Exploring the Combustion Modes of A Dual-Fuel Compression Ignition Engine

by

Jonathan Atwood Martin

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mechanical Engineering)
in The University of Michigan
2019

Doctoral Committee:

Professor André L. Boehman, Chair
Professor Christian M. Lastoskie
Research Scientist George A. Lavoie
Research Scientist Robert J. Middleton
Professor Margaret S. Wooldridge

“Only a small part of scientific progress has resulted from planned search for specific objectives. A much more important part has been made possible by the freedom of the individual to follow his own curiosity.”

— Irving Langmuir

Jonathan Atwood Martin

jonmart@umich.edu

ORCID iD: 0000-0002-1239-6300

To 晓菲

Acknowledgments

First and foremost I'd like to thank Prof. André Boehman, without whom I would not have had the opportunity to pursue my own ideas and follow a truly original research path. I could never have acquired the support I needed for this project on my own, and it is only thanks to his top-notch networking skills that I was able to find the connections I needed to see it through to the end. He is truly unflappable in the face of complex problems, and his faith in my abilities probably exceeds what I deserve.

Next I have to thank the fine gentlemen at Hino Motors, particularly Furukawa-San, Kato-San, Ichii-San, and Funayama-San. I am extremely humbled by and grateful for their keen interest in my work, the financial support they were able to provide me, and their patience when things did not go as we had planned. I hope this thesis lives up to their company's high standard of excellence. Domo arigato.

The other three gentlemen without whom I never would achieved anything close to this are William Kirkpatrick, Kent Pruss, and James Elkins, the Auto Lab's resident machine shop experts. These men are the unsung heroes of the Auto Lab, selflessly and kindly showing clueless grad students like me everything we're doing wrong, and vastly increasing the quantity and quality of the research output from not just the Auto Lab but the entire ME department.

The helpfulness of the Auto Lab faculty cannot be overstated, especially Dr. George Lavoie and Dr. Robert Middleton who always had time for me interrupt their work and unfailingly gave great suggestions. I was lucky to acquire some knowledge from Prof. Jason Martz and Dr. Stanislav Bohac before they moved on to greater heights, and I was blessed to share a lab space with Dr. John Hoard for my whole time here, whose students thankfully did not revolt against him when having to listen to the unbearable racket my engine was making. Last but not least, it was great to have Chuck Solbrig join our team as a technician for my last year or so here; his experience with engine test cells is unmatched and his excellent taste in music helped to keep me invigorated during the long, tedious data collection process near the end.

Outside of the Auto Lab, ME Prof. Margaret Wooldridge was always a consummate professional during the limited time I worked alongside her, and I wish I had more opportunity to do so. It has been great having CEE Prof. Christian Lastoskie on my committee to provide the outside perspective, and his work probably has a greater impact on climate change than I could ever hope for. Learning about turbulent combustion from Aerospace Prof. Venkat Raman was simultaneously the most challenging and most rewarding semester of my academic career, and I hope someday I can share his depth of understanding of the subject. And although my research diverged from his area of expertise, I thoroughly enjoyed my time with ME Prof. James Barber, without whom I never would've discovered Irving Langmuir and his lectures on Pathological Science, of which I hope my thesis does not prove to be an example.

And of course I have to thank my brilliant cohort of fellow students and visiting scientists, whose contributions are too numerous for this space so I will just list their names: Dr. Vickey Kalaskar, Dr. Dongil Kang, Dr. Chenxi Sun, Justin Koczak, Dr. Taehoon Han, Kwang Hee Yoo, Taemin Kim, Shuqi Cheng, Erick Garcia, Kaustav Bhadra, Aaron Garnier, Rinav Pillai, Minato Tomuro, Dr. Bo Yang, Xin Meng, Heng Chen, Dr. Adam Vaughan, Dr. Vassilis Triantopoulos, Dr. Jordan Easter, Dr. Yan Chang, Bryan Maldonado, Mitchell Bieniek, Dr. Dimitris Assanis, Dr. Ripudaman Singh, Dr. Ahmet Mazacioglu, Dr. Ivan Tivabinsky, Lucca Henrion, Allison Gibson, Angela Wu, Andrew DiMauro, Uday Subramaniam, Sumit Chopra, Rutvik Topkar, Ankit Sakuja, Akshay Shekar, and Puneet Valecha.

Finally, thanks to Dr. Phil Johnson, Kevin Ma, and Dr. T.J. Flynn. Made it up and down.

Preface

This work was primarily supported by Hino Motors, with additional funding from Volvo Technology of America and the United States Department of Energy (DOE), DOE award numbers DE-EE0006831, DE-EE0004232, and DE-EE0007745.

Table of Contents

List of Tables

Table

List of Figures

Figure

Abstract

Compression-ignition (CI) engines, also known as “diesel” engines, can produce higher thermal efficiency (TE) than spark-ignition (SI) engines, which allows them to emit less carbon dioxide (CO_2) per unit of energy generated. Unfortunately, in practice the TE of CI engines is limited by the need to maintain sufficiently low emissions of nitrogen oxides (NO_X) and soot. This problem can be mitigated by operating CI engines in dual-fuel modes with port fuel injection (PFI) of gasoline supplementing the direct injection (DI) of diesel fuel. Several strategies for doing this have been introduced in recent years, but these operating modes are usually studied individually at discrete conditions. This thesis introduces a classification system for dual-fuel CI modes that links together several previously studied modes across a continuous two-dimensional diagram. The combustion modes covered by this system include the standard modes of conventional diesel combustion (CDC) and conventional dual-fuel (CDF); the well-explored advanced combustion modes of HCCI, RCCI, PCCI, and PPCI; and a relatively unexplored combustion mode that is herein titled “Piston-split Dual-Fuel Combustion” or PDFC. The results show that dual-fuel CI engines can simultaneously increase TE and lower NO_X and/or soot emissions at high loads through the use of *Partial* HCCI (PHCCI), despite an increase in emissions of carbon monoxide (CO) and unburnt hydrocarbons (UHC). At low loads, PHCCI is not possible, but either PDFC or RCCI can be used to further improve NO_X and/or soot emissions, albeit at slightly lower TE and still higher emissions of CO and UHC. This multi-mode strategy of PHCCI at high loads and PDFC or RCCI at low loads is particularly useful when low engine-out NO_X emissions are required.

Chapter 1

Introduction

One of the biggest challenges facing engineers today is to reduce the carbon dioxide (CO_2) emitted by the world's energy systems. Throughout our history, and continuing to this day, humanity's main source of energy has been the combustion of carbon-based fuels, which inherently produces CO_2. This gas is the most prominent (but not only) "greenhouse gas" (GHG) produced by mankind that contributes to climate change, a problem that will, by overwhelming worldwide scientific consensus, generate dire consequences for humanity if our energy systems and consumption habits cannot be drastically altered in the next few decades [1].

Promising developments have been made over recent decades to reduce the amount of CO_2 emitted per unit of useful energy generated, a metric that is often referred to as "carbon intensity," reported in units of grams of CO_2 (or equivalent grams of CO_2 when considering the GHG effect of another emission) per MJ of energy generated, written as gCO_2e/MJ. Since the industrial revolution, the majority of the world's energy has been produced through the combustion of coal and petroleum, processes that have very high carbon intensities in excess of 100 gCO_2e/MJ which will be classified here as "high-carbon." Other, less conventional fuel sources such as natural gas, hydrogen, and biofuels have been estimated to significantly lower the carbon intensity of these processes into the range of 10-100 gCO_2e/MJ which will be classified here as "medium-carbon." Processes that do not involve combustion, such as nuclear and renewable electricity, can reduce overall carbon intensity below 10 gCO_2e/MJ which will be classified here as "low-carbon." It is important to note that these carbon intensities are calculated assuming current technology, and incorporate energy expenditures and losses for the entire “life cycle” of the fuel. There are still many opportunities for technological improvements to reduce the carbon intensity of even the “high-carbon” fuels and processes, which will be an integral part of overall carbon emission reductions.

When looking at the overall worldwide energy budget and its carbon intensity, it is useful to examine energy usage in the two largest sectors: electricity generation and transportation. The US Energy Information Administration has broken down the main sources of energy for each of these two sectors as of 2012, and made projections of how these breakdowns will change by 2040 [2]. By grouping the energy sources into the low-carbon, medium-carbon, and high-carbon categories as described above, Figure 1.1 shows that low-carbon sources are projected to more than double their electricity output by 2040, while output from high-carbon sources will stay around the same level as 2012. These same projections do not show any significant contributions from low-carbon energy in the transportation sector, with only a small fraction of the energy being supplied by low-carbon electricity. Although the medium-carbon sources will approximately double their output, this will not be nearly enough to cover the overall increase in demand for transportation, and high-carbon energy output in the transportation sector will continue to grow. For this reason, a high priority must be given to finding carbon intensity improvements in the high-carbon sources for transportation energy.

The majority of transportation energy is currently provided by internal combustion (IC) engines fueled with one of two petroleum derivatives: gasoline and diesel fuel. Gasoline is almost exclusively used in spark-ignition (SI) engines, while diesel fuel is almost exclusively used in compression-ignition (CI) engines, so much that the term “diesel engine” has become synonymous with “CI engine.” Although there are many steps throughout the entire “well-to-wheels” (WTW) life cycle that determine the carbon intensity of each fuel, the step that IC engine researchers focus on is the conversion of the net heat released by combustion ($Q_{c,net}$) into rotational work done on the driveshaft of the vehicle (W_{rot}). The percentage

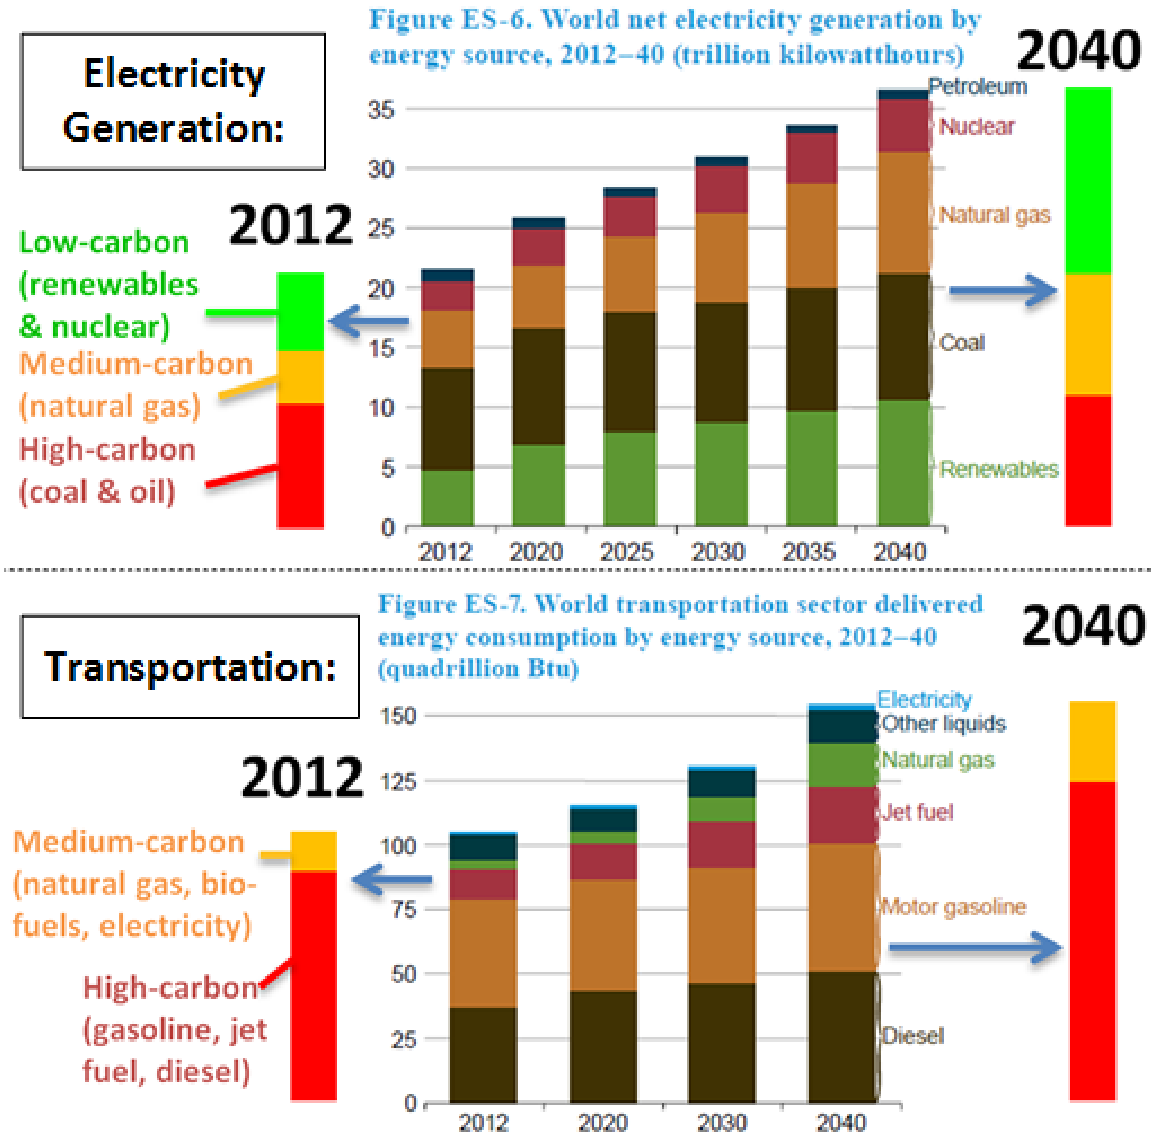

Figure 1.1 Projections of shifts in worldwide energy sources from Conti et al. [2], annotated by the author of this thesis. This shows the disparity between the electricity generation and transportation sectors in terms of carbon intensity. The use of high-carbon fossil fuels in electricity generation is expected to remain static from now until 2040, but the use of these fuels for transportation is expected to grow. This highlights the need for efficiency improvements in fossil fuel transportation systems in order to reduce CO_2 emissions.

of heat converted to work is known as Brake Thermal Efficiency (BTE), represented by $\eta_{th,brake}$ in Equation 1.1.

$$\eta_{th,brake} = W_{rot}/Q_{c,net} \tag{1.1}$$

The BTE produced by modern SI engines, depending on engine and operating condition, ranges from 18% to 30%, while the BTE of CI engines ranges from 22% to 40% [3]. The higher BTE of CI engines can be attributed to their higher compression ratio (CR), which raises the theoretical limit (known as Carnot efficiency) that can be achieved as dictated by the laws of thermodynamics, defined below in Equation 1.2. This limit is determined by the ambient temperature (T_C) and the adiabatic flame temperature of the fuel/air mixture (T_H):

$$\eta_{th,Carnot} = 1 - T_C/T_H \tag{1.2}$$

T_H will be increased by increases in the starting temperature of the fuel/air mixture when combustion is initiated. This starting temperature is determined by the compression ratio (CR) of the engine, with higher CR leading to higher T_H and thus higher Carnot efficiency. The Carnot efficiency of IC engines is generally in the range of 70-90%, much higher than the realized BTE. This difference can mostly be attributed to three main losses of heat by the engine:

- **Exhaust loss:** Unextracted heat in the hot exhaust and unburnt fuel molecules flowing out the tailpipe, which are then dispersed into the atmosphere and cannot be recovered
- **Cooling loss:** Heat transferred to the metal boundaries of the combustion chamber, which is carried away by engine coolant and dispersed to the atmosphere at the radiator
- **Mechanical loss:** Friction in the engine's moving parts and work extracted to power accessory components of the engine (such as the fuel and coolant pumps)

Despite both ideal and practical BTE being higher for diesel-fueled CI engines, the current WTW carbon intensity of diesel fuel of 102.82 gCO_2e/MJ [4] is actually slightly higher than gasoline at 100.58 gCO_2e/MJ [5]. There are many complex reasons why this is the case, but the reason that this thesis will focus on is the need to mitigate the *other* (non-GHG) emissions typically produced by CI engines: nitrogen oxides (NO_X) and soot. These emissions must be limited for the sake of public health and environmental preservation, and unfortunately conventional CI engines cannot operate at their most efficient settings without producing excessive levels of one or both of these emissions.

The high CR of CI engines raises the temperatures of combustion, which is good for efficiency (see Equation 1.2), but is bad for NO_X emissions. Most of the NO_X produced by CI engines is "thermal NO_X," meaning it is formed by the diatomic nitrogen and oxygen in air, which do not normally react with each other but become reactive at the highest temperatures seen in high-CR CI engines. This can be combated by diluting the charge via exhaust gas recirculation (EGR), but this also reduces the air/fuel ratio and leads to increased soot emissions, and the level of EGR needed can often have a detrimental effect on efficiency [6]. Even with the proliferation of EGR, most new diesel vehicles in the US also feature both a diesel particulate filter (DPF) to remove soot and use selective catalytic reduction (SCR) to convert NO_X and a diesel emissions fluid (DEF) into nitrogen, CO_2, and water. These emissions treatment methods are collectively more detrimental to the efficiency and overall carbon intensity of the entire WTW cycle than the three-way catalyst (TWC) emissions treatment systems found on most modern SI engines, which contributes to the similar carbon intensity of diesel to gasoline despite the higher BTE of diesel engines.

For this reason, much of the current research in CI engines has focused on reducing both the air/fuel ratios and temperatures of combustion to avoid NO_X and soot formation while maintaining the same high compression ratios and their associated high efficiency. In this field, air/fuel ratio is often converted into equivalence ratio ϕ, which is the stoichiometric air/fuel ratio of the fuel used divided by the actual air/fuel ratio of the combustion process. Values of ϕ below 1 indicate excess air or "lean" conditions, with a value of 0 representing pure air, while values above 1 represent insufficient air for complete combustion or "rich"

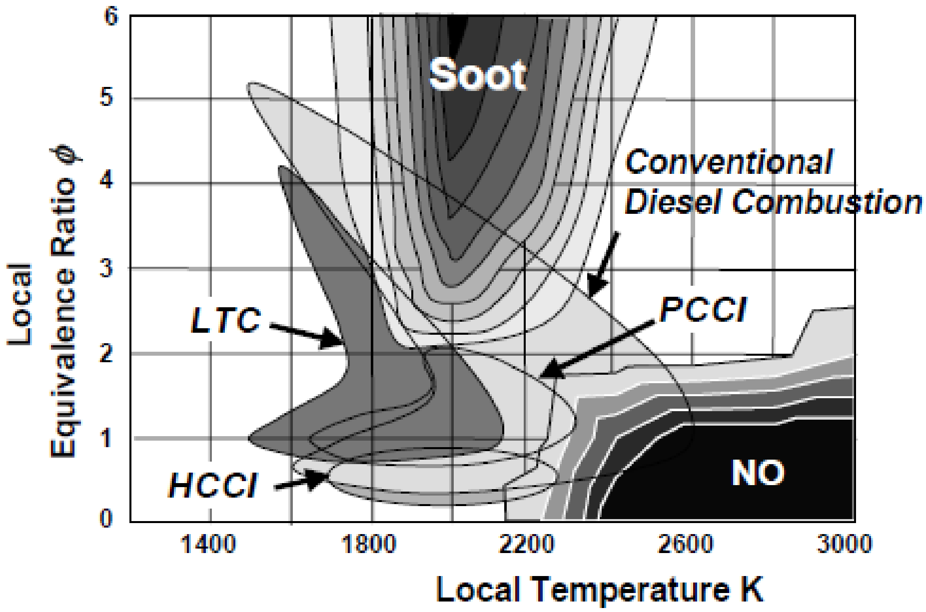

Figure 1.2 ϕ-T diagram reprinted from Neely et al. [8]. This shows the range of local conditions within a typical CI engine combustion chamber in terms of temperature and equivalence ratio ϕ. The fuel injection techniques of conventional diesel combustion create zones of both high NO production and high soot production, while the advanced combustion modes of LTC, PCCI, and HCCI avoid creating these zones.

conditions. Sometimes, especially in the CI engine field, the reciprocal of this value is used instead and represented by λ. Combustion strategies are often shown on a so-called "ϕ-T diagram" that shows the rate of formation of NO and soot at specific combinations of ϕ and T based on chemical kinetics, with the most widely used version first developed by Akihama et al. [7]. These diagrams can be overlaid with the range of ϕ-T values seen in combustion at specific points of space and time throughout the cylinder in simulations or optical studies. An example of this from Neely et al. [8] is reprinted here as Figure 1.2, which highlights the regions seen in conventional diesel combustion (CDC) as well as three advanced combustion modes: homogeneous charge compression ignition (HCCI), premixed charge compression ignition (PCCI), and low temperature combustion (LTC).

There are many strategies proposed throughout the literature to avoid the NO_X and soot "islands" seen in Figure 1.2. The strategy that will be the focus of this thesis is dual-fueling, i.e. the introduction of multiple fuels with contrasting properties into the combustion chamber. The second fuel that is combined with diesel fuel in a dual-fuel CI engine is typically more resistant to autoignition, which can be generally referred to as a "low-reactivity fuel" (LRF). LRF can therefore be premixed with air long before combustion occurs so that it achieves a lean mixture throughout the combustion chamber. Diesel fuel has low resistance to autoignition, and can be referred to along with other diesel alternatives as a "high-reactivity fuel" (HRF). Diesel and other HRF's would autoignite prematurely in a high-compression ratio CI engine if they were premixed, so HRF injection is typically withheld until just before the start of the combustion process. The HRF will then ignite before it has time to mix evenly with air, producing combustion over a wide spectrum of localized equivalence ratios and temperatures. The displacement of late-injected HRF (e.g., diesel) with premixed LRF (e.g., gasoline) displaces some of the local ϕ-T combinations down and to the left on Figure 1.2, although the HRF combustion will still occur in the conventional diesel combustion zone. This strategy will be referred to herein as "conventional dual-fuel" or CDF for short. When the premixed LRF is present, the HRF can be injected much earlier than in conventional diesel combustion, since mixing with the LRF will prevent it from autoigniting prematurely. This gives the HRF time to mix more completely with air and eliminate the areas of rich and near-stoichiometric equivalence ratio before combustion occurs. This avoids the NO_X/soot "islands" from Figure 1.2, but also increases emissions of carbon monoxide (CO) and unburnt hydrocarbons (UHC). This strategy is referred to herein as "Reactivity-Controlled Compression Ignition" or RCCI [9].

There is a trend in IC engine research of testing and comparing discrete operating modes, which are typically given names with memorable acronyms, seen above with LTC, HCCI, PCCI, and RCCI, with an abundance of other names proposed throughout the literature. However these operating modes often cannot be applied to a very wide range of speeds and loads, and must instead be incorporated as part of a "multi-mode" engine operation strategy. Another trend that can be seen in the literature is the tendency to propose multi-mode strategies that involve one discretely-defined combustion mode at certain speed/load

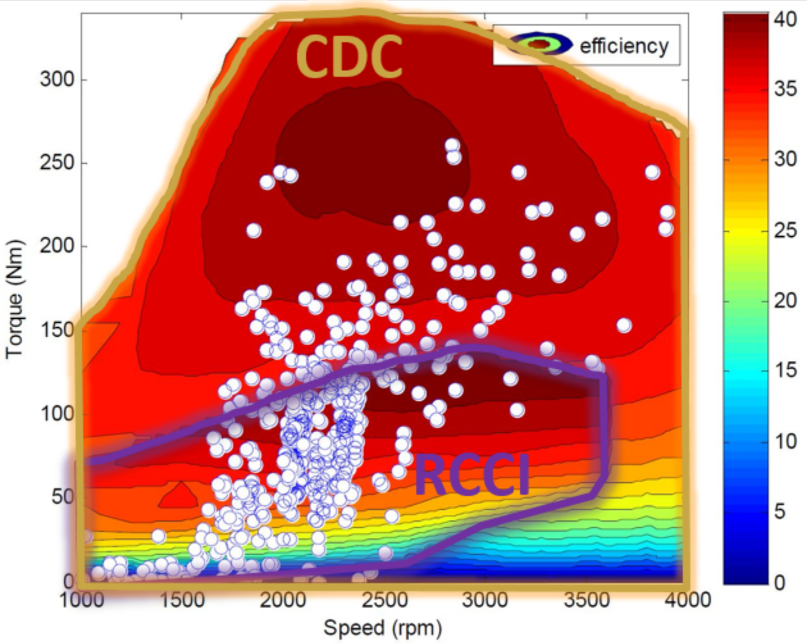

Figure 1.3 Multi-mode CDC/RCCI strategy reprinted from Gao et al. [10]. Engine dynamometer test data was used to calculate color maps of BTE for the combustion modes of RCCI (within the purple region, where it is feasible) and CDC (for the entire speed/load range of the engine). Vehicle performance was then simulated for a conventional 1,580kg mid-size passenger car with an automatic transmission put through an EPA US06 drive cycle, with the white points indicating the instantaneous speed and load of the engine at one-second intervals. This shows a possible strategy for multi-mode operation, using RCCI at low loads and switching to CDC at high loads. This thesis will investigate intermediate combustion modes that incorporate aspects of the both of the two modes (CDC and RCCI) to create a more continuous multi-mode operating strategy than the discrete two-mode strategy seen above.

conditions, and another discretely-defined mode at other conditions. An example of this can be seen in the work of Gao et al. [10], which proposes a multi-mode strategy of CDC and RCCI. In the work of Gao et al., RCCI is used only at low-to-medium loads with CDC used elsewhere, and a simulated US-06 EPA drive cycle is performed on the engine as shown in Figure 1.3.

In Figure 1.3, the white circles represent the instantaneous speed and load demanded by the drive cycle at one-second intervals. This drive cycle, if implemented in a real-world engine, would require "jumping" back and forth between RCCI and CDC quite often. The simulations performed by Gao et al. [10] were derived from steady-state engine performance measurements, and do not fully account for the transient performance changes that would occur when suddenly shifting between RCCI and CDC. In addition, there is a noticeable drop in efficiency when crossing from RCCI to CDC. For these reasons, it might be beneficial to incorporate a transition mode that combines aspects of both RCCI and CDC to “smooth out” this transition.

Given the highly transient nature of IC engine operation, “jumping” between discrete operating modes could lead to rough driving and counteract the efficiency/emissions benefits of advanced modes. This problem could potentially be solved through the application of other new IC engine technologies such as hybridization, but this comes with its own additional costs that can potentially cancel out the benefits. The solution proposed herein is to think about *smoothly shifting* along a *continuous spectrum* of operating modes instead of *jumping* between *discrete* modes. This can be achieved by creating a continuous “map” of operating modes that incorporates all of the above-mentioned operating modes, as well as any intermediate modes that may occur in between. Such a map can more fully explore the envelope of performance of dual-fuel CI engines and maximize their benefits, and this thesis will do so for one particular engine model: a light-duty, turbocharged diesel engine that has been modified to also inject gasoline.

This process will start by defining a map of operating modes as classified by the type of injections used (Chapter 2). This will be followed by an analysis of the heat release rate (HRR) patterns generated by these operating modes to interpret the combustion sub-processes occurring in each (Chapter 3). Once these processes are understood, an extensive set of experiments will “map out” the performance of a test engine across a wide spectrum of combustion modes (Chapter 4). Finally, a more focused set of test conditions will be selected for more rigorous, repeated testing to quantify the effect of the best-performing dual-fuel modes on emissions (Chapter 5).

This work has led to a journal publication in the 2018 SAE International Journal of Engines [11] which details the classification system (Chapter 2) and performs limited testing of a single speed/load condition. This original publication [11] used conventional diesel as the HRF and a mixture of propane and dimethyl ether (DME) as the LRF in a light-duty passenger vehicle turbodiesel engine (the same model used by Gao et al. [10]). The more recent work from which the data in this thesis is derived has replaced the LRF with more conventional gasoline in the same engine. Much of the HRR interpretation (Chapter 3) and performance mapping (Chapter 4) was presented at the 11th U.S. National Combustion Meeting in Pasadena, CA in March 2019. The final optimized testing from Chapter 5 occurred in the spring of 2019 and was being prepared for peer review and publication at the time this dissertation was submitted in August 2019.

Chapter 2

Classifying Operating Modes by Injection Type

This chapter introduces a new method of classifying the operating modes that are possible for a certain type of dual-fuel engine. The chapter will begin by describing the particular engine configuration in Section 2.1, move on to discussion of how to classify the types of fuel injection that can be used in Section 2.2, and finally use the percentage distribution of the fuel between these injection types to classify the operating mode of the engine in Section 2.3.

2.1 Engine Configuration

With the wide variety of potential engine configurations available (in terms of piston shape, injector design, and intake/exhaust configuration) it would be challenging to devise a dual-fuel operating mode map that applies to all possible engine configurations, and thus this map will be limited to one particular dual-fuel configuration. Given the obvious motivations to reduce carbon intensity as quickly as possible, it makes sense to select an engine configuration that is close as possible to current commercial engines, so that there are minimal design changes needed to achieve high efficiency dual-fuel combustion and thereby minimal barriers to implementation. Ideally, this engine configuration will also have successfully demonstrated efficiency improvements from dual-fuel combustion in previous studies. In this thesis, the two fuels used in the engine will be the most conventional and widely available fuels, gasoline and diesel fuel, to maximize the immediate commercial potential and minimize barriers to implementation.

The dual-fuel strategies introduced in this thesis could readily be applied to other fuel combinations as long as the two fuels have a considerable contrast in their reactivities. The gasoline could be replaced with any fuel having a similar resistance to autoignition, generally described as a "low-reactivity fuel" (LRF), such as natural gas, ethanol, or propane. Conversely the diesel fuel could be replaced with any fuel having a similar propensity to autoignition, generally described as a "high-reactivity fuel" (HRF), such as biodiesel or dimethyl ether (DME). But the work herein will focus just on gasoline and diesel.

Given the criteria above, this thesis will consider only an engine configuration that is largely unmodified from a conventional CI engine, possessing these characteristics:

- **VGT Turbocharger:** A turbine/compressor with a variable geometry turbine (VGT) that can pressurize intake air to a set point of pressure to provide enough oxygen for combustion at high loads. This is achieved by recovering heat from the exhaust flow at the turbine, with a charge air cooler installed post-compressor to reduce the temperature of the air after pressurization and maximize the density of the air (thereby maximizing the mass of air drawn into the fixed-volume cylinder).
- **High-Pressure EGR loop:** A bridge between the exhaust and intake manifolds on the high-pressure side of the turbocharger providing exhaust gas recirculation (EGR), with a cooler to reduce the temperature of the exhaust and prevent premature autoignition of premixed fuel, used to vary both air/fuel ratio and charge dilution.
- **Re-entrant Bowl Piston:** A conventional diesel engine piston with a recessed bowl in the center (into which the fuel is normally injected) and a flat outer region that closely approaches the cylinder head at the top of its range of motion.
- **DI Diesel (HRF) Injector:** A conventional high-pressure, common rail, direct-injection (DI) multi-jet diesel injector that injects into the bowl of the piston near top dead center through the center of the cylinder head.
- **PFI Gasoline (LRF) Injector**: This will be the only **new** component being introduced to the conventional CI engine configuration, and will be a conventional port fuel

injection (PFI) system that can be grafted on to the engine's intake manifold, delivering cyclic pulses of fuel into the intake port during the intake stroke, with no need to perform any complex internal modifications to the engine block or cylinder head.

To satisfy these criteria, the engine chosen for the experiments in this thesis is a General Motors (GM) 1.9 Liter ZDTH turbodiesel engine, detailed in Table 2.1 and shown in its experimental setup in Figure 2.1. This engine model is found in several light-duty passenger diesel vehicles produced by Fiat and GM, including the 2004-2008 Vauxhall/Opel Astra, and has been previously used to demonstrate high efficiency dual-fuel combustion by multiple research groups [12, 13]. The engine was installed on a brake dynamometer and instrumented to measure all the necessary temperatures, pressures, and flows needed to analyze its performance and energy balance. The original engine control unit (ECU) was replaced with a combination data acquisition (DAQ)/ECU-simulator cart from National Instruments/DRIVVEN. Besides these changes, which are standard for experimental test engines, only one change needed to be made to the engine from its original configuration as delivered by the manufacturer: the installation of the PFI gasoline injectors.

Table 2.1 GM 1.9L ZDTH engine specifications.

Layout	Inline 4	Compression Ratios	17.5 (Effective)/18.6 (Geometric)
Volume	1.909 L	Number of Valves	16 (2 intake/2 exhaust per cyl.)
Stroke	90.4 mm	Rated Power	110 kW (148 hp) @ 4000 RPM
Bore	82.0 mm	Rated Torque	315 N-m (232 ft-lb) @ 2000 RPM
Conrod	145 mm	DI Injector Model	Bosch 445-110-243
IVO	361°aTDC	DI Spray Pattern	7 jets, 148°umbrella
IVC	587°aTDC	PFI Injector Model	Bosch EV 14 CKxT
EVO	112°aTDC	PFI Rail Pressure	4 bar (58 psi)
EVC	356°aTDC	PFI Spray Pattern	30°hollow cone

This modification was inspired by a similar modification of the same engine performed by Curran et al. [12]. The modified intake manifold of Curran et al. is shown in Figure 2.1, with one of the injectors entering the intake manifold from the top in order to avoid the diesel rail pump. Curran's manifold is shown just to give the reader a better idea of how the injectors are mounted, since a suitable photograph of the modified manifold used in this

Table 2.2 Fuels used in this thesis.

	Ultra-low-sulfur diesel ("ULSD")	Unleaded Test Gasoline ("UTG-96")
Derived Cetane Number	45.1	-
Density, kg/m3	831.8	741.3
T90, °C	312	153
Lower Heating Value (MJ/kg)	42.58	42.96
Aromatic content, %	31.5	28.8
H:C Ratio	1.82	1.89
RON	-	96.5
MON	-	86.6
AKI (RON+ MON/2)	-	91.6
Sulfur Content, ppm	<15	15.9
Vapor Pressure, kPa abs	-	63.9

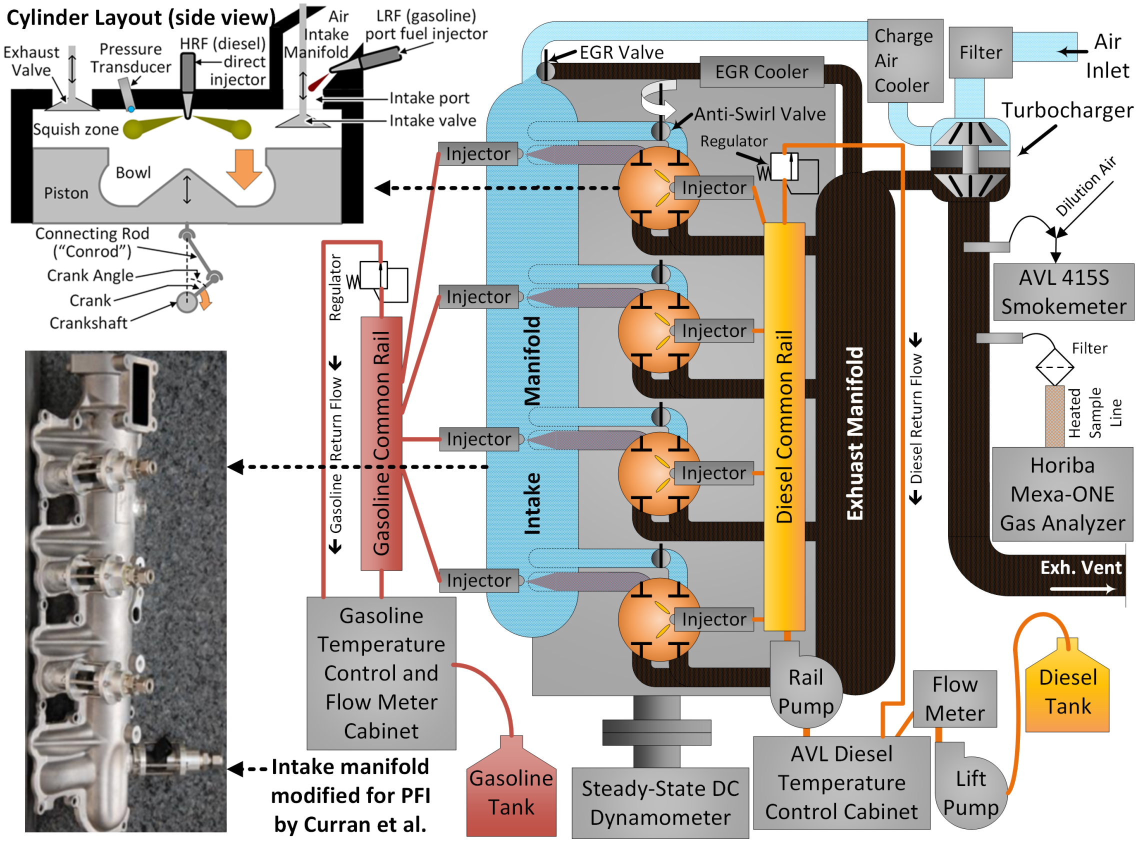

Figure 2.1 Top-down schematic view of the engine used in this thesis, with side view inside the cylinders and an example of an intake manifold modified for PFI from Curran et al. [12].

thesis was not taken before installing the manifold onto the engine, where it is difficult to get a clean photograph that clearly shows the injectors. However in thesis, all four PFI injectors were mounted directly on the passageway between the manifold and the intake port, such that the gasoline is directly aimed at the intake ports, just like the first three injectors seen in Figure 2.1.

The only feature of this engine configuration not yet mentioned is the anti-swirl valve located behind the second intake port of each cylinder. This purpose of this valve is to induce a swirling motion in the charge by throttling the air flow from the second intake valve. The engine's ECU calls for this valve to close further at high speeds and loads to induce higher swirl, likely for the purpose of improving fuel/air mixing and shortening combustion duration when the time for combustion becomes limited. Since this is a less common feature of CI engines, its effects will not be considered and it will be set at a constant value of 70% open for the duration of this study.

2.2 Primary Injection Types

Given the chosen engine configuration, the appropriate injection timing of each of the fuels must be considered. To show when the proper injection windows will occur in the engine cycle, Figure 2.2 shows how the *instantaneous* compression ratio varies throughout the engine cycle based on the crank angle θ. This is simply the ratio of the volume V_{IVC} of the combustion chamber at the time the intake valve closes (IVC) and compression begins, and the instantaneous volume $V(\theta)$ at a particular crank angle. This will roughly correspond to the pressure and temperature within the chamber before combustion occurs. To maximize torque, combustion should occur mostly within a window of -10 to 20°aTDC, and the timing at which 50% of the total heat release has occurred ("CA50") should be at approximately 5°aTDC, although this CA50 may not be the best for minimizing emissions.

The PFI injection of gasoline should only occur when the intake valve is open, and air is

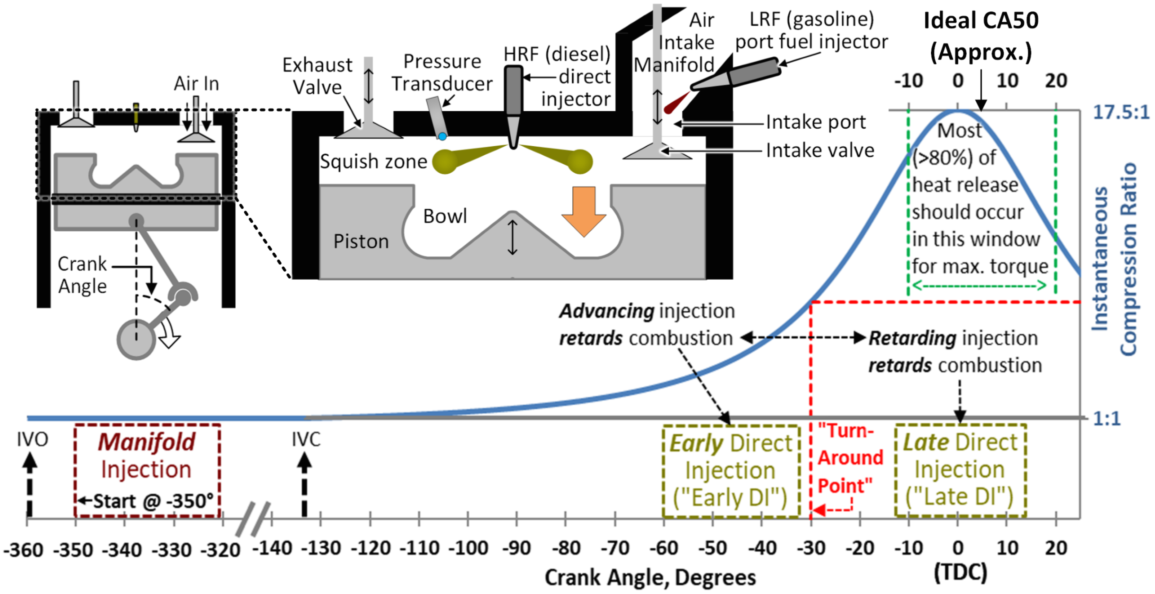

Figure 2.2 Instantaneous compression ratio versus crank angle, showing the timing of the primary injection types. The PFI injection must occur while the intake valve is open, and ideally at the top of the intake stroke to best pull this fuel into the cylinder. The Early DI will be directed into the squish zone and its timing will be negatively correlated to combustion phasing, while the Late DI will be directed into the bowl and its timing will be positively correlated to combustion phasing.

being drawn into the cylinder, so that as much as possible of the injected gasoline is drawn into the combustion chamber, and does not wet the walls of the intake passageway or diffuse back into the main body of the intake manifold. For the same reasons, the timing of the injection should maximize the suction pulling the gasoline into the chamber, which will be greatest when the piston is near the top of the stroke but descending, and the relative expansion of the combustion chamber is maximized. Another thing to consider is the valve lift, and since the valve lift profile for this engine was unknown, a buffer of 9 crank angle degrees after the stated intake valve opening (IVO) timing was given to ensure the intake valve was sufficiently lifted before starting the PFI injection. For these reasons, the start-of-injection (SOI) timing of the PFI injection was held constant at -350°aTDC.

It should be noted that PFI is not the only method possible to introduce LRF into the intake air upstream of the intake valve. Other possible methods include using a carburetor or an intake fumigation (IF) system, which mix fuel with the entire air stream continuously, rather than injecting cycle-by-cycle. An IF system was used previously in a "proof-of-concept" paper demonstrating early iterations of the ideas introduced in this thesis [11]. This earlier paper used a mixture of two fuels, propane and DME, as the LRF fuel, although they were mixed at a ratio designed to give autoignition resistance close to standard gasoline. The disadvantages of these PFI alternatives are that they do not provide cycle-by-cycle control, and cannot reach high loads without generating an air/fuel mixture above the lower explosive limit (LEL) in the intake manifold, generating a risk of backfire. However all of these methods will achieve the same result once the intake valve closes: The LRF will be distributed homogeneously throughout the cylinder. Since all of these injection methods have the same effect, they will be grouped together under the umbrella term **"Manifold Injection"**, which is the first of three primary injection types used herein.

The DI injection should occur only once the piston has ascended far enough to "catch" the injected fuel, and prevent excessive amounts of liquid diesel droplets from wetting the cylinder wall. Injecting too early will cause some of this fuel to go unburnt, reducing engine efficiency while also contaminating the engine's lubricating oil. Based on previous experiments [11], an early limit for DI start-of-injection (SOI) timing was set at -65°aTDC. However too late of an injection will also cause some fuel to go unburnt, since the piston will begin to descend and reduce temperatures and pressures before the fuel has a chance to burn. For this reason, a late limit for DI SOI timing was set at TDC.

Given these limits, one might make an uninformed guess that the ideal DI timing lies in the middle, around -30°aTDC. However, this is not the case, as illustrated in a previous study from Doosje et al. [14]. Results from this study, which featured 95% of the fuel in a manifold injection of compressed natural gas (CNG) with 5% DI diesel, are reprinted in Figure 2.3. The SOI of the diesel injection was swept from -60°aTDC to TDC, and both NO_X and soot emissions reach their maximum between -30 and -20°aTDC, so this is not a desirable injection timing. Looking at the heat release rates (HRR) reveals why: injections at these timing produce the most advanced combustion phasing, and cause most of the heat of

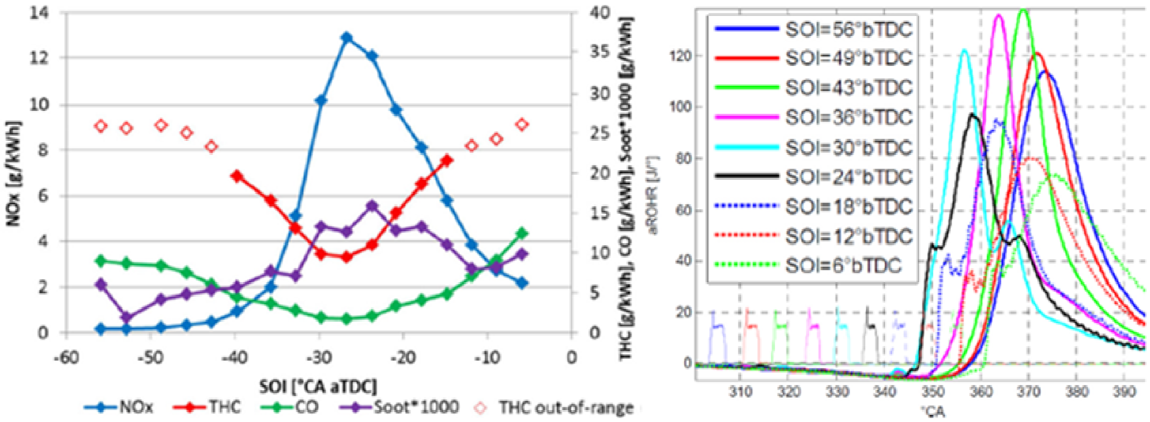

Figure 2.3 Sweep of SOI timing from early to late from Doosje et al. [14]. The left plot shows how the "Early DI" timings of 40°bTDC and earlier produce lower NO_X and soot than the "Late DI" timings of 20° and later, while the timings near the "turn-around" point producing maximum NO_X and soot. The right plot shows the trend in HRR, with the "turn-around" SOI of 30°bTDC producing the most advanced combustion phasing. The solid line traces from the Early DI timings have a different characteristic shape than the dotted line traces from the Late DI timings, indicating different combustion processes produced by each.

combustion to be released before TDC. This maximizes the peak temperatures and produces the most possible NO_X, while the delay between SOI and ignition has been minimized and there is not sufficient time for fuel-air mixing to avoid rich zones with soot production. Therefore DI diesel injections around this timing window should be avoided.

Also note in Figure 2.3 that the most advanced combustion phasing possible is achieved with an SOI of -30°aTDC. Starting at -60°, combustion phasing is initially *advanced* as the injection is retarded, but after -30°aTDC, this trend "turns around," and retarding the injection further past -30°aTDC will *retard* the combustion phasing instead. The same "turn-around point" of approximately -30°aTDC was observed in the same engine used in this thesis, and is marked in Figure 2.2. This turn-around point will be used herein to divide the DI injection window into two primary injection types: ***Early*** **DI** and ***Late*** **DI**. The precise turn-around point will change depending on the engine, injector, fuel(s) used, and specific operating conditions. For example, another study with 80% PFI gasoline with 20% DI diesel found this turn-around point at -27°aTDC [15]. Wherever the turn-around point occurs in any given condition will define the dividing line between Early DI and Late DI.

With an Early DI, the diesel (or other HRF) is being injected before the temperatures in the combustion chamber are high enough to trigger autoignition. This allows time for the fuel to atomize, vaporize, mix with air, and start to diffuse evenly throughout the combustion chamber before it begins to burn. By the time the temperature does get high enough for autoignition, all the HRF will be in a lean mixture with air (and LRF if there was a manifold injection). Although there will still be some stratification in the mixture, the highest concentration of fuel will be below the stoichiometric mixture by the time autoignition occurs, and the combustion will be completely lean and premixed, avoiding the NO_X and soot "islands" seen in Figure 1.2.

Conversely with a Late DI, the diesel is injected *after* the autoignition temperature has been reached, and will burn instead through a mixing-controlled process. There will still be a significant ignition delay between the start of injection and the start of combustion, but at moderately high loads this delay will be short enough such that autoignition will begin to occur even before the injection is over. As a result, there will be significant amount of incomplete rich premixed combustion generating soot, as well as diffusion combustion generating NO_X. The advantage, of course, is that this process is very rapid and controllable, and will produce high combustion efficiency, which is why it has historically been the main injection type employed in conventional CI engines.

An optical engine study which shows very clearly the difference between these three main types of injection was carried out by Musculus et al. [16], shown in Figure 2.4. This study did not use a PFI injector for gasoline, but instead used a side-mounted gasoline DI injector with an SOI timing of -240°aTDC, injecting the standard gasoline surrogate *iso*-octane. The HRF DI injector conformed to the specifications in Section 2.1 above, and was used to inject the standard diesel surrogate *n*-heptane. A mixture of these two fuels is called a "primary reference fuel" (PRF), with a "PRF number" corresponding to the percentage of *iso*-octane in the blend. The PRF number indicates the octane number of the blend, i.e. the resistance of the fuel to autoignition.

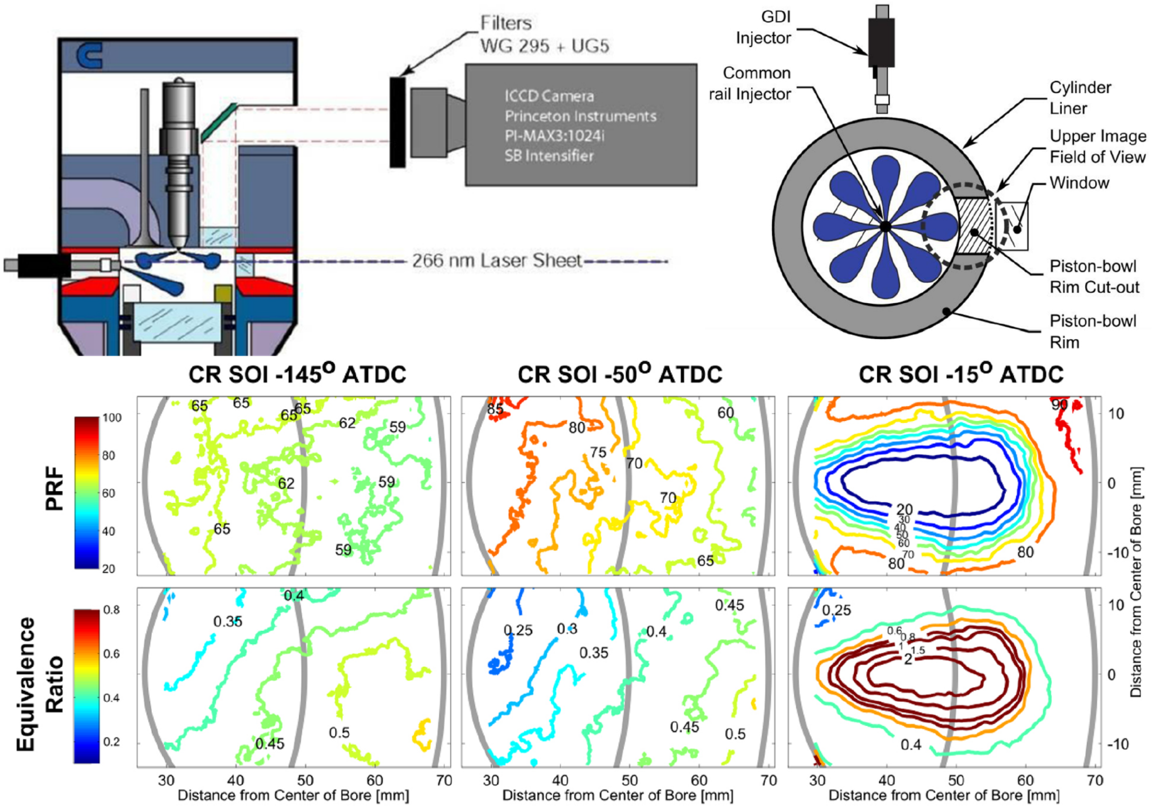

Figure 2.4 Optical study of fuel distribution and SOI timing from Musculus et al. [16], with a constant DI of LRF (*iso*-octane) during the intake stroke and three different SOI timings for a DI of HRF (*n*-heptane). The earliest SOI (left) produces a relatively homogeneous reactivity with a shallow gradient in equivalence ratio, but this timing is too early for the engine used for this thesis. The next SOI (center) could be used as an Early DI in this thesis, and produces shallow gradients in both reactivity and equivalence ratio. The last SOI (right) could be used as a Late DI in this thesis, and produces much sharper gradients in both metrics.

The goal of Musculus et al. was to determine the distributions of ϕ and PRF number throughout the cylinder when three different HRF DI timings were used: -145°aTDC, -50°aTDC, and -15°aTDC. The images shown in Figure 2.4 were taken at a timing of -5°aTDC, close to the desired ignition timing. When the HRF DI timing of -145°aTDC was used, the mixture achieves a fairly homogeneous state before ignition, with a slight gradient in equivalence ratio, but a nearly homogeneous PRF number distribution. Since the HRF *n*-heptane has had time to diffuse throughout the chamber, this mixture may not be sufficiently resistant to autoignition in an engine with a high CR, but if the right combination of CR, fuel mixture, and intake temperature is used, HCCI can be achieved.

When an HRF DI timing of -50°aTDC is used, there is only a slight increase in the stratification of equivalence ratio, but a much steeper gradient in PRF number. The entire mixture is still very lean (equivalence ratio below $\phi = 0.5$) and dominated by the LRF *iso*-octane, *except* near the edge of the combustion chamber. This should lead to better resistance to autoignition than the -145°aTDC case, especially considering the most reactive part of the mixture (ϕ closest to 1 and lowest PRF number) is near the cylinder wall, where temperatures should be somewhat reduced due to convective heat transfer to the wall. This is the strategy behind the Early DI: provide an injection of HRF that is early enough to mix down to a sufficiently lean state before maximum compression is reached, but not so early that it mixes *completely* and retains a sufficient stratification of reactivity. A combination of a Manifold Injection and an Early DI should lead to a "sequential autoignition" (SA) process, with multiple autoignition events propagating down the gradient of reactivity. This is the driving mechanism behind RCCI, and achieves fully lean and premixed combustion similar to HCCI but with a bigger envelope of operation and easier control than HCCI.

Conversely, when a HRF DI timing of -15°aTDC is used, there are *much* sharper gradients of both PRF number and ϕ. This is because the Late DI has barely had enough time to atomize and vaporize, and has only just begun diffusing throughout the combustion chamber. This will lead to a rapid process of "mixing-controlled autoignition" (MCA), wherein locally there is insufficient oxygen for complete combustion ($\phi > 1$), generating soot. The soot and other incomplete combustion products (CO and UHC) will then diffuse toward the remaining oxygen, generating a high-temperature diffusion flame that creates NO_X from the air in its outer preheat zone. When a Manifold Injection of LRF has occurred, this diffusion flame will also consume the LRF, unless the heat released by the MCA is sufficient to ignite a faster premixed flame, or even cause autoignition of the LRF mixture.

In this experiment, the proportion of fuel between the two injections was fixed at 64% in the very early LRF injection (effectively serving as the Manifold Injection despite being a DI) and 36% in the HRF injection (which was an Early DI in the second case of -50°aTDC and a Late DI in the third case of -15°aTDC). The next section will explore all the possible distributions of fuel between these types of injections, and attempt to classify operating modes base on these distributions.

2.3 Operating Modes

When considering only the three injections mentioned in the previous section (Manifold Injection of LRF, Early DI of HRF, Late DI of HRF), it is possible to plot the distribution of fuel between these three injections on a ternary phase diagram or "triangle plot," shown in Figure 2.5. This figure has been named the "MELT diagram," an acronym for "**M**anifold / **E**arly DI / **L**ate DI **T**riangle diagram. The vertices of the triangle represent the exclusive use of only one of these injections, with the percentage of fuel in that injection going down moving away from that vertex. This will serve as the basis for the "map" of dual-fuel engine performance. Although there are potentially many other multi-injection strategies that use more than one injection fitting each of the two DI types, only one Late "Main" DI and one Early "Main" DI will be considered for now. Once a sufficient understanding has been obtained regarding what goes on with just three main injections, more complex injection strategies can be considered, which will happen later in Chapter 5.

The distribution of fuel to these injections is not the only factor determining the combustion behavior, so there cannot be hard lines on this diagram that demarcate particular "combustion modes" per se. Instead, the lines on the diagram in Figure 2.5 demarcate "operating modes" or "fueling modes." For the sake of referring to particular areas of the MELT diagram, Figure 2.5 applies acronyms from the literature that are typically referred to as "combustion modes" to identify these operating modes. However there must be an investigation of the combustion behavior occurring with each of these fuel distributions while varying the other parameters in order to truly identify the "combustion mode" occurring at any given condition.

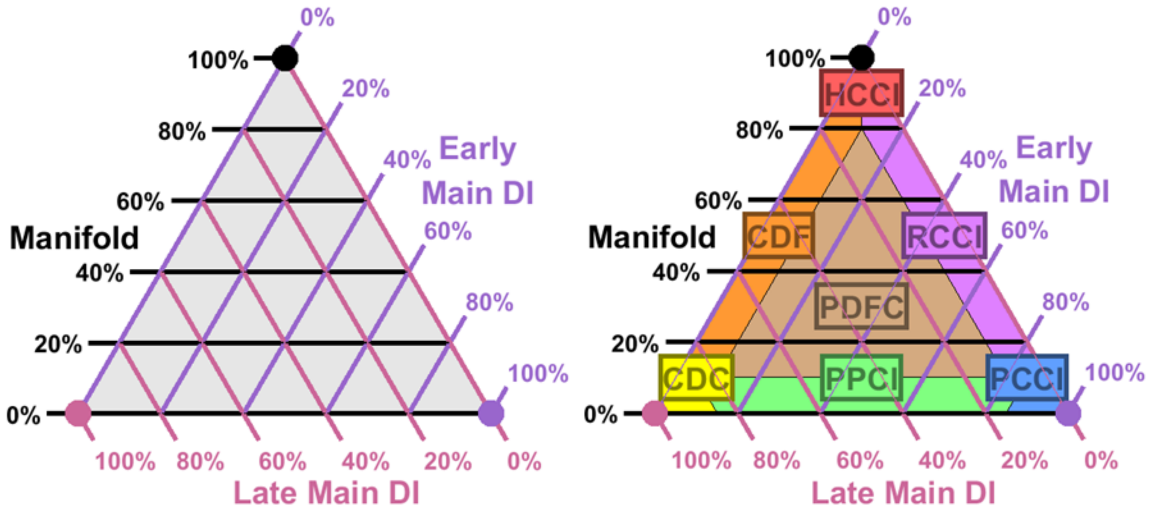

Figure 2.5 Ternary phase diagram showing distribution of fuel into 3 primary injection types (The "MELT Diagram"). On the right this diagram is overlaid with the proposed classification of operating modes, with the acronyms taken from the literature and defined below.

The acronyms used in Figure 2.5 are:

- Most fuel ($> 90\%$) in one injection:
  - **CDC:** "Conventional Diesel Combustion" (Late DI)
  - **PCCI:** "Premixed Charge Compression Ignition" (Early DI)
  - **HCCI:** "Homogeneous Charge Compression Ignition" (Manifold)
- Most fuel ($> 90\%$) in two injections:
  - **PPCI:** "Partially Premixed Compression Ignition" (Early DI & Late DI)
  - **CDF:** "Conventional Dual-Fuel" (Late DI & Manifold)
  - **RCCI:** "Reactivity-Controlled Compression Ignition" (Early DI & Manifold)
- At least 10% fuel in all three injections:
  - **PDFC:** "Premixed Dual-Fuel Combustion" or "Piston-split Dual-Fuel Combustion" (the latter is an alternative definition introduced in this thesis)

Having introduced the "MELT diagram" basis for the map of combustion modes, there can now be further investigation of what occurs at each point on the map in the following chapter. This investigation will justify the labels that have been applied to each operating mode, and introduce one more label that has not yet been considered.

Chapter 3

Interpreting Combustion Behavior from Heat Release Rate

This chapter illustrates all of the combustion modes that are possible with the dual-fuel CI engine configuration described in Section 2.1. The chapter will begin by selecting a heat release rate (HRR) model from the literature to fit the data from the test engine in Section 3.1. This is followed by the introduction of an empirical curve fit that has been commonly used by IC engine researchers to represent the HRR in Section 3.2. Next will be a detailed description of the method that was developed for fitting these curves to the HRR data in Section 3.3. The chapter will then conclude with a "tour" of all the combustion modes on the MELT diagram and show how the curve fits can be used to represent specific combustion processes through comparison with previous research in Section 3.4.

3.1 Modeling Heat Release Rate

When performing experiments in a conventional production IC engine, there is often only one measurement available taken from inside the cylinder: in-cylinder pressure, p, sampled either at a constant frequency or at fixed increments of crank angle θ. However there is a lot of information that can be obtained just from this one measurement, most notably the bulk cylinder gas temperature T_g and the HRR $\frac{dQ_c}{d\theta}$. The simplest way to do this is to assume the contents of the cylinder are fixed (no mass added or subtracted), and behave as an ideal gas. If the geometry of the engine is known, and hence its instantaneous volume $V(\theta)$ at each crank angle, the bulk gas temperature can be derived from ideal gas law, knowing the starting pressure, temperature, and volume at the time of IVC:

$$T_g(\theta) = T_{g,IVC}\frac{p(\theta)V(\theta)}{p_{IVC}V_{IVC}} \tag{3.1}$$

A simple analysis of the first law of thermodynamics yields an equation for the net heat added to the cylinder [17]:

$$\frac{dQ_n}{d\theta} = \frac{\gamma}{\gamma-1}p\frac{dV}{d\theta} + \frac{1}{\gamma-1}V\frac{dp}{d\theta} \tag{3.2}$$

The ratio of specific heats γ can be calculated from the charge composition and temperature. Since the effect of composition is outweighed by the effect of temperature in the experiments seen herein, for the sake of computational simplicity, this thesis will simply use a third-order polynomial fit of the relationship of γ to bulk gas temperature T_g:

$$\gamma = -9.967\cdot 10^{-12}\cdot T_g^3 + 6.207\cdot 10^{-8}\cdot T_g^2 - 1.436\cdot 10^{-4}\cdot T_g + 1.396 \tag{3.3}$$

If the cylinder can be considered adiabatic, then Q_n in Equation 3.2 is simply equal to the heat released by combustion Q_c. However, most researchers will at least include a term Q_{ht} to account for the heat transferred to the cylinder head, engine block, and piston, using one of many empirical correlations available in the literature. The analysis in this chapter, which is only meant to illustrate the combustion modes, will use the correlation from Chang et al. [18], referred to as the Chang model, to obtain Q_{ht} and calculate Q_c:

$$\frac{dQ_c}{d\theta} = \frac{dQ_n}{d\theta} + \frac{dQ_{ht,Chang}}{d\theta} \tag{3.4}$$

The Chang heat transfer model was chosen because it was developed using actual heat flux measurements from an HCCI engine, making it the ideal model to employ for an engine that is producing the same lean premixed autoignition type of combustion as HCCI. Heat transfer from this type of combustion is not accurately modeled by most common diesel engine models, which were developed under the assumption that all of the combustion is the same mixing-controlled type of combustion that occurs in CDC. The first model derived here will be the original and more widely used Woschni model [19], on which the Chang model is based.

3.1.1 Woschni Heat Transfer Model

The Woschni model belongs to the family of "single-zone" heat transfer models, meaning that it simplifies the cylinder contents to a single, homogeneous zone with constant properties over the entire volume. This allows for a simple relationship of the heat transfer loss to the heat transfer coefficient h, the bulk gas temperature T_g, cylinder wall temperature T_w, and, and combustion chamber surface area A:

$$\frac{dQ_{ht}}{dt} = h \cdot [T_g(\theta) - T_w] \cdot A(\theta) \tag{3.5}$$

The heat transfer coefficient h is derived differently depending on the model used. Most single-zone models rely upon the empirical relationship of two dimensionless numbers, the Nusselt number Nu (ratio of convective to conductive heat transfer) and the Reynolds number Re (ratio of inertial forces to viscous forces), with an empirical coefficients α and m that must be "tuned" to the experimental data:

$$Nu = \alpha \cdot Re^m \tag{3.6}$$

The definitions of the Nusselt and Reynolds numbers can be substituted into Equation 3.6 to produce the relationship between the heat transfer coefficient h and the bulk properties of the cylinder:

$$\frac{hL}{k} = \alpha \left[\frac{\rho v L}{\mu}\right]^m \tag{3.7}$$

This equation is then rearranged to yield a correlation for h:

$$h = \alpha \cdot L^{m-1} \cdot k \cdot \mu^{-m} \cdot \rho^m \cdot v^m \tag{3.8}$$

Density ρ can be calculated from the ideal gas law using the measured cylinder pressure and calculated bulk gas temperature:

$$\rho = \frac{m}{V(\theta)} = \frac{p(\theta)}{T_g(\theta)} \tag{3.9}$$

Characteristic length scale L is calculated differently depending on the model used. The original Woschni model simply uses the cylinder diameter or “bore” B, although it will be seen later that a better length scale can be used for when the piston gets close to the cylinder head.

The average turbulent gas velocity v, unfortunately, cannot be directly measured, and must also be “tuned” to experimental data, with empirical coefficients that will change not just with the particular engine used but also with the particular mode being used, with different combustion modes inducing different amounts of turbulence. This velocity can be broken into two components: the velocity v_p induced by the motion of the piston, and the additional velocity v_c induced by combustion of the fuel, which was first done in the highly influential Woschni model [19]. In this model, on which the Chang model is based, v_p is modeled as directly proportional to the mean piston speed $\bar{S}_p$, which is itself determined by the stroke length s and engine rotational speed ω:

$$v_p = C_1 \cdot \bar{S}_p = C_1 \cdot 2s\omega \tag{3.10}$$

Meanwhile, v_c is modeled as being directly proportional to a term representing the “intensity of combustion” as it is called by Woschni:

$$v_c(\theta) = C_2 \cdot \frac{V(\theta) T_{g,IVC}}{p_{IVC} V_{IVC}} [p(\theta) - p_{motor}(\theta)] \tag{3.11}$$

The $[p(\theta) - p_{motor}(\theta)]$ term represents the departure of the measured cylinder pressure from the theoretical “motoring” pressure that would be observed in the absence of any combustion. Although Woschni does not explain the origin of the $V(\theta)T_{IVC}/p_{IVC}V_{IVC}$ term [19], looking back on equations 3.9 and 3.1 shows that it is simply the reciprocal of the charge density ρ. This is simply normalizing the increase in pressure induced by combustion with the current charge density, since the pressure increase due to a fixed amount of combustion heat release will be greater at higher charge densities:

$$v_c(\theta) = C2 \cdot \frac{T_g(\theta)}{p(\theta)}[p(\theta) - p_{motor}(\theta)] = C_2 \cdot \frac{p(\theta) - p_{motor}(\theta)}{\rho(\theta)} \tag{3.12}$$

A value for p_{motor} can be calculated by assuming polytropic compression and expansion, i.e. in the absence of combustion the increase in pressure will be proportional to the instantaneous compression ratio taken to the power of a constant polytropic coefficient K:

$$\frac{p_{motor}(\theta)}{p_{IVC}} = \left[\frac{V_{IVC}}{V(\theta)}\right]^K \tag{3.13}$$

The best-fit polytropic exponent can be determined experimentally from the measured pressure during the compression stroke, after IVC seals the chamber and before the DI begins. This involves comparing pressure and volume at IVC with some future crank angle called θ_{end}, which will be set at -70°aTDC in order to avoid any pressure disturbance from the Early DI:

$$K = \frac{\log\left(\frac{p(\theta_{end})}{p_{IVC}}\right)}{\log\left(\frac{V_{IVC}}{V(\theta_{end})}\right)} \tag{3.14}$$

Equations 3.9, 3.10, and 3.12 can now be substituted back into Equation 3.8:

$$h(\theta) = \alpha \cdot B^{m-1} \cdot k \cdot \mu^{-m} \cdot \frac{p(\theta)}{T_g(\theta)}^{m} \cdot \left[C_1 \cdot 2s\omega + C_2 \cdot \frac{p(\theta) - p_{motor}(\theta)}{\rho(\theta)}\right]^m \tag{3.15}$$

The thermal conductivity k and dynamic viscosity μ both have exponential relationships to bulk gas temperature, with $k \propto T_g^{0.75}$ and $\mu \propto T_g^{0.62}$. Rather than calculating individual correlation coefficients for these, they can simply be rolled into the scaling coefficient α for simplicity. The resulting T_g terms can then be combined, and the θ arguments can be removed for clarity:

$$h = \alpha \cdot B^{m-1} \cdot p^m \cdot T^{(0.75-0.62m-m)} \cdot \left[C_1 \cdot 2s\omega + C_2 \cdot \frac{p - p_{motor}}{\rho}\right]^m \tag{3.16}$$

Woschni set the value of m (the exponent applied to the Reynolds number when correlating to the Nusselt number) to 0.8, since this was the value used in the fluid mechanics literature for turbulent pipe flow. The values of Woschni's best fit coefficients (when converting to SI units from the original work [19]) were $\alpha = 13 \cdot 10^{-3}$, $C_1 = 2.28$, and $C_2 = 3.24 \cdot 10^{-3} \frac{m}{s \cdot K}$. This yields the final Woschni heat transfer equation:

$$h = \alpha \cdot B^{-0.2} \cdot p^{0.8} \cdot T^{-0.53} \cdot \left[C_1 \cdot 2s\omega + C_2 \cdot \frac{p - p_{motor}}{\rho}\right]^{0.8} \tag{3.17}$$

3.1.2 Chang Heat Transfer Model

The Chang model [18] uses the same structure as the Woschni model while modifying some of the coefficients to fit the heat flux measured in an HCCI engine. There were three modifications made to generate the best fit between the heat flux measurements and the modeled heat transfer:

- A dynamic length scale L reflecting the changing combustion chamber geometry
- An adjustment to the exponent of the temperature term T
- A reduction of the coefficient of combustion-induced gas velocity C_2

Chang found that the length scale that produced the most accurate calculations was the shortest length scale of the combustion chamber at any given moment. Far from TDC, the shortest length scale is the combustion chamber's radius (half the cylinder bore B), but near TDC (when most combustion occurs) the shortest length scale is the average cylinder height, calculated from the volume and cross-sectional area A_{cs}:

$$L(\theta) = min\left(\frac{B}{2}, A_{cs}\right) = min\left(\frac{B}{2}, \frac{4V(\theta)}{\pi B^2}\right) \tag{3.18}$$

The temperature exponent was derived from the original correlation of Nusselt and Reynolds numbers in turbulent pipe flow, which is a rather general correlation, and can be further optimized to better represent the relationship between temperature and heat transfer coefficient in an IC engine. Chang optimized this value using heat flux data from the compression stroke alone, where the combustion-induced gas velocity term does not come into play. The best-fit temperature exponent was -0.73, opposed to the original value of -0.53 (from the $0.75-0.62m-m$ in Equation 3.16).

The final change to the Woschni model made by Chang was the reduction of the C_2 coefficient by a factor of 6. The reason for this was the much lower combustion-induced turbulence in HCCI as opposed to conventional CI or SI engine operation. In conventional engines, most of the heat release occurs only in the thin reaction zone of a propagating turbulent flame, which generates a high unsteady gas velocity. Whereas in HCCI, combustion is occurring simultaneously throughout a large volume of space, and much less turbulence is generated than in conventional IC engines.

With these three changes made, the resulting final Chang heat transfer model is generated, with the changes from Woschni highlighted in red:

$$h = \alpha \cdot L^{-0.2} \cdot p^{0.8} \cdot T^{-0.73} \cdot \left[C_1 \cdot 2s\omega + \frac{C_2}{6} \cdot \frac{p - p_{motor}}{\rho}\right]^{0.8} \tag{3.19}$$

The key to "tuning" the coefficients α, C_1, and C_2 in Equation 3.19 is ensure that the integrated area under the HRR curve matches the total amount of energy released by combustion, according to the measured values of fuel flow and unburnt fuel products measured in the exhaust. The correct coefficients will minimize the sum of squared errors (SSE) in the following equation:

$$SSE = \sum_{i=1}^{n_{pts}} \left[\sum_{j=1}^{n_{cyls}} \left[\int_{IVC}^{EVO} \frac{dQ_{c,ij}}{d\theta} d\theta \right] - \dot{m}_{fuel,i} Q_{net,fuel} + \sum_{k=1}^{n_{ufs}} \left[\chi_{ik} M_k Q_{net,k} \right] \frac{\dot{m}_{exh,i}}{M_{exh,i}} \right]^2 \tag{3.20}$$

In Equation 3.20, $\dot{m}$ represents flow rate, Q_{net} represents net heat of combustion, χ represents the measured molar concentration of an exhaust species, and M represents molecular weight. The subscript i represents a particular test condition or data point, j represents a particular engine cylinder, and k represents a particular unburnt fuel species. The experiments in this thesis measured the exhaust concentrations of three unburnt fuel species ($n_{ufs} = 3$): carbon monoxide (CO), methane (CH_4), and total hydrocarbons (THC). The CO and CH_4 emissions have known molecular weight and heat of combustion, but the THC emissions are measured by the use of a flame ionization detector (FID) which produces a value of moles carbon (in unburnt fuel) per mole of exhaust. The dual-fuel testing will assume that the THC emissions are a mixture of gasoline and diesel at the same ratio as the overall fuel flow rate, with the heat of combustion per mole carbon the same as the original fuel.

The experiments in this chapter will simply use Chang's original values for C_1 and C_2 (without dividing C_2 by 6, since this is not an HCCI engine), and only tune the α coefficient to match the test engine. In the following chapter, however, there will be enough data to tune all the coefficients, and alter the model to account for incremental changes in the combustion mode.

3.2 Introducing the Wiebe Function

The Wiebe function serves as a way to simplify the calculated HRR profile of an IC engine into a small set of four parameters, each of which has an intuitive meaning to IC engine researchers. The function, whose history and applications are summarized by Ghojel [20], is a modified log-normal probability distribution that is designed to fit the overall heat release profile of complex chemical chain reactions such as combustion. The function can be parameterized in multiple ways, but this thesis will use the form presented by Garnier et al. [21], which gives these four parameters:

- θ_0**:** The crank angle of ignition, i.e. when the combustion process begins
- $\Delta\theta$**:** Duration of the combustion process, in terms of change in crank angle
- $\hat{Q}$**:** The total heat released by the combustion process, in Joules
- M**:** The "shape factor" of the combustion process, with values <1 for "front-heavy" processes that start quickly then slowly taper off, and values >1 for "back-heavy" processes that build slowly then quickly come to a stop.

It should be noted that the "combustion duration" as defined here by the $\Delta\theta$ parameter is *not* the CA10-90, which is typically used to evaluate the duration of combustion during data analysis. However CA10-90 can and often is used to define the Wiebe function without changing the underlying function. That is not done here for the sake of keeping the ensuing equations as simple as possible with a minimum of constant coefficients.

With the four parameters as defined by Garnier et al. [21], the "integral" form of the Wiebe function gives the total cumulative heat release:

$$Q(\theta) = \hat{Q}\left[1 - e^{\left[-a\left[\frac{\theta-\theta_0}{\Delta\theta}\right]^{M+1}\right]}\right] \tag{3.21}$$

In Equation 3.21, a is a constant set equal to $\ln(1000) = 6.908$, so that Q will reach exactly 99.9% of $\hat{Q}$ at $\theta = \theta_0 + \Delta\theta$. Taking the first derivative of Equation 3.21 gives the "differential" form of the Wiebe function, equal to the HRR:

$$\text{HRR} = \frac{dQ(\theta)}{d\theta} = a\frac{\hat{Q}}{\Delta\theta}[M+1]\left[\frac{\theta-\theta_0}{\Delta\theta}\right]^M e^{\left[-a\left[\frac{\theta-\theta_0}{\Delta\theta}\right]^{M+1}\right]} \tag{3.22}$$

The shape of the Wiebe function in both its integral and differential forms can be seen

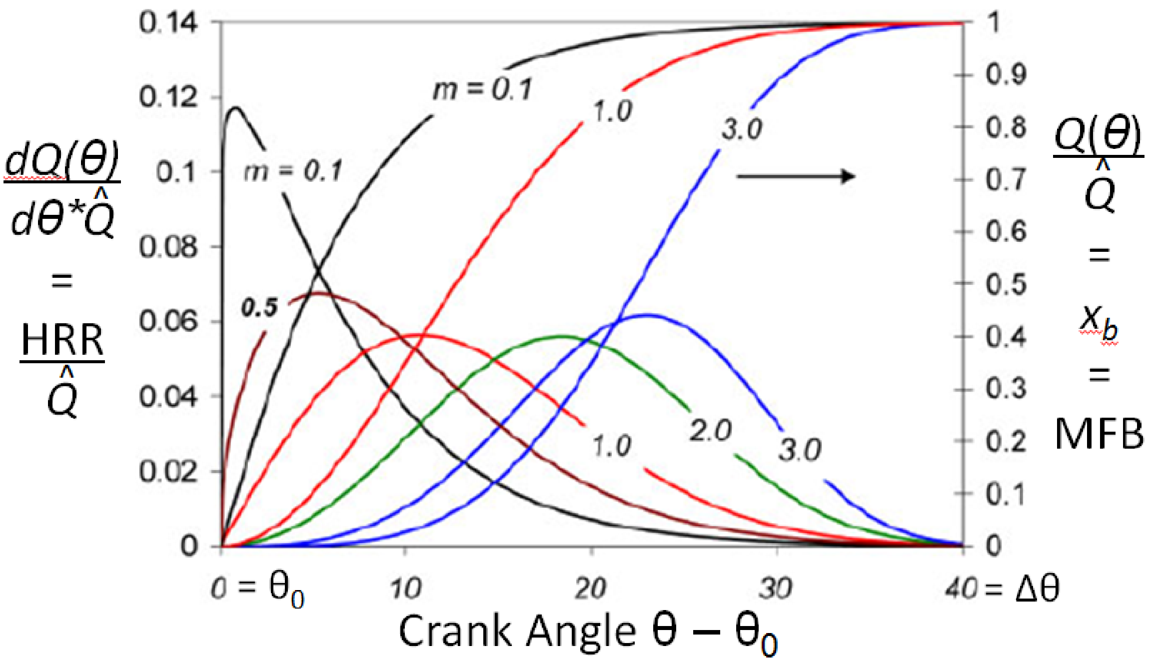

Figure 3.1 Variations in the shape of the Wiebe function taken from Ghojel [20]. The left hand y-axis is for the differential form of the Wiebe function, i.e. the bell-shaped curves, which correspond to the heat release rate (HRR) divided by the total heat release ($\hat{Q}$). The right hand y-axis is for the integral form, i.e. the s-curves, which correspond to the mass fraction burned (MFB). This shows range of possible shapes of the Wiebe function, determined by the shaping coefficient M.

in Figure 3.1. These functions are normalized by dividing by total heat release ($\hat{Q}$), and in doing so recall the log-normal probability distribution, only they are distributing the heat released by combustion rather than probability. The differential form, shaped similar to a log-normal PDF, corresponds to the heat release rate in Joules per crank angle degree ("J/CAD"), or in the normalized case of Figure 3.1 fraction of total heat release per crank angle degree. Likewise the integral form, shaped similar to a log-normal CDF, corresponds to the total cumulative heat release in Joules, or in the case of Figure 3.1 the fraction of total heat release, referred to as "mass fraction burned" (MFB) and represented by x_b. The functions plotted in Figure 3.1 have a fixed $\Delta\theta$ of 40°, and are plotted for several values of M to show how the value of this "shape factor" affects the shape of the Wiebe function.

For SI engines, the heat release profile is often simple enough to be represented by a single Wiebe function. For conventional CI engines, the heat release profile becomes more complex in shape, and it is usually necessary to divide the heat release into two

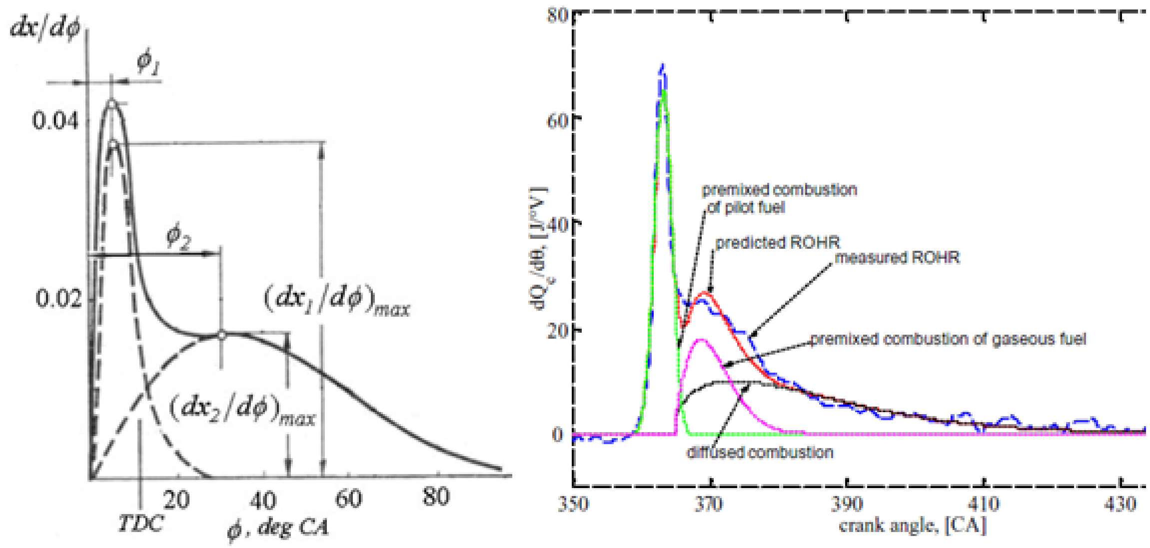

Figure 3.2 Two examples of multi-Wiebe HRR fits: a two-function CDC model from Ghojel [20] (left) and a three-function CDF model from Garnier et al. [21] (right). Both examples show how multiple simple Wiebe functions can sum together to create the full, complex HRR shape.

separate phases, a premixed phase and a diffusion phase. An example from Ghojel [20] is reprinted on the left in Figure 3.2. For dual-fuel CI engines, the heat release profile grows increasingly complex, requiring a third Wiebe function to represent the premixed combustion of a Manifold-injected LRF, such as the natural gas in the dual-fuel engine of Garnier et al. [21], reprinted on the right in Figure 3.2.

Both of these examples show how Wiebe functions can subdivide a complex heat release profile into separate, differently-behaving combustion processes, in order to understand the entire "combustion mode" that is being observed. In Section 2.3, this will be done for all the areas of the MELT diagram, identifying specific combustion mechanisms that correspond to the portion of the heat release represented by each Wiebe function. This serves two purposes: First (as will be seen in the rest of this chapter), it gives a better visual and conceptual understanding of what causes the complex, varied shapes of the HRR profiles seen throughout the MELT diagram. Second, it can be used to construct a smooth, continuous model of HRR wherein the parameters of each Wiebe function are calculated from the engine's operating conditions, to be used in optimizing an engine for dual-fuel operation. This second goal was not achieved during the course of this thesis but there is promise that the fitting routines presented in the next section will eventually be able to tune and validate such a model.

3.3 Fitting Wiebe Functions to Heat Release Rate

The most common method of fitting Wiebe functions to HRR data is to split the HRR into segments at certain "break points," fitting one Wiebe function in between each pair of break points. The best-fit Wiebe function minimizes the error between the Wiebe function and the HRR *in between those two break points*, and does not necessarily fit the HRR outside of that window. Each Wiebe function is then subtracted from the total HRR before the next function is fit between the next pair of break points. This process is depicted in Figure 3.3.

The reason that the Wiebe function is only fit to the HRR within the window between the two break points is because one combustion process can continue to occur after a second process begins. The Wiebe function should not be fit to the sum of *both* processes, only the first. Unfortunately, this means that if two combustion processes are occurring simultaneously, it is hard to tell how much of the HRR stems from one process or the other. One way around this problem is to assume that all combustion processes will produce a HRR profile with a shape similar to the Wiebe function, and find sections of the HRR curve where only one combustion process is contributing to the HRR curve.

For most functions, the Wiebe function is fit to the *start* of a particular process, and extrapolated out to the end of the process. The one exception to this is the final Wiebe function representing a diffusion flame (the final process in conventional diesel combustion). The HRR shape from this process is typically a slow decay to zero that starts while other combustion processes are contributing much more to the HRR. This function is fit to the *end* of the process, while fixing the start of the process to a previous break point. This is seen in the third row of Figure 3.3. This function must be fit *before* any of the functions occurring

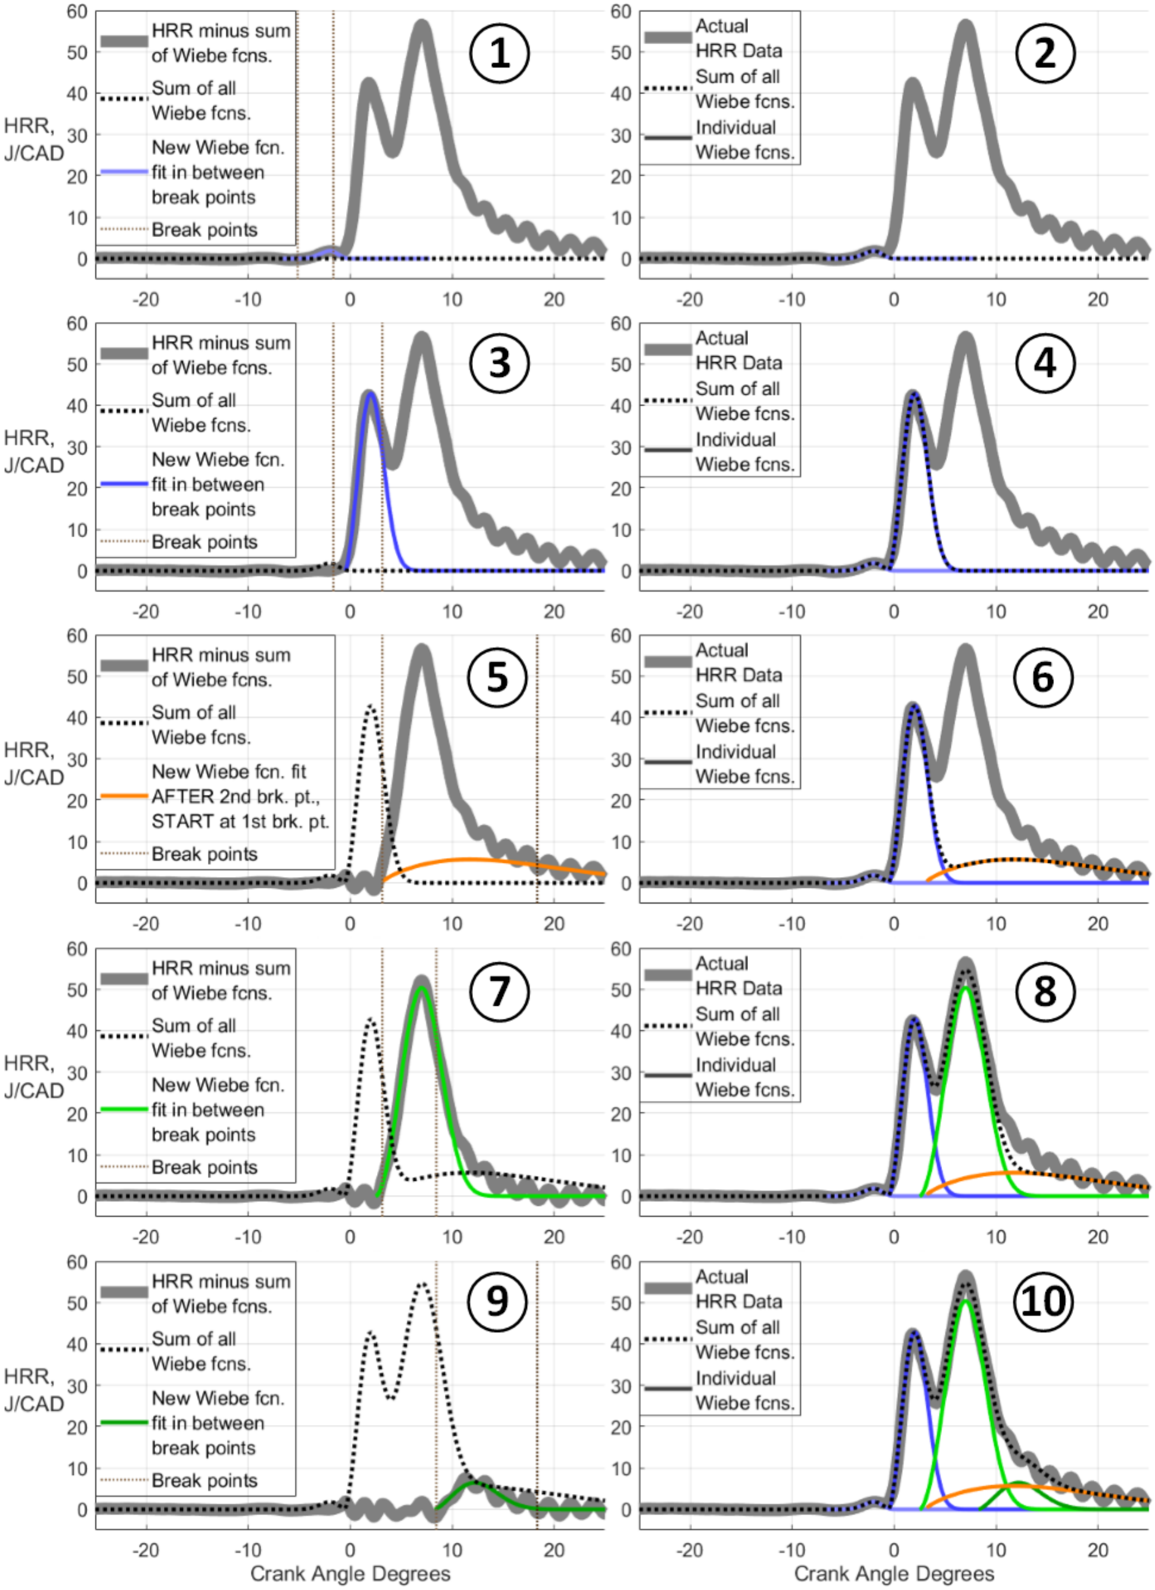

Figure 3.3 Illustration of Wiebe function fitting procedure, described on page 35.

after the earlier break point, so that the sum of all the functions will fit the overall shape of the HRR correctly.

The process of fitting each Wiebe function to the slice of HRR between the break points was a “brute force” process. This was an iterative process with the following steps in each iteration:

- Generate an initial **guess** for each parameter (from the HRR shape during the first iteration, and from the previous iterations for all subsequent iterations).
- Determine the **range** to search over, and the number (**“resolution”**) of evenly-spaced values within that range to test. The range shrinks by one half on every iteration.
- For each possible combination of Wiebe function parameters, calculate the sum of squared errors (SSE) between the Wiebe function and HRR in just the section between the two break points.
- Use the combination of parameters that produces the lowest SSE as the starting guess for the next iteration

The starting **guesses** were estimated using the first break point as the starting θ_0, the distance between the break points multiplied by 1.5 for the starting $\Delta\theta$, the integrated area between the break points multiplied by 1.5 for the starting $\hat{Q}$, and 2.3 as the starting M (this makes the point of 50% occur exactly halfway between θ_0 and $\theta_0 + \Delta\theta$). These initial guesses defined a **range** of possible values centered at the starting guess: +/- 20 CAD for θ_0, +/- 25% of the starting guess for $\Delta\theta$, +/- 12.5% of the starting guess for $\hat{Q}$, and +/- 1 for M.

For example, imagine a Wiebe function is being fit between two break points at -5 and 5 °aTDC, with an integrated area of 100 J under the HRR over that window. The starting guesses for the four parameters would be as follows:

- θ_0 = -5°aTDC (First break point)
- $\Delta\theta$ = 15 (Distance between break points times 1.5)
- $\hat{Q}$ = 150 J (integrated area under curve times 1.5)

- M = 2.3 (puts CA50 right at $\theta = 0$)

If the resolution was, for example, 5, then this would give a matrix $\overline{\overline{P_{ij}}}$ of possible Wiebe function parameters, where subscript i is the guess number (from low to high) and subscript j is the parameter number (1 is θ_0, 2 is $\Delta\theta$, 3 is $\hat{Q}$, 4 is M):

$$\overline{\overline{P_{ij}}} = \begin{bmatrix} \overline{\theta_{0i}} & \overline{\Delta\theta_i} & \overline{\hat{Q}_i} & \overline{M_i} \end{bmatrix} = \begin{bmatrix} -15 & 11.3 & 131 & 0.3 \\ -\mathbf{10} & 13.1 & 141 & \mathbf{1.3} \\ -5 & 15.0 & \mathbf{150} & 2.3 \\ 0 & \mathbf{16.9} & 159 & 3.3 \\ 5 & 17.8 & 169 & 4.3 \end{bmatrix} \begin{array}{l} \text{Starting guess - range} \\ \\ \text{Starting guess} \\ \\ \text{Starting guess + range} \end{array} \tag{3.23}$$

Every possible combination of these parameters is used to calculate a Wiebe function and the SSE with the section of HRR between the two break points. Say for example that the combination of parameters producing the best fit was the combination highlighted in bold above in Equation 3.23. This combination becomes the starting guess for the next iteration, with the range cut in half:

$$\overline{\overline{P_{ij}}} = \begin{bmatrix} \overline{\theta_{0i}} & \overline{\Delta\theta_i} & \overline{\hat{Q}_i} & \overline{M_i} \end{bmatrix} = \begin{bmatrix} -15 & 15.9 & 141 & 0.3 \\ -12.5 & 16.4 & 145 & 0.8 \\ -10 & 16.9 & 150 & 1.3 \\ -7.5 & 17.4 & 154 & 1.8 \\ -5 & 17.8 & 159 & 2.3 \end{bmatrix} \begin{array}{l} \text{Best fit - range/2} \\ \\ \text{Best fit, iteration 1} \\ \\ \text{Best fit + range/2} \end{array} \tag{3.24}$$

Trial and error showed that the best resolution to use was 11 (giving 11^4 = 14,641 possible functions per iteration), with 7 iterations. Increasing the resolution or number of iterations did not significantly improve the fits, it simply increased the computation time. With these settings, a complete Wiebe function breakdown of the HRR could be produced on an average 10 seconds on a Windows 10 PC with an Intel Xeon E3-1241 v3 CPU running MATLAB R2018b. This is clearly not the most efficient way to arrive at the best-fit Wiebe functions, but it proved quick enough that it was unnecessary to develop a computationally optimized fitting routine, at least for the purposes of this thesis.

This work sought to find a “universal” fitting routine that could be applied automatically across the MELT diagram and would produce Wiebe functions with smooth, continuous transitions between all the available combustion modes. The key step in this process was identifying a mathematical feature of the HRR that would reliably serve as a break point separating different combustion processes. Ideally, these break points would be identified universally in all combustion modes, with no mode-specific tuning required.

3.3.1 Finding a Universal Break Point

The first idea for a universal break point was to use inflection points in the HRR, i.e. points where the second derivative of HRR is equal to zero. Specifically, *upward* inflection points were used, where the the second derivative changed from negative to positive. These can be thought of as points where the heat release rate “stops decelerating and starts accelerating,” indicating the start of a new combustion process. Inflection points worked well to identify the necessary break points in the more conventional-looking HRR profiles, but did not yield all of the necessary break points in a the more complex HRR profiles such as the one seen in Figure 3.4, which has many subtle kinks where the different combustion phases blur together. Although using the upward inflection point method still produced a roughly accurate fit to the HRR in the end, it was clear that there were multiple phases of combustion occurring in between the second and third break points.

At this point, it was observed that the inflection points often occurred near-simultaneously with peaks in the *third* derivative, which is often referred to metaphorically as the “jerk” of a function. These “jerk peaks” occurred not just near the previous inflection points but also at the more subtle kinks in the HRR profile. This allows break points to be assigned at places

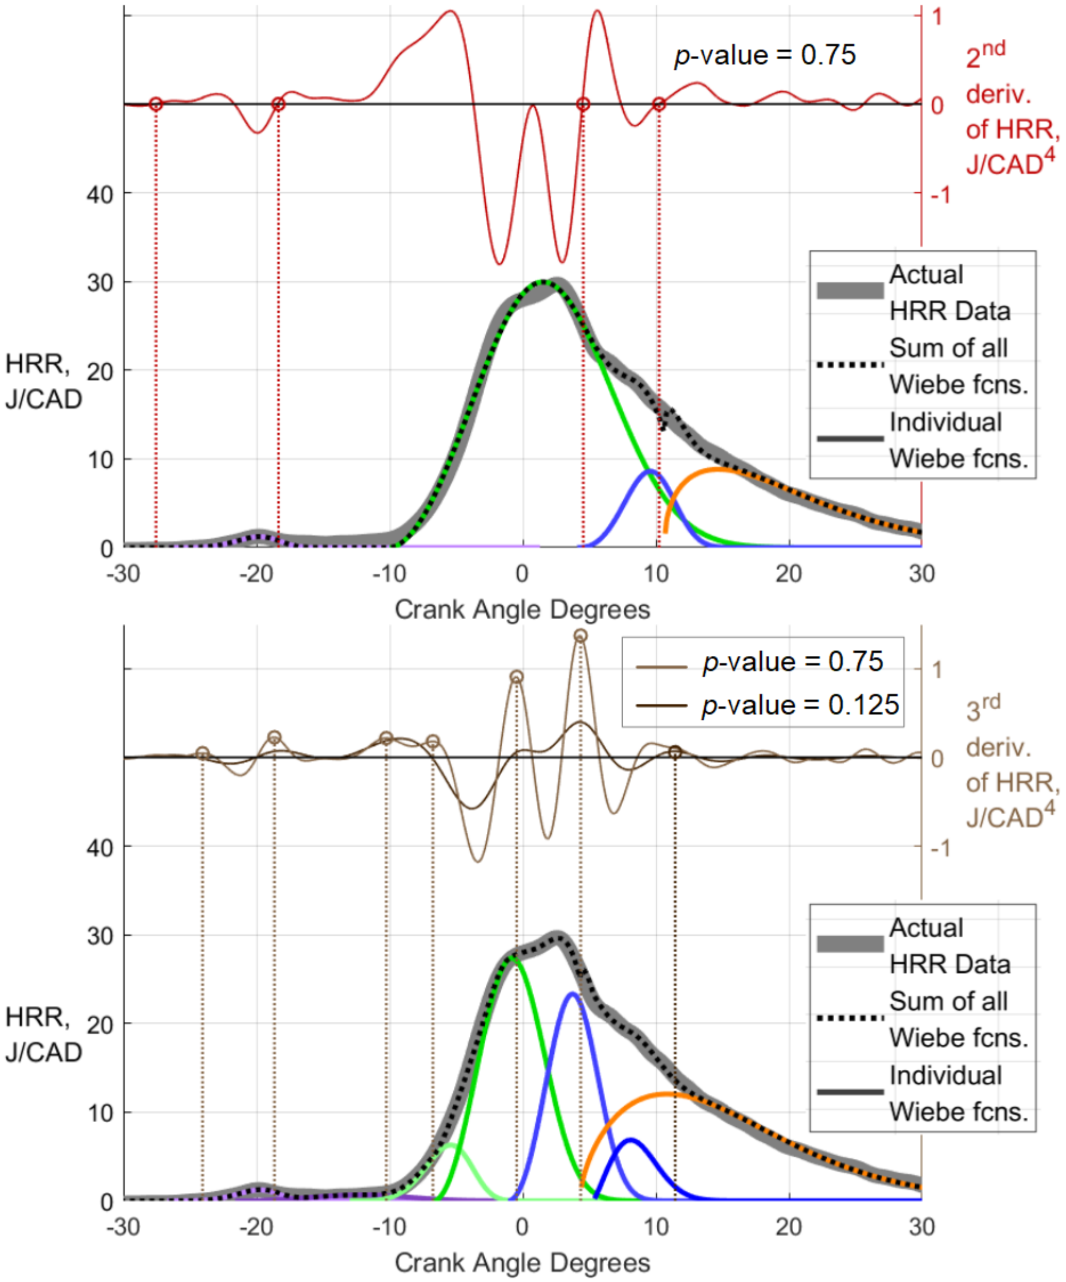

Figure 3.4 Using derivatives of the HRR to systematically find "break points" (indicated by circles). This shows how inflection points (where the second derivative of HRR crosses zero, shown in the top plot) do not sufficiently capture the break points needed to model a complex HRR pattern, but "jerk peaks" (peaks in the third derivative, shown in the bottom plot) capture all the necessary break points. The effectiveness of this method hinges on finding an appropriate *p*-value to use in the cubic smoothing spline (see Equation 3.26) applied to each subsequent derivative of HRR. A high *p*-value (low smoothing) of 0.75 is needed for most break points, but for the final break point, a much lower value (high smoothing) of 0.125 is needed given the acoustic fluctuations in cylinder pressure that occur during the final stages of combustion. These values will change depending on the engine and pressure transducer used.

like the point seen near -10° in Figure 3.4, where the HRR is “accelerating” (i.e. the second derivative is positive) all the way from -18° to -4°. Although there are no upward inflection points, there is a noticeable upward bend at -11°, where the HRR is being “jerked” upward as a new, more intense combustion process is starting. Here, there is a clear “jerk peak” that can be used to assign the break point, and doing this allows for a much more detailed HRR fit that captures all the subtleties of the HRR profile.

Of course when using such a high-order derivative calculated from noisy experimental data, it is critical to optimize the methods of numerical differentiation. Getting the right amount of fidelity in the third derivative is key to obtaining all the break points needed without producing “false” break points that are the result of noise in the data. The numerical differentiation process used was a standard five-point stencil, using data sampled at a step size h of 0.1 CAD:

$$f(\theta)' = \frac{-f(\theta+2h)+8f(\theta+h)-8f(\theta-h)+f(\theta-2h)}{12} \tag{3.25}$$

At each level of differentiation, it was necessary to apply at least a small amount of smoothing to the curve to account for the fluctuations in the original pressure trace, which become more amplified at each level of differentiation. This was done by applying the MATLAB *csaps.m* cubic smoothing spline routine. This method produces a separate cubic polynomial function (“spline”) in between each consecutive pair of data points. Since the cubic polynomial has four degrees of freedom, the splines are constrained such that both the value and the slope of the two splines on either side of a data point are equal at that particular data point. Thus when all the splines are combined into a single function $f(\theta)$, this will be a twice-differentiable function, which is important for evaluating the “smoothness” of the fit on the second derivative. In *csaps.m*, the spline fit $f(\theta)$ of the original HRR data $y(\theta)$ will minimize the following value:

$$\left(p\sum_{\theta=\text{IVC}}^{\text{EVO}} h[y-f]^2\right)+\left([1-p]\int_{\theta=\text{IVC}}^{EVO}\left[\frac{d^2f}{d\theta^2}\right]^2 d\theta\right) \tag{3.26}$$

In this value, h is the step size of θ and p is the degree of fidelity to the original data. The term on the left is the "error" term, i.e. the deviation of the spline fit from the original data, and the term on the right is the "smoothness" term, penalizing sudden changes in the slope of the fit. The value of p can be between 0 and 1, with 0 producing a flat, straight line and 1 producing an exact fit to the original data. The p-values used in this fitting routine depended on the amount of detail required in the Wiebe function breakdown. To get all the major, most obvious break points, a p-value of 0.125 could be used, with unnecessary peaks filtered out. For the detailed breakdowns presented in the next section, only slight smoothing was applied, at a p-value of 0.75. This did not shift the locations of the major peaks greatly, while capturing other, minor peaks that were not apparent with the "heavy" smoothing, such as the peak at -7°.

However, this method of finding break points was not perfect, since it would sometimes find break points in one cylinder's HRR that were not apparent in another cylinder. This led to the development of a "semi-automatic" Wiebe function fitting routine that checks with the user whether to reject certain break points, and whether to add break points at any point along the way. This code also queried the user on which combustion type each Wiebe function would be assigned to, and how many points to "skip over" when fitting the final Wiebe function. An example command window snapshot showing this process is seen in Figure 3.5.

This development of this Wiebe function fitting routine was concurrent with the development of a consistent Wiebe function breakdown where each individual function is assigned to a particular combustion process. Applying this routine to HRR's from incrementally changing fuel injection ratios shows the resulting incremental shifts between different combustion modes.

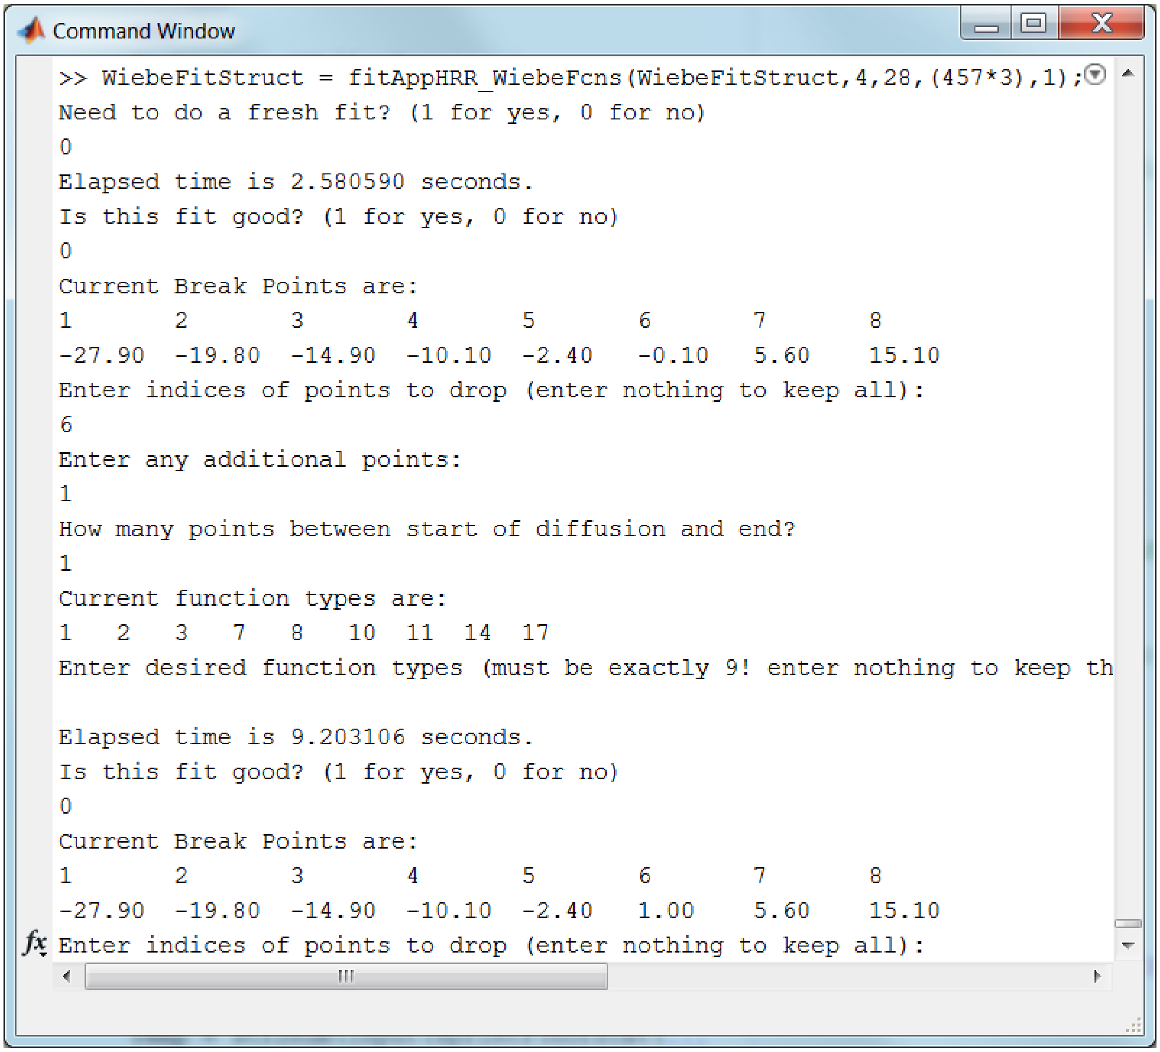

```
>> WiebeFitStruct = fitAppHRR_WiebeFcns(WiebeFitStruct,4,28,(457*3),1);
Need to do a fresh fit? (1 for yes, 0 for no)
0
Elapsed time is 2.580590 seconds.
Is this fit good? (1 for yes, 0 for no)
0
Current Break Points are:
1       2       3       4       5       6       7       8
-27.90  -19.80  -14.90  -10.10  -2.40   -0.10   5.60    15.10
Enter indices of points to drop (enter nothing to keep all):
6
Enter any additional points:
1
How many points between start of diffusion and end?
1
Current function types are:
1   2   3   7   8   10  11  14  17
Enter desired function types (must be exactly 9! enter nothing to keep th

Elapsed time is 9.203106 seconds.
Is this fit good? (1 for yes, 0 for no)
0
Current Break Points are:
1       2       3       4       5       6       7       8
-27.90  -19.80  -14.90  -10.10  -2.40   1.00    5.60    15.10
fx Enter indices of points to drop (enter nothing to keep all):
```

Figure 3.5 Sample Matlab command window showing the "semi-automatic" Wiebe function fitting code. The user is first asked if the initial fit stored in memory is good ("Need to do a fresh fit?"), and the user responds "0" for no. The code then tries an automatic fit, indicates how long it has taken to do, and asks the user if the automatic fit is correct. If not, the code displays the locations of current break points, and asks the user if it should remove any break points. The user tells the code to drop the sixth break point at -0.1 CAD, and replace it with a new break point at 1.0 CAD. Then, the user is asked how many break points to "skip over" when doing the fit of the final diffusion phase, and which combustion type to assign to each of the functions. This process repeats until the user is satisfied with the fit.

3.4 Analyzing Combustion Modes with Wiebe Functions

For this initial exploration and illustration of combustion modes, data was taken at a fixed speed (2300 RPM) and two loads: A "low load" of 4 bar BMEP (Brake Mean Effective Pressure) and a "high load" of 10 bar BMEP. These experiments explored as many points on the MELT diagram as possible, in 20% increments, and covered the entire diagram at the "low load" condition chosen. For this initial illustration, other engine parameters were varied as necessary to achieve good examples of all combustion sub-processes (a more rigorous and systematic test matrix will be investigated in Chapter 4).

Before proceeding into the Wiebe function breakdowns, it should be noted that the functions are assigned to different types of combustion based solely on comparison to various conceptual models found throughout the literature. In reality, these different combustion phases overlap in complex manners, and the segments of HRR that each Wiebe function is fit to do not overlap. The Wiebe functions can project forwards and backwards past the break points, but only in order to better fit this particular segment of the HRR. Despite this, the Wiebe functions still give a good way to "color in" the HRR to develop both a conceptual understanding and also a continuous empirical model, which only requires that the functions sum up to an accurate HRR even if they do not perfectly accurately represent the combustion stages that they are being labeled as.

The figures in the following section will each feature two examples of a multi-Wiebe function fit, from either the same point or two nearby points on the MELT diagram, shown on the right. Also on the right of these figures will be a color key to the combustion phases that each Wiebe function represents. These figures will proceed in a clockwise "tour" of the MELT diagram, adding new combustion phases into the key as they appear.

3.4.1 CDC (Conventional Diesel Combustion)

This tour starts with the easiest combustion mode to break down: conventional diesel combustion (CDC). This mode is the most heavily studied, and the most widely accepted conceptual model of CDC was devised by Dec [22]. When fuel is first injected, there is a short ignition delay of 5-10°, during which a rich air/fuel mixture develops at the head of the fuel jet. This mixture then rapidly autoignites, creating a sharp spike in the HRR. The Wiebe function corresponding to this phase (Rich Premixed Autoignition) is colored in blue in the following figures, starting with Figure 3.6, which features two examples of CDC. Since this mixture is rich, there will of course be incomplete combustion, and a diffusion flame will develop around the fuel jet. The Wiebe function corresponding to this phase (Diffusion Flame) is colored orange. There is also a “dummy” Wiebe function in gray that is fit to whatever deviation occurs from zero in the HRR prior to the onset of combustion, due to the imperfect heat transfer model used, and does not represent any combustion.

If the fuel injection is still ongoing at this point, a rich premixed flame will also develop at the point where the incoming rich fuel/air mixture meets the zone of rich premixed autoignition. At the low load (4 bar BMEP) where these initial tests were conducted, the injection was short enough to prevent this from occurring, and there are only two Wiebe functions needed, as seen in the top of Figure 3.6. At the high-load case of 10 bar BMEP, shown at the bottom of Figure 3.6, there is a third Wiebe function in dark blue corresponding to this relatively stationary rich premixed flame. This flame will occur simultaneously with the diffusion flame, but will end much more quickly once the injection ends and no more fuel is incoming. For this reason, the diffusion flame Wiebe function is fit to the section of HRR that occurs after the final break point, but the θ_0 of this function is fixed to the break point after the rich premixed autoignition.

Just to note the structure of the “combustion phases” key: the Wiebe functions that respond to autoignition occurring in multiple sites near-simultaneously throughout a volume of space (“sequential autoignition”) are indicated by thick lines in the “Combustion Phases” key in the following figures, while the functions corresponding to thin reaction zones (flames) are indicated by thin lines in the key.

The Dec model has been expanded in recent years by Musculus et al. [23], breaking the rich premixed autoignition down further into a first and second stage, with the first stage producing simpler incomplete combustion products such as H_2CO, H_2O_2, and CO, and the second stage producing more complex polyaromatic hydrocarbons (PAH) and soot. This two-stage behavior is most pronounced in low-temperature conditions, such the condition

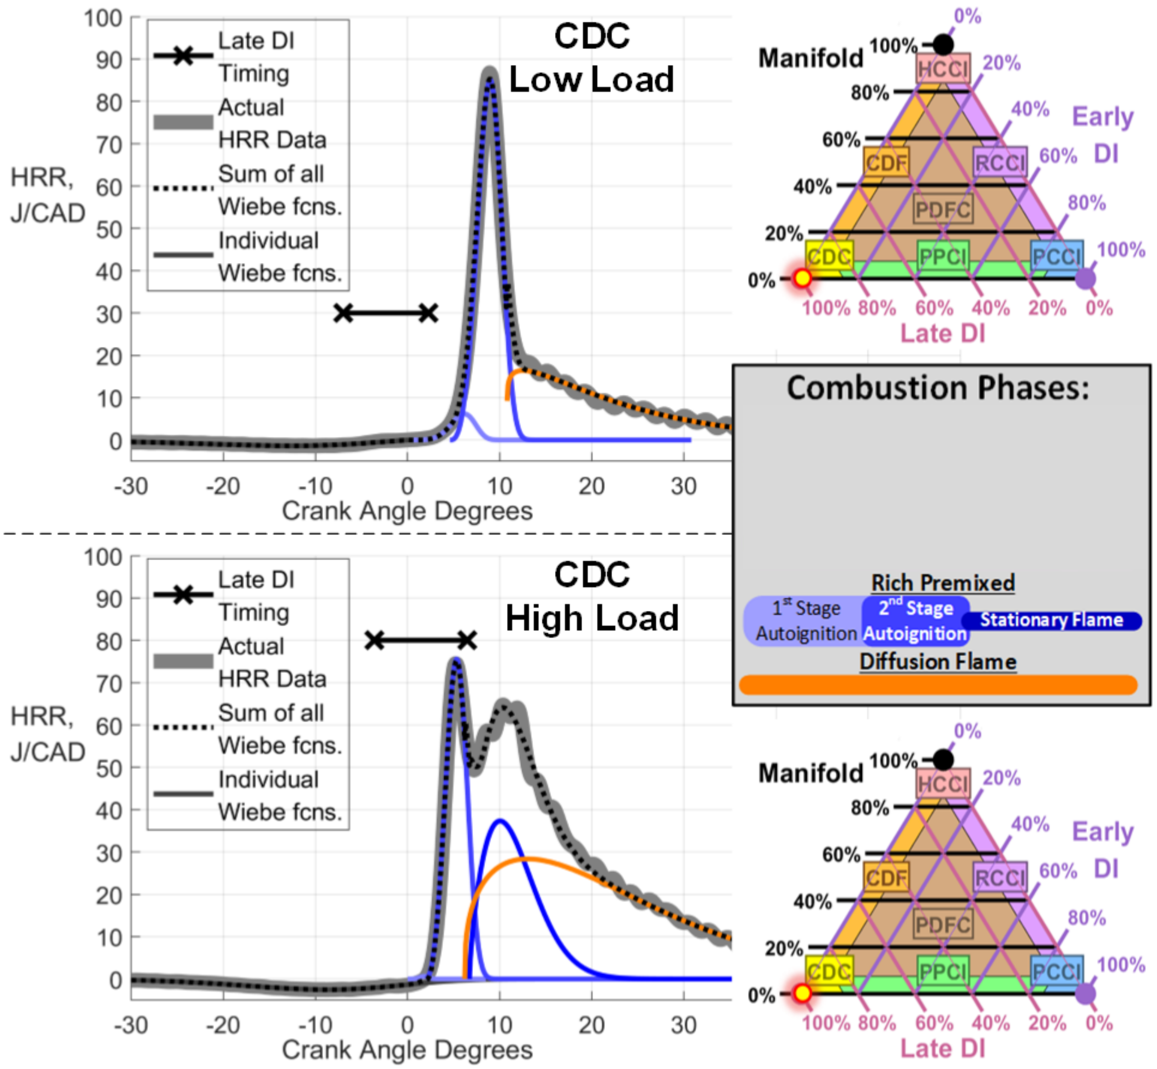

Figure 3.6 Wiebe function breakdowns of CDC, at low load (top) and high load (bottom). The individual Wiebe functions are shown by colored solid lines, with the colors corresponding to the different combustion phases in the key at the right. The thick grey line indicates the actual HRR data that was used to produce the fit, and the dotted line indicating the sum of the Wiebe functions should fall entirely within the gray line if the fit is sufficient. The X's indicate the start and end of the injection command signal (the actual injection timing may lag behind this somewhat). This shows how at low load (top), the injection ends before the rich premixed autoignition, but at high load (bottom), the injection continues through the autoignition, producing an additional phase of a rich premixed stationary flame.

shown in the bottom of Figure 3.7, which featured a very late SOI timing at -3°aTDC. This late timing caused the ignition delay to last well past TDC, and temperatures at the onset of ignition were reduced compared to the conditions in Figure 3.6. The shape of the autoignition phase is kinked and one Wiebe function cannot fit the entire shape well, but

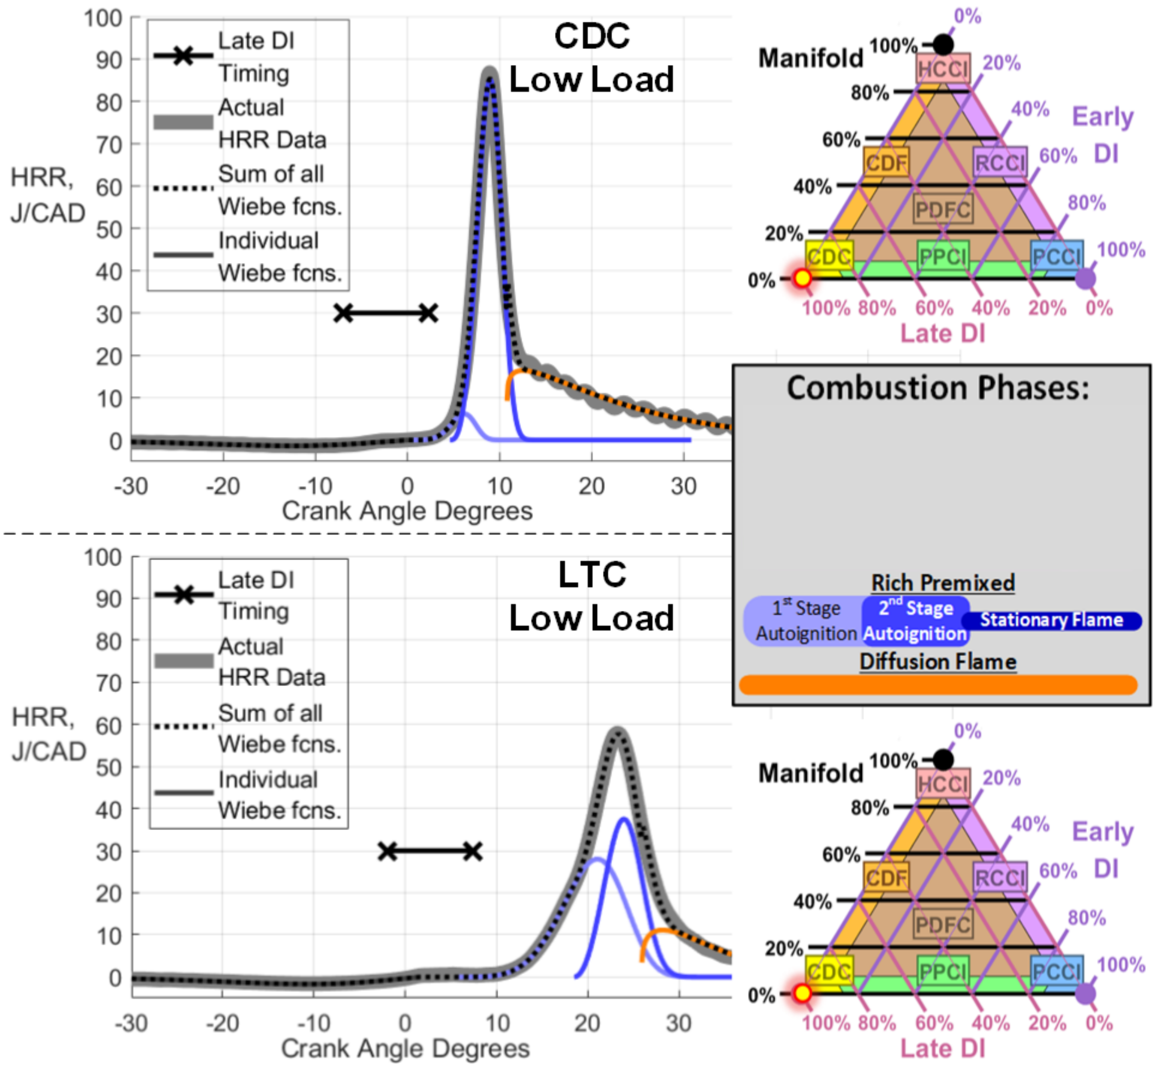

Figure 3.7 Wiebe function breakdowns of CDC and LTC, showing the large 1st stage autoignition in LTC. LTC is not shown on the MELT diagram, since it overlaps with CDC, and uses differences in injection pressure, timing, and EGR levels to maximize the amount of cooler, leaner, lower-sooting autoignition, rather than changing the types of injection used.

two Wiebe functions for the two stages of autoignition fit it quite well.

In the original CDC data point from Figure 3.6, repeated at the top of Figure 3.7, the first stage of rich premixed autoignition is much smaller. This fits well with the conceptual model that Musculus et al. [23] constructed for typical CDC operation, and also fills in what would be a small gap between the Wiebe function summation and the original HRR in this area. The conditions that feature a large 1st stage of rich premixed combustion can be labeled as "low temperature combustion" (LTC), which occurs on the same spot on the

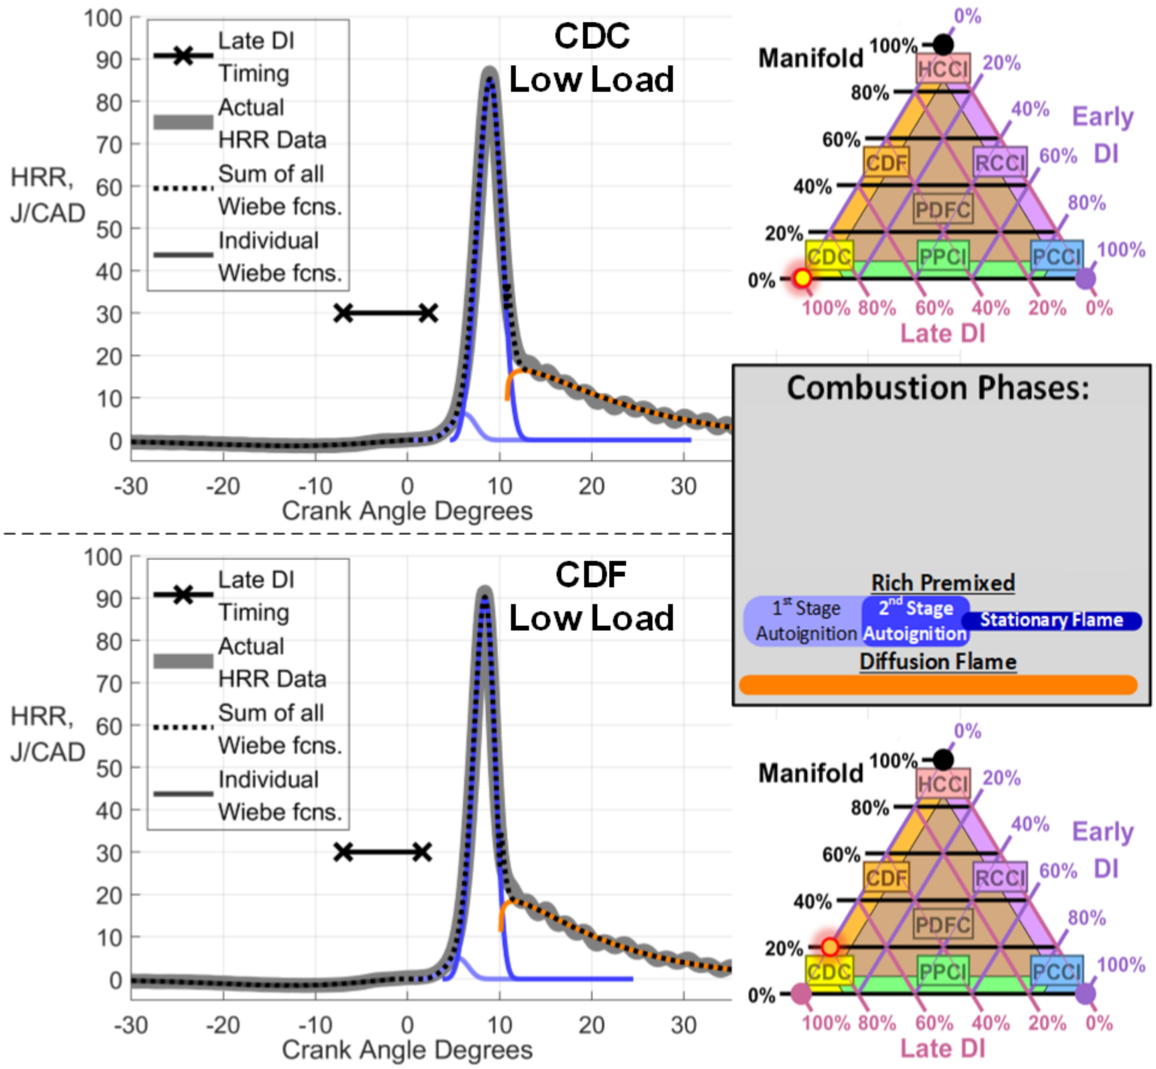

Figure 3.8 Wiebe function breakdowns for CDC (top) and "light" CDF (bottom). The "light" CDF case has only 20% Manifold fuel, and as such there is no noticeable difference in the HRR shape and no new combustion phases are introduced to model the HRR.

MELT diagram as CDC.

3.4.2 CDF (Conventional Dual-Fuel)

Moving on to conventional dual-fuel operation (CDF), in which the Late DI of diesel is partially replaced by Manifold-injected gasoline, there is initially not much of a difference in the shape of the HRR profile with just 20% Manifold fuel, as seen in the bottom of Figure 3.8. The same Wiebe functions that were used in CDC still fit the CDF HRR at this point.

The interpretation of the HRR behavior in CDF is aided greatly by the optical investigations of Dronniou et al. [24]. That work tested the dual-fueling of DI diesel surrogate (a blend of *iso*-octane and *n*-decane) with two Manifold fuels: intake fumigated (IF) methane and PFI *iso*-octane (as a gasoline surrogate). This work reveals that when the manifold fuel is at a relatively low equivalence ratio (and has adequately low reactivity), it will burn slowly through a diffusion process, without flame propagation through the premixed charge. This causes the shape of the HRR to be very similar to CDC, with the addition of premixed gasoline on the air side of the flame increasing the amount of heat released in the diffusion phase, and no obvious changes in the shape of the HRR.

However, when the percentage of Manifold fuel is increased to 60%, the HRR begins to take on a different shape, seen in the bottom of Figure 3.9. The first stage of rich premixed autoignition begins to become more pronounced and the ignition delay longer, possibly due to the entrainment of the gasoline into the rich fuel/air mixture. Despite the longer ignition delay, the spike corresponding to the second stage autoignition becomes smaller, likely due to the lower amount of Late DI fuel injected. But more notably, a “hump” begins to form subsequent to the rich premixed autoignition, on top of the diffusion phase. When this same hump was observed in the optical experiments of Dronniou et al. [24], it corresponded to an observation of flame propagation through the lean premixed charge. Since the manifold fuel in the experiment of Dronniou et al. [24] was methane (with an octane number of 120), there was a wide range of equivalence ratios at which this flame propagation could be produced. However, the 96-octane gasoline used here appears to have a small envelope in which this shape of HRR can be produced, with 60% Manifold/40% Late DI being the only point on the MELT diagram where this shape was produced in the initial experiments, and even then only with the right combination of Late DI timing and EGR.

When a higher concentration of gasoline was used at the 80% Manifold/20% Late DI point, the “hump” in the HRR profile became a second sharp “spike” immediately following the rich premixed autoignition, as seen in the bottom of Figure 3.9. There was

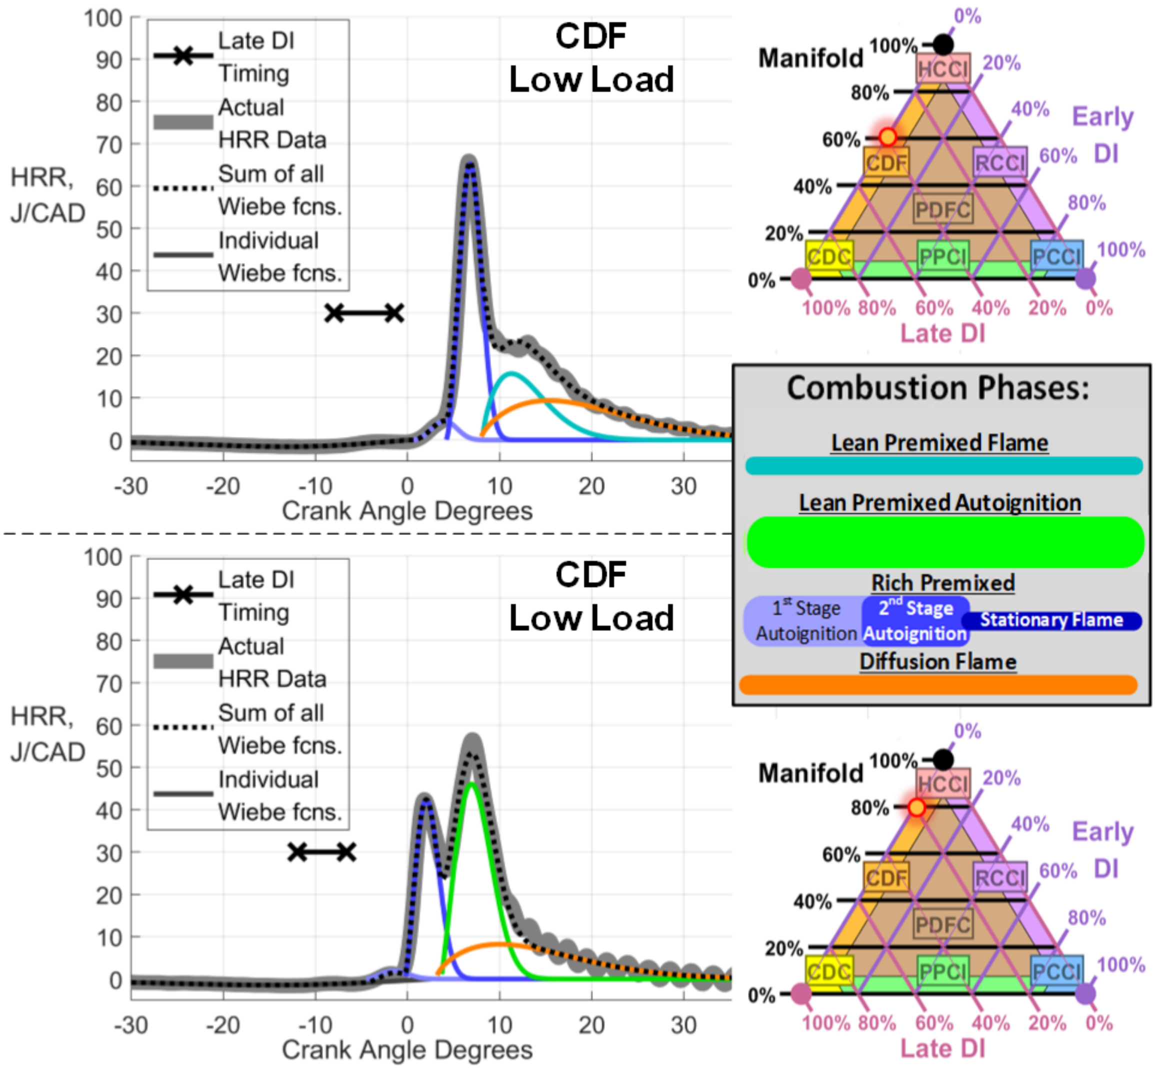

Figure 3.9 Wiebe function breakdown of "Heavy" CDF. The top case with 60% Manifold fuel appears to show lean premixed flame propagation, while the bottom case with 80% Manifold fuel appears to show lean premixed homogeneous autoignition.

also a significant decrease in the apparent amount of diffusion combustion, with the long tapering tail becoming shorter. When similar HRR patterns were occurring in the optical study of Dronniou et al. [24], it appeared that there was a rapid progression of combustion throughout the entire volume of the combustion chamber. Combustion luminosity increased simultaneously throughout the whole volume of the combustion chamber, including areas far removed from the premixed flame. This would indicate a process of sequential autoignition, with the heat released by the initial rich premixed autoignition triggering a subsequent series of lean premixed autoignition sites in the unburnt mixture, before a premixed flame could propagate to these sites. This process of end gas autoignition is essentially the same as the knocking phenomenon observed in SI engines, but unlike SI knock it does not appear to be harmful so long as it does not occur too rapidly or too early in the cycle. In fact, it is actually advantageous for the purposes of reducing combustion duration, and maximizing the area under the P-V curve for maximum indicated efficiency. There are, however, noticeable increases in the high-frequency pressure fluctuations, which are transmitted through to the calculated HRR in the diffusion phase in Figure 3.9.

This thesis will consider all of the lean premixed combustion to be of the sequential autoignition variety from here on out. Although there may be some flame propagation at select conditions, it will be impossible to differentiate between flame propagation and autoignition if they are occurring simultaneously, so this thesis will not make that distinction to make the HRR model more simple.

3.4.3 HCCI (Homogeneous Charge Compression Ignition)

This process of sequential autoignition in the lean premixed charge essentially forms the entirety of HCCI, one of the first "advanced combustion modes" proposed back in the 1980's by Thring [25]. HCCI can be highly efficient but very difficult to control, with satisfactory performance only being found within a very tight window of operating parameters. This makes sense when considering that the combustion chamber in a pure HCCI mode with 100% Manifold fuel is a perfectly-stirred reactor, with the combustion progress controlled entirely by chemical kinetics. As a result HCCI may be better suited to form part of a "dual-mode" engine operating strategy that uses other modes outside of the ideal window for HCCI [26]. One might think that HCCI is simple enough to be modeled with a single Wiebe function, but in reality, there are several non-idealities that make it hard to fit a single Wiebe function to the HCCI HRR.

In the particular case of this engine, and the speed/load condition used for this exploration

(2300 RPM/4 bar BMEP), it was necessary to use a high level of EGR (40%) to lower the equivalence ratio and raise the intake temperature to the point where stable HCCI could be achieved, shown in Figure 3.10. There was also a very high cylinder-to-cylinder imbalance, which made it necessary to slightly increase the fueling rate to cylinders 1 and 2 while reducing the fueling rate to cylinders 3 and 4 (the HRR from cylinder 1 is shown in these examples). Even when just considering a single cylinder, there is bound to be significant thermal stratification within the cylinder that will, even if the fuel and air are perfectly evenly mixed, cause autoignition to occur at different times in different parts of the cylinder and "spread out" the heat release rate [27].

A good illustration of the impact of this thermal stratification on HCCI can be found in the work of Dec et al. [28]. This study measured the vertical profile of chemiluminescence in an HCCI engine, and revealed that there were very differently-shaped profiles at different points of time. In earlier stages of combustion, when most of the heat release was occurring, the intensity was greatest farthest from the walls of the combustion chamber, where the temperatures were greatest. This produced a bell-shaped combustion intensity profile, peaking about halfway between the cylinder head and the piston surface. During the period of peak heat release, the profile became flat at the peak, with a relatively homogeneous level of intensity from about one quarter to three quarters of the distance between the cylinder head and piston. However during the later stages, when the overall level of heat release was much lower, this profile changed into an "M" shape, with two peaks near the surfaces of the cylinder head and piston. This explains the slow taper of the HCCI heat release profile, which is a different process than the diffusion phase producing the slow taper of all the previous heat release. The combustion in these "cold spots" needs to be fit with its own separate Wiebe function in order to produce a precise model of the HRR, as is done in Figure 3.10.

However there is still a poor fit at the very beginning of the HRR curve, before TDC. There appears to be a very slow initial onset of heat release before the main "spike," which

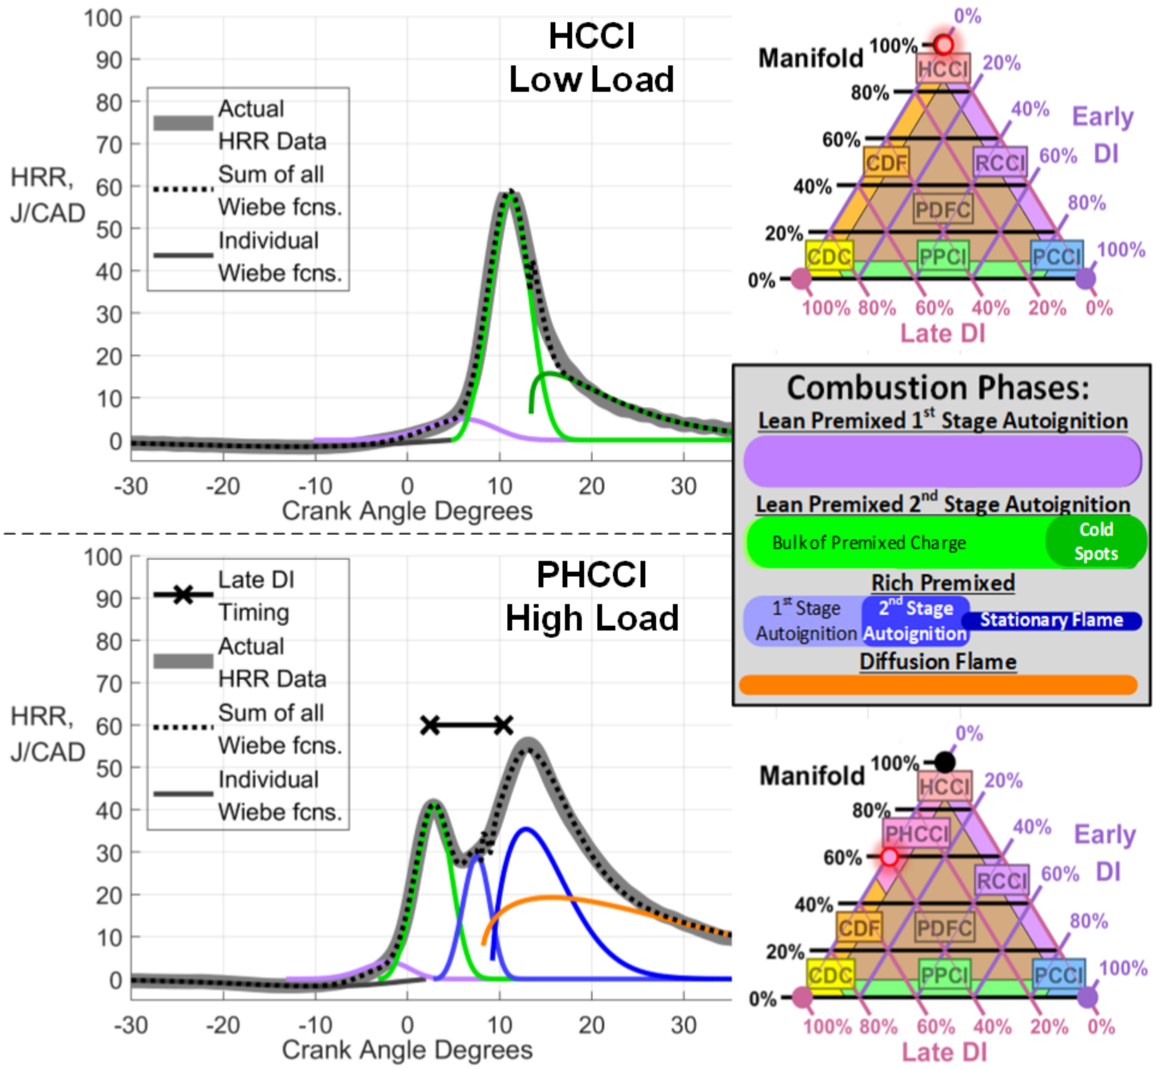

Figure 3.10 Wiebe function breakdowns of HCCI (top) and PHCCI (bottom). True HCCI (top) with 100% Manifold fuel consists entirely of lean premixed autoignition and can only be achieved at relatively low loads, while Partial HCCI or PHCCI consists of HCCI followed by a mixing-controlled CDC and can only be achieved at relatively high loads. PHCCI is fundamentally different from CDF, in which the two processes occur in the opposite order, and a new section has been created on the MELT diagram to differentiate the two.

must also be fit with its own separate Wiebe function, as is done in the top of Figure 3.10. Dec et al. [28] observed this same initial "ramp" in their HRR profiles, but did not observe any corresponding chemiluminescence. This indicates a non-luminous first stage of a two-stage autoignition process that, similar to the rich premixed autoignition, produces intermediate products that will proceed towards CO_2 and H_2O in the second stage. The long gap between the beginning of the first and second stages is a well-known feature of HCCI but highly sensitive to the fuel composition and intake air temperature, and small amounts of certain additives can be used to significantly shorten this delay [29]. For the experiments in this thesis, neither of these parameters were under control, and the unavoidable long delay before the second stage in the HCCI condition in Figure 3.10 resulted in poor TE.

This three-Wiebe function model will suit the purposes of this thesis, but it should be mentioned that much more nuanced and computationally intense models of HCCI heat release rates have been developed in recent years by research groups such as Kodavasal et al. [30] and Singne et al. [31].

3.4.4 PHCCI (Partial HCCI)

Having introduced HCCI, is important to look back to the CDF region of the MELT diagram, from which only low-load conditions have been presented so far. The bottom of Figure 3.10 shows a high-load condition with 60% Manifold fuel and 40% Early DI. In this condition, there is a sharp spike in the HRR profile beginning before the start of the Late DI. For this reason, the initial spike must come from autoignition of the Manifold fuel alone, i.e. HCCI. This is fundamentally different from CDF, where the lean premixed autoignition process happens only *after* the rich premixed autoignition of the Late DI. For this reason, a new combustion mode will be added onto the MELT diagram that only applies in high loads, called “Partial HCCI” or PHCCI.

In PHCCI, it is likely that the same final stage of autoignition of cold spots is still occurring after the autoignition of the bulk charge. However, this stage becomes “buried” below the Late DI process and it is impossible to create a Wiebe function for it with the methods used here. Similarly, it is only possible to see the main lean premixed autoignition with the Late DI retarded to start afterward. If the late DI is advanced at all, it will become impossible to distinguish the lean and rich premixed autoignition phases. For this reason, there may be some lean premixed autoignition occurring in the previous CDF conditions simultaneously with the rich premixed autoignition, blurring the line between CDF and

PHCCI.

3.4.5 RCCI (Reactivity-Controlled Compression Ignition)

Moving away from HCCI in the other direction introduces the Early DI, which will cause the premixed fuel to become at least somewhat stratified in terms of both equivalence ratio and fuel reactivity. This was first suggested by Inagaki et al. [32] as a way to better control HCCI, and later refined and christened as "RCCI" by Kokjohn et al. [33], with an example from this thesis in Figure 3.11. This produces the same sequential autoignition process that occurs in HCCI, but with a more kinked HRR shape. Autoignition first occurs in the areas of highest reactivity, with the heat released by these initial events triggering additional autoignition sites in areas of progressively lower reactivity. This chain reaction eventually reaches a central zone of relatively homogeneous reactivity, with a large "spike" indicating bulk autoignition. Like HCCI, this is followed by a slow taper corresponding to the slower combustion in the coldest parts of the combustion chamber, but this cold spot autoignition phase has a much lower magnitude than in HCCI. This is likely because the Early DI is introducing high-reactivity fuel into these cold spots, reducing the resistance to autoignition in this area.

The RCCI process is well illustrated in the optical work of Kokjohn et al. [34], where the transition from stratified to homogeneous autoignition can be clearly observed. During the initial stages, autoignition sites begin in the outer squish zone of the combustion chamber, where PLIF fuel tracer studies have located the highest concentrations of diesel from the Early DI. At first, new autoignition sites only appear slightly ahead of the burnt gas zones, progressing down the gradient of fuel reactivity and equivalence ratio (ϕ). But then, there is a transition, with autoignition sites appearing somewhat randomly throughout the bulk of the charge in the bowl region, where the reactivity and ϕ are relatively homogeneous. This transition can be attributed to the kink in the HRR profile observed in Figure 3.11, which necessitates breaking the lean premixed autoignition further into multiple Wiebe functions

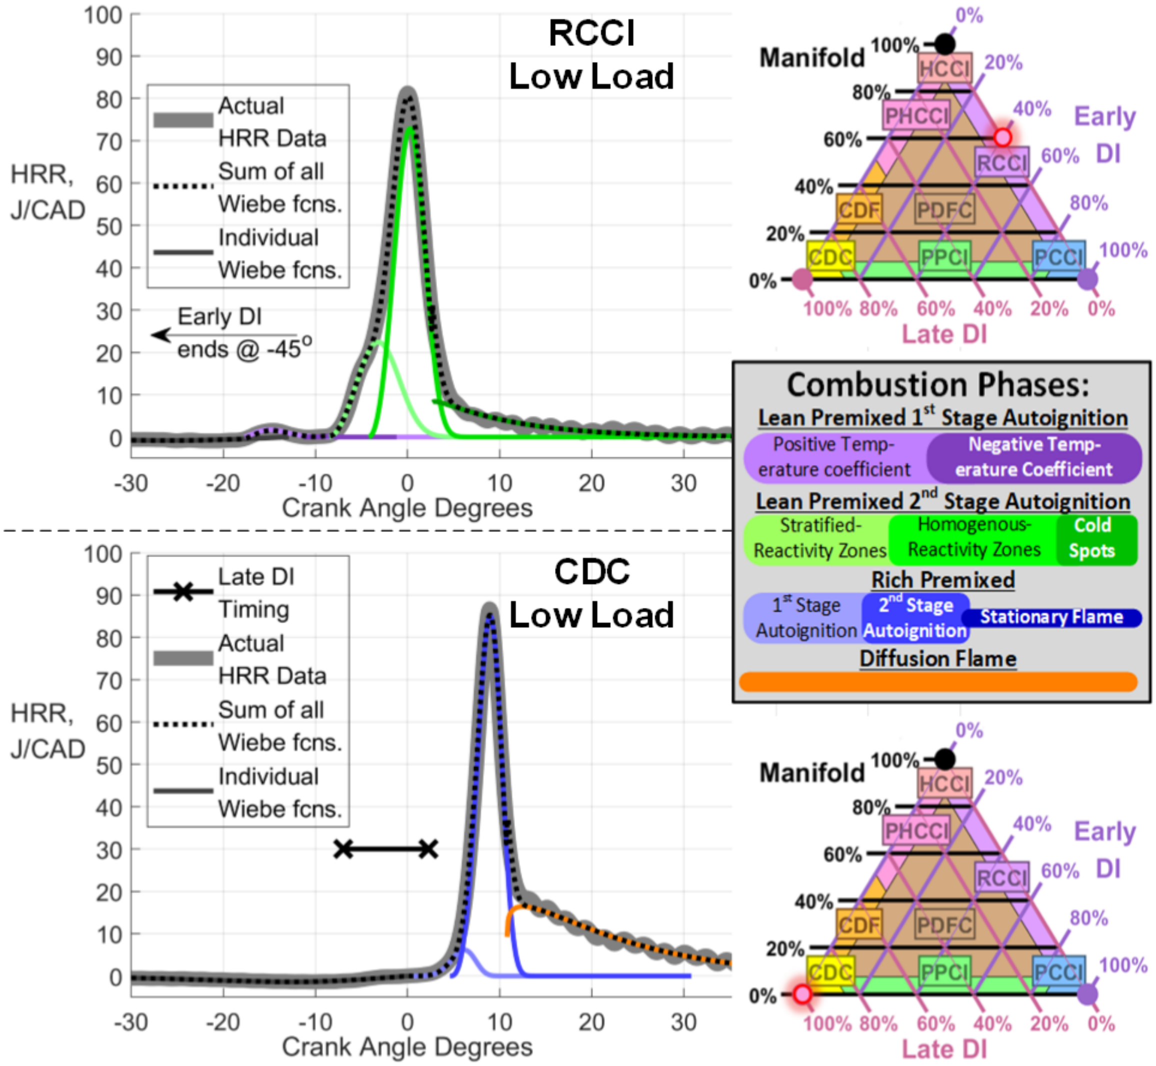

Figure 3.11 Wiebe function breakdowns of RCCI (top) and CDC (bottom). RCCI consists of two stages of autoignition with a period of negative temperature coefficient in between them, due to the premixed diesel from the Early DI. The Early DI also creates a stratification of reactivity in the premixed fuel, which divides the second stage of autoignition in to a stratified phase and a homogeneous phase. This process can occur earlier in the cycle (close to TDC) than CDC without generating diesel knock.

to maintain a good fit. The first function represents autoignition in the more stratified zones, which is more staggered and thus relatively less intense, and the second function represents autoignition in the more homogeneous zones, where it will occur relatively simultaneously and release more heat all at once.

Another combustion process that must be accounted for in RCCI is the presence of a smaller first stage of the lean premixed autoignition, seen at around -15°in the top of

Figure 3.11. This first stage is often referred to as "low temperature heat release" (LTHR) as opposed to the much more significant "high temperature heat release" (HTHR) of the second stage. This two-stage behavior is purely the result of chemical kinetics, and is a well-known phenomenon for premixed combustion of diesel/air mixtures. A comprehensive model for this behavior in the diesel surrogate *n*-heptane was developed by Curran et al. [35], showing a positive correlation between temperature and heat release between 550 K and 600 K (an area of "positive temperature coefficient") followed by a negative correlation between 600 K and 700 K (an area of "negative temperature coefficient"), then turning positive again above 700 K for the second stage of autoignition. The model of Curran et al. [35] also shows that this two-stage behavior disappears for near-stoichiometric diesel/air mixtures at an elevated pressure of 40 bar (which is why it will not be seen in the PCCI example in Figure 3.12). Although it would be feasible to use just one Wiebe function to fit the entire first stage autoignition (as seen in the top of Figure 3.10), it is useful to have a second Wiebe function to fill the "gap" encountered between the first and second stages corresponding to the negative temperature coefficient (NTC) stage, both for the purposes of yielding a more comprehensive conceptual model and also to aid the Wiebe function fitting procedure, which will otherwise try to fit this region into one of the Wiebe functions before or after it.

Having reached the opposite end of the MELT diagram, and having introduced all of the different phases of combustion associated with the different Wiebe functions, this is a good point to stop and look back at the original CDC case, shown in the bottom of Figure 3.11. This shows how, even though the peak rate of heat release in both combustion modes is about the same, RCCI heat release is almost all occurring before the CDC heat release even begins, and right around TDC, which is ideal for maximizing the area under the P-V curve. Despite this, there is about the same level of fluctuation in the heat release rate (from fluctuation in the original pressure signal) in the late stages of each condition. If CDC was advanced so that the peak HRR was near TDC like it is in RCCI, the peak would be much higher, leading to very high PRR's and high-magnitude fluctuations in the HRR, along with very high NO_X emissions. This situation is known as ”diesel knock,” and is another practical limitation on the BTE of CI engines that can be avoided with dual-fuel modes.

3.4.6 PCCI (Premixed Charge Compression Ignition)

PCCI, seen in the top of Figure 3.12, features entirely stratified premixed fuel, all Early DI diesel with no gasoline. As a result of the higher reactivity of diesel, combined with the stratification producing local mixtures closer to stoichiometric, the premixed autoignition occurs much earlier than with gasoline. Despite the stratification of equivalence ratio, this autoignition appears to occur nearly simultaneously throughout the entire combustion chamber, with almost all of the heat release concentrated in a 5° window. However the previously cited optical study [16] has shown that this fuel must be stratified, and thus this phase cannot use the same Wiebe function that was used for the homogeneous autoignition in HCCI. Instead, this phase will use the stratified autoignition function of RCCI, will be identified as the sequential autoignition sweeping through areas of stratified reactivity. This process is much less staggered than it is in RCCI, since there are only gradients in equivalence ratio, with no gradient of fuel reactivity.

It should be noted that in this particular engine, PCCI using standard diesel is not advisable except at the lowest loads. Even then, it is necessary to use large amounts of EGR to dilute the charge and slow down the autoignition process to a point that it does not cause extremely high pressure rise rates (PRR). For this reason most PCCI experiments at this load use either a lower-reactivity fuel than diesel, a lower compression ratio, or have much more heat rejection in the EGR cooler to achieve lower intake gas temperatures. This made it hard to come by an optical study that was comparable to this PCCI condition, but one example was found in the work of Kanda et al. [36]. This study observed a new mechanism for the slow taper at the end of the HRR: a film of liquid diesel on the piston and cylinder wall. This non-premixed “wall film” produces a large amount of soot during the main combustion event with very clear soot luminosity. For this reason it is ideal to use an injector with a

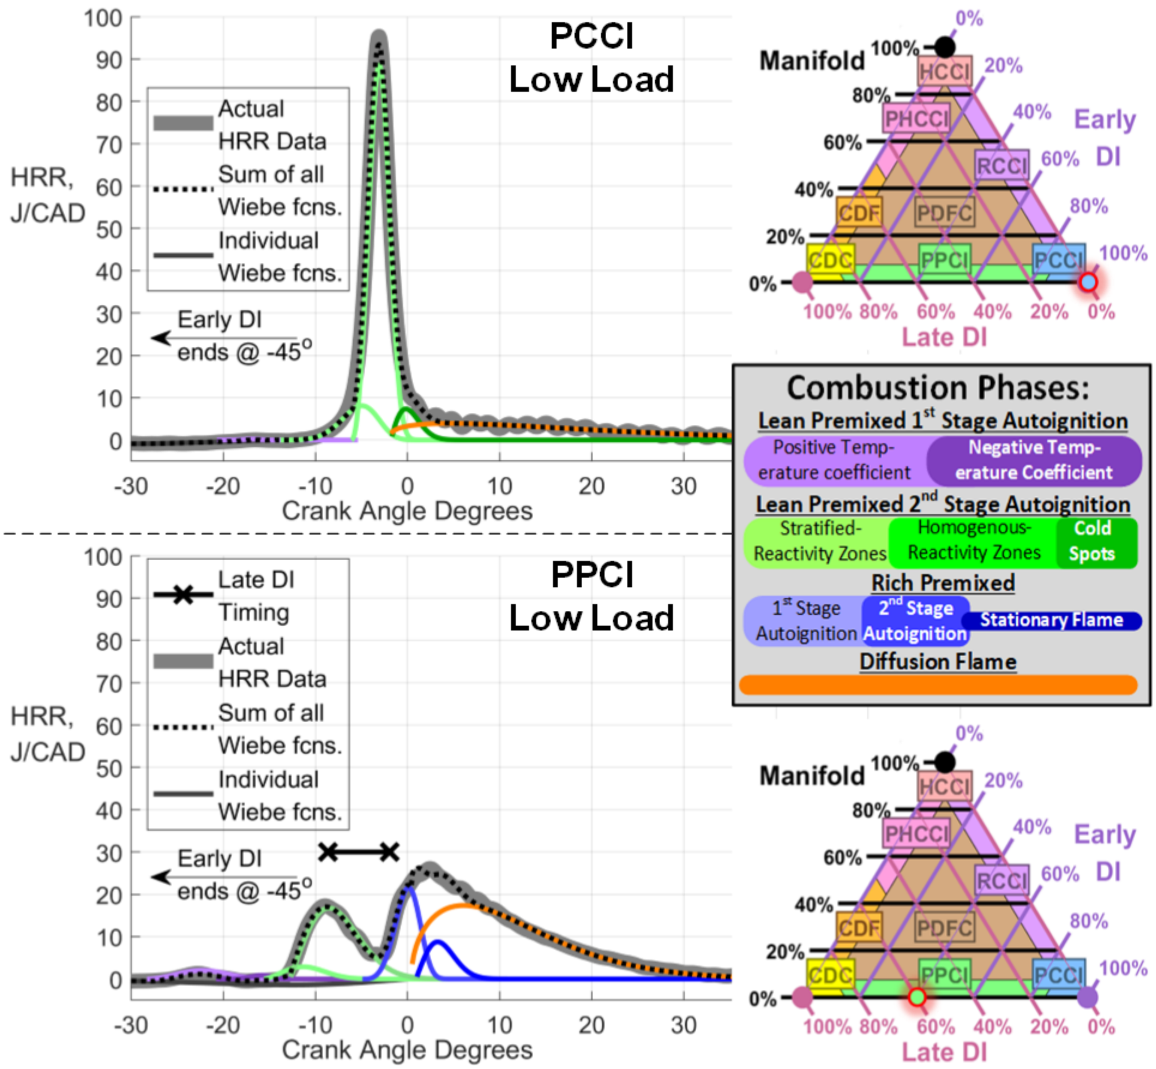

Figure 3.12 Wiebe function breakdowns of PCCI (top) and PPCI (bottom). PCCI in this engine is not advisable due to the high compression ratio and high reactivity of diesel, which causes all of the heat release to happen before top dead center. PPCI is much more applicable since this pre-TDC heat release is balanced out by the mixing-controlled process that occurs after TDC.

smaller umbrella angle when operating in modes with large amounts of Early DI diesel, so that this film is concentrated on the piston surface. On the piston, the film is exposed to the rest of the hot charge to help burn the soot produced, versus if the film winds up on the cylinder wall, it will be covered up by the piston as it approaches TDC, inhibiting oxidation of the soot and allowing fuel and soot to pollute the engine's lubricating oil. For these reasons, the rest of these experiments will avoid using more than 60% Early DI, and production dual-fuel engines that utilize the Early DI should have pistons and injectors that are designed to avoid liquid diesel collecting on the cylinder wall, as is done in the work of Kanda et al. [36]. More successful diesel PCCI results can also be achieved with engines that can provide higher EGR fractions in the range of 60-70% [37], rather than the maximum EGR rate of 40% of used in this study.

3.4.7 PPCI (Partially Premixed Compression Ignition)

In PPCI, there is essentially a combination of the PCCI and CDC processes, occurring before and after TDC respectively. There is a much more subdued version of the PCCI process occurring before TDC, and thus the pressure rise rates are not as excessive as in PCCI and it is not necessary to use excessive EGR. The Late DI occurs simultaneously with the tail end of the lean premixed autoignition, which is almost complete by the time the Late DI ignites. As a result, the rich premixed autoignition occurs in a much hotter and more dilute environment, since the lean premixed autoignition has released heat and consumed oxygen. The increase in temperature decreases ignition delay, but the lack of oxygen slows down the autoignition process once initiated, leading to a shorter and wider rich premixed peak than is observed in CDC and CDF. Another effect is that the diffusion flame is now occurring in a less oxygen-rich environment, which could inhibit the oxidation of soot that was produced by the rich premixed autoignition.

This combustion mode is less commonly studied, and there is a lack of suitable optical engine studies to confirm the accuracy of this breakdown. However there are some experiments coupled with CFD simulations that give some insight into the process. Firstly, Lechner et al. [38] evaluated PPCI using a series of narrow spray cone angle injections in order to achieve PPCI without the Early DI wetting the cylinder wall. This was successful, but then created the problem that the Early DI was being directed into the bowl rather than the outer “squish” region of the combustion chamber as shown in the CFD results. As a result this would mean that the lean premixed autoignition was occurring directly in the path of the Late DI. This would result in an even richer environment for the Late DI’s rich

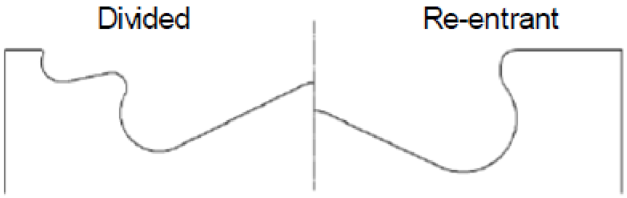

Figure 3.13 Piston redesigned by Inaba et al. [39] to optimize PPCI. The new piston design on the left is much more compatible with an Early DI, since the outer bowl will "catch" the Early DI and prevent it from reaching the cylinder wall.

premixed autoignition process, which could be the reason these authors did not produce the level of PM reduction they were seeking.

A PPCI study utilizing a wide cone angle similar to that of the injectors used in this thesis (as well as a very similarly-shaped piston) comes from Inaba et al. [39], who refer to PPCI as "Twin-Shaped Semi-Premixed Diesel Combustion". In the CFD analysis of that study it was found that, despite the heat release of the Early DI being concentrated in the squish zone, there is so little space in the squish zone near TDC that the burned gases from the Early DI must enter the bowl region. This study is in agreement with the hypothesis that, in order to reduce soot emissions in PPCI, the Late DI must be kept away from the burned gas zone created by the Early DI. This was done by altering the shape of the piston to create a "divided" combustion chamber, with a second, outer indent above the central bowl where the Early DI is aimed, seen in Figure 3.11.

This piston design dramatically improved the results of PPCI, increasing indicated thermal efficiency from 42% to 44% while cutting both UHC and smoke emissions by more than half. The reduction in smoke and UHC was partially achieved by keeping the burned gases from the Early DI in the upper bowl while the Late DI burned in the central bowl, where there was maximum available oxygen to oxidize soot. Another contributing factor was that the redesigned piston prevented the Early DI from developing a film on the cylinder wall, further reducing smoke and UHC. These reduced smoke and UHC emissions led to much higher combustion efficiency, which is directly responsible for the increase in thermal efficiency. The only drawback was that NO_X was increased due to more available oxygen for the Late DI, however this could be mitigated by using EGR, with NO_X only increasing about 20% at intake oxygen concentrations of 15% or lower, and still being below 1 g/kWh, a threshold that will be considered later for low-NO_X conditions in Chapter 4. However, PPCI still has the undesirable effect of spreading out the total heat release over a long period of time, making it more difficult to maximize indicated thermal efficiency. This is due to the high reactivity of diesel, which causes the lean premixed autoignition to occur well before TDC.

3.4.8 PDFC (Piston-split Dual-Fuel Combustion)

In PDFC, the HRR can be "compressed" into a shorter duration by adding Manifold gasoline back into the mix, seen in Figure 3.14. This lower-reactivity fuel delays the stratified lean premixed autoignition, bringing it closer to the rich premixed autoignition of the Late DI. This way, the two processes are happening near-simultaneously in different parts of the combustion chamber. With a lower concentration of Manifold fuel of 20% (top of Figure 3.14), the stratified lean autoignition will not trigger homogeneous autoignition in the bowl. The gasoline in this area will instead be consumed by the diffusion flame like it was in the low-load CDF condition. With a higher gasoline concentration of 60% (bottom of Figure 3.14), a homogeneous autoignition *will* be triggered, similar to the high-load PHCCI case seen back in Figure 3.10, but with the addition of some stratified lean premixed autoignition leading up to the homogeneous autoignition.

This operating mode, the only one involving all three of the primary injection types, is the least well-studied, and there is not currently any literature that adequately investigates the complex combustion mechanisms occurring here. The name "PDFC" is taken from the experimental work of May et al. [40], wherein it stands for "Premixed Dual-Fuel Combustion". That work was focused on the experimental performance of PDFC in an engine with no optical access, and only briefly speculated on the combustion processes occurring. This acronym will be redefined here as "Piston-split Dual-Fuel Combustion,"

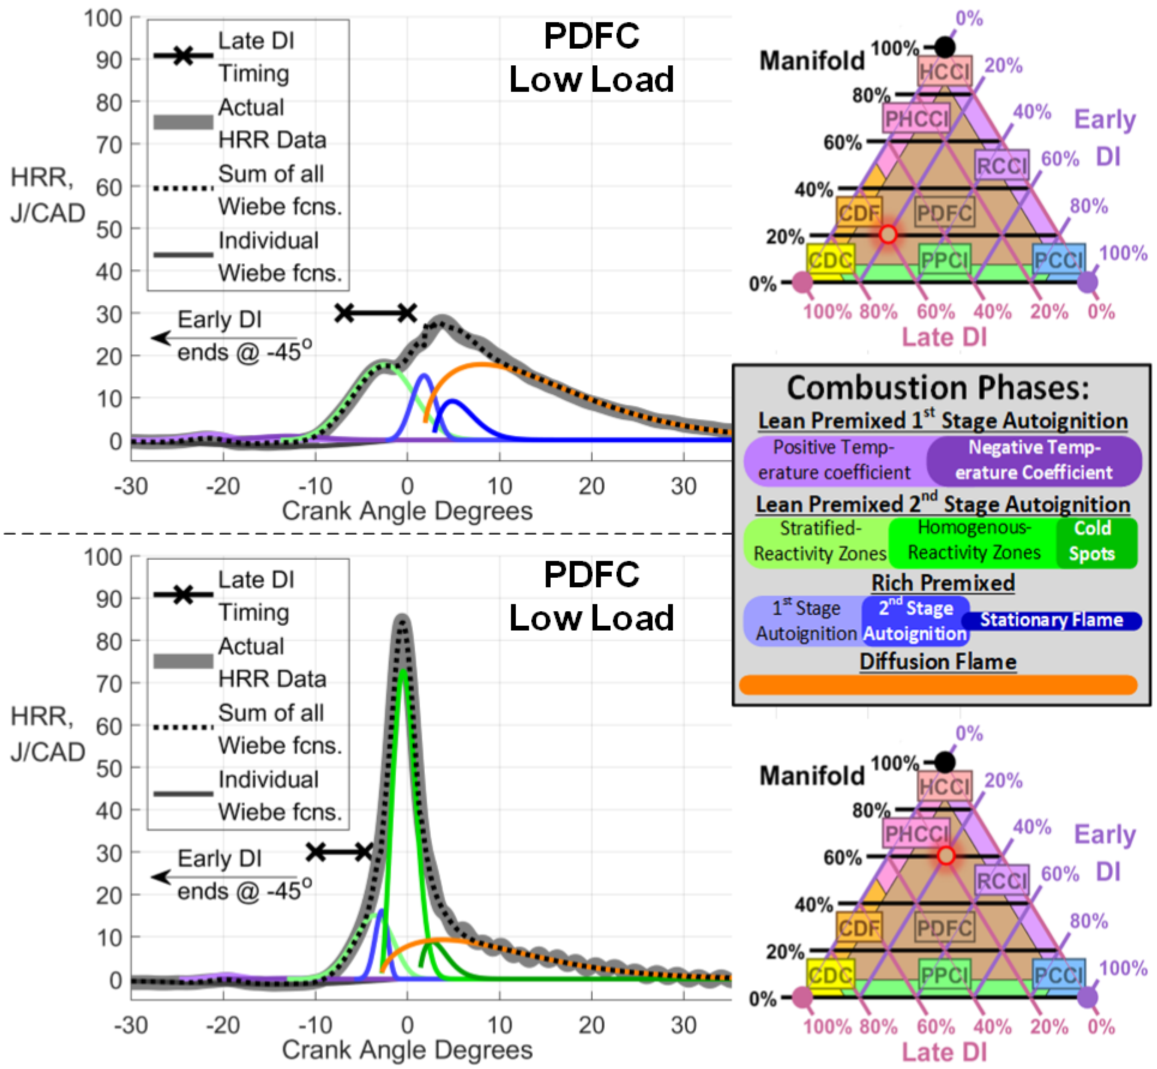

Figure 3.14 Wiebe function breakdowns of "light" PDFC (top) and "heavy" PDFC (bottom). The light PDFC case with only 20% Manifold fuel produces a much wider, flatter HRR shape with no homogeneous lean premixed autoignition "spike", similar to PPCI. The heavy PDFC case with 60% Manifold fuel produces a much more "condensed" HRR shape with a tall homogeneous autoignition "spike," similar to RCCI.

referring to the fact that the piston design splits the combustion chamber into two areas where different processes are occurring: an RCCI-like process in the squish zone and a CDC-like process in the bowl.

A rough graphical representation of this hypothesized process is shown in Figure 3.15. In this figure, the same colors are used to indicate both the rough location within the cylinder of each combustion process, as well as the Wiebe function representing the heat release by this process. Areas of unburnt fuel are indicated by mixtures of red (gasoline) and brownish yellow (diesel) while burned gases are represented by checked patterns of black and either white (for excess air from lean combustion) or yellow (for unburnt fuel and soot from rich combustion). This figure combines all of the combustion sub-processes described in this chapter, and the following chapter will map out how shifts in the magnitude and timing of these processes can affect performance.

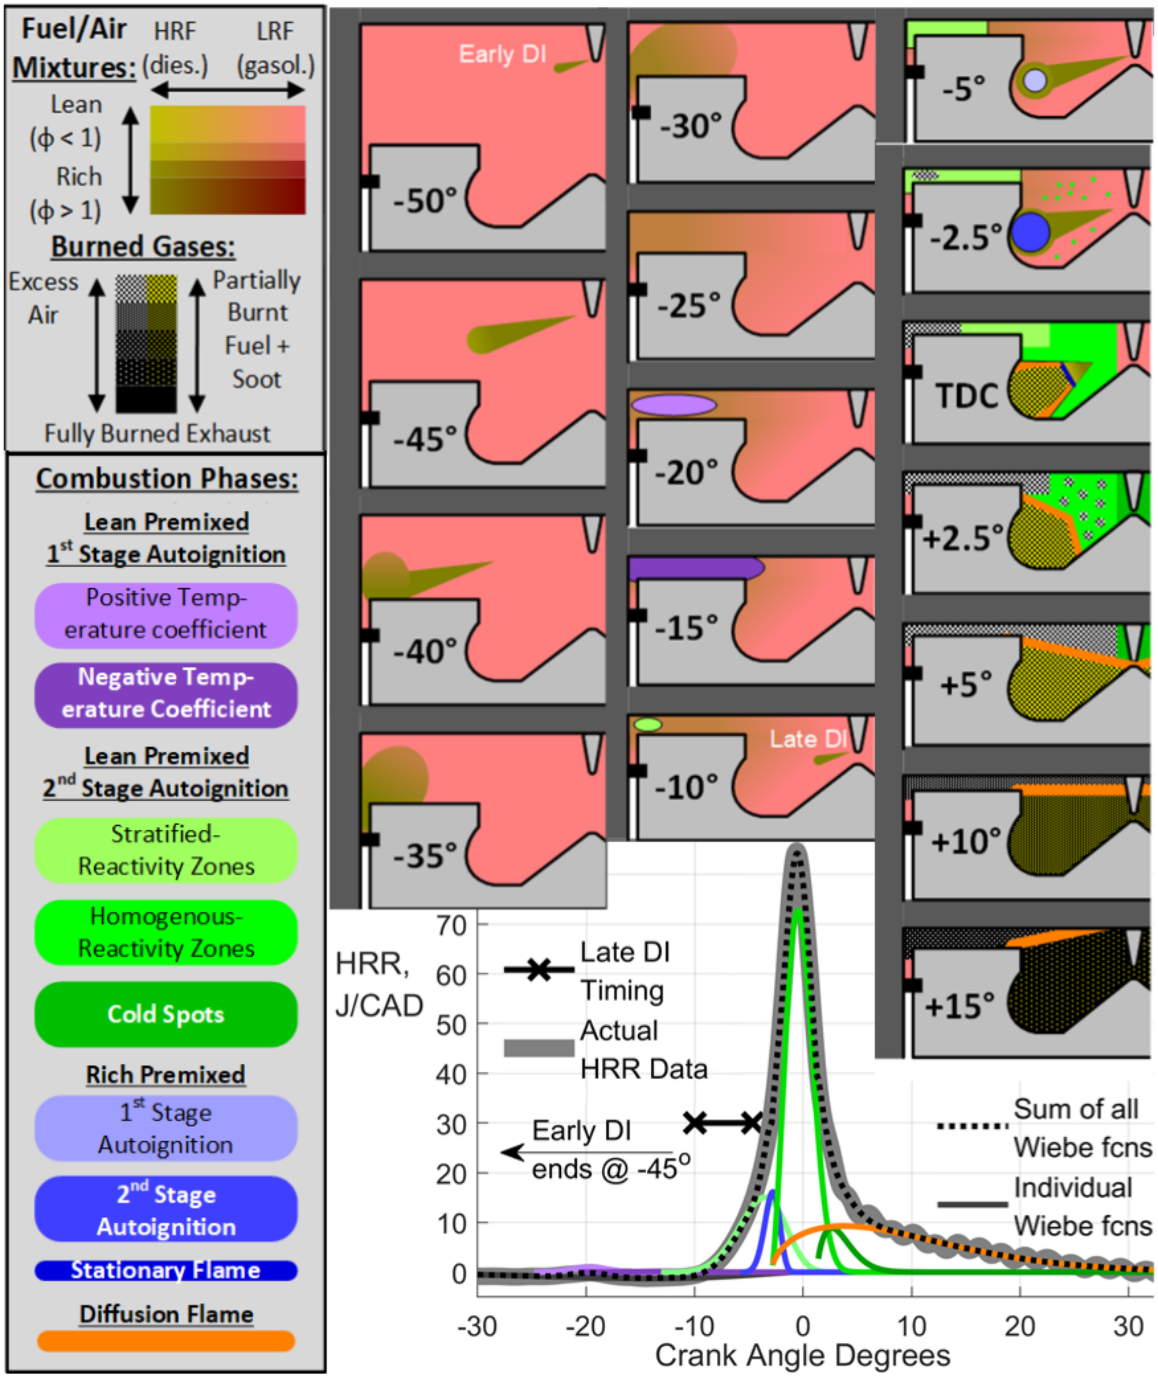

Figure 3.15 Concept diagram showing all possible combustion processes, collectively forming PDFC. Top left: color key for fuel/air mixtures, seen in the spatial diagrams to the right. Bottom left: color key for combustion phases, represented in both the spatial diagrams and the Wiebe function breakdown at the bottom right. Top right: spatial diagrams showing approximate areas where fuel/air mixtures and combustion zones will form at the given crank angle degree. Bottom right: Wiebe function breakdown of the "heavy" PDFC case seen previously in figure 3.14.

Chapter 4

Mapping Performance of Different Combustion Modes

This chapter presents an experimental study that evaluates the performance of the test engine across the widest possible region of the MELT diagram at three different speed/load conditions. Section 4.1 will outline the experimental conditions and the logic behind them, while also showing the calculations used to determine the energy balance and explain shifts in ITE. This will be followed by Sections 4.2, 4.3, 4.4, which present engine performance as mapped onto a MELT diagram for each speed/load condition. Finally, the analysis will focus in on a series of individual MELT ratios in Section 4.5 to determine which conditions to study more closely in the next chapter.

4.1 Experimental Test Plan

The performance envelope of the 1.9L engine used in this study (described previously in Section 2.1) spans a speed range of 1000-4000 RPM and a load range of 0-20 bar BMEP. Ideally, to fully evaluate the potential of dual-fuel modes, the entire performance envelope would be explored. However, there were several factors that limited the range of speed/load conditions that were actually tested:

- The frequency that each condition is demanded by drivers
- The range over which the test setup could produce reliable steady-state test results
- The desire to extrapolate the results to performance in a larger engine

In order to ensure that the conditions tested were frequently demanded by drivers, the conditions were selected with the aid of plots of simulated EPA drive cycles generated by Gao et al. [10]. These were developed using the default parameters Autonomie vehicle simulation platform for a conventional 1,580kg mid-size passenger car with an automatic transmission. Four different drive cycles were studied: the Urban Urban Dynamometer Driving Schedule (UDDS), the Highway Fuel Economy Driving Schedule (HWFET), the US06 cycle, and the New York City driving cycle (NYCDC). The widest-ranging of these cycles, the US06 cycle (seen before in Figure 1.3), forms the background of Figure 4.1, which also shows both the range of test conditions available in the 1.9L engine setup and the operating range of a Hino Motors (a primary sponsor) 5.1L engine, which generally operates at lower speeds than the 1.9L engine. Although the test setup could produce accurate results at loads below 4 bar, loads lower than this were avoided to avoid very low flowrates in the fuel flow meters (which need to be sized for the maximum possible fuel flow), leading to high uncertainty in the fuel flow measurements at these low loads. Loads higher than 10 bar BMEP also had to be avoided so that peak cylinder pressures could be kept under 150 bar in all combustion modes, the limit that was deemed safe for long-term steady-state operation. With these constraints, the six speed/load conditions seen in Figure 4.1 were selected and labeled “A” through “F”.

The other specifications of the engine are given below in Table 4.1. The test parameters were set in a National Instruments/DRIVVEN control/DAQ cart, which featured a map of the engine’s original ECU in the form of lookup tables. These lookup tables allow for a close estimate of the CDC condition that would be set up by the ECU at a given speed and load. The DI rail pressures were set directly from these lookup tables, while the boost settings were slightly lowered from the lookup table values in order to maximize the baseline CDC efficiency. This also ensured that the full range of EGR could be explored in all the available combustion modes, as preliminary experimentation found that the original boost levels from the lookup tables could not be maintained with high EGR in some conditions. The maximum EGR level was determined by estimating the equivalence ratios produced

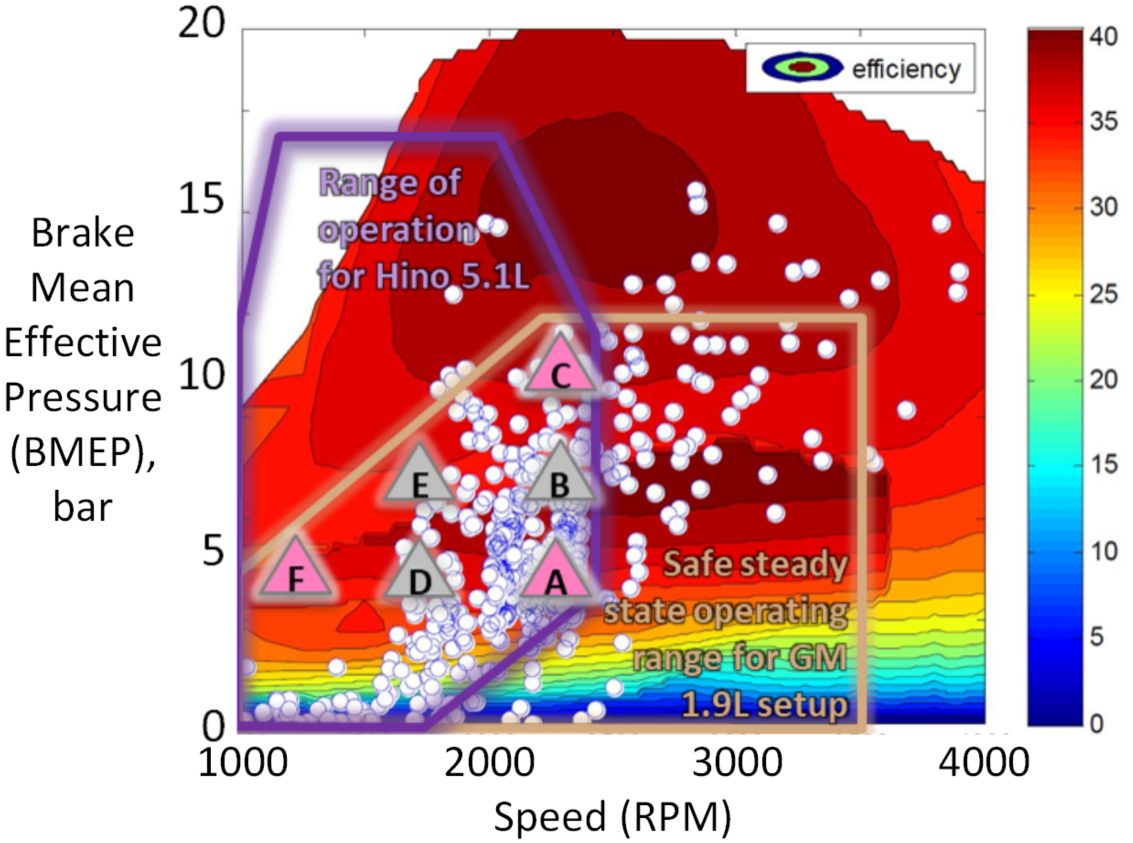

Figure 4.1 Speed/load points selected for this study, superimposed on a simulated EPA drive cycle by Gao et al. [10]. The points for which a full MELT diagram mapping was performed are highlighted in pink. Details on page 67.

by a given boost/EGR combination and ensuring that this maximum level did not produce more than 1 g/kWh of soot in CDC.

The scope of these tests had to be limited due to the amounts of time and fuel available, with each data point taking about 10-15 minutes to acquire and using on average about 1 L of fuel. This meant that one work day (including setup, warm-up, shutdown, and any troubleshooting needed along the way) would produce about 30 data points and use about 30 L of fuel. A full mapping of a particular speed/load point would require about 350 data points total, which equated to 12 full work days and 350 L of fuel. Factoring extra time for finding the limits of operation and troubleshooting unforeseen problems, it was estimated that one full month of 20 work days would be needed to map out an entire speed/load point,

Table 4.1 Speed/load conditions selected for this study. The points for which a full MELT diagram mapping was performed are highlighted in pink. Details on page 67.

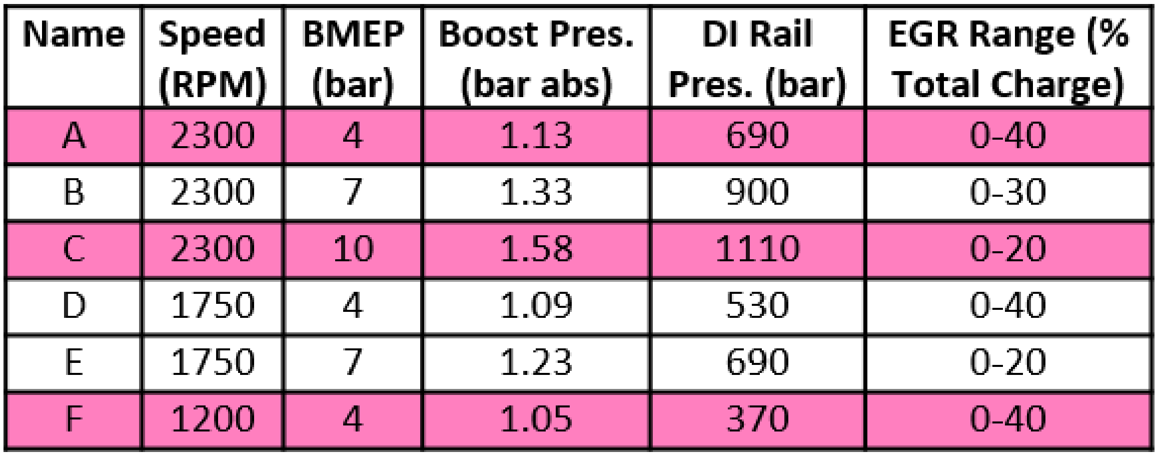

Name	Speed (RPM)	BMEP (bar)	Boost Pres. (bar abs)	DI Rail Pres. (bar)	EGR Range (% Total Charge)
A	2300	4	1.13	690	0-40
B	2300	7	1.33	900	0-30
C	2300	10	1.58	1110	0-20
D	1750	4	1.09	530	0-40
E	1750	7	1.23	690	0-20
F	1200	4	1.05	370	0-40

and about half a year would be needed to explore all six speed/load conditions in Table 4.1 completely. It was thus decided to limit the full MELT diagram exploration to just three of these speed/load conditions.

For this chapter, to map the widest area of the MELT diagram as possible, the tests will be limited to the three highlighted speed/load points of A, C, and F. The logic for this is that these three points are at the vertices of the triangle formed by all six points, and the three points that will not be mapped (B, D, and E) each lie exactly halfway between two of the other three points on the speed/load map, as seen in Figure 4.1. Thus, if maps of the A, C, and F conditions are generated, they should give an indication of where on the MELT diagram to look in B, C, and E to find optimal performance, without exploring the entire MELT diagram.

For this chapter, there is essentially a four-dimensional test matrix at each of the three speed/load points. The MELT ratio forms the first two dimensions, since even though there are three injection types, the *percentage distribution* of the total fuel across these three types of injection must always sum to 100%, and thus there are only two degrees of freedom (this is why the distribution can be plotted onto a two-dimensional MELT diagram). The third and fourth dimensions in the test matrix are EGR level and Late DI SOI timing, which exert control over the combustion phasing at a fixed MELT ratio. The test data for this chapter will be presented in the form of color/contour maps. This will plot the data onto the MELT diagram, with a specific EGR and SOI value selected for each point on the MELT diagram. These maps will show important trends for identifying optimal dual-fuel combustion modes.

4.1.1 Defining Test Parameter Limits

It would be inadvisable to try to explore the entire MELT diagram at each speed/load, because some fuel ratios, EGR levels, and injection timings will produce combustion phasing that is either much too early or much too late. Thus it is essential to identify the boundaries of the section of the MELT diagram in which decent combustion phasing can be maintained, and define criteria for what constitutes "decent combustion phasing." Within these boundaries, there must also be limits of the Late DI SOI timing at each MELT ratio. Both the MELT diagram boundaries and SOI timing limits will change at each EGR level.

Knock Limit

When combustion occurs too early, pressure is raised while the piston is still on the upward compression stroke. This is undesirable for two reasons: First, it is detrimental to efficiency by raising the compression work that the piston must exert on the charge and giving more time for heat to be lost to the combustion chamber walls before it can be extracted as useful work in power stroke. Second, it creates a phenomenon called "diesel knock," a term referring to the excessive noise and vibrations produced in diesel engines with very advanced injection timings.

Diesel knock was a primary limiting factor in the development of high-speed direct-injection engines like the engine used in this study, which to achieve their high efficiency require extremely quick mixing of fuel and air [41]. This is achieved by extremely high fuel injection pressures, which allow a large amount of rich premixed air and fuel to be created before quickly and simultaneously autoigniting, as seen in the "spikes" in the HRR profiles of CDC seen previously in Section 3.4.1. This is advantageous from the perspective of the thermodynamic cycle, since it concentrates much of the heat release and pressure rise near TDC, where it does not raise pressure (and thus negative work) during the compression stroke, but does raise pressure (and thus positive work) during the power stroke. However, higher spikes in the HRR lead to higher peak cylinder pressures (PCP) and maximum pressure rise rates (MPRR), which not only generate excessive engine noise, but also apply higher structural forces to engine components [42]. To combine these two measurements into a single indicator of knock, one can calculate the ringing intensity (RI), which was developed to predict the combustion-generated noise of HCCI [43].

$$\mathrm{RI} = \frac{1}{2\gamma p_{\max}}\left[\beta\left(\frac{dp}{dt}\right)_{\max}\right]^2\sqrt{\frac{\gamma R T_g}{M}} \tag{4.1}$$

In Equation 4.1, γ is the ratio of specific heats, R is the universal gas constant 8314 $\frac{J}{mol\cdot K}$, and M is the molecular weight of the charge, which is set equal to the molecular weight of dry air of 28.97 g/mol for these calculations for the sake of simplicity, since exhaust molecular weight will not deviate far from this value. The coefficient β correlates the maximum pressure rise rate with respect to time with the maximum amplitude of the pressure fluctuations within the cylinder, commonly set in the literature to a value of 0.05 msec [33,43–45]. This calculation produces the correlated ringing intensity in MW/m^2, and the RI knock limit typically seen in literature is 5 MW/m^2 [33,45].

However, despite this value of 5 MW/m^2 being ubiquitous in the literature, the engine used in this study produced RI in excess of 5 MW/^{2} even when operating the original ECU conditions. The recent work of Shahlari et al. [44] suggests that, while there is very good linear correlation between RI and actual engine noise, this correlation changes drastically depending on the engine being used. This justifies an increase in the RI knock limit for this thesis from 5 to 10 MW/m^2. This 10 MW/m^2 limit is based on the observed high-amplitude noise in the pressure signal after combustion (and perceived noise increase from the engine test cell) when crossing that threshold at each of the different speed/load conditions tested herein. The work of Shahlari et al. [44] also highlights the importance of filtering the raw pressure signal when calculating RI, recommending a cutoff frequency based on 75% of the lowest resonant frequency of a cylinder with the same bore as the combustion chamber, which for a bore of 82 mm would be 4 kHz. Thus, this thesis follows the recommendation of Shahlari et al. [44] to apply a 5th order Butterworth filter with a cutoff frequency of 4 kHz to the raw pressure signal when calculating RI.

Misfire Limit

Going in the other direction, when combustion occurs too late, pressure is raised while the piston has already started the downward power stroke. Again this is undesirable for two reasons: First, it fails to create positive work at the top of the power stroke, putting a "dent" into the top of the P-V diagram. Secondly, it reduces the temperature and pressure in the combustion chamber during the autoignition process, increasing the ignition delay. This increases the time with which the Late DI fuel disperses through the combustion chamber, diluting the fuel-air mixture and increasing the variability with which this fuel is distributed from cycle to cycle due to random turbulence. This variability, combined with the lowered temperatures, leads to higher cycle-to-cycle variability in the ignition delay, to the point where there is incomplete combustion on certain cycles. This phenomenon is known as "misfire", and is indicated by rapid rises in CO, UHC, and cycle-to-cycle variability, as well as a rapid decrease in overall efficiency.

The tendency to misfire in conventional diesel engines is increased by adding EGR, which further dilutes the fuel-air mixture from the Late DI and increases the ignition delay despite decreasing NO_X. As a result, much research in minimizing the NO_X emissions of CDC engines (and achieving LTC) is focused on the limits of EGR increase and SOI retardation that can be used before reaching the misfire limit [46–48]. A convenient method of quantifying misfire is to use the cycle-to-cycle variation of indicated mean exhaust pressure (IMEP), which will limit the drivability of a vehicle when it reaches 3-5% [49]. This value of 5% CoV of IMEP will serve as the misfire limit.

Efficiency Limits

For both the knock and misfire limits, there must also be a limit on the minimum gITE that must be achieved, since overly advanced or retarded combustion phasing cause a severe drop in efficiency before either of the above limits is reached. This value was set differently at each speed and load to 90% of the baseline gITE from the baseline condition of the original ECU parameters at each speed and load. This ensures that the test matrix includes only test conditions that are relevant for high-efficiency (and thus low-CO_2) engine operation.

4.1.2 Creating the Test Matrix

Using the limits specified above, the boundaries of acceptable performance on the MELT diagram were located for the three speed/load conditions of A, C, and F. Approximately 15 MELT ratios within these limits were selected per speed/load condition, shown in Figure 4.2. A wide area of the MELT diagram was accessible for speed/load A, but only a smaller area was available for speed/load F, and an even smaller area for speed/load C. For these reasons, the MELT diagrams for F and C will be stretched as shown in Figure 4.2 when making contour maps, to better show the details in the accessible region of the MELT diagram.

In speed/loads A and F, the EGR could be varied from 0% to 40%, and five levels of EGR were chosen to get the full range of possible performance (0, 10, 20, 30, and 40%). In speed/load C, maximum EGR was 20%, so only the three levels of 0, 10, and 20% were tested. At each EGR level, an SOI timing sweep was conducted to find a knock limit and misfire limit of SOI, with “hard” limits of 20°aTDC and 0°aTDC if knock or misfire limit was never reached. Then, an evenly-spaced five-point sweep of SOI between these limits was set in the test matrix. To reduce the size of the test matrix and expedite the testing

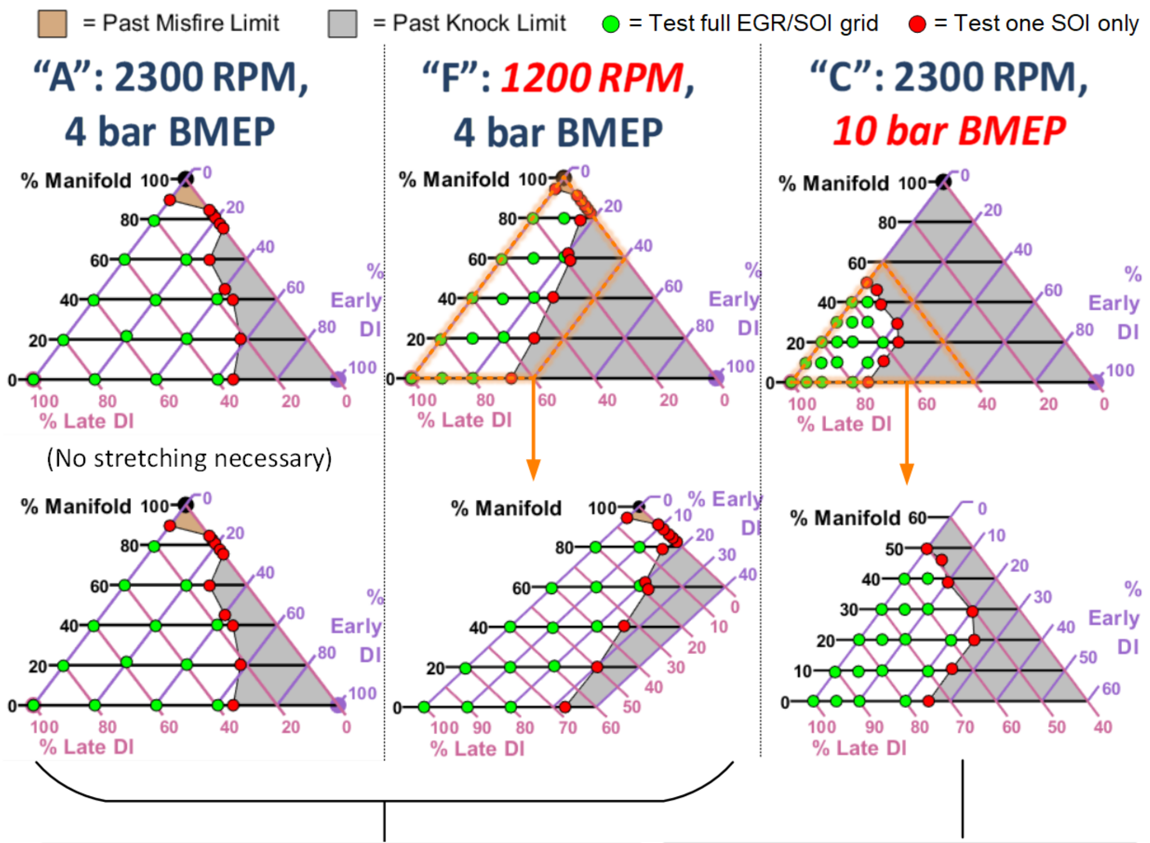

Late Main DI SOI	EGR Level				
	0%	10%	20%	30%	40%
1. Knock Limit or 20 °bTDC	✓	✓	✓	✓	✓
2. Halfway between 1 & 3	✓	X	✓	X	
3. Halfway between 1 & 5	✓	✓	✓	✓	✓
4. Halfway between 3 & 5	✓	X	✓	X	
5. Misfire limit or TDC	✓	✓	✓	✓	✓

Late Main DI SOI	EGR Level		
	0%	10%	20%
1. Knock Limit or 20 °bTDC	✓	✓	✓
2. Halfway between 1 & 3	✓	✓	✓
3. Halfway between 1 & 5	✓	✓	✓
4. Halfway between 3 & 5	✓	✓	✓
5. Misfire limit or TDC	✓	✓	✓

Figure 4.2 Test points in this chapter mapped onto the MELT diagram, showing the stretching of the MELT diagram used to show details in F and C. Green dots show the MELT ratios at which the entire table at the bottom was tested, red dots show the limits at which only one point was tested per EGR level (limits shift depending on EGR level).

process, the second and fourth points on the 10% and 30% sweeps were eliminated in A and F, for a total of 21 points per MELT ratio, while all 5 points on each sweep were kept in C, for a total of 15 points per MELT ratio. The exception to the above system was the points with no Late DI, for which there was no SOI timing to sweep. Instead, the proportions of Manifold and Early DI fuel were varied much more finely than the variations between the other MELT ratios, to produce a sweep of RCCI phasings at each EGR level.

At each test point, the data was collected via the following procedure:

- The engine was brought to a steady state at the desired operating condition, as indicated by a flat, unchanging trend in the pre-turbocharger exhaust temperature.
- High-speed synchronous data (i.e. the cyclic pressure trace) was collected in three 100-cycle "snapshots" taken approximately one minute apart.
- Medium-speed scalar data, including the temperatures and pressures measured at various points on the engine, as well as the torque measured by the dynamometer, was sampled once per cycle in each of the three 100-cycle snapshots above.
- Slow-speed data, i.e. the fuel flows, were sampled once per second over the entire two-minute test procedure.
- Gas emissions measurements in the Horiba Mexa-ONE emissions bench were sampled twice per second over a fifty-second period.
- Three filter smoke number (FSN) measurements were taken with an AVL 415S smoke meter.

It should be noted that the FSN is not a true measurement of soot mass, but rather a measurement of the change in reflectivity of a filter paper after a known "sample length" (i.e. a sample volume divided by the area of the filter paper) of raw exhaust is passed through it. AVL (the manufacturer of the smoke meter) has developed a formula for correlating the FSN to the soot concentration in the exhaust [50]:

$$\rho = \frac{1}{l} \cdot \alpha \cdot \mathrm{FSN} \cdot e^{\beta \cdot \mathrm{FSN}} \tag{4.2}$$

In Equation 4.2, ρ is the soot density in air in mg/m^3, l is the effective sample length in meters (set to 0.405 m in the AVL 415S smokemeter), and α and β are constants developed by extensive testing to compare FSN values with actual soot concentrations measured through other means. The values given by AVL [50] and confirmed in the literature [51] are 4.95 for α and 0.38 for β, and these values with be used herein to calculate soot mass.

4.1.3 Calculating Energy Balance

To understand *why* the engine's efficiency is changing in the different modes, it is necessary to calculate all of the flows of energy in and out of the engine. As previously mentioned in Chapter 1, there are three main sources of heat loss that limit the engine's thermal efficiency: exhaust loss, cooling loss, and mechanical loss. This thesis will focus on the exhaust and cooling losses while ignoring the mechanical losses. This is because the mechanical losses are determined by the overall design of the engine, including the air and exhaust passageways, EGR cooler, and turbocharger, and for the design of this engine these have all been optimized for CDC alone. In order to take full advantage of the dual-fuel modes and minimize mechanical losses, the engine would have to be redesigned. The cooling and exhaust losses, however, are more dependent on the heat transfer rate out of the cylinder and final temperature and flowrate of exhaust, respectively. Thus, the rest of this thesis will only consider the *gross indicated* thermal efficiency (gITE) of the engine, which measures the work done by combustion while the valves are closed, without subtracting out any mechanical losses that happen via throttling of the flow through the air/exhaust passageways or friction of the moving engine components.

Figure 4.3 gives the total energy balance breakdown of the engine averaged over all of the conditions taken at speed/load A, in four stacked bar charts. The bars on the left show energy *into* the engine, while the three bars on the right show energy coming *out* of the engine. The y axis is labeled for three different ways of measuring the energy flows: in terms of power (in kW), mean effective pressure (MEP, in bar), and % TE.

In the leftmost column is the available energy of the fuel in terms of its lower heating value (LHV), which represents the total chemical heat release that occurs during complete combustion, and assumes (correctly) that the water produced by combustion will remain in a gas phase throughout the process (and thus not transfer any heat to the engine via

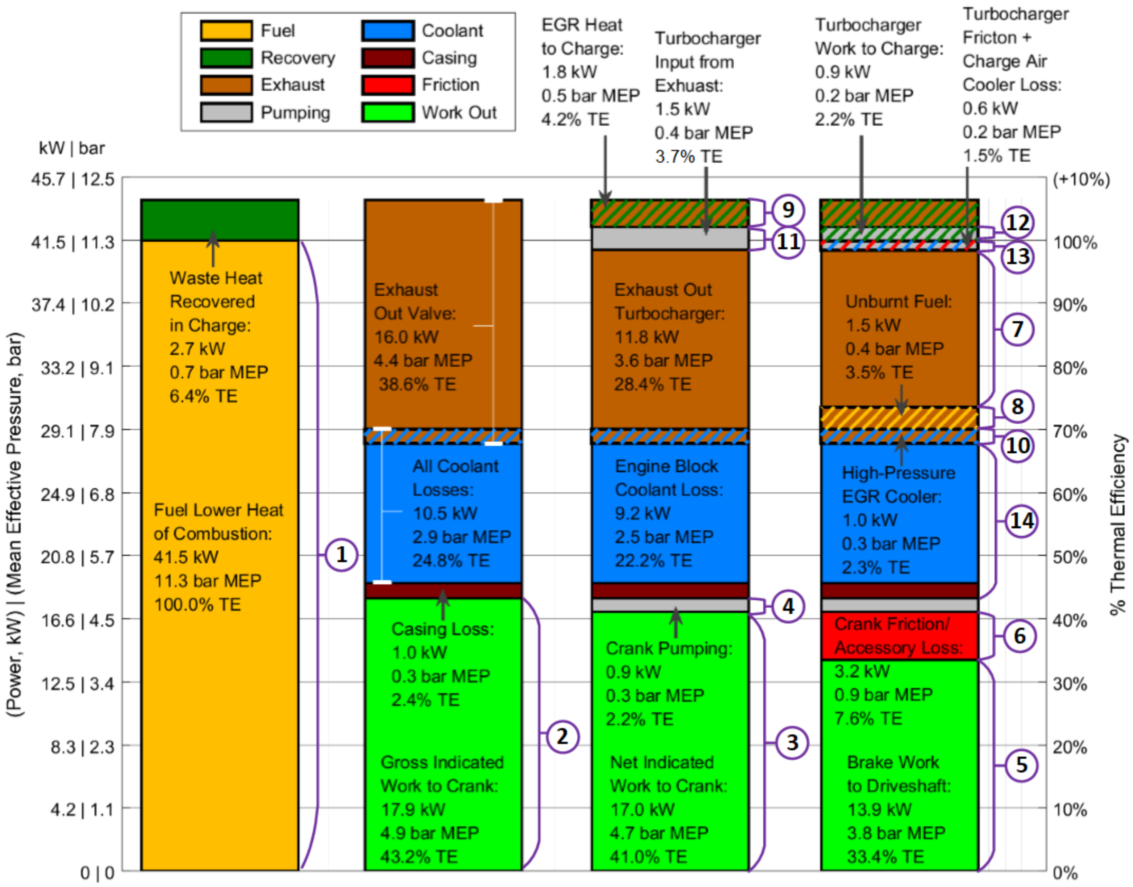

Figure 4.3 Average energy balance of speed/load A. Circled numbers correspond to subscripts of work/heat flows in equations 4.3 through 4.22. The first column shows the flows of energy *into* the engine, while the next three columns show the flows of work and energy *out* of the engine, with various degrees of detail.

condensation). This is heat flow number 1 in Figure 4.3:

$$\dot{Q}_1 = \dot{m}_{fuel} \cdot \text{LHV}_{fuel} \tag{4.3}$$

If all of this energy was converted into work, the engine would have 100% TE, but recalling Equation 1.2, this is theoretically impossible unless the environment surrounding the engine is at a temperature of absolute zero. However this is not the only heat input to the engine, since some heat is recovered from the exhaust via EGR and the turbocharger, which will be calculated later.

The work output of the engine that as measured in this thesis is the *gross indicated* work, which consists of the area enclosed by the P-V diagram while the valves are closed during the compression and expansion strokes, represented by work flow number 2:

$$\dot{W}_{grossind.} = \dot{W}_2 = \frac{\omega}{2} \sum_{k=1}^{n_{cyls}} \int_{-180^\circ}^{180^\circ} P \frac{dV}{d\theta} d\theta \tag{4.4}$$

Most of the heat that is not converted into work is lost through the exhaust and the coolant. Although some of the exhaust heat will be recovered, the second column shows the total amount of heat in the exhaust flowing out of the cylinder through the exhaust valve. None of the heat loss to the coolant will be recovered, instead being transferred to the environment at the radiator, or in the case of this test setup, a coolant/water heat exchanger. There is some overlap between the coolant and exhaust loss because some of the exhaust coming out the valve will transfer its heat to the coolant in the EGR cooler. The remaining, small fraction of heat conducts through the engine casing out to the environment instead of into the coolant.

Looking at the total work being done by the engine driveshaft, there are three sources of mechanical loss. The first is the pumping loss, which constitutes the friction between the working fluids (air/exhaust) and the engine passageways that must be overcome by the piston. If the engine is highly turbocharged and there is not too much backpressure in the exhaust from the turbine, then the pumping loss will be negative and contribute to the work done by the piston, but this is never the case for the engine and conditions tested herein. The gross indicated work minus the pumping loss is equal to the *net* indicated work, which is calculated from the *total* integrated area under the P-V curve, including all four strokes of the cycle:

$$\dot{W}_{netind.} = \dot{W}_3 = \frac{\omega}{2} \sum_{k=1}^{n_{cyls}} \int_{-360^\circ}^{360^\circ} P \frac{dV}{d\theta} d\theta \tag{4.5}$$

Pumping loss, number 4 on Figure 4.3, is calculated from the area enclosed by the P-V diagram during just the intake and exhaust strokes, which is equal to the difference of net and gross indicated work:

$$\dot{Q}_{pumping} = \dot{Q}_4 = \dot{W}_2 - \dot{W}_3 = \frac{\omega}{2} \sum_{k=1}^{n_{cyls}} \left[\int_{-360^\circ}^{-180^\circ} P \frac{dV}{d\theta} d\theta + \int_{180^\circ}^{360^\circ} P \frac{dV}{d\theta} d\theta \right] \tag{4.6}$$

The second source of mechanical loss is rubbing friction between the piston and cylinder, as well as between the crankshaft and its bearings. The third source is the work needed to power the accessories driven by belts off of the crankshaft, such as the coolant and fuel pumps. These losses cannot be directly measured individually by the test engine setup, but their sum can be calculated by the difference between net indicated work and the measured "brake" work being done by the driveshaft on the dynamometer, which is number 5 on Figure 4.3:

$$\dot{W}_{brake} = \dot{W}_5 = 2\pi\omega\tau \tag{4.7}$$

In Equation 4.7, τ is the measured torque from the dynamometer. The friction and accessory losses will be combined into one heat flow, number 6 on Figure 4.3. This flow can be calculated by the difference between the calculated net indicated work and the brake work:

$$\dot{Q}_{fric./acc.} = \dot{Q}_6 = \dot{W}_3 - \dot{W}_5 \tag{4.8}$$

To calculate the exhaust losses requires first calculating the total mass flow of exhaust, which must be equal to the sum of the mass flows of air and fuel to preserve mass balance (assuming any blow-by mass losses through the piston rings are negligible). The measured flow of fuel is very certain, since the diesel and gasoline flowmeters could be calibrated by placing the fuel tanks on a scale and measuring the change in weight. However the same cannot be done to confirm the measurements of air coming from the mass air flow (MAF) sensor of the engine, which was the only measurement of air flow available for this engine. Ideally, the setup would include a laminar flow element (LFE) that straightens the air flow and measures the pressure drop across an orifice for a highly accurate air flow measurement, but an LFE was not available for the test setup.

In order to confirm the measurements of the MAF sensor, the gas emissions data was examined to ensure that it conformed with the MAF sensor. This was done by using the Brettschneider equation [52], which uses the measured emissions concentrations to calculate the normalized air/fuel ratio (AFR) ratio, lambda (λ):

$$\lambda = \frac{AFR_{actual}}{AFR_{stoich.}} = \frac{\chi_{CO_2} + \frac{\chi_{CO}}{2} + \chi_{O_2} + \frac{\chi_{NO}}{2} + \left[\frac{H_C}{4} \cdot \frac{3.5}{3.5+\frac{\chi_{CO}}{\chi_{CO_2}}} - \frac{O_C}{2}\right] \cdot [\chi_{CO_2} + \chi_{CO}]}{\left[1 + \frac{H_C}{4} - \frac{O_C}{2}\right] \cdot [\chi_{CO_2} + \chi_{CO} + \chi_{UHC}]} \tag{4.9}$$

In Equation 4.9, the various χ are the molar concentrations of exhaust gas species and H_C and O_C are the the hydrogen:carbon and oxygen:carbon molar ratios in the fuel. Since this is a dual-fuel engine, it requires a weighted average of the ratios of each fuel based on their respective flows:

$$H_C = \frac{\dot{m}_{LRF} \cdot H_{C,LRF} + \dot{m}_{HRF} \cdot H_{C,HRF}}{\dot{m}_{fuel}} \tag{4.10}$$

The stoichiometric AFR is calculated from the measured mass concentrations y of hydrogen and carbon in the fuel (using the same weighted average as Equation 4.10) and the concentration of oxygen in the air:

$$AFR_{\text{stoich.}} = \left[\frac{y_{H,fuel}}{2M_H} + \frac{2y_{C,fuel}}{M_C}\right] \cdot \frac{M_O}{y_{O,air}} \tag{4.11}$$

This allows a calculation of the air flow that would produce the measured lambda for the measured fuel flow. Averaging this with the measured air flow from the MAF sensor yields a corrected air flow measurement:

$$\dot{m}_{air} = \frac{\dot{m}_{MAF} + \dot{m}_{fuel} \cdot \lambda \cdot AFR_{stoich.}}{2} \tag{4.12}$$

Again to preserve mass balance, the total mass of exhaust exiting the turbocharger $\dot{m}_{out}$ must be the sum of $\dot{m}_{air}$ and $\dot{m}_{fuel}$. To calculate the heat content of the exhaust in terms of enthalpy, this can be broken down into molar flows $\dot{n}_i$ of species i if the entire exhaust concentration is known. However, the gas emissions bench does not measure concentrations of H_2O or N_2 in the exhaust, so these must be calculated using stoichiometry and the composition of the air and fuel. Since the intake air's humidity was not measured, 50% relative humidity was assumed. For brevity, this section will not present the full process of calculating the molar flows in the exhaust and air, which is simply a matter of using algebra to balance the flows of carbon, hydrogen, nitrogen, and oxygen atoms from the known air and fuel flows and exhaust composition.

Once the molar flows are known, the standard enthalpy h^o of the gas mixtures can be calculated using the NIST Shomate Equation standard enthalpy correlations [53]:

$$h^o - h^o_{298} = A \cdot T + B\frac{T^2}{2} + C\frac{T^3}{3} + D\frac{T^4}{4} - \frac{E}{T} + F - H \tag{4.13}$$

In Equation 4.13, h^o_{298} is the enthalpy of formation of the gas species at the standard temperature of 298 K and A through H are coefficients reported by NIST [53] for each gas species. The relative enthalpy of a gas flow at a certain temperature is calculated by the sum of products of the molar flow rates n_i and the standard enthalpies of each species:

$$\dot{H}_{gasflow} = \sum_{i=1}^{n_{species}} \dot{n}_i \cdot h^o_i(T) \tag{4.14}$$

This can be used to calculate the heat flowing out of the engine (number 7 on Figure 4.3) using the exhaust composition and the post-turbo temperature compared to the ambient temperature:

$$\dot{Q}_{exh.out} = \dot{Q}_7 = \dot{H}_{exh.out}(T_{post-turbo}) - \dot{H}_{exh.out}(T_{ambient}) \tag{4.15}$$

Besides the unused *thermal* energy being carried away, the exhaust out of the engine contains unused *chemical* energy in the form of the unburnt hydrocarbons (UHC) and CO (number 8 in Figure 4.3). Although a precise exhaust gas composition analysis would be needed to determine the exact amount of unused heat of combustion, this can estimated by assuming that the measured UHC (which is measured in ppm Carbon) has as much heat of combustion as the original fuel on a per-mole-carbon basis:

$$\dot{Q}_{unburnt} = \dot{Q}_8 = \dot{n}_{UHC} \cdot M_C \cdot y_{C,fuel} \cdot LHV_{fuel} \tag{4.16}$$

Before the exhaust flows through the turbocharger and out of the engine, some is diverted to the EGR loop, where some of the heat is recovered in the charge. This recovered heat (flow number 9 on Figure 4.3) is calculated from the difference between the enthalpy of the air entering the intake manifold (after the compressor and charge air cooler) and the air/EGR mixture in the intake manifold:

$$\dot{Q}_{\text{EGR recovery}} = \dot{Q}_9 = \dot{H}_{\text{EGR/air mix}}(T_{\text{int. man.}}) - \dot{H}_{\text{air}}(T_{\text{post-CAC}}) \tag{4.17}$$

The heat lost in the EGR cooler (flow number 10 in Figure 4.3) is then calculated by the difference between the EGR enthalpy going into the cooler and the heat recovered:

$$\dot{Q}_{\text{EGR cooler}} = \dot{Q}_{10} = \dot{H}_{EGR}(T_{\text{pre-turbo}}) - \dot{Q}_9 \tag{4.18}$$

The heat extracted from the exhaust by the turbocharger (flow number 11 in Figure 4.3) can then be calculated by the enthalpy difference in the exhaust pre- and post-turbo:

$$\dot{Q}_{\text{ext. by turbo}} = \dot{Q}_{11} = \dot{H}_{exh.out}(T_{\text{pre-turbo}}) - \dot{H}_{exh.out}(T_{\text{post-turbo}}) \tag{4.19}$$

However, not all of the heat extracted by the turbocharger is converted into work. The useful work that is actually recovered in the charge is equal to the enthalpy difference in the charge air from before the compressor (ambient temperature) and after the CAC:

$$\dot{W}_{\text{Turbo recovery}} = \dot{W}_{12} = \dot{H}_{air}(T_{\text{post-CAC}}) - \dot{H}_{air}(T_{\text{ambient}}) \tag{4.20}$$

Some of the heat extracted in the turbine is lost to friction in the turbocharger itself, and more is lost to the cooling water in the charge air cooler. Since there was no measurement of the charge air temperature in between the compressor and the charge air cooler, it is impossible to differentiate between these two heat losses. Thus they will be lumped together into a single sum (flow number 13 in Figure 4.3) which represents the total waste heat of the turbocharging/cooling process:

$$\dot{Q}_{\text{Turbo/CAC loss}} = \dot{Q}_{13} = \dot{Q}_{11} - \dot{W}_{12} \tag{4.21}$$

The balance of the heat flow is the heat transfer out of the combustion chamber, most of which goes into the coolant with a small portion conducting out through the engine casing. This difference can be measured by measuring coolant temperatures and flowrate, but this measurement in the test setup was highly imprecise at individual data points. However the difference is not particularly relevant for this thesis, where the important quantity is the total heat transfer away from the combustion chamber. Therefore the sum of the two flows will be lumped into one heat flow (number 14 in Figure 4.3) and calculated by doing a heat balance of the entire system:

$$\dot{Q}_{\text{Cyl. heat transfer}} = \dot{Q}_{14} = \dot{Q}_{1} - \dot{W}_{2} - \dot{Q}_{7} - \dot{Q}_{8} - \dot{Q}_{10} - \dot{Q}_{13} \tag{4.22}$$

The following section will focus on three main sources of heat loss: cooling loss, exhaust loss, and unburnt fuel loss.

“Cooling loss,” as referred to in the next section, refers to the losses that generally go up as a result of increased EGR: flows # 10 and 14. Flow #10 (EGR cooler loss) goes up with increased EGR due to the increased flow of exhaust through the EGR cooler, while flow #14 (cylinder heat transfer) goes up as a result of the higher in-cylinder temperatures, caused by higher intake air temperatures from the increased EGR. “Exhaust loss” will refer to the loss that generally goes *down* as a result of increased EGR: flows # 7 and 13. Both of these flows (the exhaust heat loss and the turbocharging heat loss) go down with increasing EGR as a result of reduced mass flow of air in and exhaust out. The next section will also look at the sum of “cooling + exhaust” that shows the balance between these two trends. “Unburnt fuel loss” refers to the remaining heat of combustion trapped in the UHC and CO emissions, which roughly corresponds to a 1% loss of TE for every 2.5 g/kWh of UHC.

4.1.4 MELT Contour Maps of Performance

One of the most useful aspects of the MELT diagram is that it allows engine performance to be plotted onto a continuous contour map of operating modes, displaying the relative trends in performance in two dimensions as the ratio between injections changes. The following section will contain several maps colored relative to engine performance, with this general color scheme:

- Red = Poor performance, near the limit if applicable
- Yellow = Subpar performance, but not unacceptable
- Green = Good performance, but not the best possible
- Blue = Best performance (with exceptions)

The exceptions to this scheme are data items where performance does not always get better moving in one direction, i.e. CA50, which will be optimized at a certain value and get worse moving in either direction. In this case, green will indicate ideal performance while blue will represent subpar performance, e.g., for CA50 red and yellow mean too *late*, while blue means too *early*. There are also some parameters where one value is not necessarily "better" than another, such as lambda and SOI, which control the combustion phasing. These parameters will have different optimal values depending on which combustion mode is being used, so the colors do not give any indication of "good" or "bad."

The color maps are presented in four groups of three, from top to bottom, forming a full-page figure of 12 color maps that provides a more or less complete picture of engine performance:

- **Limits**: Data that determine the feasible operating range on the MELT diagram
  - **CoV(Gross IMEP), %**: Relative cycle-to-cycle variation in the power output, limited to 5%
  - **Ringing Intensity, MW/m^2**: Correlated knocking noise, limited to 10 MW/m^2.
  - **% Gross ITE**: Engine efficiency, limited to 90% of the original baseline value
- **Combustion Phasing**: Data that determine the timing of heat release
  - **Lambda**: The relative amount of excess air, determined by the amount of EGR
  - **Late DI SOI, °aTDC**: The timing of the Late DI, which is not applicable on the right edge of the diagram where there is no Late DI
  - **CA50, °aTDC**: The crank angle at which 50% of the total heat release over the entire cycle is reached
- **Energy Losses, %TE**: Data that show *how* the Gross ITE is going up or down
  - **Cooling Loss**: The total amount of unrecovered heat lost through the combustion chamber walls and the EGR cooler
  - **Exhaust Loss**: The total amount of unrecovered heat lost from hot exhaust flowing out the tailpipe
  - **Cooling + Exhaust**: The sum of these two losses, showing the trade-off that occurs between the two

- **Emissions, g/kWh**: The main non-CO_2 pollutants found in the untreated exhaust
  - **Nitrogen Oxides (NO_x)**: A pollutant contributing to the formation of smog and acid rain, formed by near-stoichiometric combustion
  - **Correlated Soot from FSN**: A pollutant with negative health effects for plant and animal life, formed by rich combustion
  - **Unburnt Hydrocarbons (UHC)**: A mixture of fuel and intermediate species that limits efficiency, formed by lean combustion

In each contour map, the colored circles represent actual measured data, while the color map and contours shown are a 7th-order polynomial surface fit to these points. There are multiple data points that could be used for each point on the MELT diagram: 21 or 15 for the green points in Figure 4.2, and 5 or 3 for the red points (limits). The point that is selected for each MELT diagram depends on which of two "types" of map is being presented:

"Type 1" maps in the following section are meant to give a general sense of how performance changes with a changing MELT ratio, while holding all other changes to a minimum. Thus, for these maps, a single EGR level was selected to use for all of the MELT points, depending on which EGR level produced the most "balanced" performance between ITE, NO_X, soot, and UHC emissions for each particular speed/load point. Ideally, SOI would stay the same throughout the MELT diagram, but this would produce knocking in some MELT points and misfiring in others. Thus the SOI chosen for each MELT point was the halfway point between the knock and misfire limits at that particular MELT ratio and EGR level.

The disadvantage of the "Type 1" maps is that they do not optimize performance at all, since some combustion modes may be optimized with different EGR and SOI than others. To correct this, the "Type 2" maps will each impose a certain limit of NO_X, with the EGR and SOI chosen to meet this limit. Out of the set of EGR and SOI combinations that meet the NO_X limit, the EGR/SOI combination chosen will be the one that maximizes ITE without also exceeding a certain soot limit (set to 0.1 g/kWh for all NO_X levels). If there are no

EGR/SOI combinations that conform to both the NO_X and soot limits, the combination that minimizes soot will be chosen regardless of ITE. In this way the “Type 2” maps will give a better picture of the optimized performance of each combustion mode.

4.2 Speed/Load “A”: 2300 RPM, 4 bar BMEP

In speed/load “A”, a very wide area of the MELT diagram was available, more so than in any other speed/load point. This was an intentional feature of speed/load “A”, given the results of preliminary study with a combination of propane and dimethyl ether (DME) as the Manifold fuel [11]. Only the very top of the diagram (representing “pure” HCCI) and right hand side (representing “pure” PCCI) cannot be reached within the imposed limits, but both can be achieved if the gross ITE limit is relaxed (which is why there were examples of these both in Chapter 3 at this same speed/load).

4.2.1 Constant 20% EGR

Type 1 contour maps with 20% EGR are shown in Figure 4.4, with the corresponding HRR’s and Wiebe function fits in Figure 4.5. The CoV(IMEP) is very low in the diesel-only CDC and PPCI regions, but immediately jumps up to more than 3% in the lower CDF and PDFC regions, before going back down as more Manifold fuel is added. It is possible that the CoV’s in this region are caused not by the inherent nature of the combustion modes in this region, but by fluctuations in the injected mass of gasoline from cycle to cycle in this region, where the measured flow rate of gasoline through PFI injectors was very sensitive to the injection duration used. However it cannot be ruled out that the low concentrations of gasoline cause some cycle-to-cycle feedback in how much of this gasoline is consumed via autoignition and how much is consumed by the diffusion flame. At the very top of the diagram, the CoV goes up sharply again, accompanied by a sharp rise in UHC and a drop in ITE. This region was much more clearly approaching misfire, with the observed instability of the HRR from cycle to cycle being much more apparent. When almost all of the fuel is homogeneously premixed gasoline, most of it will have to burn through homogeneous autoignition, and without the presence of much DI diesel this is difficult to achieve in such a low load. True HCCI could only be achieved at this load with the maximum 40% EGR, and an unbalanced fueling of the four cylinders to achieve the same combustion phasing with the thermal differences in each cylinder.

Ringing intensity was high (above 5 MW/m^2) with the CDC and lower CDF modes, where there is no heat release before late DI to decrease ignition delay and reduce the rich premixed "spike" that causes sharp PRR's. In the PPCI and lower PDFC modes, the lean stratified premixed autoignition before the late DI brought the RI down much lower (below 2 MW/m^2). In the higher CDF modes, the RI was reduced due to the small size of the late DI, which prevented a large rich premixed spike from forming even without any previous heat release. Conversely in the higher PDFC modes, the RI began to be affected more by the homogeneous lean premixed autoignition, which now became the bigger "spike" in the HRR. This limited access to the RCCI region, with a very tight window of Early DI sizes that would produce stable and complete autoignition without excessive ringing from the homogeneous autoignition.

The gross ITE was found to be maximized along a "ridge" running down the middle of the accessible region of the MELT diagram, from CDC to PDFC to RCCI, maintaining at least 45% ITE with two slight peaks on either ends of the PDFC region. The trend in ITE can be seen as essentially the sum of the cooling + exhaust losses and the UHC emissions. The cooling + exhaust losses explain the ITE drops towards the right vertex, and the UHC emissions explain the ITE drops toward the top vertex, with the max ITE ridge occurring along a line of ideal compromise between these two trends. The cause of these efficiency trends for the different combustion modes will be examined below.

Starting from CDC, the gross ITE went steadily downward going up into the CDF region, where the exhaust losses were slightly improved, but this was outweighed by the dramatic

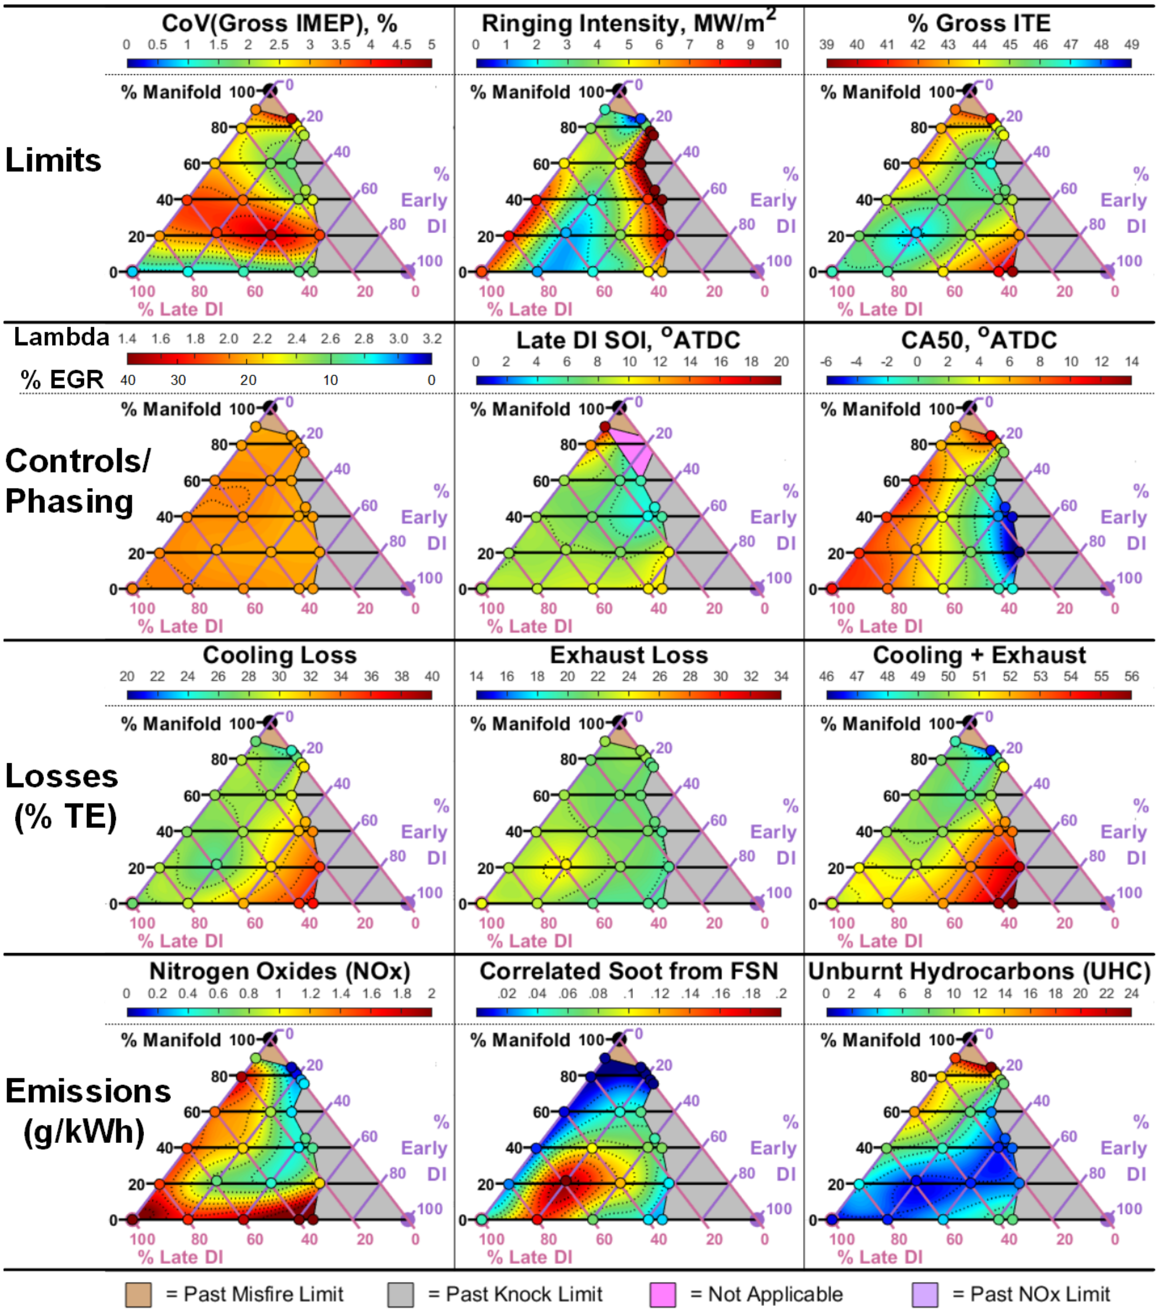

Figure 4.4 Contour maps of speed/load A (2300 RPM, 4 bar BMEP) at constant 20% EGR ("Type 1" maps). CoV is very high with small amounts of Manifold fuel, but goes back down with larger amounts, and the reason for this trend is unclear. There is a "ridge" of high gross ITE (top right plot) running from CDC (bottom left corner of MELT diagram) through PDFC (middle) to RCCI (right edge). This coincides with a NO_X "valley" but also intersects a soot "peninsula." The CA50 gets earlier going toward the right side of the diagram, but the sum of cooling + exhaust losses stays about the same traveling along the high-ITE "ridge." The ITE trend is explained by the sum of the cooling + exhaust losses and the UHC emissions, which account for about a .4% TE loss for every 1 g/kWh of UHC.

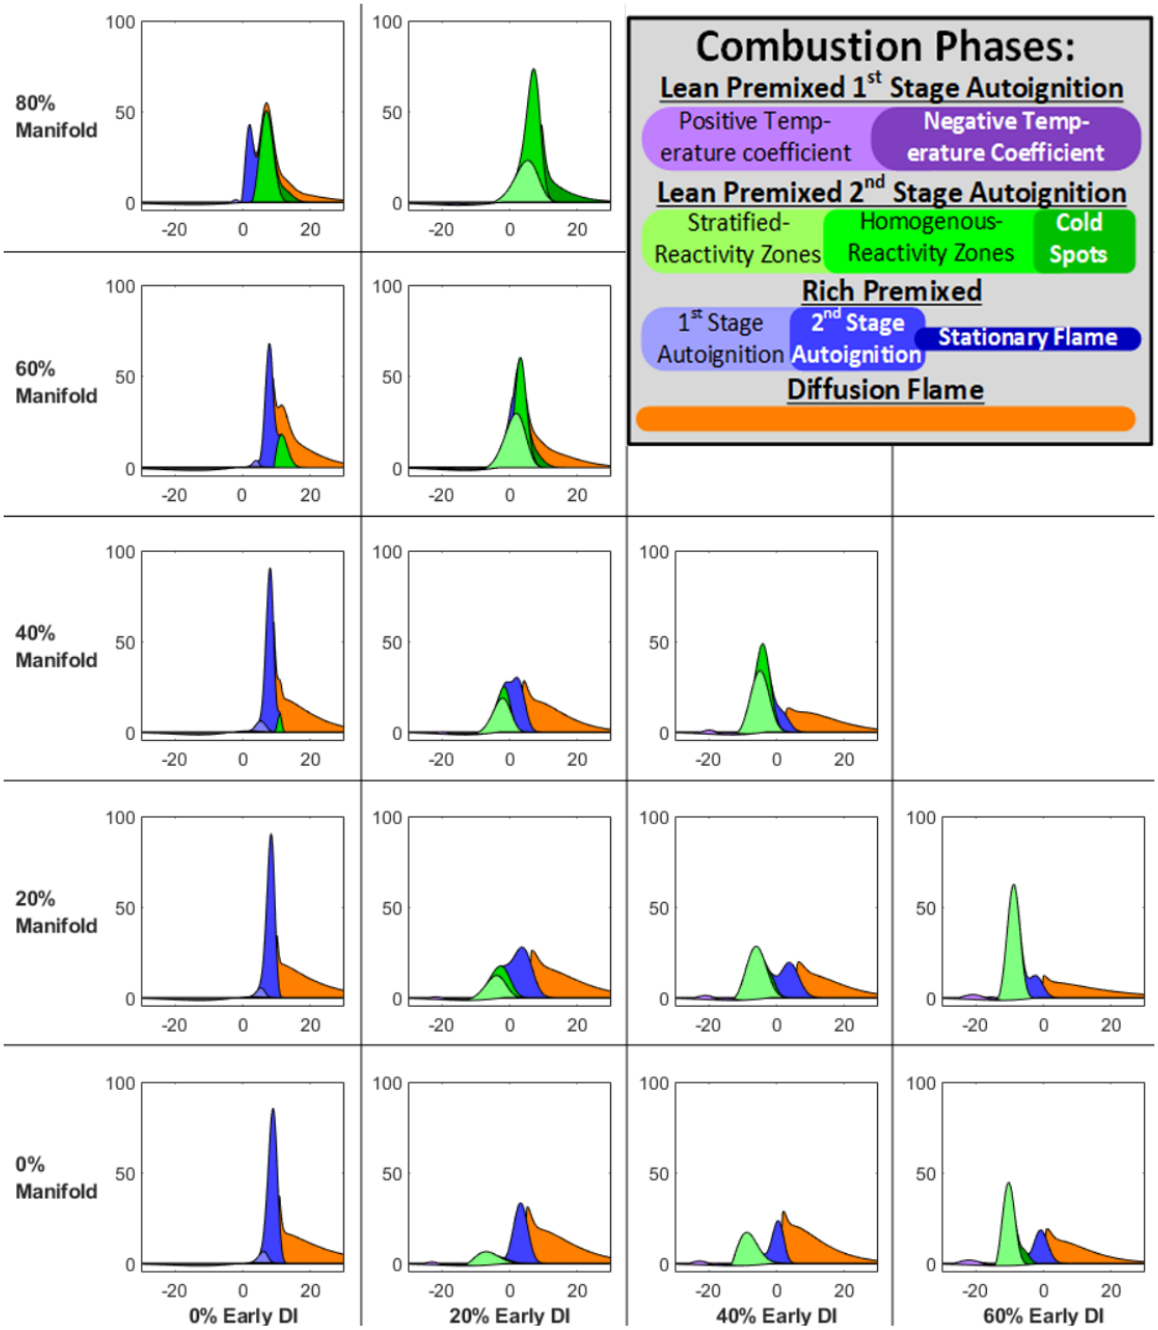

Figure 4.5 Wiebe function breakdowns of test conditions from Figure 4.4. This shows the gradual transition between combustion modes. The lean premixed combustion phases (in green) begin to dominate the overall combustion process as more Manifold fuel and Early DI are added, with more Early DI emphasizing the stratified phase (light green), and more Manifold fuel emphasizing the homogeneous phase (solid green). The points with wider, flatter HRR profiles are associated with much lower ringing intensity, compared to the points with sharp HRR "spikes." No Early DI (left side) leads to long ignition delay and a large rich premixed "spike," while too much Early DI (right side) creates a large stratified lean premixed "spike," and too much Manifold fuel (top) creates a large homogeneous premixed "spike."

increases in UHC. This can be attributed to homogeneously premixed gasoline in the squish region and especially the piston-cylinder crevice, which will not autoignite easily without any Early DI being used. This trend is slightly reversed at the point of 80% manifold fuel, where the gasoline concentration was high enough to get significant lean premixed autoignition, but even still the UHC emissions were beyond 12 g/kWh, which corresponds to a 5% loss in thermal efficiency. This unburnt loss outweighs a slight (2%) improvement in the exhaust loss, which could merely be the result of lower exhaust temperatures from the incomplete combustion. It is possible that higher gasoline concentrations could lead to further reductions in the UHC loss, but at this constant low load the concentration could not be increased before eliminating the Late DI that was triggering that lean autoignition, and at 10% Late DI the misfire limit was approached. The Late DI also must be advanced significantly in these ”heavy” CDF modes, with the high concentration of premixed gasoline leading to large increases in ignition delay.

Meanwhile on the bottom edge of the diagram, starting from CDC and moving into the PPCI region, ITE declined steadily the whole way. This was partially due to an increase in UHC, but much less so than in CDF since the premixed fuel in this case was diesel, which autoignites much more readily than gasoline. Much more dominant was the effect of the increased cooling losses, caused by the advancement of CA50. With CA50 advanced, there is more time for wall heat transfer, and although this also reduces exhaust temperatures and expands the P-V diagram, it appears that the cooling loss increase dominated the ITE trend in the PPCI mode. The lean premixed autoignition appears to happen too soon in this mode, peaking between -10 and -5°aTDC, which stretches out the total heat release over a large window of crank angles. Ideally, for optimal P-V curve utilization, this heat release would be more concentrated in a small window of crank angles. When this early peak becomes too big and too early, there is too much heat release before TDC, and even though the RI does not hit the limit, ITE begins to drop dramatically when the piston has to do too much work at the end of the compression stroke, and this can be considered part of the “knock” limit.

In the PDFC mode, the presence of low-reactivity gasoline delays the lean premixed autoignition to peak closer to TDC, which concentrates the heat release into a tighter window. This reduces the trend of increased cooling loss with advancing CA50 as compared to the PPCI mode, without significantly increasing the exhaust loss, creating a better cooling + exhaust trend with respect to CA50. This also retards the knock limit of SOI, since advancing the SOI will "stack" the lean premixed and rich premixed peaks on top of each other near SOI for high MPRR. Despite adding more premixed gasoline, the UHC emissions are improved over both PPCI and CDF for the same amounts of premixed fuel (same % Late DI). This is unsurprising when compared to CDF, since the addition of premixed HRF to replace LRF increases the combustion efficiency of the lean premixed autoignition. However, it is surprising when compared to PPCI, since premixed LRF has replaced HRF, which intuitively should increase UHC emissions. One explanation for this could be that since there is a homogeneous background concentration of fuel with the Manifold Injection, there are no areas where the mixture fraction in the premixed fuel winds up too lean for autoignition to occur, even with an HRF.

Moving toward the right edge of the diagram and the RCCI mode, CA50 starts to become too early, and for most levels of Early DI the RI limit is reached before the Late DI is completely eliminated. Only within a tight window of 17.5 to 25% Early DI can there be true RCCI with no Late DI, and within this window the trends with respect to fuel ratio are very sensitive. The cooling losses are at their absolute lowest when the lowest percentage of Early DI is used, since this distributes the temperatures of combustion the most evenly throughout the chamber, but this also causes very high UHC emissions correlating to about a 10% loss of TE. Conversely, when the highest percentage of Early DI is used, the cooling loss is increased by about 6% while the unburnt loss is reduced by a similar amount, due to the CA50 being advanced to near TDC and maximizing peak temperatures. The highest ITE from RCCI occurs in the middle, where there is an optimal trade-off between these two trends.

The NO_X emissions from CDC, CDF, and PPCI are consistently quite high, all well beyond the target of 1 g/kWh. These modes all feature a Late DI, which inherently produces a diffusion flame with near-stoichiometric combustion generating thermal NO_X. One would expect that NO_X would decline with a reduced amount of Late DI, but this does not appear to be the case for either CDF or PPCI. In the "lighter" CDF modes (with less gasoline), the premixed gasoline does not autoignite and must still be consumed through the diffusion flame, and although there is a slight drop in NO_X this may just be the result of higher UHC and thus less diffusion combustion overall. In the heavier CDF modes, one would think that the increasing lean premixed autoignition would produce lower NO_X, but in these modes the Late DI must be significantly advanced to avoid knocking, shifting the diffusion combustion to earlier in the cycle. The relatively quick lean autoignition also produces higher temperatures in the diffusion flame, and as a result NO_X from CDF does not go down below 1 g/kWh until it approaches the misfire limit.

In the lighter PPCI modes, there is a slight drop in NO_X as heat release is shifted toward the lean premixed peak, but since this peak is occurring so early, there may not have been time for the Early DI to sufficiently mix with air before autoigniting. As this peak gets earlier and larger with increasing Early DI, the NO_X goes back up, as the lean premixed combustion actually begins to generate NO_X itself. Without enough time to mix before autoigniting, there will still be regions of near-stoichiometric mixture fraction at autoignition, and these zones will produce NO_X just like the diffusion flame.

PDFC, however, shows significant decreases in NO_X for the same levels of premixed fuel as CDF and PPCI. When compared to CDF, PDFC is consuming the majority of the premixed fuel before the Late DI occurs, preventing it from contributing to the diffusion flame and thus NO_X generation. When compared to PPCI, PDFC is delaying the lean premixed autoignition and allowing more time for mixing to occur, lowering the peak mixture fraction and reducing NO_X generation from the lean autoignition. Also, compared to PPCI, PDFC is spreading out this lean premixed autoignition over a wider area, not just the squish region where the Early DI is aimed but the bowl area as well. Not only does this reduce peak temperatures, but it also decreases the concentration of oxygen in the bowl that is left behind for the diffusion flame, thus decreasing the oxygen concentration for NO_X formation and diluting the flame with CO_2.

Unfortunately, this last effect has the unwanted side effect of increasing soot emissions of PDFC compared to both CDF and PPCI. To understand this, one must consider both soot *formation* that occurs during the rich premixed autoignition and soot *oxidation* that occurs during the diffusion flame to reduce soot concentration to the final engine-out level.

Soot formation during CDF is reduced due to a decrease in the Late DI. Soot reduction may be slightly inhibited during the diffusion flame if the homogeneous lean premixed autoignition occurs and reduces the available concentration of oxygen, but this will also increase temperatures to promote soot oxidation. It appears that the overall effect is overwhelmingly positive for CDF, with soot falling continuously for higher and higher CDF.

In PPCI, soot is initially increased when the Early DI is first introduced. The heat release from the Early DI increases the temperatures during the rich premixed combustion phase, promoting soot formation. It also reduces the ignition delay and thus the available time for fuel/air mixing before autoignition, leading to a richer mixture at autoignition and thus increased soot formation. This is compounded by a reduction in the available oxygen during the diffusion flame, which leads to inhibited soot oxidation. However, after the initial rise in soot from 0 to 20% Early DI, further increases in the Early DI appear to lower soot back down to the level of CDC as more and more of the Late DI is eliminated, reducing the rich fuel mixtures available for soot formation and increasing temperatures during the diffusion flame to promote soot oxidation. This unfortunately also has the effect of increasing NO_X production, unless the Late DI is retarded. A study by Hasegawa et al. [54] on a similar PPCI condition with 50% Early DI and 50% Late DI diesel was able to achieve 1/2 the soot and 1/4 the NO_X of CDC, but only by retarding the SOI to 13°aTDC, which limits ITE.

The same soot trends for PPCI are observed in PDFC, but to a greater degree. One potential reason for this is that the lean premixed autoignition is occurring more in the bowl region, where there is Manifold fuel, as opposed to the squish zone, where there is Early DI as well. Thus when the lean premixed autoignition occurs, it consumes more oxygen and generates more heat *in the places where the Late DI is about to occur*, magnifying the effects of the temperature increase promoting soot formation and the oxygen decrease inhibiting soot oxidation. Only "heavy" PDFC with less than 20% Late DI can produce soot as low as the original CDC condition, and only RCCI or CDF can reduce soot below CDC levels. However, these heavy PDFC conditions can produce gross ITE similar to CDC with greatly reduced NO_X emissions and only slight increases in UHC compared to CDF and RCCI.

One might see these high soot levels as the "Achilles' heel" of PDFC, despite this being the only mode that can produce ITE as high as CDC. But such statements cannot be made before making a fairer comparison of the two combustion modes using the "Type 2" maps, which equalize the level of NO_X being produced across the MELT diagram.

4.2.2 1.0 g/kWh NO_X Limit

When a limit of 1.0 g/kWh NO_X is applied in Figure 4.6, PDFC actually has advantages over CDC, CDF, and PPCI in terms of soot emissions. This is because, in order to reach the NO_X limit, the CDC, CDF, and PPCI modes must all use high levels of EGR, bringing their respective lambdas down below 2 and decreasing the oxygen concentration at the diffusion flame to inhibit NO_X formation. The resulting higher intake air temperatures will lead to higher soot formation, while the reduced oxygen concentration will both increase soot formation and inhibit soot oxidation.

However, as discussed in Section 4.2.1 above, the oxygen concentration at the diffusion flame in PDFC is already reduced by the lean premixed autoignition propagating into the bowl. This allows PDFC to use much higher lambdas and thus much more excess air than

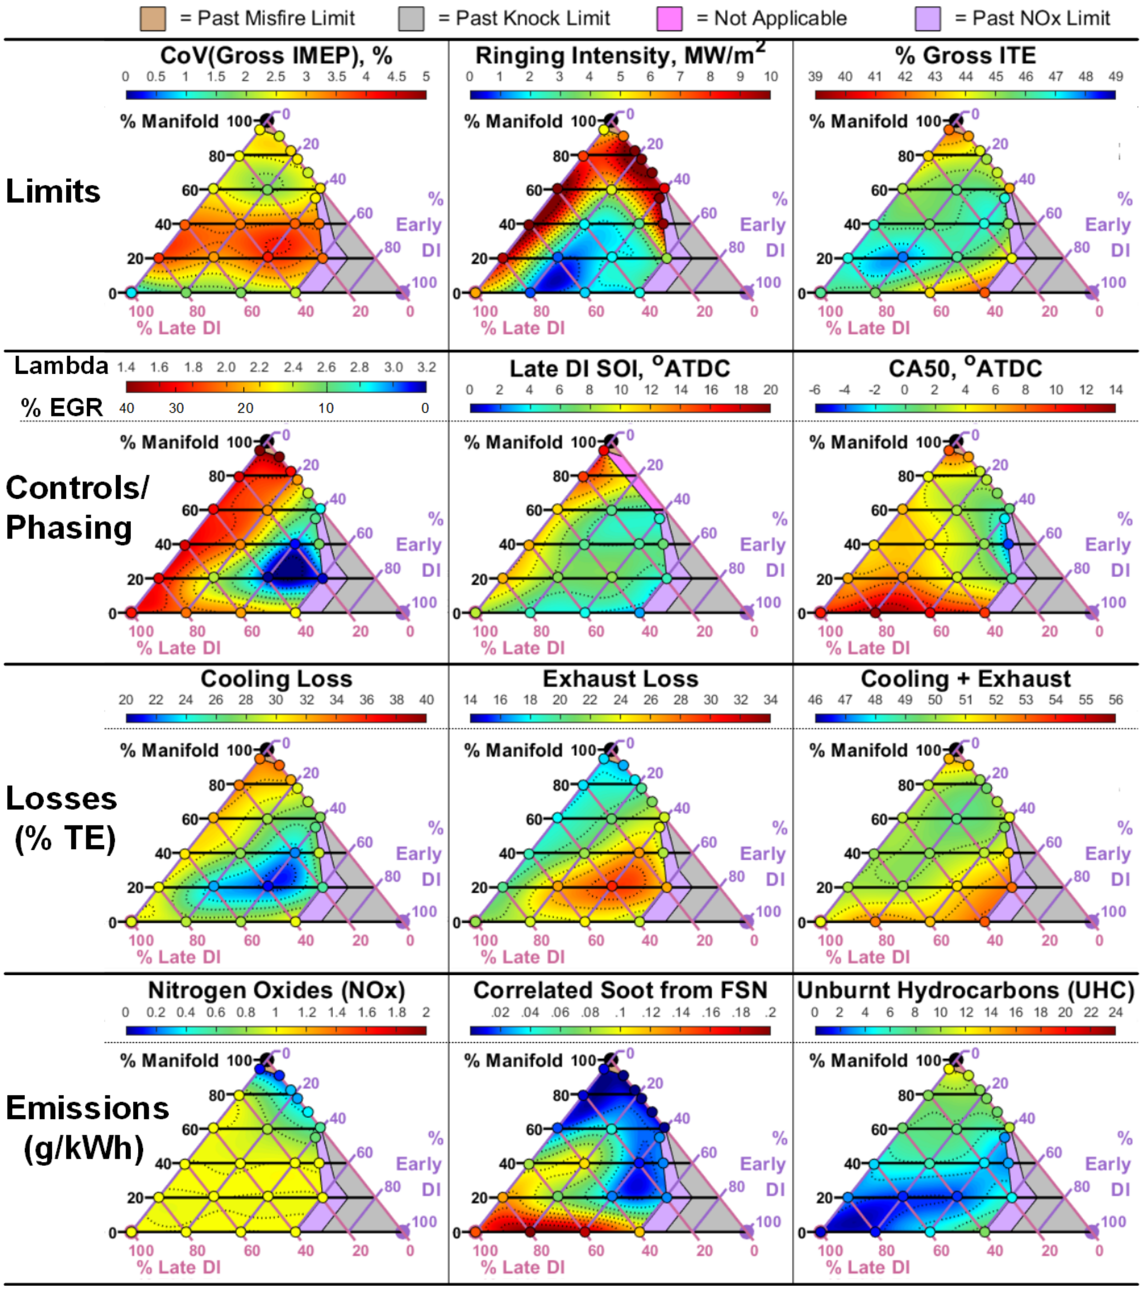

Figure 4.6 Contour maps of speed/load A (2300 RPM, 4 bar BMEP) at limit of 1 g/kWh NO_X. Trends are discussed on pages 95-97.

the other modes in order to achieve the same NO_X level, with higher lambdas available at the higher levels of Early DI. However, if lambda gets too high, this will lead to excess CO emissions and exhaust losses. For this reason, there appears to be a “sweet spot” with a lambda of about 2.5 at 20% Manifold, 20% Early DI, and 60% Late DI that maximizes gross ITE with an approximate 2% ITE improvement over CDC. At the same time, soot emissions are more than cut in half and there is only a slight increase in UHC.

The changes in lambda in the Type 2 map are closely following by the trends in cooling loss (which goes down with increasing lambda) and exhaust loss (which goes up). Thus, what is really of interest is the cooling + exhaust sum, which shows the trade-off between those two losses. Previously, in the Type 1 maps, this trend clearly favored the top of the MELT diagram, but now in the Type 2 map the best trade-offs have shifted more toward the middle of the MELT diagram, where there is less cooling loss than before due to the increased lambda. This further validates the theory that reductions in cooling loss outweigh increases in exhaust loss when trying maximize the efficiency of these dual-fuel modes.

However, one advantage of the low lambdas needed for CDF is that they promote autoignition of the homogeneous premixed fuel, whose concentration is now closer to stoichiometric. This brings the UHC emissions of CDF much closer to PDFC and increases ITE, although it does exacerbate the high RI. The lower lambda leads to a longer ignition delay, with more time for the Late DI to mix with air and hence more heat release in the rich premixed autoignition. Although the increased dilution from EGR actually serves to “slow down” the rich premixed autoignition and produce a wider and shorter spike in the HRR, it will also slow down the diffusion flame, and thus the SOI needs to be advanced much further in order to get an optimal CA50. As a result, the rich premixed spike still ends up higher in CDF than in the Type 1 maps, producing the high RI’s.

High RI is also produced in the RCCI conditions, where unfortunately in order to minimize UHC and maximize ITE the CA50 needs to occur close to TDC. Since the CA50 in RCCI closely corresponds to the peak of homogeneous premixed autoignition, this means maximizing ITE will also maximize RI. This highlights an additional benefit of PDFC: since it is “splitting” the heat release into two separate phases with two separate peaks, it is much easier to get CA50 close to TDC without producing high MPRR and high RI. This produces problems for efficiency if the two peaks become “spread out” as they are in PPCI, but PDFC generally tends to keep these two peaks closer together.

One other thing to note about RCCI is that it *cannot* produce NO_X emissions as *high* as 1.0 g/kWh. Thus, this is still not a truly fair comparison of RCCI to the other modes in Figure 4.6. To do this, the NO_X limit must be reduced further.

4.2.3 0.3 g/kWh NO_X Limit

When the NO_X limit is reduced to 0.3 g/kWh, shown in Figure 4.7, the NO_X is now more or less constant across the entire MELT diagram. For all modes except RCCI, it is impossible to do this without heavy EGR to reduce lambda below 2, which results in unacceptably high soot emissions for almost every MELT point besides the RCCI conditions. The only non-RCCI points that can do this are the very heavy CDF conditions, where the ITE is too low to make it worthwhile. This shows that, in order to meet modern NO_X emissions standards without aftertreatment, RCCI is the only option for CI engines. However, it is still worthwhile to examine the other trends in order to gain more insight into the reasons behind the performance changes.

For the low NO_X limit, the maximum ITE is actually produced by the middle CDF conditions, which produce about 3% more ITE than CDC. This is because, unlike CDC, these conditions can have a much more advanced SOI and CA50 before the strict NO_X limit is reached. The very low lambda and corresponding near-stoichiometric mixture fractions and high intake temperatures help the premixed manifold fuel autoignite completely. However, PHCCI is never produced in these conditions, since the Late DI is always needed to trigger the premixed autoignition. Surprisingly, these conditions actually feature the lowest *exhaust* losses, bucking the trend of lower cooling losses dominating the ITE trend. Furthermore, going along the line of 0% Early DI (the left side of the MELT diagram), even though lambda is constant, the *same* trend is seen in both cooling loss and exhaust loss: losses go down going up toward 60% Manifold fuel, then back up. The trend in exhaust loss makes sense given the trend in CA50, but the trend in cooling loss does not, with *advanced* CA50 causing *reduced* cooling loss; the opposite of what would be expected.

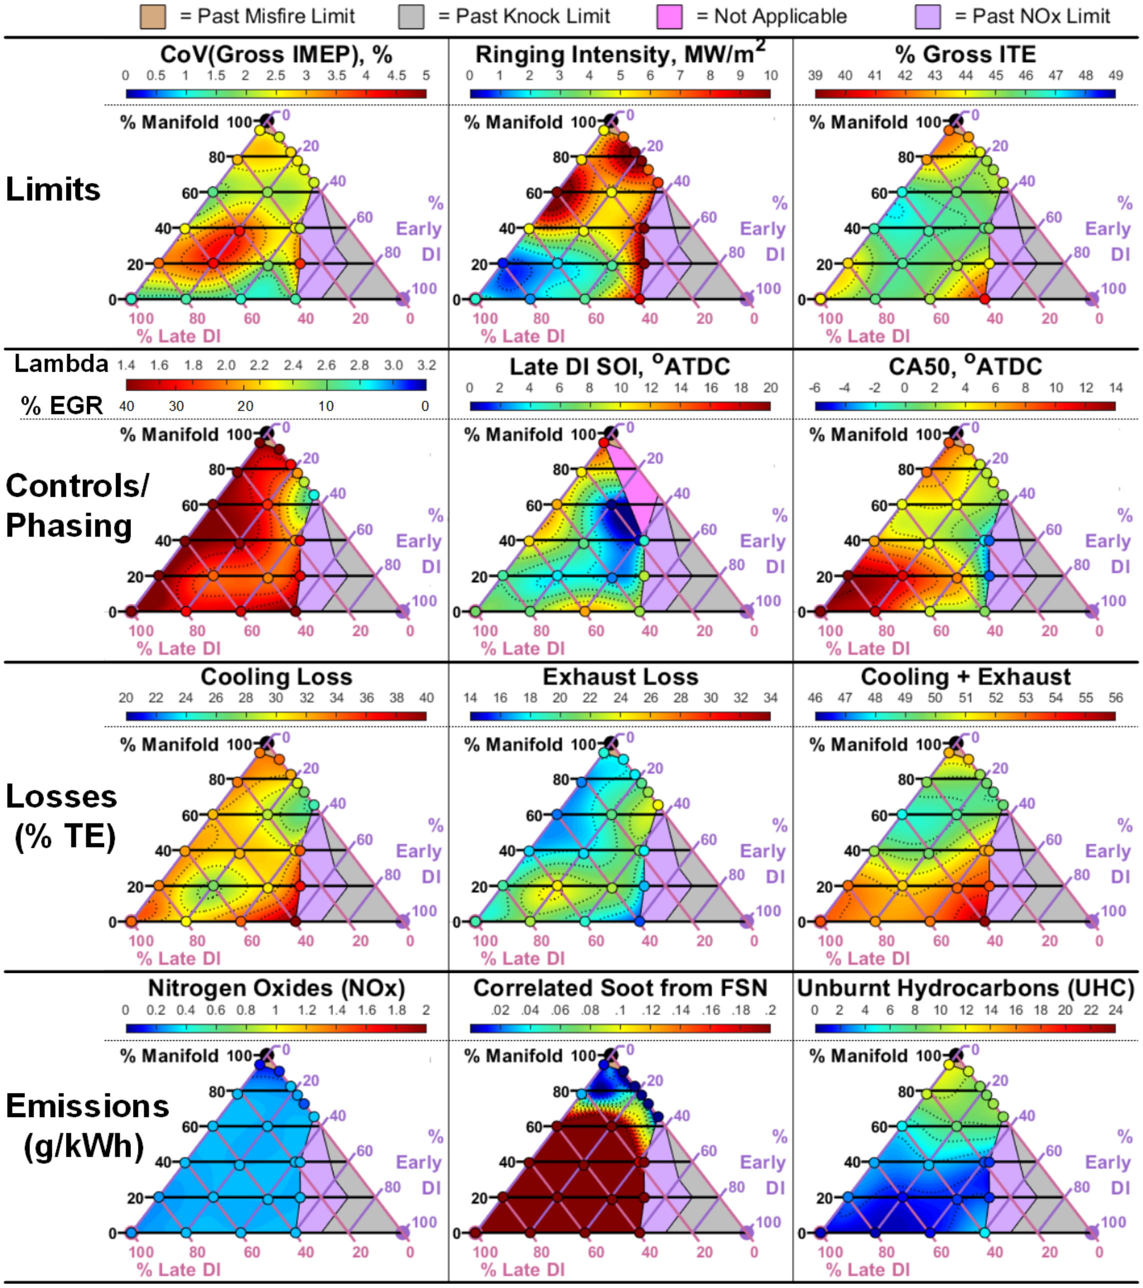

Figure 4.7 Ccontour maps of speed/load A (2300 RPM, 4 bar BMEP) at limit of 0.3 g/kWh NO_X. Trends are discussed on pages 98-100.

One possible explanation for this trend is that when the CA50 of CDF is advanced, more of the heat release is shifted into the homogeneous lean premixed autoignition. This form of combustion produces the minimal overall heat transfer losses, since it distributes the temperatures of combustion as evenly as possible throughout the combustion chamber, without concentrated high temperatures close to the wall. Thus when the Late DI is advanced in these conditions, even though the combustion phasing is shifted earlier to increase cooling losses, this is outweighed buy the reduction in cooling losses from the shift away from diffusion combustion towards lean premixed combustion. The key to taking advantage of this will be to generate the same amount of homogeneous premixed combustion without such a low lambda or less of a Late DI, so that these benefits are realized without producing excessive amounts of soot in the process.

4.3 Speed/Load “F”: 1200 RPM, 4 bar BMEP

In speed/load “F”, a much narrower area of the MELT diagram was available as compared to “A”. This is understandable when considering that the lean premixed autoignition processes of the dual-fuel modes are controlled mostly by chemical kinetics, and will be happening at mostly the same rate with respect to time as they were in “A.” However, since the engine speed has been cut nearly in half, the same processes will be happening almost twice as fast with respect to *crank angle* as in “A.” Thus, the same amount of Early DI will produce much earlier combustion phasing in terms of crank angle compared to “A,” and the maximum amount of Early DI that produces acceptable combustion phasing will be much lower than “A.” For this reason the MELT diagrams will be stretched to better show the left side, as shown in Figure 4.2.

4.3.1 Constant 30% EGR

Although the Type 1 maps for A were at a constant level of 20% EGR, the Type 1 maps for F will be at a higher level of 30% EGR, for a lambda of about 1.7 throughout. This was done to counteract the decreases in engine speed and boost, which both serve to reduce temperatures near TDC. This inhibits the lean premixed autoignition, which the additional EGR balances out by increasing intake temperature.

The CoV of IMEP does not appear to be a limiting factor in this condition, except when approaching 100% Manifold fuel. This seems to contradict what was observed in speed/load A, when there was a large increase in CoV at about 20% Manifold fuel. The increase seen at the higher engine speed had been attributed to cycle-to-cycle variation in the amount of PFI fuel injected, which was extremely sensitive to injection duration at this small injection size. But despite there being no difference in load (and thus no difference in the size of each individual injection), if anything there is a barely noticeable *decrease* in the cycle-to-cycle variation at 20% Manifold fuel in F. It will be necessary to examine this further when performing repeated testing of these two conditions in the next chapter.

RI also does not appear to be a limiting factor in F, never breaching 7 MW/m^2. This can be attributed to the lower engine speed, which lowers the turbulent gas velocities in the charge and moves to lower levels of boost and rail pressure on the engine map (see Table 4.1). With lower rail pressure, the Late DI requires much longer injection durations to reach the same amount of fuel as in A, which combined with the lower turbulent gas velocities leads to much slower fuel/air mixing. This increases the ignition delay, but reduces the amount of fuel that will be injected before autoignition occurs, while the lowered boost pressure decreases the temperature and thus the rate at which the autoignition will occur. The slower engine speed also increases the cylinder heat transfer that will occur during the compression stroke, further reducing temperature and decreasing the rate of autoignition. Although the maximum HRR spikes in Figure 4.9 appear to be about the same height as in speed/load A, these are plotted in terms of Joules per *crank angle*, and their height in terms of Joules per *second* would be much lower than in speed/load A.

The main limiting factor for speed/load F is always ITE, which goes below the limit well before RI and CoV at all of the boundaries. Just like in speed/load A, there appears to be a "ridge" of high ITE going through the middle of the available area of the MELT diagram from CDC through PDFC, although unlike speed/load A the ITE drops precipitously in RCCI. It appears that RCCI is not as useful in speed/load F as it was in speed/load A, and cannot reach an ITE comparable to the other combustion modes. The UHC emissions never

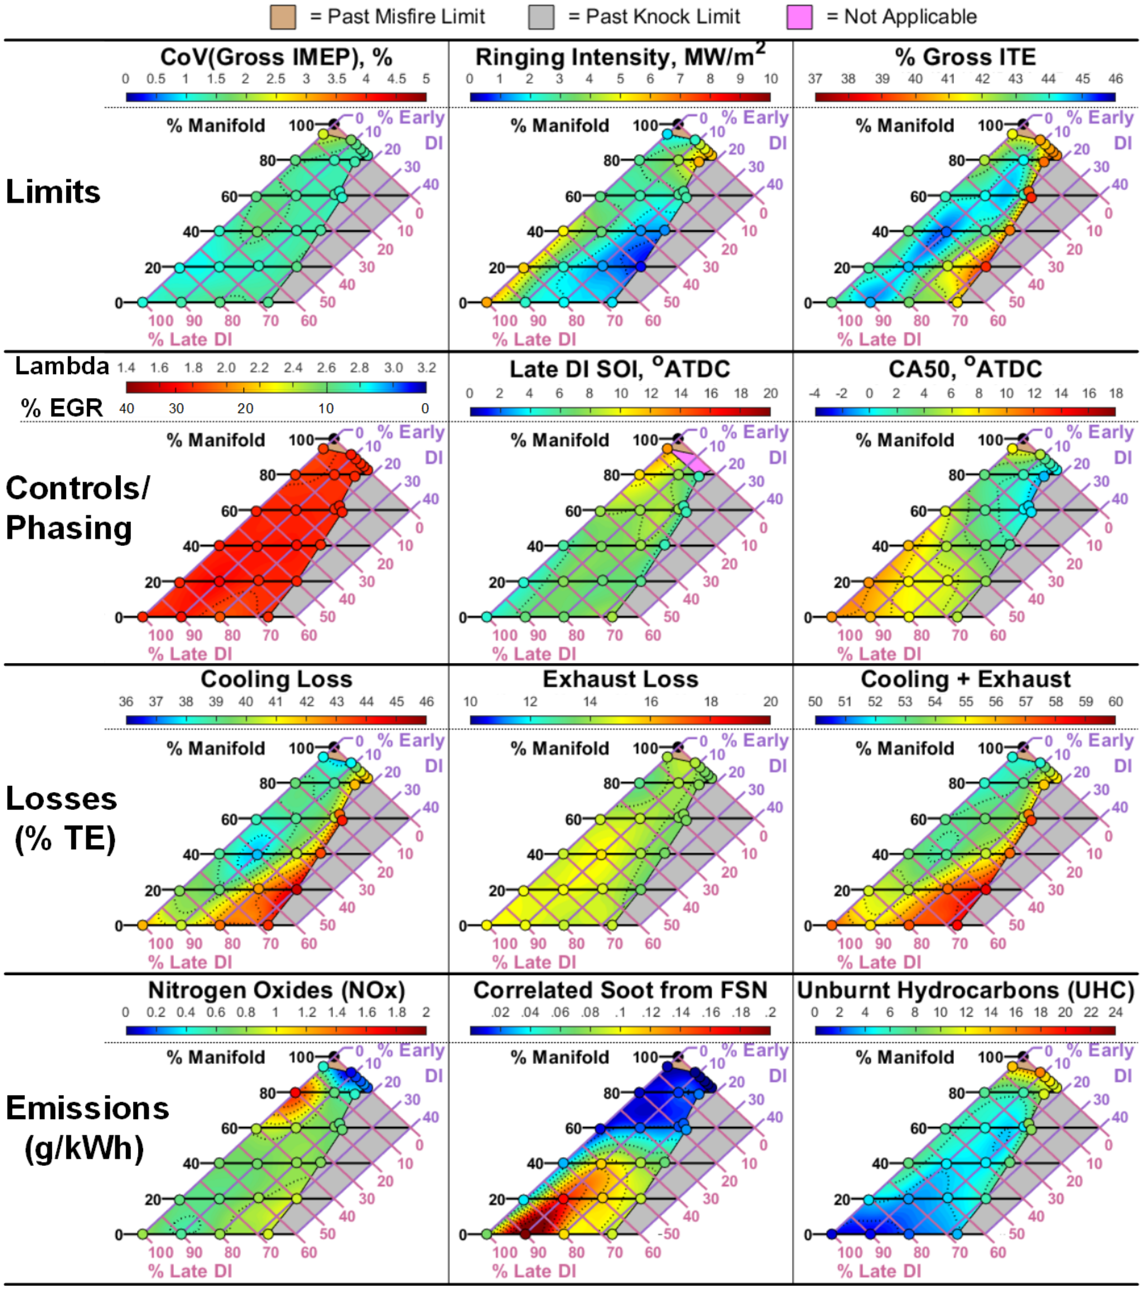

Figure 4.8 Contour maps of speed/load F (1200 RPM, 4 bar BMEP) at constant 30% EGR. Trends are discussed on pages 100-106.

reach less than 10 g/kWh, so the unburnt loss is always at least 4% of the total thermal efficiency. The reason for these unavoidable unburnt losses is likely the high cooling losses seen at this low engine speed - even the lowest cooling losses are 38% TE, whereas cooling losses in speed/load A could be as low as 26% TE. With the lower engine speed, there is more time for heat transfer away from the cylinder, thus lower temperatures near TDC, and

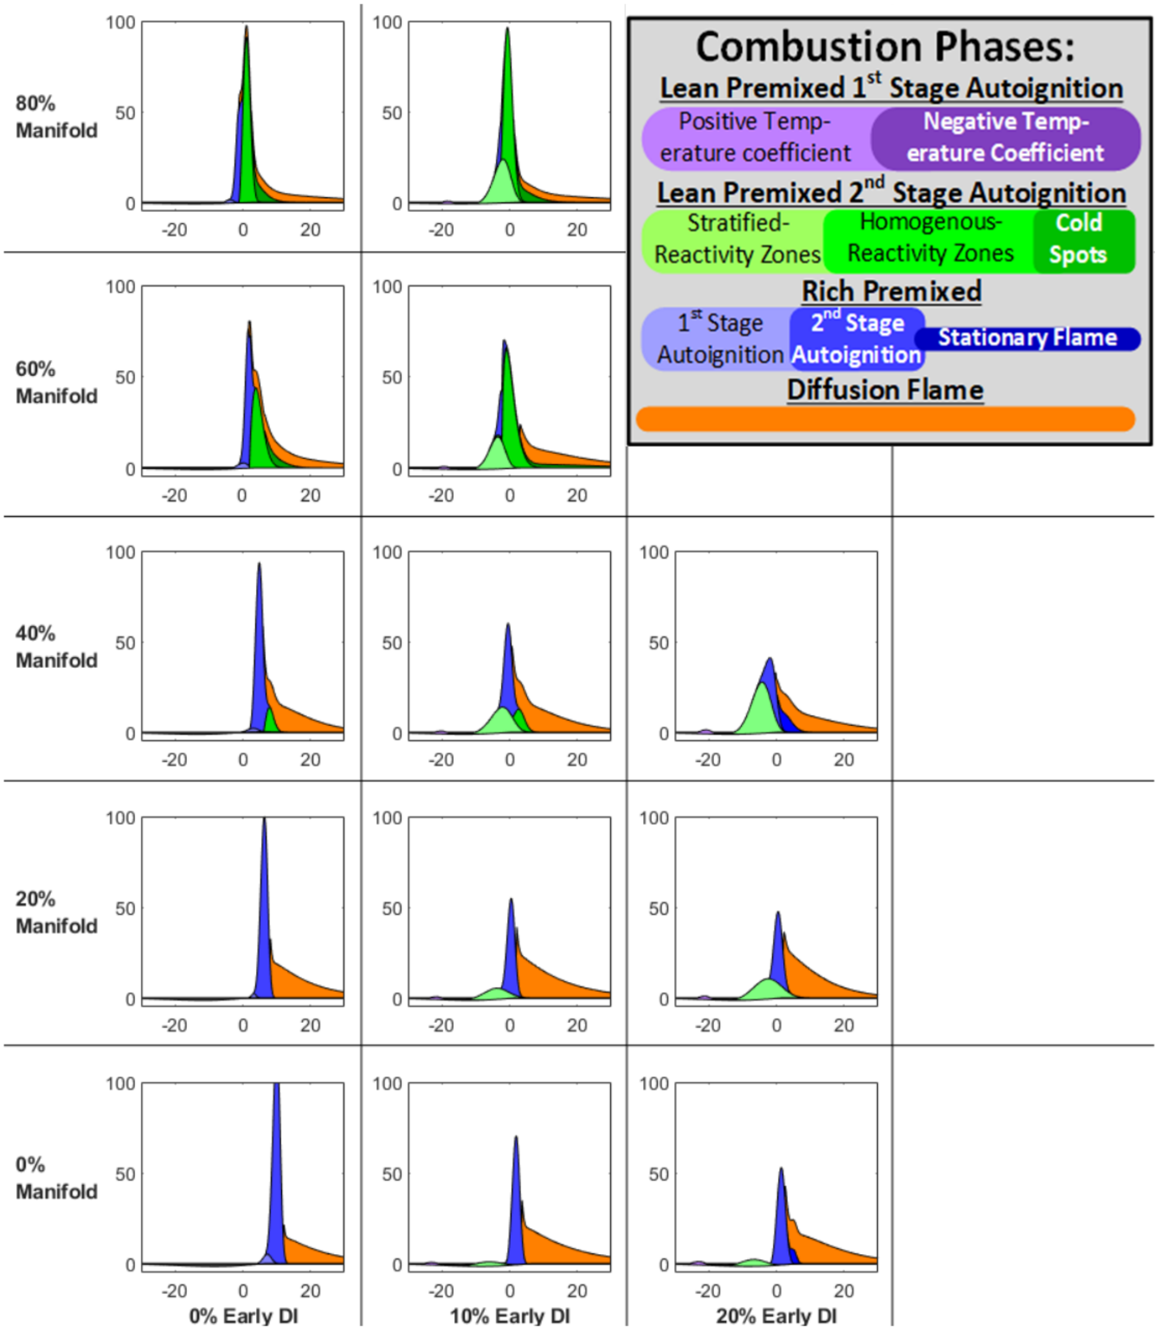

Figure 4.9 Wiebe function breakdowns of test conditions from Figure 4.8. Trends are discussed on pages 100-106.

thus less complete autoignition of the premixed fuel. This could be helped by the addition of Early DI, but this advances the combustion phasing too far before it can overcome the unburnt losses - even 20% Early DI is too much at any EGR level for RCCI in speed/load F, whereas as much as 40% Early DI could be used for RCCI in speed/load A.

Adding just 10% Late DI drastically lowers the UHC with only a slight rise in cooling loss, leading PDFC to have much higher ITE than RCCI. Just a little heat release from the Late DI helps the homogeneous lean premixed autoignition happen much more quickly and at slightly higher temperatures, and the chamber is not as cool at the end of the process, leaving fewer UHC behind. Similar to speed/load A, there appears to be a “sweet spot” with maximized thermal efficiency due to minimal cooling losses toward the middle of the diagram. However, the advantage of PDFC over CDF is not as great in speed/load F as it was in speed/load A, with fewer UHC. This may be due to the lower engine speed of speed/load F with respect to speed/load A - even though this causes lower peak temperatures at TDC, the engine spends *more time* close to TDC, and maintains its peak temperatures for longer. This allows the homogeneous premixed autoignition more time to run to completion before the piston descends and cools down the charge.

To summarize, it appears that the lower engine speed is beneficial for CDF, but detrimental to RCCI, at least in terms of thermal efficiency. However, both are still less efficient than PDFC regardless of engine speed, at least in terms of the “Type 1” maps with constant lambda.

The lower engine speed also appears to be very beneficial for NO_X emissions, which are now below 1 g/kWh at all but one point on the MELT diagram. The NO_X trend appears relatively flat except for the one aforementioned point and the RCCI and near-RCCI points, where NO_X drops to nearly zero. The one exception to this trend is the heavy CDF point with 80% Manifold fuel, which features a sharp spike of homogeneous premixed autoignition right at TDC. Although this type of combustion normally minimizes NO_X emissions, in this one case it is occurring with too much EGR and too advanced of an SOI, which maximizes the temperature and minimizes the lambda, putting it in the NO_X formation region of the phi-T diagram (recall Figure 1.2). The reason that this combustion phasing has to be so advanced compared to the other CDF points is that the combustion efficiency is very sensitive to SOI, with the misfire limit being far advanced. Heavy CDF in speed/load F is hence operating on a “knife-edge” due to the low temperatures of the low engine speed. Advance the SOI just a little bit, and the premixed combustion is occurring right at TDC, and generating NO_X. Retard the SOI just a little bit, and the autoignition will be incomplete and inefficient, leaving behind high UHC. This shows how the “Type 1” maps are not perfect for comparing combustion modes, since the ideal SOI timing for each mode may not occur exactly halfway between the knock and misfire limits, and in this case of heavy CDF the SOI may need to be moved closer to the misfire limit to avoid NO_X.

There is still a slight NO_X “valley” running through the middle of the PDFC region, but it is not as promiment as it was speed/load A. It is important to remember that, while the lean premixed autoignition processes are being sped up with respect to crank angle at the lower engine speed, the rich premixed autoignition and diffusion combustion processes are *not*. Instead, the speed of these processes correlates with the turbulent gas velocity, which scales relatively linearly to the mean piston speed (a direct function of the engine speed). So even though the lean premixed autoignition is happening within a shorter window of crank angles, the diffusion flame is happening over a relatively similar span of crank angles, if not longer, due to the higher cooling losses. This pushes the NO_X generated by each process in opposite directions - NO_X from the lean premixed process is increased by being squeezed into a tighter window close to TDC, while NO_X from the diffusion process is inhibited by being subjected to cooler temperatures on average. Even though the balance still favors the lean premixed process, the trade-off is not as advantageous as it was in A, giving PDFC less of an advantage in terms of NO_X emissions.

Unfortunately, although the NO_X “valley” has mostly disappeared, the soot “ridge” in PPCI and PDFC remains just as prominent as it was in speed/load A. Again it appears that when the Early DI is first added, there is an increase in temperature during the rich premixed spike and less of an ignition delay, producing hotter and richer local zones and much higher soot production. Adding more Early DI leads to a soot decrease as there is less fuel in the Late DI and thus less rich premixed combustion. However, unlike speed/load A, the highest soot is found with PPCI, rather than the "light" PDFC points with 20% Manifold fuel as was the case in speed/load A. The reason for this could be that the homogeneous premixed fuel is autoigniting in speed/load A at the 20% Manifold fuel points, but not in speed/load F. If the homogeneous premixed fuel does not autoignite, this leaves more oxygen in the bowl region for oxidizing the soot produced by the Late DI.

But in the end, there is really little difference in the shape of soot contours in speed/load F versus speed/load A, just that the peak soot production happens with 20% Early DI instead of 10%. However there were not any points taken with 10% Early DI in A, it is possible that this is where the true peak soot production lies. Regardless, what would be good to see in the following Type 2 maps is the same shift that happened in A - more excess air added to the PDFC conditions to achieve lower soot at a constant NO_X level.

4.3.2 1.0 g/kWh NO_X Limit

When the 1.0 g/kWh NO_X limit is applied, as seen in Figure 4.10, soot is reduced somewhat in the lower portion of the MELT diagram by decreasing EGR and increasing lambda, but not nearly to the degree seen in speed/load A. When lambda is increased too much, the lean premixed autoignition is diminished, leading to higher UHC, NO_X, and lower ITE. This was less of a problem in speed/load A, where there was the ability to use more Early DI before combustion phasing was advanced too far, and ensure that the lean premixed autoignition did not leave too many UHC even with high lambda. Thus in speed/load A the "ridge" of high ITE would overlap with a soot "valley", but here in speed/load F the ITE ridge overlaps with a soot "ridge" instead.

In the upper portion of the MELT diagram, however, the highest ITE's are actually found with NO_X emissions *below* the limit, using the highest possible EGR. This is again probably related to the lower overall temperatures seen in the low speed condition of speed/load F, and as there is more reliance on the lean premixed autoignition, the intake temperature must be raised by raising EGR to ensure complete combustion and minimal UHC. The increases

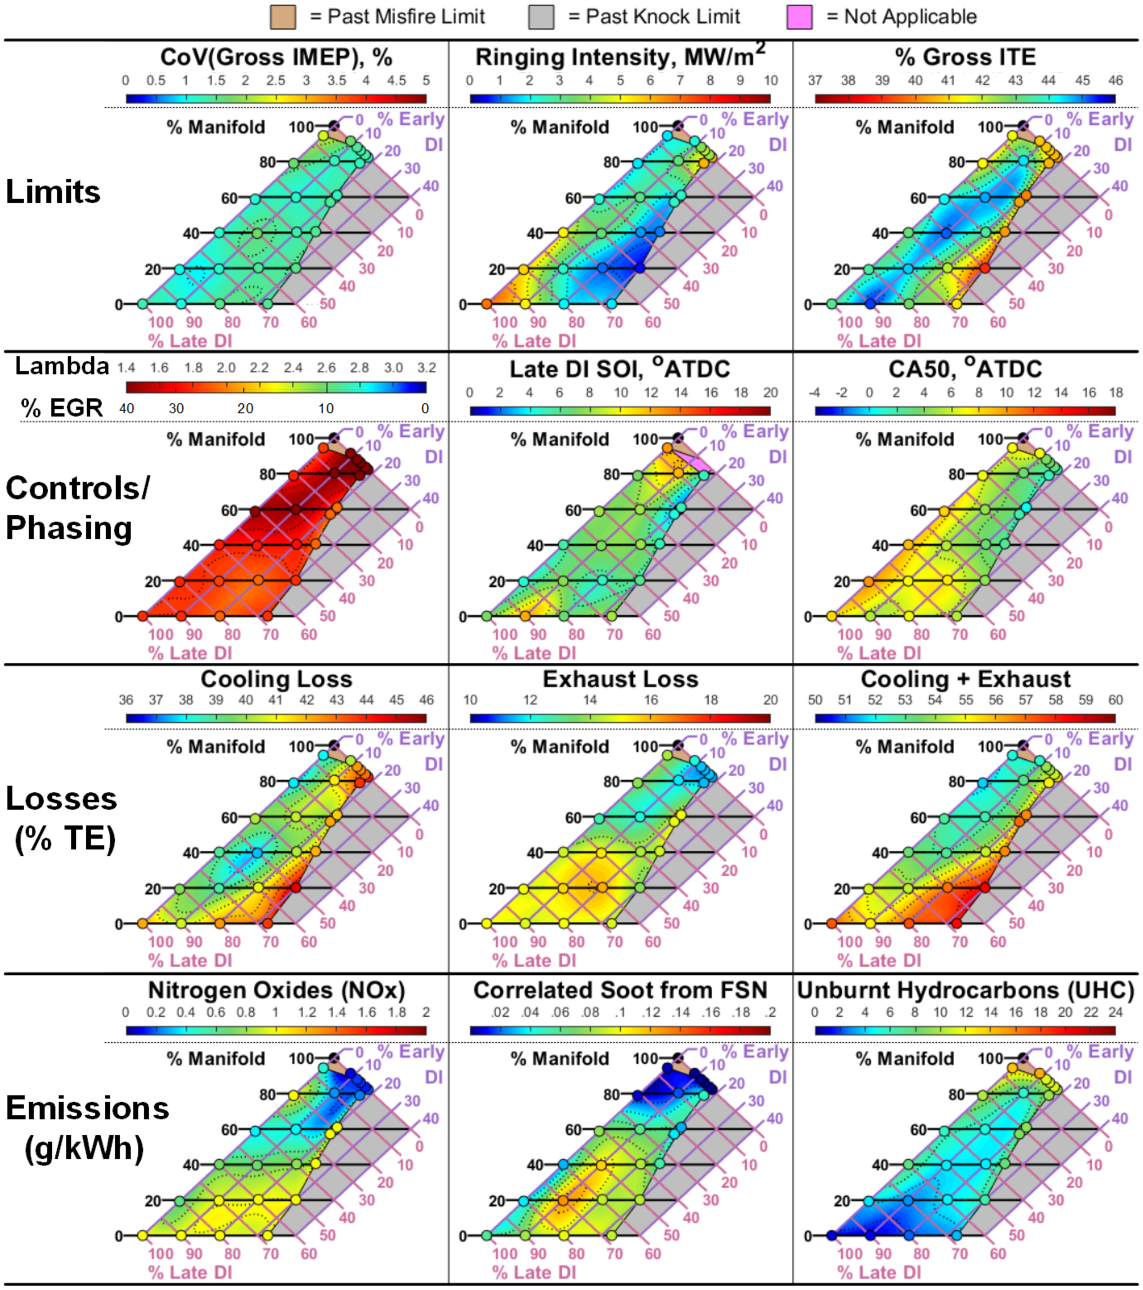

Figure 4.10 Contour maps of speed/load F (1200 RPM, 4 bar BMEP) at limit of 1 g/kWh NO_X. Trends are disccused on pages 106-108.

in cooling loss in these areas are made up for by the decreases in both NO_X and UHC, and although this also raises soot, it is occurring with low enough amounts of Late DI so that the soot produced in these areas is still less than in the lower PDFC region.

Again, this is still not a completely fair comparison since NO_X is not yet at a constant level. Before, in speed/load A, this was just because RCCI could not produce enough NO_X, but now in F there are points that *could* produce more NO_X, but get no ITE benefit from doing so. So again the NO_X limit will be dropped down further to 0.3 g/kWh.

4.3.3 0.3 g/kWh NO_X Limit

Much like in speed/load A, it is impossible to meet the lowered NO_X limit without producing unacceptable levels of soot in the majority of the MELT diagram, seen in Figure 4.11. However, unlike speed/load A, speed/load F has one non-RCCI point that can produce both low NO_X and low soot, while maintaining high ITE. This "sweet spot" of 80% Manifold fuel, 10% Early DI , and 10% Late DI appears to be the ideal operating mode for reducing engine-out emissions in speed/load F. This point seems similar to the close-by RCCI points in every respect except one: the UHC emissions and the resulting ITE. This goes to show how just a small Late DI can drastically improve the efficiency of RCCI while adding only a small amount of NO_X and soot emissions, well worth it for the sake of slashing UHC by more than half.

It is worth noting that, just as in speed/load A, there is another "sweet spot" in the cooling + exhaust loss in the CDF mode. Just as was previously shown in Figure 4.7, going along the left side of the diagram both cooling and exhaust loss are going down approaching this point. Again, the trend in exhaust loss makes sense given the trend in CA50, but the trend in cooling loss does not, which backs up the suspicion that there are rapid reductions in cylinder heat transfer losses when approaching this point since more of the heat release has been shifted into into the homogeneous lean premixed autoignition.

4.4 Speed/Load "C": 2300 RPM, 10 bar BMEP

Speed/load C has the same engine speed as speed/load A but a much higher load, which unsurprisingly results in a much smaller area of the MELT diagram being available. As a result, the following maps will only go down to a minimum of 40% Late DI and zoom in

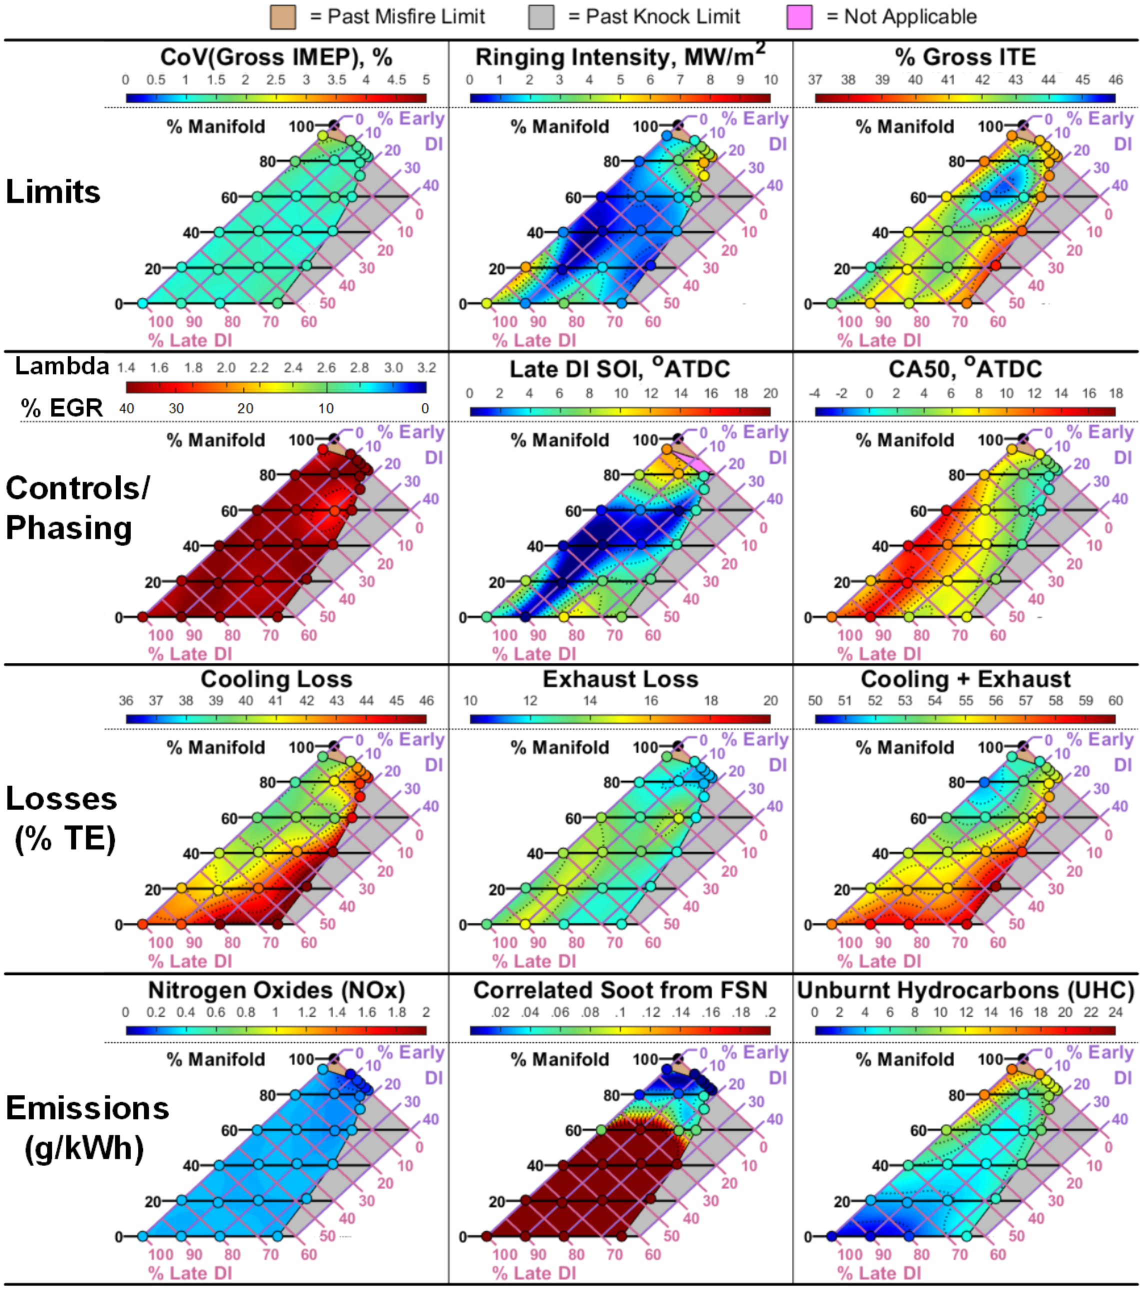

Figure 4.11 Contour maps of speed/load F (1200 RPM, 4 bar BMEP) at limit of 0.3 g/kWh NO_X. Trends are discussed on page 108.

on this area of the MELT diagram, as shown in Figure 4.2. Even though a small area is available in terms of *percentages* of Manifold and Early DI fuel, when it comes to the total *amount* of Manifold and Early DI fuel, the area available on the MELT diagram is about the same as speed/load A. This makes sense when considering both speed/load points can use about the same amount of premixed fuel before knocking.

Interestingly, an experiment on the same model of engine used in this study with very similar conditions was able to produce full RCCI (with 0% Late DI) at a very close speed/load combination [55]. However, there was one key difference in the engine where RCCI was achieved: only one of the cylinders was fired, with the other three deactivated. Likewise, another study on the same engine model with four cylinders firing found that RCCI was limited to loads of 6 bar BMEP or lower [56]. Hypothetically the single-cylinder engine had its one fired cylinder being much more rapidly cooled by the lack of heat in the other cylinders surrounding it, which slowed down the lean premixed autoignition to the point where it did not cause knock, while the same conditions would produce very harsh knock in a conventional multi-cylinder engine such as the engine used in this study.

4.4.1 Constant 10% EGR

For the Type 1 contour maps of speed/load C, seen in Figure 4.12, a constant EGR level of 10% was used, which gave a good balance between NO_X and soot emissions. At 20% EGR, all MELT points produced soot above 0.13 g/kWh, whereas at 0% EGR, all MELT points produced NO_X above 2.4 g/kWh, so 10% was chosen as a good compromise, giving a lambda of about 1.6.

The CoV of gross IMEP produced the same trend seen in speed/load A, but to a lesser degree, with an uptick in CoV at 10% Manifold fuel that goes back down at higher percentages of Manifold fuel. This backs up the suspicion that the PFI injectors produce relatively high cycle-to-cycle variations in the amount of injected of fuel at high engine speeds and small injection sizes. The reason this trend is more subdued in speed/load C can be attributed to the fact that the overall amount of injected fuel is higher, so the same cycle-to-cycle variation that produced 4% CoV in speed/load A will only produce 1.6% CoV in speed/load C, which matches the amount of variation that is observed. It remains to be seen why this variation in only being produced at 2300 RPM, and not at 1200 RPM.

Going up the MELT diagram, the RI limit is reached at about 50% Late DI. Unfortunately,

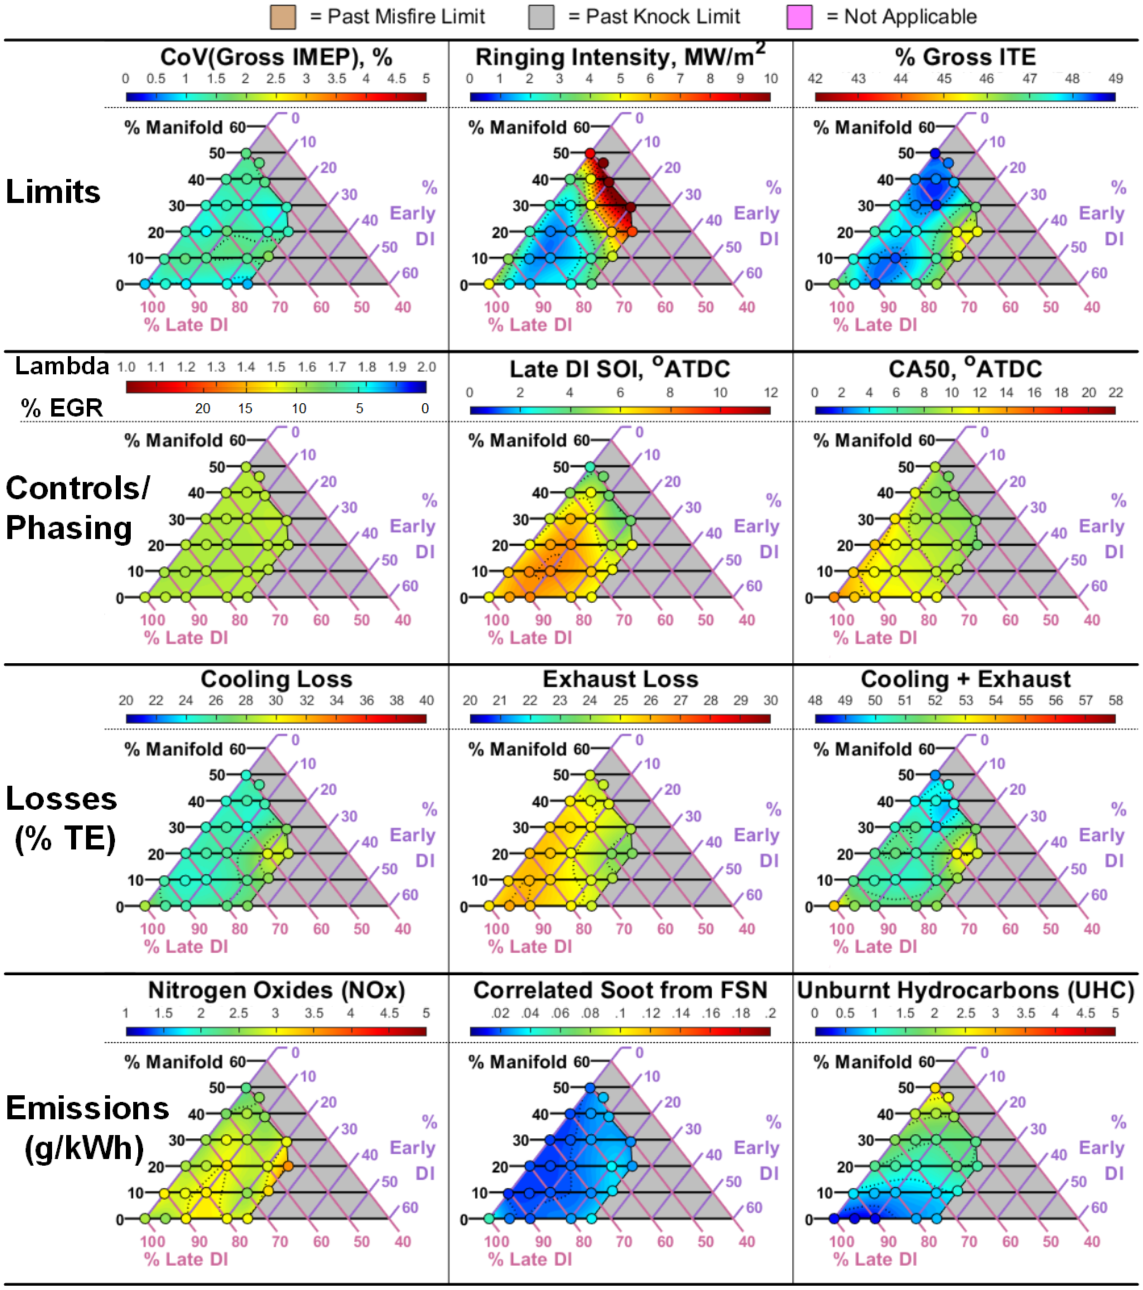

Figure 4.12 Contour maps of speed/load C (2300 RPM, 10 bar BMEP) at constant 10% EGR. Trends are discussed on pages 110-114.

at this load, must necessarily be large amounts of mixing-controlled combustion. The compression ratio of the engine and the reactivity of the fuel are simply too high to premix the full amount of fuel without it autoigniting to soon and producing excessive knock, even if only the LRF gasoline is used. Going along the bottom of the diagram, ITE rather than RI was the limiting factor. Even though ITE had not yet reached the low limit of 90% of the

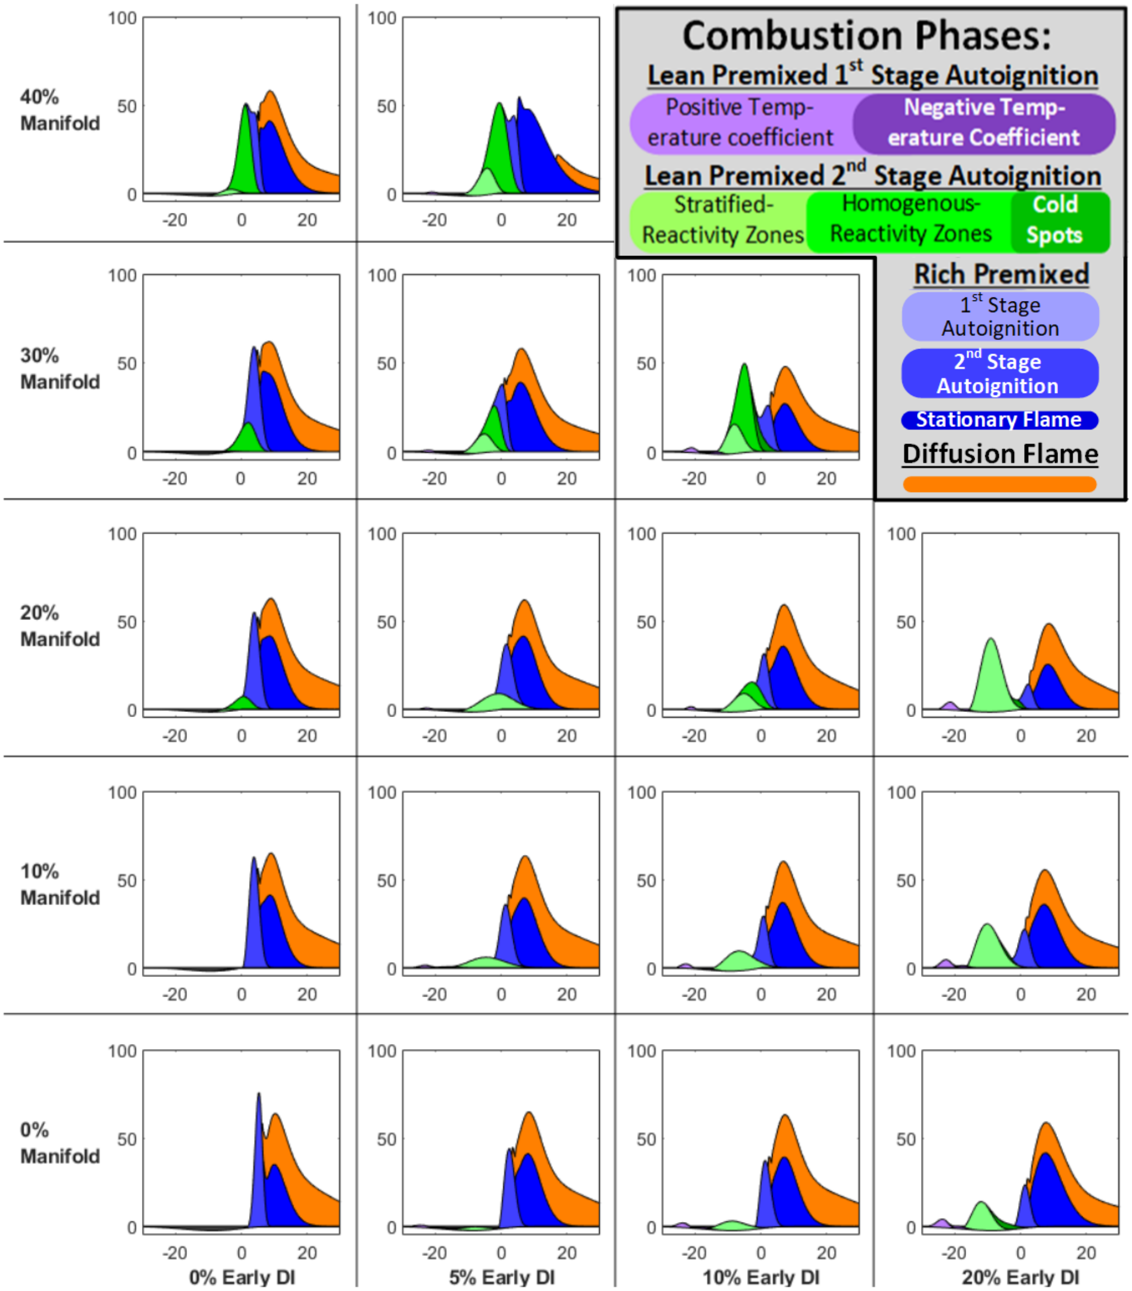

Figure 4.13 Wiebe function breakdowns of test conditions from Figure 4.12. Trends are discussed on pages 110-114.

baseline CDC ITE, it was decided to stop at 25% Early DI to focus the exploration on the higher-ITE portions of the MELT diagram. By this point the lean premixed autoignition was drifting too early and any efficiency gain from the lowered cooling loss would be offset by the extra work needed from the piston at the end of the compression stroke.

Once again, there appears a “ridge” of high efficiency running from CDC through PDFC, but unlike the lower-load conditions, this ridge turns toward the left side of the diagram, reaching peak ITE at 50% Manifold fuel and 50% Late DI, with no Early DI. Here, there is no longer CDF, but instead PHCCI. Despite the SOI being retarded, the CA50 actually advances along this line, due to a growing peak of homogeneous lean premixed autoignition that occurs before the Late DI. Along this line there is a decreasing trend of exhaust loss, which makes sense given the retarding CA50, but there is *not* a decreasing trend of cooling loss, which stays relatively constant. This leads to a decreasing trend in the cooling + exhaust loss, with the PHCCI condition having the lowest. This is somewhat balanced out by a trend of increasing UHC emissions and unburnt loss, but this UHC trend is much less severe than in the lower-load conditions (never going above about 1% TE in unburnt loss), and the overall ITE trend is positive. This contrasts with the lower-load conditions, where the unburnt loss trend dominates and causes the highest PDFC conditions to lose ITE. The reason for this change is most likely the higher temperatures and pressures experienced in this high-load 10 bar BMEP condition.

It is hard to establish much of a NO_X trend in this speed/load condition, although it does appear that NO_X decreases along the line of highest ITE. This trend is not as significant (dropping from about 3.0 to about 2.2 g/kWh) as it was in the lower speed/load conditions, which could get all the way to RCCI and drop NO_X down to less than 0.2 g/kWh. Again, the blame for this lies in the overall temperature and pressure increases at the high load, with another contributing factor being the lower overall lambda. The substitution of lean premixed combustion for diffusion combustion does not have as much of a NO_X reduction effect when the lean premixed combustion is taking place near TDC with high temperature, high pressure, and near-stoichiometric equivalence ratios, which leads the lean premixed combustion to produce NO_X itself. The NO_X reduction effect should be the greatest when there is no Early DI in the mix, thus no stratification in the equivalence ratio, keeping all of the local equivalence ratios of lean premixed combustion as far from stoichiometric as possible.

It is important to note, however, that NO_X trends appear to be closely tied to CA50 at this particular speed/load in the Type 1 maps. This is because there is much more variation in the knock limit and thus much more variation in the SOI timing (which is placed halfway between the knock and misfire limits in this map type). The NO_X trend thus roughly matches the trend of SOI, with earlier SOI producing more NO_X. A study [57] of a similar condition with varying proportions of PFI gasoline and Late DI diesel, but a *constant* SOI timing found a much flatter trend in NO_X emissions compared to the NO_X trend along the left side of the diagram, although there was a slight increase in NO_X at 60% Manifold fuel. That constant Late DI, however, results in an progressive drop-off in ITE as the manifold fuel is increased [57], whereas in this study with variable Late DI, the ITE stays relatively constant until approaching the misfire limit.

The SOI trend also effects soot in the opposite direction, with early SOI producing lower soot. The earlier SOI pushes the diffusion phase of combustion to an earlier part of the cycle, with higher temperatures and pressures that can better oxidize soot. This causes the trend of soot along the bottom edge of the diagram to be the opposite of what was observed in the low load conditions, with soot going *down* when the Early DI is first introduced, then going back up. It is possible this is due to the change in SOI rather than the change in combustion mode, and that if the SOI was held constant this trend would reverse.

Except for the CDC point with 100% Late DI, the overall trend in soot favors lower amounts of Early DI, and there is never a soot “ridge” as was observed in speed/loads A and F, although it is difficult to tell whether this is a fundamental change at the higher load or merely the result of the changing SOI. Overall, the soot emissions never go beyond 0.05 g/kWh, which is very good compared to the low load conditions. This is probably the result of the higher pressures and temperatures of the high-load condition, which encourage soot oxidation during the diffusion phase both physically (through increased turbulence and thus increased flame speed) and chemically (higher reaction rates). However, these low soot levels were only achieved with NO_X levels above 2 g/kWh throughout the diagram. If a 1.0

g/kWh NO_X limit is applied, these soot levels will likely increase dramatically.

4.4.2 1.0 g/kWh NO_X Limit

Indeed, applying the 1.0 g/kWh NO_X limit, seen in Figure 4.14, generates soot levels above 0.4 g/kWh throughout the entire accessible portion of the MELT diagram, well off the 0.2 g/kWh high end of the scale. This is due to the constant low lambda (less than 1.3) required to achieve this NO_X level, produced by the maximum 20% EGR. This produces very rich local equivalence ratios during the rich premixed autoignition, generating high levels of soot, and also leaves very low concentrations of oxygen at the end of the diffusion flame process which greatly inhibit soot oxidation.

Although the scale could be changed to show the correlated soot mass in these conditions, these numbers have very large measurement error at the high smoke numbers. Given the lack of repeated testing to produce the maps in this chapter, this produced a large amount of fluctuation in the correlated soot mass across the MELT diagram in these conditions, and the maps produced by raising the scale did not produce any meaningful trends in soot mass. It is not practical to reduce engine-out NO_X to this level at speed/load C given the high levels of soot that cannot be avoided, as well as the relatively low ITE. The NO_X limit for this condition should be at least somewhat relaxed.

Before doing that however, there are some important trends to observe at the 1.0 g/kWh NO_X level. The knock limit is contracted to 60% Late DI, with the higher EGR rate producing higher intake temperatures along with closer-to-stoichiometric equivalence ratios in the premixed fuel, advancing the autoignition. However, the knock limit overlaps with the highest observed ITE, the PHCCI condition of 40% Manifold fuel and 60% Late DI. The peak ITE was also observed in the Type 1 map, but it did not have the same advantage over PDFC and CDC that are observed here. Again, just as in speed/loads A and F, as this point is approached along the left side of the diagram, there are declines in both cooling and exhaust loss while the CA50 is advanced. The lowered exhaust loss should be expected

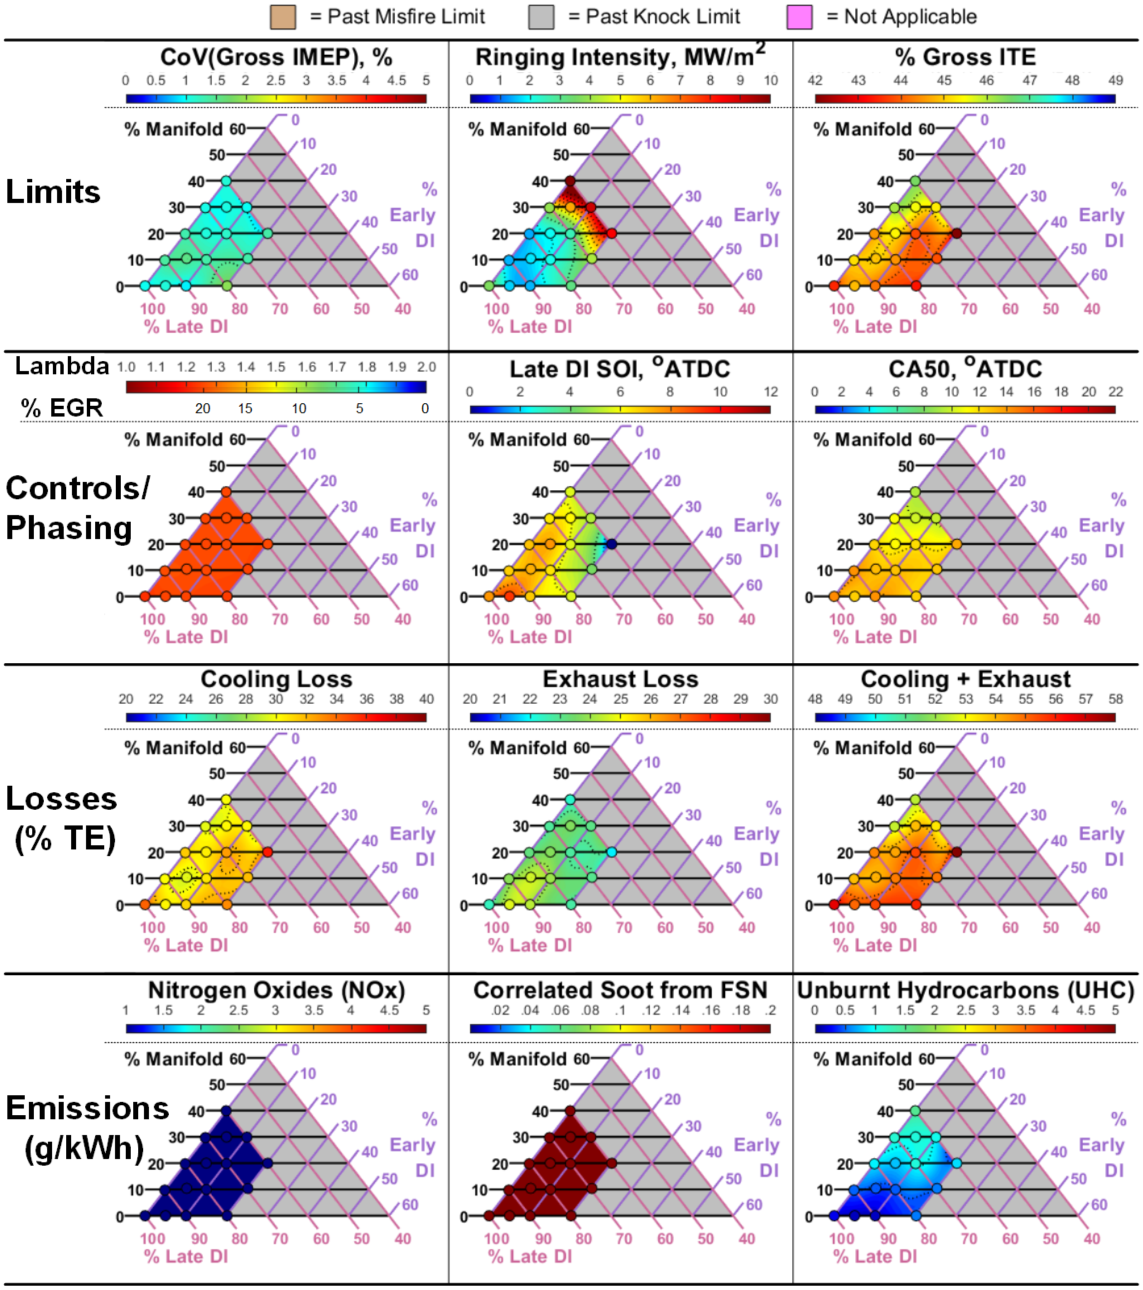

Figure 4.14 Contour maps of speed/load C (2300 RPM, 10 bar BMEP) at limit of 1 g/kWh NO_X. Trends are discussed on pages 115-117.

with the advancing CA50, but not the lowered cooling loss. This furthers the theory that homogeneous lean premixed autoignition minimizes cylinder heat transfer losses compared to mixing-controlled combustion.

It is also worthwhile to note the very low UHC emissions in this map, which never exceed 1.5 g/kWh (equal to about 0.6% TE). This is likely due to the same increases in temperature and pressure from high EGR that were the source of the shift in the knock limit. However, despite the low UHC, no condition in this map ever exceeds 47% gross ITE, which can be easily exceeded throughout most of the available area of the diagram at even a slightly higher EGR level. The main reason for this would appear to be the high cooling loss experienced with the high in-cylinder temperatures and high flow through the EGR cooler. For these reasons along with the excess soot emissions, a higher NO_X limit should be considered that allows the use of more EGR.

4.4.3 1.5 g/kWh NO_X Limit

As seen in Figure 4.15, raising the NO_X limit to 1.5 g/kWh allows for soot emissions below 0.1 g/kWh throughout a good portion of the accessible area of the MELT diagram, except for the right side and the 100% Late DI. This is due to the use of less EGR and thus higher lambda, with the low-soot areas all having lambda above 1.5 and the high-soot areas all having lambda below 1.5.

One noticeable feature of these trends is along the bottom edge of the MELT diagram, where the point with just 5% Early DI is able to achieve much lower soot than the points on either side with 0% and 10% Early DI. This seems to indicate that 5% Early DI may be just enough to reduce ignition delay and allow for a late enough SOI that EGR can be reduced while still maintaining the NO_X level, but without increasing temperatures and local equivalence ratios to the point that excess soot is generated. Not having this slight early heat release forces the SOI to be advanced to maintain efficiency and then the lambda lowered to maintain NO_X. Having more early heat release consumes too much oxygen before the rich premixed autoignition, promoting soot production. But having just the right amount of early heat release appears to allow the soot to be limited. This trend should not appear in any maps except this 1.5 g/kWh threshold, which seems to be right on the very “edge of the cliff” in the NO_X/soot trade-off in CDC.

More importantly, though, the gross ITE appears to favor the PHCCI and “heavy PDFC”

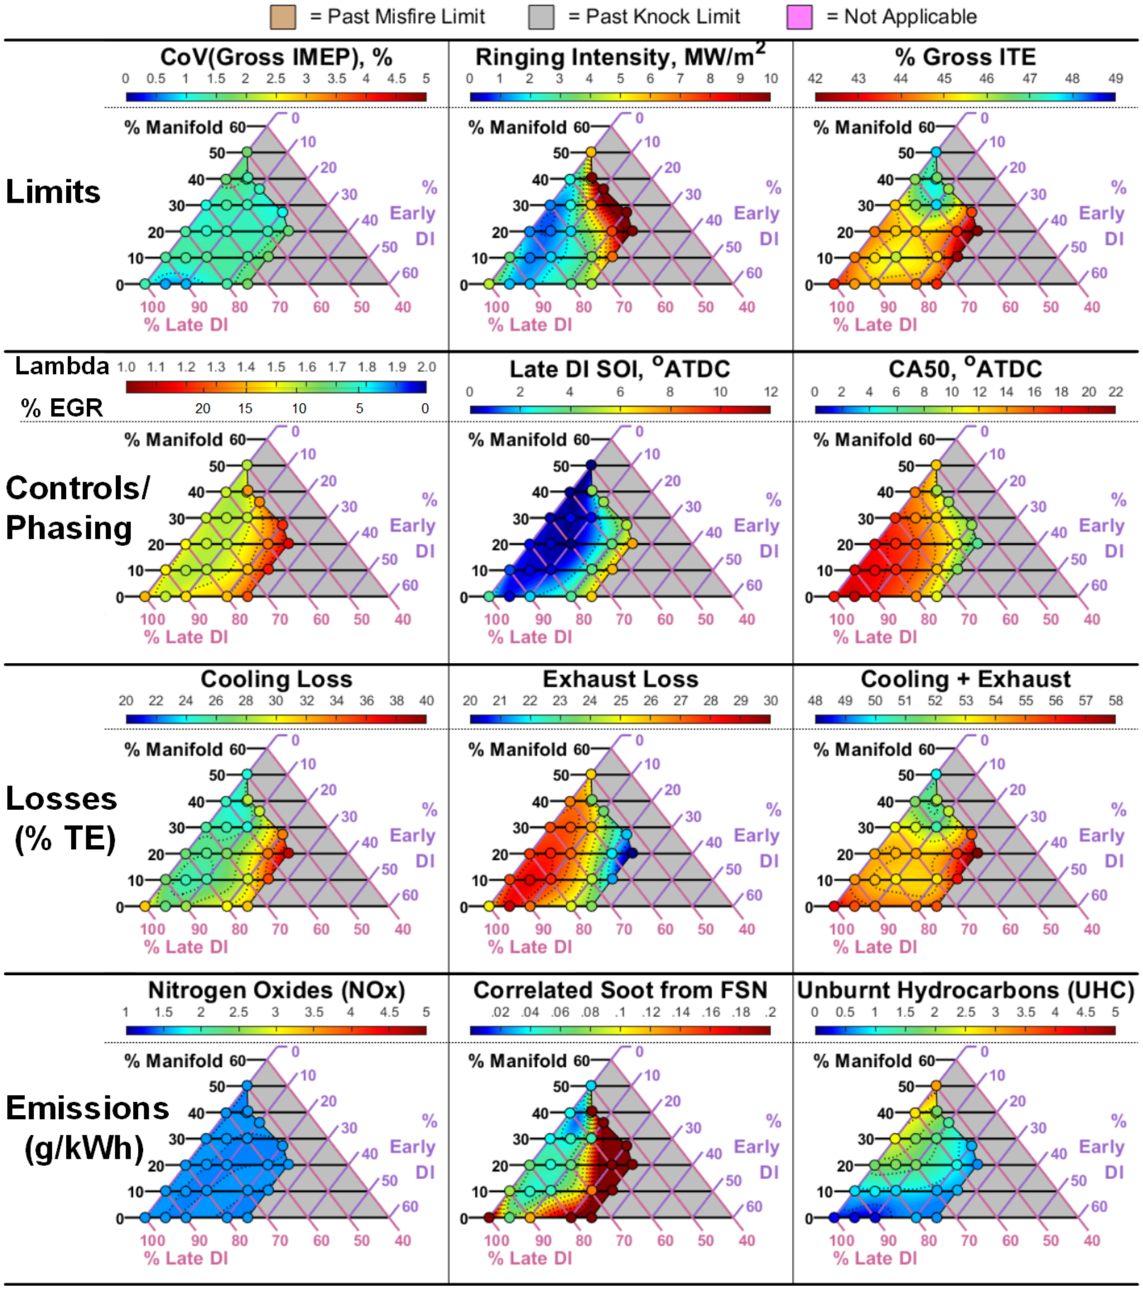

Figure 4.15 Contour maps of speed/load C (2300 RPM, 10 bar BMEP) at limit of 1.5 g/kWh NO_X. Trends are discussed on pages 117-119.

conditions, increasing ITE by about 3% TE compared to the 5% Early DI CDC condition, while soot stays relatively constant. These ITE increases occur along the same "ridge" that was observed in the Type 1 map, but now the same ridge appears to slope upward toward PHCCI. This occurs even with the UHC emissions rising along this path, since the increase in unburnt loss is outpaced by the decrease in the combined cooling + exhaust loss. Again, the trend in cooling loss appears to dominate the trend in combined cooling + exhaust loss. In the lower loads, this trend was approximately balanced out by the trend in unburnt loss, leading to a "ridge" of *constant* high efficiency running from CDC to heavy PDFC. But in this higher load, with overall higher temperatures and pressures to minimize the UHC, the ITE trend begins to favor the conditions with the more premixed fuel.

Another study that has identified this advantage of PDFC in a similar condition comes from Benajes et al. [58], who refer to this combustion mode as "Highly premixed RCCI". As part of a multi-mode engine operating strategy utilizing RCCI at lower loads, Benajes highlights how the combustion efficiency drop compared to CDC is much lower for PDFC than for RCCI, resulting in bigger efficiency gains for PDFC than RCCI, despite not having the same benefit of smoke reduction. The same study also utilizes a PHCCI mode (referred to as "Dual-fuel diffusion") at even higher loads, for even bigger efficiency gains.

4.5 Identifying Key MELT Ratios

These past three sections have mapped the performance trends of this engine across the entire MELT diagram. The following section will focus in on a few specific, optimal points ("MELT ratios") on these maps to get a better sense of what makes these points so optimal. In Chapter 5, there will be more extensive testing on these optimal points to more precisely quantify their performance, and determine how they could fit in to an overall engine operating strategy.

These key MELT ratios will be identified using the "Type 2" maps that have been optimized for reducing NO_X and soot to levels of 1 g/kWh and 0.1 g/kWh respectively, except for speed/load C in which the NO_X limit was raised to 1.5 g/kWh in order to achieve the soot limit. The Type 2 NO_X-limited maps of ITE and soot for these conditions are reprinted in Figure 4.16, with the key points that have been selected. There are 4 key points on each map, circled in a specific color:

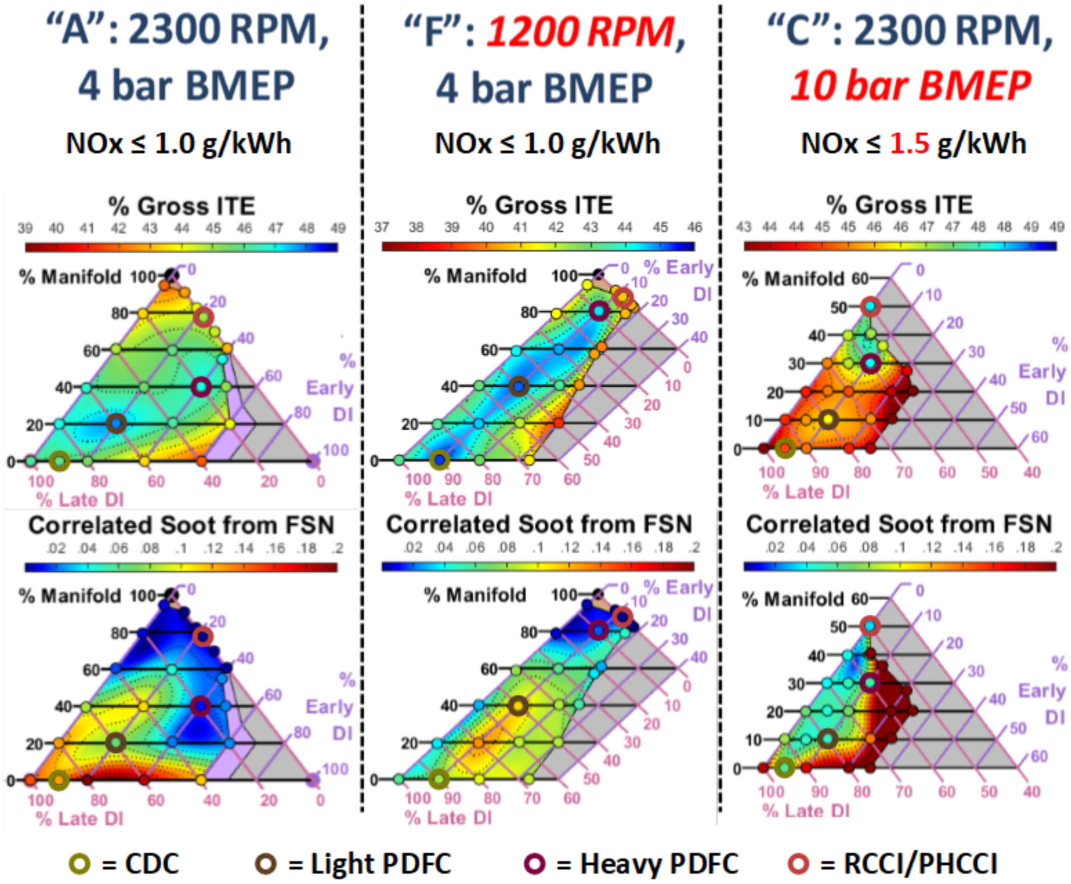

Figure 4.16 Tentative locations of key MELT ratios in speed/loads A, F, and C, discussed on pages 119-120.

- **CDC - Dark Yellow:** Conventional single-fuel operation with a slight amount of Early DI (not enough to consider it PPCI)
- **Light PDFC - Brown:** A condition in which the majority of the fuel still burns through the conventional mixing-controlled process, but with a substantial amount of lean premixed autoignition
- **Heavy PDFC - Maroon:** A condition in which the lean premixed autoignition begins to dominate the shape of the HRR, with significantly larger Wiebe functions
- **RCCI *OR* PHCCI - Light Red:** An RCCI condition in the low loads (where it is possible) or a PHCCI condition in the high load (where *it* is possible)

The HRR for these chosen conditions with the Wiebe function fits are shown in Figure

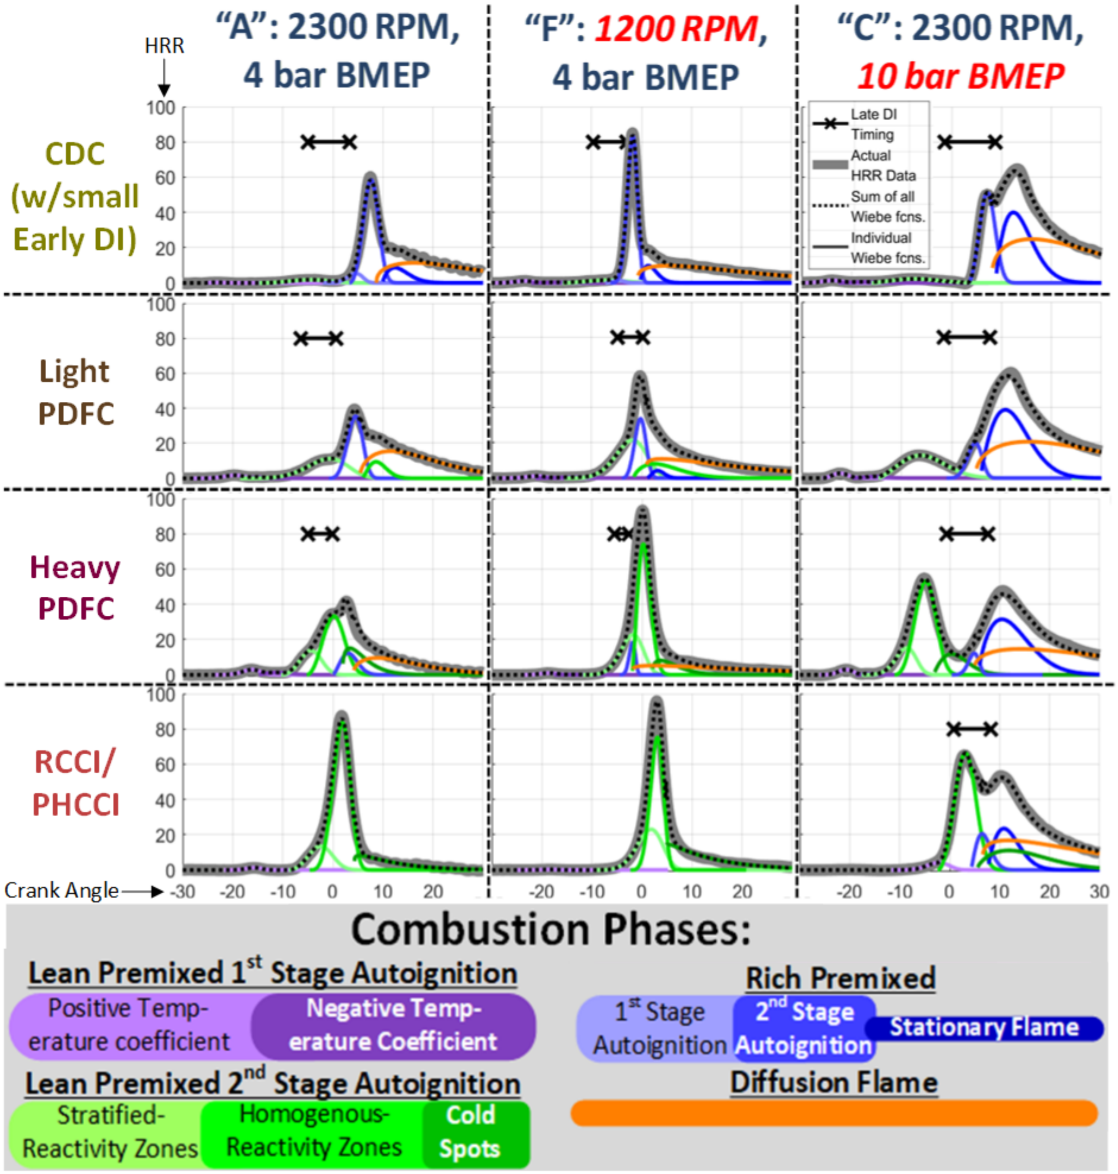

Figure 4.17 Wiebe function breakdowns of HRR from key points in Figure 4.16, discussed on pages 121-123.

4.17. This figure illustrates how the combustion mode changes to maintain the low NO_X and soot levels while substituting gasoline for diesel fuel.

The CDC conditions in the top row all feature a tall, thin spike produced by the 2nd (main) stage of rich premixed autoignition. In speed/loads A and C, the peak of this spike occurs around 7-8°aTDC, which ensures that the diffusion combustion is retarded to limit NO_X production. The same does not need to occur in speed/load F, where the spike occurs before TDC. In this condition there are lower overall temperatures that inhibit NO_X production, allowing the SOI to be advanced for the sake of higher efficiency and lower soot. It should be noted that for speed/load A, this condition with 10% Early DI was not included initially, and this data has been borrowed from the tests in the next chapter, in which it is found that the soot produced with 10% Early DI is better than both the cases with 0% and 20% Early DI at this constant NO_X level.

The Light PDFC conditions in the next row all feature a sequence of stratified lean premixed combustion, represented by a wide and short Wiebe function. In the low load conditions, this function straddles TDC, with the Late DI occurring simultaneously. Thanks to the displacement of Late DI fuel, in speed/load A this allows the SOI to be advanced and the EGR lowered without raising NO_X. Although the lean stratified combustion would be expected to increase soot over CDC, the lowered EGR cancels out this effect by increasing the available oxygen during the diffusion phase, and as a result Light PDFC has lower soot than CDC in speed/load A. The same cannot be said for speed/load F, where the SOI was already advanced as far as possible, and the EGR cannot be lowered without raising NO_X beyond the limit, resulting in higher soot in Light PDFC. Meanwhile, in speed/load C, the stratified autoignition phase is already on the downslope before the Late DI starts, resulting in a much more “split” HRR than the low-load conditions. Although this “split” HRR is not ideal for maximizing the area under the P-V curve, the Late DI must remain retarded in order to maintain low NO_X. One advantage to the “split” however is that it will keep HRR low right at TDC, which helps avoid high MPRR and thus limit ringing.

In the Heavy PDFC conditions, the lean premixed heat release now features a prominent spike, indicating a homogeneous autoignition of the bulk premixed charge. The timing of the Late DI in relation to these spikes is different in each speed/load. In speed/load F, the Late DI occurs before the spike, assisting with achieving a complete autoignition right at TDC. In the other speed/loads, this could lead to excessive ringing, but at the low speed of F the temperatures and pressures are low enough to slow down combustion and keep RI within the acceptable limit. In speed/load A, the Late DI is instead occurring *during* the homogeneous premixed spike, producing a second spike immediately following the first. This slight delay helps reduce ringing and prevents excessive NO_X production from the diffusion flame. In speed/load C, the Late DI must be retarded until after the homogeneous autoignition spike, again for the sake of achieving the NO_X limit and preventing excessive ringing.

The bottom row features RCCI for speed/loads A and F, and PHCCI for speed/load C. Both RCCI conditions feature stratified premixed autoignition beginning around -10°aTDC leading to a homogeneous autoignition spike that peaks just after TDC. This is ideal timing in terms of combustion phasing, and the lack of a Late DI is ideal for minimizing heat losses, but the absence of a diffusion flame leads to excessive unburnt emissions that limit the ITE, especially for the lower-temperature, lower-pressure speed/load F. PHCCI, on the other hand, achieves the highest thermal efficiency of all these conditions, with almost no heat release before TDC, and although the heat release is not concentrated into as tight of a window as the RCCI conditions, there are no sharp spikes to cause excessive ringing due to the Late DI spike being perfectly staggered with the homogeneous autoignition spike. There are also far less UHC than RCCI thanks to the higher overall temperatures and the presence of a diffusion flame to “clean up” any UHC left behind from the premixed autoignition.

4.6 Tuning the HRR Calculation

As previously mentioned in Section 3.1.2, the data collected in this chapter can be used to tune the coefficients of the heat release model. Ideally, the correct coefficients should cause the total integrated heat release to equal the total heat of combustion contained in the fuel, minus any unburnt emissions. Unfortunately, the single-zone models used in this work do not appear to be able to accurately model the cylinder heat transfer, and although the Chang model was an improvement over Woschni, it still gave an average error of 3.6% over all the test conditions used in this chapter.

Having fit the Wiebe functions to the data, it was worthwhile to see if any improvement in the heat transfer model could be achieved by using the Wiebe function fits. The idea was to allow the C_2 coefficient, which accounts for the turbulent gas velocity increase brought on by combustion, to shift depending on the percentage of heat release that was lean premixed autoignition (subscript “lpa”) or mixing-controlled combustion (subscript “mcc”) according to the Wiebe function fits. This of course required running the HRR calculation twice - once with the original Chang model to obtain an HRR for Wiebe function fitting, and again after the Wiebe function fitting with the new, modified model. In the “Modified Chang” model, coefficient C_2 changes at each crank angle like so:

$$C_2(\theta) = C_{2,\mathrm{lpa}} \frac{\int_{\theta_{\mathrm{IVC}}}^{\theta} \frac{dQ_{\mathrm{lpa}}}{d\theta} d\theta}{\int_{\theta_{\mathrm{IVC}}}^{\theta} \frac{dQ_{\mathrm{total}}}{d\theta} d\theta} + C_{2,\mathrm{mcc}} * \left[1 - \frac{\int_{\theta_{\mathrm{IVC}}}^{\theta} \frac{dQ_{\mathrm{lpa}}}{d\theta} d\theta}{\int_{\theta_{\mathrm{IVC}}}^{\theta} \frac{dQ_{\mathrm{total}}}{d\theta} d\theta} \right] \tag{4.23}$$

In theory, this should allow for the model to account for differences in the relative amount of turbulence induced by lean premixed autoignition and mixing-controlled combustion. Applying this model reduced the average error to 3.2%, but more importantly it shifted the distribution of these errors across the MELT diagram as seen in Figure 4.18. This figure shows the total integrated heat release as a percentage of the actual heat release for three different models: the apparent heat release (with no heat transfer modeling), the original Chang model, and the modified Chang model. These numbers are plotted onto the same Type 1 MELT diagrams that were seen previously for each speed/load condition.

In Figure 4.18, a value below 100% means that the model is under-predicting the cylinder heat transfer. This is always the case for the apparent heat release, which has no modeling of the cylinder heat transfer. What is shown by these maps is that the apparent heat release compared to the actual heat release generally goes down with the higher amount of premixed fuel, indicating that there is *more* cylinder heat transfer as more premixed fuel is added. However it is important to remember that CA50 is also advancing as the premixed fuel is

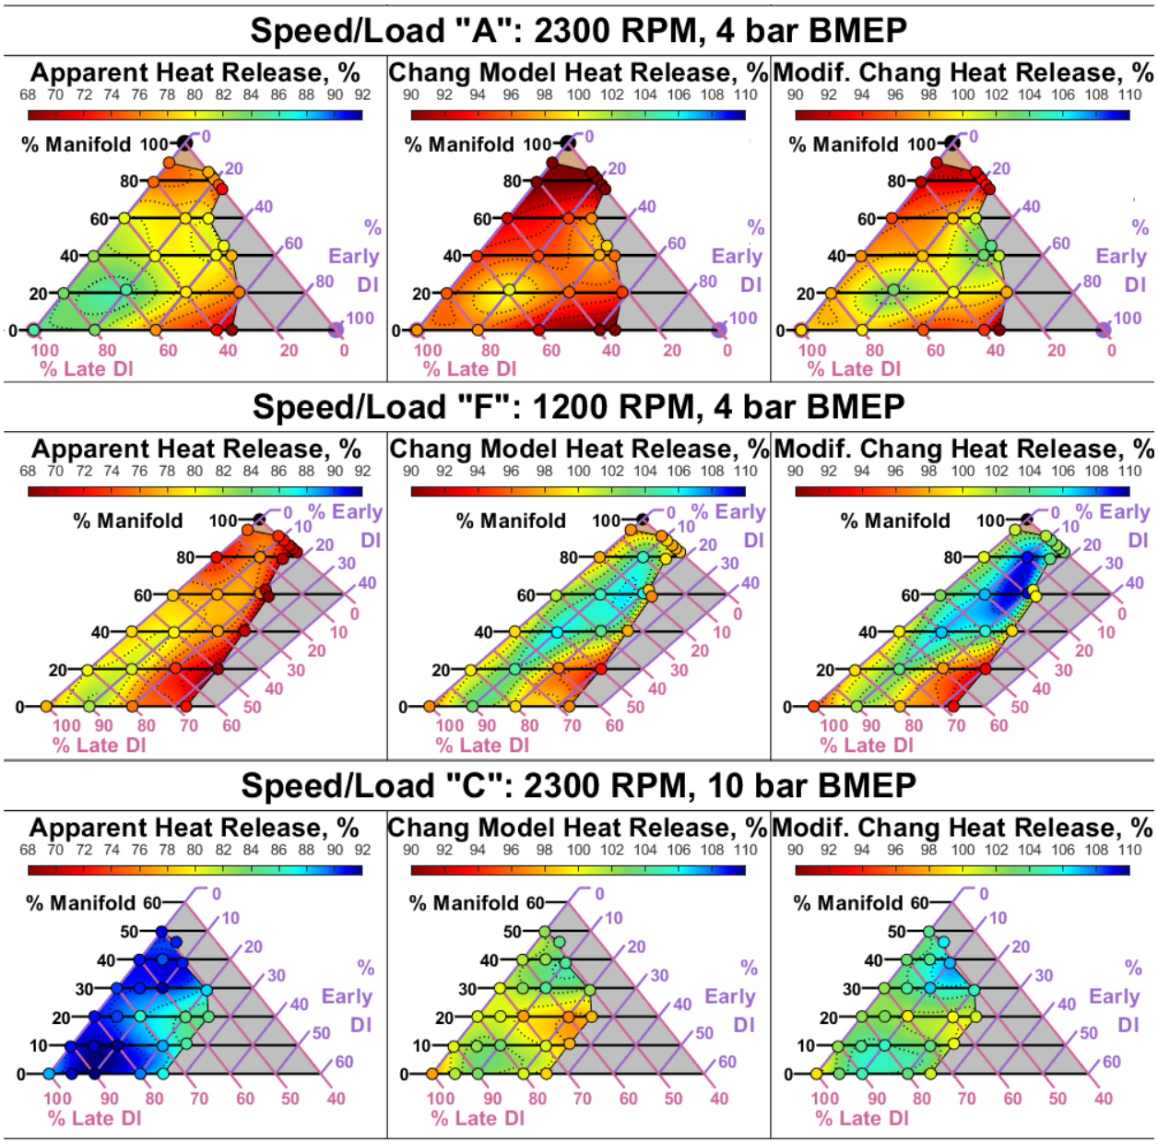

Figure 4.18 Contour maps of total integrated heat release from the calculated HRR, as a percentage of the actual heat release as measured by the fuel flow and unburnt emissions measurements. Trends are discussed on pages 124-126.

added, and this advancement of the CA50 could be the cause of the increased heat transfer, since it will generate higher peak temperatures and longer times between the initial heat release and the exhaust valve opening for heat transfer to occur.

When the Chang model is applied to account for these differences, the shape of the integrated heat release error begins to look very similar to the shape of ITE. This makes sense considering that when the cylinder heat transfer is over-predicted (giving a total integrated heat release above 100%), the actual heat transfer occurring is lower than the model expects, which would give it a higher-than-expected efficiency. However, what this means is that the model is not going to be able to accurately predict gross ITE if it is used to do a "backwards" calculation of the cylinder pressure trace from HRR. This was one of the goals of the Wiebe function modeling: to produce functions relating the test conditions to the Wiebe function coefficients, and then use the predicted HRR to back-calculate predicted cylinder pressure and thus ITE at untested conditions. But unfortunately it cannot be said that those predictions would be useful unless they use an HRR model that at least produces *random* errors across the MELT diagram, not a model that is biased toward certain areas of the MELT diagram, and certainly not one that produces errors in the same pattern as ITE.

It was hoped that the modified Chang model would improve this, but unfortunately it only "titled" the error toward the conditions with more premixed fuel, and did not eliminate the correlation with ITE. Clearly, a more complex model than the single-zone models used herein will be needed to produce a HRR calculation accurate enough to eliminate this bias, and allow for the use of a Wiebe function model for ITE prediction across shifting MELT ratios. Still, the Wiebe function HRR breakdowns produced in this thesis are useful for illustrating how different combustion sub-processes contribute to the overall HRR shape, which does not change substantially when changing the HRR model. In the future, this Wiebe function fitting method could eventually contribute to a more complex HRR model that can be used to make predictions of ITE in untested conditions. But for now, the last chapter of this thesis will focus purely on experimental determination of the performance changes in each combustion mode, with the hope that the extensive data collected will eventually be able to help validate such a model.

Chapter 5

Optimizing and Quantifying Benefits of Multi-Mode Strategy

5.1 Selecting Test Conditions

The tests in this chapter will cover all six of the speed/load conditions first introduced in Section 4.1, shown in Figure 4.1 and Table 4.1. For convenience this figure and table are repeated here in Figure 5.1 and Table 5.1. Using the maps created for speed/loads A, F, and C, the tests will first follow the path of highest ITE through the MELT diagram in Section 5.2, and then in Section 5.3 the tests will focus in on the key points identified in Section 4.5. For the path of high ITE in Section 5.2, the tests will use an EGR level and SOI timing that can be held constant along the entire path, to better show how performance changes due to the change in the fueling ratio alone. Then for each key point in Section 5.3, a small subsection of the EGR/SOI space will be identified that gives the best compromise between ITE, NO_X, soot, and UHC. By limiting the number of test conditions to only the most relevant conditions for improved engine performance, there was enough time to produce three repetitions of each test and determine the level of uncertainty in these performance changes, without using the excessive amounts of time and fuel that were needed to produce the maps in the previous chapter.

The determination of these test conditions was already begun previously in Section 4.5, which identified four key MELT ratios that gave the best performance for each of the three mapped speed/load conditions (A, C, and F). The three other conditions (B, D, and E) are conveniently located exactly halfway between two of the three already-tested

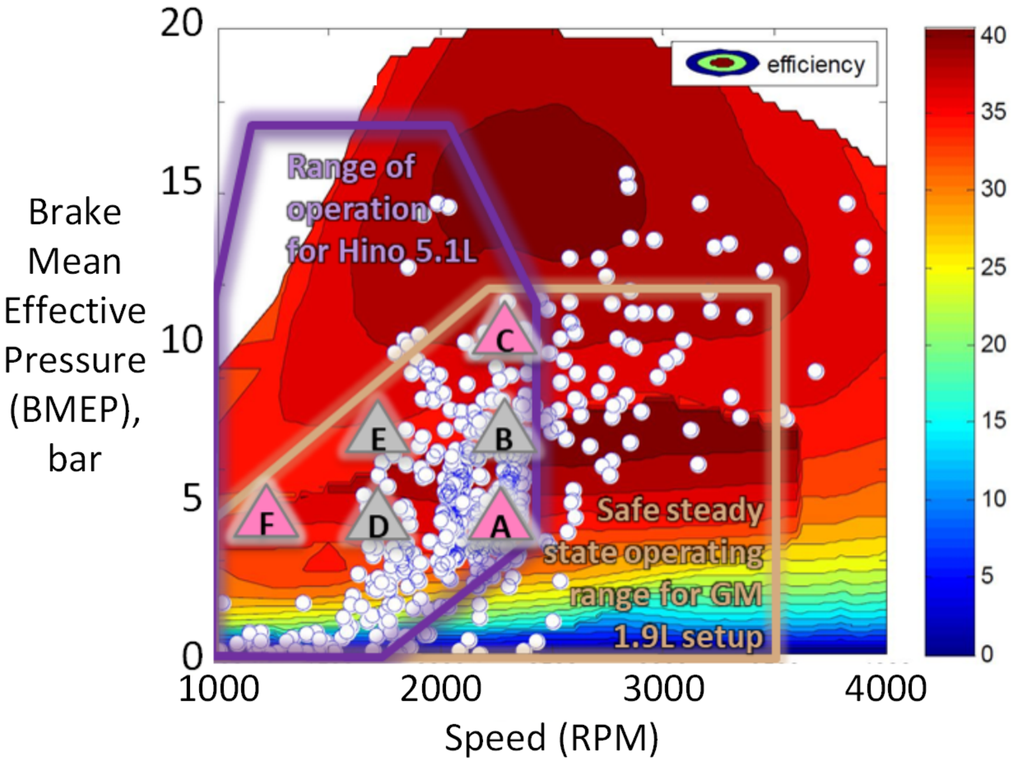

Figure 5.1 Speed/load points selected for this study, superimposed on a simulated EPA drive cycle by Gao et al. [10]. The three points that were fully mapped in Chapter 4 are highlighted in pink, and this chapter will perform a more focused set of tests on all six speed/load points.

Table 5.1 Speed/load conditions selected for this study. The three points that were fully mapped in Chapter 4 are highlighted in pink, and this chapter will perform a more focused set of tests on all six speed/load points.

Name	Speed (RPM)	BMEP (bar)	Boost Pres. (bar abs)	DI Rail Pres. (bar)	EGR Range (% Total Charge)
A	2300	4	1.13	690	0-40
B	2300	7	1.33	900	0-30
C	2300	10	1.58	1110	0-20
D	1750	4	1.09	530	0-40
E	1750	7	1.23	690	0-20
F	1200	4	1.05	370	0-40

conditions. This allows for a good initial estimate of the test parameters that will work well for the untested conditions, by simply averaging the test parameters of the two already-tested conditions that the untested condition lies halfway in between. These parameters will still need to be adjusted with some preliminary testing to confirm that they are producing satisfactory results.

5.1.1 Additional Injection Types

To get as good of an idea as possible about the range of performance, it will be prudent to consider two other common injection types that were not included in the MELT diagram. These injections do not fundamentally alter the distribution of fuel, but rather split one of the already-existing DI's into two separate injections, in order to change the combustion mode. These will be referred to as *minor* injection types, since they typically do not contain a large percentage of the total fuel.

The first minor injection type that will be considered is a "pilot injection." A pilot injection is a very common feature of most conventional diesel engines, consisting of a short DI pulse occurring before the Late DI. This injection burns in the same process as the Late DI, but because it is relatively small it does not produce a large spike in the heat release. The small spike that it does produce happens before most of the Late DI fuel has been injected, and warms up the combustion chamber. This shortens the ignition delay for the Late DI, giving it a smaller rich premixed spike, which helps in reducing the ringing intensity and allows diesel engines to run much more quietly. This also allows the SOI to be advanced further before the knock limit is reached, which can lead to higher ITE. The pilot does, however, lead to an increase in the soot emissions, for much of the same reason that the Early DI increases soot in PPCI and PDFC (as previously discussed in Sections 3.4.7 and 3.4.8). The pilot injection, however, is not an "Early DI" in the sense that the term is used in this thesis. Even though it occurs well before the Late DI, it is still occurring *after* the "Turn-Around Point" that was described in Section 2.2. The pilot is not injected early enough to achieve a fully lean mixture before it ignites, so it will still burn through the same two-phase rich premixed autoignition and diffusion flame processes as the Late DI. Thus, the pilot will be classified as a “Late Pilot DI,” and the major injection will be renamed the “Late *Main* DI.”

The second minor injection type that will be considered is an “Early Bowl DI,” which is still occurring early enough in the cycle to achieve a lean mixture before it autoignites (i.e. before the Turn-Around Point), but *after* the piston has reached the point where the rim of the piston bowl intersects with the path of the fuel jet, and the injection is directed into the bowl. This contrasts with the original Early DI, which is directed into the squish zone, and which will now be renamed the “Early *Main* DI.” The point where the fuel jet and piston bowl rim intersect, as depicted in Figure 5.2, will be called the “Cross-Over Point,” and unlike the Turn-Around Point the crank angle of this point will depend on engine geometry alone. Thus the gap between the Cross-Over Point and the Turn-Around Point will be different from engine to engine and fuel to fuel, and in some engines it may not be possible to use an Early Bowl DI. The Early Bowl DI can be advantageous for RCCI engines because it allows the stratified autoignition process to occur starting closer to the center of the combustion chamber, rather than out in the squish zone on the periphery. The relatively round center of the combustion chamber has a much higher surface area-to-volume ratio than the thin, wide squish zone, especially near TDC, so it makes more sense for the lean stratified autoignition to occur in the bowl in order to minimize heat transfer losses.

Figure 5.2 shows all five possible injections under consideration, and the crank angle windows over which they can occur in this engine. It would seem impossible to make a comprehensible map like the MELT diagram that includes all five of these injections, but it may not be advisable to use all five, or even four of these injections at once.

The Early Bowl DI, for example, does not seem like it would be an advantage in any of the conditions using a Late DI. The Late DI necessarily has to be directed into the bowl, and will produce the least soot when the charge inside the bowl has as much oxygen as possible.

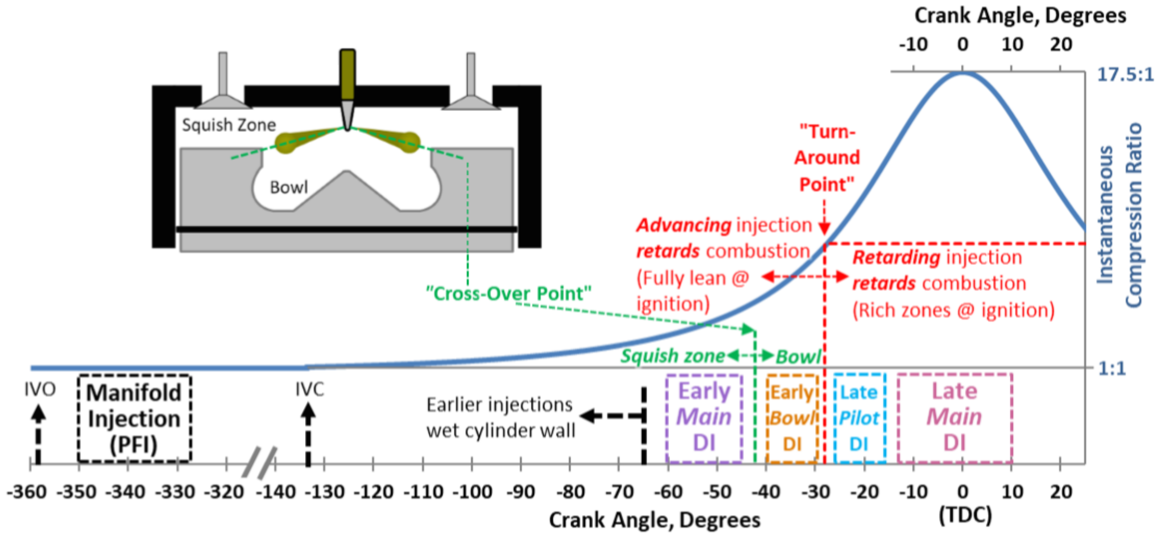

Figure 5.2 Crank angle location of all the injection types used in this study. The three "major" injections (Manifold, Early Main DI, and Late Main DI) were used to form the MELT diagram, while the two "minor" injections were excluded. This chapter considers these two minor injections, but only when one of the major injections is not included, for the reasons discussed on 130-131.

An Early Bowl DI would consume oxygen in the bowl before the Late DI has a chance to autoignite, which will make the increased soot problem of PPCI and PDFC even worse. Thus, the use of the Early Bowl DI will be limited to those conditions which have no Late DI, which in this case would be RCCI.

The Late Pilot DI, on the other hand, does not seem like it would be an advantage in any of the conditions using an Early DI (regardless of where the Early DI is directed). This is because the Early DI is already achieving the purpose of the Late Pilot DI, which is to release a small amount of heat before the Late Main DI in order to shorten the ignition delay. Thus, the Late Pilot DI will only be considered in conditions that have no Early DI, which in this case would be CDC. Although the PHCCI conditions also have no Early DI, it does not make sense to use a Late Pilot DI in this mode either since it also features an early lean premixed autoignition that accomplishes the ignition delay reduction for the Late Main DI without any need for a pilot.

Considering these two minor injections gives a more thorough examination of the possible performance of CDC and RCCI, which typically feature a Late Pilot DI and an

Early Bowl DI in most implementations. It was necessary to simplify each of these modes to a single injection for the sake of creating the continuous 2-D MELT diagram, which could necessarily only feature three types of injection. However, at least in the case of this particular engine and fuel combination, it will be shown later in this chapter that neither of these minor injections produces an improvement in performance over any of the conditions that feature only the main three injections featured on the MELT diagram.

5.1.2 Speed/Load A

The speed/load condition for which most of the Late Pilot DI and Early Bowl DI testing was performed was speed/load A. Although not all of the data will be included here, tests were performed for every possible combination of the Late Pilot DI with the Late Main DI and Manifold Injections (i.e. the CDC and CDF conditions) and of the Early Bowl DI with the Early Main DI and Manifold Injections. Using a Late Pilot DI in the CDF conditions only produced increases in soot and decreases in ITE at constant NO_X levels, so the Late Pilot DI will only be considered in combination with CDC.

In order to determine the select test conditions for repeated testing, additional contour plots of ITE were generated that, instead of being mapped onto the MELT diagram, were mapped onto an X-Y axes of EGR and SOI at a constant MELT ratio. In Figure 5.3, six of these maps are shown, featuring not just the contour map of gross ITE, but also selected contours of the four major emissions. All of the contours seen are produced using a quadratic polynomial fit, again with filled circles corresponding to the actual data, and the same color scheme used for the MELT contour maps in Chapter 4. The areas of the EGR/SOI space outside of the knock and misfire limits are colored over in gray and light brown respectively, since the interpolated contour maps will not be valid in this area. The injection distribution of these six contour maps is shown on a reprint of the Type 2 MELT contour map (with a 1 g/kWh NO_X limit) at the top left of Figure 5.3. The UHC and CO levels chosen correspond to the approximate levels of UHC and CO that were observed when the unburnt loss exceeded 2% TE, which generally limited overall ITE at or below baseline CDC levels. Meanwhile, the levels chosen for the NO_X and soot contours were based on the maximum allowable emissions for new passenger vehicles in both the US and EU in the mid-2000's (EPA Tier 2 and Euro 4 regulations, respectively). These were the last NO_X and soot emissions regulations that were achievable without aftertreatment, with the tighter regulations phased in by 2010 effectively requiring all new CI engine vehicles to employ SCR aftertreatment systems. These standards serve as a convenient target for the engine-out emissions being measured in this study, but it must be noted that emissions at these levels will still require aftertreatment to meet the much more stringent standards in place today.

It is very useful to see the emissions contours overlaid onto the ITE contour map to understand how efficiency is limited by the need to meet these emissions, as well as the knock and misfire limits. For example, in MELT point number 1 (with 100% Late DI), the NO_X limit appears to be limiting efficiency, with all of the higher ITE's occurring in the direction of increasing NO_X (as indicated by the arrows coming off of the NO_X contour). Even when the EGR is high enough to bring NO_X below the limit, the highest ITE's occur beyond the soot limit of 0.1 g/kWh (the dark brown contour). Only in the small triangle between the NO_X, soot, and misfire limits (between 25-30% EGR and 7-10°bTDC SOI) can the targeted emissions be achieved. The highest ITE is found at the tip of this triangle, and in this case it is not much lower than the maximum ITE found anywhere on the map, so there is not too much of a trade-off between emissions and efficiency when 100% Late DI is used.

The emissions/efficiency trade-off becomes much more apparent when 10% Pilot DI is applied, in MELT point number 2. The triangle in which all emissions targets are being met has been shifted to lower EGR and SOI levels, but the ITE at the tip of this triangle remains about the same. What has changed is that the knock limit has been greatly advanced by adding the Pilot DI, and the peak ITE that can be achieved is about 2.5% TE higher than the

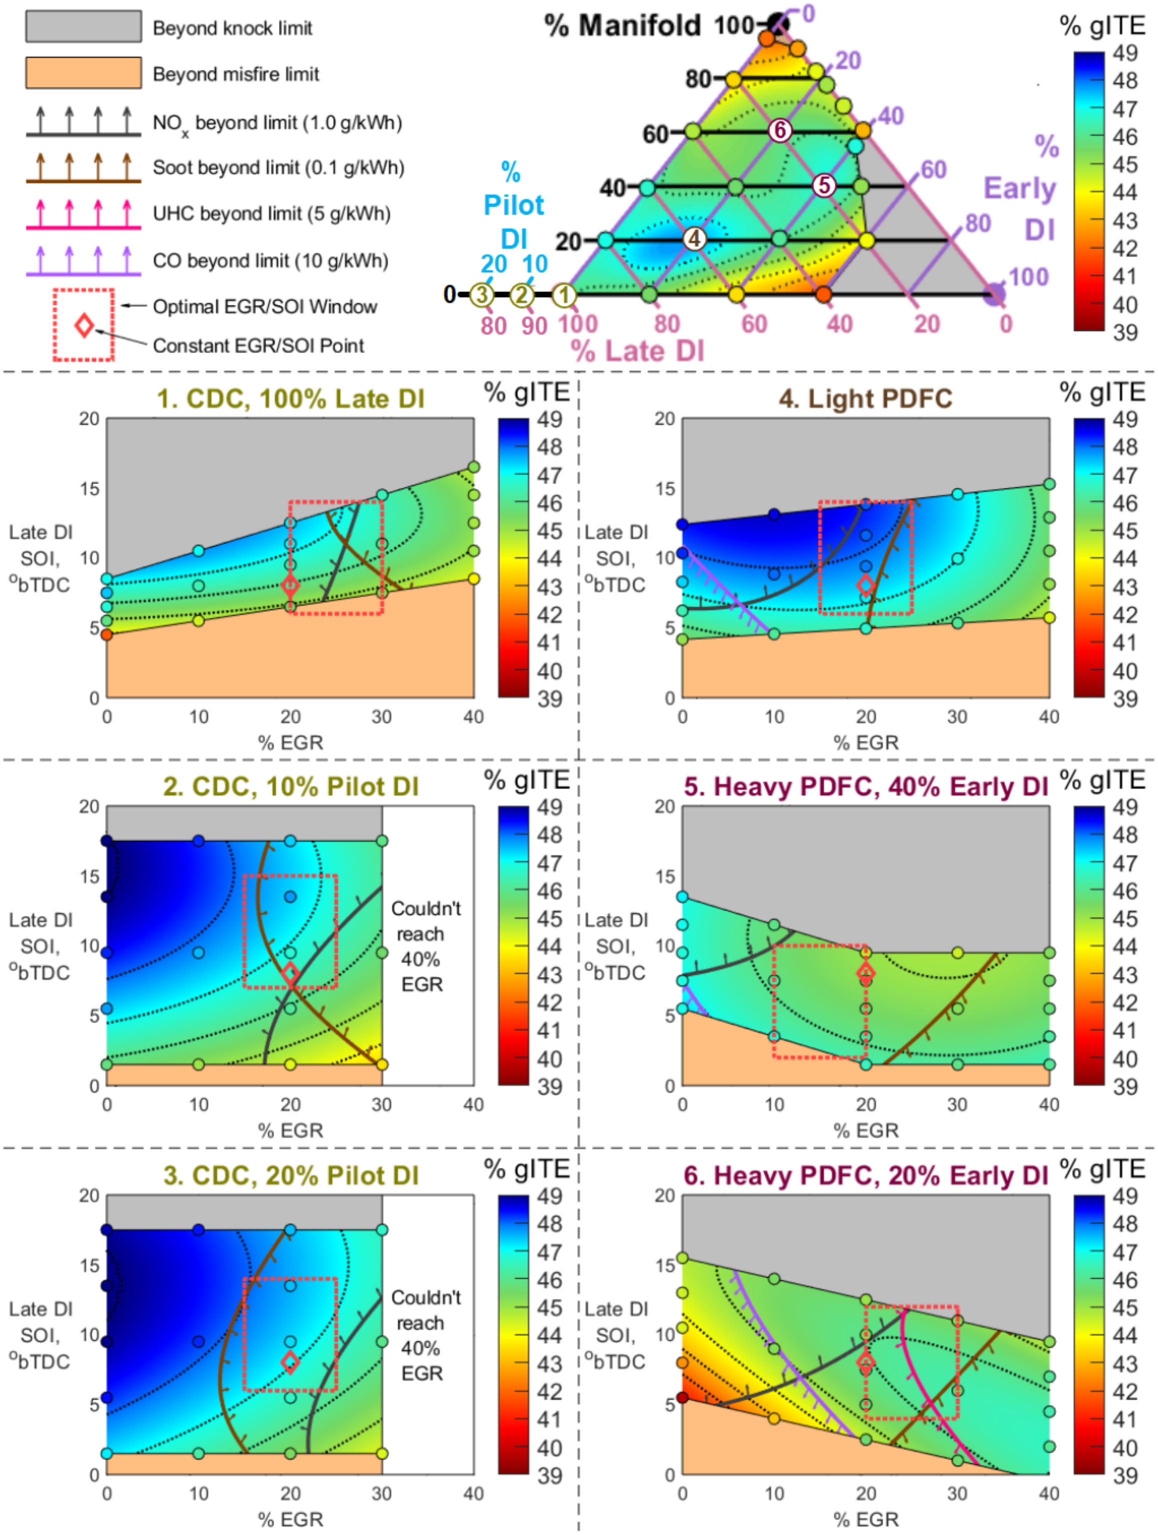

Figure 5.3 Maps of gross ITE as a function of EGR and SOI for speed/load A (2300 RPM, 4 bar BMEP), with selected emissions contours used to select a smaller EGR/SOI window for more extensive testing. Discussed on pages 132-137.

peak ITE within the emissions limits. For the repeated testing, the tests should focus in on a small window of 10% EGR and 8° SOI, and given this limitation the window cannot cover both the low-emissions and high-efficiency areas. The best compromise is to position this window to include the intersection of the emissions limits and then moving it up as far as possible, and the chosen window is shown by the red dotted box. For MELT point 2, one could argue that this window could be shifted to the left to reach higher ITE and still have the intersection of the emissions limits, but it is useful to have the window to contain equal areas below the NO_X limit and below the soot limit to be able to different test strategies that emphasize reductions in each of these emissions. It should also be noted that 40% EGR could not be reached when this testing occurred due to issues with the EGR valve that were resolved before the final testing commenced.

When using a larger Late Pilot DI of 20% total fuel, in MELT point number 3, there is no longer any overlap between the NO_X and soot limits. Moving from point 1 to point 2 and then point 3, the NO_X limit is being pushed down and to the right, since the higher temperatures and reduced ignition delay increase the amount of NO_X generated by the diffusion flame, and more EGR is needed to dilute this flame to reach the NO_X limit. Meanwhile, the soot limit is being pushed down and to the left, as the reduced ignition delay and reduced oxygen concentration in the bowl both lead to richer equivalence ratios during the rich premixed autoignition, hence more soot production. Initially, with just 10% Late Pilot DI, this is OK since the misfire limit has been retarded, and the area of overlap between the two limits shifts down to the later SOI's that were previously inaccessible. However the 10% Pilot was enough to move the misfire limit back, and the overlap between the two emissions limits is shifted down below the misfire limit. For this reason, the Late Pilot DI will be capped at 10% total fuel for the final testing.

Moving instead to the Light PDFC key point that was previously identified in Section 5.3 (MELT point number 4 in Figure 5.3), the area of overlap between the NO_X and soot limits has increased, and now includes a high ITE point of 48.1%, close to the peak ITE of 48.6%. As expected, the soot output initially increases when first moving into PDFC, and the soot limit shifts to higher EGR's. However, this is surpassed by the drop in NO_X production achieved by reducing the amount of diffusion combustion, and the NO_X limit moves further to the left than the soot limit. Meanwhile, the peak ITE is shifting to the right, indicating that EGR is more beneficial to the ITE of PDFC than it is to the ITE of CDC. This makes sense when considering that increased EGR will both increase intake temperatures and bring equivalence ratios closer to stoichiometric during the lean premixed autoignition, allowing for more complete combustion and lower unburnt loss. With no evenly-dispersed, low-reactivity gasoline, CDC does not see this same benefit from EGR. The combination of these shifts allows the optimal window to cover a large area in which the NO_X limit, soot limit, and near-peak ITE all overlap, making this MELT point an apparent "sweet spot" compared to the rest, and explaining why the efficiency peak appears here on the MELT contour map. It is not that CDC cannot produce similarly high efficiency, but that peak efficiency does not overlap with the NO_X and soot limits like it does in Light PDFC.

However, the ITE trend is apparently reversed at the original Heavy PDFC key point (MELT point number 5 in Figure 5.3). All other MELT points feature a trend of increasing ITE with advancing SOI, but this point features a trend, albeit relatively slight, of increasing ITE with *retarding* SOI. Meanwhile, more EGR is leading to lower ITE, the opposite of the trend that would be expected given the minimization of unburnt losses. The only explanation for this is that the combustion phasing has been shifted too early with this combination of injections, and as a result the combustion phasing must be retarded either through reduced EGR or retarding SOI. But, unfortunately, the peak ITE is not reached before the cycle-to-cycle variation becomes too great, and the misfire limit is approached. This is also not a desirable trend for the fact that retarding the SOI and decreasing EGR (and thus decreasing intake temperature) will increase unburnt losses. This trade-off between combustion phasing and unburnt loss results in lower peak ITE, despite an even larger area of overlap between the NO_X and soot limits. The large amount of premixed diesel from the Early DI (40% total fuel) is simply too much, and needs to be reduced in order to flip the ITE trend back toward favoring advanced SOI's and increased EGR.

This is somewhat successfully accomplished in MELT point 6, which reduces the Early DI to 20% total fuel. There is again a wide area of overlap between the NO_X and soot limits as a result of having so little fuel (20% total fuel) in the Late DI. However, it appears that this has reduced the reactivity of the premixed fuel too much, and UHC and unburnt losses are now becoming a much bigger problem. At least 25% EGR must be used to keep UHC below the limit, and even then there is still a significant amount of unburnt fuel (and increased cooling losses from the increased EGR) that limits the peak ITE to about the same level as before. The ITE peak is also now occurring with high EGR beyond the soot limit, since even with the small amount of Late DI these high EGR levels (combined with the oxygen consumption of the lean premixed autoignition) make the rich premixed autoignition especially rich. For these reasons, it seems that the ideal Early DI size for Heavy PDFC lies somewhere in the middle at around 30% total fuel: enough to ensure that the premixed fuel is sufficiently reactive without excess EGR, but not so much that it inherently advances combustion phasing too far.

With the ideal EGR/SOI window identified for each of these six points, the four key MELT points to be tested can now be finalized, as well as the EGR/SOI windows that will be used for each. There must also be a *constant* combination of EGR and SOI that will fall within each of these windows to ensure that the tests with small increments of the fueling ratio are all within reasonable levels of emissions and efficiency. This point is depicted by the red diamonds in Figure 5.3, which are all located at 20% EGR and 8°bTDC.

For the CDC key point, it was anticipated that 10% Early DI would be optimal, given the performance trends seen with 10% Late Pilot DI. Although the Early DI burns through a different process, it should still release approximately the same amount of heat as the Late Pilot DI, and have a similar effect on reducing ignition delay. With preliminary testing it was found that the window identified with the 10% Late Pilot DI also produced similar results

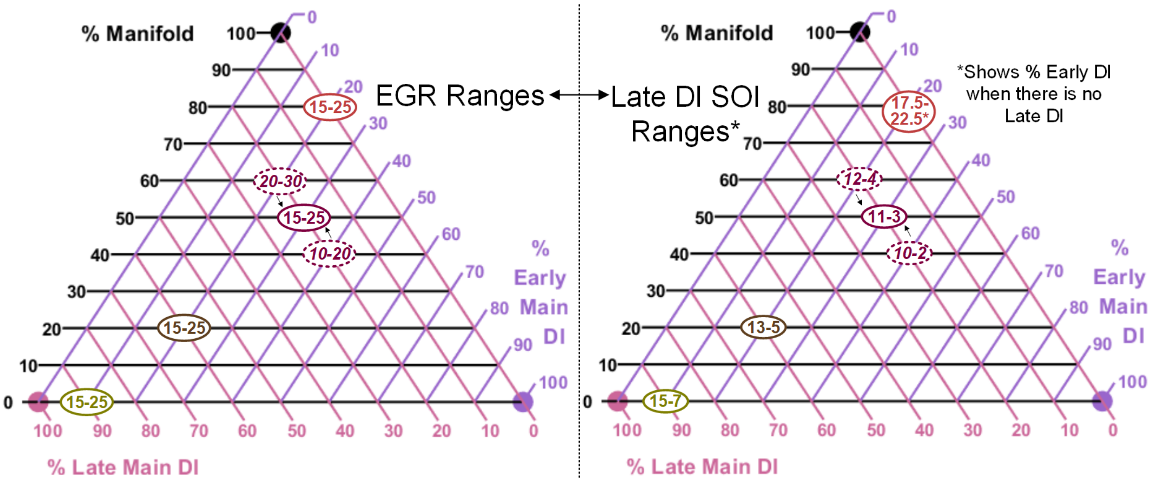

Figure 5.4 Finalized key points for speed/load A (2300 RPM, 4 bar BMEP). The left diagram shows the ranges of EGR and the right diagram shows the ranges of SOI that were identified as the ideal "window" in Figure 5.3. The Heavy PDFC key point chosen was an average between the two Heavy PDFC points shown in Figure 5.3, with the windows identified for these two points (at 20% and 40% Early DI) in the dotted circles, and the window for the finalized key point (at 30% Early DI) in the solid maroon circles. This is an example of the process used for determining all the key points at the different speed/load conditions, which required interpolating between the regularly-spaced MELT points from Chapter 4 to get a point that lies right on the path of high ITE, which does not necessarily coincide with the MELT points origianlly acquired in Chapter 4.

with 10% Early DI, and will use that same window. This is shown in dark yellow below in Figure 5.4, with the EGR window marked on the left MELT diagram and the SOI window marked on the right MELT diagram. Similarly, the Light PDFC windows are marked in brown. For the Heavy PDFC key point, it was believed the best performance would occur about halfway between the two cases that were examined, and as such those two windows (in maroon with the lines dotted and the text italicized) were averaged in order to determine the window to be used for the final Heavy PDFC key point at 30% Early DI (in maroon with solid lines and normal text). Finally, for the RCCI key point, there will not be a Late DI and so the SOI range is not applicable, and the tests will instead look at a window of Early DI size, with a range of 5% total fuel.

Although there may be some differences in the exact ideal range of EGR to use at these points, it is also convenient to be able compare results at constant EGR across the four different combustion modes. For this reason, it was decided to keep the EGR window constant rather than adjusting it for each of the four key points, even if it looks like shifting EGR slightly in one direction may be better for one of the combustion modes. The EGR window will however shift between the speed/load points out of necessity.

To determine the MELT point to use when testing small increments of fuel ratio, a straight line is plotted on the MELT diagram in between each of the key points, as seen in Figure 5.5, and 3 evenly-spaced points are taken along this line. To add in an evaluation of the Late Pilot DI for CDC, additional points are taken at 10% and 5% each of both Early Main DI and Late Pilot DI, along with a 100% Late Main DI point that features neither Early Main DI nor Late Pilot DI. This will be repeated for the other speed/load conditions, with the amounts changed in proportion to the amount of Early DI in the key CDC point in each speed/load. To add in an evaluation of the Early Bowl DI for RCCI, the Early Main DI fuel will be transferred to the Early Bowl DI in 5% increments, while maintaining a constant level of Manifold fuel. For the EGR/SOI window at each of the four key points, a 3x3 grid will be tested, with Early DI percentage replacing SOI for the RCCI key point. This completes the final test plan for speed/load A, as shown in Figure 5.5. About half of these points were tested three times on separate days, but some were tested only once, as indicated by the “3’s” and “1’s” in Figure 5.5. This same exact scheme will be applied to the other speed/load points once their key MELT points have been located, discussed in the next section.

Ideally, these points would have been tested many times each, since the test setup did not produce perfectly consistent results from day to day due to a number of factors. Some of these factors were due to the limitations of the test setup itself, which was not in a climate-controlled room, leading to unavoidable day-to-day differences in ambient temperature and intake air humidity. However, other factors were inherent to the engine, such as the continuous build-up of deposits on the EGR cooler that changed its performance over time (and hence changing the intake air temperature in relation to EGR fraction). Perhaps the biggest limiting factor for this particular test setup was the inability to get consistent injection

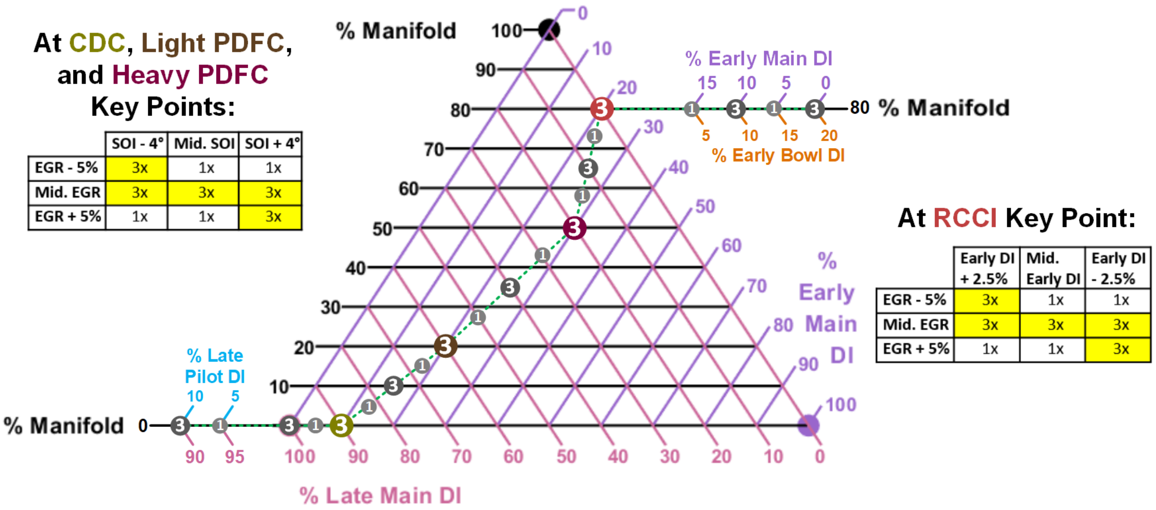
At CDC, Light PDFC, and Heavy PDFC Key Points:

	SOI - 4°	Mid. SOI	SOI + 4°
EGR - 5%	3x	1x	1x
Mid. EGR	3x	3x	3x
EGR + 5%	1x	1x	3x

At RCCI Key Point:

	Early DI + 2.5%	Mid. Early DI	Early DI - 2.5%
EGR - 5%	3x	1x	1x
Mid. EGR	3x	3x	3x
EGR + 5%	1x	1x	3x

Figure 5.5 Full finalized test plan for speed/load A. This shows the compromise between getting multiple repetitions of the same data point to quantify uncertainty and testing small increments of MELT ratio, SOI, and EGR to obtain higher resolution in the trends. The points labeled “3” or “3x” were tested three times, while the points labeled “1” or “1x” were tested only once. The results along the green dotted line (with a constant SOI and EGR) are shown in Section 5.2, where the other speed/load points will use the same method of plotting a straight line between the key points and adding the Pilot and Early Bowl injections on the end in the same manner as shown here. The results at the key points (with varying SOI and EGR) will be shown in Section 5.3.1.

sizes with the exact same commanded injection duration and injector rail pressure. This required the injection duration to be set again every time the same condition was tested through a tedious guess-and-check cycle by slightly changing the injector duration and watching the response in the measured fuel flow, torque, and IMEP (measurements which always had much more consistent correlation with each other than they did to injection duration). Not only did this cause small changes in the actual injection sizes, but it increased the amount of time required to set up a single test condition and limited the number of points that could be tested.

These and other limitations led to the final test plan in which every other point was sampled only once, while the others were sampled three times. The tests were carried out starting with testing *every* point once, then analyzing that data to confirm which points would be best to test a second and third time. The tests were scheduled such that each 3-by-3 grid and each sweep of the incremental path occurred over the same day, rather than being split across multiple days. This may have increased some risk of temporal effects on the results, but it had the advantage of ensuring that results within each sweep and within each grid were consistent. This made sure that differences from point to point within these sweeps and grids were not exaggerated by taking these points with long gaps of time in between, during which the overall engine performance may have shifted. However, when the second and third repetitions were taken for the selected points, this shifted the average result, as the result of day-to-day performance changes.

For this reason, when the results are shown, the points that were only taken once have been shifted to preserve the point-to-point differences that were observed in the first repetition of the test, which included *all* of the test points being taken on the same day. In the sweep of the path through the MELT diagram, all of the points that were only tested once (indicated by the “1’s” on Figure 5.5) were shifted by the average difference between the first repetition and the three-repetition average in the two test points on either side. Without this shift applied, the trends observed in the sweep would be “jagged,” with the results from the 1-repetition points being slightly offset from the 3-repetition points. This allows for a compromise between getting finer resolution in the sweep of the MELT diagram and not taking up too much time and fuel by repeating every point three times. The same was applied to the high and low EGR levels in the 3-by-3 grids. The difference between the first repetition and the 3-repetition average on the single point that was repeated was then applied to two other points at that EGR level. For the middle EGR level (which was designed to produce the most optimal results), all three points were repeated three times, so no correction was necessary.

This process is not perfect, and does not completely eliminate any sort of temporal bias on the results. However, for the purposes of extending the testing to as wide of a range as possible and getting a more complete understanding of the shifts in performance, but within a reasonable amount of time, it was necessary to compromise the precision of the data somewhat. The results will show that the uncertainty in the “in-between” points does not significantly affect the conclusions of this thesis.

5.1.3 Other Speed/Load Conditions

The same process used to determine the final test plan for speed/load A will applied to speed/loads F and C. Interpolation in between these speed/load points will be used to develop the plan for the remaining three speed/load points B, D, and E, which are conveniently located exactly halfway in between each combination of the original three points. Finally, some quick preliminary testing will be conducted to confirm the test plan for B, D, and E.

The same EGR/SOI contours that were shown before for speed/load A are produced again for speed/load F in Figure 5.6. Again, these contours are used to find the EGR/SOI windows that offer the best compromise between high ITE and the four emissions limits.

When all of the fuel is in the Late DI at MELT point 1, it appears that the knock limit is very restricting, with no SOI's earlier than 4°bTDC available for 20% EGR or lower. This is due to the very long ignition delay and large rich premixed spike produced in this speed/load, caused by a combination of lower injection pressure and high heat transfer losses. Adding a 10% Early DI in MELT point 2 greatly advances the knock limit and increases efficiency, with a very wide, flat peak ITE region stretching from 0 to 30% EGR. However, only a narrow slice of this region between 20% and 30% EGR meets both the NO_X and soot limits. But fortunately, and unlike speed/load A, this area overlaps with a high ITE of 45.3% at the condition of 30% EGR and 12°bTDC SOI that is very close to the peak ITE of 46.1%, although given its relation to the contour fit this point may just be a random outlier. Regardless, the overall trends show that the emissions limits pose much less of a restriction on ITE for CDC in speed/load F than they did in speed/load A.

However, when moving into Light PDFC in MELT points 3 and 4, the overlap between the NO_X and soot limits does not widen like it did in speed/load A, and while this gap still occurs near peak ITE, the peak ITE is not as high as it was before. Although there is still very high ITE close to the CDC peak in the condition of 30% EGR and 7°bTDC in MELT

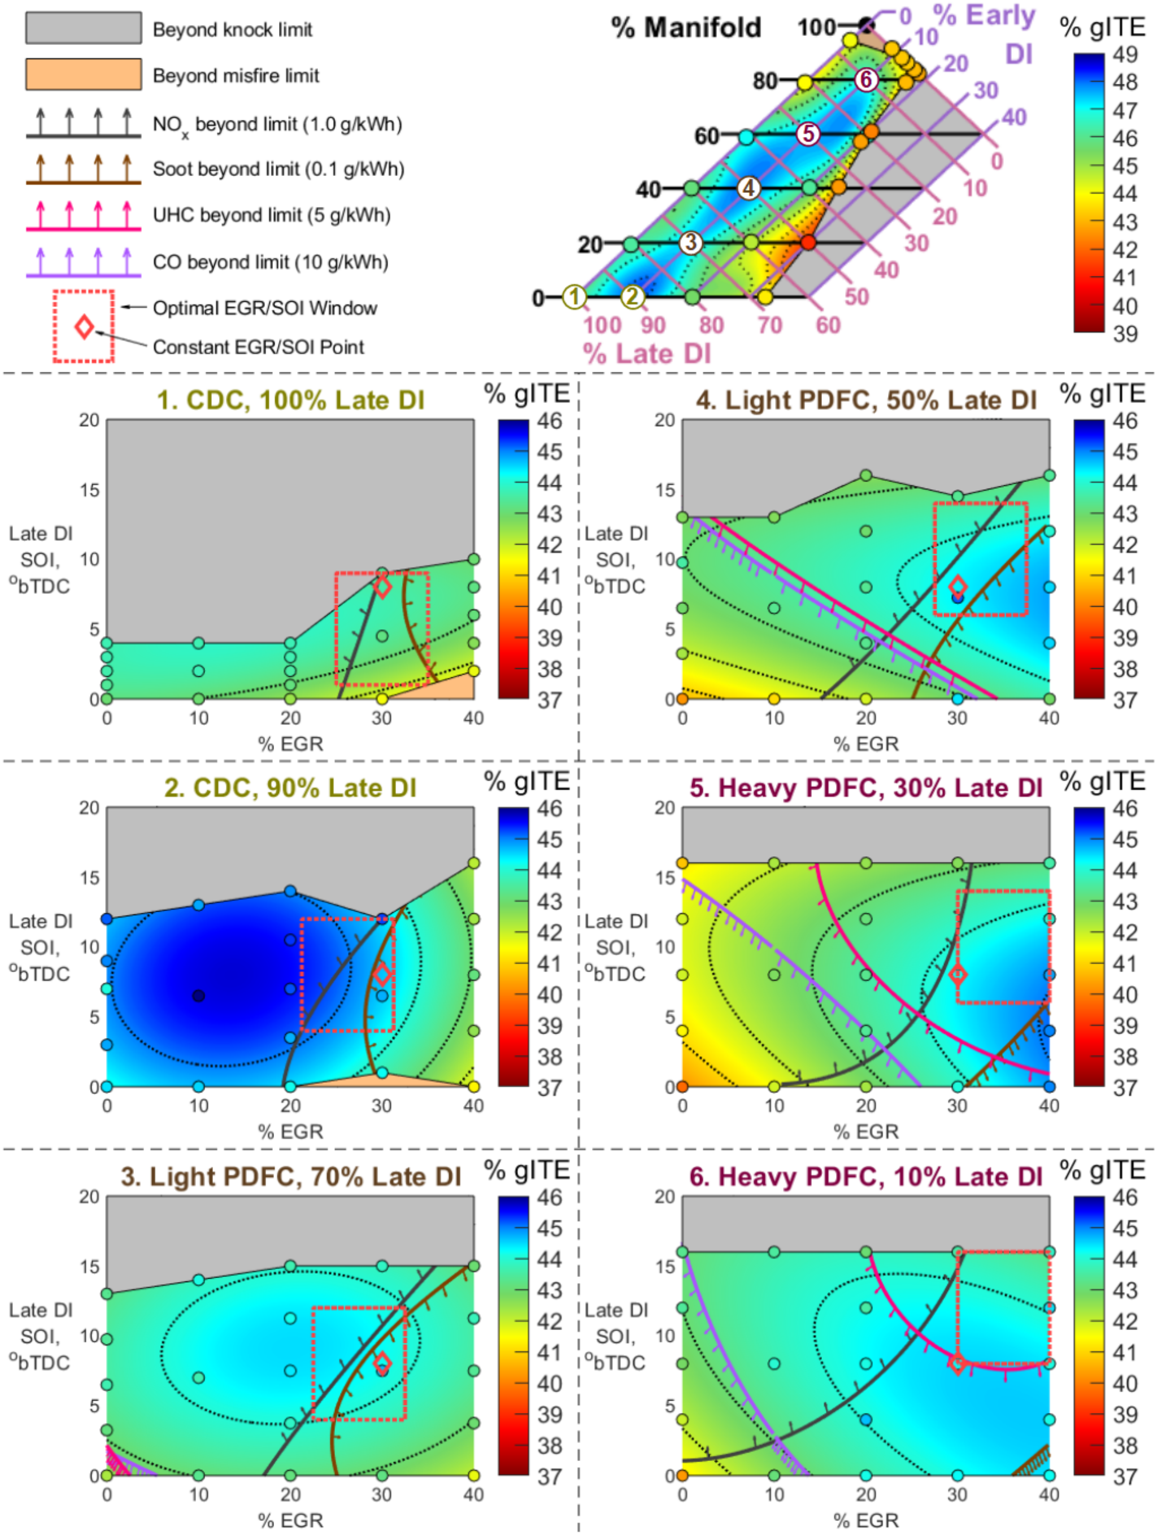

Figure 5.6 Maps of gross ITE as a function of EGR and SOI for speed/load F (1200 RPM, 4 bar BMEP), with selected emissions contours used to select a smaller EGR/SOI window for more extensive testing, discussed on pages 142-144.

point number 4, this again appears to be an outlier when seen in comparison with the other points in this EGR/SOI contour map. If this is indeed an outlier, it appears to have produced a false ITE peak in the original MELT contour map from Figure 4.10, reprinted at the top right of Figure 5.6. One trend that is preserved from speed/load A is the shifting of the peak ITE from low EGR to high EGR, as the unburnt loss becomes more and more of a factor. With the higher heat transfer losses and lower amount of allowable Early DI, the UHC losses are more dominant in determining the ITE than they were in speed/load A.

When moving into Heavy PDFC in MELT points 5 and 6, the overlap between NO_X and soot limits starts to widen, but the UHC limit creeps up further and further. In MELT point 6, with only 10% Late DI, the ITE peak interestingly appears to occur beyond UHC limit, suggesting that the combustion phasing may have become too advanced, and it is actually worth taking a high unburnt loss for the sake of retarding the combustion phasing back to a more optimal position. However, given the very small differences in ITE across this region, it cannot be said that the position of the peak ITE is accurate, and not just an outcome of noise in the data. Since near-peak ITE can be achieved while meeting the UHC limit, the EGR/SOI window will be kept behind this limit.

Moving on to speed/load C of 2300 RPM and 10 bar BMEP, shown in Figure 5.7, the UHC and CO limits never appear, due to the higher overall temperatures and pressures resulting in more complete combustion in this load, plus the fact that the conditions can never go above 50% total premixed fuel or a lambda of 2. However, the NO_X limit must also be raised to 1.5 g/kWh as previously discussed in Section 4.4, since there will never be any overlap between the NO_X and soot limits if the NO_X limit remains at 1.0 g/kWh.

Even with this raised NO_X limit, in MELT point 1 with 100% Late DI there is still only a tiny area of overlap between the NO_X and soot limits, and it occurs with less than 45% ITE, well below the peak ITE of 47.8%. This area of overlap is greatly increased with the addition of just 5% Early DI, as seen in MELT point number 2. Still, the peak ITE of 49.3% is far offset from the NO_X limit, and the EGR/SOI window is just large enough to catch

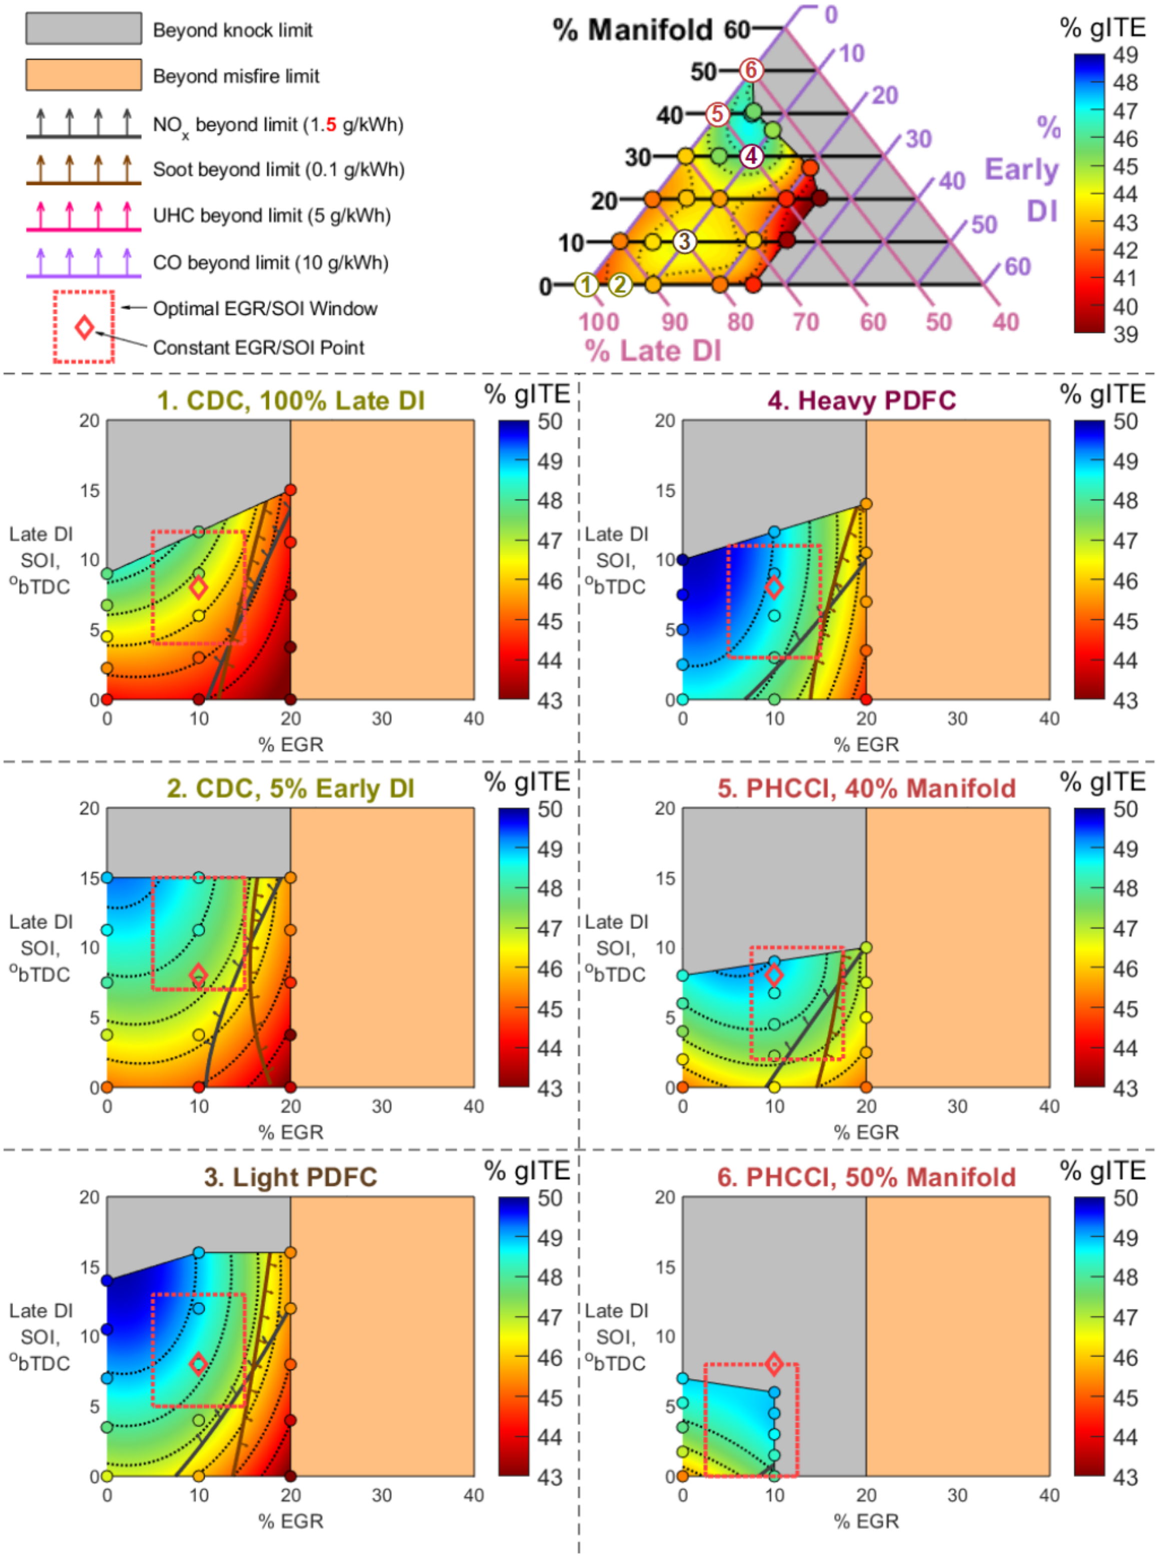

Figure 5.7 Maps of gross ITE as a function of EGR and SOI for speed/load C (2300 RPM, 10 bar BMEP), with selected emissions contours used to select a smaller EGR/SOI window for more extensive testing, discussed on pages 144-146.

part of the emissions limit overlap (with at most 46.5% ITE) in the bottom right corner, and capture near-peak ITE of 49.1% in the top left corner. The knock limit has also been advanced, although not by as great of a degree as it was in the lower loads.

Moving to Light PDFC in MELT point number 3, it appears that peak ITE has increased across the board, but the NO_X limit has also shifted downward. Both of these effects may stem from the newly-introduced lean premixed autoignition, which allows more heat release to occur earlier in the cycle without ringing (improving ITE) but increases the temperatures in the diffusion flame (increasing NO_X). However, when moving to Heavy PDFC in MELT point number 4, there is only a very slight change in the NO_X limit, while the ITE improves significantly, especially in the area of overlap between the emissions limits. This may be due to the emergence of the *homogeneous* lean premixed autoignition, which again shifts combustion phasing closer to TDC and also can reduce the heat transfer losses by distributing temperatures more evenly throughout the combustion chamber. Still, the window can only cover the emissions limit overlap in the bottom right corner and the peak ITE in the top left corner.

Moving to the PHCCI conditions in MELT point 5, the peak ITE begins to shift to the right, indicating that unburnt loss may be starting influence the ITE more than cooling loss. This makes sense considering that the Early DI has been eliminated, and the homogeneously-mixed low-reactivity gasoline now has to autoignite all on its own and will autoignite more completely with more EGR and a lower lambda. The knock limit is also retarded, the reasoning for which is complex. Even though the elimination of the Early DI retards the homogeneous premixed autoignition peak, this means that the Late DI now needs to be retarded further to prevent the rich premixed peak from being "stacked" on top of the homogeneous autoignition peak, which creates a very large peak HRR and thus high MPRR. Increasing the Manifold fuel further to 50% total fuel in MELT point 6 completes the rightward shift of the peak ITE, and should allow the NO_X limit to be achieved with closer-to-peak ITE, but unfortunately with this much gasoline the engine is now knocking with more than 10% EGR. It would seem that a good compromise would be to use 45% Manifold fuel to put the results halfway in between MELT points 5 and 6.

Having obtained these results from F and C, the four key points on the MELT diagram can now be selected, as shown in Figure 5.8. Where the key points had to be shifted away from the points that were shown in Figures 5.6 and 5.7, the original points are indicated by the dotted circles, while the finalized key points for the repeated testing are shown by solid circles. For speed/load F, the 10% Early DI CDC point produced good results and there is no need to alter the ratio; this point is the exact same as the CDC key point for speed/load A and is offset below the MELT diagram in Figure 5.8 just so that both A and F can be shown. The two Light PDFC and Heavy PDFC points shown from speed/load F produced very similar results, so the key points will be placed halfway in between both of these pairs. This also has the advantage of giving these key points the same amount of Late DI as speed/load A, for better direct comparisons. As discussed in the previous paragraph, the PHCCI condition for speed/load C was chosen to compromise between the two MELT points seen in Figure 5.7. The rest of the key points for speed/load C are unchanged from the points analyzed in Figure 5.7.

With the MELT points for A, F, and C finalized, there can be interpolation halfway in between them to produce the initial proposed key points for speed/loads B, C, and E. The exception to this is the PHCCI cases in B and E, where there cannot be interpolation between the PHCCI condition in point C and the RCCI conditions in A and F. It was reasoned that the ideal PHCCI point for B might occur with the same amount of premixed gasoline as C, plus a little extra to accommodate for the lower overall temperatures in this lower load that would allow more gasoline to be premixed before knocking. In speed/load E, which has the same load but a lower engine speed, it was reasoned that a little less gasoline could be used before knocking, since the autoignition process would be occurring at the same speed with respect to time, but more quickly with respect to crank angle.

Of course, some preliminary testing was needed to confirm these key points, but much

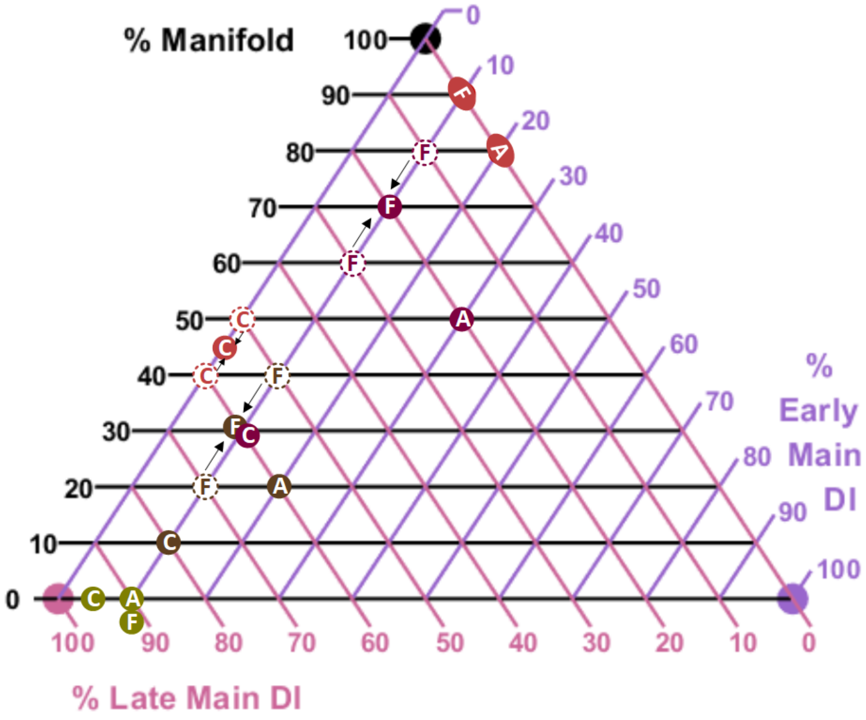

Figure 5.8 Finalized key points for speed/loads A, F, and C. The dotted circles show original points from Chapter 4, and solid circles are finalized key points that were interpolated halfway in between the original points. The Heavy PDFC point for speed/load C overlaps with the Light PDFC point for speed/load F, so these points are slightly displaced from each other to allow both of the letters to be seen, but both occur at exactly 30% Manifold, 10% Early DI, and 60% Late DI. The selection process is discussed on page 147.

less than the extensive, time-consuming mapping process used for A, C, and F. This testing found that the proposed CDC, Light PDFC, and RCCI points were well-positioned and producing results comparable to the other speed/load points. However, for the originally proposed points for Heavy PDFC and PHCCI in speed/loads B and E, it was found that they could not be reached without knocking. By trial and error, it was found that high ITE close to the knock limit could be produc by shifting theed Heavy PDFC points so that they were interpolated *three-quarters* of the way in between the two adjacent speed/load points, being closer to the high-load condition. This is shown on Figure 5.9, with the originally proposed points in dotted lines and the finalized points in solid-colored circles.

For the PHCCI points in speed/loads B and E, it was not possible to get the homogeneous premixed autoignition peak to occur completely before the rich premixed autoignition peak

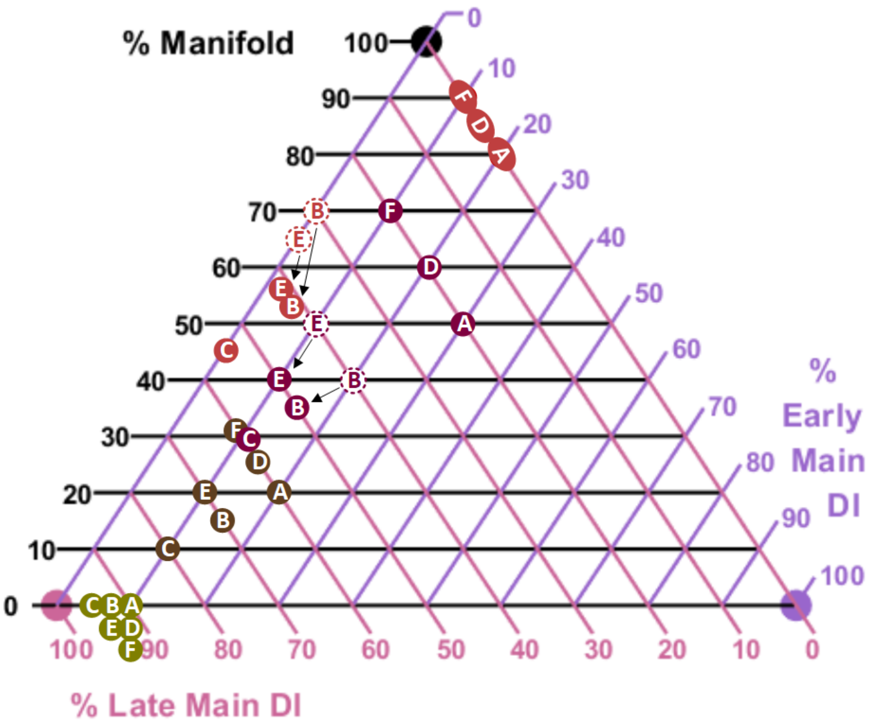

Figure 5.9 Finalized key points for all six speeds/loads A through F. The dotted circles show the original estimated key points, while the solid circles show the finalized key points that were determined after preliminary testing, as discussed in pages 147-149.

without using a small amount of Early DI, 2.5% for speed/load E and 5% for speed/load B. These small amounts of Early DI are low enough to still consider these conditions to be PHCCI and not PDFC. This actually has the effect of positioning the key points three-quarters of the way in between the RCCI points of A and F and the PHCCI point of C, keeping consistency with the way the Heavy PDFC points are arranged. This completed the test matrix design, with the EGR/SOI windows interpolated in the same manner. The exact MELT ratios used at each point are given in Table 5.2 and the results of this test matrix will be presented in the following section.

5.2 Fine Increments of Fueling Ratio

The first set of results presented here consists of the sweeps of small increments of fueling ratio that were conducted with constant EGR and constant SOI over the entire sweep. The exact injection ratios are given in Table 5.2 and the results will be presented in the form of line graphs, with one line for each speed/load condition. To make the graphs more readable, the three higher-load conditions have been offset to the top of the chart, while the three low-load conditions have been offset to the bottom. This will divide the y-axis across a dashed line, although the scale is the same on both sides. The x-axis is constant, showing the percentage of the total fuel in the Late Main DI (i.e., how much fuel is *not* premixed), except at the ends beyond the dotted lines, where the x-axis switches to showing % Pilot DI (on the left end) and % Early Bowl DI (on the right end, only valid for the low-load conditions where RCCI was achievable). Figure 5.5 shows how this sweep was conducted on the MELT diagram for speed/load A, with the same process being repeated for the other speed/load conditions using the key points marked in Figure 5.9. These key points are also circled in their respective colors for reference. Error bars are plotted for the conditions that were repeated three times, with the height of the error bars corresponding to the 95% confidence interval (CI), calculated using the standard error formula with standard deviation σ and $n = 3$ observations:

$$95\% \text{ C.I.} = \text{Mean +/- Std. Error} = \text{Mean +/-}1.96\frac{\sigma}{\sqrt{n}} \tag{5.1}$$

The results will be presented in the same order as the contour maps from Chapter 4, starting with the knock/misfire limits and ITE, followed by the combustion phasing, then the main heat losses, and finally the emissions.

5.2.1 Limits and ITE

The results for the knock/misfire limits are presented in Figure 5.10. Also shown here are the calculated tailpipe CO_2 emissions, as derived from the gross ITE, with the scale reversed to match the trends in gross ITE. These emissions are calculated by assuming that, when these combustion modes are implemented on an actual vehicle, the UHC and CO are converted to

Table 5.2 Injection ratios used for incremental MELT diagram sweeps. The colored columns correspond to the key MELT points, and match the key points indicated in Figure 5.9. The six colored header rows for the different speed/load conditions match the color of the line plots that will be presented in Figures 5.10-13.

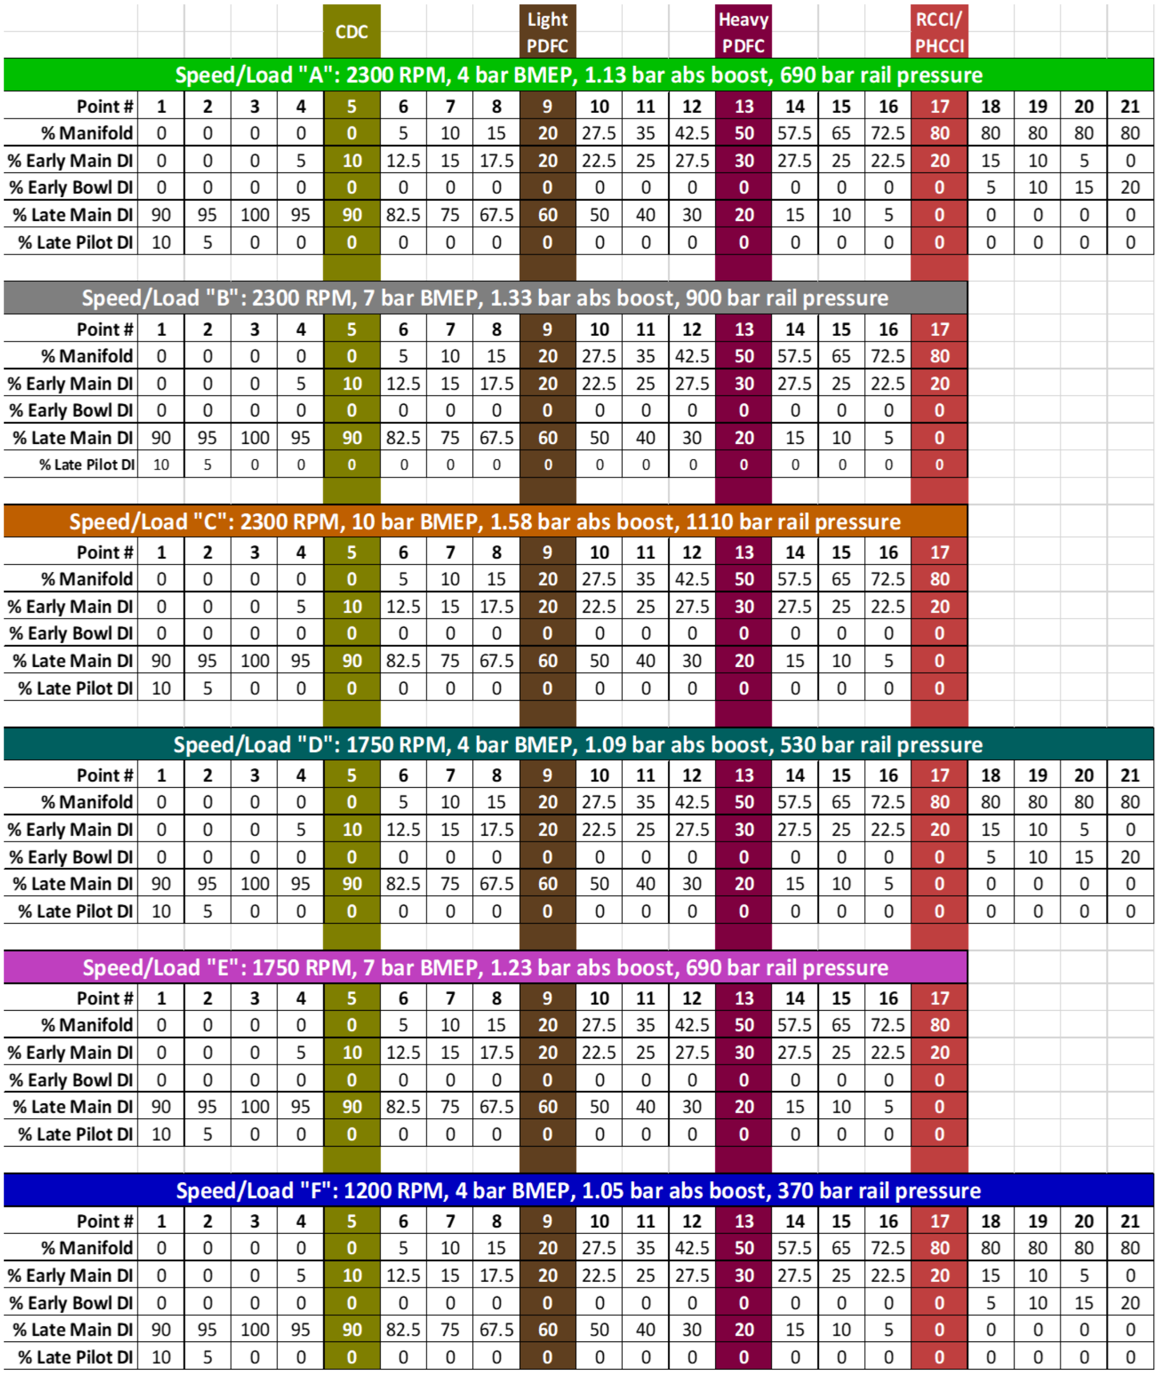

					CDC				Light PDFC				Heavy PDFC				RCCI/ PHCCI				
Speed/Load "A": 2300 RPM, 4 bar BMEP, 1.13 bar abs boost, 690 bar rail pressure																					
Point #	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**
% Manifold	0	0	0	0	**0**	5	10	15	**20**	27.5	35	42.5	**50**	57.5	65	72.5	**80**	80	80	80	80
% Early Main DI	0	0	0	5	**10**	12.5	15	17.5	**20**	22.5	25	27.5	**30**	27.5	25	22.5	**20**	15	10	5	0
% Early Bowl DI	0	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**	5	10	15	20
% Late Main DI	90	95	100	95	**90**	82.5	75	67.5	**60**	50	40	30	**20**	15	10	5	**0**	0	0	0	0
% Late Pilot DI	10	5	0	0	**0**	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**	0	0	0	0
Speed/Load "B": 2300 RPM, 7 bar BMEP, 1.33 bar abs boost, 900 bar rail pressure																					
Point #	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**				
% Manifold	0	0	0	0	**0**	5	10	15	**20**	27.5	35	42.5	**50**	57.5	65	72.5	**80**				
% Early Main DI	0	0	0	5	**10**	12.5	15	17.5	**20**	22.5	25	27.5	**30**	27.5	25	22.5	**20**				
% Early Bowl DI	0	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**				
% Late Main DI	90	95	100	95	**90**	82.5	75	67.5	**60**	50	40	30	**20**	15	10	5	**0**				
% Late Pilot DI	10	5	0	0	**0**	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**				
Speed/Load "C": 2300 RPM, 10 bar BMEP, 1.58 bar abs boost, 1110 bar rail pressure																					
Point #	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**				
% Manifold	0	0	0	0	**0**	5	10	15	**20**	27.5	35	42.5	**50**	57.5	65	72.5	**80**				
% Early Main DI	0	0	0	5	**10**	12.5	15	17.5	**20**	22.5	25	27.5	**30**	27.5	25	22.5	**20**				
% Early Bowl DI	0	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**				
% Late Main DI	90	95	100	95	**90**	82.5	75	67.5	**60**	50	40	30	**20**	15	10	5	**0**				
% Late Pilot DI	10	5	0	0	**0**	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**				
Speed/Load "D": 1750 RPM, 4 bar BMEP, 1.09 bar abs boost, 530 bar rail pressure																					
Point #	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**
% Manifold	0	0	0	0	**0**	5	10	15	**20**	27.5	35	42.5	**50**	57.5	65	72.5	**80**	80	80	80	80
% Early Main DI	0	0	0	5	**10**	12.5	15	17.5	**20**	22.5	25	27.5	**30**	27.5	25	22.5	**20**	15	10	5	0
% Early Bowl DI	0	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**	5	10	15	20
% Late Main DI	90	95	100	95	**90**	82.5	75	67.5	**60**	50	40	30	**20**	15	10	5	**0**	0	0	0	0
% Late Pilot DI	10	5	0	0	**0**	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**	0	0	0	0
Speed/Load "E": 1750 RPM, 7 bar BMEP, 1.23 bar abs boost, 690 bar rail pressure																					
Point #	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**				
% Manifold	0	0	0	0	**0**	5	10	15	**20**	27.5	35	42.5	**50**	57.5	65	72.5	**80**				
% Early Main DI	0	0	0	5	**10**	12.5	15	17.5	**20**	22.5	25	27.5	**30**	27.5	25	22.5	**20**				
% Early Bowl DI	0	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**				
% Late Main DI	90	95	100	95	**90**	82.5	75	67.5	**60**	50	40	30	**20**	15	10	5	**0**				
% Late Pilot DI	10	5	0	0	**0**	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**				
Speed/Load "F": 1200 RPM, 4 bar BMEP, 1.05 bar abs boost, 370 bar rail pressure																					
Point #	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**	**17**	**18**	**19**	**20**	**21**
% Manifold	0	0	0	0	**0**	5	10	15	**20**	27.5	35	42.5	**50**	57.5	65	72.5	**80**	80	80	80	80
% Early Main DI	0	0	0	5	**10**	12.5	15	17.5	**20**	22.5	25	27.5	**30**	27.5	25	22.5	**20**	15	10	5	0
% Early Bowl DI	0	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**	5	10	15	20
% Late Main DI	90	95	100	95	**90**	82.5	75	67.5	**60**	50	40	30	**20**	15	10	5	**0**	0	0	0	0
% Late Pilot DI	10	5	0	0	**0**	0	0	0	**0**	0	0	0	**0**	0	0	0	**0**	0	0	0	0

CO_2 in a catalytic converter (otherwise, the calculated CO_2 emissions would unfairly favor the conditions with high UHC that will ultimately be converted to CO_2). Although it may appear initially that the two graphs are identical, there is actually a slight upward shift for the CO_2 graph while moving to the right, as compared to the gross ITE. This is due to the slighter higher H:C ratio of the gasoline compared to diesel (1.89 vs. 1.82) combined with the slightly higher lower heating value (42.96 vs. 42.58 MJ/kg), which causes the gasoline to produce 1.4% less CO_2 per unit of heat released. This means that, with all else being equal, the test conditions that use more gasoline will be less damaging in terms of global warming potential. This difference is barely noticeable with the two very similar fuels of gasoline and diesel, but would be much more significant if a fuel such as natural gas was used instead of gasoline, given that its primary component is methane with a much higher H:C ratio of 4.

While examining the graph of CoV of IMEP, one should recall the contour map from speed/load A in Figure 4.4, where there was an initial “ridge” of high CoV when very small Manifold Injections were used. When this was observed, it was hypothesized that the high CoV with small Manifold Injections was due to high cycle-to-cycle variations in the amount of fuel injected by the PFI. Now, there are finer increments in the size of the Manifold Injection when moving between the CDC and Light PDFC key points. The trend seen in Figure 5.10 does not seem to match that hypothesis, since there is instead a slow upward ramp in the CoV moving from CDC to Light PDFC, which appears in all speed/load conditions to varying degrees. If the high CoV was due to the small injection size, then there should be a sudden increase in the CoV when moving away from CDC that slowly diminishes as the PFI injection gets bigger and bigger. Instead, the opposite is observed, with a slow increase toward Light PDFC followed by a sudden decrease before getting to the Heavy PDFC key point, and this trend is consistent for all six speed/loads.

This leads to a new hypothesis: perhaps the high CoV observed in Light PDFC is due to inconsistencies in the ability to achieve a homogeneous lean premixed autoignition

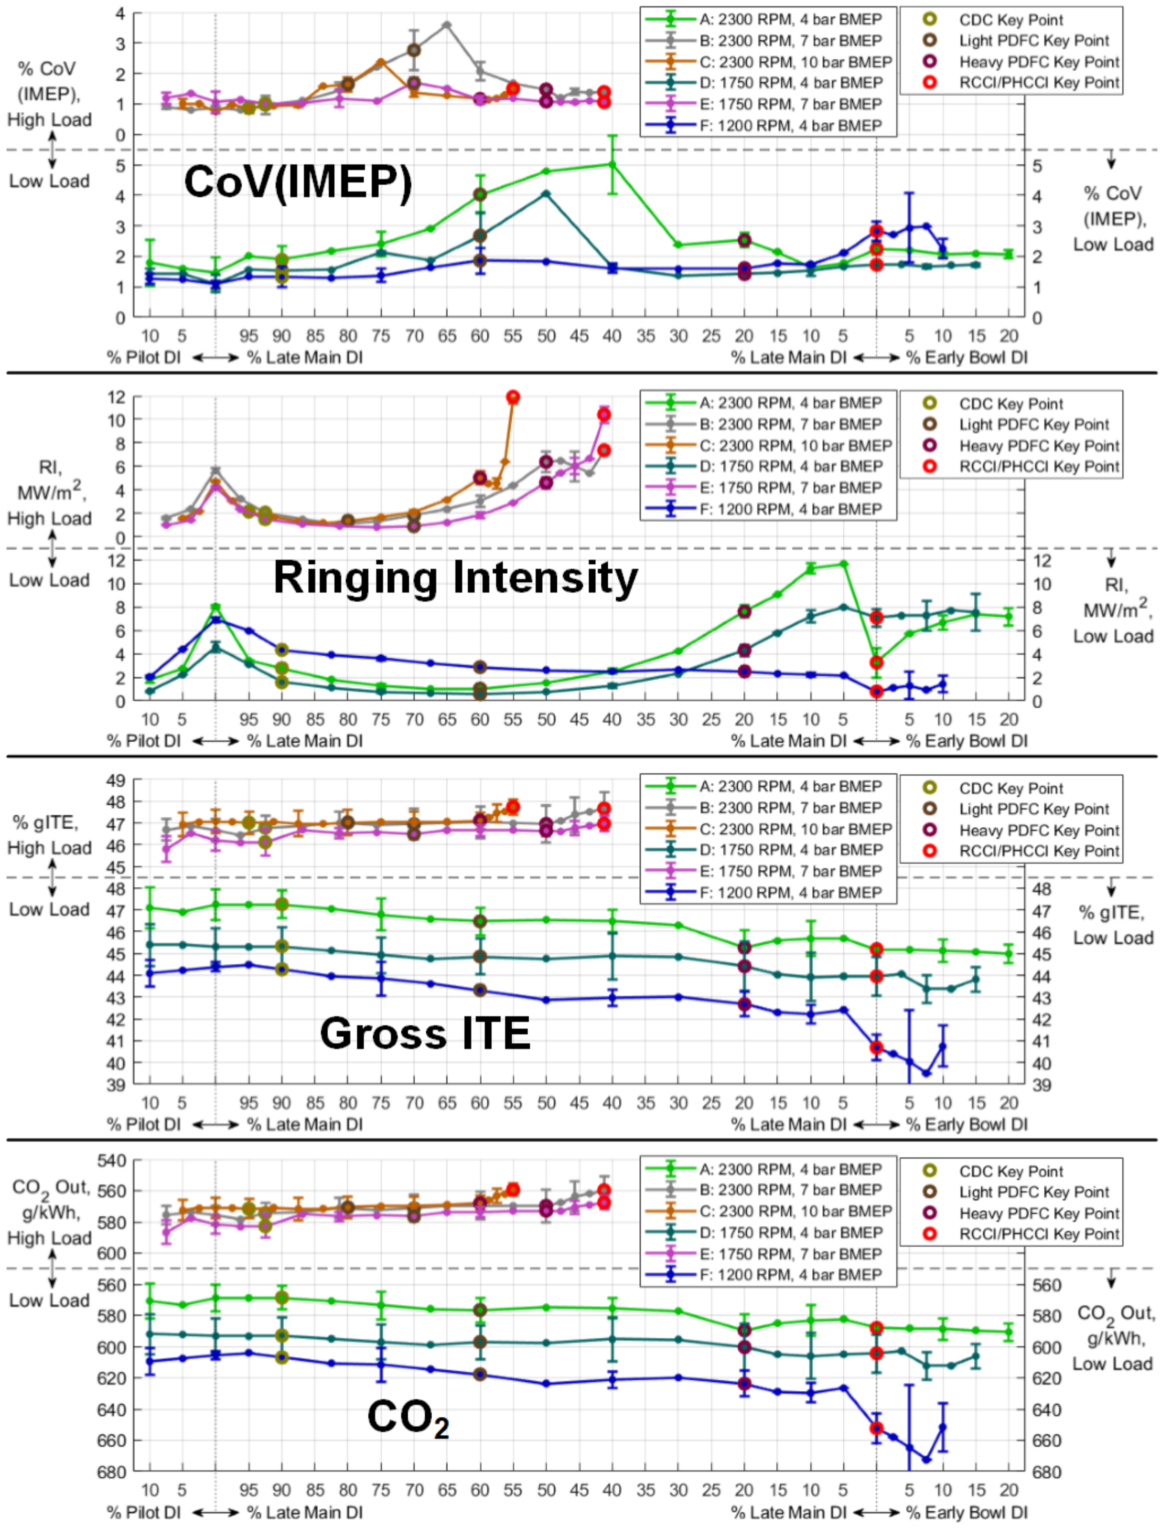

Figure 5.10 Limits and ITE/CO_2 with incremental MELT ratio shifts, discussed on pages 150-156. The exact MELT ratios used are given in Table 5.2 on page 151.

in the central bowl region. It is difficult to tell when analyzing the HRR from the Light PDFC conditions whether or not there is any homogeneous autoignition or not. The Heavy PDFC conditions tend to have a distinct “kink” the the lean premixed phase and a corresponding peak in the 3rd derivative that can be used to demarcate stratified and homogeneous autoignition. However, much more optical engine and CFD work would be needed to test this hypothesis, since the Wiebe function analysis is purely empirical. What is certain is that there is always a “hump” of high CoV that needs to be crossed when transitioning from CDC or Light PDFC to the more premixed modes. This hump appears to be smaller, however, at lower speeds and higher loads, so if such a transition was incorporated into a multi-mode engine strategy, it would be wise to make the transition at lower speeds and higher loads. The more advanced modes have slightly higher CoV’s, but never beyond 3%, much less than the 5% limit used herein.

As expected, all speed/load conditions have a peak of RI at 100% Late Main DI, with a small amount of Early DI or Pilot DI both greatly reducing the RI. It appears that only about half as much Pilot DI as Early DI is needed to produce the same drop in RI, which may indicate that the Early DI does not burn as completely as the Pilot DI, or that the heat released by the Early DI is not fully transferred to the bowl to reduce the ignition delay by the time the Late DI occurs. The RI goes down going toward Light PDFC, but back up again going toward Heavy PDFC in all speed/loads except F. This makes sense given the characteristic HRR shapes of these modes, with Light PDFC having a relatively low, flat HRR profile that does not cause sharp pressure spikes, but Heavy PDFC having a distinct spike corresponding to the homogeneous lean premixed autoignition. Although this spike is never as big as the spike from the rich premixed autoignition with 100% Late DI, it occurs earlier in the cycle and thus produces bigger spikes in pressure, causing Heavy PDFC to have about the same RI as CDC with no Pilot or Early DI.

The biggest changes in RI happen at the very end of the sweep, for both the low-load and high-load conditions. In the high-load conditions, moving from Heavy PDFC to PHCCI, there is a very sharp rise in RI at the very last point. This change is the greatest in speed/load C, where the Early DI is completely eliminated at this point, lesser in speed/load E where there is still 2.5% Early DI, and even lesser in B where there is 5% Early DI. This indicates that ringing may be a problem with "pure" PHCCI, when the premixed fuel is nearly perfectly homogeneously distributed. It may be necessary to include a tiny bit of Early DI in the mixture just to ensure some stratification in the premixed fuel and stagger the heat release to prevent ringing. In the low-load conditions, there is a sudden drop in the RI when the Late DI is completely eliminated, although this drop is much bigger in speed/load A than it is in D or F. The may be because A is the most stratified of these RCCI conditions, with the most Early DI (20%), which helps to stagger the heat release. Using a small Late DI in A causes the homogeneous lean premixed autoignition to occur more quickly and cause a much bigger PRR, and the small difference in the speed of the combustion process with respect to time is exaggerated in the high speed of 2300 RPM in A compared to the lower engine speeds of 1750 and 1200 RPM in D and F respectively. Interestingly, moving the Early DI to the bowl in speed/load A brings the RI back up, while it stays constant for D and F. It is also notable that there is much more uncertainty in the RI of the RCCI conditions than any of the other combustion modes.

When it comes to gross ITE, it appears that the high-load and low-load conditions are going in opposite directions. The high-load conditions feature an upward trend in ITE with more premixed fuel, and although this trend is only very slight going up toward Heavy PDFC, is gets steeper going to PHCCI, which raised the gross ITE by 1% TE over CDC in all three speed/load conditions. Unfortunately, the opposite is true in the low load conditions, with progressively worse gross ITE as more premixed fuel is used. Out of these three conditions, speed/load D appeared to perform the best, with the RCCI condition only being about 1% less efficient than CDC, as compared to 2% for speed/load A and 3.5% for speed/load F. Even when considering CO_2 emissions instead of gasoline, which is more favorable to the high-gasoline RCCI conditions, speed/load D still produces 10 g/kWh more CO_2 in

RCCI than CDC. However it is important to remember that these trends are not necessarily optimized for each speed/load condition, since the combustion phasing is shifting from point to point.

5.2.2 Combustion Phasing

The combustion phasing of these conditions is plotted below in Figure 5.11, in terms of the crank angle (CAXX) at which a certain percentage (XX%) of the total heat release is reached. Also shown is CA10-90, the difference between CA90 and CA10, which can be taken as a measure of the approximate duration of combustion. Of all the measurements analyzed in this chapter, these combustion phasing measurements were the most repeatable between tests, with most of the error bars too small to see, except for the RCCI conditions. This is not surprising given the nature of RCCI, where the combustion phasing is determined by the stratification of the fuel reactivity, which depends on random, turbulent charge motion.

The first plot in Figure 5.11 is CA10, which roughly indicates the onset of autoignition, either the rich premixed autoignition in CDC or the lean premixed autoignition in the advanced modes. The advancement of CA10 achieved by the Pilot DI is greater than the same size of Early DI in most cases, which combined with the trends in RI confirms that the Pilot DI is slightly more efficient at reducing ignition delay. Moving into the PDFC regions, CA10 is advanced at first slowly, but then more quickly. The first initial advancement is due to a reduction in ignition delay, which advances the CA10 as detected on the rich premixed autoignition spike. The transition to more rapid advancement of CA10 occurs when the stratified lean premixed autoignition phase becomes large enough to constitute 10% of the total heat release on its own. This transition appears to be absent from speed/load F, which indicates that the lean premixed autoignition never overtakes the rich premixed autoignition. This is acceptable for speed/load F, which is at the lowest engine speed and thus will have the kinetically-controlled lean autoignition process occurring the most quickly with respect to crank angle. The other, faster engine speed conditions need the lean autoignition process

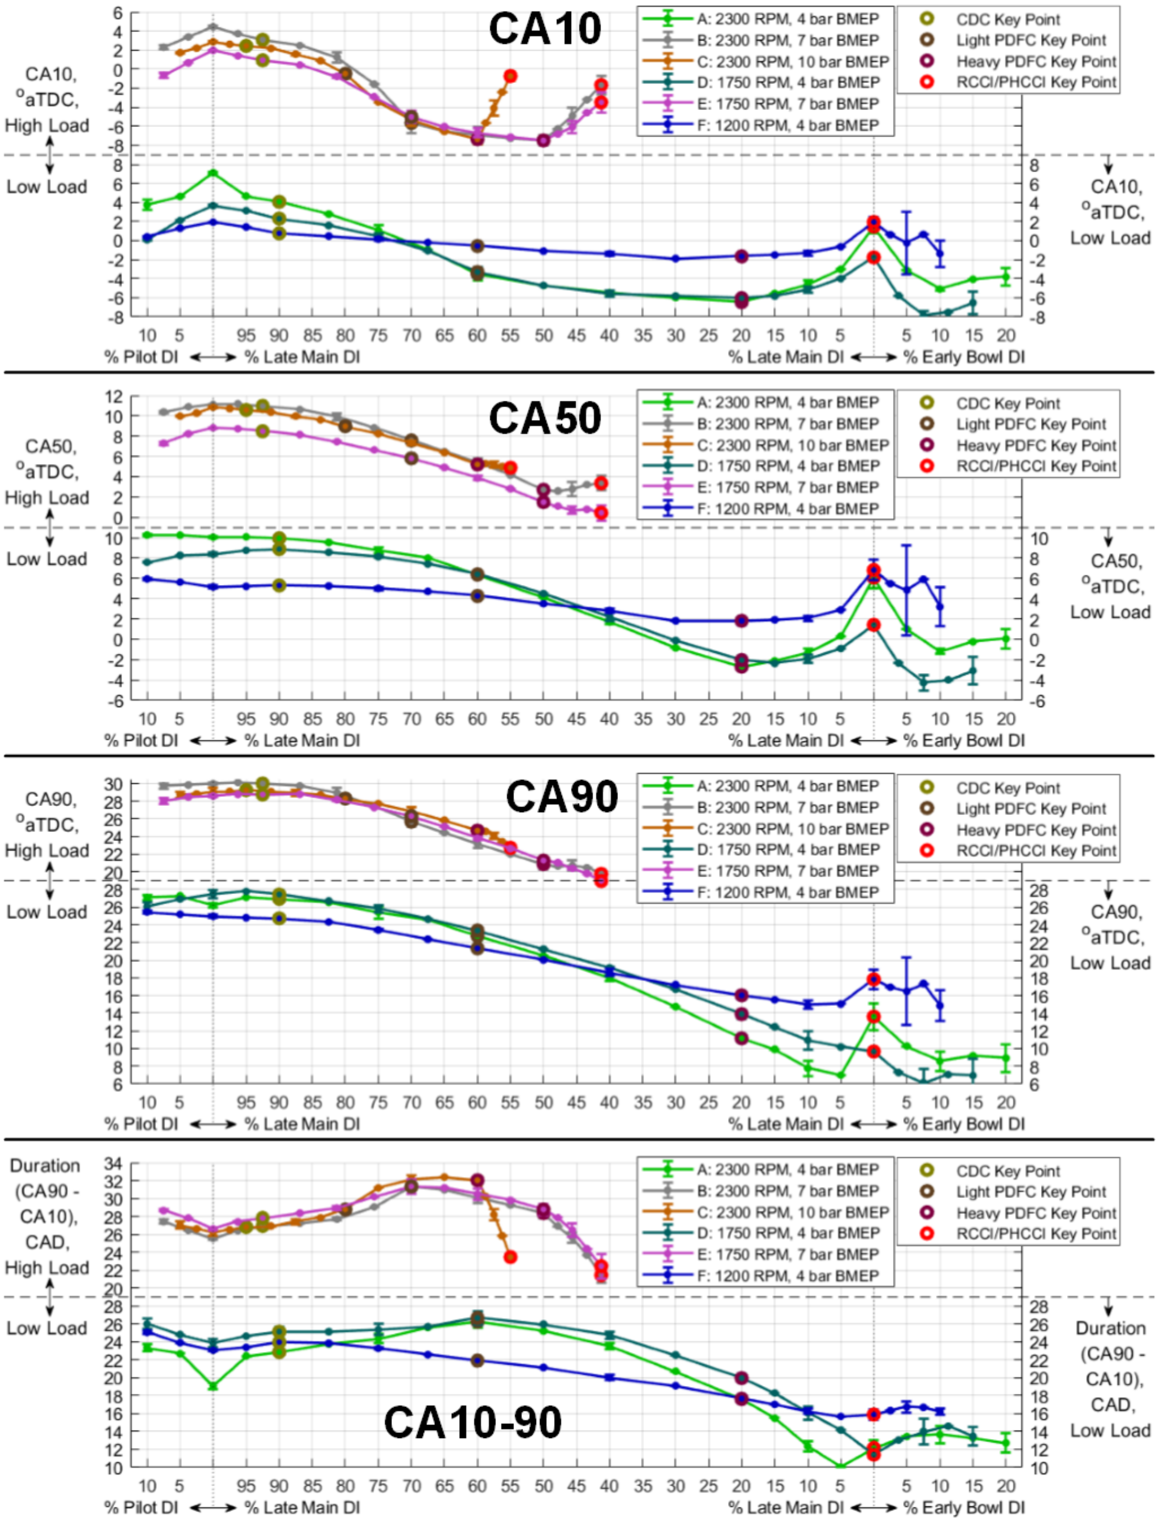

Figure 5.11 Combustion phasing with incremental MELT ratio shifts, discussed on pages 156-161. The exact MELT ratios used are given in Table 5.2 on page 151.

to start to ramp up before the rich premixed autoignition in order to get sufficiently complete combustion.

Once the lean premixed autoignition is substantial enough so that CA10 occurs on its leading edge, the advancement of CA10 slows down, but still slowly increases up to the Heavy PDFC key point. Here, the CA10 begins to retard again, at different rates in high load and low load. In the high load conditions, the path through the MELT diagram takes a sharp leftward turn toward PHCCI, with Early DI decreasing rapidly. This delays the lean premixed autoignition and increases the percentage of *homogeneous* lean premixed autoignition, with less of a slow initial ramp from the stratified lean premixed autoignition. In the low load conditions, there is a slighter leftward turn in the path through the MELT diagram to reach RCCI in speed/loads A and D, and none in speed/load F, which maintains 10% Early DI throughout. For each of these conditions, there is a slight retard in CA10 as the Late Main DI is slowly phased out, then a bigger retard at the final transition to RCCI when the Late Main DI is eliminated completely. The CA10 is advanced back to about the same place as Heavy PDFC when the Early DI is moved to the bowl, which makes sense since the Early Bowl DI occurs later in the cycle than the Early Main DI, but before the Turn-Around Point, and hence is advancing the combustion phasing. The retarding of the Early DI shortens the mixing period before autoignition, leading to equivalence ratios closer to stoichiometric at the time of autoignition.

The trend in CA50 mostly follows the trend in CA10, but much more smoothly. The advancement of CA50 appears to slowly and smoothly accelerate moving toward Heavy PDFC, as the early lean premixed autoignition begins to constitute a larger and larger percentage of the total heat release. The turn toward PHCCI in the high-load conditions, which quickly advanced CA10, does not appear to cause significant advancement in the CA50, since lean premixed autoignition never reaches 50% of the total heat release before the rich premixed autoignition occurs. Meanwhile in the low-load conditions, the CA50 is retarded more severely than the CA10 when moving to RCCI, again with a very large step change upon the final elimination of the Late Main DI. This further demonstrates how a very small Late Main DI can help to "stabilize" RCCI, and maintain the desired combustion phasing.

The trend in CA90 almost identical to the trend for CA50, except at the very end of the sweeps heading into the PHCCI and RCCI modes. Going toward PHCCI in the high-load conditions, the CA90 continues to advance, in contrast to the CA50 which stays relatively static and the CA10 which retards quickly. This can be attributed to the continuing reduction of diffusion combustion by displacing the Late DI fuel with premixed fuel. The diffusion phase is typically the slowest combustion process and governs the location of CA90, except in the low-load RCCI conditions. When the transition to RCCI occurs in speed/loads F and especially A, there is a sudden retardation of CA90, which then becomes dependent on the slow, highly variable autoignition of premixed fuel in the cold spots of the combustion chamber, which is no longer consumed by a diffusion flame. The transfer of Early DI fuel to the Early Bowl DI helps to reverse this by advancing the overall combustion phasing and increasing the temperatures during the cold spot autoignition phase.

Strangely, this jump in CA90 does not occur for speed/load D, which sees a slight advance in CA90 when the Late Main DI is eliminated completely, for which the cause cannot be determined. One possibility is that the speed of 1750 RPM is more ideal for RCCI at 4 bar BMEP than the other two engine speeds. 1200 RPM could simply be too slow of an engine speed for RCCI for the reason of increased heat transfer losses that cause peak temperatures/pressures to be too low for complete autoignition (as will be seen later with the very high UHC at this speed). When the Late DI is present to clean up the UHC and accelerate the lean autoignition this is not as much of a problem. Meanwhile 2300 RPM could simply be too *fast* of an engine speed for RCCI to achieve complete autoignition, with the piston descending and allowing expansion to reduce temperature and pressure too rapidly with respect to time. Of course it is also possible that the MELT ratio used in speed/load D was simply more optimal than those used in speed/loads and A and F, since there were only rough estimates of the best MELT ratios to use and no further optimization. This will be investigated further in Section 5.3 when multiple Early DI sizes are used for RCCI.

Although the trends described above are important for studying how the nature of the combustion phasing changes in the different modes, it is important to remember that the overall phasing can be adjusted to optimize each mode by adjusting the SOI timings. Thus, the trends in CA10, CA50, and CA90 are not necessarily inherent to those combustion modes, and all of the combustion modes can be adjusted to equalize one of these values. What is less adjustable for each combustion mode is the *duration* of combustion, represented here by CA10-90. This reflects how efficiently each combustion mode can maximize the area enclosed by the P-V curve, by focusing as much of the heat release and pressure rise near TDC as possible. The CA10-90 metric is not a perfect evaluation of this obviously, but gives a quick and easy indication of how well the heat release has been "compressed" into a small window.

For most speed/load conditions, the trend in combustion duration is upward going from CDC to Light PDFC. This is indicative of how Light PDFC generates highly "spread out" combustion phasing, with the stratified lean premixed autoignition occurring well before the mixing-controlled Late DI process. However, as more premixed fuel is added going toward Heavy PDFC, the lean autoignition does not advance so much as it gets larger, while the displacement of Late DI fuel reduces the amount of diffusion that needs to take place. Both of these processes thus "compress" their heat release into a smaller window. For the high load conditions, the Heavy PDFC duration is still greater than CDC, because there is still at least 50% of the fuel in the Late DI and thus a long diffusion process, but for the low load conditions, with only 20% Late DI, the Heavy PDFC duration is much shorter than CDC.

Going from Heavy PDFC to RCCI, the low-load combustion duration reaches its lowest level, showing how effectively RCCI can concentrate the combustion phasing into a small window when there is no slowly tapering diffusion process. However, the simultaneous retardation of the overall combustion phasing, combined with the sharp increase in UHC due to the lack of a diffusion flame leads to the RCCI conditions having a lower gross ITE. Meanwhile, the high-load combustion duration drops very sharply going from Heavy PDFC to PHCCI, since the lean premixed autoignition is retarded and compressed closer to the rich premixed autoignition as previously discussed. This sharp drop in combustion duration corresponds with little to no change in the CA50, and these two trends combined explain the previously observed rise in gross ITE. Unlike RCCI, the heat release is being compressed *without* shifting backward, resulting in maximal area under the P-V curve and optimal gross ITE. However, to fully understand the changes in ITE requires an examination of not only the combustion phasing but also the entire energy balance of the engine.

5.2.3 Energy Balance

Figure 5.12 shows the three major heat losses as described in Section 4.1.3: cooling loss, exhaust loss, and unburnt loss. This accounts for all of the difference between the gross ITE and the heat released by the fuel, except for the small (1-2%) losses in the turbocharger and charge air cooler, which did not change significantly enough to affect the ITE trend. Also shown is the sum of the cooling and exhaust loss, since these two losses will trade-off with each other as a result of combustion phasing. It is important to consider these losses in tandem with the combustion phasing, to determine whether the changes are merely the result of a change in combustion phasing or a change in the rate of cylinder heat transfer.

All of the high-load conditions show a very flat trend in the cooling loss up until the key point of Heavy PDFC, with a slight drop in cooling loss going down to PHCCI. This is a significant finding when one considers that combustion phasing, in terms of CA50, is constantly advancing. All else held equal, and advancement of combustion phasing should lead to *more* cooling loss since there is more time between combustion and the exhaust valve opening for heat transfer, so this trend seems counter-intuitive at first glance. The downward bend when turning toward PHCCI can be explained by the dramatic shortening of combustion duration, with less combustion happening before TDC for long heat transfer

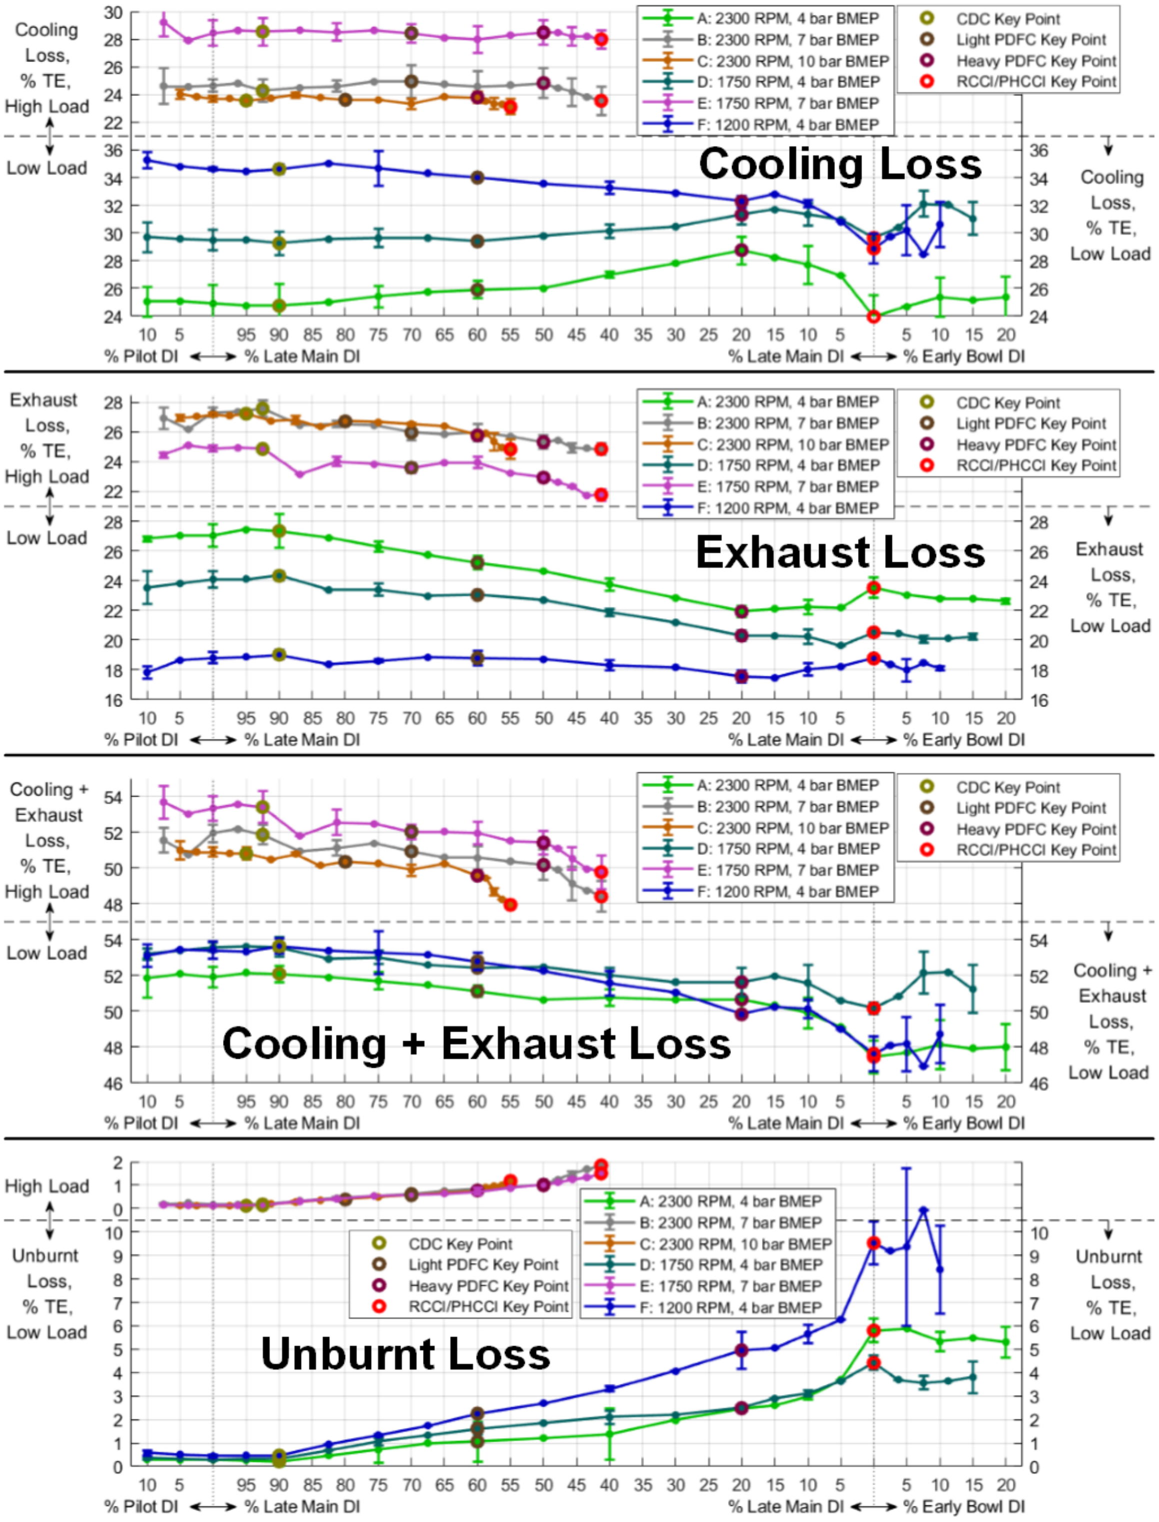

Figure 5.12 Energy balance with incremental MELT ratio shifts, discussed on pages 161-164. The exact MELT ratios used are given in Table 5.2 on page 151.

times. However the PDFC conditions that precede this turn all have longer combustion duration than CDC, as well as advanced CA10, CA50, *and* CA90, all of which should mean rising cooling loss.

There are three possible explanations for this counterintuitive trend: The first possibility, which is one of the purported benefits of lean premixed combustion, is that the cylinder heat transfer rate has been reduced. With more and more lean premixed autoignition, the temperature rise from combustion is spread out more evenly throughout the combustion chamber, and the turbulent gas velocity is relatively low, which both reduce the heat transfer rate to the combustion chamber walls. The second possibility is that the increase in unburnt loss makes up for the advancement of combustion phasing. With less total heat released, the overall cylinder temperatures will be lower, leading to less cylinder heat transfer. The third possibility stems from the fact that the cylinder heat transfer component of cooling loss is not directly measured and is instead being calculated by difference (see Equation 4.22). This means that any errors in the measurement of the other heat flows will be compounded in the final calculation of cylinder heat transfer. It is therefore very possible that the real trend in cooling loss is actually upward, but there is an error that caused one of the other heat losses to be over-estimated. In particular, the unburnt loss seems like it could be such a source of error, since it goes up with more premixed fuel and an upward bias in the unburnt loss would mean progressively more downward bias in the cooling loss with more premixed fuel. The evidence presented here is not conclusive as to which of these three possibilities is the source of the observed trend in cooling loss at high load.

The cooling loss trend at low load varies depending on the engine speed. At the highest speed of 2300 RPM (speed/load A), there is increasing cooling loss going from CDC to Heavy PDFC, which is expected given the advancing combustion phasing. At the medium speed of 1750 RPM (speed/load D) the trend is the same, but more slight, which is consistent with the fact that the combustion phasing for this condition does not advance as much as it did in speed/load A. But at the low speed of 1200 RPM (speed/load F), there is again an unexpected downward trend in cooling loss despite the advancing combustion phasing. In this case, unlike the high load conditions, there is much less ambiguity to the source of the dropping cooling loss, because this speed/load features very high unburnt losses, about double of those seen in the other two low-load conditions. By the time the Heavy PDFC point is reached, the unburnt losses have increased by about 5% TE, while the cooling loss has only dropped by 2% TE. This further highlights how speed/load F is simply too slow-speed and low-load to get decent combustion efficiency with highly premixed combustion modes. In all three low-load conditions, there is a large drop in cooling loss going to RCCI, but this is very expected given the simultaneous retarding of the combustion phasing when the Late DI is fully eliminated, as well as the sharp increase in unburnt loss.

The trends in exhaust loss are all much more intuitive, matching up well with the trends in CA50, with exhaust loss decreasing in tandem with advancing phasing. However, at the high loads, there is none of the expected trade-off with cooling loss, and as a result the combination of cooling + exhaust loss ends up trending downward with more premixed fuel, especially when going from Heavy PDFC to PHCCI. Although there is still a downward trend at low load this trend is less steep, and is outweighed by the steeper upward trend in unburnt loss. However it is important to remember that unburnt loss is calculated under the assumptions that all of the unburnt products have been measured by the exhaust gas analyzer, and that their total calorific value is equal to the original fuel molecules on a per-mole carbon basis. Since the cooling loss is calculated by a difference which includes this unburnt loss, any over- or under-estimation of the unburnt loss will lead to an equal but opposite error in the cooling loss. Either way, the unburnt loss is clearly much less of a problem for high load, where the overall temperatures and pressures are higher. However the next section will show that those higher temperatures and pressures have a drawback when it comes to the other, non-UHC emissions.

5.2.4 Emissions

In the higher loads, the overall higher temperatures and pressures appear to lead inevitably to higher NO_X emissions, which only increase as more premixed fuel is added and the combustion phasing is advanced. Despite displacing Late DI fuel and thus reducing the amount of NO_X-producing diffusion combustion, the premixed fuel is autoigniting before the diffusion phase occurs. This not only increases the temperatures of the diffusion flame, but reduces the ignition delay, thereby increasing the amount of Late DI fuel that will burn via diffusion. It is also possible that, since the lean premixed autoignition is happening at higher temperatures and richer equivalence ratios, this phase is beginning to generate some NO_X itself. This is backed up by the appearance of significant NO_X emissions from RCCI when the Early Bowl DI is used. Although the Early Bowl DI is occurring early enough to mix completely below stoichiometric (indicated by the fact that combustion phasing is advanced when this injection is retarded), it does not get lean enough to prevent NO_X emissions, while the Early Main DI *does* get lean enough. Clearly, the Early Main DI is preferable for this shape of piston, but this may be purely the result of the injector and piston design, and a design that allows for an earlier injection to enter the bowl may not have these same problems with NO_X. Regardless, the low load conditions clearly give more of a NO_X benefit than high load, with NO_X emissions going slowly down despite advancing combustion phasing, then very suddenly dropping to zero when the Late DI (and its corresponding diffusion flame) are eliminated.

The soot emissions in CDC appear to be greatly increased by both the Early DI and the Pilot DI, with much less soot being produced by 100% Late DI, despite a slight increase in NO_X emissions. As expected, the increasing premixed fuel at first increases soot emissions, with the lean premixed autoignition both decreasing ignition delay to increase soot production and consuming oxygen to inhibit soot oxidation. However, once a "tipping point" of about 20-25% premixed fuel is crossed, the soot emissions start to go back down, and in all speed/load conditions the Heavy PDFC key point produces less soot than even the 100% Late DI CDC point. By this point, the aforementioned soot-increasing effects are outweighed by the reduction in the amount of Late DI fuel. The reduction of Late DI fuel reduces the both the amount of soot produced and the amount of time and oxygen required to completely oxidize it before the temperature drops too low. The improvement in soot gets even better in both PHCCI and RCCI, when the amount of premixed fuel is maximized. For the low loads this is complemented by the simultaneous reduction in NO_X, but at the price of a sharp increase in UHC and CO emissions.

On the subject of CO emissions, there is a distinct "hump" in the trend of CO that is not apparent in the UHC trend, with CO peaking in Light PDFC before going back down in Heavy PDFC. This can be attributed to the fact that there is a very lean mixture of premixed fuel when it is first introduced, especially in the areas where the fuel is homogeneously distributed. When incomplete combustion occurs in an overly lean mixture it will tend to generate relatively more CO than UHC. These CO humps appear to correspond with the previously observed trend in CoV of IMEP, with the CO increases slowing down when high CoV starts appearing. This supports the hypothesis that the Light PDFC conditions are right on the border of generating a homogeneous lean premixed autoignition, getting inconsistent autoignition from cycle to cycle, with both the CO (from unburnt homogeneous mixture) and CoV dropping going toward Heavy PDFC, where the homogeneous autoignition is triggered much more consistently.

These trends in emissions are still, however, not necessarily inherent to each combustion mode. To see the practical emissions reductions that each combustion mode can achieve, there must be some allowance for flexibility in the combustion phasing via the EGR and SOI.

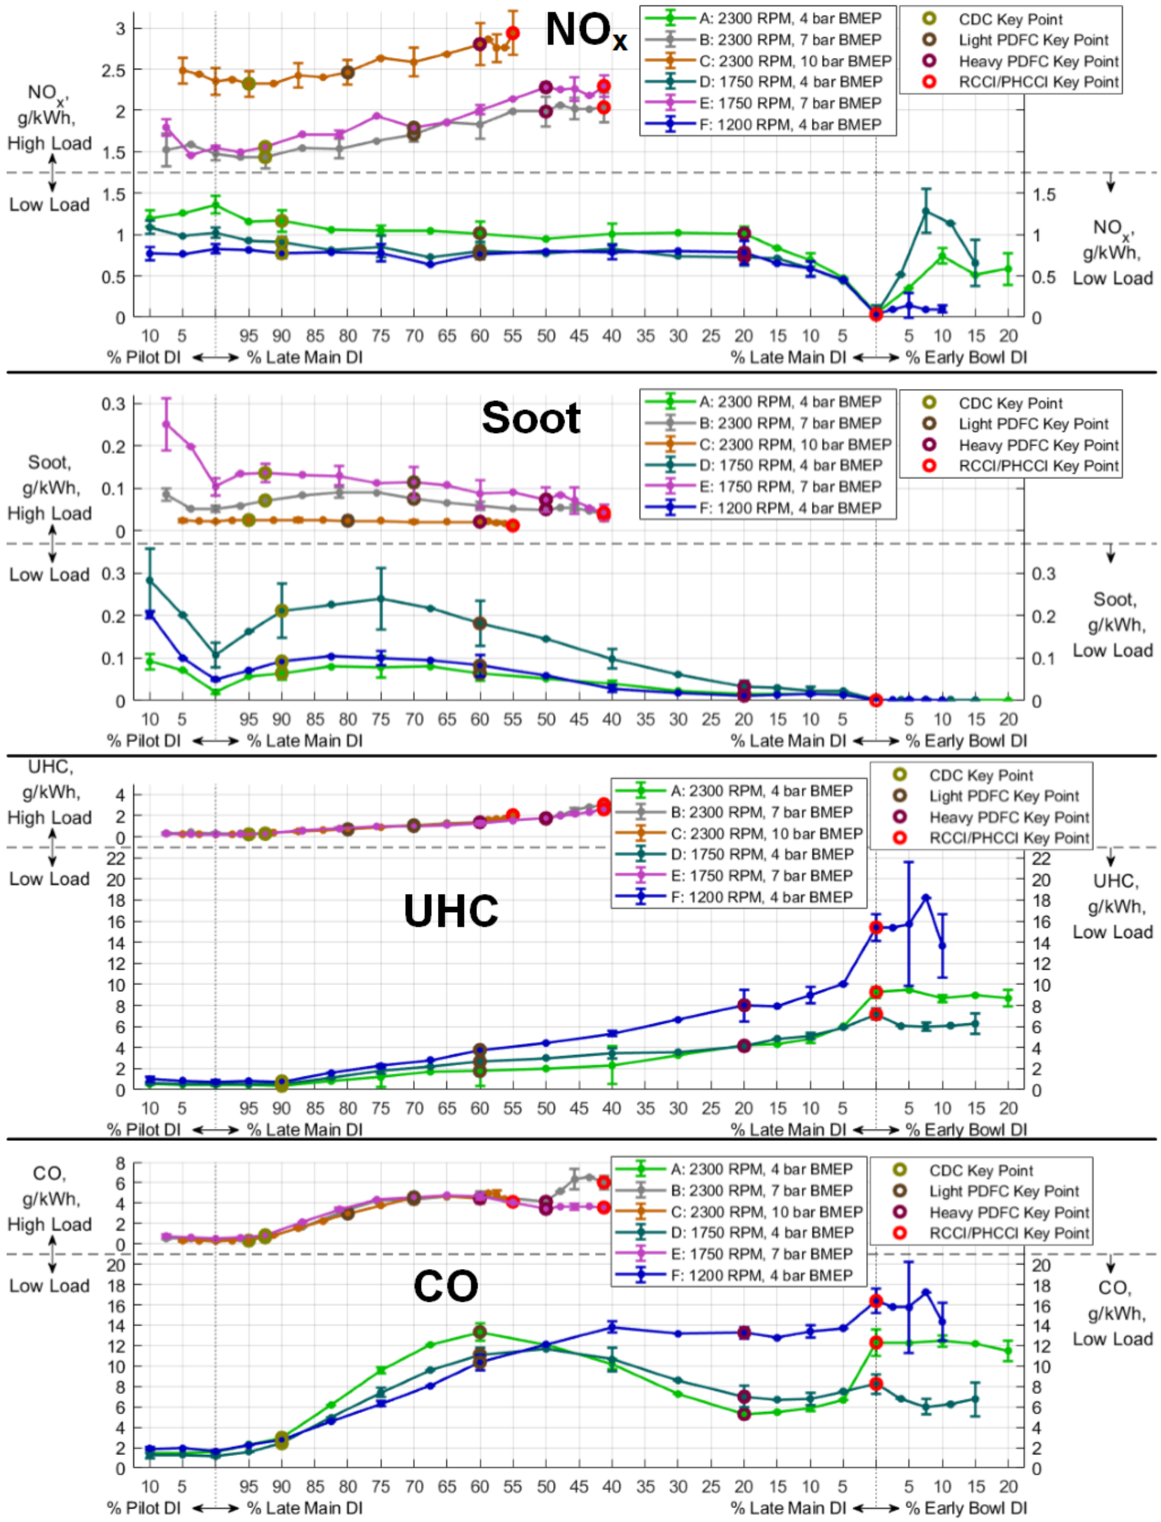

Figure 5.13 Non-CO_2 emissions with incremental MELT ratio shifts, discussed on pages 165-166. The exact MELT ratios used are given in Table 5.2 on page 151.

5.3 EGR and SOI Grids at Key MELT Ratios

The previous section presented the characteristic changes in performance when moving incrementally across different MELT ratios, with a fixed level of EGR and Late DI SOI timing. This section will instead fix the MELT ratio at the four key points, and vary the EGR and SOI to optimize the performance at each of these key points. Although there are many metrics to use to optimize the performance, this section will focus on three key metrics: ITE, NO_X, and soot emissions. As stated in the introduction, the trade-off between these three values is one of the key limiting factors in the performance of CI engines, and it is important to understand how this trade-off changes with the advanced dual-fuel modes. These three data will be plotted by using NO_X as a common x-axis in the center of the figures, with gross ITE plotted on a y-axis above, and soot plotted on a second y-axis below the x-axis. In conventional CI engines, ITE tends to increase as NO_X emissions increase, while soot emissions tend to decrease, and these plots will show both of those trade-offs at once.

5.3.1 Comparing EGR Levels

The first data presented will be a complete look at the whole data set. In the following figures 5.14, 5.15, and 5.16, each plot will contain the data from one key point, with the four key points from one speed/load condition laid out in a row. In each plot there will be three sets of three data points, corresponding to the three SOI sweeps conducted at different EGR levels.

Starting in speed/load A at the top of Figure 5.14, the typical trade-offs for a conventional CI engine are seen in the CDC plot on the left (dark yellow). As the SOI is retarded, the NO_X is reduced, while soot is increased. Gross ITE reaches a peak when a certain SOI is used, typically on the higher end of the NO_X scale. With more EGR, the soot is increased for every SOI timing, but the NO_X is reduced, and to a lesser degree the ITE. These trends stay mostly the same in Light PDFC, with the exception that soot becomes much more sensitive

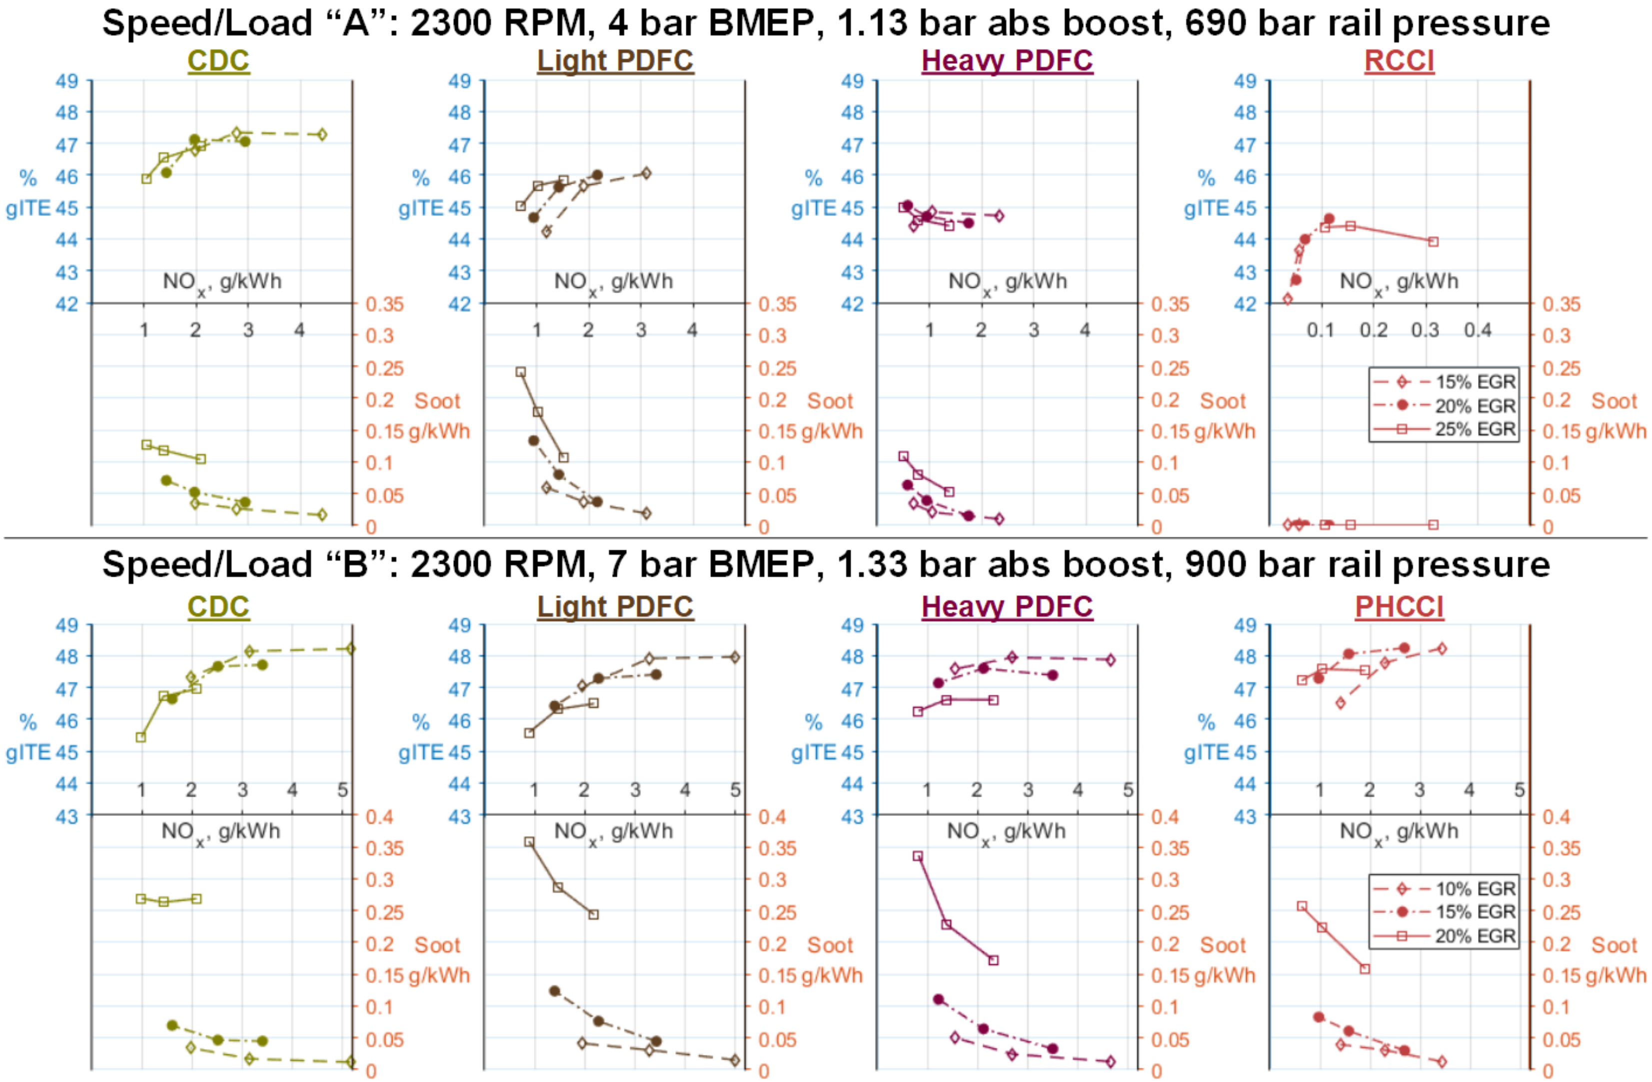

Figure 5.14 Gross ITE, NO_X, and soot at key points for speed/loads A and B, discussed on pages 168-172.

to SOI. This can be attributed to the fact that earlier SOI's in PDFC will cause the Late DI to occur when less of the lean premixed autoignition heat release has occurred, avoiding its soot-increasing effects but raising NO_X emissions. There is also an overall decrease in NO_X emissions, but this comes with an decrease in ITE as well.

The trends in ITE change substantially with Heavy PDFC. Although there is still an ITE peak at the middle SOI with the 15% EGR sweep, the other two higher EGR levels lead to increasing ITE *and* decreasing NO_X, a "win-win" situation that is almost never seen with conventional diesel engines. This indicates that the phasing of the lean premixed autoignition has been pushed so early that retarding the SOI actually results in better utilization of the P-V curve, a phenomenon discussed previously in Section 5.1.2. The soot/NO_X trade-off remains the same as Light PDFC but with the overall soot drastically decreased and the NO_X slightly decreased. In RCCI, on the other hand, the NO_X is so low that the scale must be increased by a factor of 10, and the soot levels are all below the detection limit of the smoke meter. In this mode, recall that there is no Late DI SOI to vary, so the sweeps shown here are sweeps of the Early DI/Manifold fuel ratio, with more Early DI increasing the NO_X by advancing combustion phasing.

When moving to the higher load of speed/load B at the bottom of Figure 5.14, some of these trends stay the same, while others do not. In CDC, the ITE becomes more sensitive to EGR, and the penalty to ITE for reducing NO_X becomes steeper. The difference in soot levels between the second and third EGR levels also becomes quite stark. All of these same trends apply to Light PDFC and Heavy PDFC as well, but just as in speed/load A, the sensitivity of soot to SOI becomes much greater. The ITE is also progressively becoming more sensitive to EGR than SOI, as the Late DI has less influence on the combustion phasing, but not to the point where the ITE trend reverses direction. This indicates that, unlike speed/load A, in speed/load B the chosen level of Early DI for Heavy PDFC is not bringing the lean premixed autoignition so early that a later SOI gives better P-V curve utilization. There is also no perceptible drop in ITE across the first three combustion modes

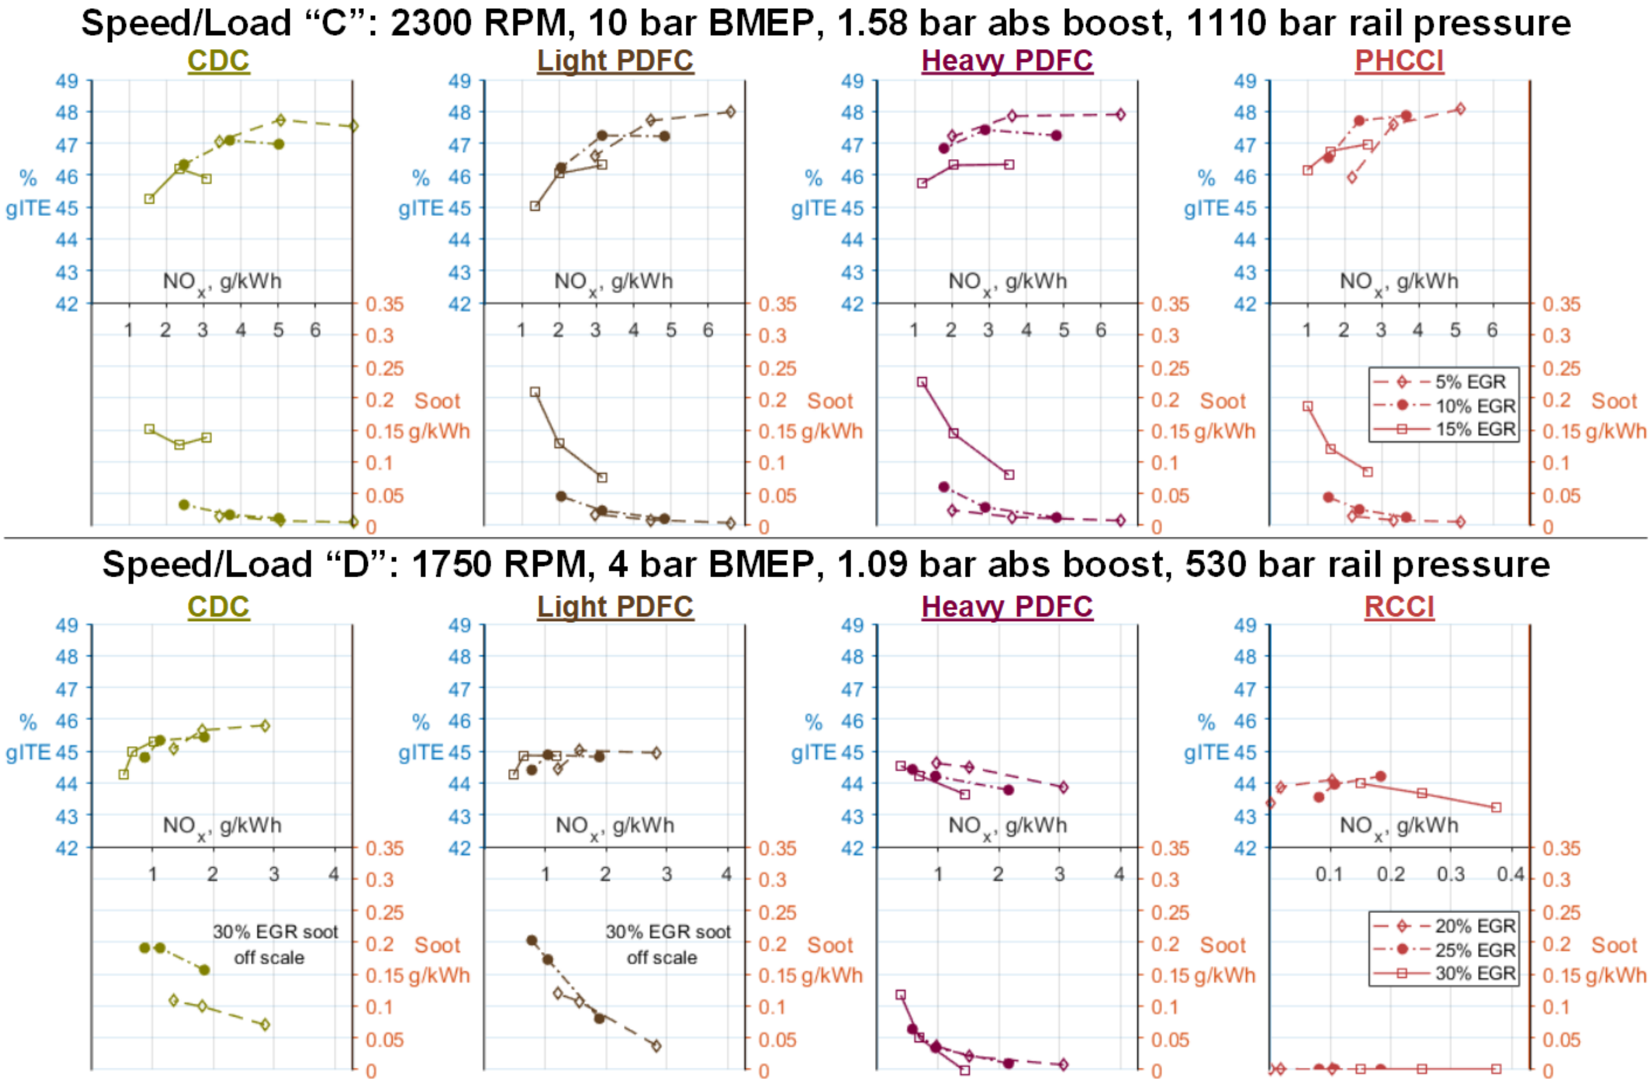

Figure 5.15 Gross ITE, NO_X, and soot at key points for speed/loads C and D, discussed on page 172.

as there was in speed/load A. Finally, in PHCCI, the ITE produced by the higher EGR levels is brought up to nearly the level found in the lowest EGR levels. As a result, *all* of the PHCCI conditions are generating lower NO_X, higher ITE, and nearly identical soot compared to the CDC conditions. The was not seen in the previous section because the SOI had been held constant, but in the targeted EGR/SOI windows, the ideal SOI range for PHCCI is retarded 6 degrees. This suggests that PHCCI may be a "win-win" for speed/load B, whereas in speed/load A there were no clear winners, only trade-offs. These results stay largely the same in speed/load C at the top of Figure 5.15, where again PHCCI simultaneously improves NO_X and ITE while soot stays about the same, albeit with much more sensitivity to SOI than CDC.

Going back to a low-load condition in speed/load D at the bottom of Figure 5.15, the trends become more similar to speed/load A, with the ITE/NO_X trend in Heavy PDFC again reversing. The main difference in speed/load D is that the highest EGR level chosen appears to be too high for reasonable soot emissions in CDC and Light PDFC, pushing these soot emissions beyond the high end of the same scale that was used for the first three speed/load conditions. However, once Heavy PDFC is reached, the soot from this highest EGR range drops precipitously, down to the levels of soot that were seen with the *lowest* EGR level in the first two modes. RCCI again drops NO_X by an order of magnitude, but with the ITE's significantly lower.

Going back up to a higher load in speed/load E, seen at the top of Figure 5.16, the trends are largely the same as in speed/loads B and C, with PHCCI clearly the optimal mode as far as ITE, NO_X, and soot are concerned. Unfortunately, much like speed/load D, the highest EGR level appears to be a bit too high, putting soot off of the high end of the scale for *all four* modes, not just CDC and Light PDFC as was the case in speed/load D. This is because the Late DI cannot be shortened enough to eliminate the soot production like it could in the lower load of D. The drop in ITE at this highest EGR level is also more pronounced than it was in B and C, indicating that the 20% EGR level is simply too dilute to get efficient

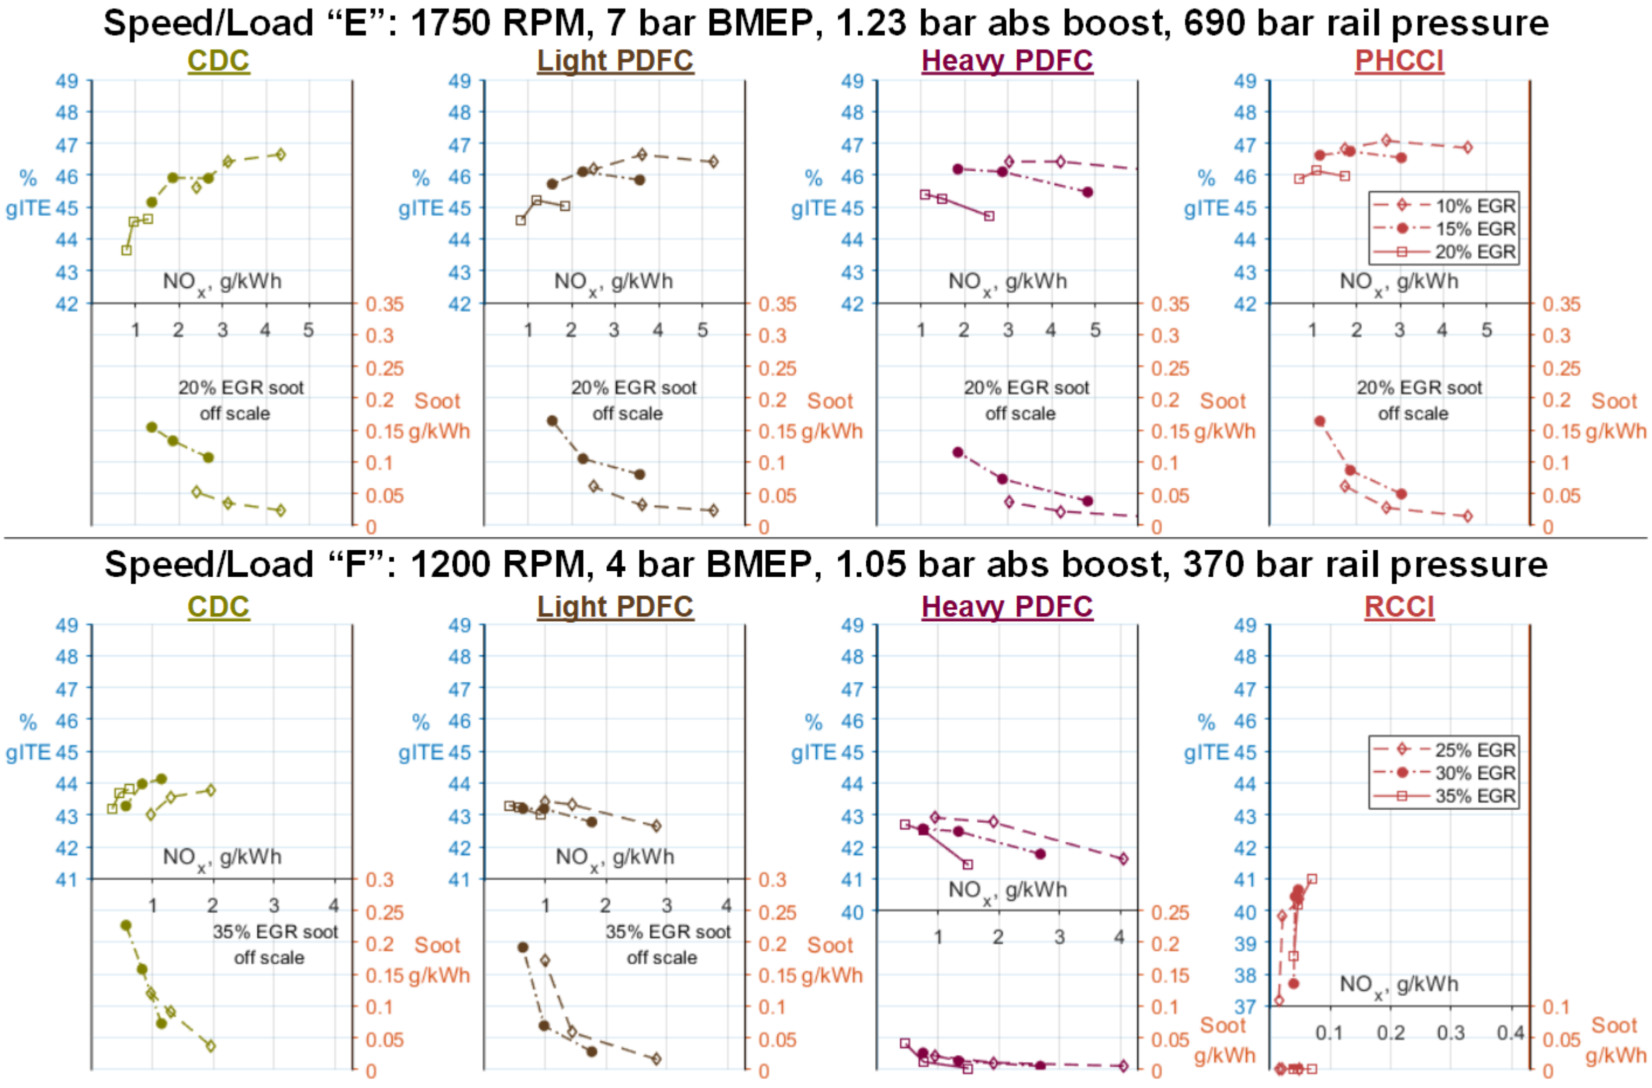

Figure 5.16 Gross ITE, NO_X, and soot at key points for speed/loads E and F, discussed on pages 172-174.

combustion at this speed/load combination.

Finally, in speed/load F at the botom of Figure 5.16, there are similar trends as the other two low-load conditions, with a few exceptions. Firstly, unlike speed/loads A and D, the ITE of CDC is slightly improved with greater EGR, although as the next section will show this improvement is small compared to the uncertainty bounds. The high sensitivity of soot to SOI begins in CDC rather than Light PDFC, which may be due to the higher overall EGR levels, lower engine speed, or reduced injection pressure, all of which inhibit oxidation of soot. Again the highest EGR level produces soot off-scale for CDC and Light PDFC, but suddenly drops down in Heavy PDFC before being completely eliminated in RCCI. The NO_X levels of RCCI are extremely low, almost two orders of magnitude below the other modes, but this only comes at the expense of very poor ITE. The ITE of speed/load F in general is relatively low, and it may be that the engine speed is simply too slow for this small engine size, which gives the combustion chamber a high surface-to-volume ratio and thus high cooling losses compared to larger engines that can run efficiently at low speeds.

Although the plots in this section are useful for observing performances changes at different EGR levels, it is hard to see the differences between modes when they are on separate axes. The following section will select one EGR level for each combustion mode, and plot multiple combustion modes on the same axis for comparison.

5.3.2 Comparing MELT Ratios

To start, a comparison of all four combustion modes at the same EGR level is shown in Figure 5.17. This illustrates the two distinct trends seen at low load (speed/loads A, D, and F) and at the higher loads (speed/loads B, C, and E). At low load, Light PDFC reduces ITE while giving only slight improvements (if any) in soot and NO_X. Heavy PDFC at low load greatly reduces soot across the board, and while it reduces ITE further at the advanced SOI's, it can improve ITE and reach lower NO_X levels at the retarded SOI's, except in speed/load F. RCCI almost completely eliminates NO_X and soot, but at the expense of ITE,

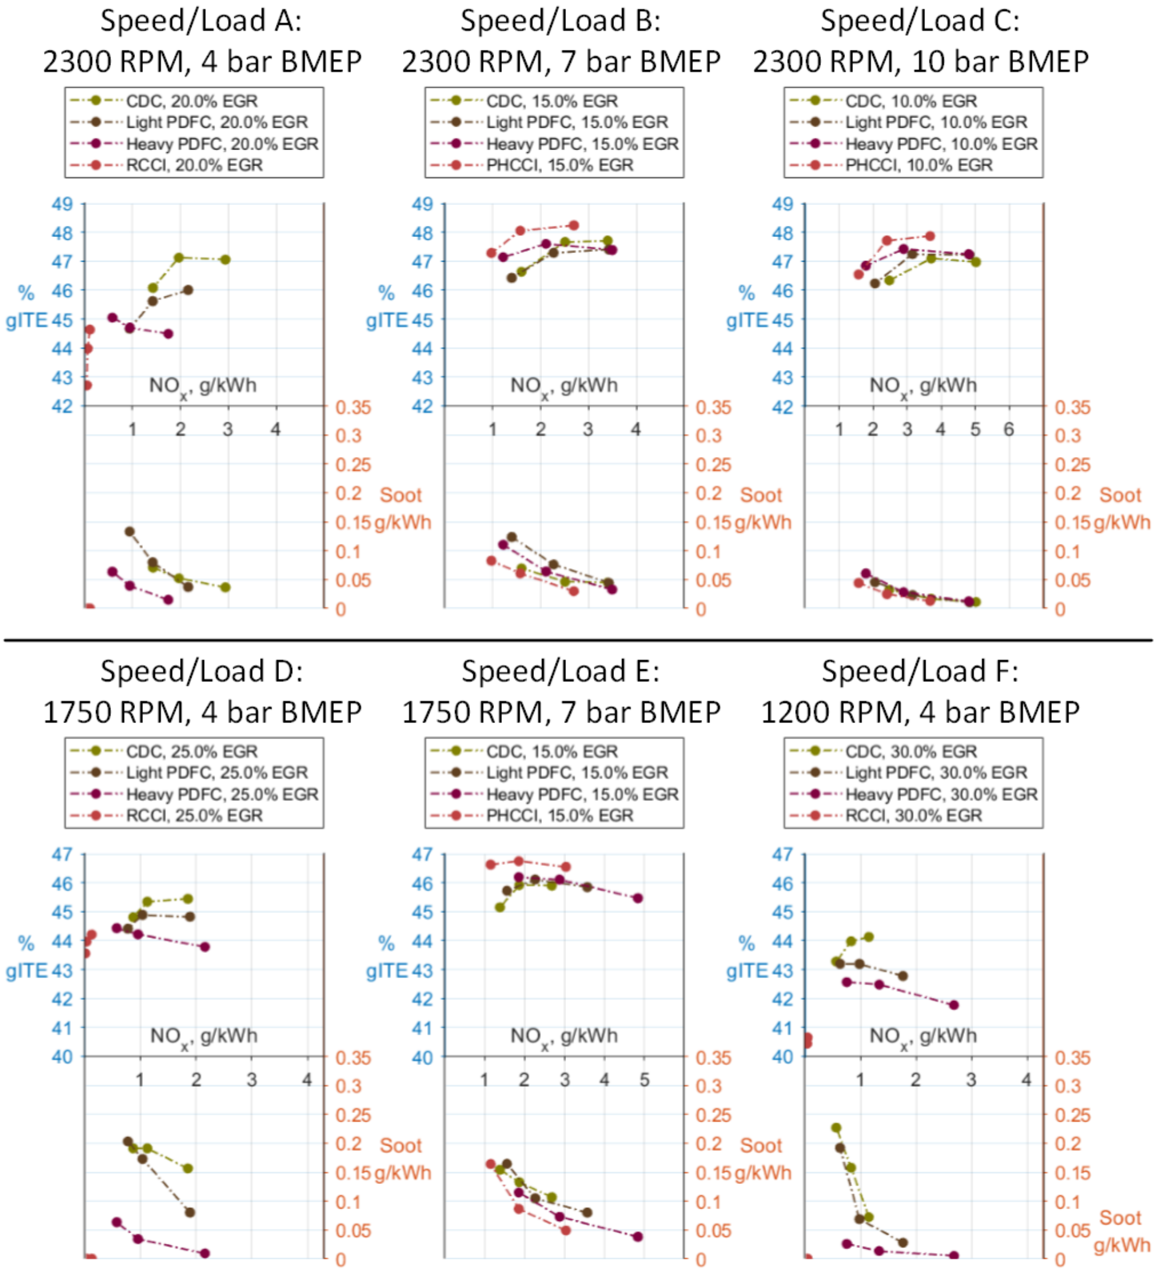

Figure 5.17 Gross ITE, NO_X, and soot at key points with middle EGR level, discussed on pages 174-175.

especially in speed/load F. Meanwhile, at the high loads, the four modes are clustered more tightly together, but it is clear that PHCCI is performing somewhat better than the rest, with improved efficiency, lower NO_X, and similar or reduced soot in every condition.

Since this middle EGR level is the level where all three SOI's were tested repeatedly, the uncertainty bounds can be plotted to show just how significant these trends are. Since such a plot would be difficult to read if error bars were plotted for all four modes, they will only be shown for two modes at each speed/load. These two modes will be the baseline CDC mode and a subjectively selected “best” advanced dual-fuel mode. For high load, it is easy to see that PHCCI gives the best performance of the three dual-fuel modes, but for low load it is debatable whether Heavy PDFC or RCCI is giving the best performance. For these plots, Heavy PDFC is chosen for low load simply because the RCCI points are too clustered together at the end of the scale. The result is shown in Figure 5.18.

At low load, the error bars show that there is very high certainty in the NO_X and soot reductions produced by Heavy PDFC, with the one exception of the highest-sooting condition in speed/load A, which saw very high trial-to-trial variation. This indicates that this condition is right on a “tipping point” of exponential soot increase. The 95% confidence intervals of ITE do not quite clear each other in every case, there is moderate certainty that Heavy PDFC has to significantly reduce ITE in order to achieve these emissions improvements. But at high load, there is moderate certainty that PHCCI is producing ITE improvements in all three speed/loads, as well as NO_X reductions in speed/loads B and C. There may also be some soot improvements in these high loads, but this is low certainty given the many overlapping error bars. These same high-load improvements were seen to a lesser degree with PDFC, but in those cases the difference was slight enough that the error bars were consistently overlapping, and as such there is low certainty in those improvements.

The ITE improvements from PDFC have an unfortunate side effect of raising the RI in some cases, something that is much less of an issue with Heavy PDFC. This can be seen in Figure 5.19, where the x-axis has been changed to CA50 instead of NO_X to show the relationship of both ITE and RI to combustion phasing. In speed/load B, the RI of CDC and PHCCI are quite similar, but for speed/loads C and E, the RI of PHCCI is much higher, especially in the highest load of C. However, as seen previously in Figure 5.10, RI from the key point PHCCI conditions can be reduced greatly with just a slight reduction in the amount of premixed fuel and only a slight reduction in the other benefits. Clearly,

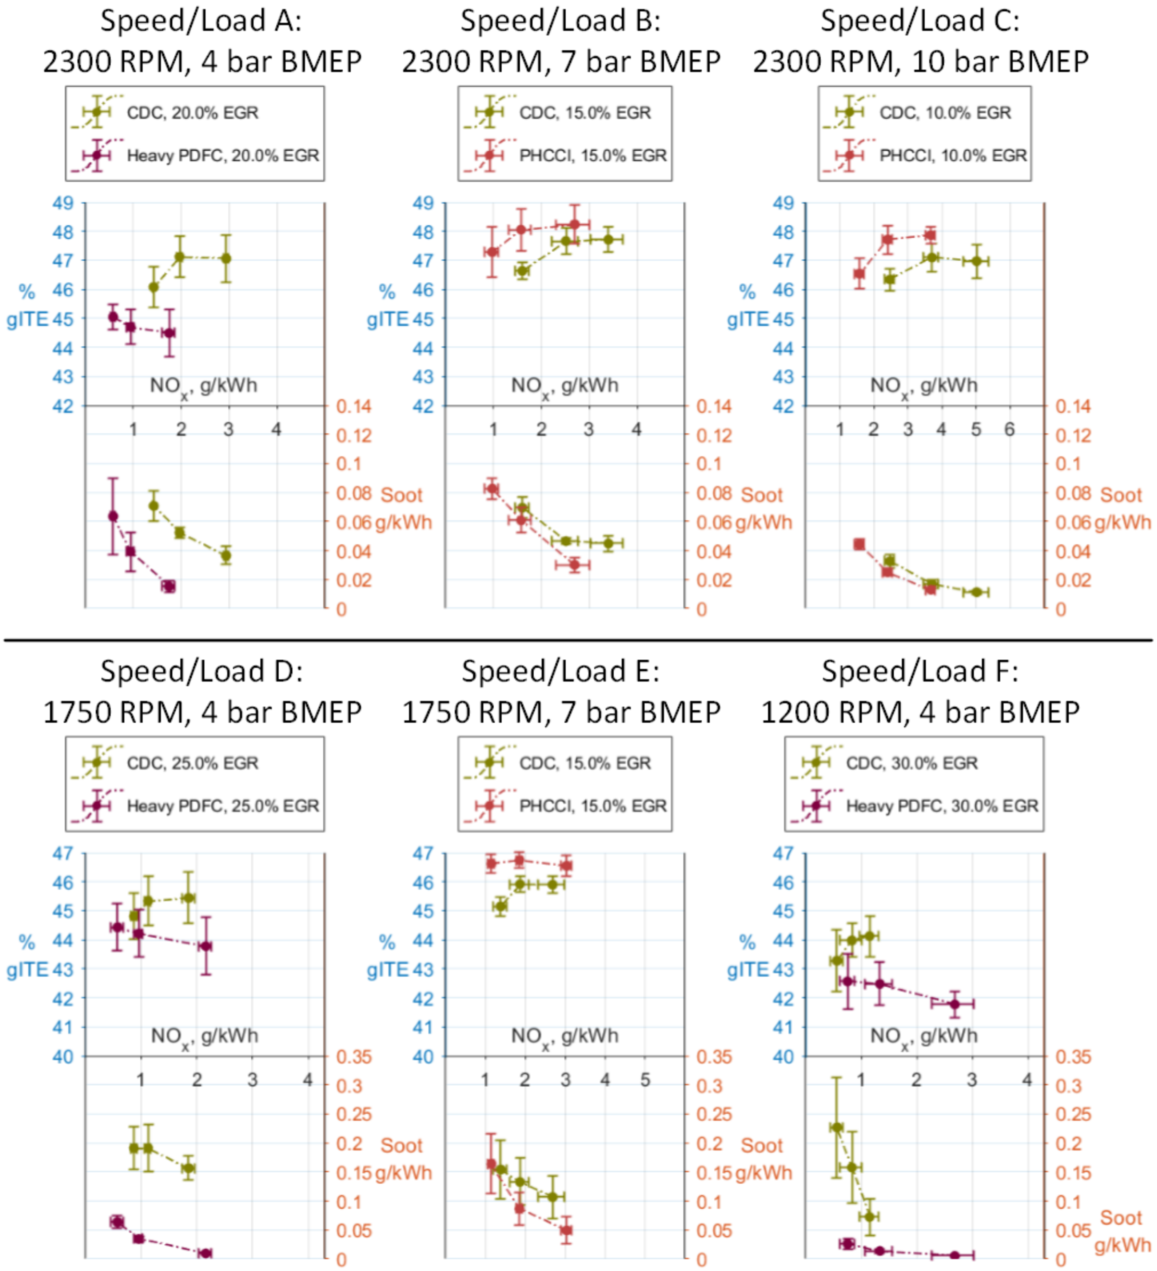

Figure 5.18 Gross ITE, NO$_X$, and soot at key points with 95% confidence intervals, discussed on page 176.

any implementations of PHCCI will have to carefully balance the ITE gains with noise restrictions by further optimizing the fueling ratio. These same limitations do not appear to apply to PDFC in most cases, where the increased stratification of premixed fuel staggers the lean autoignition process down a gradient of reactivity and reduces the rate of pressure rise to avoid ringing.

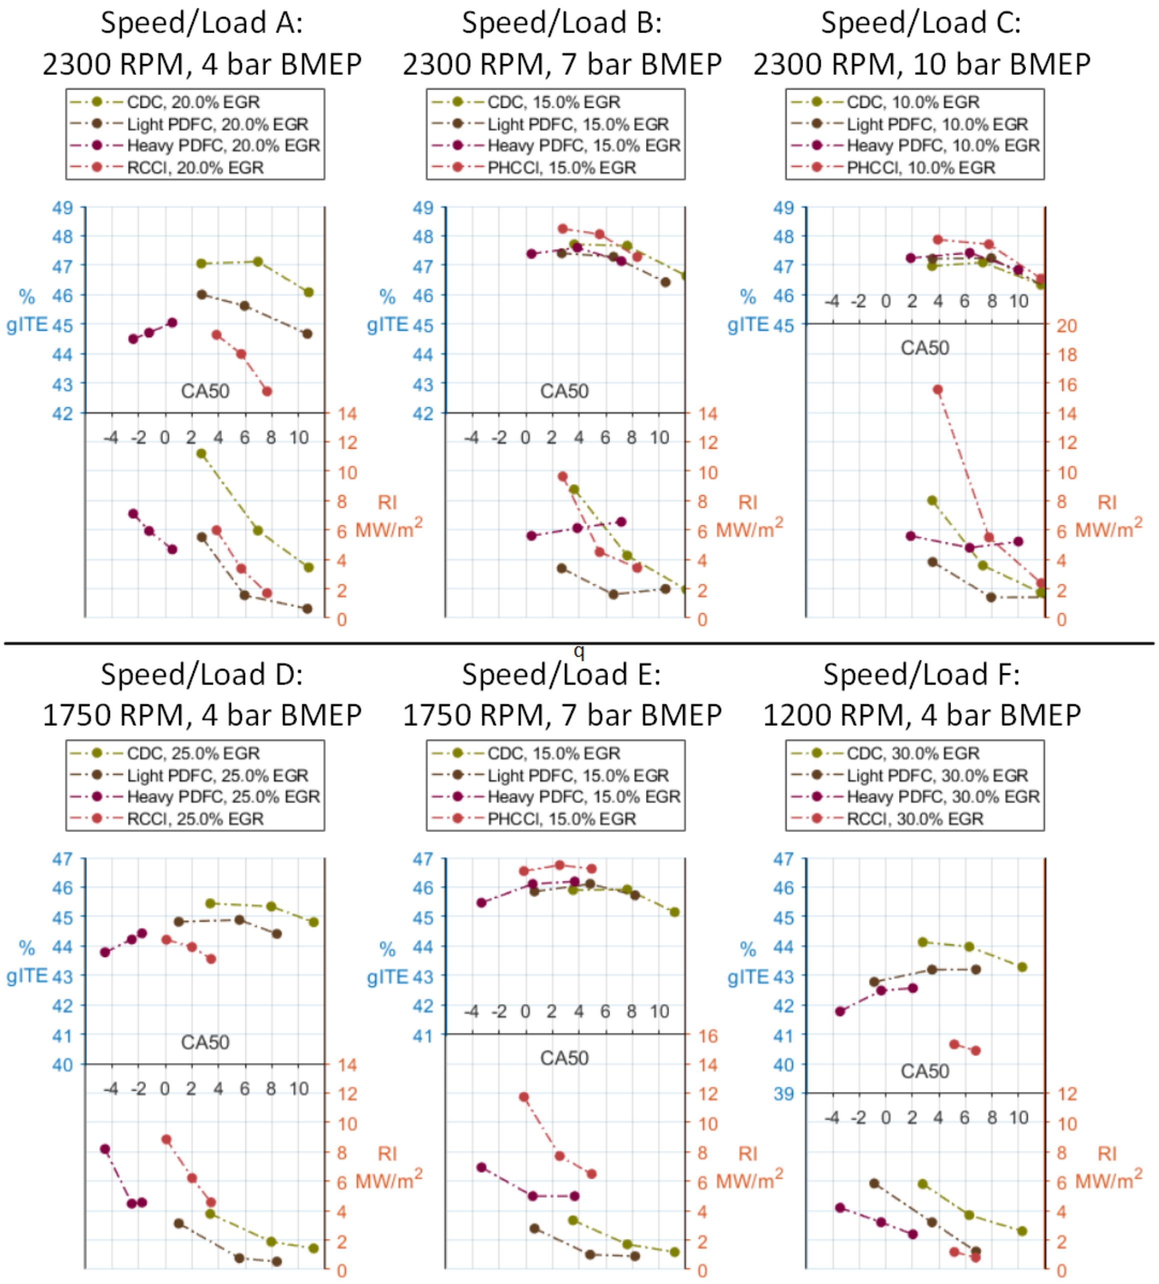

Figure 5.19 Gross ITE, CA50, and RI at key points with middle EGR level, discussed on pages 176-179.

Another phenomenon that can be observed in Figure 5.19 is the way the ITE usually reaches a peak at a CA50 of around 2-4°aTDC. If noise and NO_X emissions were not a factor, then optimizing each combustion mode would simply be a matter of bringing CA50 as close to this peak as possible, but as it stands such a CA50 often results in unacceptable

RI and NO_X due to the maximized peak pressure rise rates and temperatures, *except* for the case of Heavy PDFC in low loads. In this mode, the lean premixed autoignition forms the bulk of the heat release, but does not produce much NO_X because of the locally lean mixture fractions and relatively low temperatures due to low load. This allows the CA50 to be within this peak zone of 2-4°aTDC without producing excessive NO_X or ringing. Unfortunately, it appears that the selection process may have given these points slightly too much Early DI, and may have overshot the optimal CA50 for Heavy PDFC in speed/loads A, D, and F.

However, CA50 is just one component of combustion phasing. Figure 5.20 shows not just the CA50, but the CA10 and CA90 as well. At the low loads, the slope of the relationships of CA50 to CA10 and CA90 are both fairly consistent, but with a slightly shallower slope in the CA10 of the Light PDFC conditions, and a slightly steeper slope in the trend of CA90 in the Heavy PDFC conditions. This makes sense when considering that in Light PDFC, the phasing of the lean premixed autoignition occurs early and relies more on the fuel ratio than the SOI timing, so retarding the SOI should not slow this process down by very much. In Heavy PDFC, when the lean premixed autoignition consumes more oxygen throughout the combustion chamber, the soot left over from the Late DI must burn through a very dilute exhaust/air mixture, which slows down the last phase of combustion (diffusion). This effect is compounded when the Late DI is retarded giving more time for the autoignition process to consume the oxygen and for the piston to descend and cool the chamber before the Late DI soot can burn completely. Both of these trends have the effect of extending the duration of PDFC compared to CDC at low loads. In high loads, the same effect that applied to Light PDFC in low loads is seen in *both* Light and Heavy PDFC, which makes sense considering that the "Heavy" PDFC in the high loads has similar percentages of premixed fuel as the "Light" PDFC in the low loads. The extended duration of combustion limits ITE by pushing more of the heat release away from the peak window of 2-4°aTDC, with earlier combustion increasing cooling loss and later combustion increasing exhaust loss. This was seen previously in Figure 5.11, but Figure 5.20 shows additionally that the

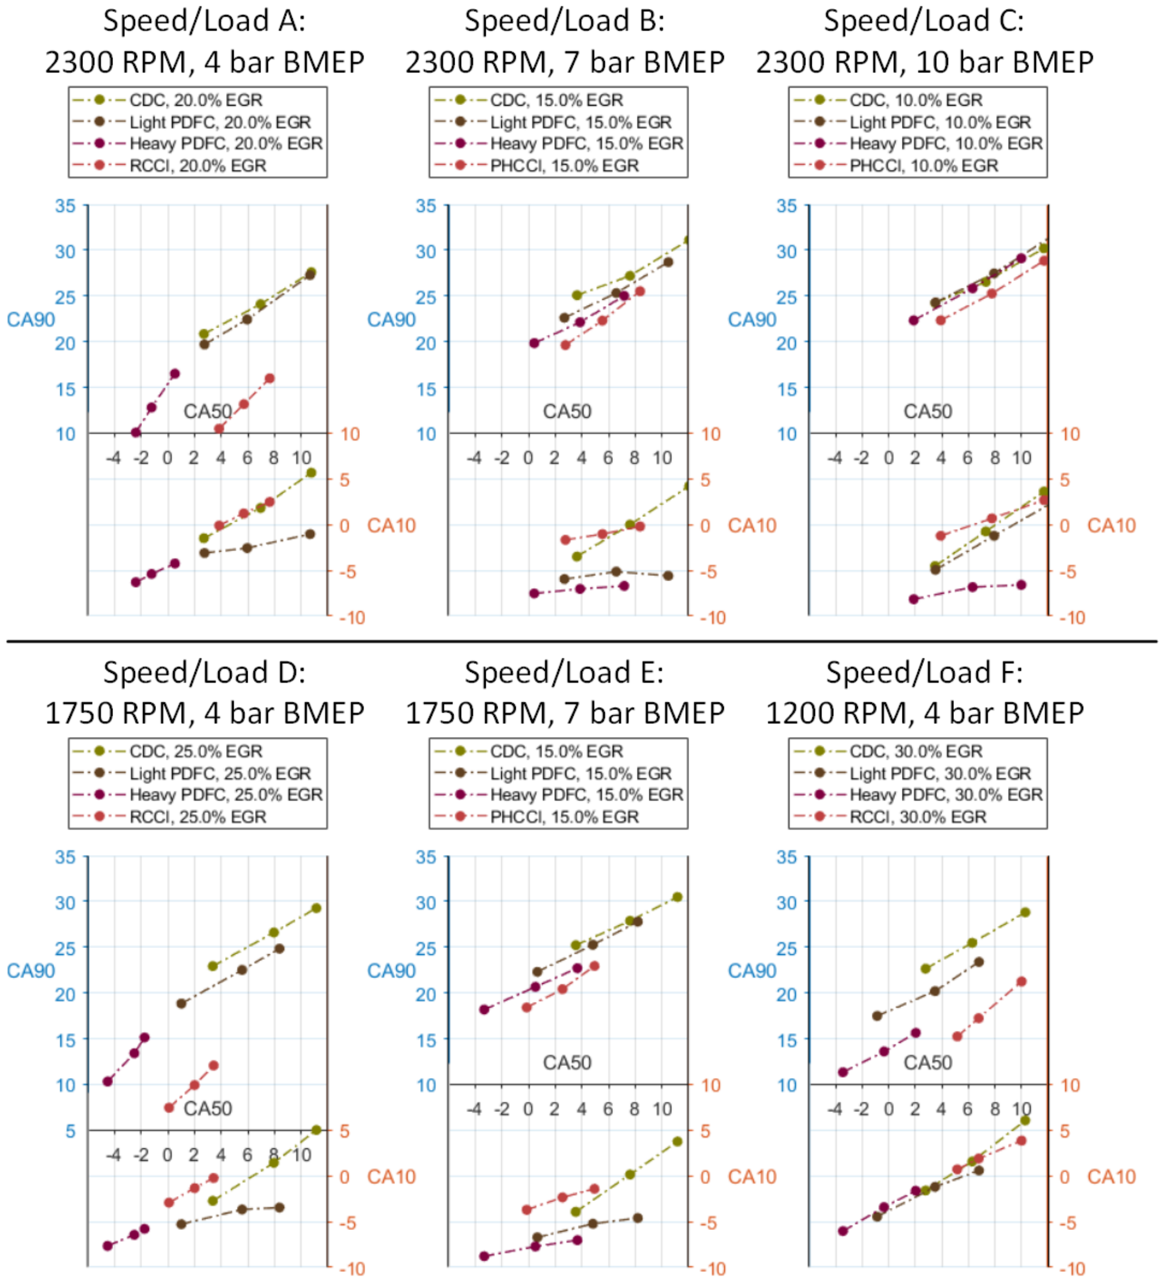

Figure 5.20 Combustion phasing at key points with middle EGR level, discussed on pages 179-181.

increased duration gets worse with later SOI's, and these later SOI's will be needed to limit NO_X emissions. This is unlike CDC, where the duration stays constant with different SOI's.

However, the durations of both RCCI and PHCCI are much shorter than the the other modes, and even though the durations get longer with later CA50's, they never exceed those of CDC. In RCCI, this can be attributed to the fact that there is no slow diffusion combustion process that drags out the end of the heat release, although this may also be the source of high unburnt emissions in RCCI. PHCCI, on the other hand, contracts the combustion duration by “squeezing” the total heat release shape closer together, as seen in Figure 4.17. Compared to PDFC, the lack of stratification delays the onset of lean premixed autoigntion but makes it much more rapid when it does occur, helping to concentrate more of the heat release closer to that ideal 2-4°aTDC window. This does however contribute to the high RI issue with PHCCI, since concentrating so much heat release close to TDC makes for a very rapid pressure increase.

To see the effect of combustion phasing changes on the heat balance, one can look at the trends of cooling loss and exhaust loss as related to CA50 in Figure 5.21. This figure shows that, as a rule, when combustion phasing is retarded, cooling losses go down, but exhaust losses go up. This is consistent with less time between the heat release and the exhaust valve opening, and thus less time for heat transfer to the combustion chamber walls, but also less time for heat conversion into work before the pressure seal is broken. For the most part, this appears be an equal trade-off, with cooling losses going down and exhaust losses going up at roughly the same rate.

The main trends that stands out in these graphs are the departure of RCCI and PHCCI from the trend of the other combustion modes. For RCCI, the exhaust loss is drastically lower than the other combustion modes with the same CA50 in speed/loads A and D, and although the cooling loss is on par with CDC and Light PDFC at the earliest CA50’s, it too drops below the other modes at later CA50’s. This is partially a result of the reduced combustion duration, which optimizes the trade-off between these two losses and improves ITE, but is also partially the result of the high unburnt losses, which lowers ITE. PHCCI also sees a simultaneous drop in exhaust and cooling loss compared to the other modes, but unlike RCCI, the drop in cooling loss is more significant than the drop in exhaust loss.

The reason for this difference can be attributed to the different ways in which each mode is reducing combustion duration, as seen in Figure 5.20. RCCI is mostly reducing

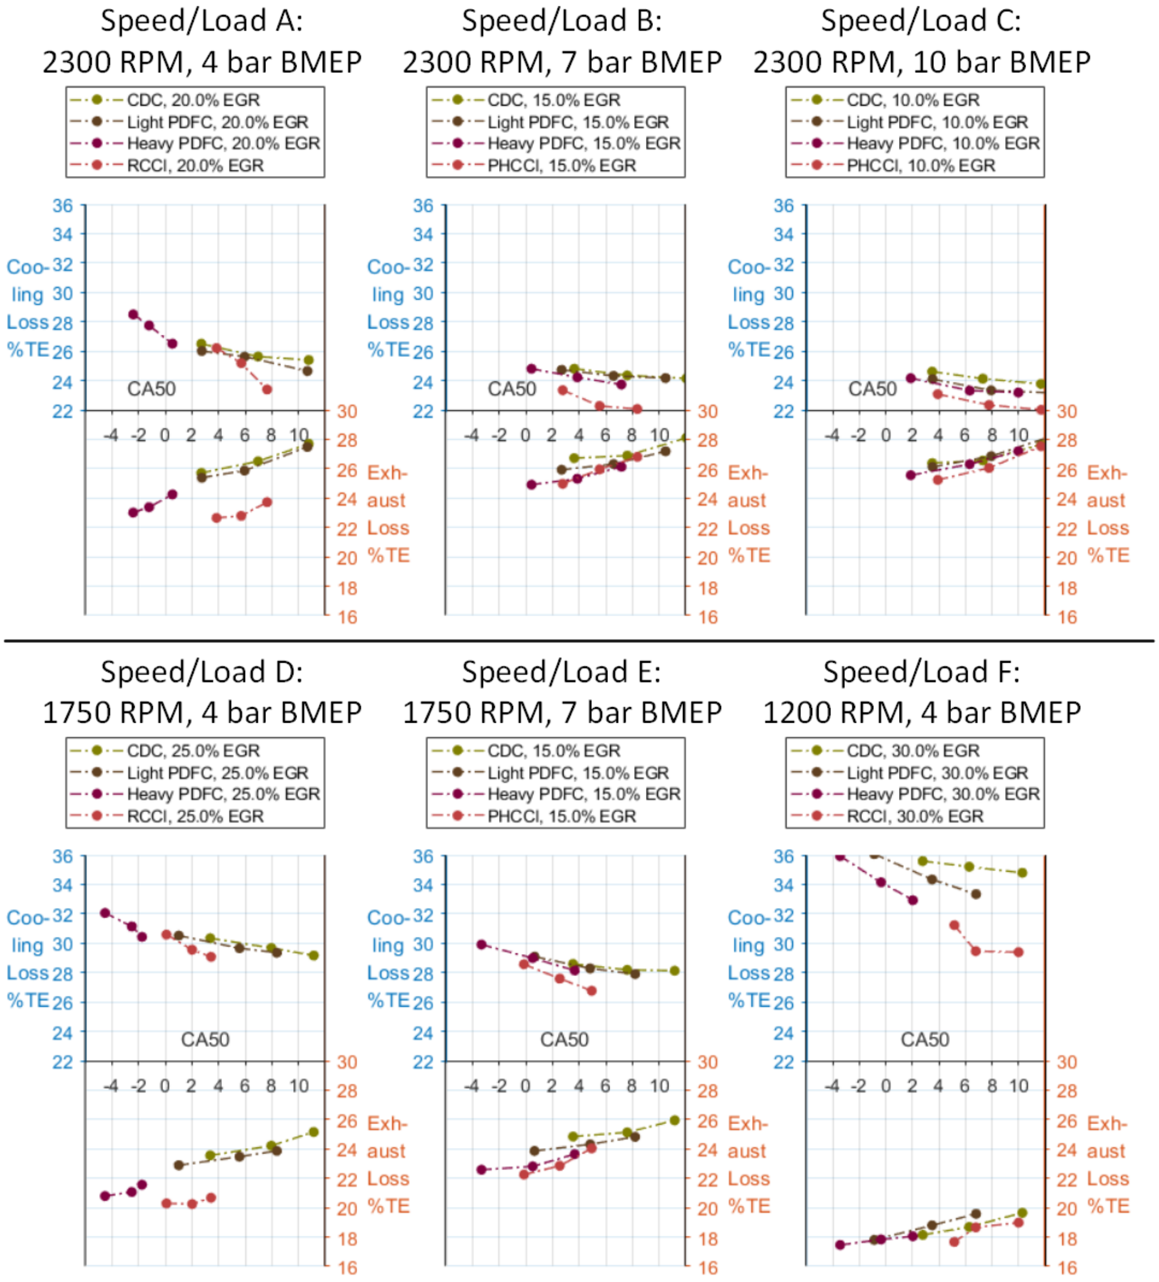

Figure 5.21 Cooling/exhaust loss trade-off at key points with middle EGR level, discussed on pages 181-183.

combustion duration by advancing the CA90, i.e. advancing heat release that was once occurring *after* the ideal window of 2-4° aTDC. This should increase cooling loss, but will create a larger reduction in exhaust loss. Meanwhile the combustion duration reduction in PHCCI is more “even” on both sides, with CA10 retarding and CA90 advancing in more equal measure. All else held equal this would reduce cooling and exhaust losses equally, but in actuality the cooling loss reduction is greater. As previously mentioned in Section 5.2.3, there are three possible reasons for the lower-than-expected cooling loss in both cases: reduced cylinder heat transfer, high unburnt loss, or compounded measurement errors affecting the cooling loss calculation, which is calculated by difference.

Figure 5.22 shows both the sum of the cooling and exhaust losses and the unburnt loss plotted with respect to CA50. This figure makes it clear the sum of cooling and exhaust losses is always going down with more premixed fuel. Given that the PDFC modes have increased combustion duration compared to CDC, this indicates that combustion phasing alone is not responsible for the improvement in this trade-off, and there must also be reductions in the rate of cylinder heat transfer, due to either the overall decrease in temperature from more unburnt fuel, or the more even dispersion of temperatures throughout the combustion chamber due to the increasing amount of lean premixed autoignition. At low load, the unburnt loss is increasing more than the cooling + exhaust losses are reducing with each change in combustion mode, but the opposite is true at higher loads. This shows that in the low-load conditions, any reduction in cylinder heat transfer loss is not worth the increase in unburnt emissions from the standpoint of CO_2 emissions alone (it may still be worth it when considering other emissions). But for high-load conditions, the cylinder heat transfer improvements outweigh and are worth the increase in unburnt emissions. This is especially true for PHCCI, when there is minimal stratification of temperatures and thus minimal heat transfer loss. When combined with the shortened combustion duration, this gives PHCCI the highest efficiency of the modes studied here despite the significant unburnt fuel loss.

After confirming the significance and cause of the performance changes at constant EGR levels, these combustion modes can now be optimized in a manner similar to what was done with the “Type 2” contour maps in the previous chapter: by comparing results while maximizing ITE and limiting emissions. First, the results will be compared when limiting NO_X, and then compared at similar levels of soot. These results will eventually be

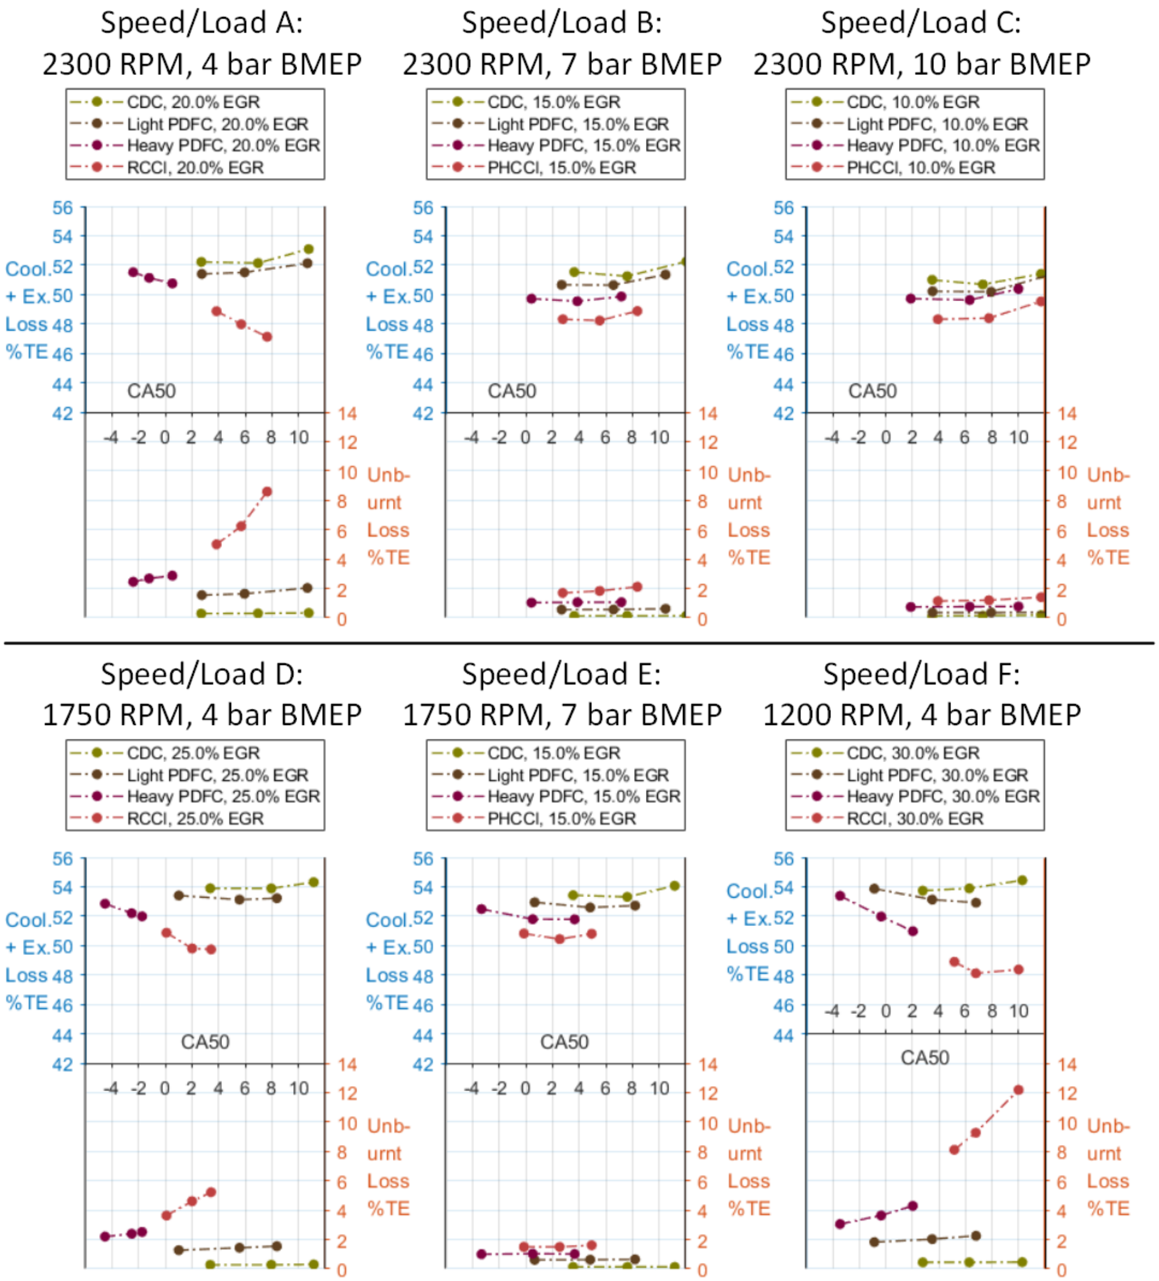

Figure 5.22 Heat/unburnt loss trade-off at key points with middle EGR level, discussed on page 183.

used to compare a “Low-NO_X” multi-mode operating strategy with a “Low-Soot” operating strategy.

Thanks to the careful selection of the EGR/SOI windows, the CDC point with the highest EGR and latest SOI produced NO_X of approximately 1 g/kWh at each speed/load point except C, where the lowest NO_X was about 1.5 g/kWh. This is convenient considering that this high EGR/late SOI point is one of the points that was chosen to repeat three times (see Figure 5.5), increasing the certainty in these results. Unfortunately, the highest EGR levels that were selected for speed/loads D, E and F produced excessive soot, and for a more fair comparison these results should be averaged with the results from the middle EGR level when selecting the optimal "Low-NO_X" CDC condition. The ITE and NO_X have been averaged together, but given the exponential relationship between soot and EGR at this high level, it was decided to calculate a logarithmic average of the soot levels:

$$\text{Soot}_{\text{Low-NO}_\text{X}} = e^{[\ln(\text{Soot}_{\text{Mid. EGR}})+\ln(\text{Soot}_{\text{Max. EGR}})]/2} \tag{5.2}$$

This brings the CDC soot level for speed/loads D, E, and F down into the same range as speed/loads A, B, and C. These points are seen in Figure 5.23.

For the other combustion modes, the goal is to reach the same NO_X levels while getting maximum efficiency and minimal soot. To get to the same NO_X levels with Light PDFC, the same high amount of EGR as CDC is required. However, for the Heavy PDFC (at low load), and PHCCI (high load) conditions, the same approximate NO_X emissions can be achieved with the *middle* EGR level. In every case, this means that these advanced modes can produce the same amount of NO_X for less than half the amount of soot. These soot improvements are highlighted by the downward green arrows in Figure 5.23. Although there is a corresponding drop in ITE in speed/loads A and F, there is no change in ITE in speed/load D, and these are outweighed by larger gains in ITE in the PHCCI modes of all three high-load conditions.

This highlights a major advantage of the highly premixed modes of PHCCI and Heavy PDFC: they can achieve low NO_X emissions with less EGR (and hence more excess air) than CDC. Since these modes generate homogeneous lean premixed autoignition that consumes oxygen before the diffusion flame is established, they can tolerate a higher level of oxygen concentration in the charge without increasing the total amount of NO_X generated by the

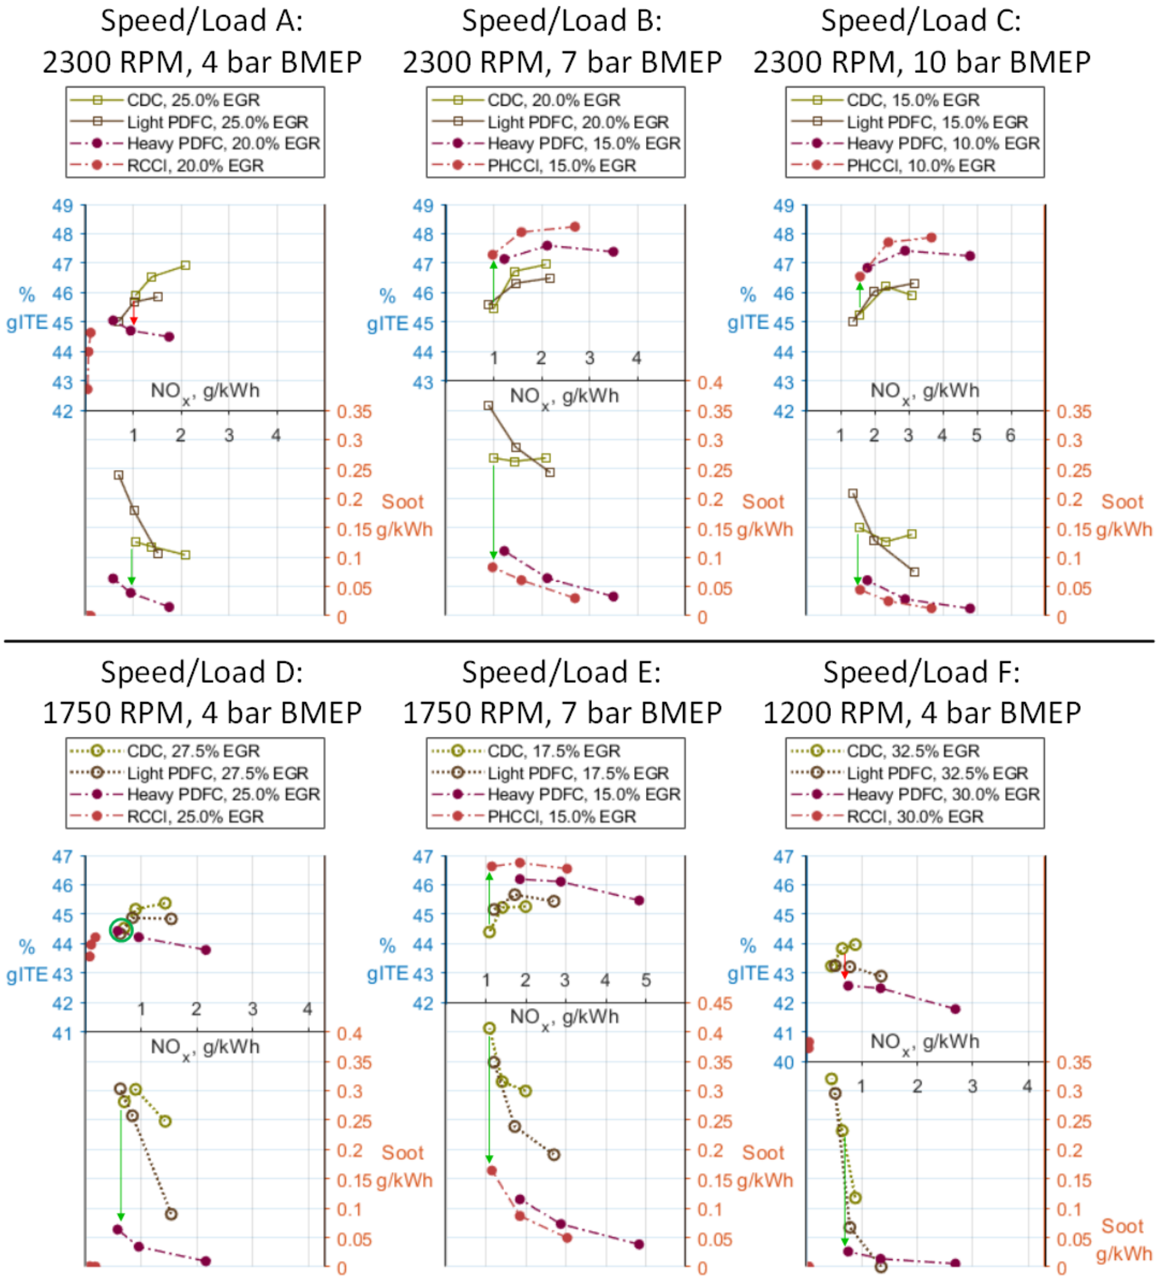

Figure 5.23 ITE and soot shifts at a fixed NO_X level between CDC and advanced modes, discussed on pages 185-187.

diffusion flame. This excess oxygen enhances soot oxidation, and with soot production already reduced by the shortened Late DI, these two effects combine to produce massive reductions in the eventual engine-out soot level. At the low loads, with Heavy PDFC, there may be some ITE reduction compared to CDC, but the upward trend in ITE with reduced

NO_X minimizes the size of this drop. At the high loads, with PHCCI, the decreased EGR requirement helps to increase the ITE gains further over the already significant gains seen with equal EGR levels. As a result, if one were to pursue a Low-NO_X multi-mode strategy for this engine, the ideal mode combination would be Heavy PDFC at low load and PHCCI at higher loads. RCCI could also be substituted for Heavy PDFC at low load to reduce NO_X and soot even further, but at the expense of a slight drop in ITE.

To find the "Low-Soot" CDC conditions, the point with the lowest EGR level and earliest SOI timing was used, which is again one of the points that was repeated three times for heightened accuracy. This point was found to coincide with the maximum ITE in speed/loads D, E, and F. Although the middle SOI timing was found to produce slightly higher ITE in speed/loads A, B, and C, the ITE from the earliest SOI's was never outside the uncertainty bounds of the middle SOI's, and using the middle SOI instead would have led to soot increases of about 50% in each case. For these Low-Soot conditions, the y-axis of soot will be increased by various degrees to show the detail in Figure 5.24.

At high load, to achieve the same approximate level of soot, all of the advanced combustion modes needed to use the same low level of EGR as the CDC condition. In the high-load PDFC modes, there were no significant NO_X or ITE improvements, but in PHCCI, there were some slight (low certainty) ITE improvements along with highly significant and certain NO_X reductions. Meanwhile at low load, the Heavy PDFC mode was able to achieve the same soot levels as CDC while using 5% more EGR and more retarded SOI's, which cut NO_X by more than half in each case. These NO_X improvements at constant soot are highlighted by the leftward-pointing green arrows in Figure 5.24. However, in the low-load cases, these NO_X reductions resulted in a drop of about 2% ITE. In speed/loads A and D, RCCI can be used instead of Heavy PDFC for even further improvements in NO_X and soot without significant change in ITE, but ITE drops by another 2% if RCCI is used in speed/load F. Once again, in low load there is a trade-off between ITE and emissions going from CDC to the more advanced modes. But in high loads, PHCCI is a win-win situation,

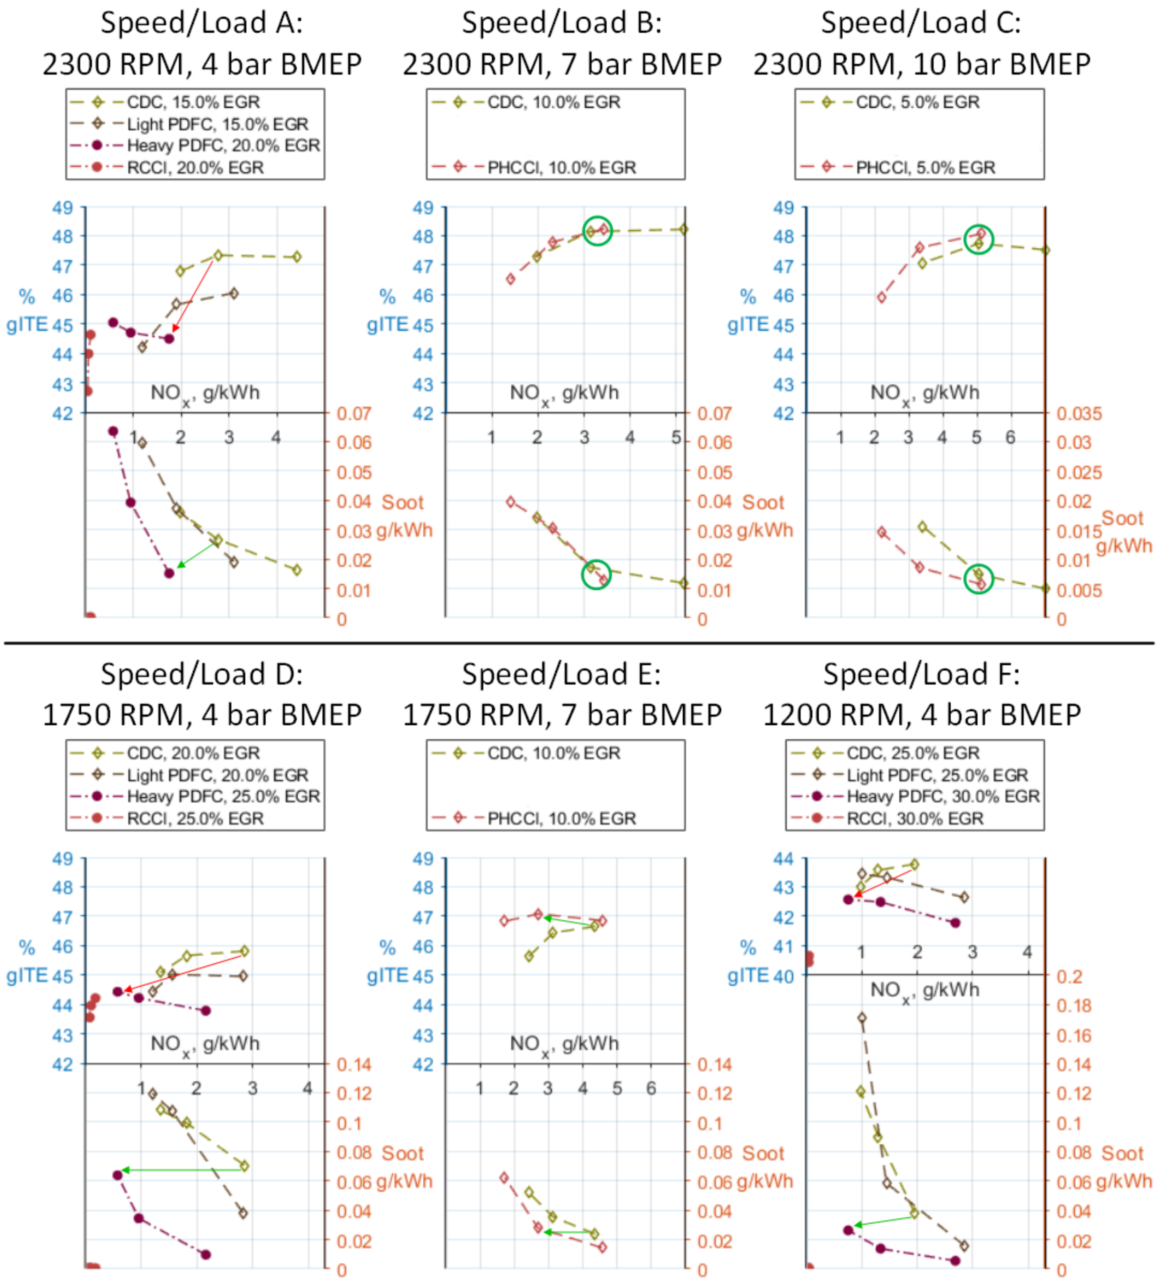

Figure 5.24 ITE and NO_X shifts at a fixed soot level between CDC and advanced modes, discussed on page 187.

lowering emissions and *possibly* raising ITE, although these ITE improvements are smaller and less certain than in the Low-NO_X strategy.

Now that engine performance has been optimized for multiple modes at certain speed/load conditions, a more holistic analysis can be performed on how certain multi-mode strategies would affect overall engine performance.

5.4 Estimating Benefits Over Entire Drive Cycle

The results from the previous section give a clear indication of possible changes in *steady-state* performance at six select speed/load points. Before using this data to estimate how the results will translate to overall drive-cycle performance, it is important to take note of the inherent limitations of the test setup. First and foremost, all of the data collected is at steady-state, with the engine operating constantly at a single speed and load and stable temperatures observed throughout the entire system. This is, of course, necessary in order to have accurate calculations of the energy balance in a system where the response times of certain sensors are very different from others. In any actual implementation, one would have to consider the effects of transient operation, such as the fact that suddenly increasing load would cause the cylinder wall temperature to lag behind the heat released by combustion, leading to lower cylinder temperatures (and inhibited lean premixed autoignition) in such a transient operation than in steady state.

Also, the six speed/load points tested herein are not enough to provide a comprehensive map of the entire engine operating range. They are, however, chosen to represent the speed/load areas most commonly seen during the simulated drive cycle shown in Figure 4.1, with the exception of speed/load F. This point was chosen despite its distance from the drive cycle points for a combination of relevance to larger, lower-speed engines as well as enabling the interpolation of these results to speed/loads D and E. The most important areas not covered by the selected points are the very low-load region, where temperatures may be too low and equivalence ratios too lean to generate efficient lean premixed autoignition, and the very high-speed region, where the piston may be moving too fast for the lean premixed autoignition to occur within a reasonably short crank angle window.

All of that notwithstanding, it is still worthwhile to calculate a weighted average of performance with the data that was collected for the sake of demonstrating potential benefits and encouraging more research to fill the aforementioned knowledge gaps. From the density of the second-by-second US06 drive cycle points in Figure 4.1, it appears that points A and B are inside a denser cluster of points visited by the drive cycle, and will thus be weighted twice as much as the other speed/load points. Speed/load F is irrelevant to this drive cycle and will not be included in the averages.

5.4.1 Selecting Operating Strategies

The results in the previous section make it clear that, when mainly considering ITE, NO_X, and soot, there is one advanced dual-fuel combustion mode (PHCCI) that is clearly advantageous at high load, and two modes (Heavy PDFC and RCCI) that give mixed results at low load. Thus, it was decided to consider four different options for which combustion mode to use at each load:

- **"CDC/CDC":** Use CDC for both low (4 bar BMEP) and high (7 and 10 bar BMEP) loads. This is the "baseline" to compare to the dual-fuel strategies.
- **"CDC/PHCCI":** Use CDC at low load, but PHCCI at high loads. This will improve efficiency over the baseline while also lowering NO_X and soot emissions.
- **"PDFC/PHCCI":** Use the Heavy PDFC mode at low load and PHCCI at high loads. This will improve NO_X and soot emissions further but with lower efficiency than the CDC/PHCCI option.
- **"RCCI/PHCCI":** Use RCCI at low load and PHCCI at high loads. This will give the best NO_X and soot emissions, but with still lower efficiency than the PDFC/PHCCI option.

With each of these combustion mode combinations, there are nine combinations of EGR and SOI to choose from, and a wide range of NO_X, soot, and CO_2 emissions that can be achieved. For RCCI, the NO_X and soot are much lower than the other modes and the point that gives the highest ITE can be used regardless of the emissions levels targeted. To get fair comparisons between each of the other modes, there will be three different methods of choosing each EGR/SOI combination:

- **"Low CO_2/Soot":** In CDC, use the earliest SOI and lowest EGR level to maximize ITE (thus minimizing CO_2 emissions) and minimize soot. In the advanced modes (PHCCI, PDFC, and RCCI), use the EGR/SOI combination that matches the level of soot produced by CDC. These will be the same points highlighted in Figure 5.24.
- **"Balanced":** Simply use the middle SOI and middle EGR level for *all* combustion modes. With the CDC/CDC baseline strategy, this will produce emissions comparable to those produced with the original ECU settings, and yield a good balance between NO_X, soot, and CO_2.
- **"Low NO_X":** In CDC, use the latest SOI and highest EGR level to minimize NO_X emissions. In the advanced modes, use the EGR/SOI combination that matches the level of NO_X produced by CDC. These will be the same points highlighted in Figure 5.23.

All in all, this will yield 12 different sets of results: 3 single-fuel baselines and 9 dual-fuel alternatives. This may seem like a lot of strategies to compare, but it is important to see all these options to understand how the desirability of the dual-fuel modes changes depending on what goals are set with the emissions.

5.4.2 Results from Selected Strategies

The results of these comparisons will be given in the form of 5-axis radar plots. These plots display the various strengths and weaknesses of each strategy simultaneously. The scales have been chosen so that poor results will occur toward the outside of the plot, while ideal results will be located close to the center. Each point on the radar will be given a numeric label that is colored using the same color scheme that was used for the contour maps in

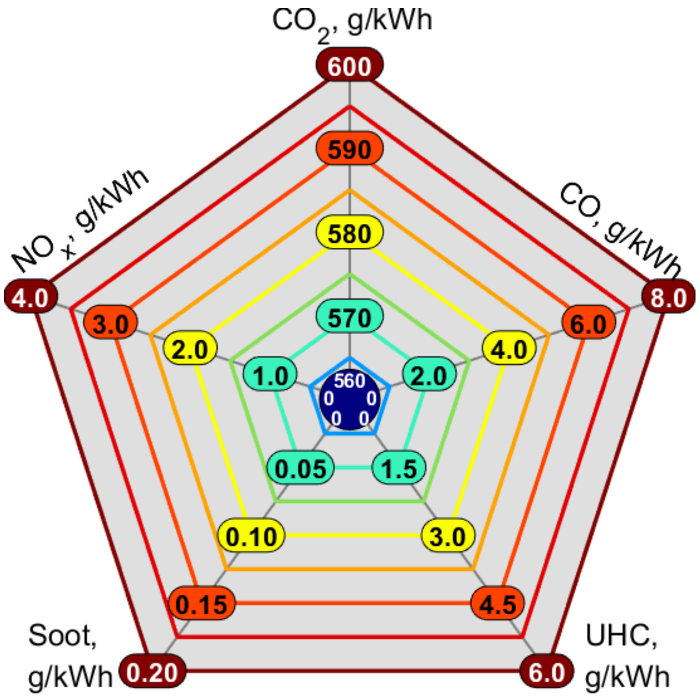

Figure 5.25 Key for emissions radar plots. All emissions are shown on a gross indicated basis.

Chapter 4. These radar plots will be shown in two forms: In the first, the five points of the radar will be the five major emissions produced: CO_2, NO_X, soot, UHC, and CO. The key to these plots is given in Figure 5.25. In the second form, the UHC and CO emissions measurements will be replaced with CoV(IMEP) and RI, to get a rough picture of how the emissions improvements must be paid for in terms of engine drivability and noise/durability respectively. The key to these plots is in Figure 5.27.

When showing the emissions alone in Figure 5.26, it is apparent that there is no strategy that minimizes all of the emissions simultaneously, only various trade-offs between the different emissions. The left column shows the results from the baseline strategies using CDC in all conditions. If taking the approach of minimizing CO_2 and soot emissions, there will also be very low CO and UHC emissions, but very high NO_X emissions of 3.12 g/kWh. By taking a more balanced approach, there is no increase in CO or UHC required to cut NO_X to 2.13 g/kWh, but as a consequence the CO_2 emissions rise by 6 g/kWh and the soot emissions nearly triple. But this pales in comparison to the soot and CO_2 increases required

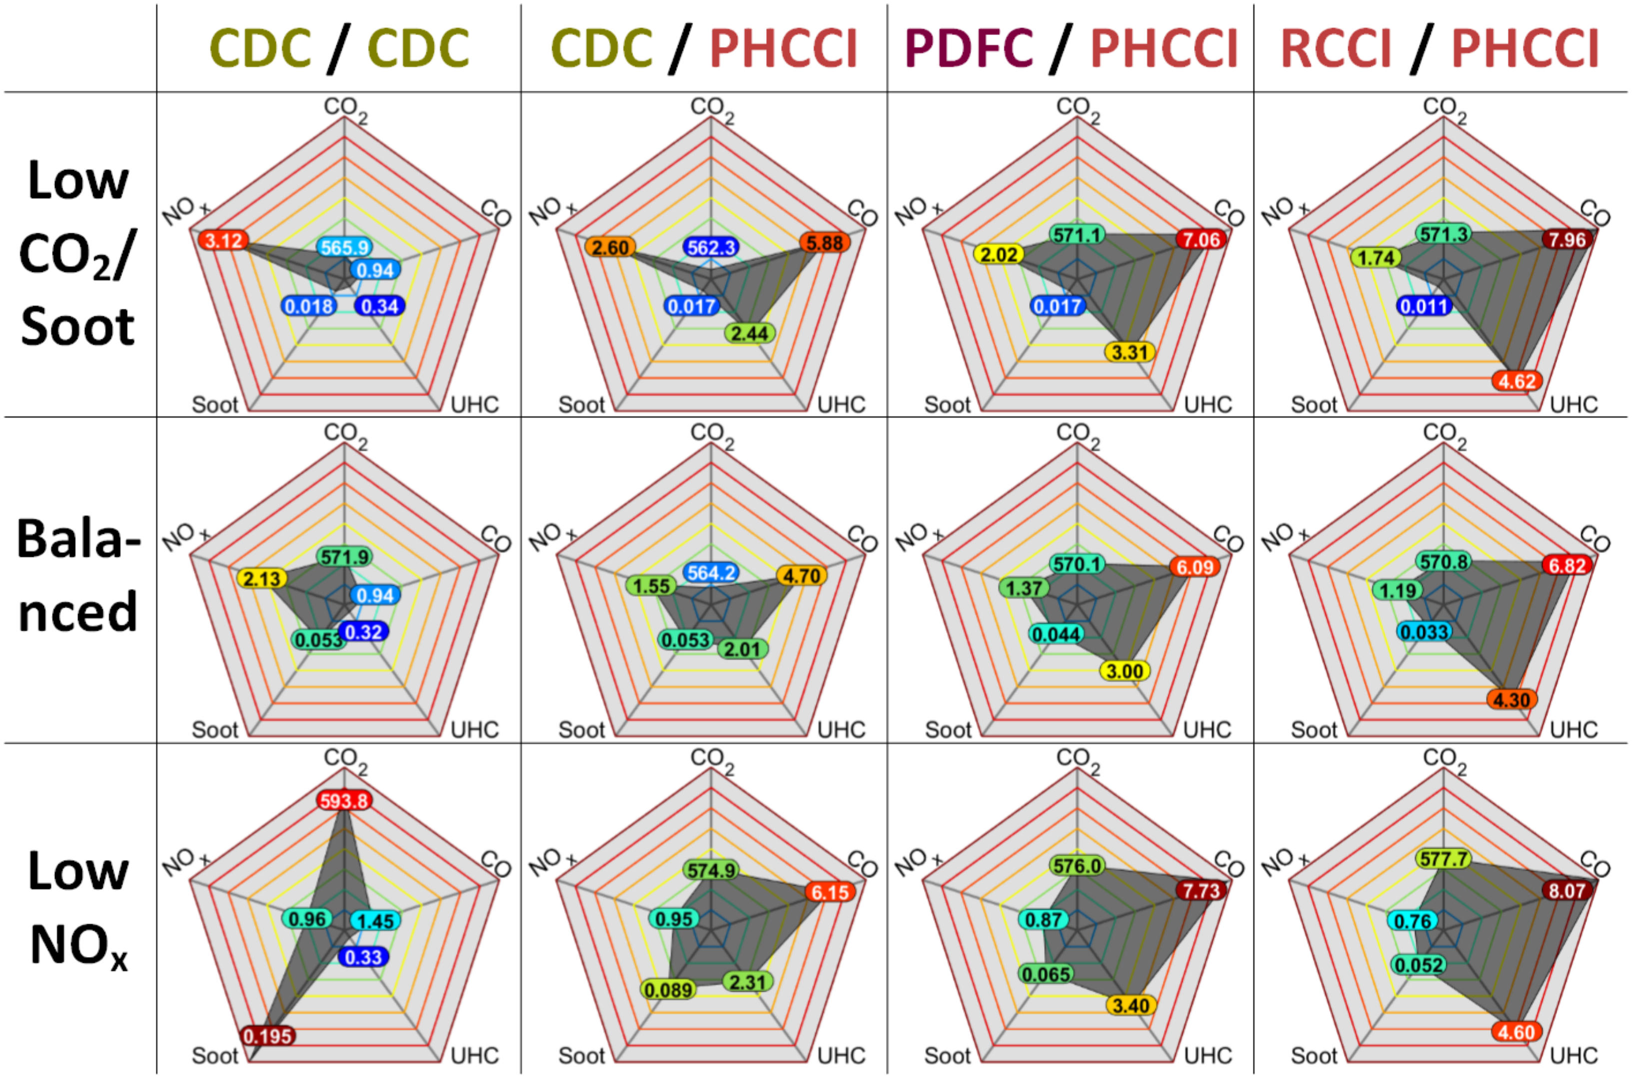

Figure 5.26 Emissions radar plots for different multi-mode combinations (Low load/high load in top row) and emissions targets (left column).

to reduce engine-out NO_X down to less than 1 g/kWh. The CO_2 emissions are more than 13 g/kWh above any of the other strategies, while soot emissions are more than double any of the other strategies. Clearly, the low-NO_X strategy is not viable for a standard diesel engine, and these engines must use aftertreatment to meet NO_X emissions standards.

If the engine instead uses PHCCI at high loads, while still operating on CDC at low loads, there is a significant decrease in CO_2 emissions over the baseline no matter which strategy is chosen, but only at the price of sharp increases in CO and UHC emissions. The magnitude of this CO_2 reduction depends on the strategy - for the Low-CO_2 strategy the reduction is only 3.6 g/kWh, versus 7.7 g/kWh for the Balanced strategy and a massive 18.9 g/kWh reduction for the Low-NO_X strategy. The Low-CO_2 and Balanced strategies also achieve more than a 0.5 g/kWh reduction in NO_X while soot stays constant, whereas the Low-NO_X strategy holds NO_X constant while cutting soot emissions by more than half. However, each of these strategies sees the CO and UHC emissions more than quadruple, although the Balanced strategy is slightly better than the other two in this regard. The question that remains is whether these unburnt emissions can be easily handled with a catalytic converter, and are worth the reduction in CO_2 and aftertreatment consumables (DEF and DPF) needed to treat the NO_X and soot emissions.

Unfortunately, the CO_2 emissions go back up if the engine uses advanced dual-fuel modes in the low load conditions as well (PDFC in the third column, and RCCI in the rightmost column in Figure 5.26), although this increase depends heavily on the strategy that is used. When using PDFC in low load as well as PHCCI at high load, the CO_2 emissions from the Low-CO_2/Soot strategy go up to 5.2 g/kWh above the baseline, while the Balanced strategy manages to keep them 1.8 g/kWh below the baseline, and the Low-NO_X strategy still retains a large 17.8 g/kWh improvement. There are additional improvements in NO_X and soot, with the most significant of these being a 0.58 g/kWh reduction in NO_X for the Low-CO_2/Soot strategy, while the other improvements in NO_X and soot are much less dramatic than when initially switching to PHCCI in high load. However on the positive

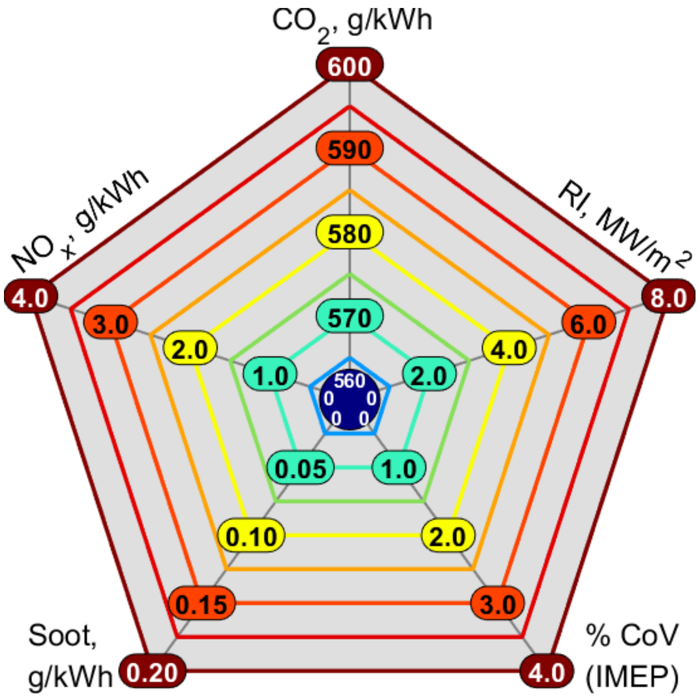

Figure 5.27 Key for final radar plots, with CO and UHC emissions replaced with RI and CoV(IMEP).

side the increases in CO and UHC are also less severe, although the reduced exhaust gas temperatures at low load could make these emissions more difficult to treat than those produced at high load. If the engine uses RCCI at low load, there is a significant further decrease in both NO_X and soot emissions for each strategy, but a much steeper penalty in terms of both CO and especially UHC emissions, along with a slight increase in the CO_2 emissions.

The trend in CO and UHC emissions is very clear and consistent - the more lean premixed combustion there is, the more unburnt emissions will be left to clean up in the engine-out exhaust. As previously mentioned, these emissions are relatively easy to remove from the exhaust with a catalytic converter, and this is why they are not as big of a problem as CO_2, NO_X, and soot. What may be a bigger issue with the advanced dual-fuel modes are the increases in ringing and cycle-to-cycle variation. For this reason the two axes of CO and UHC on the right side of the radar plots will be replaced with RI and CoV of IMEP, as shown in the updated key in Figure 5.27.

The updated radar plots in Figure 5.28 show that there is a trade-off between RI and CoV that mirrors the trade-off between NO_X and soot. The Low CO_2/Soot strategy will produce higher RI, while the Low-NO_X strategy will produce higher CoV. Much like the CO and UHC emissions, both of these values increase with the advanced dual-fuel modes, but this increase is not always monotonic from left to right like it was with the CO and UHC. With the Low-CO_2/Soot strategy, RI is actually reduced by using PDFC and RCCI, since this strategy requires early SOI and low EGR which increase the RI of CDC much more than they do for PDFC. Also, in the Balanced and Low-NO_X strategies, CoV actually goes down when switching from PDFC to RCCI at low load, although this CoV is still significantly higher than using CDC at low load.

When going down each column in Figure 5.28, it can be seen that the trade-off of RI and CoV greatly favors the Low-NO_X strategy, which brings RI down below 5 MW/m^2 no matter which combustion modes are used, without incurring much of an increase in CoV, which never exceeds 2%. This, combined with the fact that all of the advanced mode combinations have much lower CO_2 and soot emissions than the all-CDC baseline, shows that the dual-fuel modes are the most advantageous when pursuing a Low-NO_X engine-out emissions strategy. Thus, the dual-fuel multi-mode strategies will be more advantageous if there is a high priority on lowering engine-out NO_X emissions to reduce the burden on aftertreatment systems. It remains to be seen whether there are lower life-cycle CO_2 emissions from using a more balanced strategy and relying on aftertreatment to achieve acceptable tailpipe emissions, or whether it is more beneficial to use the dual-fuel strategy to keep the engine-out NO_X emissions to a minimum.

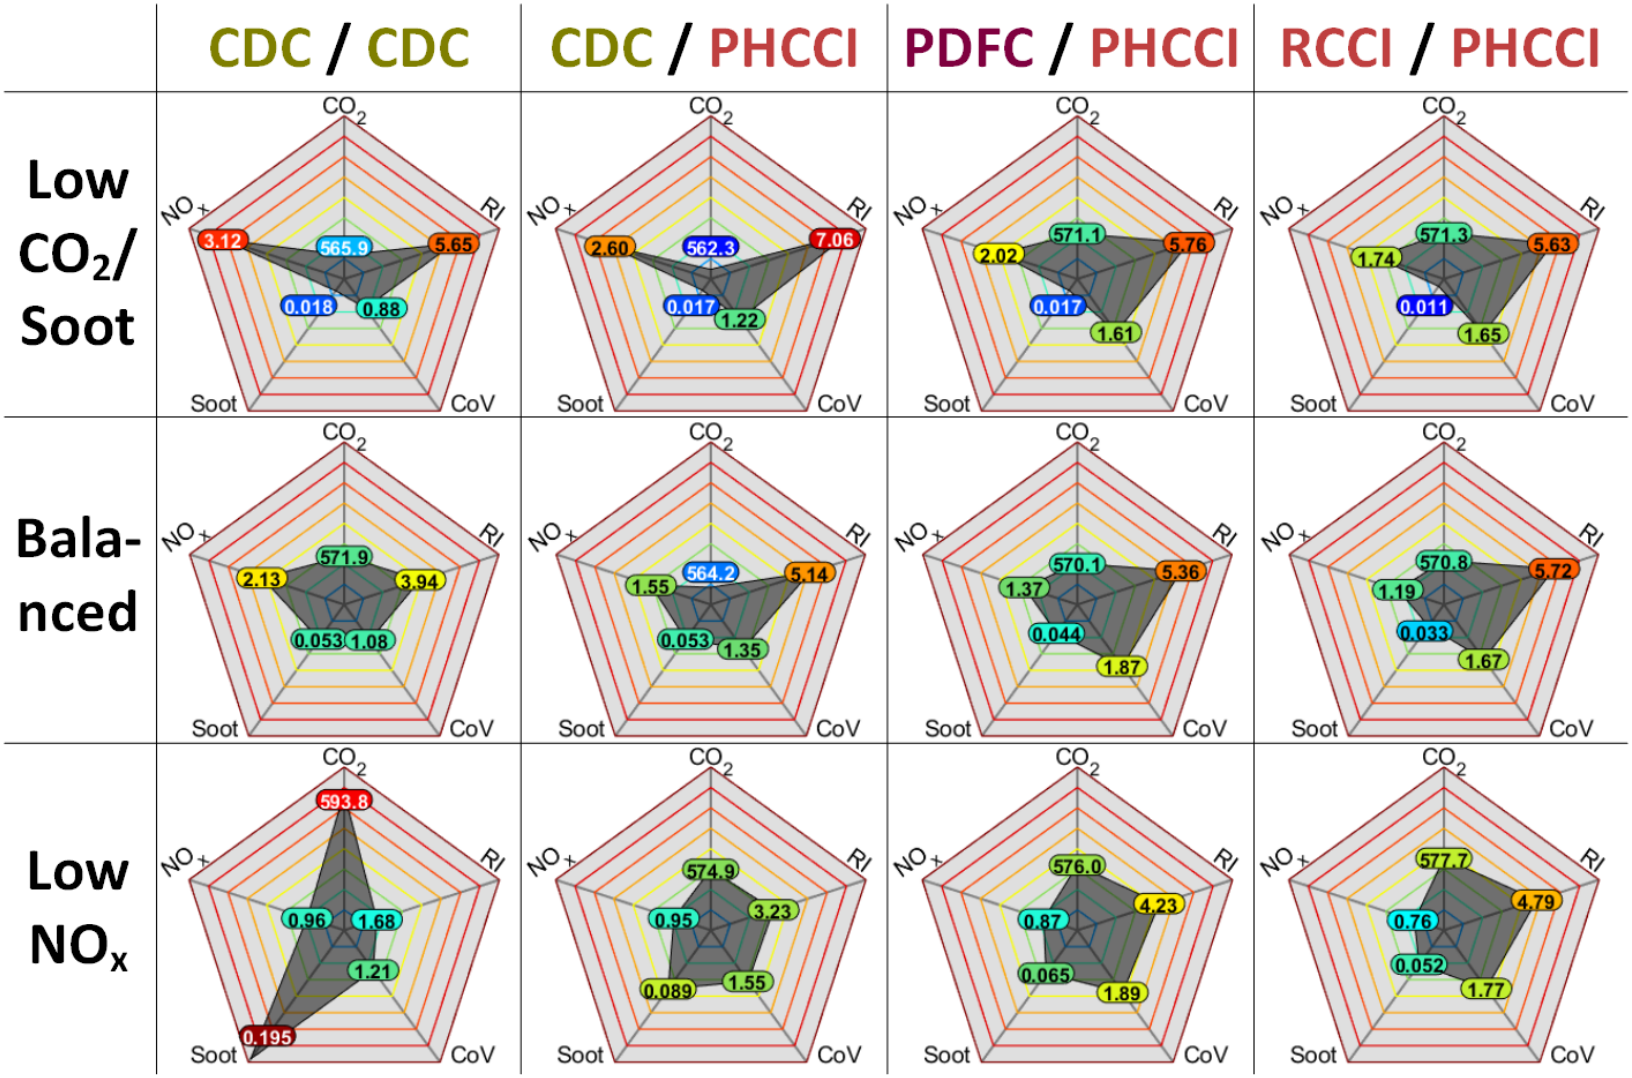

Figure 5.28 Final radar plots for different multi-mode combinations (Low load/high load in top row) and emissions targets (left column).

Chapter 6
Summary and Conclusions

This thesis has presented an experimental investigation of the possible the improvements in performance that can be achieved by converting a standard CI diesel engine into a dual-fuel gasoline/diesel engine, with no other mechanical modifications to optimize the engine besides to addition of the PFI gasoline injectors to the intake manifold. The results show that this engine configuration, using various multi-mode strategies, can improve steady-state engine-out emissions of CO_2, NO_X, and soot at the price of increasing UHC and CO emissions, as well as higher noise and cycle-to-cycle variability. Along the way to these "headline" findings, this thesis has made substantial contributions to multiple areas of knowledge in the IC engine research field.

6.1 Relationships Between Advanced Combustion Modes

The first area of knowledge to which this thesis has contributed is the understanding of advanced combustion modes, in particular dual-fuel combustion modes and how the different modes are related to each other. As described in Chapter 1, there has heretofore been a lack of study on the continuous relations between these modes, which are inconsistently defined across different studies by a confusing "alphabet soup" of acronyms. These issues have been remedied by the creation of the MELT diagram in Chapter 2, which breaks down the fuel injection into three key injection types: Manifold, Early DI, and Late DI. The Manifold injection uses a low-reactivity fuel (LRF), while the two DI's use a high-reactivity fuel (HRF). Plotting the distribution of these injection types creates a continuous, two-dimensional coordinate system that links together many combustion modes that have previously been compared mostly on a discrete, “side-by-side” basis, but can now be studied along a continuous spectrum.

Studying the continuous changes in the HRR pattern along incremental steps through the MELT diagram (as was done in Chapter 3) gives a much more cohesive understanding of how all the different advanced combustion modes are related to each other. By using Wiebe function fitting, these HRR patterns can be broken up or “de-convoluted” into the combustion sub-processes that occur at different times and at different magnitudes in the different combustion modes. These Wiebe function fits illustrate the gradual shifts in the breakdown of combustion between the rich premixed (soot-generating) and diffusion (soot-oxidizing, NO_X-generating) combustion of conventional diesel combustion and the lean premixed (soot- and NO_X-free) combustion of the advanced combustion modes. By using less Late DI and more Early DI and Manifold fuel, the amount of lean premixed combustion is increased, while the rich premixed and subsequent diffusion processes are decreased. The ratio of high-reactivity Early DI to low-reactivity Manifold fuel controls the size, shape, and timing of the HRR profiles from this chemically-controlled lean premixed process, which can be seen in the Wiebe function fits. These fits also illustrate the more subtle differences between the different combustion modes, which are described in detail in section 3.4.

At the time of this writing, the Wiebe function fits have only been used to illustrate the spectrum of combustion modes, but they have the potential to do much more. One key goal in the Wiebe function fitting was to ensure that there was always continuity between the fits at different test conditions, so that all the fits could be linked together into a continuous model. The next step to take with this work is to develop correlations between the test conditions (in terms of fuel properties and injection settings) and the Wiebe function parameters. This would allow for predictions of shape and timing of the different combustion processes at untested conditions without the need for experimental measurements, and help to “screen” different proposed test conditions to determine which conditions will produce an optimal combustion mode. This would involve creating accurate heat transfer and emission calculations from the Wiebe functions, and it remains to be seen whether this can be done using the fits presented in this thesis.

6.2 Performance Trade-Offs Between Combustion Modes

The MELT diagram was used not just to illustrate combustion modes, but also to create maps of their relative performance, which was the focus of chapter 4. A key trend that was immediately apparent in these maps was the “path” of high efficiency (in terms of gross indicated thermal efficiency or gITE) that ran through the MELT diagram at multiple speed/load combinations. This path showed the optimal ratio of Early DI and Manifold injections to use depending on the amount of Late DI being used, to control the phasing of the lean premixed autoignition process. When too much Early DI was used, the phasing of the lean premixed heat release became too advanced, causing the piston to do more work at the end of the intake stroke and allowing more time for this heat to transfer out of the combustion chamber, leading to higher cooling losses. Conversely when too little Early DI was used, the lean premixed heat release became too retarded and could not achieve full conversion of the homogeneously premixed Manifold fuel, leading to higher unburnt losses. Although there were many other complicating factors that are detailed in the body of the thesis, these were the two biggest contributions to the gITE trend seen in the MELT diagram maps, and the path of high efficiency coincides with the best trade-off between cooling and unburnt losses.

At low loads of 4 bar BMEP, this path of high efficiency starts in the Conventional Diesel Combustion (CDC) mode where mostly Late DI is used and ends in the Reactivity-Controlled Compression Ignition (RCCI) mode, where there is no Late DI and only a combination of Early DI and Manifold fuel. To get to RCCI from CDC, this path traverses the intermediate Piston-split Dual-Fuel Combustion (PDFC) mode, where all three injection types are used in significant amounts. When all else is held equal, PDFC produces an increase in soot compared to CDC. However, when the test conditions are allowed to change to suit the different combustion modes, this reveals that PDFC can achieve low NO_X emissions with less EGR than CDC, and this reduction in EGR gives PDFC more excess air to help promote soot oxidation. This effect is most significant in the “Heavy” PDFC modes that feature smaller Late DI’s, when there is less soot being produced initially by the Late DI.

Moving along the path of high efficiency from CDC to PDFC and then RCCI, the NO_X emissions and cooling losses relative to combustion phasing go down, but the unburnt losses increase rapidly. This trade is detrimental to gITE at low load. The low overall temperatures and equivalence ratios in these conditions inhibit the autoignition of the premixed fuel more than they reduce the heat transfer out of the combustion chamber, reducing the gITE of PDFC and RCCI relative to CDC. At higher loads of 7 and 10 bar BMEP, however, this trade is instead beneficial to gITE. The increased overall temperatures and equivalence ratios of the higher loads promote autoignition of the premixed fuel more than they increase the cooling loss, and the gITE of the advanced modes is increased relative to CDC.

At these higher loads, the path of high efficiency is found with less Early DI relative to Manifold fuel, since less high-reactivity Early DI is needed to advance the phasing of the lean premixed autoignition. However, the amount of premixed fuel that can be used in these conditions is limited, and RCCI cannot be reached. Instead, the path of high efficiency ends in the mode of Partial Homogeneous Charge Compression (Partial HCCI or PHCCI), when no more Manifold fuel can be added without knocking, even without any Early DI. The PHCCI mode has the additional advantage of produces the shortest duration of combustion, which combines with the reduced cooling loss relative to combustion phasing to give PHCCI the highest gITE of any of the modes studied in this thesis.

One of the main limitations of this study was that it only mapped the performance of a single engine, with a fixed injector design and piston shape that were optimized for CDC, and only the two most conventional fuels of diesel fuel and gasoline. This same mapping could be carried out for different injector designs that narrow the umbrella angle of the HRF direct injector, so that the Early DI is not directed far out onto the combustion chamber walls. This could be complemented by a change in the shape of a piston, so that the relatively flat squish zone where the Early DI is directed is changed to a secondary outer bowl like the piston shown in Figure 3.13. Changing the piston could also change the compression ratio, which would have a great impact on the timing of the lean premixed autoignition and the optimal ratio of Early DI and Manifold fuel. The same can be said for changing the reactivity of the fuels used, perhaps replacing the fossil fuels of gasoline and diesel with less carbon intensive fuels such as natural gas and biodiesel. It is important to investigate whether or not the trade-offs observed in this thesis are inherent to the combustion modes that were tested, and how these trade-offs change with changes to the engine and fuels used.

6.3 Proposed Multi-Mode Strategies

Once the performance mapping of this engine was complete, it was possible to identify key points on the MELT diagram that could produce changes in performance at different engine speeds and loads. At the higher-load conditions, it was apparent that PHCCI was the ideal mode to use, whereas at the lower-load conditions, it was less apparent whether CDC, PDFC, or RCCI was the best mode for this engine.

PHCCI can be used at high loads to improve efficiency and hence lower CO_2 emissions by reducing combustion duration and lowering the heat transfer losses out of the combustion chamber. This is achieved by displacing a large portion of the mixing-controlled combustion of conventional CI engines with lean premixed autoignition of the homogeneously premixed gasoline. With combustion occurring at locally lean equivalence ratios, these temperatures of combustion are lower than at the near-stoichiometric reaction zone of a diffusion flame generated with CDC, although the overall average equivalence ratio remains the same. This lowers the heat transfer rate out of the combustion chamber and avoids the creation of thermal NO_X, and also reduces the amount of locally rich combustion generating soot. By using just the right amount of gasoline and a small amount of early-injected diesel to get the timing right, this lean autoignition will occur just before the combustion of the late-injected diesel and minimize the amount of time needed for the total combustion process to occur. This reduced combustion duration helps to optimize the trade-off between cooling and exhaust loss, and although it can lead to an increase in engine noise, decreasing noise pollution is a lower priority at the current moment than reducing CO_2 pollution to mitigate the imminent threat of climate change. The increases in CO and UHC are also less of a threat from a climate change perspective and can be counteracted by using a catalytic converter.

The PDFC and RCCI dual-fuel modes can be used at low loads when PHCCI is not achievable, although they do not provide the same CO_2 benefit as PHCCI due to lower overall temperatures that reduce combustion efficiency and increase the leftover unburnt emissions of CO and UHC. PDFC performs better than RCCI in this respect because it generates a diffusion flame that can consume these unburnt products at the end of the combustion process, but RCCI provides the greatest reductions of NO_X and soot by completely eliminating the mixing-controlled combustion that generates these emissions.

These modes can be used in conjunction with PHCCI in a multi-mode strategy that will reduce overall cumulative CO_2, NO_X, and soot emissions simultaneously, as long as they are compared to a CDC baseline that is limited to producing 2 g/kWh NO_X or less. The best results occur when the engine must generate very low engine-out NO_X emissions of less than 1 g/kWh, which the dual-fuel modes can achieve with up to 3.2% lower CO_2 emissions and less than half the soot of a standard diesel engine. However if NO_X emissions are not a concern, CDC should be used instead at the low loads with PHCCI at high loads, and this CDC/PHCCI combination still produces a 1.0% reduction of CO_2 over a conventional diesel engine.

Further work should be done to understand how these multi-mode strategies will perform under transient conditions with various exhaust aftertreatment systems to determine their potential impact on overall “well-to-wheels” emissions. The NO_X level of 1 g/kWh is low in terms of engine-out emissions, but still would not come close to meeting modern emissions standards without aftertreatment. What must be determined is whether the reduced NO_X levels can lead to reductions in energy and resource use by these aftertreatments. It is also important to consider whether the additional resources and energy (and the CO_2 produced while generating said energy) used to produce dual-fuel engines are worth the benefit that these engines have been shown to capable of in this thesis.

6.4 Closing Statement

Having thoroughly explored the possibilities of dual-fuel CI combustion with this engine, I will conclude by saying that there is no clear “winner” in terms of which combustion mode or emissions reduction strategy is best. No matter what we do, there will always be trade-offs. For this reason it is important to look at the full spectrum of possibilities, not limiting oneself to discrete test conditions but exploring along a continuum of options to identify the local optima. Mitigating climate change is a process of slowly chipping away at CO_2 emissions wherever we can with incremental changes and there may inevitably be a price to pay in term of other emissions. However, by thoroughly exploring all of our options and weighing all of their advantages and disadvantages, we will be able to identify the best possible trade-offs and reduce the environmental impact of the world’s vast network of IC engine transportation systems. Although it is important to maintain optimism that we can eventually replace these engines with combustion-free alternatives, in the meantime we must make every possible effort to reduce the CO_2 produced by already-existing engines, and buy more time for the alternatives to develop.

References

[1] Rajendra K Pachauri, Myles R Allen, Vicente R Barros, John Broome, Wolfgang Cramer, Renate Christ, John A Church, Leon Clarke, Qin Dahe, Purnamita Dasgupta, et al. *Climate change 2014: synthesis report. Contribution of Working Groups I, II and III to the fifth assessment report of the Intergovernmental Panel on Climate Change*. IPCC, 2014. doi: 10013/epic.45156.

[2] John Conti, Paul Holtberg, Jim Diefenderfer, Angelina LaRose, James T Turnure, and Lynn Westfall. International energy outlook 2016 with projections to 2040. Technical report, USDOE Energy Information Administration (EIA), Washington, DC (United States). Office of Energy Analysis, 2016. doi: 10.2172/1296780.

[3] Rolf Reitz. Internal combustion engines: I: Fundamentals and performance metrics, combustion institute/princeton cefrc summer school, June 2018. https://cefrc.princeton.edu/sites/cefrc/files/reitz_princeton-day1-hour1-2018.pdf.

[4] Chan Pham and Wes Ingram. Detailed california-modified greet pathway for ultra low sulfur diesel (ulsd) from average crude refined in california, December 2014. https://www.arb.ca.gov/fuels/lcfs/121514ulsd.pdf.

[5] Chan Pham and Wes Ingram. Detailed ca-greet pathway for california reformulated gasoline blendstock for oxygenate blending (carbob) from average crude refined in california, December 2014. https://www.arb.ca.gov/fuels/lcfs/121514carbob.pdf.

[6] Ming Zheng, Graham T Reader, and J Gary Hawley. Diesel engine exhaust gas recirculation: a review on advanced and novel concepts. *Energy conversion and management*, 45(6):883–900, 2004. doi: 10.1016/S0196-8904(03)00194-8.

[7] Kazuhiro Akihama, Yoshiki Takatori, Kazuhisa Inagaki, Shizuo Sasaki, and Anthony M Dean. Mechanism of the smokeless rich diesel combustion by reducing temperature. *SAE Transactions*, 110(3):648–662, 2001. doi: 10.4271/2001-01-0655.

[8] Gary D Neely, Shizuo Sasaki, Yiqun Huang, Jeffrey A Leet, and Daniel W Stewart. New diesel emission control strategy to meet us tier 2 emissions regulations. *SAE Transactions*, 114(4):512–524, 2005. doi: 10.4271/2005-01-1091.

[9] Rolf D Reitz and Ganesh Duraisamy. Review of high efficiency and clean reactivity controlled compression ignition (rcci) combustion in internal combustion engines. *Progress in Energy and Combustion Science*, 46:12–71, 2015. doi: 10.1016/j.pecs.2014.05.003.

[10] Zhiming Gao, Scott Curran, C Daw, and R Wagner. Light-duty drive cycle simulations of diesel engine-out exhaust properties for an rcci-enabled vehicle. In *8th US National Combustion Meeting, University of Utah, May*, pages 19–23, 2013. https://www.osti.gov/biblio/1082102.

[11] Jonathan Martin, Andre Boehman, Rutvik Topkar, Sumit Chopra, Uday Subramaniam, and Heng Chen. Intermediate combustion modes between conventional diesel and rcci. *SAE International Journal of Engines*, 11(6):835–860, 2018. doi: 10.4271/2018-01-0249.

[12] Scott Curran, Vitaly Prikhodko, Kukwon Cho, C Scott Sluder, James Parks, Robert Wagner, Sage Kokjohn, and Rolf D Reitz. In-cylinder fuel blending of gasoline/diesel for improved efficiency and lowest possible emissions on a multi-cylinder light-duty diesel engine. Technical report, SAE Technical Paper 2010-01-2206, 2010. doi: 10.4271/2010-01-2206.

[13] Adam B Dempsey, N Ryan Walker, Eric Gingrich, and Rolf D Reitz. Comparison of low temperature combustion strategies for advanced compression ignition engines with a focus on controllability. *Combustion Science and Technology*, 186(2):210–241, 2014. doi: 10.1080/00102202.2013.858137.

[14] Erik Doosje, Frank Willems, and Rik Baert. Experimental demonstration of rcci in heavy-duty engines using diesel and natural gas. Technical report, SAE Technical Paper 2014-01-1318, 2014. doi: 10.4271/2014-01-1318.

[15] Shuaiying Ma, Zunqing Zheng, Haifeng Liu, Quanchang Zhang, and Mingfa Yao. Experimental investigation of the effects of diesel injection strategy on gasoline/diesel dual-fuel combustion. *Applied Energy*, 109:202–212, 2013. doi: 10.1016/j.apenergy.2013.04.012.

[16] Sage Kokjohn, Rolf Reitz, Derek Splitter, and Mark Musculus. Investigation of fuel reactivity stratification for controlling pci heat-release rates using high-speed chemiluminescence imaging and fuel tracer fluorescence. *SAE International Journal of Engines*, 5(2):248–269, 2012. doi: 10.4271/2012-01-0375.

[17] John B Heywood et al. *Internal combustion engine fundamentals*. Mcgraw-hill New York, 1988.

[18] Junseok Chang, Orgun Güralp, Zoran Filipi, Dennis Assanis, Tang-Wei Kuo, Paul Najt, and Rod Rask. New heat transfer correlation for an hcci engine derived from measurements of instantaneous surface heat flux. *SAE transactions*, 113(3):1576–1593, 2004. doi: 10.4271/2004-01-2996.

[19] Gerhard Woschni. A universally applicable equation for the instantaneous heat transfer coefficient in the internal combustion engine. *SAE transactions*, 76(4):3065–3083, 1968. doi: 10.4271/670931.

[20] J I Ghojel. Review of the development and applications of the wiebe function: a tribute to the contribution of ivan wiebe to engine research. *International Journal of Engine Research*, 11(4):297–312, 2010. doi: 10.1243/14680874JER06510.

[21] C Garnier, A Bilcan, Olivier Le Corre, and C Rahmouni. Characterisation of a syngas-diesel fuelled ci engine. Technical report, SAE Technical Paper 2005-01-1731, 2005. doi: 10.4271/2005-01-1731.

[22] John E Dec. A conceptual model of dl diesel combustion based on laser-sheet imaging. *SAE transactions*, 106(3):1319–1348, 1997. doi: 10.4271/970873.

[23] Mark PB Musculus, Paul C Miles, and Lyle M Pickett. Conceptual models for partially premixed low-temperature diesel combustion. *Progress in energy and combustion science*, 39(2-3):246–283, 2013. doi: 10.1016/j.pecs.2012.09.001.

[24] Nicolas Dronniou, Julian Kashdan, Bertrand Lecointe, Kyle Sauve, and Dominique Soleri. Optical investigation of dual-fuel cng/diesel combustion strategies to reduce co2 emissions. *SAE International Journal of Engines*, 7(2):873–887, 2014. doi: 10.4271/2014-01-1313.

[25] Robert H Thring. Homogeneous-charge compression-ignition (hcci) engines. Technical report, SAE Technical paper 892068, 1989. doi: 10.4271/892068.

[26] Rudolf H Stanglmaier and Charles E Roberts. Homogeneous charge compression ignition (hcci): benefits, compromises, and future engine applications. *SAE transactions*, 108(3):2138–2145, 1999. doi: 10.4271/1999-01-3682.

[27] Magnus Sjöberg, John E Dec, and Nicholas P Cernansky. Potential of thermal stratification and combustion retard for reducing pressure-rise rates in hcci engines, based on multi-zone modeling and experiments. *SAE transactions*, 114(3):236–251, 2005. doi: 10.4271/2005-01-0113.

[28] John E Dec, Wontae Hwang, and Magnus Sjöberg. An investigation of thermal stratification in hcci engines using chemiluminescence imaging. *SAE Transactions*, 115(3):759–776, 2006. doi: 10.4271/2006-01-1518.

[29] Shigeyuki Tanaka, Ferran Ayala, James C Keck, and John B Heywood. Two-stage ignition in hcci combustion and hcci control by fuels and additives. *Combustion and flame*, 132(1-2):219–239, 2003. doi: 10.1016/S0010-2180(02)00457-1.

[30] Janardhan Kodavasal, SeungHwan Keum, and Aristotelis Babajimopoulos. An extended multi-zone combustion model for pci simulation. *Combustion Theory and Modelling*, 15(6):893–910, 2011. doi: 10.1080/13647830.2011.578663.

[31] Prasad S Shingne, Robert J Middleton, Claus Borgnakke, and Jason B Martz. The effects of boost pressure on stratification and burn duration of gasoline homogeneous charge compression ignition combustion. *International Journal of Engine Research*, 20(3):359–377, 2019. doi: 10.1177/1468087417754177.

[32] Kazuhisa Inagaki, Takayuki Fuyuto, Kazuaki Nishikawa, Kiyomi Nakakita, and Ichiro Sakata. Dual-fuel pci combustion controlled by in-cylinder stratification of ignitability. Technical report, SAE Technical Paper 2006-01-0028, 2006. doi: 10.4271/2006-01-0028.

[33] Sage Kokjohn, Reed Hanson, Derek Splitter, John Kaddatz, and Rolf Reitz. Fuel reactivity controlled compression ignition (rcci) combustion in light-and heavy-duty engines. *SAE International Journal of Engines*, 4(1):360–374, 2011. doi: 10.4271/2011-01-0357.

[34] Sage L Kokjohn, Mark PB Musculus, and Rolf D Reitz. Evaluating temperature and fuel stratification for heat-release rate control in a reactivity-controlled compression-ignition engine using optical diagnostics and chemical kinetics modeling. *Combustion and Flame*, 162(6):2729–2742, 2015. doi: 10.1016/j.combustflame.2015.04.009.

[35] Henry J Curran, Paolo Gaffuri, William J Pitz, and Charles K Westbrook. A comprehensive modeling study of n-heptane oxidation. *Combustion and flame*, 114(1-2):149–177, 1998. doi: 10.1016/S0010-2180(97)00282-4.

[36] Tomohiro Kanda, Takazo Hakozaki, Tatsuya Uchimoto, Jyunichi Hatano, Naoto Kitayama, and Hiroshi Sono. Pcci operation with early injection of conventional diesel fuel. *SAE transactions*, 114(3):584–593, 2005. doi: 10.4271/2005-01-0378.

[37] William L Hardy and Rolf D Reitz. A study of the effects of high egr, high equivalence ratio, and mixing time on emissions levels in a heavy-duty diesel engine for pcci combustion. Technical report, SAE Technical Paper 2006-01-0026, 2006. doi: 10.4271/2006-01-0026.

[38] Guntram A Lechner, Timothy J Jacobs, Christos A Chryssakis, Dennis N Assanis, and Robert M Siewert. Evaluation of a narrow spray cone angle, advanced injection timing strategy to achieve partially premixed compression ignition combustion in a diesel engine. *SAE transactions*, 114(3):394–404, 2005. doi: 10.4271/2005-01-0167.

[39] Kazuki Inaba, Yuto Ojima, Yoshimitsu Kobashi, Gen Shibata, and Hideyuki Ogawa. Twin shaped semi-premixed diesel combustion with a combustion chamber dividing the fuel sprays (first report). *Transactions of Society of Automotive Engineers of Japan*, 49(2):205–210, 2018. doi: 10.11351/jsaeronbun.49.205.

[40] Ian May, Vinícius Pedrozo, Hua Zhao, Alasdair Cairns, Steve Whelan, Hoi Wong, and Paul Bennicke. Characterization and potential of premixed dual-fuel combustion in a heavy duty natural gas/diesel engine. Technical report, SAE Technical Paper 2016-01-0790, 2016. doi: 10.4271/2016-01-0790.

[41] MF Russell and R Haworth. Combustion noise from high speed direct injection diesel engines. *SAE transactions*, 94(4):810–831, 1985. doi: 10.4271/850973.

[42] PW Schaberg, T Priede, and RK Dutkiewicz. Effects of a rapid pressure rise on engine vibration and noise. Technical report, SAE Technical Paper 900013, 1990. doi: 10.4271/900013.

[43] JA Eng. Characterization of pressure waves in hcci combustion. Technical report, SAE Technical Paper 2002-01-2859, 2002. doi: 10.4271/2002-01-2859.

[44] Arsham J Shahlari, Chris Hocking, Eric Kurtz, and Jaal Ghandhi. Comparison of compression ignition engine noise metrics in low-temperature combustion regimes. *SAE International Journal of Engines*, 6(1):541–552, 2013. doi: 10.4271/2013-01-1659.

[45] John E Dec and Yi Yang. Boosted hcci for high power without engine knock and with ultra-low nox emissions-using conventional gasoline. *SAE International Journal of Engines*, 3(1):750–767, 2010. doi: 10.4271/2010-01-1086.

[46] Hideyuki Ogawa, Tie Li, Noboru Miyamoto, Shingo Kido, and Hejime Shimizu. Dependence of ultra-high egr and low temperature diesel combustion on fuel injection conditions and compression ratio. Technical report, SAE Technical Paper 2006-01-3386, 2006. doi: 10.4271/2006-01-3386.

[47] Naeim A Henein, A Bhattacharyya, Jeremy Schipper, A Kastury, and Walter Bryzik. Effect of injection pressure and swirl motion on diesel engine-out emissions in conventional and advanced combustion regimes. Technical report, SAE Technical Paper 2006-01-0076, 2006. doi: 10.4271/2006-01-0076.

[48] Andrea Catania, Roberto Finesso, and Ezio Spessa. Real-time calculation of egr rate and intake charge oxygen concentration for misfire detection in diesel engines. Technical report, SAE Technical Paper 2011-24-0149, 2011. doi: 10.4271/2011-24-0149.

[49] Richard D. Atkins. *An Introduction to Engine Testing and Development*. SAE International, apr 2009. https://saemobilus.sae.org/content/R-344.

[50] Smoke value measurement with the filter-paper-method: Application notes. online, June 2005. https://www.avl.com/documents/10138/885893/Application+Notes.

[51] PA Lakshminarayanan and S Aswin. Estimation of particulate matter from smoke, oil consumption and fuel sulphur. Technical report, SAE Technical Paper 2017-01-7002, 2017. doi: 10.4271/2017-01-7002.

[52] Johannes Brettschneider. Berechnung des luftverhältnisses lambda von luft-kraftstoff-gemischen und des einflusses von messfehlern auf lambda. *Bosch Tech Ber*, 6(4), 1979. https://trid.trb.org/view/1052620.

[53] Donald R. Burgess Jr. *Thermochemical Data*. National Institute of Standards and Technology, Gaithersburg, MD 20899, 2019. doi: 10.18434/T4D303.

[54] Ryo Hasegawa and Hiromichi Yanagihara. Hcci combustion in di diesel engine. *SAE transactions*, 112(3):1070–1077, 2003. doi: 10.4271/2003-01-0745.

[55] Adam B Dempsey, N Ryan Walker, and Rolf Reitz. Effect of piston bowl geometry on dual fuel reactivity controlled compression ignition (rcci) in a light-duty engine operated with gasoline/diesel and methanol/diesel. *SAE International Journal of Engines*, 6(1):78–100, 2013. doi: 10.4271/2013-01-0264.

[56] Scott Curran, Reed Hanson, Robert Wagner, and Rolf D Reitz. Efficiency and emissions mapping of rcci in a light-duty diesel engine. Technical report, SAE Technical Paper 2013-01-0289, 2013. doi: 10.4271/2013-01-0289.

[57] Chao Yu, Jian-xin Wang, Zhi Wang, and Shi-jin Shuai. Comparative study on gasoline homogeneous charge induced ignition (hcii) by diesel and gasoline/diesel blend fuels (gdbf) combustion. *Fuel*, 106:470–477, 2013. doi: 10.1016/j.fuel.2012.10.068.

[58] Jesús Benajes, Antonio García, Javier Monsalve-Serrano, and Vicente Boronat. Achieving clean and efficient engine operation up to full load by combining optimized rcci and dual-fuel diesel-gasoline combustion strategies. *Energy conversion and management*, 136:142–151, 2017. doi: 10.1016/j.enconman.2017.01.010.