$\begin{aligned} C_A(t=0) &= 0.5 \\ C_B(t=0) &= 1.0 \\ C_R(t=0) &= 0.0 \\ C_S(t=0) &= 0.0 \end{aligned} \tag{4.35}$$

The set of equations Eq. (4.4)and Eq.(4.34) with the list of parameters and initial values Eq. (4.35), are solved using the Runge-Kutta method. The concentration profiles with respect to time for different species are shown in Figure 4.1 and 4.2 for a heating rate of 1K/min . Figure 4.3 shows the profiles in terms of conversion of B. The process is then simulated at 4 different heating rates of 1, 2, 4, 8 K/min. Figures 4.3 and 4.4 show the results.

4.5 Experimental apparatus and procedures

Iso-conversional methods require a series of experiments on the same sample to be conducted at different heating rates. In our study, ramped temperature oxidation experiments with the analysis of effluent gas are conducted using the apparatus described in the experimental section.

4.6 Experimental Results – Hamaca oil

Experiments were conducted with Hamaca crude oil (10.5˚API) at four different heating rates: 2.5 K/min, 3 K/min, 3.5 K/min, 4.5 K/min. All the experiments were carried out at the same pressure of 100 psia and air injection rate of 2 SLPM. The only parameter that differs is the heating rate. The sample consists of 42 g of fine sand, 3 g of clay, 4g water and 2 g of Hamaca oil. Some trial and error was necessary to determine the optimal mass of oil. The mass of the oil should be small enough to minimize the temperature deviations resulting from exothermic reactions and should be large enough to produce sufficient effluent gas to be analyzed. The sample is put into kinetic cell and the kinetic cell is placed in the furnace. The temperature is increased linearly and the change in oxygen consumption is recorded.

Temperature within the cell and the concentration of the effluent gas is recorded. Figure 4.5 shows oxygen consumption with respect to time for different heating rates.

Figure 4.6 shows the temperature change with respect to time. Notice that each curve deviates from the temperature program at two specific temperature ranges. Those points correspond to LTO at 500-620 K and HTO at 640-750 K, respectively. As a result of the exothermic reactions occurring in LTO and HTO, temperature deviates from the pre-programmed history.

In the experiments, we measure the composition of the effluent gas and in our analysis. One can relate rate of oxygen consumption with rate of carbon consumption by the following relation.

$$\frac{dC_C}{dt} = \nu \frac{dC_{O_2}}{dt} \tag{4.36}$$

Data is processed before the analysis. Change in oxygen concentration is calculated by using Eq. (4.37).

$$\Delta C_{O_2} = O_2 - \frac{C_{N_2} C_{O_2}}{1 - C_{CO_2} - C_{CO} - C_{O_2}} \tag{4.37}$$

Figure 4.7 shows experimental oxygen consumption as a function of temperature. The two humps on O_2 uptake on Figure 4.7 are typical characteristics of oil combustion. Oxygen uptake peaks at two different points. Those correspond to LTO and HTO. Figure 4.8 indicates that LTO reactions occur between 500 K and 620 K and between 620 K and 640 K MTO is observed. At 640 K HTO reactions begin and these reactions occurred until 720 K.

It would be inconvenient to calculate directly the change in oxygen concentration by just subtracting inlet oxygen concentration from outlet gas

concentration. During the experiments compositions change. Therefore, we need to apply a correction for the outlet gas that is represented by Eq. (4.37).

Before proceeding, we first analyze the data with the method proposed by Fassihi [1] by assuming a first order elementary reaction. Tables 4.2 and 4.3 show the activation energies estimated from each experiment for each group of reactions. Note that this method appears to give significantly different activation energies depending on the heating rate

Now we apply the iso-conversional method proposed by Vyazovkin (2001). This approach provides a model-free method and also accounts for the variations in the temperature as well as the various heating rates. Figure 4.9 shows the estimated activation energy as a function of conversion. At different conversion levels, the different activation energies are estimated (circled data). This result clearly indicates that the reaction mechanism changes during the experiment. This is as we expected because different groups of reactions are known to be occurring during oil oxidation.

In order to relate temperature with conversion, we average temperature of the four different temperature histories to estimate one set of activation energies. This procedure is consistent with the iso-conversional approximation. Figure 4.10 shows activation energy average versus temperature.

Results exhibit a clear HTO region with a single activation energy of 51000 J/mol between 650 and 700 K and MTO with 30000 J/mol. Additionally, reactions occurring in LTO are not similar in magnitude to be grouped together. Figure 4.10 appears to teach us that lumping the LTO reactions into a single reaction may be a gross over simplification for Hamaca oil. Also the conversion range between levels of 0.7 and 0.8 there may be another group of reactions in the temperature range 720 and 750 K.

Figure 4.10 indicates a region of the graph that shows zero activation energies. The routine used to calculate activation energy by using iso-conversional method has a positivity constraint and estimated activation energy is not allowed to be negative. Therefore, 0 is the minimum value. The reactions corresponding to the region $0.25<\alpha<0.4$ seem not to follow Arrhenius behavior.

Now we focus on predictions using the iso-conversional methods. In fact, we can simulate the process by using the same objective function used for calculating activation energies (Eq. (4.22)) without knowing the pre-exponential factor or the reaction model. Here we minimize the same objective function by using the experimental data and the change of activation energy with respect to conversion that we acquired previously. This leads us to calculate the time required for a given conversion level to be reached at a specified temperature program. Figure 4.11 shows the prediction results. The predictions are made for the case with a heating rate of 3 K/min. The dark line shows the experimental results and the light one is the estimate.

4.7 Experimental Results – Cymric oil

It is known that water soluble metallic additives improve the burning characteristics of some crude oils (Shallcross et al., 1991; Castanier et al, 1992; He et al., 2005). In this section the effect of metallic additives is investigated on the

combustion of Cymric light oil. Three experiments were conducted at heating rates of 2, 2.5, and 3 K/min at the same pressure of 100 psia and flow rate of 2 SLPM. The sample consists of 42 g of fine sand, 3 g of clay, 4g water and 3 ml of Cymric oil (34˚API). In the runs with metallic additives 0.5 g of $Fe(NO_3)_3$ was added. Figures 4.12 and 4.13 show oxygen uptake as a function of time for the case without and with metallic additives respectively. Also Figures 4.14 and 4.15 show temperature and conversion profiles.

The iso-conversional analysis was applied to both cases and the results are given by Figure 4.16. Figure 4.14 shows that for the case without any additives, LTO reactions occur between 550 and 600 K and HTO reactions occur between 640 and 710 K. As it is shown by Figure 4.15, for the case with metallic additives, LTO reactions occur again between 550 and 600 K on the other hand HTO reactions occur between 600 and 650 K. In Figures 4.14 and 4.15 temperature profiles deviate from the pre-programmed history at 400 K. These effects are due to boiling of water at 100 psi and do not correspond any reaction.

As Figure 4.16 clearly indicates, at the end of LTO reactions a significant activation energy barrier exists. That energy barrier might be the reason for the reported poor burning quality of Cymric light oil (He al, 2005). On the other hand for the case with metallic additives, the figure does not exhibit this energy barrier. As a result, addition of metallic additives removed this energy barrier. Figure 4.16 indicates a clear distinction between the cases with and without metallic additives. The addition of metallic additives changes the path completely resulting in better burning quality. He et al. (2005) identified the creation of activated sites on the solid matrix with the addition of metallic salt additives. It was inferred that these sites contributed to improved combustion performance, Both, curves in Figure 4.16 exhibit a region with zero activation energy. These belong to neither LTO nor HTO groups. The reactions during these regions seem not to follow Arrhenius law behavior. As an observation, the zero apparent activation energy region widens and moves to greater conversion with the addition of metallic additives. Additional work is clearly needed to understand how the iso-conversional – RTO method differentiates ISC kinetics among good and bad performers.

4.8 Summary

Iso-conversional methods of crude oil reaction kinetic analysis are developed to provide a model-free procedure to find activation energy as a function of conversion. Iso-conversional approaches avoid any uncertainties or errors in the parameter estimates associated with the reaction model assumption. The iso-conversional are not assumption free however. All iso-conversional methods are based on the principle that at the same extent of conversion, the reaction rate is only a function of temperature. It is important to understand these assumptions and the limitations associated. It must be evaluated in another study to understand thoroughly the iso-conversional principle applied to oxidation. In general, energetic materials are consistent with the iso-conversional principle (Burnham and Dinh, 2007). It is also possible to make kinetic predictions using the proposed method.

The non-linear iso-conversional method proposed by Vyazovkin (2001) not only accounts for the variations in temperature but also account for the variations in

activation energy. The method provides activation energy as a function of conversion or temperature. In this method a complex series of reactions are treated as a single step. Therefore, estimated activation energies are apparent. Nevertheless, we believe that representation of apparent activation energies as a function of conversion gives a picture about the complex physical system and is promising to characterize the burning quality of particular oils.

There are some issues regarding the analysis. First of all, consistency of experiments is required for the analysis. The method of analysis requires that the only difference between experiments should be the heating rate. It is not always easy however, to control these types of experiments. Therefore, we need a way to evaluate the inconsistencies between the experiments to be used in the analysis. As it is already shown, it is possible to make predictions using the method described. In that case, the experimental data and the activation energy as a function of conversion are given as input and for a given conversion the time required is estimated. Making predictions in such a way makes it difficult to be implemented for practical reasons. First, it is not appropriate to input the experimental data as it restricts us to make predictions in the correlation range. Second, it is inconvenient to estimate the time required for a given conversion rather estimating conversion for a given time or temperature. For the last, although it is possible to make predictions without estimating the Arrhenius constant, the method does not provide a way to estimate Arrhenius constant (pre-exponential factor). We continue to work to develop a self-consistent method to obtain the Arrhenius constant in a model-free fashion.

Observations associated with the addition of water-soluble metallic additive:

- the large activation energy barrier disappeared
- path of reactions changed favorably,
- the peak consumption of oxygen decreased,
- the temperature region that LTO reactions occurred is unchanged on the other hand the temperature at which HTO reactions began decreased, and
- the region with zero activation energies extended.

Table 4.1 – Parameters used in the synthetic example.

C.	**k_1(A, E)**	**k_2(A, E)**
Arrhenius Constant, 1/min	10^8	10^{16}
Activation Energy, J/mol	10^5	2.1×10^5

Table 4.2 – Activation energy of LTO from different experiments, conventional analysis.

	2.5 K/ Min	**3.0 K/ Min**	**3.5 K/ Min**	**4.5 K/ Min**
Slope, (E/R, K)	6838	4502	5488	7491
E, J/mol	56851	37430	45627	62280

Table 4.2 – Activation energy of HTO from different experiments, conventional analysis.

	2.5 K/ Min	**3.0 K/ Min**	**3.5 K/ Min**	**4.5 K/ Min**
Slope, (E/R, K)	4587	6992	4985	6248
E, J/mol	38136	58132	41445	51946

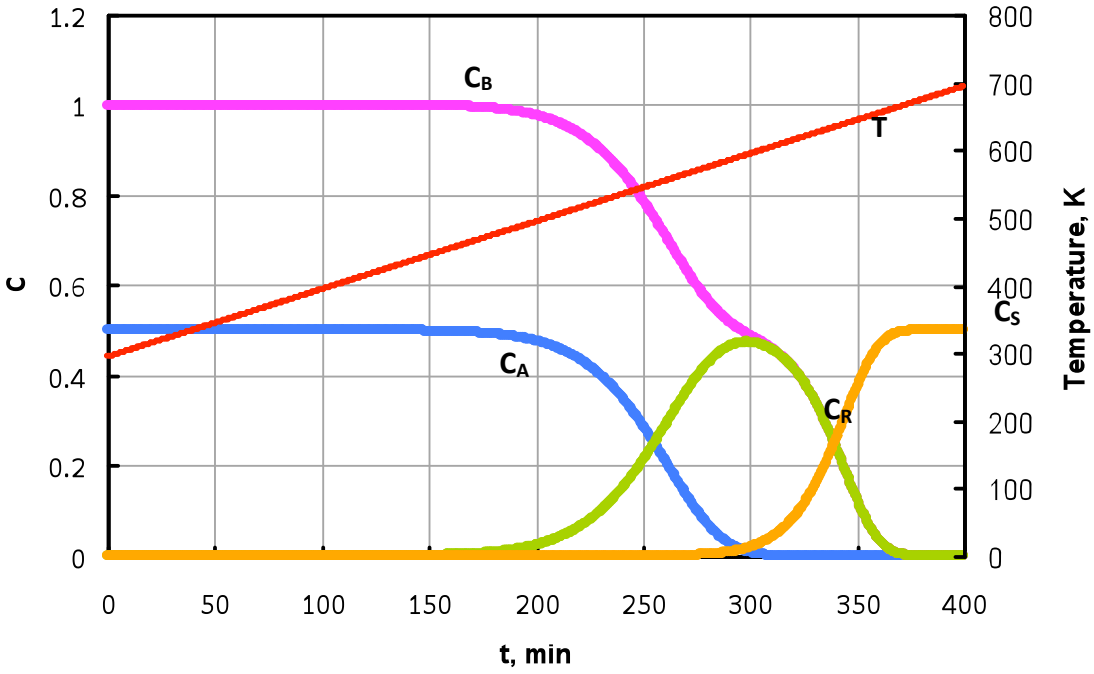

Figure 4.1 – Concentration profiles for the case with heating rate of 1 K/min.

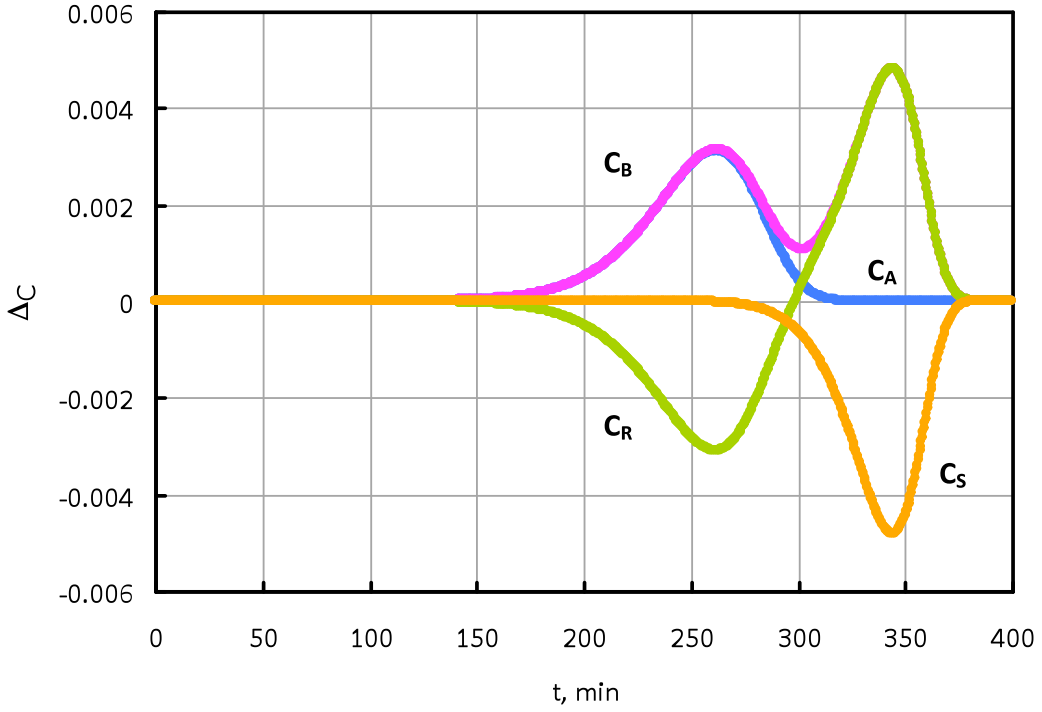

Figure 4.2 – Concentration profiles (ΔC) for the case with heating rate of 1 K/min.

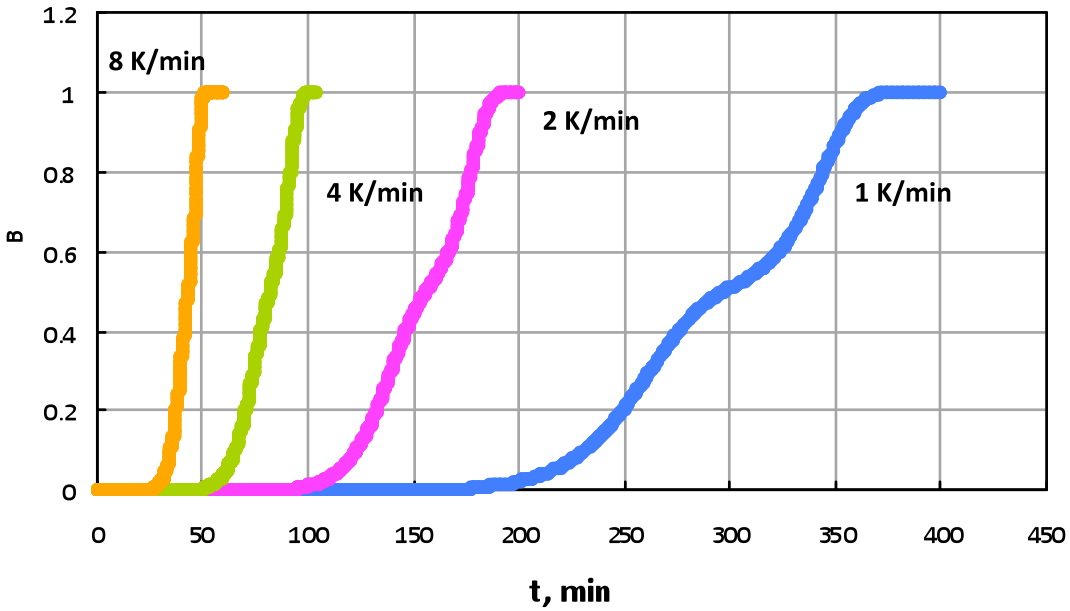

Figure 4.3 – Conversion profiles for species B, synthetic case.

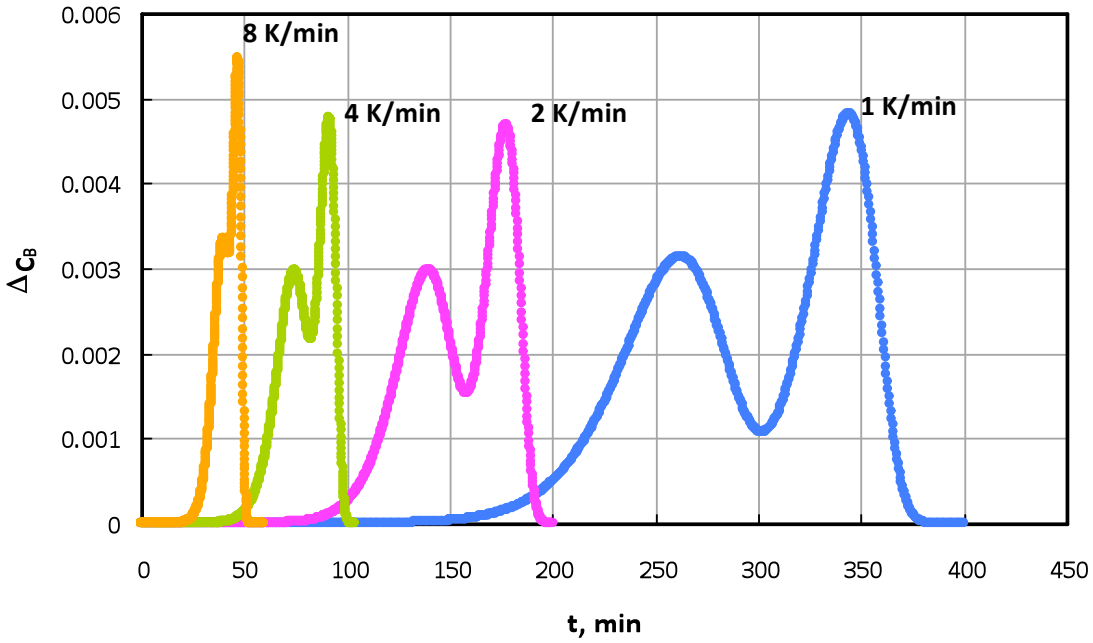

Figure 4.4 – Concentration profiles (ΔC) for B, synthetic case.

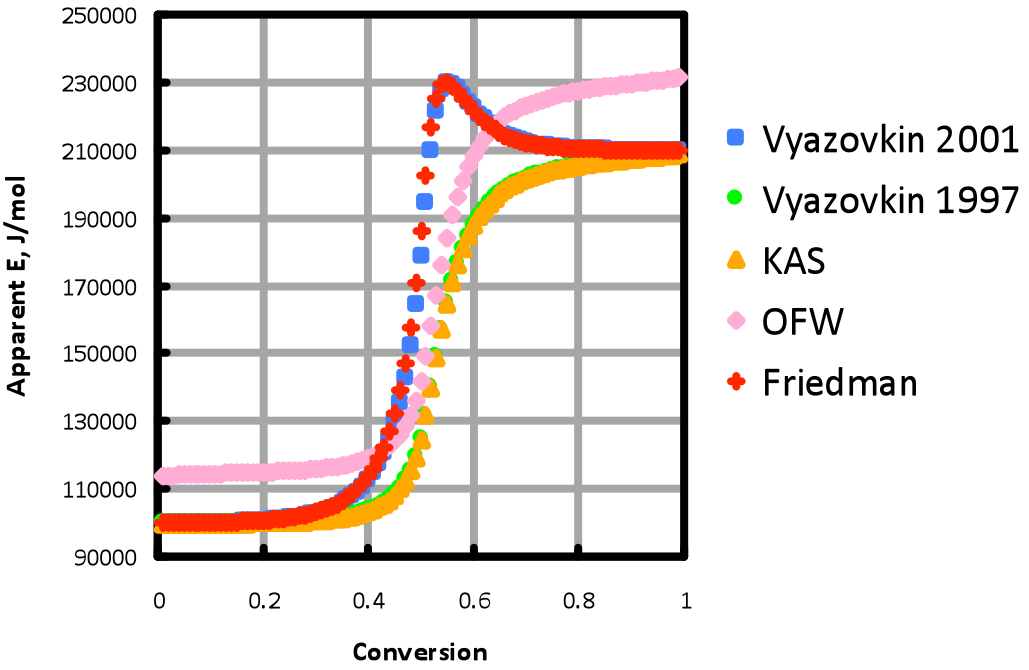

Figure 4.5 – E dependencies evaluated for the simulated process by different methods, synthetic case.

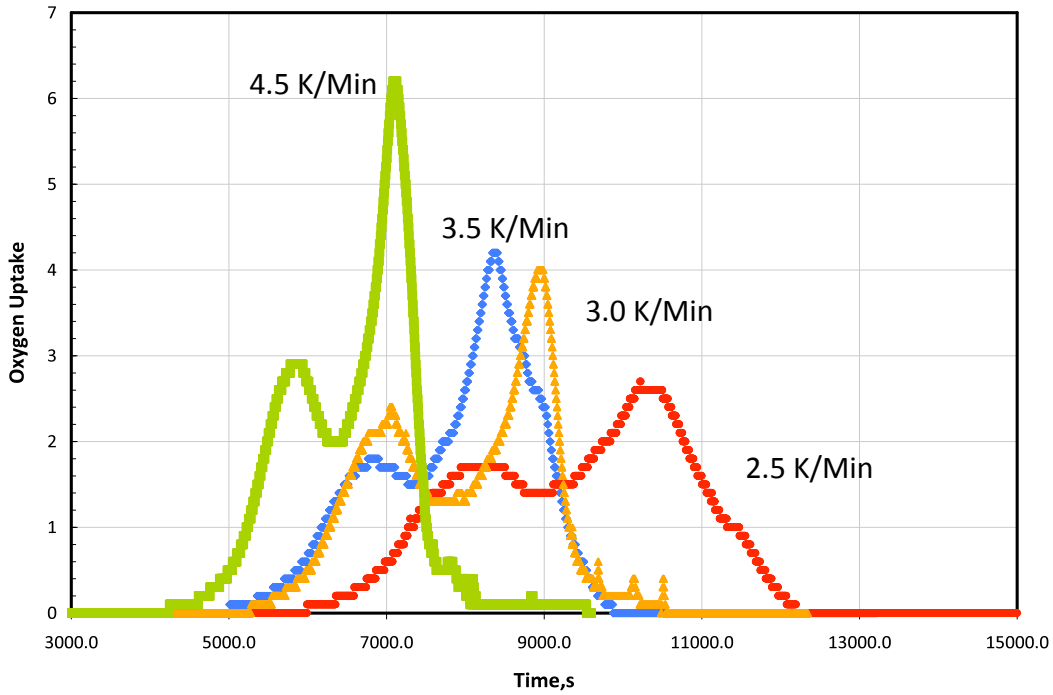

Figure 4.6 – Oxygen consumption with respect to time from RTO of Hamaca crude oil.

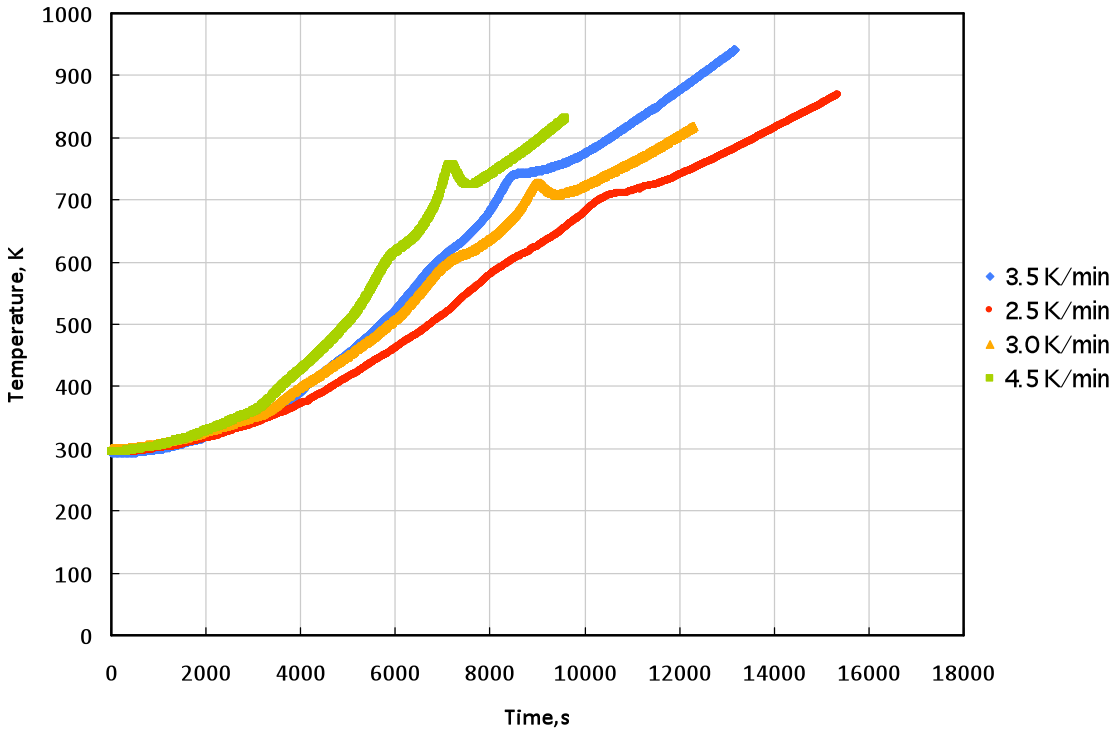

Figure 4.7 – Temperature with respect to time from RTO of Hamaca crude oil.

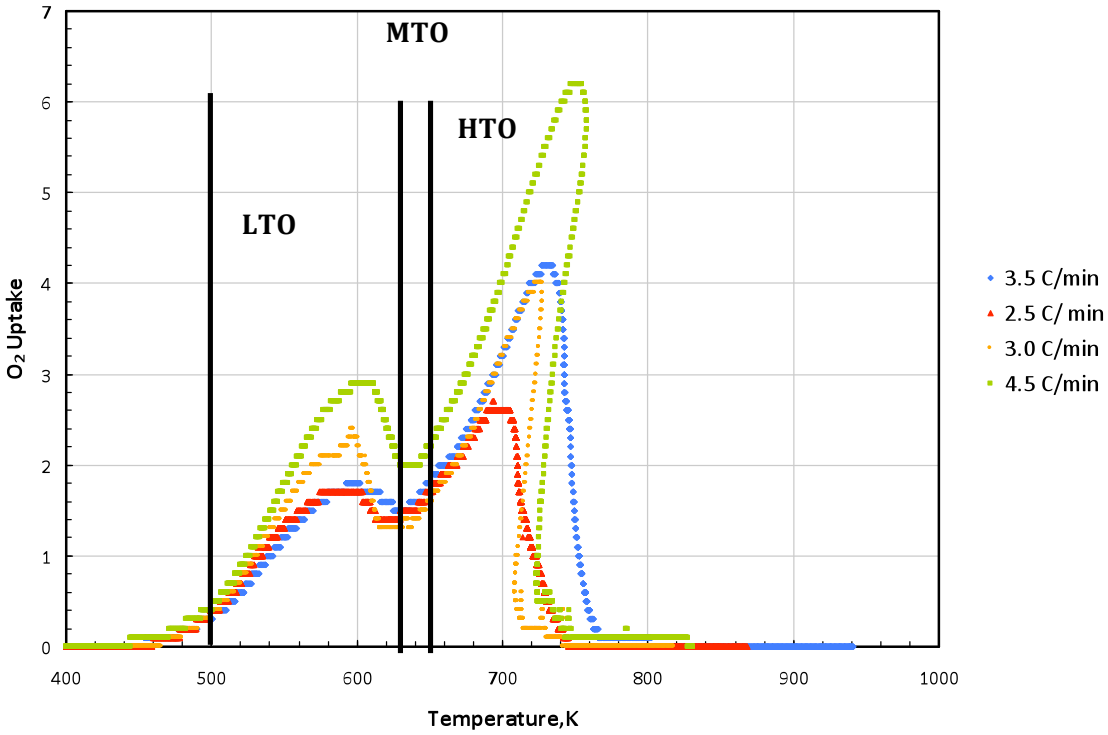

Figure 4.8– Oxygen consumption with temperature from RTO of Hamaca oil.

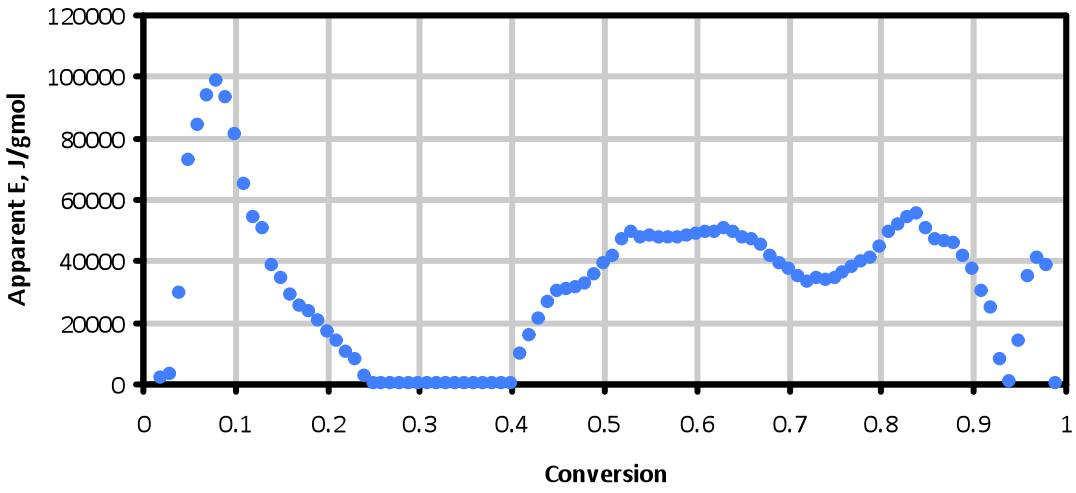

Figure 4.9 – Apparent activation energies from iso-conversional method, Hamaca crude oil.

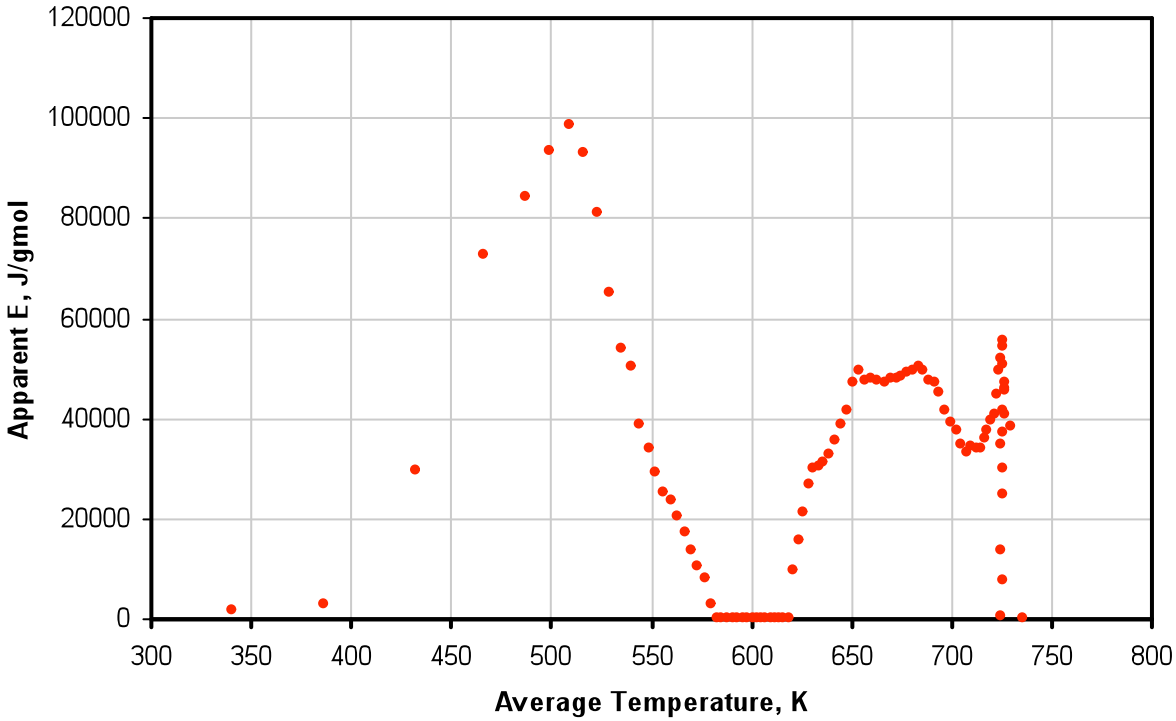

Figure 4.10– Apparent activation energy & average temperature, Hamaca crude oil.

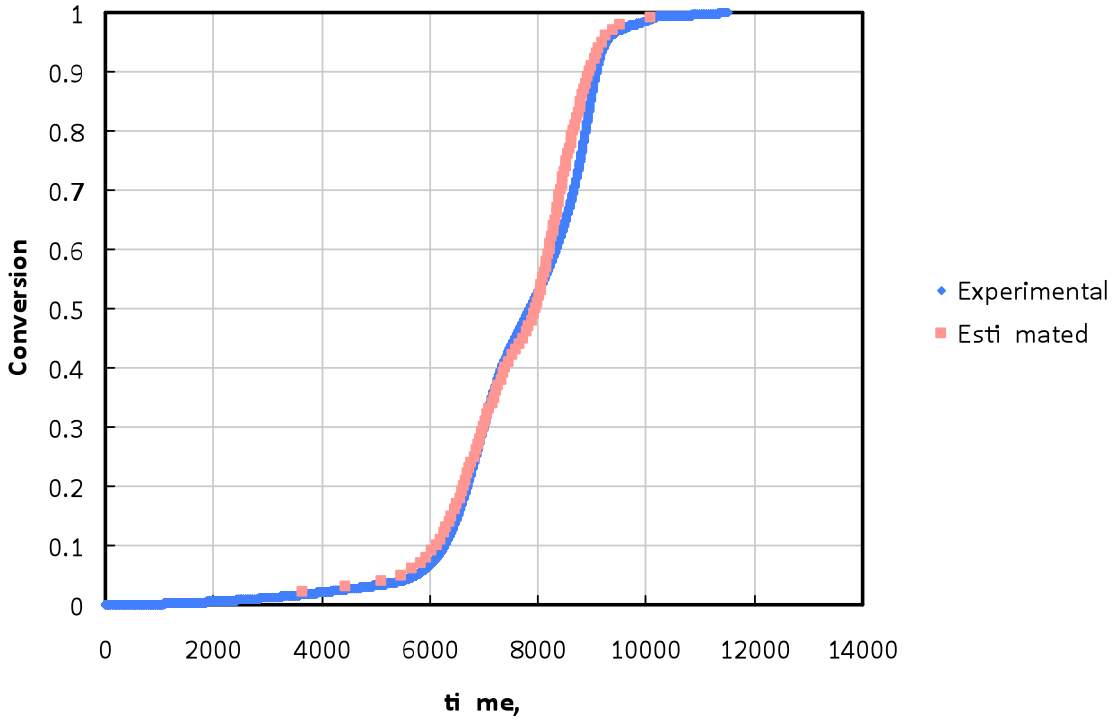

Figure 4.11 – Kinetic predictions for the case with β = 3 K/min with positivity constraint in optimization of activation energies.

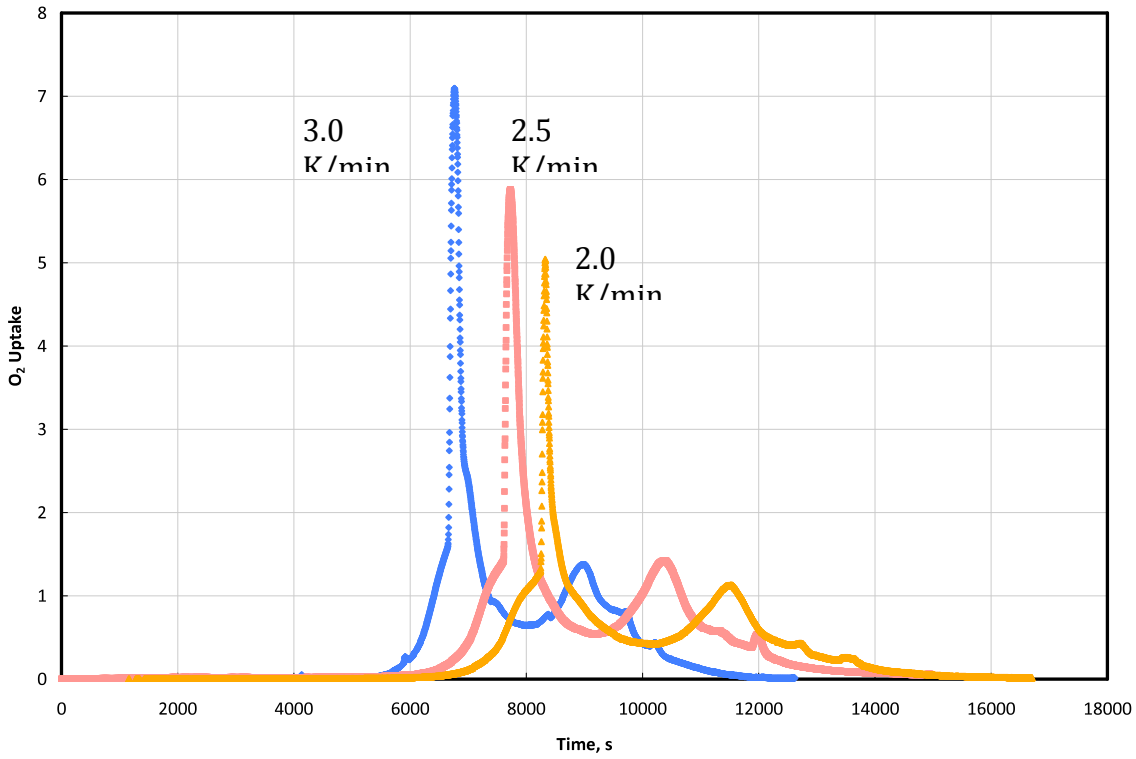

Figure 4.12 –Oxygen uptake vs. time for the case without metallic additives, Cymric cruce oil.

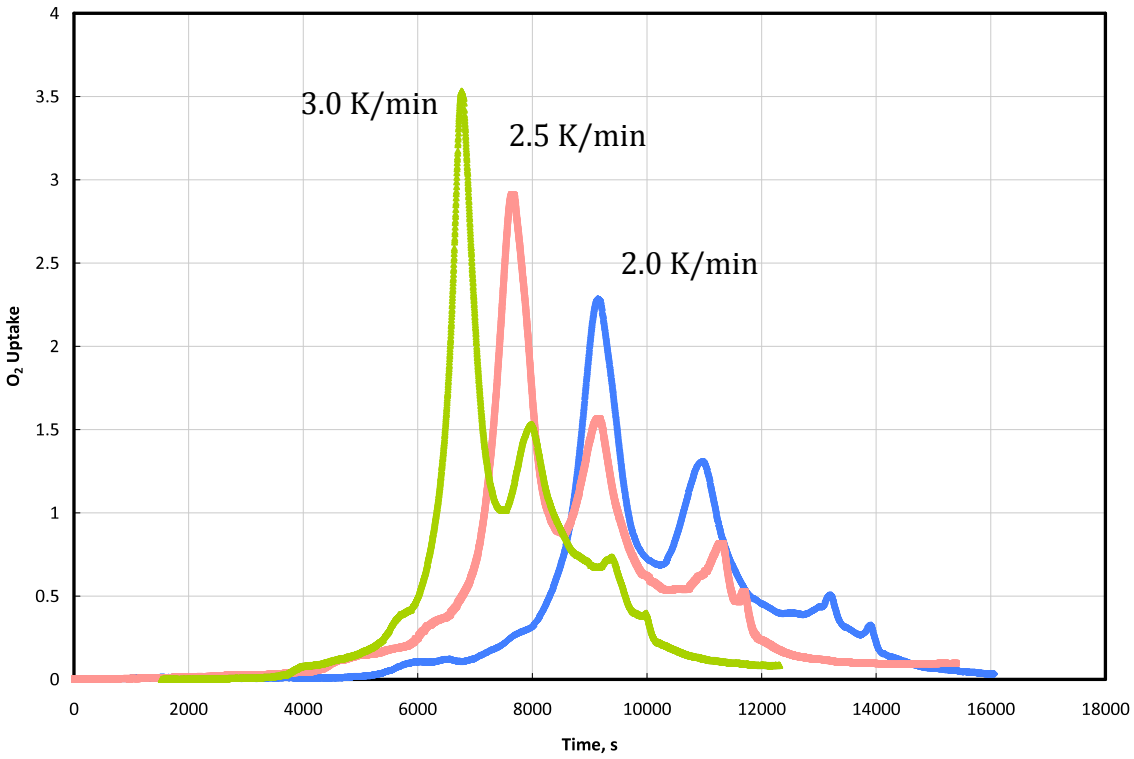

Figure 4.13 –Oxygen uptake vs. time for the case with metallic additives, Cymric crude oil.

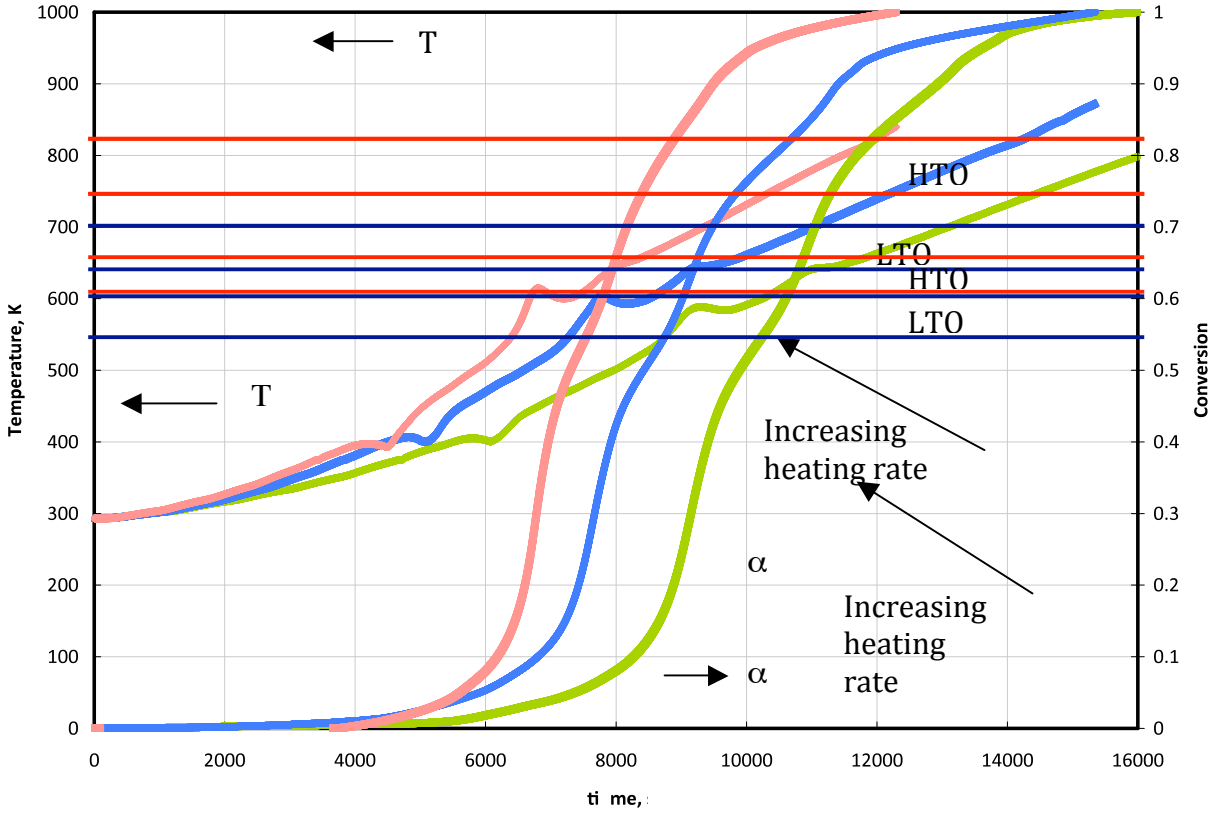

Figure 4.15 –Temperature and conversion profiles for the case with metallic additives, Cymric crude oil.

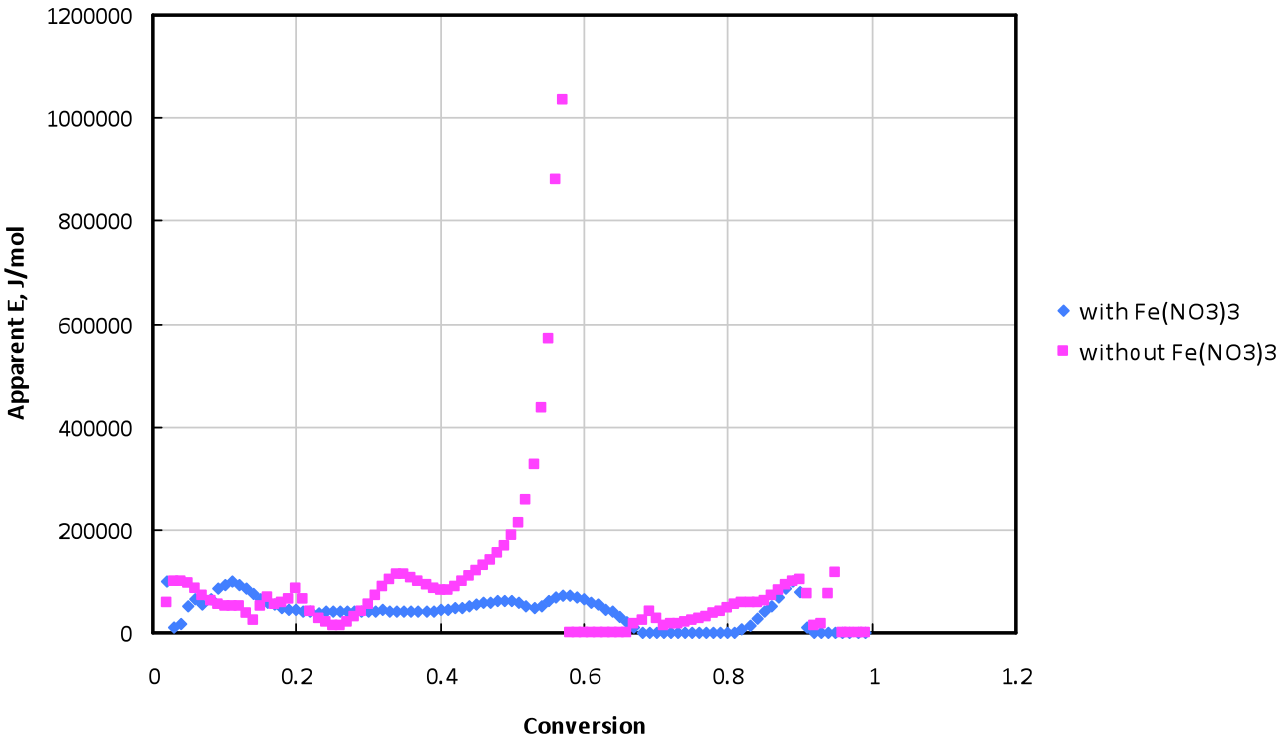

Figure 4.16–Apparent E vs. conversion for Cymric crude oil with and without metallic additives.

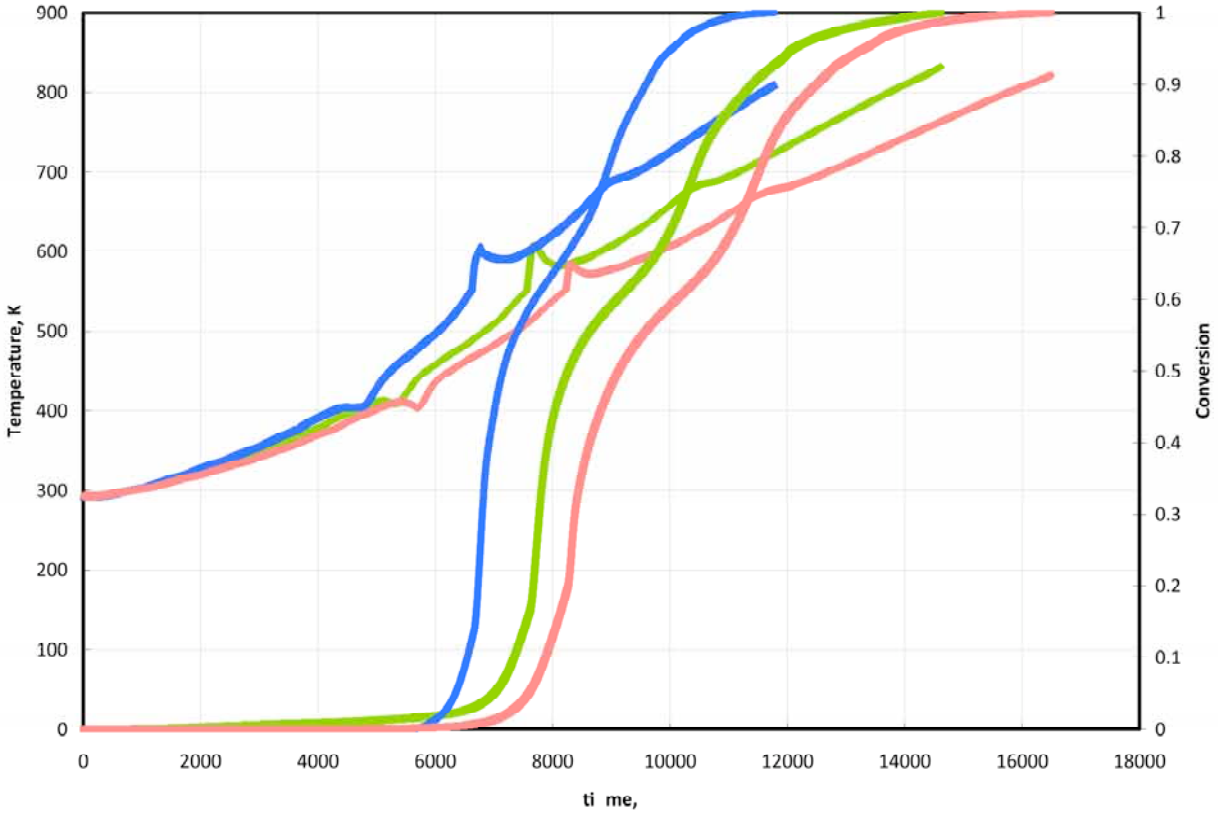

Figure 4.17 –Temperature and conversion profiles for the case without metallic additives, Cymric crude oil.

5. Combination of Solvent Injection and In-Situ Combustion for In-Situ Upgrading and EOR

Application of cyclic solvent injection into heavy and viscous crude oil followed by in-situ combustion of heavy residues is explored from a laboratory perspective. The solvent reduces oil viscosity in-situ and extracts the lighter crude-oil fractions. Combustion cleans the near well region and stimulates thermally the oil production. Both solvent injection and in-situ combustion are technically effective. The combination of the two methods, however, presents new potential for in-situ ugrading and enhanced oil recovery.

Hamaca (Venezuela) and West Sak (Alaska) crude oils were employed in this study. First, ramped temperature oxidation studies were conducted to measure the kinetic properties of the oil prior to and following solvent injection. Pentane, decane, and kerosene were the solvents of interest. Second, solvent was injected in a cyclic fashion into a 1 m long combustion tube. Then, the tube was combusted. Hamaca oil presented good burning properties, especially following pentane injection. The pentane extracted lighter components of the crude and deposited preferentially effective fuel for combustion. On the other hand, West Sak oil did not exhibit stable combustion properties without solvent injection, following solvent injection, and even when metallic additives were added to enhance the combustion. We were unable to propagate a burning front within the combustion tube. Further investigation of the West Sak system may reveal a combination of metallic additives and oxygen partial pressure that allows combustion. The experimental results do show that this combined solvent-combustion method is applicable to the broad range of oil reservoirs with properties similar to Hamaca.

5.1 Introduction

This study investigates the effect of solvent injection on the subsequent performance of in-situ combustion. The work is based on experimental results obtained by a combination of these two successful in-situ upgrading processes for viscous oils. It is envisioned that application in the field occurs first by a cycle of solvent injection, a short soak period, and subsequent oil production using the same well (Castanier and Kovscek, 2005) . By mixing with oil, the solvent decreases the oil viscosity and upgrades the crude by extracting in-situ the lighter ends of the crude oil. The heavy ends, that are markedly less interesting, are left behind. Injection of solvent and oil production occurs for a number of cycles until the economic limit is reached or until the deposition of crude oil heavy ends damages production. The solvent injection phase is followed by in-situ combustion that burns the heavy ends left from the solvent injection. By switching from air to nitrogen injection the combustion is extinguished. Again, oil is produced by the same well used for injection in a cyclic fashion. Combustion enhances the production by decreasing thermally the oil viscosity and adding energy to the reservoir through the formation

of combustion gases. The combustion also upgrades the oil through thermal cracking (Castanier and Brigham, 2003).

For our experiments, two oils of particular interest were used. The first experiments employed crude oil from Hamaca (Venezuela), where the field location requires important costs of transporting crude to upgrading facilities. The second set of experiments was conducted with viscous West Sak oil (Alaska) where steam injection currently appears to be unsuitable because of heat losses to permafrost.

While the presence of oil in the Orinoco heavy-oil belt, in Central Venezuela, was discovered in the 1930's, the first rigorous evaluation of the resources was made in the 1980's . It contains between 1.2 and 1.8 trillion recoverable barrel (Kulhman, 2000) of heavy and extra-heavy oil. The 9-11 °API density crude is processed at the Jose refinery complex on the
northern coast of Venezuela. The cost of transporting heavy oils to the northern coast provides an incentive to investigate in-situ upgrading. In 2003, the total production from these projects was about 500,000bbl/d of synthetic crude oil. This figure was expected to increase to 600,000bbl/d by 2005 (Acharya, 2004).

West Sak is a viscous oil reservoir located within the Kuparuk River Unit on the North Slope of Alaska. It is part of a larger viscous oil belt that includes Prudhoe Bay. The estimated total oil in place ranges from 7 to 9 billion barrels with an oil gravity ranging from 10 to 22 °API. The reservoir depth ranges from 2,500 to 4,500 feet with gross thickness of 500' and an average net thickness of 90'. The temperature is between 45 and 100 °F and there is a 2,000 ft (600m)-thick Permafrost layer. In March 2005, 16,000 BOPD were produced and 40,000 BOPD were planned for 2007 (Targac. 2005), although it is not reported if this target was met. Within the scope of this study, West Sak is of particular interest because there are technical difficulties with steam injection that include (Gondouin and Fox, 1991):

- surface generated steam passing through a thick permafrost layer
- the well would sink if the permafrost melted,
- the reservoirs consists of thin, medium permeability layers,
- the formation may contain swelling clays that reduce the rock permeability when exposed to steam condensate,

Solvent injection and in-situ combustion are effective in a variety of fields. Both techniques upgrade the oil directly in the reservoir thereby making heavy resources easier to exploit. The combination of these two processes is applicable at large scale to recover viscous oil, or in-situ combustion could be applied on an ad hoc basis to clean the wellbore region, increase the permeability, and thus to act as a stimulation process.

5.2 Procedure

Two types of experiments were performed: ramped temperature oxidation (RTO) studies to measure combustion kinetics and tube experiments to determine combustion front propagation. This combination of experiments is effective for screening of combustion candidates. The apparatus as described in section A. Experimental was used. We did not try to reproduce the reservoir pressure and temperature conditions. Instead, we focused on the feasibility of combining solvent

injection with in-situ combustion and representative porous media were created to conduct these experiments.

Our interpretations were developed by comparing results from combustion of hydrocarbons remaining after solvent extraction to combustion of crude oil without prior solvent injection. Both the kinetics cell and the combustion tube were packed with mixtures of sand, clay, oil, and water. Mixtures were prepared in proportions of 8500 g of sand, 450 g kaolinite, 400 g water, and 450 g crude oil. In some experiments, 29 g of the metallic additive ferric nitrate $Fe(NO_3)_3$ is added in an attempt to enhance combustion properties. The sand employed is identical to that described in section B. Water Soluble Additives to Improve ISC

The kinetics experiments begin by packing about 50 grams of the oil/water/sand/clay mixture into a kinetics cell placed vertically to minimize the effect of gravity. The outlet (upper in a vertical sense) of the kinetics cell was connected to a back pressure regulator to pressurize the cell to 40 psi during the solvent injection phase. At the inlet, the desired volume of solvent was first injected followed by nitrogen injection at ambient temperature to displace upgraded oil. Once, oil production ceased, the injected gas was switched to air and the kinetics cell was placed in a furnace. The furnace was programmed to increase temperature linearly with time (approximately 180°C per hour) up to 700°C The area between the measured temperature versus time trend and the line representing the constant heating versus time represents the amount of combustion. During kinetic runs of the residues remaining after solvent extraction, the pressure at the entrance of the cell was about 120 psi and exit pressure was maintained at 90 psi by the backpressure regulator. When the sample reached 650 °C, the system was turned off and left to cool down to room temperature. According to the data, temperatures at which HTO and LTO reactions are dominant were determined, as well as the heat created during HTO reactions.

A tube run began by placing the thermowell in the middle of the tube and anchoring it to the bottom flange. Roughly, 2 cm height of clean sand was added at the bottom of the tube. The tube was filled with the same sand mixture packed in the kinetics cell. After each incremental addition of the oil-sand mixture, the sand in the tube was tamped with a plunger that was perforated in the middle to allow the thermowell to pass. After pressure testing and placement in the insulating jacket, the tube was subjected to solvent extraction in a fashion similar to that employed for kinetics experiments. Following extraction, nitrogen was injected to displace any remaining solvent and to establish a pathway for the air to follow. The top of the tube was gradually heated to 450°C with an electrical resistance band heater placed at the top of the tube. Nitrogen was injected through the tube to establish permeability. When the temperature was great enough to trigger HTO reactions, nitrogen injection was stopped and air was injected at 2.5 l/min. The band heater was turned off at the same time. Ignition was always observed immediately upon switching to air injection. The thermocouple always recorded a significant increase in temperature when switching from nitrogen to oxygen injection. This observation confirmed confirmed ignition of the tube as did the presence of carbon oxides in the effluent gas.

The produced gas composition was continuously monitored and the displaced liquids were collected through the separator. Temperature measurements made during a tube run were used to establish temperature profiles and to monitor the propagation of the combustion front. Particular attention was paid to the temperature at the combustion front. If it was below roughly 350°C, then the combustion front was no longer creating enough heat and the experiment was stopped. Also the CO_2 and CO concentrations in the effluent gas were used to confirm whether the high temperature combustion reactions were active or not.

5.3 Results from Kinetics Experiments

Kinetics experiments were conducted on Hamaca, Venezuela and West Sak, Alaska crude oils. The samples were approximately 11 and 21 ° API and had asphaltene contents of 11.3 and 3.0 wt %, respectively. We report here only single sets of results for reasons of brevity. Duplicate experiments are found in the thesis of Cristofari (2006).

5.3.1 Hamaca

Figure 5.1 represents effluent gas compositions profiles for oxidation of the Hamaca oil and no solvent injection. Figure 5.2 presents O_2 effluent concentration and temperature profiles from the same test. The range of temperature for LTO and HTO reactions is determined by oxygen uptake and positive deviation from the linear heating schedule. In agreement with the literature, LTO reactions were first observed at 350°C. For this oil, only a small consumption of oxygen took place. Heavy oils are in fact less susceptible than light oils to LTO (Sarathi, 1999; Mamora and Brigham, 1995). The HTO reactions began at 400 °C and the temperature increased immediately by 290 °C Again, positive deviation from the linear heating schedule is taken as evidence of oxidation. HTO reactions were characterized by a great consumption of oxygen and production of CO, CO_2, and CH_4 in the effluent gas. These reactions are quite exothermic and create the heat necessary to sustain combustion. The oxygen uptake versus temperature from the kinetics cell is analyzed with standard techniques (Shallcross et al. 1991, Moore et al.,1995) to obtain the activation energy of high temperature oxidation. The activation energy is found to be 83 kJ/mol.

Figure 5.3 represents the temperature histories during the combustion phase, after different solvents were first injected into the kinetics cell. HTO reactions consist of burning coke (the fuel for these reactions) that is formed primarily from the crude oil heavy ends, in particular asphaltenes. Generally, the more heavy ends present in the oil, the greater is the fuel deposition and the greater the total heat generated by HTO. Because oxygen is always in excess in our experiments, the limiting factor for the temperature to rise is the amount of fuel. Positive deviations from the linear heating schedule of the kinetics cell are directly related to the amount of heavy ends present in the kinetics cell. Thus, the rise in temperature during HTO reactions is an indication of the amount and quality of fuel deposited in the kinetics cell resulting from the injection of different solvents.

The activation energies for HTO instruct about the ease of combustion. As the solvent progresses from pentane to decane to kerosene, the activation energy

increases from 100 to 105 to 150 kJ/mol. Hence, these tests teach that pentane is a suitable solvent and that successful in-situ combustion for Hamaca crude oil following solvent extraction is challenged as the solvent becomes heavier. The heat created during HTO reactions is the energy necessary to sustain in-situ combustion. If the HTO reactions do not produce enough heat, no fuel is deposited and ultimately the combustion is not self sustaining. The most exothermic HTO reactions occur for the test without any solvent extraction. Because no solvent has been injected previously, nothing has been extracted and so maximum fuel deposition is expected in the kinetics cell. All the other tests with prior solvent extraction display smaller positive deviations from the linear heating schedule indicating that less fuel has been deposited.

There is a difference, however, between the kinetics tests of the residues remaining after each solvent injection. When pentane is injected, HTO reactions of the solvent residues display greater heat production in comparison to the tests of the residues remaining after kerosene or decane injection. The temperature increases by 250 °C as HTO reactions start (as gauged by the deviation from the linear heating schedule). In the case of kerosene or decane, the increase in temperature is smaller. It is insignificant for kerosene and temperature rises only by 150 °C in the case of decane. Thus, the amount of heat created in these latter two cases is likely insufficient to sustain an in-situ combustion process. All four experiments were performed with the same amount of sand/oil mixture at the beginning, but the extent of HTO reactions are different. Because HTO reactions of deposits remaining after pentane extraction result in a greater temperature increase than combustion of deposits from decane and kerosene injection, we surmise that more fuel is deposited during combustion of the residues following pentane injection than during the combustion after decane or kerosene injection. Pentane is a lighter hydrocarbon solvent and has extracted a relatively greater fraction of light components than heavy components from the kinetics cell, whereas kerosene and decane, heavier hydrocarbon solvents, dissolved a significant fraction of the heavy components. Recalling also that pentane dissolves resins but not asphaltenes, more asphaltenes precipitate when pentane is injected (Speight, 2007). Subsequently, heavier fractions that result in fuel for combustion are deposited in the kinetic cell. Whereas asphaltenes are soluble in decane and to an even greater extent in kerosene, there is less precipitate when decane or kerosene are injected. Because asphaltenes are less soluble in decane than in kerosene, the reactions with decane result in greater temperature increases in comparison to measurements following kerosene injection, as shown in Table 5.1. That is, the temperature increases from solvent to solvent are consistent with the solvent action of each injectant.

Figure 5.4 compares the temperature profiles during the combustion phase for two different amounts of pentane (100ml and 50 ml) injected to solubilize the heavy oil. From the temperature profiles presented in Figure 5.4 HTO reactions are equivalently exothermic: approximately the same amount of fuel has been burned in the two experiments. Injecting more than 50 ml of pentane does not dissolve additional heavy ends. With the Hamaca oil, pentane injection still results in vigorous in-situ combustion because the heavy ends necessary for the fuel

deposition were precipitated and remain in the cell in significant amounts. The amount of pentane to inject is thus closer to 50 ml than 100 ml for the kinetics cell.

5.3.2 West Sak

Figure 5.5 presents the composition of the effluent gas in O_2, CO, CO_2 and CH_4 during a kinetics test without any previous solvent injection. Compared to Hamaca oil, the concentration of O_2 decreases not only during the HTO reactions but also during the LTO reactions: light oils differ from the heavy oil by consuming substantially more oxygen during LTO reactions (Sarathi, 1999). LTO reactions are believed to transform the low molecular weight fraction of crude oil into higher molecular weight products (Shallcross et al. 1991).

Although West Sak did not seem to burn efficiently, pentane should only extract the light ends and leave the heavy ends involved in the fuel formation for HTO reactions. By injecting pentane prior to the combustion, we should obtain a heavier oil in the kinetic cell and it should not affect the increase in temperature during HTO reactions. From Figure 5.6 and Table 5.1 the amount of heat produced was reduced when 25 ml of pentane at 40 psi was injected prior to the combustion. From these results, we decided not to use either decane or kerosene because they would have dissolved even more heavy ends and lead to less exothermic reactions.

Because we know that metallic additives enhance combustion, we did an in-situ combustion experiment with metallic additives. The results are presented in Figure 5.7. The metallic additives enhance the fuel deposition by creating activated sites on the rock for the fuel deposition (see section B.). HTO reactions were more exothermic as expected, but the temperature increase is not as significant in comparison to Hamaca oil. Metallic additives (iron in our case) did greatly enhance combustion; the heat created in the kinetic experiments was 3 times larger in comparison to the control test as shown in Table 5.1. Similar to the experiments with Hamaca, the oxygen uptake versus temperature was analyzed to obtain the activation energy for HTO of each of the tests above. The West Sak crude alone had an activation energy of 62 kJ/mol, whereas the value decreased to 60 kJ/mol and 58 kJ/mol for tests of the hydrocarbons remaining after extraction and the test with metallic additive. The HTO activation energies for all tests with West Sak are less than the corresponding tests using Hamaca. This comparison suggests that HTO of West Sak crude is challenged by a lack of fuel and not necessarily overly large activation energies.

5.4 Results from Tube Experiments

Similar to the presentation of the kinetics results, we discuss first Hamaca and the West Sak results. Again, duplicate experiments are compiled in the thesis of Cristofari (2006).

5.4.1 Hamaca

Figure 5.8 presents O_2, CO, CO_2 and CH_4 concentrations during a tube experiment with Hamaca oil. The sequence of injection consisted of injecting 500 ml of pentane, recovering most of it and launching an in-situ combustion process. The combustion

was very successful because O_2 appeared very late and CO and CO_2 were produced throughout the experiment. These observations are consistent with effective combustion. The high content of CH_4 in the effluent gas is due to the presence of pentane in the tube. It has been cracked into methane. From Figure 5.9, temperature profiles during the combustion show very well the front propagation; we notice also the formation of the steam plateau at 130 °C ahead of the front. According to the steam tables, the saturation temperature is 160 °C at 80 psi. Light components present in the gas phase decrease the saturation temperature.

5.4.2 West Sak

Figure 5.10 presents the results from the tube combustion of West Sak oil without solvent injection. The measurement started when we switched from nitrogen to air injection. Based on the results, the experiment is divided in three parts: first from the beginning until approximately twenty minutes, second from twenty minutes to approximately seventy minutes, third, from seventy minutes until the end of the experiment. Air injection started a t = 0. In the first period, methane was produced due to the cracking of the oil during the pre-heating phase. Injected air had not yet reached the gas analyzer. A rapid increase in temperature, however, was observed at the top of the tube immediately after injecting air. The second period started when CO appeared in the effluent gas. It confirmed that the oil had ignited. Only a few minutes later, however, O_2 appeared in the effluent signaling that it was not being consumed in the combustion front. The oxygen flowed through the entire tube with only a small portion contributing to oxidation of the unaffected oil. LTO reactions were operative in the tube and we assume that the LTO reactions contributed to increased oil viscosity and density. After seventy minutes, the third phase began. Carbon oxide, CO and CO_2 concentrations in the effluent started to decrease steadily whereas O_2 increased. Furthermore, the decrease in CO and CO_2 reflect that the combustion front was extinguished. The combustion did not propagate and the combustion was not self-sustaining.

From the temperature profiles presented in Figure 5.11 we notice that after 27 minutes of air injection, the highest temperature was 420 °C which is great enough to trigger HTO reactions. The requirements to sustain the combustion were satisfied. That is, oxygen was present and the temperature was high. However, at 75 minutes, the front temperature had declined to 412°C, which is very close to the minimum required to launch HTO reactions. This minimum was determined by the kinetic experiments, above. At 120 minutes, the front extinguished and only LTO reactions occurred. We stopped the reaction by switching the injection to nitrogen.

Figure 5.12 presents the concentration profiles for the tube experiment where we injected 500 ml of pentane before the combustion. Air was injected at the same time as gas concentration measurements began. The results are consistent with the kinetic experiments. Pentane injection did not improve oil combustion properties. Without any previous pentane injection (Figure 5.10) there was only 3% of oxygen in the effluent gases after 55 minutes of air injection, whereas there was already approximately 10% of oxygen in the effluent gas after 55 min of air injection when 500 ml of the solvent pentane was injected. The combustion front extinguished 3 times faster in comparison to the no solvent injection case. No front

propagation was noticed after ignition and so we did not record any temperature profiles.

The kinetic experiments showed that 3 times more heat was created when metallic additives had been dissolved in the water. From Figure 5.13, we notice the greater temperatures reached during a combustion tube experiment with metallic additives and the characteristic temperature profile was formed. The combustion front and the steam plateau were well defined whereas nothing comparable happened without metallic additives (Figures 5.10 and 5.11). Although the initial combustion following ignition was enhanced, the combustion front did not propagate through the entire tube. The front extinguished at 20 cm from the top but still propagated twice as far as the front in the ``pure combustion" experiment. By comparing the times in the temperature profiles (Figure 5.11 and Figure 5.13) we notice that they are very similar. Thus, the combustion front with metallic additives propagated approximately twice as fast as the front in the ``pure combustion" experiment. In both cases, the heat losses were equal because the procedure was identical. Even though the addition of metallic additives should result in more total heat of combustion, and the front clearly propagated after ignition, not enough heat was released to sustain the process through the entire tube. Although heat losses in these experiments are relatively small, one ought to keep in mind that heat losses are even less significant in the field (Sarathi, 1999).

From the composition profiles presented in Figure 5.12 similar to the ``pure combustion" experiment, we distinguish three different phases occurring during the experiment. The first period started at t = 0 and finished at t = 60 min. It is marked by high oxygen consumption and significant production of carbon dioxide. The second phase stretched from t = 60 min to t = 107 min. Here, carbon dioxide production peaks and begins to decline. Third, from t = 107 min to the end of the experiment the consumption of oxygen declines as does the production of carbon dioxide. Overall, more oxygen was consumed in this experiment because O_2 appeared later in the effluent gases. This suggests that more fuel was deposited and the metallic additives did enhance the combustion process to an extent. The combustion was still not self-sustaining.

5.5 Summary

The primary observations are summarized as follows:

- Solvent injection prior to combustion recovers most of the light fractions of the oil by extraction. Thus, these valuable ends are not degraded by combustion.
- The most effective solvent tested, for the cyclic solvent-combustion process, was pentane.
- Injecting decane and kerosene, heavier hydrocarbonsolvents in comparison to pentane, prior to combustion displaced a greater fraction of crude oil heavy ends. Ultimately, the extraction of heavy components decreased the fuel deposition during combustion. In subsequent combustion tube tests there was little heat created during HTO reaction after injecting kerosene. In the case of crude oil that experiences excess fuel deposition during in-situ

combustion, solvents such as decane or kerosene injection should be considered before beginning in-situ combustion.

- In-situ combustion burns the heavy ends that are less markedly interesting and creates energy in the reservoir in order to upgrade the oil through thermal cracking (pyrolisis).
- In agreement with the literature\Mamora, in-situ combustion was successfully applied with the Hamaca oil. Furthermore, we were able to inject a large amount of pentane in the media, recover the diluted oil containing the crude oil light ends, ignite the heavy ends remaining in the media and create enough heat to conduct successfully in-situ combustion.
- West Sak oil did not present good burning properties and in-situ combustion could not be conducted efficiently. Little heat was created during HTO reactions. Prior pentane injection worsened the effect. On the other hand, metallic additives appeared to enhanced combustion properties somewhat, yet combustion was not self sustaining.

The cyclic injection of solvent followed by combustion has been presented above largely as an enhanced recovery process. If this combination cannot be successfully applied at large scale, combustion, or a combination of solvent injection and combustion, may be considered to clean the near wellbore region if asphaltenes precipitate preferentially and result in decreasing production rates.

An economic attractive feature inherent to this combination is the in-situ upgrading of heavy oil, not only by deasphalting resulting from the solvent injection but also by thermal cracking during the combustion. This results in the production of higher gravity oil and potential savings during transport and further processing of the oil, especially in refineries. This process could make heavy oil much more attractive by reducing the price differential between heavy and light crude and by enhancing production rates.

Table 5.1 Increase in temperature during HTO reactions of Hamaca crude oil in kinetics cell.

Experiment	Solvent Injected	Asphaltene solubility in solvent	Increase in Temperature (°C)
"Pure combustion"	No solvent	n/a	290
Pentane + combustion	Pentane (C5)	~0	250
Decane + combustion	Decane (C10)	+ (moderate)	150
Kerosene + combustion	Kerosene (C9-C16)	++ (significant)	0

Table 5.2 Table D.1 Increase in temperature during HTO reactions of West Sak crude oil in kinetics cell.

Experiment	Solvent Injected	Increase in Temperature (°C)
"Pure Combustion"	n/a	82
Pentane + combustion	Pentane (C5)	56
Metallic additives	n/a	225

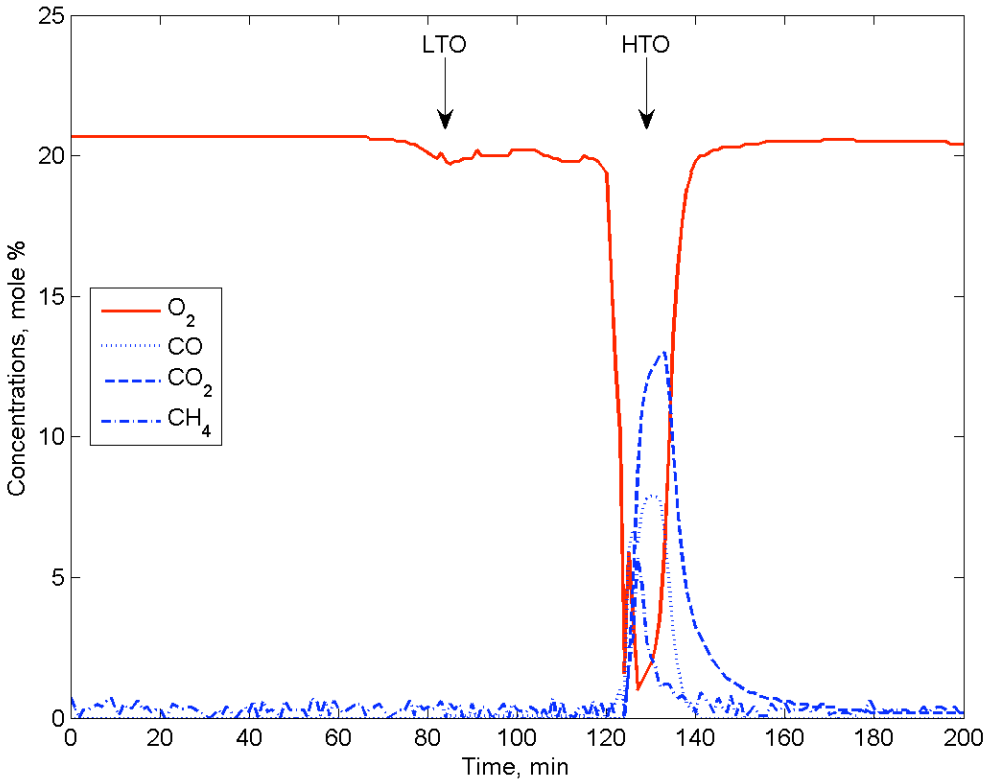

Figure 5.1 Concentration histories during kinetics testing of Hamaca crude oil.

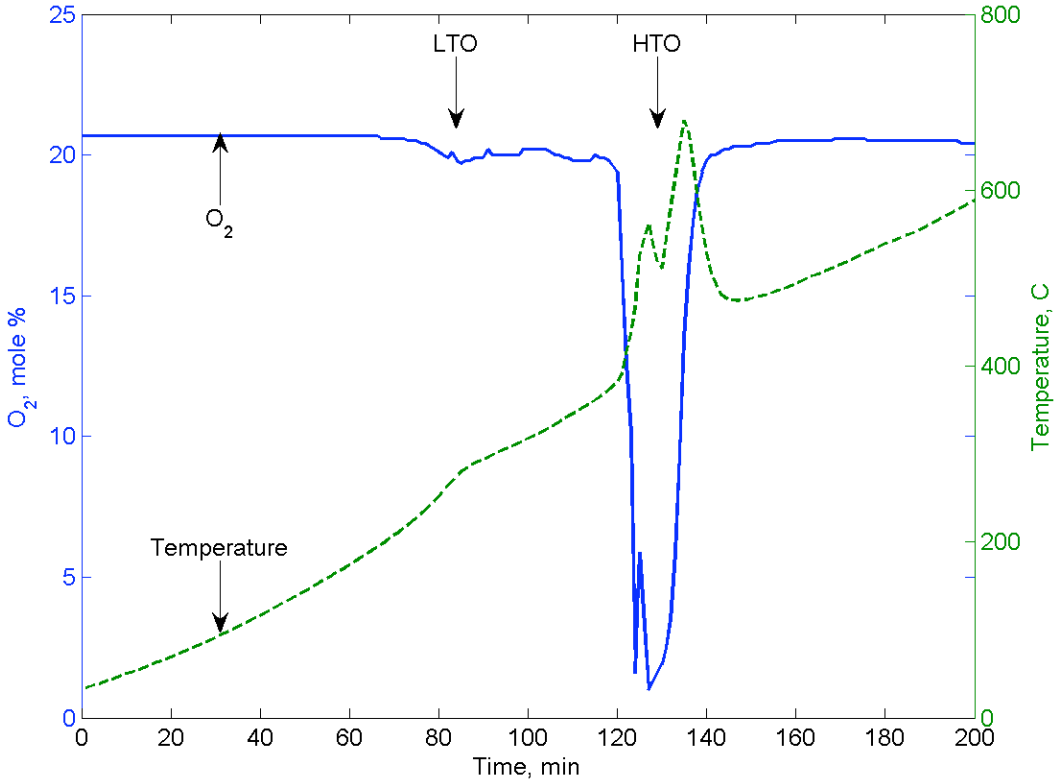

Figure 5.2. Concentration and oxygen temperature histories during kinetics cell tests of Hamaca crude oil.

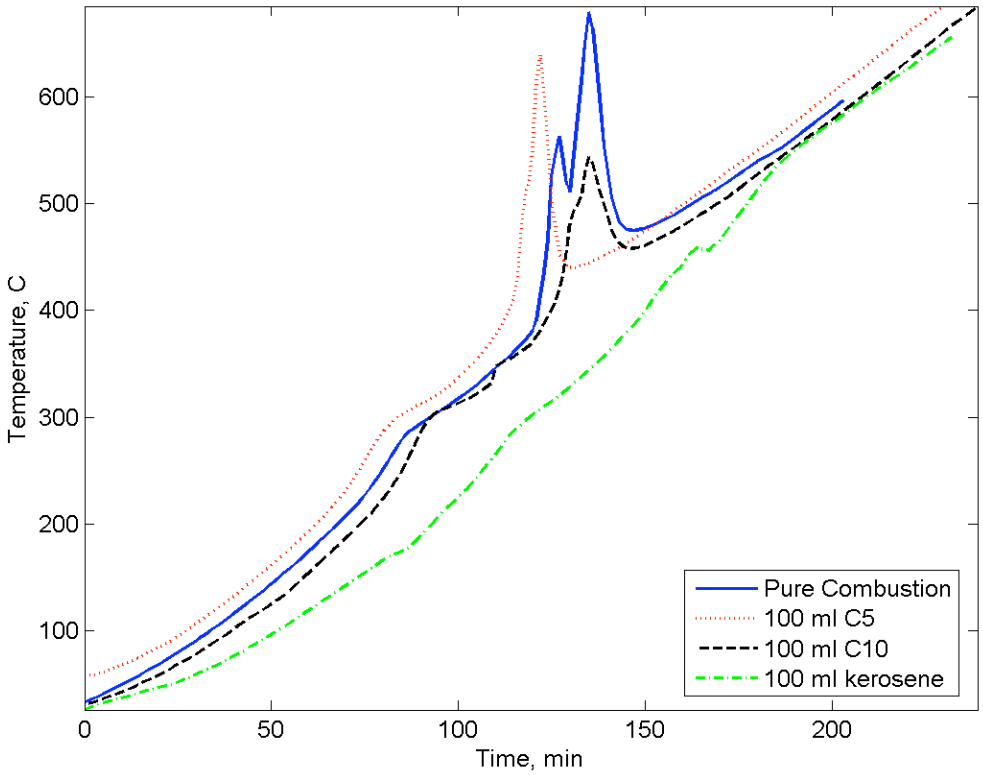

Figure 5.3. Temperature histories from kinetics cell testing of different solvents with Hamaca crude oil.

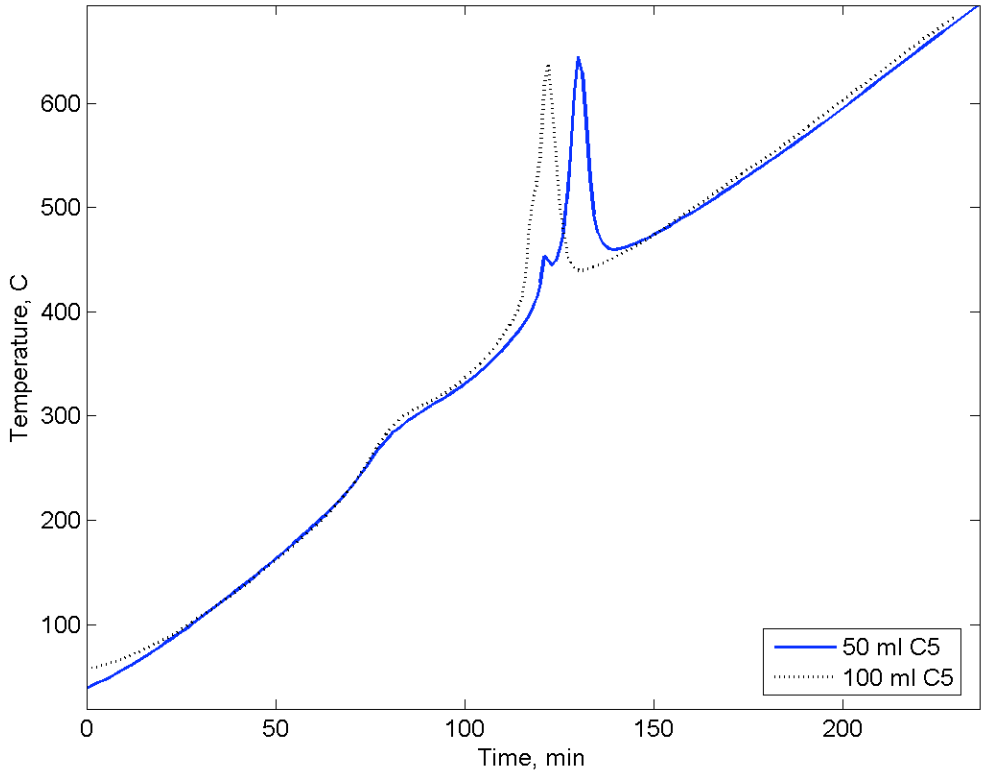

Figure 5.4. Temperature histories from kinetics cell tests after injecting various amounts of pentane into Hamaca crude oil.

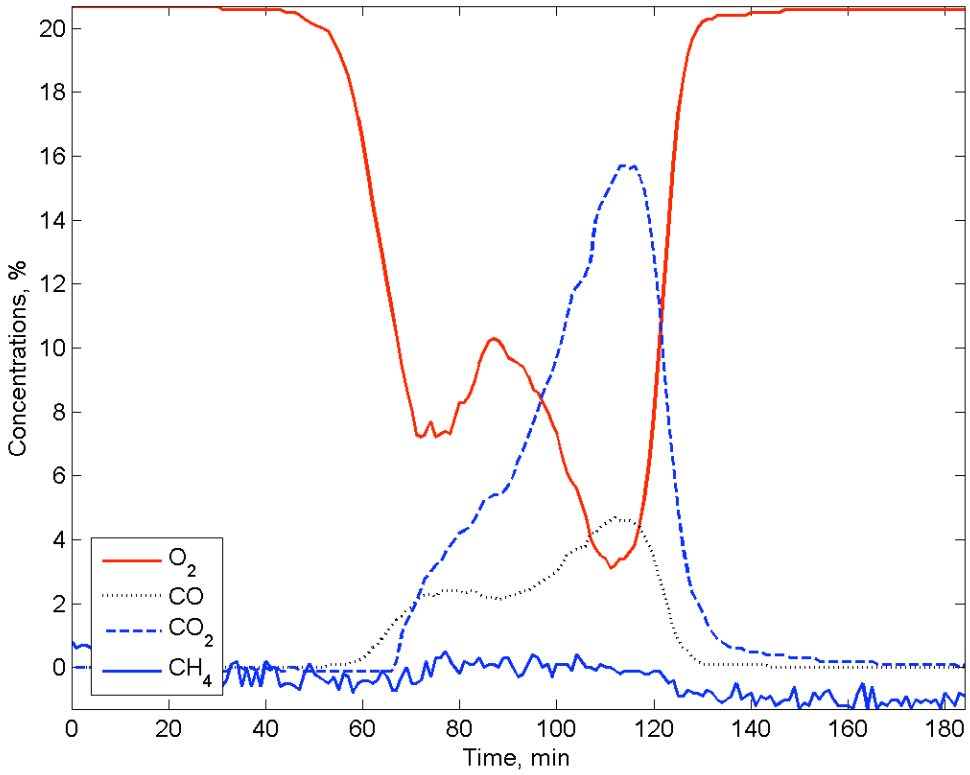

Figure 5.5. Concentration histories during kinetics cell testing of West Sak crude oil.

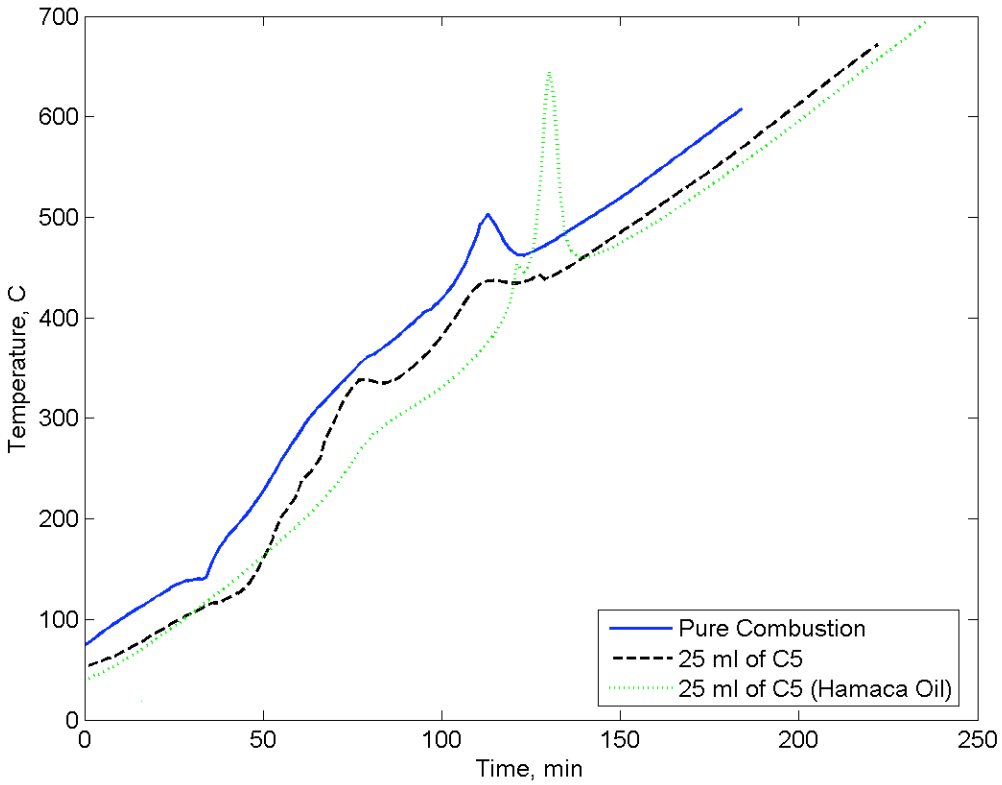

Figure 5.6. Temperature histories from kinetics cell testing of West Sak crude oil following pentane injection.

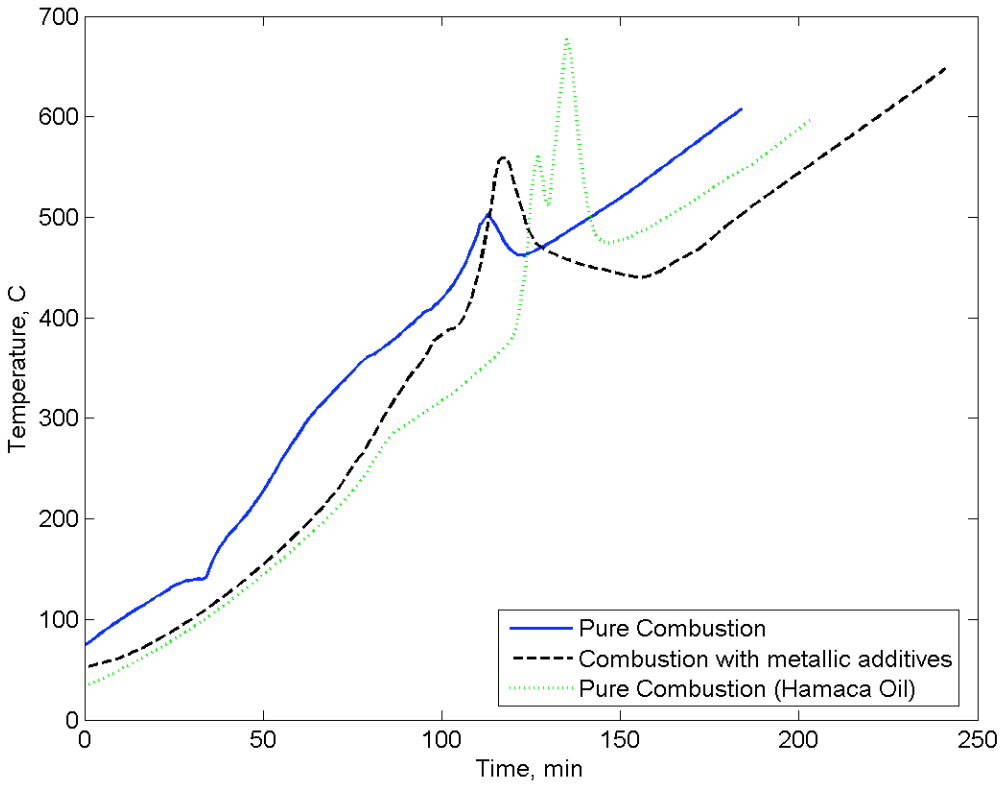

Figure 5.7. Temperature histories from kinetics cell testing of West Sak crude oil combined with metallic additives.

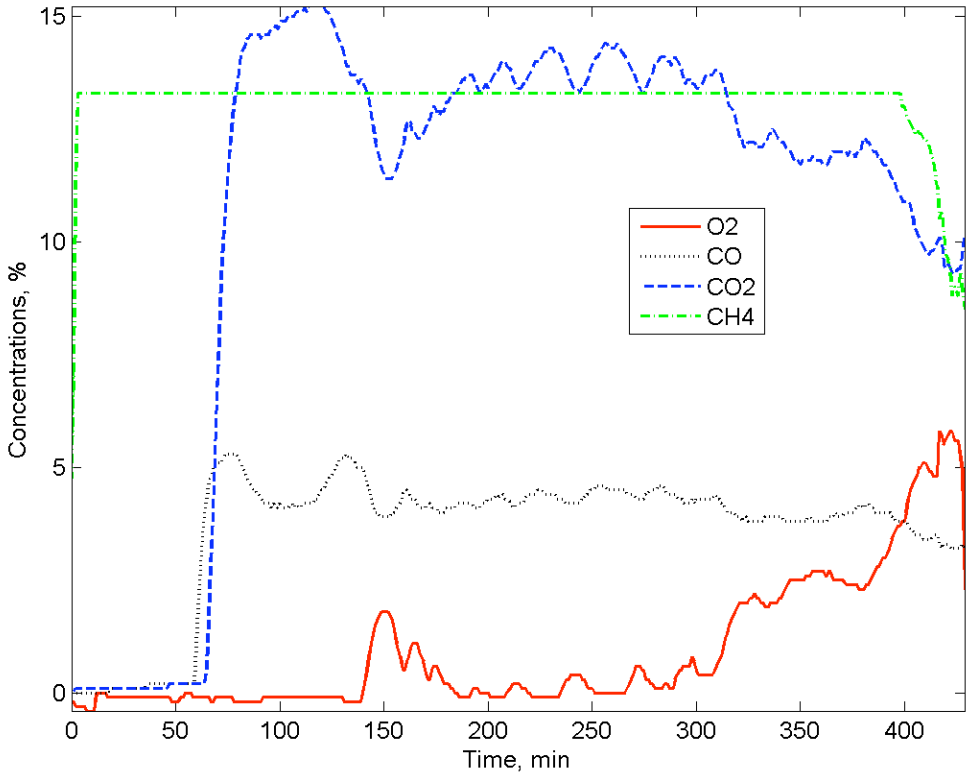

Figure 5.8. Concentration histories during combustion tube experiment of Hamaca crude oil.

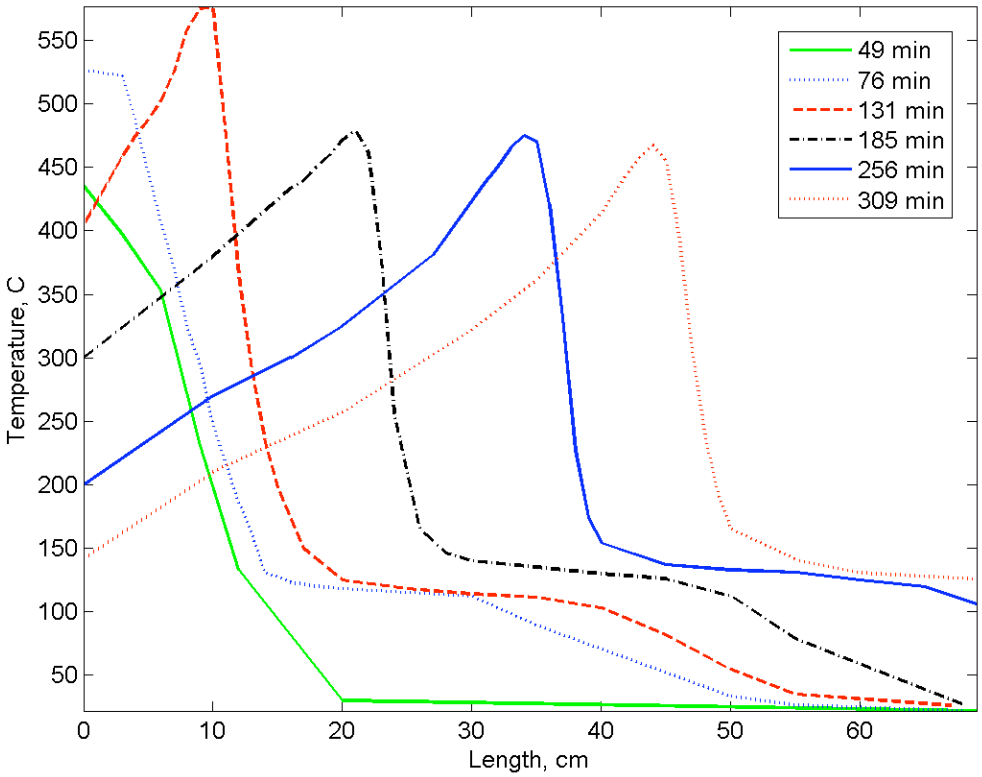

Figure 5.9. Temperature profiles during combustion tube experiment of Hamaca crude oil.

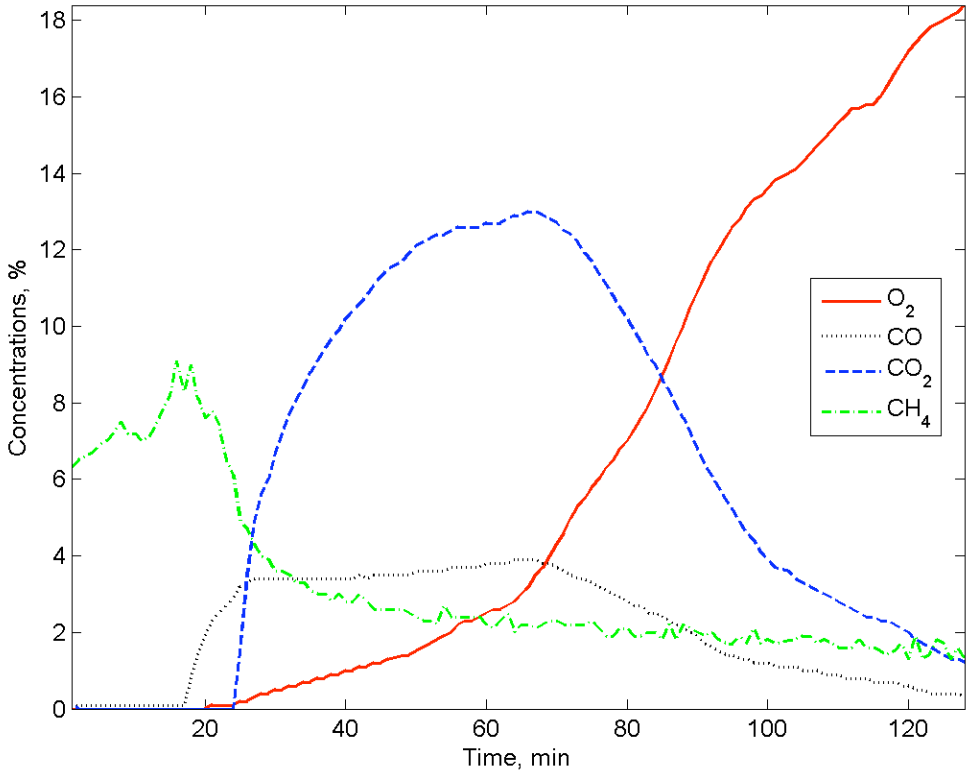

Figure 5.10. Concentration histories during combustion tube experiment with West Sak crude oil.

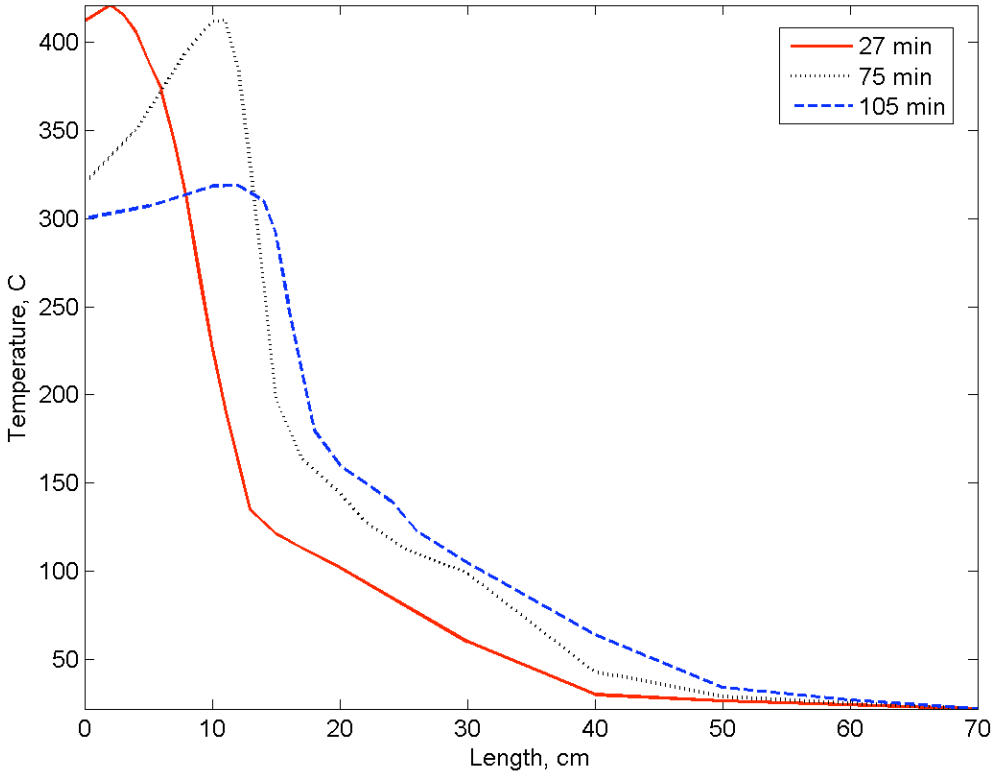

Figure 5.11. Temperature profiles during tube experiment with West Sak crude oil.

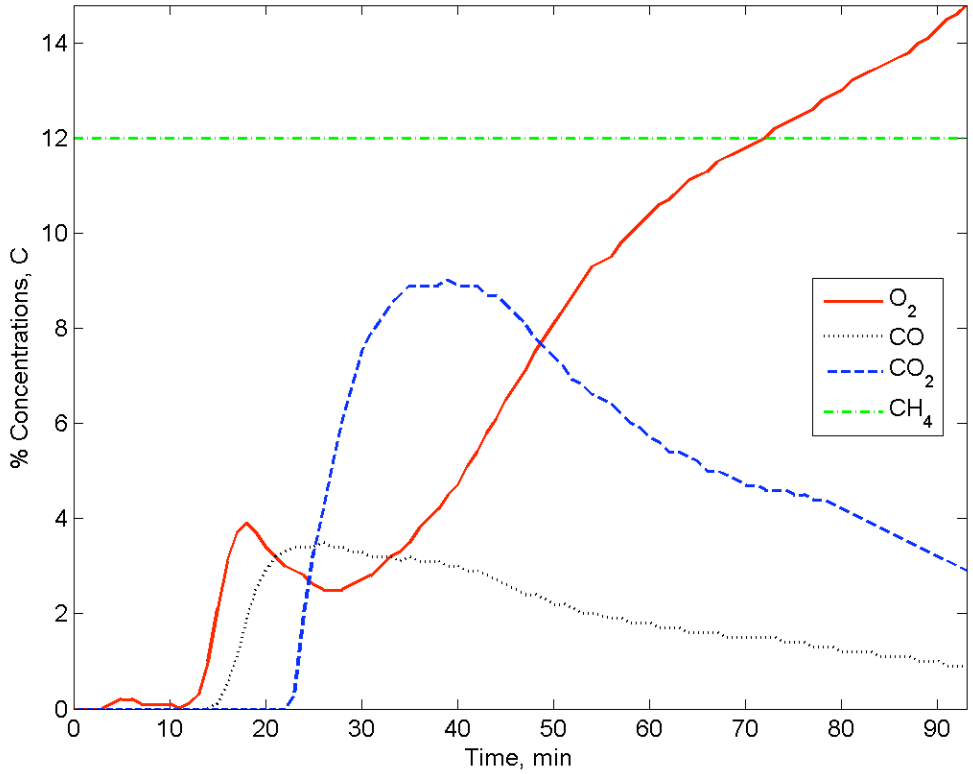

Figure 5.12. Concentration histories during combustion tube experiment with West Sak crude oil after injecting 500 ml of pentane.

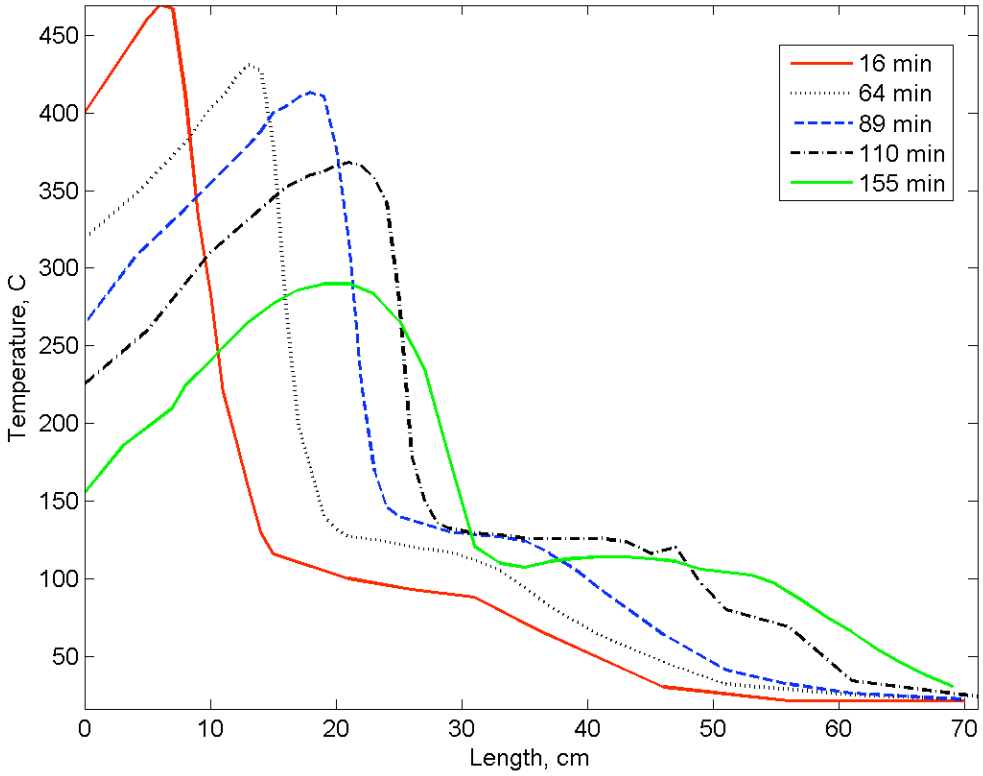

Figure 5.13. Temperature profiles during tube experiment with West Sak crude oil and metallic additive ($Fe(NO_3)_3$).

6. In-Situ Combustion in Heterogeneous Porous Media

An extensive literature exists on the nature of the reactions that occur when air comes into contact with crude oil during in-situ combustion (ISC); however the effect of reservoir heterogeneity on the ISC process has not been experimentally addressed in prior studies. This section presents the results of our experimental efforts to shed more light on ISC. We probe the effect of pressure, temperature, injection flow-rate, and matrix properties on the combustion of a 19.7 °API heavy oil in a tight-formation sand. The experimental perspective includes both oil oxidation kinetics and combustion-front dynamics in a 1-m long combustion tube. Most importantly, geological effects on the ISC process are investigated by incorporating various degrees of porous medium heterogeneity in the combustion tube runs. From the experimental data, pre and post burn CT scans of the combustion tube, and postmortem analyses, we deduce the effect of both small and large-scale heterogeneities on the spatial and temporal propagation of the combustion front.

Specifically, kinetic runs show typical ISC behavior consistent with oxidation and good reactivity of the oil-rock system. It was observed that the supply pressure, injection flow-rate, the reservoir matrix and a metallic additive significantly affected the performance of this particular system. Better combustion resulted from greater oxygen partial pressures and by using the reservoir rock as compared to an artificial sand-clay mixture. The combustion tube runs suggest that the process of in-situ combustion can be successfully implemented in a reservoir with small-scale heterogeneities (a few inches) and may be challenged by much larger scales (feet).

6.1 Introduction

A better understanding of the role of heterogeneity as it relates to fluid flow improves the efficiency of hydrocarbon extraction processes. Despite many decades of laboratory experimentation and some successes in the field, in-situ combustion applications have not increased (Sarathi, 1998). Aging fields, heavy oil, oil shales and dwindling prospects for finding new, significant reserves may, however, lead to a renewed and vigorous interest in the process of in-situ combustion.

In-situ combustion (ISC) is the oldest yet one of the least understood of the existing and proven enhanced oil recovery processes. The research reported here was conducted by laboratory experiments of dry, forward combustion processes. Fig. 6.1 indicates a number of identifiable zones at a given instant during an in-situ combustion process (Castanier and Brigham, 2007). Of the zones in Fig. 6.1, the most important to the research presented here is the combustion-front (zone 3). The ISC process is effective as long as the combustion front is exothermic and in motion. By tracking the spatial and temporal propagation of such a fully-formed burning front in representative laboratory models, we probed the effects of geological heterogeneity on the ISC process.

The illustration of the different zones as shown in Fig. 6.1 gives much insight into the combustion process but it is not necessarily a true representation of what happens in the field. The heterogeneous nature of the reservoir causes overlaps between zones and may make their relative locations and sequence of occurrence

slightly different than as described (Sarathi, 1998). Compared to other recovery techniques, in-situ combustion involves the additional complexity of exothermic chemical reactions, temperature-and pressure-dependent reaction kinetics as well as an oxygen-transport-dependent burning front. All of these factors may be influenced by heterogeneity.

6.2 Experimental

For this study, two levels of heterogeneity are considered: small-scale (a few inches) and large-scale (feet). The distinction is drawn from the degree of randomness. A highly-random setting, with geological changes observed every few inches, is classified as small-scale. If some significant lateral or vertical extent is described and some continuity or consistency, of about tens of feet, observed in the unconformity, then it is classified as large-scale heterogeneity. These scales assume typical pay-zone thicknesses of tens of feet and lateral extents of hundreds to thousands of feet. The two pre-defined spatial scales are represented as described in the following section.

To represent heterogeneities on the smaller scale, a porous medium was packed into the combustion tube in the manner illustrated in Fig. 6.2. The remainder of the experimental apparatus is as illustrated in Fig. 2.1. A region of constant permeability adjoins another region consisting of a random distribution of less-permeable "chunks" of reservoir cores embedded in a high-permeable matrix of crushed core material.

The small-scale heterogeneous section is an idealized dual-porosity medium. The less-permeable chunks represent the matrix of high porosity but low permeability and the high-permeable crushed material represents the fracture of much higher permeability in a typical dual-porosity medium. The average size of the chunks is one inch as compared to the combustion tube of three inches inside diameter. The total volume taken in by the chunks is an estimated 0.038 ft^3 as compared to 0.072 ft^3 for the volume of the section of the combustion tube containing the chunks.

The line of inquiry remains the fate of the combustion front as it travels from left to right in such a geological setting. To represent heterogeneities on the larger scale, the porous medium was packed as illustrated in Fig. 6.3. A single region of constant permeability except for a less permeable core is arranged in such a way as to allow for a fracture, or a thief zone, in vertical alignment with it generating a zone of contrasting permeabilities as may be found in layered porous media or fractured reservoirs.

6.2.1 Materials

The specimens used for the experiments include field core and a 19.7 °API crude oil from a moderately tight sandstone reservoir. The core samples have significant residual hydrocarbon on them. For comparison purposes, a 16°API crude from the Wilmington field in Southern California is also used.

An analysis of the reservoir grains from the 19.7 °API reservoir was performed using SEM. First, residual hydrocarbon was extracted with toluene. The analysis indicates that it is a poorly-sorted, silt-to-coarse grained silicate sand. This

sand is composed of quartz, potassium feldspar ($KAlSi_3O_8$), and other heavy minerals, particularly pyrite. This characterization of the grains is drawn from the information illustrated in Figures 6.4 to 6.8. Figures 6.4 and 6.5 show SEM images and graphs that indicate clearly the presence of quartz and potassium feldspar. Figure 6.7 shows a quartz grain and an unidentified material with great angularities. And lastly, Figure 6.8 shows a pyrite grain at magnification 500X.
This unidentified material is possibly calderite (a garnet) or pumpellyite.

Other solid material includes Berea sandstone core, clean, graded, kiln-dried sand that is 60 mesh in size, and clay, in the form of kaolinite. House distilled water and Ferric Nitrate, 9-Hydrate, Crystal, Fe(NO3)3.9H2O are also used.

6.2.2 Samples for Kinetic Testing

The kinetic runs were performed on mixtures of sand, kaolinite, oil, and water or on crushed reservoir material. "Mixture", as will be used from here on, refers to the mixture of sand, kaolinite, water, and oil that simulates an oil-bearing porous medium. The mixture was made in proportions of 85g sand, 45g clay, 4g water and 4.5g – 12g crude oil, and 0.4g metallic additives, if any. The range for the crude oil content depends on the required oil saturation.

6.2.3 Combustion Tube Runs

The packing material (either the mixture or the crushed cores) for the combustion tube was prepared in the same proportions and manner as for the kinetics cell except for the larger quantities needed. The thermal well was vertically placed in the center of the tube before packing. The mixture or crushed reservoir material was packed into the combustion tube. About 5ml of linseed oil was added to the top of the pack, at the location of the heater, to enhance ignition. On the top of the pack, about 5cm of clean, coarse sand was added to help uniformly distribute the injected gas. The bottom was also packed with another 5cm of coarse sand to facilitate production.

One aspect of this study that differs from previous laboratory examinations of in-situ combustion via tube runs is that we combined pre and post combustion X-ray CT scanning of the tube. Scanning allowed us to visualize the position and distribution of heterogeneities in the combustion tube as well as to understand details of combustion in low permeability versus greater permeability heterogeneities.

The CT scanner used is a Picker 1200 SX fourth generation scanner. The spatial resolution is 0.35 by 0.35 by 5 mm. Raw CT images are analyzed without further image processing.

In all, a total of 10 experiments were conducted: 7 kinetic experiments and 3 tube runs. The type, nature and number of experiments performed depended on the availability of resources, particularly field samples. Table 6.1 gives a summary of the conducted experiments.

6.3 Experimental Results

6.3.1 Kinetics

Data analyses of the results from ramped temperature oxidation of the experimental samples under the operating conditions tabulated in Table 6.1 yielded the following. Experiment #1, was carried out on a mixture of sand, kaolinite, 19.7°API oil and water. The air supply pressure was 50psig. The effluent gas profiles for the mixture are presented as gas composition versus temperature in Fig. 6.9.

The data for oxygen consumed implies that low-temperature oxidation (LTO) occurs at about 350°C and high-temperature oxidation (HTO) at about 485°C. An analysis of the effluent gas data using the Arrhenius equation (see sections 3 and 4) and numerical integration results in an E/R (ratio of activation energy to gas constant) value of about 4500K for the LTO and about 6700K for the HTO. These values, as well as the gas profiles, are consistent with others in the literature for viscous oil (Bousaid and Ramey, 1968; Fassihi et al, 1984; He et al, 2004).

Experiments #2, 3 and 4 were performed on the crushed reservoir cores from the 19.7 °API oil reservoir at supply pressures of 50psi, 200psi and 290psi and air flux of 0.5SLPM, 0.5SLPM and 1.5 SLPM, respectively. Analyses of the results show that the oxygen partial pressure and air flux affect the kinetics of crude oil. The oxygen consumption profiles for the different experiments are summarized in Fig. 6.10.

The data for oxygen consumed at 50psi reaches a plateau around 350°C. This likely arose because the air flux was low at 0.5 SLPM considering the large percentage of combustible hydrocarbons present within the crushed cores. At the lower partial pressure was increased from 50 to 200 and then to 290 psi and the air flux from 0.5 SLPM to 1.5 SLPM, the oxygen uptake reduced accordingly, until an exhibition of the classic kinetic profile for combustible viscous oils occurred at the highest pressure value of 290psig and air injection rate of 1.5L/min. These results indicated that these samples required greater pressure and air flux for satisfactory combustion because of the significant amount of fuel in the matrix.

The results from experiments #1 and 2 are used to compare the effect of the reservoir matrix on oxygen consumption. Experiment #1 was performed using clean sand and kaolinite as the matrix while the native reservoir material was crushed and used in Experiment #2. Both experiments were conducted under similar operating conditions of pressure and air flux. The oxygen-uptake profiles for the two experiments are shown in Fig. 6.11.

Figure 6.11 shows very high oxygen uptake and ultimately incomplete combustion for the crushed reservoir material and less oxygen uptake with an exhibition of the classic kinetic profile and a clean, complete combustion for the mixture of sand, clay, oil from the same field and water. This occurred because of the differences in the molecular make-up and structure of the reservoir matrix and the artificial mix as well as the differences in hydrocarbon content.

Figures 6.12 and 6.13 summarize the post mortem analysis. Fig. 6.12 indicates a clean burn in the case of the mixture and Fig. 6.13 shows incomplete combustion with coke-type particles in the residue in the case of the crushed reservoir material.

These results suggest a possible inherent catalytic capability of the native reservoir material and/or additional hydrocarbon deposits residual on the reservoir cores resulting in a greater hydrocarbon content and consequently, greater oxygen uptake.

Toluene was then used to extract all of the hydrocarbons on the crushed reservoir cores. The resulting mixture of grains (referred to as "washed core") was placed in the kinetics cell, and the grains tested to find if the matrix had any reactive capability. This experiment showed no oxygen consumption, confirming that the grain material alone was inert. The small amount of CO_2 produced, as shown in Fig. 6.14, resulted from small amounts of toluene that remained on grains following the washing procedure.

The CO_2 production profiles, Fig. 6.14, suggest that the residual organic phase present on the grains in addition to the oil in the reservoir matrix brings about greater oxygen consumption and CO_2 production during the prominent reactions. Tests with reservoir material showed a far greater CO_2 production than that with the mixture.

Experiments #3, 5 and 6 were used to investigate the effects of oil saturation and metallic additives on the kinetics by evaluating the activation energies for the LTO and HTO reactions. The values of E/R reported in Fig. 6.15 are generally consistent with those found in the literature for combustible viscous oils in porous media (Bousaid and Ramey, 1968; Fassihi *et al*., 1984; He *et al*., 2004). It is interesting to note the trends in the obtained values. By increasing the oil content and adding a metallic additive ($Fe(NO_3)_3.9H_2O$) the activation energy for the high temperature oxidation increased quite significantly, and the activation energy for the low temperature oxidation decreased, although to a lesser degree. This can be attributed to the catalytic effects of these additives on the reactions. Due to the importance of these effects, it is important that any laboratory study of in-situ combustion should be made on the native core material to evaluate the feasibility of this process in the field as the core material itself may contain some metallic ions.

Experiment #7 was performed on the crude from the Wilmington field. This was done to establish baseline burning characteristics and suitability for combustion-tube testing to verify that a combustion front could be propagated through this particular sand/kaolinite/oil and water system. The results as presented in the gas profiles in Fig. 6.16 indicate typical kinetic profiles with well defined LTO and HTO peaks and fine burning characteristics.

E.3.2 Combustion Tube

After performing the kinetics experiments, combustion tube runs were used in the investigation of geological effects on the process of in-situ combustion. Three runs, represented by experiments 8, 9 and 10, were performed as shown in Table 6.1. The makeup of the porous media and oil contents mirrored that of the kinetics (RTO) tests. In Experiment #8, the combustion tube was packed with crushed reservoir material only. Initially, nitrogen was injected to achieve gas communication across the tube and the temperature was increased to about 400°C at the inlet. Then injection was switched to air at a supply pressure of 200 psi. The air (20.7% O_2) flux was 3 SLPM. As soon as air injection began, the heater was turned off.

The temperature profiles along the pack and the effluent gas profiles are shown in Figs. 6.17 and 6.18. As long as carbon oxides are produced (and oxygen is consumed), there is an active and mobile combustion front. The motion of the peak temperatures in each temperature profile implies that a combustion front is traversing the medium. A smooth, well behaved propagation of the combustion front, as illustrated by the temperature profiles in Fig. 6.17, is typical of a successful combustion tube experiment conducted on a uniformly packed porous medium.

Experiment #9 was made with the objective of investigating the effect of small-scale heterogeneities on the process. The design was fashioned after the model described for small-scale heterogeneous media (Fig. 6.2). It was also designed to reduce oil banking and allow increased air flux in the system if such a need arose. The packing of the tube is unique: the first 10 cm of the core adjacent to the igniter were packed with crushed reservoir material only, the next 16 centimeters were packed with oil/water/sand and kaolinite mixture similar to that used in experiment 8, the next 60 centimeters were packed with pieces (chunks) of reservoir core with crushed core/oil mixture packed around the pieces, the bottom of the tube was packed with coarse, quartz sand to facilitate oil production. The sections containing reservoir material were 97.1% rock and 2.9% oil by weight. The initial oil saturation (with no consideration of residual hydrocarbon on core material) was roughly 0.23. After ignition, the air flux was set to 3 SLPM. After 60 min, the air flux was increased to 6 SLPM for the duration of the experiment.

The effluent gas production is presented in Fig. 6.19. The variations seen on the gas composition plot are directly related to the air flux, namely when the air rate was low, the amount of oxygen consumed decreased as did the production of carbon oxides. The reverse behavior was experienced when the air flux was increased. Some of the flat portions in Fig. 6.19 arose because gas flow to the gas analyzer was suspended periodically when the pressure drop across the analyzer exceeded 3psi. After relief of pressure, flow to the analyzer was resumed.

The temperature data are summarized on Fig. 6.20. The heterogeneities in the matrix did not seem to affect the front movement. Despite the scattered nature of the peak temperature on Fig. 6. 20, the combustion front did progress. Note the progression of peak temperature between 244 and 329 minutes.

Another observation is that the amount of air consumed is quite large. The air used is estimated as follows for experiment 9. At say 45 cm from the ignition point, the volume swept is the surface of the tube (44 cm2) multiplied by the distance traveled by the front (45 cm). Hence the volume is 1980 cm^3. The amount of air injected to this point is 1.74 x106 cm^3. The average concentration of O_2 in the effluent gas is 4%, by mole. Hence, approximately 16.7 % of O_2 is utilized. This equates to an air requirement of 860 standard cubic meters to burn a cubic meter of rock. This is substantially more air than usually needed for oils of similar gravity. For instance, the South Belridge field results showed air requirements of about 117m^3 of air per cubic meter of reservoir (Kok, 2001). This large air consumption likely resulted from the presence of significant residual hydrocarbon on the reservoir matrix provided for the tests. This extra hydrocarbon results in a greater amount of fuel that in turn consumes a larger volume of oxygen for proper combustion.

Digital images and CT scans of various sections of the packed tube taken before and after the tube runs also tell the same story, as shown in Fig. 6.21. Fig. 6.21 is divided across the middle by a picture of the combustion tube. The images on the left represent the experimental samples before the run at their approximate position along the tube. The images on the right represent the post-burn samples. Images as viewed from top to bottom also represent the vertical (top-to-bottom) packing of the combustion tube. The arrow on the image of the combustion tube indicates the direction of air flux and fire front travel.

CT scans show the arrangement and the orientation of packed samples within the tube and indicate that this orientation is relatively unaffected by the ISC process, Figure 6.22. Pre-burn digital images of the top-section of the tube show porous media with significant oil saturation and a constant, homogeneous permeability of about 4 D. The lower section contains less permeable chunks (approximately 300 md) of reservoir material interspersed with the crushed material, as found in the upper section. Also, the image on the lowest left corner is that of clean, coarse, quartz sand included at the very bottom of the combustion tube to facilitate oil production.

Post-burn digital images from the post-mortem analyses indicate a nice propagation of the combustion front through the homogeneous section as well as undeterred travel through the dual-porosity section identified by the samples in that region being drained of oil. The image in the lowest right corner is that of the quartz sand after the passage of produced oil.

Experiment #10 was made with the objective of investigating the effect of large-scale heterogeneities on the ISC process. The design was fashioned after the model described above for large-scale heterogeneous media (see Fig. 6.3). The packing of the tube was also unique: the first 15 cm of the tube adjacent to the igniter was packed with oil/water/sand and kaolinite mixture, the next 60 centimeters was packed as two different vertical sections – one half of this section was filled up with the mixture and the other half with a less-permeable (700md) semi-cylindrical, oil-saturated Berea sandstone core of 60cm in length and 3.8cm radius. The remaining 15cm was then filled with the mixture again and the bottom packed with coarse quartz sand as before.

The semi-cylindrical sandstone core could have been cemented to the inner walls of the tube to create only two vertical regions of contrasting permeabilities. However, for the purpose of this experiment, an intentional spacing was allowed as a simple representation of a fracture. This helped in introducing the large-scale heterogeneity that we wished to explore. The oil saturation was about 0.3 on the average and the supply pressure was 200psi. The air flux was about 4SLPM.
The effluent gas production is presented in Fig. 6.23.

The effluent gas data show that the production of carbon oxides started declining at about 70mins after ignition. Similarly, the oxygen content at the exit started increasing at about the same time. The increase was continuous until the end of the experiment. The plots indicate the increasing levels of oxygen bypass with time. The temperature data, summarized in Fig. 6.24, show peak temperatures decreasing from a value of about 520°C at 37mins after ignition to 480°C at 70mins and about 390°C at 99mins. The peak temperatures reduced at about the same time

the gas profiles indicated increasing levels of oxygen bypass. This suggests that as the fire front approached the contrasting permeability interface, increasing quantities of air quickly passed through the highly permeable fracture and appears to have caused the combustion front to be starved of the much needed oxygen. This appears to have led to the effective extinction of the combustion front.

As before, the images in Fig. 6.25 are divided across the middle by a picture of the combustion tube. The images on the left represent the experimental samples before the run with the images placed at roughly their position in the tube. The images on the right represent the post-burn samples. Images as viewed from top to bottom also represent the vertical (top-to-bottom) packing of the combustion tube. The arrow on the image of the combustion tube indicates the direction of air flux and fire front travel.

Pre-burn digital images of the top-section of the tube show porous media with significant oil saturation and a constant, homogeneous permeability of about 4 D. The lower section consists of two different vertical zones – one half, the mixture similar to that found in the upper section and the other half, the less permeable Berea sandstone core. The arrangement leaves a fracture as a third zone of high permeability in the permeability contrast region. Post-burn digital images indicate a nice propagation of the combustion front through the homogeneous section, evidence of combustion activity at the interface of the heterogeneous zone, but an undisturbed region afterwards. The combustion front did not propagate much farther than the interface of different permeabilities (the zone containing the large scale heterogeneity).

6.4 Summary and Conclusions

This study investigated the effect of geological features, mostly in the form of varying scales of heterogeneity, on the process of in-situ combustion.
The research involved carrying out a variety of experiments in a systematic manner. Kinetic experiments were first conducted for the two oil samples contained within either a mixture of sand and kaolinite or in their native reservoir matrix. We probed the effects of pressure, flow rate, metallic ions and the matrix on the nature of the reactions that occurred. Combustion tube experiments were then conducted to investigate the effect of geological heterogeneities on the ISC process by monitoring the spatial and temporal propagation of a fully developed combustion front, in porous media with two different scales of heterogeneities, defined as small (a few inches) and large scale (feet).

The results of this study are summarized as follows:

1. Novel methods for introducing heterogeneity into combustion tube experiments were investigated and employed. Heterogeneity is introduced by the layering of zones of different permeabilities and by the distribution of less permeable material in a more permeable matrix. In general, this adds to the knowledge base of in-situ combustion and in particular, to the limited methods of incorporating heterogeneity into laboratory experiments.
2. The unhindered travel of the combustion front through the dual porosity section of the small-scale heterogeneous set-up indicates that in-situ

combustion can be successfully implemented in a reservoir with small-scale heterogeneities.

3. The ISC process may be challenged by much larger scales of heterogeneity as a result of possible bypassing of the much needed air through the thief zones effectively starving the burning front and extinguishing the process. This is an area where more research is warranted.
4. Ramped temperature oxidation studies of the two geographically-separated viscous oils used indicate favorable oxidation behavior consistent with the combustion of heavy oil.
5. There are conditions of pressure and air flux below which the classic kinetic profiles with well-defined peaks of LTO and HTO are not observed, and only what can be described as plateaus of the reactions occur, leading to very significant oxygen uptake and incomplete combustion. These conditions depend on the oil saturation and the reactivity of the matrix.
6. The inclusion of hydrated iron nitrate ($Fe(NO_3)3.9H_2O$) as a metallic additive to the crushed Middle-Eastern reservoir material reduced the activation energy for the LTO reaction and increased the HTO activation energy. This is a characteristic signature of an effective metallic ion. Further analysis, such as combustion tube runs, may be conducted with the inclusion of this metallic compound to check on the feasibility of using it in field pilots.
7. The matrix, in whose pore space the oil is contained, has been shown to be very critical in the assessment and study of an ISC project. Tests are better if performed both on the native reservoir material as well as on an artificial mix of porous media. This is because of the differences between the two, as regards composition, catalytic ability, and hydrocarbon content.

6.5 Recommendations

The study presented in this report includes quantitative work on the kinetics of crude oil combustion in porous media and a more general, qualitative description of what can happen to an ISC process conducted in two broad categories of different geological settings. More quantitative analysis of these generalized descriptions should be the future lines of inquiry. Steps to be taken in that direction include answering the following questions:

1. Exactly how large is large-scale and how small is small scale, and even how random is random? More quantitative analyses of the extent and scope of heterogeneities and their effects on the ISC process should be performed. Dimensions (or dimensionless quantities) and ratios would be much better tools in assessing the feasibility of ISC on a given geology.
2. How much of a contrast in permeability actually challenges the ISC process? This may be investigated by using layered porous media settings of the classic Dykstra-Parsons style (Lake, 1989) to check the possibility of the decoupling of the combustion front, its extinction or the possible coherence of the fronts within the different sections as proposed by Akkutlu and Yortsos (2005).
3. The possibility of the formation of multiple fronts should also be researched. In adiabatic systems, bypassed air may lead to the spontaneous ignition of

previously undisturbed regions within the reservoir, generating more fire fronts and resulting in some regions of the reservoir being swept and others completely bypassed.

4. More elaborate, three-dimensional laboratory and field scale reservoir simulations of the ISC process in different geological settings should be conducted. History-matching of laboratory experiments is of paramount importance. If inconsistencies persist, simulation codes may be revisited to incorporate experimentally-observed phenomena and hence allow for more accurate simulation of this highly-complex thermo-chemical process.

Table 6.1 Summary of experimental conditions.

Expt. #	***Expt. Type***	***Sample***	***Operating conditions***		
			Pressure	***Air Flux***	***Back pressure***
1	Kinetics	Mixture, using 19.7·API crude	50psig	0.5SLPM	40psig
2	Kinetics	Crushed reservoir cores from the 19.7·API crude reservoir	50psig	0.5SLPM	40psig
3	Kinetics	Crushed reservoir cores from the 19.7·API crude reservoir	200psig	0.5SLPM	150psig
4	Kinetics	Crushed reservoir cores from the 19.7·API crude reservoir	290psig	1.5SLPM	150psig
5	Kinetics	Crushed reservoir cores from the 19.7·API crude reservoir	200psig	0.5SLPM	150psig
6	Kinetics	Crushed reservoir cores from the 19.7·API crude reservoir+ Oil from the same field + Metallic additive	200psig	0.5SLPM	150psig
7	Kinetics	Mixture, using the oil from the Wilmington field, California	50psig	0.5SLPM	35psig
8	Combustion tube	Homogeneous pack of crushed reservoir cores from the 19.7 °API crude oil reservoir	200psig	3SLPM	100psig
9	Combustion tube	Heterogeneous pack, as described in small-scale model (section), using field cores 19.7 °API crude oil reservoir and a sand/clay/water and 19.7 °API crude oil mixture	200psig	3-6SLPM	100psig
10	Combustion	Heterogeneous pack, as described in	100-200psig	3-6SLPM	100psig

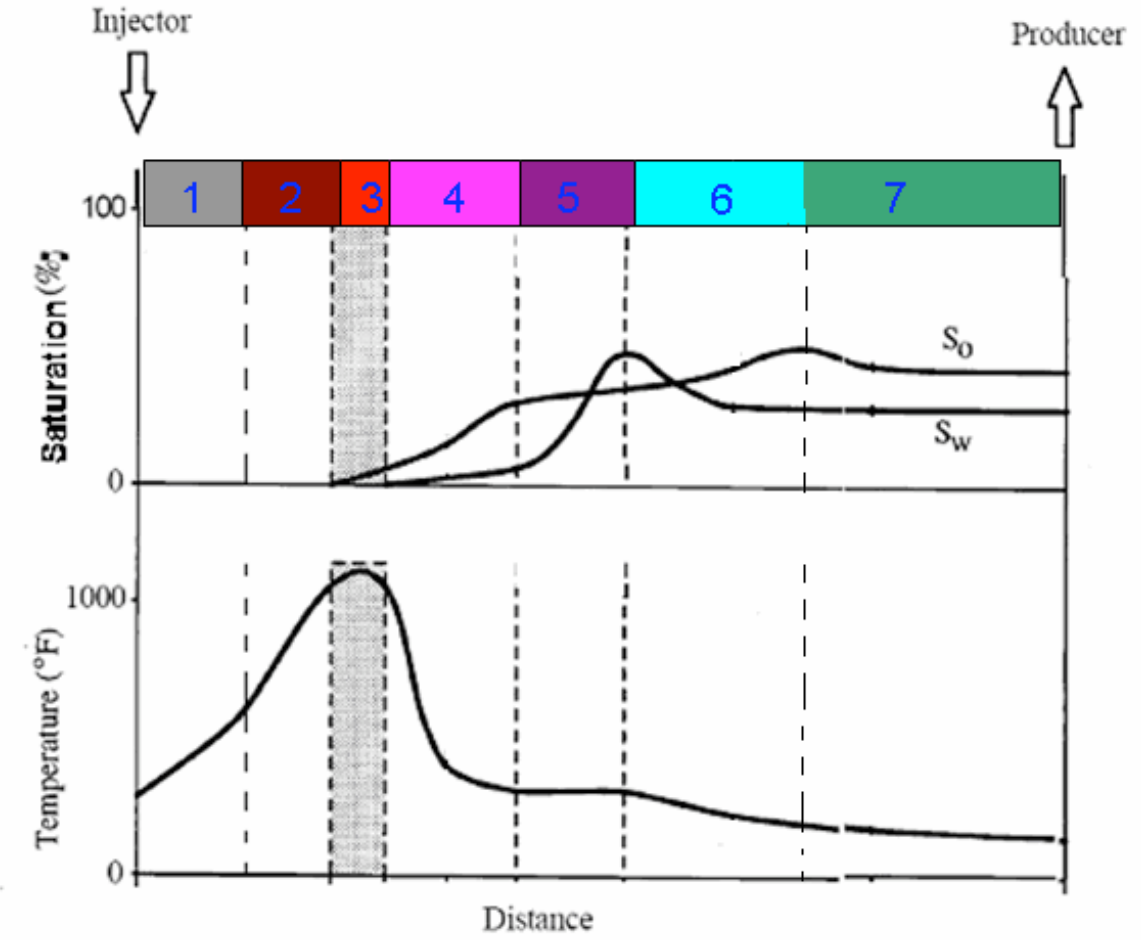

Figure 6.1. Illustration of temperature and saturation profiles in a typical in-situ combustion process.

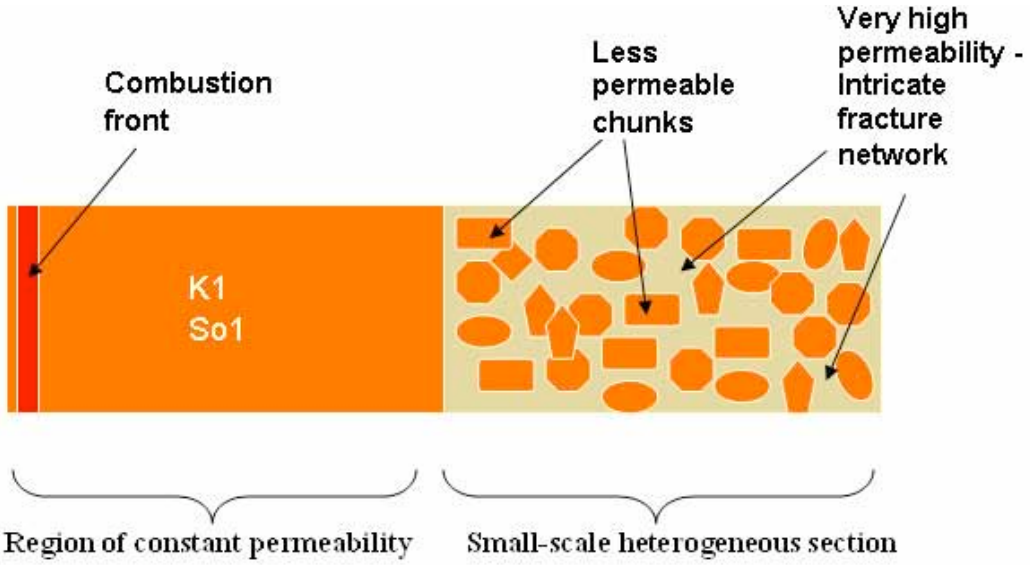

Figure 6.2. Experimental representation of small-scale heterogeneities.

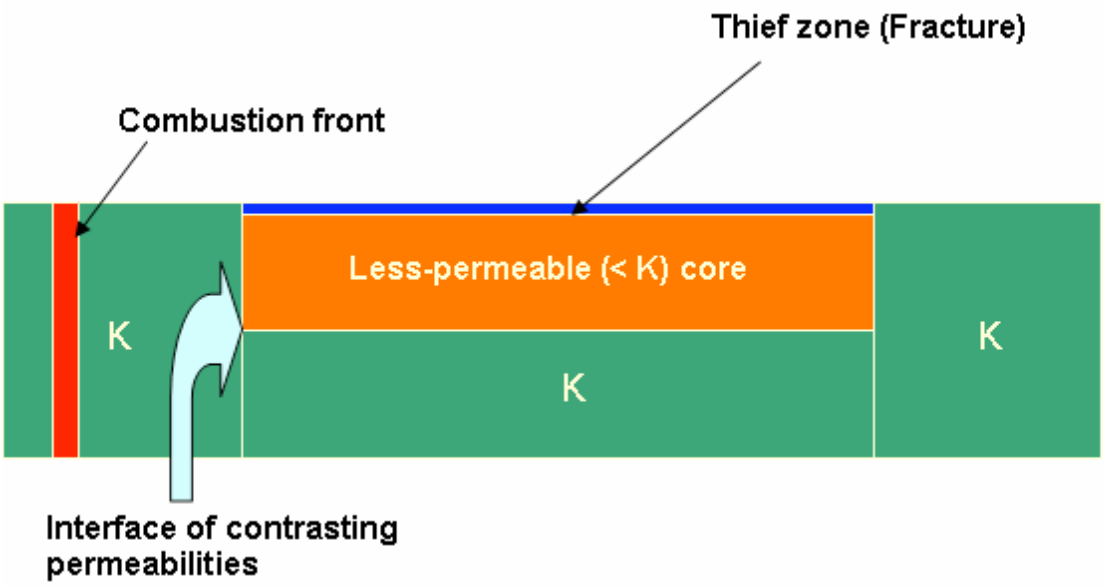

Figure 6.3. Experimental representation of large-scale heterogeneities.

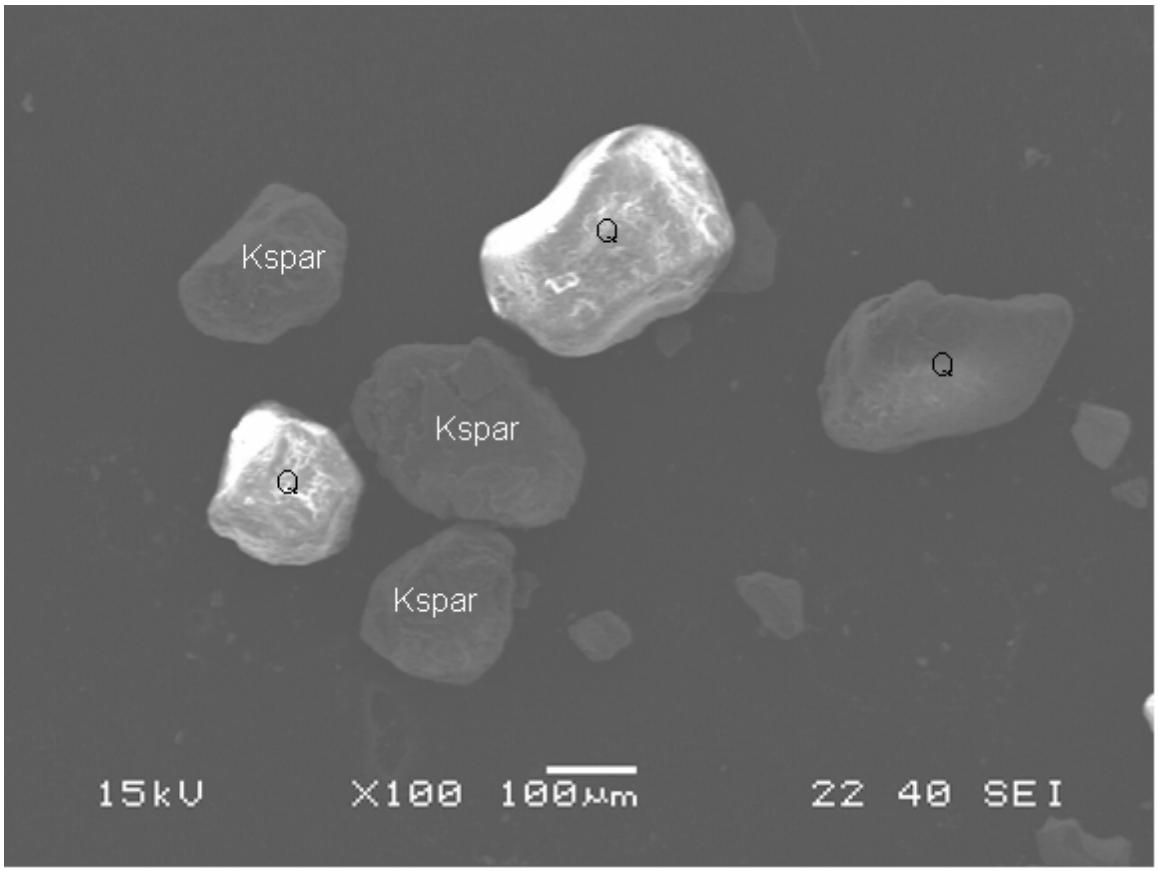

Figure 6.4. Scanning electron microscope images (100X) of reservoir grains illustrating quartz (Q) and potassium feldspar (Kspar).

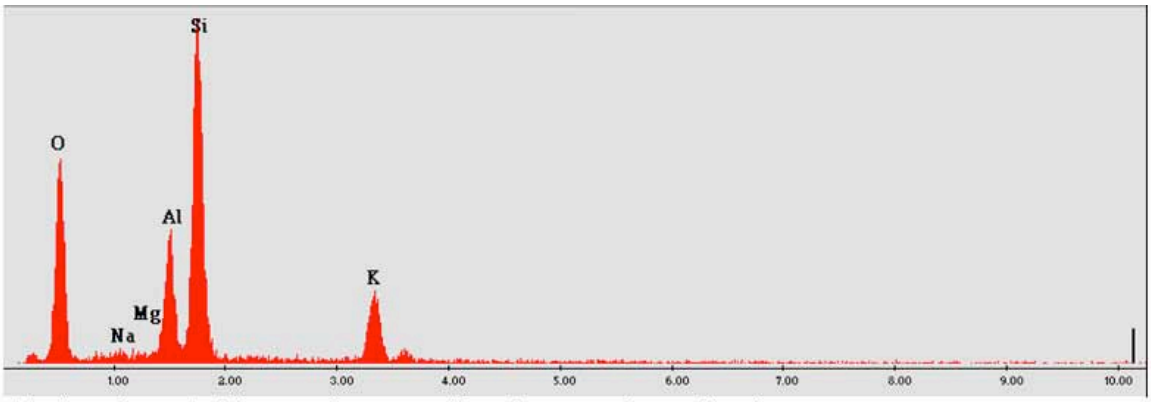

Potassium feldspar (e.g., orthoclase, microcline)

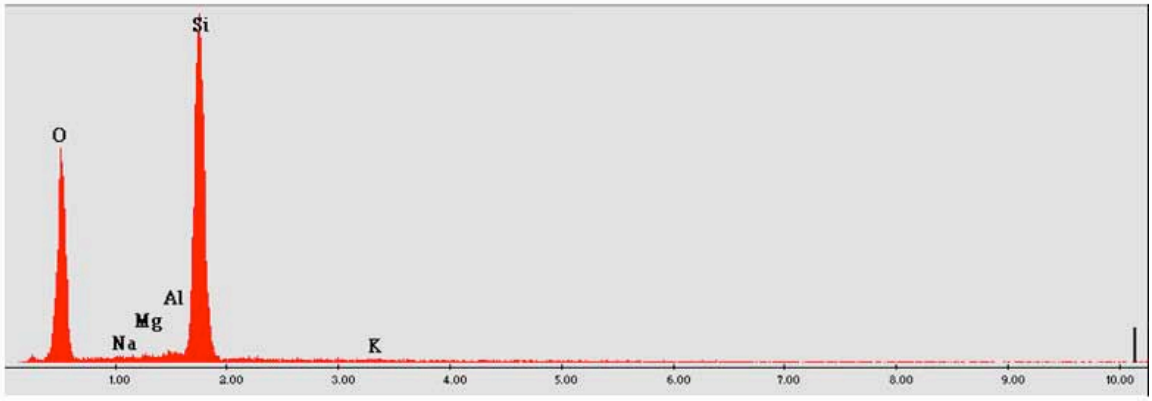

Silica (quartz and perhaps chert)

Figure 6.5. Elemental analysis of minerals found in reservoir grains.

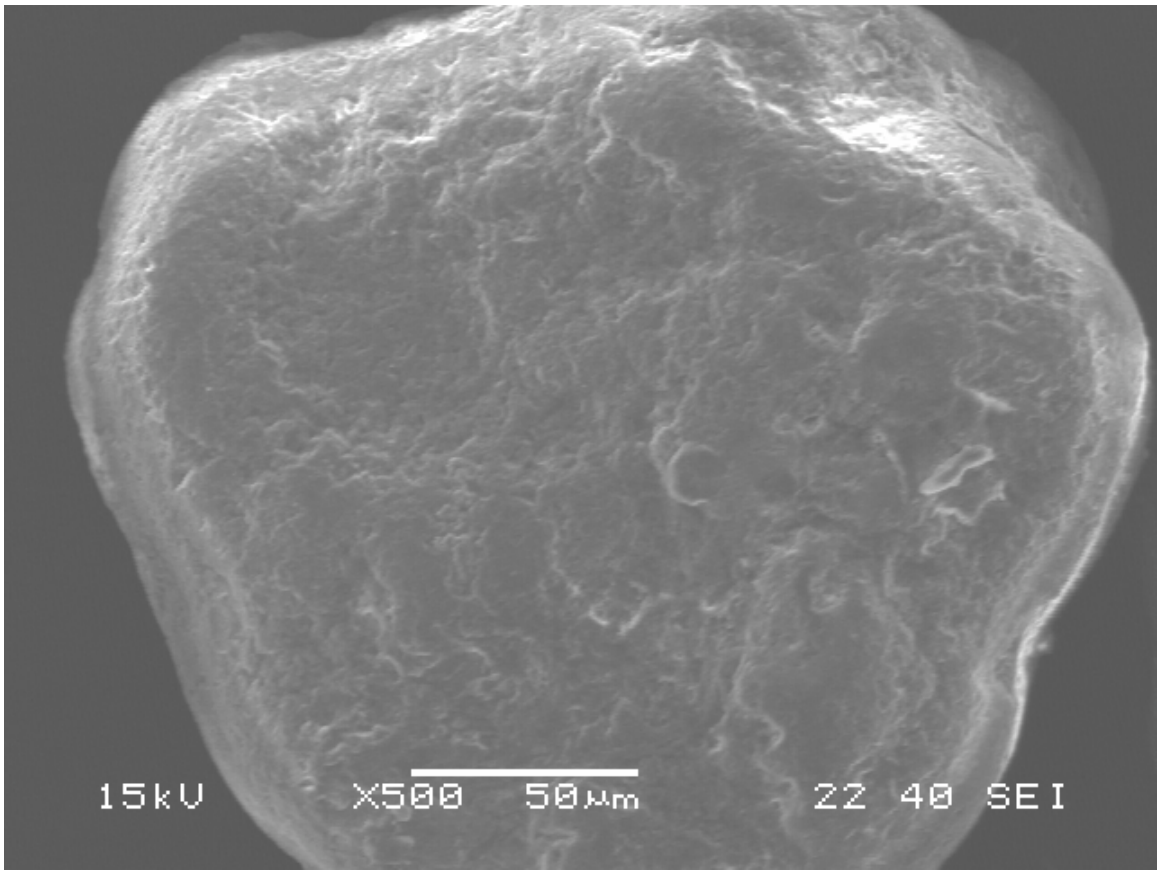

Figure 6.6. SEM image (500X) of the rough surface of a quartz grain.

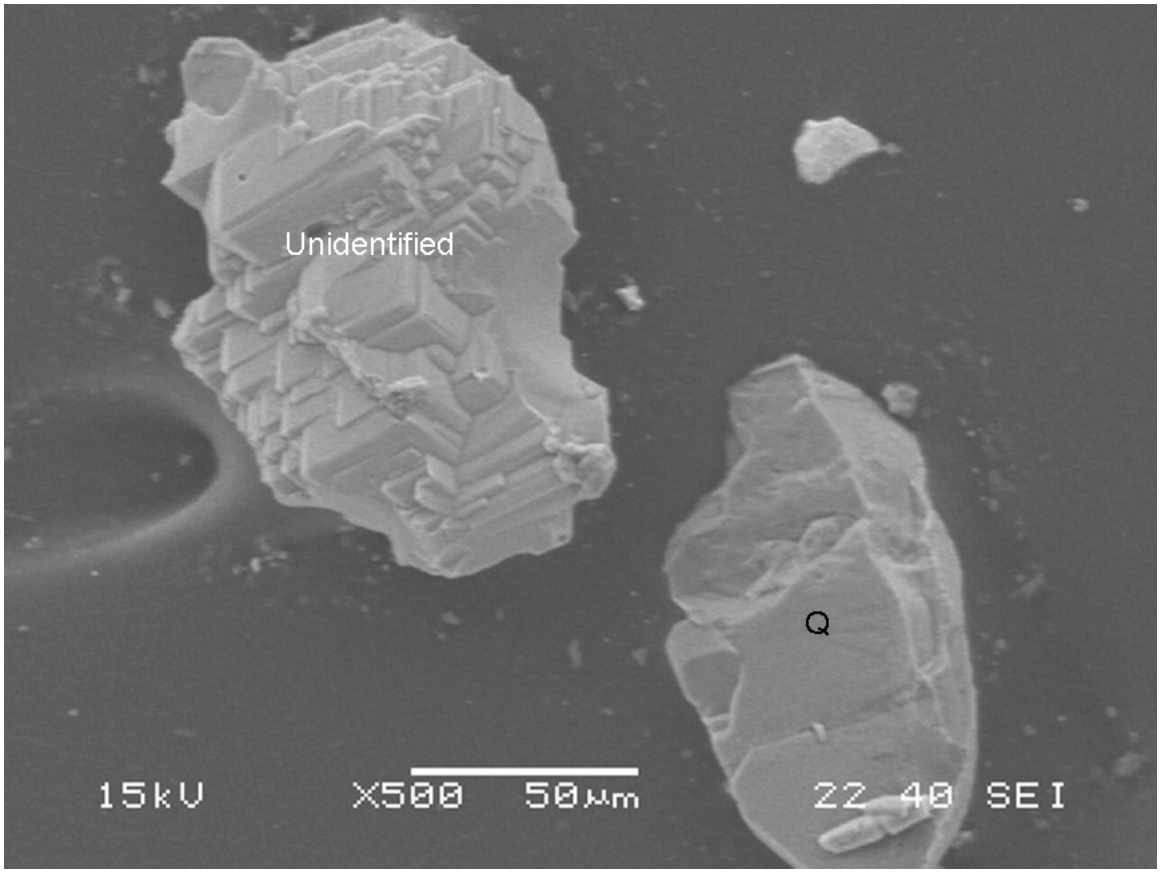

Figure 6.7. Large magnification (500X) sometimes shows material that is not identifiable such as angular grain on the left.

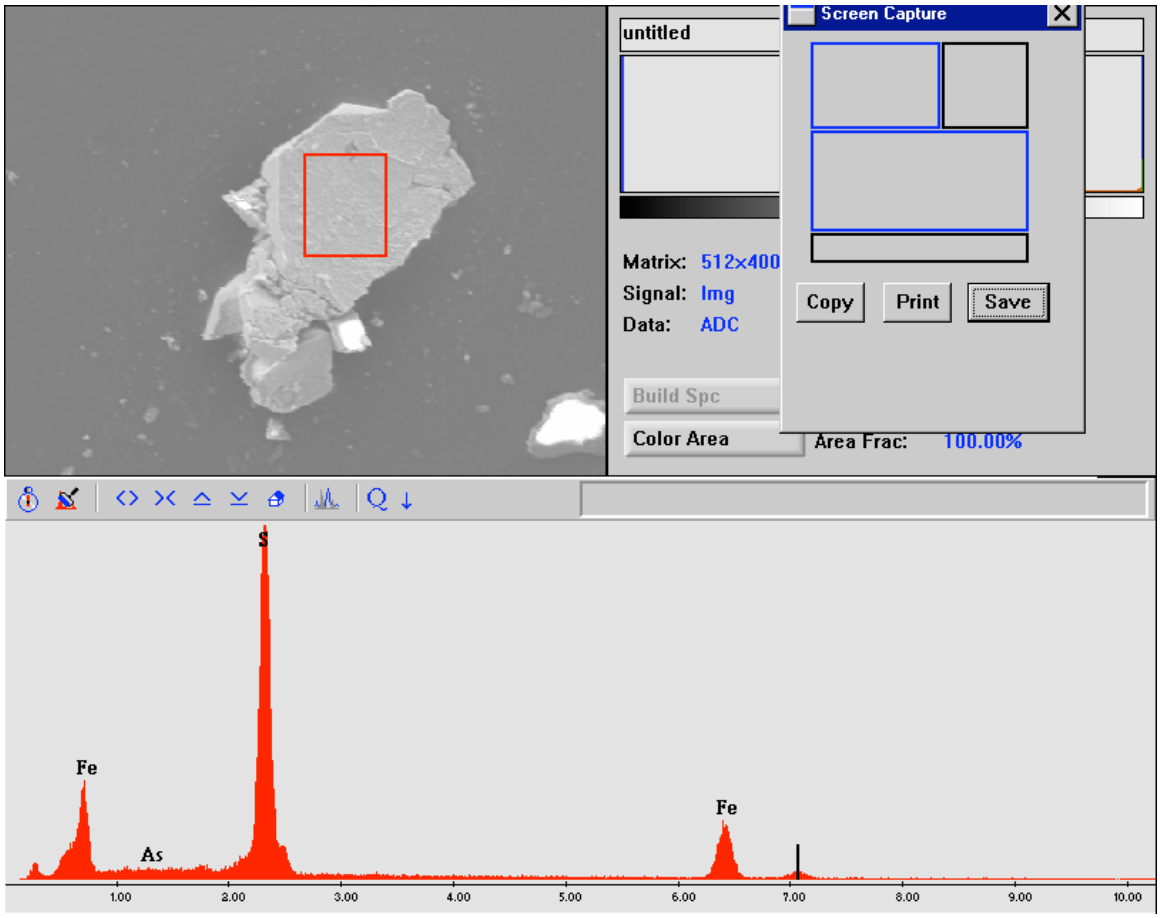

Figure 6.8. SEM image of a pyrite grain (500X) and elemental analysis (lower image) showing the abundance of iron and sulfur.

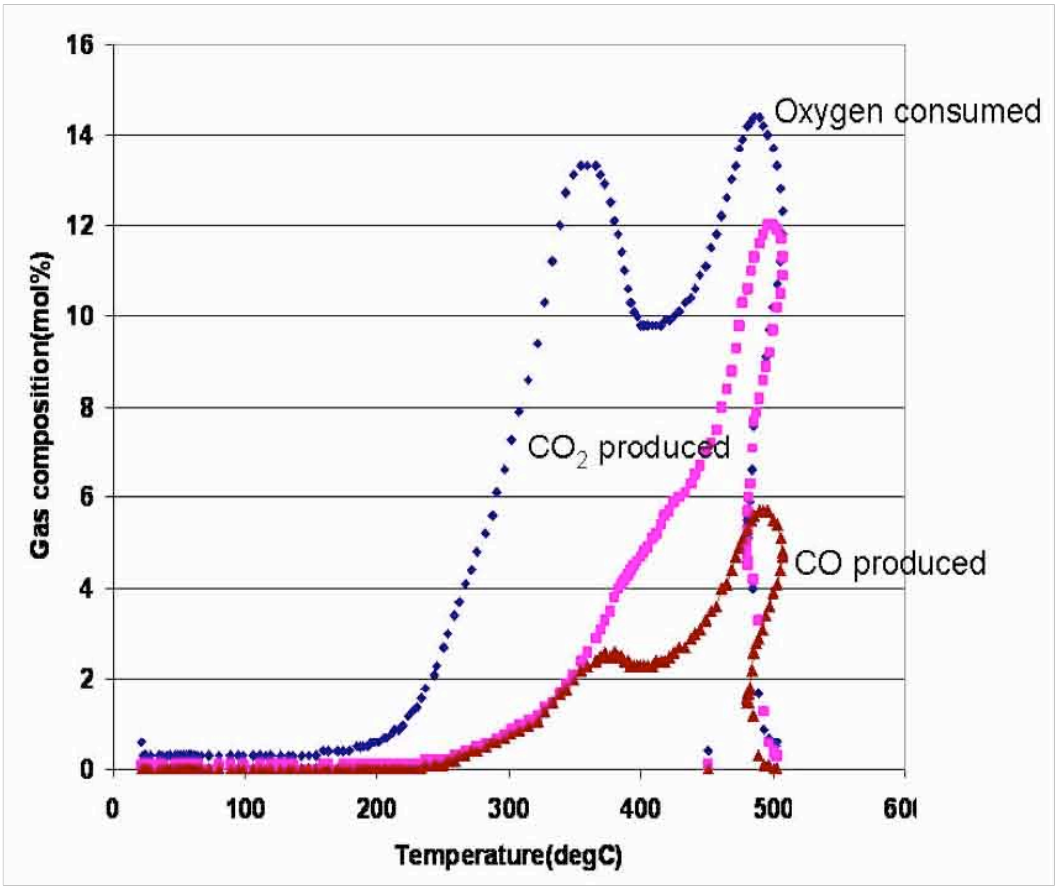

Figure 6.9 Effluent gas profiles at 50 psig (mixture), experiment 1.

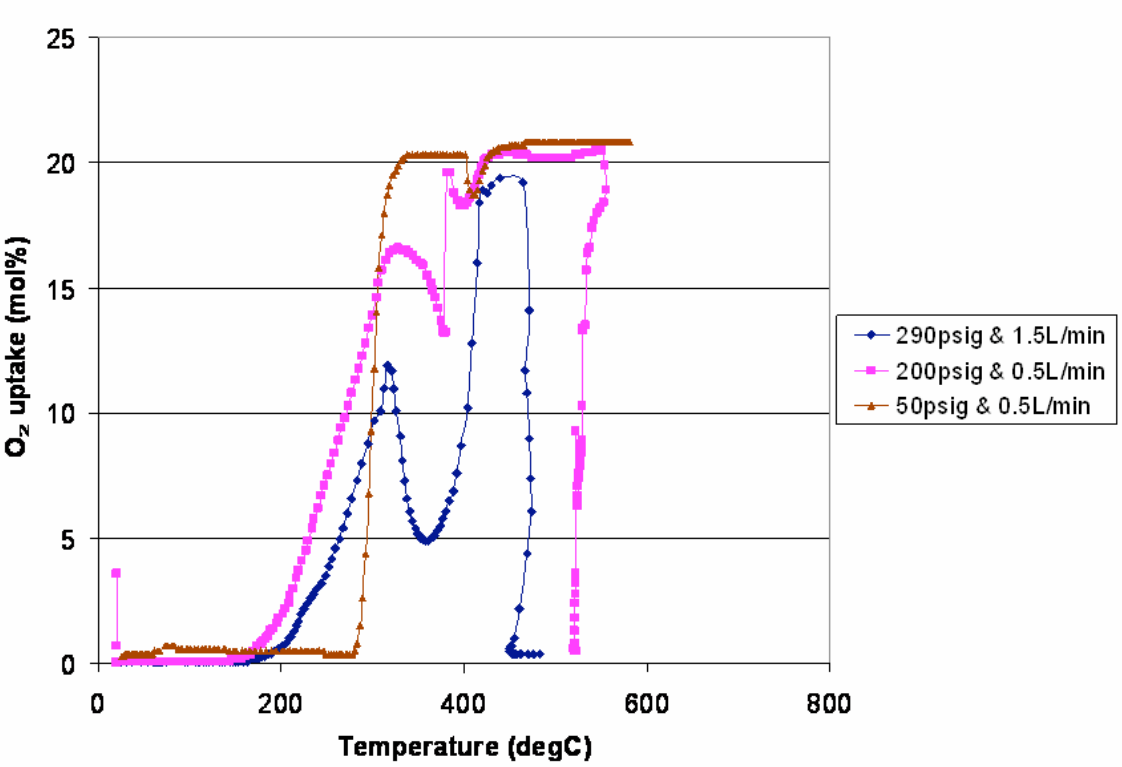

Figure 6.10. Oxygen uptake for the crushed reservoir samples at supply pressures of 50, 200 and 290psig and air flux of 0.5, 0.5 and 1.5L/min, experiments 2, 3, and 4, respectively.

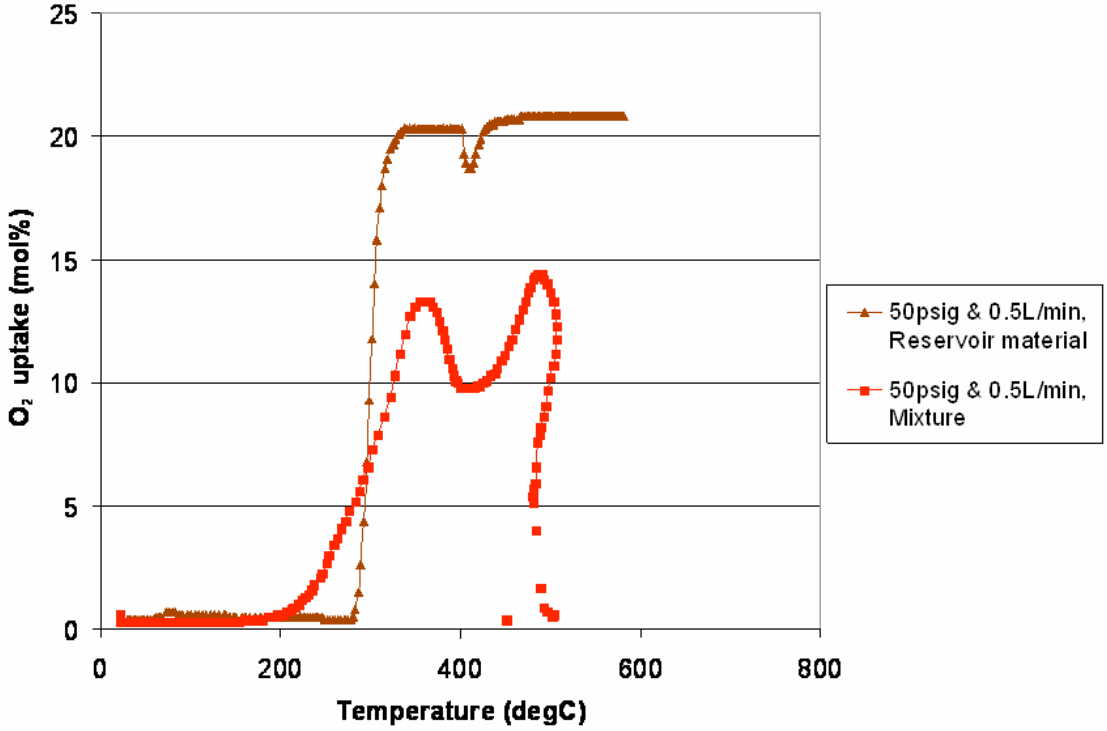

Figure 6.11. Oxygen uptake for the mixture and crushed reservoir samples at supply pressures of 50psig and air flux of 0.5L/min, experiments 1 and 2.

Figure 6.12. Top: Sand, clay, and water mixture before experiment. Bottom: residue after experiment.

Figure 6. 13. Left - Crushed reservoir material before experiment at 50psig. Right - Residue after experiment at 50psig supply pressure. Note some coke-type particles in the residue.

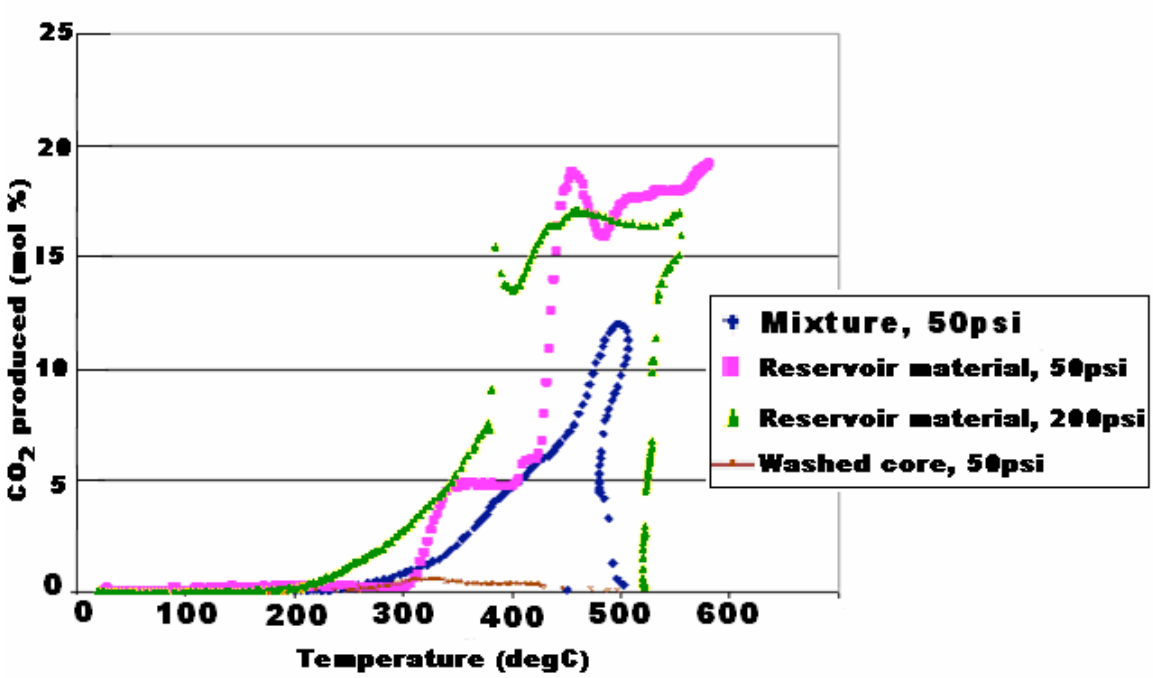

Figure 6.14. CO_2 production profiles for the mixture @ 50psig, reservoir material @ 50psig, reservoir material @ 200psig and the washed core also @ 200psig.

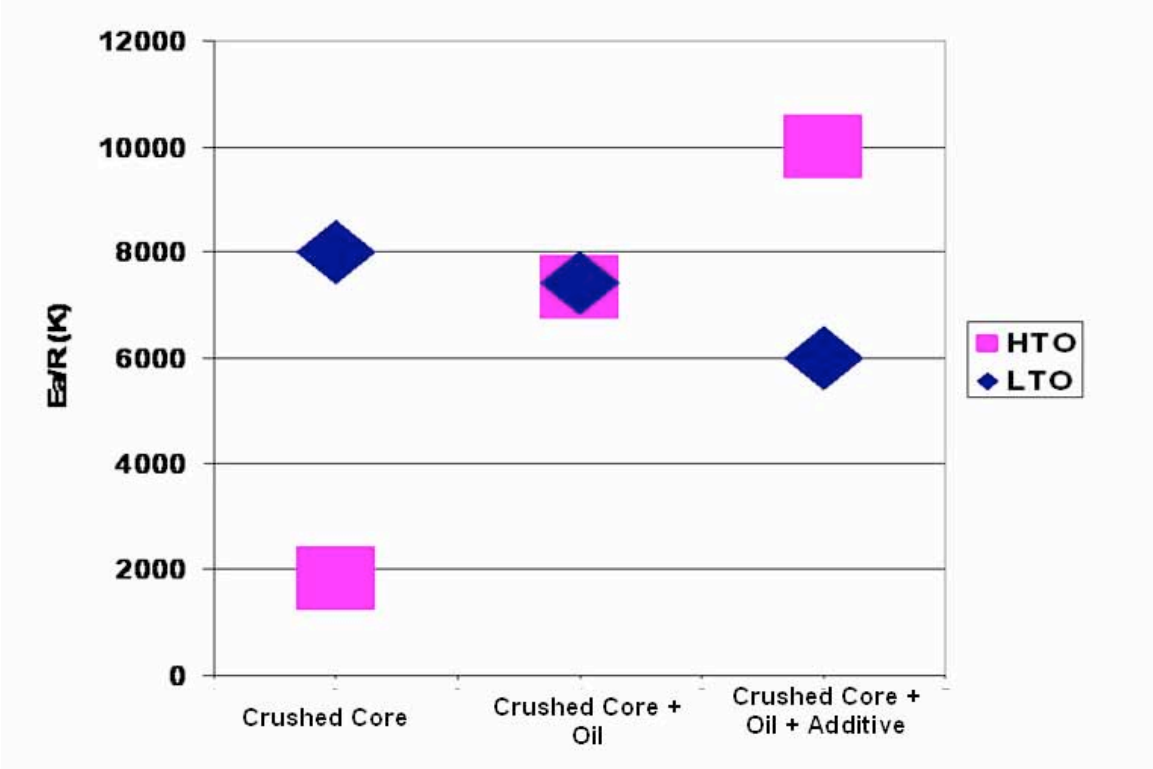

Figure 6.15. Summary of activation energies for LTO and HTO obtained during kinetics testing for the 19.7 °API heavy crude oil.

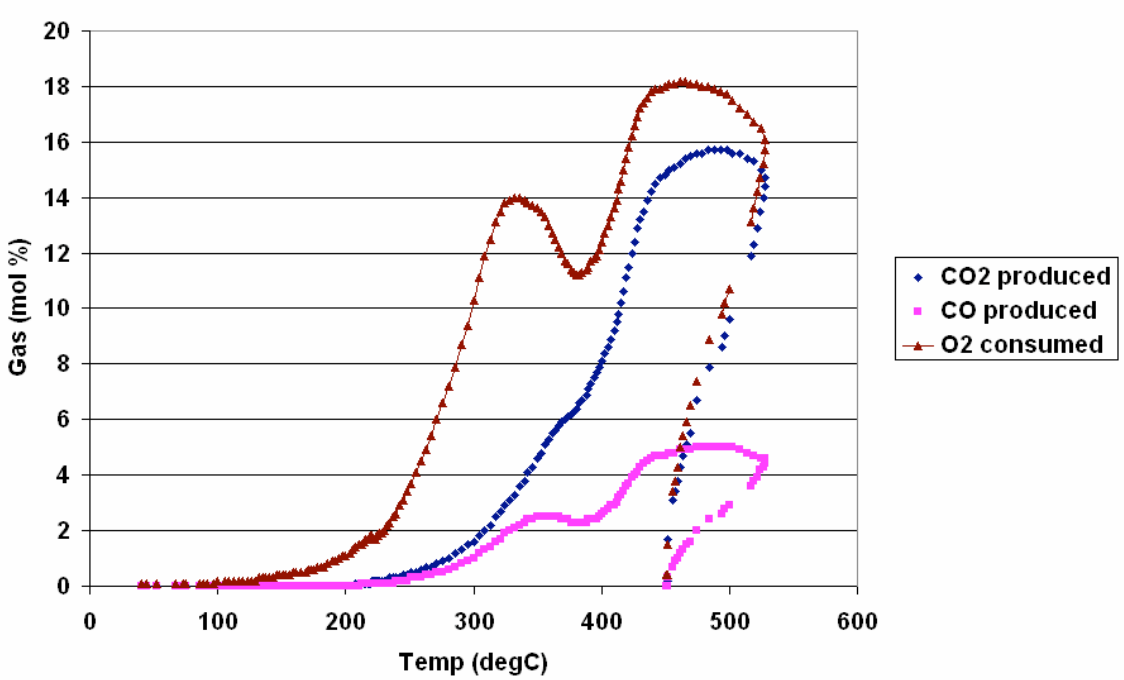

Figure 6.16. Effluent gas profiles for Wilmington (CA) crude oil at 50 psig.

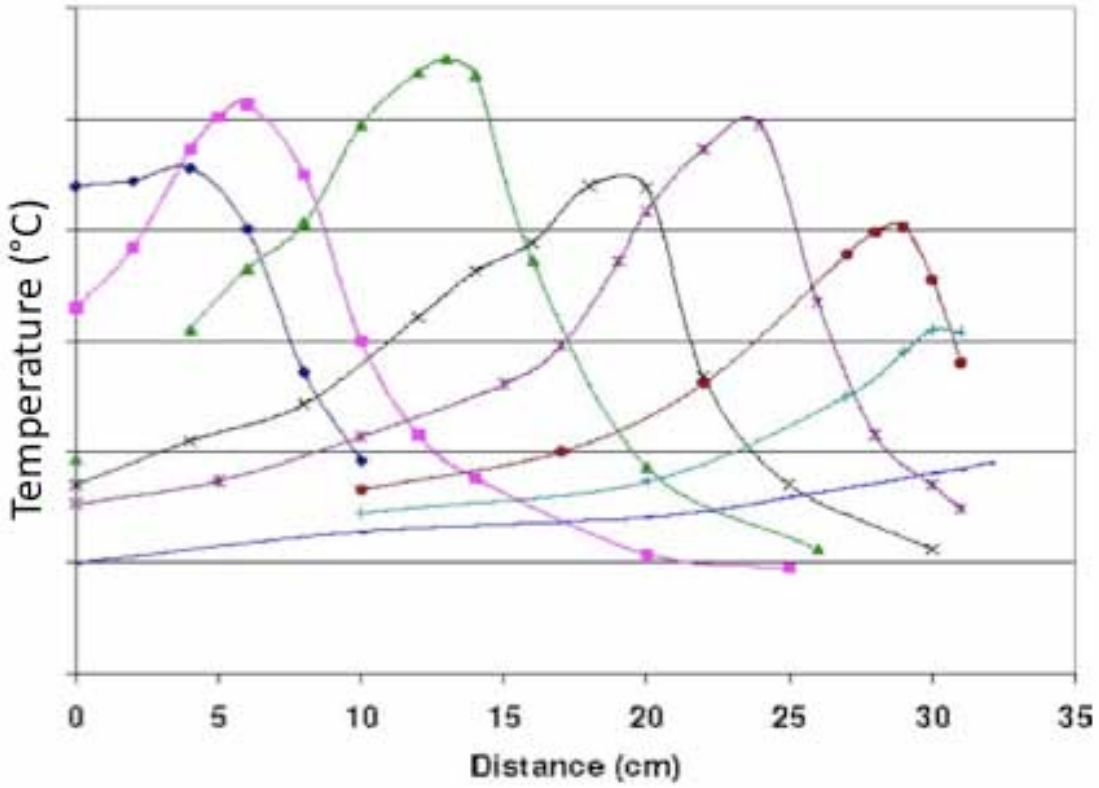

Figure 6.17. Temperature profiles, showing spatial and temporal propagation of combustion front; experiment #8. Tube is packed with crushed reservoir material only. Legend indicates times.

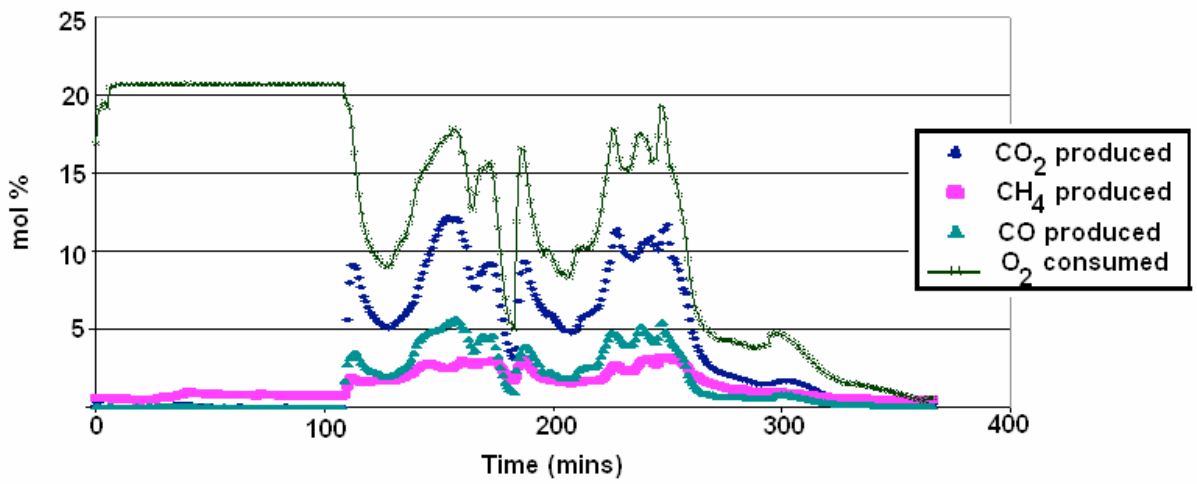

Figure 6.18. Gas production history, experiment #8. Tube is packed with crushed reservoir material only.

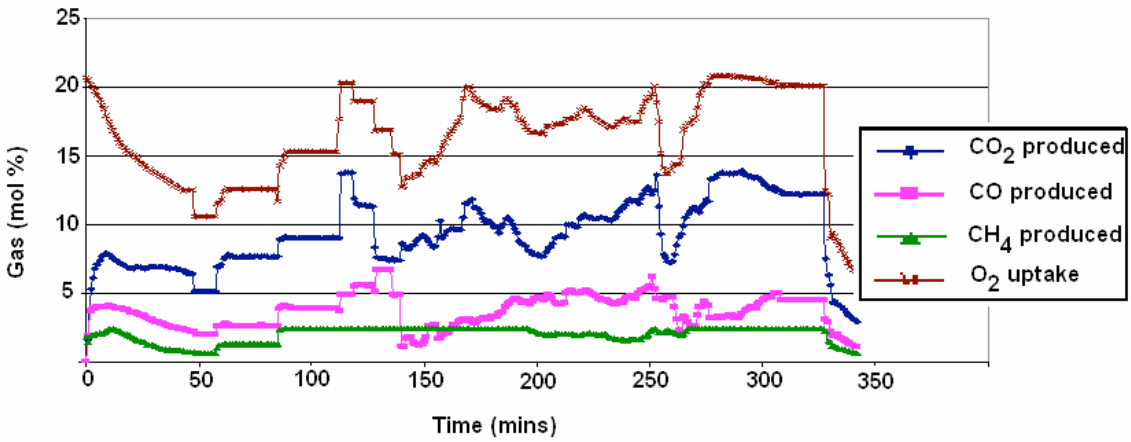

Figure 6.19. Effluent gas history: small-scale heterogeneous pack, experiment #9.

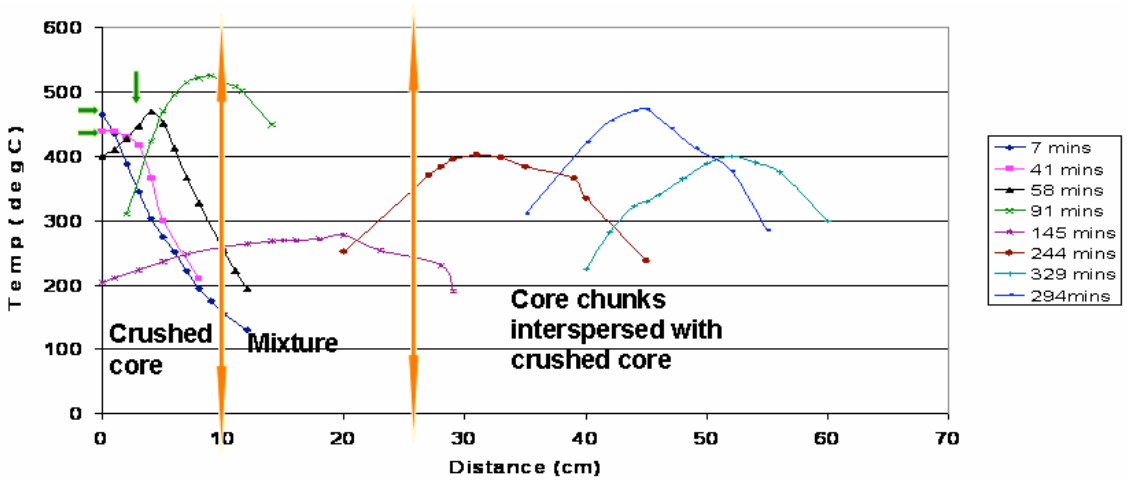

Figure 6.20. Axial temperature profiles from experiment 9, small-scale heterogeneous pack. Figure indicates that combustion front propagation is undeterred by small-scale heterogeneities.

Figure 6.21. CT scans and digital images of pre and post-burn small-scale heterogeneous pack; experiment 9.

Figure 6.22. Left - CT scan of a homogeneous section within the tube, experiment #9. Right - CT scan of a small-scale heterogeneous section within the tube.

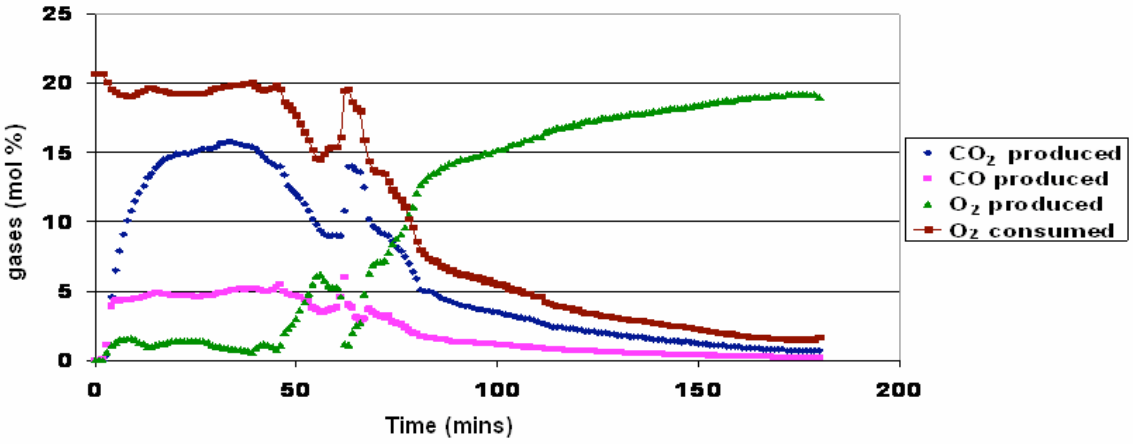

Figure 6.23. Effluent gas history from Experiment 10: large-scale heterogeneous pack.

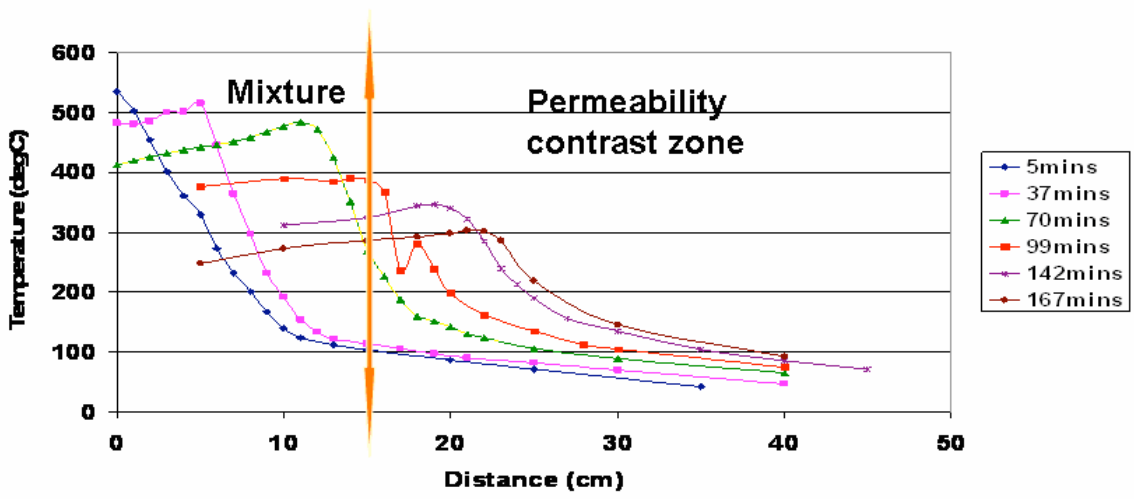

Figure 6.24. Axial temperature profiles from experiment 10, large-scale heterogeneous pack. Note decreasing front temperature as the as the front reaches the heterogeneous zone.

Figure 6.25. CT scans and digital images of pre and post-burn large-scale heterogeneous pack, experiment #10.

7 Simulation model I - Virtual Kinetic Cell: Specialized Integration Methods for Kinetics

In this chapter we discuss the development of a specialized algorithm for the integration of stiff chemical reactions in ISC processes. Because kinetics is the fastest process that must be resolved in ISC simulation, the formulation of specialized methods for kinetics is an essential first step in the design of an improved solution strategy. These methods are combined with the operator splitting approach discussed in this report to form the basis of a new class of ISC solvers.

Physical subprocesses, such as mass and heat transport, have associated temporal, and often also spatial, scales that are typically much larger than those for kinetics. Operator splitting methods, briefly discussed in the chapter on adaptive temporal methods, are attractive for solving processes with disparate temporal scales. They allow subprocesses to be essentially decoupled so that each subprocess can be solved using tailored numerical methods and time steps. For example, stiff kinetics is best treated by implicit methods and requires relatively small time steps, whereas convective transport can often be integrated explicitly with much larger time steps.

In a reaction substep, each grid block is effectively treated as a small kinetic cell with homogeneous pressure and temperature and well mixed fluids. We focus on the design of efficient numerical integration of the stiff kinetics. High computational efficiency is very desirable as a typical simulation may involve millions of reaction substeps. Phase behavior, that is the transfer of components between phases, adds further complexity to compositional and thermal processes. For non-reacting systems, it is reasonable to assume that the time scales for component transfer between phases to reach an equilibrium state are faster than for other subprocesses. The validity of the assumption for reactive systems is unclear, but in lack of better alternatives we shall assume thermodynamic equilibrium conditions. How to numerically treat the interaction of phase behavior with kinetics is, as yet, not fully understood in processes such as in-situ combustion. We expect that predictions are very sensitive to the approach chosen, because phase behavior and kinetics both affect the amount and composition of the combustion fuel, the heat released and the combustion gases produced.

7.1 The Mathematical Model

7.1.1 Reaction Models

The minimal model

This model includes a minimal realistic set of components and reactions to represent ISC behavior. The components are in the model are light oil (LO), heavy oil (HO), coke (C), oxygen ($\mathrm{O_2}$), water (W), and inert gas (IG). The model reactions are given by

$$\begin{aligned} \mathrm{LO} + \mathrm{O_2} &\rightarrow \mathrm{IG} + \mathrm{W} \\ \mathrm{HO} + \mathrm{O_2} &\rightarrow \mathrm{IG} + \mathrm{W} \\ \mathrm{HO} &\rightarrow \mathrm{LO} + \mathrm{C} + \mathrm{IG} \\ \mathrm{C} + \mathrm{O_2} &\rightarrow \mathrm{IG} + \mathrm{W}. \end{aligned} \tag{7.1}$$

The first two reactions model oxidation of the light and heavy oil, respectively. The cracking of heavy oil is modeled with the third reaction, while the last reaction models complete oxidation of the deposited fuel (coke). This minimal model is the common model used in many ISC simulations and it is the default model available in the commercial simulator CMG STARS.

The SARA based model

In the SARA (Saturates, Aromatics, Resins, and Asphaltenes) approach, chemical reactivity and phase behavior is taken into account when grouping the components. Saturates are saturated hydrocarbons with straight or branched chains, little ring-structure, and only little nitrogen, sulfur, and oxygen content. Aromatics contain one or more aromatic rings. Resins are the second heaviest fraction with high polarity due to considerable nitrogen, sulfur, and oxygen content. Asphaltenes are defined as the fraction of the crude insoluble in n-heptane. Within each fraction, groups are distinguished based on reactivity. A further subdivision of the fractions based on boiling point may be used to obtain improved representation of phase behavior.

The model we consider in this work contains 14 (pseudo) components, listed in Table 7.I., and the 14 reactions given following the table. Details of this model can be found in [Freitag2005, Freitag2006, Freitag2005b]. The reactions include pyrolysis, LTO, and HTO of the SARA fractions along with HTO of partially oxidized LTO residue and the coke formed by pyrolysis. The LTO residues of resins, aromatics, and saturates are non-volatile oil components whereas the LTO residue of asphaltenes is solid.

Table 7.I.: List of (pseudo) components for the SARA based model. The components can exist in the phase(s) listed in the last column.

Component Name	Abbreviation	Phase(s)
Water	H_2O	water
Inert oil	InertOil	oil
Oxidized resins/aromatics	OxdResAr	oil
Oxidized saturates	OxdSat	oil
Asphaltenes	Asph	oil
Resins	Resins	oil
Aromatics	Arom	oil/gas
Saturates	Sat	oil/gas
Light oil	Lites	oil/gas
Carbon dioxide	CO_2	oil/gas
Nitrogen	N_2	gas
Oxygen	O_2	gas
Oxidized asphaltenes	OxdAsph	solid
Pyrolysis coke	PyrCoke	solid

$$
\begin{aligned}
&\text{cracking/}\\
&\text{pyrolysis}\ \begin{cases}
\text{Asph} \rightarrow \mathrm{H_2O} + \text{Arom} + \text{Sat} + \text{Lites} + \text{PyrCoke}\\
\text{Resins} \rightarrow \mathrm{H_2O} + \text{Asph} + \text{Arom} + \text{Sat} + \text{Lites}\\
\text{Arom} \rightarrow \text{Asph} + \text{Lites}
\end{cases}\\
&\text{LTO}\ \begin{cases}
\text{Asph} + \mathrm{O_2} \rightarrow \mathrm{H_2O} + \mathrm{CO_2} + \text{OxdAsph}\\
\text{Resins} + \mathrm{O_2} \rightarrow \mathrm{H_2O} + \mathrm{CO_2} + \text{OxdResAr}\\
\text{Arom} + \mathrm{O_2} \rightarrow \mathrm{H_2O} + \mathrm{CO_2} + \text{OxdResAr}\\
\text{Sat} + \mathrm{O_2} \rightarrow \mathrm{H_2O} + \mathrm{CO_2} + \text{OxdSat}
\end{cases}\\
&\text{HTO}\ \begin{cases}
\text{PyrCoke} + \mathrm{O_2} \rightarrow \mathrm{H_2O} + \mathrm{CO_2}\\
\text{OxdAsph} + \mathrm{O_2} \rightarrow \mathrm{H_2O} + \mathrm{CO_2}\\
\text{OxdResAr} + \mathrm{O_2} \rightarrow \mathrm{H_2O} + \mathrm{CO_2}\\
\text{OxdSat} + \mathrm{O_2} \rightarrow \mathrm{H_2O} + \mathrm{CO_2}\\
\text{Asph} + \mathrm{O_2} \rightarrow \mathrm{H_2O} + \mathrm{CO_2}\\
\text{Resins} + \mathrm{O_2} \rightarrow \mathrm{H_2O} + \mathrm{CO_2}\\
\text{Arom} + \mathrm{O_2} \rightarrow \mathrm{H_2O} + \mathrm{CO_2}
\end{cases}
\end{aligned}
$$

7.1.2 Building the Virtual Kinetic Cell (VKC)

Each grid block in a reaction substep is treated as a small chemical reactor, or kinetic cell, in which only chemical kinetics and phase behavior are taken into account, see Figure 7.1. We have extended the model with inflow and outflow boundary conditions for mass and energy, as well as external heating and cooling terms to investigate possible interactions in the reservoir between gridblocks. For example, we can investigate the sensitivity to oxygen availability by varying the oxygen inflow rate. We can also apply the VKC in an inverse mode to determine kinetic parameters from data obtained through laboratory kinetic cell experiments.

We model a closed system consisting of three fluid phases (oil, water, and gas), an immobile solid phase, and the porous medium. Temperature, pressure, and component concentrations are assumed to be uniform in the cell. For convenience, we introduce the molar concentration c_i^j of a component $i, i = 1, \ldots, n_c$, in a phase j, $j \in \{o, w, g\}$, and the overall molar concentration C_i of component i. They are given by

$$c_i^j = \phi_f x_i^j \xi^j S^j, \tag{7.3}$$

$$C_i = \sum_{j \in \{o,w,g\}} c_i^j, \tag{7.4}$$

where S^j are the phase saturations, ξ^j the molar phase densities, and x_i^j the mole fraction of component i in phase j. ϕ_f is the fluid porosity which varies with solid fuel concentration in the pore space according to

$$\phi_f = \phi_v - \frac{C_s}{\xi^s}, \tag{7.5}$$

where C_s and ξ^s are the density and concentration of the solid component, respectively. The solid concentration is defined with respect to total volume. The mass conservation of each chemical component can be expressed as

$$\frac{dC_i}{dt} = \sum_\gamma A_{i\gamma} R_\gamma + \frac{F_i^{in} - F_i^{out}}{V_T}, \qquad i = 1, \ldots, n_c. \tag{7.6}$$

Here, F_i^{in} and F_i^{out} are the molar flow rates of component i in and out of the cell, respectively, V_T is the cell volume, R_γ the kinetic expression for the γ^{th} chemical reaction, and $A_{i\gamma}$ is the stoichiometric coefficient for component i in reaction γ.

The energy conservation equation is given by

$$\frac{dU}{dt} = \sum_\gamma (-\Delta H_\gamma^r) R_\gamma + \frac{H^{in} - H^{out} + Q^{ext}}{V_T}, \tag{7.7}$$

where H^{in} and H^{out} are the fluxes of enthalpy in and out of the cell, and Q^{ext} is a heat source/sink term due to external heating or cooling, which is given by

$$Q^{ext} = -u_a (T - T^{ext}), \tag{7.8}$$

where T^{ext} is the heating/cooling temperature and u_a is an overall heat transfer coefficient.

Finally, the total internal energy of the system per unit volume, U, is given by

$$U = (1 - \phi_v) U^r + (\phi_v - \phi_f) \xi^s U^s + \phi_f \sum_{j \in \{o,w,g\}} \xi^j S^j U^j. \tag{7.9}$$

Since the cell volume is assumed constant, gases are allowed to leak out of the cell when volume changes occur due to chemical reactions and temperature changes. When only components in the gas phase can leave the cell, the molar flow rates in (7.6) may be expressed as

$$F_i^{out} = Q^{out} \xi^g x_i^g, \tag{7.10}$$

where ξ^g is the gas density and Q^{out} is the volumetric flow rate, given by

$$Q^{out} = k_v \left(P - P^{ext} \right) \tag{7.11}$$

Here, k_v denotes a valve coefficient.

Since pressure and temperature variations and chemical reactions cause volume changes, a constraint is needed in order for the fluid and solid volumes, V^f and V^s, to match the void pore volume, V_p. We set

$$V_p = V^f + V^s. \tag{7.12}$$

The chemical reactions occurring are assumed to be kinetically driven. They are modeled using the Arrhenius rate relations

$$K_\gamma = \alpha_\gamma \exp\left(-\frac{E_a^\gamma}{R_g T}\right), \tag{7.13}$$

where α_γ and E_a^γ are the pre-exponential factor and activation energy for reaction γ, respectively, and R_g is the universal gas constant. We will assume that the reactions are first order in all reactants. Oxidation reaction rates for components in phase j are modeled as

$$R_\gamma = K_\gamma \cdot \left(P x_{O_2}^g\right) \cdot \left(\phi_f \xi^j S^j x_i^j\right) \tag{7.14}$$

For cracking/pyrolysis reactions, we use

$$R_\gamma = K_\gamma \cdot \left(\phi_f \xi^j S^j x_i^j\right) \tag{7.15}$$

Reaction kinetic parameters for ISC models are normally obtained from laboratory experiments. It is still an open question how to upscale kinetics from the laboratory scale to the reservoir scale. Lacking better alternatives, we will use the laboratory parameters directly, as is generally done.

We assume that components partition into at most two phases, that components in the solid phase exist exclusively in this phase, and that the water phase consists of water only. For simplicity, the equilibrium between hydrocarbon components in the gas and oil phases is described by

$$K_i = \frac{x_i^g}{x_i^g}, \tag{7.16}$$

where the equilibrium factors are assumed to vary only with pressure and temperature according to

$$K_i = \frac{a_1}{P} \cdot \exp\left(\frac{a_2}{T - a_3}\right), \tag{7.17}$$

with a_1, a_2 and a_3 correlation constants. Using this simplified phase equilibrium description the flash calculation reduces to solving the Rachford-Rice equation for the molar gas phase fraction, β^g, given by

$$\sum_{i \in \{hc\}} (x_i^g - x_i^o) = \sum_{i \in \{hc\}} \frac{C_i (K_i - 1)}{1 - \beta^g + \beta^g K_i}. \tag{7.18}$$

The assumption that the equilibrium factors do not vary with composition is generally valid away from critical regions. The simplified phase description is a good starting point for model and algorithm development. In future work, we plan to extend the model with an equation-of-state (EoS) based phase equilibrium description. The challenges associated with this extension to a full phase equilibrium description are discussed the next chapter.

As primary variables we choose temperature, pressure and the n_c overall component concentrations. The VKC equations, (7.6), (7.7) and (7.12), comprise a

set of differential-algebraic equations (DAEs). Including the total internal energy as a variable and Eqn. (7.9) as an extra constraint, we can write the DAE in semi-explicit form. In this formulation, temperature is aligned with the energy constraint (7.9) and pressure with the volume constraint (7.12). In multiphase regions Eqn. (7.18) is appended to the equation system and aligned with the gas phase fraction. The two ISC reaction models discussed in the previous section are used for testing and benchmarking.

We note at this point that the mass balance equations (7.6) are solved using overall component concentrations as dependent variables. For batch reactor problems in general, a reformulation is often used based on extent of reaction [Aris1989], which changes the number of ODEs from n_c to n_r. Thus, for problems with more components than independent reactions this may be advantageous. The extent of reaction concept has been generalized to open systems involving flow [Friedly1991, Friedly1992]. For certain reactive porous media transport problems in water resources modeling this generalized extent of reaction formulation leads to a reduced and simpler set of equations to be solved [Friedly1992, Krautle2005]. For the ISC problem, however, with a nonlinear transport part and component transfers between phases due to phase behavior, the reformulation is not directly applicable. We solve the ISC equations using overall quantities (such as total component concentrations) as primary variables. We therefore formulate the kinetic cell model in the same set of variables.

7.2 The Numerical Approach

7.2.1 Choice of integration method for reactions

The kinetic cell equations consist of coupled ordinary differential equations (ODEs) and algebraic equations (AEs). We will consider the general form of the resulting differential algebraic equation (DAE) system

$$\mathbf{M}\frac{d\mathbf{u}}{dt} = \mathbf{f}(t,\mathbf{u}), \qquad \mathbf{u}(t_0) = \mathbf{u}_0, \tag{7.19}$$

where $\mathbf{u} \in \mathrm{R}^m$ is a vector of state variables depending on t, and $\mathbf{f}$ is a vector function mapping $\mathrm{R} \times \mathrm{R}^m$ into R^m. For the kinetic cell model the mass matrix, $\mathbf{M} \in \mathrm{R}^m \times \mathrm{R}^m$, is simply a diagonal matrix. The diagonal elements M_{ii} are one when the i^{th} equation is a differential equation, and zero when it is algebraic. The initial conditions are assumed consistent with the algebraic constraints. In the VKC model, the right-hand-side functions may be discontinuous because of phase changes. The number of elements in the state vector is typically between 10 and 20. We require results to have between 2 and 4 correct significant digits (relative accuracy between $10^{-2} - 10^{-4}$). This requirement is based on our expectations of the temporal and spatial accuracies in other parts of the equations. As we will show later, the

performance of the proposed methods is not sensitive to small variations in the required accuracy.

The short integration intervals between operator splitting updates in the VKC and the stiff reaction kinetics are the primary characteristics guiding our choice of integration method. When integrating from t^n to t^{n+1} one-step methods only use state information from t^n whereas multi-step methods use information from several previous integration steps, t^{n-j}, with $j = 0,1,\ldots,s$ [Wanner1992]. Multi-step methods are often implemented in a variable order formulation, in which they start out at low order and with small time steps. They build up to higher order as more information becomes available, and are, at that stage, generally outperforming one-step methods. This suggests that multi-step methods are attractive for problems with long and smooth integration intervals, whereas one-step methods are preferable on shorter intervals and on problems with frequent discontinuities in the solution. Hence, we will focus here on one-step methods. Of the one-step methods, we limit ourselves to the family of Runge-Kutta (RK) methods, which are well studied in the literature. Other classes of methods could be considered also, such as the Rosenbrock methods [Rosenbrock1963] or the extrapolation methods [Wanner1992], but we will consider Runge-Kutta methods exclusively here. For stability reasons we shall further limit our search to implicit RK methods that allow efficient integration of the stiff kinetics.

A general s-stage Runge-Kutta scheme for solving the DAE system (7.19) may be expressed as:

$$\mathbf{M}\mathbf{U}_i = \mathbf{M}\mathbf{u}_n + \Delta t^n \sum_{j=1}^{s} a_{ij}\mathbf{f}(t^n + c_j\Delta t^n, \mathbf{U}_j), \tag{7.20}$$

$$\mathbf{M}\mathbf{u}_{n+1} = \mathbf{M}\mathbf{u}_n + \Delta t^n \sum_{i=1}^{s} b_i\mathbf{f}(t^n + c_i\Delta t^n, \mathbf{U}_i). \tag{7.21}$$

Here, $\mathbf{U}_i$ denotes the solution at the i^{th}, $(i = 1,\ldots,s)$, internal stage of integration step n and Δt^n denotes the time step length.

$$\begin{bmatrix} c_1 & a_{11} & a_{12} & \ldots & a_{1s} \\ c_2 & a_{21} & a_{22} & \ldots & a_{2s} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ c_s & a_{s1} & a_{s2} & \ldots & a_{ss} \\ \mathbf{u}_{n+1} & b_1 & b_2 & \ldots & b_s \end{bmatrix} = \begin{bmatrix} \mathbf{c} & \mathbf{A} \\ & \mathbf{b}^T \end{bmatrix} \tag{7.22}$$

Runge-Kutta methods are classified according to the structure of their Butcher tableau as illustrated in Eqn. (7.22). For explicit methods (ERK), the matrix $\mathbf{A}$ is strictly lower triangular, which implies that all the internal stages (F.20) can be calculated explicitly, making ERK methods computationally fast and straightforward

to implement. However, ERK methods generally have poor stability properties, which make them unsuitable for stiff problems [Wanner1992]. The four remaining subclasses of Runge-Kutta methods are all implicit, that is, the values of the internal stages are no longer calculated explicitly from the values of the previous stages. Each integration step of an implicit method requires the solution of a system of ms nonlinear equations. Normally, an iterative method, such as Newton's method, is applied. For diagonally implicit RK (DIRK) methods, the stage values can be calculated sequentially. Hence, the computational costs of DIRK methods are lower than those of fully implicit RK (FIRK) methods, for which all ms equations must be solved simultaneously. If all diagonal elements of $\mathbf{A}$ are identical and its upper diagonal elements are zero, the method is said to be a singly diagonally implicit (SDIRK). If the first stage of an SDIRK method is explicit, the method is said to be explicit singly diagonally implicit (ESDIRK).

For our purposes, the ESDIRK methods, for which the general form of the Butcher tableau can be found in e.g. [Kvaernoe2004], are attractive. The diagonal structure of these methods allows sequential evaluation of the internal stages, and because the diagonal elements are equal, the iteration matrix for solving the nonlinear stage equations (7.20) need only be evaluated and factorized once per integration step. ESDIRK methods can be constructed such that they are both A- and L-stable as well as stiffly accurate [Hairer1996]. Stiffly accurate methods avoid the order reduction phenomenon observed by [Prothero1974] when applied to stiff ODEs. Also, the explicit first stage of ESDIRK methods ensures high stage order (≥ 2) which is important for the order of accuracy in the algebraic components of the DAE [Hairer1996]. An additional advantage of having high stage order is that it allows the construction of high order interpolants to be used for generating output between mesh-points. We will exploit these interpolants when constructing an algorithm for locating the discontinuities that occur due to phase changes.

The Butcher tableau for ESDIRK methods takes the following form

$$\begin{array}{ccccccc} 0 & 0 & 0 & & & & \\ c_2 & a_{21} & \gamma & 0 & & & \\ c_3 & a_{31} & a_{32} & \gamma & 0 & & \\ \vdots & \vdots & \vdots & \ddots & \ddots & \ddots & \\ c_{s-1} & a_{s-1,1} & a_{s-1,2} & a_{s-1,3} & \cdots & \gamma & 0 \\ 1 & b_1 & b_2 & b_3 & \cdots & b_{s-1} & \gamma \\ \mathbf{u}_{n+1} & b_1 & b_2 & b_3 & \cdots & b_{s-1} & \gamma \end{array} \tag{7.23}$$

[Alexander2003] proposed a four stage ESDIRK method of order 3 with an embedded error estimator of order 4. [Cameron2002] also published a four stage ESDIRK method suited for index 2 DAEs, but with orders 2 and 3 for the error estimator and advancing method, respectively. Finally, [Kvaernoe2004] recently published a range of ESDIRK methods of orders 3 to 5 emphasizing strong stability properties of both the error estimator and the advancing method. ESDIRK methods

have been applied as time integrators in fully implicit fluid dynamics simulations [Bijl2002] as well as in operator splitting methods for integrating the reaction part of convection-diffusion-reaction equations [Kennedy2003].

7.2.2 Efficient implementation

We implemented four ESDIRK methods of orders 2-5. The lowest order method is simply the trapezoidal rule. When combined with the implicit Euler method we have an embedded pair of methods of orders 1 and 2. We refer to this method as ESDIRK12 and use a similar notation for higher order ESDIRK schemes. The bulk of the computational costs in an ESDIRK method is in the solution of the nonlinear algebraic equations that arise in each internal stage. Efficient control of the iterative scheme applied to these equations and of the discretization error by step size adjustments are the two main design challenges.

Error and convergence control

The step size is adapted based on an error estimate for the lowest order method obtained by subtracting the two solutions of orders p and $p-1$. This error estimate is essentially free, since it involves no additional function evaluations or system solves.
Adjusting the step size to meet an accuracy requirement is, in essence, a control problem. We implemented the predictive controller

$$\Delta t^n = \frac{\Delta t^{n-1}}{\Delta t^{n-2}} \left(\frac{\varepsilon}{r_n} \right)^{k_2/(p+1)} \left(\frac{r_{n-1}}{r_n} \right)^{k_1/(p+1)} \Delta t^{n-1} \tag{7.24}$$

suggested by [Gustafsson1992]. In Eqn. (7.24) k_1 and k_2 are the gain parameters for the controller, while ε is the desired tolerance (including a safety factor). [Gustafsson1992] suggests using $k_1 = k_2 = 1$ based on experiments, where the number of integration steps is measured for a given problem as a function of k_1 and k_2. The number of steps required is not sensitive to k_1 and k_2 for the interval $[k_1, k_2] \in 0.5,1.5\,;\,0.5,1.5]$. Within this interval, variations in k_1 and k_2 result in changes in the number of integration steps of less than 5%. Our experiments using the VKC model show similar results, so we choose $k_1 = k_2 = 1$. The variable r_n is the norm of the estimated local error at t^n, given by

$$r_n = \sqrt{\frac{1}{m} \sum_{i=1}^{m} \left(\frac{err_i}{atol_i + rtol_i \cdot |u_i|} \right)^2} \tag{7.25}$$

in which *atol* and *rtol* are (componentwise) absolute and relative error tolerances specified by the user. Experiments have shown that the controller (7.24) gives a small reduction in the number of failed integration steps along with a smoother

variation of step sizes compared to the conventional control law $\Delta t^n = \left(\varepsilon / r_n\right)^{1/(p+1)} \Delta t^{n-1}$ that is commonly implemented in ODE solvers.

Stage value predictions

The nonlinear equations arising in each internal stage of the ESDIRK methods are solved using a modified Newton's method given by

$$\left[\mathbf{M} - \Delta t^n \gamma \frac{\partial \mathbf{f}}{\partial \mathbf{u}}\right] \Delta \mathbf{U}_i = \mathbf{M}\mathbf{u}_n + \Delta t^n \sum_{j=1}^{i-1} a_{ij} \mathbf{f}(t^n + c_j \Delta t^n, \mathbf{U}_j^{(k)}) - \mathbf{M}\mathbf{U}_i^{(k)}, \tag{7.26}$$

$$\mathbf{U}_i^{(k+1)} = \mathbf{U}_i^{(k)} + \Delta \mathbf{U}_i. \tag{7.27}$$

The iteration matrix $\mathbf{M} - \Delta t^n \gamma \partial \mathbf{f} / \partial \mathbf{u}$ is evaluated and factorized once per integration step. Because the state dimension is low, modified Newton approaches, where the Jacobian is reused over several consecutive steps, will not have a significant impact. Hence, we re-evaluate and factorize the iteration matrix in each integration step. We construct an initial guess for the stage value at each internal stage by continuous extension of the method as suggested in [Thomsen1986]. The continuous extension is given by

$$\mathbf{u}(t^n + \theta \Delta t^n) = \mathbf{u}_n + \Delta t^n \sum_{i=1}^{s} b_i^*(\theta) \mathbf{f}(t^n + c_i \Delta t^n, \mathbf{U}_i), \tag{7.28}$$

where the quadrature weights, b_i^*, depend on θ according to

$$b_i^*(\theta) = b_{i,1}^* \theta + b_{i,2}^* \theta^2 + \ldots + b_{i,p}^* \theta^p. \tag{7.29}$$

The coefficients for the continuous extension are determined by requiring that the set of order conditions used to determine the coefficients of the ESDIRK method itself is satisfied. A prediction of the stage value at the i^{th} internal stage is then obtained by substituting $\theta_i = 1 + c_i \Delta t^{n+1} / \Delta t^n$ in (7.28).

The iterations are terminated when

$$\left\| \mathbf{R}^{(k)} \right\|_2 \le \kappa \cdot rtol, \tag{7.30}$$

with $\mathbf{R}$ the scaled vector of residuals, $rtol$ the relative error tolerance for the local discretization error. The parameter κ is normally chosen between 0.5 and 0.01. For reasons of robustness we chose $\kappa = 0.01$, such that the errors allowed in the iteration process are much smaller and do not interfere with the estimation and control of the local integration error. Experiments show that the cost of the extra iterations needed using $\kappa = 0.01$ compared to, for example, $\kappa = 0.1$, is minor, since the tighter value of κ also reduces the total number of integration steps.
During iterations the convergence rate is estimated as

$$\alpha_k = \frac{\left\| \mathbf{R}^{(k)} \right\|_2}{\left\| \mathbf{R}^{(k-1)} \right\|_2}. \tag{7.31}$$

If for some k during iterations $\alpha_k > 1$, the iterations are terminated and the step size decreased. The estimated convergence rate is also used by the step size controller to limit the step size if convergence in the previous step was too slow. A target minimum convergence rate of 0.3 is used, which from experiments on small scale problems similar to the VKC was found to balance the trade-off between many short integration steps, but with rapid convergence, and longer steps requiring more iterations per step.

7.2.3 Discontinuities due to phase changes

Phase changes in the VKC cause discontinuities in the right-hand-side functions of Eqns. (7.6) and (7.7). Straightforward integration across these discontinuities may lead to non-physical phase changes, poor convergence and repeated step failures. In this section, we propose an algorithm for robust detection and location of phase changes by considering the kinetic cell as a discrete event problem. The appearance or disappearance of a fluid phase marks the occurrence of a discrete event, e.g. a change from a single phase region to a two-phase region, or vice-versa. The time of the phase change can not be determined a priori. The detection of a phase change and subsequent location of the exact time of change are the main components of the proposed discrete event algorithm.

Discrete event systems

We will represent the discrete event DAE system in the following form:

$$\frac{d\mathbf{v}}{dt} = \mathbf{f}(t, \mathbf{v}, \mathbf{w}), \qquad \mathbf{v}(t_0) = \mathbf{v}_0 \tag{7.32}$$

$$0 = \mathbf{g}(t, \mathbf{v}, \mathbf{w}), \qquad \mathbf{w}(t_0) = \mathbf{w}_0 \tag{7.33}$$

$$0 < q_j(t, \mathbf{v}, \mathbf{w}), \qquad j = 1, \ldots, n_{ev} \tag{7.34}$$

where the DAE system, originally given in Eqn. (7.19), has been reformulated in semi-explicit form. Now, $\mathbf{v} \in \mathrm{R}^{m_d}$ is the vector of differential variables and $\mathbf{w} \in \mathrm{R}^{m_a}$ is the vector of `algebraic' variables. $\mathbf{f}$ and $\mathbf{g}$ are vector functions mapping $\mathrm{R} \times \mathrm{R}^{m_d} \times \mathrm{R}^{m_a}$ into R^{m_d} and R^{m_a}, respectively. We will assume that $\partial \mathbf{g} / \partial \mathbf{v}$ is non-singular (index one DAE). The q_j's are event functions associated with the current system state and n_{ev} denotes the number of event functions. In general, the system may change between many different states. If the attainable states are indexed as $p \in \mathrm{I}_{state}$ then each state is represented with $\{m_d, m_a, n_{ev}\}_p$.

To illustrate the concepts we consider the discrete event formulation of the VKC. We neglect changes associated with the water phase and consider only the two-phase equilibrium between hydrocarbon components in the oil and gas phases. Three

attainable states exist for this system: (1) single phase oil, (2) single phase gas, and (3) two-phase oil-gas. In the two-phase region, the active equations for the VKC are

$$\frac{dC_i}{dt} = \sum_\gamma A_{i\gamma} R_\gamma + \frac{F_i^{in} - F_i^{out}}{V_T}, \qquad i = 1, \ldots, n_c, \tag{7.35}$$

$$\frac{dU}{dt} = \sum_\gamma (-\Delta H_\gamma^r) R_\gamma + \frac{H^{in} - H^{out} + Q^{ext}}{V_T}, \tag{7.36}$$

$$0 = U - \left[(1-\phi_v) U^r + (\phi_v - \phi_f) \xi^s U^s + \phi_f \sum_j \xi^j S^j U^j \right], \tag{7.37}$$

$$0 = V_p - V^f - V^s, \tag{7.38}$$

$$0 = \sum_{i \in \{hc\}} \frac{C_i (K_i - 1)}{1 - \beta^g + \beta^g K_i}. \tag{7.39}$$

Here, $m_d = n_c + 1$, $m_a = 3$ and $n_{ev} = 2$. The two event functions associated with the two-phase region are

$$0 < \sum_{i \in \{hc\}} C_i (1 - K_i), \tag{7.40}$$

$$0 < \sum_{i \in \{hc\}} C_i \left(1 - \frac{1}{K_i} \right), \tag{7.41}$$

corresponding to disappearance of the gas phase and the oil phase, respectively. When a change to a single phase region occurs, Eqn. (7.39) is omitted from the formulation. Each single phase state then has one associated event function defining the criterion for a switch back to the two-phase region.

Discrete event algorithm

Several general purpose event detection algorithms have been suggested in the literature, see, for example, [Park1996] and [Thomsen2006]. Our event detection algorithm is based on the principle of discontinuity locking [Park1996]. Within each integration step the system is locked in its current state, even if one or more event functions are satisfied. The event functions are evaluated at the end of each step, and if any of them are satisfied, the exact time of occurrence is located. The approach is based on the assumption that the system of equations is mathematically well behaved in a small neighborhood of the phase changes.

Since the changes often occur between mesh points, we need to interpolate the differential and algebraic variables of the DAE in order to evaluate the event functions. For differential variables we use the continuous extension (7.28) provided by the ESDIRK methods. The algebraic variables could, in principle, also be interpolated by a suitable polynomial. However, after location of the event, a consistent initialization calculation is required to restart the integration in the new system state. That is, based on values for the differential variables, the algebraic constraints are solved to find consistent initial values for the algebraic variables. This consistency calculation may result in the numerical phenomenon referred to as *discontinuity sticking* [Park1996]. The problem is illustrated in Figure 7.2.

Assume that the event function $q^*(t, \mathbf{v}, \mathbf{w}) > 0$ was satisfied in the integration step $[t^n, t^{n+1}]$. The exact event time, t^*, may than be located by solving the scalar equation

$$q^*(t^*, \mathbf{v}^p(t^*), \mathbf{w}^p(t^*)) = \delta, \tag{7.42}$$

where $\mathbf{v}^p(t)$ and $\mathbf{w}^p(t)$ are interpolating polynomials for the differential and algebraic variables, respectively. δ is a small tolerance employed to ensure $q^* > 0$ at t^*. However, a consistent initialization at t^* may result in values of the algebraic variables that are different from $\mathbf{w}^p(t^*)$ if t^* does not coincide with the mesh points or the internal quadrature nodes of the ESDIRK scheme

$$\mathbf{w}^p(t^*) \neq \mathbf{w}(t^*), \qquad \text{if} \qquad t^* \neq t^n + c_i \Delta t^n, \quad i = 1, \ldots, s. \tag{7.43}$$

Consequently, the value of the event function may have changed

$$q^*(t^*, \mathbf{v}^p(t^*), \mathbf{w}^p(t^*)) \neq q^*(t^*, \mathbf{v}^p(t^*), \mathbf{w}(t^*)), \tag{7.44}$$

and the event function may no longer be satisfied ($q^*(t^*, \mathbf{v}^p(t^*), \mathbf{w}(t^*)) < 0$). If this is the case, the same event is detected again immediately in the next integration step.

To avoid discontinuity sticking problems, we locate the phase changes by solving the system of equations

$$\mathbf{g}(t^*, \mathbf{v}^p(t^*), \mathbf{w}(t^*)) = 0 \tag{7.45}$$

$$q^*(t^*, \mathbf{v}^p(t^*), \mathbf{w}(t^*)) = \delta, \tag{7.46}$$

where the unknowns are the algebraic variables, $\mathbf{w}$, and the event time, t^*. Only few Newton iterations are required to solve the system, since good starting guesses are available.

7.3 Results

7.3.1 VKC simulation of a ramped temperature experiment

As a first example, we illustrate the VKC by simulating a ramped temperature experiment using the minimal reaction model. Ramped temperature experiments are often carried out in the laboratory to determine burning characteristics for different oils.

Table 7.2: Initial and operational conditions for simulating the ramped temperature experiment using the VKC.

Initial and operational conditions	
Water	[0.00]*moles*
Light oil	[0.00]*moles*
Heavy oil	[0.55]*moles*

Oxygen	$[0.00]moles$
Inert gas	$[0.45]moles$
Coke	$[0.00]moles$
Temperature	$[373]K$
Pressure	$[13.88]MPa$
Total cell volume	$[0.0014]m^3$
Air feed rate	$[2.78\cdot10^{-6}]m^3s$
Rock porosity	0.40
Heat transfer coefficient	$[60]kJmole\cdot K$
Valve coefficient	$[2.74\cdot10^{-11}]m^3s\cdot Pa$

The simulated setup consists of an oil sample placed in the kinetic cell which is heated externally from $[373]K$ to $[873]K$ over a period of 10 hours with a constant feed of air of $[2.78\cdot10^{-6}]m^3s$ $[[10]Lhr]$. Initial and operational conditions are summarized in Table 7.2. Figure 7.3 shows the cell temperature, oxygen consumption rate, and concentrations of light oil, heavy oil, coke, and oxygen at three different concentrations of oxygen in the feed.

The minimal reaction model used in this experiment does not include LTO reactions in the traditional sense of oxygen addition reactions, but we still observe two peaks in oxygen consumption rate originating from direct oxidation of the oil based components and oxidation of coke which occurs at a higher temperature. The variation with oxygen feed concentration shows that, as expected, a low oxygen concentration promotes cracking of heavy oil whereas a high oxygen concentration favors the direct oxidation which leads to a significant increase in temperature.

In the second example we study the influence of the activation energy. We again simulate the ramped temperature experiment with an oxygen feed concentration of 21%. We perform a base simulation using the original reaction parameters and two perturbed simulations, in which we first increase and then decrease activation energies for all reactions by a factor of two. We choose the pre-exponential factors such that the nominal and perturbed reaction rate constants coincide at a specified temperature, so that

$$\alpha_{perturbed} = \alpha_{nominal} \cdot \exp\left(\frac{1}{R_g T_{coin.}} \left(E_{perturbed} - E_{nominal} \right) \right). \tag{7.47}$$

Here, $E_{nominal}$ and $E_{perturbed}$ refer to the nominal and perturbed activation energies, respectively, and $T_{coin.}$ is the temperature at which the two rate expressions coincide. We choose $T_{coin.} = [623]K$, which corresponds to the mean temperature during the simulation.

Figure 7.4 shows the cell temperature, oxygen consumption rate, and concentrations of light oil, heavy oil, coke, and oxygen at the three different activation energies. Again, we see two peaks in oxygen consumption rate originating from direct oxidation of the oil components and oxidation of coke. Increasing the activation energies results in an initial delay of the reaction onset time, but a more rapid oxidation of oil and coke after onset, which is also reflected in the increase in oxygen consumption rate.

7.3.2 Phase changes

For the case with 20% oxygen in the feed, a change from single phase oil to two-phase oil-gas occurs at $t = [3.15] hrs$. Attempting to integrate directly across the phase change results in repeated step failures in the solver as illustrated in Figure 7.5. The figure shows the step size sequence as selected by the controller (7.24) for a section of the integration interval near the phase change. Each time the solver attempts a step across the phase boundary, a convergence or error test failure results forcing the solver to reduce the step size. The solver fails eight times and reduces the step size by a factor of 100 before stepping across successfully. Using the phase change detection algorithm, the change is detected and consistently located without failed steps. Initialization in the new two-phase region, however, results in 4 step failures and a step size reduction by a factor of approximately 30. Completely avoiding step failures and step size reductions is difficult since a change into the two-phase region changes the dynamics of the problem.

7.3.3 Performance comparison

To evaluate the performance of the new, tailored ESDIRK implementations we have compared them to two stiff ODE solvers that are commonly used: DASSL [Petzold1982] and LSODE [Hindmarsh1983]. Both solvers are based on backward differentiation formulas (BDF), and are variable order. We show results for four ESDIRK methods, namely

- Two stage, second order trapezoidal rule with implicit Euler as error estimator.
- Three stage method, described in [Kvaernoe2004]. Second order method with third order error estimator.
- Four stage method, described in [Alexander2003]. Optimal third order method with fourth order error estimator.
- Seven stage method, described in [Kvaernoe2004]. Fourth order method with fifth order error estimator.

Coefficients for the methods can be found in the cited references. We split the total integration interval into N subintervals. We assume the pressure is constant throughout the entire integration. The temperature is kept fixed in each subinterval.

Its value is obtained by sampling the characteristic ISC temperature profile as shown in Figure 7.6. Thus, we mimic implementation of the solvers within an operator splitting environment for the full ISC equations. In each subinterval the VKC equations reduce to the set of ODEs given by Eqn. (7.6), along with the phase equilibrium constraint (7.18). For DASSL and LSODE, which are not equipped with functionality for handling discontinuities due to phase changes, we solve the implicit equation (7.18) whenever the ODE right-hand-side functions are evaluated.

Table 7.3: Initial overall oil composition for the SARA based model.

Component	Mole fraction
Water	0.0000
Inert oil	0.0005
Oxidized resins/aromatics	0.0000
Oxidized saturates	0.0000
Asphaltenes	0.0229
Resins	0.0914
Aromatics	0.2680
Saturates	0.5625
Light oil	0.0547
Carbon dioxide	0.0000
Nitrogen	0.0000
Oxygen	0.0000
Oxidized asphaltenes	0.0000
Pyrolysis coke	0.0000

We use the SARA based reaction model. In terms of number of components and reactions we think that this model best represents the requirements in realistic ISC simulations. We assume ideal fluid phases and use a pressure and temperature correlation for the equilibrium K-factors. Initial compositions are listed in Table 7.3. A pressure of [2.023]MPa is used in all subintervals. Air is cycled through the cell at a constant rate. The total simulation time is [100]hrs and two experiments are carried out using $N = 25$ and $N = 100$ subintervals, respectively.

The numerical results are compared to a very accurate reference solution computed by ESDIRK34 using $atol = rtol = 10^{-14}$. The measure of accuracy is based on the max-norm of the relative error at the end of the integration interval. The accuracy is represented as the minimum number of significant correct digits, SCD, in the solution, defined as

$$\mathrm{SCD} := -\log_{10}\left[\max_i \left|\frac{u_i(t_{end}) - u_i^{ref}(t_{end})}{u_i^{ref}(t_{end})}\right|\right]. \tag{7.48}$$

Solver comparisons are presented as work-precision diagrams where efficiency is measured using CPU time. Thus, to produce the diagrams the problem is solved using a range of input tolerances for each of the solvers tested. All codes are compiled using the Compaq Visual Fortran compiler with the same optimization settings. The CPU time for each case is averaged over 25 runs on a Pentium 4, 3GHz PC with 512MB RAM. The input tolerances are chosen as $atol = rtol = 10^{-(1+j/3)}$, $j = 0,\ldots,21$. The initial integration step size is computed internally by the solvers. For this purpose, the ESDIRK solvers use the algorithm suggested by [Hairer1996]. Between subintervals the last step of the current interval is passed as a guess for the first step in the next interval. All solvers require the Jacobian of the ODE right-hand-side functions. These are computed analytically using the automatic differentiation tool Tapenade ver. 2.2 [Tapenade].

Figure 7.7 shows the work-precision diagrams when solving the SARA based reaction model using 25 and 100 subintervals. The resulting accuracies in the solution are, of course, different from the local tolerances provided as input to the solvers that govern the local error and step size control. In the high accuracy range (4-6 significant correct digits) the methods perform according to their order: The high order methods (DASSL, LSODE and ESDIRK45) outperform the lower order methods (e.g. ESDIRK12 and ESDIRK23). However, we are mainly interested in the accuracy range corresponding to 2-4 significant correct digits, since we expect the accuracies obtained in other parts of the equation to fall in this range. The low and intermediate order methods perform best in this range. For example, with 100 subintervals and three significant correct digits the ESDIRK23 solver is two times faster than LSODE and 3-4 times faster than DASSL.

The work-precision curves for the ESDIRK solvers settle at a constant level in the low accuracy range instead of continuing the linear trend found for high accuracies. Moreover, the behavior in this range is more erratic. This can be partly attributed to the error and convergence controller. Consider, for example, ESDIRK45. An input tolerance level of $atol = rtol = 10^{-1}$ gives better than expected results (SCD = 2), but this comes at the cost of increased computation time. Inspection shows that the step size is limited by convergence in the nonlinear solver and not by the local error. The relatively large step sizes allowed by the error controller lead to convergence failures in the nonlinear solver, which then forces a step size reduction. This shows that ODE error control devices perform best when operating at low tolerance levels.

Comparing the results for $N = 25$ and $N = 100$ shows that all solvers need more time when the integration is interrupted frequently. The BDF methods have the largest restart overhead because, when restarted, they revert to first order requiring small time steps, and slowly build up higher order information. Both DASSL and LSODE manage to integrate across the phase changes. Inspection of the process shows, however, that both methods experience a high number of step failures, and lower their order before successfully stepping across. The loss in efficiency associated with these step failures alone is difficult to isolate but is

reflected in the results shown in the figure. The differences in performance are relatively small when using only 25 subintervals. The advantage of one-step methods over multi-step methods is, however, clearly observed when the number of subintervals is increased, the difference between DASSL and the ESDIRK solvers being the most notable. None of the solvers are very sensitive to the required accuracy. Although with some irregularities in the low accuracy range, the computational cost increases linearly with the required accuracy in the solution.

7.4 Discussion

The work presented in this chapter addresses time integration of reactive, compositional, and thermal porous media processes, and focusses in particular on the design of specialized solvers for integration of stiff chemical kinetics subject to phase equilibrium constraints.

To facilitate the algorithmic development we constructed a virtual kinetic cell. The VKC allows us to study the kinetics and phase behavior of reactive transport processes in an isolated setting. Through user prescribed inflow and outflow boundary conditions, interactions with the reservoir can be studied. Two ISC reaction models with 6 and 14 components, respectively, were implemented and tested. The VKC was demonstrated with a ramped temperature experiment using the 6-component model. Ramped temperature experiments are routinely carried out in the laboratory. In future work we will validate the simulation results against experimental data. Although not explored in this work, the VKC can also be applied to help determine optimal experimental conditions in the design phase. For example, parametric sensitivity studies can be used to determine regions of the parameter space where additional laboratory experiments should be carried out. A more rigorous approach involves optimal model based design, where the experimental conditions (such as oxygen flow rate and concentration in ramped temperature kinetic cell experiments) are optimized in an iterative loop. The optimality criterion is then related to the statistical quality of the parameter estimates that results from using the experimental data to estimate model parameters. Optimal experimental design is explored in e.g. [Bock2000].

Because of the stiff chemical kinetics, the low to medium accuracy required in the solution, and the short integration intervals experienced in an operator splitting environment, we selected the class of ESDIRK methods for the temporal integration of the chemical reactions. To increase computational efficiency, we select the time step size with a predictive controller for the local integration error, and generate starting guesses for the modified Newton iterations using continuous extensions of the ESDIRK methods. Performance comparisons between the ESDIRK solvers and the popular stiff ODE solvers DASSL and LSODE, show that the higher order ESDIRK solvers are at least comparable if not superior in terms of computational speed, especially over the relatively short integration intervals that are generally required in an operator splitting environment. If we aim for two to four correct digits in the

solution, which we expect to be of practical interest in ISC simulations, the low and medium order ESDIRK methods (ESDIRK12, ESDIRK23 and ESDIRK34) are two times faster than LSODE and three or four times faster than DASSL. Apart from the ESDIRK methods, other classes of integration methods could be of interest, such as the Rosenbrock methods [Sandu1997a,Sandu1997b], which we will consider in future work.

We extended the ESDIRK methods to handle the discontinuities that arise due to phase changes. Phase changes are detected by monitoring sign changes of special event functions. The proposed algorithm proved robust in detecting and locating phase changes, and lowered the number of convergence and error test failures by more than 50%. The DASSL and LSODE solvers that are not equipped with a phase change algorithm will, in most cases, successfully integrate across the change after repeated step size reductions caused by convergence and error test failures. Overall efficiency gains from using the phase change algorithm when measured over a long integration interval with only one or two phase changes occurring are modest, but the improved robustness near phase boundaries is certainly valuable.

The VKC model uses a simplified constant K-value phase equilibrium description. We are currently working on the extension to full equation-of-state based phase equilibrium. In a rigorous flash algorithm the phase state depends on the outcome of a stability analysis which cannot be expressed in closed form and, hence, does not directly fit into the framework of event functions. For transitions from a two-phase region to a single phase region, we can still detect changes by monitoring phase fractions. In the reverse situation the problem is now more complicated. In principle, this would require a stability analysis to be performed each time conditions change to check if the new single phase state is unstable. A possible solution is to exploit shadow regions as in [Michelsen2006] to speed up flash calculations, which will reduce the number of stability analyses required.

7.5 Summary

The virtual kinetic cell (VKC) is a novel and useful tool for analyzing kinetics and phase behavior in reactive, compositional, and thermal processes. It can provide simulation support for laboratory kinetic cell experiments, and be used for studying interactions between kinetics and phase behavior as well as interactions with the reservoir through specialized boundary conditions.
ESDIRK methods are well-suited for integration of the stiff kinetics due to their strong stability properties. In particular, when implemented in an operator splitting environment, the ESDIRK methods outperform stiff multi-step methods. Experiments show that the methods lead to 50-75% reductions in computational costs compared to multi-step methods.
Phase change detection using a discrete event system approach is an attractive tool which significantly reduces integration step failures and therefore computational time.

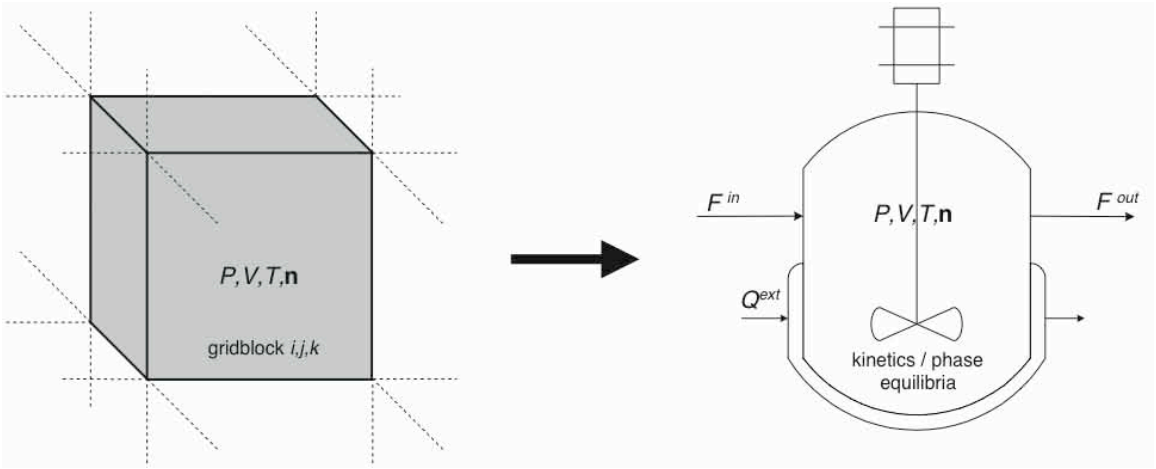

Figure 7.1: Using an operator splitting approach, each gridblock is effectively treated as a small chemical reactor.

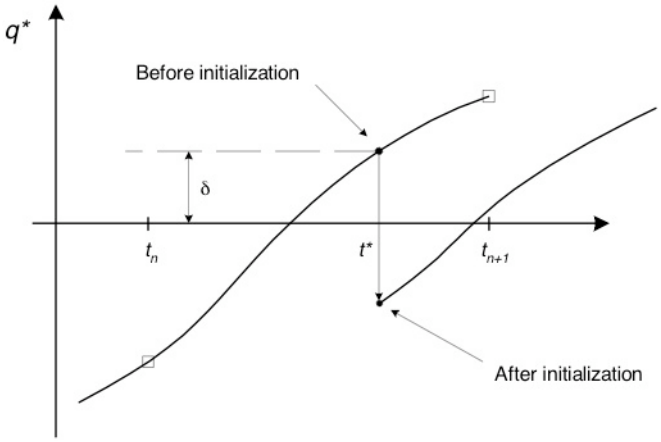

Figure 7.2: Illustration of discontinuity sticking resulting from inconsistent algebraic variables. (Figure adapted from [Park1996]).

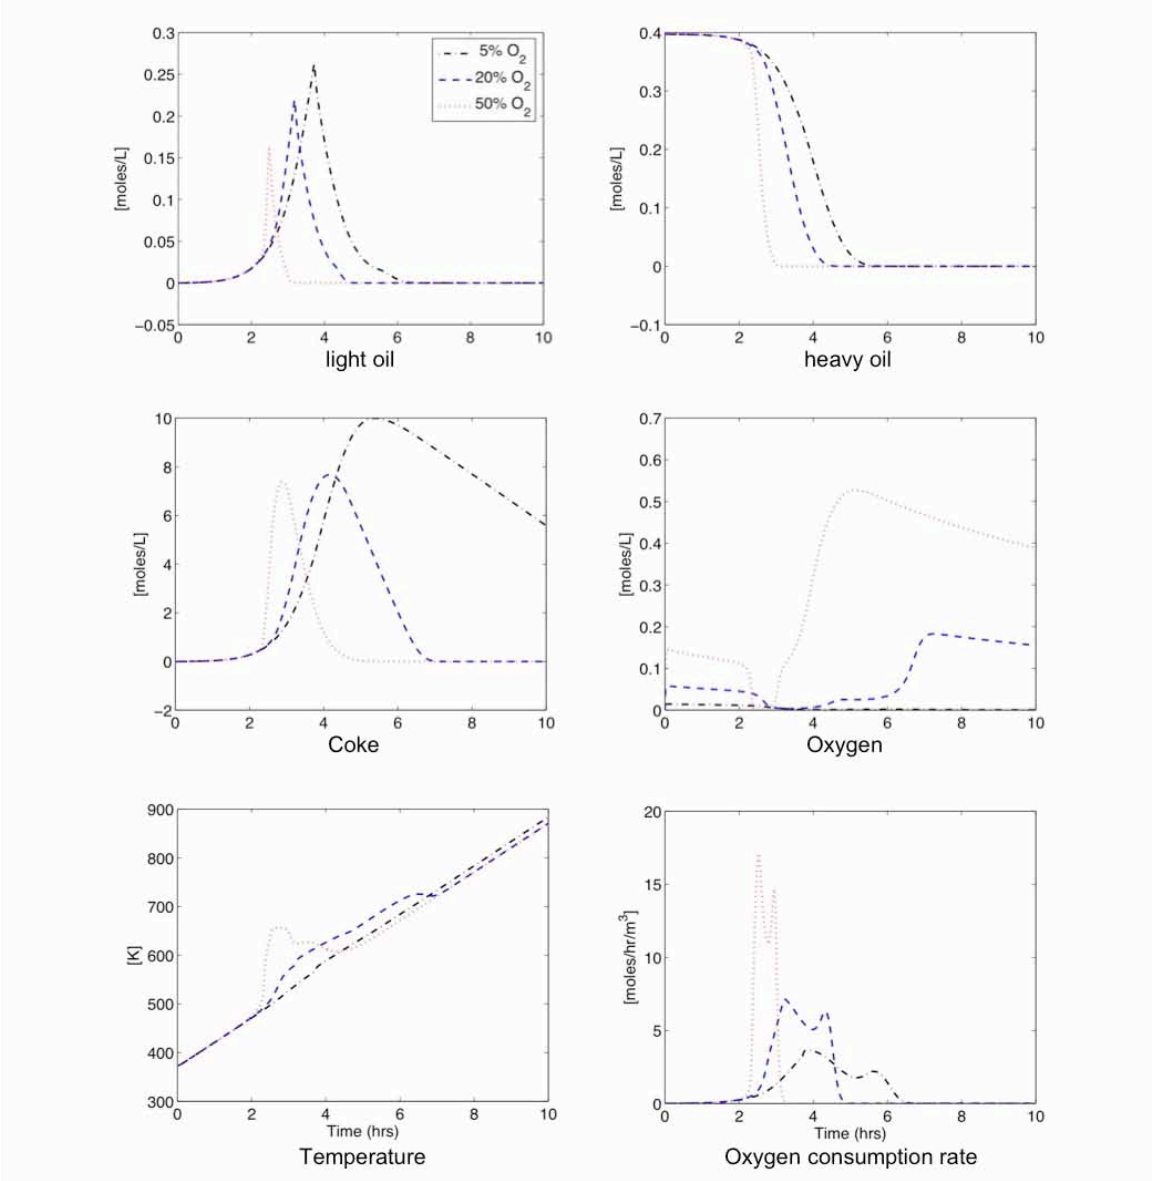

Figure 7.3: Simulation results from a ramped temperature experiment using the minimal reaction model. The temperature is raised from $[373]K$ to $[873]K$ over a period of 10 hours. Cell temperature, oxygen consumption rate, and component concentrations are shown for 3 different oxygen feed concentrations.

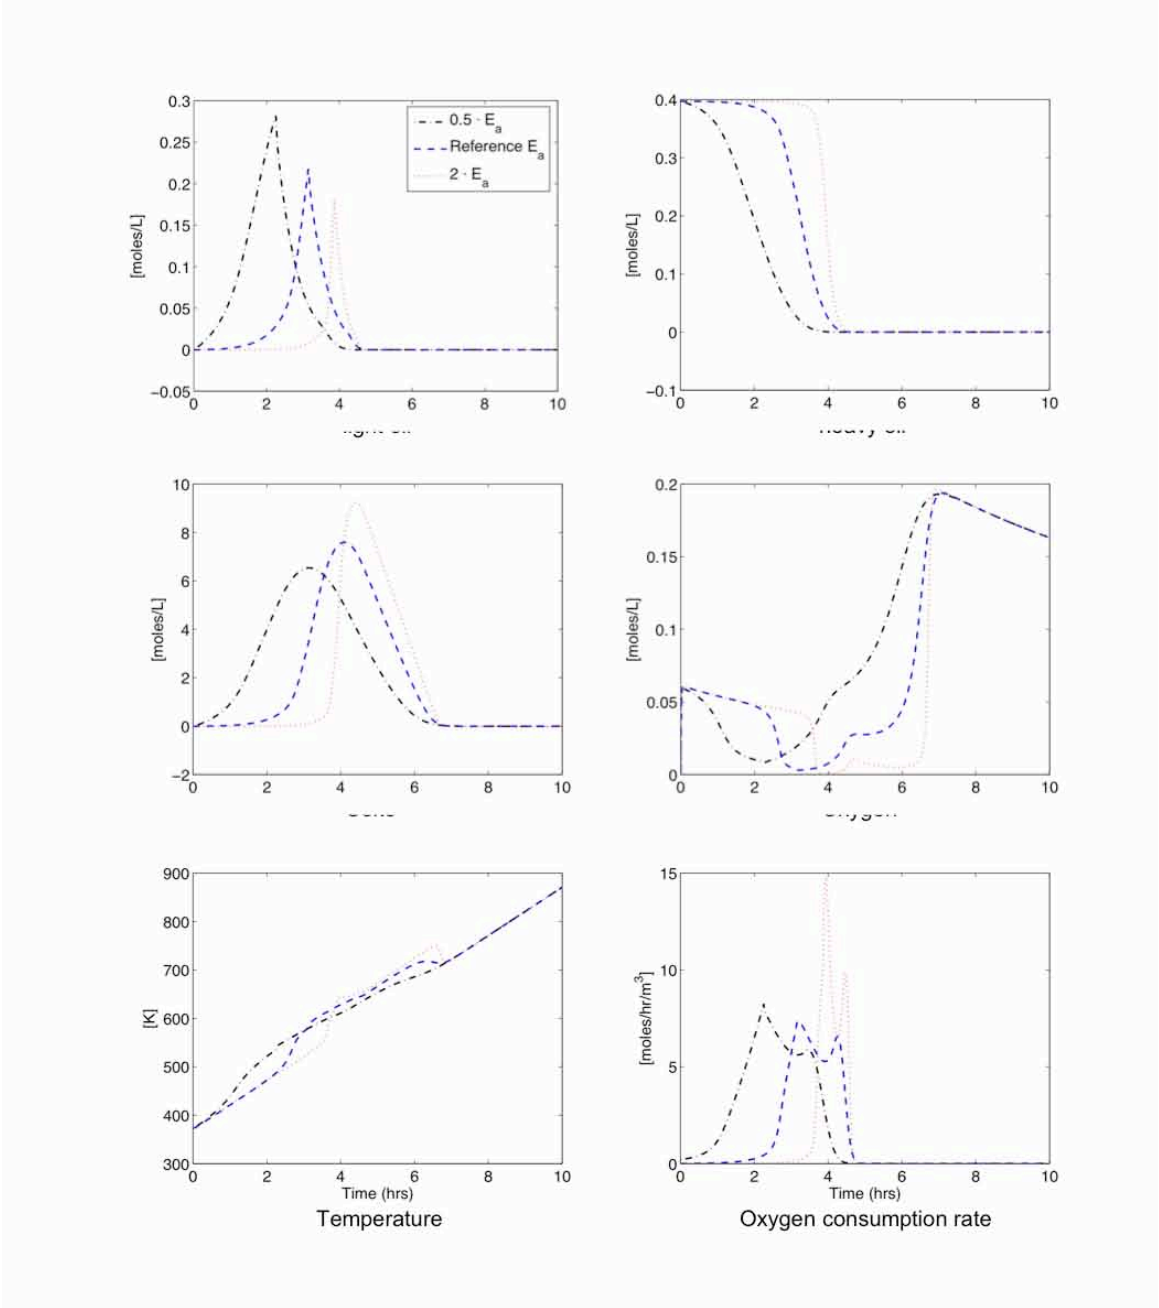

Figure 7.4: Simulation results from a ramped temperature experiment using the minimal reaction model. The temperature is raised from [373]K to [873]K over a period of 10 hours. Cell temperature, oxygen consumption rate, and component concentrations are shown for at 3 different reaction activation energies.

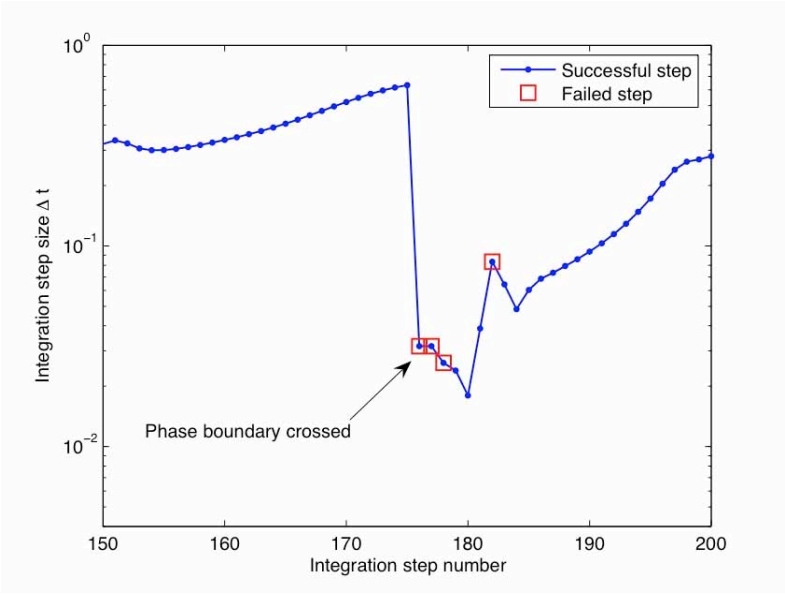

Without phase change detection

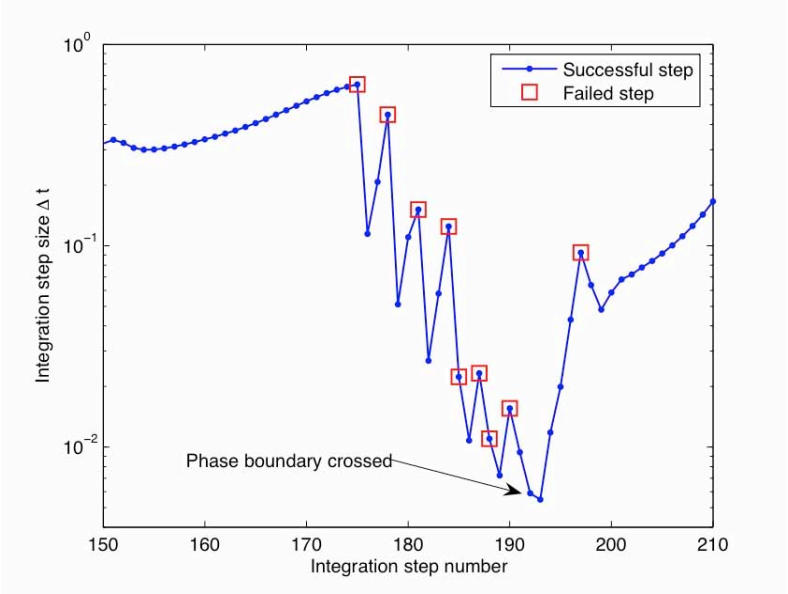

With phase change detection

Figure 7.5: Step size sequences for the ESDIRK23 solver near a phase change when simulating the ramped temperature experiment. The phase change algorithm improves the integration robustness when crossing phase boundaries.

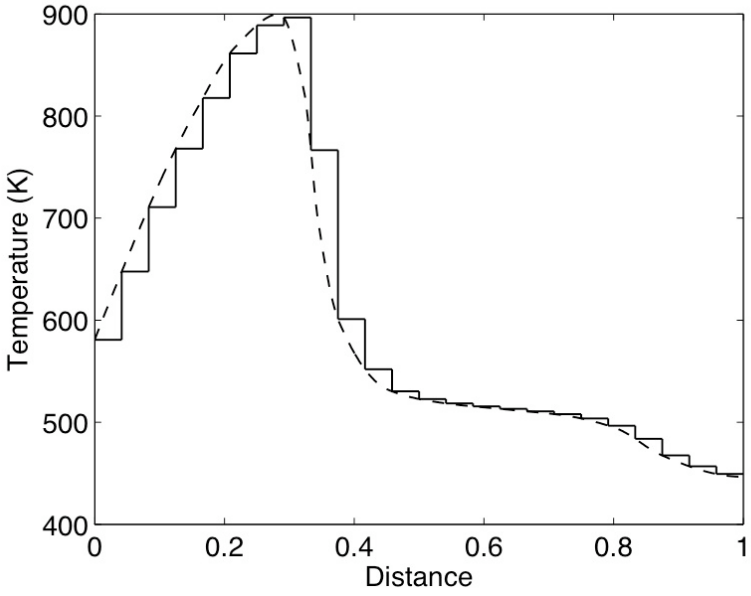

Figure 7.6: The temperatures for the integration subintervals in the performance comparison are obtained by sampling the characteristic ISC temperature profile. The continuous temperature profile is shown along with a zero-order parametrization corresponding to $N = 25$ equidistant subintervals.

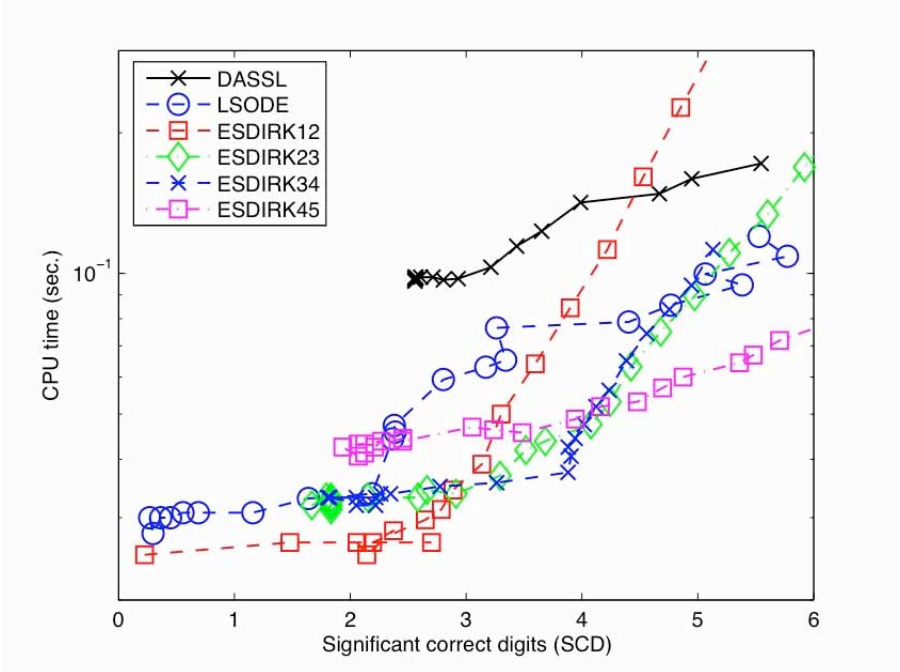

100 subintervals

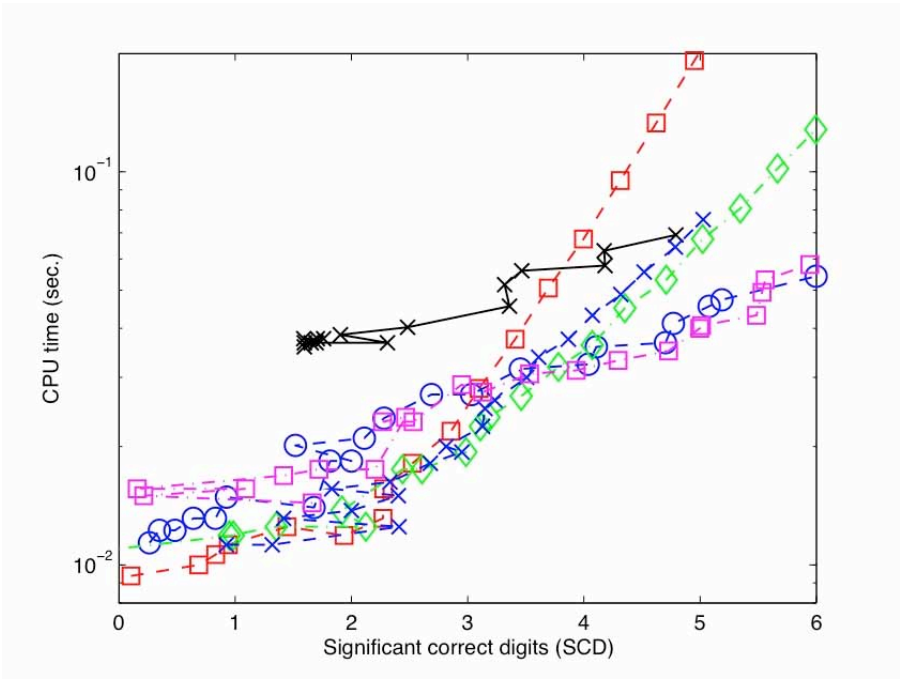

25 subintervals

Figure 7.7: Work-precision diagrams for the SARA based ISC reaction model. The tailored ESDIRK implementations are compared to the widely used DASSL and LSODE codes. Comparisons are made for 25 and 100 subintervals

8 Simulation model II - Virtual Kinetic Tube and Sensitivity to Phase Behavior

In the previous chapter we presented tailored integration methods for stiff kinetics along with algorithms for robust handling of phase changes, which we tested using a Virtual Kinetic Cell (VKC). In this chapter, we extend our study to the influence of phase behavior treatment on kinetics and on performance predictions. Compositional effects and their impact on reaction paths, and hence on overall process performance, are not fully understood for processes like ISC. There are two main reasons why we expect kinetics and performance prediction to be sensitive to phase behavior models. First, chemical reactions in ISC processes are assumed to be heterogeneous gas-liquid or gas-solid reactions, where reaction rates depend on phase component concentrations, which are sensitive to phase behavior. Second, in ISC processes a limited number of pseudo-components is used to represent the very large number of liquid and gas components in reservoir settings. The pseudoization is generally driven by the kinetics, not the phase behavior. For example, a pseudo-component may be used to model kinetically similar hydrocarbons with widely varying boiling points. Reliable phase behavior models are therefore difficult to develop, and it is natural to expect sensitivities to the chosen formulation. We again use the VKC to study this interaction between phase behavior and kinetics, as well as a novel one-dimensional, in-house ISC simulator, which we refer to as the Virtual Combustion Tube (VCT). Both the VKC and the VCT are equipped with a family of phase behavior models, including a Peng-Robinson equation of state [Peng1976], and two types of equilibrium K-value correlations described below.

To test the sensitivity, we developed a one-dimensional combustion tube model (VCT). The VCT is fully compositional with a dual strategy for describing the phase behavior, using either equilibrium K-value correlations or the Peng-Robinson Equation of State (EoS). The VCT is the first EoS based ISC simulator published in the literature. As in the VKC, two reaction models are implemented: the so-called minimal model, used in STARS and most commercial applications, and the more involved SARA fractions model, suggested by e.g. [Freitag2005]. The VCT and the VKC provide a flexible platform for the study of laboratory scale combustion experiments.

We perform extensive grid sensitivity tests on the VCT to better understand the grid dependency of ISC simulations. The test results are used to guide the selection of grid step sizes to ensure reliable, grid independent, numerical solutions.

We investigate the interaction between kinetics and phase behavior models using both the VKC and the VCT. In particular, we analyze ignition/extinction phenomena, and show that in certain operation regimes the choice of phase behavior model will critically change the reaction dynamics and the performance of the ISC process.

8.1 A Virtual Combustion Tube

We model a porous media system containing n_c chemical components distributed among three mobile fluid phases (oil, water, and gas) and one immobile solid phase. Components existing in the solid phase are assumed to do so exclusively. The mobile components are assumed to exist in at most two phases, with the water component partitioning into the water and gas phases, and the remaining components partitioning into the oil and gas phases. The chemical components are indexed as

$$i \in \mathsf{I} = \mathsf{I}^{ol} \cup \mathsf{I}^{w} \cup \mathsf{I}^{s},$$

where

$$\begin{aligned} \mathsf{I}^{ol} &= \{1, \ldots, n_{ol}\}, \\ \mathsf{I}^{w} &= \{n_{ol} + 1\}, \\ \mathsf{I}^{s} &= \{n_c - n_s + 1, \ldots, n_c\}. \end{aligned}$$

Thus, the total number of components is $n_c = n_{ol} + n_s + 1$.

8.1.1 Governing equations

Let the porous medium be represented by Ω with boundary S, then the governing equations for ISC can be derived from n_c component conservation equations, an energy conservation equation, and a volume conservation equation. The i-th component conservation equation can be written in integral form as

$$\frac{d}{dt}\int_{\Omega} C_i \, d\Omega + \int_{\mathsf{S}} \mathbf{q}_i^m \cdot \mathbf{n} \, d\mathsf{S} = \int_{\Omega} \left(Q_i^{m,reac} + Q_i^{m,well} \right) d\Omega, \tag{8.1}$$

where $i \in \mathsf{I}$, C_i is the total concentration of component i per bulk volume, $\mathbf{q}_i^m$ is the mass flux, and $Q_i^{m,reac}$ and $Q_i^{m,well}$ are the mass source densities due to chemical reactions and wells, respectively. We note that

$$\mathbf{q}_i^m = \mathbf{0}, \qquad i \in \mathsf{I}^s.$$

Neglecting capillary forces, we can express the component flux term in (8.1) using the standard multi-phase extension of Darcy's law

$$\mathbf{u}^j = -\frac{k_r^j}{\mu^j} \mathbf{k} \left(\nabla P - \rho^j g \nabla D \right) \tag{8.2}$$

where $\mathbf{u}^j$ is the phase flow velocity, $\mathbf{k}$ is the absolute permeability, k_r^j is the relative permeability, μ^j is the viscosity, ρ^j is the mass density, g is the gravitational constant, and D is the reservoir depth. The total component flux is the sum of the fluxes in the individual phases, given by

$$\mathbf{q}_i^m = \sum_j x_i^j \xi^j \mathbf{u}^j, \tag{8.3}$$

where $j \in \{o, g\}$ for $i \in \mathsf{I}^{ol}$, and $j \in \{w, g\}$ for $i \in \mathsf{I}^{w}$. x_i^j are the mole fractions of components i in phase j, and ξ^j the molar densities of phase j.
The energy conservation equation is

$$\frac{d}{dt}\int_\Omega U d\Omega + \int_\mathsf{S} \left(\mathbf{q}^{h,adv} + \mathbf{q}^{h,cond}\right)\mathbf{n} d\mathsf{S} = \int_\Omega \left(Q^{h,reac} + Q^{h,well}\right) d\Omega, \tag{8.4}$$

where U is the total internal energy of the system, $\mathbf{q}^{h,adv}$ is the heat transport due to advection, $\mathbf{q}^{h,cond}$ is the heat transport due to conduction, and $Q^{h,reac}$ and $Q^{h,well}$ are the heat source densities due to chemical reactions and wells, respectively. The advective heat transport is the sum of the heat carried by each mobile phase

$$\mathbf{q}^{h,adv} = \sum_j \xi^j \mathbf{u}^j h^j, \tag{8.5}$$

where h^j is the molar enthalpy. Fourier's law is used to express the transport due to conduction

$$\mathbf{q}^{h,cond} = -\mathbf{k}_c \nabla T, \tag{8.6}$$

where the tensorial quantity $\mathbf{k}_c$ is the effective thermal conductivity of the saturated medium.
Finally, the principle of volume conservation expresses that the fluid and solid phases must fill the pore space exactly

$$V_p = \sum_j V^j, \tag{8.7}$$

where V_p is the void pore volume and V^j is the volume of phase j for $j \in \{o, w, g, s\}$.

Phase equilibrium

In addition to the conservation requirements, we assume that the system is in thermodynamic equilibrium as expressed by the equality between chemical potentials, or equivalently fagacities, of a component in all phases. For components partitioning between the oil and gas phases we have

$$f_i^o = f_i^g, \qquad i \in \mathsf{I}^{ol}, \tag{8.8}$$

where f_i^j is the fugacity of component i in phase j. We assume that the partitioning of the water component between the gas and water phases can be modeled using a simple K-value correlation.

8.1.2 The Numerical Model

We use the n_c mass balances and the energy balance as primary equations, and as variables we choose the n_c overall component concentrations, C_i, and temperature and pressure. The advantage of working with overall quantities as variables, as opposed to intensive variables such as mole fractions and saturations, is that we avoid the switching of variables in implicit methods associated with phase

disappearance/reappearance. For thermal-compositional models the thermodynamic state (intensive and extensive) of a system is uniquely described in terms of $n_c + 1$ independent primary variables. For convenience, we work with the full set of component concentrations and include the volume constraint as an additional equation. Thus, we solve simultaneously a set of $n_c + 2$ equations per gridblock, where temperature is aligned with the energy balance, pressure is aligned with the volume constraint, and the component concentrations are aligned with the respective component balances. For the purpose of the work presented in this paper we are concerned more with robustness than computational efficiency, but we note that the volume constraint and component balances for all immobile solid components can be eliminated from the primary equation set, since both types of equations are local to each gridblock.

Temporal Discretization

For stability reasons, the time discretization is performed using the backward Euler scheme treating all flow, well and reaction terms implicitly. This fully implicit, control-volume based scheme is common in reservoir simulation, and it is also implemented in the STARS simulator STARSmanual, which allows us to validate our code with this commercial code. The discretized equations comprise a set of nonlinear algebraic equations, which are solved by Newton-Raphson iterations. We perform exact flash calculations in each Newton-Raphson iteration step.

We use the timestep formula due to [Aziz1979] given by

$$\Delta t^{n+1} = \Delta t^n \min_{i,k} \left[\frac{(1+\lambda)\Delta u_{ik}^n}{\Delta u_i^* + \lambda \Delta u_{ik}^n} \right], \tag{8.9}$$

where Δu_{ik}^n is the change in variable i in gridblock k over the previous timestep, Δu_i^* is the target change, and λ is a tuning factor. For the simulations presented in this paper we use $\Delta P^* = [1]atm$, $\Delta T^* = [5]K$, $\Delta z_i^* = 0.05$ and $\lambda = 0.75$, where $z_i = C_i / \sum_i C_i$ is the overall mole fraction of component i.

Control-Volume Discretization

The governing equations for the VCT are discretized using a control-volume approach on a block-centered Cartesian grid with standard two-point flux approximations.

Consider the 1D case with equidistant grid-spacing, and let $S_k = \{S_{k+12}, S_{k-12}\}$ denote the surfaces of gridblock k. Introducing the flow potential

$$\nabla \Phi^j = \nabla P - \rho^j g \nabla D,$$

the flow terms can then be discretized as

$$\begin{aligned}\int_{S_k} \mathbf{q}_i^m \cdot \mathbf{n} dS &= -\int_{S_k} \sum_j \frac{x_i^j \xi^j k_r^j}{\mu^j} \mathbf{k} \nabla \Phi^j \cdot \mathbf{n} dS \\ &\approx \sum_j \Lambda_{i,k+12}^j T_{k+12} \Delta^+ \Phi_k^j - \Lambda_{i,k-12}^j T_{k-12} \Delta^- \Phi_k^j,\end{aligned} \tag{8.10}$$

where $T_{k\pm12}$ is the transmissibility for the gridblock connections, $\Lambda_{i,k\pm12}^j$ is the generalized mobility

$$\Lambda_{i,k\pm12}^j = \left(\frac{x_i^j \xi^j k_r^j}{\mu^j} \right)_{k\pm12},$$

which is evaluated in the upstream direction, and the discrete flow potentials are

$$\begin{aligned}\Delta^+ \Phi_k^j &= P_{k+1} - P_k - \rho_{k+12}^j g\left(D_{k+1} - D_k\right) \\ \Delta^- \Phi_k^j &= P_k - P_{k-1} - \rho_{k-12}^j g\left(D_k - D_{k-1}\right),\end{aligned} \tag{8.11}$$

where the mass density in the gravity term is evaluated as a saturation weighted average

$$\rho_{k\pm12}^j = \frac{S_{k\pm1}^j \rho_{k\pm1}^j + S_k^j \rho_k^j}{S_{k\pm1}^j + S_k^j}. \tag{8.12}$$

The heat flux terms are discretized in a similar way

$$\begin{aligned}\int_{S_k} \left(\mathbf{q}^{h,adv} + \mathbf{q}^{h,cond}\right) \mathbf{n} dS = -\int_{S_k} \left(\sum_j \frac{h^j \xi^j k_r^j}{\mu^j} \mathbf{k} \nabla \Phi^j + \mathbf{k}_c \nabla T \right) \cdot \mathbf{n} dS \\ \approx \sum_j \Lambda_{k+12}^{j,h} T_{k+12} \Delta^+ \Phi_k^j - \Lambda_{k-12}^{j,h} T_{k-12} \Delta^- \Phi_k^j \\ + \frac{S_{k+12} k_{c,k+12}}{\Delta x} \Delta^+ T_k - \frac{S_{k-12} k_{c,k-12}}{\Delta x} \Delta^- T_k,\end{aligned} \tag{8.13}$$

where the fluid dependent parts of the interface heat conductivity $k_{c,k\pm12}$ are evaluated upstream in. The fluid dependent part in the discretized advective term is

$$\Lambda_{k\pm12}^{j,h} = \left(\frac{h^j \xi^j k_r^j}{\mu^j} \right)_{k\pm12}.$$

Phase equilibrium

In the K-value approach we assume ideal mixing of pure component properties, whereas for the EoS approach we use the EoS itself to model phase densities and enthalpies. In all the sensitivity studies in the paper we have used the simplified set of property models, so that the differences observed reflect only differences in phase behavior treatment.

For the K-value approach, we use the Wilson correlation

$$\ln K_i = \frac{P_i^{crit}}{P} + 5.373\left(1 + \omega_i\right)\left(1 - \frac{T_i^{crit}}{T}\right), \tag{8.14}$$

where P_i^{crit}, T_i^{crit} and ω_i are the critical pressure and temperature and accentric factor, respectively, for component i.

In the EoS approach, the fugacities are obtained from an EoS. We use the Peng-Robinson [EoS Peng1976]. Details on the EoS and computation of fugacities can be found in e.g. [Michelsen2004] or [Whitson2000].

Chemical reactions

The chemical reactions are assumed to be kinetically driven, and are modeled using standard Arrhenius rate relation, where the rate depends on component concentrations in a specific phase. Generally, two different kinds of reactions dominate ISC processes: cracking/pyrolysis reactions and oxidation (full/partial) reactions. The thermal cracking of component i in phase j is modeled as

$$R_\gamma = K_\gamma \cdot \left(\phi_f S^j \xi^j x_i^j\right)^{m_\gamma}, \tag{8.15}$$

where γ is used to index the reaction number and m_γ is the reaction order in the component concentration. Oxidation reactions are generally heterogeneous reactions between oxygen in the gas phase and an oil component in the oil phase, and they are modeled using the oxygen partial pressure

$$R_\gamma = K_\gamma \cdot \left(P x_{o_2}^g\right)^{n_\gamma} \cdot \left(\phi_f S^j \xi^j x_i^j\right)^{m_\gamma}. \tag{8.16}$$

The rate constants are temperature dependent according to Arrhenius's law

$$K_\gamma = \alpha_\gamma \exp\left(-\frac{E_a^\gamma}{R_g T}\right), \tag{8.17}$$

where α_γ is the frequency factor, E_a^γ is the activation energy and R_g is the universal gas constant. From the reaction rates we can express the mass and heat source densities. The net production of component i in chemical reactions is

$$Q_i^{m,reac} = \sum_\gamma A_{i\gamma} R_\gamma, \tag{8.18}$$

where $A_{i\gamma}$ is the stoichiometric coefficient for component i in reaction γ (negative for reactants, positive for products). Similarly, the net heat generation is

$$Q^{h,reac} = \sum_\gamma (-\Delta H)_\gamma^r R_\gamma, \tag{8.19}$$

where $-\Delta H^r$ is the heat of reaction.

Rock and fluid properties

For comparison purposes we have implemented the same set of property models as used by the commercial simulator CMG STARS [STARSmanual]. In addition, when using the EoS phase behavior approach, we have included an option to use densities and residual enthalpies from the EoS.

Relative permeabilities are modeled using the normalized Stone II model [Stone1973].

The porosity available to fluid flow depends on solid concentration and is modeled as

$$\phi_f = \phi_v - \sum_{i \in s} \frac{C_i}{\xi^s}, \tag{8.20}$$

where C_i is the concentration per unit reservoir volume.

Phase densities are evaluated from the EoS. We may write a general cubic EoS in the form

$$\left(Z^j\right)^3 + c_2\left(Z^j\right)^2 + c_1\left(Z^j\right) + c_0 = 0, \tag{8.21}$$

where Z^j is the phase compressibility factor, and where the coefficients c_2, c_1 and c_0 depend on pressure, temperature and phase composition. Solving Eqn. (8.21) for Z^j we can express the molar and mass phase densities as

$$\xi^j = \frac{P}{R_g T Z^j}, \qquad \rho^j = \xi^j \sum_i x_i^j M_i. \tag{8.22}$$

In the ideal case we assume $Z^g = 1$, and pure component liquid densities are obtained from

$$\xi_i^o = \xi_{i,ref}^o \exp\left(c_i^p (P - P_{ref}) - c_i^t (T - T_{ref})\right) \tag{8.23}$$

where $\xi_{i,ref}^o$ is the density at the reference conditions, and the coefficients c_i^p and c_i^t are the compressibility and thermal expansion, respectively. Linear mixing of molar volumes is assumed

$$\frac{1}{\xi^o} = \sum_i \frac{x_i^o}{\xi_i^o}. \tag{8.24}$$

Water density is calculated according to a correlation similar to (8.23).

Accurate modeling of oil phase viscosity is important for ISC processes, since the reduction in oil viscosity with temperature is the primary mechanism for enhanced recovery. Pure component viscosities can be entered either through tables or obtained from the correlation

$$\mu_i^o = a_i^o \exp\left(\frac{b_i^o}{T}\right), \tag{8.25}$$

where a_i^o and b_i^o are correlation constants. The oil phase viscosity is then obtained by logarithmic mixing of the pure component viscosities

$$\ln \mu^o = \sum_i x_i^o \ln \mu_i^o. \tag{8.26}$$

Accuracy in gas phase viscosity is generally less important. The correlation used here is

$$\mu_i^g = a_i^g T^{b_i^g}, \tag{8.27}$$

and the mixing rule used is

$$\mu^g = \sum_i \frac{x_i^g \sqrt{M_i}\,\mu_i^g}{x_i^g \sqrt{M_i}}. \tag{8.28}$$

Water viscosity is obtained from a correlation similar to (8.25).

Phase enthalpies are computed from ideal gas heat capacity correlations and vaporization enthalpy correlations. For a pure component gas we have

$$h_i^{g,id} = \int_{T_{ref}}^{T} C_{p,i}(\bar{T})\,d\bar{T}, \tag{8.29}$$

where T_{ref} is the temperature of the reference state. In the ideal case the vaporization enthalpy is obtained from the correlation [STARSmanual]

$$h_i^{vap} = hv_i^1 \left(T_i^{crit} - T\right)^{hv_i^2}, \tag{8.30}$$

where T_i^{crit} is the component critical temperature and hv_i^1 and hv_i^2 are correlation constants. Thus, for the ideal case we have

$$h^g = \sum_i x_i^g h_i^{g,id} \tag{8.31}$$

for the gas phase and

$$h^o = \sum_i x_i^o \left(h_i^{g,id} - h_i^{vap}\right) \tag{8.32}$$

for the oil phase. In the non-ideal case the residual enthalpies are obtained from the EoS leading to

$$h^g = \sum_i x_i^g h_i^{g,id} + h^{g,res} \tag{8.33}$$

$$h^o = \sum_i x_i^o h_i^{g,id} + h^{o,res}, \tag{8.34}$$

where $h^{j,res}$ is the residual enthalpy of phase j. The water phase enthalpy is obtained from correlations similar to (8.29) and (8.30).

The total internal energy can be expressed as

$$U = \left(1-\phi_v\right)U^r + \left(\phi_v - \phi_f\right)\xi^s U^s + \phi_f \sum_j \xi^j S^j U^j, \tag{8.35}$$

where U^r is the volumetric internal energy of rock, and U^j is the molar internal energy of phase j, which is obtained from the enthalpy by subtraction of the mechanical work

$$U^j = h^j - \frac{P}{\xi^j}. \tag{8.36}$$

The mobile and immobile phases have associated with them a scalar thermal conductivity k_c^j. Moreover, the rock has an associated tensorial thermal conductivity denoted $\mathbf{k}_c^r$. The effective thermal conductivity of the saturated porous medium is computed as a weighted volume average

$$\mathbf{k}_c = (1-\phi_v)\mathbf{k}_c^r + \left[(\phi_v - \phi_f)k_c^s + \phi_f \sum_j S^j k_c^j\right]\mathbf{I} \tag{8..37}$$

In this paper we assume the rock conductivity to be isotropic.

Well models

We consider wells completed in a single gridblock and two types of well operations: fixed single-phase injection rate for injectors and fixed back pressure for producers. Typically, a fixed gas rate (air or oxygen enriched air) is specified, or both water and gas rates are specified as for wet combustion processes. The component source density due to wells can be written

$$Q_i^{m,well} = -\sum_j \frac{x_i^j \xi^j k_r^j}{\mu^j} WI(P - P_w) \tag{8.38}$$

where WI is the well index, P is the block pressure and P_w is the well bottom-hole pressure. Expressions for the well index for standard block-centered cartesian grids can be found in [Aziz1979]. The heat source density due to wells can be written as

$$Q^{h,well} = -\sum_j \frac{h^j \xi^j k_r^j}{\mu^j} WI(P - P_w) \tag{8.39}$$

The outer boundaries of the reservoir are assumed closed, and mass is only transported in and out of the system through wells. However, in general heat losses will occur to the surrounding formation through heat conduction. Since in this paper we are primarily concerned with the simulation of laboratory scale combustion tubes, we will model heat losses using an external proportional heat controller

$$Q^{h,htr} = -u_a(T - T^{ext}) \tag{8.40}$$

where u_a is a heat transfer coefficient (or controller gain) and T^{ext} is the external heating/cooling temperature (controller set point).

8.2 Results

8.2.1 Sensitivity Studies

In order to validate our implementation we carry out spatial and temporal grid convergence studies and perform careful comparisons with STARS. We simulate a laboratory dry air combustion tube test using the minimal reaction model. Air is

injected at a rate of $[0.016]m^3hr$ into an oil saturated core, which is preheated to $[311]K$. During the first 30 minutes of simulation the injection end is heated by a constant energy input of $[211]kJhr$, which establishes a combustion front. We refer to [Kristensen2008] for a complete list of input parameters for the simulation. We simulate the system until $t^* = [8]hrs$ using the time step strategy outlined above. The simulation is repeated for varying grid densities using $N_{blk} = 10 \cdot 2^n$, $n = 0,1,\ldots,7$, and compared to a reference solution generated using $N_{blk} = 2560$. The error is measured as

$$e_p = \frac{\mathsf{P}\mathbf{u}(t^*) - \mathbf{u}^{ref}(t^*)\mathsf{P}_p}{\mathsf{P}\mathbf{u}^{ref}(t^*)\mathsf{P}_p}, \qquad p = 1,2. \tag{8.10}$$

Figure 8.1 shows the errors (8.10) as a function of grid density for three of the key variables: temperature, coke concentration and oxygen concentration. The estimated asymptotic convergence rates are indicated in the figure. To illustrate the grid sensitivity of the solution, the corresponding temperature, coke and oxygen profiles are plotted in Figure 8.2. Super-linear convergence is observed in all variables and in both norms. In order to reduce the error to, say, 5% we observe that approximately 160 gridblocks are required. At coarse grids the smoothing of temperature leads to excessive coke formation, which affects the combustion front propagation. Using 80 gridblocks, the location of the combustion front is underestimated by almost 20% compared to the converged result. With the ignition heater settings used in these simulations, a combustion front is established at all grid densities, but the front propagation speed varies dramatically. The oil recovery is affected both in terms of rate of recovery and ultimate recovery, because at coarse grids a larger fraction of the oil is burned as shown in Figure 8.3.

The convergence study was carried out using the K-value option for phase behavior. We have carried out an additional study using the EoS option, which indicated similar observations and convergence rates. In general, the number of gridblocks required to reduce the error below a certain threshold accuracy will be process and oil dependent.

8.2.2 Comparison with STARS

To further validate the VCT we compare our simulation results with results obtained using STARS. As above, we simulate dry air injection into an oil saturated core. Comparisons are shown in Figure 8.4 using 200 gridblocks in both simulators. Care was taken to make sure that all input data, treatment of boundary conditions, and numerical settings were the same in both simulations. However, with limited implementational details available for STARS, some differences may still exist. In general, good agreement is observed, and together with the convergence study we take this as sufficient validation of the VCT.

8.2.3 Ignition/Extinction Dynamics

For optimum performance we are interested in operating ISC processes at conditions corresponding to combustion in a high temperature regime, where full oxidation reactions take place. For a simplified 1D problem [Akkutlu2003] has shown the existence of multiple (pseudo) steady states in terms of combustion front temperature as a function of the air injection rate and heat loss rate. When operating near ignition/extinction branches it is likely that the system behavior will be very sensitive to the phase behavior model. In this section we carry out simulation studies with both the VKC and the VCT to explore the impact of phase behavior on ignition/extinction dynamics.

Kinetic cell experiments are routinely carried out in the laboratory, typically as ramped temperature experiments, where an oil sample placed in the cell is heated, while air is circulated through the cell. Since the kinetic cell is a batch-type experiment, the steady state always corresponds to all the oil being consumed and only the injection gas left in the cell. The paths taken to reach the steady state can be very different. In the simulations discussed below we heat up the cell to a specified temperature, and then keep a constant cooling temperature at this level. The cell heat controller is modeled as

$$Q^{h,htr} = -u_a\left(T - T^{ext}\right) \tag{8.11}$$

where u_a is a heat transfer coefficient (controller gain) and T^{ext} is the external heating/cooling temperature (controller set point). We define ignition as the state, where sufficient heat is evolved from reactions to increase the cell temperature above a threshold temperature as illustrated in Figure 8.5. The ignition/extinction behavior of the system, as defined in this way, is then studied as a function of air injection rate and heat loss rate (u_a). The ignition regime is the area of the parameter space, where the maximum temperature during a simulation is above the threshold temperature. The boundary of the regime corresponds to a level curve for the maximum temperature (where the maximum temperature is equal to the threshold temperature), and we will refer to this curve as an ignition/extinction curve.

We solve the continuation problem in the two parameters (air injection rate and heat loss rate) and map out the ignition/extinction regimes for both the minimal and the SARA reaction models, and using both K-value based and EoS based phase behavior descriptions. In both phase behavior approaches we use the same property models (phase densities, enthalpies, etc.), and make sure that the numerical solution is converged, so that the differences observed reflect only differences in phase behavior treatment.

We consider first the minimal reaction model. Figure 8.6 shows the ignition/extinction regimes for the EoS and K-value approaches, respectively, using a heat controller temperature of $[473]K$ and a threshold temperature of $[500]K$. At low air flow rates combustion is not sustained, whereas at high air rates the oil components are stripped from the oil phase before they react. Intermediate flow rates lead to ignition when the heat loss rate is sufficiently low. Increasing the heat loss rate leads to extinction. At the lower part of the ignition regime the phase behavior approaches predict similar behavior. Figure 8.7 shows ignition/extinction regimes for the SARA reaction model. In this case, the heat controller temperature is $[573]K$ and the threshold temperature $[650]K$. For both reaction models we have carried out sensitivity studies to assess the dependence on the choice of threshold temperature. In general, only small variations are observed, and only on the lower part of the ignition regime, when the threshold temperature is varied $[\pm 50]K$. Hence, the conclusions drawn from Figure 8.6 and 8.7 are independent of threshold temperature.

The K-value approaches are based on an ideal assumption of composition independent K-values, which is generally valid away from critical regions. When increasing the air flow rate, the pressure in the cell increases significantly above the initial pressure. A probable explanation for the differences in ignition regimes observed at higher air rates is therefore the increased pressure which leads to less ideal conditions. The Wilson correlation is a generalization of Raoult's law. Since the correlation extends the pure component vapour pressure curves beyond their critical temperature, it tends to overestimate the size of the two-phase region. For both the minimal and the SARA reaction models we observe an extended region of ignition when using Wilson K-values. Since oil components are assumed only to react in the oil phase, underestimation of oil volatility will lead to ignition for a wider range of operating conditions. Thus, our results agree with the expected behavior of the Wilson K-values.

The range of air injection rates typically used in laboratory kinetic cell experiments is 1 to 30 standard cell volumes per minute. The range covered in Figures 8.6 and 8.7 is approximately 0.1 to 100 standard cell volumes per minute (the cell volume is $[1]L$), so the sensitivities observed occur in the practical range of operating conditions. Tabulation of K-values optimized to an EoS is the common approach when working with commercial simulators for ISC. Hence, in subsequent sections we will focus on the comparison of optimized K-values with the full EoS.

8.2.4 One-dimensional Combustion Tube Experiments

Kinetic cell experiments do not reflect the quantitative behavior of one-dimensional ISC tube tests. In order to properly translate the observed phase behavior sensitivity to actual ISC processes involving the coupling of flow, transport and reactions, we need to carry out ignition/extinction experiments using the combustion tube model.

Ultimately, we are interested in the effect of phase behavior treatment in ISC processes on displacement efficiency, and thereby on the total recovery. Ideally, by extending to multi-dimensional problems, we will be able to assess the effect of phase behavior sensitivity on actual field ISC processes. We prefer, however, a step-wise approach, where the model is gradually extended, which allows us to carefully study the behavior at each stage and isolate the important parts.

We consider a dry air VCT. Air is injected into an oil saturated core, which is preheated to $[311]K$. The injection end is heated for the first half hour of the simulation by a constant rate energy input. We study the ignition/extinction behavior as a function of air injection rate when varying the phase behavior model. We simulate the system for $[10]hrs$ and measure the front temperature as the maximum temperature in the system at the end of the simulation. For an ignition heater rate of $[70]kJhr$, Figure 8.8 shows the front temperature as a function of air injection rate. Comparison is made between the EoS and the optimized K-value based approaches. To ensure that the observed differences are not due to numerical errors, the results are presented at two different grid resolutions. For both phase behavior models, the critical air rate, which triggers ignition, does not change when the grid is refined. We take this an indication that the solution has converged. At low air rates, a stable combustion front cannot be sustained, and the system remains in a non-ignited state after the heater is shut off. At a certain point, when the air rate is increased, the system jumps onto a high-temperature ignition branch. This behavior is common to many reaction engineering problems, where two (pseudo) stationary, stable branches exist, corresponding to ignition and extinction. For a certain interval of air injection rate, an intermediate, unstable branch exists, which is, however, not readily captured in our numerical experiment.

In the optimized K-value based approach, a higher air injection rate is required to reach ignition compared to the EoS based approach. Figure 8.8 also shows the iginition/extinction behavior when also varying the ignition heater rate. The observed behavior is in qualitative agreement with the results from the VKC, in that the EoS approach predicts the largest area of ignition in both cases. Consequently, the optimized K-values may overestimate the required air injection rate to sustain combustion. To illustrate the potentially dramatic difference in system behavior when operating near ignition/extinction branches, we simulate the VCT with an air injection rate of $[0.012]m^3hr$ and ignition heater rate of $[70]kJhr$ (see Figure 8.9). With the full EoS, a stable combustion front is established, whereas in the K-value approach the system never reaches ignition.

As a final investigation, we consider the SARA reaction model and study ignition/extinction behavior with the VCT. Figure 8.10 shows the front temperature as a function of air injection rate. Again, a higher rate is required to sustain combustion when using the K-value approach, but the difference is smaller compared to the minimal model, and combustion can generally be sustained at very low injection rates. The shape of the curves indicates that there are several

operation regimes, and the change to the high-temperature combustion regime is more gradual. The SARA model accounts for low-temperature oxidation reactions which dominate when the air supply is insufficient to reach high-temperature combustion.

8.3 Discussion

We demonstrated that numerical solutions for ISC processes are very sensitive to numerical and modeling errors. The sensitivity to numerical errors is directly evident from the convergence plots shown. At least 200 gridblocks are needed for the numerical errors in space and time not to strongly affect the solution quality. The errors observed are not only due to numerical discretization, but are also a result of upscaling of the kinetics. For the minimal reaction model, coarsening of the grid leads to excessive coke formation, which critically affects the front propagation speed and, hence, the oil recovery. The grid sensitivity study provides guidelines for the number of gridblocks required between wells in a practical reservoir setting in order to achieve a certain level of accuracy in the solution. In fact, for actual reservoirs, upscaling errors will be much larger, and the results here only indicate a lower bound on the number of gridblocks required. For the combustion tube problem we carried out the simulations with a minimum of 200 blocks to reduce numerical errors and isolate phase behavior model differences. In general, the number of gridblocks required to reduce the error below a threshold accuracy will be both process and oil dependent. Hence, it is essential to always conduct a grid convergence study.

The sensitivy to modeling errors is analyzed here for phase behavior models. For the VKC simulations we observe an extended area of ignition when using a K-value approach based on Wilson's correlation. This behavior is expected, since the correlation extends the pure component vapour pressure curves beyond their critical temperature, and, hence, tends to overestimate the size of the two-phase region. Since ISC kinetics is commonly modeled as gas-liquid or gas-solid (in case of reactions involving coke) reactions with rates typically depending on oxygen partial pressure and the concentration of an oil component in the oil phase, underestimation of oil volatility will lead to prediction of ignition for a wider range of operating conditions. Thus, our results for Wilson K-values agree with the expected behavior. The predicted ignition regimes from the EoS approach and the optimized K-values differ only slightly for the VKC, but when coupled with flow and transport in the one-dimensional VCT, the sensitivity seems stronger. For example, for a particular ignition heater setting, there is a 17% difference in the critical air flow rate required to sustain combustion between the optimized K-value approach and the EoS approach. The optimized K-value approach leads to a smaller ignition regime, which is in agreement with the observed behavior in the VKC. Thus, while

the VCT simulates actual ISC processes, the same qualitative conclusions concerning impact of phase behavior models can be obtained from VKC simulations. The ranges of air injection rates in both the VKC and VCT simulations are representative for practical ISC laboratory experiments.

It is not surprising that we find this sensitivity, given the assumption of phase specific reactions, which makes the kinetics sensitive to the distribution of components among phases. The API gravity of the two oils considered in this work ($[25]^{\circ} API$ and $[16]^{\circ} API$, respectively) is in the medium to heavy range. The heavy, SARA based oil shows less sensitivity to phase behavior when simulated in the VCT. It seems intuitive that lighter oils will be more sensitivity to phase behavior, in which case the assumption of only liquid based reactions for hydrocarbon components is questionable. Current industry simulators allow specification of reactions in all phases, but we are not aware of any published work on ISC kinetics for gas phase reactions.

The sensitivity to phase behavior generally means that it is advisable to conduct experiments using a VKC or VCT before any field experimentation to assess how strongly results may be affected. Based on these experiments a decision can be made concerning which phase behavior representation to use. In general, the experience gained from VKC and VCT simulations will reduce the work required and improve understanding of the results when extending simulation to the field.

In the Whitson characterization procedure used for the Peng-Robinson EoS, critical parameters are estimated from liquid density and normal boiling point data through a set of correlations. Several characterization procedures exist, and although not addressed in this work we expect to see sensitivity also to the choice of characterization procedure. Moreover, the pseudoization of the oils considered here was done for kinetics and not for phase behavior. The SARA fractions are kinetically different, but each fraction may span a range of components with very different phase behavior. For example, the saturates pseudo component covers basically all saturated hydrocarbons with straight or branched chains. A further subdivision of the fractions based on boiling point may improve the representation of phase behavior, but it is outside the scope of this paper to study effects of pseudoization. We note, however, that the development of systematic pseudoization techniques for reactive-compositional processes, which take into account both reactivity and phase behavior, is an important area for future research and one of the current bottlenecks in the development of reliable models for ISC.

A natural extension of the work presented in this paper is to consider multi-dimensional problems, which introduces additional parameters affecting the behavior of ISC processes, and which cannot be studied in a one-dimensional model. Of particular interest is the sensitivity to reservoir heterogeneity, the important question being how the formation and sustained propagation of an ISC combustion front is affected by heterogeneities in the reservoir rock. For multi-dimensional problems we need improved numerical methods, and we are currently working on a computational framework, which leverages on adaptive gridding and operator splitting techniques to address the multi-scale nature of ISC processes [Younis2006,Nilsson2005a,Kristensen2007].

8.4 Summary

The Virtual Combustion Tube (VCT) is a novel, equation of state compositional simulator for in-situ combustion (ISC). The simulator is validated through careful convergence studies as well as comparisons with the industry standard simulator STARS. The VCT is implemented with both equation of state based and simplified K-value based phase behavior descriptions, which enables studies of differences to improve understanding of phase behavior for ISC processes.
The sensitivity to grid step size shows that careful grid convergence studies are needed in each reservoir study.

The choice of phase behavior model impacts ISC process performance. Different models lead to different regimes of ignition. Driven by chemical reactions, ISC processes rely critically on reaching proper ignition conditions. Although difficult to rigorously quantify, there are areas of operating conditions (such as air injection rate), where a change from a simple K-value based phase behavior description to an equation of state based description will change the system dynamics, and therefore performance predictions, dramatically (from ignition to extinction). The oils tested here are in the medium to heavy range. More volatile oils are likely to be more sensitive to phase behavior. Strict guidelines cannot be given at this stage as to which phase behavior model to use, but the decision must be based on a sensitivity study for the particular oil.

The VCT and VKC provide a reliable and flexible platform for the study and support of laboratory scale combustion experiments. Although limited in physics, the VKC sensitivity results agree qualitatively with the VCT, and the VKC can thus be used as a fast screening tool when assesssing different kinetics and phase behavior models. Simulations using the VCT and VKC are the natural first steps before going to full-scale models.

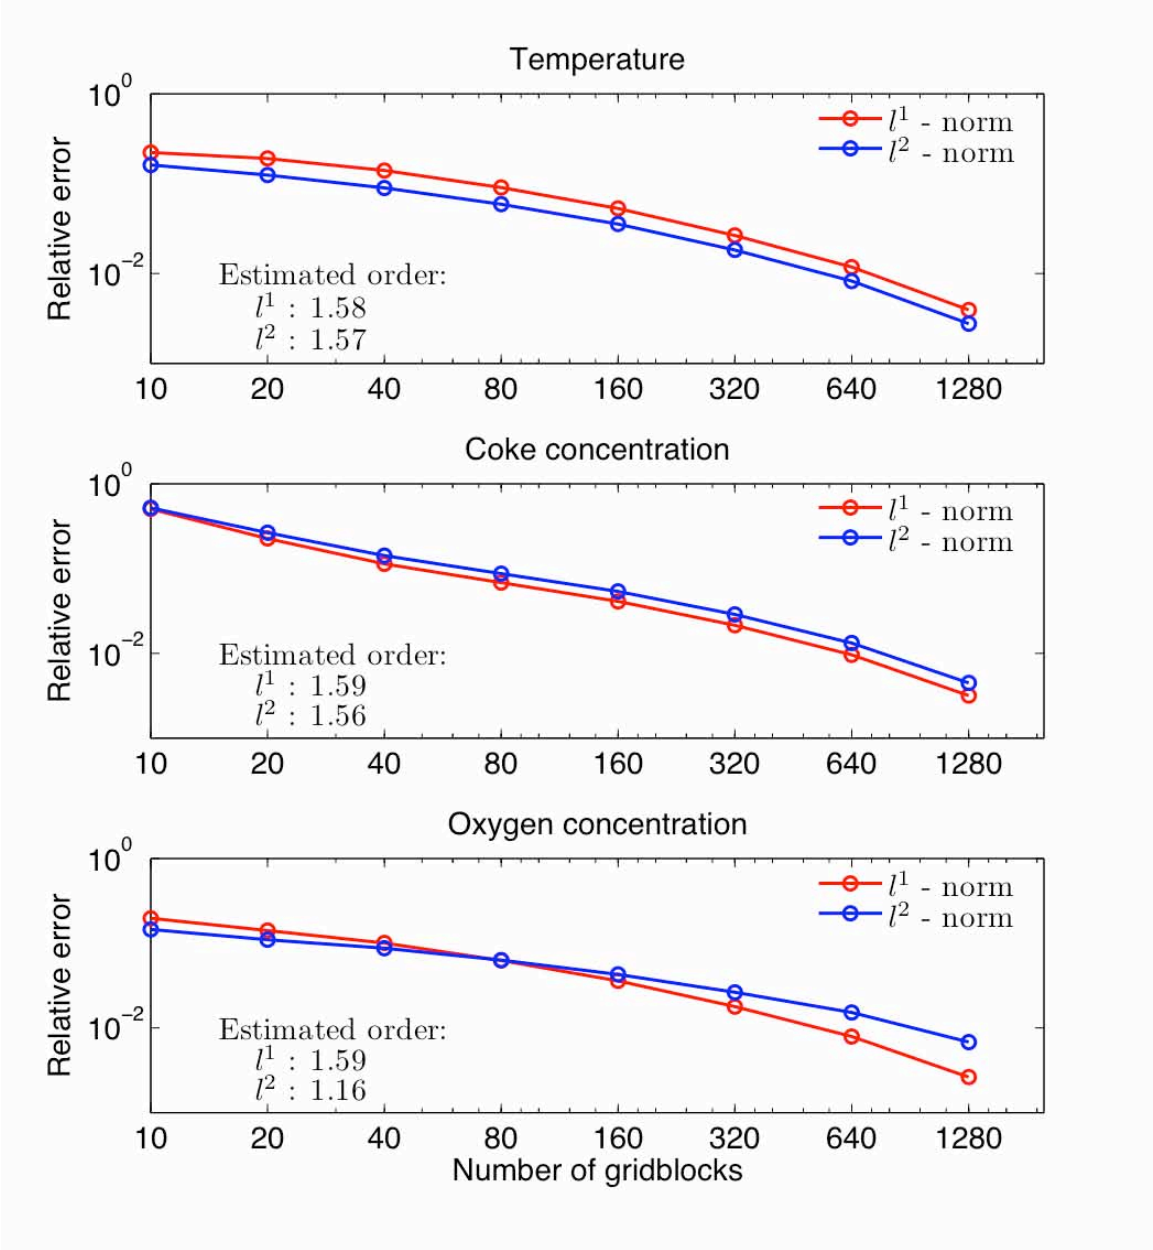

Figure 8.1: Estimated relative error in temperature, coke concentration, and oxygen concentration as a function of the number of gridblocks.

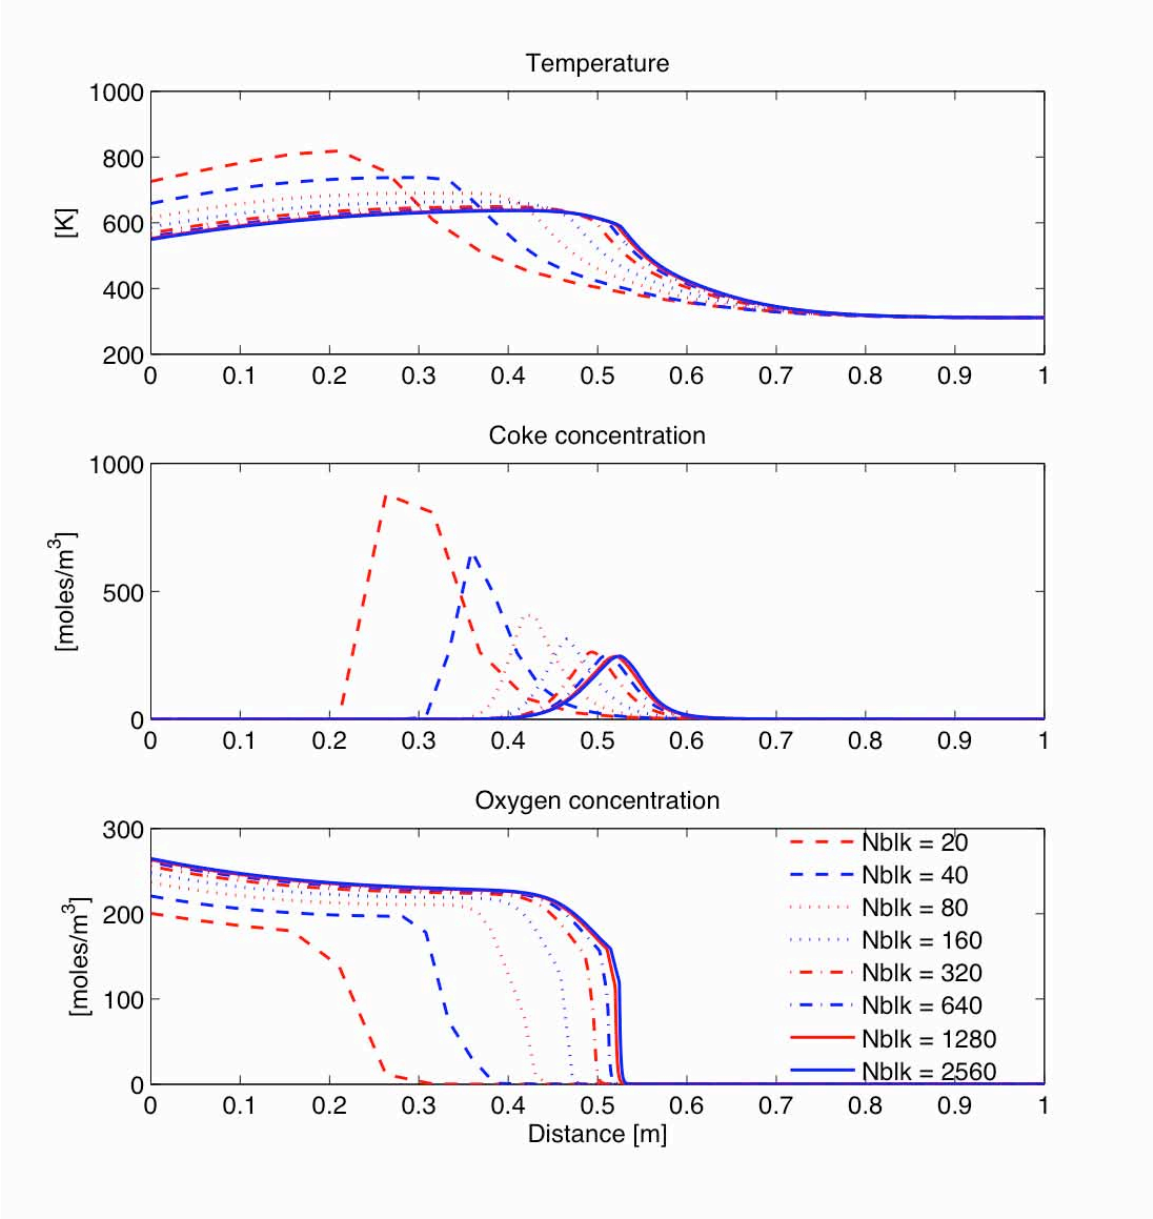

Figure 8.2: Temperature, coke concentration, and oxygen concentration profiles at different grid densities for the combustion tube simulation.

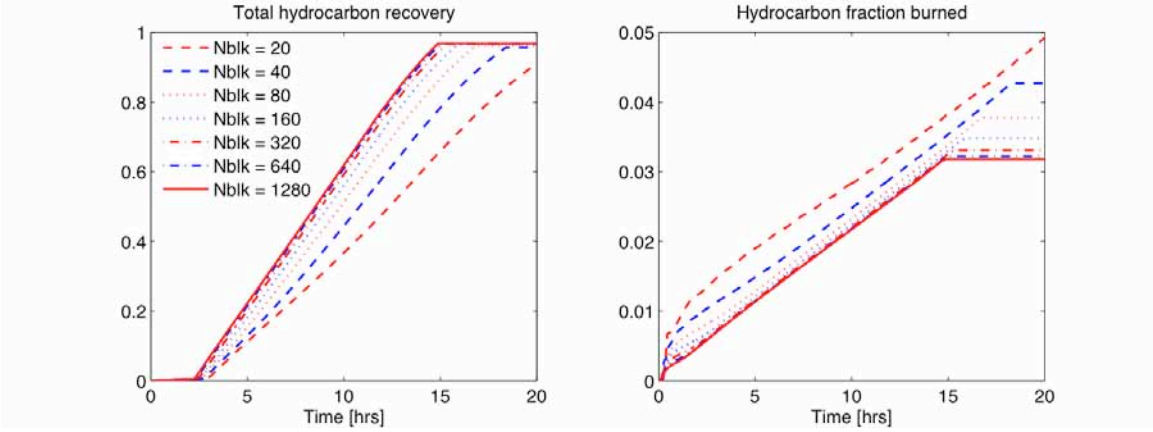

Figure 8.3: Cumulative oil recovery and fraction of oil burned at different grid densities for the combustion tube simulation.

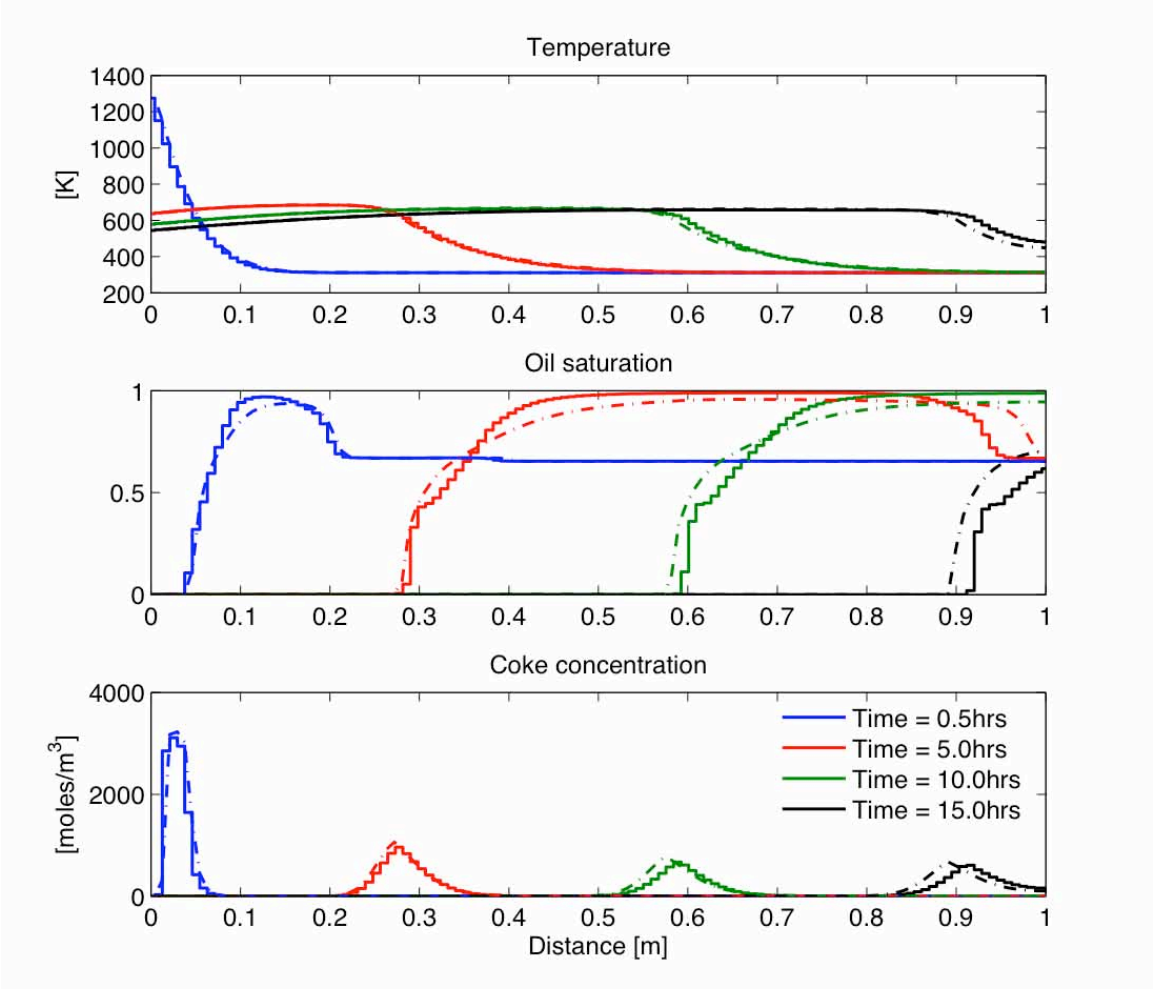

Figure 8.4: Simulation comparison between the VCT and the commercial simulator STARS. Temperature, oil saturation, and coke concentration profiles are shown. Solid lines : STARS, dashed lines : VCT.

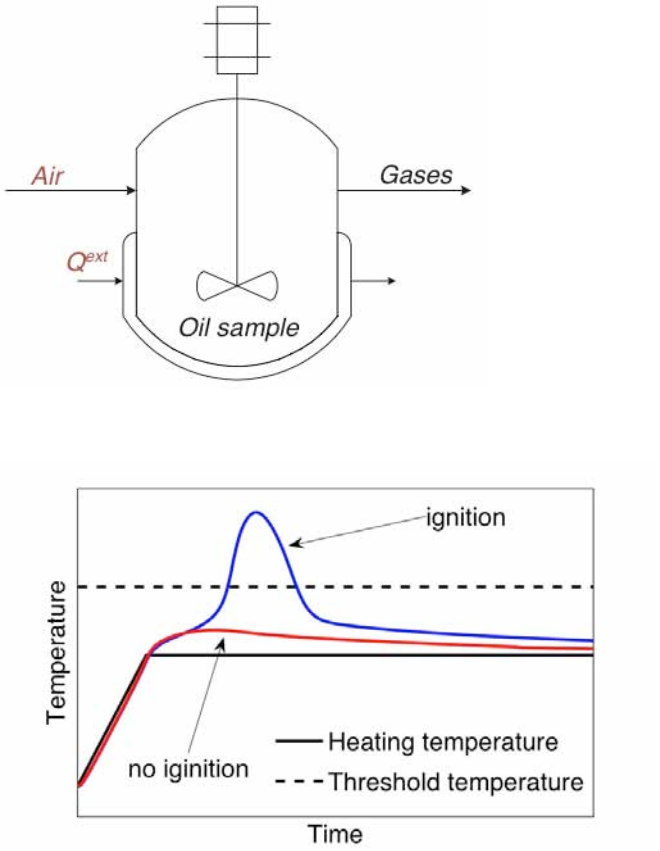

Figure 8.5: The VKC is modeled as a batch-type reactor, but with continuous flow of air. Ignition is defined as the state, where the onset of combustion in a ramped temperature experiment releases enough heat to increase the cell temperature above a certain threshold temperature.

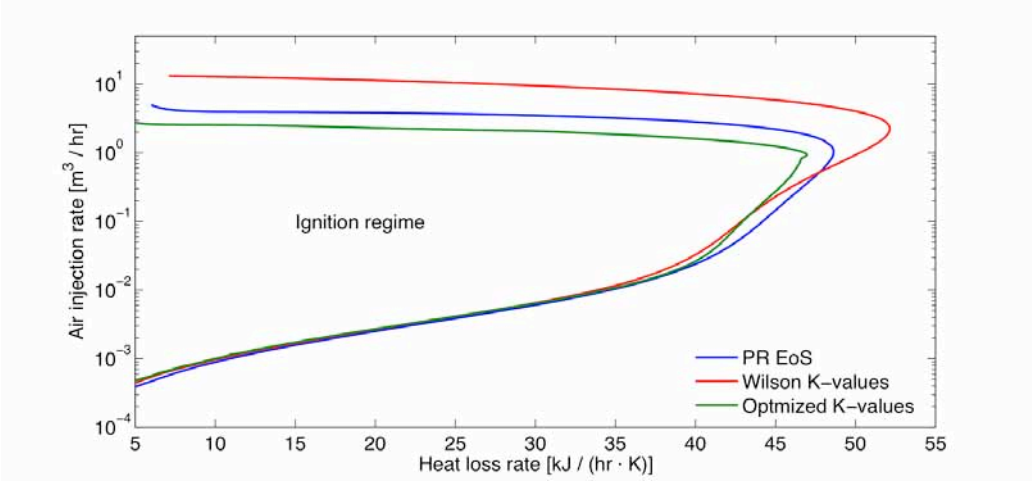

Figure 8.6: Ignition/extinction regimes for the minimal reaction model using the VKC. Comparison is made between EoS based and K-value based phase behavior representations.

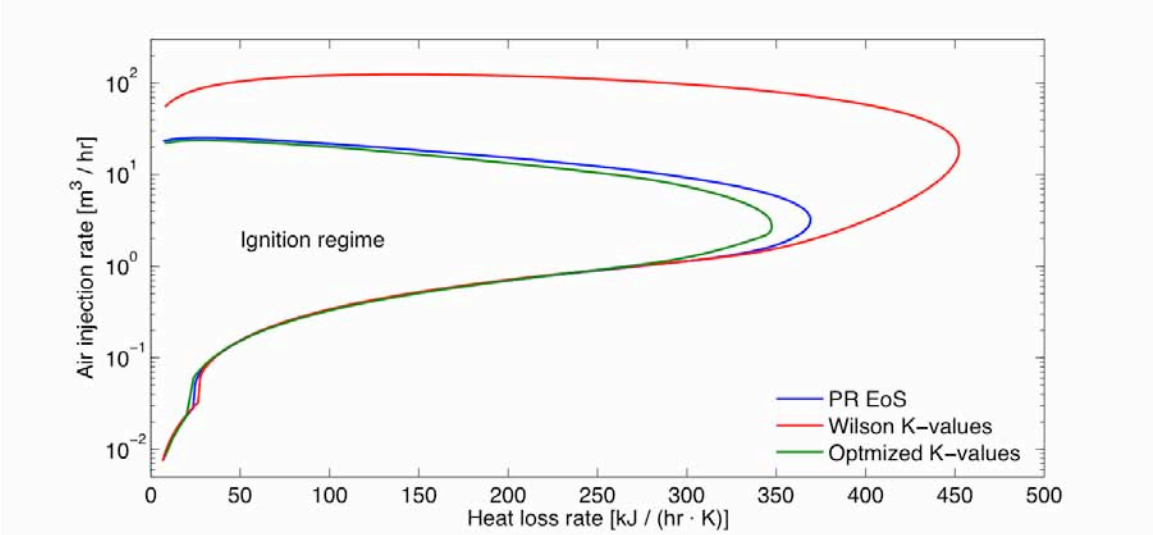

Figure 8.7: Ignition/extinction regimes for the SARA reaction model using the VKC. Comparison is made between EoS based and K-value based phase behavior representations.

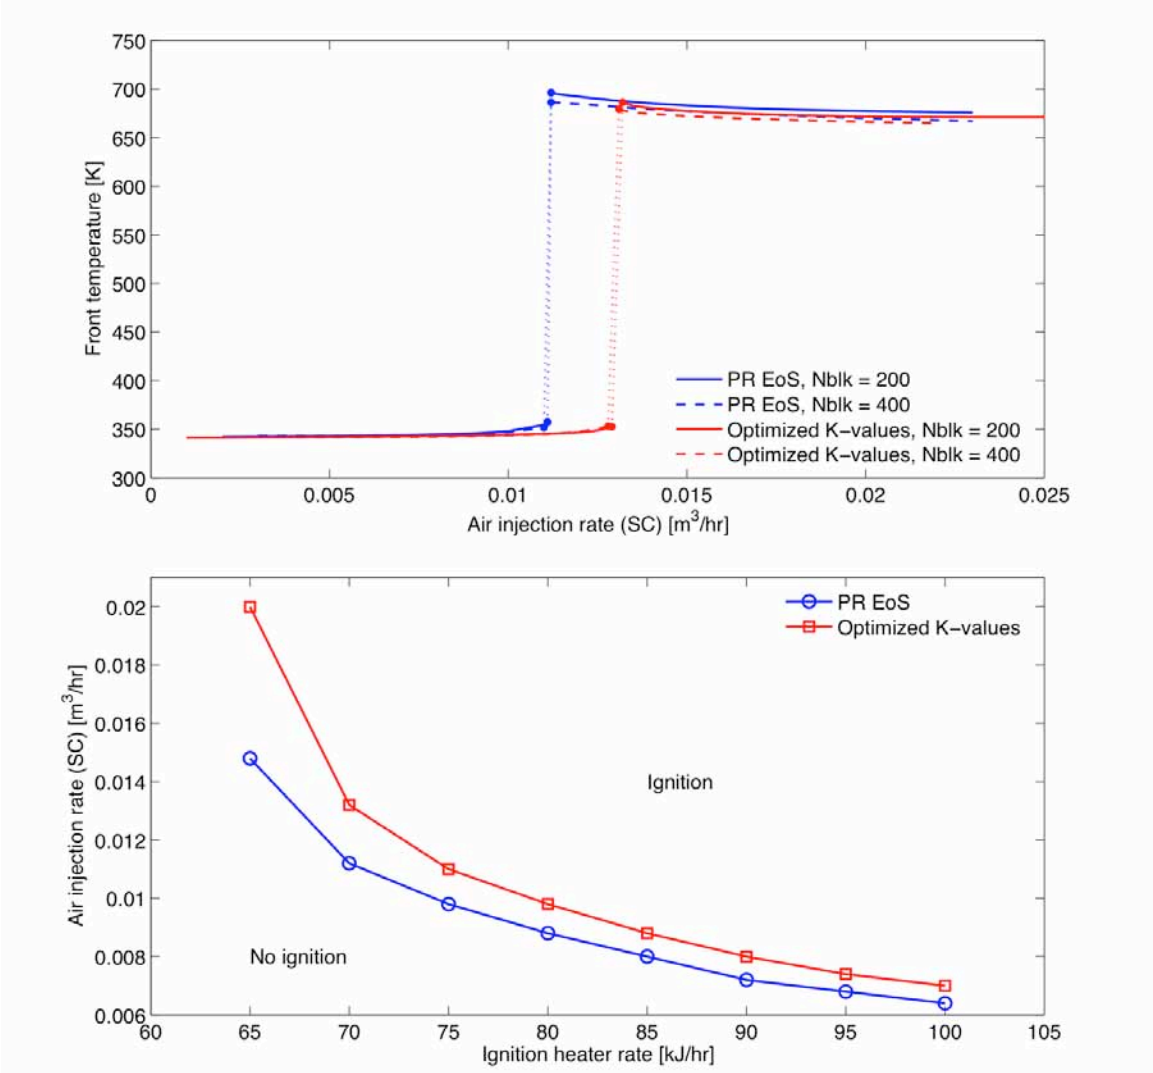

Figure 8.8: Ignition/extinction behavior for the 1D VCT using the minimal reaction model. Above: combustion front temperature as a function of air injection rate for a specific ignition heater rate. Below: Ignition/extinction curves for varying air rates and ignition heater rates.

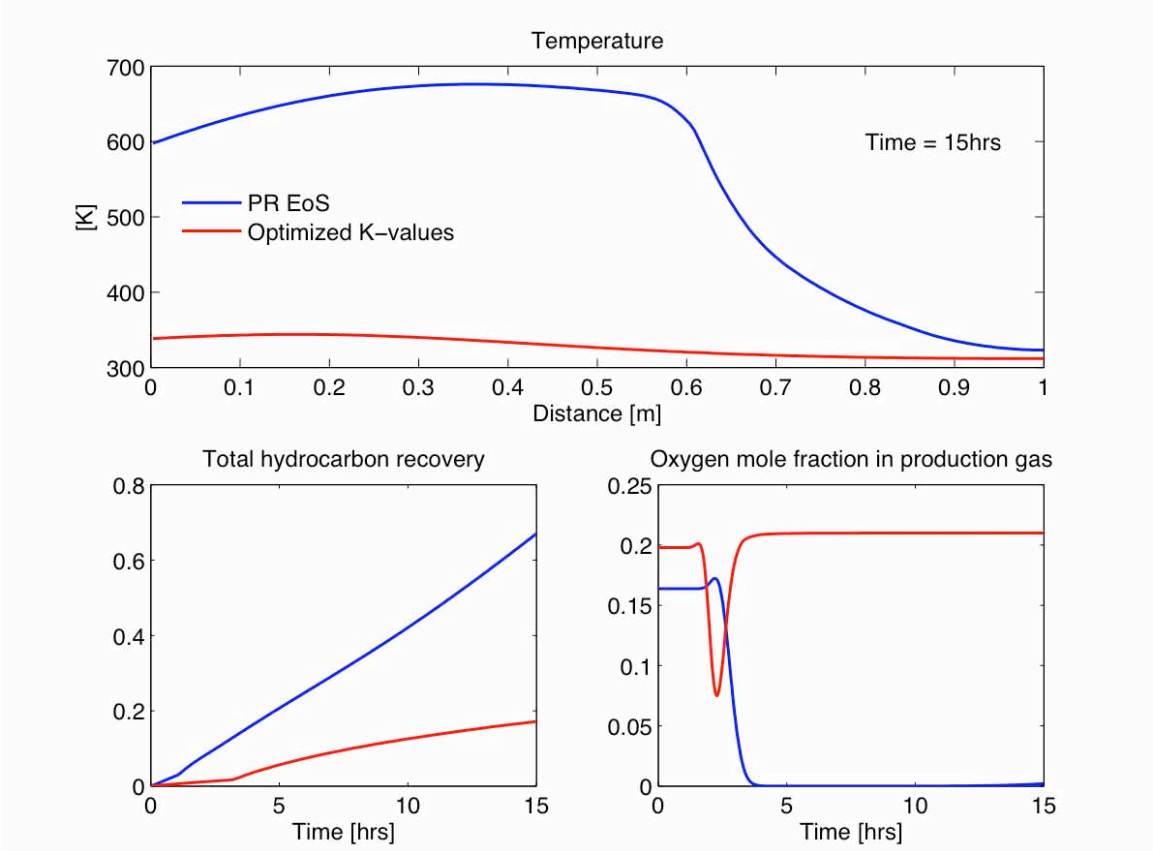

Figure 8.9: VCT simulation comparison for the minimal reaction model when operating near and ignition/extinction branch. Temperature profiles are shown along with total cumulative hydrocarbon recovery and oxygen mole fraction in the production gas.

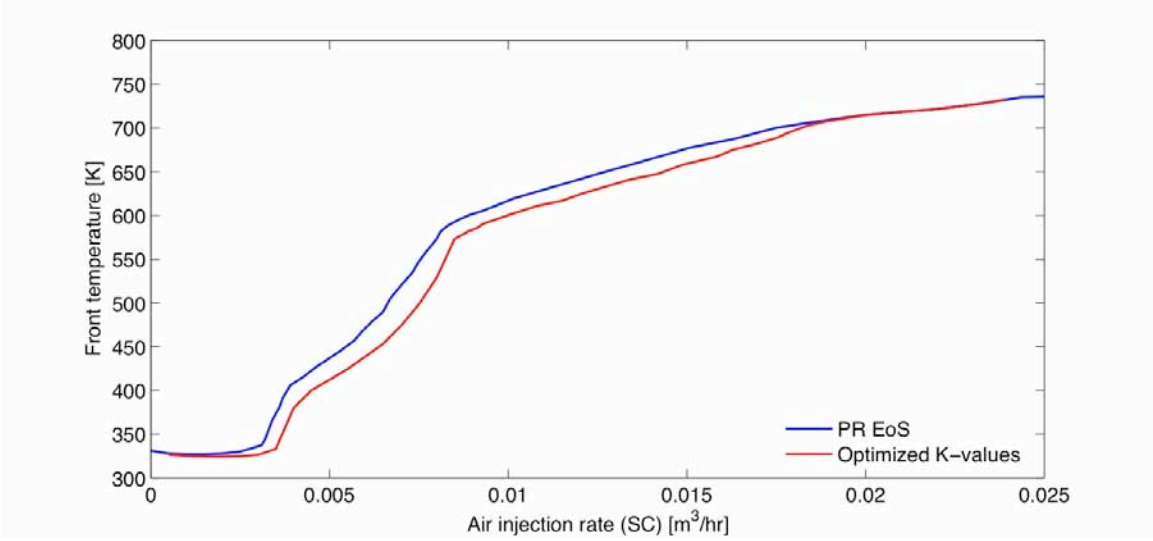

Figure 8.10: Ignition/extinction behavior for the 1D VCT using the SARA based reaction model. The combustion front temperature is shown as a function of air injection rate.

9. The ISC Code

In this chapter we discuss briefly important implementational details for both the fully coupled and operator splitting schemes. Implementation of the ESDIRK schemes is discussed in Chapter 7. The simulator was developed using the Fortran 90/95 language and the Compaq Visual Fortran compiler. In the implementation we have emphasized a clear and logical code structure with flexibility to add new features, and to a lesser extent computational aspects such as optimal memory usage and time efficiency.

9.1. Code Structure

The sequence of tasks carried out by the simulator can be summarized in the following pseudo-code:

```
Read user input
Perform preprocessing operations
Initialize gridblocks in equilibrium
While time less than final time do
      Compute global timestep
      if FIM scheme then
          Advance one FIM step
          While residual larger than tolerance do
              Flash all gridblocks
              Compute properties
              Evaluate residual and Jacobian
              Solve linear system and perform update
          end while
      else
          Advance one SPLIT step
          Perform half step on flow/transport
          Perform full step on reactions
          Perform half step on flow/transport
      end if
      if step converged then
           Update time and solution
      else
          Reset solution and reduce time step
      end if
end while
```

The input data include

- Static grid properties: block size, void porosity, and permeability.

- Component and phase properties: critical parameters, binary interaction coefficients, liquid reference densities, heat capacity correlation constants, viscosity correlation constants, and thermal conductivities.
- Reaction kinetic data: reaction stoichiometry, activation energies, pre-exponential factors, and reaction enthalpies.
- Initial conditions: temperature, pressure, and overall composition.
- Well data: injection rates, temperature, and composition. Production well bottom-hole pressure.
- Heat controller specifications: heater size, set point, gain, and timing (on/off switches).
- Numerical specifications: time integration scheme (fully coupled or splitting) and convergence tolerances.

The preprocessing operations include calculation of grid transmissibilities, well indices and initialization of the gridblocks from initial conditions. The gridblocks are initialized in thermodynamic equilibrium. Initialization for simultaneous gravity and thermodynamic equilibrium has not been implemented.

Timesteps in the fully coupled scheme and global timesteps in the splitting scheme are calculated from the timestep formula [Aziz1979]

$$\Delta t^{n+1} = \Delta t^n \min_{i,k} \left[\frac{(1+\lambda)\Delta u_{ik}^n}{\Delta u_i^* + \lambda \Delta u_{ik}^n} \right], \tag{9.1}$$

where Δu_{ik}^n is the change in variable i in gridblock k over the previous time step, Δu_i^* is the target change, and λ is a tuning factor.

The termination criterion for the Newton-Raphson iterations is based on the max-norm of the normalized residual

$$e_{res}^{(l)} = \mathsf{P}\tilde{\mathbf{R}}^{(l)}\mathsf{P}_\infty, \tag{9.2}$$

where the components of the normalized residual are

$$\begin{aligned} \tilde{R}_{ik}^{(l)} &= \left(\frac{R_{ik}}{C_{ik}} \right)^{(l)}, \qquad i = 1, \ldots, n_c, \\ \tilde{R}_{ik}^{(l)} &= \left(\frac{R_{ik}}{U_k} \right)^{(l)}, \qquad i = n_c + 1, \\ \tilde{R}_{ik}^{(l)} &= \left(\frac{R_{ik}}{\phi_v V_b} \right)^{(l)}, \qquad i = n_c + 2, \end{aligned} \tag{9.3}$$

i.e. on a gridblock basis the component balances are normalized by the total component amounts, the energy balance is normalized by the total internal energy, and the residual of the volume constraint is normalized by the pore volume.

Although only 1D problems are considered here, the simulator is designed in such a way that a different gridding module can easily be included. In the computation of interblock flux terms, the simulator kernel only sees a list of gridblock connections

with associated transmissibilities, and no assumptions are made regarding grid geometry at this level.

The derivatives required in the Jacobian of the fully coupled, implicit system, as well as the Jacobian of the implicit flow/transport substep in the splitting scheme, are approximated by numerical differences. This approach is flexible when adding new features or changing property models, but computationally more expensive than analytical derivatives. The perturbations in the numerical derivatives are chosen according to

$$\delta u_{ik} = \sqrt{\varepsilon_m \max\left(10^{-5}, |u_{ik}|\right)} \tag{9.4}$$

where ε_m denotes the machine precision. The grid structure is exploited in the computation of Jacobians, so that several variables are perturbed simultaneously. The overall computational cost scales linearly with the number of gridblocks. We note that STARS also uses a numerical approximation to the Jacobian [STARSmanual].

Exact flashes are performed in each Newton-Raphson iteration using information from the previous iteration as starting guess. Stability analysis is performed in single phase gridblocks to determine if the current state is stable.

Two linear solvers are included for solution of the linear equation systems arising in each Newton-Raphson iteration:

- MA48: a direct solver based on sparse Gaussian elimination [Duff1996].
- GMRES: an ILU preconditioned GMRES solver from the SPARSKIT library [Saad1994].

9.2 Results: Simulation Examples

To illustrate the behavior of ISC in combustion tube experiments, and as a supplement to the results presented in the research papers, we consider in this chapter a number of simulation examples. We study the influence of boundary and initial conditions and conclude with an example illustrating the splitting scheme proposed in the previous chapter.

9.2.1 Influence of boundary and initial conditions

As a first example we consider simulation of a dry combustion tube experiment. The reaction and component model is introduced in Chapter 7. The model has six components: two oil pseudo components, oxygen, solid coke, water, and inert gas. The reactions include direct oxidation of the oil components, cracking of the oil into solid coke, and oxidation of the coke component. The oil has an API gravity of 25.

Air is injected at a rate of $[16]Lhr$ (surface conditions) into an oil saturated core, which is preheated to $[311]K$ and pressurized to $[137]atm$. The injection end is

heated for the first half hour of the simulation by a constant rate energy input. For the remainder of the simulation the combustion tube is operated adiabatically. Figure 9.1 shows the temperature, oil saturation, water saturation, and solid coke concentration profiles at four different times during the simulation. As air is injected and the tube heated, a combustion front is established, which propagates slowly into the core and displaces the oil. No initial water is present, but a small amount of water is formed in reactions, which moves ahead of the combustion front.

In the next example we consider a similar setup as above, but we study the influence of the initial water saturation. Figure 9.2 shows temperature, saturation, and coke concentration profiles at $t = [6]hrs$ for four different initial water saturations. Figure 9.3 shows the corresponding production curves in terms of the total cumulative hydrocarbon recovery and the fraction of hydrocarbons burned in reactions. As observed from Figure 9.2, the initial water saturation does not influence the amount of coke lay-down and thereby the front propagation speed. In all cases a combustion front is established. An oil bank is formed downstream of the combustion front, and further downstream the water is displaced by the oil bank. Since the same amount of coke is formed, the relative amount of hydrocarbons burned in reactions increases with decreasing initial oil saturation, and the ultimate recovery is therefore lower with low initial oil saturation.

With no initial water present we consider in Figure 9.4 and 9.5 a variation in the oxygen feed concentration. The injection is fixed at $[16]Lhr$, and the oxygen mole fraction in the feed is varied from 0.1 to 0.7. The feed concentration has direct influence on the front propagation speed as seen from Figure 9.4. An oxygen mole fraction of 0.1 is too low to sustain combustion. The initial amount of coke formed, when heating the injection end of the combustion tube, is never consumed. For higher feed concentrations a combustion front is established. This example clearly illustrates the inefficiency of the process when proper ignition conditions are not reached.

A large part of the energy generated in dry combustion processes remains behind the combustion front as heat absorbed in the reservoir rock. Because of the low heat capacity of air, the heat in the burned zone is not transported efficiently by the injected air, as observed in e.g. Figure 9.4. To overcome this inefficiency water is sometimes injected simultaneously with air, the risk being that the water will quench the combustion and lead to extinction. Figure 9.6 and 9.7 show results from wet combustion simulations with varying amounts of water. By injecting water, the energy behind the combustion front is swept forward more efficiently than for dry air injection, and the combustion front is propagated faster. However, at a water rate of $[0.05]pvihr$ the process is effectively quenched. We note also that coke is left behind the combustion front due to the rapid decrease in formation temperature caused by water injection.

9.2.2 A heavy oil example

The second reaction and component model considered in this thesis is based on a characterization of the oil in terms of its SARA (Saturates, Aromatics, Resins and Asphaltenes) fractions. The model has 14 components and 14 reactions. The oil has an API gravity of 16. Complete input data for the simulation is given in [Kristensen2008].

We consider a simple dry combustion simulation. Air is injected at a rate of $[45]Lhr$ (surface conditions) into an oil/water saturated core. The initial core temperature and pressure are $[323]K$ and $[14.5]atm$, respectively, and the initial water saturation is 0.25. Again, to initiate combustion, the injection end of the combustion tube is heated externally for the first half hour of the simulation. For the remainder of the simulation a proportional heat loss along the core is accounted for according to Eqn. (5.24) with a heat controller gain of $[4]kJm^2 \cdot K \cdot hr$ and controller set point of $[323]K$.

Figure 9.8 shows the temperature, oil saturation, water saturation, and pressure profiles at four different times during the simulation. As the combustion front propagates into the core, the oil is banked ahead of the front creating an almost fully oil saturated zone. The temperature in this zone is still at the initial level, and the mobility is therefore very low, which causes the pressure to increase behind the front. This phenomenon has been observed in practical laboratory combustion tube tests for heavy oils [Moore1995]. In laboratory tests, as well as in the field, equipment capabilities will limit the injection pressure, and formation of a low-mobility oil bank, as observed in Figure 9.8, will effectively plug the system and stall the combustion front. The break-down in oxygen supply to the front shifts the reactions to the low-temperature regime with partial oxidation reactions, which is generally an unfavorable operating regime. Plugging of the system is a common cause of failure in laboratory combustion tests [Moore1995].

9.3 Summary

Our ISC simulator in 1D contains operator splitting, specialized solvers for kinetics and a full equation of state. It has been validated by comparison to STARS. The tests described in this chapter show its accuracy. Additional results can be found in Chapter 8, which is based on the simulator as described in this section.

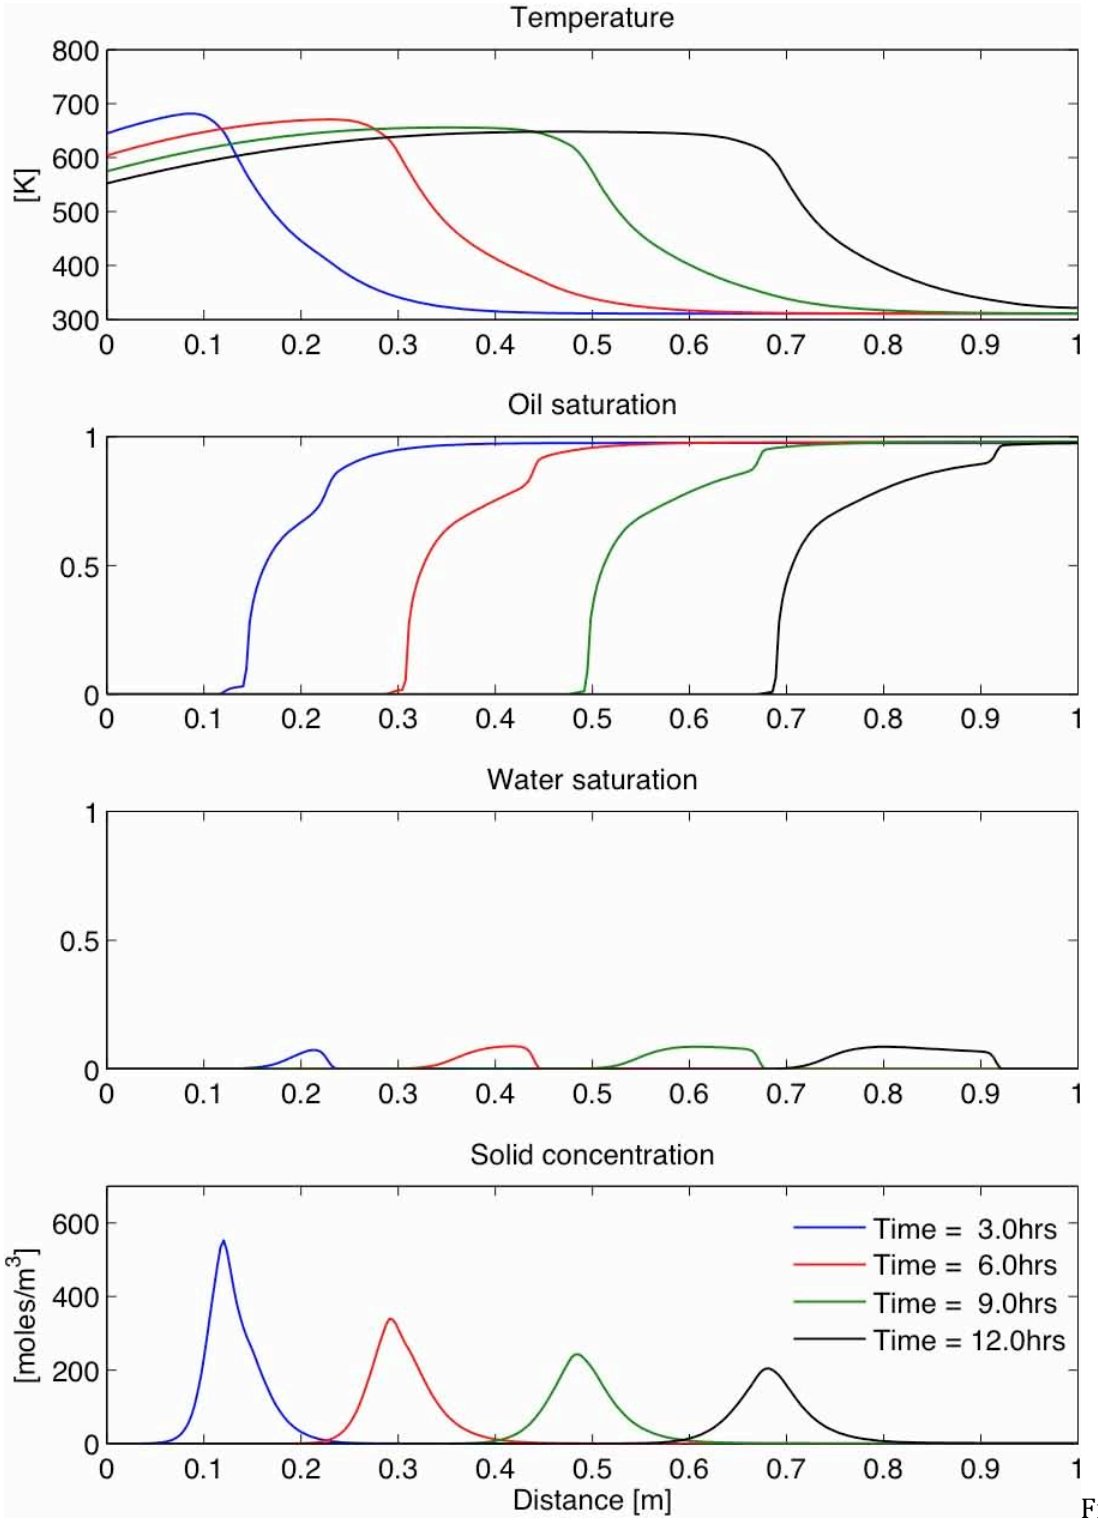

Figure 9.1: Dry combustion simulation results for the minimal reaction model. Temperature, saturation, and solid coke concentration profiles are shown at four different times during the simulation.

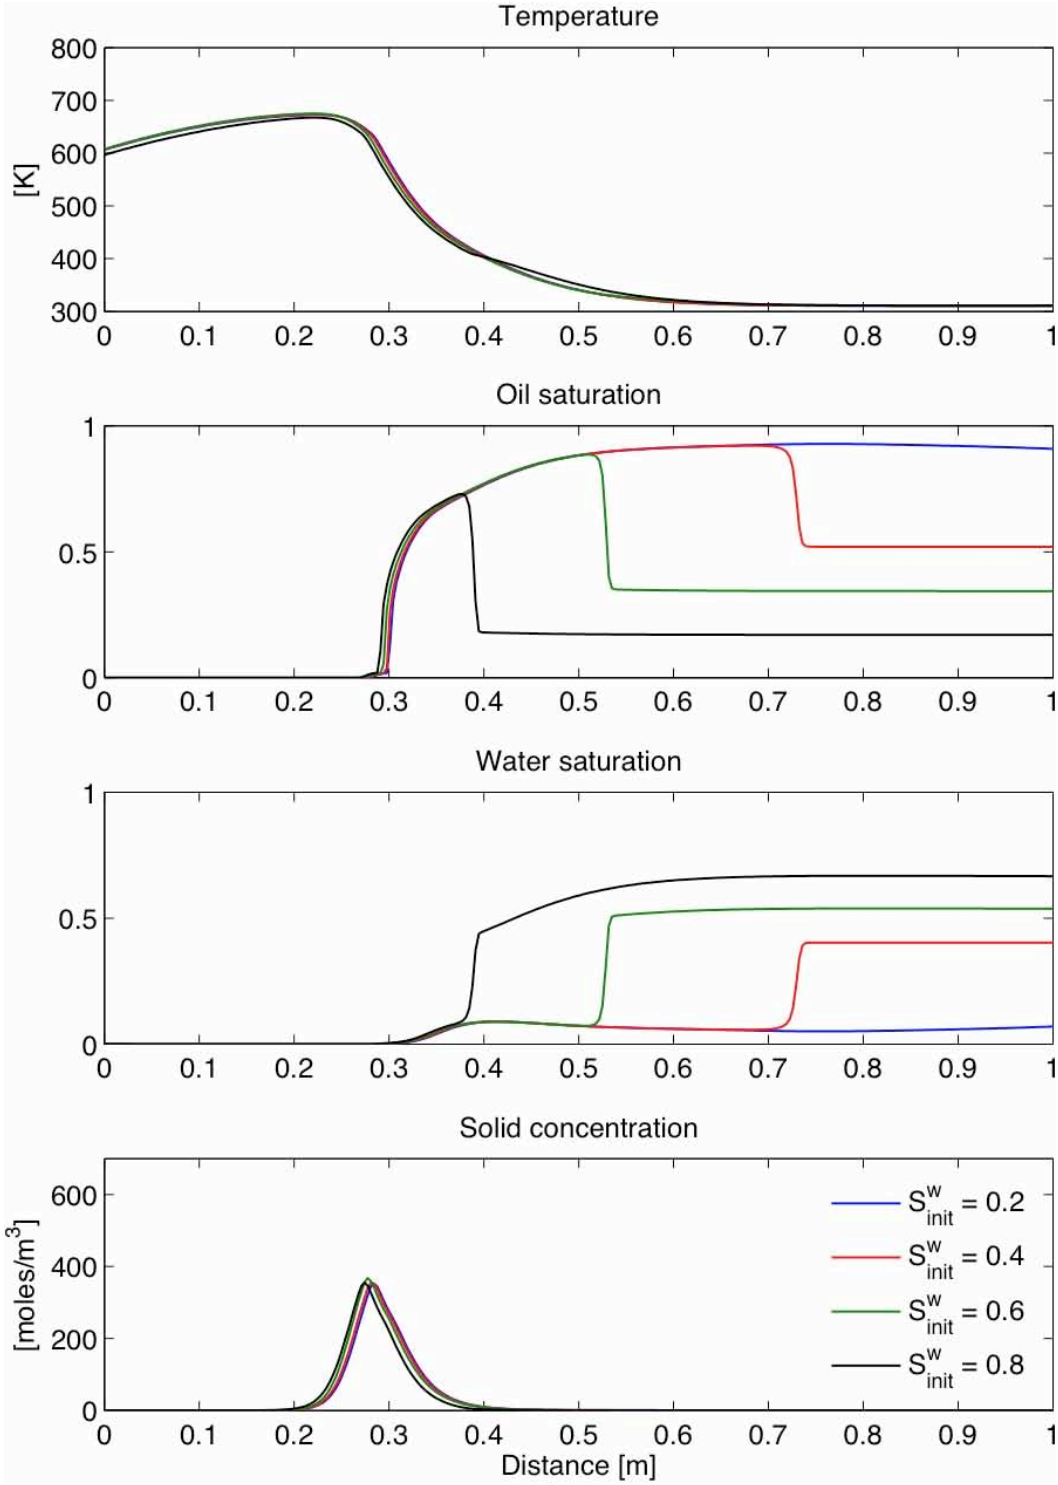

Figure 9.2: Dry combustion simulation results with varying initial water saturation. Temperature, saturation, and solid coke concentration profiles are shown at $t = [6]hrs$ and four different initial water saturations.

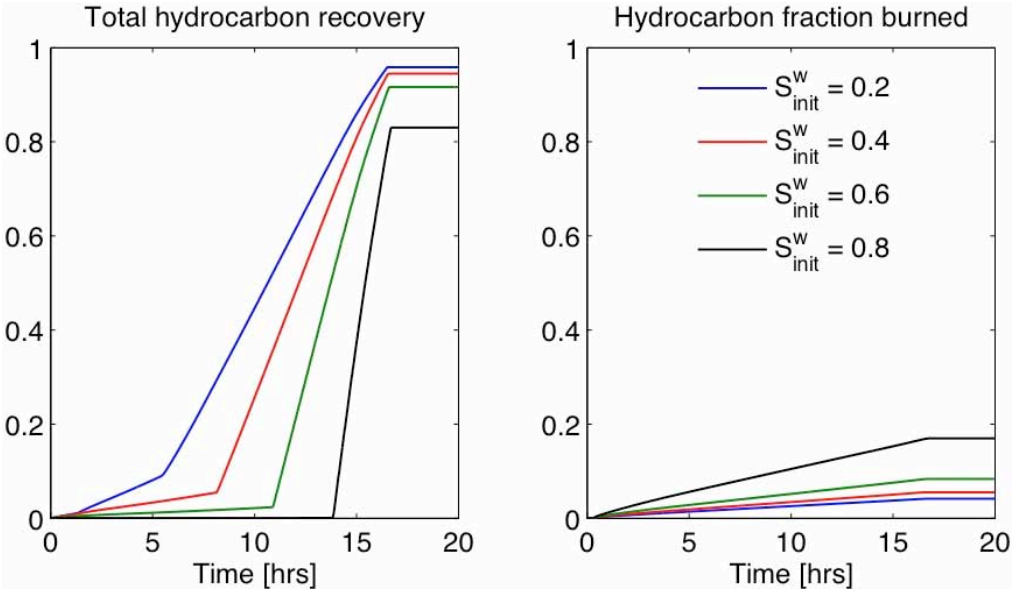

Figure 9.3: Production curves for dry combustion simulation with varying initial water saturation. The total cumulative hydrocarbon recovery is shown along with the fraction of hydrocarbons consumed in reactions.

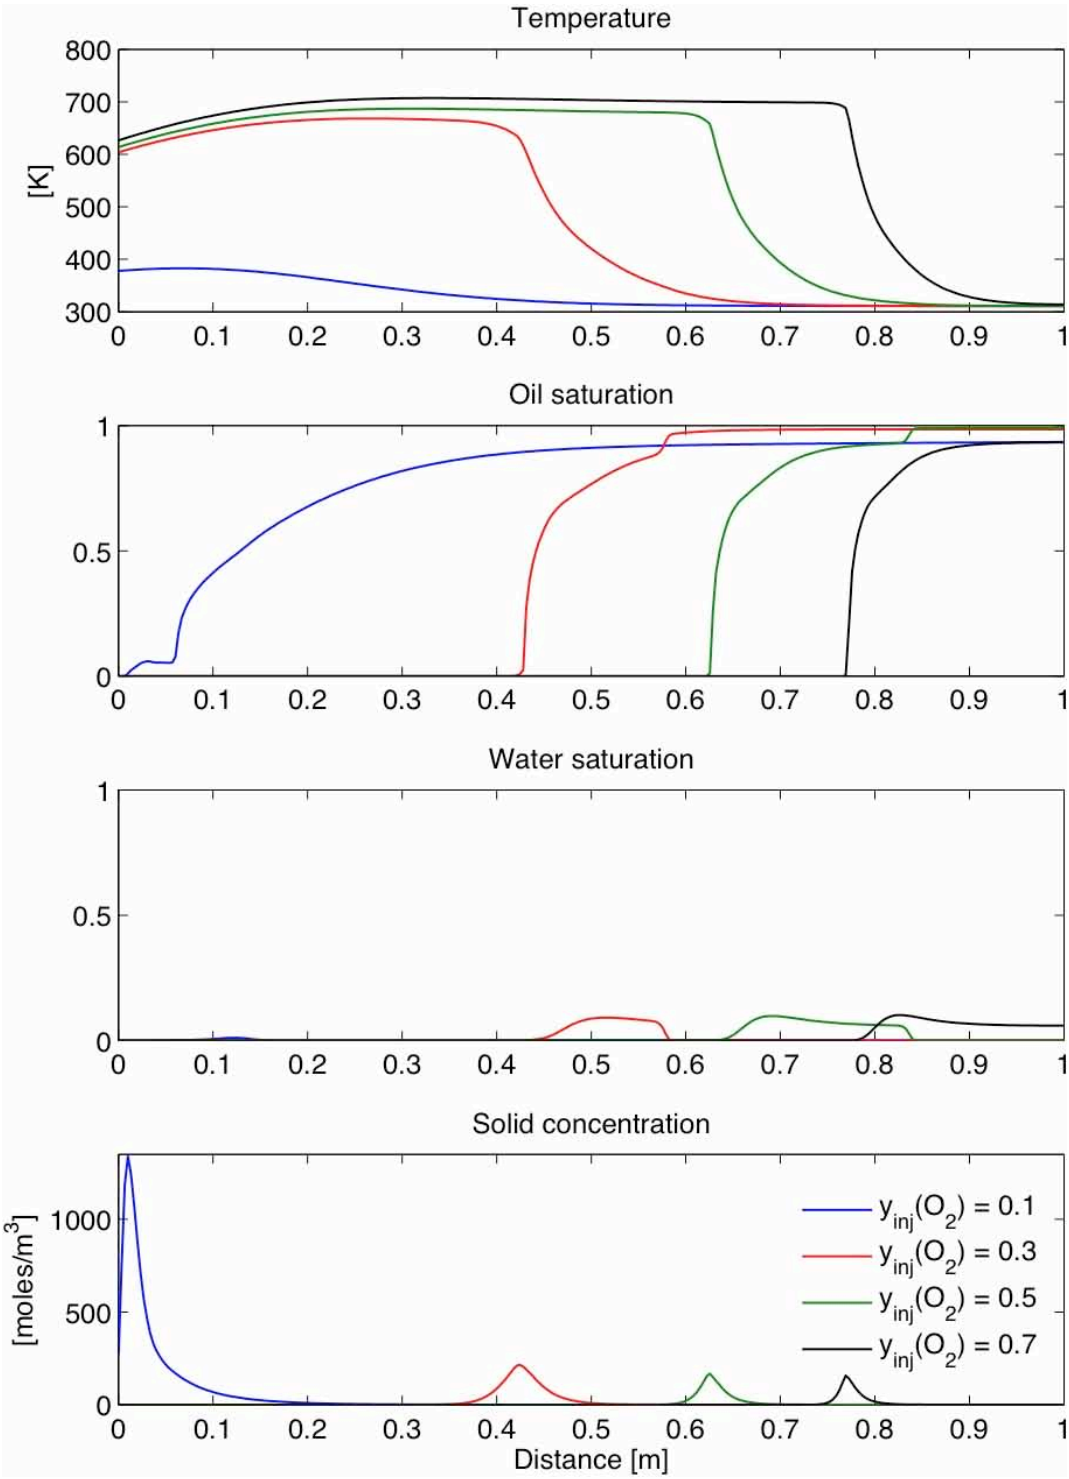

Figure 9.4: Dry combustion simulation results with varying oxygen feed concentration. Temperature, saturation, and solid coke concentration profiles are shown at $t = [6]hrs$ and four different oxygen feed concentrations.

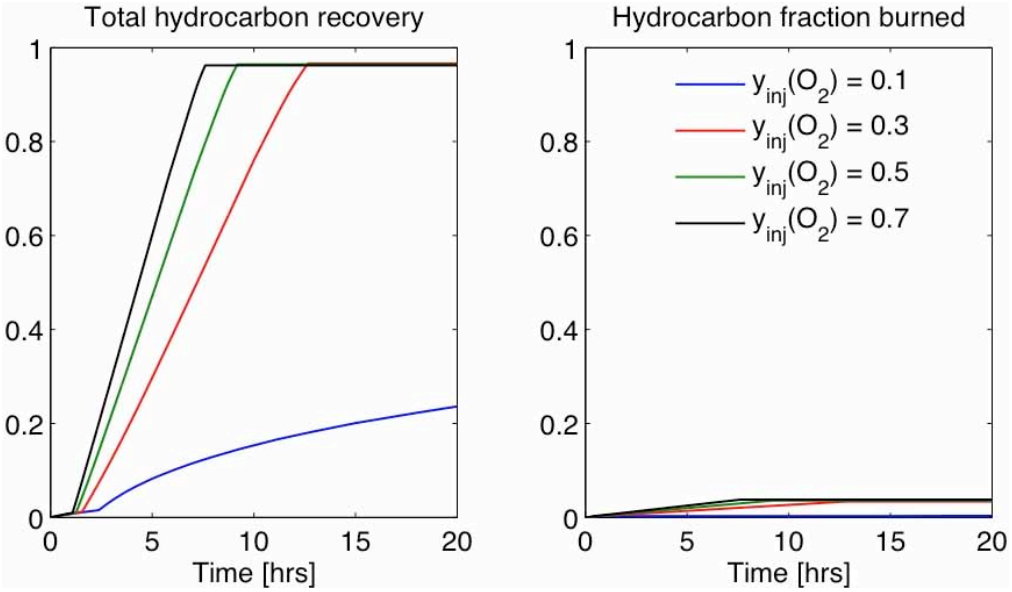

Figure 9.5: Production curves for dry combustion simulation with varying oxygen feed concentration. The total cumulative hydrocarbon recovery is shown along with the fraction of hydrocarbons consumed in reactions.

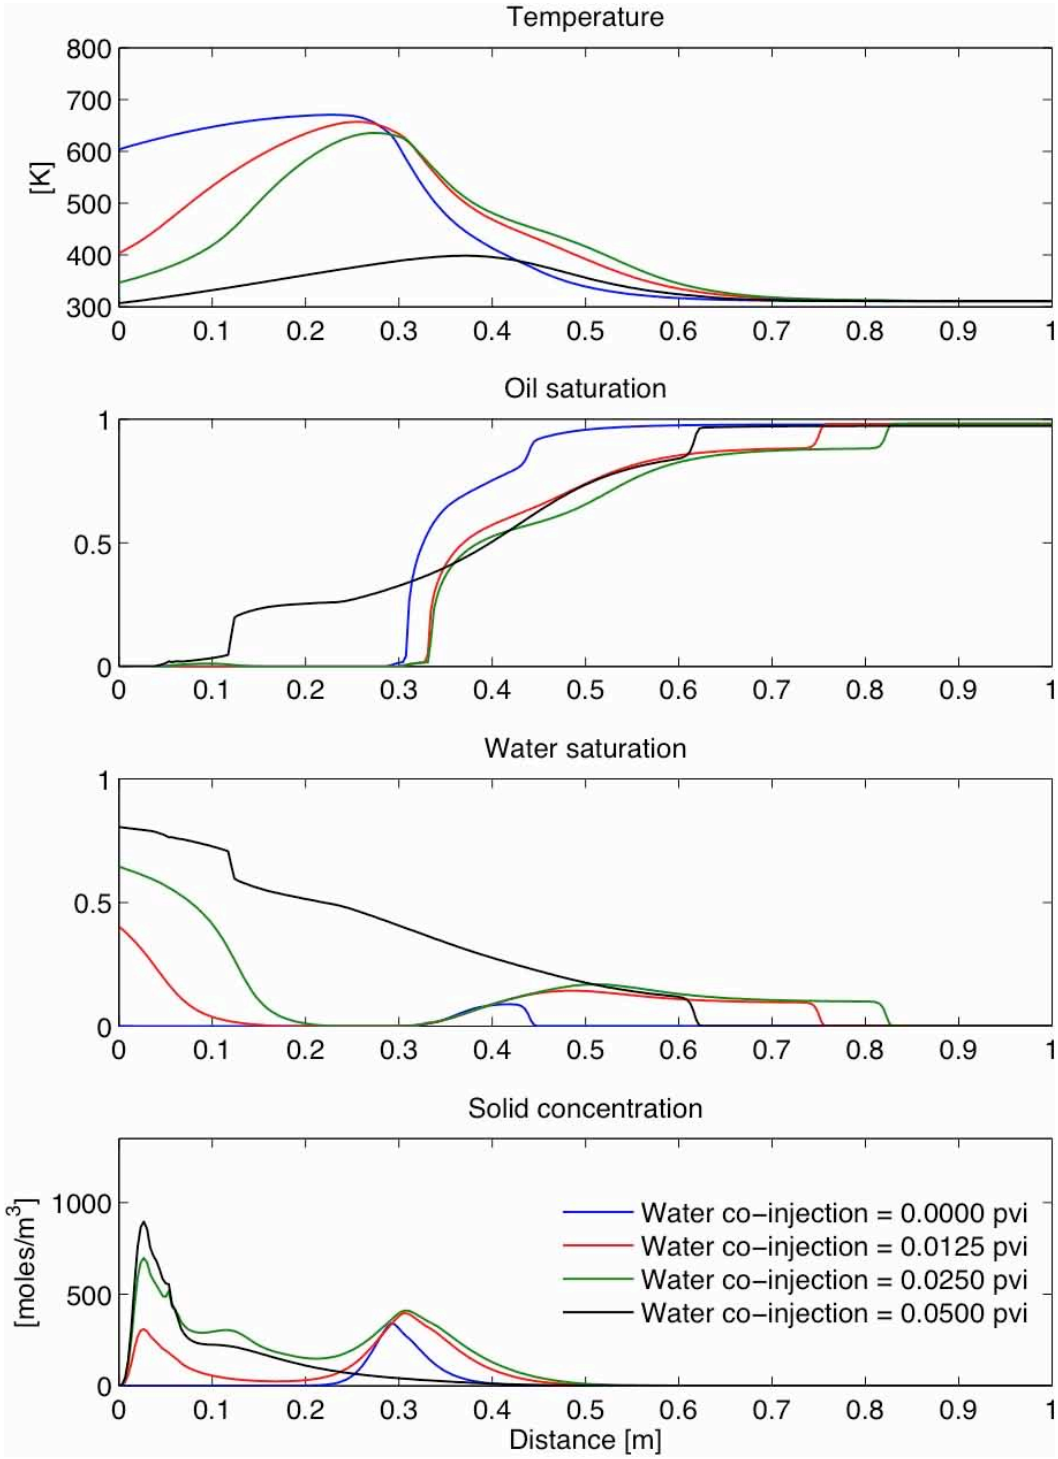

Figure 9.6: Wet combustion simulation results with varying water co-injection rate. Temperature, saturation, and solid coke concentration profiles are shown at $t = [6]hrs$ and four different water rates. The rates are given in pore volumes injected per hour.

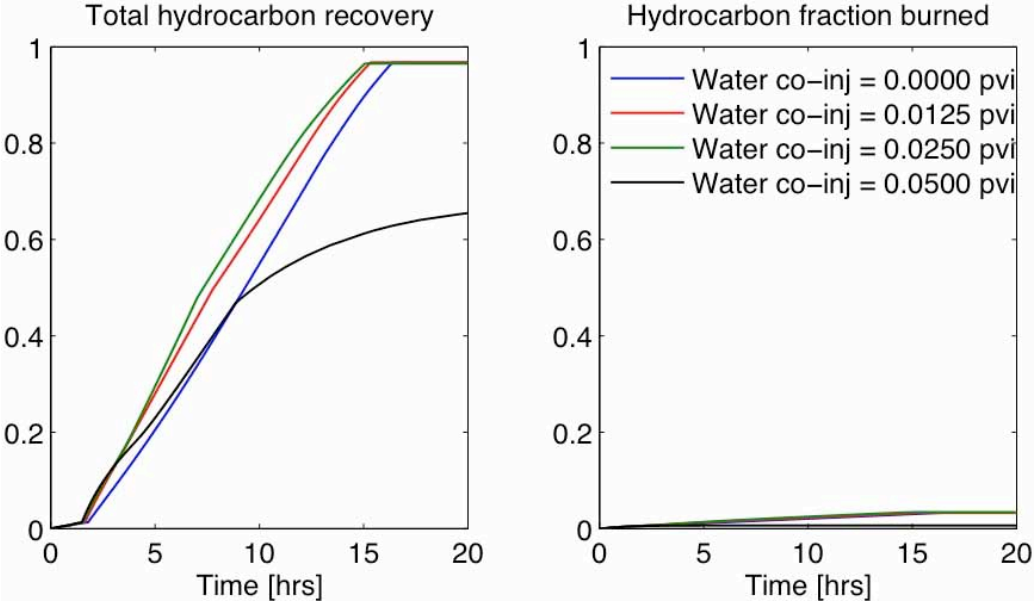

Figure 9.7: Production curves for wet combustion simulation with varying water co-injection rates. The total cumulative hydrocarbon recovery is shown along with the fraction of hydrocarbons consumed in reactions.

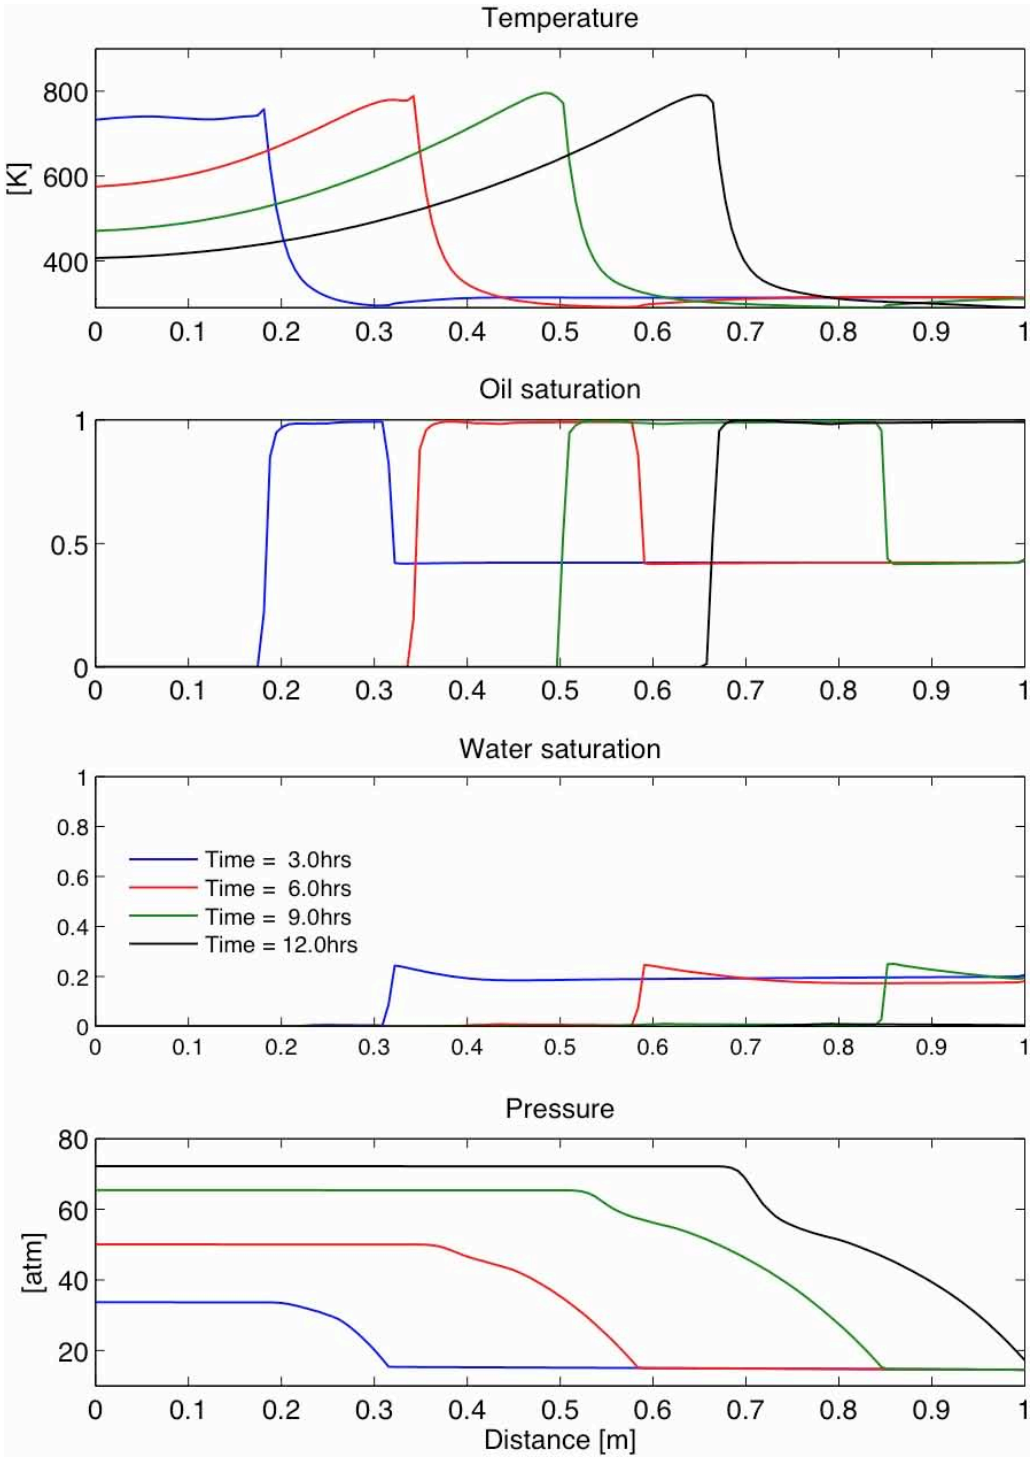

Figure 9.8: Dry combustion simulation results for the SARA based reaction model. Temperature, saturation, and pressure profiles are shown at four different times during the simulation.

10. Adaptive Spatial Grids and Specialized Upscaling

To capture the processes that occur on small spatial scales, and to accurately represent reservoir heterogeneity, we developed an spatially adaptive gridding strategy. In this chapter we discuss the adaptive grid we employ as well as two new transmissibility upscaling methods that we developed specifically for adaptive grids.

In this section, we use single phase, steady and incompressible flow in a heterogeneous reservoir. The governing dimensionless pressure equation is obtained by combining the continuity equation $\nabla \cdot \mathbf{u} = \mathbf{0}$ and Darcy's law given by $\mathbf{u} = -\mathbf{k} \cdot \nabla p$, which leads to

$$\nabla \cdot (\mathbf{k} \cdot \nabla p) = 0. \tag{10.1}$$

Here p is the pressure, $\mathbf{u}$ is the Darcy velocity and $\mathbf{k}$ the permeability tensor, all of which are non-dimensionalized by appropriate reference values. We ignored any sources or sinks in the domain. This equation is valid on the fine scale, at which we assume that the permeability tensor as given by geostatistical methods is diagonal, constant in each grid cell and may be highly variable in space. Although homogenization of equations with variable coefficients generally yields coarse scale equations of a different form than the original fine scale equations, it is common practice to keep the same equation as 1 for the coarse pressure, with the permeability tensor replaced by the coarse scale permeability tensor $\mathbf{k}^*$, also referred to as the effective permeability tensor, and p by the coarse scale pressure. Hence, in the remainder of the chapter we will use equation 10.1 indiscriminately of scale.

For more complex multi-phase flow problems, the system of governing equations generally includes a pressure equation similar to Eqn. 10.1, based on the commonly used multi-phase extension to Darcy's law, that is coupled to a system of transport equations describing the movement of either phase saturations or multicomponent species. Although we focus on this relatively simple single-phase system, any transmissibility upscaling method, including the method developed here, is applicable to more general multi-phase systems.

10.1 Cartesian Cell-based Anisotropic Refined (CCAR) Grids

In local grid adaptation strategies the grid density is increased where needed to improve solution accuracy. Specificially to ISC processes, we desire it to improve the representation of heterogeneity in regions with high permeability contrasts, and to zoom in on critical areas (propagation fronts, and combustion zones). We think of adaptive grid strategies as part of the class of multi-scale methods, as discussed in,

for example, [Kipp2007] and references therein. Multi-scale methods, designed to capture multiple scales involved in the reservoir fluid flow processes, include approaches that upscale to a coarse simulation grid and reconstruct the solution on a finer scale within each coarse grid cell using a functional representation, but also approaches based on grid refinement. Whether representing the flow on the fine-scale functionally or discretely, the success of the multi-scale method depends to a great extent on the chosen boundary conditions on the local fine regions.

When designing an efficient, adaptive simulation approach, a grid topology must ideally facilitate fast transitions from dense to coarse grids without sacrificing grid and solver quality, allow efficient computation despite potentially high levels of refinement, and lead to fast grid generation. We have found a Cartesian grid topology with anisotropic cell-based refinement and coarsening quite attractive because it allows aggressive grid adaptation while maintaining grid quality. We use adaptivity, rather than flow-based gridding, to align the grid with important flow and geological features and assess its effectiveness in ensuring flow accuracy. The CCAR grid is formed by a number of anisotropic refinements of its cells; i.e., splitting a cell in half along a coordinate direction, to construct a suitably refined mesh. This is illustrated for three refinement levels in Figure 10.1.

We impose the restriction that no cell can have more than four neighbors in each direction (or two neighbors in two dimensions) to maintain solution accuracy across interfaces [Berger1989]. The CCAR grids are stored using the unstructured data approach first suggested in [Ham2002], which leads to efficient look-ups. In this work, we denote cells in a CCAR grid, G, by $C_i, i = 1,\ldots,N_{cells}$, with N_{cells} the number of cells in G, and cell faces by $F_i, i = 1,\ldots,N_{faces}$, with N_{faces} the total number of faces in G.

For heterogeneous media, the computation of the flux through a face with a hanging node requires careful consideration. Consider Figure 10.2, showing three cells $C_i, i = 1,2,3$, with dimensions h_x^i and h_y^i, and diagonal permeability tensors with diagonal elements k_x^i and k_y^i. We will construct a flux, similar to a two-point flux, for the face of cell 1 bordering cells 2 and 3. The area of the face of cell i that borders this face is denoted by A_i. The average pressures in the cells i are denoted by p_1, p_2 and p_3.

Requiring continuity of the numerical flux across the half-faces, we obtain the equations

$$\frac{k_x^j}{h_x^j}(p(m_j) - p(c_j)) = \frac{k_x^1}{h_x^1}(p(d_j) - p(m_j)), \quad j = 2,3, \tag{10.2}$$

where the points m_j, c_j and d_j are as given in the figure, and the corresponding fluxes

$$q_j = \frac{2h_y^j}{\frac{h_x^1}{k_x^1} + \frac{h_x^j}{k_x^j}} (p(d_j) - p(c_j)), \quad j = 2,3. \tag{10.3}$$

We represent the pressure at the auxiliary points d_2 and d_3 in terms of the available pressures at nearby cell centers. To that end, we use a convex combination of two interpolations, one with a vertical orientation and one with a horizontal orientation, which are illustrated in Figure 10.3. For the vertically-oriented interpolation in d_2, we use bilinear interpolation at c_1, and the two closest cell centers besides c_2 and c_3, that form a triangle containing d_2. For the horizontally-oriented interpolation of d_2, we use bilinear interpolation at c_1, c_2, and the closest cell center that completes a triangle containing d_2. Similar triangles are constructed for d_3. We combine the horizontal and vertical interpolations with a weighting determined by the angle of the pressure gradient through the point d_i, $i = 2$ or 3, at which we are interpolating. This introduces a natural flow-based interpolation, and the resulting, wider, stencil at hanging nodes is reminiscent of the L-scheme suggested in [Aavatsmark2006].

The resulting matrix remains sparse with a reasonable condition number thanks to the compact finite volume discretization. Because we do not have a strict TPFA discretization at hanging nodes, the matrix is not guaranteed to be an M-matrix and hence positivity is not ensured, but the number of matrix elements with the wrong signs is typically very small. Solution of the system using an algebraic multigrid solver [BoomerAMG] is fast with almost linear convergence properties for our test cases.

10.2 Transmissibility Upscaling for Adapted Grids with Two-point Flux Approximations

A large variety of single-phase upscaling techniques exist that are used to find representative permeability or transmissibility fields on the coarse grids. In purely local methods, the upscaled parameters are computed over fine-scale regions that correspond to the target coarse block under study. In extended local upscaling methods, regions neighboring the coarse block are included. In both approaches, the local flow is driven by generic Dirichlet, Neumann or periodic boundary conditions imposed on the fine-scale regions, see [Durlofsky1991, Pickup1994]. Although local methods may give satisfactory results for some reservoirs, they generally do not perform well when the spatial scales of important reservoir heterogeneities, such as high-permeability flow paths, are larger than the simulation grid, unless these grids are aligned with such paths. Global methods, which use global fine-scale flow simulations to determine coarse-scale parameters, can reduce errors in the presence of large connected flow paths, but they are generally expensive

[Holden2000]. An attractive and relatively novel method to improve large-scale connectivity in upscaled permeabilities or transmissibilities is the local-global (LG) upscaling method [Chen2003]. Here, a standard extended local method is used to obtain an approximate coarse permeability or transmissibility field. Then, a coarse global solve is performed with generic boundary conditions. Interpolation of the coarse solution provides the Dirichlet boundary conditions for the next iteration of local solves. The process is repeated until the coarse field converges, which generally takes only a couple of iterations.

Discretization strategies can also strongly impact the quality of the coarse solutions. With two-point flux approximations (TPFA), a flux through a given cell face is computed using pressure values from the cells that share the face, whereas with multi-point flux approximations (MPFA), additional cells may be involved. TPFA methods are attractive because of their simplicity, compactness and monotonicity. In many cases, the TPFA approximation is reasonable, but in cases with significant full-tensor effects, that may be introduced by the permeability field directly or by grid topology, larger errors should be expected, as assuming a TPFA is similar to neglecting the cross terms in upscaled permeability tensors. Interestingly, and perhaps surprisingly, TPFA methods combined with local-global procedures have been shown to provide good upscaled transmissibilities even in challenging cases [Chen2003, Chen2006]. It is generally agreed that when strong full-tensor effects are present, MPFA methods are preferable.

In this section we address how to combine local-global upscaling and local grid adaptivity to effectively represent large-scale heterogeneities on computational grids for single-phase flows, how to control the local grid adaptation during the upscaling process to improve the upscaling accuracy. For concreteness and simplicity, we apply the upscaling strategy to the cell-based anisotropic Cartesian refinement strategy discussed in the previous section. For similar reasons, we limit ourselves to TPFA methods. The proposed methodology is general enough, however, to be applicable to alternate gridding and discretization strategies. We focus on finite volume upscaling methods that lead to a cell-centered, coarse-scale pressure equation and conservative fluxes.

We consider Eqn. 10.1 in two dimensions. In case of diagonal permeability tensors, the standard finite volume discretization on Cartesian grids computes the flux f across the cell face shared by the neighboring cells (i, j) and $(i+1, j)$ as

$$f = -T_{i+1/2,\,j}\left(p_{i,\,j} - p_{i+1,\,j}\right) \tag{10.4}$$

and similar for (i, j) and $(i, j+1)$. The transmissibility T can be computed from upscaled permeability values, as in $T = \dfrac{k_{i+1/2} H \Delta y}{\Delta x}$, where H is the model thickness, and where the interblock permeability $k_{i+1/2}$ is computed via a harmonic average of the upscaled permeabilities k_i and k_{i+1}. In this work, we will instead compute the transmissibility of each coarse cell face T directly using a local or extended local

method, as discussed in the introduction. In the finite volume discretization the numerical fluxes across all faces of the cell are summed and equated to zero. This leads to a matrix-vector equation of the form $A\mathbf{p} = \mathbf{b}$, where A is the matrix containing the transmissibilities and any coefficients resulting from the prescribed boundary conditions, $\mathbf{p}$ the vector containing all unknown p_{ij} and $\mathbf{b}$ a vector containing boundary data.

10.2.1 Multi-Level Local-Global

As mentioned in the introduction, it has been shown that for Cartesian grids, the local-global (LG) method leads to improved results as compared to local, or extended local, methods [Chen2003]. For each coarse grid cell, effective permeabilities or face transmissibilities are computed by solving a local flow problem on the fine geo-cellular grid directly around the coarse cell. This local flow problem is driven by boundary conditions interpolated from a coarse global solve. The resulting coupling between the local flow problems and the coarse global solve leads to an iterative scheme that generally converges rapidly. Because of its effectiveness, LG upscaling serves as the cornerstone of the local-global upscaling algorithm for CCAR grids, which we now present.

In permeability upscaling, we associate with each cell $C_i, i = 1,$ in the CCAR grid G an extended region $E_i = C_i \cup B_i$, where B_i is a border region of constant width r around C_i. In transmissibility upscaling the (extended) local region is centered around a coarse cell face F_i, rather than the cell itself. In each region E_i we discretize the pressure equation on the uniform geocellular grid, which here we assume to be Cartesian. We note that in the local-global upscaling method in [Chen2003] the width of each dimension of the border region is chosen to be a multiple of the dimension of the corresponding grid cell instead. This model works well on a uniform grid, but in a CCAR grid, the coarsest cells do not such a large border because the fine-scale permeability within such cells tends to be smoother relative to the coarse grid scale. In permeability upscaling our border region scaling is also attractive for efficiency reasons: if a cell C_i with corresponding extended region E_i should be bisected during refinement, the resulting new cells will have extended regions whose union coincides with E_i. This allows pressure values on the boundary of E_i to be re-used for the new regions.

The general MLLG algorithm is given as:

Construct the base CCAR grid G using static refinement indicators
while grid adaptation not completed **do**
while not converged **do**
for each cell C_i (permeability upscaling) or each face in G **do**
Select boundary conditions on E_i:
if this is the first iteration **then**
Use generic boundary conditions: constant pressure in one
coordinate direction, no-flow in the other coordinate direction
else
Use bilinear interpolation of global solutions p_x and p_y to
approximate pressure on the boundary of E_i
Use Dirichlet boundary conditions using these pressure values
end
Solve the pressure equation on E_i with selected boundary conditions
Compute fluxes between the fine grid cells in E_i and pressure values at their cell
centers
Use these fluxes and pressure values on E_i to compute new diagonal permeability
tensor K_i for C_i, or a new transmissibility for F_i
end
Solve generic flow problems on G to obtain pressure fields p_x and p_y
Check for convergence
done
Check if any of the cells in G must be refined
done

The first step in the MLLG algorithm is the construction of an initial base CCAR grid. This base grid is designed with the aim of preserving important high permeability flow paths. After the base grid is generated, we apply local-global upscaling, as in the Cartesian case, to find representative permeability values on the cells C_i of the

CCAR grid, or transmissibility values for the faces F_i of the CCAR grid. Once the local-global step has converged, which generally only takes one or two iterations, we check if further grid refinement is necessary, and repeat the local-global iteration on the updated grid. Thus, the MLLG algorithm has an inner iteration, which is a single application of local-global upscaling for CCAR grids, and an outer iteration with the number of outer iteration steps bounded above by the number of levels of refinement between the coarsest scale of the base grid, and the scale of the fine grid.

10.2.2 Creation of the Base CCAR Grid

To create a base CCAR grid, we apply a static refinement indicator based on the method proposed by Younis and Caers in [Younis2002]. The method is based on a ``leaky'' connected-set approach that delineates clusters of fine cells with similar permeability. Heuristic measures are used to color each connected set, indicating the level of the importance of capturing the set by refinement. Refinement flags are then obtained by explicit boundary detection. An illustration of the process appears in Figure 10.4. In the current implementation, a minimum level of refinement is set based on the resolution of the geocellular grid and efficiency considerations, and a CCAR grid G is generated by iteration until these level constraints are met. Such a grid is illustrated in Figure 10.5.

In our experience, this approach has given us reasonable CCAR grids to start our MLLG process. We note that the grid is further adapted during MLLG.

10.2.3 Boundary Data for Local Solves

For each cell C_i of the CCAR grid G, we must solve local flow problems (in two dimensions) over the corresponding extended region E_i. In order to obtain the Dirichlet boundary data for each such problem, we use bilinear interpolation, with the interpolation points obtained from each of the global coarse-scale solutions computed during the previous iteration of local-global upscaling. While it is possible to use a higher-order interpolation scheme, it is not advisable to do so because the interpolant would exhibit oscillatory behavior, as opposed to the monotonicity we would expect from a pressure field.

We illustrate the interpolation for permeability upscaling. In Figure 10.6, a cell C_i is shown, along with its extended region E_i. For each of the points p_{ij}, $j = 1,2,\ldots,8$ along the boundary of E_i, we find the three cells of G whose centers are as close to p_{ij} as possible, and form a triangle containing p_{ij}. We then use barycentric coordinates [Watson1994] to approximate the pressure at p_{ij}.

10.2.4 Local Flow-based Refinement of the Base Grid

The fluxes across a face, from each global pressure solve, can be used to determine whether cells bordering the face should be refined. Specifically, let F_i be a face with area A_i, and let q_i be the flux across the face obtained by a pressure solve on the CCAR grid with flow driven in the direction perpendicular to the face. If Q is the global flow, and A is the area of the face of the global flow domain perpendicular to the direction of the flow, then we determine that the cells sharing the face should be refined if

$$|q_i| > \alpha \frac{A_i}{A} |Q|, \tag{10.5}$$

where α is a proportionality constant that depends on the ratio of areas, i.e. $\alpha = \alpha\left(\frac{A_i}{A}\right)$. The term $A_i |Q|/A$ can be interpreted as the face's ``fair share'' of the global flow.

To choose $\alpha(A_i/A) > 1$, we use the following heuristics. First, when A_i is the area of a face of a cell at the coarsest scale in the grid, then $\alpha(A_i/A)$ should be approximately 1, so that refinement will always take place except in areas that are receiving less than their fair share of flow. This low tolerance is necessary to avoid overlooking high-flow paths that occur within large regions that otherwise have very low flow. On the other hand, as A_i approaches the size of a face in the fine grid, then $\alpha(A_i/A)$ should be approximately A/A_i, so that refinement will not take place unless there is very high flow; that is, nearly all of the global flow is crossing the face F_i. This progression from low to high tolerance is used to avoid superfluous refinement while not missing any finer-scale regions of high flow. Under the assumption that $\alpha(A_i/A) = C(A_i/A)^{-p}$ for some positive real number p, we obtain

$$\alpha(A_i/A) = \frac{\delta}{(A/A_c)^p}\left(\frac{A_i}{A}\right)^{-p},$$

where A_c is the area of the largest face in the grid parallel to F_i, $\delta \approx 1$ is the chosen value of $\alpha(A_c/A)$, and p is given by

$$p = \frac{\ln \frac{A}{\delta A_f}}{\ln \frac{A_c}{A_f}},$$

where A_f is the area of each fine grid face that is parallel to F_i.

If the face F_i borders only two cells, then both are refined. Otherwise, the face borders two hanging cells and one non-hanging cell; only the non-hanging cell is refined. All refinement is in the direction that is orthogonal to the face.

We can use additional refinement criteria to reduce the number of ``hanging'' nodes that occur in high-flow regions. First, we apply the above flow-based criterion to each ``half-face'' that is shared by a hanging cell and a non-hanging cell. If the flow across this half-face is more than its fair share, then the non-hanging cell is refined along the direction that is orthogonal to the face. This refinement also takes place if too much of the flow across the entire face (75%, in practice) is through one hanging node.

We note that in transmissibility upscaling, we also refine if the estimated transmissibility for a face is negative. This is further discussed in the next subsection.

10.2.5 A New Local Gradient-based Weighting of Transmissibilities

Referring back to Figure 10.2 we choose the extended region E_i centered at the face F_i for which we compute the flux to be large enough to contain the cells centered at the points c_1, c_2 and c_3. We solve the pressure equation on E_i with Dirichlet boundary conditions as in permeability upscaling. The computation of the fluxes from the local solution, along with the approximate pressure differences $\Delta p_j = p(d_j) - p(c_j)$ obtained by averaging the local pressure field over the appropriate regions, can be used to compute the transmissibilities T_2 and T_3 corresponding to the fluxes q_1 and q_2. These transmissibilities are then used in the stencils centered at c_2 and c_3. In the stencil centered at c_1, we use the transmissibility $T_1 = T_2 + T_3$.

The original local-global upscaling algorithm in [Chen2003] computed upscaled transmissibility across a face by considering only data obtained from a global flow driven by a pressure difference in the direction perpendicular to the face. But, since we are using two generic flow problems, we actually obtain two fluxes for each face. These fluxes should indeed not have equal weight: a flow that is nearly parallel to the face should not contribute much, if anything, to the computed face transmissibility. If it does, this may result in a stronger process dependency, which we would like to avoid. Rather than choosing only one flow, we designed a new weighting process for the transmissibilities, which we now describe.

Consider the simple case of a face shared by only two cells. Let $q^{(1)}$ and $q^{(2)}$ denote the fluxes across the face obtained from the two local pressure solves, obtained either in an extended local or a local-global method, and let $\Delta p^{(1)}$ and $\Delta p^{(2)}$ denote the differences in the cell averages of the pressure from these solves. From these values, we obtain two approximations for the transmissibility,

$$T^{(j)} = -\frac{q^{(j)}}{\Delta p^{(j)}}, \; j = 1,2. \tag{10.6}$$

Let n be a unit vector that is perpendicular to the face. We find a unique transmissibility T using weighted least squares, with the weights determined by the average pressure gradients near the face, $\nabla p^{(1)}$ and $\nabla p^{(2)}$, where the averages may be cell or line averages (it remains to be determined which approach is best). This approach yields the system $WAT = W\mathbf{b}$, where

$$W = \begin{bmatrix} \omega^{(1)} & 0 \\ 0 & \omega^{(2)} \end{bmatrix}, \quad A = \begin{bmatrix} \Delta p^{(1)} \\ \Delta p^{(2)} \end{bmatrix}, \quad \mathbf{b} = -\begin{bmatrix} q^{(1)} \\ q^{(2)} \end{bmatrix},$$

and

$$\omega^{(i)} = \frac{|\nabla p^{(i)} \cdot n|}{\mathsf{P}\nabla p^{(i)}\mathsf{P}}, \quad i = 1,2.$$

The rationale for choosing the weights $\omega^{(i)}$ in this manner is to determine, in an average sense, the orientation of the flow from the two local solves. If either of these flows is determined to be nearly parallel to the face, then it is not be assigned much significance in the computation of the transmissibility across the face. On the other hand, if the flow is nearly perpendicular to the face, then the ratio of the negative flux to the pressure difference is believed to represent the coarse-scale transmissibility accurately.

The weighted least squares problem has the solution

$$T = \frac{A^T W^T W \mathbf{b}}{A^T W^T W A} = -\frac{[\omega^{(1)}]^2 q^{(1)} \Delta p^{(1)} + [\omega^{(2)}]^2 q^{(2)} \Delta p^{(2)}}{[\omega^{(1)}]^2 [\Delta p^{(1)}]^2 + [\omega^{(2)}]^2 [\Delta p^{(2)}]^2}.$$

In the case of a face adjacent to hanging nodes, we apply this same approach to compute the transmissibilities across the half-faces. If either of the equations in (5) would yield a negative transmissibility, then it is not included in the computation of T. If both equations would yield negative transmissibilities, then we set T equal to the sum of the fine-scale transmissibilities along the face, and flag the cells sharing the face for refinement.

We are able to prevent most instances of negative transmissibility from arising in the first place by computing flux across each face as follows. Assume that a face F_i in the coarse grid consists of n fine-scale faces f_j, $j = 1, \ldots, n$. We compute the flux using

$$q_i = \sum_{j=1}^{n} T_j \Delta p_j,$$

where, for each j, T_j is the transmissibility across f_j, and Δp_j is the pressure difference across the face. To compute this pressure difference, we first use the pressure values in the two fine-scale cells c_j and d_j bordering f_j. Then, we compute a pressure difference $\Delta p_{j'}$ using the pressure values in the two cells $c_{j'}$ and $d_{j'}$ that are one cell removed from f_j along the direction orthogonal to f_j. If Δp_j and $\Delta p_{j'}$ have opposite sign, then we assume that Δp_j is an isolated fluctuation

and compute the flux across f_j using $\frac{1}{3}\Delta p_{j'}$; otherwise we use Δp_j. The cells involved in the computation are shown in Figure 10.7.

10.3 Transmissibility Upscaling for Adapted Grids with Multi-Point Flux Approximations

We aim to construct a multi-point finite-volume scheme that has four desirable properties. First, we would like our scheme to be "close" to a two-point scheme. The reasons for this criterion are simplicity and efficiency (by maximizing matrix sparsity), robustness (as the scheme is more likely to generate an M-matrix), and consistency (the scheme should reduce to a two-point scheme in the case of homogeneous permeability on a Cartesian grid). Second, the scheme should be easily applicable to adaptive grid strategies, such as the CCAR. Third, the scheme should be very accurate for smooth pressure fields. If the pressure field is not smooth, improved accuracy can be achieved by local grid refinement. Fourth, the scheme should perform well for grids with a high aspect ratio, such as are commonly used in reservoir simulators.

10.3.1 Construction of VCMP for Cartesian Grids

To achieve these goals, we allow the stencil used in the multi-point flux approximation to vary per cell face. Figure 10.8 depicts an interior face in a Cartesian grid. Our multi-point flux approximation will use a subset of or all of the six pressure values $p_j, j = 1,\ldots,6$ indicated in the figure. We will use two generic flow problems on an (extended) local region around each face. The multi-point flux approximations are required to honor the two flow problems at the face, as well as two other constraints on the stencil weights that are given below. The remaining degrees of freedom are used to drive the scheme as close as possible to a two-point approximation. We now discuss this process in more detail.

For each $j = 1,\ldots,6$, we let t_j denote the weight that will be assigned to point j in the flux approximation. The flux approximation has the general form

$$f = -\mathbf{t}^T\mathbf{p}, \tag{10.7}$$

where

$$\mathbf{t} = \begin{bmatrix} t_1 & \cdots & t_6 \end{bmatrix}, \quad \mathbf{p} = \begin{bmatrix} p_1 & \cdots & p_6 \end{bmatrix}^T. \tag{10.8}$$

As in local upscaling methods, we solve the pressure equation on the local region of the fine grid containing the six points with two generic Dirichlet boundary conditions. We let $p^1(x,y)$ and $p^2(x,y)$ be the solutions of these local problems, and p_j^i denote the value of $p^i(x,y)$ at point j.

The boundary conditions are chosen so as to satisfy the following criteria: (1) the pressure is linear along the boundary, (2) the pressure varies between 0 and 1, and (3) the gradients for the two problems are parallel to each of the coordinate axes. The pressure field p^1 is computed using boundary values chosen so that flow is across the face, and p^2 is obtained from boundary values chosen so as to drive flow parallel to the face.

For extended local upscaling, we add a border region around the coarse grid cells centered at points 1 through 6. We then prescribe the same boundary values as in the purely local case, except that these values are prescribed at the corners of the extended region, rather than at points 1 through 6. Then, the pressure values at points 1 through 6 are obtained from the computed pressure fields.

For both problems, the values of the pressure at all other fine grid boundary points are obtained by linear interpolation of the values at the corners of the local (or extended local) region. Figure 10.9 illustrates the pressure fields p^1 and p^2 for an example heterogeneous permeability field, along with the fine-scale fluxes across the face for both solutions.

For $i = 1,2$, we let f_i denote the coarse-scale flux (sum of fine-scale fluxes) across the face obtained from the local solution $p^i(x, y)$. To compute the weights $\{t_j\}_{j=1}^6$, we solve a general optimization problem in which we wish to compute

$$\min_{\mathbf{t}} \sum_{i=1}^{2} \alpha_i^2 \, | \mathbf{t}^T \mathbf{p}^i - f_i |^2 + \sum_{j=3}^{6} \beta_j^2 t_j^2 \tag{10.9}$$

subject to the essential linear constraints

$$\sum_{j=1}^{6} t_j = 0, \quad t_{2j-1} \le 0, \quad t_{2j} \ge 0, \quad j = 1,2,3. \tag{10.10}$$

In the current implementation, the weights α_i are chosen to be $| f_i |$ and the weights β_j are chosen to be equal to $(| f_1 | + | f_2 |)/M$, where M is a tuning parameter. The larger the value of M, the more closely the flows are honored. For small M we will obtain a two-point flux with $t_j = 0$ for $j = 3,4,5,6$. For simplicity in this paper, we use $M = 1000$. Depending on the main orientation of the local heterogeneity, the fine-scale fluxes of one of the flow problems (e.g., p^1 in Figure 10.9) may have the greater role than the other in determining the upscaled transmissibilities.

We solve this problem using the lsqlin function from MATLAB's Optimization Toolbox. This function employs an active set method, which is also a projection method. It finds an initial feasible solution by first solving a linear programming problem. If a minimizing set of weights $\{t_i\}_{i=1}^6$ is found, then we examine the weights and determine whether any of them are negligibly small. If so, then we solve the

minimization problem (10.9) again, with the corresponding variables excluded from consideration. We repeat this process until we obtain a solution for which no weights are negligibly small, or until we are left with a 2-point flux. The resulting weights are then used to approximate the flux in the relevant finite-volume stencils.

It is worth noting that in the limiting case $M \to \infty$, the fluxes f_i must be honored exactly, and the weights t_3 through t_6 must be minimized subject to this constraint as well as the constraints (10.10). In this case, the optimization problem (10.9) can easily be solved analytically. First, we compute four three-point stencils t^i, for $i = 3,4,5,6$. Each stencil t^i involves points 1, 2 and i, and it can be obtained by solving a system of three linear equations

$$\begin{bmatrix} 1 & 1 & 1 \\ p_1^1 & p_2^1 & p_i^1 \\ p_1^2 & p_2^2 & p_i^2 \end{bmatrix} \begin{bmatrix} t_1^i \\ t_2^i \\ t_i^i \end{bmatrix} = \begin{bmatrix} 0 \\ -f_1 \\ -f_2 \end{bmatrix}. \tag{10.11}$$

Then, the stencil t is given by

$$t = \sum_{i=3}^{6} w_i t^i \tag{10.12}$$

where the weights w_i, $i = 3,\ldots,6$, are given by

$$w_i = \frac{\prod_{k \neq i} [t_k^k]^2}{\sum_j \prod_{k \neq j} [t_k^k]^2} \tag{10.13}$$

where all products and sums only consider those three-point stencils t^i for which t_i^i has the correct sign, according to (10.10). It follows that the cost of computing the spatially varying stencils can be bounded independently of the behavior of the flow around the face. Furthermore, this bound is reasonable; the solution of the optimization problem (10.9), (10.10) does not add undue computational cost to the upscaling process. It follows that the computational expense of VCMP is comparable to that of extended local upscaling, and is therefore negligible compared to that of an entire reservoir simulation.

In the rare case that this approach breaks down (e.g. if no correctly-signed three-point stencils are available), and no minimizing set of weights can be found that also satisfies the constraints (10.10), we use a two-point flux with points 1 and 2, with the transmissibility chosen so that, as closely as possible, both flows are honored in a weighted least-squares sense. This is accomplished using the flux-based weighting of transmissibilities, discussed in the previous section. In our experiments, we encountered a very small number of negative transmissibilities (the resulting stencil does not have the correct signs). Almost all of these cases were for the channelized problem discussed below. In the adaptive grid framework, we use grid refinement to remove the problem cells, which is an effective strategy as shown below. On

Cartesian grids, we again use a two-point flux, with the transmissibility equal to the sum of the fine-scale transmissibilities across the face.

10.3.2 Extension to CCAR Grids

VCMP can easily be extended to CCAR grids. Figure 10.10 depicts a face shared by a hanging node, centered at point 2, and its parent, centered at point 1.

For such a face, the flux approximation is constructed from at most five points, rather than six. We then compute (10.9) subject to the constraints (10.10), as before, with obvious modifications due to the absence of any of the points 3, 4, 5 or 6.

In the rare cases that no solution can be found to (10.9), our approach is similar to the Cartesian case. Here, we then fall back on a compact three-point stencil including points 1 and 2 that honors the local flow across the problem face. If that were to fail as well, but we have not encountered any such cases in our experiments, we flag this face for refinement, or resort, as before, to using a transmissibility equal to the sum of the fine-scale transmissibilities across the face.

We use four refinement criteria to guide grid adaptivity during the upscaling process. The first two criteria are the flow-based criteria introduced in the previous section for MLLG, which flag a face for refinement if it has a disproportionately large flux, and reduce the number of "hanging" nodes that occur in high-flow regions. The third criterion is to refine if the computed transmissibility across the face is negative, as we discussed in a previous section. This refinement is isotropic. Finally, we examine the nonzero elements of the matrix A obtained from the flux approximations. Suppose that in the $i-th$ row, there exists an index $j \neq i$ such that A_{ij} has the same sign as A_{ii}. If cell j shares a face with cell i, then the face is flagged for refinement. The reason for this refinement is that the sign of A_{ij} is contrary to the goal of obtaining an M-matrix, and therefore is an indication that the discretization is not sufficiently robust. However, if there is low flow across this face in both global coarse pressure fields, the refinement is cancelled.

10.4 Results

10.4.1 Multi-Level Local-Global upscaling

We demonstrate the effectiveness of the MLLG upscaling method on a variety of permeability fields and boundary conditions.

Table 1: Abbreviations used to refer to upscaling algorithms in Section 10.4.

Abbreviation	Description
EL	Extended local, on uniform Cartesian grid
EL-W	EL with new transmissibility weighting
LG	Local-global, on uniform Cartesian grid
LG-W	Local-global with new transmissibility weighting, on uniform Cartesian grid
EL-CCAR	Extended local, on CCAR grid
MLLG	Multi-level local-global, on CCAR grid
MLLG -W	Multi-level local-global with new transmissibility weighting, on CCAR grid

The Test Suite

Our test suite consists of three types of fine permeability fields that are used in the experiments described in this section. All have dimension $L_x \times L_y$, where $L_x = 256$ and $L_y = 64$, and uniform grid step size 1 in both spatial dimensions. The fields are:

- 10 realizations with long correlation lengths of $\ell_x = 50$, $\ell_y = 10$, aligned with the grid.
- 10 realizations with long correlation lengths of $\ell_x = 75$, $\ell_y = 5$, not aligned with the grid
- Four channelized domains, given by layers 44, 48, 76 and 80 of the SPE10 Comparative Project (see [Christie2001]) that are modified to fit the domain dimensions $L_x = 256$ and $L_y = 64$.

Examples of realizations for the permeability fields with long correlation lengths are shown in Figure 10.11. For these cases, which are generated using two-point statistics, we use mean 3.0 and standard deviation 1.735, which corresponds to a log-normal distribution. The fine-scale permeability fields of the channelized domains are shown in Figure 10.12.

The main purpose of the realizations with long correlation lengths is to show the effectiveness of the new transmissibility upscaling obtained with the new weighting scheme. The results also show the accuracy of local-global upscaling vs. extended local upscaling, but we refer to [Chen2003] for more extensive comparisons. For these cases, we use the upscaled transmissibilities to solve the pressure equation on a Cartesian grid with 32 cells in the x-direction, and 16 cells in the y-direction. The dimensions of the cells are $\Delta x = 8$ and $\Delta y = 4$.

The channelized cases are used to assess the effectiveness of MLLG and MLLG-W. Channelized reservoirs are generally difficult to upscale accurately because of the strong and large-scale flow paths. For the channelized cases, we use a CCAR grid obtained by iterative refinement of an initial 8×4 uniform Cartesian grid, with $\Delta x = 32$ and $\Delta y = 16$. The refinement is performed according to the criteria described above.

For each permeability field and each grid, we solve the pressure equation with the following six sets of boundary conditions, in order to measure how effective each algorithm is in reducing process dependency resolving the fine-scale velocity:

- Constant pressure/no-flow boundary conditions in the x- and y-directions, respectively. For the first set, we prescribe $p(0, y) = 1$, $p(L_x, y) = 0$, and $p_y(x,0) = p_y(x, L_y) = 0$. For the second set, we prescribe $p(x,0) = 1$, $p(x, L_y) = 0$, and $p_x(0, y) = p_x(L_x, y) = 0$. These are the generic boundary conditions used for the global coarse solves in local-global upscaling, and also for the local fine solves in local or extended-local upscaling.

- Flow from each of the four corners of the domain. For example, for flow from the lower left corner, we prescribe $p(0, y) = \exp(-(2y/L_y)^4)$ (shown in Figure 10.13), $p(x,0) = \exp(-(2(L_x - x)/L_x)^4)$, and $p(L_x, y) = p(x, L_y) = 0$. The boundary conditions for the other three corners are defined similarly.

We will compare the total flow to that obtained from a fine-scale solution, for each set of boundary conditions, and for each upscaling method. For cases generated from 2-point statistics, we present the mean and standard deviation of relative errors in total flow across all realizations. In selected cases, we will also examine the velocity fields, compared to the fine-scale velocity field averaged over each coarse grid cell.

Long Correlation Length, Aligned with Grid

The relative errors in total flow are given in Table 10.2 for the 10 realizations. Clearly, both methods outperform EL. This is expected, and in agreement with what is reported in [Chen2003], as EL cannot accurately represent permeability contrasts at a scale larger than the coarse grid size. We see that LG and LG-W are quite comparable, due to the fact that the gradient-based weighting is designed to behave like LG in regions where the flow is oriented with the grid, which is the case with these realizations.

Table 10.2: Mean percentage errors in total flow, and corresponding standard deviations, after applying various upscaling methods to ten realizations with long correlation length, aligned with the grid, using a 32×16 coarse grid.

	Fine		EL		LG		LG-W	
Flow	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
In x-direction	53.97	6.11	1.36	0.31	0.45	0.36	0.44	0.39
In y-direction	53.97	6.11	0.50	0.19	0.58	0.22	0.63	0.19
From lower left	18.83	6.00	1.00	0.56	1.17	0.71	1.16	0.78
From upper left	−19.62	7.67	1.58	0.72	1.55	0.84	1.50	0.73
From lower right	16.97	5.69	1.39	0.69	1.09	0.66	1.16	0.62
From upper right	−17.63	6.36	0.80	0.42	0.80	0.46	0.85	0.40

Long Correlation Length, Not Aligned with Grid

Table 10.3 lists the relative errors in the total flow for the 10 realizations with long correlation lengths $\ell_x = 75$ and $\ell_y = 5$, not aligned to the grid. This time, both EL and LG perform poorly. LG-W has the best overall performance, and is quite accurate for all boundary conditions, except for the low-flow case in which flow is driven in the x-direction. Such sensitivity of the upscaling accuracy in weaker flow cases is fairly typical for TPFA methods on permeability fields that are not grid-aligned. In such cases, it is indeed preferable to go to MPFA methods, as discussed in section 10.3 and below.

Table 10.3: Mean percentage errors in total flow, and corresponding standard deviations, after applying various upscaling methods to ten realizations with long correlation length, not aligned with the grid, using a 32×16 coarse grid.

	EL		LG		LG-W	
Flow	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
In x-direction	6.38	4.58	22.10	5.26	29.82	5.91
In y-direction	14.94	5.95	6.35	3.72	3.15	1.74
From lower left	26.84	7.13	16.07	4.52	4.53	2.94
From upper left	18.94	6.73	9.59	3.88	3.63	2.77
From lower right	12.99	4.76	6.03	3.38	3.27	4.02
From upper right	26.58	5.95	15.52	4.84	4.69	2.21

Channelized cases

Table 10.4 summarizes the results for layer 44 for the global flow, comparing multi-level local-global upscaling with the new transmissibility weighting (MLLG-W) to extended local upscaling on a Cartesian grid (EL), local-global upscaling on a Cartesian grid (LG), local-global upscaling with the new weighting on a Cartesian grid, (LG-W), and multi-level local-global upscaling without the new weighting (MLLG). We see that in all cases, the new transmissibility weighting improves the results significantly. Not only are the errors in total flow reduced strongly, process dependency is also greatly reduced. MLLG yields comparable accuracy to the Cartesian methods, while using a CCAR grid that contains significantly fewer cells due to coarsening in low-flow regions.

Table 4: Fine-scale total flow (first column) and percentage errors in total flow for modified SPE10 Layer 44 for various upscaling methods and boundary conditions. The EL, LG and LG-W methods use a 32×16 Cartesian grid, MLLG uses a 365-cell CCAR grid, and MLLG-W a 371-cell CCAR grid. The difference in these CCAR grids is due to the differences in transmissibility calculations, thus differences in local flow calculations, which impact the refinement criteria.

	Fine-scale					
Flow	total flow	EL	LG	LG-W	MLLG	MLLG-W
In x-direction	7.8544	13.68	3.55	3.61	0.09	3.59
In y-direction	35.7709	7.20	2.68	3.57	2.31	3.51
From lower left	7.4063	8.51	1.78	0.44	1.57	1.46
From upper left	-1.346	0.96	12.53	0.62	20.31	1.73
From lower right	21.0316	3.95	3.97	6.54	3.40	5.93
From upper right	-28.0695	4.86	3.27	4.64	2.60	4.26

Table 10.5 shows the percentage errors in total flow for layer 76. This layer poses severe challenges to TPFA methods. Indeed, none of the methods are able to match the fine-scale flow with reasonable accuracy in the cases of flow from the upper left or lower right corner, so they are omitted here. For the remaining boundary conditions, we observe that (1) using a CCAR grid preserves the accuracy obtained using a Cartesian grid with more cells, and (2) gradient-based weighting reduces the process dependency for both Cartesian and CCAR grids, making MLLG-W the best all-around choice.

Table 10.5: Errors in total flow for modified SPE10 Layer 76, for various upscaling methods and boundary conditions. The EL, LG and LG-W methods use a 32×16 Cartesian grid, LG-CCAR uses a 425-cell CCAR grid, and MLLG uses a 422-cell CCAR grid.

Flow	Fine scale	EL	LG	LG-W	MLLG	MLLG-W
In x-direction	4.2935	32.91	5.69	3.39	5.93	3.99
In y-direction	1.5661	2.46	4.20	4.63	3.83	4.52
From lower left	0.42436	38.23	36.81	26.79	33.69	26.25
From upper right	-0.099821	9.30	16.96	19.46	17.26	20.50

Table 10.6 shows the percentage errors in total flow for the same upscaling methods for layer 80. In this layer, the main channel is oriented along the grid. Because of this orientation, the upscaling methods that employ gradient-based weighting do not provide as much benefit, but they still exhibit reduction in process dependency. In particular, note the smaller errors for the low-flow cases.

Table 10.6: Errors in total flow for modified SPE10 Layer 80, for various upscaling methods and bounday conditions. The EL, LG and LG-W methods use a 32×16 Cartesian grid, MLLG uses a 431-cell CCAR grid, and MLLG-W uses a 419-cell CCAR grid.

Flow	Fine scale	EL	LG	LG-W	MLLG	MLLG-W
In x-direction	14.6725	12.23	6.22	4.87	5.88	6.10
In y-direction	3.0696	4.07	0.11	1.24	0.02	1.29
From lower left	13.2104	6.81	6.69	6.77	4.28	6.21
From upper left	-8.9824	6.53	5.02	7.08	0.37	4.60
From lower right	-7.8925	5.31	9.66	6.83	10.05	8.28
From upper right	-0.60702	0.30	17.33	12.39	15.27	14.24

10.4.2 Variable Compact Multi-Point Approximation

Short Correlation Length

Table 10.7 lists relative errors in total flow for one of the realizations with short correlation lengths $\ell_x = \ell_y = 5$. We see that for the case in which $p(x, y) = x + y$ on the boundary, EL seems the most accurate, with VCMP-EL being the most accurate, on average, across all other cases.

Table 10.7: Relative errors in total flow for domain with short correlation length, for various upscaling methods and boundary conditions, on a 32×16 coarse grid.

Flow	Fine scale	EL	ELW	MLLG	VCMP	VCMP-EL
In x-direction	4.13	0.04	0.20	0.29	0.05	0.04
In y-direction	77.38	0.03	0.00	0.01	0.04	0.01
From lower left	29.85	0.01	0.01	0.01	0.04	0.01
From upper left	– 32.57	0.02	0.00	0.02	0.06	0.04
From lower right	34.24	0.03	0.02	0.03	0.05	0.02
From upper right	– 32.47	0.03	0.02	0.02	0.06	0.02
$p = x + y$ on bdry	– 6731.42	0.01	0.04	0.11	0.09	0.08
$p = x - y$ on bdry	3918.35	0.02	0.02	0.03	0.06	0.02

To determine whether EL really is the most accurate for permeability fields of this type and with the boundary pressure equal to $x + y$, we examine the flow errors for all ten realizations. The results of this examination are given in Table 10.8. We see that the superior accuracy of EL for this case was coincidental, and that VCMP-EL in fact performs most consistently across all realizations.

Table 10.8: Mean errors in total flow, and corresponding standard deviations, after applying various upscaling methods to ten realizations with short correlation length, using a 32×16 coarse grid.

	EL		MLLG		VCMP-EL		
Flow	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	
In x-direction	0.04	0.01	0.28	0.03	0.05	0.01	
In y-direction	0.03	0.01	0.01	0.00	0.01	0.00	
From lower left	0.02	0.01	0.01	0.01	0.01	0.01	
From upper left	0.02	0.01	0.01	0.01	0.01	0.01	
From lower right	0.03	0.02	0.01	0.01	0.01	0.01	
From upper right	0.03	0.01	0.02	0.01	0.01	0.01	
$p = x + y$ on bdry	0.04	0.02	0.06	0.03	0.02	0.02	
$p = x - y$ on bdry	0.03	0.02	0.06	0.02	0.01	0.01	

Long Correlation Length, Not Aligned with Grid

Table 10.9 lists relative errors in total flow for one of the realizations with long correlation lengths $\ell_x = 75$ and $\ell_y = 5$, not aligned to the grid. This time, EL performs quite poorly compared to other methods. MLLG is quite accurate except for the cases in which flow is driven in the x-direction, and for $p(x, y) = x - y$ on the boundary. On the other hand, VCMP-EL is very accurate for all boundary conditions. This superiority is confirmed by examining the flow errors for all ten realizations. as shown in Table 10.10.

Table 10.9: Errors in total flow for domain with long correlation length, not aligned with the grid, for various upscaling methods and boundary conditions, on a 32×16 coarse grid.

Flow	Fine scale	EL	ELW	MLLG	VCMP	VCMP-EL
In x-direction	2.92	0.05	0.48	0.54	0.13	0.08
In y-direction	97.11	0.18	0.02	0.04	0.06	0.02
From lower left	55.85	0.37	0.07	0.09	0.03	0.00
From upper left	– 52.73	0.33	0.06	0.10	0.02	0.00
From lower right	41.88	0.21	0.01	0.04	0.06	0.02
From upper right	– 58.36	0.30	0.02	0.01	0.04	0.01
$p = x + y$ on bdry	– 16132.22	0.39	0.04	0.04	0.00	0.01
$p = x - y$ on bdry	2105.36	0.44	0.07	0.07	0.12	0.03

Table 10.10: Mean errors in total flow, and corresponding standard deviations, after applying various upscaling methods to ten realizations with long correlation length, not aligned with the grid, using a 32×16 coarse grid.

	EL		MLLG		VCMP-EL	
Flow	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
In x-direction	0.03	0.02	0.49	0.10	0.06	0.02
In y-direction	0.23	0.06	0.03	0.02	0.02	0.01
From lower left	0.35	0.08	0.05	0.05	0.01	0.01
From upper left	0.27	0.08	0.06	0.06	0.01	0.01
From lower right	0.21	0.05	0.03	0.03	0.02	0.02
From upper right	0.36	0.06	0.04	0.02	0.01	0.01
$p = x + y$ on bdry	0.41	0.04	0.05	0.03	0.01	0.00
$p = x - y$ on bdry	0.45	0.13	0.09	0.05	0.04	0.02

Channelized Domain

Table 10.11 lists relative errors in total flow for a channelized domain, using a 32×16 coarse grid. For this particularly difficult test case, we see that no one method performs more accurately than the others for all boundary conditions. If the case where $p(x, y) = x + y$ on the boundary is excluded, then MLLG is clearly the most accurate.

Similar results are obtained using a CCAR grid with each method, as shown in Table 10.12. However, MLLG, EL and ELW do not perform nearly as well on the $p(x, y) = x + y$ case, likely due to insufficient refinement. For these methods, the refinement is determined by analyzing local flow for the two generic problems in which flow is driven in the coordinate directions, so it is not surprising that the resulting grid would not be suitable for flow driven across the portion of the domain through which there is little flow in the generic problems. On the other hand, the variations of VCMP use additional refinement criteria based on matrix properties, resulting in additional refinement and greater accuracy for this case.

Table 10.11: Errors in total flow for a channelized domain, for various upscaling methods and boundary conditions, on a 32×16 coarse grid.

Flow	Fine scale	EL	ELW	MLLG	VCMP	VCMP-EL
In x-direction	7.85	0.19	0.08	0.04	0.08	0.11
In y-direction	35.77	0.10	0.02	0.01	0.08	0.02
From lower left	7.41	0.15	0.04	0.00	0.04	0.01
From upper left	–1.35	0.34	0.02	0.01	0.21	0.18
From lower right	21.03	0.07	0.05	0.01	0.09	0.01
From upper right	– 28.07	0.09	0.04	0.01	0.07	0.00
$p = x + y$ on bdry	–16696.89	0.52	0.19	0.27	0.07	0.05
$p = x - y$ on bdry	– 2163.40	0.00	0.10	0.01	0.01	0.03

Table 10.12: Errors in total flow for a channelized domain, for various upscaling methods and boundary conditions, on CCAR grids.

Flow	Fine scale	EL	ELW	MLLG	VCMP	VCMP-EL
In x-direction	7.85	0.19	0.08	0.04	0.08	0.11
In y-direction	35.77	0.10	0.02	0.00	0.08	0.02
From lower left	7.41	0.15	0.04	0.01	0.05	0.03
From upper left	–1.35	0.35	0.08	0.02	0.21	0.18
From lower right	21.03	0.07	0.04	0.01	0.08	0.01
From upper right	– 28.07	0.09	0.04	0.01	0.07	0.01
$p = x + y$ on bdry	–16696.89	0.63	0.36	0.28	0.07	0.05
$p = x - y$ on bdry	– 2163.40	0.28	0.40	0.17	0.00	0.03

10.5 Discussion

10.5.1 Upscaling for Two-point Flux Approximations

This work shows that it is possible to obtain good upscaling accuracy on challenging permeability fields, such as the channelized cases used in this paper, on Cartesian adapted grids (CCAR grids) and two-point flux approximations (TPFA) using the proposed multi-level local-global method (MLLG). We compared global flow errors for channelized systems, which are generally challenging for upscaling methods, produced by upscaled CCAR grids to errors produced by local-global upscaling on Cartesian grids of similar density. In [Chen2003] it was shown that the uniform local-global method on these channelized systems strongly outperformed local or extended local upscaling methods. For the generic boundary conditions used in the upscaling process the upscaled CCAR grids cut these global flow errors again, often as much as 50% or more. Process dependency is also reduced significantly. These results suggest that refinement is an effective means to control upscaling errors and reduce process dependency. A large contributor to the improved accuracy presented in this paper is the new gradient-based weighting of transmissibilities, which we strongly recommend for local-global upscaling strategies.

We are especially encouraged by the accuracy we obtain with our method in the local flow. The combination of adaptivity and local-global upscaling retains important connected flow paths and accurately compute local flow directions and speeds. This is promising for displacements with high mobility contrasts, such as gas injection processes, where the mobile fluid will seek these high permeability paths. We are implementing the CCAR methods in a compositional solver to further investigate this potential benefit.

The computational overhead of MLLG is relatively low, because global solves are performed on coarse grids, the sizes of extended local regions are limited, and generally very few steps are required for the local-global iteration to converge.

The effectiveness of MLLG is good news, as it means that it is perhaps not as essential to align computational grids, thus simplifying gridding strategies, or to always use more expensive multi-point flux approximations. Multi-point flux approximations are necessary in cases with very strong permeability tensor effects and/or weak flows. Indeed, our results with VCMP, reported above, show that multi-point flux approximations in those cases are preferable.

Future work will include the investigation of other possible adaptivity criteria, including criteria for coarsening, and alternative methods of handling occurrences of negative transmissibilities. The upscaling approach presented here, while suitable

for CCAR grids, can just as easily be applied to other types of grids. Future work will include generalization of our algorithm to other gridding strategies such as patched refinements, unstructured grids, or curvilinear grids.

10.5.1 Upscaling for Multi-point Flux Approximations

The results suggest that VCMP(-EL) has the best overall performance for our suite of test problems, and is a very promising approach. In the cases exhibiting mild effects of full-tensor anisotropy, the EL and Weighted EL methods performed well with few exceptions. The local-global method and VCMP also performed well for these permeability fields with only one exception for each method. Local-global upscaling does not perform as well in cases where the total flow is low, or the flow direction is not adequately represented by the generic global flow problems used to determine the boundary conditions for the local solves. VCMP, on the other hand, is less accurate in cases where the pressure field is rough. In these cases the current choice of local boundary conditions may not be sufficiently good. But, overall, our results indicate that VCMP is less sensitive to changes in the magnitude or direction of flow than the local-global method.

For the more challenging permeability fields with stronger full-tensor effects VCMP clearly outperformed the other methods overall and showed the least process dependency. However, for both permeability fields there were certain choices of boundary conditions for which other methods obtained superior results. In particular for the channelized example, LG gave very accurate total flow rate results, which is to be expected as it was designed for such cases specifically.

For the channelized field, we also observed a couple of instances in which the pressure fields computed using VCMP and VCMP-EL lacked monotonicity. This can occur because there is no strict guarantee with VCMP that the pressure matrix will be an M-matrix. The variability of the VCMP stencil ensures that connections between diagonally neighboring cells will have the correct sign, but not for direct neighbors that share a face. We introduced a modification of VCMP in which fluxes that contribute to incorrectly-signed elements are recomputed, with additional constraints to control the magnitude of selected weights. By exploiting the flexibility in the minimization problem (10.9), (10.10), we are able to guarantee an M-matrix with minimal loss of accuracy.
We are also investigating a local-global version of VCMP that we hope will combine the strengths of LG and the VCMP method presented here. Global information can be incorporated into VCMP in several ways. For example, we can simply replace the linear boundary conditions used for both local flow problems with the boundary conditions used by LG which are interpolated from the two coarse generic flow problems. However, care must be taken to ensure that the local flow problems are complimentary. Our initial implementation of VCMP-LG avoids this complication by using global information from only one of the generic coarse flows. The choice is made independently for each coarse face based on the magnitude of the fluxes

across the face. The second local flow is defined using linear boundary conditions with the gradient chosen to complement the other flow. Preliminary results using this approach are shown in Table 10.13. We note that a similar strategy can be used to incorporate specific (well-driven) flow into VCMP, as well as MLLG.

Table 10.13: Errors in total flow for domain with long correlation length, aligned with the grid, for variations of VCMP and various boundary conditions, on a 32×16 coarse grid.

Flow	Fine scale	VCMP-EL	VCMP-LG
In x-direction	9.41	0.00	0.00
In y-direction	39.22	0.00	0.00
From lower left	14.16	0.02	0.02
From upper left	−12.87	0.05	0.03
From lower right	13.86	0.02	0.01
From upper right	−16.49	0.01	0.01
$p = x + y$	−5854.41	0.02	0.02
$p = x - y$	81.16	0.26	0.15

While our focus has been on the two-dimensional case, VCMP has a straightforward extension to Cartesian grids in three dimensions that involves three local flow problems for each coarse face. In 3D, the VCMP stencil only includes points corresponding to cells that share faces and cells that share edges. Hence, the VCMP stencil will contain between 7 and 19 point whereas a full multipoint stencil contains 27 points. The extension of VCMP to 3D CCAR grids is also manageable since there is no need to individually consider each possible configuration of cells around hanging nodes when determining transmissibility.

10.6 Summary

10.6.1 Upscaling for Two-point Flux Approximations

For highly heterogeneous (e.g., channelized) systems, the proposed integration of grid adaptivity and upscaling is shown to consistently provide more accurate coarse-scale models for global flow, relative to reference fine-scale results, than do

existing upscaling techniques using TPFA discretizations, including local-global upscaling, applied to uniform grids of similar densities.

Adaptation allows the use of a coarse-scale model that contains fewer grid points than would be required when using a Cartesian grid to achieve the same accuracy during simulation.

Apart from more accurate global flow results, the proposed method also computes improved local flow results. Connected high permeability flow paths are preserved, and local velocity directions and magnitude in these important paths are computed with good accuracy.

We developed a refinement indicator based on local fluxes that guides the grid adaptation process during upscaling and further reduces solver and upscaling errors.

Process dependency is strongly reduced, that is, the approach computes accurate global flow results also for flows driven by boundary conditions different from the generic boundary conditions used to compute the upscaled parameters.

Upscaling accuracy can be significantly improved and process dependency reduced with the use of a gradient-based weighting of transmissibilities computed from multiple local flow problems in the local-global method.

10.6.2 Upscaling for Multi-point Flux Approximations

We have introduced a new (extended) local transmissibility upscaling scheme that we are calling the variable compact multi-point method or simply VCMP. Because a multi-point flux approximation is used, VCMP is able to accommodate full-tensor anisotropy which is generally present in coarse scale flow problems. However, VCMP is designed to be as close as possible to a two-point flux approximation scheme, and in particular, the flux approximation reduces to two points for homogeneous cases and other cases lacking full-tensor effects. In 2D on structured grids, the VCMP stencil for the coarse pressure equation has between five and nine points. The stencil is variable in the sense that it adapts to the orientation of the underlying fine permeability distribution. In the case that adaptive grids are used, hanging nodes that occur at grid interfaces can easily be handled due to the unstructured formulation of the transmissibility calculations. The combination of these features makes VCMP a very promising scheme for upscaling flow through general heterogeneous media.

As expected, for permeability fields with significant full-tensor anisotropy VCMP outperforms local and extended local methods that are based on a two-point flux approximation. In addition, VCMP compares favorably with local-global upscaling methods (also based on a two-point flux approximation) although for the

channelized permeability field the local-global scheme produced superior results in three out of eight flow problems. This is not surprising as local-global methods were designed for channelized systems. Future improvements and extensions of VCMP including a local-global version were discussed in the previous section. In addition to accurately predicting global flow rates, upscaled transmissibility fields computed with VCMP locally agree well with averaged fine scale results and hence are suitable for use in multi-phase flow simulations.

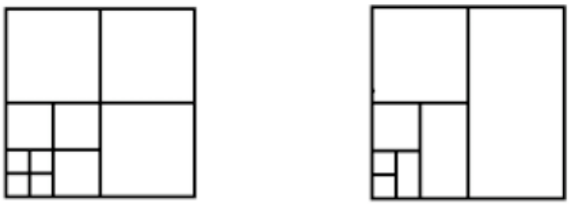

Figure 10.1: CCAR grid cells with 3 levels of refinement. Left: isotropic refinement. Right: anisotropic refinement.

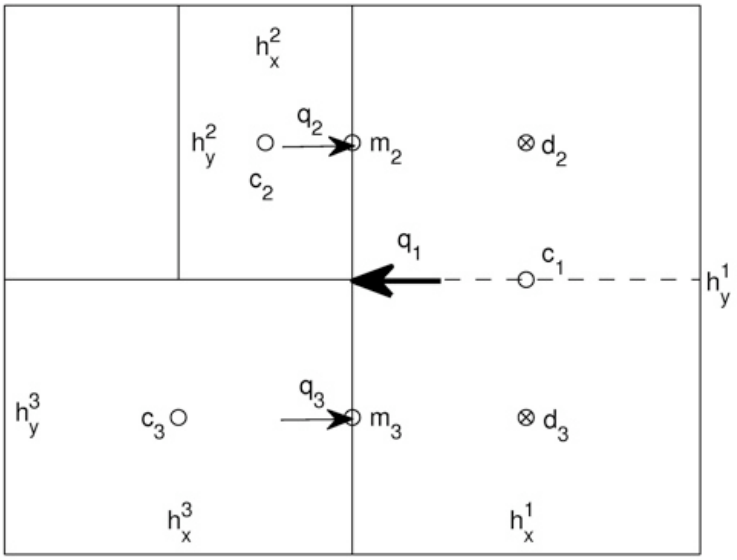

Figure 10.2: The points c_j, $j=1,2,3$, d_j, $j=2,3$ and m_j, $j=2,3$, are used to compute the fluxes q_2 and q_3 across the half-faces between cell 1 and cell 2, and between cell 2 and cell 3.

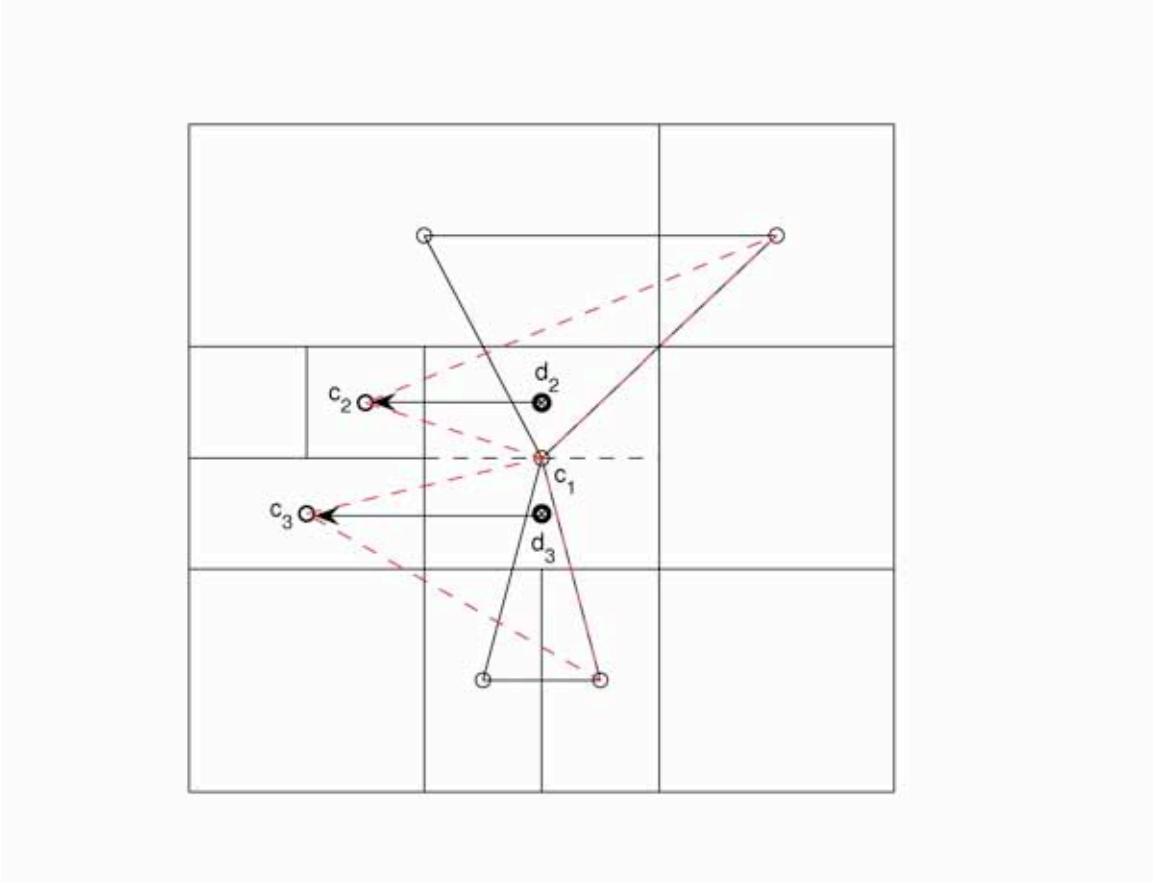

Figure 10.3: Triangulations used to represent the pressure at d_2 and d_3 in terms of cell centers. The values at these points are obtained by bilinear interpolation of the values at the vertices of the horizontally, or vertically, oriented triangles containing them.

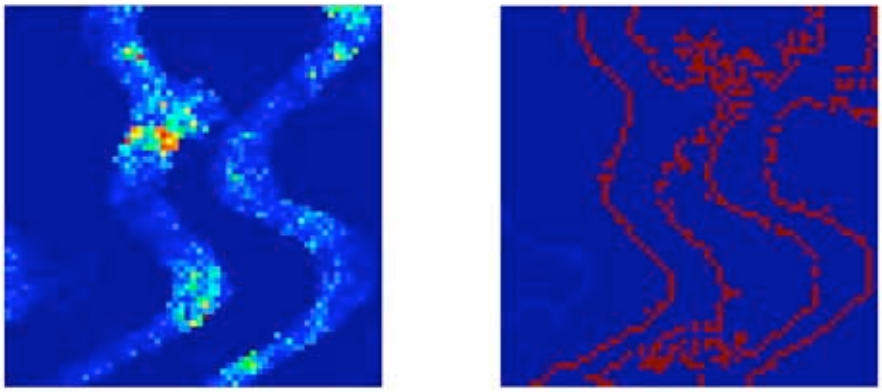

Figure 10.4: Permeability field (left), and boundaries of connected flow paths detected by leaky connected-set approach (right)

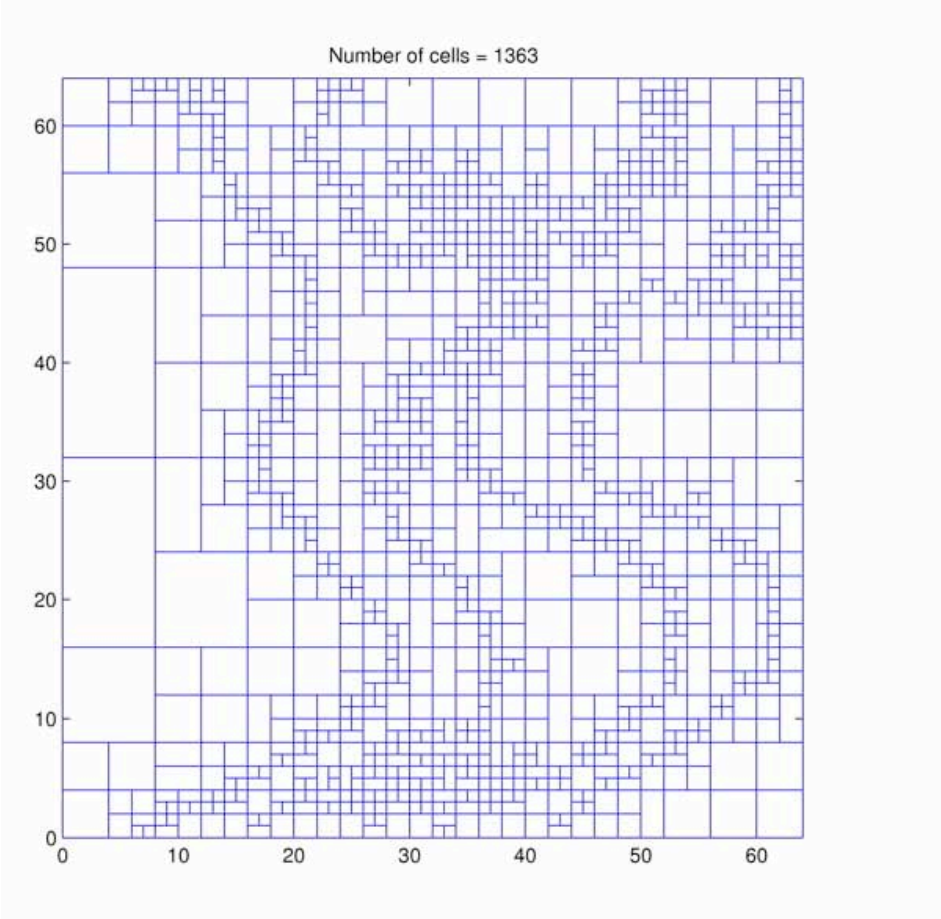

Figure 10.5: base CCAR grid generated from permeability field and connected sets from Figure 4

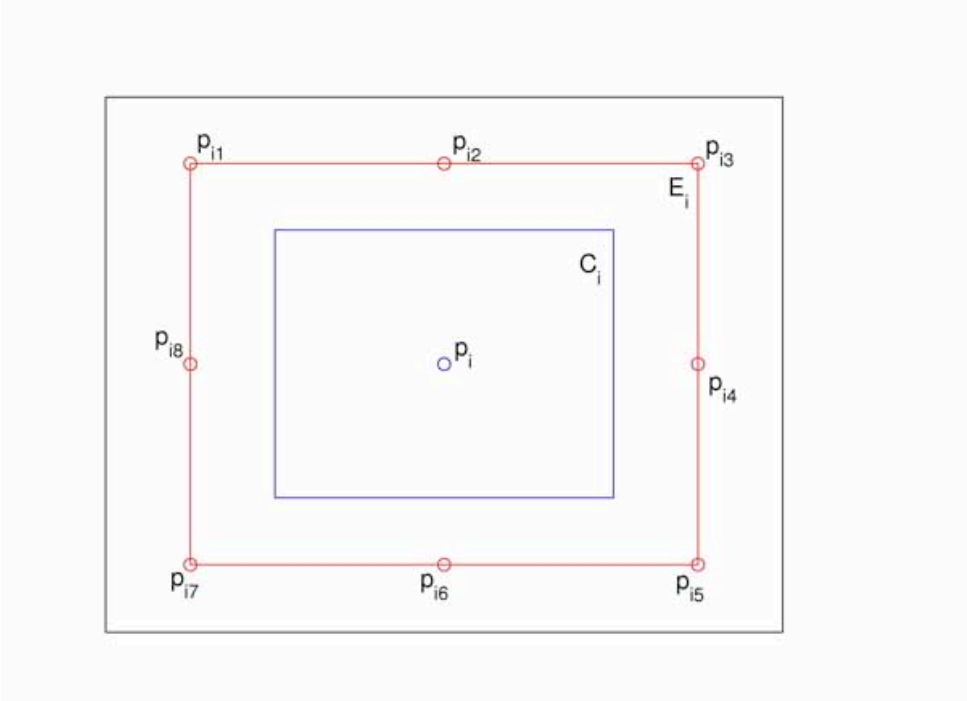

Figure 10.6: points p_{ij}, $j = 1,2,\ldots,8$ on the boundary of an extended region E_i, shown for permeability upscaling, at which global coarse-scale pressure is interpolated. The region E_i corresponds to the cell C_i with center p_i.

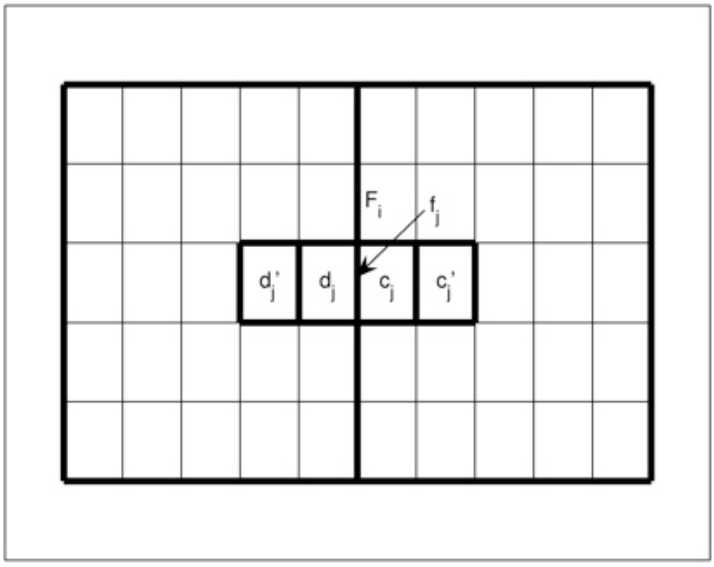

Figure 10.7: The cells c_j, d_j, $c_{j'}$ and $d_{j'}$ used to compute the pressure differences Δp_j and $\Delta p_{j'}$ across the face f_j that is a portion of the face F_i, the flux across which is used to compute the upscaled transmissibility across F_i.

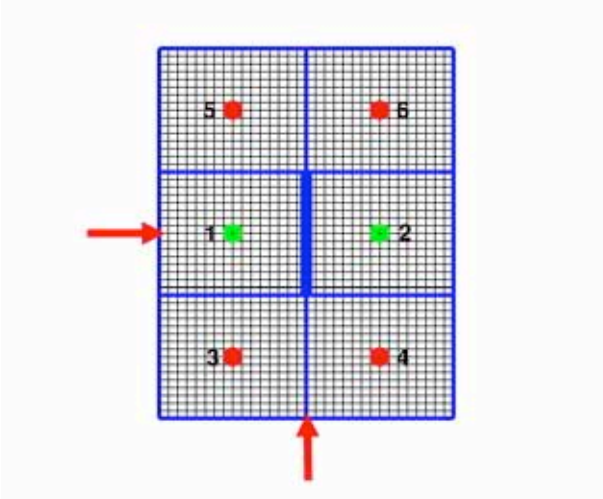

Figure 10.8: The red and green dots denote the centers of cells of a Cartesian grid that are used to construct a VCMP approximation to the flux across a face, indicated by the thick blue line. The green dots correspond to pressure values whose weights that are to be dominant; the red dots correspond to values whose weights are to be minimized.

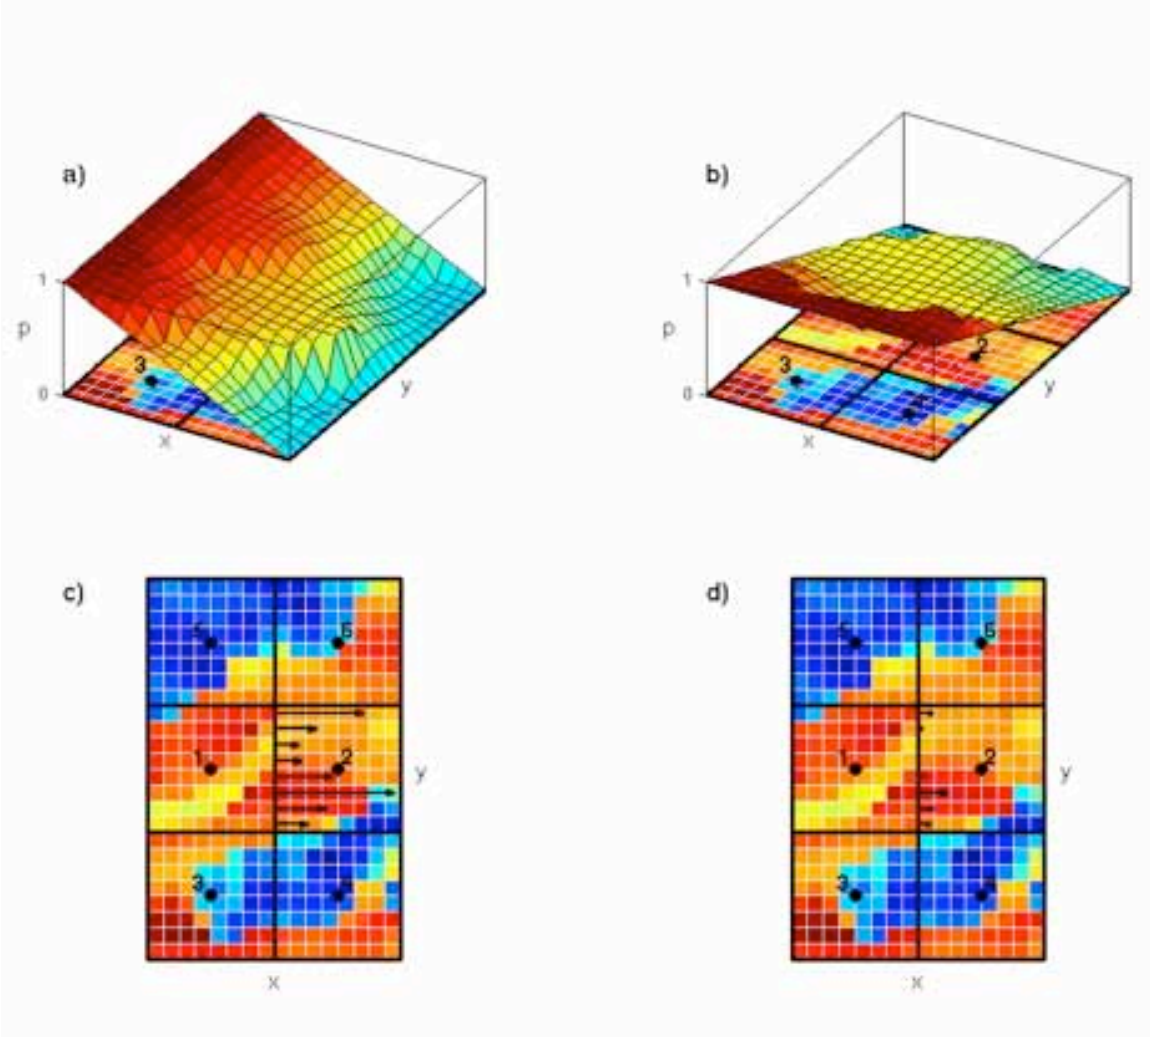

Figure 10.9: Plots (a) and (b) show solutions p^1 and p^2 to local fine-scale problems with generic boundary conditions as described in the text. The black arrows in plots (c) and (d) denote the fine-scale fluxes from p^1 and p^2, respectively, across the coarse grid face at the center of the local region.

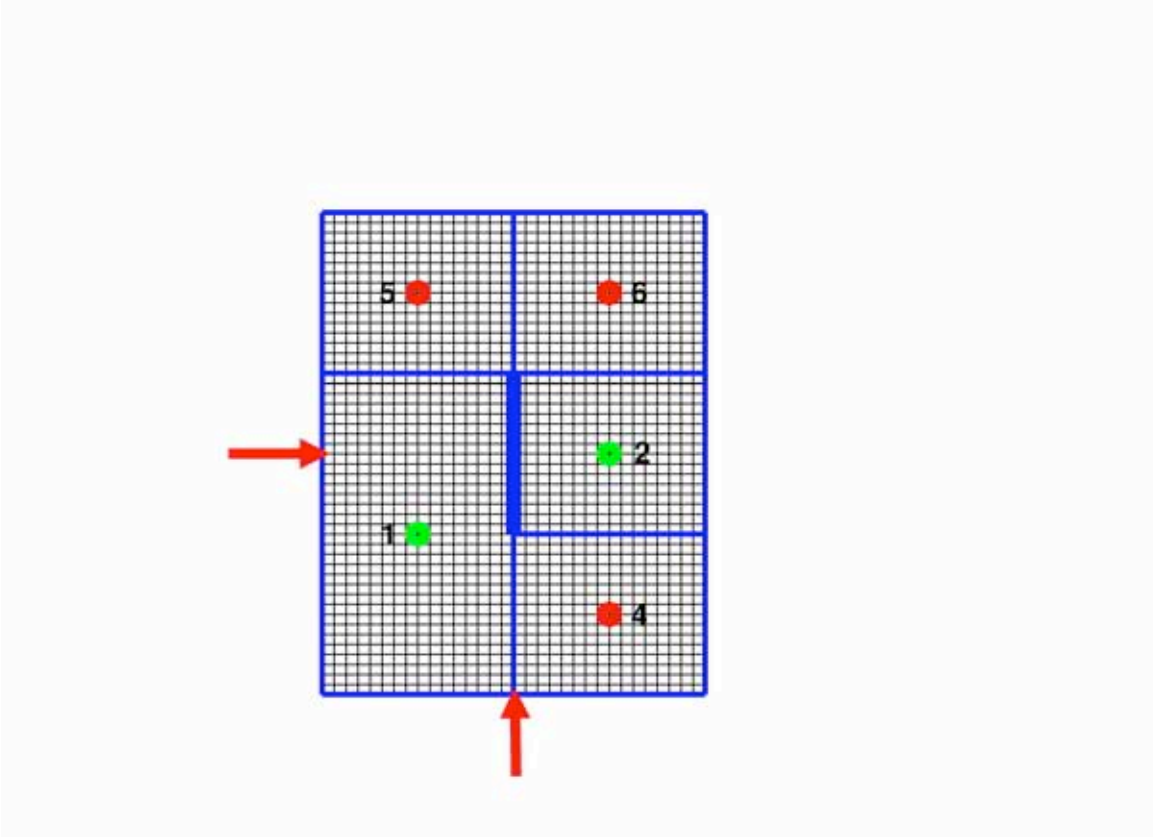

Figure 10.10: The red and green dots denote the centers of cells of a CCAR grid that are used to construct a VCMP approximation to the flux across a face, indicated by the thick blue line. The green dots correspond to pressure values whose weights that are to be dominant; the red dots correspond to values whose weights are to be minimized. Points 2 and 4 are centers of hanging nodes that, together, share a face with the cell centered at point 1.

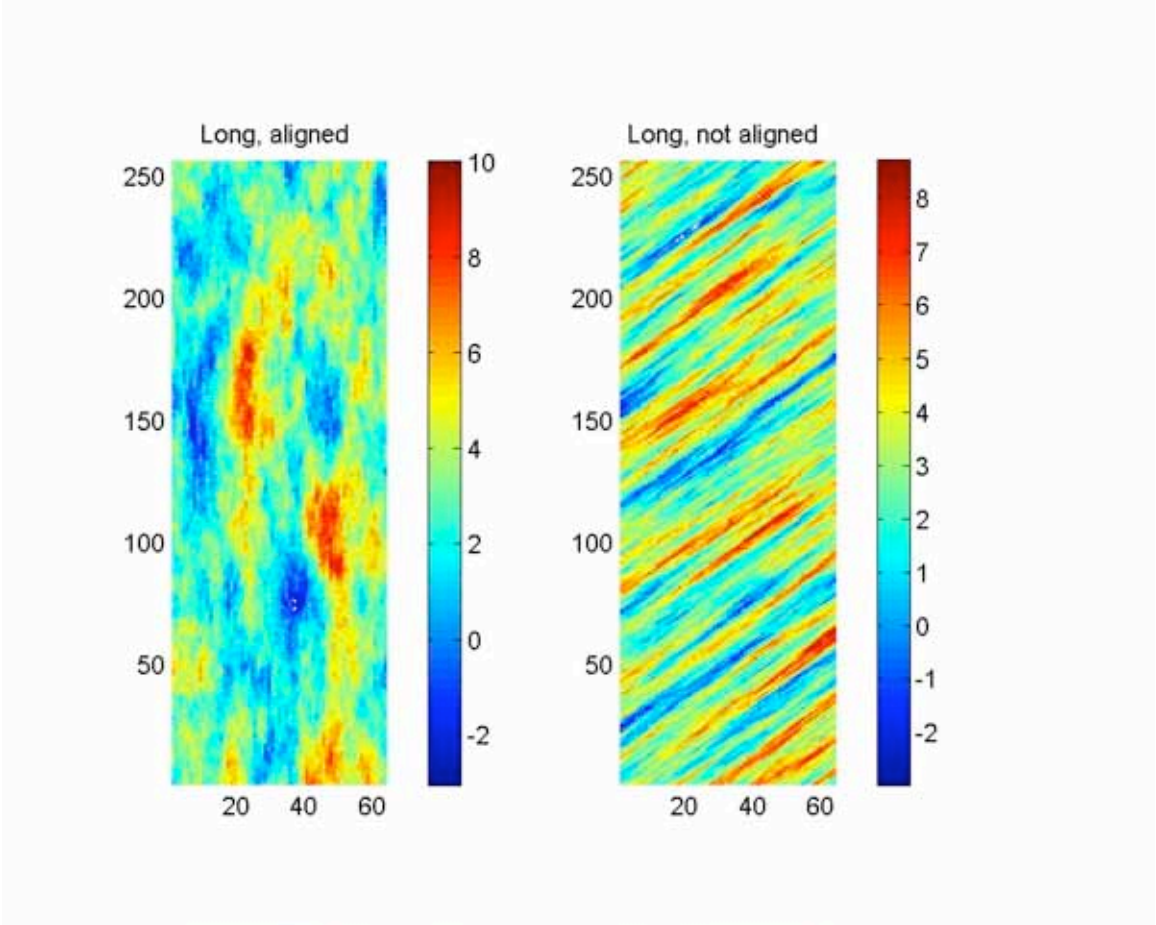

Figure 10.11: Fine-scale permeability fields used as test cases. Left: long correlation length, aligned to grid. Right: long correlation length, not aligned to grid. The grids contain 256×64 cells, with $\Delta x = \Delta y = 1$.

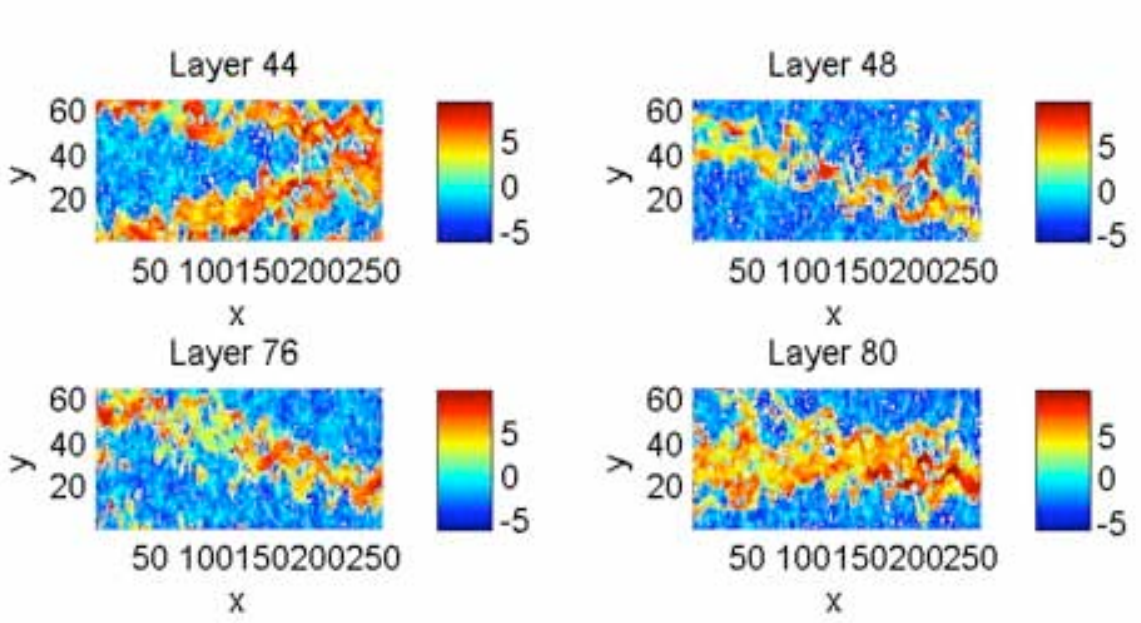

Figure 10.12: Fine-scale permeability fields for modified version of SPE10 layers 44, 49, 76 and 80. The grids contain 256×64 cells, with $\Delta x = \Delta y = 1$.

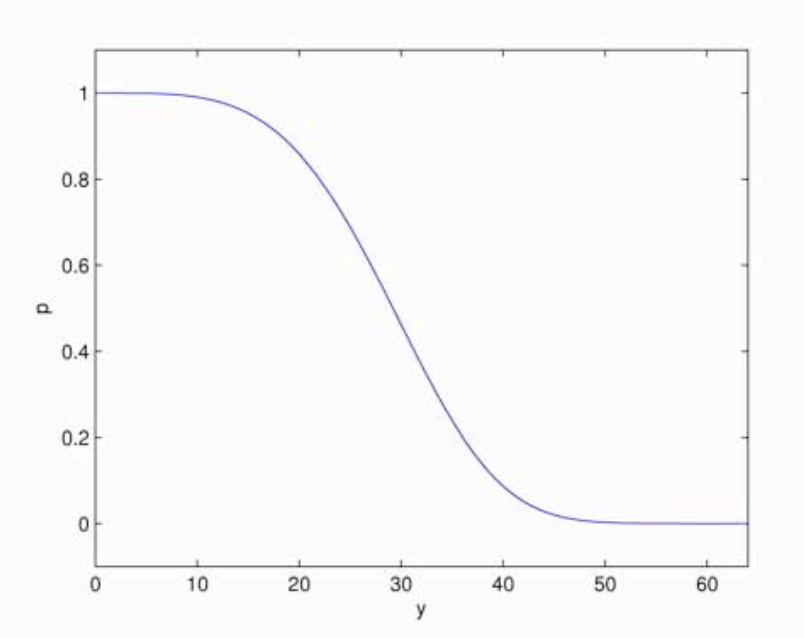

Figure 10.13: The pressure $p(0, y) = \exp(-(2y/L_y)^4)$ prescribed on the left boundary, for the set of boundary conditions in which flow emanates from the lower left corner of the domain.

11 Adaptive Methods for Temporal Grids

11.1 Operator Splitting

In-situ combustion is a multi-scale process. That is, physical subprocesses present act on a range of spatial and temporal scales. Solving for these processes in the standard fully coupled fashion therefore leads to numerical challenges: in order to simulate all processes accurately the fastest spatial and temporal scales dictate spatial and temporal grid size selection leading to extremely high computational loads. Therefore, it makes great sense to investigate the use of decoupling techniques that allow subprocesses to numerically evolve at their own scales with their own specialized solvers. Naturally, this leads to splitting errors that must be understood and controlled.

Numerical splitting techniques are essentially factored forms of numerical operators. They have been extensively used in other areas to reduce multi-dimensional problems into a sequence of simpler one-dimensional problems (e.g. ADI methods), to solve problems with more than one physical mechanism of transport with appropriate schemes for each mode independently, or to couple combustion kinetics to transport in open-air combustion, amongst others.

The development of a split-step method for the ISC process has the potential to address the following points:

- Separation of scales. The splitting allows the integration of chemical reactions to occur over a finer sampling than a single step for flow or transport, leading to a more efficient treatment of the sources of stiffness.
- Specialized numerical schemes to treat different processes. For example diffusion and advection terms appearing in the same equation can be treated independently and then combined.
- Projection methods which could allow one part of say pressure to be solved at each time-step followed by a correction of another part attributed to change in chemistry over the time-step.

First, we give a brief summary of factored methods, followed by a highlight of the forms specialized to ISC that will be studied, and finally, by a summary of the most important considerations to analyzed.

11.1.1 Operator Splitting Methods

Following [Hundsdorfer2003] we can illustrate the splitting idea by considering the semi-discrete PDE

$$\frac{d\mathbf{w}}{dt} = \mathbf{f}(\mathbf{w}_1, \mathbf{w}_2) = \mathbf{f}^A(\mathbf{w}_1, \mathbf{w}_2) + \mathbf{f}^R(\mathbf{w}_1, \mathbf{w}_2), \tag{11.1}$$

where $\mathbf{f}^A$ and $\mathbf{f}^R$ denote discrete advection and reaction operators, respectively, and $\mathbf{w}_1$ and $\mathbf{w}_2$ denote two different variables, for example pressure and saturation. Generally, operator splitting methods can be grouped in two categories: Implicit-Explicit (IMEX) methods and Additive Splitting (ADS) (or Locally One-Dimensional (LOD)) methods.

ADS methods can be applied to additively separable operators, such as advection and reaction in the ISC problem. The global integration step is decomposed into a number of substeps, where in each substep only one operator (i.e. one physical process), or a subset of operators, is considered. The solutions to each substep are not consistent with the full equation. A consistent approximation is only formed at the end of all substeps.

In addition to numerical errors introduced when integrating each substep in the scheme, a splitting method also introduces an approximation error, referred to as a *splitting error*, due to the splitting itself. A simple, first order accurate ADS splitting for (11.1) is

$$\begin{aligned} \frac{d\mathbf{w}^{(1)}}{dt} &= \mathbf{f}^A\left(\mathbf{w}_1^{(1)}, \mathbf{w}_2^{(1)}\right) \qquad t \in \left[t^n, t^{n+1}\right], \quad \mathbf{w}^{(1)}(t^n) = \mathbf{w}^n, \\ \frac{d\mathbf{w}^{(2)}}{dt} &= \mathbf{f}^R\left(\mathbf{w}_1^{(2)}, \mathbf{w}_2^{(2)}\right) \qquad t \in \left[t^n, t^{n+1}\right], \quad \mathbf{w}^{(2)}(t^n) = \mathbf{w}^{(1)}(t^{n+1}), \end{aligned} \tag{11.2}$$

giving $\mathbf{w}^{n+1} = \mathbf{w}^{(2)}(t^{n+1})$ as the approximation at t^{n+1}. A symmetrical, second order accurate splitting, known as *Strang splitting* [Strang1968], can be achieved in the following way

$$\begin{aligned} \frac{d\mathbf{w}^{(1)}}{dt} &= \mathbf{f}^A\left(\mathbf{w}_1^{(1)}, \mathbf{w}_2^{(1)}\right) \quad t \in \left[t^n, t^{n+12}\right], \quad \mathbf{w}^{(1)}(t^n) = \mathbf{w}^n, \\ \frac{d\mathbf{w}^{(2)}}{dt} &= \mathbf{f}^R\left(\mathbf{w}_1^{(2)}, \mathbf{w}_2^{(2)}\right) \quad t \in \left[t^n, t^{n+1}\right], \quad \mathbf{w}^{(2)}(t^n) = \mathbf{w}^{(1)}(t^{n+12}), \\ \frac{d\mathbf{w}^{(3)}}{dt} &= \mathbf{f}^A\left(\mathbf{w}_1^{(3)}, \mathbf{w}_2^{(3)}\right) \quad t \in \left[t^{n+12}, t^{n+1}\right], \quad \mathbf{w}^{(3)}(t^{n+12}) = \mathbf{w}^{(2)}(t^{n+1}), \end{aligned} \tag{11.3}$$

giving $\mathbf{w}^{n+1} = \mathbf{w}^{(3)}(t^{n+1})$ as the approximation at t^{n+1}. Here, the advection step is symmetrically split in half. This concept of introducing *fractional steps* at the splitting scheme level can be generalized to design higher order splitting methods

Hundsdorfer2003. Fractional steps can also be introduced at the substep level. ADS methods define a global timestep, which is the overall timestep of the splitting scheme, but subprocesses can be integrated with smaller timesteps, which are appropriate for each physical process. In addition to individual timestepping, the substeps in (6.25) and (6.26) can be integrated with specialized numerical methods. As such, ADS methods provide a flexible framework for designing tailored integration schemes for complex, multiscale processes.

In IMEX methods the overall equation is not decoupled in the same sense as for ADS methods. A single integration scheme is applied, but the operator variables are treated either explicitly or implicitly. The choice of explicit or implicit treatment can be made for the entire equation or for each individual operator as illustrated by

$$\mathbf{w}^{n+1} = \mathbf{w}^n + \Delta t \mathbf{f}\left(\mathbf{w}_1^{n+1}, \mathbf{w}_2^n\right)$$
$$\mathbf{w}^{n+1} = \mathbf{w}^n + \Delta t\left[\mathbf{f}^A\left(\mathbf{w}_1^{n+1}, \mathbf{w}_2^n\right) + \mathbf{f}^R\left(\mathbf{w}_1^n, \mathbf{w}_2^{n+1}\right)\right] \qquad (11.4)$$

The familiar IMPES (Implicit Pressure Explicit Saturation) scheme in reservoir simulation is an example of an IMEX method of type (6.27), where pressure is implicit and saturations are explicit. The explicit treatment of certain variables reduces computational cost, but introduces stability restrictions. Each stage of an IMEX scheme is a consistent approximation to the full equation. Although IMEX methods do not allow specialized treatment of subprocesses, the main advantage of the methods is that they can be applied to problems with non-separable operators, which is not the case for ADS methods. The flow/transport coupling in the ISC equations is non-separable, and IMEX methods are the only feasible alternative to full coupling.

11.1.2 Applications of operator splitting

Operator splitting techniques are used in a number of applications areas including atmospheric chemistry [Zlatev1995], open-air combustion [Smith2005] and mathematical biology [Hundsdorfer2003]. In the context of reactive porous media transport, splitting techniques are extensively used in subsurface contaminant transport problems [Arbogast1996]. The chemical components in contaminant transport models are often present in small amounts in an aqueous phase, but undergo a range of chemical processes including oxidation-reduction, dissolution-precipitation, adsorption-desorption, and ion exchange [Walter1994]. The transport equations usually account for advection, dispersion and molecular diffusion, and the chemical reactions are either equilibrium controlled or mixed equilibrium and kinetically controlled [Walter1994,Yeh1991]. The chemical processes typically have little or no influence on the flow, and the flow and transport problems are therefore naturally decoupled. The coupled transport and reaction problem is often solved using splitting techniques similar to those given above. The transport terms are spatially connected, but uncoupled between chemical components. Reaction terms, on the other hand, are spatially disconnected, but coupled through chemical components. A solution methodology based on decoupling the transport and

reaction terms therefore results in a transport step, where the (linear) component transport equations are solved individually, and a reaction step, where the chemical equilibrium equations are solved gridblock by gridblock. In the groundwater hydrology literature this approach is referred to as the Sequential Non-Iterative Approach (SNIA) [Walter1994]. To reduce splitting error, a Sequential Iterative Approach (SIA) is often used, where the transport and reaction steps are iterated until convergence [Yeh1991].
Operator splitting techniques for contaminant transport models have been implemented in several simulators including HYDROGEOCHEM [Yeh1991], THOUGHREACT [Xu2004] and PARSIM [Arbogast1996].

11.1.3 A splitting scheme for the ISC equations

Compared to subsurface contaminant transport models, the ISC equations are, in a sense, more tightly coupled. Most importantly, there is a strong feedback on flow from reactions. The heat released in reactions increases oil mobility, and the gases evolved act as a gas drive ahead of the combustion front. Solid deposition from cracking reactions reduces the rock porosity, which therefore affects flow. In general, volume changes due to reactions as well as phase behavior have a strong impact on flow and transport.

We propose a Strang splitting for the flow/transport and reaction terms in the ISC equations, where the flow/transport substep is split in half. The flow/transport substep is solved fully implicitly using a similar scheme as for the fully coupled problem. The reaction substep is solved independently in each gridblock using a subclass of Runge-Kutta methods denoted ESDIRK (Explicit Singly Diagonally Implicit Runge-Kutta), which are tailored for stiff chemical kinetics problems.

In principle, the volume and phase behavior constraint equations are active at all times, i.e. in both the flow/transport and reaction substeps. However, if the volume constraint is enforced in the reaction substep, and the pressure is allowed to increase as gases evolve from reactions, the resulting pressure field after the reaction substep will be non-physical with higher values in gridblocks, where significant reaction took place. Instead we propose a relaxed volume formulation, where the volume constraint is omitted and the pressure is kept constant in the reaction substep. As a result, the phase volumes do not match the pore volume after the reaction step. This volume inconsistency is corrected during the final half-step on flow/transport, such that a consistent solution is obtained at the end of the splitting step. The phase equilibrium constraints are active in both substeps.

Let $\mathbf{f}^t \in \mathrm{R}^{N_b(n_c+2)}$ denote the vector of discrete fluxes, and let $\mathbf{Q}^w \in \mathrm{R}^{N_b(n_c+2)}$ and $\mathbf{Q}^r \in \mathrm{R}^{N_b(n_c+1)}$ denote the vectors of source/sink terms and reaction terms, respectively, as given by

$$\mathbf{f}^t = \begin{bmatrix} \mathbf{f}_1^t \\ \vdots \\ \mathbf{f}_{N_b}^t \end{bmatrix}, \qquad \mathbf{f}_k^t = \begin{bmatrix} f_1^m \\ \vdots \\ f_{n_c}^m \\ f^h \\ 0 \end{bmatrix}_k, \tag{11.5}$$

$$\mathbf{Q}^w = \begin{bmatrix} \mathbf{Q}_1^w \\ \vdots \\ \mathbf{Q}_{N_b}^w \end{bmatrix}, \qquad \mathbf{Q}_k^w = \begin{bmatrix} Q_1^{m,well} \\ \vdots \\ Q_{n_c}^{m,well} \\ Q^{h,well} + Q^{h,htr} \\ V_p - \sum_j V^j \end{bmatrix}_k, \tag{11.6}$$

$$\mathbf{Q}^r = \begin{bmatrix} \mathbf{Q}_1^r \\ \vdots \\ \mathbf{Q}_{N_b}^r \end{bmatrix}, \qquad \mathbf{Q}_k^r = \begin{bmatrix} Q_1^{m,reac} \\ \vdots \\ Q_{n_c}^{m,reac} \\ Q^{h,reac} \end{bmatrix}_k, \tag{11.7}$$

where the residual of the volume constraint is included with the well terms. We define in addition the block diagonal matrix $\mathbf{M} \in \mathsf{R}^{N_b(nc+2)\times N_b(n_c+2)}$

$$\mathbf{M} = \begin{bmatrix} \mathbf{M}_1 & & \\ & \ddots & \\ & & \mathbf{M}_{N_b} \end{bmatrix}, \qquad \mathbf{M}_k = \begin{bmatrix} 1 & & & \\ & \ddots & & \\ & & 1 & \\ & & & 0 \end{bmatrix}_k. \tag{11.8}$$

The discrete fluxes in (11.5) are

$$f_i^m = \frac{1}{V_b}[\sum_j \Lambda_{i,k-12}^j T_{k-12} \Delta^- \Phi_k^j - \Lambda_{i,k+12}^j T_{k+12} \Delta^+ \Phi_k^j], \tag{11.9}$$

$$\begin{aligned} f^h = \frac{1}{V_b}[&\sum_j \Lambda_{k-12}^{j,h} T_{k-12} \Delta^- \Phi_k^j - \Lambda_{k+12}^{j,h} T_{k+12} \Delta^+ \Phi_k^j \\ &+ \frac{S_{k-12} k_{c,k-12}}{\Delta x_{k-}} \Delta^- T_k - \frac{S_{k+12} k_{c,k+12}}{\Delta x_{k+}} \Delta^+ T_k]. \end{aligned} \tag{11.10}$$

We note that $f_i^m = 0$ for $i \in \mathsf{I}^s$.

The splitting steps can now be formulated as

1. Solve flow/transport for $t \in \lfloor t^n, t^{n+12} \rfloor$

$$\mathbf{M}\frac{d\mathbf{u}^{(1)}}{dt} = \mathbf{f}^t(\mathbf{u}^{(1)}) + \mathbf{Q}^w(\mathbf{u}^{(1)}) \qquad \mathbf{u}^{(1)}(t^n) = \mathbf{u}^n. \tag{11.11}$$

2. Solve reactions for $t \in \lfloor t^n, t^{n+1} \rfloor$ and $k = 1, \ldots, N_b$

$$\frac{d\tilde{\mathbf{u}}_k^{(2)}}{dt} = \mathbf{Q}_k^r\left(\tilde{\mathbf{u}}_k^{(2)}\right) \qquad \tilde{\mathbf{u}}_k^{(2)}\left(t^n\right) = \tilde{\mathbf{u}}_k^{(1)}\left(t^{n+12}\right) \tag{11.12}$$

where $\tilde{\mathbf{u}}$ denotes the vector of primary variables, but without pressure, which is assumed constant during the reaction step.

3. Solve flow/transport for $t \in \left[t^{n+12}, t^{n+1}\right]$

$$\mathbf{M}\frac{d\mathbf{u}^{(3)}}{dt} = \mathbf{f}^t\left(\mathbf{u}^{(3)}\right) + \mathbf{Q}^w\left(\mathbf{u}^{(3)}\right) \quad \mathbf{u}^{(3)}\left(t^{n+12}\right) = \mathbf{u}^{(2)}\left(t^{n+1}\right) \tag{11.13}$$

which gives $\mathbf{u}^{n+1} = \mathbf{u}^{(3)}\left(t^{n+1}\right)$ as the final approximation.

A nontrivial issue with splitting methods relates to the treatment of boundary conditions. Boundary conditions are imposed on the whole system and conditions for the substeps are missing. In the proposed scheme, the well terms are included with the flow/transport step, but could, in principle, also have been included with the reaction step. The choice made here is based primarily on physical intuition regarding which operators in the equations are influenced the most by the well terms. Clearly, the treatment of boundary conditions requires further analysis, but we shall not pursue it here.

Although the Strang splitting is formally second order accurate, [Hundsdorfer2003] notes that possible order reductions can occur due to boundary conditions or stiff terms in the equations. Here, we apply Strang splitting primarily due to its symmetry, where the last half-step on flow/transport can be regarded as a corrector-step to correct for the volume inconsistency after the reaction step.

A natural next step in the development of a splitting scheme would be to reduce the implicitness by introducing an IMEX splitting for the flow/transport step, where pressure is solved implicitly, but component concentrations are updated explicitly (i.e. an IMPEC scheme -- Implicit Pressure Explicit Composition). In fact, Younis2006 have designed a hybrid IMEX-ADS scheme with this feature, but here we limit the discussion to the simpler ADS scheme.

11.2 Asynchronous Time Stepping

When simulating reservoir fluid processes on fast time scales, explicit time stepping methods may be the best compromise between accuracy and efficiency. Explicit time stepping is often applied in reservoir modeling in high resolution compositional simulation, for example, to obtain the desired accuracy in the transient solution. When applying an explicit time stepping method, the time-step chosen depends on a global stability criterion, which generally takes the form of a CFL constraint. The global time step is severely restricted due to the influence of only a few computational grid cells. This is certainly the case for In-Situ Combustion processes where the stiff kinetics sets very strict CFL limitsm and advective

processes operate at a much faster scale than diffusive processes, for example. We introduced operator splitting techniques in previous sections to alleviate this problem. Still, for each individual process, the chosen time step is much smaller that what would be required from local considerations of stability and accuracy for large portions of the domain. It would therefore be attractive to have a robust methodology for selecting time step sizes locally. For this purpose, we propose the Flux-conservative Asynchronous Time Integration method (FATI) described in this section.

Asynchronous time integration has its origin in Discrete Event Simulations (DES), but has recently been applied to continuous systems by Omelchenko and Karimabadi in [Omelchenko2006]. Omelchenko and Karimabadi choose their local time step sizes based on a predicted growth of the solution, potentially leading to violation of the local CFL constraint. Although in the test cases they investigated this is acceptable, we prefer to satisfy the local CFL constraint.

11.2.1 Asynchronous Time Stepping in 1D

To introduce the concepts behind asynchronous time stepping, it is insightful to first consider a simplified case. We take the linear advection equation

$$u_t + u_x = 0, \quad 0 \le x \le 1, \tag{11.14}$$

with periodic boundary conditions. We discretize the domain using N grid cells C_i, with $i = 1, \ldots, N$ of uniform width $h = 1/N$, and apply the standard first order upwind method. The continuous fluxes f at the cell interfaces of a C_i are denoted by $f_{i\pm 1/2}$. We use $F_{i\pm 1/2}$ to denote the current numerical interface fluxes. Here, $F_{i+1/2} = u_i$. As a first introduction to asynchronous time stepping, we assume that the all $C_i, i \le k$ are stepped forward in time with a time step size of dt, whereas cells $C_i, i > k$ are integrated in time with a time step size of dt/p with p integer. We require dt to satisfy the CFL condition, that is $\frac{dt}{dx} \le 1$.

We move the solution forward from time t^n to $t^{n+1} = t^n + dt$. For $i < k$, we have

$$u_i^{n+1} = u_i^n - \frac{dt}{h}\left(F_{i+1/2} - F_{i-1/2}\right) = u_i^n - \frac{dt}{h}\left(u_i^n - u_{i-1}^n\right) \tag{11.15}$$

For $i > k$, we integrate using a time step of dt/p. Intermediate values $u_i^{n+m/p}$ are generated by

$$u_i^{n+m/p} = u_i^{n+(m-1)/p} - \frac{dt}{ph}\left(F_{i+1/2} - F_{i-1/2}\right), \quad m = 1, \ldots, p,$$

with $F_{i+1/2} = u_i^{n+(m-1)/p}$. We note that the numerical diffusion introduced by the first order upwind scheme for a time step dt/p domain is equal to $\frac{1}{2}(h - dt/p)u_{xx}$. It is therefore smallest for $p = 1$ and grows with p.

To allow the given discretization in cell C_{k+1} at the first intermediate time $t^{n+1/p}$, the flux at the $k+1/2$ interface at this intermediate time step is required. This flux is equal to $u_k^{n+1/p}$, and it is computed in the synchronization step of the algorithm. We get

$$u_k^{n+1/p} = u_k^n - \frac{dt}{pdx} D_- u_k^n, \tag{11.16}$$

where we introduced the notation $D_- u_k = u_k - u_{k-1}$. To ensure conservation and to ensure that $u_{k+1}^{n+2/p}$ can be computed in the next intermediate time step, the algorithm now changes the flux at the interface $k+1/2$ to $F_{k+1/2} = u_k^{n+1/p}$. In this example, the flux $F_{k-1/2}$ at the left boundary of cell C_k is not changed and still equal to u_{k-1}^n. The second intermediate solution is now constructed as

$$u_k^{n+2/p} = u_k^{n+1/p} - \frac{dt}{pdx}(u_k^{n+1/p} - u_{k-1}^n). \tag{11.17}$$

This is equivalent to

$$u_k^{n+2/p} = u_k^n - \frac{2dt}{pdx} D_- u_k^n + \left(\frac{dt}{pdx}\right)^2 D_- u_k^n. \tag{11.18}$$

The flux $F_{k+1/2}$ is again updated, and this procedure is repeated for the remaining intermediate stages. Finally, we get

$$u_k^{n+1} = u_k^n - \frac{pdt}{pdx} D_- u_k^n + \left(\frac{dt}{pdx}\right)^2 D_- u_k^n - \left(\frac{dt}{pdx}\right)^3 D_- u_k^n + \ldots + \left(-\frac{dt}{pdx}\right)^p D_- u_k^n,$$

or

$$u_k^{n+1} = u_k^n - \frac{dt}{dx} D_- u_k^n + \left(\sum_{q=2}^{p} \left(-\frac{dt}{pdx}\right)^q\right) D_- u_k^n. \tag{11.19}$$

The above expression can be interpreted as a numerical update for cell C_k in the conservative form

$$u_k^{n+1} = u_k^n - \frac{dt}{dx}\left(G_{k+1/2} - G_{k-1/2}\right), \tag{11.20}$$

with $G_{k-1/2} = u_{k-1}^n$, as the flux at interface $k-1/2$ did not change, and $G_{k+1/2} = (1-\beta)u_k^n + \beta u_{k-1}^n$. Here, $\beta = \frac{1}{p}\left(\frac{\alpha}{1+\alpha}\right)\left(1-(-\alpha)^{p-1}\right)$, and $\alpha = \frac{dt}{pdx}$.

It is clear, that this discretization is inconsistent for $p \neq 1$. For $p = 1$, we get $\beta = 0$ as expected. For $p \geq 2$, we find that $0 < \beta < 1$. It has a maximum of 0.25 for $p = 2$ and

goes to 0 as p increases. We note however that as p increases, the inconsistency propagates further downwind of the interface: it starts to effect the cells C_{k+1} through C_{k+p-1} at, respectively, the time steps $t^{n+2/p}$ through t^{n+1}. Therefore, we do not expect the comprised effect of this inconsistency to be strongly dependent on p. Numerical solutions shown below also verify this. This behavior is very similar to that observed in [Riaz2008], where interfaces in an Adaptive Implicit Method are analyzed. Despite the inconsistency, the scheme remains monotone: for $\frac{dt}{dx} \le 1$, we get $\frac{\partial u_k^{n+1}}{u_k^n} = 1 - \frac{dt}{dx}(1-\beta) \ge 0$ and $\frac{\partial u_k^{n+1}}{u_{k-1}^n} = (1-\beta)\frac{dt}{dx} \ge 0$, as desired.

In the above description, we assumed that the flux $F_{k-1/2}$ is not changed throughout the time step dt, for simplification. In practice, we do change this flux at an intermediate time $t^{n+m/p}$ if $|u_k^{n+m/p} - u_k^n|$ computed as above exceeds the expected solution change $\frac{dt}{dx}|F_{k+1/2}^n - F_{k-1/2}^n|$, computed from fluxes at time t^n. This may occur, for example, in regions with high rates of change. In that case, we accept the updated solution $u_k^{n+m/p}$ and recursively synchronize the neighboring cell C_{k-1}. The net effect is simply that the interface location is shifted one cell per recursive call.

The simplified algorithm can be generalized. We allow each grid cell C_i to have its own time step size dt_i. In our implementation, the dt_i is determined by a local CFL constraint. At initialization, we compute dt_i for each cell C_i, and store its value as the next update time for the cell. We now select the cell, or cells, with the minimal update time, and move the solution in this cell forward. We update the fluxes at its cell interfaces in a synchronization step, as we did above for the simpler case, to ensure conservation. Synchronized cells are in turn synchronized recursively with their neighbors if needed, that is, if the change of the solution in a synchronized cell exceeds that predicted by the interface fluxes of this cell at its previous update time. For each fully updated cell, we compute a new update time. We now start the algorithm again by selecting the next cell(s) with the minimal update time. For a more thorough discussion and implementation details, see [Omelchenko2006].

11.2.2 Asynchronous Time Stepping For Reservoir Simulation in 2D

We will study incompressible Darcy flow in an oil reservoir coupled to simple models for immiscible water injection and miscible gas injection. In the absence of gravity, the multi-phase extension of Darcy's law represents the total fluid flux by $v = -\lambda_T k \nabla p$ where λ_T is the total fluid mobility, k is the permeability of the porous medium, and p is the pressure. This expression leads to the pressure equation,

$$\nabla \cdot (\lambda_T k \nabla p) = q_w, \tag{11.21}$$

where q_w represents sources and sinks due to wells. No flow conditions are applied at the boundaries of the 2D domain. For both the miscible and immiscible models, advective transport can be modeled using the saturation equation,

$$\frac{\partial S}{\partial t} + \nabla \cdot (vF(S)) = q_w S_w. \tag{11.22}$$

Here, S denotes the saturation (volume fraction) of the injected fluid. For injection wells $S_w = S_{inj} = 1$ and the saturation $S_w = S_{prod}$ varies in time for producing wells. $F(S)$ is the fractional flow which for the immiscible model is given by,

$$F(S) = \frac{S^2}{\lambda_T (1 - S_{or})^2} = \frac{S^2}{S^2 + \frac{1}{M}(1 - S - S_{or})^2}. \tag{11.23}$$

We use a residual oil saturation of $S_{or} = 0.2$ and an end point viscosity ratio $M = \frac{\mu_{oil}}{\mu_{inj}} = 4$. For the miscible injection problem we simply use $F(S) = S$ and the total mobility is given by the quarter-power mixing rule,

$$\lambda_T(S) = ((1 - S) + M^{0.25} S)^4, \tag{11.24}$$

where in this case $M = 8$.

We solve the coupled system for flow and transport with an IMPES (implicit pressure, explicit saturation) strategy using a general polygonal grid of control volumes. Given the initial mobility distribution, the pressure equation is discretized using the MPFA L-method [Aavatsmark2008] and is solved for the cell centered pressures and the total fluid flux through each cell face. While keeping the flux field fixed, the saturation equation is discretized in space using single-point upwinding and integrated forward in time for a predetermined global time step. This global time step is chosen so that 95 percent of the grid cells satisfy the CFL constraint. Generally, this global time step will be much larger (10-20 times) than the time step required to satisfy the CFL condition in every cell.

We use the FATI algorithm described in the previous section to perform the integration over a global time step. Hence, at least five percent of the cells are required to take multiple local time steps, and due to the synchronization procedure multiple time steps are also taken in other neighboring cells. For the sake of comparison we also use a synchronous forward Euler time integration. In order to satisfy CFL constraints, this method requires several local time steps to be taken within a single global time step. After advancing saturations to the new time level, the mobility field is updated and the entire process is repeated until the end time for the simulation is reached.

For our two examples we have chosen to use a Cartesian grid with nested local refinements around the wells. It is important to note that the simulation algorithm does not contain explicit information about the three refinement levels used in the

grid. This allows the time integration and the spatial discretizations to be treated separately, which greatly simplifies the implementation as compared to Cartesian adaptive methods that require specialized treatment at adaptive grid interfaces. For simplicity, permeability is taken to be a homogeneous scalar. Wells are modeled by fixing the pressure in the designated grid cells so that a unit pressure drop is imposed between injectors and producers. All simulations were run until one pore volume had been injected.

11.3 Results

11.3.1 Operator Splitting

For a fixed grid with 100 cells we simulate a dry combustion experiment using the minimal reaction model. We compare solutions obtained from the fully coupled scheme (FIM) and the splitting scheme. To allow comparison we then continue with a fixed (global) timestep of $\Delta t = [0.05] hrs$ when a stable combustion front is established. Thus, the time step used in the FIM scheme, from this point onwards, corresponds to the global timestep in the splitting scheme. The reaction substep is integrated using ESDIRK34 (see Chapter 7) with absolute and relative error tolerances set to 10^{-4} and 10^{-3}, respectively. The solutions from the FIM and splitting schemes are compared to a solution obtained on the same grid with the FIM scheme, but with timesteps of $\Delta t = [0.001] hrs$.

Figure 11.1 shows temperature, oil saturation, and coke concentration profiles for the FIM scheme, the splitting scheme, and the reference solution. The bottom plot in Figure 11.1 shows the number of fractional reaction steps required in the splitting scheme in each cell to meet the integration tolerances. We see that the ESDIRK solver adapts to the reaction dynamics in each cell. Only 5-10% of the cells near the combustion front require more than one integration step. The computational effort is thereby localized, which is one of the main advantages of a splitting based method.

We have verified that the reference solution in Figure 11.1 (solid line) is temporally converged. By comparing the solution profiles we observe that the splitting based solution is slightly more accurate than the FIM solution. The splitting scheme uses two fully implicit half-steps on flow/transport and is, hence, more expensive than the FIM scheme, although fewer nonlinear iterations are required in each half-step compared to a full step in the FIM scheme. We note that no effort has been made to optimize the splitting scheme with respect to choice of order of the ESDIRK method and error tolerances for the reaction integration.

11.3.2 Linear Advection in 1D

For illustration, we numerically solve the linear advection equation (11.1) for $p = 2$ and $p = 8$ on a periodic domain. We take a Gaussian bump as initial condition, and select two interfaces mid-way each slope. At the first interface, we transition from a time step dt/p to a time step dt, and at the second interface from dt back to dt/p. We take $dt/dx = 0.95$ and move the interface along with the solution in time, so that the same points on the solution graph are hit by the inconsistency over and over again. Figure 11.2 depicts the solution after one revolution for $N = 200$, for both refinement factors. The shape of the solution can be easily understood from the above discussion. On the right slope, the numerical flux at the interface is too large, and the solution decreases too fast ahead of the interface. The opposite occurs at the left interface. The figure also shows the strong damping in the dt/p region. Although, as we remarked above, the inconsistency itself goes down with p, the region of influence of this inconsistency grows. Indeed, we find that the solution is degraded in a similar way for $p = 2$ and $p = 8$.

Although the inconsistency clearly degrades the solution accuracy in the above example, we are not concerned. The example is one that was fabricated specifically to exaggerate the errors introduced. In practice, the inconsistency seems to have little discernible effect. This becomes immediately apparent when the location of the interface is perturbed so that errors are not accumulated as before. In Figure 11.3 we show the solution after one revolutions for $N = 200$ for $p = 2$, where at each time step $t^n = n * dt$, the interface is randomly shifted by $0, +1$ or -1 grid cells. Clearly, the effect of the inconsistency is significantly reduced. Figure 11.3 shows grid convergence results for the original and perturbed interface problem for $p = 2$. We see that with a moving interface the order of accuracy is sub-linear. In the perturbed case, we restore the first order accuracy. Similar behavior is observed for $p = 8$. This behavior, of large errors when an interface moves with the solution and negligible errors when it does not, is similar to the behavior due to lower order boundary closures, see [Gustafsson2008].

11.3.3 Results for Two-dimensional Reservoirs

Figure 11.5 shows the saturation distributions for the immiscible and miscible injection problems as computed by the FATI algorithm. The results for the synchronous scheme are visually identical to the asynchronous results and have

therefore been omitted. The performance of the FATI algorithm differs slightly for the two problems due to the form of the transport equations. The miscible transport problem involves linear advection. Hence, the local time step sizes are chosen only with regard to spatial variations in the total flux field. Local time steps for the immiscible problem also must take into account spatial variations in $\frac{dF}{dS}$.

Our implementation is not suitable for making direct comparisons of timings. However, to give an initial indication of the potential for speed-up with the FATI algorithm we report the total number of cell saturation updates and the total number of flux evaluations for the synchronous and asynchronous versions. For both methods we only count an update or an evaluation when an actual change to a saturation or flux is recorded.

For the immiscible injection problem the FATI algorithm required a total of 2.3 million cell updates and 4.1 million flux evaluations. For comparison the synchronous algorithm requires 39.9 million cell updates and 83.4 million flux evaluations. Hence the savings with FATI are a factor of 17 for cell updates and 20 for flux evaluations. The miscible results yield 1.9 million updates and 4.0 million evaluations with the FATI algorithm. These numbers represent savings factors of 12 and 13 for the cell updates and flux evaluations, respectively. Its important to note that the reduction in the number of cell updates and flux evaluations is not directly indicative of the savings in computation time as the execution of these tasks with FATI will require some additional overhead as compared to a synchronous implementation. A more thorough investigation of computational efficiency will be conducted in future work.

11.4 Discussion

The suggested operator splitting technique has so far only been tested on a simplified splitting step: separate integration of the kinetics. Applying the full splitting as discussed here requires a good understanding, for each ISC process, of characteristic length and time scales, and hence a dimensional analysis. This work is in progress.

We have performed an initial analysis of a flux-conservative asynchronous time integration scheme (FATI) and investigated the application of this scheme to 2D reservoir simulation problems. Our analysis revealed that when applied to linear advection in 1D the FATI algorithm is monotone despite being formally inconsistent. In a contrived example it was shown that errors due to the inconsistency can accumulate and lead to a significant deterioration of solution accuracy. However, for more realistic problems, these errors remained small and did not influence the observed order of accuracy. Future investigations will determine if the inconsistency can be removed or if the synchronization procedure can be modified to avoid introducing errors in critical regions.

2D examples of immiscible and miscible displacements revealed that saturation distributions computed with FATI were nearly indistinguishable from those computed using synchronous integration. However, the FATI algorithm allowed more than an order of magnitude reduction in the number of cell updates and flux evaluations for the transport solve.

Overall our results indicate that FATI is a promising simulation strategy for reservoir simulation. However, a detailed study of the true computational efficiency of the algorithm remains to be performed for practical simulations. Although we achieved an order of magnitude reduction in cell updates and flux evaluations for 2D problems, the ultimate savings in computational time for the transport solve may be less due to differences in memory accesses and overheads associated with managing the queue for cell updates. However, we expect that the advantages of FATI will increase significantly for heterogeneous 3D problems and that the overall strategy will be highly efficient for many simulations.

11.5 Summary

To improve the efficiency of explicit, as well as implicit, temporal integration schemes we suggest the use of operator splitting for the ISC equations. The ISC equations are highly suitable candidates as the process they model consists of subprocesses at various spatial and temporal scales. We tested a first splitting method in which the kinetics was split from the other processes. The splitting scheme is more accurate, but computationally more expensive on a per time step basis than the FIM scheme. The maximum global splitting step is limited by convergence in the flow/ transport substeps. It is comparable in length to the maximum FIM step. The flow/transport substeps dominate the computing time. Approximately 10% of the time is spent on reaction integration.
We developed an asynchronous time stepping method for reservoir modeling, that shows great promise. We are in the process of implementing this in compositional and thermal simulators.

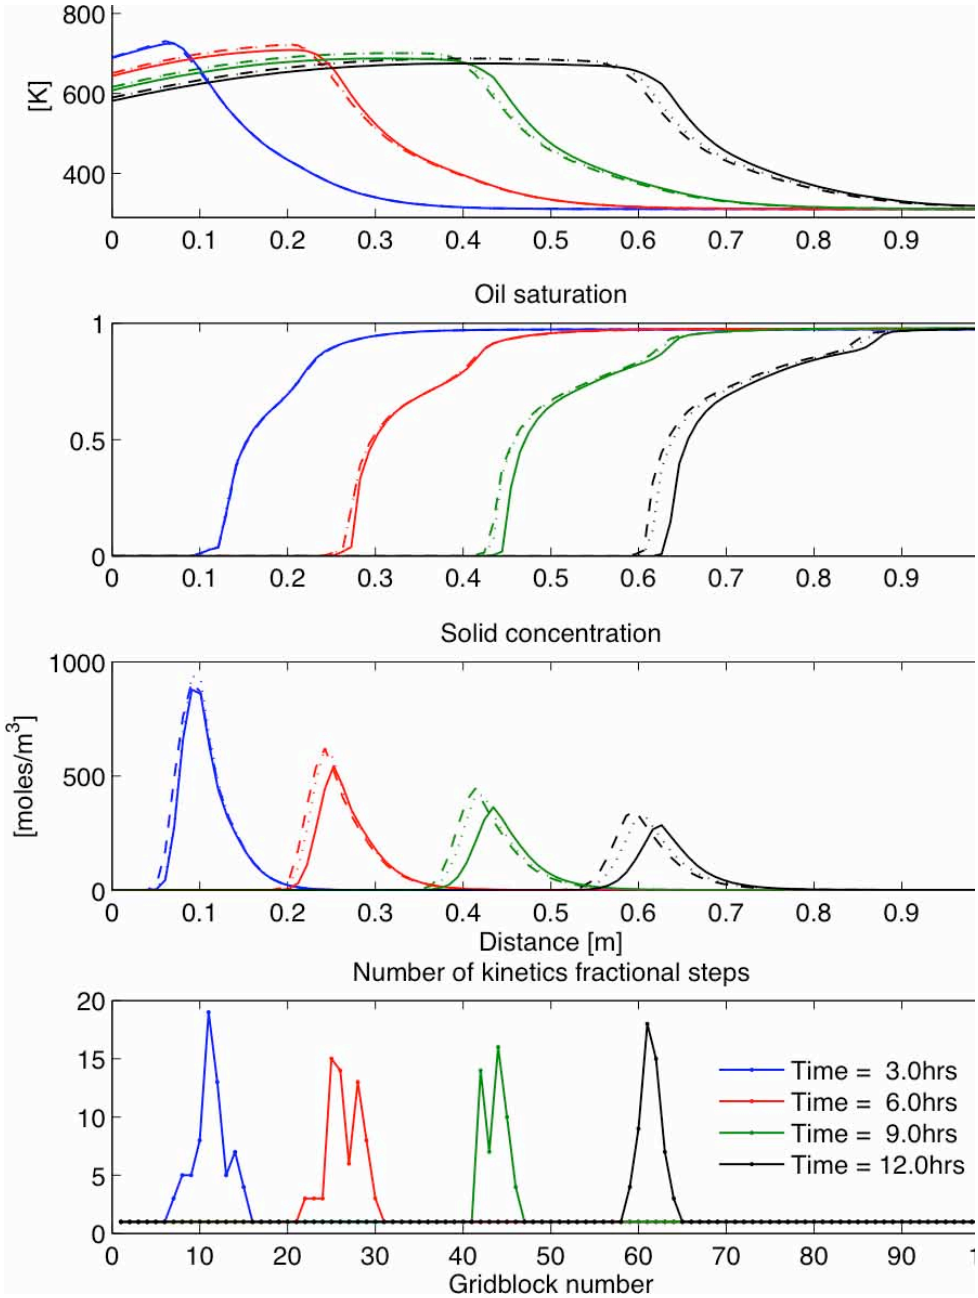

Figure 11.1: Temperature, oil saturation, and solid coke concentration profiles. Comparison is made between the FIM scheme (dashed line), the Strang splitting scheme (dotted line), and a temporally converged FIM solution (solid line). The bottom plot shows the number of fractional steps on kinetics for each cell in the grid.

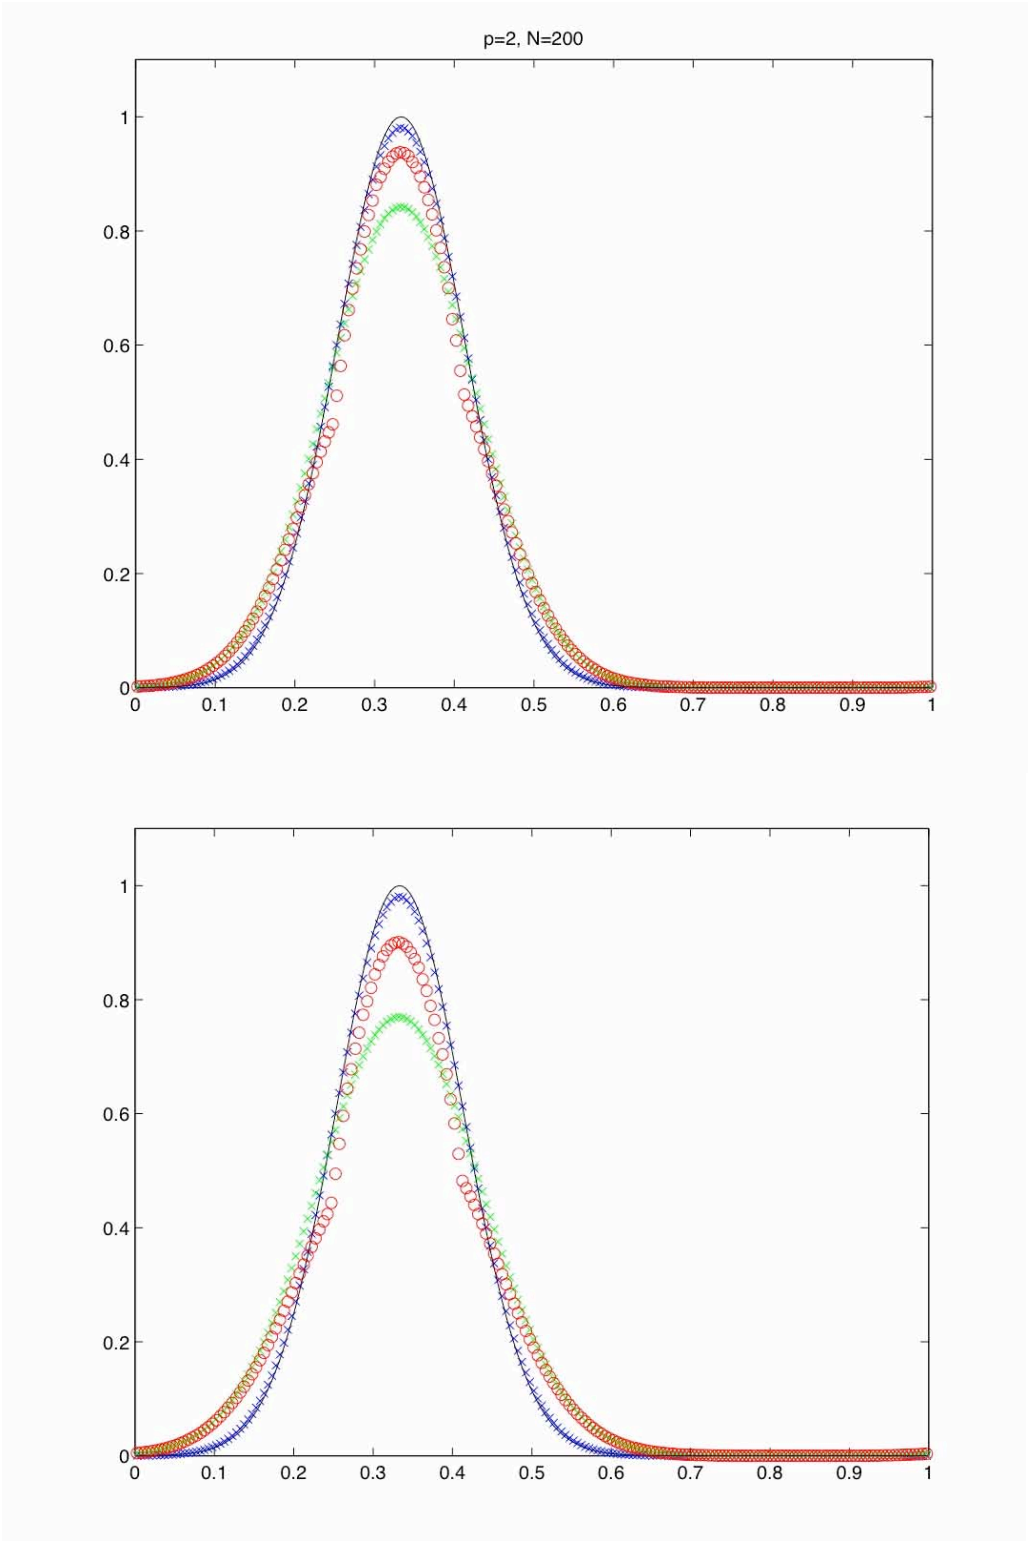

Figure 11.2: The solution to the linear advection equation $u_t + u_x = 0$ with time step dt in between the moving interfaces and dt/p outside. Figure a shows the solution for $p = 2$ and b for $p = 8$, for $N = 200$ after one full revolution in the periodic domain.

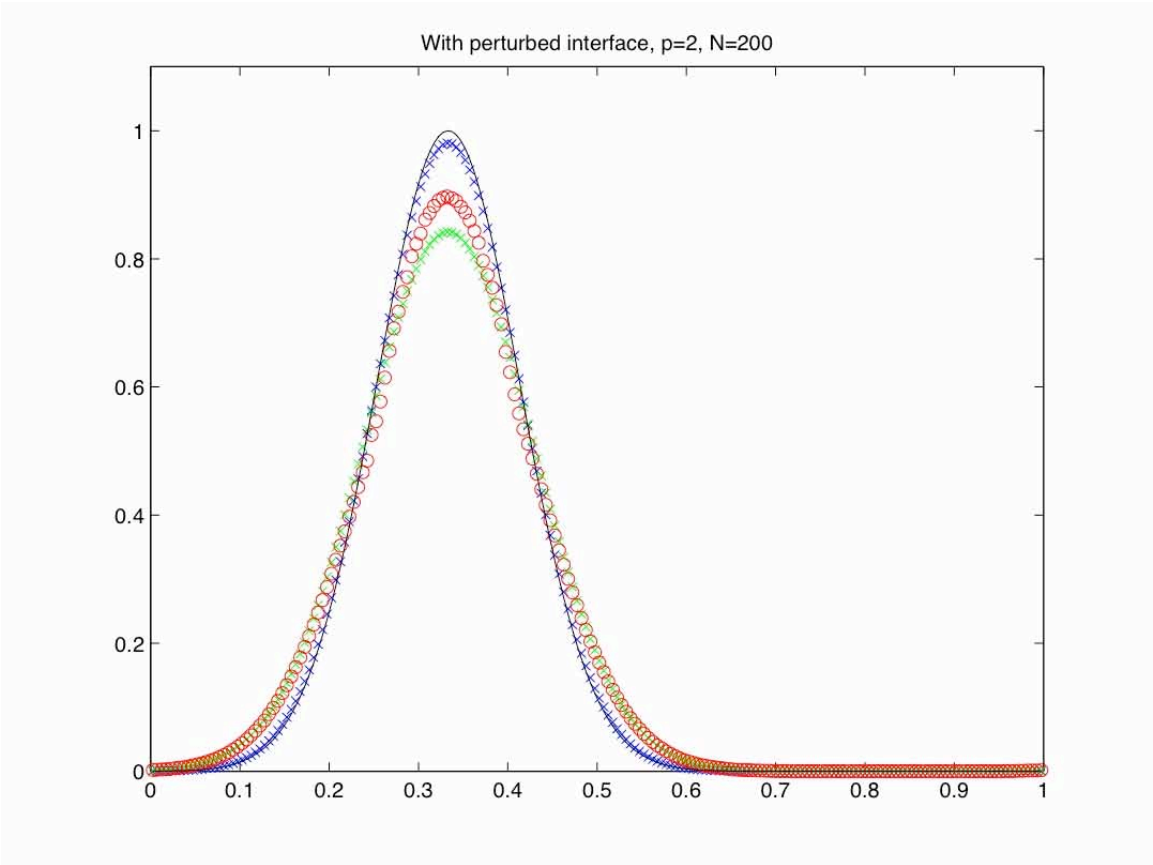

Figure 11.3: The solution to the linear advection equation $u_t + u_x = 0$ with time step dt in between the moving interfaces and $dt/2$ outside for $N = 200$ after one revolution in the periodic domain. The interfaces are perturbed randomly by one grid point at each time step.

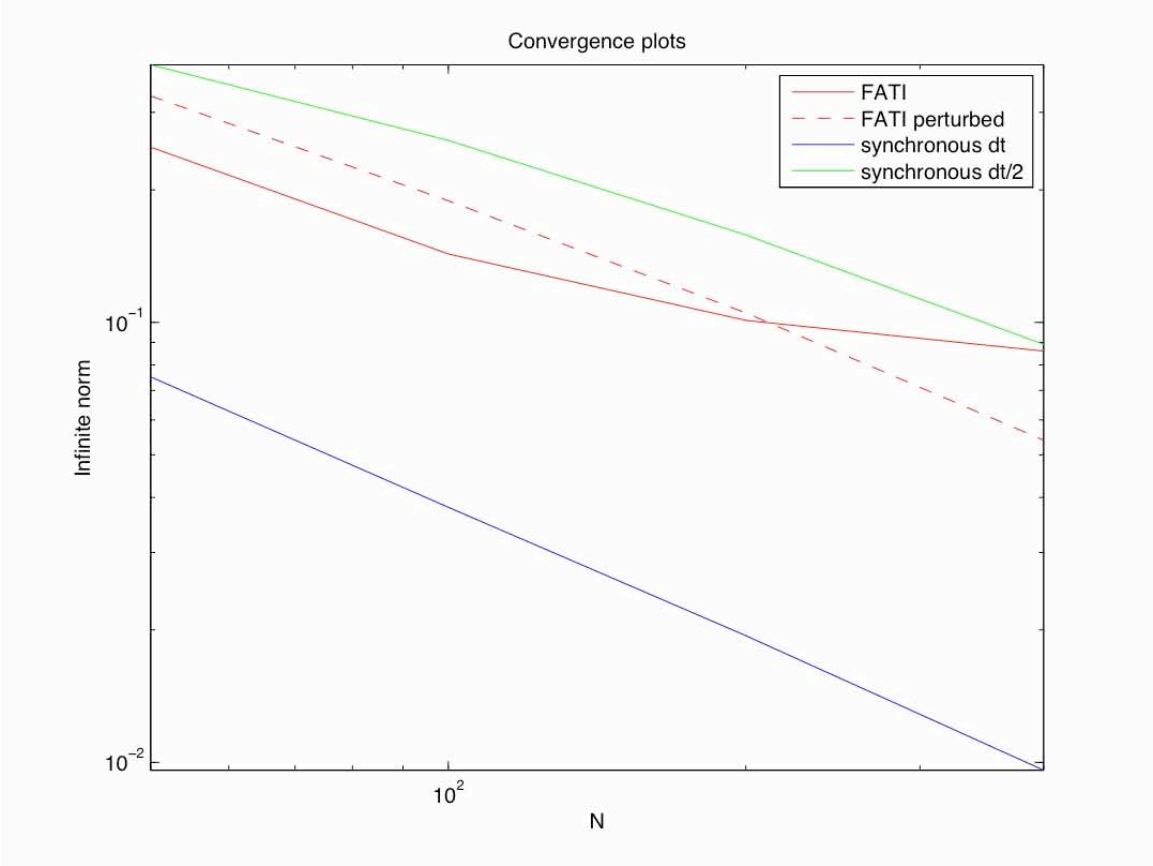

Figure 11.4: Convergence behavior for $p = 2$ on a log-log scale. The blue and green curves give the results for synchronous time stepping with a constant dt and dt/p, respectively. The solid red curve shows convergence behavior with a moving interface, and the dashed curve with a perturbed moving interface

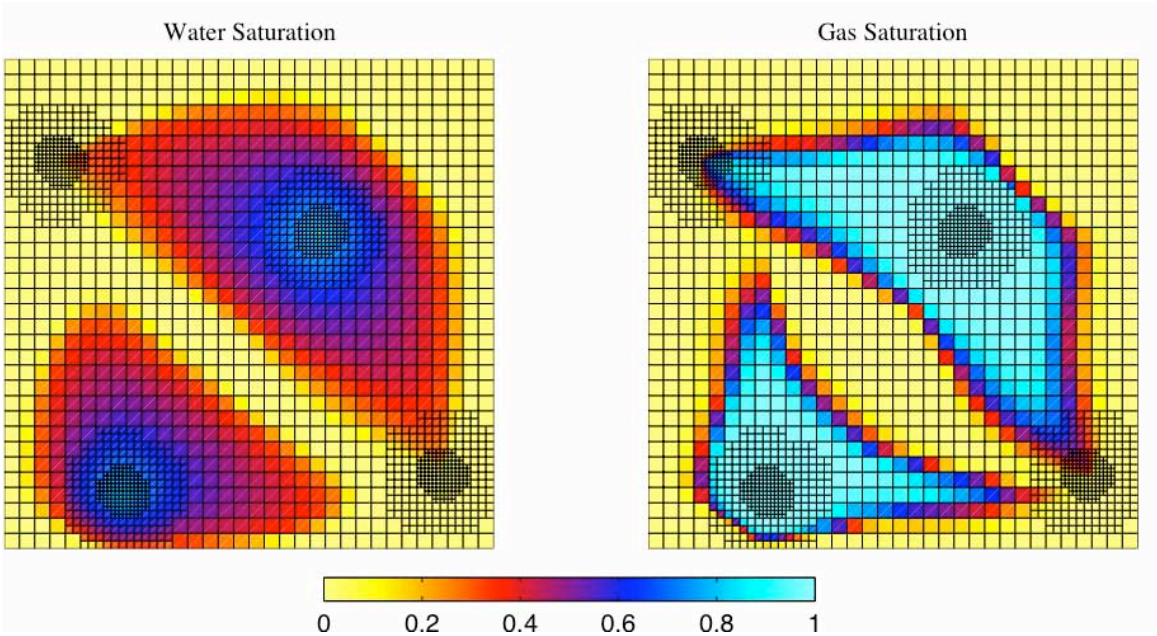

Figure 11.5: Saturation distributions computed using the FATI algorithm. The solution on the left shows the water saturation for the immiscible displacement after 0.25 pore volumes have been injected. The solution on the right shows the gas saturation after 0.4 pore volumes of miscible gas injection. Solutions computed using the synchronous version of the algorithm appear identical and hence have been omitted.

References

Aavatsmark I, Barkve T, Mannseth T. Discretization on Non-Orthogonal, Quadrilateral Grids for Inhomogeneous, Anisotropic Media. Journal of Computational Physics, 127:2-14, 1996.

Aavatsmark I, Eigestad GT and Nordbotten JM. A compact MPFA method with improved robustness. Proceedings of the 10th European Conference on the Mathematics of Oil Recovery, 2006.

Abu-Khamsin, S. A.: The Reaction Kinetics of Fuel Formation, PhD dissertation, Stanford University, Stanford, CA (Oct 1984).

Acharya, U. K., Kent, A., Waite, M. W., Tankersley, T., Johansen, S., and Robertson, C., 2004. Subsurface Challenges in Reservoir Modeling for Hamaca Project, paper SPE 86971 presented at the SPE International Thermal Operations and Heavy Oil Symposium and Western Regional Meeting held in Bakersfield, California, U.S.A., 16-18 March.

Adegbesan, K. O., Donnelly, J. K., Moore, R. G. and Bennion, D. W. (1987), Low-temperature-oxidation kinetic parameters for in-situ combustion: Numerical simulation, Society of Petroleum Engineers. SPE 12004: Presented at the 58th Annual Technical Conference and Exhibition, San Francisco, CA, October 5-8.

Adegbesan, K.O., Donnelly, J.K., Moore, R.G., Bennion D.W., 1997. "Low Temperature-Oxidation Kinetic Parameters for In-Situ Combustion: Numerical Simulation", SPERE, November 1987.

Akahira, T. and Sunose, T. , 1971 . Res. Report, Chiba Inst. Technol., Vol.16, 22.

Akkutlu, I. Y. and Yortsos, Y. C. (2002), The effect of heterogeneity on in-situ combustion: The propagation of combustion fronts in layered porous media, Society of Petroleum Engineers. SPE 75128: Presented at the SPE/DOE Thirteenth Improved Oil Recovery Symposium,Tulsa, OK, 13-17 April.

Akkutlu, I. Y. and Yortsos, Y. C. (2003), `The dynamics of in-situ combustion fronts in porous media', Combustion and Flame 134, 229--247.

Akkutlu, I. Y. and Yortsos, Y. C. (2004), Steady-state propagation of in-situ combustion fronts with sequential reactions, Society of Petroleum Engineers. SPE 91957: Presented at the 2004 SPE International Petroleum Conference, Puebla, Mexico, 8-9 November.

Aldushin, A. P., Matkowsky, B. J. and Schult, D. A. (1997), `Buoyancy driven filtration combustion', Combustion Science and Technology 125, 283--349.

Alexander, R. (2003), `Design and implementation of DIRK integrators for stiff systems', Applied Numerical Mathematics 46, 1--17.

Ambastha, A.K., and Kumar, M., 1999. "New Insight into In-Situ Combustion Simulation for Heavy Oil Reservoirs", SPE 56543.

Arbogast, T., Bryant, S., Dawson, C., Saaf, F., Wang, C. and Wheeler, M. (1996), `Computational methods for multiphase flow and transport problems arising in subsurface contaminant remediation', Journal of Computational and Applied Mathematics 74, 19--32.

Aris, R. (1989), Elementary Chemical Reactor Analysis, Dover Publication, Inc., Mineola, New York.

Aziz, K. and Settari, A. (1979), Petroleum Reservoir Simulation, Elsevier Applied Science Publishers, London.

Baes, C. F., Mesmer, J. R. E., 1977. The Hydrolysis of Cations. Wiley: New York.

Bagci, S. 2005. Effect of Pressure on Combustion Kinetics of Heavy Oils, Energy Sources, 27:887 – 898.

Bauer, I., Bock, H. G., Körkel, S. and Schlöder, J. P. (2000), `Numerical methods for optimum experimental design in DAE systems', Journal of Computational and Applied Mathematics 120, 1--25.

Berger MJ, Colella P. Local adaptive mesh refinement for shock hydrodynamics. Journal of Computational Physics, 82:64-84, 1989.

Berger MJ, Oliger JE. Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations. Journal of Computational Physics, 53:484-512, 1984.

Bijl, H., Carpenter, M. H., Vatsa, V. N. and Kennedy, C. A. (2002), `Implicit time integration schemes for the unsteady compressible navier-stokes equations: Laminar flow', Journal of Computational Physics 179, 313--329.

Bourgeat A. Homogenized Behavior of Two-Phase Flows in Naturally Fractured Reservoirs with Uniform Fractures Distribution. Computational Methods in Applied Mechanics and Engineering, 47:205-16, 1984.

Bousaid, I.S. and Ramey, H.J. Jr.,1968. Oxidation of Crude Oil in Porous Media, Soc. of

Pet. Eng. J. (June 1968) 137-148.

Brigham, W.E. and Castanier, L. M.: “In-Situ Combustion”, Chapter 17, Petroleum Engineers Handbook, Society of Petroleum Engineers, 2008.

Burger, J.G. and Sahuquet, B., 1972. Chemical Aspects of In-Situ Combustion - Heat of Combustion and Kinetics, Society of Petroleum Engineers Journal. (Oct.) 410-422.
Burger,J., Sourieau, P.,Combanarous, M., 1985. Thermal Recovery, Editions Technip, Paris, 1985.

Burnham A. K. and Dinh L. N., 2007. A Comparison of Iso-Conversional and model-fitting Approaches to Kinetic Parameter Estimation and Application Predictions”, Journal of Thermal Analysis and Calorimetry, (July 2007), Vol. 89, Num 2, 479-490.

Butcher, J. C. and Chen, D. J. L. (2000), `A new type of singly-implicit runge-kutta methods', Applied Numerical Mathematics 34, 179--188.

Butler, R. M., McNab, G. S. and Lo, H. Y. (1981), `Theoretical studies on the gravity grainage of heavy oil during in-situ steam heating', Canadian Journal of Chemical Engineering 59(4), 455--460.

Cao, H. (2002), Development of techniques for general purpose simulators, PhD thesis, Stanford University, USA.

Castanier, L. M. and Brigham, W. E. (2004), In-situ combustion. Society of Petroleum Engineers Handbook.

Castanier L.M., Baena, C.J., Holt, R.J., Brigham W.E., Tavares, C., 1992. "In Situ Combustion with Metallic Additives" SPE 23708 presented at the Second Latin Petroleum Engineering Conference of the SPE, Caracas, Venezuela, March 8-11, 1992.

Castanier, L. M. and Brigham, W. E., 2003. Upgrading of Crude Oil via In Situ Combustion, Journal of Petroleum Science & Engineering 39, 125-136.

Castanier, L. M., and Kovscek, A. R.: " Heavy Oil Upgrading In–Situ via Solvent Injection and Combustion: A "New" Method" Proceedings of the EAGE 67th Conference & Exhibition — Madrid, Spain, 13 - 16 June 2005.

Chen Y, Durlofsky LJ, Gerritsen MG, Wen XH. A coupled local-global upscaling approach for simulating flow in highly heterogeneous formations. Advances in Water Resources, 26:1041-1060, 2003

Chen Y, Durlofsky LJ. Adaptive Local--Global Upscaling for General Flow Scenarios in Heterogeneous Formations. Transport in Porous Media 62:157-185, 2006.

Christensen, J. R., Darche, G., Déchelette, B., Ma, H. and Sammon, P. H. (2004), Applications of dynamic gridding to thermal simulations, Society of Petroleum Engineers. SPE 86969: Presented at the SPE International Thermal Operations and Heavy Oil Symposium, Bakersfield, CA., March 16-18.

Christie MA, Blunt MJ. Tenth SPE comparative solution project: a comparison of upscaling techniques. SPE Reservoir Evaluat Eng 2001;4:308-17.

Christofari, J., Castanier, L. M. and Kovscek, A. R. (2006), Laboratory investigation of the effect of solvent injection on in-situ combustion, Society of Petroleum Engineers. SPE 99752: Presented at 2006 SPE Symposium on Improved Oil Recovery, Tulsa, OK, April 22 - 26.

Clara, C., Durandeau, M., Quenault, G., Nguyen, T.H., 1999. "Laboratory Studies for Light-Oil Air Injection Projects: Potential Application in Handil Field" SPE 64272 presented at the SPE Asia Pacific Oil and Gas Conference and Exhibition, Jakarta. April 20-22.

Clara, C., Durandeau, M., Quenault, G. and Nguyen, T.-H. (2000), `Laboratory studies for light-oil air injection projects: Potential application in handil field', SPE Reservoir Evaluation and Engineering 3(3), 239--248.

CMG (2004), STARS, Advanced Process and Thermal Reservoir Simulator, 2004 edn.

Coats, K. H. (1980), `In-situ combustion model', Society of Petroleum Engineers Journal 269, 533--554.

Cristofari, J, 2006. Laboratory Investigation of the Effect of Solvent Injection on In-Situ Combustion for Viscous Oils, M.S. Thesis Stanford University. online: http://ekofisk.stanford.edu/pereports/web /default.htm

Crookston, R. B., Culham, W. E. and Chen, W. H. (1979), `A numerical simulation model for thermal recovery processes', Society of Petroleum Engineers Journal pp. 37--58.

Dabbous, M. K. and Fulton, P. F. (1974), `Low-temperature-oxidation reaction kinetics and effects on the in-situ combustion process', Society of Petroleum Engineers Journal 255, 253--262. SPE 4143.

Deutsch C, Journel AG. GSLIB: Geostatistical Software Library and User's Guide, 2nd. Edition, Oxford Press, 1998.

Donaldson, E.C., Chilingarian, G.V., and Yen, T.F: “Enhanced Oil Recovery, II – Processes and Operations, Developments in Petroleum Science. Amsterdam, the Netherlands: Elsevier Science Publishers,

Duff, I. S. and Reid, J. K. (1996), `The design of MA48: A code for the direct solution of sparse unsymmetric linear systems of equations', ACM Transactions on Mathematical Software 22(2), 187--226.

Durlofsky LJ. Numerical Calculation of Equivalent Grid Block Permeability Tensors for Heterogeneous Porous Media. Water Resources Research, 27:699-708, 1991.

Edwards MG. Elimination of Adaptive Grid Interface Errors in the Discrete Cell Centered Pressure Equation. Journal of Computational Physics, 126:356-372, 1996.

Edwards MG, Rogers CF. Finite Volume Discretization with Imposed Flux Continuity for the General Tensor Pressure Equation. Computational Geosciences, 2:259-290, 1998.

EIA-DOE (2007), International energy outlook 2007, Technical report, Energy Information Administration, U.S. Department of Energy.

Enright, W. H., Jackson, K. R., Nørsett, S. P. and Thomsen, P. G. (1986), `Interpolants for runge-kutta formulas', ACM Transactions on Mathematical Software 12(3), 193--218.

Farouq Ali, S. M. (1977), Multiphase, multidimensional simulation of in-situ combustion, Society of Petroleum Engineers. SPE 6896: Presented at the 52nd SPE Annual Fall Technical Conference and Exhibition, Denver, Colorado, Oct. 9-12.

Fassihi, M. R., Brigham, W. E. and Ramey Jr., H. J. (1984 a), `Reaction kinetics of in-situ combustion: Part 1 - observations', Society of Petroleum Engineers Journal pp. 399--407.

Fassihi, M. R., Brigham, W. E. and Ramey Jr., H. J. (1984 b), `Reaction kinetics of in-situ combustion: Part 2 - modelling', Society of Petroleum Engineers Journal pp. 408--416.

Fassihi, M.R.: Analysis of Fuel Oxidation in In-situ Combustion Oil Recovery, PhD dissertation, Stanford University, Stanford, CA (May 1981).

Flynn, J. H. and Wall, L. A., 1966. A Quick Direct Method for the Determination of Activation Energy from Thermogravimetric Data", Polym. Lett., 4 323.

Freitag, N. P. and Exelby, D. R. (2006), `A SARA-based model for simulating the pyrolysis reactions that occur in high-temperature EOR processes', Journal of Canadian Petroleum Technology 45(3), 38--44.

Freitag, N. P. and Verkoczy, B. (2005), `Low-temperature oxidation of oils in terms of

SARA fractions: Why simple reaction models don't work', Journal of Canadian Petroleum Technology 44(2), 54--61.

Friedly, J. C. (1991), `Extent of reaction in open systems with multiple heterogeneous reactions', AIChE Journal 37(5), 687--693.

Friedly, J. C. and Rubin, J. (1992), `Solute transport with multiple equilibrium-controlled or kinetically controlled chemical reactions', Water Resources Research 28(6), 1935--1953.

Friedman, H., L., 1964. Kinetics of Thermal Degradation, Journal of Polymer Science: Part C, 1964, No 6: 183-195.

Gates, C.F. and Ramey H.J., 1980. "A Method for Engineering In-Situ Combustion Oil Recovery Project" J. Petr. Tech. 285.

Gautier Y, Blunt MJ, Christie MA. Nested Gridding and Streamline-Based Simulation for Fast Reservoir Performance Prediction. SPE Reservoir Simulation Symposium, 1999, SPE 51931.

Gerritsen MG, Jessen K, Mallison BT, Lambers JV. A Fully Adaptive Streamline Framework for the Challenging Simulation of Gas-Injection Processes. SPE ATC, 2005, SPE 97270.

Gerritsen MG, Lambers JV. Solving Elliptic Equations in Heterogeneous Media using Cartesian Grid Methods with Anisotropic Adaptation. To be submitted to Journal of Computational Physics, 2007.

Gillham, T. H., Cerveny, B. W., Turek, E. A. and Yannimaras, D. V. (1997), Keys to increasing production via air injection in gulf coast light-oil reservoirs, Society of Petroleum Engineers. SPE 38848: Presented at the SPE Annual Technical Conference and Exhibition, San Antonio, TX, October 5-8.

Gondouin, M. and Fox III, J. M., . The Challenge of West Sak Heavy Oil: Analysis of an Innovative Approach", paper SPE 22077 Proceedings of the the International Arctic Technology Conference, Anchorage, AK

Grabowski, J. W., Vinsome, P. K., Lin, R. C., Behie, A. and Rubin, B. (1979), A fully implicit general purpose finite-difference thermal model for in-situ combustion and steam, Society of Petroleum Engineers. SPE 8396: Presented at the 54th Annual Fall Technical Conference and Exhibition of the SPE, Las Vegas, NV, September 23-26.

Greaves, M., Xia, T. X. and Ayasse, C. (2005), Underground upgrading of heavy oil using thai: ``toe-to-heel air injection'', Society of Petroleum Engineers. SPE 97728: Presented at 2005 SPE Thermal Operations and Heavy Oil Symposium, Calgary, Canada, November 1 - 3.

Grim, R. E., 1968. Clay Mineralogy, McGraw Hill: New York.

Gustafsson, K. (1992), Control of Error and Convergence in ODE Solvers, PhD thesis, Department of Automatic Control, Lund Institute of Technology.

Hairer, E., Nørsett, S. and Wanner, G. (1996), Solving Ordinary Differential Equations I, second revised edn, Springer.

Hairer, E. and Wanner, G. (1996), Solving Ordinary Differential Equations II, second revised edn, Springer.

Ham FE, Lien FS, Strong, AB. A Cartesian Grid Method with Transient Anisotropic Adaptation. Journal of Computational Physics, 179:469-494, 2002.

Hascoët, L. and Pascual, V. (2004), Tapenade 2.1 User's Guide, INRIA, France.

Haukås, J. (2006), Compositional reservoir simulation with emphasis on the IMPSAT formulation, PhD thesis, University of Bergen, Norway.

He C. Structured Flow-Based Gridding and Upscaling for Reservoir Simulation. PhD Thesis, Department of Petroleum Engineering, Stanford University.

He, B., 2004. The Effect of Metallic Salt Additives on In-Situ Combustion Performance," M.S. Report Stanford University. Online at http://ekofisk.stanford.edu/pereports/web /default.htm

He B., Chen, Q., Castanier L. and Kovscek A., 2005. "Improved in-situ Combustion Performance with Metallic Additive", paper SPE 93901, presented at the SPE Western regional meeting held in Irvine, CA.

Henson VE, Yang UM. BoomerAMG: a Parallel Algebraic Multigrid Solver and Preconditioner. Applied Numerical Mathematics, 41:155-77, 2002.

Hindmarsh, A. C. (1983), ODEPACK, a generalized collection of ODE solvers, in R. S. Stepleman, ed., `Scientific Computing (IMACS Transactions on Scientific Computing, Vol. 1)', North-Holland, Amsterdam, pp. 55--64.

Holden L, Nielsen BF. Global Upscaling of Permeability in Heterogeneous Reservoirs: the Output Least Squares (OTL) Method. Transport in Porous Media, 40:115-43, 2000.

Holt R. J. 1992.: "In Situ Combustion with Metallic Additives". M.S Report. Stanford University, 1992. Online at http://ekofisk.stanford.edu/pereports/web /default.htm.

Hornung R, Trangenstein J. Adaptive Mesh Refinement and Multilevel Iteration for Flow in Porous Media. Journal of Computational Physics, 136:522-545, 1997.
Hundsdorfer, W. and Verwer, J. G. (2003), Numerical Solution of Time-Dependent Advection-Diffusion-Reaction Equations, Springer Series in Computational Mathematics, Springer.

Islam, M. R., Chakma, A. and Farouq Ali, S. M. (1989), State-of-the-art of in-situ combustion modelling and operations, Society of Petroleum Engineers. SPE 18755: Presented at the SPE California Regional Meeting, Bakersfield, CA., April 5-7.

Jørgensen, J. B., Kristensen, M. R. and Thomsen, P. G. (2008), `A familiy of esdirk integration methods', Submitted to SIAM Journal on Scientific Computing .

Kennedy, C. A. and Carpenter, M. H. (2003), `Additive runge-kutta schemes for convection-diffusion-reaction equations', Applied Numerical Mathematics 44, 139--181.

Kippe V, Aarnes JE, Lie KA. A Comparison of Multiscale Methods for Elliptic Problems in Porous Media Flow. To be published in Computational Geosciences, Special issue on Multiscale Methods, 2007.

Kissinger, H. E.:, 1957. Reaction Kinetics in Differential Thermal Analysis" Anal. Chem., 29(11), 1702-1706.

Kok, M.V., 2001. Performance Prediction of In Situ Combustion Processes, Energy Sources, Part A: Recovery, Utilization, and Environmental Effects, 23:10, 937 – 943.

Kräutle, S. and Knabner, P. (2005), `A new numerical reduction scheme for fully coupled multicomponent transport-reaction problems in porous media', Water Resources Research 41.

Kristensen, M. R. (2008), Development of Models and Algorithms for the Study of Reactive Porous Media Processes, PhD thesis, Technical University of Denmark.

Kristensen, M. R., Gerritsen, M., Thomsen, P. G., Michelsen, M. L. and Stenby, E. H. (2007 a), Coupling chemical kinetics and flashes in reactive, thermal and compositional reservoir simulation, Society of Petroleum Engineers. SPE 106218: Presented at 2007 SPE Reservoir Simulation Symposium, Houston, TX, February 26 - 28.

Kristensen, M. R., Gerritsen, M., Thomsen, P. G., Michelsen, M. L. and Stenby, E. H. (2007 b), `Efficient integration of stiff kinetics with phase change detection for reactive reservoir processes', Transport in Porous Media 69(3), 383--409.\

Kulhman, M., 2000. The Benefits of In Situ Upgrading reactions to the Integrated

Operations of the Orinoco Heavy Oil Fields and Downstream Facilities, paper SPE 71615 presented at 2000 SPE/AAPG Western Regional Meeting held in Long Beach, California, 19-23 June.

Kumar, V.K., Fassihi, M.R., and Yannimaras, D.V. 1994. "Case History and Appraisal of the Medicine Pole Hills Unit Air Injection Project", SPE 27792 presented at the SPE/DOE Symposium on Improved Oil Recovery held in Tulsa, OK, April 17 – 20.

Kumar, V. K., Gutierrez, D., Moore, R. G. and Mehta, S. A. (2007), Case history and appraisal of the west buffalo red river unit high-pressure air injection project, Society of Petroleum Engineers. SPE 107715: Presented at the SPE Hydrocarbons Economics and Evaluation Symposium, Dallas, TX, April 1-3.

Kutta, W. (1901), `Beitrag zur näherungsweisen integration totaler differentialgleichungen', Zeitschrift für Mathematik und Physik 46, 435--453.

Kværnø, A. (2004), `Singly diagonally implicit runge-kutta methods with an explicit first stage', BIT Numerical Mathematics 44, 489--502.

Lake, L.W., 1989. Enhanced Oil Recovery, Prentice Hall, Englewood Cliffs: New Jersey, 1989.

Lambers JV, Gerritsen MG, Mallison BT. Accurate Local Upscaling with Variable Compact Multi-point Transmissibility Calculations. To be published in Computational Geosciences, Special issue on Multiscale Methods, 2007.

Law, C., K., 2006. Combustion Physics, Cambridge University Press .

Lee, B. I. and Kesler, M. G. (1975), `A generalized thermodynamic correlation based on three corresponding states', AIChE Journal 21, 510--527.

Lee SH, Tchelepi HA, Jenny P, DeChant LJ. Implementation of a Flux-Continuous Finite-Difference Method for Stratigraphic, Hexahedron Grids. SPE Journal, 7:267:277, 2002.

Levenspiel, O., 1999. Chemical Reaction Engineering, 3rd edition,John Wiley and Sons.

Mamora, D. D.: "The Effect of Low-Temperature Oxidation on Kinetic Tube Experiments", SPE 29614, SPE Memorial Series, Vol. 1, 1995.

Mamora, D. D.: Kinetics of In-situ Combustion, PhD dissertation, Stanford University, Stanford, CA (May 1993).

Mamora, D. D. and Brigham, W. E., 2005. The Effect of Low-Temperature Oxidation on the Fuel and Produced Oil During In-Situ Combustion, IN SITU (Nov.) 19(4), 341-

365.

Mamora, D. D. and Brigham, W. E 1995. "Implications of Low-Temperature Oxidation in Kinetic and Combustion Tube Experiments," presented at UNITAR Int. Conf. on Heavy Oil and Tar Sands, Houston, TX USA., Feb. 12-17, 1995.

Mamora, D. D. 1995. "New Findings in Low-Temperature Oxidation of Crude Oil," SPE 29324, presented at SPE Asia-Pacific Oil and Gas Conf., Kuala Lumpur, Malaysia, 20-22 Mar..

Michelsen, M. L. (1982), `The isothermal flash problem. part 1: Stability', Fluid Phase Equilibria 9, 1--19.

Michelsen, M. L. and Mollerup, J. M. (2004), Thermodynamic Models: Fundamentals and Computational Aspects, Tie-Line Publications, Holte, Denmark.

Mittal R, Iaccarino, G. Immersed Boundary Methods. Annual Review of Fluid Mechanics, 37:239-61, 2005.

Moore, G., Mehta, R. and Ursenbach, M. (2002 a), `Air injection for oil recovery', Journal of Canadian Petroleum Technology 41(8), 16--19.

Moore, R. G., Laureshen, C. J., Belgrave, J. D. M., Ursenbach, M. G. and Mehta, S. A. (1995), `In-situ combustion in canadian heavy oil reservoirs', Fuel 74(8), 1169--1175.

Moore, R. G., Laureshen, C. J., Ursenbach, M. G., Mehta, S. A. and Belgrave, J. D. M. (1999), `Combustion/oxidation behavior of athabasca oil sands bitumen', SPE Reservoir Evaluation and Engineering 2(6), 565 -- 572.

Moore, R. G., Mehta, S. A. and Ursenbach, M. G. (2002 b), A guide to high pressure air injection (HPAI) based oil recovery, Society of Petroleum Engineers. SPE 75207: Presented at the SPE/DOE Improved Oil Recovery Symposium, Tulsa, OK, April 13-17.

Nightingale, E. R., Jr., 1959. Phenomenological Theory of Ion Solution—Effective Radii of Hydrated Ions, Journal of Physical Chemistry. 63, 1381–1387.

Nilsson, J., Gerritsen, M. and Younis, R. (2005), A novel adaptive anisotropic grid framework for efficient reservoir simulation, Society of Petroleum Engineers. SPE 93243: Presented at 2005 SPE Reservoir Simulation Symposium, Houston, TX, January 31 - February 2.

Nordbotten JM, Aavatsmark I, Eigestad GT. Monotonicity of Control Volume Methods. Submitted to Numerische Mathematik, 2006.

Orr Jr., F. M. (2007), Theory of Gas Injection Processes, Tie-Line Publications.

Ozawa, T.:, 1965. A new method of Analyzing Thermogravimetric Data", Bull. Chem. Soc. Japan, 38, 1881.

Park, T. and Barton, P. I. (1996), `State event location in differential-algebraic models', ACM Transactions on Modeling and Computer Simulation 6(2), 137--165.

Pascual, M., Lacentre, P. and Coombe, D. (2005), Air injection into a mature waterflooded light oil reservoir. laboratory and simulation results for barrancas field, argentina, Society of Petroleum Engineers. SPE 94092: Presented at the SPE Europec/EAGE Annual Conference, Madrid, Spain, June 13-16.

Peng, D. Y. and Robinson, D. B. (1976), `A new two constant equation of state', Industrial Engineering Chemistry Fundamentals 15, 59--64.

Petzold, L. R. (1982), `DASSL: A differential/algebraic system solver'. 10th IMACS World Congress on System Simulation and Scientific Computation.

Pickup GE, Ringrose PS, Jensen JL, Sorbie KS. Permeability Tensors for Sedimentary Structures. Mathematical Geology, 26:227-50, 1994.

Prats, M. (1986), Thermal Recovery, Vol. 7 of SPE Monograph Series, Society of Petroleum Engineers.

Press, H. et al., 1986. Numerical Recipes-The Art of Scientific Computing, Cambridge University Press, Cambridge 329.

Prothero, A. and Robinson, A. (1974), `On the stability and accuracy of one-step methods for sloving stiff systems of ordinary differential equations', Mathematics of Computation 28(125), 145--162.

Rangel-German, E. R., J. Schembre, C. Sandberg, and A. R. Kovscek, 2004. "Electrical-Heating-Assisted Recovery for Heavy Oil," J. Pet Sci. & Eng. 45(3-4), 213-233.

Rasmussen, C. P., Krejbjerg, K., Michelsen, M. L. and Bjurstrøm, K. E. (2006), `Increasing the computational speed of flash calculations with applications for compositional, transient simulations', SPE Reservoir Evaluation & Engineering pp. 32--38.

Ren, Y., Freitag, N. P. and Mahinpey, N. (2005), A simple kinetic model for coke combustion during an in situ combustion (ISC) process. Presented at the 6th Canadian International Petroleum Conference, Calgary, Alberta, June 7-9.

Rosenbrock, H. H. (1963), `Some general implicit processes for numerical solution of differential equations', Computer Journal 6(4), 329--330.

Runge, C. (1895), `über die numerische auflösung von differentialgleichungen', Mathematische Annalen 46, 167--178.

Saad, Y. (1994), SPARSKIT: A Basic Tool Kit for Sparse Matrix Computations, University of Minnesota, MN.

Sammon, P H, "Dynamic Grid Refinement and Amalgamation for Compositional Simulation", SPE RSS, 2003, SPE 79683.

Sandu, A., Verwer, J. G., Bloom, J. G., Spee, E. J., Carmichael, G. R. and Potra, F. A. (1997), `Benchmarking stiff ODE solvers for atmospheric chemistry problems: II. rosenbrock solvers', Atmospheric Environment 31(20), 3459--3472.

Sandu, A., Verwer, J. G., Loon, M. V., Carmichael, G. R., Potra, F. A., Dabdub, D. and Seinfeld, J. H. (1997), `Benchmarking stiff ODE solvers for atmospheric chemistry problems: I. implicit vs. explicit', Atmospheric Environment 31(19), 3151--3166.

Sarathi, P. S. (1999), In-situ combustion handbook -- principles and practices, Technical report, U.S. Department of Energy. DOE/PC/91008-0374.

Schult, D. A., Matkowsky, B. J., Volpert, V. A. and Fernandez-Pello, A. C. (1996), `Forced forward smolder combustion', Combustion and Flame 104, 1--26.

Senum G., I., and Yang, R., T., J.:" Rational approximations of the integral of the Arrhenius function", Thermal Analysis, 1979, Vol. 11, 445.

Shallcross, D.C., Rios, C.F. De Los, Castanier, L.M., 1991. Modifying In-Situ Combustion Performance by the Use of Water-Soluble Additives, Society of Petroleum Engineering Reservoir Engineering Aug, 287-294.

Speight, J.G., 2007. The Chemistry and Technology of Petroleum, 4th Ed., CRC Press, Boca Raton Fl, 2007.

Stokka, S., Oesthus, A. and Frangeul, J. (2005), Evaluation of air injection as an ior method for the giant ekofisk chalk field, Society of Petroleum Engineers. SPE 97481: Presented at the SPE International Improved Oil Recovery Conference, Kuala Lumpur, Malaysia, December 5-6.

Stone, H. L. (1973), `Estimation of three-phase relative permeability and residual oil data', Journal of Canadian Petroleum Technology 12(4), 53--61.

Strang, G. (1968), `On the construction and comparison of difference schemes',

SIAM Journal of Numerical Analysis 5(3), 506--517.

Targac, G. W., Redman, R. S., Davis, E. R., Rennie, S. B., McKeever, S. O., and Chambers, B. C., 2005. Unlocking the value in West Sak Heavy Oil", paper SPE/PS-CIM/CHOA 97856 PS2005-390 Proceedings of the 2005 SPE International Thermal Operations and Heavy Oil Symposium held in Calgary, Alberta Canada, 1-3 November.

Tadema, H . J, 1959. Mechanism of oil production by underground combustion, Proceedings of the 5th World Petroleum Congress section 11, paper 22, 279-287.

Thomsen, P. G. (2006), Discontinuities in ODEs - systems with change of state, Technical Report IMM-2006-07, Informatics and Mathematical Modelling, Technical University of Denmark.

Tingas, J., Greaves, M. and Young, T. J. (1996), Field scale simulation study of in-situ combustion in high pressure light oil reservoirs, Society of Petroleum Engineers. SPE 35395: Presented at the SPE/DOE Improved Oil Recovery Symposium, Tulsa, OK, April 21-24.

Trangenstein J, Bi Z. Large Multi-Scale Iterative Techniques and Adaptive Mesh Refinement for Miscible Displacement Simulation. SPE/DOE Improved Oil Recovery Symposium, 2002, SPE 75232.

Twu, C. H. (1984), `An internally consistent correlation for predicting the critical properties and molecular weights of petroleum and coal tar liquids', Hydrocarbon Processing 16, 137--150.

Venuto, P.B., 1989. "Tailoring EOR Processes to Geologic Environments," World Oil 209 (November) 61 – 68.

Vyazovkin, S., 1997. "Evaluation of Activation Energy of Thermally Simulated Solid-State Reactions under Arbitrary Variation of Temperature", Journal of Computational Chemistry, Vol. 18 No 3, 393-402, 1997.

Vyazovkin S. and Wight C. A., 1999. Model-free and model-fitting approaches to kinetic analysis of isothermal and nonisothermal data, Thermochimica Acta, , 43-68.

Vyazovkin, S., 2001. Modification of the Isoconversional Method to Account for Variation in the Activation Energy, Journal of Computational Chemistry, Vol. 22 No 2, 178-183, 2001.

Walter, A. L., Frind, E. O., Blowes, D. W., Ptacek, C. J. and Molson, J. W. (1991), `Modeling of multicomponent reactive transport in groundwater. 1. model development and evaluation', Water Resources Research 30(11), 3137--3148.

Watson DF. Contouring: A guide to the analysis and display of spacial data.

Pergamon, 1994.

Wen XH, Durlofsky LJ, Edwards MG. Use of Border Regions for Improved Permeability Upscaling. Mathematical Geology, 35:521-547, 2003.

Westbrook, C. K., Mizobuchi, Y., Poinsot, T. J., Smith, P. J. and Warnatz, J. (2005), `Computational combustion', Proceedings of the Combustion Institute 30, 125--157.

Whitson, C. and Brule, M. (2000), Phase Behavior, Society of Petroleum Engineers. SPE Monograph, volume 20.

Whitson, C. H. and Michelsen, M. L. (1989), `The negative flash', Fluid Phase Equilibria 53, 51--71.

Williams, R., Burrage, K., Cameron, I. and Kerr, M. (2002), `A four-stage index 2 diagonal implicit runge-kutta method', Applied Numerical Analysis 40, 415--432.

Xu, T., Sonnenthal, E., Spycher, N. and Pruess, K. (2004), THOUGHREACT User's Guide: A Simulation Program for Non-Isothermal Multiphase Reactive Geochemical Transport in Variably Saturated Geologic Media, Lawrence Berkeley National Laboratory, Berkeley, CA.

Yeh, G.-T. and Tripathi, V. S. (1991), `A model for simulating transport of reactive multispecies components: model development and demonstration', Water Resources Research 27(12), 3075--3093.

Younis, R. and Gerritsen, M. (2006), Multiscale process coupling by adaptive fractional stepping: An in-situ combustion model, Society of Petroleum Engineers. SPE 93458: Presented at the 2006 SPE/DOE Symposium on Improved Oil Recovery, Tulsa, OK, April 22-26.

Younis RM, Caers J. A Method for Static-based Upgridding. Proc. of the 8th European Conference on the Mathematics of Oil Recovery, Sept. 2002

Zlatev, Z. (1995), Computer Treatment of Large Air Pollution Models, Kluwer Academic Publishers.

National Energy Technology Laboratory

626 Cochrans Mill Road
P.O. Box 10940
Pittsburgh, PA 15236-0940

3610 Collins Ferry Road
P.O. Box 880
Morgantown, WV 26507-0880

One West Third Street, Suite 1400
Tulsa, OK 74103-3519

1450 Queen Avenue SW
Albany, OR 97321-2198

539 Duckering Bldg./UAF Campus
P.O. Box 750172
Fairbanks, AK 99775-0172

Visit the NETL website at:
www.netl.doe.gov

Customer Service:
1-800-553-7681