Power Output of the Stirling Engine

A Thesis
Presented to
The Division of Mathematics and Natural Sciences
Reed College

In Partial Fulfillment
of the Requirements for the Degree
Bachelor of Arts

Luke Liechty

May 2014

Approved for the Division
(Physics)

Johnny Powell	Mary James

Acknowledgements

What a wild ride this has been, a flurry of unrelenting insanity only possible at a place like Reed. Writing this thesis has been a trying task and numerous people deserve my wholehearted appreciation. Firstly, I would like to give a huge, astronomically even, thanks to my two superb advisors, Mary James and Johnny Powell. Each took me through a very different stage in the development of this work, and I have learned much from both. Thank you for trusting me and tolerating my pride and stubbornness. An equally big thanks to Jay Ewing for teaching me literally everything I know about machining and helping me immensely throughout the build process, and to Bob Ormond for sharing motorcycle tales and answering my incredibly random questions even when posed at a time when Evil was in desperate need of defeat. Many thanks as well to Nicholas Wheeler for generously providing me with the original of the engine used in this study.

To my friends through it all, Kevin and Eva—you two are the loves of my life here. It feels weird to thank the best of friends for being the best of friends, so I'll settle instead for reiterating just how much I love you two. To my parents through it *all*, continued words of thanks just won't cut it. Everything that I have done, at Reed or elsewhere, within academia or without, has been made possible by you two, and I will never forget that. I love you both so, so much.

And lastly, a standing ovation to MartyParty for the MVP Purple Trap Mixtape, the soundtrack to my thesis and the best hourglass I know.

Table of Contents

List of Figures

Abstract

Presented here is a general analysis of the power output of the split-cylinder gamma-type Stirling heat engine. More specifically, this work offers an analysis of the validity of the mathematical model developed by Gustav Schmidt in 1871, and reproduced by J. R. Senft in 2002, in predicting the power output of such a Stirling engine given a set of test conditions and engine parameters [1, 2]. A stock Stirling engine manufactured by Solar Engines of Phoenix, AZ, was obtained and heavily modified to allow for manipulation of the engine parameters κ, the swept volume ratio, and α, the phase angle, so that their effects on output power, measured by means of a de Prony brake, could be independently determined and compared to theory. The mechanics of the Stirling engine is first described in detail, both conceptually and in terms of the underlying thermodynamic principles governing the Stirling cycle, before the model and its key assumptions are presented. The engine used is also thoroughly described, along with the modifications performed to it that enabled manipulation of κ and α. Due to a combination of insufficient data and the inherent difference between the *indicated* power predicted by the model and the *brake* power measured experimentally, the results of this study are largely inconclusive, yet do suggest that the Schmidt model can to a degree, and especially at low engine speeds, predict broad trends.

Introduction

Human beings are lazy. So lazy, in fact, that we routinely and ironically devote our cleverness to developing mechanisms that accomplish some goal in our stead. Out of laziness, one might then say, arose the engine, the literal heart enabling so many modern machines. From the internal combustion engines of our beloved motorcycles to the herculean gas turbines of jumbo jets and the miniscule electric motors of model airplanes, the ubiquitous engine permeates our lives almost without notice or due appreciation of its diversity. Most engines today are powered by electricity or by exploding gasses, but engines do exist that are powered by nothing more than the heat energy provided by some fuel source. Engines of this type, appropriately called heat engines, absorb heat energy and convert a portion of that energy into mechanical work, or force exerted over some distance, with the size of that portion dictated by the engine's efficiency [3, 4]. In the case of the steam engine, certainly the most well-known of all heat engines, water within a closed boiler absorbs heat energy from a fuel source, vaporizes, and builds pressure within the boiler. Through an intricate medley of pipes and valves, this pressurized water vapor is made to do work on a piston which in turn can be made to do work on crankshafts, wheels, gears, propellers, pumps—whatever is required. In short, heat energy is absorbed by some medium (water), and through an elaborate mechanical process some of this energy is converted into work, precisely the mission statement of the heat engine.

Even though it once dominated the world of engines, the steam engine does not have exclusive rights to the classification of heat engine. The Stirling engine, developed in 1816 by the Scottish minister Rev. Robert Stirling, also fits the heat engine prescription but in a manner more simple and elegant than steam [5, 4]. Perhaps the most curious difference is that while a steam engine needs a heat source to boil its water, a Stirling engine only requires a temperature *differential* between what are called the hot and cold spaces of the engine. Hence, an appropriately designed engine may in fact run off of a heat *sink* (or cold "source") instead of a heat *source*, provided said heat sink produces a sufficiently large temperature differential. Even more

remarkable is that by no means must this temperature differential be substantial; a high-quality, low-temperature engine can operate quite ethereally off of a differential on the order of a mere half-degree Celsius (0.5 C°) [6, 7]! Indeed, relatively inexpensive, low-temperature Stirling engines can easily be sourced online that run contentedly and indefinitely off the heat from one's own hand, or the heat-sapping capacity of a single ice cube.

The astounding low-temperature capability of the Stirling engine is not the only feature distinguishing it from the steam engine, however. Even higher-powered Stirling engines, capitalizing on a more extreme temperature differential to produce 150 horsepower or more, differ from steam engines in that they do not expel their working fluid, or the fluid which actually does work on a piston, during operation [5]. Gone are the iconic vapor clouds and perpetual hissing of the steam engine discharging steam, the working fluid, from its pistons for the Stirling engine contains its working fluid in a nearly hermetically sealed engine body. Hence, unless a burner is used as a heat source, the Stirling engine expels no exhaust [8]. And if a burner is used, as is typically the case, combustion is not confined to occur, for example, over the brief time periods allotted by, and in the often hypoxic environment of, the internal combustion engine and its cylinders. As a result, the *external* combustion allowed by the Stirling engine is continuous and complete, eliminating the residual unburned gas fumes and particulate matter characteristic of internal combustion engines [4]. Clean exhaust, near silent operation, no valves, timing belts, carburetors, fuel injectors or spark plugs—the Stirling produces captivating motion and output power seemingly out of nowhere.

But of course we know better. Power never comes from nowhere, and in the case of the Stirling engine it is derived from heat energy alone through the controlled, cyclic expansion and compression of the contained working fluid. Numerous physical models have been proposed to describe how the Stirling engine cycle produces this power, most notably that proposed in 1871 by the German professor Gustav Schmidt [1], then of the German Polytechnic Institute in Prague, and reproduced by James R. Senft of the University of Wisconsin [9, 2]. This model is designed to predict the power output of an idealized Stirling engine given a set of known parameters of the engine in question, and was chosen for this investigation due to the fact that it is the only notable model that gives a closed-form expression for power [9]. As such, it takes as its input a set of parameters derived from the physical dimensions of certain engine components, and predicts the power output of the engine as configured without requiring any computational analysis. Testing the accuracy of this model is

the primary focus of this thesis. To accomplish this goal, a stock Stirling engine, manufactured by Solar Engines of Phoenix, AZ, has been heavily modified in the Reed College machine shop to allow for manipulation of relevant engine parameters, specifically the swept volume ratio κ and the phase angle α, so as to compare their effects on power output to those predicted by the model.

Chapter 1

Physics of the Stirling Cycle

1.1 Conceptual Outline

As was mentioned earlier, the Stirling engine derives its power from the controlled, cyclic expansion and compression of the working fluid within the engine space. How this expansion and contraction is achieved is the genius of the Stirling engine, and is best understood in the context of the gamma Stirling engine shown in Fig. 1.1, one of three engine configurations. The other two have been imaginatively named alpha and beta engines respectively, and the differences between the three reside in the physical layout of a few key engine components. As alpha and beta Stirling engines teach us nothing of the Stirling cycle that can't be more easily elucidated with reference to a gamma engine, they will be largely ignored.

Regardless of configuration, all Stirling engines share a few common design features, most notably the presence of two pistons which, for a gamma engine, are referred to as the power piston, labeled "P" in Fig. 1.1, and the displacer piston, "D." As is characteristic of the gamma engine, the power piston resides in a separate cylinder from the displacer piston, though the two cylinders are connected such that the working fluid, generally air, can pass freely between them. The two pistons themselves are also connected, though indirectly, by any manner of mechanical linkage to a crankshaft on which a flywheel is mounted, whose purpose is to store rotational momentum to help smooth the engine's rotation. A possible linkage configuration, reminiscent of the connecting rods on the drive wheels of old steam locomotives, is shown in Fig. 1.2. Crucial to the Stirling cycle is that however this linkage is accomplished, it constrains the two pistons to move out of phase by some phase angle α, labeled as such in Fig. 1.2. Though Fig. 1.1 does not show this linkage, it does assume a phase angle of 90° so that as one piston is in the middle of its stroke, or its

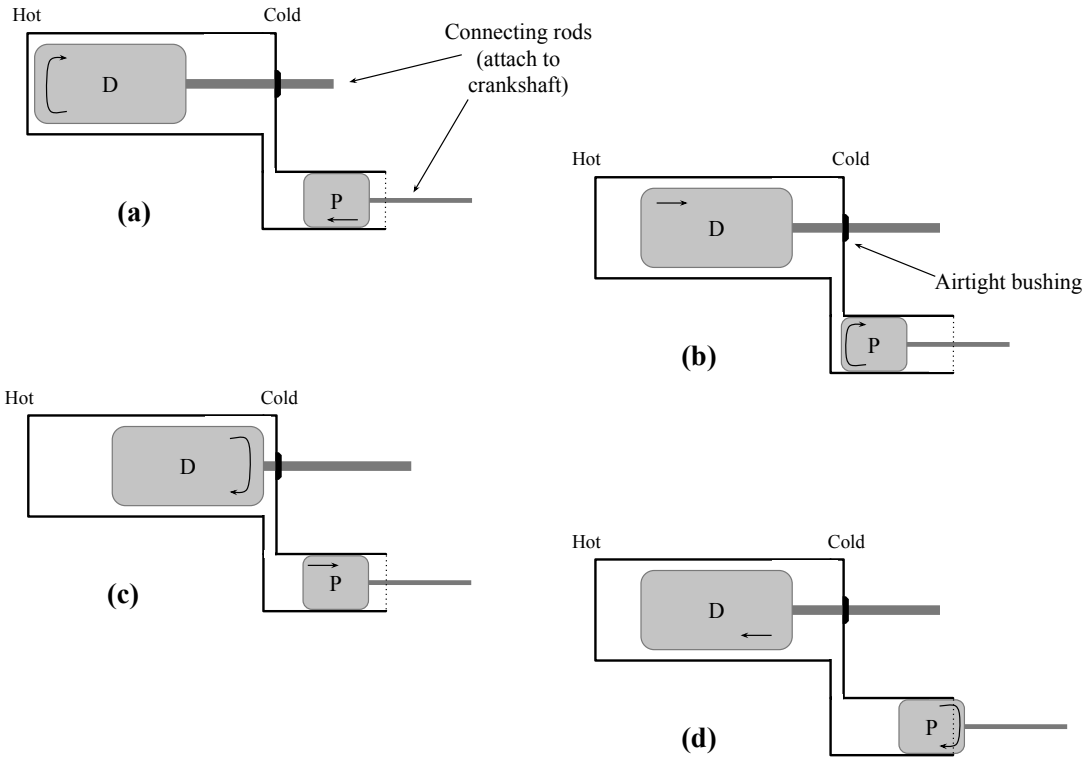

Figure 1.1: Schematic diagram illustrating the Stirling cycle as performed by a gamma-type stirling engine. The straight arrows drawn on the pistons indicate current direction of motion; the curved arrows indicate the piston is changing direction from the previous step. Note the gap between the displacer piston and its cylinder, through which air is shuttled between the hot and cold spaces.

total back-and-forth travel, the other is at the end.

Also crucial to the Stirling cycle is that the power piston be made to reside snugly within its cylinder, creating a nearly airtight seal, but for the displacer piston to fit loosely enough within its cylinder for air to pass freely through the resulting annular gap between the outer surface of the piston and the inner wall of the cylinder. As such, the displacer piston exists literally to displace air; wherever it is, air cannot be. Step (a) in Fig. 1.1 shows the displacer piston positioned all the way in the hot space of the engine, that portion exposed to the heat source, thereby displacing most of the air within the engine to the cold space, that portion exposed to the heat sink (atmosphere, liquid-cooled environment, etc.). Assuming the engine is running, flywheel angular momentum rotates the crankshaft, moving the displacer piston towards the cold space, displacing air through the annular gap and towards the hot space [step (b)]. As the air is shuttled from the cold to hot spaces, it absorbs heat energy and begins to expand. This expansion pushes on the power piston, forcing it out and beginning a power stroke of the engine [step (c)]. The pistons being out of phase by an angle α ensures that the air has time to heat up and expand before the power piston has reached the end of its stroke and is ready to begin a new cycle. It further ensures that before the power stroke is complete, the displacer piston begins moving back to the hot space, forcing air to the cold space in anticipation of the power piston beginning a compression stroke [step (d)]. As the hot air in the cold space releases energy to the heat sink, it contracts, and along with flywheel angular momentum moves the power piston back into its cylinder, carrying out the compression stroke and completing the Stirling cycle [back to step (a)].

1.2 Thermodynamic Analysis

In practice, the Stirling cycle is easy to implement. As suggested by the diagrams in Fig. 1.1, Stirling engines are mechanically very simple, requiring far fewer parts than the typical internal combustion or steam engine, and recreational hobbyists can easily build simple, working engines using household materials alone [5, 6]. In part due to this simplicity, Stirling engines tend to be impressively efficient, validating the thermodynamic theory of the Stirling cycle. Efficiency in the context of Stirling engines, and indeed *all* heat engines, regardless of how they work or their physical characteristics, is defined as the ratio of the work performed by the engine W to the heat absorbed Q_H, so that the efficiency e is given by $e = \frac{W}{Q_H}$ [3]. Hence, the efficiency of an engine indicates how much work one would expect an engine

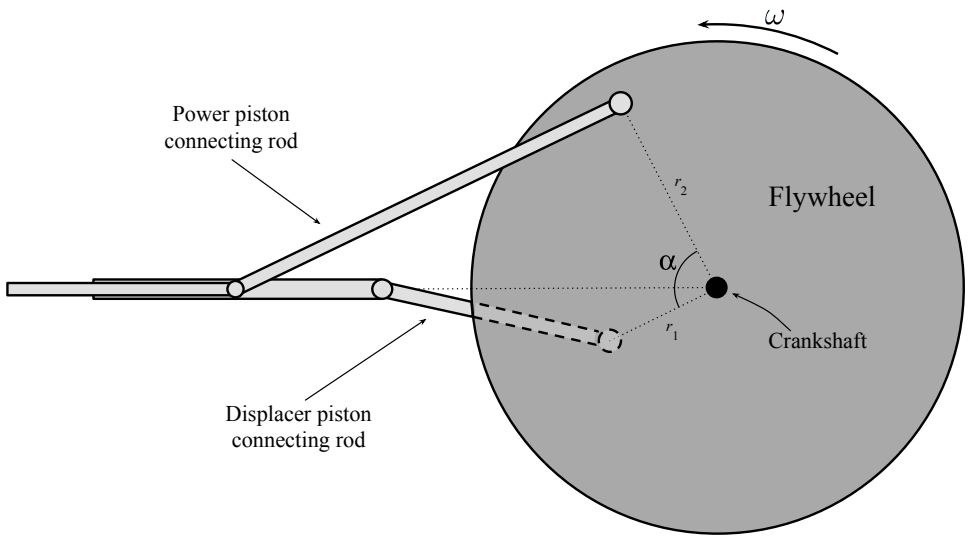

Figure 1.2: Side-view diagram illustrating a possible linkage configuration. The lower linkage rod is drawn with dotted lines to indicate it being attached to a *second* flywheel, mounted to the crankshaft behind and concentric to the one visible. The phase angle α is shown, and the end of the crankshaft is visible as a black circle in the center of the flywheel. The angular velocity ω of the flywheel is in a counterclockwise direction, meaning that the displacer piston leads the power piston by the angle α at all points in the cycle. In this configuration, where $r_1 < r_2$, the stroke of the displacer piston will be shorter than the stroke of the power piston.

to perform given a certain amount of input energy. Curiously, as W and Q_H are dependent only on the underlying thermodynamic processes of the Stirling cycle, the theorized efficiency of a Stirling engine does not depend on the physical characteristics of any real engine. The theoretical efficiency, therefore, does not take into account engine configuration (alpha, beta, or gamma), piston linkage configuration, piston size, total engine volume, or other such physical parameters. Evidently, the cycle to which the efficiency pertains is different from, though intimately related to, the one whose steps are sequenced in Fig. 1.1.

To understand the thermodynamics of the Stirling cycle as well as its efficiency, we do away with everything descriptive of a specific Stirling engine and focus entirely on the changes in pressure, volume, and temperature of the working fluid as it is cycled back and forth within the engine space. To simplify matters, we make a few assumptions. Firstly, we treat the working fluid as having the properties of an ideal gas, meaning that the individual gas molecules are identical, the volume occupied by the gas as a whole is much larger than the combined volume of the constituent molecules, the motion of each molecule is random yet in accordance with Newton's laws, no forces act on the molecules except during collisions with the container or other molecules, and these collisions are perfectly elastic, meaning no kinetic energy is lost in the process [10, 11]. A gas having these properties abides by the ideal gas law

$$pV = nRT, \tag{1.1}$$

where p is the pressure of the gas, V its volume, T its temperature (in Kelvin), and n and R are both constants. At this point, the salient conclusion of the ideal gas law is simply that the pressure, volume, and temperature of an ideal gas (or another gas modeled as such) are related by the simple constant nR. Hence, if the temperature of the gas increases, for example, so must the product pV. Similarly, if the volume increases so must the quotient T/p.

The ideal gas assumption is relatively innocuous with negligible drawbacks since real gases behave much like ideal gases so long as they are not put under conditions (of pressure, volume, and temperature) that would bring them near their condensation point, or the point at which a gas ceases to be a gas [12]. As the conditions within a Stirling engine are not at all conducive to the condensation of the working fluid, the ideal gas assumption can have few objections. What's more, it allows us secondly to assume that the working fluid expands (doing work on the power piston) and compresses isothermally, or without a change in temperature, but necessarily with a

change in pressure by the ideal gas law [13]. This assumption is difficult to realize in practice, but the alternative approach of accounting for non-isothermal expansion and compression processes produces diabolically complicated theoretical models that are beyond the scope of this paper [10] and can only be solved numerically [14].

Treating the working fluid as an ideal gas further allows us to assume that when in the hot space, the working fluid absorbs heat and increases in temperature isovolumetrically, or without any change in volume. Similarly, we assume that when in the cold space, the gas releases energy isovolumetrically [13]. This last assumption is somewhat problematic because in real engines, the working fluid will never heat or cool isovolumetrically. This is a necessary consequence of the (assumed) purely sinusoidal movement of the pistons; since the pistons are in constant motion (except instantaneously at rest at either end of their respective strokes), the volume of gas is constantly changing throughout all points of the real Stirling cycle [4]. Hence an isovolumetric change is not realistic. But as we are concerned with the ideal Stirling cycle, and its idealized efficiency, we make this assumption nonetheless.

We can use these assumptions of isovolumetric heating and cooling, and isothermal expansion and contraction, to represent the ideal Stirling cycle on a simple pressure-volume (p-V) diagram, with the pressure of the working fluid plotted against its volume. Moving radially away from the origin indicates an increase in temperature of the working fluid. Fig. 1.3 shows the p-V diagram for the idealized Stirling cycle, with the indicated steps roughly corresponding to those depicted in Fig. 1.1. At point (a), we find the working fluid to be at its highest volume of V_2 and lowest pressure of P_1, similar to step (a) in Fig. 1.1 where the power piston is extended resulting in a high volume, and the working fluid has been displaced to the cold space, producing a low gas temperature and pressure. Moving from point (a) to point (b) in Fig. 1.3, we find the gas isothermally decreasing in volume (to V_1) while increasing in pressure (to P_3), all the while staying at temperature T_C. This effect is caused by the power piston moving in and decreasing total engine volume, resulting in an increase in gas pressure. Notice, however, that between steps (a) and (b) in Fig. 1.1, the displacer piston also moves, shuttling some (but not yet all) of the cool air to the hot space, where it comes into contact with the heat source. In direct disagreement with our assumption of isothermal expansion and contraction, this *would* result in an increase in average gas temperature between steps (a) and (b) in either figure, though no such increase is evident in the p-V diagram.

Nevertheless, From point (b), the gas absorbs energy isovolumetrically as the displacer piston forces it to the hot space, increasing the gas pressure from P_3 to P_4

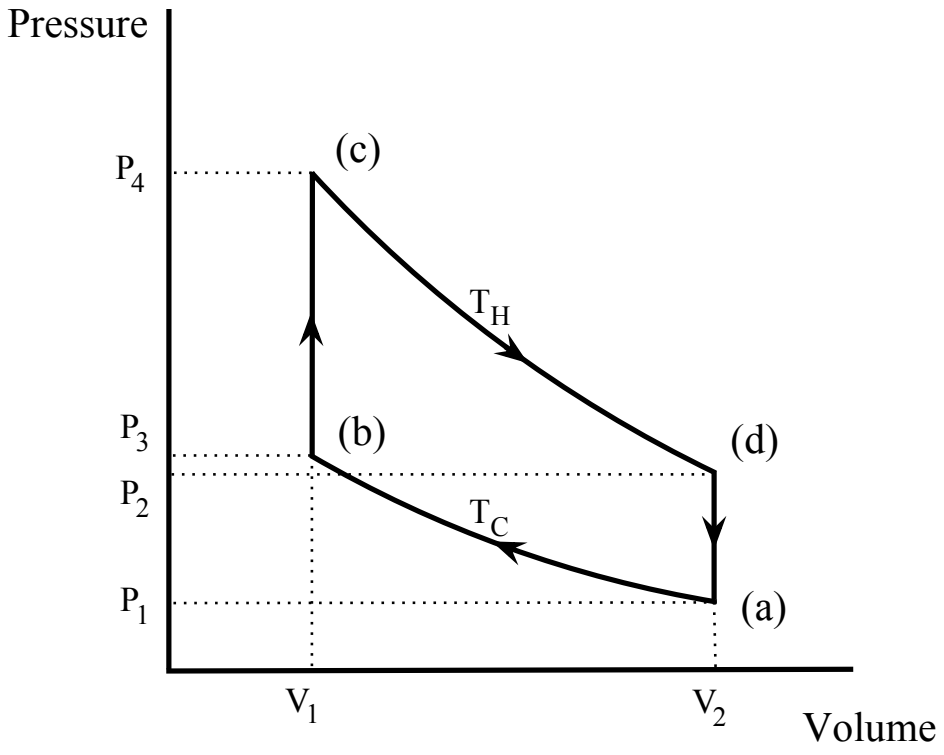

Figure 1.3: A pressure-volume diagram for the idealized Stirling cycle, with the arrows indicating the direction of the cycle. The transitions from (b) to (c) and from (d) to (a) are assumed to be isovolumetric ($\Delta V = 0$), and the transitions from (a) to (b) at temperature T_C and from (c) to (d) at temperature T_H are assumed to be isothermal ($\Delta T = 0$).

and bringing it to point (c) at a higher temperature T_H. Step (c) in Fig. 1.1 shows that as the gas absorbs energy in the hot space, the power piston is extending, simultaneously increasing total engine volume and exposing the flaw in the isovolumetric expansion/compression assumption. From point (c) to point (d), the gas moves along another isotherm, simultaneously decreasing in pressure to P_2 and increasing in volume to V_1. This is the power stroke of the engine, or the step where the working fluid does work on the power piston, forcing it outwards as in step (d) of Fig. 1.1. To complete the cycle, the gas, now in the cold space, undergoes an isovolumetric change in temperature and pressure, bringing it back to point (a) [or step (a) in Fig. 1.1] at temperature T_C.

Despite its loose correspondence to the Stirling cycle undergone by a real engine, the ideal Stirling cycle model, and its p-V diagram, are useful for determining the theorized efficiency of the Stirling engine. Recall that the efficiency of any heat engine is given by

$$e = \frac{W}{Q_H}\,, \tag{1.2}$$

where W is the work performed by the engine and Q_H is the absorbed heat. In the Stirling cycle, work is only performed during the two isothermal steps, where the general form for work done during such a process is given by

$$\begin{aligned} W &= \int_{V_i}^{V_f} p\,\mathrm{d}V \\ &= nRT \int_{V_i}^{V_f} \frac{1}{V}\,\mathrm{d}V = nRT \ln\left(\frac{V_f}{V_i}\right). \end{aligned} \tag{1.3}$$

Here we have used our assumption that the working fluid is an ideal gas, and the ideal gas law $pV = nRT$, to rewrite p as $p = \frac{nRT}{V}$, with p the pressure of the working fluid, T its temperature (in Kelvin), V its volume, n the number of moles of gas present, and R the ideal gas constant [15]. The terms V_i and V_f refer to the initial and final volumes of the gas during the isothermal process, so that with reference to the steps outlined in Fig. 1.3, we determine the work done by the gas on the piston between (c) and (d) to be

$$W_{cd} = nRT_H \ln\left(\frac{V_2}{V_1}\right),$$

and the work done by the piston on the gas between (a) and (b) to be

$$W_{ab} = nRT_C \ln\left(\frac{V_1}{V_2}\right) = -nRT_C \ln\left(\frac{V_2}{V_1}\right).$$

To determine Q_H for the engine, we need to consider when in the cycle the working fluid is absorbing heat from its surroundings. Clearly heat absorption occurs from (b) to (c) in Fig. 1.3, when the working fluid is undergoing an isovolumetric temperature increase, but also from (c) to (d), when the gas is isothermally expanding. Indeed, in order for the gas to expand but have its temperature remain constant, it *must* absorb energy from its surroundings [4]. Hence, we have that $Q_H = Q_{bc} + Q_{cd}$, with Q_{bc} and Q_{cd} being the heat absorbed during steps (b)-(c) and (c)-(d) respectively. Now, step (c)-(d) indicates an isothermal expansion, which by definition means that the internal energy U of the gas remains constant. Working from the first law of thermodynamics, which states that $\mathrm{d}U = \mathrm{d}Q - \mathrm{d}W$, we have that for an isothermal process $\mathrm{d}U = 0$ and hence $\mathrm{d}Q = \mathrm{d}W$ so that $Q_{cd} = W_{cd}$ [16]. Since (b)-(c) is isovolumetric, no work is done on the gas, hence $\mathrm{d}W = 0$ and the first law reduces to $\mathrm{d}Q = \mathrm{d}U = mC_v\Delta T$, where m is the mass of the gas, C_v is its specific heat capacity, and ΔT is the change in temperature undergone during the process [12]. Therefore, $Q_{bc} = mC_v\,(T_H - T_C)$, and Eq. (1.1) becomes:

$$\begin{aligned} e = \frac{W_{ab} + W_{cd}}{Q_{bc} + Q_{cd}} &= \frac{W_{ab} + W_{cd}}{Q_{bc} + W_{cd}} \\ &= \frac{nR\,(T_H - T_C)\ln\left({}^{V_2}\!/_{V_1}\right)}{mC_v\,(T_H - T_C) + nRT_H \ln\left({}^{V_2}\!/_{V_1}\right)}. \end{aligned}$$

Making the substitutions $\beta = {}^{T_C}\!/_{T_H}$, $\gamma = {}^{V_2}\!/_{V_1}$, and $mC_v = {}^{1}\!/_{2}\,\eta R$, with the integer η indicating the degrees of freedom of a molecule of the gas, our expression for e simplifies to

$$e_{\text{stirling}} = \frac{1-\beta}{1 + \frac{\eta\beta}{2\ln(\gamma)}}. \tag{1.4}$$

Evidently, then, the efficiency of a Stirling engine is entirely dictated by the ratio of expanded gas volume to compressed gas volume (the compression ratio γ), the ratio of the heat sink temperature to the heat reservoir temperature (the temperature ratio β), and by the chemical properties of the working fluid (whose molecular structure dictates η). Note that Eq. (1.2) does not consider frictive or other mechanical losses

of efficiency. But even so, since γ, β, and η are all real numbers greater than zero (volumes are positive-valued so $\gamma > 0$, temperature is measured in Kelvins so $\beta = {}^{T_C}/_{T_H} > 0$ for all values of T_C and T_H, and η is defined as a positive integer), it follows from Eq. (1.2) that $e < 1$ for all V_1, V_2, T_C, T_H, and η. Even without considering heat losses or mechanical friction, Eq. (1.2) predicts a maximum efficiency always lower that 100%, a frustrating and unavoidable consequence of the second law of thermodynamics and the hugely important concept of entropy [3].

Colloquially put, entropy is a measure of the amount of disorder in a system. Hence, a cluttered desktop is in a higher state of entropy than the same desktop whose clutter has been organized into stacks and piles. Perhaps a more relevant example pertains to the distribution of energy among gas molecules in a sealed container (the interior of a Stirling engine, perhaps?): such a chamber containing a large number of energetic molecules on one side and very few on the other is in a lower state of entropy than if the energy in the gas were dispersed evenly among all constituent molecules. We observe the latter case in nature because it is a state of higher probability than the former [15]. Imagine allowing a few billion gas molecules, half having some high energy and the rest having a lower energy, to suddenly and freely enter a closed, evacuated chamber and redistribute themselves randomly for some amount of time. The likelihood of finding, at the end of this redistribution period, all the high-energy molecules on one side of the chamber and all the low-energy molecules on the other is truly miniscule, negligible; vastly more favorable is some relatively uniform distribution of high- and low-energy molecules, and this, appropriately, is what we observe in nature.

The facts that every potential state of the components of a thermodynamic system, such as the gas molecules in a Stirling engine, has a corresponding statistical probability of existing, and that nature tends towards those states having high probability, introduce a definite 'arrow of time' into the evolution of thermodynamic systems. Such systems will tend to evolve in time according to which potential states are most likely. And since the entropy of a state is directly related to its relative probability of existing (relative to the other available states, that is), we deduce that thermodynamic systems *always* evolve in time in such a way as to increase the entropy of the system and its surroundings [3].[1] This conclusion, in fact, is the most general form of the second law of thermodynamics, and applies directly to Stirling engines when we realize that the (ideal) Stirling engine is precisely one of the thermodynamic systems

[1] This statement, generally taken to be true, has nevertheless successfully been challenged by I. Prigogine and T. Petrosky in [17].

to which the second law pertains. Since a Stirling engine operates in a cycle, starting in one state, evolving in time, and ultimately returning to that same state, the entropy of the gas in the engine cannot change over one complete cycle. However, according to the second law, the entropy of the engine *and its surroundings* will increase over time. If we consider the heat source and heat sink as part of the surroundings (reasonable because a heat source isn't really part of the engine itself, though it acts upon it, and is at least distinct and physically separated from the working fluid), we may ascertain that entropy is produced during the heating and cooling of the working fluid [3].

To see where this entropy comes from, we must turn to a more precise and technical (yes, mathematical) interpretation of entropy. Defined in terms of the infinitesimal quantity of heat $\mathrm{d}Q$ and the temperature T, we have that the *change* in entropy ΔS of a system is given by

$$\Delta S = \int \frac{\mathrm{d}Q}{T} \geq 0\,, \tag{1.5}$$

where the integral is taken along the path of change (moving, for instance, from Q_a to Q_b), and where the condition $\Delta S \geq 0$ is required by the second law of thermodynamics. In real engines, such as that diagramed in Fig. 1.1, as the displacer moves from the hot space to the cold space, the air that is consequently forced to the hot space is at a lower temperature T_C than the the heat source itself, running at a balmy T_H. Hence by Eq. (1.5) the entropy *lost* by the heat source as it gives up a total amount of heat heat Q_H to the cold(er) working fluid is simply $\Delta S_{\text{lost}} = {}^{Q_H}\!/_{T_H}$, whereas the entropy *produced* during the heat exchange is $\Delta S_{\text{prod}} = {}^{Q_H}\!/_{T_C}$. But since $T_C < T_H$, we have that $\Delta S_{\text{prod}} > \Delta S_{\text{lost}}$, meaning that the entropy produced during the heat exchange is greater than the entropy lost [3]. But looking at Eq. (1.5), we see that if we integrate over a complete *cycle*, for instance from Q_1 back to Q_1, we find that $\Delta S = 0$. So for one complete Stirling cycle [point (a) to point (a) in Fig. 1.3], we should have that $\Delta S_{\text{cycle}} = 0$ [11]. But so far we have a net *increase* in entropy $\Delta S_{\text{net}} = \Delta S_{\text{prod}} - \Delta S_{\text{lost}} > 0$, so satisfying Eq. (1.5) requires the expulsion of entropy equivalent to ΔS_{net} during the cooling stage of the cycle. The result is that the engine must dump some amount of waste heat into its surroundings during the cooling stage. This waste heat, as the name implies, is what constrains real heat engines to run at less than 100% efficiency: a significant amount of the absorbed heat Q_H passes through the engine without being converted into mechanical work.

1.3 Maximizing Efficiency and The Carnot Limit

We know from the second law of thermodynamics, and consequently from Eq. (1.5), that no heat engine (or engine in general, for that matter) can be perfectly efficient under any condition. This constraint then begs the question: how efficient can a heat engine be, and how will such a maximally efficient engine operate? As suggested in the previous section, the game we play in designing a maximally efficient engine is one of minimizing the production of entropy. To do so, we first focus on the expansion and compression processes inherent to any heat engine, and try to minimize the entropy produced during expansion, thereby also minimizing the entropy (manifested as waste heat) that must be expelled during compression. We know that the entropy lost by the heat source during expansion is Q_H/T_H, whereas the entropy gained by the gas is Q_H/T_{gas}, where T_{gas} is the temperature of the gas immediately prior to heating. The entropy produced during expansion is simply the difference between the two such that

$$\Delta S_{\text{expand}} = \frac{Q_H}{T_{\text{gas}}} - \frac{Q_H}{T_H}\,. \tag{1.6}$$

Now if $T_{\text{gas}} = T_H$, we see immediately that Eq. (1.6) reduces to zero, indicating that no new entropy has been produced during the heating stage under this condition. We seem to have satisfied our goal of minimizing increases in entropy by setting $T_{\text{gas}} = T_H$, but unfortunately such a temperature condition precludes any heat transfer between the heat source and the working fluid, heat transfer without which expansion cannot occur and the engine cannot run at all [3]. So instead, we make T_{gas} *infinitesimally* less than T_H so that heat transfer can still occur. By letting the gas expand infinitesimally slowly during heat exchange, we can keep T_{gas} both constant and infinitesimally less than T_H, thereby keeping $\Delta S = 0$ during the expansion step, as desired. To keep T_{gas} constant, we require the gas to expand isothermally during the heating stage. Similarly, during the compression step when the gas is releasing energy to the heat sink, we want to keep T_{gas} constant and infinitesimally *higher* than T_C to avoid producing new entropy while still allowing for heat exchange. These conditions mandate an isothermal compression.

Though the expansion and compression processes are most susceptible to producing new entropy, the intermediate stages analogous to steps (b)-(c) and (d)-(a) in Fig. 1.3, are still avenues for potential entropy production. After all, entropy is a function of temperature [see Eq. (1.5)], and during these two remaining steps the gas temperature does change. To minimize the entropy produced during these steps, we need the gas to move from T_H to T_C (more correctly, from just below T_H to just

above T_C) and back without heat being transfered to or from the gas. That is, we want $\Delta Q = 0$ during these processes, meaning that the gas must move adiabatically, or without heat transfer, between the two temperature extremes [16]. This completes the cycle: we have an isothermal expansion and heat transfer at T_H, then an adiabatic expansion dropping T_{gas} to T_C, then an isothermal compression at T_C, and finally an adiabatic compression bringing T_{gas} back to T_H. These are the four steps of our entropy-minimizing cycle, known as the Carnot cycle, proposed by the French scientist Sadi Carnot in 1824, and are shown on a p-V diagram in Fig. 1.4 [3].

The natural question at this point is, well, how did we do? That is, how efficient *is* the Carnot cycle? Proceeding from Eq. (1.2) by following an argument similar to that which led to Eq. (1.4), we determine that the efficiency of an engine operating in a Carnot cycle e_{carnot} is given by

$$e_{\text{carnot}} = 1 - \frac{T_C}{T_H}\,. \tag{1.7}$$

This is the famous Carnot efficiency, and is curiously dependent *only* on the temperatures of the heat source and heat sink. This surprising result allows us to alter the efficiency of a Carnot engine simply by manipulating the temperature ratio T_C/T_H; decreasing T_C or increasing T_H, or both, will serve to increase the efficiency of the engine. Furthermore, we note that e_{carnot} approaches 1 (or 100%) as T_C approaches zero, but as T_C cannot actually be zero (remember, temperatures here are measured in degrees Kelvin, so $T_C = 0$ implies T_C is at *absolute* zero), we find that $e_{\text{carnot}} < 0$ no matter what. Again, as always, the second law of thermodynamics exerts its dominance.

Even though it still predicts efficiencies resolutely less than 100%, the Carnot efficiency in fact places an upper bound on the efficiency of any heat engine. That is, for any heat engine, be it a Stirling engine, steam engine, or other, its efficiency is less than that of a Carnot engine operating under the same heat source and heat sink temperatures [15]. Indeed, by looking at Eq. (1.4), which gives the efficiency of a Stirling engine, we see that the numerator is precisely the Carnot efficiency, and that the denominator is always a number greater than 1. Hence, $e_{\text{stirling}} < e_{\text{carnot}}$ for all temperatures T_C and T_H.

1.4 Heat Regeneration and Recovering Carnot Efficiency

The reason Stirling engines are less efficient than Carnot engines is that during the two isovolumetric heat transfer stages, not all the heat absorbed or expelled comes from heat sources or heat sinks at temperature T_H or T_C, respectively [18]. Since heat transfer cannot occur instantaneously, we can think of the isovolumetric heat exchange processes as though the gas were being exposed to a series of progressively warmer (or cooler) thermal reservoirs, with the transition between each having an associated entropy change. The closer the temperature of one of these reservoirs is to T_{gas}, the less efficient the heat transfer between the gas and the reservoir [18]. As the Carnot cycle assumes only two thermal reservoirs, it escapes these efficiency losses entirely.

In order to improve the efficiency of the Stirling engine, a beautifully clever device known as as a regenerator can be incorporated into the engine design which essentially takes the place of those extra thermal reservoirs implicit in the isovolumetric heat exchanges. A regenerator is any device which can store waste heat from the cooling stage for use in the subsequent warming step, thus minimizing heat loss to the environment [14]. Typically the regenerator takes the form of a wire mesh or some other porous, thermally conductive medium through which the working fluid flows in moving between the hot and cold spaces of the engine. Fig. 1.5 offers a modification to the engine schematic presented in Fig. 1.1 whereby a regenerative channel has been included between the hot and cold spaces of the engine. Notice that the diameter of the displacer piston "D" has been increased, eliminating the annular gap between the displacer and its cylinder, so that the displacer now forms an airtight seal within its cylinder. Consequently, as the displacer piston moves through its stroke, the working fluid is forced from one side of the engine to the other *through* the regenerative channel and the regenerator itself. Air coming from the hot space warms the regenerator as it moves to the cold space, whereupon the air cools the regenerator as it returns back to the hot space, thus establishing a temperature gradient across the regenerator. That is, this process of being alternatively heated and cooled eventually makes the temperature of one end of the regenerator (the end nearest the hot space) close to T_H, and the other end close to T_C.

The effect of this established temperature gradient is that as the cool working fluid is forced to the hot space, its temperature T_{gas} gradually increases from T_C to just slightly less than T_H (call it $T_{\text{gas}} = T_H - \epsilon$, with $\epsilon \ll T_H$) so that when the gas is

absorbing heat from the heat source, the entropy produced, by Eq. (1.6), is simply

$$\begin{aligned}\Delta S_{\text{heat}} &= \frac{Q_H}{T_H - \epsilon} - \frac{Q_H}{T_H} \\ &= \frac{\epsilon\, Q_H}{T_H\,(T_H - \epsilon)}\,,\end{aligned}$$

which approaches zero as ϵ goes to zero. Similarly, as hot air moves to the cold space, it gives off heat to the regenerator so that by the time it enters the cold space its temperature is $T_C + \epsilon$ so that only minimal heat (entropy) must be expelled as waste. If this regenerator functions perfectly, that is, captures *all* the heat given off by the gas as it cools to T_C, and returns *all* that stored heat, without loss, to the gas during the heating stage, then $\Delta S_{\text{heat}} = \Delta S_{\text{cool}} = 0$, and we have, in effect, reduced the Stirling cycle to a Carnot cycle by eliminating all thermal reservoirs except those two at T_H and T_C [18]. The adiabatic heating and cooling processes of the Carnot cycle have been replaced by corresponding isovolumetric transitions during which perfect regeneration ensures that $\Delta Q = 0$, or that no heat is lost to, or absorbed from, the surroundings, precisely the condition for an adiabatic transition [18]. Hence, under the condition of perfect heat regeneration, the efficiency of the Stirling cycle reduces to that of the Carnot cycle [10, 19].

While perfect regeneration is impossible to achieve in practice, partial regeneration is easily implemented and is often implicit in many engine designs. For instance, even the simple engine diagramed in Fig. 1.1, and, indeed, the actual engine used in this study, utilize an albeit simplistic regenerator: the very displacer cylinder itself. The annular gap between the displacer piston and the inner cylinder wall serves as the regenerative channel, and the metal walls of the cylinder as the regenerator itself. As air is forced back and forth through the annular gap, a temperature gradient is set up between the hot and cold ends of the *cylinder* just as one was established between opposite ends of the *regenerator* in Fig. 1.5, with more or less the same result: the gas is warmed before it reaches the hot space by the very heat energy it transfered to the regenerator in moving to the cold space. The regenerative effect may be increased by shrinking the size of the annular gap, thereby forcing more air into contact with the cylinder wall when passing around the displacer, and increasing total heat transfer to and from the gas.

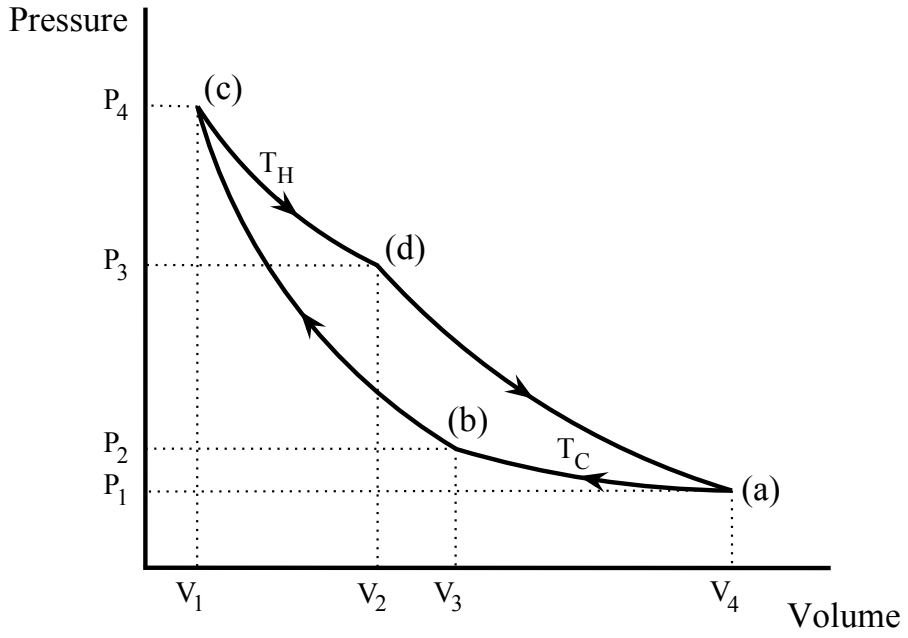

Figure 1.4: A pressure-volume diagram for the Carnot cycle, with the arrows indicating the direction of the cycle. The transitions from (b) to (c) and from (d) to (a) are adiabatic ($\Delta Q = 0$), and the transitions from (a) to (b) at temperature T_C and from (c) to (d) at temperature T_H are isothermal ($\Delta T = 0$).

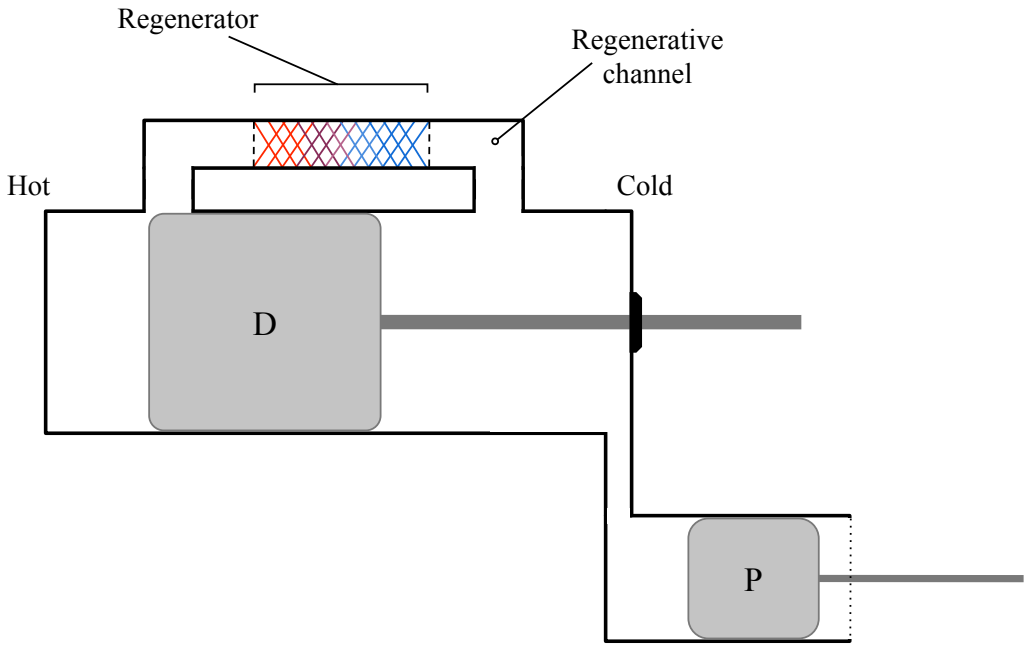

Figure 1.5: A schematic diagram illustrating a possible means of implementing a regenerator into a gamma-type Stirling engine. The regenerator itself is assumed to consist of some porous, thermally conductive material such as wire mesh, and resides within a regenerative channel connecting the hot and cold spaces of the engine. As the displacer piston shuttles the working fluid back and forth, it passes through the regenerator, alternately heating and cooling it, thereby establishing a temperature gradient illustrated as a color spectrum.

Chapter 2

Power Output and Predictive Modeling

So far we have largely been considering the efficiency of the Stirling engine, an approach that has been useful for exploring and giving a thermodynamic formalism to the Stirling cycle, and for introducing the core concept of entropy. However, our analyses and conclusions to this point are insufficient for elucidating the central question of this thesis: How much power can a Stirling engine produce, and how well can this value be predicted? To begin tackling this question, we leave behind our discussions of efficiency and we abandon our conception of the Stirling engine as carrying out a repeating cycle of thermodynamic processes (not that this isn't true). Instead, we develop a mathematical formalism for describing changes in the pressure of the working fluid during engine operation, and from this determine how much work the fluid does on the pistons, ultimately allowing us to derive expressions predicting how much power the engine can produce. Unlike the methods used to derive the efficiency of the Stirling engine, this method for predicting power output *does* take into account actual physical properties of the engine, namely relevant dimensions such as piston stroke, piston cross-sectional area, and phase angle. This is a marked transition as these properties played no role in the purely thermodynamic analysis of Chapter 1, which treated the engine as some arbitrary entity whose sole requirement was to *somehow* perform the Stirling cycle outlined in Fig. 1.3. That is, our methods so far have completely ignored everything pertaining to how the engine actually works: its layout, total contained volume, the sizes (and number) of its pistons, its compression ratio, phase angle, even the magnitude of the temperature differential to which the engine is subjected [3]. Bringing these quantities (and others, as shall be seen) back into the mix means that while the equations we derive for predicting power output can

apply to all Stirling engines, by specifying values for these relevant quantities we can collapse our equations into describing the power output of a single, *real* engine. The thermodynamic analysis of before could not accommodate this degree of specificity as it treated all Stirling engines as literally identical.

2.1 Defining Power

When we speak of the power output of an engine, what we're in truth referring to is how much work that engine can perform in a given amount of time. That is, power P is defined as the rate at which work W is performed, giving

$$P = \frac{\Delta W}{\Delta t}, \tag{2.1}$$

where Δt is the amount of time, in seconds, taken to perform an amount of work ΔW [12]. More broadly, power can be defined as the rate at which energy E is transformed, yielding the more general expression $P = \frac{\Delta E}{\Delta t}$ [20]. But the change in energy ΔE is simply the work W performed, if any, plus the heat H transferred, if any. As the Stirling engine does not induce or perform any heat transfer (though it certainly *depends* on heat transfer to operate), we have that $H = 0$ and $\Delta E = \Delta W$, reducing the above general expression for power to the more specific case of Eq. (2.1).

For the Stirling engine, power is extracted at the crankshaft, meaning that the crankshaft is the component of the engine used to do work on some other body: a generator, a fan, a gearbox, etc. As the engine makes the crankshaft rotate, we must be dealing with rotational power stemming from the performance of rotational work, or work arising from a force being exerted through some angular span. But in the context of rotational dynamics, the concept of force needs a slight revision as not all forces exerted on a body that is free to rotate will actually induce a rotation. Imagine trying to rotate a bicycle wheel by pushing on the tire in a direction parallel to the axle. In this case, the force applied won't induce rotation since no component of the force acts in a direction in which the wheel is free to rotate. Furthermore, imagine trying to rotate the same wheel by pushing on the tire radially inwards towards the axle. This method, too, will fail, this time because the force applied has no leverage over the wheel. To induce rotational motion, then, the force applied to a body that is free to rotate must have a component perpendicular to the axis of rotation, and this component must act at a point on the body some nonzero distance perpendicularly away from the axis of rotation [12]. This nonzero distance constitutes the lever arm,

and along with the magnitude of the applied force, determines the angular acceleration α resulting from that force. The longer the lever arm for a given force, the larger the angular acceleration, with the classic example of this phenomenon involving the opening of a door: the further away from the hinges one pushes, the faster the door will open (angularly accelerate) for a force of a given magnitude and direction.

So we see that the angular acceleration of a rotating body is dependent on both the magnitude of the applied force and the magnitude of the lever arm, and is in fact proportional to the product of the two, a product known as a torque τ. For a given lever arm R at whose end a force, having a component $F_\perp$ perpendicular to the axis of rotation, is exerted, the resulting torque is given by

$$\tau = RF_\perp, \tag{2.2}$$

so that the angular acceleration $\alpha \propto RF_\perp$. As a torque is what induces an angular acceleration, it is seen as the rotational analogue to the force F in Newton's second law $F = ma$, which induces the *linear* acceleration a. Hence, work in a rotational sense is dependent on the torque instead of directly on the force F, and, if the torque is constant, is equal to that torque times the angular sweep θ through which the torque acts. That is, for a constant torque, $W_{\text{rotation}} = \tau\theta$. Rotational power then becomes the rate at which this rotational work is performed, indicating that for a constant torque τ, the power associated is just

$$\begin{aligned} P = \frac{\mathrm{d}W}{\mathrm{d}t} &= \frac{\mathrm{d}\tau}{\mathrm{d}t}\theta + \tau\frac{\mathrm{d}\theta}{\mathrm{d}t} \\ &= 0 + \tau\frac{\mathrm{d}\theta}{\mathrm{d}t} = \tau\omega, \end{aligned} \tag{2.3}$$

with ω being the angular velocity of the body, or simply a measure of how quickly it is rotating [12]. What Eq. (2.3) tells us is that if we know (or can measure) how much torque is being produced by our Stirling engine, or rather how much torque the engine is exerting on the flywheel, and if we further know how quickly that torque is causing the crankshaft to rotate, we can compute the output power of the engine at that engine speed ω. This is a crucial result and will be used later on.

2.2 Terminology of the Senft Model

We now know how the Stirling engine produces power: work is done by the power piston exerting a torque on the crankshaft (with, incidentally, the lever arm for the torque equal to the radial distance from the crankshaft to the connecting rod mounting point, or half the piston stroke) causing it to rotate; the rate at which this work is performed is the power generated. As simple as Eq. (2.3) is, it masks the practical difficulties associated with measuring the output power of an engine (these difficulties will be outlined in lavish detail and with no small degree of palpable frustration in Chapter 4), as well as the mathematical complexity involved in predicting the output power. Despite the complexity, numerous predictive models, most of them necessarily computerized and having only numerical solutions, have been developed to do just this. The model presented here, an iteratIon of the first mathematical model describing the power output of the Stirling engine proposed in 1871 by Gustav Schmidt, is special in that it offers a closed-form expression for engine power that can be solved analytically with well-known methods. The Schmidt model as presented by Senft in [2] depends on three quantities determined by the physical dimensions and layout of the engine itself: the dead volume ratio (and dead volume itself), the swept volume ratio, and the phase angle α. In the following sections, these quantities will be outlined with reference to the specific engine used during testing.

2.2.1 The Phase Angle α

Out of these three key quantities, the phase angle α is perhaps the simplest to conceptualize as it only refers to the degree to which the displacer and power pistons are out of sync. In practice this translates into the two pistons being in different stages of their respective strokes at any given point in time. Consider, for instance, the common bicycle, a "curious vehicle" because "it's passenger is its engine" [21]. In the case of this superior vehicle, the pedals and crank arm assembly constitute the crankshaft of the passenger-as-engine, and the angle between the two pedals constitutes the phase angle between the legs-as-pistons. That is, treating the cyclist as an engine and her legs as pistons, we note that her piston-legs are constrained to be out of sync, or out of phase, by 180°, meaning that when one foot is at the top of the stroke (pedal at 12 o'clock), the other is at the bottom (pedal at 6 o'clock). The reason for this phase difference is to smooth out power transfer. Imagine the jerky sensation that would result from both legs moving perfectly in phase: the power stroke, when both legs are pushing down on both pedals, would be ridiculously powerful compared to the shared

return stroke, causing sporadic acceleration. Instead, with both legs perfectly out of phase, the weak return stroke of one leg is compensated for by the power stroke of the other, resulting in relatively constant power transfer to the wheels.

The phase angle on the Stirling engine, however, is not present to smooth the power transfer from the power piston to the crankshaft. Instead, it exists to allow time for the working fluid to expand when in the hot space and contract when in the cold space before the effects of this expansion or contraction are utilized by the power piston. Heat transfer from the heat source to the working fluid, or from the working fluid to the heat sink, cannot happen instantaneously, and hence these processes must be given time to occur before the power piston begins its power or compression strokes respectively. The phase angle, along with the engine speed ω, determine the allotted time period for expansion or compression, so that for a given engine speed, the larger the phase angle α, the more time the working fluid has to undergo heat transfer before the power piston is in a position to capitalize on the resulting expansion or compression.

With this in mind, it may be tempting to make the power and displacer pistons move 180° out of phase for this configuration would maximize the amount of time available for heat transfer both to and from the working fluid. However, this configuration would result in the power piston beginning its compression stroke with the working fluid still in the cold space, so that only when the power piston has reached the end of its stroke has the displacer piston managed to force the working fluid entirely to the hot space where it can *begin* to heat up. Hence the power piston will begin its subsequent compression stroke while much of the working fluid is still busy *expanding*. Similarly, the power piston will begin its power stroke while much of the working fluid is in the process of dissipating heat energy to the heat sink, thereby undergoing compression, or the exact opposite of what we need to induce a power stroke. Setting $\alpha = 90°$ then seems like a reasonable compromise, and as it turns out this is a common phase angle for many Stirling engines, and is the optimal angle for some.

2.2.2 The Swept Volume Ratio κ

The swept volume ratio, which we shall follow Senft [2] in calling κ, refers to the ratio of the volumes swept out by the power and displacer pistons. We will follow the Schmidt convention by more specifically defining κ as the ratio of the volume swept out by the power piston to that swept out by the displacer piston. Fig. 2.1 provides

a visual of what a piston's swept volume is for the specific case of the displacer piston of the engine used throughout this study. As seen in the figure, a pistons swept volume is simply its cross-sectional area multiplied by its stroke, or the total volume of working fluid the piston can move out of the way in the course of its stroke. In the automotive and motorcycle industry, when one talks of an engines displacement, one is actually referring to the total swept volume of its pistons, *not* the total volume of the engine cylinders as one might be led to think.

Despite how mathematically simple the swept volume, and the swept volume ratio, are, Fig. 2.1 nevertheless highlights a few aspects of the displacer piston of the engine used here that could cause confusion in understanding its swept volume. The first is that the piston is entirely enclosed within its cylinder, most unlike the power piston or pistons of motorcycle engines, all of which close off their respective cylinders at one end like the plunger of a syringe. The fact that the displacer piston resides entirely within its cylinder invites the question of whether we are to include the total volume of the piston itself in calculating its swept volume. We certainly don't in the case of the power piston, for from the point of view of the working fluid within the power cylinder, the piston appears as nothing more than a circular plane moving back and forth, having no length or volume. But in the case of the displacer in Fig. 2.1, the piston itself occupies a volume within the engine space, so it seems reasonable to suppose that it sweeps out a volume equal to its cross-sectional area multiplied by the sum of its length *and* its stroke. To do so would be incorrect, however. Taking swept volume to mean, again, the total volume of working fluid the piston can move out of the way in the course of its stoke, we note that the volume enclosed by the displacer piston itself *cannot* factor into its swept volume as the working fluid is permanently barred from that volume by virtue of it being enclosed, and hence can't ever be displaced from it. In other words, the volume enclosed by the piston itself is never available to the working fluid, regardless of the position of the piston along its travel. Instead of including that volume in a calculation of the displacer pistons swept volume, we instead need to eliminate it from the total volume of working fluid in the engine at any given point in time.

The second characteristic of the displacer piston that makes determining its swept volume difficult is the presence of the piston rod to which the displacer is attached. This rod complicates the swept volume of the displacer piston for a similar reason as the piston being fully enclosed within its cylinder: part of the rod resides within the engine space at all times and at all points along the stroke, and hence occupies a volume never available to the working fluid. If we take the piston in Fig. 2.1 to be

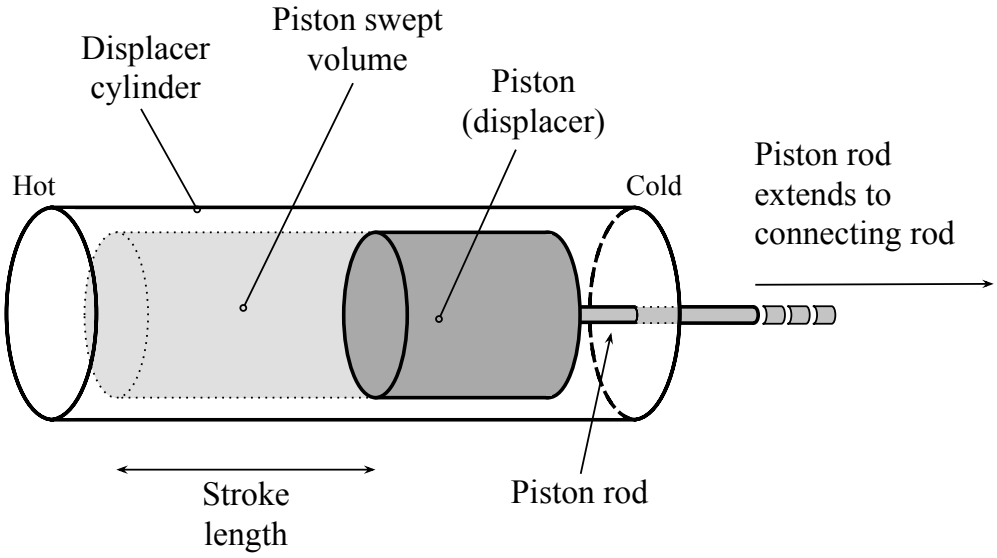

Figure 2.1: A diagram illustrating the swept volume for the type of displacer piston found in the engine used in this work. Note, particularly, the piston rod which joins to the displacer connecting rod at one end, and, passing through an airtight bushing, enters the engine space and fixes to the displacer piston itself at the other. This rod sweeps out a non-negligible volume in much the same way the piston itself does. The swept volume for the displacer piston is simply the cross-sectional area of the piston multiplied by the length of the stroke (the length of the piston, perhaps surprisingly, is not relevant). The swept volume for the rod is its cross-sectional area multiplied by the difference between the length of the stroke and the length of rod remaining in the chamber when the piston is at the top of its stroke (as it is in the figure).

at the top end of its stroke, we see that even with most of the piston rod extending out of the engine space, some small amount always remains within. This remaining nub is needed to prevent the piston from colliding with the end of its cylinder at the extreme top of its stroke. The volume occupied by the nub is analogous to the volume occupied by the piston itself in that it constitutes a volume never available to the working fluid at any point in the engine cycle, and hence can't factor in to a calculation of the rods swept volume. As with the piston itself, then, the swept volume of the rod is simply its cross-sectional area multiplied by the stroke of the piston. It follows that the swept volume for the entire displacer piston assembly is the stroke of the piston multiplied by the sum of the cross-sectional areas of the piston and piston rod; the volumes of the the nub and piston itself are not relevant.

2.2.3 The Dead Volume Ratio χ

The dead volume of an engine is a very dramatic term for the volume within the engine that is not swept out by the pistons. Hence it is equal to the total engine volume available to the working fluid (so this does not include the volume of the displacer piston or of the small nub of piston rod remaining in the engine space at the top of the stroke) minus the total swept volume of the pistons. The dead space ratio χ is defined by Senft in [2] as being the ratio of the dead space volume to the swept volume of the displacer piston. The descriptor 'dead' indicates that this volume is not actively traversed by the pistons. In the case of the gamma Stirling diagrammed in Fig. 1.1, the dead volume includes all the volume in the annular gap, the volume of the chamber connecting the power cylinder to the displacer cylinder, the volume remaining in the hot space in step (a), the volume remaining in the cold space in step (c), and the volume remaining at the bottom (left-hand side) of the power cylinder in step (b). In internal combustion engines, the dead space volume is the volume of the combustion chamber, or the volume in the cylinder when the piston is at the top of its stroke, the gas in the cylinder is maximally compressed, and the spark plug is about to fire, beginning a power stroke. The compression ratio, or ratio of the expanded volume to compressed volume (V_2/V_1 from Fig. 1.3), for the internal combustion engine can then be written in terms of the dead space volume as

$$\text{compression ratio} = \frac{\text{power piston swept volume} + \text{dead space volume}}{\text{dead space volume}}.$$

For this equation to apply to the gamma Stirling engine with a displacer piston of the

type shown in Fig. 2.1, we must modify the numerator to include the swept volume of the displacer piston rod and the fact that the piston rod and power piston are out of phase so that the minimum or maximum engine volume is not necessarily reached when the power piston is at the extreme bottom or top of its stroke, respectively. This is due to the displacer piston *rod* (but not the displacer piston itself) affecting the pressure of the working fluid, albeit to a much lesser degree than the power piston.

2.3 Deriving the Model Equations

2.3.1 Assumptions

Reducing the convoluted, interrelated workings of the Stirling engine to a set of mathematical equations is no mean feat. Getting a closed-form expression for power out of the deal is even more impressive, and requires that we make a number of simplifying assumptions, some of them similar to those made in Section 1.2 when we formalized the ideal Stirling cycle. In [2], Senft outlines the assumptions of the Schmidt model presented therein, which are:

(1) The motion of both the power and displacer pistons is assumed to be perfectly sinusoidal;

(2) The working fluid behaves like an ideal gas in accordance with the ideal gas law, Eq. (1.1);

(3) The working fluid expands isothermally in all engine spaces;

(4) The regenerator is ideal so that heat regeneration is perfect with no waste heat [10];

(5) The instantaneous pressure of the working fluid is the same across all engine spaces;

(6) No air leaks out of the engine space through imperfect seals or otherwise.

Whether or not these assumptions are reasonable (that is, realizable in practice) is an important consideration, and will be discussed presently.

2.3.2 The Senft Model

The basic idea of the Senft model for predicting the power output of gamma Stirling engines is to develop a mathematical relationship between the power output of the engine (work W per unit time) and two salient characteristics of the working fluid: its volume V and pressure p. To that effect, deriving the Senft model involves first determining the volume of working fluid in the engine at any point during the engine cycle, then using the ideal gas law, Eq. (1.1), to determine the resulting instantaneous pressure which can be incorporated into Eq. (1.3) to determine the work done per cycle, as desired. The work done per cycle multiplied by the number of cycles per second gives the power output of the engine at that speed. Following this prescription, we begin by developing equations that describe the total, instantaneous volume of working fluid in the engine (that is, the volume at any given moment in time). In his derivation, Senft first divides the engine into a hot space and a cold space, each one having a bit of a strange definition. As seen in Fig. 2.2, the hot space is the half of the swept volume of the displacer piston closest to the heat source, and the cold space is the half of the displacer swept volume closest to the heat sink, plus the swept volume of the power piston. Hence, the hot space volume V_H and the cold space volume V_C are what we may call *active* volumes, or volumes traversed (swept out) by at least one of the pistons. This implies that the total engine volume V_{tot} is the sum of the cold space and hot space volumes, plus any volume not swept through by either piston, namely the dead volume V_N, so that $V_{\text{tot}} = V_H + V_C + V_N$.

As for actually determining V_H and V_C, Fig. 2.2 may be of aid. Senft 'sets his clock' at $t = 0$ when the power piston is at the top of its stroke, or at top dead center, as reflected in the figure. This way, the instantaneous angle of the power piston crank is simply ωt, with ω the engine speed in cycles per second (measured in Hertz, Hz, with units of s^{-1}). Due to the phase separation α between the two pistons, the instantaneous angle of the displacer crank with respect to the horizontal (the plane in which the power piston crank arm resides when the piston is at top dead center) becomes $\omega t + \alpha$. Then with reference to the figure, one should be able to convince one's self that

$$V_H = \frac{V_D}{2}\left(1 + \cos(\omega t + \alpha)\right). \tag{2.4}$$

Certainly, this equation makes sense in the two extremes of $(\omega t + \alpha) = 0$ and $(\omega t + \alpha) = \pi$. In the former case, with the instantaneous angle of the displacer crank equal to zero, the displacer piston must be at top dead center and we have that

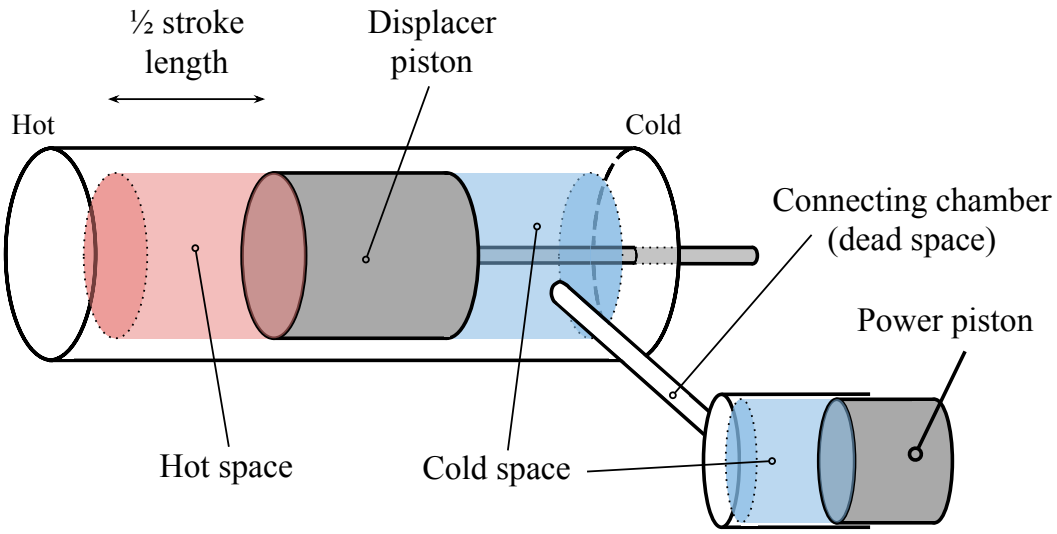

Figure 2.2: A diagram illustrating the hot and cold spaces, as defined by Senft, of a gamma type Stirling engine. Note that the two spaces consist only of volumes swept out by the pistons; the remaining volume (the white space within the black-outlined engine space) is the dead volume V_N. Not shown is the contribution to the cold space volume from the swept volume of the piston rod. As shown, the power piston is at the top of its stroke ($t = 0$), and the displacer, leading by 90°, is halfway through its travel.

Quantity	Symbol
Hot space volume	V_H
Cold space volume	V_C
Total engine volume	V_{tot}
Displacer piston swept volume	V_D
Power piston swept volume	V_P
Total swept volume	V_T
Dead volume	V_N
Engine speed	ω
Temperature ratio (T_C/T_H)	τ
Swept volume ratio (V_P/V_D)	κ
Dead volume ratio (V_N/V_D)	χ

Table 2.1: A table providing key notation used throughout the Senft model.

$V_H = {}^{V_D}/_2(1+1) = V_D$, indicating that the hot space now consists of the entirety of the displacer swept volume. If we imagine that the displacer in Fig. 2.2 were to move completely to the cold side of the cylinder (top dead center), the swept volume occupying the cold space would be transfered to the hot space, V_H would indeed equal V_D, and the figure would look more like Fig. 2.1. In the latter case, corresponding to the displacer piston being at bottom dead center, or all the way towards the hot side of the cylinder, we have $V_H = {}^{V_D}/_2(1+(-1)) = 0$, which makes sense as all the swept volume is now in the cold space, thus emptying the hot space.

The portion of the cold space within the displacer cylinder is then the swept volume of the displacer V_D minus the volume in the hot space V_H, whatever it happens to be (remember, it depends on t). The total cold space volume V_C is this portion plus the contribution from the swept volume of the power piston, giving

$$V_C = V_D - V_H + \frac{V_P}{2}(1+\cos\omega t), \tag{2.5}$$

with the last term being the power piston contribution. This equation makes sense at top dead center with respect to the power piston ($\omega t = 0$) because in this case, it predicts the power cylinder contribution to the cold space to be the total power piston swept volume V_P, as desired. Equation (2.5) also checks out at bottom dead center ($\omega t = \pi$) since it predicts the power piston contribution to V_C to be zero, which we would expect since with the piston slammed into its cylinder, no working fluid can be present in the cylinder that could contribute to V_C. As the sum of the hot, cold, and dead space volumes, the total instantaneous engine volume $V(t)$ can be written as

$$V(t) = V_D + \frac{V_P}{2}(1+\cos\omega t) + V_N,$$

which can be rewritten in terms of the swept volume ratio κ and the dead space ratio χ (see Table 2.1) as

$$V(t) = V_D\left(1 + \frac{\kappa}{2}(1+\cos\omega t) + \chi\right). \tag{2.6}$$

In order to compute the work per cycle of the engine with instantaneous volume $V(t)$, we are ultimately going to have to use Eq. (1.3), which calls for the quantities $\mathrm{d}V$ and p. The latter is the pressure of the working fluid within the engine space and will be computed momentarily. To determine the former, or the differential volume of working fluid within the engine space, we need only take the time derivative of $V(t)$ giving $\frac{\mathrm{d}V}{\mathrm{d}t}$ from which we can extract $\mathrm{d}V$. Then, in terms of the total swept volume $V_T = V_D + V_P = V_D(1+\kappa)$, we have

$$\begin{aligned}\frac{\mathrm{d}V}{\mathrm{d}t} &= \frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{V_D\kappa}{2}(1+\cos\omega t)\right)\\ &= -\frac{V_D\,\omega\kappa}{2}\sin\omega t\\ &= -\frac{V_T\,\omega\kappa}{2(1+\kappa)}\sin\omega t,\end{aligned}$$

giving

$$\mathrm{d}V = -\frac{V_T\,\omega\kappa}{2(1+\kappa)}\sin(\omega t)\,\mathrm{d}t. \tag{2.7}$$

To derive an expression for the instantaneous pressure of the working fluid from $V(t)$, we return to a key assumption from Section 2.3.1, namely that the working fluid can be modeled as an ideal gas. This assumption allows us to relate $V(t)$ to the instantaneous pressure $p(t)$ by way of the ideal gas law. However, we will not use the form of the law given in Eq. (1.1), but instead will use the form where $pV = NkT$, with N the number of particles (or the number of molecules of working fluid) and k Boltzmann's constant [11]. Rearranging this equation slightly and substituting values gives

$$\begin{aligned}p(t) &= Nk\frac{T}{V}\\ &= Nk\left(\frac{T_H}{V_H}+\frac{T_C}{V_C}+\frac{T_N}{V_N}\right)\\ &= \frac{NkT_C}{\tau V_H + V_C + V_D\dfrac{2\tau}{1+\tau}\,\chi},\end{aligned}$$

with τ the temperature ratio defined in Table 2.1. Inserting our expressions for V_H and V_C, Eqs. (2.4) and (2.5), we obtain

$$
\begin{aligned}
p(t) &= \frac{NkT_C}{\frac{\tau V_D}{2}\left(1+\cos(\omega t+\alpha)\right)+V_D-\frac{V_D}{2}\left(1+\cos(\omega t+\alpha)\right)+\frac{V_P}{2}(1+\cos\omega t)+V_D\frac{2\tau}{1+\tau}\,\chi} \\
&= \frac{NkT_C}{\frac{\tau V_D}{2}+\frac{V_D(\tau-1)}{2}\cos(\omega t+\alpha)+\frac{V_D}{2}+\frac{V_P}{2}(1+\cos\omega t)+V_D\frac{2\tau}{1+\tau}\,\chi} \\
&= \frac{NkT_C}{\frac{V_D}{2}\left[\tau(1+\cos(\omega t+\alpha))-\cos(\omega t+\alpha)+1+\kappa(1+\cos\omega t)+V_D\frac{4\chi\tau}{1+\tau}\right]}.
\end{aligned}
\tag{2.8}
$$

We can tidy up this monstrosity and put it into a more suggestive form by making a few substitutions to condense terms. Letting

$$Y = \tau + 1 + \kappa + V_D\frac{4\chi\tau}{1+\tau},$$

$$A = \kappa - (1-\tau)\cos\alpha,$$

and

$$B = (1-\tau)\sin\alpha,$$

we can reduce Eq. (2.8) to

$$p(t) = \frac{NkT_C}{{}^{V_D}\!/_2(Y + A\cos\omega t + B\sin\omega t)}. \tag{2.9}$$

As Senft shows, however, by defining the quantities

$$\theta = \cos^{-1}\frac{A}{\sqrt{A^2+B^2}},$$

and

$$X = \sqrt{A^2+B^2}$$

such that

$$\begin{aligned} X\cos(\omega t-\theta) &= \sqrt{A^2+B^2}\,[\cos(\omega t)\cos(\theta)+\sin(\omega t)\sin(\theta)] \\ &= \sqrt{A^2+B^2}\left[\cos(\omega t)\frac{A}{\sqrt{A^2+B^2}}+\sin(\omega t)\sqrt{1-\frac{A^2}{A^2+B^2}}\right] \\ &= A\cos(\omega t)+B\sin(\omega t), \end{aligned}$$

we can further reduce Eq. (2.9) to the more tractable form:

$$p=\frac{NkT_C}{{}^{V_D}/_2\,[Y+X\cos(\omega t-\theta)]}. \tag{2.10}$$

Now, the phase angle α is confined to the range $0<\alpha<\pi$, because if $\alpha>\pi$ or $\alpha<0$, we encounter a situation where the displacer piston no longer leads the power piston in the engine cycle. As a result, the Stirling cycle of Fig. 1.3 will not be carried out, and in fact the engine will not run in its normal direction no matter the temperature ratio τ between its hot and cold sides. It will, however, run contentedly in reverse because then the displacer piston will once more be leading the power piston (by an angle equal to $\alpha-\pi$). Amazingly, if the displacer leads the power piston in the forward direction of rotation, and kinetic energy is input into the crankshaft of the engine so as to drive it in the opposite direction, the Stirling engine will carry out the steps of the thermodynamic cycle of Fig. 1.3 *in reverse order*, and will actually act as a refrigerator [3]! With an appropriately designed Stirling engine, it is even possible to drive the crankshaft in reverse (with the power piston leading the displacer) with an electric motor thus refrigerating one side of the displacer cylinder and heating the other to the point where the cool end is iced over and the hot end is scalding. If the electric motor is then disconnected, the Stirling engine will run in the opposite direction, its typical forward direction with the displacer leading the power piston, off of the temperature differential just established [5]. This unique ability results from the fact that the Stirling engine is a true *reversible* heat engine [15].

Since the quantity $\cos(\omega t-\theta)$ in Eq. (2.10) ranges from -1 to 1, we find that the maximum and minimum values of Eq. (2.10) are given by

$$p_{\max}=\frac{NkT_C}{{}^{V_D}/_2[Y-X]}$$

and

$$p_{\min} = \frac{NkT_C}{{}^{V_D}\!/_2[Y+X]},$$

respectively, so that the average pressure $\bar{p}$ of the working fluid (note the loss of dependency on t—this is a time average, after all) throughout one engine cycle is

$$\bar{p} = \sqrt{p_{\max}\, p_{\min}} = \frac{NkT_C}{{}^{V_D}\!/_2\sqrt{Y^2 - X^2}},$$

which allows us to put our expression for the instantaneous pressure $p(t)$ into its final form:

$$p(t) = \frac{\bar{p}\sqrt{Y^2 - X^2}}{Y + X\cos(\omega t - \theta)}. \tag{2.11}$$

So far we used geometrical arguments to determine the total instantaneous engine volume $V(t)$, Eq. (2.6), from which we derived the instantaneous pressure of the working fluid $p(t)$, Eq. (2.11), using the ideal gas law and a number of simplifying substitutions. We are now fully equipped to derive an expression for the work per cycle W_{cyc} of the Stirling engine. Substituting Eqs. (2.7) and (2.11) into Eq. (1.3), and noting that the definition of $\mathrm{d}V$ given by Eq. (2.7) changes the integral from one over volume to an integral over time, we get that

$$\begin{aligned} W_{\text{cyc}} &= \int p\,\mathrm{d}V \\ &= -\frac{V_T\,\bar{p}\,\omega\kappa\sqrt{Y^2 - X^2}}{2(\kappa+1)} \int_0^{{}^{2\pi}\!/_\omega} \frac{\sin\omega t}{Y + X\cos(\omega t - \theta)}\mathrm{d}t, \end{aligned}$$

where the upper limit of integration ${}^{2\pi}\!/_\omega$ is simply the time taken for the engine crankshaft to sweep out 2π radians, or one full revolution, beginning with $t = 0$ when the power piston is at top dead center. Integrating the above expression gives us the final form for the work per cycle:

$$W_{\text{cyc}} = \frac{V_T\,\bar{p}\,\kappa\,\pi\,(1-\tau)\sin\alpha}{(\kappa+1)\left[Y + \sqrt{Y^2 - X^2}\,\right]}. \tag{2.12}$$

Note that this expression depends only on V_T, $\bar{p}$, and the quantities τ, κ, χ, and α, all of which are determined by engine dimensions and the temperature differential to which the engine is subjected. Further note that as W_{cyc} indicates the work per cycle (in Newton-meters, $\mathrm{N}\cdot\mathrm{m}$), the power the engine produces must equal this work times

the cycles per second ω. That is,

$$P_{\text{stirling}} = W_{\text{cyc}}\,\omega. \tag{2.13}$$

This equation makes sense dimensionally since

$$[W_{\text{cyc}}] \cdot [\omega] = \text{N} \cdot \text{m} \cdot \left(\text{s}^{-1}\right) = \frac{\text{N} \cdot \text{m}}{\text{s}} = [\text{Power}]\,.$$

Lastly, it is worth mentioning that Eq. (2.12) specifically predicts the cyclic *indicated* work performed by the engine, where indicated work is defined as being calculated from pressures and volumes in the engine space under the assumption that no losses (thermal or frictive) occur in converting the energy of the expanding working fluid into mechanical motion at the crankshaft. This type of work (or, really, this type of work *measurement*) earned its name from the indicator device developed by James Watt to record and plot pressure and volume variations within the cylinders of steam engines. This device essentially produced a real, non-idealized p-V diagram for the engine from which work as well as and power and efficiency could be readily calculated. Hence, when we ultimately use Eqs. (2.12) and (2.13) to calculate theoretical power outputs, we will more specifically be calculating theoretical *indicated* power outputs. This is problematic because when we physically measure the power output of the engine (described in the next chapter), we will be measuring the power output at the crankshaft by applying a controlled load or brake against which the engine will have to fight. Hence we will be measuring the *brake power* of the engine, not its indicated power. Apart from how they're measured, the difference between the two is that brake power *does* take into account all losses in the engine because it is measured at the output shaft of the engine. To compute brake power from indicated power, then, one would have to factor in the friction between pistons and cylinders, connecting rods and pivots, and the crankshaft and its pivots, and deduct the effects of these frictive forces, among other sources of power loss, from the indicated power. Hence, for any real engine, the brake power is necessarily less than the indicated power. We will have to account for this difference in our analysis.

Chapter 3

The Test Rig

3.1 Engine Design

In order to test the validity of the predictive model outlined in Chapter 2, a simple, gamma-type Stirling engine, originally manufactured by Solar Engines of Phoenix, AZ, was heavily modified to allow for variation of both the phase angle α and the swept volume ratio κ. To that effect, new flywheels were made that allowed the stroke of both the displacer piston and the power piston to be adjusted independently, and for the phase angle to be manipulable independently from the strokes. Additional modifications included the manufacture of new connecting rods that could accommodate variations in stroke, a new crankshaft, an upgraded heat source, and a new mounting block for the engine and its heat source. These modifications allow the parameters α and κ to be varied independently and repeatedly so that their effect on the power output of the engine can be recorded consistently and accurately, and ultimately be compared to the theoretical predictions of the Senft model.

Fig. 3.1 provides a labeled image of the engine in its original form, showing most of its components. Hidden from view are the displacer piston, which resides inside the displacer cylinder, the power piston connecting rod, and most of the crankshaft, which supports and connects the two flywheels and is the component from which power output is measured. A few design features of the engine deserve special notice: firstly, both flywheels only have one point of attachment for each respective connecting rod. While perfectly adequate for the proper function of the engine, a single mounting point does not allow for the stroke of either piston to be adjusted. After all, the stroke is *entirely* determined by the radial distance between the center of the flywheel (where it mounts to the crankshaft) and the mounting point itself, and is in fact *equal* to twice that radial distance. Secondly, the set screws, one per flywheel, are the

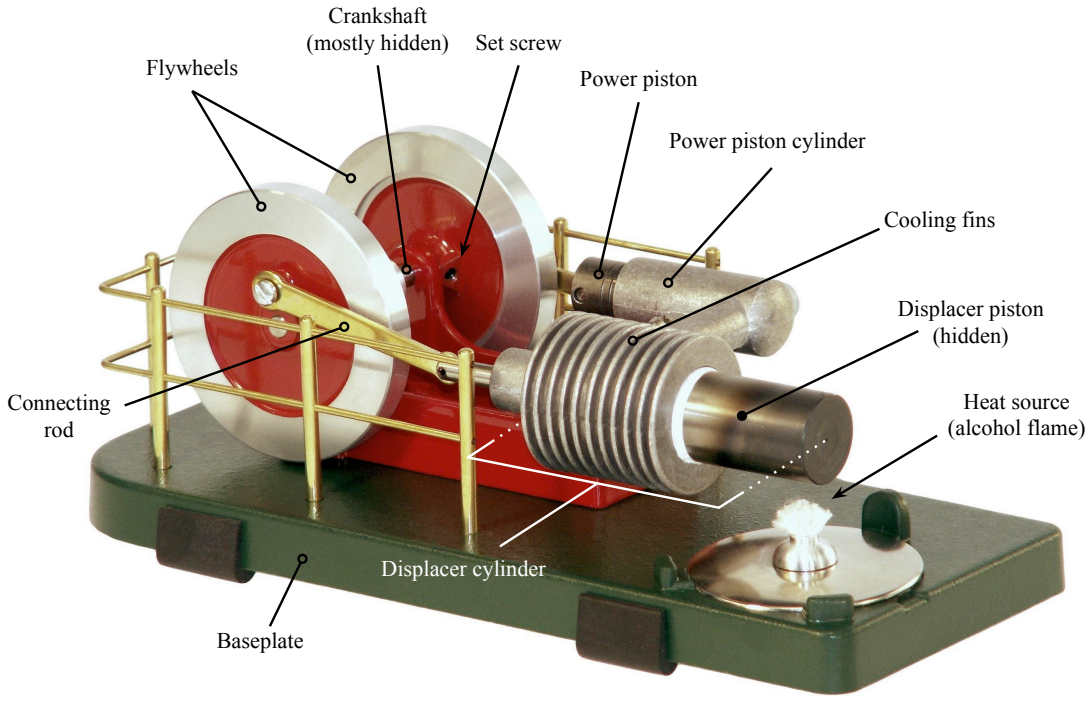

Figure 3.1: The engine, a Solar Engines of Phoenix, AZ, Stirling #1, in its original, unmodified form. Note the set screws on the flywheels (only one is visible) that allow for phase angle adjustment, and the solitary mounting point on each flywheel for its respective connecting rod. This latter feature of the original engine design does not permit adjusting the stroke of either piston. (Original image taken from: http://www.bullnet.co.uk/shops/test/images/solar1.jpg)

sole element anchoring the flywheels to the crankshaft; the fit between the flywheels and crankshaft is ever so slightly loose, a sure connection between the two requiring pressure from the set screw. The set screws are furthermore wholly responsible for maintaining the phase angle α. Loosening a set screw and rotating its corresponding flywheel about the crankshaft while holding the other stationary changes the angle between the connecting rod mounting points of the two flywheels, which is precisely the angle α in Fig. 1.2. Lastly, the heat source for the original engine is a simple alcohol flame, the wick appearing as a white tuft protruding from the chromed-out fuel reservoir cap. While sufficient for casual operation of the engine as a display piece, such a heat source is too inconsistent and uncontrollable for the strict rigors of this experimental test rig.

3.1.1 Flywheels and Connecting Rods

The original engine flywheels constrained the strokes of both the power and displacer pistons to 15.5 mm, giving a permanent swept volume ratio of $\kappa = 0.58$. (Though each piston had the same stroke, the difference in piston radii meant that each piston swept out a different volume during its stroke, giving $\kappa \neq 1$.) Engine dimensions, most notably the depths of the power and displacer cylinders, kindly introduced a further constraint to the range of permissible strokes. No matter what clever arrangement of flywheels and connecting rods, the fact remains that the piston stroke cannot be so large that the piston bottoms out in its cylinder (pushed in too far) or is pulled from it completely at the extreme ends of its travel. Similarly, the stroke can't be made short to the point where the connecting rod mounting point on the flywheel interferes with the junction between the flywheel and the crankshaft. Since the mounting point is simply a threaded hole drilled into the flywheel into which a screw, serving as a pivot for the connecting rod, is inserted (see Fig. 3.3), the radial distance between the hole for the crankshaft and the connecting rod mounting hole (r_1 or r_2 in Fig. 1.2) must be sufficiently large to ensure the two holes do not overlap. What's more, this radial distance must be large enough so that sufficient material exists between the two holes to ensure structural integrity of the flywheel.

To establish the range of permissible strokes, we need to determine how far the piston can travel into its cylinder before bottoming out, and how far it can travel out of its cylinder while still leaving enough inside to preserve the hermetically sealed engine environment and retain structural integrity. Pull the piston out too far and not only could some working fluid potentially escape, but side loads (any force acting

antiparallel to the piston's velocity vector) exerted on the piston by the connecting rod as it attempts to force the piston back in its cylinder could cause the piston to bind up, halting and perhaps destroying the engine. To avoid such a catastrophe, we require some minimum percentage, call it the retention percentage σ, of the piston to remain in the cylinder even at the extreme top of its stroke. For this engine, we set $\sigma \approx 0.25$, meaning that roughly 25% of the piston remains within the cylinder at the top of its stroke.

As suggested by Fig. 3.2, a number of engine dimensions must first be measured before maximum stroke can be calculated: the lengths of the piston ℓ_{pis}, the cylinder ℓ_{cyl}, and the connecting rod ℓ_{rod}, as well as the radial distance r between the center of the flywheel and the connecting rod mounting pivot, and the total distance from the crankshaft to the bottom of the cylinder ℓ_{tot}. Since the stroke is equal to $2r$, we want to find r in terms of the other quantities. With reference to Fig. 3.2, we see that configuration (a) requires

$$r + \ell_{\text{rod}} + \ell_{\text{pis}} < \ell_{\text{tot}} \,, \tag{3.1}$$

and from (b) we ascertain the restriction

$$(\ell_{\text{rod}} - r) + (1 - \sigma)\,\ell_{\text{pis}} + \ell_{\text{cyl}} < \ell_{\text{tot}} \,. \tag{3.2}$$

Due to time constraints, it was not feasible to alter ℓ_{cyl} , ℓ_{tot} , or ℓ_{pis} as adjusting any one of these parameters would require extensive, high-tolerance machining of various engine components. Hence, those were all taken as fixed quantities with r and ℓ_{rod} variable. The length of the connecting rod was chosen to remain variable because manipulating r and ℓ_{rod} in concert allows for a wider range of permissible strokes than the manipulation of r alone. Recall that the stroke is solely dependent on r, so increasing r while decreasing ℓ_{rod} by the appropriate amount to avoid bottoming out the piston still results in a longer stroke. After grouping our constant terms ℓ_{cyl} , ℓ_{tot} , and ℓ_{pis}, it follows from Eqs. (3.1) and (3.2) that

$$\ell_{\text{rod}} + X < r < Y - \ell_{\text{rod}} \,, \tag{3.3}$$

where $X = (1 - \sigma)\,\ell_{\text{pis}} + \ell_{\text{cyl}} - \ell_{\text{tot}}$ and $Y = \ell_{\text{tot}} - \ell_{\text{pis}}$ are both constants determined by measurable quantities.

Unfortunately, this equation masks the true limitations on r, namely the depth of the cylinder ℓ_{cyl} and the desired retention percentage σ. The latter quantity determines how far the piston can be pulled out of its cylinder, and the former determines

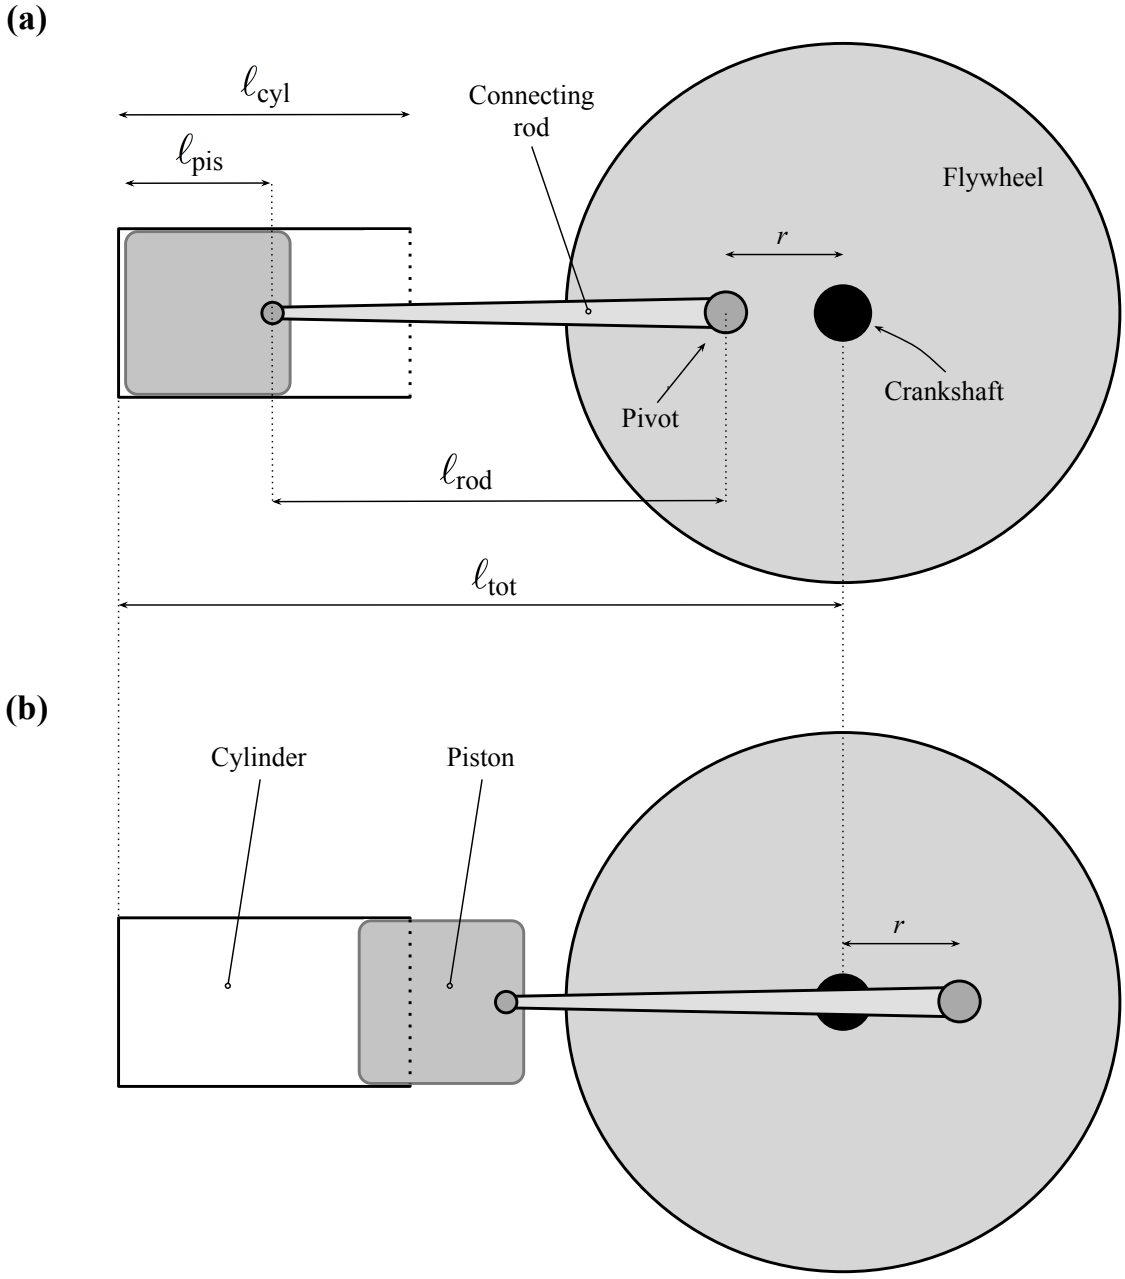

Figure 3.2: Schematic detailing the relevant lengths for determining the maximum stroke of either the power or displacer piston. Configuration (a) shows the position of the piston and connecting rod at the extreme bottom of the stroke (just before the piston would "bottom out" on the end of the cylinder), whereas configuration (b) shows the piston and rod positions at the extreme top of the stroke. Note in (b) how part of the piston remains within the cylinder even at max extension; the length of piston remaining within the cylinder is equal to $(1 - \sigma)\,\ell_{\text{pis}}$, where σ is the retention percentage.

how far it can be pushed in, so together they determine the piston's maximum travel. This, naturally, is also the maximum permissible stroke. In practice, then, determining the maximum allowed stroke involved first measuring ℓ_{cyl} then deciding on a value for σ. This decision involved guesswork, the question under consideration being: how far can the piston extend from its cylinder before it *seems* like the piston will experience excessive side-loads when being pushed back in? With σ and ℓ_{cyl} established, simple calculation gives the maximum stroke, and dividing this number by two gives the maximum value for r. From Eq. (3.3), the required ℓ_{rod} can then be calculated with ease. Furthermore, it follows from Eq. (3.3) that this ℓ_{rod} is also the maximum length necessary to accommodate all strokes less than or equal to the maximum stroke. This fact, which may be intuitive, is convenient as it means that only one connecting rod (per piston) is necessary so long as it accommodates the maximum stroke of the piston.

Now, the whole aim behind making new flywheels was to enable manipulation of the stroke of either piston so that the swept volume ratio κ of the engine could be varied, and its effects on power output measured. Given that new flywheels needed to be made, the goal became to make these flywheels in such a way as to maximize the amount by which κ could be varied. At the very least, the new flywheels should allow κ to be varied from some value less than one to some value greater. This way it would be possible to test the performance of the engine with the displacer sweeping out more volume than the power piston, as well as with the power piston sweeping out more volume than the displacer. The limits imposed on r by ℓ_{cyl} and σ constrained r to a range insufficient for producing the desired κ range from variation of the stroke of one piston alone, so variable strokes for *both* pistons was absolutely necessary. To that effect, two flywheels were made, one for each piston, each with multiple connecting rod mounting holes positioned at different radial distances from the center of the flywheel, and hence corresponding to different, discreet stroke settings. After deciding which stroke range was possible for the power piston, its flywheel was designed to have six mounting holes, giving the piston a range of strokes from 9 mm to 19.15 mm, in increments of 2.03 mm. Because of the original design of the displacer piston and its cylinder, the stock displacer stroke is the maximum permissible without extensive modification to the main engine body. Hence the displacer flywheel was designed with only four connecting rod mounting holes, corresponding to displacer piston strokes ranging from 9 mm to 15.49 mm in increments of 2.17 mm. As the stroke of each piston is independently adjustable, these distinct stroke settings give a respectable 24 different values for κ, ranging from 0.36 to 1.34, thereby satisfying our desire to test

values of κ both less than and greater than one.

In an effort to smooth out the rotational motion of the original engine, the glorious new flywheels were made to be more massive than the originals, enabling them to store more angular momentum and therefore to better dampen slight irregularities in the angular velocity of the crankshaft. (Slightly jerky crankshaft rotation at low engine speeds was the largest concern.) A heavier flywheel has a larger moment of inertia I than a lighter one, so that for a given engine speed $\vec{\omega}$, its angular momentum $\vec{L}$ is also larger, in accordance with the relation $\vec{L} = I\vec{\omega}$ [12]. Hence, it takes a larger amount of force on the new flywheel (from the piston and connecting rod) to alter its angular velocity $\vec{\omega}$ than it does for the originals. While the stock flywheels were made from steel, in an effort to increase their moment of inertia, the new flywheels were turned on a lathe from solid brass bar stock, a more dense material than steel that aside from providing the requisite heft, looks absolutely stunning when machined. They were machined thicker and of a larger diameter than the originals to further increase their mass and moment of inertia over the stock flywheels, and were transfered from the lathe to the milling machine for drilling and tapping of the pinch bolt and connecting rod mounting holes.

In contrast to the design goal for the flywheels, the aim for the new connecting rods was to make them as lightweight as possible. They, too, have a mass-dependent moment of inertia I_{rod} that we want to minimize as it will only absorb mechanical energy from the engine (energy that could otherwise do work on the crankshaft) without any benefit such as smoothing out engine rotation. While its motion is quite complicated, the center of mass of each connecting rod does have some time-dependent angular velocity $\omega(t)$ that is *not* constant. This is a consequence of the fact that the center of mass of the connecting rod does not trace out a circle, and hence the point about which the angular velocity is measured is constantly changing. This means that there is some nonzero angular acceleration associated with each connecting rod. By the rotational analogue of Newton's second law, $\tau = I\alpha$, where τ is a torque, I is the moment of inertia of the body in question, and α is its angular acceleration, this nonzero angular acceleration requires a torque, and hence a force, and hence energy, which must be provided by the engine [22].

To reduce the moment of inertia of the rods, thereby boosting efficiency by minimizing the amount of energy wasted on their constant acceleration, the new rods were made from aluminum sheet metal, giving them a lower mass than the original brass items. The displacer connecting rod was kept the same length as the original because engine parameters precluded elongation of the displacer stroke from its stock

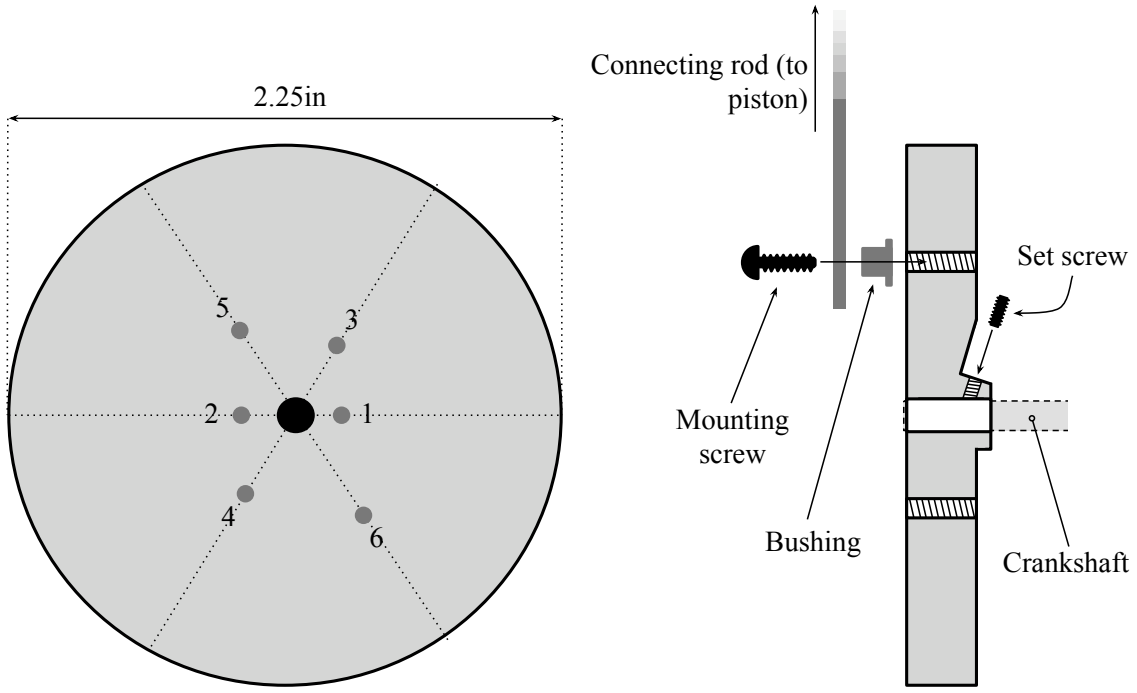

(a) Front view

(b) Side view

Figure 3.3: Diagram showing the front and side views of the power piston flywheel (the displacer piston flywheel, not shown, is identical, but with four connecting rod mounting holes instead of six). The mounting holes are positioned such that those of similar radial distance from the center are positioned opposite of each other to improve flywheel balance. Looking at the side view, one can see how the connecting rod is fixed to the flywheel. A hole (not visible) drilled through one end of the connecting rod fits over the bushing and rests against its flange, then a screw is inserted through the bushing and is screwed into the flywheel, holding the bushing in place but allowing the connecting rod to rotate freely.

setting, but the power piston connecting rod was made slightly shorter to compensate for the longer-than-stock strokes of the new power piston flywheel. Using a large shears, the rods were rough-cut from aluminum sheet as rectangles whose edges were rounded and scalloped out into a dog-bone shape using an oscillating spindle sander to increase elegance and further reduce weight. They were then given a honed finish by being manually wet-sanded on all sides for that added sultry gleam and for the more practical purpose of eliminating any remaining rough edges (burrs) that could catch on pivot components and impair proper and consistent engine function. Fig. 3.4 provides a diagram of the power piston connecting rod, showing by means of a dotted boundary the size of the rough cut before sanding and honing.

3.1.2 Upgrading the Heat Source

As shown in Fig. 3.1, the heat source for the original engine was a simple alcohol flame. On one end, the green baseplate incorporated a small reservoir capable of holding enough rubbing alcohol (or Everclear, or kerosene, or gasoline—the Stirling isn't picky) for 7-10 minutes of engine operation. A cloth wick, with one end anchored in the metallic fuel reservoir cap visible in Fig. 3.1 and the other submerged in the fuel, sustained the flame and centered it beneath the very end of the hot space at the edge of the displacer cylinder. Though this type of heat source made the engine both highly portable and completely self-contained, thereby legitimizing its existence as a functional, table-top showpiece, it was woefully inconsistent and impossible to control. The flame would flutter incessantly, sporadically die, and hence could not come close to maintaining the hot space at a constant temperature. So instead of forevermore fiddling with a fiercely fickle flickering flame, the original alcohol burner was substituted for an electric heating element made from nichrome wire looped around the hot side of the displacer cylinder.

Nichrome is an alloy of typically 80% nickel and 20% chromium that is notable for its very high electrical resistance R. It is highly temperature resistant, with a melting point of $\sim$1400 °C, and is readily available as a wire in a variety of thickness gauges. Due to this high resistance, if a current I is passed through a length of nichrome wire, the wire will radiate an impressive amount of heat Q_{rad}. This is due to collisions between moving electrons (which produce the current I) and the atoms of the wire, collisions which transfer some of the kinetic energy of the electrons onto the atoms, exciting them and causing the wire to heat up [12]. Hence, the nichrome wire (or any resistive conductor) will convert electrical energy to thermal energy, and will

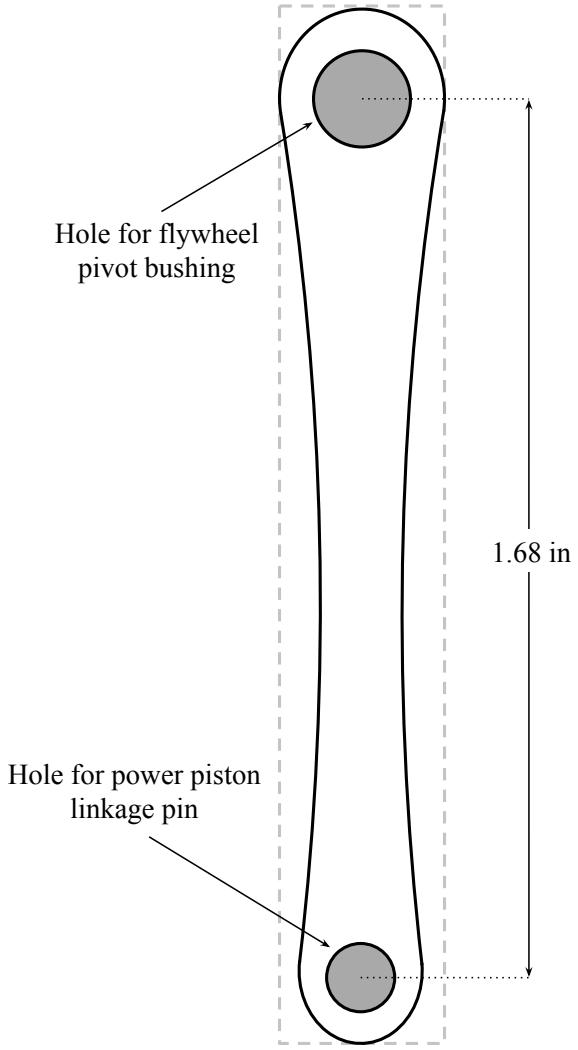

Figure 3.4: Diagram detailing the power piston connecting rod. The top hole fits around the bushing through which the flywheel mounting screw is inserted. The bottom hole fits around the power piston pivot pin forming a simple, low-friction joint. The dotted rectangle shows the size of the connecting rod after an initial rough cut from aluminum sheet; after the holes were drilled, the edges were rounded into the final dog-bone shape using an oscillating spindle sander.

do so at a rate equal to that of electrical energy production and delivery within the circuit. But this rate is simply the electrical power P available, which is given by the relation $P = IV = I^2R$, where V is the voltage drop across the resistive material, R is its resistance, and Ohm's law, $V = IR$, has been used to eliminate dependence on V [20]. The total radiated heat, then, is proportional to the electric power available, giving

$$Q_{\text{rad}} \propto I^2R,$$

with the *rate* of heat radiation—but not the total *amount* of heat radiated—equal to I^2R. (In fact, and perhaps intuitively, the total amount of heat radiated is also a function of time, or how long the energy transformation, from electrical to thermal, is allowed to occur at the rate specified by the power available P.) The large resistance of nichrome wire thus indicates that a substantial amount of heat will be radiated for a given current being passed through the wire. This property makes nichrome an ideal material for use as an electrical heater, and indeed it is commonly found as the heating element in the common toaster or in electronic cigarette lighters.

Using nichrome wire as the heat source for the engine was appealing as the nichrome won't flicker, and won't grow more intense, fade, or die completely so long as a constant current I, provided by a power source, is made to pass through the wire. What is more, nichrome wire can be molded into the desired shape beforehand, then anchored into a fixed position, thereby directly heating only a controlled area of the engine. (The fact that heat conduction occurs between different engine components means that, despite the fact that the nichrome only directly heats a small portion of the engine, other components will necessarily heat up as well.) The name of the game here is the consistent application of a steady amount of heat, and the properties and nature of nichrome wire made it a clear choice over an alcohol flame. Hence, the original heat source was scrapped and the new one made by wrapping a length of nichrome wire around the hot end of the displacer cylinder, then carefully spreading the loops of wire to ensure they would not contact each other and short-circuit. To ensure the nichrome wouldn't also short out on the conductive steel displacer cylinder, a length of woven fiberglass cloth (both highly heat-resistant and non-electrically conductive) was sandwiched between the nichrome coils and the cylinder. The heating element was held in place around the displacer cylinder by means of a ceramic clamp bolted to the baseplate. The fact that the clamp was made from ceramic is crucial: the clamp material had to be an excellent insulator so that it wouldn't short out the heating element, thereby preventing current from passing through and heat-

ing it, and it had to be capable of withstanding the high temperatures (~ 500 °C) of red-hot nichrome. The ceramic clamp worked perfectly, holding the heating element secure so that leads running to a 12V/10A power supply could be attached to and removed from the nichrome wire without altering the position of the coils relative to the displacer cylinder.

Configured as such, the nichrome heating element offered a simple and effective alternative to the temperamental alcohol heat source, allowing precise moderation of temperature (by way of increasing or decreasing the current passing through the wire), and, more than anything, ensuring consistent application of heat. It is not within the scope of this thesis to be able to accurately control *how much* heat is being applied, as that would only be important if we were concerned with measuring the efficiency of the engine. It is, however, imperative to have a *consistent* heat source so that, regardless of the precise amount of heat radiated, the temperature of the hot space T_H can be kept constant. Upgrading the heat source marked the final upgrade to the engine itself, which can be seen in its final, ready-for-testing form in Fig. 3.5.

3.2 Power Measurement

Since the ultimate goal of this work is to measure the power output of a Stirling engine and compare the results to those predicted by the theoretical model outlined in Section 2.3, all of the modifications described above are useless without some way of actually measuring the power the engine can produce. The measurement of this key quantity is in practice more difficult than Eq. (2.3), $P = \tau\omega$ (constant τ), may seem to suggest. According to this equation, all we need to do is measure the torque produced by the engine and multiply that by the engine speed ω, measured in Hertz, at which that torque measurement was taken. But while the engine speed is a simple quantity to measure, torque is not so cooperative, especially considering how little torque and power the miniature engine used here is capable of producing. Making torque and power measurements on a more substantial engine would be a relatively simple process as devices do exist today that are designed specifically for that purpose. The torque and horsepower ratings for motorcycle and automobile engines are today all made using a dynamometer, or dyno for short. These devices work by applying a known load to the engine, usually by way of hydraulics or magnetic fields, which is equivalent to applying a torque to the engine that it will try to combat. By measuring the changes in engine speed as increasing load (torque) is applied, output power can be computed. Since no production dyno could be found that was sensitive enough

to measure the very low power output of the engine used here, a very simple form of dynamometer called a de Prony brake was made instead that uses frictive forces to create a controlled load.

3.2.1 The de Prony Brake and its Design

Proposed in 1822 by the French military engineer Baron Riche de Prony, the de Prony brake, or simply Prony brake, is a very rudimentary dynamometer that clamps onto the output shaft of an engine and uses friction between the shaft and brake to exert a load on the engine [23]. As shown in Fig. 3.6, the brake consists of two opposing plates, one long arm and a similar but shorter piece, that are shaped to fit snugly around the crankshaft when joined by two screws. The screws pass loosely through the top arm and screw into threads tapped into the lower plate so that as the screws are tightened, the two plates are pulled together, pinching down on the crankshaft. When in operation, the opposite end of the long arm, the end featuring a pointed metal tip, is made to rest on the pan of a scale so that the arm remains horizontal during testing.

When the engine is running with the Prony brake partially engaged, friction between the brake plates and the crankshaft will exert a torque on the brake arm, forcing its metal tip into the scale. In response, the scale will 'push back' with the same amount of force, an amount that can be determined easily from the mass reading on the scale (the force exerted and scale reading differ by a factor of g since the brake exerts a force $F = ma = mg$ with respect to the scale and its calibration). Since the scale exerts a force back on the brake which is free to pivot around the crankshaft, the scale must in fact exert a *torque* on the brake with a lever arm equivalent to that marked in Fig. 3.6. Since the arm remains stationary, it must be the case that the torque exerted by the *scale* on the brake is equivalent to the torque exerted by the *engine* on the brake at that engine speed. Since the two torques are equivalent, the reading on the scale must also correspond to the torque exerted by the engine, where in fact

$$\tau_{\text{engine}} = mg\ell, \tag{3.4}$$

with m the mass recorded by the scale (in kilograms) and ℓ the length (in meters) of the brake lever arm. Increasing the brake clamping force serves to increase the friction forces between the brake and crankshaft, slowing the engine but driving the point of the brake into the scale with more force, corresponding to a larger torque.

The Prony brake was machined from Delrin (polyoxymethylene), a stiff, low-friction plastic, so that frictive forces wouldn't immediately overwhelm the engine and so that the brake wouldn't scratch and damage the crankshaft as it would were it made from metal. Wood was also considered as a material for similar reasons, but the idea was rejected as it is a difficult material to machine precisely. After all, precision was required as the lever arm length needed to be known exactly (to facilitate correct calculations of torque from a mass reading on a scale), and the shaped surfaces conforming to the crankshaft needed to be cut very smoothly to reduce vibration along the brake arm when the brake is engaged.

With the engine as low-powered as it is, it isn't sufficient just to have the brake made out of a low-friction material, however. Also necessary was that the mechanism for tightening the brake be capable of increasing the brake force extremely gradually. The power output of low-power engines is *generally* difficult to measure because even a small increase in brake force has a pronounced effect on engine speed [24]. Increasing the sensitivity, or fine adjustment capabilities, of the Prony brakes tightening mechanism helps facilitate accurate torque measurement by allowing for more, and for more closely spaced, data points to be collected.

To make the brake force as minutely controllable as possible, the brake force adjustment screw was chosen to have a high thread count per inch, and as seen in Fig. 3.6, was made to act on the upper arm through a spring. The high thread count means that for each revolution of the screw, it will only sink into its hole by a small amount equal to the distance between threads. That is, in the case of the 40 threads per inch (tpi) adjustment screw used in the brake, one complete revolution will sink the screw by $1/40$ of an inch. The spring was used to negate the sharp increase in brake force that would result from any tightening of the adjustment screw beyond the point where the head of the screw has first contacted the upper arm (the point where the brake first engages). Due to the rigidity and incompressibility of Delrin, tightening sensitivity beyond this point would be wholly dictated by the thread count of the adjustment screw. With the spring, however, while there is a point where the brake first engages, tightening the adjustment screw beyond that point just serves to compress the spring. The spring in turn exerts a force on the upper arm given by Hooke's Law, $F_{\text{spring}} = -kx$, where k is the relevant spring constant and x corresponds to how much the spring has been compressed (or stretched) from its equilibrium length [12]. Tightening the adjustment screw increases x, and therefore the force applied by the spring, but that force is metered by k; the smaller the value of k, the smaller the increase in brake force for a given turn of the adjustment screw. The overall effect of

incorporating a spring into the adjustment screw, then, is that the brake behaves as though the thread pitch on the adjustment screw were fantastically and unrealistically high. It was thought, too, that the spring may have the additional benefit of absorbing potential vibrations that could cause jumpy readings on the scale, but this has not been confirmed.

3.2.2 Measuring Engine Speed ω

With the Prony brake we can measure the torque the engine produces, but to determine its output power, we also need to know how fast the engine is running when a given torque measurement is taken. That is, we need a device that can measure how quickly the engine performs one cycle. But as the engine used in this study is decidedly lacking in power, this device cannot be required to contact the engine in order to take a measurement as this would necessarily absorb power, likely of a non-negligible amount. The solution was to use an optical tachometer, a device that capitalizes on nothing but reflecting light and clever circuitry to take its measurements. It works by shining laser light at an opaque surface and using some form of photodetector, often a phototransistor, to detect if any light reflects back. By painting over the gleaming sheen of the outer rim of one of the flywheels to make the rim opaque, but sticking a single small patch of reflective tape somewhere on the painted rim, we have a flywheel that will reflect light back into the photodetector of the tachometer precisely once per rotation provided the tachometer remains stationary with respect to the flywheel. With only one reflective point per revolution (or per cycle), if we shine the tachometer at the rotating flywheel, all it has to do is count the elapsed time (in seconds) between successive photodetections, and divide by that time to give the engine speed in revolutions per second.

Originally, an optical tachometer circuit was built with an LED as a light source and a phototransistor as a detector, but when it was realized that a ready-made optical tachometer could be sourced cheaply online, the original circuit was scrapped. A black Sharpie was used to cover the shiny finish of one of the machined flywheels, and the reflective tape included with the tachometer was applied, ultimately with good effect, to the opaque rim following the recommendations of the instruction manual.

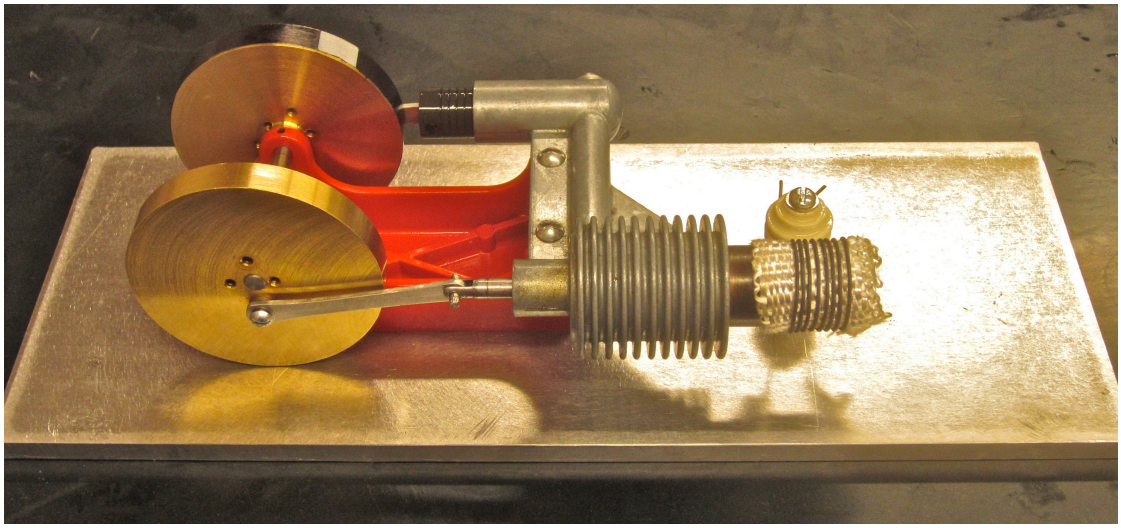

Figure 3.5: Here we see an image of the engine in its final, fully modified form. Visible in the image are the upgraded brass flywheels, the new aluminum connecting rods, the aluminum base, nichrome heating coil, and ceramic heater clamp. The rest of the engine was kept stock.

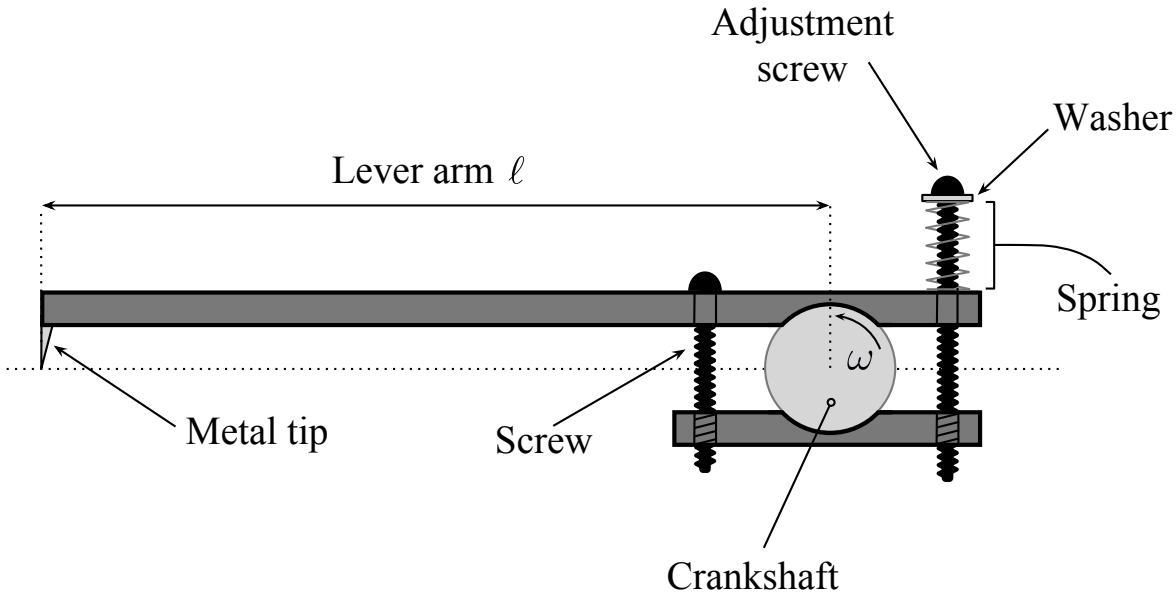

Figure 3.6: A side-view diagram of they type of Prony brake considered in this study. The two screws pass loosely through the long, upper plate and screw into the bottom plate to create pressure around the crankshaft. While the screw on the left is set to a certain height and left alone, the longer screw on the right is turned throughout testing to gradually increase the brake) on the crankshaft. The spring wrapped around this longer screw helps enable the brake force to be increased in the smallest possible increments; if it weren't there, the screw would tighten down onto the upper plate too quickly for brake force to be applied gradually. The metal tip provides a precise point of contact between the brake arm and the scale so that the lever arm can be measured accurately.

Chapter 4

Experimental Methods

In the last chapter we assembled a testable engine, built a device with which we can measure torque and power, and with this we are armed to see how the remaining collection of components can converge as a functional engine test rig. Recall that the aim is to measure the power output of the engine across a range of swept volume ratios κ, and phase angles α, and compare the results to the predictions of the Schmidt model. By virtue of the fact that the engine speed ω is easily variable with the Prony brake, the relationship between engine speed and brake power will also be investigated. Testing the model's ability to accurately predict the relationships between engine power and κ, engine power and α, and engine power and ω will require numerous data runs with many data points per run. In fact, a total of 24 swept volume ratios, with one data run per ratio, are possible with the flywheels configured as they are, and a continuous range of phase angles are possible from 0 to 180°. This translates into a large number of data points across which some measure of general consistency must be maintained, but a glance at the model reveals just what needs to be kept consistent. With reference to Eq. (2.12), we see that given set of engine dimensions (like the radius of the power piston) that collectively determine V_T, $\bar{p}$, τ, α, κ, Y, and X, if we supply the model with the appropriate temperatures T_C and T_H and the engine speed ω, it will provide us with a single indicated power output. As we have designed it so far, we will be able to determine the brake power output of the engine at a given engine speed if we can measure that engine speed and the corresponding torque value. Doing so will provide a means to calculate a single brake power output for that engine speed that can be compared against the theoretical indicated power. Testing each of the aforementioned three power relationships hence requires that we be able to maintain constant T_C and T_H throughout each trial to control the variable τ while determining the effects of ω on power output. To test

the ability of the model to predict the first two relationships, those between power and κ and power and α, however, we must further ensure that T_C and T_H are kept constant *between* trials.

4.1 Ensuring Temperature Consistency

Redesigning the heat source and converting it from a protean alcohol flame to an electrically powered heating element was the first step to introducing a measure of consistency in engine temperatures both during a single trial run and between multiple. The nichrome heater uses a power supply to produce the current required for heating, and with this power supply it is a simple matter to keep this current constant; after all, the power supply used was *designed* to produce a constant current. Supplied as such with a constant current, the nichrome heater will maintain a constant temperature, and T_H will hence remain constant, as desired, throughout a trial run. Provided the power supply is set to deliver the same amperage per trial, the heater will reach the same temperature per trial, keeping T_H consistent *between* trials as well. And So long as no significant ambient air currents are present around the heating element and so long as the nichrome isn't in contact with any thermally conductive material (which would cause heat convection and heat conduction, respectively, away from the heater), T_H will remain as constant during and between each trial as can reasonably be expected. As the tests took place in a ventilated room, some minimal air currents were inevitably present, and it was hoped they would be either negligible or else steady enough to still permit a constant T_H, even between trials.

Upgrading to a nichrome heater allows consistency in T_H, but can't be relied on for also maintaining a constant T_C. Due to heat conduction within the engine itself, heat from the nichrome and hot side of the engine will eventually distribute itself across the other various components of the engine, including the cold space. To get the most power and, as Eq. (1.4) shows, the highest efficiency, out of the engine, we want to keep T_C as low as realistically possible relative to T_H; that is, we want to keep the temperature ratio $\tau = {}^{T_C}/_{T_H}$ as low as possible. We of course also want to keep it as constant as possible, and to that effect will incorporate a cooling fan into the test rig. A small, 12 V box fan was chosen for this application due to its satisfactory airflow and because box fans in general are inherently free-standing and protected. Aiming the fan directly at the cooling fins surrounding the cold space of the engine (see Fig. 3.1) will maximize the effectiveness of the fan, but in this configuration we

also run the risk of sending unwanted air currents to the heating element, impairing its ability to maintain a constant temperature. To combat this, we set up a simple, metal wind barrier between the hot and cold sides of the engine to contain airflow to the cold side.

4.2 Final Assembly

With consistent engine temperatures ensured by the upgraded heat source and air-cooling system, we are at last ready to complete our experimental set up. The modified engine was mounted to a new, solid aluminum baseplate into which a hole was drilled to accept the ceramic heating element clamp. The new baseplate was also made significantly larger than the original to make room for clamps anchoring the engine to the experiment table. As one might expect, the rapid, reciprocating motion of the displacer and power pistons produces a significant amount of lateral vibration, sufficient to literally vibrate the engine off of the table, a catastrophic result precluded by the use of clamps. Large C-clamps were used as they are both strong and simple to install and remove. With the engine in position, the windscreen was installed between the hot and cold sides of the displacer cylinder, and ended up doing double duty by also serving as a retaining bracket for the cooling fan. Next, the Prony brake was assembled around the crankshaft, and a digital mass scale, bought online, was moved into position under the metal tip of the brake arm. A thin piece of metal was slipped between the metal tip and the scale to help protect the scale pan from scratches caused by the sharp brake arm tip.

Proper positioning of the brake arm relative to the scale pan is imperative to an accurate determination of engine torque. As is the norm among scales, the pan of the scale used in this study is confined to move vertically and hence can only respond to forces that act vertically downward. This means that in order to accurately translate the reading on the scale into the torque to which it corresponds, we need to ensure that all the force $\vec{F}$ exerted by the tip of the brake arm on the scale pan acts vertically downward. That way, the force exerted *by the scale* on the brake arm $\vec{F}_{\text{scale}}$ will act perpendicularly to the lever arm with a magnitude equivalent to that of $\vec{F}$. Since torque τ is the product of the length of the lever arm R and the perpendicular force [Eq. (2.2)], to properly convert the reading on the scale into a torque we must therefore require that

$$\vec{F} = -F_z\,\hat{z}, \tag{4.1}$$

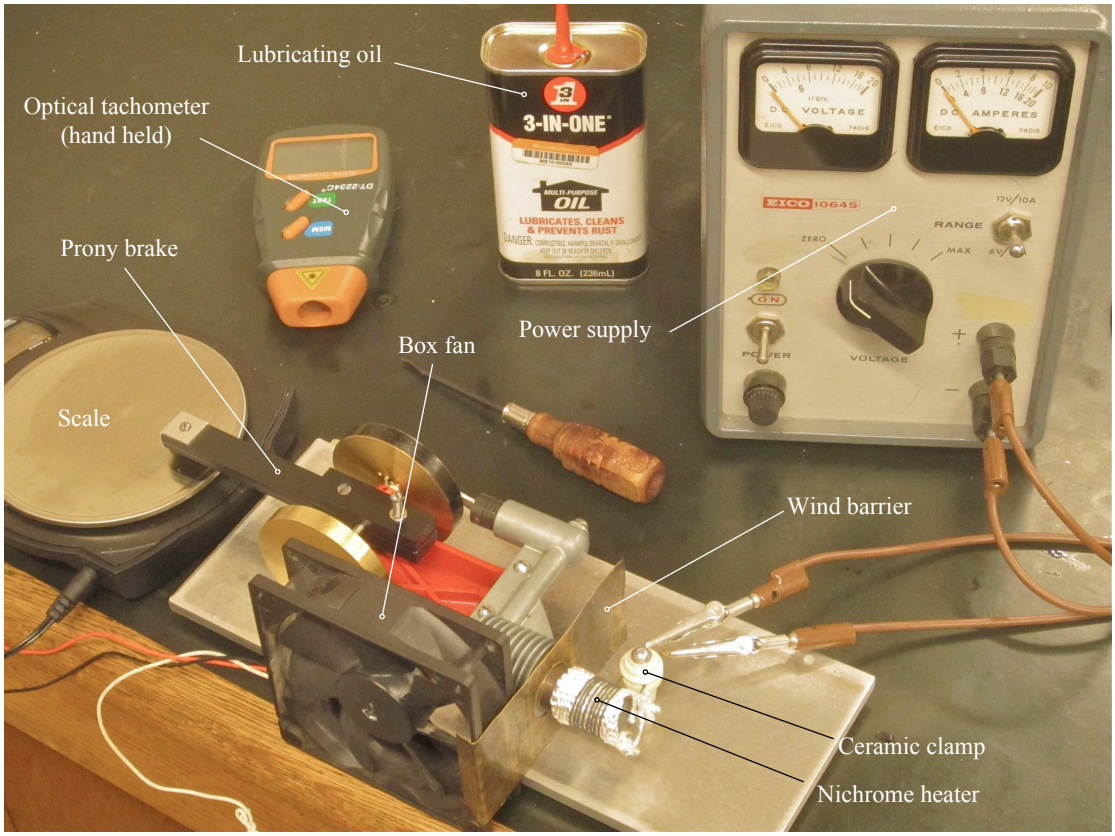

Figure 4.1: This image shows all the key components of the experimental setup as they were positioned during testing. The engine itself, the Prony brake, fan, scale, optical tachometer, power supply and nichrome heater are all visible. Not visible, however, are the C-clamps that generally anchor the engine baseplate to the table. They were removed prior to the image being taken to make the other components more apparent.

where F_z is the component of $\vec{F}$ along the z-axis, and $+\hat{z}$ is the usual basis vector pointing vertically upwards from the scale pan as seen in Fig. 4.2. Since $\vec{F}$ acts perpendicularly to R, satisfying Eq. (4.1) requires that the lever arm rest in the x–y plane so that it is perpendicular to $\hat{z}$. That is, for the scale read-out to properly correspond to engine torque, we need the brake arm to be parallel to the table when fixed to the crankshaft at one end and resting on the scale at the other. In such a configuration, we have $|\vec{F}| = |\vec{F}_{\text{scale}}|$, as desired, due to the fact that $\vec{F}$ is acting entirely in a direction in which the scale can and will respond. Fig. 4.2 shows the different forces acting between the scale pan and the tip of the brake arm, illustrating the required equivalence of $\vec{F}$ and $F_z\,\hat{z}$. In configuration (a), where the brake is *not* properly positioned, not all the force exerted on the scale by the brake is registered by the scale because $\vec{F} \neq -F_z\hat{z}$. This is due to a component of $\vec{F}$, namely F_x, that acts in the $-\hat{x}$ direction and hence cannot be detected by the scale.

As it happened, the height of the pan of the scale used for testing demanded that the engine baseplate be propped up slightly in order for the Prony brake to lie level as needed. To that effect, two quarter-inch-thick metal spacer plates were sandwiched between the engine baseplate and the experimental table, one at either end of the baseplate, with the clamps applying pressure to the ends of the baseplate, over the spacers. Elevating the engine relative to the scale pan so as to optimally orient the Prony brake marks the final step in preparing the experimental apparatus, or engine test rig. Fig. 4.1 shows an image taken of the test rig used in this study. The engine, heat source, heat source power supply, cooling fan, wind barrier, Prony brake, and scale are all visible and are positioned as they were during testing.

4.3 Conducting Trials and Obtaining Data

As generally expected, a number of preparatory steps preceded each trial run. Firstly, either the phase angle α or the swept volume ratio κ had to be set, depending on which of the two parameters was being measured in a particular trial. Adjusting the swept volume ratio was a trivial matter of moving each connecting rod from one hole on its respective flywheel to another. A protractor was used to set the phase angle, with $\alpha = 90°$ for all swept volume ratio tests.

After fixing α and κ, the first step to taking measurements of the engine's performance was to give the engine adequate time to reach a stable operating temperature and speed with the Prony brake only loosely attached (with the bottom plate hanging on the screws, not contacting the crankshaft). As the Stirling derives its power from

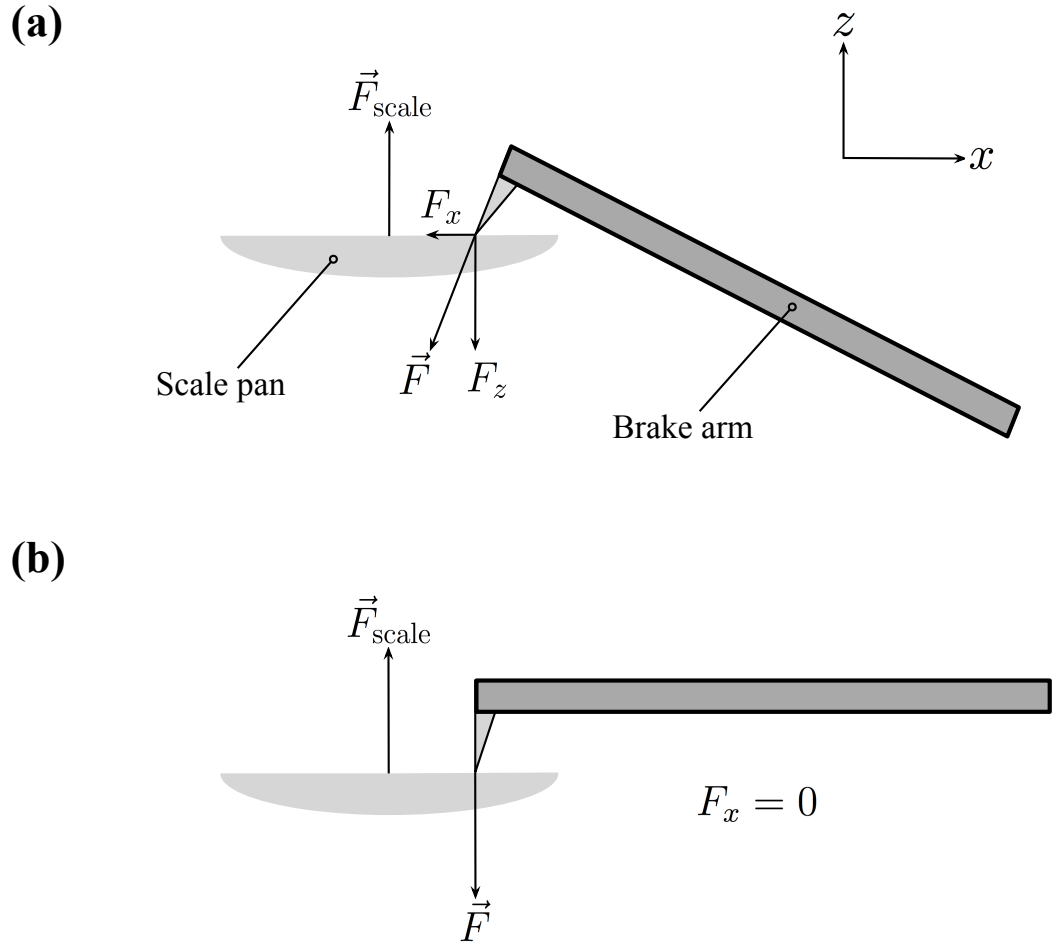

Figure 4.2: A diagram showing the forces acting between the metal tip of the Prony brake arm and the scale pan. Configuration (a) shows how improper brake arm orientation (that is, not horizontal with respect to the table) can lead to an improper calculation of torque. With the arm at an angle, the force exerted by the scale on the brake $\vec{F}_{\text{scale}}$ is not equal and opposite to the total amount of force $\vec{F}$ with which the arm is pushing on on the scale because of the component of force acting in the $-\hat{x}$ direction, F_x. Due to the scale being constrained to movement solely along the z-axis, F_x is a force not combatted and hence not registered by the scale. Inadequate scale readings would would cause torque calculations to return values that are too low. Configuration (b), in contrast, displays the brake arm in the necessary and proper position. Here, $\vec{F} = -\vec{F}_{\text{scale}}$, and the reading on the scale properly reflects the perpendicular force required to produce a torque on the brake *about* the crankshaft sufficient to oppose the torque on the brake *from* the crankshaft.

a temperature differential across the engine body, it was imperative to let the engine run untouched until a stable differential had been reached. Otherwise, the engine would be literally more powerful by the end of the trial run than at the beginning. At least five minutes were dedicated to allowing the engine temperatures to stabilize; stability was confirmed by monitoring engine speed, with a constant engine speed indicating a steady temperature differential. Engine speed, not engine temperature, was chosen as an indicator of the engine having reached a stable state because engine speed, not engine temperature, is directly involved in our calculation of power. The optical tachometer also proved more consistent and controllable than the thermocouple thermometer used to measure engine temperatures as the thermometer tended to record different values for, say, the hot space, when contacting different parts of the hot cylinder.[1]

After reaching a stable state, the engine was stopped briefly so that the scale, with the brake resting on it, could be zeroed so its readings would not reflect the force on the scale from the weight of the brake arm itself. Once restarted from this pause, the engine was again given time to reach a steady state. As the engine was warm when momentarily stopped, and as the heat source was left on while the scale was zeroed, this second waiting period only lasted thirty seconds or so. At this point, data taking could begin. The first data point was taken with the brake still only loosely attached, and even still the friction between the brake arm and the crankshaft produced a torque sufficient to register on the scale. The reading on the scale and the engine speed at which the reading was taken were both recorded.

Tightening both screws at the same rate, the lower plate of the brake was gradually brought into contact with the crankshaft with care being taken to ensure that, once fully contacting the crankshaft, the lower plate was parallel to the upper plate. This helped ensure that the brake, when tightened further, exerted its brake force on the crankshaft as evenly as possible in hopes of preventing any compromising vibrations between the crankshaft and the brake [24]. Any increase in brake force from this point on was accomplished using the adjustment screw alone (see Fig. 3.6). Each time brake force was increased, the engine was given time to respond and stabilize at a new, lower speed, at which point the engine speed and scale reading were both recorded. This

[1] A thermocouple is a device made of two dissimilar metal wires connected directly together at one end (the measuring end) and connected at the other through a circuit. The non-measuring junction point is kept at a reference temperature so that a temperature *difference* between the two ends is measured. Heating the measuring end produces electromotive forces of different magnitudes in the two wires causing electrons to jump from one wire to the other across the junction point, thereby creating a temperature-dependent voltage drop across the junction. This voltage can then be accurately correlated back to its corresponding temperature.

at times hypnotic process of increasing brake force, waiting for stability, then making and recording measurements was continued until the engine speed was decreased to the point where any subsequent increase in brake force would, despite even the most sensitive touch at the adjustment screw, slow the engine to a stall. As engine performance did in fact change with variation in κ and α, the engine speed at which the test was concluded differed from trial to trial, but this does not significantly affect our results. The temperatures T_C and T_H, as well as the ambient temperature T_{amb} were recorded at regular intervals during each trial.

To help maintain proper engine function throughout the entire testing period, the engine was disassembled, cleaned, re-lubricated, and reassembled after every three trial runs, or after every 3-4 hours of engine operation. Though Stirling engines tend to run very clean because at no point does engine lubricating oil contact or combine with combustion bi-products (namely particulate matter, or soot), cleaning was performed with high regularity as a precaution against premature wear of engine components. To keep the engine running smoothly throughout each test, the crankshaft pivots, connecting rod pivots, power piston, and displacer piston rod were lubricated with a lightweight, multipurpose oil at regular intervals during the trial. Periodic lubrication was especially important during trials involving a long piston stroke as side-loading on the pistons is highest for longer strokes, and the consequences of side-loading, namely excessive wear between the piston and the cylinder, or else the piston binding up completely within the cylinder (a catastrophic failure), are most effectively precluded by the presence of a healthy film of oil between the piston (or piston rod) and its cylinder.

Chapter 5

Presentation and Analysis of Data

As suggested in the previous chapter, the data at the end of each trial presented itself as a list of paired numbers: a mass (in grams) and the engine speed (in RPM) at which the mass measurement was taken. This data was not immediately useful to our cause of testing the applicability or accuracy of the Senft model, however, as it makes no direct mention of cyclic work or engine power. As is so often the case, the quantities we measure are not the quantities we seek, and we must arrive at the latter through calculations performed on the former. While relatively simple, these calculations are tedious in our case and hence a Mathematica code was developed to perform them quickly and easily. The code began with a list of raw data points of the form {RPMs, mass} from which it would generate plots of engine torque and power versus engine speed. The section of code that produces these plots is reproduced in Appendix B, Fig. A.2. To facilitate judgement of the effectiveness of the Senft model in predicting the relationship between power output and engine speed ω, additional code was written to plot both measured power output (versus RPM) and theoretically predicted power output on the same axes so that differences, or similarities, between the two curves could be readily identified. To accomplish this, the program would run the Senft model under the same conditions of the particular trial run being analyzed; that is, it would calculate W_{cyc} from Eq. (2.12) across a range of engine speeds ω with the swept volume ratio κ, the temperature ratio τ, the dead volume ratio χ, and the phase angle α all equal to the values they took during the trial. The code would generate a list of data of the form {RPM, predicted power output} which could then be plotted with ease. The code used for this operation is reproduced, with explanatory comments, in Appendix B, Fig. A.1.

5.1 The Data, Plotted

With the code written thus, analyzing the data from each trial run was a simple matter of entering temperature and phase angle values as well as the stroke settings (n1 and n2 in the code of Fig. A.1) into the Mathematica programs and pressing enter. The true labor was found in the data collection process, which turned out to be very time consuming with each trial taking on the order of one hour to complete. The time required to service the engine every three trials added up quickly as well, with the effect that not as many trials were performed as initially intended. As designed, the upgraded engine flywheels, offering four displacer piston stroke settings and six power piston settings, enabled testing of 24 distinct swept volume ratios as well as an infinite range of phase angles. However, due to the ineluctable presence of deadlines and time constraints, only ten swept volume ratios and six phase angles were tested. That said, the swept volume ratios tested still ranged from 0.48 to 1.0, giving a more or less satisfactory spectrum broad enough for the relationship between κ and power output to be clearly visible. Similarly, though the number of phase angles tested was relatively small, their range, from 45° to 120° in increments of 15°, allowed testing values of α both less than and greater than 90°, and, crucially, was sufficient for clearly exposing the relationship between α and power output.

The processed data is displayed in Figs. 5.1 through 5.8. In Figs. 5.1, 5.2, 5.5 and 5.6, the torque and power curves are color-coded with power displayed in red and torque in blue. As the dimensions of power are not the same as those for torque ($[\text{Power}] = \text{N} \cdot \text{m} \cdot \text{s}^{-1} \neq \text{N} \cdot \text{m} = [\text{Torque}]$), two scale axes are required for each plot, with the power curve scaled to the left vertical axis and the torque curve to the right. Similarly, Figs. 5.3, 5.4, 5.7 and 5.8 also display two curves per plot, but this time both correspond to power and hence a single scale axis could be used. However, due to the sheer difference in range between the measured power curve (displayed again in red) and the theoretical, indicated power curve (in orange), two scale axes were chosen with the measured power scaled to the left axis, the indicated power to the right. Despite the slightly confusing nature of having two scale axes for a single plot, if both curves had been scaled to the same axis, either the indicated power curve would be appropriately displayed with the measured power curve inconveniently reduced to an indistinguishable set of points along the RPM axis, or the measured power curve would be displayed with the indicated power curve not even visible on the plot.

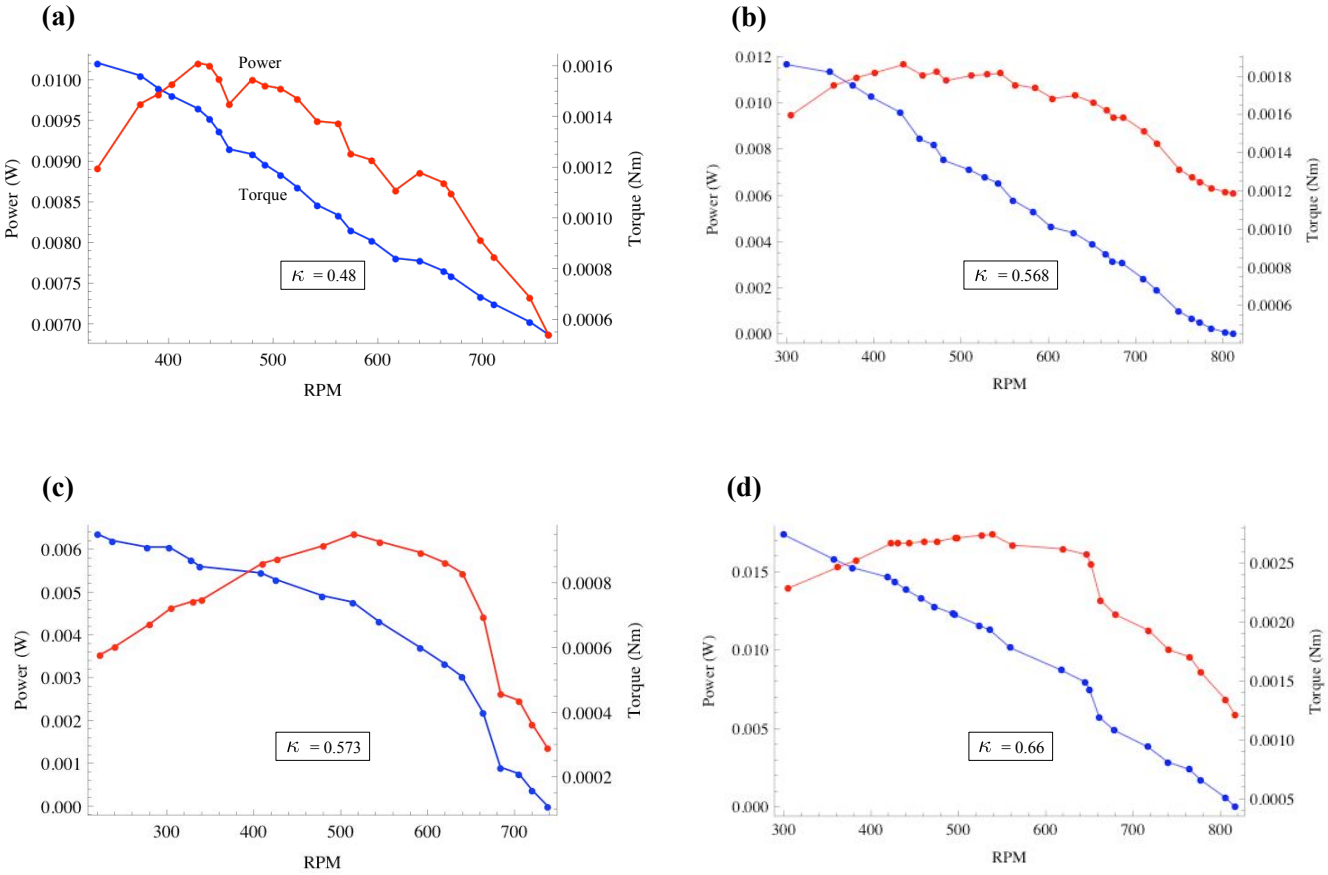

Figure 5.1: Shown above are the torque and power versus RPM plots for the first four swept volume ratios tested. The power curves are shown in red and scale to the left vertical axis while the torque curves, in blue, scale to the right. Each trial was conducted with the phase angle $\alpha = \pi/2$, and the swept volume ratio κ for each trial is shown on its respective plot. Worth noting is the difference in maximum power output between plot (b) and plot (c). While the swept volume ratios for those two trials differs by less than 1%, and the temperature ratio τ by 7%, the maximum power output in plot (b) is nearly twice that of plot (c). This impressive difference cannot be accounted for by the Senft model, which would predict roughly equal power outputs for the two trials given the slightly dissimilar κ and τ between the two.

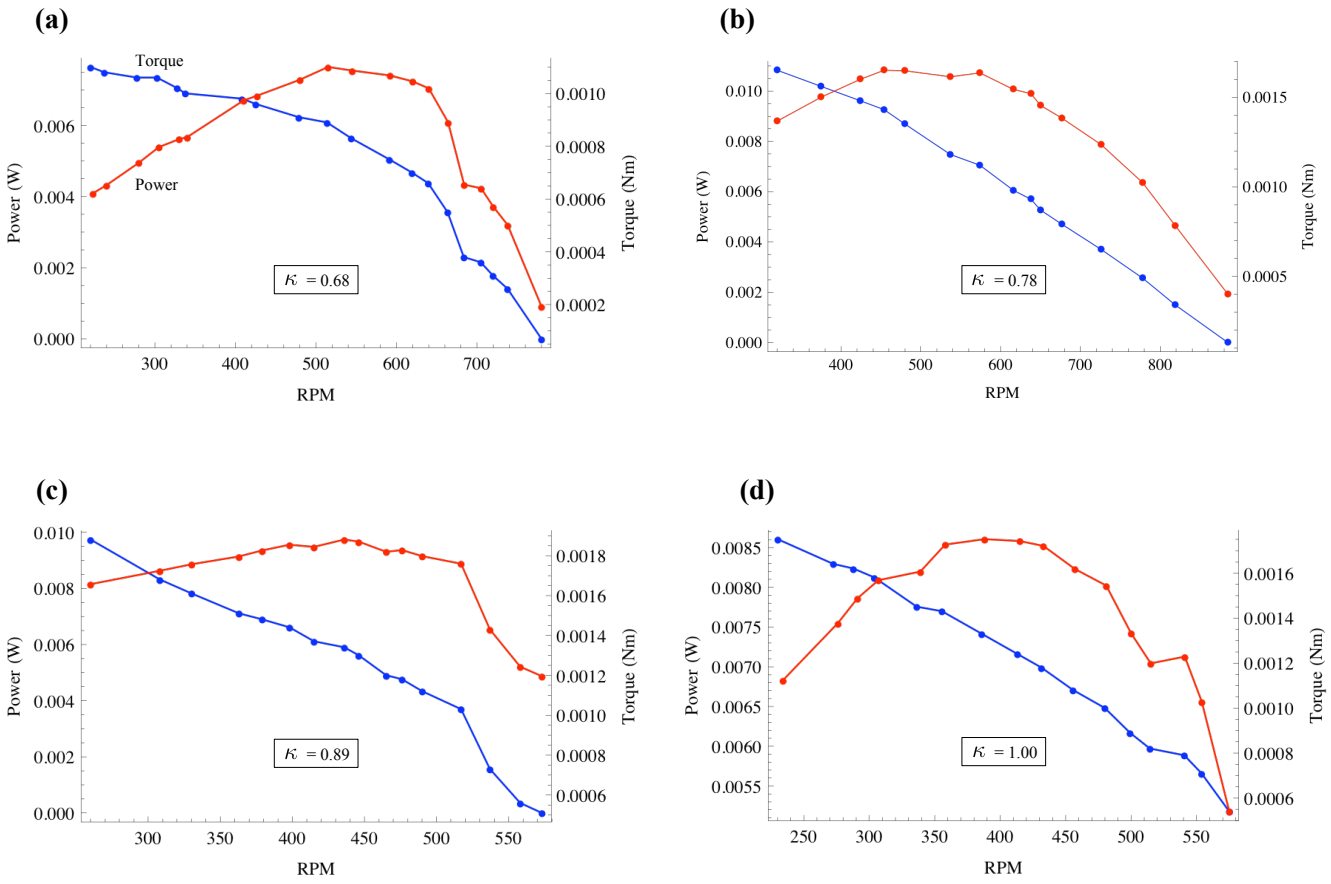

Figure 5.2: Shown above are the torque and power versus RPM plots for the last four of eight swept volume ratios tested. The power curves are shown in red and scale to the left vertical axis while the torque curves, in blue, scale to the right. For all trials, $\alpha = \pi/2$, and the swept volume ratio for a given trial is shown on its respective plot. The sharp kink at roughly 440 RPM in the power curve of plot (b), and the kink at 525 RPM in the power curve of plot (d) are more likely due to experimental error than to any curious power delivery characteristics of the engine.

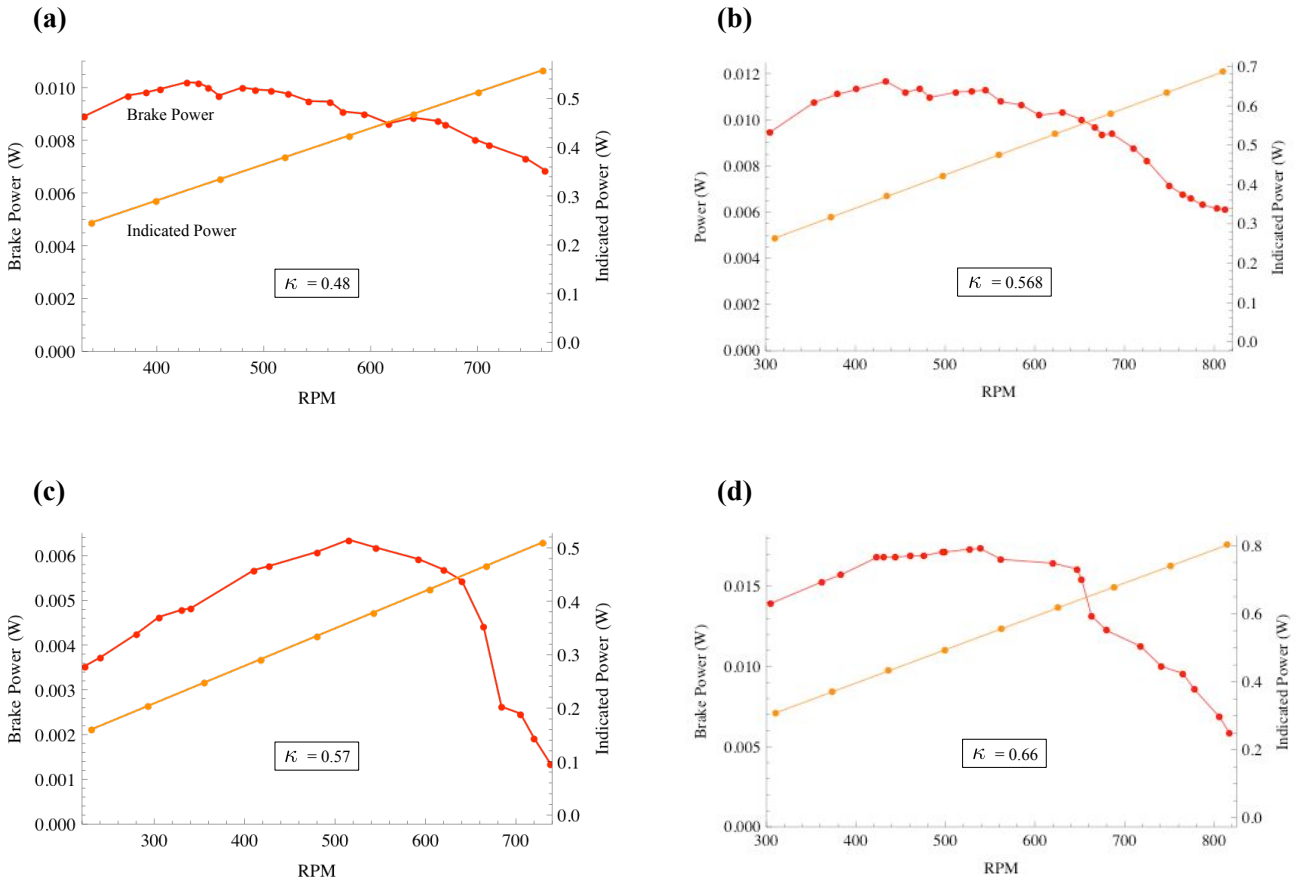

Figure 5.3: Shown above are plots comparing measured output power (brake power) to the indicated output power predicted by the Senft model for the first four swept volume ratios tested. The measured (brake) power curves are shown in red and scale to the left vertical axis while the theoretical (indicated) power curves, in orange, scale to the right. For each trial, $\alpha = \pi/2$. Note that for each trial, for a given RPM the indicated power output is orders of magnitude greater than the measured power output. Further note that while the Senft model predicts that output power will increase continually and linearly with RPM at a rate determined by W_{cyc} (since $P_{\text{indicated}} = W_{\text{cyc}}\omega$), the measured power output falls off after a certain maximal value. Lastly, though the slopes of the measured power curves appear to agree with the slopes of the indicated power curves until near the point of maximal power output (which would suggest, falsely, that W_{cyc} is the same experimentally as theoretically), this is merely a coincidental consequence of having two different scale axes on the same plot.

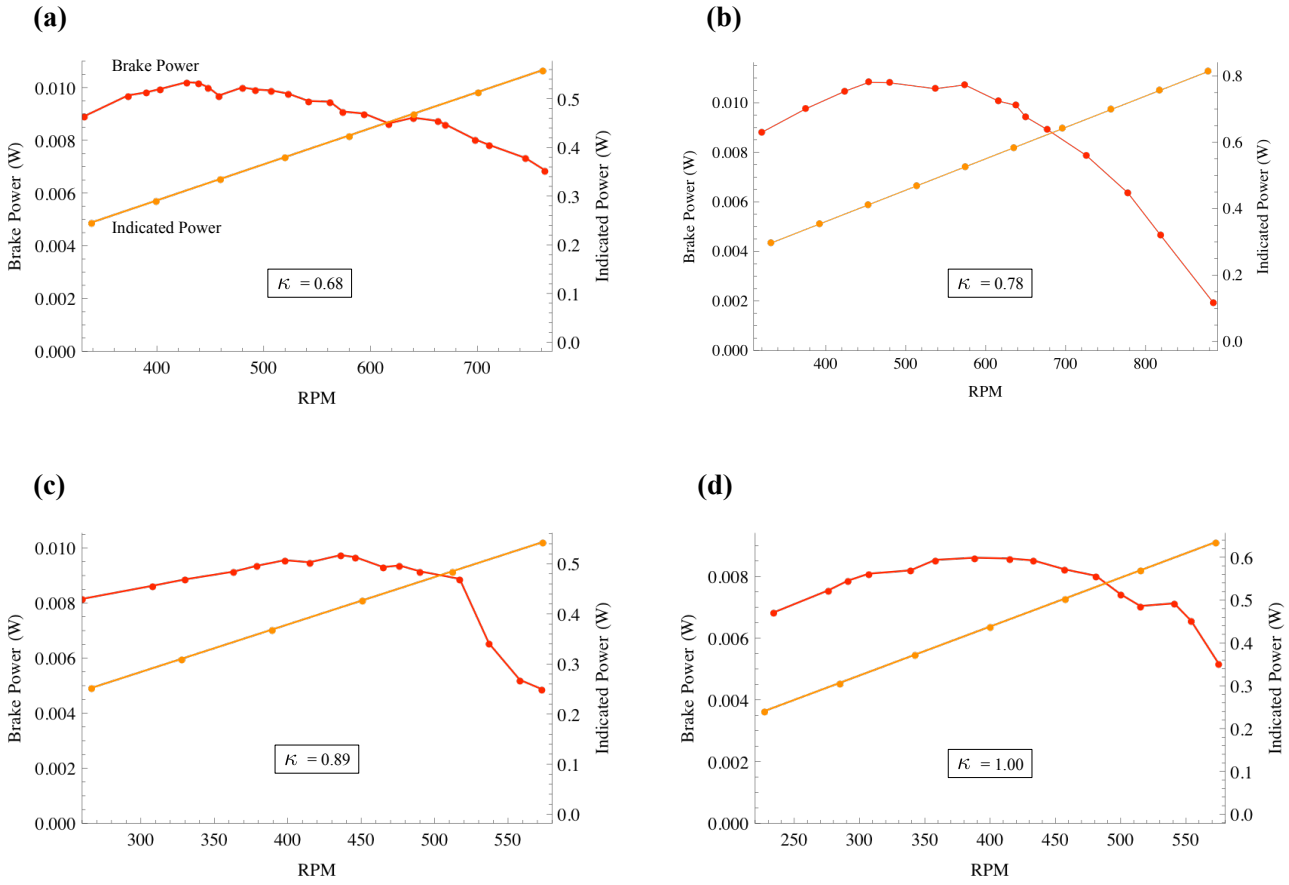

Figure 5.4: Shown above are plots comparing measured output power (brake power) to the indicated output power predicted by the Senft model for the final four of the eight swept volume ratios tested. The measured (brake) power curves are shown in red and scale to the left vertical axis while the theoretical (indicated) power curves, in orange, scale to the right. For each trial, $\alpha = \pi/2$. Note that for each trial, for a given RPM the indicated power output is orders of magnitude greater than the measured power output. Further note that while the Senft model predicts that output power will increase continually and linearly with RPM at a rate determined by W_{cyc}, measured power falls off after it reaches a certain maximal value.

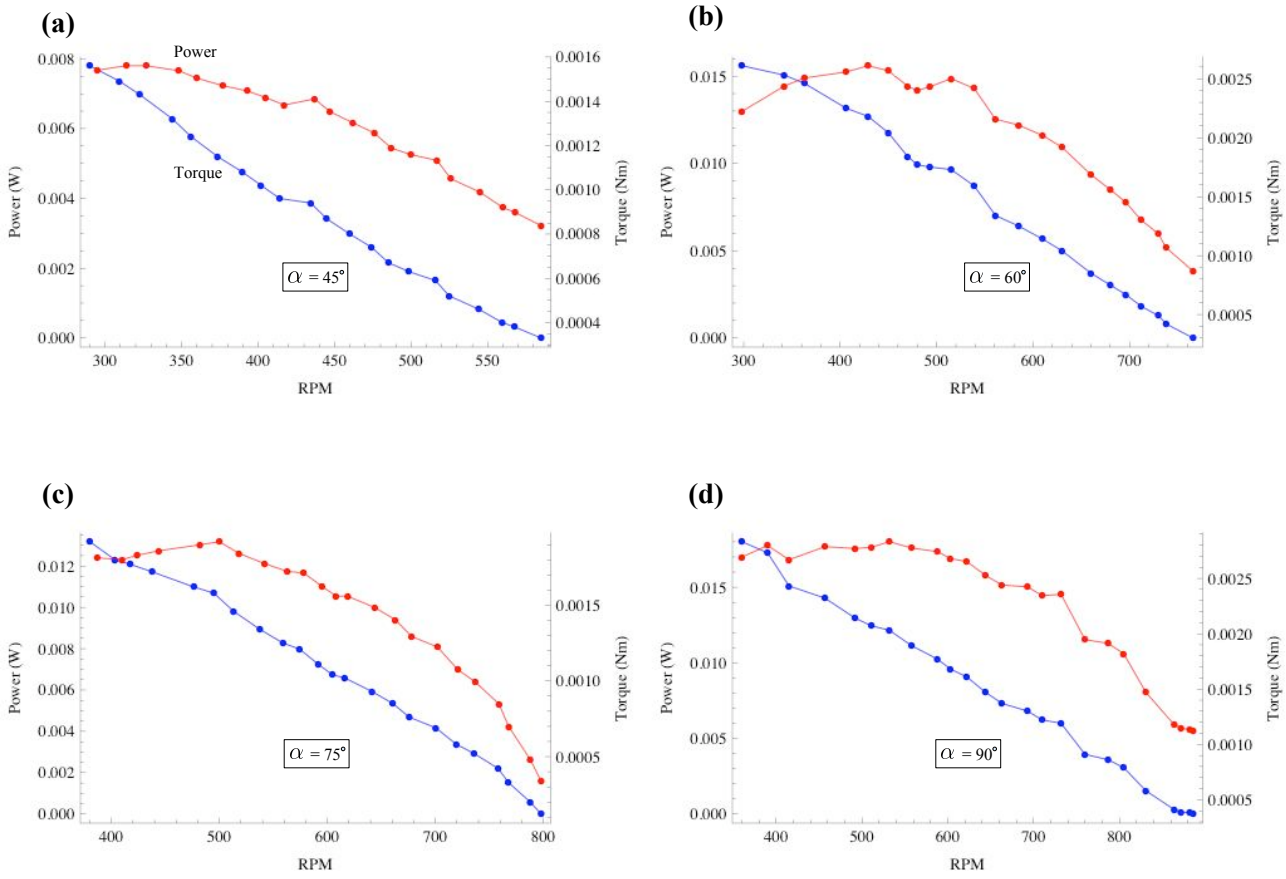

Figure 5.5: Torque and power versus RPM plots for the first four phase angles tested. The power curves are shown in red and scale to the left vertical axis while the torque curves, in blue, scale to the right. Each trial was conducted with the swept volume ratio $\kappa = 0.57$; the phase angle α for each trial is shown on its respective plot. Generally, output power appears to increase with increasing α.

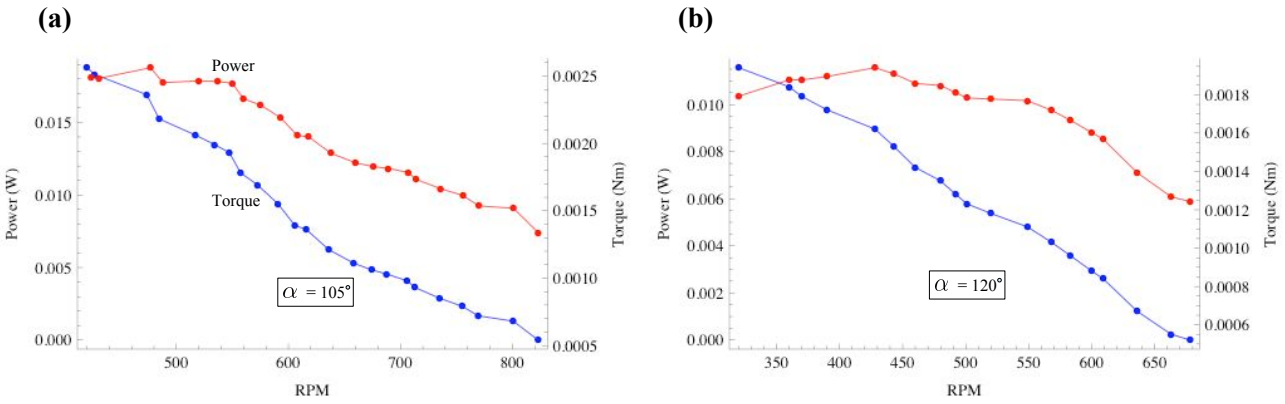

Figure 5.6: Torque and power versus RPM plots for the final two phase angles tested. The power curves are shown in red and scale to the left vertical axis while the torque curves, in blue, scale to the right. Each trial was conducted with the swept volume ratio $\kappa = 0.57$; the phase angle α for each trial is shown on its respective plot. The 'trend' in these final two plots is for power output to decrease as α is increase past $\approx 105°$.

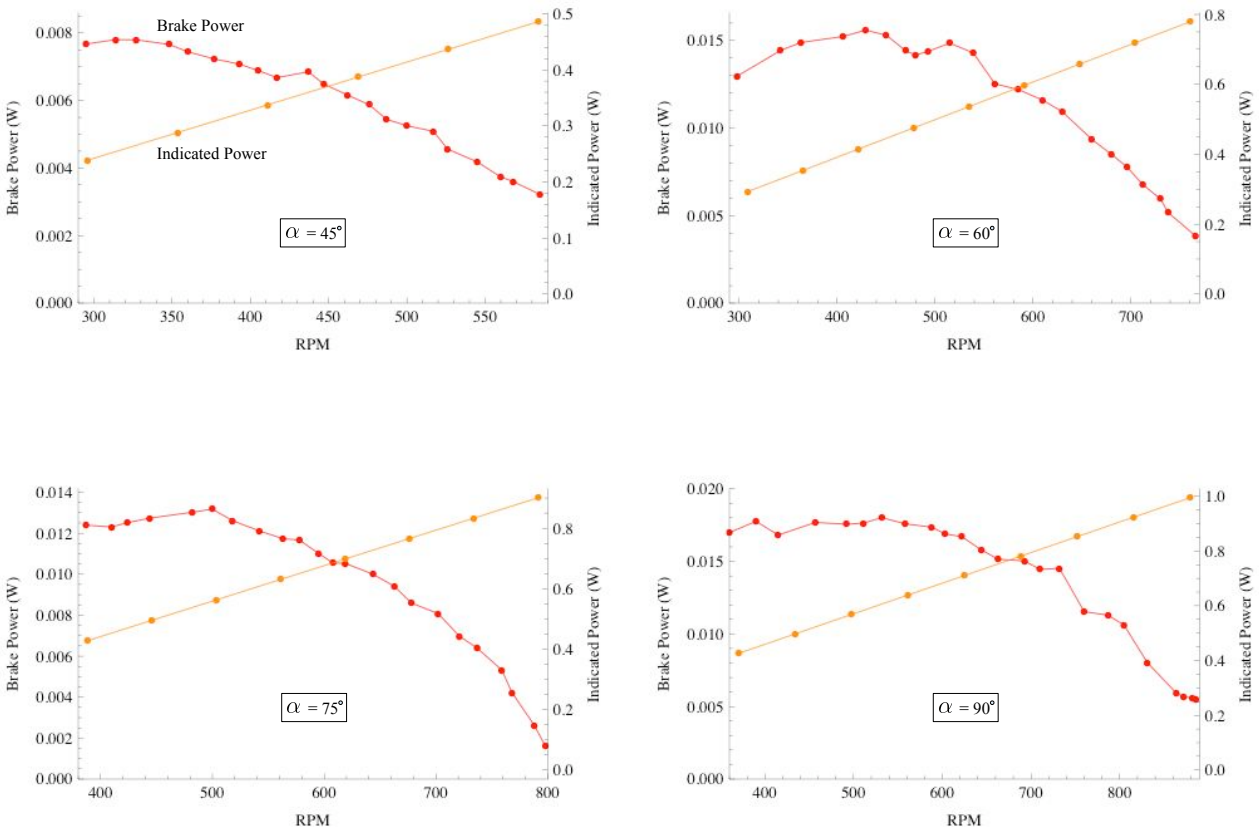

Figure 5.7: Shown above are plots comparing measured output power (brake power) to the indicated output power predicted by the Senft model for the first four of six phase angles tested. The measured (brake) power curves are shown in red and scale to the left vertical axis while the theoretical (indicated) power curves, in orange, scale to the right. For each trial, $\kappa = 0.57$. Note that for each trial, for a given RPM the indicated power output is orders of magnitude greater than the measured power output. Further note that while the Senft model predicts that output power will increase continually and linearly with RPM at a rate determined by W_{cyc}, measured power falls off after it reaches a certain maximal value.

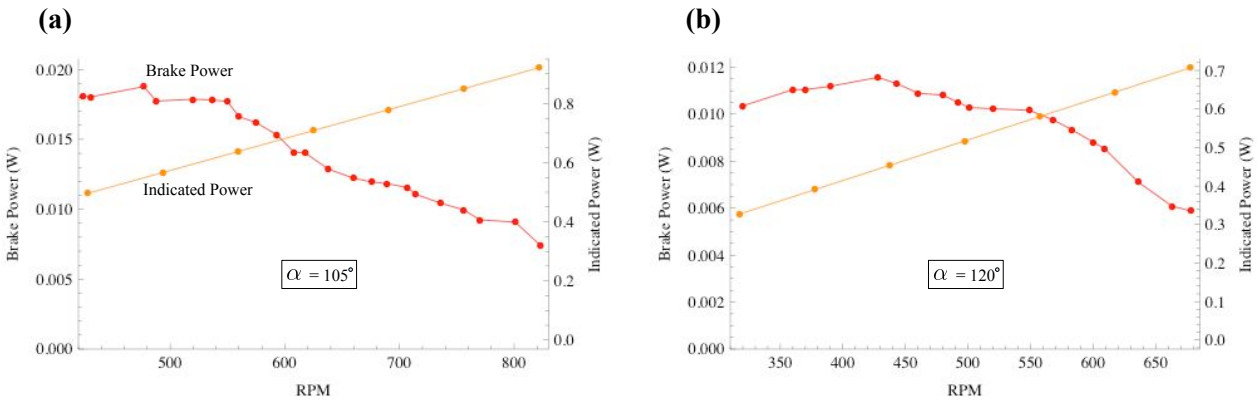

Figure 5.8: Shown above are plots comparing measured output power (brake power) to the indicated output power predicted by the Senft model for the final two of six phase angles tested. The measured (brake) power curves are shown in red and scale to the left vertical axis while the theoretical (indicated) power curves, in orange, scale to the right. For each trial, $\kappa = 0.57$. Note that for each trial, for a given RPM the indicated power output is orders of magnitude greater than the measured power output. Further note that while the Senft model predicts that output power will increase continually and linearly with RPM at a rate determined by W_{cyc}, measured power falls off after it reaches a certain maximal value.

5.2 Analysis

What should be immediately apparent from the figures in the previous section is that the data obtained does not obviously agree with theory. The task at hand, then, is to determine whether this absence of evident agreement is the result of a poor experimental design or poor execution of the design, or if it is the consequence of a flawed model. Even making simple, qualitative observations of the data presented in Figs. 5.1–5.8 can prove quite telling here, and is as good a place to start as any.

5.2.1 Brake Power vs. Indicated Power

Despite the jaggedness of some of the measured power curves in Figs. 5.1–5.8, for all of them there exists an RPM or a range of RPMs (meaning an ω or range of ω) at which the engine is maximally powerful. Even plot (a) in Fig. 5.5, the trial most dubiously in agreement with this statement, shows the beginnings of a downward trend away from a point of maximum power as RPMs decrease from approximately 310. Moreover, increasing engine speed beyond the optimal RPM in all trials tends to result in a *reduction* in power commensurate with the simultaneous reduction in torque with increasing RPMs. The Senft model, however, cannot account for this trend between power and ω. As seen in Figs. 5.3–5.4 and 5.7–5.8, the power output predicted by the model increases linearly, and indefinitely, with RPM. This disparity arises out of the difference in where on the engine the model assumes the power measurement is being taken, and where on the engine power was actually measured in this study. Torque measurements, from which power was calculated, were made using the Prony brake which was connected to the engine at the crankshaft. Hence the output power derived from the torque measurements taken at the crankshaft can only reflect the power output available *at the crankshaft*. What we have to this point been calling the brake power is occasionally referred to as the shaft power purely because it is measured at the final output shaft of the engine.

The model, however, assumes a very different point of measurement. As the derivation of the model presented in Section 2.3 shows, the model ignores the mechanisms that are responsible for, and that capitalize on, the changes in pressure within the engine space. Recall from Section 2.3.1 that the model takes *as an assumption* the sinusoidal motion of the pistons on which the volume and pressure of the working fluid depend without any consideration of how this motion is achieved. Eq. (2.6),

$$V(t) = V_D \left(1 + \frac{\kappa}{2}(1 + \cos\omega t) + \chi\right),$$

which gives the total instantaneous engine volume, certainly makes no mention of the connecting rods, flywheels, and pivots that make possible the time-dependent, sinusoidal piston motion evident in the $(\cos \omega t)$ term. Instead of predicting the power output at the crankshaft, which doesn't exist as far as the model is concerned, the model predicts the power available at the junction between the working fluid and the piston, the latter of which it treats as nothing more than a moveable wall of its container, the engine body. Measuring power at this junction gives the *indicated power* for an engine, and entails recording the pressure and volume variations within the engine cylinders, and using that data to create a p–V diagram for the real engine cycle. The indicated work per cycle can then be determined easily as it is equal to the area of the closed loop in the diagram that marks the path taken by the working fluid. This is a consequence of the relation $W = \int p \, \mathrm{d}V$ [Eq. (1.3)], as well as *where* during the cycle work is actually performed. With reference to the p–V diagram for the ideal Stirling cycle, Fig. 1.3, the working fluid does some amount of work on the piston as it follows the isotherm from (c) to (d), and has an amount of work done on it as it is compressed along the isotherm from (a) to (b). The cyclic indicated work is the net difference between these two amounts, which by the above integral relation, is equivalent to the area beneath the upper isotherm (total work performed *by* the fluid) minus the area beneath the lower isotherm (total work performed *on* the fluid), or simply the area enclosed by the loop in the diagram.

5.2.2 Thermodynamic Power Losses

The salient consequence of indicated power and brake power requiring measurements at different points on the engine is that indicated power does not take into account the myriad frictional and other thermodynamic losses that occur between the working fluid doing work on the piston, and that work being transformed into a torque about the output shaft. Indicated power is a measurement of engine power before the thermodynamic losses in whatever mechanisms (pistons and seals, linkage rods and pivots, etc.) existing between the engine cylinders and the output shaft of the engine have taken effect. Hence, brake power equals indicated power minus all thermodynamic losses from relevant engine components that stand between the output shaft and the cylinders in the powertrain.[1] Hardly a surprise, then, that the measured

[1] As a natural extension, rear wheel power, a useful measurement for motorcycles or rear-wheel-drive cars, equals indicated power minus losses due to engine components, *minus* losses due to final drive components (such as a driveshaft, differential, or chain) existing between the output shaft of the engine and the rear wheel itself.

power output was so much less than the theoretically predicted power output across all trials. The Senft model wasn't accounting for frictional and other thermodynamic losses that inevitably accompany a real engine.

The thermodynamic losses that the model does not consider can further explain the earlier conundrum of the measured power curves having distinct maxima despite the model predicting a steady, linear increase in power with RPM. Notable among these losses are those that arise from heat transfer from the gas to the engine body, from working fluid leakage around seals, and from aerodynamic friction, or viscous fluid, losses [10, 25]. The losses from heat transfer are exacerbated by the fact that the engine body, typically made from iron, steel, or aluminum, is thermally conductive and hence readily absorbs heat. These thermal losses are more consequential for a study of maximum engine efficiency, but they nevertheless influenced this study as they effectively made T_H lower, and T_C higher, than the measured. Recall that T_H and T_C reflect working fluid temperatures, but were recorded by measuring the surface temperature of the engine at the hot and cold ends respectively. Assuming ideal heat transfer, these engine temperatures should be the same as the gas temperatures, but the real effects of thermal losses preclude such a convenient equivalence. Hence, the values for T_H and T_C input into the Senft model indicated temperatures that were higher or lower, respectively, than those the working fluid actually reached. Engine speed likely has an effect on the magnitude of these thermal losses for with increased engine speed comes a decrease in the amount of time available for heat transfer.

Somewhat ironically, the losses due to air leakage stem from an effort to keep another thermodynamic loss, friction, to a minimum. Instead of o-rings or the tight metal ring seals like those on the pistons of internal combustion engines, the seals between the power piston and its cylinder, and between the displacer piston rod and the bushing at the end of the cylinder, consist solely of oil films occupying the minute clearance gap (on the order of 0.001 inch thick) between the components. While certainly low-friction, this type of seal is (consequently) not the most airtight, permitting leakage of the working fluid leading to engine pressures that are lower than theoretically expected. Such leakage is visible if the power piston is manually inserted into its cylinder: the pressure that builds in the engine space as the piston is pushed in dissipates after a few seconds if the piston is held in. In addition to mucking with engine pressures, working fluid leaking into and out of the engine through these seals produces a heat pumping effect that can, like the thermal losses discussed above, make T_H lower than measured [26].

The viscous fluid losses, or flow losses, are the final major form of thermodynamic

loss not accounted for by the Senft model. These losses originate mostly from the working fluid needing to flow through the narrow annular gap between the displacer and the inner cylinder wall with each pass of the piston, as well as from the fluid being forced back and forth between the thin chamber connecting the displacer and power cylinders (see Fig. 2.2). Since the working fluid, though a gas, is nevertheless viscous, forcing it through these tight gaps and passages with each engine cycle requires energy that ends up being taken from the engine. As one can imagine, flow losses are highly dependent on engine speed for as ω increases, the frequency with which the working fluid must move back and forth through the annular gap and the connecting chamber increases [25]. Since energy is expended in forcing the working fluid through these tight gaps, as engine speed increases the rate of expending energy increases, and hence the power absorbed also increases.

Due to the velocity-dependent nature of flow losses, it is likely that beyond a certain RPM, the power losses incurred may overwhelm what would otherwise be an increasing power output resulting from increasing engine speed ω following Eq. (2.3). This statement is deliberately weak since even attempting to describe the effects of dynamic losses such as leakages and flow losses on power output requires third-order numerical modeling techniques beyond the scope of this thesis [25]. Even friction, a more familiar impediment to mathematical descriptions of real systems, is generally difficult to accurately represent in models. Indeed, perhaps due to the scarcity of experimental data on the subject of friction losses in Stirling engines, there exists no equation in the literature correlating frictional forces and power losses [26]. However, friction tends to be dependent on the surface area of the body in question, so for a small engine such as the one used here, one that necessarily has a large piston-to-cylinder contact patch relative to the small internal volume, the power loss effects of friction can be substantial, even exceeding 20% of the engines indicated power [26].

5.3 Conclusions

We have established that the Schmidt model presented here is fundamentally incompatible with the Prony brake method of power measurement as the former assumes a completely different measurement point than the latter method entails. But how can we tell whether or not the measurements made with the Prony brake are reasonable? Could it not be the case that output power did in fact increase steadily with RPM, as the model predicts, but errors in the experimental design and execution of the same skewed the results into the clearly non-linear plots of Figs. 5.1–5.2, and

Figs. 5.6–5.6? Unlikely for a number of reasons. Firstly, the plots from the 14 trial runs all agree in that they exhibit a trend of an increasing, peaking, then decreasing power output as RPMs increase.[2] With torque curves as smoothly trending as they are [with the possible exceptions of plot (c) in Fig. 5.1, and plots (a) and (b) in Fig. 5.2], it is difficult to make the argument that the experimental design and methods were ineffective to the point where the results acquired would indicate a relationship between output power and engine speed so contrary to reality. Based on the clear and certainly repeated trends visible in the plots, the data obtained seems sufficiently self-consistent to warrant *not* being discounted due to its disagreement with the theoretical model presented. Indeed, results from studies performed by Karabulut *et. al.* in [27] and Senft in [2] on the torque and power characteristics of the Stirling engine agree with those presented here in that they, too, show peaking power curves that fall off with increasing RPM, and torque curves that fall similarly. It is worth noting that Karabulut *et. al.* also used a Prony brake type dynamometer in their tests.

Apart from being supported in the literature, the trend of decreasing torque, responsible for decreasing power, makes intuitive sense as well. It seems reasonable to argue that torque would be highest at lower engine speeds, for at such speeds the working fluid is allowed more time to heat up and expand (or cool and compress) during the power (compression) stroke, yielding increased force on the piston and connecting rod during the stroke, and hence more torque on the crankshaft. This explanation isn't quite adequate, however, for the measurements made by Karabulut *et. al.* indicate *peaking* torque curves, just like the power curves [27]. The data presented here does not display peaking torque curves, likely because it became very difficult to take consistent measurements of engine torque at speeds much lower than 250 RPM. Improving the experimental design by modifying the Prony brake and scale system could solve this measurement problem. The scale could be replaced with a strain gauge for more precise, and digitally recordable, force measurements, and and the brake clamp made electronically controllable so that brake force could be applied in the smallest, most controllable increments. Furthermore, the clamp area could be shrunk so that a given turn of the brake adjustment screw would yield a smaller increase in brake force, thereby increasing brake sensitivity even at low RPMs.

Self-consistency and agreement with the literature, except in certain aspects of the torque curves, lends validity to the measured data, and the fact that brake power and

[2]The possible exception to this trend is plot (a) in Fig. 5.5. While the data doesn't show with absolute certainty a point of maximum power (though it suggests one, at around 325 RPM), it does, like the other trials, conclucively indicate a decreasing power output with increasing RPM.

indicated power are by definition different quantities means that the disagreement between measurements and theory is insufficient to discount the model as fundamentally flawed. By taking both the model and the experimental data as at least viable, then, and ignoring the formidable difference in magnitude between the brake and indicated power across all RPMs, we may draw some conclusions from the data as to the accuracy of the model. We are only equipped to make observations of broad trends, but this is sufficient for elucidating how well the model can describe the relationships between output power and swept volume ratio, and between output power and phase angle. This is different from evaluating the extent to which the model can accurately predict the *magnitude* of the power output for a given RPM, which we understand from the previous section to be an incorrect standard of judgement.

As for the relationship between phase angle α and power, the model predicts that, for a given engine speed, if the temperature ratio τ is kept constant but α is allowed to vary, output power will increase until $\alpha \approx 92°$, at which point power will decrease with any subsequent increase in α. With the exception of the trial shown in plot (c) of Fig. 5.5, a similar trend is observed experimentally, with brake power at 500 RPM (arbitrarily decided) increasing from $\alpha = 45°$ to $\alpha \approx 100°$, then falling off with increasing α. Fig. 5.9 plots output power against phase angle α (in radians) according to the model and to the measured data. Despite the difference in scale between the two curves, a similar relationship between α and output power is clearly evident. If time had permitted, more trial runs for each phase angle setting would have been performed to better confirm or deny the existence of this trend. As it stands, however, it appears as though the model predicts a relationship between α and brake power that is in fact visible experimentally.

The validity of the model in predicting the relationship between κ and power output, however, is more questionable. The model predicts that, for a given RPM, engine brake power will increase with increasing κ, all else being equal, but this is not at all what happens in practice. This incongruity is readily visible in Fig. 5.10, which plots both indicated power and measured brake power against κ. Though eight different swept volume ratios were tested, only the data from five different trials could be used in making this plot. The reason for this pertains to *how* κ is adjusted, as well as what variables in the model, Eq. (2.12), need to be kept constant. The swept volume ratio is adjustable by changing the strokes of the power and/or displacer pistons, but with reference to Table 2.1, we can see that the dead volume ratio χ is also indirectly dependent on the displacer piston stroke by virtue of being directly dependent on the swept volume of the displacer piston. Hence, any change

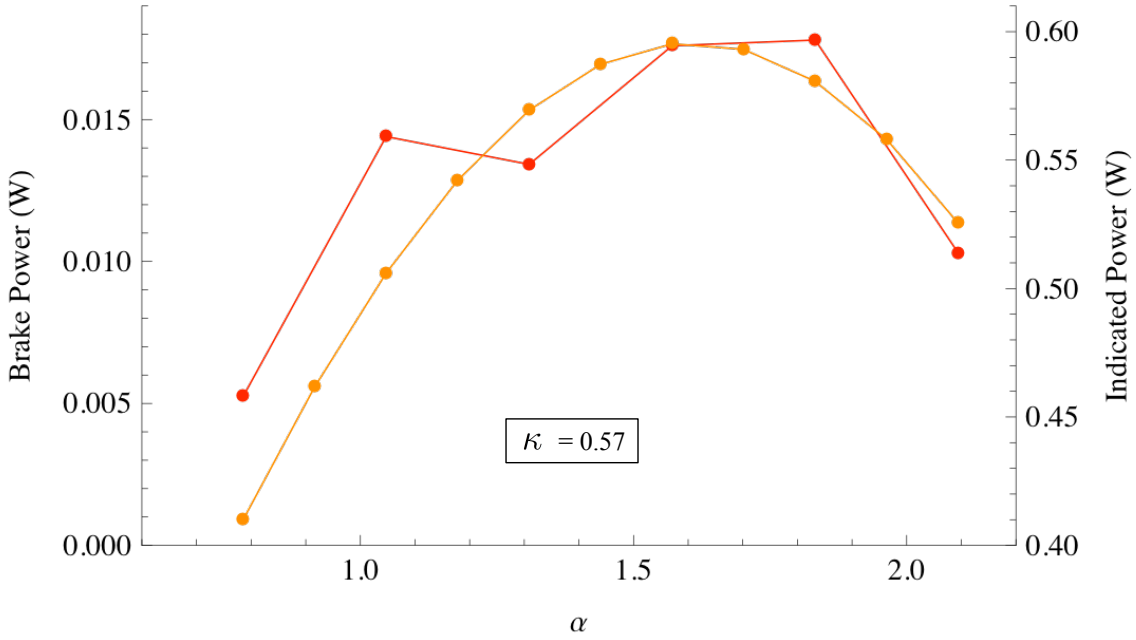

Figure 5.9: A plot of the relationship between phase angle α and output power at 500 RPM according to the Senft model (orange curve) and to experimental data (red curve). The model predicts an increasing power output as α increases to $\approx 90°$, and a decreasing power output as α increases further. A very similar trend is visible in the experimental data, indicating that the model does in fact predict the relationship between α and power output despite it being unable to accurately predict the magnitude of the output power.

in κ requiring a change in displacer piston stroke also results in a change in χ. To determine the relationship between *just* κ and power, we must, however, keep the variables α, τ, and χ all constant across trials. Of the eight swept volume ratio trials, five of them used the same displacer stroke setting, with κ set by the power piston stroke alone. Though the displacer piston stroke was kept constant during these remaining three trials, it was kept at a different value than the other five, and hence the data from all eight trials could not be plotted together. Since three data points, one each from the three remaining trials, hardly constitutes enough data for a plot (though, admittedly, five points isn't much better), these trials were ignored in this section of the analysis.

While the model predicts increasing power with increasing κ, Fig. 5.10 shows that the real engine did not behave in the same way. Instead, power output was observed to increase with κ until reaching a peak at $\kappa \approx 0.8$, at which point power would decrease with any subsequent increase in κ. Plots presented in [2] and [10] show a similar trend, with brake power increasing, peaking, then decreasing with increasing κ. A likely reason for this disagreement between measured data and the model's predictions is that an increase in κ *mandates* an increase in power piston stroke since the need to keep χ constant precludes adjusting the displacer piston stroke to manipulate κ. The larger the value of κ, then, the larger the power piston stroke necessarily, and the faster the piston must move back and forth in its cylinder for a given engine speed ω. This results from the fact that for each engine cycle at a given engine speed, the piston must traverse a greater linear distance with a longer stroke than with a shorter one. Velocity-dependent flow losses, therefore, become more significant at higher κ, perhaps, at least in part, accounting for the decrease in brake power with κ increasing past its power maximizing value.

Also an affect of larger κ and correspondingly longer power piston strokes is that more of the piston is pulled from its cylinder per cycle at the top of its stroke. The more the piston is drawn from its cylinder, the more the working fluid has a chance to escape through imperfect oil seals (because less of the seal is present), thereby reducing engine pressures and perhaps causing a heat pumping effect as well. Both of these consequences of long power piston strokes could explain the peaking brake power curve when a steadily increasing curve is theoretically expected. With longer strokes also comes increased side-loading on the piston because of the extreme angles (with respect to the horizontal) achieved by the connecting rod in the course of a cycle. These extreme angles mean that more of the forces acting between the piston and its connecting rod are directed perpendicularly to the piston's velocity vector. Increased

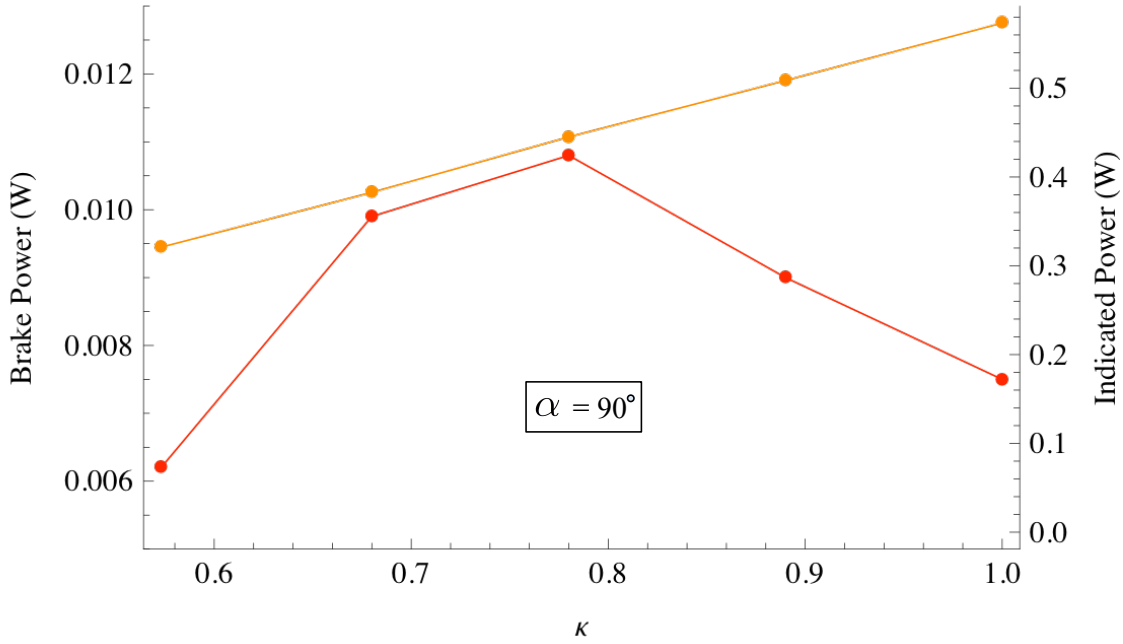

Figure 5.10: A plot of the relationship between swept volume ratio κ and output power at 500 RPM according to the Senft model (orange curve) and to experimental data (red curve). The model predicts an increasing power output as κ increases, but a very different trend is visible in the experimental data, indicating that the model does not in fact accurately predict the relationship between κ and power output.

side-loading kinks the piston in its cylinder, decreasing the general effectiveness of the oil-film lubricant (by displacing it), thus increasing friction between the piston and its cylinder. Increased side-loading would also serve to progressively chip away at output power with increasing κ. As the consequences of a long power piston stroke—increased piston velocity yielding increased flow losses, decreased effectiveness of oil seals, increased opportunity for heat pumping, increased side-loading—all fall under the umbrella of thermodynamic losses that the model, by definition and by design, cannot account for, the model's disagreement with measurements in the case of the relationship between power and κ is not sufficient to discount the model as inherently flawed.

To compile the conclusions from this section and the previous subsection into a more succinct statement, it appears, from comparison to measured data, that the Schmidt model, presented by Senft in [2], is generally a poor fit for real engines. Since it yields indicated power for a set of engine variables whereas the measurement technique employed yields brake power, the measured data cannot be expected to agree with the model, and vice-versa. That said, the model did seem capable of predicting very broad, qualitative trends at least at low engine speeds. All the plots in Figs. 5.1-5.8 exhibit increasing brake power with increasing RPM up to a power maximizing RPM. The increase in power up to that RPM is mirrored in all cases by the theoretical predictions. The model also successfully predicted the relationship between output power and phase angle α with RPM set to 500. (It would be an interesting extension to explore the model's ability to predict the relationship between power and phase angle α at higher, or lower, RPMs as well.) The model was also successful in predicting increasing power output with increasing swept volume ratio κ, but could not account for the observed decrease in brake power accompanying an increase in κ beyond its power maximizing value. Overall, though, and especially considering the marked difference between indicated power and brake power, the Schmidt model performed reasonably well, and could at least certainly prove *useful* in designing new Stirling engines.

Appendix A

Code

```
DataMuncher[n1_, n2_, a_, Tlow_, Thot_, Tamb_, RPMs_, output_] :=
 Module[{pstroke, dstroke, Vcc, Vdc, Vpis, Vnub, VDp, VSp, VSd, SVR, VDead,
   DVR, Vtot, Vstot, S, Tratio, A, B, theta, X, Y, PbarInst, Work, w, power},
  If[n1 == 1, dstroke = .00899];
  If[n1 == 2, dstroke = .01117];
  If[n1 == 3, dstroke = .01336];
  If[n1 == 4, dstroke = .01549];
  If[n2 == 1, pstroke = .00899];
  If[n2 == 2, pstroke = .01102];
  If[n2 == 3, pstroke = .01306];
  If[n2 == 4, pstroke = .01509];
  If[n2 == 5, pstroke = .01712];
  If[n2 == 6, pstroke = .01915];
  (* n1 is the hole number of the stroke setting on the displacer flywheel,
  and n2 is the hole number of the stroke setting on the power piston
   flywheel. This section is merely an extension of a labeling system used
   to organize the data files and is irrelevant to the physical model. *)

  (* Here begins the model itself. The first 14 definitions
   describe engine and experimental parameters; the remaining definitions
   use the result of the Senft model to compute work and power. *)
  Vcc = (Pi * (.00268 / 2) ^ 2) * .0448;
  (* Vcc ([m]^3) is the volume of the
   connecting chamber between the power and displacer cylinders. *)
  Vdc = (Pi * (.01728 / 2) ^ 2) * .02408 + (Pi * (.01685 / 2) ^ 2) * .03345;
  (* volume of the displacer chamber *)
  Vpis = (Pi * (.01589 / 2) ^ 2) * .03859;
  (* volume of the displacer piston *)
  Vnub = (Pi * (.00465 / 2) ^ 2) * (.00813 - (dstroke / 2));
  (* volume of piston rod stub *)
  VDp = (Pi * (.01264 / 2) ^ 2) * (.01239 - (pstroke / 2));]
(* power cylinder dead volume *)
```

```
VSp = (Pi * (.01262 / 2) ^ 2) * pstroke;
(* power piston swept volume *)
VSd = (Pi * (.01589 / 2) ^ 2) * dstroke + (Pi * (.00465 / 2) ^ 2) * dstroke;
(* displacer piston swept volume *)
SVR = VSp / VSd;
(* swept volume ratio κ, dimensionless *)
VDead = Vcc + VDp + (Vdc - VSd) - Vnub - Vpis;
(* dead space volume*)
DVR = VDead / VSd;
(* dead volume ratio χ, dimensionless *)
Vtot = VSp + VSd + VDead;
(* total internal engine volume *)
Vstot = VSp + VSd;
(* combined total piston swept volume *)
S = ((101 325 * Vtot) / (Tamb + 273.15));
(* Here, S=Nk from the ideal gas law PV=NkT, Vtot is in [m]^3,
temperature is in [K], so pressure is in pascals (Pa). Approximating
  lab air pressure as sea level air pressure gives p = 101325 Pa. *)
Tratio = (Tlow + 273.15) / (Thot + 273.15);
(* hot-to-cold temperature ratio τ, dimensionless *)
A = SVR - (1 - Tratio) * Cos[a];
B = (1 - Tratio) * Sin[a];
theta = ArcCos[A / Sqrt[A ^ 2 + B ^ 2]];
X = Sqrt[SVR ^ 2 - 2 * SVR * (1 - Tratio) * Cos[a] + (1 - Tratio) ^ 2];
Y = 1 + Tratio + SVR + (4 * Tratio * DVR) / (1 + Tratio);
(* A, B, theta, Y, X all dimensionless *)
PbarInst = (S * (Tlow + 273.15)) / ((VSd / 2) * Sqrt[Y ^ 2 - X ^ 2]);
(* PbarInst in pascals: [kg]/([m]*[s]^2) *)
Work = N[(Pi * (1 - Tratio) * Vstot * PbarInst * SVR * Sin[a]) /
    ((SVR + 1) * (Sqrt[Y ^ 2 - X ^ 2] + Y))];
(* Work given in Joules *)
w = RPMs / 60;
power = Work * w;
(* Power computed in J/s, or Watts W *)
If[output == "power", Return[power]];
If[output == "work", Return[Work]]]
```

Figure A.1: This string of code, continued from the previous page, takes as its inputs the numbers of the flywheel mounting holes used (n1 and n2), the phase angle α, the temperatures $T_{\text{low}} = T_C$ and $T_{\text{hot}} = T_H$ as well as the ambient room temperature T_{amb}, the engine speed in revolutions per minute, and the desired output of the program, either the power generated or cyclic work performed. These inputs are sufficient for determining κ, τ, and χ, which, with the addition of α and ω, are sufficient for calculating W_{cyc} and power following the Senft model. This code was used to generate plots of theoretical power output versus engine speed to compare with plots of *measured* output power versus engine speed.

```
DataSlayer[list_, offset_] :=
 Module[{rpm, rps, force, torque, TorqueCurve, A, B, a},
  rpm = Table[list[[i, 1]], {i, 1, Length[list]}];
  rps = Table[N[rpm[[i]] / 60], {i, 1, Length[list]}];
  (* Converts RPM into revolutions per second ω *)
  force = Table[(list[[i, 2]] - offset) * 0.00981, {i, 1, Length[list]}];
  (* Calculates force from the mass measurements,
   taking g = 9.81 m s^(-2). The variable 'offset'
     compensates for any issue in effectively zeroing the scale *)
  torque = Table[force[[i]] * .102, {i, 1, Length[list]}];
  (* Calculates torque from forcy assuming a lever arm R = 0.102m *)
  power = Table[torque[[i]] * rps[[i]], {i, 1, Length[list]}];
  TorqueCurve = Table[{rpm[[i]], torque[[i]]}, {i, 1, Length[list]}];
  PowerCurve = Table[{rpm[[i]], power[[i]]}, {i, 1, Length[list]}];
  a = ListLinePlot[TorqueCurve, PlotRange -> All,
   PlotStyle -> Blue,
   PlotMarkers -> {Automatic, Tiny},
   ImagePadding -> 50,
   Axes -> False,
   Frame -> {False, False, False, True},
   FrameTicks -> {None, None, None, All},
   FrameLabel -> {None, None, None, "Torque (Nm)"}];
  (* Generates a plot of torque versus RPM
    with the torque scale axis appearing on the left of the plot *)
  b = ListLinePlot[PowerCurve, PlotRange -> All,
   PlotStyle -> Red,
   PlotMarkers -> {Automatic, Tiny},
   ImagePadding -> 50,
   Frame -> {True, True, False, False},
   FrameTicks -> {All, All, None, None},
   FrameLabel -> {RPM, "Power (W)", None, None}];
  (* Generates a plot of power versus RPM with
    the power scale axis appearing on the right of the plot *)
  curves = Overlay[{a, b}];
  (* Combines the torque and power plot into one with multiple
    vertical scale axes. This is possible because the plots a and b,
   despite having different ranges, have the same domain *)
  Print[curves];
  Print["torque   ", torque];
  Print["power   ", power]]
```

Figure A.2: Here we find the code used to produce the torque and power curves presented in Chapter 5. The program DataSlayer takes as inputs a list of data of the form (RPMs, mass), the raw data from a trial run, and a value labeled as "offset." The purpose of this latter value is to scale the mass readings down (or up) if the scale couldn't be properly zeroed after positioning of the Prony brake before commencing the trial. Getting the scale to tare effectively was difficult; the offset variable solves this problem. The code assumes that $g = 9.81\,\mathrm{m\,s^{-2}}$, and uses a lever arm R of 0.102 m.

References

[1] G. Schmidt, Zeitschrift des Vereines Deutschen Ingenieus **15**, 1 (1871).

[2] J. R. Senft, Int. J. Energy Res. **26**, 1087 (2002).

[3] D. V. Schroeder, *An Introduction to Thermal Physics* (Addison-Wesley, San Francisco, 1999).

[4] G. Walker, Sci. Am. **229**, 80 (1973).

[5] A. Ross, *Stirling Cycle Engines* (Solar Engines, Phoenix, AZ, 1977).

[6] G. Aragón-González, M. Cano-Blanco, A. Canales-Palma, and A. León-Galicia, Rev. Mex. Fis. **59**, 199 (2013).

[7] J. R. Senft, *Ringbom Stirling Engines* (Oxford UP, Cambridge, MA, 1993).

[8] P. T. Gaynor, R. Y. Webb, and C. C. Lloyd, Low temperature differential stirling engine based power generation, Institute of Electrical and Electronics Engineers, 2008.

[9] C. M. Invernizzi, *Closed Power Cycles: Thermodynamic Fundamentals* (Springer-Verlag, New York, 2013).

[10] D. W. Kirkley, J. Mech. Eng. Sci. **4**, 204 (1962).

[11] P. T. Landsberg, *Thermodynamics: With Quantum Statistical Illustrations* (Interscience Publishers, New York, 1961).

[12] D. C. Giancoli, *Physics for Scientists & Engineers*, 3rd ed. (Prentice Hall, Upper Saddle River, NJ, 2000).

[13] T. Finkelstein, Optimization of phase angle and volume ratio for stirling engines, Society of Automotive Engineers, 1960.

[14] P. C. T. de Boer, Int. J. Energy Res. **33**, 813 (2009).

[15] K. S. Krane, D. Halliday, and R. Resnick, *Physics* (Wiley, New York, 2001).

[16] K. Kondepudi and I. Prigogine, *Modern Thermodynamics: from Heat Engines to Dissipative Structures* (Wiley, New York, 1998).

[17] T. Petrosky and I. Prigogine, Adv. Chem. Phys. **99** (1997).

[18] C. Salter, J. Chem. Educ. **77**, 1027 (2000).

[19] C. H. Cheng and H. S. Yang, Energy **36**, 5899 (2011).

[20] R. H. Romer, *Energy: an Introduction to Physics* (W. H. Freeman & Co., San Francisco, 1976).

[21] Quote by John Howard.

[22] H. Crew, *General Physics: An Elementary Text-book for Colleges* (Macmillan Company, New York, 1908).

[23] E. W. Constant, *The Origins of the Turbojet Revolution* (Johns Hopkins UP, Baltimore, MD, 1980).

[24] T. G. Bullen, Am. J. Phys. **25**, 382 (1957).

[25] L. B. Gordon, Loss terms in free-piston stirling engine models, NASA Lewis Research Center, 1992.

[26] R. K. Bumataria and N. K. Patel, Int. J. Res. Eng. Tech. **2**, 433 (2013).

[27] H. Karabulut, C. Cinar, E. Ozturk, and H. S. Yucesu, Renew. Energ. **35** (2010).