DEVELOPING MAXIMUM-EFFICIENCY ARCHITECTURES FOR STEADY-FLOW CHEMICAL ENGINES

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF MECHANICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Rebecca Zarin Pass
June 2015

This dissertation is online at: http://purl.stanford.edu/dp252vs9919

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Chris Edwards, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Reginald Mitchell, Co-Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Adam Brandt

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

Improving the efficiency of natural gas-based power plants is essential for significantly curtailing future greenhouse gas emissions. Historically, improvements to power plant thermodynamics have been implemented as incremental changes from existing systems. Only recently has work been done to systematically develop optimally-efficient combustion engines using existing technologies without artificial constraints on the system's size, components, or shape. This systematic process was used to optimize steady-flow combustion engines for work and heat transfer.

This dissertation builds off of this recent work by additionally optimizing systems for matter transfer. The best flow paths for air and fuel are determined, as well as the flow path for an additional environmental (non-exergetic) fluid: water. As an alternative to the highly-entropic combustion process, a solid-oxide fuel cell is also incorporated into the design scope. Both numerical and analytical approaches to optimizing these architectures are examined. Through a systematic expansion in the scope of candidate systems, this thesis develops optimal thermodynamic schematics that achieve efficiencies more than 15 percentage points better than those for the most efficient existing power plants.

Acknowledgments

In the worst doldrums of the PhD, I cheered myself up by fantasizing about writing these acknowledgments and reminding myself how lucky I am to have all of these people in my life. Now that it is actually time to put these thank yous in writing, it is surprisingly hard to do, but here goes:

Thank you to the Global Climate and Energy Project for funding this work. Thank you to Adam Brandt and Reginald Mitchell for your thoughtful guidance and time spent reading this thesis. To Hai Wang for your kind support and service on my thesis committee. To Jennifer Schwartz Poehlmann, Robyn Dunbar, and Sheri Sheppard for being excellent role models and coaches in teaching. I always left your company reinvigorated.

To Julie Kennedy, who has been my longest-standing mentor. You have been central to my entire Stanford experience and supported me through every risk, failure, and success along the way. The first day we met you promised we would have a life long friendship. I'm so excited to continue it.

To my exceptional PhD advisor, Chris Edwards, for making the PhD experience about my own learning. You opened my (skeptical) eyes to the exciting and powerful world of thermodynamics and I have never looked back. You took me in not despite my unusual background, but because of it. And you also set up an incredible community of people who care deeply about teaching, which led to a wonderful set of labmates.

I cannot imagine a better group of coworkers than I had throughout my PhD.

Thanks to the older generation of Paul Mobley, J.R. Heberle, and Sankaran Ramakrishnan. All of you were excellent mentors and friends. Sankaran in particular was crucial to this entire process and most of the proofs in this thesis were written on the floor of his apartment in Los Angeles. Sankaran, through discussions of culture, food, thermodynamics, and friendship, my happiest memories of producing this work occurred with you.

To my incredible peers in "The Wad": Julie Blumreiter, Adelaide Calbry-Muzyka, BJ Johnson, Greg Roberts, and Ben Kessel. You have all taught me so much about thermodynamics, work-life balance, and becoming a more self-assured and open-minded person. I love being on your team. To Adelaide in particular for being my academic life partner (ALP).[1] We were hired as a dynamic duo and will continue to act as one throughout life. This would never have happened without you, nor could I have imagined it any other way. I'm so glad you did your laundry that day.

To Mark Donohue, John Fyffe, Carol Regalbuto, and Natt Oliver for continuing the wonderful lab culture and adding a breath of fresh air. A serious benefit of remaining in the area after my PhD will be the ability to continue to see you.

To all of the other friends of exergy, including Lena Perkins, Ben Jensen, Sebastian Wigström, Chloe Mai, Jacob Alvarez, and Chris Cameron.

To my WISE group members for your support and friendship. And thank you to Stanford VPGE for funding such wonderful programs.

Outside of lab, I had a wonderful network of friends and family who supported me in this process. In particular, thanks to Aparna Bhaduri, my childhood soul-mate who I met in adulthood. To Sayuri Yapa for being my indefatigable cheerleader through math quals and life. To Rachel Danford who is the best second I could ever ask for in a duel. To Jessamyn Edra for keeping me grounded through hikes and literature. To Sacha Verweij for answering my personal ad and becoming the most delightful math partner and friend. To Brook Barajas, Keyur Shah, and Kate Parkinson for wonderful times away from work.

[1]Thanks to Jackie Crespo for her foresight in coining this incredible term.

Finally, I'd like to thank my entire wonderful extended family. To the entire Milrod/Guss clan for endless warmth and hospitality. To Gerald, Judy, and Erica who cheered and commiserated from afar. To Jessica who has been my west coast family since I arrived at 17 and continues to delight with surprise visits to campus at just the right times. To Jeff and mom for whom any finite list of reasons would do injustice. To Hannah and Alex for unwavering love and support. To Margot for being my first and best friend through life. You have always been my hero. And to Jesse, for taking a risk and sticking around for this wild ride. I'm so thankful you did and can't wait for the next stage.

I love you all.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Currently the global electric power sector accounts for 40% of carbon dioxide emissions. This is largely through the combustion of fossil fuels, which the International Energy Agency predicts will remain the dominant supplier of electricity for at least another 25 years [1]. In that time period, world electricity demand is projected to increase by 80%, with renewable energy resources accounting for only half of the growth [1]. Thus, the use of fossil fuels will in fact increase in the next few decades. A significant percentage of this growth will come from the burgeoning supply of natural gas world-wide [1].

In order to significantly curtail emissions going forward, the efficiency of natural-gas-based power plants needs to significantly increase. Power plants are built to last several decades. To improve the efficiency of infrastructure thirty years from now, more efficient systems have to be built today, which means they need to use existing technologies.

A vast amount of research exists offering partial solutions to this issue through improved plant design, construction, and operation. From the thermodynamic perspective, most of the proposed improvements have been designed as incremental changes from existing systems. Others are more systematic optimizations, but within a confined subset of possible systems.

Only recently has work been done developing the thermodynamic theory behind optimally efficient combustion engines without artificial constraints on the system's size, components, or shape. Teh first did this for piston-cylinder engines [2]. Ramakrishnan translated and evolved Teh's results to the realm of steady-flow engines, more typical of power plants [3]. Their findings are reviewed in more detail in Chs. 2 and 4.

This thesis expands on Ramakrishnan's work to determine the maximum possible energy efficiency achievable from natural-gas-based, steady-flow combustion engines using existing devices. Through a systematic expansion in the scope of candidate systems, optimal thermodynamic schematics are proposed that achieve efficiencies more than 15 percentage points better than those for the most efficient existing power plants.

This research does not optimize for other factors that affect infrastructure projects, such as economics, dynamic operation, or ease of maintenance. The thermodynamic optimization is complex and not currently well-understood. First determining the space of highly-efficient solutions is a crucial step after which these other constraints can be applied to narrow down the range of architectures to those that are both efficient and economically viable.

1.1 Measuring Efficiency

Energy is conserved, but it can be transformed between different modes of storage and transfer. An engine is a system that performs a series of energy transfers and transformations that take a readily available energy resource (e.g., chemical bond energy in fossil fuels) and converts it to a more fungible mode, work (or its equivalent, electricity). Thus, even though the term *engine* is often associated with mobile transportation, it is also appropriate for power plants. To be useful, the starting resource has to differ from the environment in pressure, temperature, or composition. A device can then be inserted between the resource and environment that extracts

work while the relevant equilibration occurs.

For any given resource and environment, there is a quantifiable value for the maximum possible work that could be extracted. This value is called the *exergy* (or sometimes *availability*) of the resource. The most relevant efficiency metric for the purposes of comparing systems that produce work is the *exergy efficiency*, defined for an engine as the net work produced divided by the total exergy invested. This is the efficiency that will be used throughout this dissertation.

It is worth noting that it is common in the power industry to instead use a first-law, or heating-value-based, efficiency for an engine. This assumes that the value of the fuel is based solely upon the energy released during the combustion of that fuel. However, combustion is only one possible reaction mechanism and as discussed above, fuels can have mechanical or thermal exergy as well as chemical. For this reason, heating-value-based efficiencies are not as translatable between different conversion processes. That said, for hydrocarbons fuels, heating-value and exergy-based efficiencies usually yield similar values. (E.g., for methane, the main component in natural gas, the lower heating value is 97% of the exergy value).

1.2 Framing the Optimization Problem

Developing maximally efficient engines requires determining the optimal set of devices and their transfer types. This process is referred to as *architecture* optimization. In contrast is *parametric* optimization, which presumes an existing, fixed architecture and varies the device properties to alter the thermodynamic states of the fluid. Deciding to include a solid-oxide fuel cell in an engine is an architectural decision. Deciding that the fuel cell should operate at 1000℃ is a parametric decision. In this thesis, the primary goal is to determine optimal architectures. However, the corresponding parameters can affect the choice of architecture and therefore also have to be determined.

Optimization problems have an objective function that needs to be minimized

or maximized, a set of variables that affect the objective function, and a set of constraints on those variables. Here, the objective function is the exergy efficiency, which is to be maximized. The variables are the device types, parameters, and their sequencing. The problem is constrained to choose from a set of pre-specified devices, to operate those devices in a plausible manner, and to obey the laws of physics. The set of allowable devices presupposes that the fuel is chemical. The use of exergy efficiencies requires specifying the environment, which is set to be 100% humidity air at 25℃ and 1 bar.

1.3 Structuring the Problem

There are only four possible ways to transfer energy: with matter, as heat, as work, and with radiation. Non-thermal radiation (e.g., lasers) is set aside as irrelevant to practical steady-flow engines, leaving three remaining transfers. These can occur internally (within the system) or externally (between the system and the outside environment). Heat, work, and every fluid (matter) each allow two possible degrees of freedom: one associated with internal transfers and one associated with external transfers. A depiction of a chemical engine with examples of each possible transfer is shown in Figure 1.1.

The solid rectangle indicates the extent of the engine and the associated thermodynamic system boundary. Outside of this boundary is the environment, as specified above. Any transfers (denoted by black arrows) that occur entirely within the engine are internal; any that cross the engine boundary are external.

There are three external *inflows* to the engine: fuel, air, and, optionally, water. The fuel is chemical and serves as the only exergetic input to the engine. As such, there are no external transfers of work or heat into the engine. Once the fuel, air, and water are in the engine, they proceed through a series of energy transformations and transfers via the devices, denoted by the dashed rectangles.

There are also three external *outflows* of energy from the engine: matter, work,

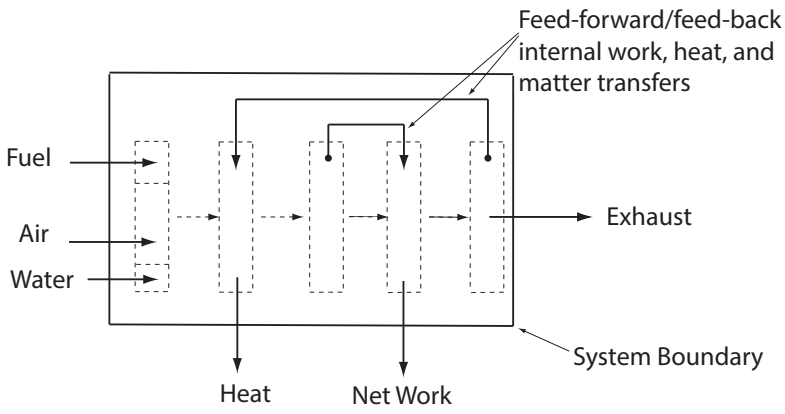

Figure 1.1: An example of a matter-, heat-, and work-regenerative engine.

and, optionally, heat. The engine is operating in steady-state and must exhaust matter. Without loss of generality, this is done as a single mixed stream.[1] Because the system is an engine, by definition, work must also be transferred out. Optionally, thermal energy may be exhausted to the environment as heat transfer.

In addition to these external transfers in and out of the engine, there are also optional internal transfers of work, heat, and matter between devices within the engine. Relative to the primary working fluid stream, these transfers can be categorized as either feed-forward or feed-back. Feed-forward transfers move energy in the same direction as the predominant working fluid stream while feed-back transfers do the opposite. A feed-back work transfer occurs, for example, when the rotational energy produced by a turbine is coupled back to a compressor via a shaft. An example of a feed-forward matter transfer is when some compressed air bypasses the combustor to supply cooling fluid to a downstream turbine.

Ramakrishnan developed a systematic approach to optimize steady-flow combustion engines starting from a bare-minimum architecture and slowly adding complexity

[1]Entropy is produced when the exhaust chemically mixes into the environment. The entropy is accounted for regardless of whether it occurs physically internal or external to the engine. This will be discussed further in Ch. 4.

through additional allowed transfers (degrees of freedom) [3]. He first optimized an engine fed by only air and fuel, and using only compressors, burners, and turbines. This allowed for internal transfers of work (via rotating shafts), but no heat transfer. Internal heat exchangers were then added and the system was re-optimized. Finally, he additionally allowed for external heat exchange from the engine to the environment. This resulted in an optimized air-fuel, steady-flow combustion engine with both internal (regenerative) and external work-and-heat transfers. His optimal architecture will be further discussed in Ch. 6.

This thesis will build off of Ramakrishnan's results to additionally optimize for matter transfers. The best flow paths for air and fuel will be determined, as well as the flow path for an additional environmental (non-exergetic) fluid: water. At that point, any further improvement in efficiency will require a new device. As an alternative to the highly-entropic combustion process, a solid-oxide fuel cell is the next device added.

Additionally, alternative optimization methods will be explored. Once the engine optimization problem has been structured, it could be solved either analytically (as Ramakrishnan did) or computationally. This thesis will explore both paths in an effort to determine optimally-efficient engines.

1.4 Organization of the Dissertation

This dissertation involves several separate approaches to expanding the scope of steady-flow combustion engine optimization. Staring with Ch. 3, each chapter presents a new methodology for optimization or a new degree of freedom to add to the optimization. The relationship between the expansions in optimization scope (and the associated chapters) is shown in Fig. 1.2.

The relevant literature associated with optimizing engines is reviewed in Ch. 2 of

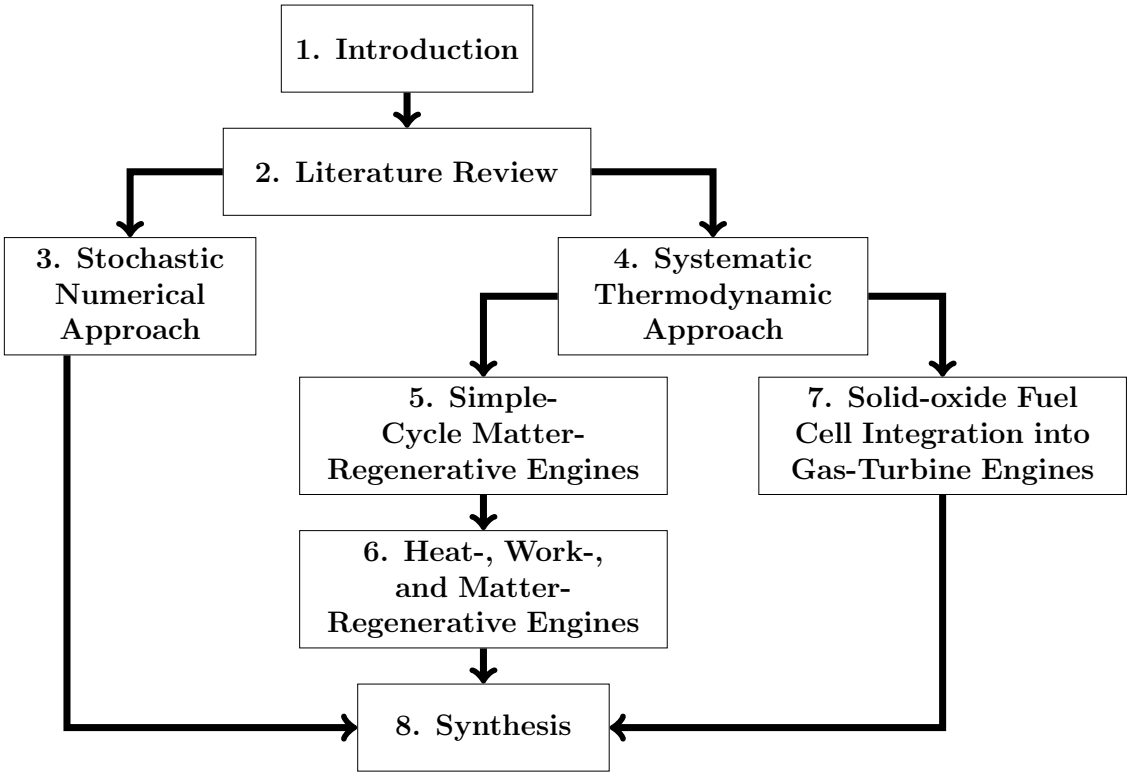

Figure 1.2: Relationship of dissertation chapters.

this dissertation. In Ch. 3, a stochastic numerical approach to designing maximally-efficient systems is presented. An alternative approach based on a systematic thermodynamic methodology is then discussed in Ch. 4. This approach will be used in Chs. 5 and 6 to incorporate new modes of matter transfer into the optimal architecture determined by Ramakrishnan. At this point, the steady-flow combustion engine will have been optimized for work-, heat-, and matter-regeneration. The next relevant expansion in the scope of possible engine architectures is to include an alternative chemical reaction device to the combustor. A solid-oxide fuel cell is incorporated into the optimal architecture in Ch. 7. Finally, in Ch. 8 lessons learned about architecture optimization are synthesized, and promising paths for further improvement in steady-flow engine efficiency are identified.

Chapter 2

Literature Review

Previous research has been conducted to develop optimally-efficient steady-flow engines. This dissertation expands on that work by adding new degrees of freedom in the form of devices (fuel cells), modes of energy transfer (matter regeneration), and fluids (water). Additionally, an alternative numerical methodology is also explored.

This literature review covers both methodologies for optimization and optimal design of steady-flow combustion engine architectures. The thermodynamic and numerical methodologies will be reviewed first, followed by the engine designs. The engine design review will be centered on innovative approaches to moving matter through power plants, followed by integration of gas-turbine engines with solid-oxide fuel cells.

2.1 Methodology 1: Thermodynamic Optimization

Structural optimization problems (of which engine optimization is a type) can be characterized as either parametric or architectural. Parametric optimizations search for the best operating points of a system, such as temperatures, pressures, and flow rates. In contrast, architectural optimizations search for the best set and ordering

of devices to perform the requisite energy transfers and transformations. (Usually this requires also optimizing the parameters for each architecture.) The parametric optimization space has been well-examined [4–6][1] and this work is more focused on the architectural optimizations.

Architectural optimizations of combustion engines have been attempted using heat-engine models. For a system that receives heat from a hot thermal reservoir and exhausts to a cold reservoir, the Carnot limit dictates the maximum possible efficiency. This is applicable when examining an *external* combustion process, in which the engine receives heat from an outside fluid. This assumes that there is no control over the architecture associated with the combustion process, and only the heat engine can be optimized. However, the *internal* combustion process involves transforming chemical bond energy into sensible (thermal) energy. This is fundamentally different than the process of receiving heat from an external reservoir. As such, the Carnot limit does not apply.

Potentially more appropriate to chemical engines, there have been attempts to define a physically significant thermodynamic length metric that is related to the efficiency of a process [7–11]. Salamon and Berry found that the Weinhold length[2] could be interpreted as irreversibility and therefore be minimized [12]. However, this only holds for unrestrained transformations with small deviations from equilibrium. It does not apply to combustion processes or any process involving energy transfers [3].

Optimal control is another approach that allows the determining of both architecture and process parameters. Teh first applied this method to maximize work

[1]Broadly, these studies find that maximizing turbine inlet temperature optimizes efficiency for work-regenerative engines. When internal heat transfer is introduced, the peak pressure of the system drops in order to maintain thermal energy in the post-turbine flue gas. When external heat transfer is added (i.e., intercoolers) the peak pressure trends up again because of the reduced penalty of compression work. Adding excess combustion moderator as air, water, or product gas also leads to an increase in system pressure.

[2]Weinhold used a thermodynamic state space defined as $U = U(S, V, N_1, ..., N_n)$. The Weinhold length is defined as the square root of the sum of all second partial derivatives of U with respect to the independent variables.

out of a simple piston-cylinder engine [13]. However, this method is not viable in the case of multiple controls, such as in work-and-heat regenerative architectures, because singular control arcs result, which cannot be solved for a unique optimal solution. [3].

In his work, Teh came up with a criterion for the minimization of entropy generation due to the combustion process itself. Because the combustion process conventionally occurs in an adiabatic fashion, the difference in inlet and outlet entropies is entirely due to entropy generation. By examining the shape of the thermodynamic property surfaces for common reactants and products of internal combustion engines, Teh discovered that these inlet and outlet entropies approach each other (i.e., are less different from each other, and result in less entropy generation) when the reactant system is at a highly energetic state [14]. In his work with piston-cylinder engines, the relevant property was internal energy. In flowing systems, this translates to enthalpy [15]. Teh expressed this idea of minimizing combustion irreversibility by maximizing reactant energy as the *extreme-states principle*. This principle will be used extensively in this thesis and referred to in short as "extreme states."

Building off of the extreme-states principle, Ramakrishnan developed a methodology for optimizing the efficiency of steady-flow combustion engines [3]. This methodology begins with a bare-minimum engine architecture and systematically adds a single degree of freedom at a time after which the entire system is re-optimized. Ramakrishnan's methodology will be used for the majority of this thesis, and as such, is presented in significant detail in Chapter 4.

2.2 Methodology 2: Numerical Optimization

Numerical methods are used to find an optimal point given an objective function, a set of variables and their domains, and a set of constraints. It is possible that the steady-flow combustion engine optimization problem could be framed in a manner that is solvable by a numerical method. Many methods exist for numerical

optimization, but only some are relevant to the problem of power plant efficiency optimization. The most commonly used and promising algorithms for this type of problem are reviewed below.

2.2.1 Mixed-Integer Nonlinear Programs (MINLP)

In many ways mixed-integer nonlinear programs are more a class of problems than they are a method [16]. They refer to systems that have both continuous and discrete variables and require nonlinear constraints or a nonlinear objective. This is typical of power plant optimizations. Device types are categorical and are distinguished using discrete variables while their parameters are continuous. If mass flow rates or temperatures are variables, then the conservation equations are nonlinear.

In the context of energy optimization, a mixed-integer nonlinear program typically refers to an algorithm by which the search space is in some way partitioned into smaller problems that are either convex or in another way easier to solve [16]. The results of these partitioned problems are coalesced, pruned, and optimized for a global result. A common outer algorithm to perform the partitioning and pruning is called branch-and-cut [17].

In structural optimization, MINLP algorithms have been used to choose between a finite number of potential architecture configurations. This is typically done using a set of binary variables that determine whether a component of the architecture is *on* or *off*. Because the binary variables have to be known ahead of time, the set of possible architecture decisions is also known. The set of architectures resulting from these decisions is often called the *superstructure* [18–21].

Researchers have used MINLPs to do superstructure optimizations of energy systems. In chemical engineering, MINLPs have been used for optimization of separation systems. These algorithms are used to make decisions such as how many trays to use in distillation columns, where the feed streams are located, and where heat exchange should occur, in order to minimize cost or maximize efficiency [19, 20, 22]. MINLP has also been used to determine optimal partitioning of energy systems between

multiple goals, such as joint electricity and methanol production [23], electricity and fresh water production [18], and electricity and heat production [24–26]. Additional optimizations have been done to minimize cost or fuel usage with variable loads as a function of time [27, 28].

The main weakness of MINLP is the limitation on the size of the problem. Highly nonlinear, constrained problems require time-consuming calculations, which severely limit the possible number of variables. Importantly, the algorithm cannot adapt from the pre-defined superstructure, restricting the number of possible architecture decisions significantly [29, 30]. Finally, a global optimum is not guaranteed, depending on the shape of the search space [30].

In order to avoid the superstructure problem, algorithms have been developed that can adapt by taking candidate solutions and altering them to produce new candidates. Holland first laid out the criteria for an adaptive system as: 1) an ability to produce representative samples (candidates) from the environment (search space), and 2) an ability to develop higher quality new candidates [31]. Such algorithms tend to be stochastic, meaning that some element of their decision making is done probabilistically, not deterministically. The next several algorithms fall into this category.

2.2.2 Particle Swarm Optimization (PSO)

Particle swarm optimization (PSO) is used to mimic animal swarming behavior. Individual agents (e.g., ants) are initially placed on the objective function surface randomly. As the agents collect data on the best objective value points they have found so far, their search directions are updated. Each individual's search direction is determined by a weighted combination of the individual's best found location (the cognitive parameter), the swarm overall's best found location (the social parameter), and the swarm's prior location (the memory parameter). [32].

Garcia-Nieto, et.al., found PSO to be the fastest algorithm for determining optimal timing of red, yellow, and green lights for a collection of intersections in two

cities, with the goal of maximizing vehicle throughput [33]. Onwunalu, et.al., used PSO to optimize well placement for oil recovery to maximize net present value [34]. Del Valle, et. al., reviewed the use of PSO in electrical power systems, such as for reactive power and voltage control. They note the potential for using this method for topology optimization, but give no examples of current use [35].

In order to use this method, variables must lie in a search space with a meaningful sense of direction. This is problematic for categorical variables, like device type. Consider optimizing the third device in a power plant: The best efficiency found by the swarm might have the third device as a compressor, but the best efficiency found by an individual search agent might have the third device as a burner. There is no meaningful weighted average of these two choices for further investigation.

2.2.3 Simulated Annealing

Based on a sampling algorithm by Metropolis [36], simulated annealing begins with a number of randomly chosen starting states on the objective surface. With specified probabilities, each state is then moved to an immediately neighboring state on the surface with either a better or worse objective value (i.e. going up or down slope). In time, the probability of choosing worse objective values decreases. In this way, the algorithm allows for the possibility of escaping locally (but not globally) optimal states to eventually hone in on the global solution. It is named for an analogy to annealing in metallurgy in which the rate of cooling is kept slow to allow for ultimately reaching the thermodynamic equilibrium point. In simulated annealing the rate of changing the probability of accepting a worse solution is therefore called the *temperature*. Kirkpatrick, et.al., used this method to solve the traveling salesman problem and to determine the optimal layout of circuits between two chips with a minimum number of inter-chip connections [37].

Simulated annealing could potentially be used for engine optimization. However, its reliance on a single algorithm parameter (temperature) makes it unlikely to successfully explore a rough, non-continuous surface.

2.2.4 Evolutionary Algorithms (EAs)

There are many types of evolutionary algorithms (EAs), but they have certain characteristics in common. All evolutionary algorithms include a population of individual candidate solutions to the problem. The candidates are evaluated for their objective value, called *fitness*, and then new candidates, called *children*, are created from attributes of promising existing candidates, called *parents*, in a process called *recombination*. *Mutations* (small random changes to individuals) are used to introduce diversity into the system. With both recombination and mutation, evolutionary algorithms are designed to mimic biological evolution, in which the individual candidates here are representative of an individual's genotype.

Genetic algorithms (GAs) are arguably the most common type of evolutionary algorithm. In genetic algorithms, the individual is represented by a binary vector. Mutations are bit flips and recombination is done by taking segments of multiple parent vectors to create new child vectors [31].

At the same time that genetic algorithms were developing in the United States, evolutionary strategies (ES) were being developed in Germany, largely by Schwefel and Rechenberg [38, 39]. In evolutionary strategies, the individuals are real-valued vectors. Additionally, individual candidates include *strategy variables* that dictate how the algorithm works (e.g., mutation rate). These strategy variables are thus evolved along with the individuals. Evolutionary programs (EPs) are similar to evolutionary strategies, but have a continuous arrival of new individuals rather than discrete generations. They were first used by Fogel [40]. The final common category is genetic programs (GPs), which were originally developed by Koza [41]. In genetic programs, the solutions are computer programs whose fitness is the ability to solve a given problem.

Evolutionary algorithms have been used for a wide variety of parametric optimization problems in infrastructure management. Toffolo used a genetic algorithm to optimize parameters for the S-Graz cycle (an oxyfuel power plant) [42]. Reehuis optimized pipe sizes for a water distribution network [43]. Larsen minimized waste

heat on a marine diesel [44]. Nieman minimized pollutants for a heavy-duty engine [45]. And Behbahani developed a controller for the removal of fish heads [46].

Evolutionary algorithms have also been used for limited architectural decisions in the vein of the superstructure optimizations by MINLPs. People use evolutionary algorithms to minimize weight or size of structures for a given load by turning potential parts of the foundation on/off [47, 48]. In power systems, this is sometimes done to make decisions about heat exchanger configurations [29, 49–51]. Other applications include electronic chip layout [52], distillation column configuration [21, 53], and gas turbine steam injection [54]. Like the MINLP problems, these tend to be limited in scope with typically fewer then ten variables.

Open-ended architectural optimizations have been done with evolutionary algorithms without a constraining superstructure. This has occurred in the realm of electronic circuit design, in which the method is sometimes called *evolvable hardware* [55]. McConaghy used small expert-designed *building blocks* to then synthesize larger circuits for filters and antennas [56]. Koza applied similar techniques to both circuits and metabolic pathways [41]. Grimbleby used resistors, capacitors, and inductors as building blocks to design passive filters with desired frequency response characteristics. Two novel designs were found by the algorithm that met all criteria [57]. Miller successfully used an evolutionary algorithm to make a 3-bit multiplier using 20% fewer gates than the best known human design [55]. Fitness was defined as the number of correct entries in a truth table.[3]

Convergence Theory

Unlike many of the other algorithms considered above, evolutionary algorithms can be designed to converge to the globally optimal point with probability 1. To do so, mutation rates need to be strictly bounded above zero such that every possible individual in the search space is reachable. Additionally, the best individual cannot

[3] A truth table lists all possible circuit inputs and their corresponding desired outputs as binary values.

be lost in future generations. *Elitism* must be included, which insures that the best individual is guaranteed to carry over to subsequent generations. These conditions are sufficient for proving an evolutionary algorithm's convergence properties [58–62]. Unfortunately, there is no guarantee of reaching the global optimum in finite time [60].

There are various proofs that provide bounds on the speed of convergence. These are typically done by representing the algorithm as a Markov chain with a probability transition matrix from one generation to the next. Each element (i,j) in the probability transition matrix is the probability that the ith possible individual will be included in the jth possible population [63]. Using ergodic theory, convergence can be bounded in terms of the eigenvalues of the transition matrix and the mutation probability rate [59]. Thus, for any specified set of mutation, crossover, and selection parameters, an upper bound on the convergence time can be calculated. This is useful for estimating the time-to-solve viability for a proposed algorithm even before running sample trials.

This global convergence property, along with the ability to operate on non-continuous surfaces with categorical variables led to the decision to try using evolutionary algorithms for engine optimization in Ch. 3. The next sections of this literature review will discuss engine architectures pertinent to the desired optimization scope.

2.3 Engine Designs: Matter-Regenerative Cycles

The simplest gas-turbine engines have air and fuel compressed in parallel, mixed and combusted, and then expanded as a joint product stream. There are however many other possible paths that the air and fuel could take. Combustion engines can also use additional fluids. Water can be injected into the engine to create a *wet cycle*. Alternatively, water or another fluid can circulate in a closed loop, which only interacts thermally with the main engine. Such engines are referred to as *combined*

cycles. A brief summary of previous studies on complex matter flows in power systems follows.

2.3.1 Feed-Forward Matter Transfer

Feed-forward refers to the process of splitting some fluid off from the main stream, bypassing some number of devices, and then mixing the fluid back in. An example of this occurs when post-compressor air bypasses the combustor to provide blade cooling in a gas turbine. Additionally, fuel can be fed-forward relative to the air in the case of a reheat system where combustion happens in stages [64, 65]. While this process is sometimes done to reduce NO_x emissions, it also allows for more of the combustion to happen at an extreme state and increases the work extracted from the turbines.

2.3.2 Feed-Back Matter

Feed-back matter transfer refers to secondary matter flows in a direction reverse to that of the primary air-fuel or product stream. Feed-back along with feed-forward has been proposed in the context of blade cooling by Chiesa and Macchi [66]. The feed-forward, compressor-bleed air that is heated during turbine blade cooling is fed back into the combustor. It is shown that such a feed-forward/feed-back air cooling scheme provides an increase in efficiency over open-loop air cooling. Feed-back transfer of exhaust gas (i.e., exhaust gas recirculation) has been studied in combined cycles for post-combustion CO_2 separation by Bolland and Mathieu [67], Botero et. al [68], Sipöcz and Tobiesen [69], and others. All these studies indicate negligible change in the gas-turbine efficiency with exhaust-gas recirculation, but a slight improvement of combined-cycle efficiency with carbon separation (due to marginally-improved heat-integration of the absorption system re-boiler with the steam cycle). Solid-oxide fuel cell systems use matter feed-back to recycle some of the anode exhaust stream back to the anode inlet [70–72]. In this case, the recirculation of matter

provides thermal energy and necessary reactants to initiate methane reformation and the resulting electrolytic hydrogen oxidation. This results in an improvement of the dynamic response time to varying loads [25], as well as a reduction in fuel cell irreversibilities [73], ultimately resulting in higher efficiency. There are also cases of using cathode-side recirculation in an SOFC to improve thermal management [74,75].

2.3.3 Wet Cycles

Combustion engines that use injected water are typically referred to as *wet* cycles. These cycles differ in their quantity of water, location of injection, and amount of water pre-heating. When the water is injected as a liquid, the system is generally referred to as a *humid air turbine* (HAT) cycle. The basic HAT cycle uses a humidification tower to mix liquid water and air after compression. Some HAT cycles (called TOPHAT) inject water directly into compressors, which cools the air through the latent heat of evaporation and therefore decreases compression work [76]. Both of these cycles can go through all of the same possible intercooling and recuperative heating devices as dry cycles. Finally, some cycles (called steam-injected gas-turbines, STIG) heat the water to a vapor and inject it directly into the combustor.

Researchers consistently find that the cycle is more efficient the earlier the water is injected [6, 77–79]. Kavanagh numerically optimized the parameters for various wet cycles and for a 1573 K peak temperature, found the maximum efficiencies for TOPHAT, HAT, and STIG, to be 53%, 52%, and 49%, respectively [77]. This efficiency benefit of early water injection is attributed to smaller temperature differences of water and air, improved compressor performance, and the lack of a poorly temperature-matched boiler present with STIG [6].

All of the wet cycles result in significant power gains over their respective dry cycles. Kavanagh found the dry, simple-cycle, gas-turbine engine to have an air-specific work of approximately 360 kJ/kg, compared to 564 kJ/kg for STIG, 617 kJ/kg for HAT, and 706 kJ/kg for TOPHAT [77].

2.3.4 Combined Cycles

Combined cycles use the hot flue gas from a gas turbine to power a bottoming, closed-loop, heat engine. Conventionally the heat engine's working fluid is water or an organic compound that undergoes a liquid-to-vapor phase change, however it can also be a fluid that remains a gas throughout the cycle [4, 80, 81]. The most efficient existing power plants are natural-gas combined cycles. As of 2015, Siemens, GE, and Mitsubishi Heavy Industries all have combined cycles with greater than a 60% LHV efficiency [82–84]. While combined cycles do not get the cooled compression benefit that wet cycles do, phase-change working fluids allow for pumping a liquid to a very high pressure with minimal work and then expanding a gas to a condensor near ambient temperature but with much lower than ambient pressure. Compared to gas cycles, steam cycles allow for another order of magnitude in pressure decrease (down to approximately 7 kPa), which adds significantly to the overall work output.

Combined cycles consistently have higher operating efficiencies than any of the wet cycles. However, the heat recovery steam generator (HRSG) makes combined cycle systems physically large. As such, many see the main role of wet cycles as improving the efficiency of smaller systems (< 100 MW) at lower marginal capital cost [6, 85]. Despite current efficiency trends, there is no known proof of the highest possible efficiency of either wet or combined cycles. This will be addressed in Chs. 6 and 8 of this thesis.

2.4 Engine Designs: Solid-Oxide Fuel Cells

Fuel cells keep the air and fuel separated by a semi-permeable membrane while performing oxidation. Electricity is directly produced in the process. Solid-oxide fuel cell (SOFC) systems, such as the Bloom Box, are cited at 50% LHV-based efficiency, making them highly efficient single-cycle systems with strong potential for future development [71]. This dissertation examines the potential efficiency gains from integrating solid-oxide fuel cells into steady-flow engines. This is not a new

idea, and researchers have modeled efficiencies between 50 and 70% for such systems [86–90]. Others have looked at more complex integrations of solid-oxide fuel cells with humidified gas-turbine cycles [91, 92], bottoming steam cycles [93], adsorption refrigeration [92], and energy storage systems [94]. As will be discussed further in Ch. 7, Mitsubishi Heavy Industries has proposed an industrial-scale solid-oxide fuel cell, gas turbine system [70].

What has not yet been done is a rigorous thermodynamic analysis of the best possible integration strategy of SOFCs and gas-turbines engines. Zhang et.al. and Zhao et.al. derived analytical expressions, but treated the gas turbine as a heat engine with a maximum efficiency dictated by Carnot [95, 96]. This is fundamentally incorrect since it misidentifies combustion irreversibility and removes any discussion of affecting that irreversibility through integration of SOFCs and gas turbines.

2.4.1 Methane Reformation

Solid-oxide fuel cells oxidize hydrogen to produce electricity. However, hydrogen is not a naturally available energy resource, so it is either formed through electrolysis from water or reformed from methane in natural gas. The breakdown of methane is traditionally done by steam methane reformation (SMR), which can be described by the global reaction

$$CH_4 + H_2O \rightarrow 3H_2 + CO \tag{2.1}$$

Steam methane reformation is slow relative to other chemical reactions in the fuel cell, and is traditionally modeled with an Arrhenius-form kinetic expression. As summarized by Aguiar, et al., potential kinetic expressions vary widely in both their rate and pressure-dependence [97]. Ultimately, the kinetic expression given by Achenbach [98] is used here

$$k = k_0 P_{CH_4} e^{\frac{-E_a}{RT}} \tag{2.2}$$

where the activation energy E_a is 82 kJ/mol and the pre-factor k_0 is 4274 $mol/(s - m^2 - bar)$. This expression is used commonly by several other researchers [87, 99].

The origin of the steam for the steam methane reaction also varies. It is sometimes provided by pumping and boiling an external water stream [87,100,101], by recycling water-rich anode products to the entrance of the fuel cell [97, 99, 102, 103], or a combination of the two [25,104]. The need to boil a significant quantity of water tends to decrease the efficiency of the systems using external water. However, the concern with not using external water is that insufficient steam during methane reformation can lead to carbon deposition on the anode. Experimental and numerical studies differ widely in their reported minimum quantity of steam necessary to avoid this coking effect–from a water-methane molar ratio of 1 [105] to 5 [106]. In addition to varying water quantity, several researchers have successfully introduced catalysts or other dopants to the anode to reduce coking [101, 103, 107].

2.5 Conclusion

This chapter surveyed both optimization methodologies and optimized engine designs. The wide range of both methodologies and designs indicates a lack of consensus on the best strategy to maximize steady-flow engine efficiency. The first step in clarifying strategy is taken in the next chapter on numerical engine optimization.

Chapter 3

Stochastic, Numerical Optimization Methodology

Most of this thesis focuses on an analytic, thermodynamic approach to optimization. While this approach works, it requires specialized knowledge and can only examine one degree of freedom at a time. In this chapter an alternative approach is explored using a stochastic numerical methodology. This approach has the potential to optimize for many degrees of freedom simultaneously in a very large search space, while using less (albeit, still significant) expert judgment. If this algorithm is successful in finding efficient engines, it may reveal expert biases or blind spots and new thermodynamic strategies. This would in turn improve how experts think about the thermodynamics of such engines.

3.1 Problem Characterization

Table 3.1 summarizes the optimization problem treated here. The variables in optimizing engine efficiency are the device types, operating parameters, and sequencing. Device types are of a finite set that can be labeled by integers. Such non-binary, categorical variables mean that there is not a continuous, smoothly varying design

max:	exergy efficiency
varying:	device types, parameters, sequencing
given:	a set of possible devices and their performance metrics air, water, and chemical fuel at 298 K, 1 bar
such that:	the laws of thermodynamics are obeyed

Table 3.1: Engine optimization problem.

surface that can be optimized using a derivative-based method. It also dampens the possibility of vector-based methods like particle swarm optimization (PSO) that require a meaningful sense of direction in thermodynamic state space.

Unlike device types, device parameters are continuous within a specified domain. For example, a pressure ratio could vary between 1 and a maximum, such as 200:1. Similarly, the turbine inlet temperature could vary between ambient and a maximum, such as 1800 K. The combination of integer and continuous variables is sometimes handled using a mixed-integer, nonlinear program (MINLP) algorithm.[1] Unfortunately, this requires a finite, pre-defined list of possible connections. This is not something that can be done for the current purposes, which makes MINLP algorithms inapplicable.

The algorithm also has to be capable of finding global optima, as opposed to merely local optima.[2] This combination of mixed integer-continuous, non-smooth, and global convergence requirements led to the decision to use an evolutionary algorithm. As discussed in Chap. 2, in an evolutionary algorithm, candidate architectures (individuals) are assessed for their fitness (exergy efficiency). Mimicking biological

[1] As discussed in Chap. 2, the term MINLP is used to refer to both a program class and a type of algorithm. The engine optimization problem is of MINLP class, but is not well-suited to the same-named algorithm.

[2] Hypothetically, multiple architectures could each reach the maximum efficiency, resulting in global optima, not just a single optimum.

evolution, the more fit individuals are more likely to pass on their traits (design values) to the next *generation* of individuals. Additionally, random mutations to design variables are used to increase diversity in the population.

3.2 Variable Structure and Storage

There are many types of evolutionary algorithms and the appropriate choice is partly determined by how the variables are structured and stored. Each candidate power plant (hereafter referred to simply as an *individual*) has to be distilled to a set of defining design variables that are then stored in memory. The structure that stores the associated variables has to be compact while still allowing for complex and large systems. As suggested by Emmerich, a graph notation is used where each node is a device that performs an energy transfer or transformation, and each edge is a matter flow stream [108,109]. In this study, the graph is directed from specified source nodes of air and fuel (methane) to a sink node representing the exhausted mixture of air, unburned fuel, and reaction products. An arbitrary example of a graph of this type is shown in Fig. 3.1. Circles represent nodes and arrows represent edges.

A function can then *decode* this graph and calculate the thermodynamic properties of each edge using the known thermodynamic properties of the source air and fuel nodes and the device properties. The work transfers associated with relevant devices are not stored as edges, but rather computed as dependent variables to yield a net work output.

The structure above is sufficient for graphs without heat transfer or feed-back matter transfers. Such simple problems ultimately proved enough of a challenge for

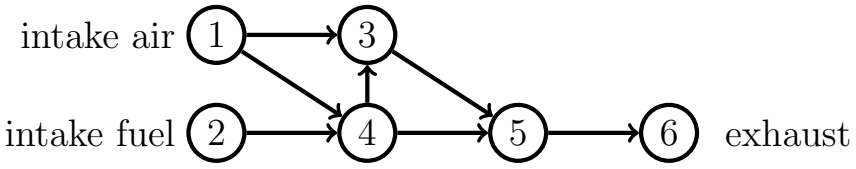

Figure 3.1: Generic directed graph for source nodes to exhaust node.

the algorithm that it wasn't necessary to allow for these additional degrees of freedom associated with heat and matter transfer. The focus in the remainder of this chapter will be on simple architectures without heat or feed-back matter transfers. While an algorithm was developed that successfully incorporated those transfer modes, its coding details are different enough that it will be discussed separately in Section 3.5.

There are multiple ways of storing a graph in memory. A binary vector was ultimately chosen for its compactness and ease of recombination and mutation. The first few bits of the vector indicate the number of nodes in the individual graph. With this information, a decoder can then interpret the rest of the vector. The source air and fuel nodes are fixed so their node types and parameters are not stored for each individual. All that needs to be stored is their edge connections (matter transfers). One bit is used for each possible edge. A 1 indicates the presence of an edge, while a 0 indicates its absence. For example, the coding 10100 would be used to express the fact that a node connects to the first and third of the five other available nodes.

After the bits encoding for the edge connections of the source node comes the encoding for the *active* nodes. Each active node has device specifications in addition to its edges. Each device type is assigned a categorical code. The number of bits required for designating device type is dependent on the number of device categories. For example, to use compressors, burners, and turbines, two bits are required to represent the three possible categories. The category bits are followed by edge bits, just as was done for the source nodes. The parameters that specify the *length* of the device (e.g., pressure ratio or length of cooling load in an intercooler) are not stored as part of the individual, but are determined within the objective function to optimize overall efficiency. This will be described in more detail in Section 3.3.3.

To demonstrate this storage method, the simple-cycle gas-turbine engine will be coded into its corresponding binary vector. This engine can be depicted by the graph in Fig. 3.2. Each circle represents a device/node and each arrow represents an edge/mass transfer. The nodes are assigned unique identifying numbers, here

indicated in the circles. The first step is to translate each device type into a corresponding binary value using Table 3.2.[3]

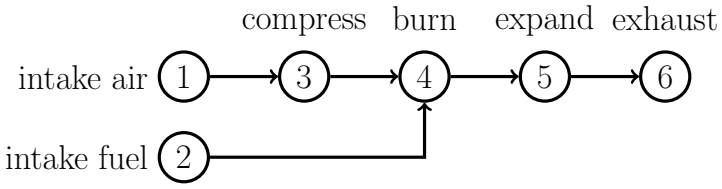

Figure 3.2: Graph for a simple-cycle gas-turbine engine.

Code	Device
00	Compressor
01	Burner/Mixer[2]
10	Turbine

Table 3.2: Device type encodings

In addition to its type each device/node is associated with a bit string representing its matter transfers/edges. Because there are no feed-back matter transfers, each device can only have edges that connect to higher-numbered nodes. Furthermore, the highest-numbered active node is the only one that connects to the exhaust sink. So Node 3 could connect to Nodes 4 or 5 only. To indicate that it only connects to Node 4, the bit string 10 is used. Node 5 *must* connect to the sink, Node 6, so no bits are needed to store this information for each individual graph.

[3]A mixer and a burner can share a code because they both act as adiabatic devices in which the inputs are chemically equilibrated. The device is more appropriately labeled as a burner when fuel and air are present such that combustion occurs, and as a mixer otherwise.

Ultimately the binary vector representing the entire graph in Fig. 3.2 is as follows:

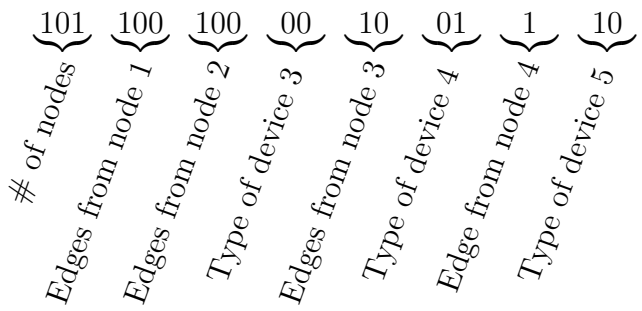

The first three bits indicate how many nodes will be in the graph. The fact that three bits are used for this information limits the number of possible nodes in this example to eight. Because Nodes 1 and 2 are always associated with air and fuel, respectively, their device types are not stored. Additionally, there are no edges connecting the air and fuel inlet nodes. Therefore only three bits are needed to store the possible edges (to Nodes 3, 4, and 5) from each of the two source nodes. The non-source nodes are additionally specified by the two-bit strings given in Table 3.2, followed by their edge connections.

3.3 Algorithm Details

Unlike the architectural variables, operating parameters are continuous and smoothly varying within their valid domain, and could be optimized by a wider range of algorithms. For this reason, many researchers recommend a tiered algorithmic approach in which an outer evolutionary algorithm develops candidate architectures and an inner loop optimizes the parameters [46, 108, 110, 111]. The fitness corresponding to the optimal parameters associated with a specific architecture is then sent back to the outer evolutionary algorithm. This approach will be taken here, and the overall outer and inner algorithm are outlined in pseudo-code below. In the following sections, each of the pseudo-code steps will be discussed in detail.

I. create initial population

II. while more generations allowed

   A. for each individual in the population

      i. check if fitness is already known

      ii. while not converged

         a. guess new parameters

         b. calculate the fitness

      iii. save individual if high fitness

   B. update elite

   C. select parents

   D. perform recombination

   E. mutate children

3.3.1 Initialization

Several algorithmic and thermodynamic properties were fixed for all run cases. These are shown in Table 3.3. In each generation, 100 individuals were computed. Of these, two were called *elite* individuals, which means they were duplicates of the two best-found solutions in the previous generation. This insures that the best individuals are not randomly lost and is a necessary condition for ultimate global convergence. All computation was done in Matlab [112]. All thermodynamic properties were calculated using Cantera, based on the GRI 3.0 mechanism [113].

3.3.2 Creation

At the onset of the algorithm, a number of individuals are needed. It is possible to seed the initial population with known functional designs. However, the goal of this study was to test the ability of the algorithm to discover functional designs with

Parameter	Value
Population Size	100
Number of Elite Individuals	2
Inlet Temperature	300 K
Inlet Pressure	1.01325 bar
Excess Air Coefficient	2
Polytropic Efficiency of Turbo-machinery	0.9
Fuel Composition	100% CH_4
Air Compositon	79% N_2, 21% O_2

Table 3.3: Algorithmic and thermodynamic properties kept as constant.

minimal input so this option was not pursued. The initial population was instead computer generated.

There is a trade-off in creating individuals between a desire for more randomness (increased independence of the individuals) and more coverage of the search space (more well-distributed individuals). Maaranen makes the distinction between pseudo-random number generators, which do the former, and quasi-random number generators, which do the latter [114]. Common attempts at maximizing diversification include the use of Latin hypercubes, algorithms that maximize the hamming distance (number of unmatched bits), and the quasi-random generator by Niederreiter [115, 116].

Both pseudo- and quasi-random creation methods were used to seed the initial population. Initially, the vectors were determined pseudo-randomly with a 50/50 probability of each bit being a 1 or 0. After observing the binary vector representation for known power plants, it was determined that the 50/50 split would generally correspond to overly connected graphs. In order to be more representative of typical power plants, the probability of each bit being zero was adjusted to 2/3. The quasi-random Latin hypercube function in Matlab was also implemented in some trials.

3.3.3 Objective Function

The next section of code involves all of the procedures required to evaluate the objective function value (fitness) associated with the optimization problem. It is responsible for translating the individual bit strings into the associated architectures, determining the order in which to process the devices, passing this information to the inner algorithm, and then returning the resulting efficiencies.

Outer Architectural Calculations

In evaluating the objective function, the first task is to analyze the node linkages and determine the appropriate order in which to calculate device outputs such that all inputs are known in time. This is done by making a dependency matrix that catalogs the input edges each node needs before its outlet edges can be computed. Any node that has no unknown dependencies is ready to be calculated.

The calculations of node outputs from inputs depend on the device type:

Compressor Using the polytropic efficiency specified in Table 3.3, the work of a differential pressure increase is added to the enthalpy of the input fluid. This process is repeated until one of the following three conditions are met: 1) the specified pressure ratio is reached, 2) the maximum allowable pressure ratio is reached, or 3) the maximum allowable temperature is reached. From the composition, final pressure, and final enthalpy, the remaining thermodynamic properties are calculated. If no compression occurs, the compressor devolves into a frictionless, adiabatic pipe.

Burner Using Cantera, the input fluid is equilibrated at constant pressure and enthalpy to determine the product composition, temperature, and entropy. If the product temperature is greater than the maximum allowable temperature, then the burner is "turned off", and the input fluid is returned.

Turbine The turbine is processed equivalently to the compressor. The fluid expands

to the maximum of the specified pressure ratio and the ambient pressure. If no expansion occurs, the turbine devolves into a frictionless adiabatic pipe.

In this way, the laws of thermodynamics are obeyed without having to write them as explicit constraints. Whenever a device has multiple outlet streams, the total mass outflow is equally apportioned between the streams.[4] The net work output for the entire plant is summed across all individual devices' work requirements. This number is then divided by the extensive lower-heating value of the fuel, which gives a first-law efficiency.

After a device has been successfully processed, the dependency matrix is updated to reflect which edges are newly calculated, and therefore which nodes are newly ready to be processed. This procedure has the benefit of never calculating unconnected portions of the graph.

To expedite processing, individuals that do not contain the bare minimum set of a compressor, burner, and turbine in their device list are assigned a negative efficiency proportional to the number of missing essential devices (2 or 3) and their individual states are not processed.

The restrictions on the operating space of the devices allows the optimization to proceed free of explicit constraints. However, they add to the roughness of the surface space for the parameters. As will be discussed in more detail below, this increases computation time per individual. An alternative method of allowing a constrained optimization problem was considered. The key constraint would be on the temperature of the system. However, this would require using temperature as an independent variable in the optimization problem. In particular, any burner outlet temperature would have to be explicitly written as an independent variable. However, no direct control exists in a burner, so a pre-determined exit temperature

[4]While apportioning the mass flow equally between edges is arbitrary, in the limit of many-node graphs, it is not restrictive. Sequential mass splits can result in finer subdivisions. For example, to assign 1/4 of a flow to a compressor and 3/4 to a burner, an initial split assigns 1/2 to a burner and 1/2 to a mixer. From the mixer, 1/2 of the flow goes to the compressor and the other half goes to the burner.

would require artificially altering pressures (and temperatures) of earlier edges. This would significantly increase the total number of variables (and associated search time).

A third option would use a penalty function that detracts from the system efficiency based on the extent of the temperature violation. This option will be revisited in Sec. 3.6.

Inner Parametric Optimization

The parameters associated with the devices (e.g., pressure ratio, fuel cell length) can vary continuously within specified bounds. As such, they can be optimized by more efficient methods than a stochastic algorithm. Additionally, removing these variables from the evolutionary algorithm significantly decreases the number of bits needed to encode an individual, which improves the performance of the evolutionary algorithm.

This inner problem is continuous and unconstrained (because of the nature of the objective function). Thus, the first attempt was made with Matlab's *fminsearch* algorithm. However, this algorithm did not give consistent results for different estimated starting conditions. This is most likely due to the large discontinuity in the slope of efficiency with respect to the parameters as a result of the burner switching on/off. Additionally, the optimal set of parameters will most likely lie immediately adjacent to this discontinuity, which would maximize the turbine inlet temperature within bounds.[5] The next algorithm tried was *fmincon*. While the problem is unconstrained, *fmincon* does a better job of searching the space in such unsmooth regions. *fmincon* returns similar, but not identical, parameters for various starting points. Because the goal is to find a global optimum, this prompted the use of a multi-start method, which makes several parallel calls to *fmincon* with different starting points, and uses this data to determine an overall optimal point. Matlab's multistart algorithm consistently found the optimal parameters for known problems, such as the

[5]While this is certainly true for the pressure ratio in the work-regenerative engine, it is less obviously true when there is also heat transfer. With recuperative heat transfer, the optimal pressure will likely be a moderate value.

two-stage reheat turbine.

3.3.4 Selection

Selection is the process of choosing which individuals will be part of the parent group that recombines to produce the next generation. Generally, more-fit individuals should be chosen as parents. There is a trade-off between only choosing the very best individuals, which speeds the convergence and limits the size of the search space, versus choosing more broadly, which slows convergence but expands the size of the search space and diversity of candidate solutions.

In order to balance these competing aims, a size-three deterministic tournament selection was used [62]. In this process, three individuals are chosen randomly from the population. Within these individuals, the highest fitness individual becomes a parent.

This process is repeated as many times as parents are needed. As will be discussed in the next section, two parents are used to produce two children.[6] Some variants of the algorithm allow for subsequent generations to be chosen from *both* the parent and child populations. However, as Beyer recommends for unbounded search spaces, only the child population was used here [58]. As such, the number of three-way tournaments is equal to the number of non-elite individuals in a generation.

3.3.5 Recombination

Recombination is the method by which attributes of multiple parents are combined to create a child. While not strictly necessary for global convergence, many researchers have found that it improves convergence speed over algorithms that solely use mutation [57]. In its simplest variation, called one-point crossover, a random number, x is generated. The first x bits of the new individual are taken from Parent 1. The

[6]There is nothing stopping the algorithm from choosing the same parent twice. Similarly, there is a small probability that the same individual will be represented multiple times in the three-way tournament.

remaining bits are taken from Parent 2. In two-point crossover, two random numbers are used to determine switching points from one parent's data to the next. Both of these strategies are forms of discrete crossover in which each bit is taken from one parent or the other.[7]

In situations in which some variables are known to be highly correlated with each other, it is common to group these variables and only allow crossover between, but not within, groups [117,118]. In this application, there were no obvious groups to be made, so simple one-point crossover was used. Two children were made from each crossover process by flipping which parent is used for the first segment of bits and which parent is used for the remaining bits. For example, if the parents were:

Parent 1	Parent 2
0 1 0 0 0 0	1 1 1 1 1 1

and the random number indicating the crossover location was 4, then Child A has the first four bits from Parent 1 and the last 2 bits from Parent 2. Child B has the first four bits from Parent 2 and the last two bits from Parent 1. This results in

Child A	Child B
0 1 0 0 1 1	1 1 1 1 0 0

3.3.6 Mutation

Mutation is the process by which, at a set probability, variables within an individual are changed. With binary data this is a simple bit flip.[8] It is common to set the probability of bit mutation to a rate of $1/n$, where n is the length of the bit string [58].

[7] In real-valued individuals (more typical of evolutionary strategies), there are also intermediate methods in which the values for the new individual are averages of the values of the parents [62]. These algorithms also tend to have strategy variables (algorithmic parameters as variables), which can also be part of the crossover process [53].

[8] With real-valued data, mutations can be specified by a distribution function with the mean equal to the existing value [58, 62].

The mutation rate is often adjusted as the algorithm progresses in time. Mainly this occurs through a gradual decrease in the probability of mutation as the algorithm starts to converge on a region in the search space [108]. However, to avoid algorithms getting trapped in local optima, many researchers also employ a mutation *shock* in which the mutation rate is raised significantly (between 25 and 50% probability) for a few generations and then returned to the baseline level [21]. Sometimes mutation shocks are employed at a fixed rate (e.g., every 100 generations) and other times based on convergence criteria when the population loses diversity [119].

In this algorithm, various mutation strategies were investigated. As a baseline, the $1/n$ strategy was used. In Mutation shocks were applied at a fixed rate (e.g., every 50 generations) and at a convergence-based rate (e.g., after 200 generations without finding a new elite individual). Each mutation shock temporarily increased the mutation probability to 25%.

3.3.7 Stopping Criteria

Stopping criteria can be set by an absolute period of time (or number of generations), or it can be based on convergence [21, 62]. In this work, the algorithm tended to converge on a locally- (but not globally-) optimal solution fairly quickly. As such, the convergence-based ending criteria tended to stop the process prematurely. A generation-based criteria was used instead, with a number of generations varying up to 20 times the observed time frame for first convergence.

3.4 Tools to Monitor Algorithm Success

Fine-tuning many of the algorithmic parameters is somewhat of an art. This section describes some of the commonly used techniques to monitor algorithm progress and make those parameter adjustments.

Because the code is stochastic, minimal information can be discerned from any individual trial. Instead, trends between trials with the same initial criteria (but

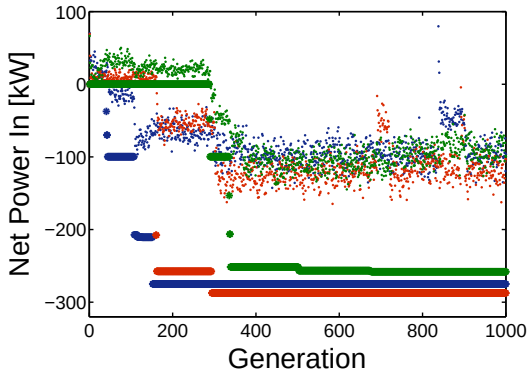

Figure 3.3: Algorithm progress of four different trials (distinguished by color). Small dots represent the average value and large circles represent the best (elite) value each generation.

different random number generator seeds) can be monitored. Figure 3.3 shows an example of algorithm progress over time for four separate trials. The vertical axis displays net work into the system, so the most fit individual is actually the most negative on the graph.[9] Each color represents a different trial. The small dots correspond to average individual fitness for that generation. The larger circles correspond to be the best individual fitness for that generation.

Figure 3.3 indicates that progress may be stalled after a relatively early point in the overall time frame. The average individual in each generation has a zero or positive net work in, indicating that most of the graphs represent non-functional power plants. Upon seeing this information, the selective criteria might be made more stringent to encourage more reproduction of fit individuals while at the same time implementing mutation shocks to attempt to dislodge the algorithm from local optima.

However, an average value per generation does not give a great sense of the overall distribution of individuals. For this reason, graphs like Fig. 3.4 were also used for individual runs to get a sense of the quartiles for fitness, in addition to the mean.

[9]This is in keeping with convention that despite the specific application, optimization problems are framed as minimization problems.

(Plotting multiple runs on top of each other was ultimately too crowded to decipher on this kind of graph). This plot confirms what the previous plot hinted at: Even the top quartile of individuals has no net power output and a fitness score of zero. In order to improve from this result, a stronger selection pressure may be applied.[10]

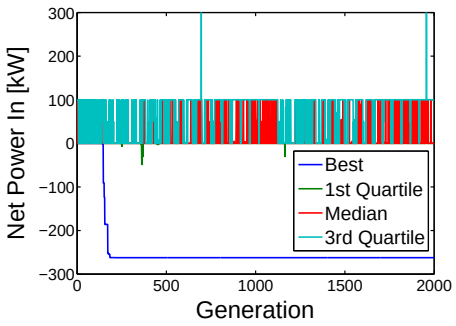

Figure 3.4: Distribution of individuals over time.

Another useful visualization is that shown in Fig. 3.5. Here the entire population at a specific generation is visualized as a black and white picture. Each row is the binary vector for a single individual. For a specific trial of five-node individuals, the figure at left shows Generation 1 and at right shows Generation 400. It is relatively easy to see that most of the individuals have converged on the same or a small subset of possible vectors by Generation 400. This kind of quick check is much easier to do visually than it is to do with the raw binary data.

[10]Various attempts to increase selection pressure through larger tournament size or changed probability of choosing the best individual within the tournament were tried. Unfortunately, none of these led to a significant improvement to the graph in Fig. 3.4.

Generation 1 Generation 400

Figure 3.5: Visualizing diversity in binary matrices of five-node individuals.

3.5 Inclusion of Heat and Matter Transfer

Adding heat transfer to the evolutionary algorithm complicates both the variable storage and the fitness evaluation. Heat has to flow downhill in temperature and the irreversibility of the process depends on the size of the temperature difference.

Heat transfer can be added as edges, like mass transfer. This requires adding at least one variable per edge to describe whether its type is matter or heat. The variable increase results in a prohibitively expensive increase in the size of the search space. It also leads to complications of where exactly the source and sink for the heat transfer are located. If the heat transfer edge originates from a node that is designated, for example, as a compressor, is the heat transfer performed before, during, or after compression?

For these reasons, an alternative approach was taken in which a heat transfer device, or node type, was added. For intercoolers, this is straightforward: a stream enters the device, is cooled to some set temperature, and the entropy is treated as having been rejected to the environment.[11] For internal heat exchange, the approach taken was to split the process into two devices: a receiving side and a donating side. As such, heat exchange devices have to work in pairs. Steps were taken in creation,

[11]Whether this entropy is rejected to water or air in the environment is immaterial.

crossover, and mutation to ensure that graphs had an even number of heat exchange units (0 being an acceptable number). If exactly two units existed, then the pairing was obvious. For algorithms that allowed more than two heat exchangers, each heat exchanger had to store the number of its partner node. This was done in the same way that pressure ratios were stored for compressors and turbines.

Internal heat exchange adds an iterative loop to the calculation of the thermodynamic states in the graph edges. The beginning and end of the list of devices that need to be iterated are both marked by heat exchangers. When the first heat exchanger is processed by the objective function, a vector is started that stores every subsequent node that depends on the output of that heat exchanger. Once the partner end of the heat exchanger is found, the vector of nodes that needs to be iterated over is complete. The heat exchange is then modeled as occurring in a counter-flow, adiabatic, and isobaric chamber. In the same way that turbomachinery has characteristic polytropic efficiencies, heat exchangers have a fixed effectiveness (typically set to 0.9).

This approach works equally well for regenerative matter transfers. Mixers and splitters work in pairs in the same way that heat exchangers do and any feed-back loops can be iterated in the same manner. Allowing matter feed-back loops removes the acyclic aspect of the graph. This means that every node can hypothetically connect to any other active node, not just those numbered higher than it. This increases the number of bits required to store edges for each node. It also slightly increases the computation necessary for determining the dependency matrix of which nodes depend on the outputs of other nodes.

Ultimately, any extremely efficient architectures are going to require regenerative heat and/or matter transfers. Having a procedure for handling them was an essential part of the evolutionary algorithm development process. However, usefully incorporating these transfer mechanisms will require resolving the currently severe computational limits discussed below.

3.6 Results

While the algorithm can operate for any number of graph nodes, it seemed prudent to first test its ability on small individuals. The smallest functional architecture would be a Brayton cycle (compressor, burner, turbine), which as shown previously requires six nodes, including the source and sink nodes. The code consistently derived the Brayton architecture with the associated optimal pressure ratio when operated over a range of algorithmic parameters.

Results will now be shown for a specific set of parameters, as specified below.

Creation 100 individuals of 15 bits created using Matlab's Latin Hypercube calculator, *lhsdesign*.

Selection Size-3 Deterministic Tournament

Mutation 1/15 probability per bit with shocks to 1/4 for 2 generations every 12 generations without finding a new elite

Crossover 0.8 probabilty of 1-point crossover. Location chosen uniform randomly.

Memory Fitness and bitstring saved for all individuals with positive work out (negative work in).

Parametric Calculation Matlab's *fmincon* was used with *MultiStart*. The minimum number of starts was set to 1. The pressure ratio was allowed to be between 1 and 200, but this was normalized to a scale of 0 to 1. On the 0 to 1 scale, the initial guess was set to 0.5.

Using these specifications, fifty trials were run and the CPU time for each subprocess was averaged, as shown in Table 3.4.[12] The bolded rows represent the three largest categories of processes.

[12]Matlab's *cputime* function was used rather than *tic/toc* to get a more accurate assessment of processing time.

Process		Time (s)
Optimize Parameters		**116**
Evaluate Structure		**58.7**
	Burn	21.7
	Compress	9.63
	Expand	7.54
Evolve Individuals		**0.597**
	Create	0.313
	Memory	0.227
	Mutate	0.0237
	Select	0.0185
	Crossover	0.0156

Table 3.4: Average CPU Process Times over 50 Trials

Optimize Parameters refers to all processes involved in running *fmincon* with *MultiStart* minus the time spent within the structural calculations. This accounted for the vast majority of the overall algorithm time, doubling the next most time-intensive process, which was the structural evaluations.

Within Evaluate Structure, the most intensive functions were those associated with the individual devices. The chemical equilibration process makes the burner evaluations slow. This process already was streamlined to include only eight possible molecules: CO_2, CO, CH_4, H_2O, H_2, O_2, N_2, and Ar. In reality, other minor species would also be created. It is unclear how to reasonably expedite this process further. The compressor time is greater than the turbine time because the compressor has to obey temperature peaks, which requires additional calculation.

The other processes are bundled into Evolve Individuals. These account for only a small fraction of the overall time. Any significant algorithm speed-ups will have to result from improvements to the previous two sections. Conversely, it is possible that more sophisticated process could be used here to improve the efficiency of the algorithm without a large overall time penalty.

Figure 3.6 shows the CPU time and algorithm generations required to derive the

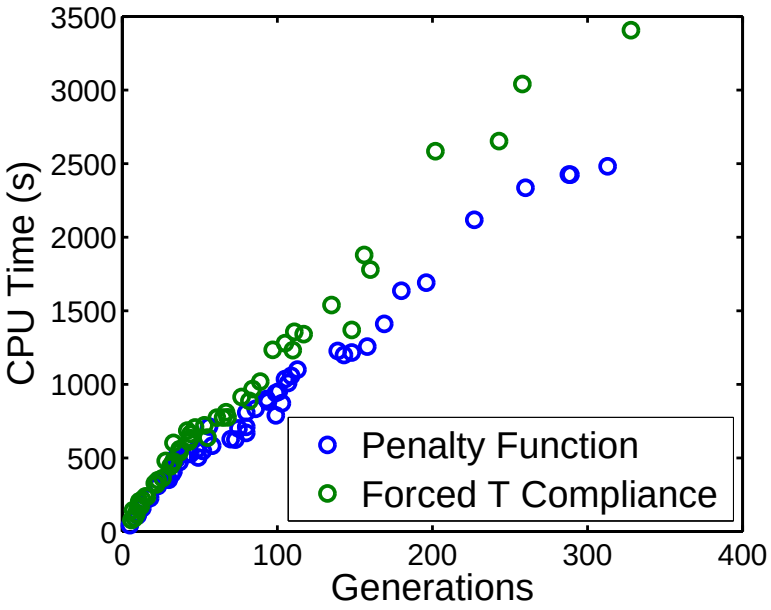

Figure 3.6: Algorithm progress for an objective function with an explicit penalty function vs. one with internally enforced temperature compliance.

Brayton over 50-trial studies. The green circles represent the data for the algorithm as described above with no penalty function and device processing that forces peak temperature compliance. The blue circles show an alternative in which the devices do not enforce temperature constraints, but a penalty function is used such that the optimization problem becomes:

$$\text{min: } \dot{W}_{netin} + 10000(\text{Realized Peak T} - \text{Allowable Peak T}) \tag{3.1}$$

where the temperatures are measured in Kelvin and the work is measured in Watts. The optimal Brayton has approximately -300,000 W in. The scalar of 10,000 was chosen such that each degree of violation results in an approximate 3% penalty in net work and efficiency.

For both the penalty function and forced temperature compliance cases, the

CPU time is approximately linearly correlated with the number of computed generations. (The flattening out of the blue series at high generations is likely a statistical anomaly.) The slope of the green series is steeper, meaning that on average it takes more time to compute a generation with the forced temperature compliance. This is largely due to the time-intensiveness of calculating the maximum allowable pressure in the compressor when the temperature limit has been reached. In the penalty function series, this calculation is avoided and the compressor outputs the fluid state solely as a function of the desired pressure ratio.

While the penalty function trials had faster processing per generation, on average they required more generations to determine the optimal Brayton architecture. This resulted in the penalty function being slower overall. The median time to completion of the algorithm was 712 seconds with the penalty function and 650 seconds without. If this relationship holds for larger graph sizes, then the 9% additional time required for the penalty function will be quite costly. However, it is unclear if this relationship will hold, particularly when larger systems will on average involve more compressor calculations per individual.

Using only compressors, burners, and turbines, Ramakrishnan determined that the next-smallest architecture that would improve efficiency would be a reheat Brayton cycle ($CBTBT$) [3]. This cycle allows for more of the combustion to occur near the maximum temperature and therefore reduces irreversibility. The architecture requires eight devices including the source and sink nodes. However, when allowed to make up to eight-node architectures, no attempted variant of the algorithm ever produced this cycle. Instead, it produced redundant Brayton systems. The redundancies varied. For example, any burners with only a single fluid input devolved into frictionless pipes, as did any turbines whose input was at atmospheric pressure. Other redundancies took the form of serial compressors or burners rather than a single stage with a higher overall pressure ratio.

When heat exchangers were added as potential devices, the algorithm successfully evolved a recuperative Brayton cycle only once out of dozens of trials. In this

cycle, thermal energy is recovered from the turbine outlet and transferred to the post-compression air. In all other runs, only non-recuperative Brayton cycles were developed. These seven-node Brayton cycles could involve any of the redundancies discussed above, as well as redundancies with heat exchangers between equal temperature streams (like inlet air and fuel), or sequential cold-side and hot-side heat exchangers.

3.7 Future Potential

This problem scales quite poorly with the number of devices in the architecture. The minimal five-device problem requires 15 bits to store if done acyclically with an assumed problem size (i.e, only storing device types and edges). Using the forced-temperature compliance data above, this problem was typically solved with 100 individuals in approximately 70 generations. In the implausible best-case scenario where all individuals were unique, a single trial searched 100 individuals/generation * 70 generations = 7000 individuals before finding the optimum. In the realm of 15-bit binary vectors, there are 2^{15} possible individuals. Thus, examining 7000 individuals required searching approximately 20% of the total space.

If this efficiency scales linearly with problem size, then at each vector length 20% of the search space might have to be searched. Table 3.5 gives some estimates of how many individuals this will require searching for differently sized problems. This is a fairly optimistic view, since in reality there will be repeat individuals, inconsistencies between runs, and potentially greater-than-linear time scaling. Reality is likely significantly worse than the already unreasonable times reported in Table 3.5. This starts to explain why the current algorithm was unsuccessful in finding the eight-node reheat Brayton solution or consistently finding a thermally recuperative engine when that option was allowed.

All of the timing data comes from computations performed on a Lenovo T430 ThinkPad with an Intel i5 core running at 2.80 GHz. While attempts were made

Devices	Bits to store	Individuals in space	Individuals to search	Time to search
5	15	3.28e4	6.55e3	O(minutes)
8	39	5.50e11	1.10e11	O(centuries)
10	60	1.15e18	2.31e17	O(million years)
15	130	1.36e39	2.372e38	O(age of universe)

Table 3.5: Scaling of evolutionary algorithm computation with architecture size

to optimize the efficiency of the code, the programmer is not an expert in high-speed computing. In collaboration with more expert computer programmers, the code could easily be implemented in parallel, and with sufficient cores, computed in a fraction of the time. However, to have the potential to find novel architectures, the algorithm will have to allow graphs with upwards of 10 or 15 nodes. Even with the massive speed increase afforded by a supercomputer, this would be a daunting task.

The only possibility of reducing the search times in Table 3.5 would be to develop an algorithm that needs to search a significantly smaller fraction of the search space. While this seems difficult, a few possible ideas that could help are briefly discussed here.

First, there is the possibility of using algorithms that create random graphs that share key characteristics with given target graphs. This has been done extensively in certain fields, particularly with efforts to model the structure of the internet [120–122]. Computer libraries exist to do such graph generation [123]. However, the characteristics tend to be probability distributions of types of nodal connections (e.g., the probability that a node with three edges is connected to a node with one edge). Because even highly efficient power plants are likely to have less than 50 nodes, such probability distributions start to become overly prescriptive.

Second, there is a broad literature on methods of handling constraints. While the choice was made here to produce an unconstrained optimization, it came with the consequence of a strangely shaped, unsmooth search space. This is not necessarily the best approach. A larger variety of penalty functions could be explored to see if a

significant speedup is possible. Some researchers have shown that an unconstrained, overly restricted search space did not improve their search times over the explicitly constrained version of their problem [49]. Paredis has done interesting work using predator-prey models to co-evolve individuals in an evolutionary algorithm alongside a set of constraints [124,125]. Both the individuals and the constraints are evaluated for their own fitness metrics. Individuals that meet the most constraints are deemed more fit whereas constraints that are harder to meet are more fit. This process, sometimes called a filter approach, prioritizes the evolution of individuals that can meet the toughest constraints.

As discussed in Sec. 3.4, almost all individuals in each generation tend toward no work output. This indicates that in the process of crossover and mutation, more functional graphs are broken than formed. This is typical of graph-based evolution and several researchers have experimented with methods to preserve functional groups (or subsets of graphs) during crossover [126–128]. For example, compressors could be forced to always be attached to a single inlet node and at least one outlet node. Genetic repair mechanisms could be used that force newly reproduced individuals to meet some level of functionality [58]. However, these tend to be somewhat arbitrary and reflect the biases of the programmer. They also add computing time per generation.

A final idea is to group sets of devices into larger building blocks. There could be a single building block that performs first adiabatic and then isothermal burning/expansion when the maximum temperature was reached. On the plus side, this could help the algorithm find more interesting architectures. On the negative side, the more building blocks that exist, the more bits that are required to uniquely code them, which adds to the length of the individual and the size of the search space. Additionally, because the goal of this algorithm was to build architectures with minimal prescription, the use of more sophisticated building blocks eventually defeats the purpose and risks becoming overly prescriptive.

Ultimately, this project was part of a larger study to determine optimal engine

architectures by any viable methodology. As will be shown in the coming chapters, a systematic, thermodynamic approach was much more successful in expanding the bounds of known optimal cycles to larger, more complex architectures. Building off of knowledge about optimal work-and-heat integration, the best use of matter transfer involving air, water, fuel, and products thereof was found. Architectures were also developed with solid-oxide fuel cells. These architectures require many more devices than the simple systems examined here. The success of the thermodynamic approach for the current purposes obviated the need for further analysis with this evolutionary methodology.

Chapter 4

Systematic, Thermodynamic Methodology for Engine Architecture Optimization

The remaining majority of this thesis uses a systematic, thermodynamic methodology to optimize engines. In this chapter the approach is explained in detail. A framework is introduced that is capable of representing any possible steady-flow chemical engine. This framework is used to expand to more complex optimal engines from simpler, more-constrained systems. The chapter concludes with two conceptual tools that help with the optimization process.

4.1 Representation of an Engine

4.1.1 Engine Nomenclature

It will be convenient to have a shorthand form for identifying different engine architectures. This is accomplished by referring to an engine as an acronym for its sequence of devices. Each device is designated by a single letter as given in Table 4.1. For example, the recuperative simple-cycle gas-turbine engine shown in

Device	Shorthand
Compressor	C
Turbine	T
Burner	B
Solid-oxide Fuel Cell	F
Matter Mixer	M
Matter Splitter	S
Intercooler	I
Heat Exchanger (receiving side)	X_{in}
Heat Exchanger (donating side)	X_{out}

Table 4.1: Shorthand nomenclature used to refer to common devices in engine architectures

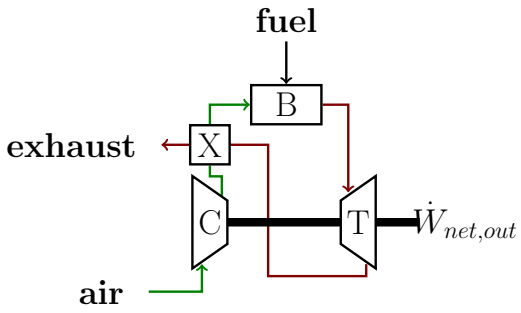

Figure 4.1: A recuperative simple-cycle gas-turbine engine.

Fig. 4.1 will be referred to as

$$CX_{in}BTX_{out}$$

which shows the basic sequence of devices in the appropriate order along the primary flow path.

4.1.2 Device Parameters

All devices have constraints related to their physical limitations, and some also have energy transfers that can be controlled through *length* parameters. A constraint is a restriction on the ability to build a device and restricts the operating range

Constraint	Value	Citation
Turbine polytropic efficiency	0.9	[129]
Compressor polytropic efficiency	0.9	[129]
Pump polytropic efficiency	0.75	[129]
Peak turbine blade temperature	1800 K	[84]
Heat exchanger effectiveness	0.9	[130]
Peak heat exchanger temperature	1100 K	[131]
Solid-oxide fuel cell YSZ conductivity	6.5 S/m	[132]

Table 4.2: Device constraints used in numerical illustrations.

or capability of the device. A maximum allowable turbine inlet temperature, a polytropic efficiency, and a heat exchanger effectiveness are all constraints. Because the goal is to design optimal systems, constraint values are chosen to reflect state-of-the-art technology. Constraints can change in time as device manufacturing and materials improve, but at any given time, the current constraints provide hard limits for the respective device's operation. The methodology presented here is independent of the numerical values of these constraints, but the resulting architectures may not be. When architectures are modeled with numerical illustrations, their constraints take on the values given in Table 4.2.

A device whose extent of an energy transformation is directly coupled to a transfer can be thought of as having a *length* associated with its process. The length is an operating parameter that a device can be manufactured or manipulated to meet. Pressure ratios for turbo-machinery, utilization for a fuel cell, or cooling load of an intercooler are all examples of device parameters that indicate the desired length of their service. This is in contrast to a burner, which has no coupling between the extent of the chemical transformation and an external work or heat transfer. Once the burner receives the fluid inputs, the output is fixed (i.e., a burner cannot be specified to go to 40% extent of reaction). This lack of a controlling transfer and its associated length parameter is also true for mixers. With no heat or work controls, these devices transform the entering fluid solely as a function of the fluid and device constraints.

4.2 The Systematic Approach

The goal of this dissertation is to develop maximally-efficient, steady-flow, chemical engines. By the Gouy-Stodola theorem, this is equivalent to optimizing for the minimum engine entropy generation. The systematic entropy generation minimization approach developed by Ramakrishnan is used here [3]. This process involves four steps:

I. Begin with a bare-minimum engine architecture.

II. Introduce one new degree of freedom.

III. Insert the option for this degree of freedom between every existing device in the sequence.

IV. Minimize total entropy generation of the complete system by optimizing the length of each device in the new candidate sequence.

Steps II-IV are then repeated as many times as required to extend the complexity of the engine. Each of these steps will now be discussed in more detail below.

4.2.1 A Bare-Minimum Sequence

A bare-minimum engine sequence must include a work-extraction mode, and for a chemical engine, a chemical reactor to release the bond energy of the fuel. Ramakrishnan used a turbine as the work extractor and a burner as the chemical reactor. Additionally, because the turbine requires a pressure difference, for steady-flow, this minimal sequence also needs an upstream compressor. Thus, his minimal sequence is the compressor, burner, turbine, which will hereafter be referred to as CBT.

4.2.2 A New Degree of Freedom

A degree of freedom can be a new energy transfer mode, a new device, or a new fluid. All three are explored in this thesis. Energy transfer modes are specified by the type

of transfer (i.e., work, heat, or matter), the flow rate of energy, and the locations where the transfer occurs. Devices are specified by their purpose, constraints, and operating parameters as previously discussed. The choice of a new fluid is limited. By definition, anything out of equilibrium with the environment has exergy and is therefore a "fuel".[1] Thus, almost all fluids that could be introduced would fall in this category and would be enveloped in the existing fuel label. With an environment composed of air at 100% relative humidity, both air and liquid water are non-exergetic fluids that can be used in the engine. Air is essential as an oxidizer for the chemical fuel. Thus, the only optional fluid is water.

Ramakrishnan already examined the optimization of systems that allow internal work and heat transfer as well as external heat transfer. In Chs. 5 and 6 of this thesis, his optimal sequence will be expanded by introducing internal and external matter transfers. In doing so, these chapters will explore the addition of matter mixers and splitters as new devices, and water as a new fluid. Chapter 7 will start with the bare-minimum sequence and introduce solid-oxide fuel cells as a new device.

4.2.3 Insertion of the Degree of Freedom

The new degree of freedom could occur anywhere within the existing sequence. For example, if the original sequence was CBT and the new degree of freedom was a fuel cell, F, the new candidate sequence would be

$$(CF)_k(BF)_m(TF)_n$$

where k, m, and n are arbitrary integers allowing for repetition of sequences of devices. This representation allows for all possible uses of the new degree of freedom without changing the baseline sequence. If the use of a fuel cell is removed, the sequence must degenerate to the previously existing best sequence, the CBT. Thus,

[1] In this thesis when a fuel is strictly chemical, it will be specified as such.

the compressor, burner, and turbine cannot be rearranged relative to each other.[2]

4.2.4 Minimization of Entropy Generation

Every single restrained device in the candidate architecture must have its length parameters re-optimized for the new sequence. This includes the previously optimized devices, which may now function differently given the new overall architecture. For example, the optimal pressure ratio for a compressor will change going from a CBT to a system that also includes fuel cells, such as a $CFBT$. Any device whose contribution always increases the entropy generation of the entire sequence will have an optimal length of 0, which removes the device from the sequence. In this way, the domain of all possible sequences is whittled down to a specific sequence that likely involves fewer than the 2(k+m+n) devices allowed. The length of each device is limited by the plausible domain of the relevant operating parameter.

4.2.5 Implications of the Systematic Approach

The above procedure enables the broadening of the scope of engine architecture to more and more complex systems, as desired. What remains unclear is how to accomplish the parameter optimizations in Sec 4.2.4 under a wide range of cases. Different strategies are used in different cases. Two methodologies that are particularly helpful are discussed in Sec. 4.3 and 4.4.

[2]Consider introducing a new device, like a fuel cell, and taking it to the limit of almost zero length (for example, zero residence time and fuel utilization). The other devices in the sequence should not rearrange relative to each other as this length approaches, and then reaches, zero.

4.3 First Tool for Parameter Optimization: Attractor Minimization

Ramakrishnan developed the *attractor* method for use in optimizing steady-flow combustion engines. This method has two key concepts. First is that for many systems, the entropy generation is proportional to the exhaust entropy. Second is that it is possible to map every device process to the effect it will ultimately have on the exhaust entropy. These two points are explained in detail below.

4.3.1 Using Exhaust Entropy as a Proxy for Entropy Generation

The total engine entropy generation can be split into the internal component that occurs as a result of non-reversible processes and the external component, which is the result of the exhaust products equilibrating in the environment.

$$\dot{S}_{gen} = \dot{S}_{gen,int} + \dot{S}_{gen,ext}$$

Applying an entropy balance around the system boundary, this becomes

$$\dot{S}_{gen} = \frac{\dot{Q}_{out}}{T_0} - \frac{\dot{Q}_{in}}{T_0} + \dot{m}(s_{out} - s_{in}) + \dot{S}_{gen,ext}$$

Using three specifying criteria that will often be valid in this thesis, the expression above can be simplified. The first is the use of a strictly chemical (not thermal) fuel. This implies that there is no heat transfer in to the engine. Therefore, the $\dot{Q}_{in}$ term can be eliminated

$$\dot{S}_{gen} = \frac{\dot{Q}_{out}}{T_0} + \dot{m}(s_{out} - s_{in}) + \dot{S}_{gen,ext}$$

Second, in many of the systems analyzed here, there is no heat transfer from the engine to the environment (i.e. they can be idealized as adiabatic). These cases

allow further simplification to

$$\dot{S}_{gen} = \dot{m}(s_{out} - s_{in}) + \dot{S}_{gen,ext}$$

Third, if the inlet state is fully specified, $\dot{m}$ and s_{in} are fixed. Thus, the total entropy generation varies only with the outlet entropy and the external entropy generation. By Guoy-Stodola, the external entropy generation is directly proportional to the exhaust exergy. The mass-specific exhaust exergy transfer ψ_{out} is written

$$\psi_{out} = (h_{out} - h_{out}(T_0, P_0)) - T_0(s_{out} - s_{out}(T_0, P_0)) + \psi_{out,chem}(T_0, P_0) \tag{4.1}$$

where T_0 and P_0 are the environmental temperature and pressure, which together comprise the restricted dead state. The first two terms on the right of Eq. 4.1 account for the thermo-mechanical exergy, and the final term accounts for the chemical exergy. The exhaust is at the environmental pressure and, as specified in Ch. 1, has already been chemically equilibrated internally. Thus for fixed inlet matter, all of the properties in Eq. 4.1 at the restricted dead-state are constants.

The chemically-equilibrated exhaust has to exit the engine at ambient pressure. Taking the derivative of Eq. 4.1 with respect to the outlet entropy at constant pressure and chemical affinity gives

$$\left.\frac{\partial \psi_{out}}{\partial s_{out}}\right|_{P=P_0, \mathcal{A}_{out}=0} = \frac{\partial h_{out}}{\partial s_{out}} - T_0 = T_{out} - T_0$$

where $\mathcal{A}_{out}$ is the chemical affinity of the flue gas. The right side $T_{out} - T_0$ is always greater than or equal to zero. This indicates that ψ_{out} is a monotonic function of s_{out}. Therefore, minimizing the total entropy generation of the system is equivalent to minimizing the outlet entropy under the three given conditions. When the chemical system does not involve external heat transfer or variations in external mass transfer, it will only be necessary to examine variations in s_{out}. This will be helpful surprisingly often. When these conditions are not met, the parameter optimization will require

a different approach to entropy generation minimization.

4.3.2 Mapping the Thermodynamic State to its Effect on Exhaust Entropy

Ramkrishnan showed that it is possible to map every thermodynamic state and process through a device to the effect it will ultimately have on s_{out}. Air, water, fuel, and products of reaction are tracked as they move through the engine. At the entrance and exit of each device, the fluids traveling in parallel are mapped to a state in which they are thermally, mechanically, and chemically equilibrated into a single product state, called the *attractor* state.[3] For flowing systems, the attractor is at the same enthalpy, pressure, and atomic composition as the original state.[4] This attractor state is essentially what would occur if the engine were to terminate at the associated point in the system and be allowed to equilibrate. Thus, the attractor's entropy is equivalent to the desired tracking variable s_{out}. With this tool, the addition, removal, or change in operating conditions of a device can be mapped to its ultimate effect on exhaust entropy, and thereby entropy generation and overall engine efficiency.

Table 4.3 illustrates how a system of arbitrary species A and B are mapped to an attractor state. The attractor state has the same extensive enthalpy, pressure, and atomic composition as the original combined gas state. Together, A and B have an extensive enthalpy of 5 Joules. If, hypothetically, chemical equilibrium dictates $A + B \rightarrow AB$ with a very high extent of reaction, then the final composition is 50% AB, 50% B. Thus, the attractor state has 2 moles whereas the actual pre-reaction state has 3. The intensive enthalpy for the attractor is then 5/2 J/mol.

[3]For formulating the attractor, it is convenient to always keep the parallel fluids in mechanical equilibrium. This can be envisioned by inserting hypothetical isentropic compressors and turbines in any secondary flow streams to bring them to the same pressure as the main air stream. This equal pressure representation is explained in detail by Ramakrishnan. [3]

[4]The natural analog for a piston-cylinder engine is to keep the volume and internal energy constant between each state and its attractor.

This enthalpy, composition, and given system pressure constitute a fully defined thermodynamic state. Entropy, and all other desired thermodynamic properties, can then be calculated using a chosen equation of state.

Property	A	B	Attractor State
Moles	1	2	2
Enthalpy [J/mol]	1	2	5/2
Pressure [bar]	5	5	5

Table 4.3: Simplified Example of Attractor Calculation.

4.3.3 Visual Attractor Analysis

The attractor methodology can be used to optimize device length parameters either visually or mathematically. Here the visual approach is demonstrated using the example of a Brayton cycle [3]. The Brayton cycle is implemented as a CBT: a compressor, a burner, and a turbine. In Fig. 4.2 this simple cycle is shown on a total-mass-specific, enthalpy-entropy diagram.

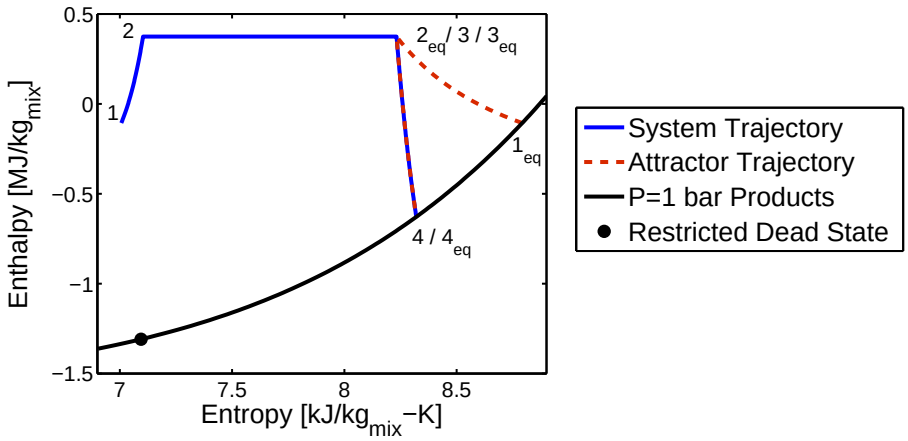

Figure 4.2: System-attractor plot (h-s) for a Brayton cycle. (Natural gas-air, $\phi = 0.43$, $\eta_C = \eta_T = 0.9$.)

The actual system trajectory is shown in blue and the attractor trajectory is shown in red. At State 1, the system is composed of unreacted natural gas and air at atmospheric temperature and pressure. The attractor is at State 1_{eq} and is at the same enthalpy and pressure as State 1, but now comprised of the adiabatically equilibrated products of natural gas and air at atmospheric pressure. State 1_{eq} is at a much higher entropy than State 1 due to the entropy generation of the atmospheric chemical equilibration. As the reactants are compressed, the system increases in enthalpy and entropy to State 2. The attractor actually decreases in entropy to State 2_{eq}, while again matching the enthalpy and pressure of the actual system. As discussed in Ch. 2, the decrease in attractor entropy occurs because the entropy generation due to chemical reaction decreases as the reactant's energy increases [14]. Because the gases are burned at constant enthalpy and pressure, the attractor does not change across a combustor. Thus, State 2_{eq} is the same as State 3_{eq}. It is also the same as State 3, because now the fluid is in chemical equilibrium and there is no distinction between the attractor state and the actual state. The fluid then expands through the turbine to State $4/4_{eq}$ on the atmospheric isobar for the chemically equilibrated composition, shown in black. The black circle marks the restricted dead state for the chemically-equilibrated system. Because it is at ambient temperature and pressure, exhaust at this state would have no thermal or mechanical exergy. As such, it is the best (lowest entropy) exhaust state possible for the given composition.

While visualizing the enthalpy-entropy plot may be illuminating for an individual system, it is particularly helpful in quickly comparing the merits of multiple possible trajectories. Figure 4.3 has the same Brayton cycle, but now for three different pressure ratios: 20:1, 40:1, and 60:1. All three cases use the same fuel, equivalence ratio, and polytropic efficiencies for the turbine and compressor. Maximizing efficiency involves minimizing the final attractor entropy. Of the three pressure ratio cases shown here, the 60:1 case ends at the lowest entropy, and therefore is the most

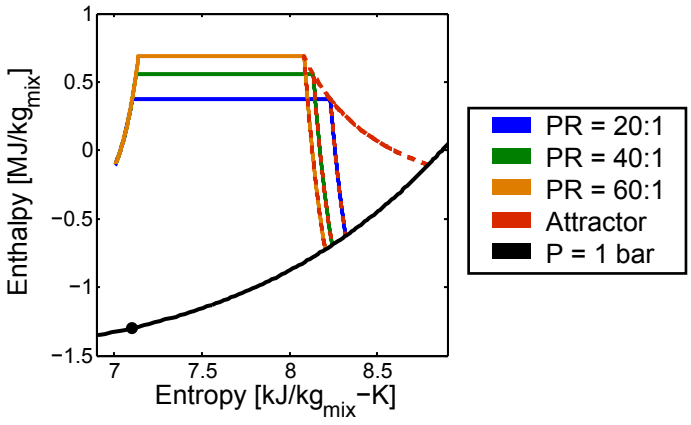

Figure 4.3: Using the attractor to compare three candidate pressure ratios for a Brayton cycle. (Natural gas/air, $\phi = 0.43$, $\eta_C = \eta_T = 0.9$.)

efficient cycle. If this inquiry were extended to all possible pressure ratios, the optimal pressure ratio would be the one associated with the lowest (farthest left) final entropy.

The lowest entropy point could be found by simply examining the system trajectory without the attractor. However, the attractor supplies information about what each device's effect is on the exhaust entropy. The shape of the red attractor curve indicates that added compression improves the potential efficiency of the system by pulling the outlet state toward lower entropy, but with decreasing marginal returns. At some pressure ratio the slope of the attractor during compression will be equal and opposite to the slope of the attractor during expansion. At this point, the marginal decrease in exhaust entropy provided by compression is canceled by the marginal increase in exhaust entropy caused by expansion. Finding the point of equal attractor slope is equivalent to finding the optimal pressure ratio. This ability to determine the effect of any marginal design change on the overall outlet entropy will be critical in optimizing device lengths.

4.3.4 Mathematical Attractor Analysis

Ramakrishnan showed that the optimal device lengths can be determined analytically, as well as visually. The exhaust entropy can be written as

$$s_{out} = s_{in} + \Delta s$$

Numbering the engine states 1 through n, this can be rewritten as

$$s_n = s_1 + \Delta s$$
$$s_n = s_1 + \sum_{i=1}^{n} \Delta s_{i,i+1}$$

The inlet is equivalent to State 1 and the outlet is equivalent to State n. Expanding the right side gives

$$s_n = s_1 + \Delta s_{1,2} + \Delta s_{2,3} + \cdots + \Delta s_{n-1,n}$$

The attractor entropy is equal to the system entropy at the outlet state. Using the eq subscript to indicate the attractor, the above equation can be rewritten

$$\begin{aligned} s_{neq} &= s_1 + \Delta s_{1,2} + \Delta s_{2,3} + \cdots + \Delta s_{n-1,n} \\ &= s_1 + \Delta s_{1,1_{eq}} + \Delta s_{1_{eq},2_{eq}} + \Delta s_{2_{eq},3_{eq}} + \cdots + \Delta s_{n-1_{eq},n_{eq}} \\ &= s_1 + \Delta s_{1,1_{eq}} + \sum_{i=1}^{n} \Delta s_{i_{eq},i+1_{eq}} \end{aligned}$$

Minimizing the exhaust entropy as a proxy for total entropy generation only works if the inlet state is fixed. In this case, s_1 and $\Delta s_{1,1eq}$ are also fixed. Therefore

$$\min\ s_{out} \propto \min \sum_{i=1}^{n} \Delta s_{i_{eq},i+1_{eq}} \tag{4.2}$$

Finally, $\Delta s_{i_{eq},i+1_{eq}}$ for any device can be written mathematically as

$$\Delta s_{i_{eq},i+1_{eq}} = \int_{h_i}^{h_{i+1}} \frac{ds_{eq}}{dh} dh$$

because h is the same as h_{eq}.

Ramakrishnan derived expressions for $\frac{ds_{eq}}{dh}$ for several devices [3]. For all devices, the derivation begins with Gibbs equation for the chemically-equilibrated attractor

$$T_{eq} dS_{eq} = dH_{eq} - V_{eq} dP_{eq} - \sum_i \mu_{i,eq} dN_{i,eq}$$

The attractor is defined as being at the same enthalpy h, pressure P, and atomic composition as the actual state, but brought to chemical equilibrium. Therefore, $H_{eq} = H$ and $P_{eq} = P$. Because the system is in equilibrium, the chemical affinity is zero. Thus the $\sum_i \mu_{i,eq} dN_{i,eq}$ term is identically zero. The equation simplifies to

$$T_{eq} dS_{eq} = dH - V_{eq} dP \tag{4.3}$$

Compressors & Turbines

Both compressors and turbines are modeled as adiabatic with set polytropic efficiencies. For a compressor

$$dh = \frac{v dP}{\eta_C}$$

where η_C is the polytropic efficiency. Substituting the above expression into Eq. 4.3 and eliminating dP results in

$$\frac{ds_{eq}}{dh} = \frac{1}{T_{eq}} (1 - \frac{v_{eq} \eta_C}{v})$$

In order to minimize entropy production, compressors should only be used when $\frac{ds_{eq}}{dh}$ is negative. This occurs when $v_{eq} \eta_C \geq v$, which is true for the gas stream before chemical equilibration, confirming the understanding that compressors are

used before burners and fuel cells. This criteria dictates the change in enthalpy through the compressor, and therefore the device length.

Similarly it was found that turbines should only be used when $v_{eq} \geq v\eta_C$. For many fuels, this is true starting very near the end of the oxidation process. Assuming that through a turbine the fluid remains chemically equilibrated, v and v_{eq} are the same and the change in the attractor entropy with enthalpy is written

$$\frac{ds_{eq}}{dh} = \frac{\eta_T - 1}{T_{eq}\eta_T}$$

where η_T is the turbine polytropic efficiency. (The greater the polytropic efficiency, the less s_{eq} increases with decreases in enthalpy.) The turbine always adds to the entropy generation, but significantly less so than a system that mechanically equilibrated with the environment in the absence of a turbine.[5]

Burners

Burners are modeled as isobaric and isenthalpic.[6] Thus, the entire right side of Eq. 4.3 is zero. Therefore dS_{eq} is also zero. The attractor for a burner is represented as a single point on the h-s diagram that is coincident with the actual state point. This is in keeping with the notion that burners have no length parameter.

Heat Exchangers

Again the analysis starts with

$$T_{eq}ds_{eq} = dh - v_{eq}dP$$

[5]The turbine entropy expression can be equivalently written in terms of dP instead of dh. This leads to $\frac{ds_{eq}}{dP} = \frac{v\eta_T - v_{eq}}{T_{eq}}$. In the absence of a turbine, there is no change in enthalpy, but there is still a pressure drop to equilibrate with the environment. This mechanical equilibration has an entropy increase dictated by $\frac{ds_{eq}}{dP} = \frac{-v_{eq}}{T_{eq}}$, which is always larger than the entropy increase for a turbine.

[6]The incorporation of friction or imperfect insulation in any of the devices does not change the methodology. In his dissertation, Ramakrishnan showed that introducing these imperfections changes the optimal number of stages of each part of the architecture (i.e., fewer stages of isothermal burning), but not the types of devices or their relative sequencing [3].

Heat exchangers are assumed to be isobaric, externally adiabatic, and involve frozen chemistry within each stream.[7] For the attractor of a single side of the heat exchanger

$$\frac{ds_{eq}}{dh} = \frac{1}{T_{eq}}$$

For the side being heated, h is increasing as the fluid progresses through the device, and therefore so is s_{eq}. The attractor's temperature increases as energy is added along the heat exchanger's flow path. The temperature of the attractor may be shifted from that of the actual system due to the endo/exothermicity of the equilibration reactions. The slope $\frac{1}{T_{eq}}$ is decreasing as energy is added to the fluid. This leads to the conclusion that less entropy generation will occur when the cold side of the heat exchanger is as hot as possible. For the hot side of a heat exchanger, h is decreasing, and so is s_{eq}. Again, the derivative is $\frac{1}{T_{eq}}$ so lower temperatures will result in the greatest entropy decrease.[8] This result–that the hot side of the heat exchanger should be as cold as possible and the cold side as hot as possible–is in line with intuition that the temperature difference across a heat exchanger should be minimized.

Figure 4.4 depicts the trajectories of the system and attractor for the cold side of a heat exchanger. The system is composed of two unmixed fluids at the same pressure: methane and air, both being heated. As expected, both the system and attractor increase in entropy with enthalpy. The attractor has a steeper slope on an h-s diagram because $T_{eq} \geq T$ for an unreacted methane-air system.

Solid-Oxide Fuel Cells

The incorporation of fuel cells into this methodology is new. The analytical expression for the change in attractor entropy with respect to enthalpy can be developed similarly to what was done previously by Ramakrishnan. A solid-oxide fuel cell tube

[7] Regardless of reactive chemistry or lack there of in the heat exchanger, the attractor is in equilibrium.

[8] Note that T_{eq} may be different for the two sides of a heat exchanger, based on compositional differences.

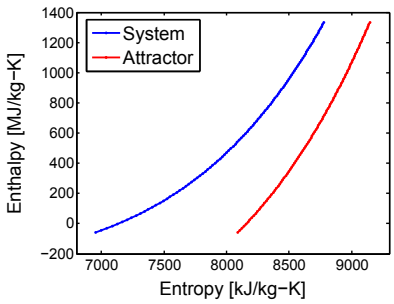

(a) Enthalpy-entropy trajectories for the system state and attractor.

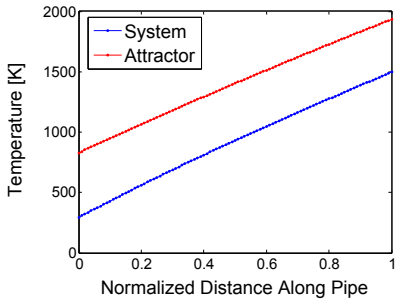

(b) Temperature profiles for accompanying H-S diagram.

Figure 4.4: Model for a heated, but unreacted air/methane system. (The cold side of a heat exchanger)

(the membrane electrode assembly along with the electrode-side channels) can be modeled as isothermal and isobaric.[9] As was true with the heat exchanger, under isobaric conditions Eq. 4.3 simplifies to

$$\frac{ds_{eq}}{dh} = \frac{1}{T_{eq}}$$

For a fixed membrane temperature, an exothermic reaction in a fuel cell tube must lead to a decrease in enthalpy. Combining this with the extraction of work requires that dh, and therefore ds_{eq}, is negative. For the system, as opposed to the attractor, as the oxidation progresses, the chemical affinity of the gases decrease. The combination of decreased affinity and decreased enthalpy lead to a drop in T_{eq}. Thus the slope of the attractor on an h-s plot is positive but decreasing, as shown in Fig. 4.5. Note that because ds_{eq} is negative for exothermic fuel cells, but zero for burners, it will always be preferable to use an exothermic fuel cell over a burner for chemical reaction where possible. Endothermic regions of the fuel cell behave differently. The kinetics at the inlet of the internally-reforming solid-oxide fuel cell are dominated by the endothermic breakdown of methane. In this region, both

[9]The validity of these assumptions is further discussed in Ch. 7.

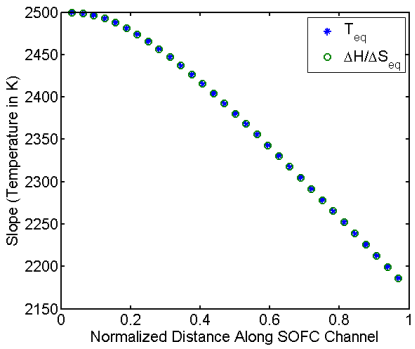

Figure 4.5: Confirmation that $\frac{ds_{eq}}{dh} = \frac{1}{T_eq}$

dh and ds_{eq} are positive. Exothermic regions should dominate the overall fuel cell energetics, so as to result in a negative change to s_{eq} overall.

When anode recycling is used, there is a stage of matter transfer just before the entrance to the anode-side channel. This matter adds enthalpy and entropy to the system, thus $\frac{ds_{eq}}{dh}$ is positive. Similarly, at the end of the fuel cell channel, matter is removed from the main path to be recirculated back around to the SOFC inlet. Removal of matter results in a downward jump on an h-s diagram.

4.3.5 Attractor Summary

In systems where there is no heat interaction with the environment and the inlet matter streams are fixed, minimizing outlet entropy is equivalent to minimizing total entropy generation. The attractor maps the effect that each device's length has on the final exhaust entropy. This information can be used visually or mathematically to determine the device lengths that minimize the overall system entropy generation. For a given degree of freedom, finding these optimal lengths is the final step in the systematic thermodynamic methodology. The attractor method will be used extensively to optimize engines for inclusion of solid-oxide fuel cells in Ch. 7.

4.4 Second Tool for Parameter Optimization: Concentric Cycle Analysis

In Chs. 5 and 6 of this thesis, the effect of optimizing an engine for variable matter transfers is examined. Changing the relative quantities of fluids entering the engine and where they merge and split from each other changes the attractor chemistry. The derivations in Sec. 4.3 no longer apply. An alternative method for determining optimal device lengths was developed for use in this thesis. This method, which is referred to as *concentric cycle analysis*, has three key ideas. First is realizing that matter transfers of interest here are dominated by diffusive and thermal equilibrations. Second, that it is possible to conceptualize a *concentric* system that separates out the thermal from the diffusive effects of matter transfer. Finally, that an analysis of the concentric system can show that many forms of matter transfer are not beneficial (i.e., that their process lengths should be zero). These concepts are discussed in detail below.

4.4.1 Equilibrations Associated with Matter Transfer

This study will include matter transfers that proceed directly from a splitter to a mixer with no intervening devices. Thus, all matter transfers are unrestrained and will involve mechanical, thermal, diffusive, and reactive equilibration. These modes of equlibration are not always equally significant, as will be demonstrated with the following two examples.

Figure 4.6a shows the compositional surfaces of an air-fuel system. Initially, the air and fuel are considered to be at mechanical, but not thermal, diffusive, or reactive equilibrium. Thus, they have the same pressure, but not necessarily temperature or composition. They are then brought into thermal equilibrium with the jump from the blue surface to the green surface. (The green surface is not clearly visible, but exists between the blue and orange surfaces.) The distance between these surfaces on the entropy axis is a measure of the entropy generation associated with the thermal

equilibration. A second jump in entropy occurs when the air and fuel physically mix, but do not yet react, thus incurring diffusive entropy generation. A much larger jump occurs between the diffusively equilibrated (mixed with frozen chemistry) and the reactively equilibrated surfaces. From this, it is evident that the majority of the entropy generation associated with the equilibration of fuel and air is due to chemical reaction.

Figure 4.6b shows the compositional surfaces associated with an air-combustion products system. While the same jumps exist between these surfaces as in Fig. 4.6a, the entropy generation is differently proportioned between the steps. The diffusive and thermal equilibrations significantly outweigh the consequences of reactive equilibrium. Teh and Ramakrishnan have previously optimized air-fuel combustion engines for minimizing the reactive entropy generation in the burner. The additional matter transfers discussed in this thesis are similar to those shown in Fig 4.6b in which thermal and diffusive entropy generation dominate.

This study involves no devices that can take advantage of diffusive matter differences (such as concentration cells), so diffusive equilibration will always be entropic. Thermal differences can be used, however, and could provide a reason to move thermal energy within a system. This is why heat transfer was already examined by Ramakrishnan. In addition to heat transfer, matter transfers can also move thermal energy. However, matter transfers that involve mixing fluids of different composition will inevitably incur diffusive entropy generation. Thus, a central question is when to use a heat or a matter transfer, if either, in an optimally efficient system.

4.4.2 Conceptualizing the Concentric Cycle

A new tool that helps distinguish between the thermal and diffusive effects of matter transfer is introduced here and will be referred to as *concentric cycle analysis*. As an introductory example, the concentric cycle is applied to the study of a simple-cycle gas-turbine engine with exhaust gas recirculation, as depicted in Fig. 4.7. In the device nomenclature, this would be $MCBTS$, where M is a flow mixer and S is a

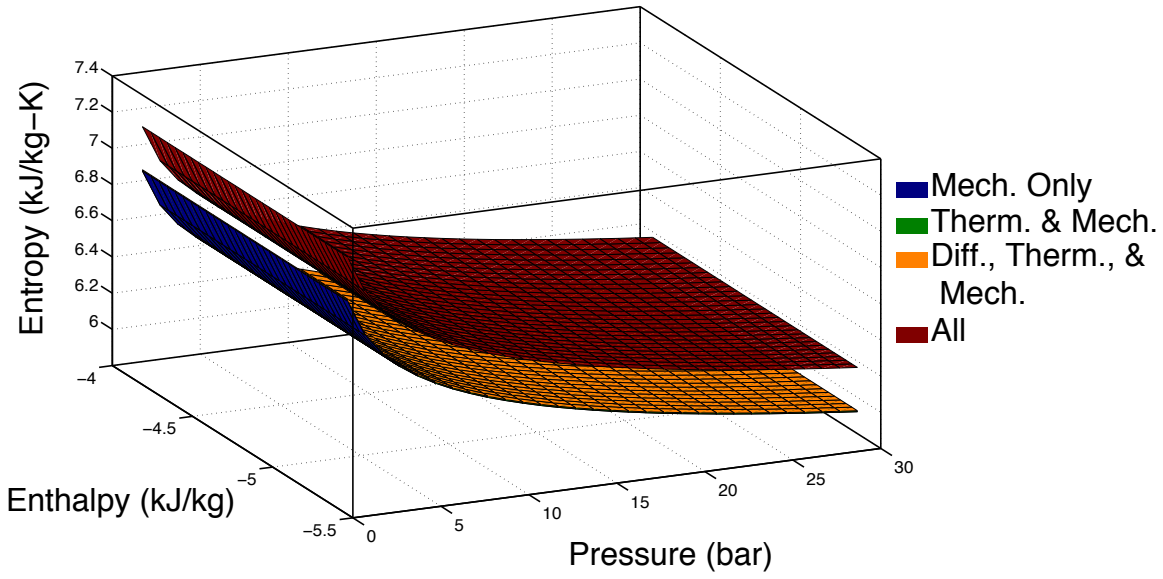

(a) Equilibration surfaces for a methane-air system

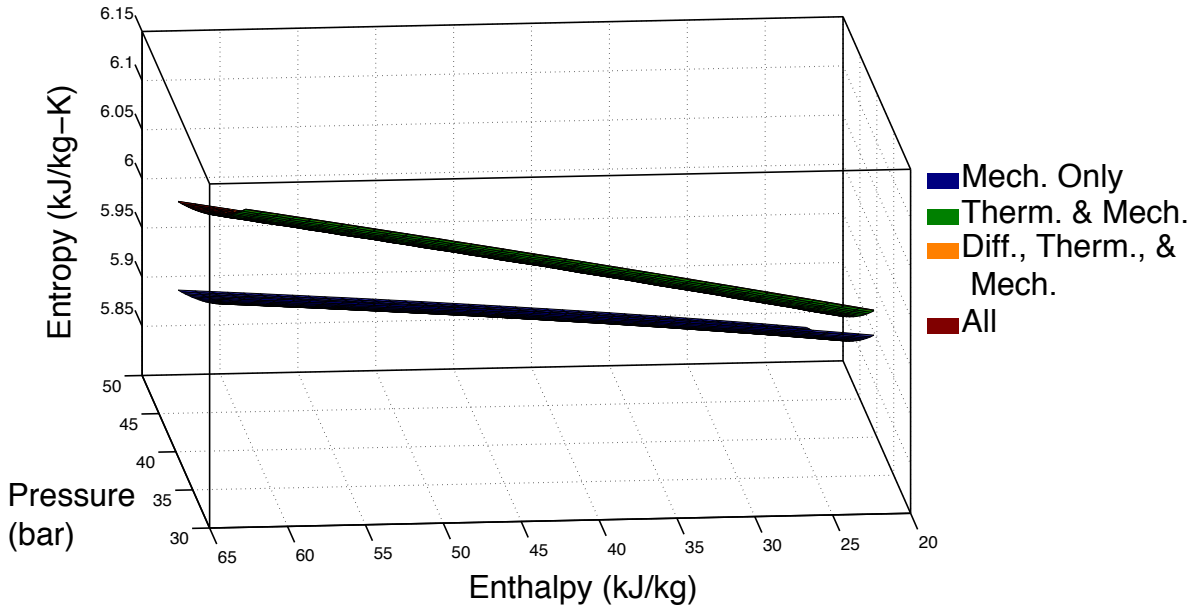

(b) Equilibration surfaces for a combustion products-air system

Figure 4.6: Equilibration surfaces.

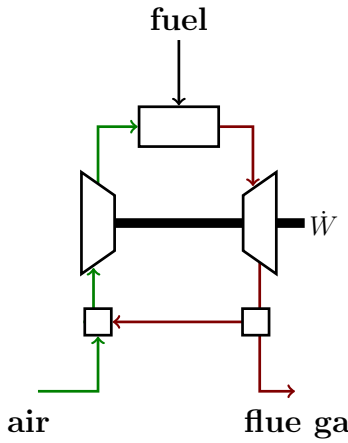

Figure 4.7: A simple-cycle gas-turbine engine with exhaust gas recirculation ($MCBTS$).

flow splitter. Exhaust gas recirculation (EGR) is used to recuperate thermal energy that would otherwise be rejected from the engine. By providing moderator in the burner, it can also allow the system to go to a higher pressure. A hotter, more pressurized system could result in less entropy generation in the burner.

Consider the following thought experiment: Instead of mixing the exhaust gas and the air, let the exhaust gas travel in a separate stream, but through an exact parallel set of devices as the air, and at matched pressure. At every point in the cycle, the recirculated gas is maintained in thermal (but not diffusive) equilibrium with the air. This is conceptualized by considering differential stages of compression and perfect, isobaric heat exchange between the recirculated gas and air streams, as depicted in Fig. 4.8.[10] The air, fuel, and recirculant flow rates, as well as the overall pressure ratio and turbine inlet temperature, are the same for the mixed-fluid and concentric-cycle systems shown in Figs. 4.7 and 4.8.

In order to be a useful model of the energy efficiency of the original mixed system in Fig. 4.7, the concentric system in Fig. 4.8 and the original mixed system have to have the same work transfers. This is shown to be true below.

[10]Because the concentric cycle is by intention ideal, it would have perfect, isobaric heat exchange even if the original mixed cycle did not.

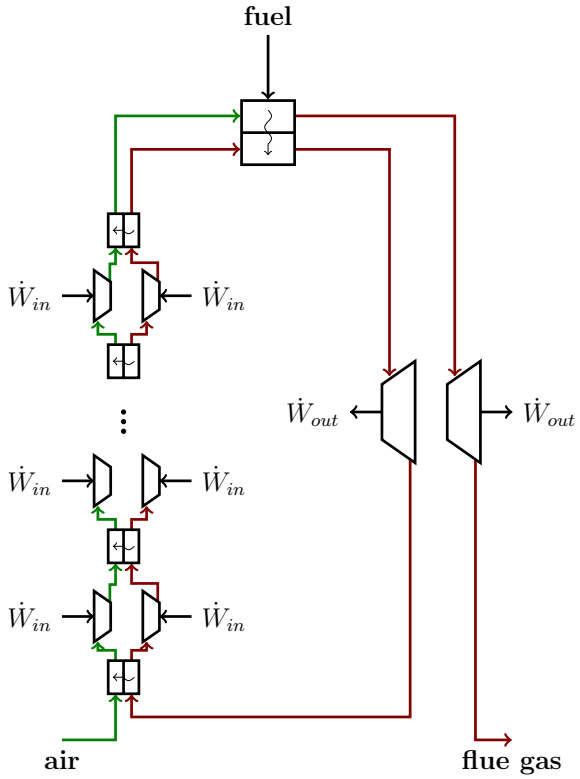

Figure 4.8: Concentric version of an exhaust gas recirculation system.

The exhaust gas is assumed to have been fully oxidized such that no significant further chemical reaction occurs upon mixing with fresh air. Both the exhaust and air are reasonably modeled as ideal gases, in which case there is no excess volume of mixing and

$$\dot{m}_m v_m = \dot{m}_a v_a + \dot{m}_r v_r \tag{4.4}$$

where subscript m is the mixed exhaust-air system, a is the air, and r is the recirculating exhaust gas.

The mixed system of air and recirculating flue gas can be thought of as an ideal solution obeying the Amagat model of two fluids with matching temperature and pressure undergoing compression, and with identical polytropic efficiencies. With unspecified polytropic efficiency eta_m, the work required to compress the original mixed system is written

$$\dot{W}_{comp,m} = \dot{m}_m \int_{P_{low}}^{P_{high}} \frac{v_m dP}{\eta_{comp,m}}$$

Using the ideal solution result in Eq. 4.4 to substitute in for the mixed volume gives

$$\dot{W}_{comp,m} = \int_{P_{low}}^{P_{high}} \frac{\dot{m}_a v_a + \dot{m}_r v_r}{\eta_{comp,m}} dP$$

Splitting the fraction apart into separate terms for each fluid gives

$$\dot{W}_{comp,m} = \int_{P_{low}}^{P_{high}} \frac{\dot{m}_a v_a}{\eta_{comp,m}} dP + \int_{P_{low}}^{P_{high}} \frac{\dot{m}_r v_r}{\eta_{comp,m}} dP \tag{4.5}$$

The polytropic efficiency η_m can be chosen such that

$$\frac{\dot{m}_a v_a}{\eta_{comp,m}} dP + \frac{\dot{m}_r v_r}{\eta_{comp,m}} dP = \frac{\dot{m}_a v_a}{\eta_{comp,a}} dP + \frac{\dot{m}_r v_r}{\eta_{comp,r}} dP \tag{4.6}$$

Combining Eqns. 4.5 and 4.6 gives

$$\dot{W}_{comp,m} = \int_{P_{low}}^{P_{high}} \frac{\dot{m}_a v_a}{\eta_{comp,a}} dP + \int_{P_{low}}^{P_{high}} \frac{\dot{m}_r v_r}{\eta_{comp,r}} dP$$

This expression on the right is exactly the expression for the work in the concentric representation of the recirculating flue gas system. Therefore,

$$\dot{W}_{comp,m} = \dot{W}_{comp,concentric}$$

The concentric cycle representation of the system requires the exact same work transfers as the mixed system.

The concentric-cycle version of the EGR system allows for the separate consideration of thermal and diffusive equilibration processes. The concentric cycle in Fig. 4.8 is equivalent to the mixed cycle in Fig. 4.7 except that the concentric cycle does not suffer entropy generation due to diffusion. Thus, analysis of the concentric cycle representation of EGR gives an upper bound on the efficiency of the actual EGR system. By serving as an upper bound on efficiency, concentric-cycle analysis can help identify which types of matter transfer have the potential to improve engine efficiency and which can be ruled out immediately as non-optimal. In particular, if the concentric-cycle version of a proposed system cannot improve upon the equivalent system without the matter transfer in question, then there is no merit in using that matter transfer.

4.4.3 Concentric Cycle Summary

The concentric cycle analysis is an alternative to the attractor methodology in circumstances involving changes to the system matter transfers and therefore exhaust composition. It is used to determine whether a proposed matter transfer could, in the best case, be beneficial over the equivalent system without matter transfer (but allowing heat transfer). When the proposed matter transfer is not beneficial, its flow

rate is set to zero length and the associated mixer and splitter are removed from the candidate architecture during the optimization process. Thus, the concentric cycle is a binary tool in determining optimal device lengths: it can only determine whether a mixer and splitter should be removed or whether they should stay. If the devices should stay, then further analysis will be required to determine the splitter length (mass flow rate).

Chapter 5

Matter Transfers in a Simple-Cycle Gas-Turbine Engine

The next goal is to determine how the optimal work-and-heat regenerative architecture changes when degrees of freedom associated with internal and external passive matter transfers are opened. Before doing this, it will be instructive to examine the matter transfers possible with a simple-cycle gas-turbine engine, as show in Fig. 5.1. This engine only has stages of compression, burning, and expansion. Studying the simple cycle will allow for isolating the thermodynamic principles and irreversibility trade-offs involved in making decisions about matter transfer. Chapter 6 will then return to the optimal heat-and-work regenerative engine to see how the principles developed in this chapter can be applied.

5.1 Model Problems

Matter transfers are classified as being of one of three different modes: feed-forward, feed-back, or external. These are briefly described below and fully explored in Sec. 5.2-5.4.

I. Feed-Forward Transfer

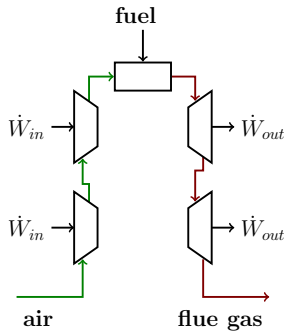

Figure 5.1: Simple-cycle gas-turbine engine, shown with 2 stages of compression and expansion

A fraction of the main fluid stream can be split off, bypass some number of subsequent devices, and then rejoin the main stream. This occurs in blade cooling, which is used to maintain the turbine blades at an acceptable temperature. To do this, some post-compressor air bypasses the combustor and mixes with the hot fluid entering a turbine stage.

II. Feed-Back Transfer

Some matter can be split off from the main fluid stream and recirculated to enter an earlier device in the sequence and remixed. This occurs in flue gas recirculation in which hot post-combustion products are fed backward into the colder pre-combustion reactant stream in an effort to improve performance or reduce emissions.

III. External Transfer

For a fixed flow rate of fuel, the relative proportion of air can be changed to alter performance. Additionally, any non-exergetic fluid can be added. (An exergetic fluid would be counted as part of the fuel stream). For an air environment (dead state) at 100% relative humidity, this allows inclusion of liquid water. Water can be injected into the air stream at various points with the potential to improve efficiency and/or power.

In this analysis matter transfers proceed directly from a splitter to a mixer. There are no intermediate processes that alter the enthalpy of the matter. Pumps, diffusers, and heat exchangers are thus excluded, while adiabatic valves are allowed. Such intermediate processes could allow for combined cycles, chemical separations, and other complex phenomena. While these processes are important to eventually analyze and incorporate into the optimal engine architecture, they are distinctly more complicated than passive matter transfers. Thus, matter with intermediate processes will be treated as a further degree of freedom beyond passive matter transfer, and is not in the scope of this dissertation.

Without intermediate processes, momentum conservation requires matter to flow in the direction of a favorable pressure gradient. This will be modeled by an adiabatic and isenthalpic throttling processs. When examining idealized upper bounds on efficiency, matter will be treated as transferring between points of equal pressure. Fuel is assumed to enter the system from a pressurized line. For convenience of illustration, the fluid is required to maintain a single phase at all times. This restricts the quantity of liquid water introduced to an amount less than saturation. This requirement simplifies the analysis without changing or restricting the fundamental conclusions. Chapter 6 will readdress the impact of this restriction.

The analysis of each model problem has three subsections: 1) conceptual analysis, 2) numerical illustration, and 3) thermodynamic interpretation. Because this chapter is examining a simple-cycle gas-turbine engine, it will focus on developing the optimal process sequence and delay a discussion of optimal process lengths until the reintegration into the optimal architecture with work-and-heat regeneration in Ch. 6.

5.1.1 Numerical Models

The model problem analyses do not assume any specific fuel type, and the results are applicable to any chemical fuel-air system. For illustration however, fuels are chosen that are used commonly in gas turbine engines: natural gas (for stationary power)

Property	Specification
Air Composition	20.50% O_2, 76.99% N_2, 0.036% CO_2, 1.56% H_2O, 0.91%Ar
Temperature	298 K
Pressure	1.01325 bar

Table 5.1: The environmental dead state used for numerical and graphical results

and kerosene (for aircraft). Both of these fuels span a range of compositions, so reasonable surrogates are needed. Natural gas is modeled as $90\%CH_4, 5\%C_2H_6, 5\%N_2$. For kerosene, the choice is made to use 100% dodecane ($C_{12}H_{26}$). While this is overly simplistic, it makes a 2nd-law analysis possible and still provides a large liquid molecule as a contrast to a small gaseous fuel. The simulations are completed in Matlab [112]. The thermodynamic software package Cantera is used for the gaseous properties in the simulations [113]. It employs the database GRI-Mech 3.0 for evaluating the state of the working fluid [133]. Properties for dodecane are taken from the NIST Chemistry WebBook [134, 135]. Properties for the environment are in Table 5.1. Device constraint values are those from Table 4.2.

5.2 Feed-Forward Matter Transfer

5.2.1 Conceptual Analysis

A feed-forward matter transfer involves splitting off some of the main fluid stream, bypassing a set of downstream devices, and then remixing with the main stream at a later point. Doing this passively in the direction of a favorable pressure gradient limits the possible locations that such a transfer could occur. In a simple-cycle gas-turbine engine, as shown in Fig. 5.1, a feed-forward transfer could occur either from a high-pressure product gas to a lower-pressure product gas, or from the compression-side air to the expansion-side product gas.

Feed-Forward within the Expansion Process

Feed-forward within the expansion process would involve some of the matter skipping a turbine stage, going through a throttle, and remixing with the main flow, as shown in Fig. 5.2b. This will be compared to the equivalent system without a feed-forward transfer, shown in Fig. 5.2a. The two systems have the same inlet state, same overall mass flow rate, and same final pressure. All devices in both systems are adiabatic. The systems in Fig. 5.2 have no external transfers of heat. Therefore, the results

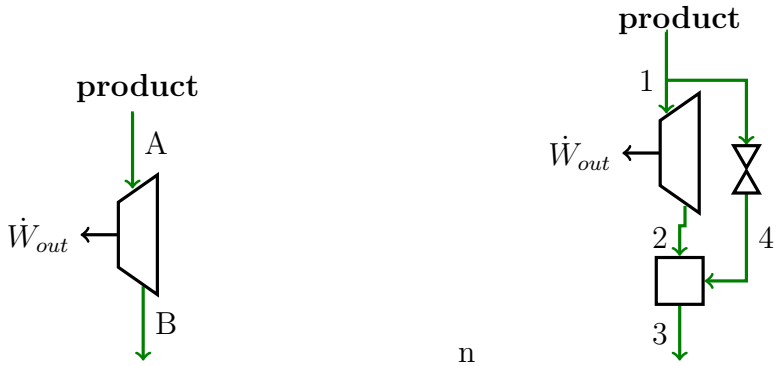

(a) No feed-forward transfer n (b) Feed-forward transfer

Figure 5.2: Comparison of expansion with and without a feed-forward matter transfer.

from Ch. 4 hold, and the irreversibility of each systems is monotonic with its outlet entropy. Whichever system has the higher outlet entropy is more irreversible. This is equivalent to determining which system has the higher outlet attractor entropy. States 2 and B are at the same pressure and composition. Removing the pressure and chemical terms, Gibbs equation for the attractor can be integrated between these two states.

$$T_{eq} ds_{eq} = dh - \cancel{v_{eq} dP} \tag{5.1}$$

$$\int_{s_{B,eq}}^{s_{2,eq}} T_{eq} ds_{eq} = \int_{h_B}^{h_2} dh \tag{5.2}$$

Because temperature is always positive, whichever state (2 or B) has the higher enthalpy will also have the higher attractor entropy. Now all that is left is to determine which system has the higher outlet enthalpy.[1]

The outlet enthalpies are compared using mass and energy balances. An energy balance around the system with a feed-forward system is written

$$\dot{m}_3 h_3 = \dot{m}_1 h_1 - \dot{m}_2(h_1 - h_2)$$

Subtracting this equation from the energy at the outlet of the other system, $\dot{m}_B h_B$, gives

$$\dot{m}_B h_B - \dot{m}_3 h_3 = \dot{m}_B h_B - \dot{m}_1 h_1 + \dot{m}_2(h_1 - h_2)$$

Here $h_B = h_2$ and $\dot{m}_1 = \dot{m}_3 = \dot{m}_B$. Therefore the above equation can be rewritten in terms of only States 1, 2, and 3.

$$\begin{aligned}\dot{m}_3(h_2 - h_3) &= \dot{m}_3 h_2 - \dot{m}_3 h_1 + \dot{m}_2(h_1 - h_2)\\ \dot{m}_3(h_2 - h_3) &= (\dot{m}_2 - \dot{m}_3)(h_1 - h_2)\end{aligned}$$

$\dot{m}_2$ is a fraction of $\dot{m}_3$ so the difference in masses on the right hand side is always negative. The difference in enthalpies on the right hand side is always positive because work is extracted between States 1 and 2. The product of a negative term with a positive term is always negative. Therefore,

$$\begin{aligned}\dot{m}_3(h_2 - h_3) &< 0\\ \dot{m}_3(h_B - h_3) &< 0\end{aligned}$$

This indicates that the enthalpy, and therefore attractor entropy, are greater exiting

[1]This procedure of canceling constant terms in Gibbs and integrating between two states to relate their entropies to their enthalpies will be used repeatedly in this chapter. Going forward, it will simply be referred to as "integrating Gibbs between two states."

the system with the feed-forward transfer. The feed-forward system is more irreversible. The throttling produces entropy and removes some of the ability of the engine to produce work. It also adds entropy of mixing when the relatively hotter bypass gas is remixed with the colder main stream gas.

The simple cycle only allows compression, followed by burning, followed by expansion. Thus, the only possible additional devices after States $2/B$ before exhaust are more stages of expansion. Regardless of further expansion, a higher entropy at State B than State 2 requires a higher outlet entropy for the feed-forward system than for the simple system. Therefore, feed-forward transfers within the expansion side of a simple-cycle gas turbine engine always add unnecessary irreversibility to the system.

Feed-Forward from Compression to Expansion

Feed-forward from the compression process to the expansion process can either be done at constant pressure (in an ideal case), or with a throttle to a lower pressure stage. As just discussed, the throttled case destroys exergy and increases the entropy of mixing due to increased thermal differences. The isobaric transfer represents an idealized efficiency upper bound on a realistic case with friction and inevitable pressure drops. However, these real effects only need to be considered if the upper-bound system proves more efficient than an equivalent system with no feed-forward transfer. If the best-case frictionless system with feed-forward transfer is still of lower efficiency than a simple-cycle, then feed-forward matter transfers can be eliminated without having to delve into more detailed models.

Figure 5.3 shows a comparison between a simple-cycle without feed-forward transfers to one with a feed-forward transfer at a constant-but-arbitrary pressure. Requiring the two different systems to have all of the same operating parameters is impossible. One parameter has to remain free between the two systems while all others are held fixed for a reasonable comparison. The choice of the free parameter is arbitrary, and the analysis yields the same results regardless of choice. Here the choice

is made to allow the higher-pressure turbo-machinery to vary in pressure between the two systems. The fuel flow rates, turbo-machinery polytropic efficiencies, and turbine inlet temperatures are held constant between the two systems. Additionally, the inlet air flow rate is the same at States 1 and A.

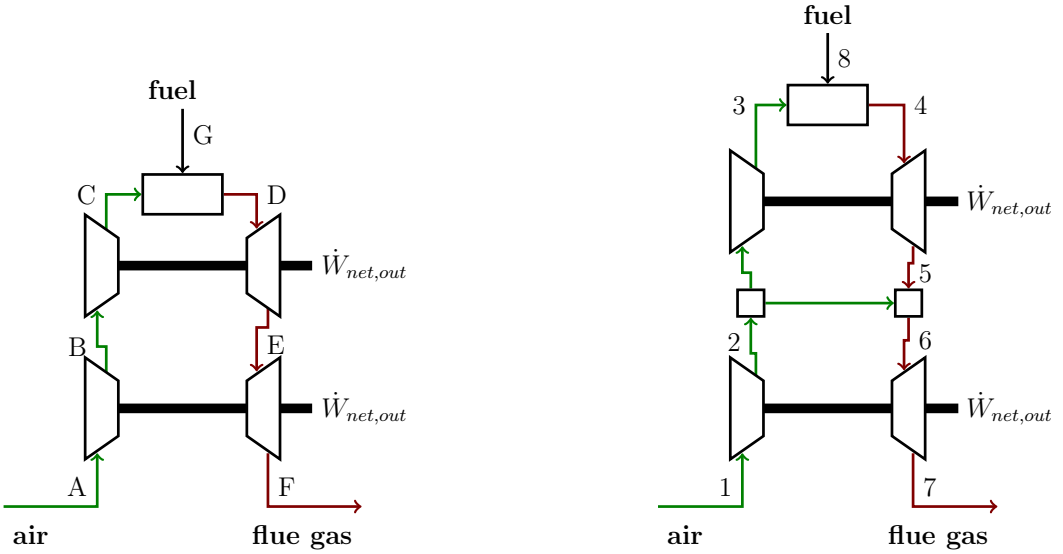

(a) Simple-cycle without feed-forward

(b) Cycle with feed-forward air

Figure 5.3: Insertion of a feed-forward transfer.

The intensive thermodynamic state is identical for the two systems up through the entrance to the burner. In the feed-forward case, there is a smaller mass flow rate of air entering the burner. This results in a richer mixture in the burner than in the simple-cycle case. Assuming that the simple-cycle burner operates lean, the feed-forward burner with less air is closer to stoichiometric, and therefore has a higher adiabatic flame temperature. However, the turbine inlet temperature is constrained ($T_D = T_4$), so the maximum system pressure must be lower in the feed-forward case than in the simple-cycle case. Therefore $T_4 = T_D, P_4 < P_D$.

State E is an arbitrary point along the expansion pathway in the simple cycle. It is chosen to match the pressure of the feed-forward transfer stream. Therefore, $P_E = P_5 = P_6$. It is unclear how T_6 relates to T_E. Because the combustion was done in a lean environment in both systems, it is reasonable to assume that State 6

and State E have negligible differences in composition (near-complete combustion). Since they are also at the same pressure, by integrating Gibbs between the two states, whichever state is at lower enthalpy also has lower entropy. Since the final expansion to atmospheric exhaust is the same for the two systems, whichever has a lower entropy entering the turbine will also have a lower entropy exiting the turbine, and therefore a higher system efficiency.

The next step is to determine the relationship between the enthalpy at States E and 6. The higher enthalpy will be the state of higher entropy and therefore associated with lower overall system efficiency.

Notation

Subscripts will refer to the states numbered and lettered in the accompanying diagrams. The symbol $\widetilde{v}$ refers to the specific volume of the simple cycle as opposed to v, the specific volume of the feed-forward cycle. Specific volumes will also be subscripted by either a for air or p for products to clarify composition.

Derivations

First an energy balance is done around the high-pressure turbo-machinery and burners for the two systems shown in Fig 5.3.

$$\dot{m}_E h_E = \dot{m}_B h_B + \dot{W}_{BC} + \dot{m}_G h_G - \dot{W}_{DE}$$
$$\dot{m}_6 h_6 = \dot{m}_2 h_2 + \dot{W}_{23} + \dot{m}_8 h_8 - \dot{W}_{45}$$

States 2 and B are the same. Additionally, it is assumed that the fuel entering the burner is over-pressurized to the same state regardless of the air pressure such that $h_G = h_8$. Similarly, $\dot{m}_1 = \dot{m}_2 = \dot{m}_A = \dot{m}_B = \dot{m}_C$ and $\dot{m}_6 = \dot{m}_7 = \dot{m}_D = \dot{m}_E = \dot{m}_F$ and $h_2 = h_B$. Subtracting the second equation from the first yields

$$\dot{m}_6(h_E - h_6) = \dot{W}_{BC} - \dot{W}_{23} + \dot{W}_{45} - \dot{W}_{DE}$$

Writing out the work terms as integrals gives

$$\dot{m}_6(h_E - h_6) = \int_{P_B}^{P_C} \dot{m}_B \widetilde{v_a} \frac{1}{\eta_C} dP - \int_{P_2}^{P_3} \dot{m}_3 v_a \frac{1}{\eta_C} dP + \int_{P_5}^{P_4} \dot{m}_4 v_p \eta_T dP - \int_{P_E}^{P_D} \dot{m}_E \widetilde{v_p} \eta_T dP$$

Noting that $P_2 = P_5 = P_B = P_E$, $P_3 = P_4$, and $P_C = P_D$, the above equation can be rewritten in terms of fewer unique pressures. Additionally, $v_a = \widetilde{v_a}$ because the intensive states in the two systems are identical during compression.

$$\dot{m}_6(h_E - h_6) = \int_{P_2}^{P_C} \dot{m}_1 v_a \frac{1}{\eta_C} dP - \int_{P_2}^{P_3} \dot{m}_3 v_a \frac{1}{\eta_C} dP + \int_{P_2}^{P_3} \dot{m}_5 v_p \eta_T dP - \int_{P_2}^{P_C} \dot{m}_6 \widetilde{v_p} \eta_T dP$$

The integrals are rearranged such that some go over the pressure bounds of the feed-forward system (P_2 to P_3) and the remaining ones cover the span between the maximum pressure of the feed-forward system and the simple system (P_3 to P_C).

$$\dot{m}_6(h_E - h_6) = \int_{P_2}^{P_3} (\dot{m}_1 - \dot{m}_3) v_a \frac{1}{\eta_C} dP + \int_{P_3}^{P_C} \dot{m}_1 v_a \frac{1}{\eta_C} dP + \int_{P_2}^{P_3} (\dot{m}_5 v_p - \dot{m}_6 \widetilde{v_p}) \eta_T dP - \int_{P_3}^{P_C} \dot{m}_6 \widetilde{v_p} \eta_T dP$$

Combining the integrals with the same bounds gives

$$\dot{m}_6(h_E - h_6) = \int_{P_2}^{P_3} \left((\dot{m}_1 - \dot{m}_3) v_a \frac{1}{\eta_C} + \dot{m}_5 v_p \eta_T - \dot{m}_6 \widetilde{v_p} \eta_T \right) dP + \int_{P_3}^{P_C} (\dot{m}_1 v_a \frac{1}{\eta_C} - \dot{m}_6 \widetilde{v_p} \eta_T) dP \quad (5.3)$$

The sign of this result depends upon the choice of fuel, the turbo-machinery polytropic efficiencies, and the equation of state. If the turbo-machinery is isentropic ($\eta_T = \eta_C = 1$),

$$\dot{m}_6(h_E - h_6) = \int_{P_2}^{P_3} \left((\dot{m}_1 - \dot{m}_3) v_a + \dot{m}_5 v_p - \dot{m}_6 \widetilde{v_p} \right) dP + \int_{P_3}^{P_C} (\dot{m}_1 v_a - \dot{m}_6 \widetilde{v_p}) dP \quad (5.4)$$

In this case, the second integral is negative because $\dot{m}_1 < \dot{m}_6$ and $v_a < \widetilde{v_p}$.[2] A system of air and combustion products can reasonably be modeled as an ideal solution. Using the ideal solution assumption the first integral is zero:

$$\dot{m}_6 \widetilde{v_{p6}} = (\dot{m}_1 - \dot{m}_3)\widetilde{v_a} + \dot{m}_5 \widetilde{v_p}$$

Thus, with isentropic turbo-machinery and ideal solutions, $\dot{m}_6(h_E - h_6) < 0$, which implies that $h_F < h_7$, which indicates that the simple cycle is more efficient than the feed-forward cycle. For other assumptions about property methods and turbo-machinery, Eq. 5.3 can be evaluated to determine relative efficiencies.

5.2.2 Numerical Illustration

Variations with Bleed Mass Fraction

As more air is fed-forward, the burner sees a relatively richer mixture. Given that gas-turbine systems operate lean, this results in a burner operating closer to stoichiometric with less air to act as moderator to absorb the exothermicity of the combustion reaction. As a result, for a fixed turbine inlet temperature, an increased feed-forward mass flow rate requires a decreased peak pressure. Figure 5.4 shows this phenomenon for systems fueled by natural gas and dodecane. For both systems, the feed-forward transfer is done at 5 bar.

Figure 5.5 shows the consequences of this mass transfer and lowered system pressure on both systems. These plots indicate that increased feed-forward mass leads to a significant decrease in system efficiency and an increase in exhaust exergy. The increased exhaust exergy is a result of the decreased pressure ratio associated with feed-forward transfer and the resulting reduction in the system's ability to extract work.

[2] This is only true when $\frac{\partial v}{\partial T}|_P > 0$, which holds for most conceivable air-fuel mixtures in a gas-turbine system.

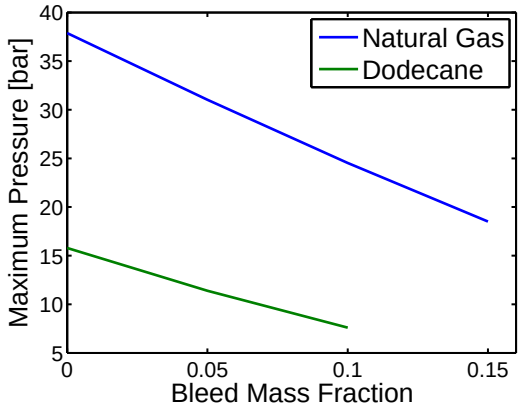

Figure 5.4: Effect of feed-forward transfer mass flow rate on system pressure.

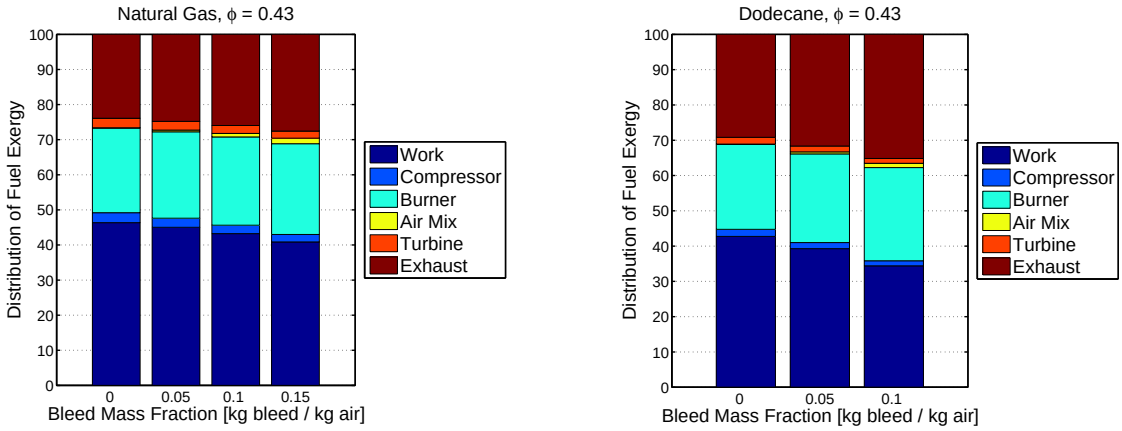

Figure 5.5: Exergy distribution for common air-fuel systems. Feed-forward transfer at 5 bar. This graph shows the same systems as Fig. 5.4

Variations with Bleed Pressure

Figure 5.6a shows that the feed-forward transfer pressure has only a minimal effect on overall exergy efficiency. Any choice of transfer pressure results in a lower efficiency than the equivalent system without any feed-forward transfer, as indicated by the dashed red line. Figure 5.6b breaks down these efficiency changes by the individual destruction terms. Once the quantity of bleed mass is set, there is a unique maximum system pressure, regardless of bleed pressure. Thus, the destruction associated with fuel mixing and combustion is independent of bleed pressure, and not plotted here. The exergy destruction in the compressor increases with bleed pressure because more mass is being compressed. The turbine destruction also increases with bleed pressure because a larger portion of the expansion is done with a colder fluid. Meanwhile the air mixer shows a slight improvement with bleed pressure due to the decreased temperature difference in the mixer as the bleed air is further compressed before mixing. The trade-off between the turbo-machinery effects and the mixing effect lead to the ultimate variation in exergy efficiency with transfer pressure.

5.2.3 Thermodynamic Interpretation

It is unsurprising that the dodecane and natural gas systems behave similarly with a feed-forward transfer. They have similar combustion products, dominated by water and carbon dioxide.[3] Because water and carbon dioxide are triatomic, the product streams from these reactions have higher specific heats than the air stream, which is dominated by diatomic nitrogen and oxygen. Similar compositions also means the product gases of the various combustion reactions have similar molar masses. Given that these systems are well-modeled as ideal solutions of ideal gases, similar specific heats and molar masses indicate that the gases would have similar turbo-machinery/mixing trade-offs. Thus, using standard fuels and an understanding of

[3]While this is somewhat less true when comparing very rich to very lean combustion systems, it is reasonable to assume that gas turbines will operate lean such that minimal chemical exergy is wasted in the flue gas.

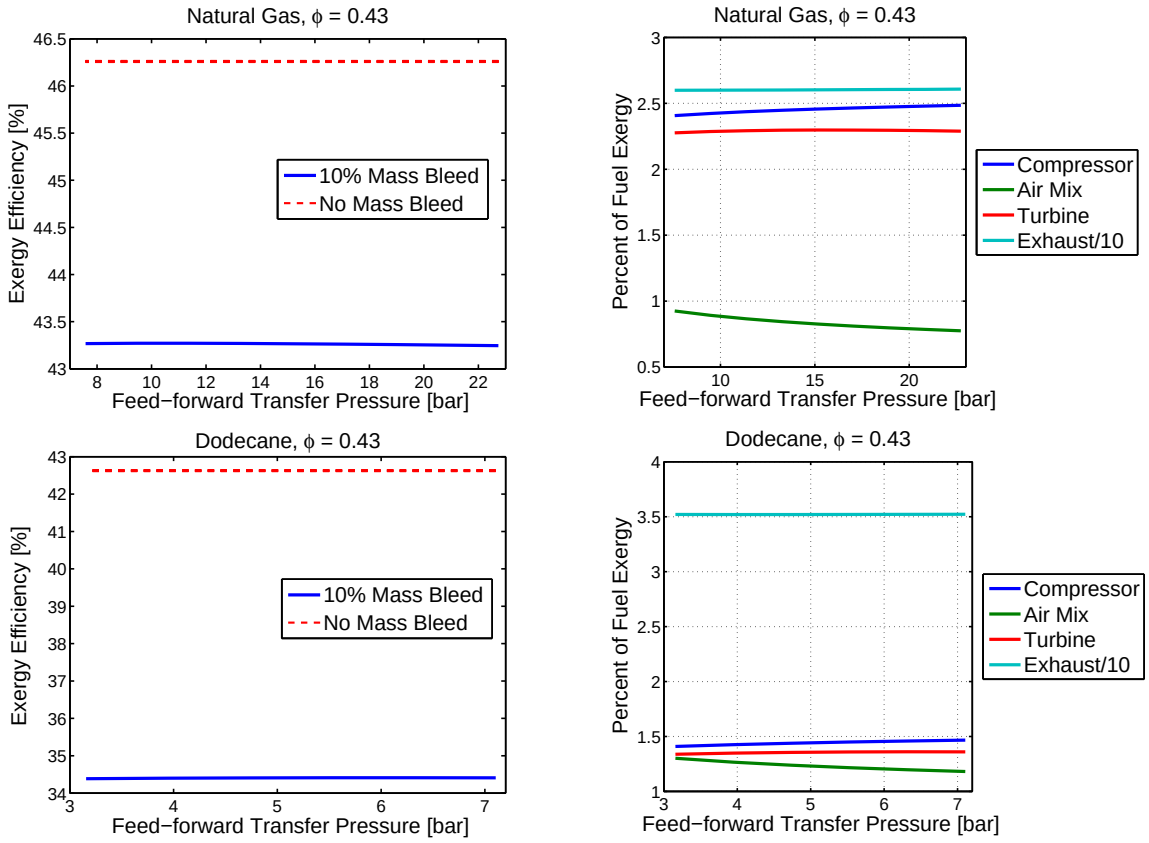

Figure 5.6: Effect of feed-forward bleed pressure on exergy efficiency. Bleed mass flow rate is 10% of overall air mass flow rate.

their common properties, it is shown that feed-forward air transfer in a gas-turbine engine will result in decreased efficiency for practical applications.

5.3 Feed-Back Matter Transfer

5.3.1 Conceptual Analysis

Feed-back matter transfer involves splitting off some of the main fluid stream, sending it upstream to an earlier device in the flow and remixing it in. In a simple-cycle gas-turbine engine this could happen either by throttling some of the compressed air and mixing it in with lower-pressure air or by recirculating some of the product gas in with the air during compression. Both of these options are examined below.

Feed-Back within the Compression Process

Figure 5.7 shows the simple cycle and a variant with a feed-back transfer within the compression process. As was done with the feed-forward comparison, the two systems have the same inlet mass flow rates, turbine inlet temperature, and polytropic efficiencies. The pressure ratio varies between the two systems in order to accommodate these fixed variables while allowing for the architectural change.

Neither system has external heat transfer, so again comparing outlet entropies is sufficient to determine relative efficiencies. Because they have the same atomic composition, they will have the same equilibrium product composition. With the same product composition, turbine inlet temperature, and mass flow rates, whichever system can expand through a larger pressure ratio will have the lower final entropy. Individually numbered and lettered states in Fig. 5.7 are now compared to assess which system has a higher pressure ratio.

States 1 and A are equal, but State 2 is hotter than State 1 due to the mixing in of hot air from State 4. Thus the lower-pressure compressor in the feed-back case sees a hotter, less dense fluid than its counterpart in the simple system. This results

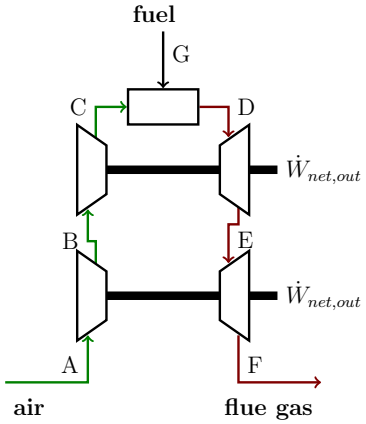

(a) Gas-turbine system without feed-back.

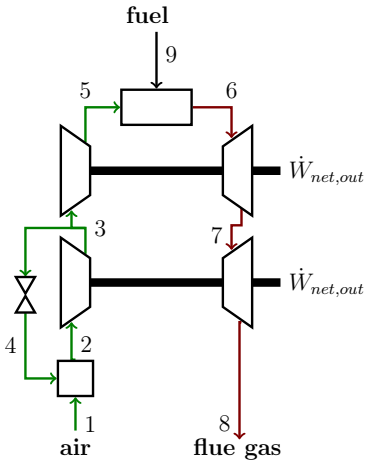

(b) Gas-turbine system with feed-back within compression.

Figure 5.7: Comparison of compression with and without a feed-back matter transfer.

in more compressor work per unit pressure change in the feed-back case. State B is an arbitrary point along the compression chain, so its pressure is chosen to match that of State 3. This results in State 3 being hotter than State B. If the two systems went to the same peak pressure, then $P_C = P_5$, $T_C < T_5$, and $h_C < h_5$. However, in order to hit the fixed turbine inlet temperature, the burner air inlet must have the same enthalpy in both systems. The only way this is possible is if $P_5 < P_C$. Thus, the system with the feed-back transfer has a lower maximum pressure than the system without a feed-back loop.

This analysis made no assumptions on the conditions at the inlet State $1/A$ or the value of the pressure at State B. As such, this ambiguous architecture covers the scope of all possible compression-side feed-back systems in a simple-cycle gas turbine engine. Thus, this form of feed-back transfer is identified as non-optimal.

Feed-Back from Expansion to Compression

Some of the partially expanded products could be recirculated either to the equivalently-pressurized compressor stage or to a lower-pressure compressor stage. The ideal

situation would be to transfer isobarically, thereby minimizing throttling and the resulting hot recompression. This case is examined now.

First, the system pressure of the open feed-back matter system is compared to that of the equivalent system without feed-back. A numerical simulation using standard hydrocarbon fuels then demonstrates that the feed-back system has a lower peak pressure. Finally, the concentric cycle analysis is used to evaluate the efficacy of this lower-pressure feed-back system.

Derivation

The first task is to determine how the inclusion of feed-back changes the overall system pressure in Fig 5.8b relative to the original system pressure in Fig 5.8a. A set of energy balances are used to define the enthalpy at the turbine inlet state as a function of known parameters. For the original simple system, an energy balance around the combustor gives

$$\dot{m}_D h_D = \dot{m}_G h_G + \dot{m}_C h_C$$

Writing out the enthalpy at State C in terms of the compression work from State B results in

$$m_D h_D = \dot{m}_G h_G + \dot{m}_B (h_B + \int_{P_B}^{P_D} \frac{v}{\eta_C} dP) \tag{5.5}$$

Only the first law of thermodynamics has been used here, so this statement is true for arbitrary fluids and parameters. For the current problem, the fuel and air flow rates and composition, the intermediate pressure (P_B), the turbine inlet temperature, and the polytropic efficiencies are fixed.[4] Any choice of P_D leads to a unique value of h_D. There is only one combination of h_D and P_D that will coincide with the turbine inlet temperature requirement for the given composition. Thus the above equation allows for the unique determination of the system pressure.

[4]The choice of parameters to fix is again arbitrary, but consistent with the decisions made in Sec. 5.2.1

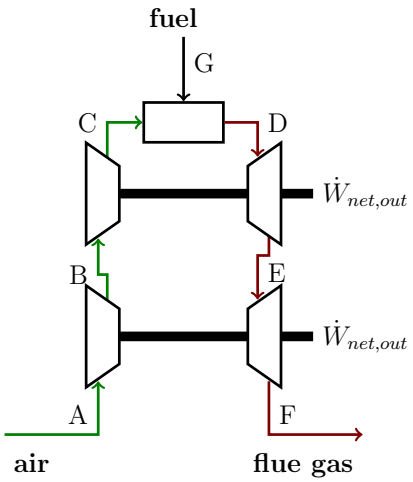

(a) Simple cycle without feed-back.

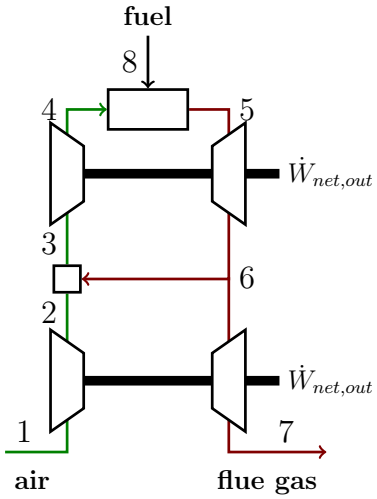

(b) Feed-back matter at arbitrary pressure.

Figure 5.8: Schematics for a simple-cycle gas-turbine engine without and with a feed-back transfer.

The same type of analysis is done for the feed-back system:

$$\begin{aligned}
\dot{m}_5 h_5 &= \dot{m}_8 h_8 + \dot{m}_4 h_4 \\
\dot{m}_5 h_5 &= \dot{m}_8 h_8 + \dot{m}_4 \left(h_3 + \int_{P_2}^{P_5} \frac{v}{\eta_C} dP \right) \\
\dot{m}_5 h_5 &= \dot{m}_8 h_8 + \dot{m}_2 h_2 + (\dot{m}_2 - \dot{m}_3) h_6 + \dot{m}_4 \int_{P_2}^{P_5} \frac{v}{\eta_C} dP \\
\dot{m}_5 h_5 &= \dot{m}_8 h_8 + \dot{m}_2 h_2 + (\dot{m}_2 - \dot{m}_3) \left(h_5 - \int_{P_5}^{P_2} v \eta_T dP \right) \\
&\quad + \dot{m}_4 \int_{P_2}^{P_5} \frac{v}{\eta_C} dP \\
(\dot{m}_5 - \dot{m}_2 + \dot{m}_3) h_5 &= \dot{m}_8 h_8 + \dot{m}_3 h_2 - (\dot{m}_2 - \dot{m}_3) \int_{P_5}^{P_2} v \eta_T dP \\
&\quad + \dot{m}_4 \int_{P_2}^{P_5} \frac{v}{\eta_C} dP
\end{aligned} \tag{5.6}$$

Given the same fixed parameters as above, as well as the feed-back flow rate, this equation uniquely determines the system pressure.

Unfortunately Eqns. 5.5 and 5.6 are implicit and cannot be used to easily demonstrate the relative pressures between the simple and feed-back systems. However, this pressure comparison can be made numerically. The feed-back stream increases the temperature and the mass flow rate of the compressed fluid. This effect decreases the maximum system pressure before the products reach the fixed turbine inlet temperature. On the other hand, the feed-back stream increases the amount of moderator in the burner, which leads to an increase in the burner inlet temperature and therefore pressure. The trade-off between the increased compressor work and increased combustion moderator is the key in determining the change in system pressure upon introduction of feed-back. Thus the important thermodynamic properties to track are the specific volume of the compressed fluid and the specific heat of the recirculated fluid.

Figure 5.9 shows the effect of feed-back pressure on the overall system pressure

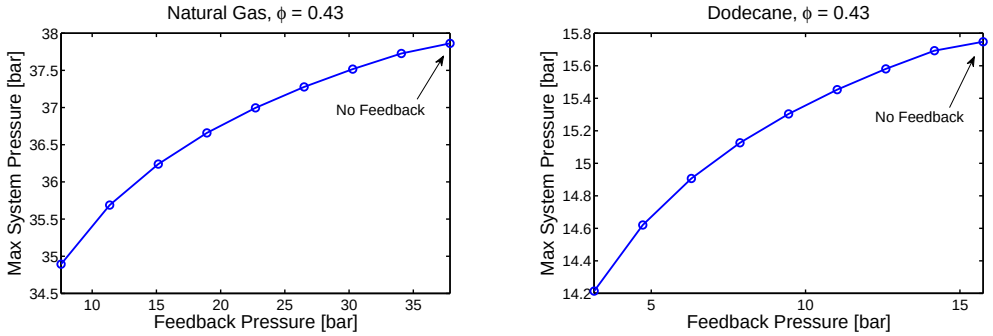

Figure 5.9: Effect of feed-back pressure on maximum system pressure. The highest pressure point (indicated with an arrow) is for a system with no feed-back. Feed-back flow rate = 10% of air flow rate.

for two different air-fuel systems. In both cases, increased feed-back pressure is associated with increased system pressure. However, the case without any feed-back always resulted in the highest system pressure. This indicates that the inclusion of a feed-back matter stream reduces the system pressure from what it would otherwise have been. This analysis is continued assuming that the introduction of feed-back matter reduces the pressure of the system. All that is left is to determine the efficacy in having such a lower-pressure system. To do this, a concentric-cycle analysis will be used.

Concentric-Cycle Analysis of a Feed-Back System

As detailed in Ch. 4, the concentric cycle is a conceptual tool that separates the thermal and diffusive effects of matter transfer. Instead of directly mixing in feed-back matter with the main air stream, it is compressed in parallel with the air, while maintaining perfect thermal contact. This is visualized as a series of differential compressors and heat exchangers, as shown in Fig. 5.10b. This allows all of the potential thermal benefits of the feed-back, but none of the inherent entropy generation due to mixing fluids of different composition. Because the concentric cycle removes the entropy of mixing, it serves as an upper bound on the possible efficiency of the openly mixed system. This upper bound will be compared to the simple cycle

with no feed-back, as depicted in Fig. 5.10a. Demonstrating that the simple cycle can achieve a higher efficiency than the concentric cycle indicates that no feed-back system will surpass the simple cycle in efficiency. The turbine inlet temperature, air

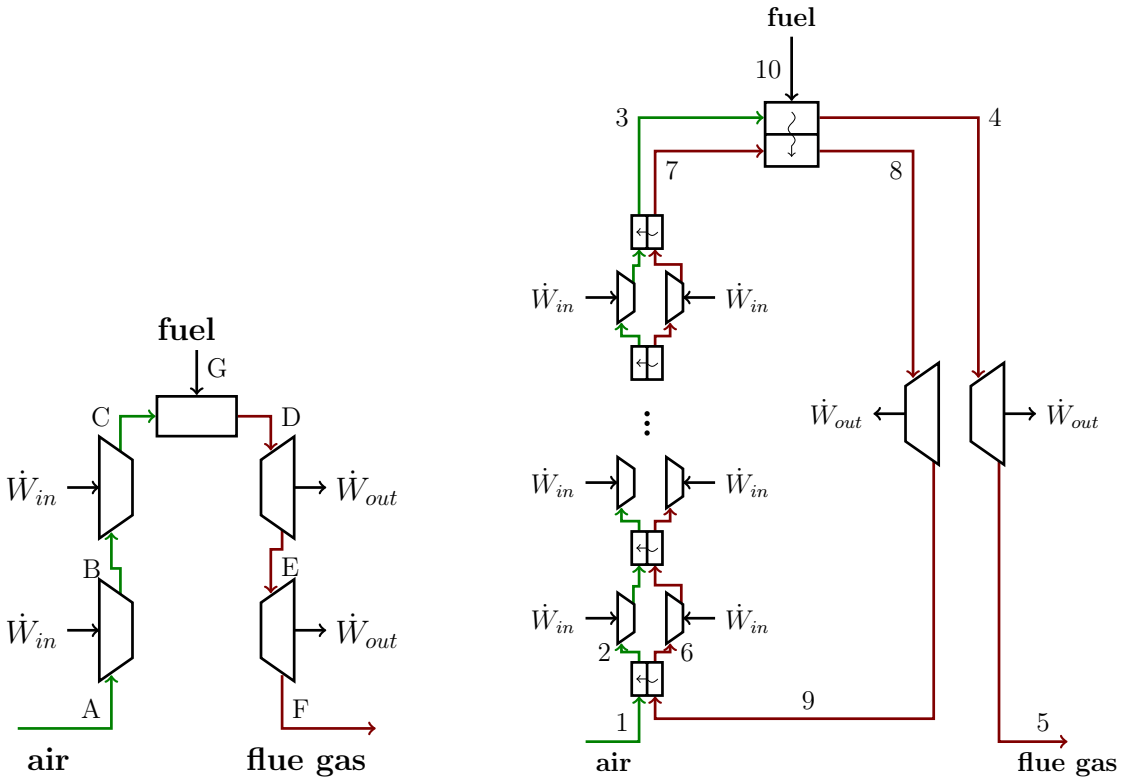

(a) Simple-cycle without feed-back (b) Concentric representation of feed-back

Figure 5.10: Mapping a cycle to its concentric representation.

and fuel flow rates, and polytropic efficiencies are again fixed between the simple cycle and the concentric cycle. The inlet pressure is arbitrary, but the same for the two systems. The burner/main heat exchanger and the differential heat exchangers are adiabatic, isobaric, and have no minimum temperature difference (other than $T_H > T_C$) between the hot and cold streams. The independent variable is the ratio of air mass flow rate in the outer cycle to feed-back mass flow rate. The dependent variable that changes in response to the independent variable is the maximum system

pressure.

As discussed above, the pressure in the concentric-cycle must be lower than that in the simple-cycle. It is assumed that the enthalpy at the combustor outlet does not change with small pressure changes (or the associated minor chemical composition change). This requires the enthalpy of the combustor outlet to be fixed because temperature is fixed.

Fluid states are referred to as numbered in Figs 5.10a and 5.10b. States 4 and D have the same temperature, but State 4 has a lower pressure. By the same argument made in Sec. 5.2.1, the integration of Gibbs between States 4 and D indicates that State 4 has a higher entropy than State D.

The pressure of State E is arbitrary and is chosen to match the pressure of State 4. State 4 is at the turbine inlet temperature whereas State E is the result of the partial expansion of the simple-cycle system at State D to the pressure of State 4, using the defined polytropic efficiency of the turbo-machinery. This state must have a temperature less than that of State 4.

$$T_E < T_4, P_E = P_4$$

Now State E and State 4 are both expanded through the same pressure ratio with the same turbo-machinery to reach States F and 5, respectively. This results in

$$T_F < T_5, P_F = P_5$$

Then by integration of Gibbs equation between States 5 and E,

$$s_F < s_5$$

Because the simple cycle and concentric cycle have the same inlet entropy and no transfers as heat to or from the environment, the more efficient system is the one that minimizes the exhaust entropy. Therefore, the simple-cycle is more efficient.

Any system with feed-back matter transfer can be mapped to a concentric-cycle version of that system that is of higher efficiency than the original system with an open transfer. And for any concentric cycle, there exists a simple cycle operating at the same main-stream equivalence ratio and turbine inlet temperature that is of higher efficiency than the concentric cycle. Therefore there is a simple cycle that is more efficient than any possible feed-back system operating at a lower pressure than the original system.

5.3.2 Numerical Analysis

Variations with Feed-Back Pressure

The concentric cycle analysis is confirmed by the results in Fig. 5.11. Feed-back transfers reduce the efficiency of the gas-turbine system. The lower the pressure of the feed-back transfer, the worse this effect is. Interestingly, the burner exergy destruction is lower in the presence of feed-back moderator than in its absence. The extreme states principle dictates that combustion entropy generation is minimized when the energy of the reactants is maximized [14]. By adding moderator, the feed-back system allows the burner to operate at a higher temperature, thereby reducing entropy generation. This positive effect of moderation does not outweigh the negative effects in the mixing chambers or the decreased ability to extract work in a lower-pressure system.

Variations with Feed-Back Mass Flow Rate

Increasing the mass flow rate of the feed-back transfer exacerbates the unfavorable exergetic distribution discussed above. As a larger percentage of the main flow rate is fed back to the compressors, the compressor work and feed-back mixer get more destructive. While it is possible that feed-back matter could be useful for controlling emission of certain pollutants, it is not helpful in maximizing system efficiency. Thus feed-back matter transfer on a simple-cycle gas turbine are dismissed as non-optimal.

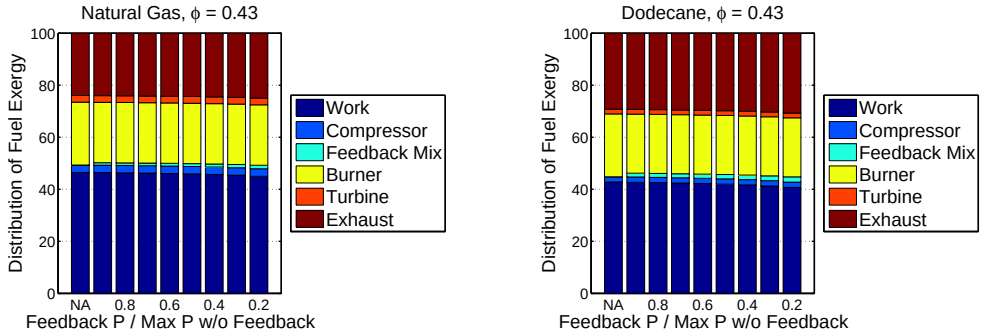

Figure 5.11: Exergy distribution for feed-back in common air-fuel systems. NA refers to the system without any feed-back. These graphs correspond to the same systems as in Fig. 5.9

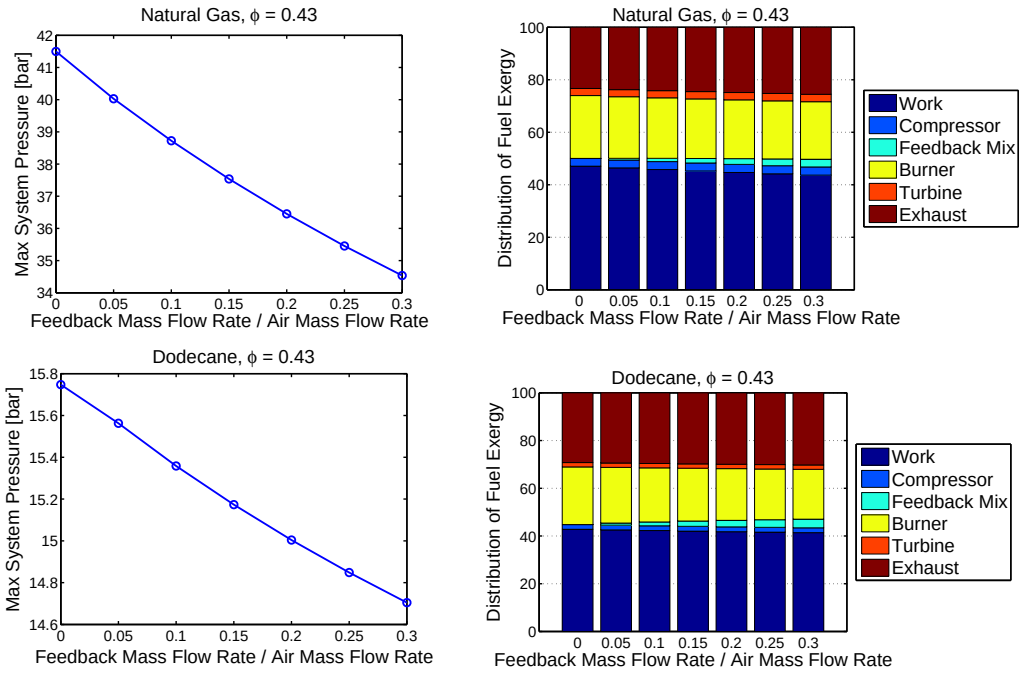

Figure 5.12: Effect of feed-back mass fraction on overall system pressure and exergy distribution. Feed-back occurs at 10 bar in all cases.

5.3.3 Thermodynamic Interpretation

The concentric cycle is a valuable conceptual tool in understanding matter transfer. By separating the effects of diffusive and thermal equilibrium, it is seen that the feed-back stream operates as a heat engine. The feed-back stream receives energy in the combustor and exhausts heat to the compressing air. As a heat engine, the feed-back loop naturally produces entropy. This entropy is transferred with the exhaust heat to the compressing air stream. This is unlike a typical heat engine, which would exhaust entropy to the environment. This inability to remove entropy from the system is the critical failure of the feed-back system. A clear contrast to this case is explored next in the external matter analysis.

5.4 External Matter Transfer

5.4.1 Conceptual Analysis

The final mode of matter transfer involves varying the types and quantities of fluids entering the engine. All fluids enter the engine separately and exit as a single fluid in the exhaust. For this reason, there is no independent control of an external matter transfer *out* of the engine, and only variations to the inflow of matter are examined. For a fixed flow rate of fuel, the flow rates of environmental flows into the engine can be varied. So far only air has been used, but now the possibility of introducing water at ambient temperature and pressure is introduced. The scope of this study remains engines dominated by air and fuel, so the flow rate of water is required to remain significantly smaller than the flow rate of air.

Many different wet cycles have been proposed to increase the efficiency and power of a gas turbine engine. In Ch. 2 examples were given of systems that vary in the amount of water used, the location of the water mixing, and the use of heat exchangers. Here the fundamental engine services provided by water are explored using only a simple-cycle gas turbine. These will be synthesized and incorporated

into the optimal matter-, heat-, and work-regenerative engine in Ch. 6.

As discussed in Ch. 4, only passive matter transfers without chemical separations are in the current purview. Once mixed, the water cannot be separated into a purified stream. This eliminates the possibility of either a water feed-back or feed-forward transfer. (A further degree of freedom associated with more active matter transfers would allow for these possibilities.) While the water could be transferred in these two manners as part of a mixed air or product stream, the previous two sections indicated this was unwise. Thus, the water flows in parallel with the mainstream air system until a point at which it is deemed time to mix the two streams. If no mixing location is chosen, it occurs automatically in the exhaust, as shown in Fig. 5.13.

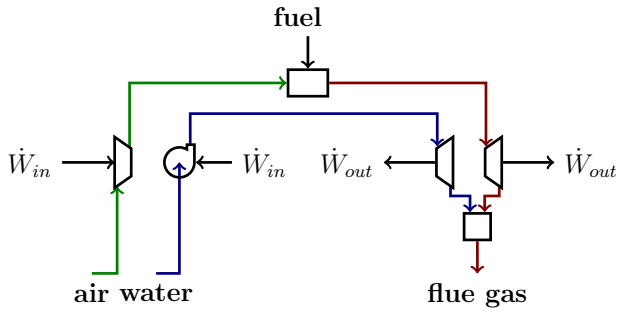

Figure 5.13: Parallel flow of air and water through a gas-turbine engine.

The Utility of Introducing Water

The water is pumped alongside the compressed air. If the water remains unmixed, then it does not enter the combustor where it could serve as moderator. Instead, it flows straight into a turbine and undergoes the same expansion ratio as the product gas. This flow plan for the water is not useful. With isentropic turbo-machinery, this process simply returns water to the environment at the same conditions it entered.[5] With non-ideal turbo-machinery, this process requires net work input and does nothing to alter the main air-fuel cycle, thus reducing the efficiency of the overall system.

[5] A gas turbine would be inappropriate here, but it would be possible to use a reaction turbine, or other device more suitable to liquid water power generation.

For these reasons this flow plan is not used in practice.

If water is going to be helpful in the engine, it will have to mix with the main air stream before or during combustion. The water will absorb thermal energy from the combustion process, allowing the water to behave as if it were an open heat engine. Now the water contributes to the net work output of the total cycle and lowers the overall back-work-ratio. Additionally, in order to maintain a set turbine inlet temperature, the air in the wet cycle would have to enter the combustor at a higher enthalpy than in a system without water. This increased reactant enthalpy will reduce the entropy generation of combustion. Additionally, it will require an increase in the air pressure. In net, this results in a water-based heat engine embedded within a higher-pressure, more extreme-state combustion engine. This appears advantageous for efficiency and warrants further study.

It is important to note that this use of an embedded heat engine is crucially different from that seen in Section 5.3. In the case of feed-back matter, the heat engine rejected its entropy to the main air stream. Here, the water rejects its entropy to the environment.

Optimal Water Mixing Location

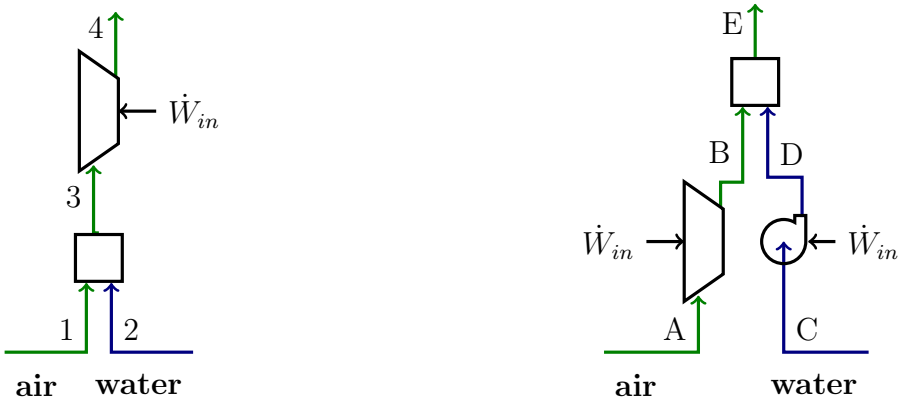

(a) Case I: Mixing before compression. (b) Case II: Compression before mixing.

Figure 5.14: Comparison of mixing location.

The optimal location of water mixing needs to be determined. It could occur before, during, or after compression. The consequences of water and air mixing will be diffusive and thermal. Here the irreversibility associated with mixing water and air before compression or after separate pressurization is compared. These two possibilities are illustrated in Fig. 5.14.

States 4 and E are at the same composition and pressure. Again by integrating Gibbs equation between these two states it is seen that the lower entropy corresponds to the lower enthalpy and temperature. The two systems have the same inlet matter streams, so the only difference in final temperature comes from the differences in compression/pumping work. Writing out the expressions for work for each system will allow for a determination of which has the greater final temperature. Specific volumes have subscripts to reference their fluid: a for air, w for water, and aw for wet air.

The work for the mixed system (Case I) is written

$$\dot{m}_3 w_{3,4} = \dot{m}_3 \int_{P_3}^{P_4} \frac{v_{aw}}{\eta_C} dP \tag{5.7}$$

Similarly, for the separate pressurization system (Case II), the total work can be written as

$$\dot{m}_A w_{A,B} + \dot{m}_C w_{C,D} = \dot{m}_A \int_{P_A}^{P_B} \frac{v_a dP}{\eta_C} + \dot{m}_C \int_{P_C}^{P_D} \frac{v_w dP}{\eta_P}$$

and because $P_A = P_C$ and $P_B = P_D$,

$$\dot{m}_A w_{A,B} + \dot{m}_C w_{C,D} = \int_{P_A}^{P_B} \left[\frac{\dot{m}_A v_a}{\eta_C} + \frac{\dot{m}_C v_w}{\eta_P} \right] dP \tag{5.8}$$

Subtracting Eq. 5.7 from Eq. 5.8 gives

$$\begin{aligned} \dot{m}_A w_{A,B} + \dot{m}_C w_{C,D} - \dot{m}_3 w_{3,4} = \int_{P_A}^{P_B} \Bigg[& \frac{\dot{m}_A v_a - (\dot{m}_A + \dot{m}_C) v_{aw}}{\eta_C} \\ & + \frac{\dot{m}_C v_w}{\eta_P} \Bigg] dP \end{aligned} \tag{5.9}$$

If the pump and compressor have the same polytropic efficiency, then the sign of Eq. 5.9 is determined by whether v_{aw} is larger or smaller than the mass average of v_a and v_w. This relationship isn't immediately obvious. Whereas the parallel systems have the benefit of pumping a liquid as opposed to gaseous water, the mixed system has the cooling effect of vaporization and therefore compressing an overall colder, denser fluid.

Figure 5.15 shows the value of the integrand in Eq. 5.9 for various ratios of the polytropic efficiency of the compressor to the pump. In all cases, the integrand is positive, indicating that it takes more work to pressurize the fluids separately than to mix and then compress. The decrease in the specific volume of air through the evaporation of the water is greater than the increase in the specific volume of the water through vaporization. Thus, when possible, it is more efficient to mix the fluids and then compress.

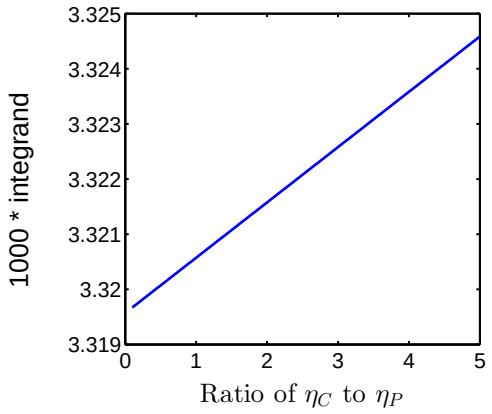

Figure 5.15: Evaluation of the integrand in Eq. 5.9 for varying ratios of compressor-to-pump polytropic efficiency.

On a broad level, mixing before compressing is in keeping with the strategy to prepare extreme state reactants. Because it is advantageous to have combustion reactants as hot as possible, the cooling is only valuable when it improves the compression process and therefore allows the system to go to even higher pressures. Cooling after

compression is of no value. Additionally, the quantity of available cooling is highest at low pressures and temperatures when the enthalpy of vaporization is greatest.

5.4.2 Numerical Analysis

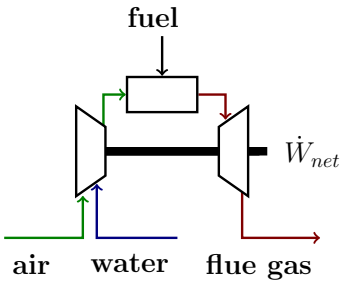

(a) Water injection into the compressor.

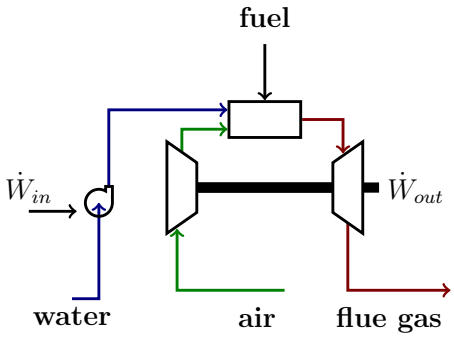

(b) Water injection into the combustor.

Figure 5.16: Schematics for possible water injection strategies

Figure 5.16a shows a schematic of a simple-cycle gas turbine with liquid water injection into the compressor with the air. Compare this to the system in Fig. 5.16b in which water is injected directly into the combustor. Figure 5.17 shows the exergy distribution for both systems operating with the same mass flow rates and turbine inlet temperature. Even when only injecting a mass of water equivalent to one thousandth of the mass of the air, the difference in water injection placement has over a 1% difference in exergy efficiency in the two cycles. This difference is primarily

seen in the decreased exhaust and mixing irreversibilities associated with compression injection, as discussed in Sec. 5.4.1.

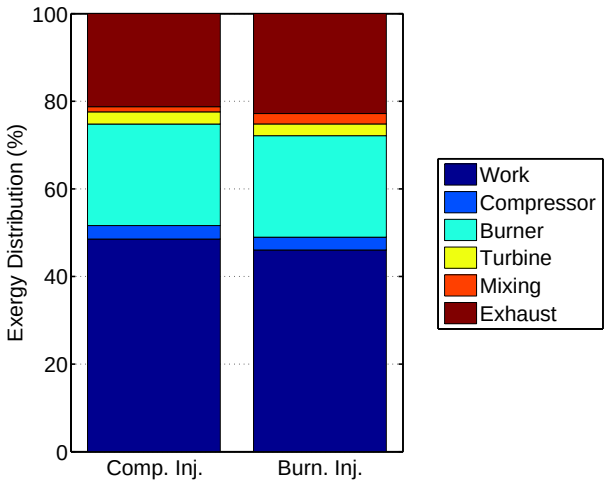

Figure 5.17: Comparison of irreversibilities with change in location of water injection

For the compression-injection schematic, a parametric study sweeping out different mass-flow ratios of air-to-water and air-to-fuel is shown in Fig. 5.18. The efficiency of the system increases as the excess air coefficient is increased by allowing the system to go to higher pressure ratios for the same fixed turbine inlet temperature. However, efficiency increases faster with respect to pressure as water is added to the system. The red line in the diagram indicates the systems operating at a 200:1 pressure ratio. Restricting the system to operate below this pressure ratio results in a maximum lower heating value efficiency just above 60%, at a minimum pressure ratio of 120:1.

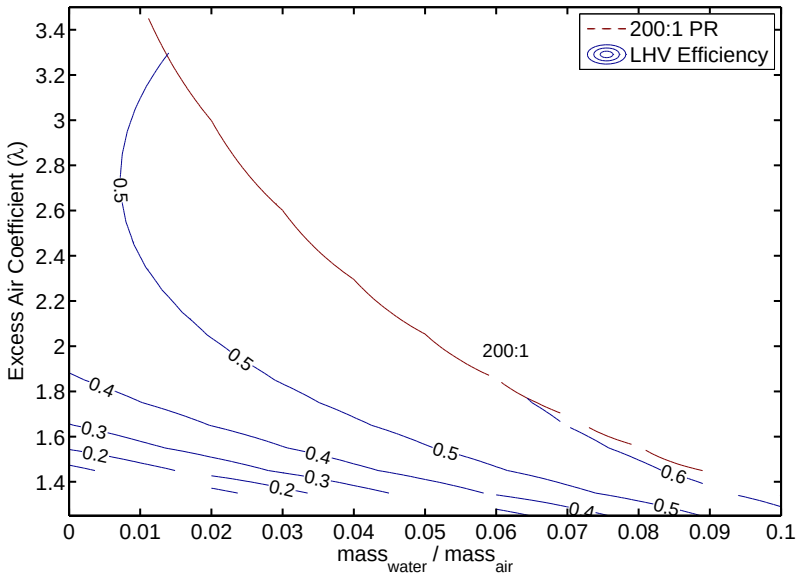

Figure 5.18: LHV efficiency contours for various air-water-fuel ratios. Contours are cut off beyond the 200:1 pressure ratio line. Turbine inlet temperature 1800K.

5.4.3 Thermodynamic Interpretation

The water acts as an open heat engine embedded into the gas-turbine system. Because of its phase change, it improves the overall system back-work ratio. Additionally, by providing moderator in the combustor, the water forces the engine to operate at an overall higher pressure ratio for the same turbine inlet temperature. This leads to an increased ability to extract work from the hot products and a lower exergy of the exhaust.

5.5 Conclusions

Neither feed-back nor feed-forward matter transfers are helpful in optimizing the energy efficiency of a simple-cycle gas-turbine engine. Water injection during compression is helpful, however. This strategy has the possibility of achieving a lower

heating value efficiency of 60%. This very simple system matches the efficiency of the best existing combined-cycle gas turbine engine [82]. In the following chapter, the optimal matter strategies developed here are integrated in with the previously determined optimal heat-and-work regenerative strategies. This will yield the best possible efficiency of a single-cycle combustion-based gas-turbine engine.

Chapter 6

Optimal Matter Transfer in a Steady-Flow, Combustion Engine

In this chapter, the previously developed principles of optimal matter transfer are applied to Ramakrishnan's optimally efficient work-and-heat regenerative architecture for steady-flow combustion engines [3]. Ramakrishnan's architecture was derived using a series of differential stages, as indicated in the Fig. 6.1b. For practicality, these differential series are displayed in the schematic in Fig. 6.1 instead using one stage each of finite-sized devices.[1] Numerical examples are calculated using the full differential architecture.

Initially, air is compressed with intercooling to reject entropy without rejecting significant exergy. A series of adiabatic compressors and heat exchangers are then used to increase the enthalpy and pressure of the air. These heat exchangers are sized to be capacity-matched with the expanding products, thereby minimizing heat transfer irreversibility. Because heat exchangers have an upper temperature limit that is lower than the desired combustor inlet temperature, the air is further compressed after heat exchange to maximize the extreme state of the reactants entering

[1] Although not essential to what follows, Ramakrishnan found that the single-stage, finite version of the architecture achieves the vast majority of the efficiency benefit of the fully differential, idealized version [3].

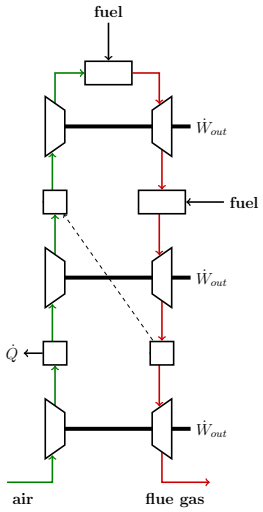

(a) Schematic

Process	Abbreviation
Intercooled Compression	$(CI)_k$
Adiabatic Compression	C
Recuperative Heating	$(X_{in})_m$
Post-Heating Compression	C
Adiabatic Combustion	B
Isothermal Combustion	$(TB)_n$
Capacity-matched Heat Exchange	$(TX_{out})_m$
Final Expansion	T

(b) Architecture processes

Figure 6.1: Optimal work-and-heat regenerative cycle. One stage of each process is shown in the schematic. [137].

the combustor[2]. The first part of the combustion is performed adiabatically until the maximum turbine temperature is reached. At this point, the maximum temperature is maintained through a series of burners and turbines. After all reaction is complete, the products are further expanded while donating thermal energy to the reactants as permitted by temperature and capacity matching.

In this chapter, degrees of freedom associated with matter transfer are introduced into this architecture. This will require re-optimizing both the architecture (device sequence) and the associated length parameters, (e.g., pressure ratio). Using the principles outlined in Ch. 5, decisions related to overall architecture are first discussed. This is followed by a parametric optimization analysis, which is used to illustrate the potential efficiency gains achievable by the optimal architecture.

[2]GE cools the entrance of their H-class gas-turbines down to 1430℃, implying that uncooled turbines (and by extension, compressors) could operate at temperatures up to approximately 1400℃ [136].

6.1 Architecture Optimization

Ramakrishnan optimized the flow paths of the main air and fuel streams to minimize entropy generation due to chemical reaction [138]. The architectural segment associated with combustion is thus already known. However, there are still architecture decisions to be made regarding each of the three matter-transfer modes: 1) feed-forward, 2) feed-back, and 3) external.

6.1.1 Feed-Forward Matter Transfer

In Ch. 5 it was shown that the addition of feed-forward matter transfer in the simple-cycle architecture resulted in increased entropy generation upon mixing. In a fully optimized heat-and-work regenerative architecture, there are more possible locations for a feed-forward transfer. These locations can again be categorized as either within the product stream or transferring from the reactant to the product stream.

Feed-Forward within the Product Stream

Within the product stream, a feed-forward transfer could bypass a turbine stage, a heat exchange stage, or a combination of the two. As mentioned above, the combustion process has already been optimized and will not be altered here.[3] The proof in Sec. 5.2.1 showed that there is no value in using feed-forward to bypass a turbine stage. Feed-forward could still be used to bypass a heat exchanger stage, as shown in Fig. 6.2a. A throttle is shown for the realistic case where there is a pressure drop due to fluid friction in the heat exchanger. However, the results are equivalent in the limit in which the heat exchanger is isobaric and the throttle disappears. The product side of the heat exchanger has a higher extensive thermal capacity than the reactant side since it has a greater mass flow rate (air and fuel) and is relatively enriched in triatomic products, as opposed to diatomic air. Thus, by reducing the

[3]Interestingly, the combustion process already involves feed-forward matter transfers of fuel relative to the air in order to create a near-isothermal process.

(a) Feed-forward around a HX stage. (b) No feed-forward stream.

Figure 6.2: Insertion of an HX feed-forward transfer.

mass flow rate on the product side, the feed-forward process could potentially achieve better thermal capacity matching between the hot and cold streams. But in doing so, it would incur an irreversibility associated with mixing the thermally mis-matched streams, labeled 2 and 4 in Fig. 6.2a.

The outlet state of the heat exchanger is equivalent to its attractor state. While there is heat transfer out of the product stream, there is no heat transfer out of the engine, so the method of minimizing the attractor entropy to minimize the overall engine irreversibility applies.[4] To determine if the feed-forward transfer is worthwhile, the product outlet entropy of the system in Fig. 6.2a is compared to that in Fig. 6.2b with the same inputs.

Integration of Gibbs equation between States 3 and B, which have the same pressure and composition, indicates that whichever state has the higher enthalpy also has the higher entropy. To fix the hot-side inlet state between Figs. 6.2a and 6.2b (States A and 1), the extensive heat transfer rate $\dot{Q}$ has to be fixed. With a fixed overall energy transfer out of the product stream, the output enthalpies are the

[4]If external heat transfer was introduced with intercoolers, the transfer would be entirely completed before the entrance to the cold-side heat exchanger. Thus, it would not differ between the two systems (with and without feed-forward transfer around a heat exchanger) and would not affect the result presented here.

same for both schematics. The overall hot-side pressure drop is assumed to be the same between the two cases. Thus, $h_B = h_3$, $P_B = P_3$, and therefore, $s_B = s_3$. No change occurred to the overall system efficiency.

With no benefit of bypassing a turbine or a hot-side heat exchanger, feed-forward transfer within the product stream is dismissed.

Feed-Forward from the Reactant Stream to the Product Stream

In Ch. 5 the possibility of feed-forward air transfer around stages of compression, burning, and expansion was examined. This was shown to increase mixing and exhaust irreversibility while detracting from the potential to produce work. The same situation occurs in the more complex optimal heat-and-work regenerative cycle.

For fixed mass flow rates and turbine inlet temperature, the system with feed-forward transfers will have less excess air in the burner than a system without feed-forward. This requires increased intercooling (and the associated irreversibility) or a decrease in the overall pressure ratio for the feed-forward system. With a decreased peak pressure, the final stage of combustion will also occur at a lower pressure. This combustor outlet state has a higher entropy than it would have if it were at a higher pressure for the same turbine inlet temperature. Thus, the feed-forward system either has more intercooling irreversibility or more exhaust matter exergy destruction, if not both.

Fundamentally, the increased mixing and exhaust irreversibilities associated with the feed-forward process are the same in both the simple-cycle and the work-, heat-, and matter-regenerative cycle. Feed-forward matter is therefore shown to be non-optimal for exergy efficiency.

6.1.2 Feed-Back Matter Transfer

Section 5.3 indicated that feed-back matter transfer increases the entropy in the main system and decreases the efficiency of the simple-cycle architecture. As with feed-forward, there are more potential locations for a feed-back transfer in the work-and-heat regenerative cycle. These locations are again categorized by the participating fluids.

Feed-Back within the Reactant Stream

In Ch. 5 it was determined that air feed-back around a compression stage deteriorates system performance. In the larger optimal cycle, the reactant stream also undergoes intercooling and recuperative heat exchange. However, a feed-back directly around either of these devices would require re-compression due to the inevitable (if small) pressure loss in the heat exchangers. (This analysis is restricted to *passive* matter transfers, so such a feed-back loop is technically out of the current scope.) But even if that were not so, and the compression work was allowed, the transfer would still not be advisable. Feed-back around a heat exchanger would increase the cold-side entrance temperature, which would reduce the overall quantity of thermal recuperation. In the process, it would increase entropy generation due to mixing at different temperatures. Feed-back around an intercooler would similarly increase the entropy generation of mixing and add to the mass flow rate that needs to be irreversibly cooled. Therefore, the optimal cycle does not include feed-back within the reactant stream.

Feed-Back from the Product Stream to the Reactant Stream

Feed-back matter from the product to reactant stream will increase the intensive enthalpy of the reactant stream once mixed. Depending on where this occurs, one of three consequences could result:

I. Increased intercooling load.

If the feed-back stream mixes in before the intercooling ends, then the additional thermal energy could simply be removed from the system. This adds to the intercooling irreversibility. If the feed-back happens at a higher pressure point, or the intercooling is kept constant, one of the remaining enumerated possibilities occurs.

II. Decreased thermal recuperation.
Without increased intercooling, if the same pressure ratios are kept constant, the mixed feed-back and air stream will be hotter entering the cold-side of the heat exchanger. This reduces the total amount of thermal recuperation possible and increases the exhaust exergy of the products. This situation could be avoided if instead of fixing the pressure ratios, the heat exchanger inlet temperature were fixed and the corresponding pressure ratio reduced.

III. Decreased pressure ratio.
If either the feed-back matter stream mixes in with the reactants at a higher pressure than the recuperative heat exchange, or the temperature range in the heat exchanger is kept constant, the pressure at the end of combustion will decrease. Because the reactant stream is hotter with the feed-back fluid, it cannot accept as much energy before reaching the turbine inlet temperature. Therefore, for fixed exothermicity of the combustion reaction, less work can be added in compression (or more of the combustion has to be done after some expansion). The result is that combustion ends at the turbine inlet temperature, but at a lower pressure than it would have otherwise. This is the same situation as occurred for the simple-cycle architecture. The final combustor outlet state is at a higher entropy than it would have been at a higher pressure. This results in a higher exhaust entropy and therefore system irreversibility.

In the simple-cycle system, feed-back matter from products to reactants resulted in increased compressor and mixer irreversibilities, which were not sufficiently counteracted by a slight decrease in combustion irreversibility (due to a more extreme

reactant state). In the work-and-heat regenerative cycle, the compressor and mixer irreversibilities remain a problem. Meanwhile, the benefit to the combustor is largely unrealized. The mixed adiabatic-isothermal combustion strategy allows for the reaction to take place at a high enthalpy regardless of the pressure. The moderator effect of the feed-back matter is less significant as a result.

Despite more candidate locations, feed-back matter is again seen to be disadvantageous for optimizing efficiency in the heat-and-work regenerative cycle.

6.1.3 External Matter Transfers

Optimal Location of Water Transfer

In Ch. 5 it was determined that the process of mixing air and water produces the least entropy when done at the system inlet. This gives the best ability to receive the thermal moderation benefit throughout the compression process and minimizes the temperature difference in the initial mixing. However, the quantity of water that can evaporate into the air is limited at ambient (low) temperature. This limitation remains in effect during the intercooling process, which results in cool temperatures at high pressures. Thus, the water mixing will have to occur after the intercooled compression process. As saturation conditions allow, the water should be introduced as soon as possible.

Figure 6.3 shows the schematic associated with this water injection strategy. Air is compressed with intercooling between States 1 and 3 after which water at equal pressure can be mixed in. If the water cannot all be mixed into the air at State 3, then a small amount of compression is done to further heat the air before mixing in more water. This is repeated as many times as necessary. At a specified pressure (here State 5), recuperative heat transfer occurs from the products to the reactants. Hypothetically water could continue to be mixed in after this point, as indicated by the mixing of streams 16 and 6. After all thermal recuperation is complete, the wet air enters the final adiabatic compressor at State 7.

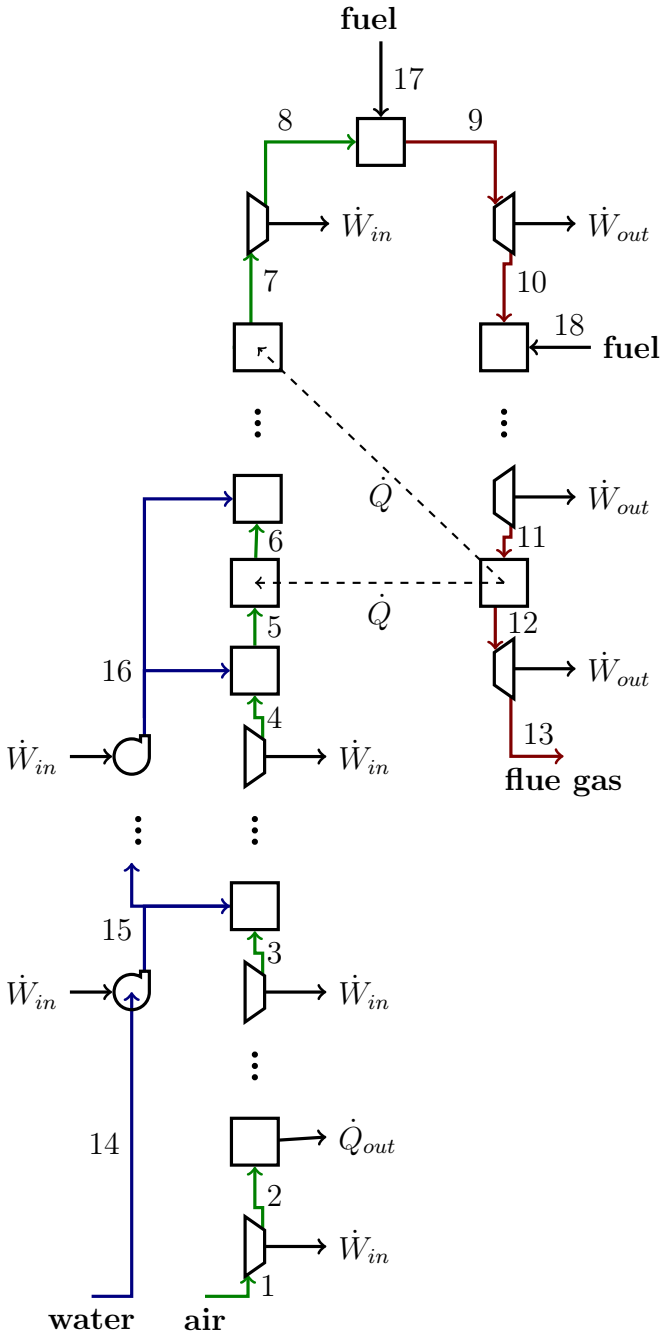

Figure 6.3: Full architecture with water injection after intercooling. As discussed in Sec. 6.1.3, the hot-side heat exchange is done in a single stage.

Optimal Location of Heat Exchangers

It is possible that the alteration of external matter transfer will require moving the hot-side heat exchangers to reduce irreversibility. Heat transfer irreversibility originates from temperature differences between the hot and cold streams. Choosing the location of the hot-and-cold sides of the heat exchanger to minimize temperature differences reduces entropy generation. Ramakrishnan showed that this is accomplished in the work-and-heat regenerative architecture by matching the extensive heat capacities of the two streams. With the dominant heat capacity on the hot side, counter-flow heat exchangers are best-matched in temperature at the hot-end of the heat exchanger. In order to maintain this temperature matching, the hot-side heat exchange should not happen in one stage, but rather in small stages alternating with stages of expansion to reduce the product enthalpy and keep the temperature well matched with the cold-side reactants.[5]

The mismatch in heat capacities of the products and reactants, while present, is a small effect. Ramakrishnan showed that using 20 optimized stages of hot-side heat exchange and expansion gained less than a 0.5% system efficiency over a system that delayed all heat exchange until a single stage located after expansion [3]. When water is added to the reactant stream, the heat capacities of the products and reactants are better matched, and the need to expand between heat exchangers becomes less significant. Thus in practice, the architecture would most likely be simplified to have all of the heat recuperation occur after expansion. (This single-stage of recuperation is shown in Fig. 6.3 between states 11 and 12.

Figure. 6.4 compares the heat recuperation process in the optimal work-and-heat regenerative architecture with and without water injection. While air typically has some water in it, to emphasize the small effect, a completely dry (0% water) case is compared to a case with a water mass flow rate equal to 1% of that of air. The 1% water case results in slightly better capacity matching in the heat exchanger. This

[5]If a system could be conceived where the reactants had the higher heat capacity, the cold-side heat exchangers would each be small, and alternated with compressors, and the hot-side heat exchange could be done in a single step.

is more easily visible in Fig. 6.4b, which shows the temperature difference across the heat exchanger at every streamwise position. Additionally, when many stages of intercooling are used during compression, the wet system enters the cold side of the heat exchanger colder, and therefore increases the overall amount of heat exchange possible.

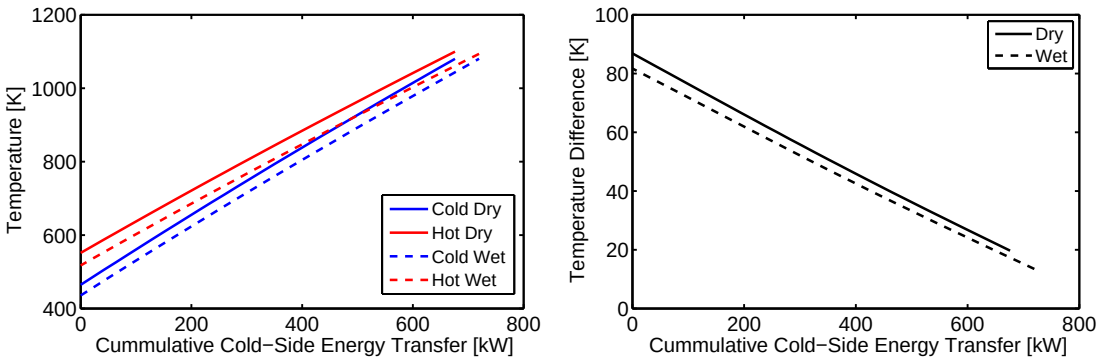

(a) (b)

Figure 6.4: Heat exchanger performance for systems with and without water injection. Intercooling to 20 bar. Dry system has air at 0% relative humidity. Wet system has water mass flow rate as 1% air mass flow rate. Turbine inlet temperature fixed at 1800 K.

6.2 Parameter Optimization

There are five parameters that need to be optimized for the current architecture. Three are the reactant-stream pressures associated with the end of intercooling, the beginning of recuperative heat exchange, and the combustor inlet. The other two are the mass flow ratios of air-to-fuel (or equivalence ratio) and air-to-water. Only the combustor inlet pressure can be reasonably determined without jointly optimizing for all five.

Combustion Inlet Pressure

The recuperative heat exchanger has a maximum temperature lower than that of the combustor, here 1100 K. Thus the remaining pressure ratio of compression and the fluid composition dictate the combustor inlet temperature. Because the fluid is well-modeled as an ideal gas, it is actually the temperature, not the pressure that dictates the extreme state of the fluid.

For all air-water compositions examined, combustion irreversibility was minimized when the burner inlet temperature was between 1200 and 1300 K. The difference between having combustion begin at 1200 vs 1300 K resulted in less than a 0.1 percentage-point change in system efficiency in all cases. Thus, for ease of modeling, the combustor inlet temperature was fixed at 1200 K for all cases.

Joint-parameter Optimization

Removing the combustion inlet pressure/temperature, there are four remaining independent parameters that affect the system efficiency in combination. The optimization can be conceptualized as occurring over a surface in a five-dimensional space composed of the four independent parameters and the dependent efficiency metric. Due to the complexity of the system, the surface is not smooth and it has many local efficiency maxima. Examining the effect of an individual parameter is equivalent to taking a transect across the surface along one of the five-axes that does not necessarily correlate with efficiency contours. As such, there is minimal information to be gained in optimizing the four parameters individually or sequentially. Instead, the entire surface must be scanned to find the globally optimal efficiency point. The roughness of the surface proved a significant difficulty in using built-in numerical optimizers in Matlab. Ultimately the surface was scanned by calculating the plant efficiency within four nested loops varying each of the parameters.

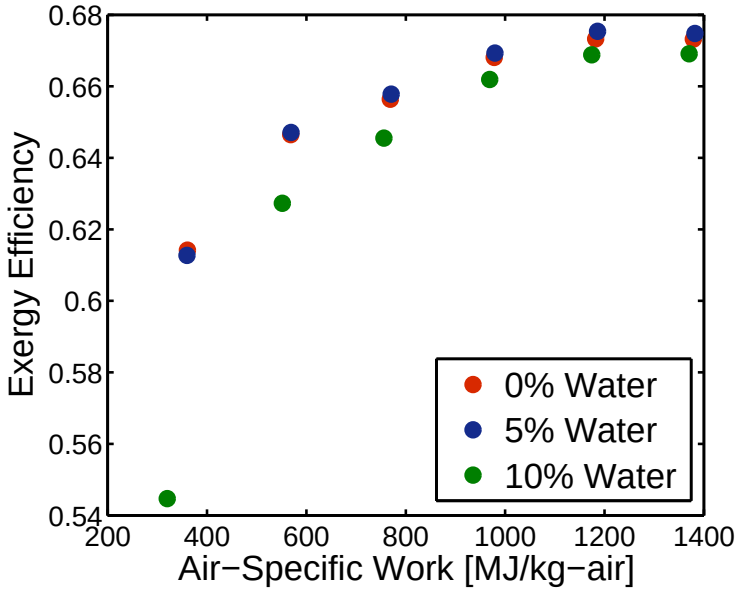

Figure 6.5: Efficiency-work curves of architectures with optimized parameters over a range of equivalence ratios and water quantities.Water percentages defined as mass of water with respect to mass of air. Turbine inlet temperature is 1650 K in all cases.[6]

6.3 Results

Figure 6.5 shows the final optimized curves of efficiency vs. air-specific work for the work-, heat-, and matter-regenerative system.[6] Each point on a curve is for a different air-fuel equivalence ratio, ranging in steps of 0.1 from 0.2 at the far left to 0.7 at the right. Table 6.1 lists the pressures that were key independent parameters for all of the points in Fig. 6.5. Ultimately, the addition of a small amount of water makes almost no difference to the peak system efficiency compared to the optimal work-and-heat architecture. (All of the optimal work-and-heat architectures, with and without water, are more efficient and have higher air-specific work than existing gas-turbine systems.) When 10% of the air flow rate is added as water, the efficiency

[6]For purposes of comparison, the numerical results in this section will use the same turbine inlet temperature as Ramakrishnan rather than the 1800 K value used elsewhere.

Process Ending	**Equivalence Ratio**					
	0.2	**0.3**	**0.4**	**0.5**	**0.6**	**0.7**
0% Water						
Intercooling	5	9	17	33	60	113
Heat Exchange	6	12	31	47	104	227
Final Compression	9.2	18.4	47.5	72.1	160.0	348.1
5% Water						
Intercooling	1	2	5	7	15	26
Heat Exchange	6	10	31	47	121	24
Final Compression	9.2	15.3	47.5	72.1	185.6	380.3
10% Water						
Intercooling	1	1	2	3	5	9
Heat Exchange	21	25	45	70	138	285
Final Compression	33.7	38.3	69.0	107.4	211.6	437.1

Table 6.1: Pressure (in bar) at end of listed processes for various equivalence ratios and air-water ratios. All points correspond to those in Fig. 6.5.

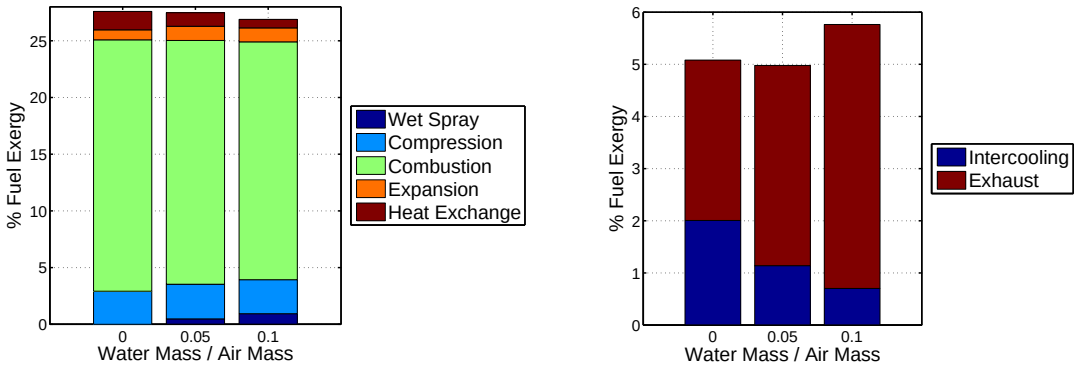

(a) Internal destruction. (b) External destruction.

Figure 6.6: Exergy distribution across different amounts of water injection. $\phi = 0.5, TIT = 1650$

of the system actually decreases below the original optimized dry case. This is particularly noticeable at low equivalence ratios when there is excessive moderator in the system.

Figure 6.6 depicts some of the trade-offs that occur when water is added to the system. The internal irreversibilities shown in Fig. 6.6a decrease with added water flow rate, as the water results in improved extensive heat capacity matching in the heat exchanger. However this alternatively could have been accomplished with staged heat exchangers and turbines, as shown by Ramakrishnan [3]. Additionally, water reduces combustion irreversibility. These positive effects of increased water outweigh the associated entropy production of the water-air mixing, labeled Wet Spray. In the dry case, combustion losses account for 22.2% of the fuel exergy, compared to 21.5% and 21.2% for the 5% and 10% water cases, respectively. Since the inlet combustor temperature is fixed, this is largely a result of the changed chemical composition. The entropy generation in an adiabatic and isobaric combustor is

$$\dot{S}_{gen} = \frac{\dot{m}_{fuel}}{\hat{M}_{fuel}} \int_{\xi=0}^{\xi=1} \frac{1}{T} \sum_{i=1}^{n} \mu_i \frac{\nu_i}{\nu_{fuel}} \tag{6.1}$$

where i is the chemical species, the ν_i are the stoichiometric coefficients of the global reaction from reactants to products, $\hat{M}_{fuel}$ is the fuel molar mass, and ξ is the extent of reaction. While the temperature profile will vary between water quantity cases, it is always 1200 K at the inlet of the combustor and quickly increases to a maximum of 1650 K.[7] Water is one of the main combustion products. By adding water to the reactants, the difference in chemical potentials between the reactants and products decreases, resulting in an overall decrease in entropy generation. These reductions in irreversibility are partially countered by an increase in the destruction associated with the compression and humidification processes.

The addition of water first decreases and then increases the external exergy destruction. The decrease comes from a reduction in intercooling. The increase comes from a hotter and more massive exhaust stream that is at a composition further from the dead state than the products of dry combustion.

The optimal amount of water results from when the sum of the internal and external destructions is at a minimum. This occurs at approximately 5% water with an exergy efficiency of 67.5% (69.8% LHV efficiency). This trade-off between internal and external irreversibilities is characteristic of optimizing an external transfer. Ramakrishnan first discussed this trade-off with the optimization for external heat transfer in the form of intercooling [3]. Whereas an internal transfer between devices can improve both the internal and external destructions, this is not true for an external transfer.[8]

The addition of water may not have changed the efficiency significantly from the dry case, but it does affect the operating parameters. These are discussed briefly below.

[7]The combustors staged with turbines are also each approximated as adiabatic and isobaric and Eq. 6.1 still applies.

[8]The differences between optimizing for internal versus external transfers are discussed in detail by Ramakrishnan [3].

Intercooling Pressure

Water injection and intercooling both reduce the temperature of the compressed fluid. The temperature drop leads to a decrease in specific volume and an associated decrease in compression work. Because water injection and intercooling compete to deliver this benefit, an increased water injection rate results in a decrease in intercooling. Whereas the dry system with an equivalence ratio of 0.6 was intercooled in many stages to keep the temperature within 20 K of ambient up to a pressure of 60 bar, the wet systems were only cooled to 15 and 5 bar for the 5% and 10% water cases, respectively.

Cold-Side Heat Exchange Pressure

The cold-side heat exchanger pressure is optimized for an indirect goal. Because the intercooling can be separately adjusted, any temperature can be achieved at the inlet of the heat exchanger regardless of its pressure. However, the higher the pressure of the heat exchanger, the higher the pressure of the combustion system. As such, the heat exchanger pressure is actually determined to avoid expansion after entrance to the single-stage heat exchanger on the *product* stream. The product stream is expanded such that the ambient pressure is reached at the hot-side heat exchanger inlet temperature of 1100 K. This allows for maximal recuperative heat exchange.

Combustor Inlet Pressure

As discussed above, the combustor inlet pressure is chosen to take the system from the 1100 K heat exchanger limit to the 1200 K combustor inlet. The compression to accomplish this temperature change naturally increases with increased water content and extensive heat capacity. At 0.6 fuel-air equivalence ratio, the peak pressure for the optimized 0%, 5%, and 10% water cases was 160 bar, 186 bar, and 212 bar, respectively.

Relative Fluid Flow Rates

Air and water both serve as thermal moderator so a significant influx of water should reduce the quantity of air. However, there was not an efficiency incentive to use significant amounts of water. Additionally, the efficiency curves in Fig. 6.5 are quite flat between 0.6 and 0.75 equivalence ratio. While the optimal fuel-air equivalence ratio drifted upward as more water was added, the difference was not particularly large (e.g. from 0.7 to 0.74 moving from 0% to 10% water).

6.4 Ramifications for Optimal Water Use

6.4.1 Water for Thermal Energy Transfer

Ultimately, the addition of external water makes a negligible difference in the optimal system efficiency. This is at first surprising, given the great success of water to improve efficiency and power in the simple cycle in Ch. 5 as well as in the literature in STIG and HAT configurations. However, none of those cycles were already optimized for regenerative heat transfer. The optimal heat-and-work regenerative engine without water already surpasses the peak power and efficiency of all of the wet cycles reported by Horlock [4].[9] With optimal heat transfer strategies, there is essentially nothing to be gained through water. Water is only being used here for its thermal properties. It is not considered to be significantly reactive, nor is there value in the diffusive mixing with air. Using water in lieu of heat transfer at various points changes the operating parameters, such as the peak pressure ratio, but not the system efficiency.

Because the benefits of water are only thermal, there is no inherent reason that the fluids need to be mixed. Using concentric cycle analysis, the best cycle would be the same as above but with water moving in parallel to the air/product streams always in perfect thermal equilibrium. This would avoid entropy generation with

[9]The best reported wet cycles reaches an efficiency of approximately 55% at an air-specific work of 780 MJ/kg-air.

diffusion. Now the water is used in an open heat engine. As a heat engine with fixed ambient temperature, it is most efficient when the hot side is as hot as possible. Routing the water such that it has a section of thermal contact with the burner then allows for the greatest possible temperature difference in the heat engine. However, in reality, keeping the fluids in parallel would require using closed heat exchangers that constrain temperatures to be lower than in the burner (here 1100 K). So the penalty for transitioning from the mixed to concentric cycles is the loss of the ability to exchange directly with the burner.

This could potentially change if the reactive, diffusive, or mechanical properties of water were considered. Allowing chemical separations or the addition of new devices that utilize the chemical attributes of water could change this result. Additionally, any system that allows for circulation of pure water (either through a separate steam cycle or through the separation from a mixed stream) allows for the potential to take advantage of the mechanical exergy associated with expanding to a condenser at ambient temperature and therefore below ambient pressure. However as discussed below, the potential gains in efficiency from utilizing those mechanical and diffusive degrees of freedom for water appear to be small.

6.4.2 Water for Thermal and Mechanical Energy Transfer

Once the water is separated from the air, it no longer needs to be used in an open cycle. By closing its cycle, the water can be expanded below dead state pressure to the point of condensation at ambient temperature. However, by containing it in a separate cycle, it loses the ability to receive thermal energy at the burner temperature and is instead restricted to the closed heat exchanger temperature of 1100 K. This closed water cycle–concentric with the air-fuel cycle–is shown in Fig. 6.7.

The closed concentric system achieves slightly higher efficiencies than the injected water system, although these values are likely within the uncertainty of the model. (For mass flow rates of water equal to 5%, 10%, 50%, and 100% that of the air, the exergy efficiency was 67.4%, 67.6%, 68.0%, and 67.8%, respectively.) Such a small

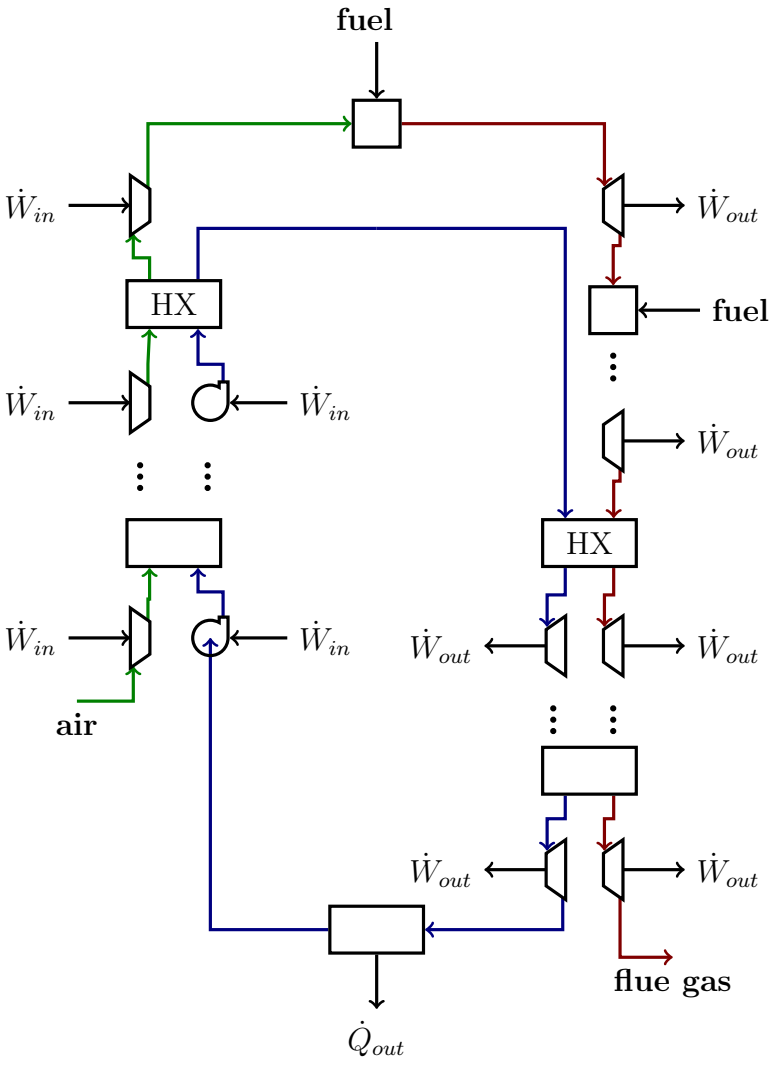

Figure 6.7: Closed water system.

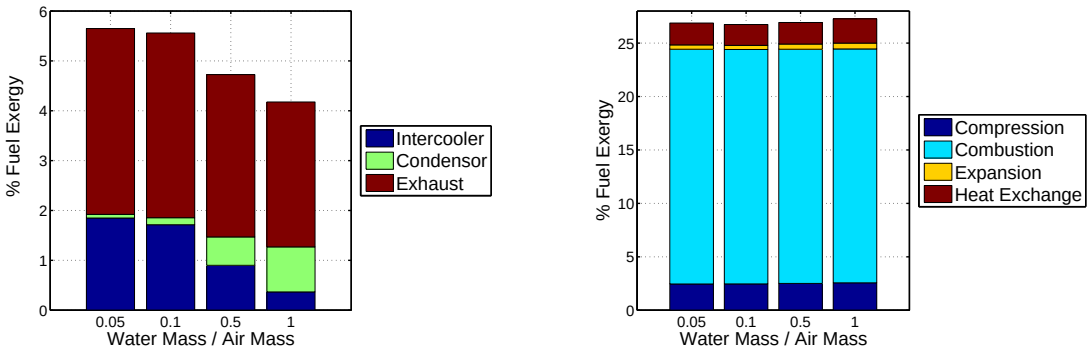

(a) External exergy destruction (b) Internal exergy destruction

Figure 6.8: Exergy distribution for the closed concentric cycle.

increase in efficiency will not likely justify building the complex concentric version of the system.

Figure 6.8 shows the internal and external destruction for the concentric closed cycle with 5%, 10%, 50%, and 100% of the mass flow rate of air as water. While the water is no longer an external energy transfer, it requires a condenser, which has an external transfer of heat. As such, there is again a trade-off between internal and external destruction.

Increased water mass flow rate decreases the external exergy destruction in three ways. First, the amount of intercooling decreases because the water cools the air during compression. The peak pressure of intercooling decreases from 81 bar for the 5% water case to 4 bar in the 100% water case. Second, the temperature of the flue gas decreases from 456K for the 5% water case to 405K for the 100% water case. Finally, the cases with more water also have a higher fuel-air equivalence ratio. This means that there is less mass in the flue gas, which also reduces the external destruction. The condenser exergy destruction increases with water mass flow rate. However, the efficiency of the system peaks before this destruction overtakes the other improvements to reduce external exergy destruction.

Meanwhile, increased water first decreases and then increases the internal exergy destruction. This is a combination of several small changes to the compression,

combustion, expansion, and heat exchange irreversibilities. Increased water slightly decreases the combustion irreversibility (due to a change in equivalence ratio and extensive mass in the combustor) and compressor irreversibility. The turbine and heat exchange irreversibilities first decrease with water mass, but the water adds extensively to the destruction in these devices and eventually this extensive effect trumps any positive intensive effect. The advantage of the low pressure steam expansion does increase the optimal quantity of water to a value closer to 50% that of the mass flow rate of air.

6.5 Conclusions

Energy in gas-turbine engines can be transferred by heat, work, or matter. While matter could in theory involve mechanical, thermal, diffusive, and chemical equilibrations, ultimately only the the thermal and diffusive equilibrations were used here. Because the diffusive equilibrations resulted in entropy generation, air's and water's only positive contribution was thermal. As a result, the optimization for matter resulted in negligible improvement over the system previously optimized for work and heat.

Closing off a concentric water system enabled the production of work from expansion below the ambient pressure to the condenser temperature. This use of a mechanical equilibration resulted in a very slight improvement in efficiency. The fact that closing the water system and removing the potential to exchange thermal energy in the burner improved (or at least didn't deteriorate) the engine efficiency indicates that combined cycles may have an inherent benefit over wet cycles. This can be attributed to the ability to use the mechanical component of water's exergy.

The inclusion of a concentration cell would allow for extracting work during diffusion. However, the potential efficiency gain by incorporating this device is inherently small. As discussed in Ch. 4 (and shown in Fig. 4.6), the diffusive component of chemical exergy is dwarfed by the reactive component.

To realize a significant efficiency gain with matter, more work has to be extracted from the reactive chemical exergy. This could be done with a new device, such as a fuel cell. (This will be explored in the next chapter.) Additionally, more active matter transfers with devices that allow for chemical or phase separations could add to the opportunity to extract energy from chemical differences. Without such a change in devices or active matter transfer, the optimal work-, heat-, and matter-regenerative architecture presented here provides the best possible energy efficiency for a steady-flow combustion engine.

Chapter 7

Solid-oxide Fuel Cell, Gas Turbine Engines

So far the thermodynamic methodology has only been applied to steady-flow combustion engines. In this chapter, an alternative to the unrestrained combustion reaction is introduced with a solid-oxide fuel cell. Unlike in a combustor, in a fuel cell the reaction progress is directly coupled to electricity production. This *restrained* mechanism has different irreversibility characteristics than the *unrestrained* combustion and is explored in depth here.

This optimization will start back with the minimum CBT architecture and add in the solid-oxide fuel cell (SOFC). After the new sequence is optimized, internal and external heat transfer will also be added as new degrees of freedom. In order to do this, it is necessary to have a computer model of a solid-oxide fuel cell, gas-turbine engine that captures the operation and performance of representative systems. A proposed Mitsubishi Heavy Industries (MHI) system is used as the model system for this purpose. This chapter will first discuss the building and validating of the computer model of the MHI system. It will then apply the systematic thermodynamic procedure from Chap. 4 to optimize engines with solid-oxide fuel cells.

7.1 Building and Validating the SOFC-GT Model

Ultimately a model is needed that can examine combinations of solid-oxide fuel cells and gas-turbine engines. Several such cycles have been proposed and modeled [86–88,93]. Mitsubishi Heavy Industries has proposed such a double-cycle system and is now working on a variant that includes a bottoming steam cycle [139]. They report that their double and triple cycles have LHV efficiencies greater than 52% and 70%, respectively. The 70% value is remarkably higher than the most efficient existing natural gas combined cycles, which reach 60% [82–84]. Such a high value is intriguing and was chosen as the target to model, understand, and adapt for optimality. For brevity, double-cycle systems of solid-oxide fuel cells and gas turbines are hereafter referred to as SOFC-GT's. Similarly, the triple cycle variant with the bottoming steam cycle is called a SOFC-GT-ST.

7.1.1 Reverse Engineering the MHI System

A computer model was built of the SOFC-GT-ST system in Matlab. The model was also implemented in segments to analyze a GT and a GT-ST system. Because it encapsulates all of the cycles studied, only the triple cycle will be explained in full, but with an understanding that the same processes and parameters were applied within the various sub-models. Schematics and device parameters were taken from MHI where possible, and approximated from engineering literature otherwise.

The triple-cycle model was designed to match the exergetic operation of the MHI plant. The schematic and operating parameters were based on a presentation on the MHI system at the World Engineers Convention in 2011 [139]. Figure 7.1 is taken directly from this presentation. Accompanying the figure was text stating that this cycle could achieve greater than a 70% LHV efficiency. From this graph and efficiency value, many operating parameters can be reverse engineered.

The most straight-forward parameter to determine is the mass flow rate of fuel.

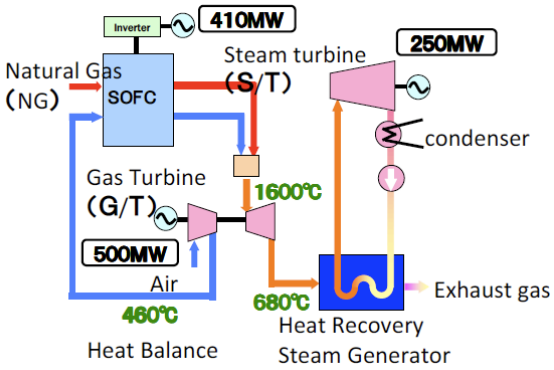

Figure 7.1: Triple cycle schematic taken directly from Yamada [139].

First, the conservative estimate of 70% efficiency (η_{LHV}) is supposed. Then, a reasonable natural gas composition is chosen and the corresponding LHV calculated. At this point, Eqn. 7.1 is solved for $\dot{m}_{fuel}$.

$$\eta_{LHV} = \frac{\dot{W}_{net}}{\dot{m}_{fuel} LHV_{fuel}} \tag{7.1}$$

This and all subsequent calculations on the MHI plant were done with both methane as the fuel and separately with a more realistic six-component natural gas mixture. The results differed by less than 2% for the calculated parameters. Given this agreement, the simpler methane system was used for all subsequent analyses. Using methane results in a fuel flow rate of 33 kg/s.

The gas-system pressures can also be estimated from Fig. 7.1. The outlet temperature of the compressor is given as 460℃. Air is assumed to enter the compressor at ambient conditions with a composition of 76.6% N_2, 20.4% O_2, and 3% H_2O. If a high efficiency compressor is used (polytropic efficiency 89%, estimated), the pressure ratio is found to be 18.4:1 [129]. An additional assumption of a 4% pressure drop across the combustor then defines the turbine pressure ratio.[1]

The key remaining unknown is the mass flow rate of air. This can be estimated

[1]These polytropic efficiencies and pressure drops will be used in all further calculations for gaseous streams.

in two separate ways from Fig. 7.1. The first involves an energy balance around the compressor and gas turbine. The temperatures bounding the turbine are given in the diagram, as is the net electrical work out of the gas turbine. The electrical work specification can be converted back to a mechanical work estimate assuming a 99% efficient generator. The turbine's composition is still unknown, but must be the products of complete oxidation of the fuel with air. Depending on exactly how much work the fuel cell does, the product composition may be more enriched with H_2O and CO_2 versus excess O_2. Regardless, the vast majority of the products will be N_2 and the specific heat capacity is unlikely to change significantly with the oxidation partitioning between the fuel cell and combustor. The heat capacity will depend on the quantity of air, which is guessed iteratively until an energy balance around the turbine and compressor yields the known net work output. The resulting air mass flow rate of 862 kg/s corresponds to a turbine outlet temperature of 705℃–within 25 degrees of the prescribed value of 680℃. This near-convergence in temperature partially validates the assumption that the mass flow rate, not the composition, dominates the product heat capacity. This air-to-fuel mass flow rate ratio corresponds to an excess air coefficient λ of 1.49.

The second method to estimate the air mass flow rate similarly involves an energy balance with iteration, but this time around the fuel cell and combustor. As discussed above, the compressor outlet state is known. Since no pressurization is indicated for the natural gas in Fig. 7.1, it seems reasonable to assume the fuel enters the plant at the same pressure as the compressed air and at ambient temperature, as if it were delivered in a high-pressure pipeline.[2] This assumption results in both inlet states to the fuel cell being fully defined. With a reasonable inverter efficiency of 95%, it is possible to find the work directly from the fuel cell as 432 MW. The post-burner composition is again the products of complete combustion of a presumed amount of air with the known amount of fuel. The correct guess is the one that results

[2]While the fuel had to be compressed at some point, it seems unlikely that MHI incorporated the associated energy penalty in the 70% efficiency number. As such, in an effort to reproduce their numbers, this analysis shall also neglect the energy cost of delivering methane to the fuel cell.

in consistency with the first law around the fuel cell and combustor. This method results in an air mass flow rate of 705 kg/s, which leads to an excess air coefficient of 1.25.

Clearly these two methods of calculating the mass flow rate of air do not agree perfectly. This is not surprising, given the number of assumptions that must be made and the likely use of rounding the numbers on MHI's diagram. The two values of the excess air coefficient were taken as plausible bounds within which the computer model would investigate. The final value of λ used in the model was determined based upon best agreement with MHI's overall work breakdown. It is also worth noting that the shift in λ between 1.25 and 1.49 should not radically change the exergy distribution in the plant, which is the present focus. As such, exact matching in the excess air coefficient is not crucial to the subsequent analysis.

While analysis of Fig. 7.1 resolved the work flows, gas pressure ratio, and approximate gas flow rates, many more parameters, and in the case of the steam cycle, full schematics, remain uncertain. These were judiciously estimated from the literature, as discussed in the following sections on each of the constituent power cycles. The validity of the ultimate schematic used in the model is demonstrated by its ability to replicate the known power distribution and efficiency of the MHI plant.

7.1.2 Solid-Oxide Fuel Cell Model

Four key chemical processes are modeled in the fuel cell: the two half-reactions occurring at the anode and cathode, the water-gas shift reaction, and the steam-methane reformation (SMR) reaction. The half-reactions are driven by the electrochemical potential difference across the electrolyte and are shown in Eqn. 7.2 and Eqn. 7.3 for the anode and cathode, respectively. The water gas shift reaction, shown in Eqn. 7.4, is significantly faster than the other chemical processes involved and is taken to be in equilibrium at all times. The SMR reaction (Eqn. 7.5) is kinetically limited, and an empirical formula for the reaction rate is needed. Various Arrhenius fits have been proposed [97]. The one chosen is shown in Eqn. 7.6. It is first-order in the partial

pressure of CH_4 and zeroth order in all other components. The activation energy E_a and the prefactor k_0 are reported to be 82 kJ/mol and 4274 mol/(s-m^2-bar), respectively, for an anode 80% ZrO_2 and 20% Ni cermet by weight [98].

$$O^{2-} + H_2 \rightarrow H_2O + 2e^- \tag{7.2}$$

$$\frac{1}{2}O_2 + 2e^- \rightarrow O^{2-} \tag{7.3}$$

$$H_2O + CO \rightleftharpoons H_2 + CO_2 \tag{7.4}$$

$$CH_4 + H_2O \rightarrow CO + 3H_2 \tag{7.5}$$

$$k = k_0 P_{CH_4} e^{\frac{-E_a}{RT}} \tag{7.6}$$

Each of these reactions is advanced in differential time steps moving flow-wise through the channels in a 1+1D model [88]. Exergy and other performance-related thermodynamic properties are tracked at every step.

The solid-oxide fuel cell model attempts to match important power criteria of the MHI system. This is done by adjusting the operating parameters while still adhering to conventional specifications for a yttria-stabilized zirconia (YSZ) electrolyte and nickel-based anode. The cell is assumed to be steady, internally reforming, isobaric, and isothermal.[3] While real fuel cells are not perfectly isothermal, they need to maintain approximately uniform temperature distributions in order to manage thermal stresses. MHI's plots of temperature over the cell length are relatively flat and centered at 920℃ when operating at .775 volts [139]. This voltage and temperature are set as constants in the code. The model also includes two significant processes for SOFC functioning: thermal management and anode recycling. While neither is explicitly included in MHI's schematic (Fig. 7.1), there is significant supporting evidence to believe they are required. The modeling efforts for these two processes are discussed next.

[3]Start-up scenarios are not modeled for any of the plants discussed in this paper, as they are not central to the analysis of exergy efficiencies.

The MHI SOFC operates in the vicinity of 920℃, but the air entering the fuel cell box in Fig. 7.1 is only at 460℃. Additionally, no external transfer of heat is indicated on the MHI diagram. These data indicate that the compressed air must be preheated before entering the cathode and that this heat transfer is internal to the overall SOFC subsystem. This is not uncommon, as the exothermic fuel cell has to be cooled to maintain a reasonable temperature and could be a source of heat for the reactants. While it is possible to speculate about the layout of the SOFC tubes and the reactant and product support systems, it is not necessary for the purposes of this study. This analysis assumes that the fuel cell system has the ability to exchange heat between reactant streams, product streams, and the tubular structures themselves, in whatever manner results in net adiabatic behavior consistent with the overall schematic. This fulfills the goals of preheating the inputs to an appropriate temperature, maintaining the cell at this temperature, and cooling the products to a temperature appropriate for adiabatic combustion and turbine expansion. This abstracted model of the heat exchange processes is depicted in Fig. 7.2.

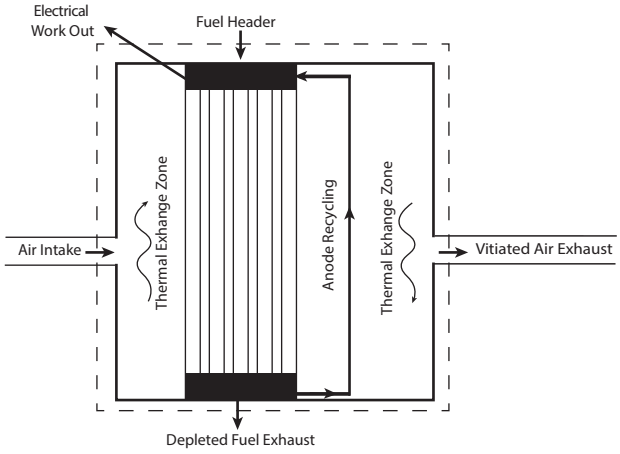

Figure 7.2: Depiction of a SOFC tube bundle enclosed in an adiabatic, pressurized vessel.

Anode recycling is another process assumed to occur in the MHI system. It

involves taking some of the partially-oxidized fluid from the anode outlet and recirculating it to mix with the fresh fuel at the anode entrance. This provides water and partially-reacted fuel to initiate the electrochemical reactions in the fuel cell. Additionally, anode recycling distributes the fuel reforming along the entire cell length, which evens out the regions of exothermicity and endothermicity. This helps in maintaining a more uniform temperature distribution. The recycling process requires a convergence loop on the mole fractions wrapped around the basic fuel cell computations. The quantity of anode recycling in the model is varied until the work output of the fuel cell matches that of the MHI system for a given excess air coefficient.

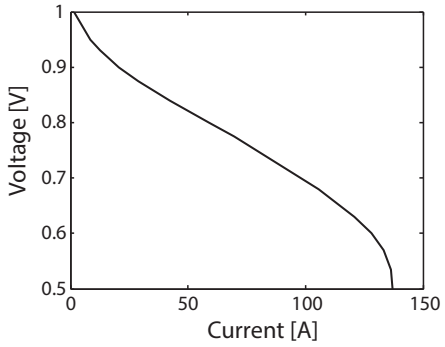

Figure 7.3: Voltage-current curve for a single tubular cell at 920℃, 20 bar

With the above set of modeling specifications, a voltage-current (VI) curve was made and is shown in Fig. 7.3. The plot indicates that the fuel cell losses are primarily ohmic, meaning it is mostly limited by the conductivity of the YSZ. This is in keeping with several studies that have shown that the conductivity of YSZ is severely decreased in an SOFC as compared to measurements made in isolated YSZ [132, 140]. This decrease is attributed to NiO from the anode dissolving in the zirconia, thereby driving an irreversible transition from the cubic phase of zirconia to a mixture of cubic and tetragonal phases. This mixture of phases was shown to reduce the conductivity to approximately 50-60% of the original value measured in the absence of Ni [132]. Accordingly, for this analysis, the conductivity of the YSZ was set to 3.5 S/m at 920℃.

7.1.3 Gas Turbine Engine Model

The reverse engineering process accomplished from Fig. 7.1 was successful in determining the gas turbine engine states. However, the exact location of the given 1600℃ stream and the probable use of blade cooling are unspecified. Most likely, the burner outlet at the given 1600℃ is cooled substantially in a nozzle before entering the turbine. Resolving this ambiguity is only critical in the current analysis if it will significantly impact the exergy distribution in the triple cycle. To test the sensitivity to the presence of blade cooling, the model was tested with both an uncooled turbine with an inlet at 1600℃, and with a steam-cooled turbine with an inlet at 1450℃. In the latter case, the heat recovery steam generator (HRSG) was used as the source of steam. The two cases ultimately differed by less than 2% overall efficiency on an LHV basis with no substantial differences in exergy distribution. Given this limited impact, blade cooling was neglected for the remainder of the analysis.

7.1.4 Steam Turbine and HRSG Model

Unlike in the gas turbine analysis, Fig. 7.1 was largely unhelpful in specifying the steam cycle. As such, the basic schematic and device specifications were modeled using a typical, state-of-the-art design currently being built in California [141]. This design involves a modified Rankine cycle with a three-pressure-level HRSG. Turbines and pumps are assigned 85% polytropic efficiencies. Pressure drops through the HRSG and condensor are estimated from the California plant's schematics. Their effects can be seen in the model's resulting thermodynamic state data summarized in Table 7.1.

The simplified schematic used for the model is seen in Fig. 7.4. All of the water enters the HRSG at 2.5 bar in the low pressure (LP) economizer. Of this water, 9.6% continues through the HRSG at this pressure into the LP evaporator and superheater. After the LP superheater, the water mixes with output from the intermediate pressure (IP) turbine and is expanded in the LP turbine to 6.8 kPa. The outlet of

the LP turbine is then sent through a condensor where it is assumed to be cooled against a non-exergetic stream until reaching a saturated liquid state. At this point, the water is pumped back up to be fed into the LP economizer.

The next 12.6% of the water entering the HRSG is further pressurized in the IP pump to 50 bar, while the remaining 77.8% of the water is pressurized in the high pressure (HP) pump to 150 bar. At these elevated pressures, the fluid is fed through the appropriately pressurized economizers, evaporators, and superheaters. The HP economizer is split into two sections in order to minimize the irreversibility of the heat exchange with the combustion products. The IP system includes two reheat sections that are interlaced with the two HP superheater sections in order to most effectively extract the enthalpy of the flue gas. Upon leaving the HRSG, the HP stream is expanded in the HP turbine where it is mixed with the output of the second IP reheater before expanding again in the IP turbine. Mass, energy, and exergy flows were tracked through the entire steam cycle and integrated into the quantification of the overall system performance.

While the California plant served as a useful template for the model, it was not configured for the same gas inlet state as exists in the triple cycle. As such, an enthalpy-temperature diagram, shown in Fig. 7.5, was used as a design tool to minimize heat exchange irreversibility. The pressures and mass flow rates chosen and reported above were set in order to minimize temperature differences between the GT flue gas and the water to the extent possible, given a 10K minimum temperature difference and reasonable material constraints. These final values reflect higher IP and HP pressures, and a greater percentage of the overall flow in the HP system than seen in the California plant. The higher pressures result in smaller enthalpies of vaporization, which reduce the extent to which the flue gas and water slopes diverge and create larger temperature mismatches. The flow rate shifts from the IP to the HP system, which results in smaller temperature differences in the high-temperature region of the heat exchanger. Both the pressure and flow rate changes result in increased turbine work over the baseline plant as well.

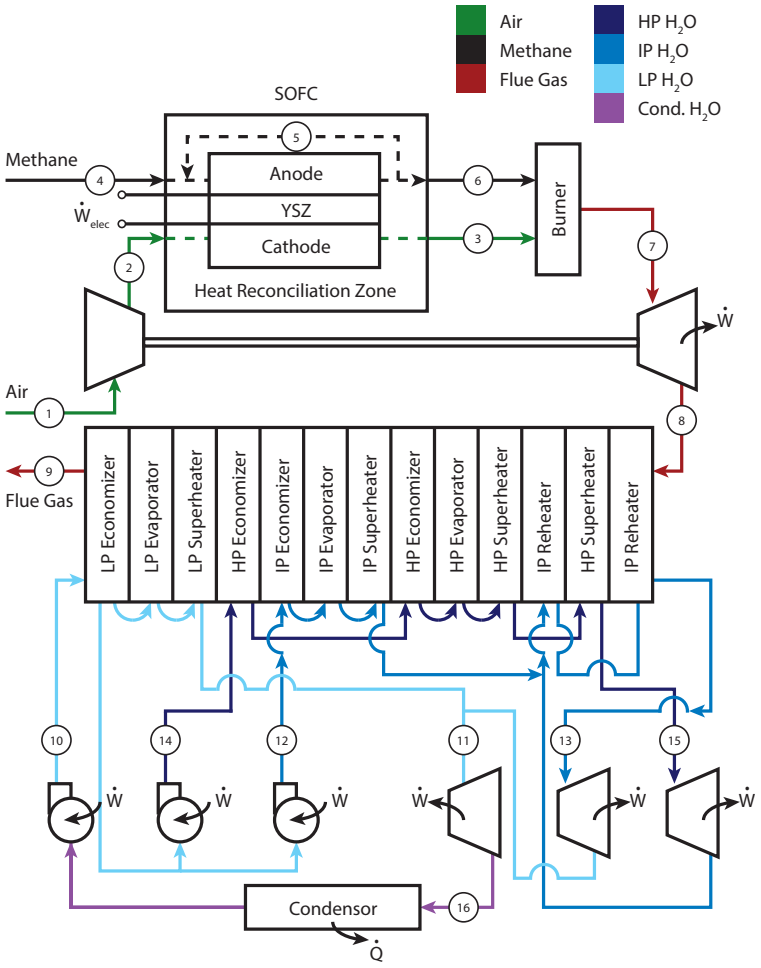

Figure 7.4: Triple cycle schematic for modeling the MHI system.

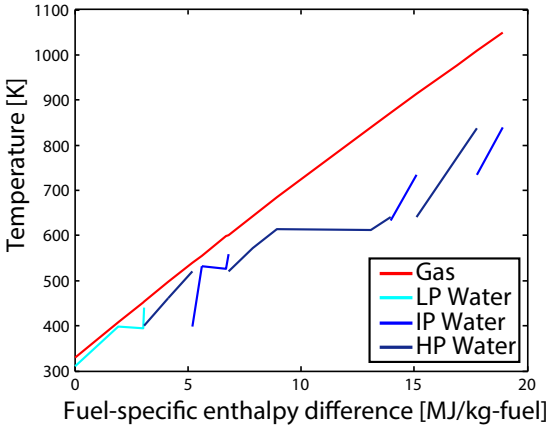

Figure 7.5: Temperature vs. fuel-specific enthalpy difference for the HRSG.

7.1.5 Model Validation

State	$\dot{m}$ kg/s	T K	P Pa	h kJ/kg	s kJ/kg-K
1	760.07	298.15	101325	-254.29	6.9929
2	760.07	733.15	1840100	207.30	7.0973
3	713.85	690.35	1840100	145.83	7.0387
4	33.126	298.15	1840100	-4650.1	10.115
5	6.6252	1193.15	1840100	-5468.3	12.715
6	79.345	690.35	1840100	-6712.5	11.370
7	793.19	1825.9	1766496	-540.22	8.6523
8	793.19	1048.8	101325	-1618.2	8.7470
9	793.19	328.63	101325	-2486.5	7.3706
10	189.93	310.02	250000	-15816	4.0504
11	189.93	482.62	198904	-13081	11.069
12	23.947	398.52	5000000	-15441	5.1008
13	171.67	840.00	3295407	-12366	10.895
14	147.72	399.82	15000000	-15429	5.1055
15	147.72	838.00	14016335	-12472	10.129
16	189.93	311.55	6760	-13559	11.297

Table 7.1: Thermodynamic data for select states, as numberd in Fig. 7.4

Thermodynamic state data for critical streams in the model are summarized in Tables 7.1 and 7.2. All mass flow rates are scaled so the mass flow rate of fuel is consistent between the model and the MHI plant. While two possible excess air coefficients were found in the reverse engineering of the MHI plant, ultimately a value of $\lambda = 1.25$ was used in the model. This coefficient was found to be more consistent with reasonable quantities of anode recycling while matching the work output of the fuel cell. It did result in a turbine inlet temperature approximately 47℃ colder than the burner temperature given in Fig. 7.1. However, given the uncertainties about the exact location of the reported 1600℃ and the mechanisms for blade cooling, this imprecision was considered to be within the overall model tolerance. Table 7.2 shows that the anode recycling (State 5) successfully introduces critically needed H_2O and

State	CH_4	CO	CO_2	H_2	H_2O	N_2	O_2
2					3.00	76.63	20.37
3					3.17	81.02	15.81
4	100						
5	15.51	18.50	9.66	33.06	23.27		
6	15.51	18.50	9.66	33.06	23.27		
7			7.19		17.17	71.02	4.42

Table 7.2: Molar composition (%) at key states, as numbered in Fig. 7.4

H_2 to the front end of the cell. Streams 3 and 6 in Table 7.2 indicate that the fuel cell utilized approximately 69% of the incoming methane's LHV, implying that the combustor is then responsible for 31% of the fuel energy conversion.

Table 7.3 shows the work breakdown for the SOFC-GT-ST model. For comparison, the MHI work flows are also listed, as made available in Fig. 7.1. The model was scaled to have the same fuel cell capacity as the MHI system so that the two columns are a direct comparison. The overall efficiency and power distribution for each cycle is nearly identical between the model and the MHI system. This indicates that the modeling assumptions were not unreasonable for the purposes of tracking energy flows. Additionally, it implies that the cycle schematics, internal processes, and device parameters chosen for the model are likely similar to those used by MHI. This is particularly remarkable for the HRSG, for which there was almost no information.

7.1.6 Exergy Analysis

An exergy analysis aids in further understanding the significant efficiency gains in the SOFC-GT-ST over its constituent systems. Figure 7.6 shows the ultimate distribution of the incoming fuel exergy for each of the systems analyzed. All plants are fueled with methane, have a maximum gas system pressure of 18.4 bar, and a maximum temperature of 1553℃, as was dictated by the triple-cycle model. The dead state for all exergy analyses was taken to be 75.65% N_2, 20.29% O_2, 3.12% H_2O, 0.90% Ar and 0.04% CO_2 at To = 298.15 K and Po = 101325 Pa.

Component	Power (MW)	MHI Power (MW)
SOFC		
Fuel Cell	432	
Inverter	-21.6	
Net	410	410
Gas Turbine		
Compressors	-351	
Turbine	855	
Generator	-5.04	
Net	499	500
Steam Cycle		
Pumps	-2.92	
Turbines	259	
Generator	-2.59	
Net	253	250
Overall	1162	1160
Efficiency	70.2%	70+%

Table 7.3: Work flow for the triple-cycle plant.

Gas-Turbine Engine

The stand-alone gas turbine engine has 24.3% of the exergy of the incoming fuel lost to combustion irreversibility. The primary way to reduce this entropy generation is to further preheat, and to a lesser extent, pre-pressurize, the reactants. There is a nearly linear relationship between combustor reactant and product temperatures. Thus, temperature limits on turbine blades impose a limit on how much the reactant temperature can be increased without also increasing the flow of a non-reacting moderator. This moderator is most often excess air, which has to be compressed, thereby reducing the power of the system. An even larger problem than the combustion irreversibility in the GT engine is the 31.0% of the fuel exergy exiting the system in the exhaust. This wasted resource is largely thermal and is the basis for adding a bottoming heat engine in the GT-ST system.

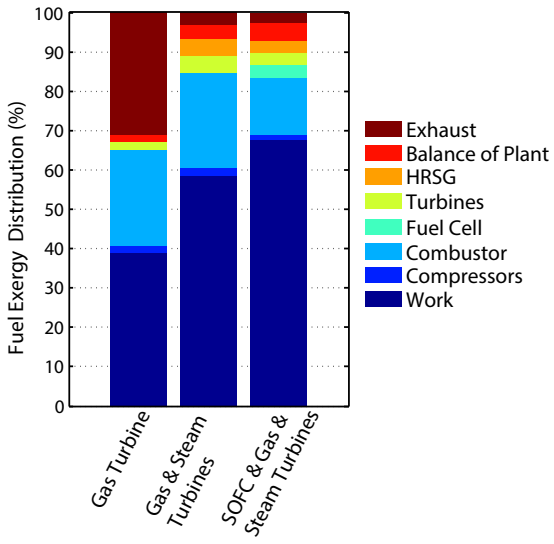

Figure 7.6: Exergy distribution for analyzed plants.

Gas-Turbine, Steam-Turbine Engine

The GT-ST system effectively utilizes the GT's exhaust as an exergetic resource to power a bottoming steam cycle. Thus the LHV efficiency is approximately 20% higher in the GT-ST than in the GT engine. However, the GT-ST has the same problem with combustion irreversibility, and also introduces new sources of irreversibility in the steam turbines, condensor, and HRSG.[4] The three pressure-level HRSG shows a significant improvement over its single-pressure counterpart, destroying only 4.3% of the total fuel exergy. As such, it is not a primary source of irreversibility. Another 3.1% of the fuel exergy exits the system in the exhaust. Imperfect heat exchangers that must avoid pinch limit the extent to which this can be further decreased. The overall GT-ST LHV efficiency reported here is 61.3%. This is a little high for the state-of-the-art today and can be explained by the absence of accounting for blade cooling in the turbine [82]. In a real system, steam from the HRSG or excess air from the compressor could be used for cooling with a small, but non-negligible efficiency

[4]The condensor, mixers, generators, and the fuel cell inverter, in the case of the triple cycle, are grouped in the "Balance of Plant" bar in Fig. 7.6.

penalty.

Solid-oxide Fuel Cell, Gas-Turbine, Steam-Turbine Engine

The SOFC-GT-ST system essentially manages to put another engine upstream of a natural gas combined-cycle plant. Not only does the SOFC avoid interference with the NGCC, but it actually manages to improve the performance of its linked cycles in several significant ways. The temperature requirements of the fuel cell result in a high reactant temperature for the combustor. Typically, this high temperature would be problematic for a combustor, as it would lead to a dangerously high turbine inlet temperature. However, the fuel cell accomplishes three critical tasks that allow the burner to react without violating material constraints of the downstream turbine. First, the fuel cell removes energy from the reacting fluid by extracting work. This lowers the enthalpy in the combustor. Second, because the fuel cell has already partially oxidized the fuel, the combustor has less exothermic reaction to accomplish. Third, the highly dilute fuel cell products provide moderator for the combustor. These three features allow the combustor to successfully operate in a highly energetic regime without harming the turbine. In keeping with the extreme states principle, this results in only 14.5% of the fuel exergy being destroyed in the combustor, as opposed to 24.2% for the GT-ST [14].

Additionally, because the fuel cell creates moderator for the combustor in the form of water and carbon dioxide, this system requires significantly less excess air than a standard gas-turbine system. Thus, the SOFC-GT-ST operates much closer to stoichiometric than a standard GT system. Where an excess air coefficient of 1.25 was used in the triple cycle, a value closer to 2.5 would be typical for a gas turbine. By achieving effectively the same outlet pressure and temperature from the combustor, but at a lower air/fuel ratio, the SOFC-GT-ST also has the lowest back-work ratio of the systems analyzed. The reinvested compressor and pump work is only 31.7% of the triple-cycle turbine work, as opposed to 33.0% and 42.3% for the double and single cycles, respectively. By reinvesting less work in air compression,

the triple cycle produces more power per unit air flow rate.

While the fuel cell process itself is not perfect, it is a restrained process, unlike the unrestrained combustion process. This means that marginal changes in work can be achieved by marginal changes in load. This is in direct contrast to the combustor for which there is no control on the outputs once the reactants enter the device. Thus, the fuel cell irreversibility accounts for only 3.4% of the incoming fuel exergy, while utilizing 69% of the LHV of the fuel. This is less than a quarter the destruction of the combustor for more than twice the chemical energy conversion.

7.1.7 Modeling Conclusions

A SOFC-GT-ST triple-cycle power plant is in development by Mitsubishi Heavy Industries. This plant is remarkable for both its 70% reported efficiency and its combination of fuel cell and combustion technologies. An analysis was carried out to understand how the various subsystems interact to produce a net result 10 percentage points more efficient than the current state-of-the-art NGCC plant. A computer model of the triple cycle was built, incorporating insights gained from reverse engineering MHI schematics. The model accurately reproduces the overall efficiency and power of the proposed system. The air systems were found to be integrated in a manner fundamentally different than simple cascading of independent systems, as is done with the natural gas combined cycle. In many ways, the integrations mutually enhance the linked systems' thermodynamic performance. Most significantly, the fuel cell's combination of high temperature and partially oxidized products prove to be an excellent reactant stream for the combustor. This system linkage also results in an excess air coefficient near unity, which reduces the back-work ratio, thereby increasing the overall system air-specific power. Because the gas turbine is inlet temperature-limited, its output thermal exergy is not significantly altered by the SOFC. Thus the bottoming steam cycle is effectively unchanged from a stand-alone NGCC plant.

7.2 Systematic Addition of SOFCs

While the MHI system's reported 70% efficiency is impressive, it is not necessarily optimal. The thermodynamically optimal configurations and limitations of this cycle are yet to be fully understood. Using the systematic thermodynamic approach from Chap. 4, the SOFC-GT engine is built up from a bare-minimum sequence (CBT), adding one system component as a degree of freedom at a time. After the SOFC is added to the minimum sequence, subsequent degrees of freedom will allow internal and external heat transfer. The specification of all devices ultimately considered in this analysis are listed below. Because the concern of this work is maximizing exergy efficiency, the analysis is restricted to devices critical for following the net energy flows in the system. This eliminates support systems that would be essential to a physical plant, but are minimally important from an overall energy standpoint (e.g., air filters, start-up support systems).

C Compressors are used for both fuel and air mixtures. They are modeled as adiabatic with a polytropic efficiency of $\eta = 0.9$

F Solid-oxide fuel cells are described extensively in the previous section.
Fuel utilization is defined as 1 -($LHV_{fuel,orig}$))/($LHV_{SOFC,out}$).
For operational constraints, 800℃ $\leq T_{cell} \leq$1000℃, and 10% $\leq$ utilization $\leq$ 90%.

I Inter-coolers reject heat with a 95% effectiveness to the atmosphere between stages of compression. If multiple intercoolers are used, they are implemented in evenly spaced pressure increments.

B Burners are adiabatic, isobaric, and completely oxidize the fuel.

T Turbines, like compressors, are adiabatic and have polytropic efficiencies of 0.9. The maximum inlet temperature for a turbine is taken from MHI to be 1600℃.[5]

[5]This ultimately is irrelevant for the optimal system because the vast majority of the reaction

X Heat exchangers are specified *in situ* as either X_{in} or X_{out} depending on whether the stream is receiving or donating energy. All heat exchangers will be considered two-stream, counterflow devices with an effectiveness of 90%. A minimum 10℃ pinch is enforced.

7.2.1 Adding a Degree of Freedom: A Solid-oxide Fuel Cell

As discussed in Chap. 4, the starting engine sequence for optimization is the compressor, burner, turbine (CBT). Now the possibility of adding solid-oxide fuel cells is introduced. Fuel cells could be interspersed anywhere in the existing sequence so that the new sequence is

$$(CF)_k(BF)_m(TF)_n$$

where k, m, and n are arbitrary integers to allow for device repetition. Many permutations can be immediately removed due to the various device constraints and operating parameter domains. There is no value in having a solid-oxide fuel cell downstream of a burner.[6] The combustion reaction has a very high affinity and equilibrium constant, so the products out of the burner will be essentially fully oxidized. There is nothing for the fuel cell to do and it essentially behaves as a pipe. Thus the fuel cell stages should only exist before the burner and the candidate sequence simplifies to

$$(CF)_kBT$$

The solid-oxide fuel cell membrane needs to operate between 800℃ and 1000℃. This membrane is encompassed in a larger chamber with thermal exchange so the inlet does not have to be quite that hot, but it still needs preheating from ambient conditions. Since heat exchange is not yet allowed, this preheating can only be provided by

occurs in the SOFC and the remaining minimal exothermicity in the burner does not result in such high temperatures.

[6]Recall that the degree of freedom associated with matter streams splitting and remixing is not open. If it were, then it would be possible to have a fraction of the air or fuel go into the burner and a fraction bypass it and enter the fuel cell directly.

compression, which then couples the pressure ratio to the fuel cell inlet temperature. With at least one mandatory stage of upfront compression, the candidate sequence can be written as

$$C(FC)_kBT$$

Because the fuel cell has to operate in a relatively narrow temperature range, there is limited ability to add adiabatic compressors between fuel cell stages. Any compression after a fuel cell stage would start quite hot with a high specific volume and therefore a low efficiency. It would end even hotter and likely violate the upper bound on the fuel cell temperature. There is no advantage to interspersing compression between fuel cell stages over doing all of the compression beforehand. The sequence is then

$$CFCBT$$

The final question to ask is whether there is value in having a compressor between the fuel cell and burner. The argument in favor of compression would be to further increase the energy of the reactants and thereby reduce entropy generation in combustion. However, the energy of the reactants is limited by the turbine inlet temperature constraint. Starting with nearly 1000℃ gas leaving the fuel cell, there is little ability to add further energy to the system through work without violating this constraint. Additionally, any compression done at this point would be inefficient due to the high specific volume. Thus, the sequence has been simplified to

$$CFBT$$

There are no heat interactions allowed between this engine and the environment. And at least initially, the inlet flow rates of fuel and air can be fixed. These were the same criteria in Chap. 4 for allowing the use of exhaust entropy as a proxy to track total entropy generation. The attractor method is now applied to the $CFBT$ system to optimize parameter lengths.

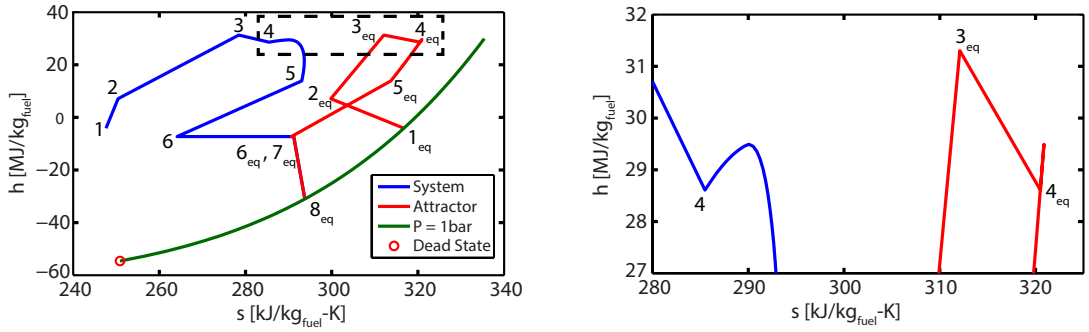

(a) System-attractor trajectory for a $CFBT$ (b) Detailed view of boxed region.

Figure 7.7: System-attractor trajectory with detailed view at entrance to fuel cell.

Trajectory Shape

Figure 7.7a shows a typical system-attractor trajectory for the $CFBT$ configuration. Progressing through it starting with the ambient air-fuel system in blue at the far left, the processes are as follows:

I. Ambient air and fuel enter the system at State 1 on the blue line. If they were immediately chemically equilibrated without further energy interactions, they would land at the corresponding State 1_{eq} on the red line.

II. Compression of the air and fuel occurs between States 1 and 2. The compressor looks linear here because only the end points are plotted. The attractor progresses to State 2_{eq}.

III. State 3 denotes the end of the pre-heating of reactants through internal heat exchange in the F. Again, this heating is not a defined process, but rather the result of imposing an energy balance for the fuel cell subsystem.

IV. State 4 marks the inlet to the fuel cell channel. This state does not align with state 3 because mass-transfer through anode recycling occurs in between. State 4 is thus the result of mixing the fresh fuel stream at State 3 with the recycled mass stream from State 5.

V. State 5 occurs at the end of the fuel cell channel.

VI. After the required heat exchange within the fuel cell and mass extraction for recycling, the system is at state 6.

VII. State 7 and 7_{eq} denote the outlet of the combustor. At this point the attractor (red) and system (blue) lines converge.

VIII. Finally, the system and attractor end at State 8_{eq} at atmospheric pressure following the turbine expansion.

Figure 7.7b shows a detailed view from Fig. 7.7a centered at the entrance to the fuel cell channel. In this diagram, it is clear that the attractor entropy does not immediately drop, but rather rises at the front end of the fuel cell channel and then starts to fall. This is due to internal reforming of methane. The initial few centimeters of the fuel cell are largely dominated by the rate-limiting endothermic process of breaking down methane into hydrogen and carbon monoxide. Only after this has been initiated can the exothermic electrolytic reaction occur, which produces work. Thus, the fuel cell enthalpy rises and then falls with the accumulated extraction of heat and work.

Several operating parameters of the $CFBT$ were examined for their effect on fuel cell performance and overall system efficiency. These are discussed individually in the following sections.

Temperature

The investigation of optimal temperature begins with the attractor diagram in Fig. 7.8. Note that this graph is of the same form as Fig. 7.7a, but the system trajectory was eliminated to emphasize that the analysis is based solely on the attractor. The attractor indicates a significant effect of temperature on efficiency, where the highest end of the allowable temperature range has the lowest final entropy. Because all cases started with the same initial entropy, the lowest final entropy is equivalent to

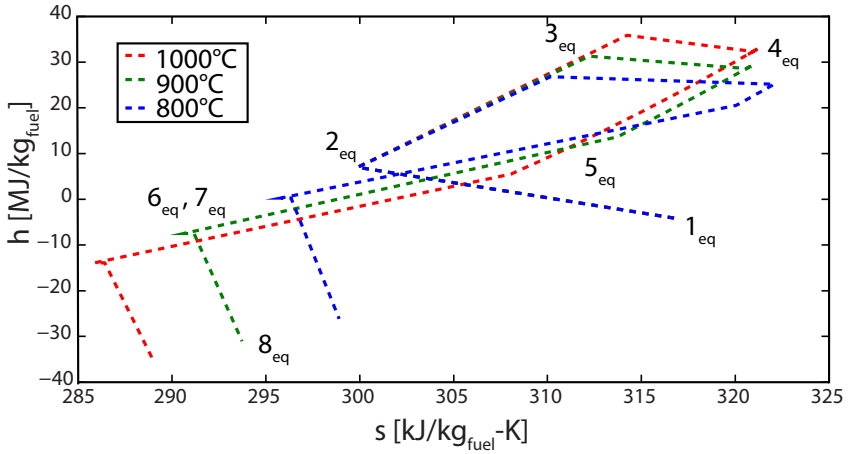

Figure 7.8: Temperature sensitivity on an *h*-*s* attractor diagram

the lowest entropy production. Interestingly, the geometry of the graph indicates that there are two primary reasons for this result. The first is that the fuel cell trajectory from State 4_{eq} to 5_{eq} lengthened. The fuel cell physical length was held constant at 0.5 m for all cases. Thus, a lengthened fuel cell trajectory indicates more work per unit length is extracted from the fuel cell at higher temperature. This is expected given that the internal resistivities and the chemical kinetics of the fuel cell are improved with increased temperature.

The fuel cell exergy destruction plots in Fig. 7.9 further elucidate the matter. More oxidation is occurring in the fuel cell at higher temperatures. This increased oxidation leads to proportionally higher levels of destruction, despite being beneficial to the overall system efficiency. At 800℃, hydrogen traverses the anode-side gas diffusion layer slower than at 1000℃.[7] At 1000℃, methane breaks down faster, forming hydrogen faster, than at 800℃. As such, at the highest temperature, the cathode gas diffusion overtakes the anode-side diffusion as a more predominant issue.

The second effect the attractor indicates is that the entropy generation due to

[7]The "wiggles" in Fig. 7.9a are result of the numerical handling of the steam methane reformation kinetics at very slow reaction rates. This effect is not associated with a significant macrosopic change in energy distribution in the system.

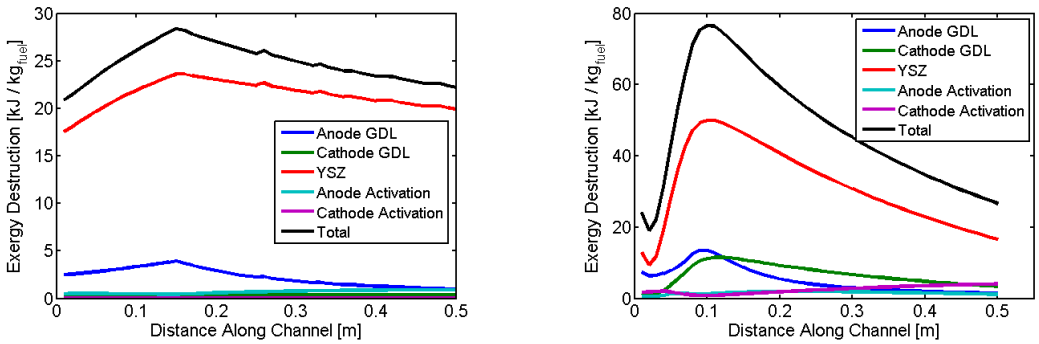

(a) T = 800℃ (b) T=1000℃

Figure 7.9: Effect of fuel cell temperature on internal exergy destruction profile. P = 10 bar, λ=2, recycle ratio = 0.5, length = 0.5 m

mixing of the fresh fuel and the recycled anode output is lower at higher temperature. This is indicated by the path length from State 3_{eq} to 4_{eq}. All three temperature cases have the same recycling ratio of 0.5: 1 kg/s recycled mass per 2 kg/s fresh fuel. Thus, this difference in length of process is an intensive effect of the dependencies of chemical equilibration on temperature [142].

Figure 7.10 confirms the attractor-based findings. Figure 7.10a shows that increased temperature improves fuel utilization regardless of system pressures. Because the fuel utilization is so low for the 800℃ case, Fig. 7.10b shows that the system behaves mainly like a gas turbine engine, and thus improves with pressure over the plotted range. Figure 7.10b also confirms the link between fuel cell utilization and overall LHV efficiency, and their mutual improvement with increased temperature. However, this graph does not make it clear that the system is also improving by the reduction in mixing entropy at the entrance to the anode channel.

Pressure

At left, Fig. 7.11 shows the attractor trajectory for three different system pressure ratios: 5:1, 15:1, and 30:1, all at a fixed fuel cell channel length of 0.5 m. The lowest pressure case has a significantly lower final entropy and correspondingly higher

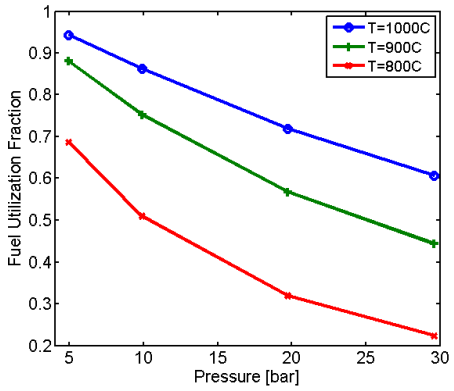

(a) Effect of temperature, pressure, and excess air coefficient on utilization

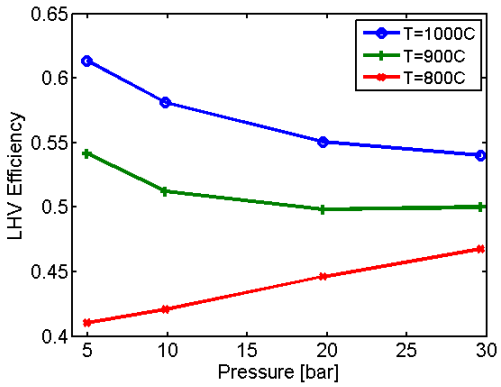

(b) First-law efficiencies of displayed systems

Figure 7.10: Examination of $CFBT$ sensitivity to temperature, pressure, and excess air coefficient at fixed cell length. λ=2, recycle ratio = 0.5, length = 0.5 m

efficiency than the other cases. (62.37% LHV efficiency versus 56.16% for the 15 bar, and 53.67% for the 30 bar cases). This is largely due to fuel utilization, seen as the length of the attractor path between States 4_{eq} and 5_{eq}. At fixed channel geometry, an increase in pressure, and therefore density, can either lead to a greater mass flow rate at constant velocity or a slower velocity at constant mass flow rate. Here the velocity is fixed at 0.1 m/s, so the mass flow rate changes with pressure.[8] An increased pressure and mass flow rate decrease the utilization over a fixed fuel cell length. (If alternatively the mass flow rate had been fixed and the velocity reduced with increased pressure, the utilization would have increased.)

The pressure also affects the temperature of the post-compressor reactants. The lower pressure cases rely more on the internal thermal energy exchange within the fuel cell to bring the reactants up to the membrane temperature. This reduces the enthalpy at the outlet of the fuel cell system, since more energy is donated to the inlet of the system. As a result, the line from State 5_{eq} to 6_{eq} is longest in

[8]As pressure changes, the tube diameter for a set velocity changes. Once the thermodynamics of the architecture are fully understood, the logistics and economics of this issue could then be assessed.

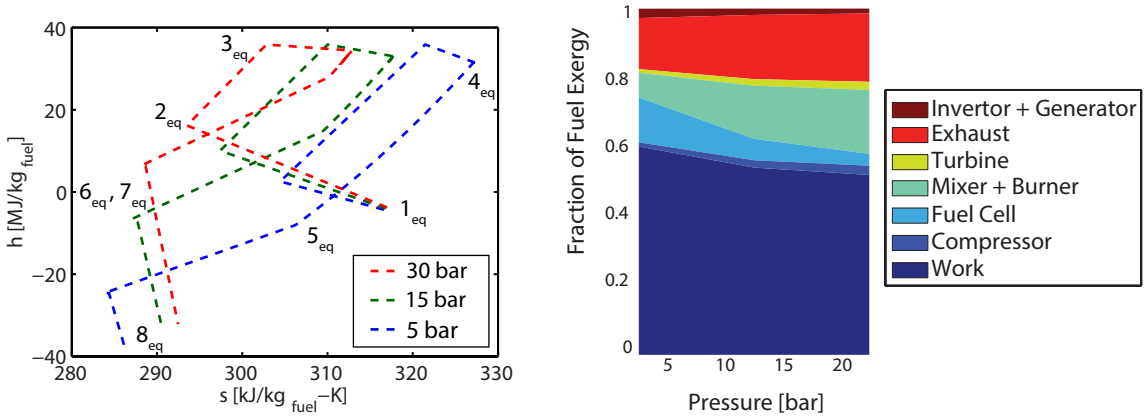

Figure 7.11: Variation of efficiency with pressure. T = 1000℃, $\lambda = 2$, Length = 0.5 m, $R = 0.5$

the 5 bar case. The reduced enthalpy out of the fuel cell lowers the turbine inlet temperature, and therefore the turbine outlet temperature. This reduces the exhaust exergy destruction. These positive consequences of lowered pressure outweigh the preference of the gas turbine system to operate at higher pressure ratios.

The bias toward improving the fuel cell over the gas turbine is highlighted in the graphic on the right of Fig. 7.11. This figure shows the distribution of the fuel exergy as the pressure changes for the system. At low pressures, the fuel cell is doing the majority of the oxidation and has correspondingly more destruction associated with it than the burner does. As the pressure is increased, this balance changes and the burner does proportionally more chemical reaction. The marginal increase in exergy destruction of the burner is greater than the marginal decrease for the fuel cell, leading to an overall less efficient system at higher pressure. This result, without any of the explanation, was evident in the parametric plots of Fig. 7.10.

If the fuel cell channel length increased, the balance between pressure preferences to reduce entropy generation in the fuel cell versus in the gas turbine would change. The effect of channel length is shown next.

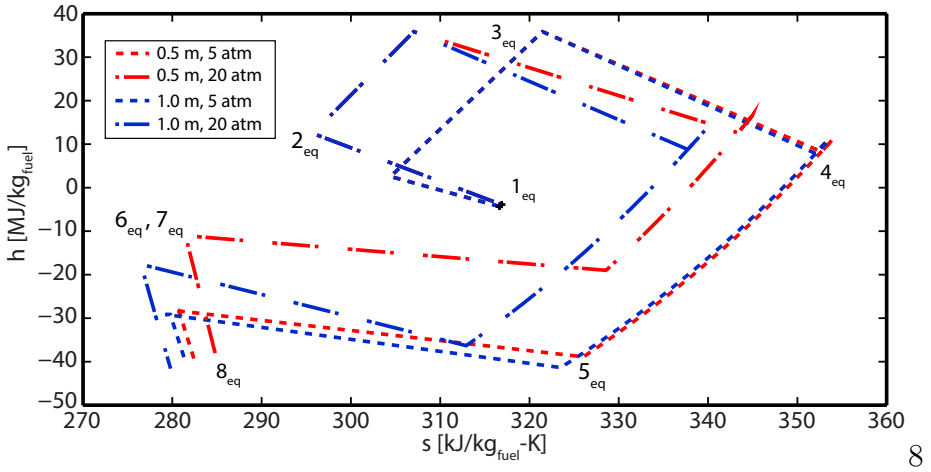

Figure 7.12: Attractor trajectories. T = 1000℃, λ=2, recycle ratio = 3

Channel Length

At fixed flow velocity, increased SOFC channel length increases residence time in the fuel cell. Longer residence time leads to greater fuel utilization in the SOFC. In keeping with the derivation in Chap. 4 that dS_{eq} is negative for an exothermic fuel cell and zero for a burner, increased length has a profound effect on overall system efficiency.

Figure 7.12 shows an enthalpy-entropy plot of the attractor trajectories for various fuel cell lengths and peak pressure combinations. At any pressure, increased length lowers the final entropy state and increases efficiency. This difference is more significant at high pressure where the fuel utilization is decreased. At low pressure, the vast majority of the fuel is oxidized within the 0.5 m fuel cell. Thus the additional 0.5 m has minimal benefit. At high pressure, the added length has a more significant effect on fuel utilization and efficiency. The optimal pressure depends on the chosen length. Here, the 5 bar case is more efficient for a short membrane, but the 20 bar case is more efficient at the 1.0 m length. Overall, the combination of high pressure and long fuel cell has the greatest efficiency because it allows high oxidation in the fuel cell and significant expansion in the turbine.

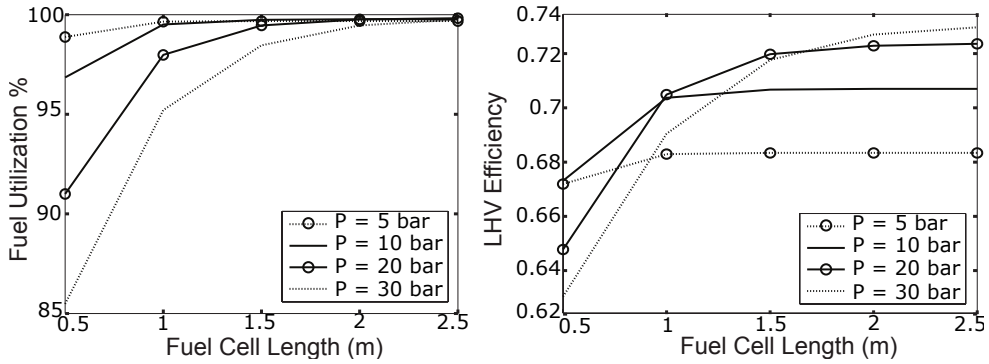

Figure 7.13: Optimal $CFBT$ varying pressure. T = 1000℃, λ =2, recycle ratio = 3.

These effects are confirmed in the parametric plots in Fig. 7.13. At short lengths, the cell is utilization-limited, and this is the driving force for overall efficiency. At long lengths, the preferences of the gas turbine engine dictate the relative efficiency of the different pressure cases. As the fuel cell lengthens, there are decreasing marginal gains in efficiency because the vast majority of the fuel has already been oxidized. There will be a practical limit based on size and cost in a real system. Thus, for most of the remaining analysis, a 0.5 m length system is assumed, typical of a Siemens Westinghouse SOFC [102].

Excess Air

Conventionally, gas turbine systems have significant quantities of excess air beyond the stoichiometric requirement to oxidize the fuel. In addition to allowing for complete oxidation of the fuel, the air serves as a moderator to keep peak temperatures down in the combustor and thereby avoid damaging the turbine blades. For the SOFC-GT, this is less of a concern, as the reactions in the fuel cell provide sufficient water and carbon dioxide as moderator for the burner. There is a trade-off in selecting the optimal excess air coefficient. On the one hand, adding air raises the back-work ratio and associated energy penalty. On the other hand, additional air improves fuel cell efficiency by keeping the chemical potential of oxygen at the cathode high.

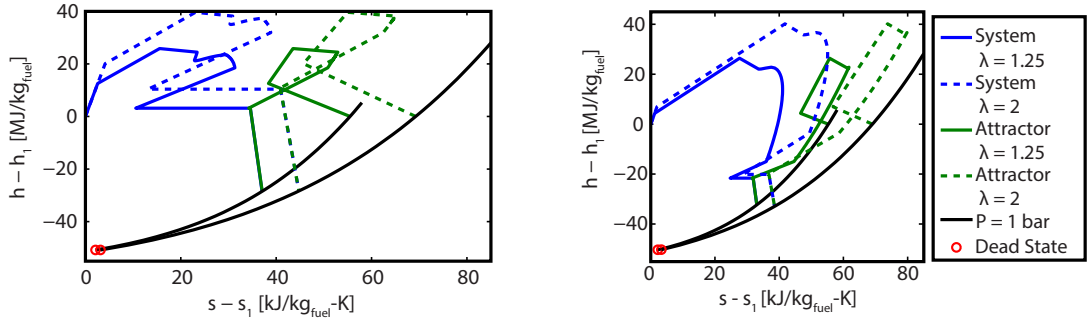

(a) P = 30 atm (b) P = 5 atm

Figure 7.14: Variation of system efficiency with excess air coefficient λ. T=1000℃, recycle ratio=0.5, L=0.5 m

Analyzing excess air adds a small complication in the graphical attractor analysis. As discussed previously, minimizing internal irreversibility, as done visually with the attractor, no longer guarantees minimization of external irreversibility. This is a result of the changing composition at the inlet state with λ. As seen in Fig. 7.14, the attractors for λ=2 and λ=1.25 no longer end on the same atmospheric pressure isobar. Minimizing the entropy difference between final and initial states no longer guarantees optimality. This problem could be dealt with analytically by expanding the entropy generation equation to include external losses [3]. For the current purposes, Fig. 7.14 still provides sufficient information to compare trajectories. The difference between the initial and final enthalpy on the vertical axis is a measure of total work output per unit mass fuel. With fixed fuel mass, the greatest difference $h-h_1$, corresponds to the most efficient system. In the plots shown, this corresponds to ending at the lowest possible $h_{final} - h_1$.

Figure 7.14a and b show the effect of varying excess air coefficient at 30 atm and 5 atm, respectively. While the difference between the λ =2 case and the λ=1.25 case is quite small on both graphs, the final plotted points indicate that the lower excess air coefficient is superior at higher pressure and the higher excess air coefficient is superior at lower pressure. Figure 7.15 verifies the attractor results parametrically. At the lower pressure ratio, the penalty for compressing excess air is minimal, so it is

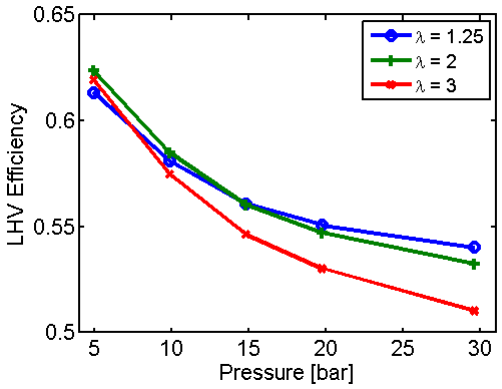

Figure 7.15: Parametric variation of efficiency with λ. T = 1000℃, recycle ratio = 0.5, length =0.5 m

worth sending excess air to the fuel cell. As the pressure increases, the balance point changes, and the system benefits from reducing the air flow rate. For this study, even if the pressure dictates a lower λ, the minimum excess air coefficient is restricted to 1.25 to ensure that complete oxidation occurs.

Recycling Ratio

Recycling is used in the SOFC by extracting some mass from the outlet of the anode-side channel and redirecting it to the inlet of the fuel cell anode to be mixed with fresh fuel. Recycling brings water necessary for the decomposition of methane to the inlet of the fuel cell. It also provides partially decomposed fuel (i.e., hydrogen) that can be directly oxidized at the anode, unlike the methane which must first decompose. This allows the first few centimeters of membrane length to be used effectively.

In Fig. 7.16, it is seen that adding recycled mass makes a significant difference in the attractor trajectory. The attractor segments associated with mass transfer into the fuel cell, 3_{eq} to 4_{eq}, and out of the fuel cell, 5_{eq} to 6_{eq}, become longer for the higher recycling ratios. This is dominated by the addition of the recycling mass to the system. This mass ultimately leaves the fuel cell for recirculation at the same temperature, pressure, and composition it entered with, so the transfer itself is not

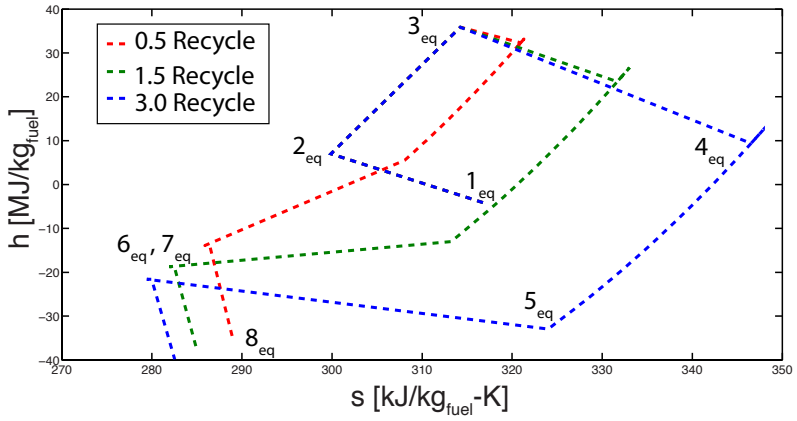

Figure 7.16: Changes to the attractor with recycling ratio, defined as $\frac{m_{recycled}}{m_{newFuel}}$. T = 1000℃, P = 10 bar, $\lambda = 2$, length = 0.5 m

the cause of change in overall efficiency. The long segment from 4_{eq} to 5_{eq} associated with the fuel cell attractor in blue, versus those in green and red, indicates that the fuel utilization has gone up significantly. Efficiency goes up accordingly, with the 3 kg recycled mass per kg of new fuel resulting in significant decreases in final state entropy over the lower recycling rates.[9]

Large amounts of recycling provide partially decomposed fuel at the front end of the cell. This smooths out the chemical affinity of reaction along the length of the channel. Since $\dot{S}_{gen}$ is proportional to the affinity, removing extremes tends to reduce entropy generation. This is an analog to the advantage of a counterflow heat exchanger over a co-flow heat exchanger in that reducing peak temperature differences increases efficiency of thermal energy exchange. Through the O_2 and H_2 mole fractions, Fig. 7.17a,b visualizes this difference. At low recycling ratio, the fuel exists almost entirely as methane at the front end of the channel. As this methane starts to break down, there is a large swing in the mole fraction of hydrogen, peaking above 50% . Compare this to the high recycling ratio case in which there is

[9]The estimated pressure drop along the length of the fuel cell is estimated at less than 0.01 bar. As such, even large mass flow rates of recycled fluid can be recirculated with an ejector and a negligible pumping penalty.

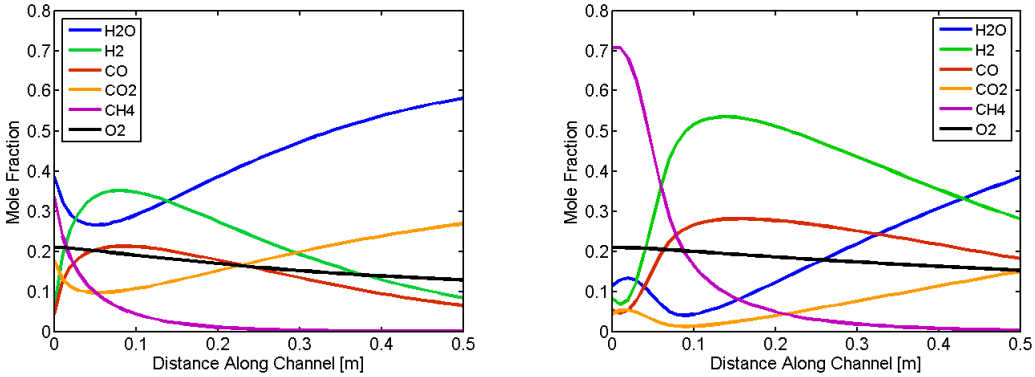

Recycling ratio = 3 Recycling ratio = 0.5

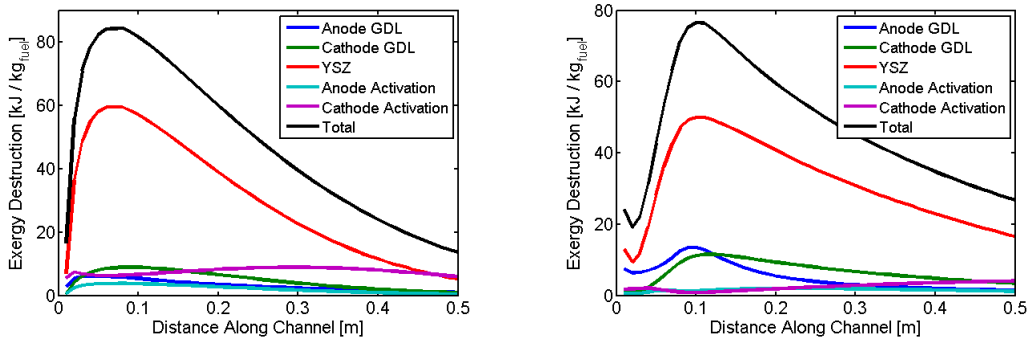

Recycling ratio = 3 Recycling ratio = 0.5

Figure 7.17: Effect of anode recycling on molar flow rates and on internal destruction.T = 1000℃, P = 10 bar, $\lambda = 2$, length = 0.5 m

less methane upfront, hydrogen is produced faster, and peak hydrogen mole fraction is below 40%. The higher recycling ratio with the less dramatic variation in hydrogen mole fraction results in a smaller average difference between the hydrogen and oxygen concentrations. Thus, while the same amount of methane is available in both cases, the chemical affinity across the cell has a smaller peak in the high recycling case.

Figure 7.17c,d show that different processes are limiting the efficiency in the two recycling cases. This is similar to the effect observed with changes in temperature in Fig. 7.9. In the 300% recycling case, significantly more oxidation occurs, and the overall quantity of destruction is correspondingly higher. When there is very little recycling, the inlet at the anode-side channel is almost entirely methane. As such, the cell is limited in its ability to move species across the anode gas diffusion layer and to move an oxygen ion through the YSZ. Over time (channel length), some of the methane reforms and the system is less restricted by the anode GDL and more restricted by activation at the cathode. This becomes the most limiting process toward the end of the fuel cell channel. The high recycling case essentially jump-starts Fig. 7.17c to begin where Fig. 7.17d ended. There is plenty of partially oxidized fuel at the anode-side inlet and the system is no longer limited in transporting molecules across the anode gas diffusion layer. It is now limited by the rate at which it can break down oxygen and move it through the electrolyte to react with hydrogen.

Again, the parametric analysis in Fig. 7.18 corroborates the attractor analysis, with both fuel utilization and LHV efficiency increasing with recycling ratio. As was the case with the fuel cell length, the marginal benefits of increased recycling are decreasing. Practical limits will dictate the ultimate ratio used. For this study, a 3:1 recycling ratio seems a reasonable upper bound for feasible systems.

Optimal $CFBT$

From the above analyses, some characteristics of the optimal $CFBT$ system are deduced.

I. The fuel cell should operate at the high end of the allowable SOFC temperature

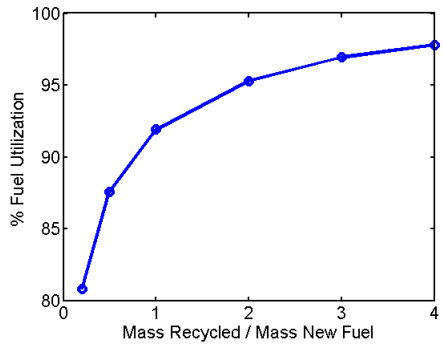

(a) Effect of recycling ratio on SOFC utilization

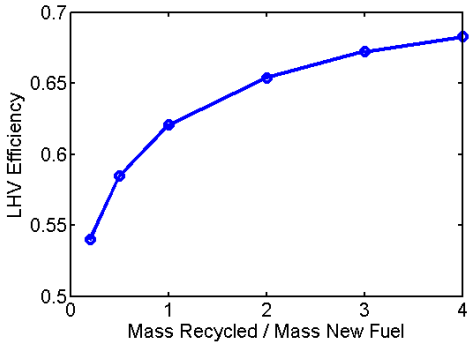

(b) First-law efficiencies of displayed systems

Figure 7.18: Examination of $CFBT$ sensitivity to recycling ratio. T=1000℃, P = 10 bar, $\lambda = 2$, length = 0.5 m

range.

II. The recycling ratio should be high.

III. Fuel cell length should provide sufficient residence time. The more fuel that can be oxidized in the fuel cell before being sent to the burner, the better.

IV. If sufficient fuel cell length were provided for near complete oxidation, then the pressure ratio would be optimized for the gas turbine system. However, at any given fuel cell length and flow velocity, lowering the pressure increases fuel oxidation. Thus, assuming the system is constrained by membrane length, the pressure should be kept low.

V. Again assuming limited length, a higher excess air coefficient will result in an increased chemical potential of oxygen on the cathode side of the membrane, which will speed the overall reaction. However, this benefit is counterbalanced by the cost of compressing additional air. As such, the appropriate choice for the air-to-fuel ratio will depend upon the chosen system pressure and the oxidation rate in the fuel cell.

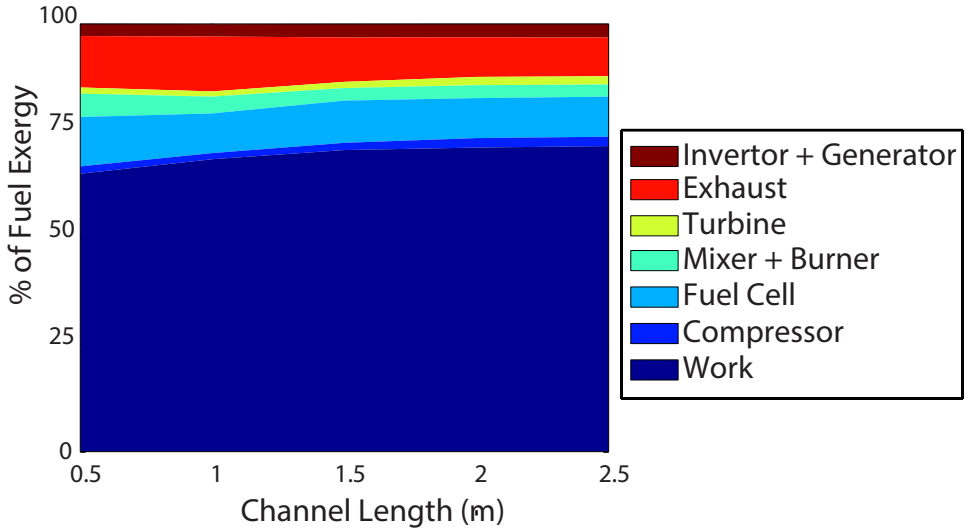

Figure 7.19: Optimal $CFBT$ varying pressure, λ over length. T = 1000℃, recycle ratio = 3.

The optimal operating conditions are largely dictated by the available fuel cell length. The length then determines the appropriate choice of pressure ratio and excess air coefficient. Figure 7.19 shows how the maximum system efficiency changes with membrane length. By decreasing the excess air coefficient and increasing the pressure ratio with increasing length, the fuel cell utilization can be kept at a near-constant, high value. Hence, the destruction in the fuel cell is consistent across all lengths. The burner improves with increased fuel cell length through a decreased chemical affinity. The turbine destruction increases slowly with increased expansion ratio, but this effect is less significant than the decrease in combustor destruction, leading to an overall efficiency improvement with increased length. It is significant to note at this point that this double cycle can achieve greater than 70% LHV efficiency. This is the efficiency reported by MHI for their triple cycle–an SOFC-GT with a bottoming Rankine cycle. To surpass this efficiency benchmark with the use of fewer engines is a significant achievement. However, the exhaust destruction for the optimal $CFBT$ cycle is still sizable. Figure 7.20 shows the exhaust temperature

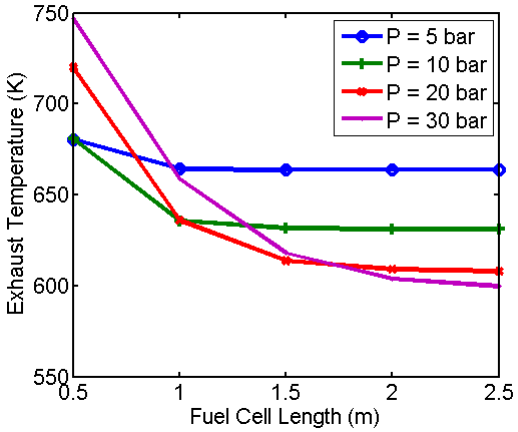

Figure 7.20: Exhaust temperature for various $CFBT$ scenarios

over a range of pressures and fuel cell lengths. The optimal pressure for each length is associated with the lowest exhaust temperature at each length. (Compare this to the same set of graphs in Fig. 7.13.) This implies that even in the most efficient scenario, there is significant thermal exergy in the exhaust that could be utilized to produce more work. The next architecture attempts to address this issue.

7.2.2 A New Degree of Freedom: Internal Heat Exchange

Heat exchangers are newly introduced to the system. These devices can be placed anywhere in the architecture sequence and are distinctly different than the internal thermal balancing within the fuel cell device. Heat exchangers are limited by a maximum temperature and therefore cannot be used between the chemical reaction devices.[10] Thus, the sequence can be written

$$(CX_{in})_k FB(TX_{out})_m$$

where k and m are arbitrary integers. As discussed in Chap. 4, the temperature difference between the exchanging fluids should be as low as possible to reduce entropy

[10] As discussed by Ramkrishnan, there is also no value to the engine in doing so [3].

generation. Given that the hot and cold sides are reasonably capacity matched, all of the heat exchange should be done between the post-compressor air and the post-turbine products, yielding

$$CX_{in}BTX_{out}$$

Figure 7.21 shows the H-S plot for this configuration operating over a range of viable SOFC pressures. The final entropy values for all cases are lower than those for the $CFBT$ systems. Thus, the recuperator does improve system efficiency. Interestingly, the efficiency increases from a 2 bar to a 5 bar system, and then decreases again. The attractor shows several competing pressure dependencies that lead to this result. First, the fuel cell utilization between States 4_{eq} and 5_{eq} decreases with pressure, as already discussed. This has the effect of reducing efficiency with pressure. Second, as the pressure is raised, the heating accomplished by the recuperator decreases, as shown from State 2_{eq} to 3_{eq}, and from 9_{eq} to 10_{eq}. This would also lead to reduced efficiency with pressure. Third, as a result of the decrease in recuperator usage at high pressure, the fuel cell must provide more of the thermal energy to preheat its reactants, lengthening the segments from 3_{eq} to 4_{eq}, and from 5_{eq} to 6_{eq}, which reduces entropy, thereby increasing efficiency with pressure.

The parametric plots in Fig. 7.22 corroborate this finding. As the pressure is raised, there is less of a temperature difference between the hot and cold sides of the recuperator (Fig. 7.22c). Eventually, this temperature difference goes below the minimum allowable pinch of 10K and the recuperator shuts itself off, as seen at the 12-bar point in Fig. 7.22b.

Figure 7.22d shows the trade-off that is occurring to cause a local maximum. As the pressure is raised, the exergy destruction in the recuperator decreases because the temperature difference between the hot products and cold reactants decreases. However, the fuel utilization in the SOFC decreases with pressure, leading to increased destruction in the burner with pressure. There is also a small contribution at higher pressures from turbine destruction.

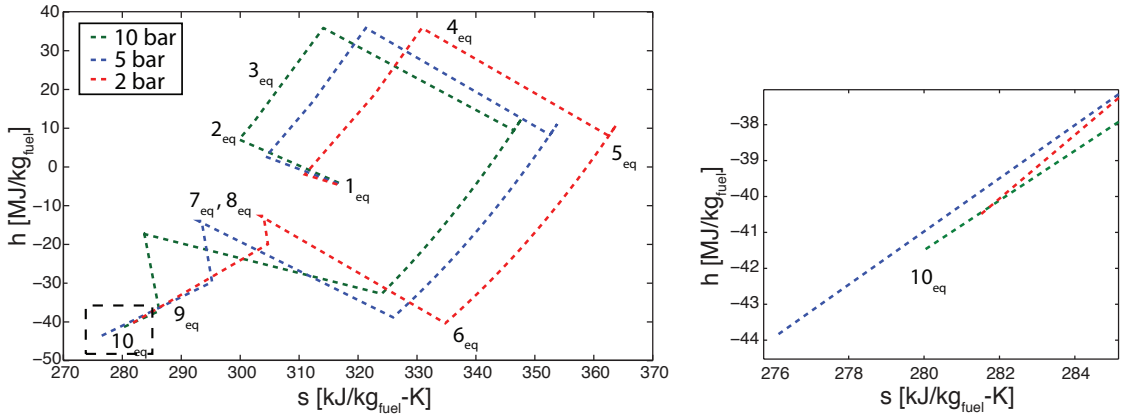

Figure 7.21: Attractor trajectory for $CX_{in}FBTX_{out}$ with 10, 5, and 2 bar peak pressures. Zoomed view at right.

The addition of the recuperator should not change the optimization principles laid out for the $CFBT$. The system should still be optimized for fuel cell utilization despite the fact that less oxidation in the fuel cell would result in a higher turbine inlet and outlet temperature, which would then improve the range of the heat exchanger. Figure 7.23 depicts this with a system operating over a range of fuel cell channel lengths, and therefore fuel cell utilization percentages. As the zoomed graph on the right indicates, decreasing fuel cell utilization was not beneficial. The work extracted in the fuel cell outweighs the improvements to the heat exchanger. The gains are most significant in jumping from the 83.0% utilization at 0.1 m to the 94.9% utilization at 0.25 m.

Optimal $CX_{in}FBTX_{out}$

Figure 7.22 indicated a peak efficiency at 0.5 m channel length of 75%. Unlike in the case of the $CFBT$, there is only minimal change with fuel cell length, as Fig. 7.24a shows. By the time the fuel cell length increases to 1.0 m, the fuel utilization remains above 95% even at relatively high pressures. As such, the efficiency of the fuel cell system drops slowly with increased pressure. Thus, the turbine inlet temperature

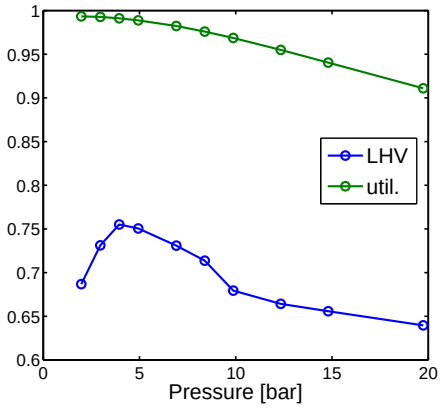

LHV efficiency and fuel utilization as functions of pressure

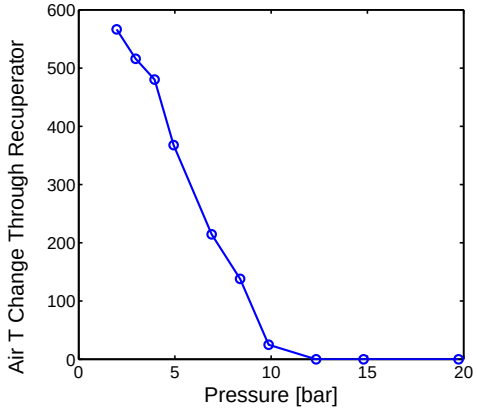

Temperature gain by cold air through the recuperator

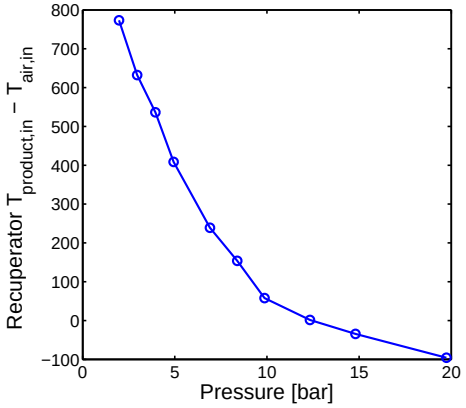

Temperature difference between inlets of the recuperator

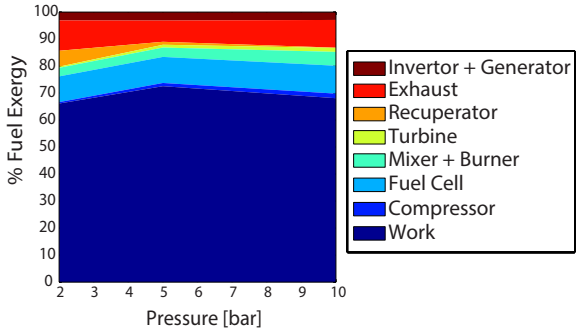

Exergy distribution over various pressures

Figure 7.22: Effects of pressure variation on the $CX_{in}FBTX_{out}$. T=1000℃, λ=2, recycle ratio=3, L=0.5 m

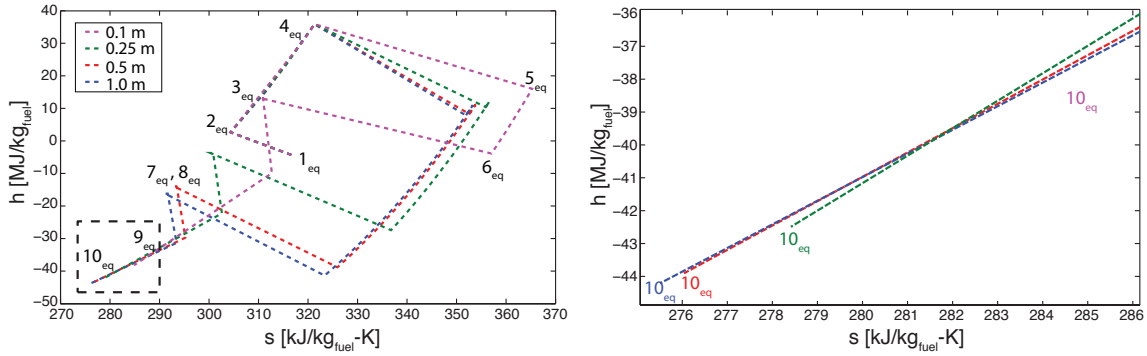

Figure 7.23: Effect of SOFC length on $CX_{in}FBTX_{out}$ system. Zoomed view at right.

rises slowly with increased pressure. This leads to a greater drop in turbine outlet temperature at longer length with increased pressure ratio. Figure 7.24b shows the comparison of the turbine outlet temperature between the two length cases. As this temperature decreases, the recuperator is more limited in its ability to heat the reactants, and it turns itself off at a lower pressure than when it was operated with a shorter fuel cell length.

In the limit at which the system is not restricted by fuel cell utilization, the pressure of the system should be optimized for the gas-turbine engine, which will be higher than the present optimum at 5 bar. However, as just discussed, a recuperator cannot be employed at higher pressures. While limitless membrane length is an unrealistic, idealized case, one final architecture will be introduced to address this potential problem.

7.2.3 Another Degree of Freedom: External Heat Transfer

The option of external heat transfer out of the engine to the environment is now considered. This will be accomplished through the introduction of the intercooler devce, I. The new general sequence is of the type

$$(CI)_k(X_{in}I)_m(FI)_n(BI)_o(TI)_p(X_{out}I)_q$$

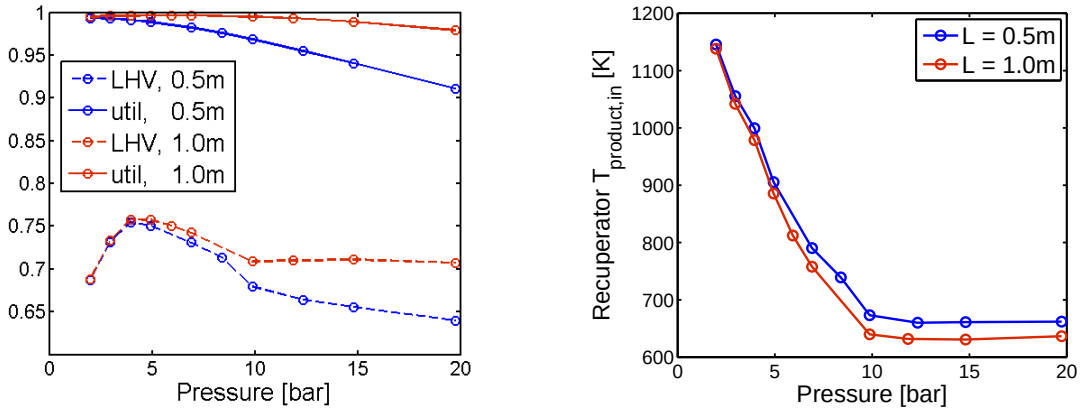

(a) LHV efficiency and fuel cell utilization.

(b) Recuperator hot-side inlet T.

Figure 7.24: Variation of optimal efficiency with channel length. T = 1000℃, λ = 2, recycle ratio = 3

where k, m, n, o, p, and q are arbitrary integers. Most of the possible sequences can be eliminated immediately. First, there is no reason to intersperse heat exchangers with intercoolers. On the receiving side, the thermal energy being recuperated is then immediately rejected. This results in removal of both the heat exchanger and intercooler. On the donating side, the intercooler limits the amount of energy that can be recovered. Similarly, interspersing intercoolers with the burner or fuel cell increases entropy generation of the chemical reaction. Interspersing intercoolers with turbines removes the ability to produce work and reduces the ability to recover thermal energy, as already discussed. Thus, the only potential candidate is intercooling between compression stages; with heat recuperation after completing compression.

$$(CI)_k X_{in} FBTX_{out}$$

The previous section indicated that at high pressures, there was not a sufficient temperature difference between the post-compressor air and the post-turbine products to add a recuperator. Here the worthiness of the recuperator is revisited when

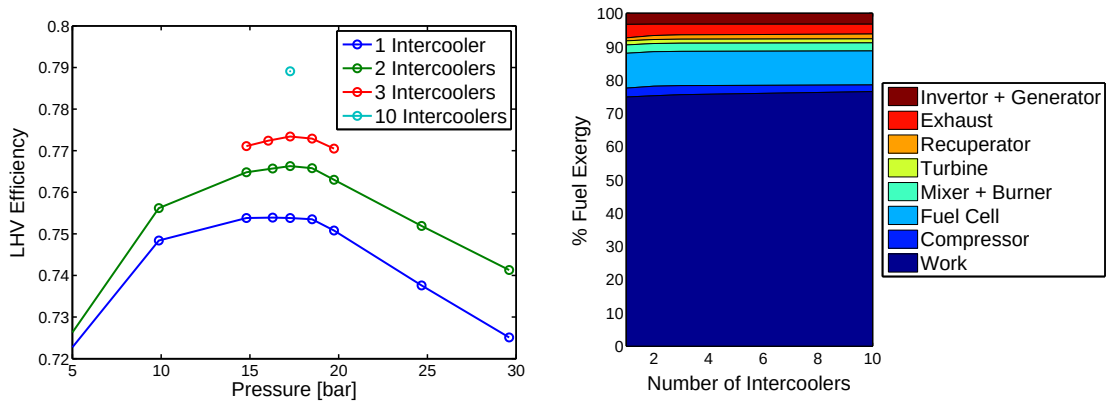

(a) Efficiency over range of pressures, number of intercoolers

(b) Exergy distribution over number of intercoolers

Figure 7.25: Optimal $(CI)_n CX_{in} FBTX_{out}$. T=1000℃, λ=1.25, Length=3m, R=3

the air is compressed near isothermally, as opposed to adiabatically. Intercoolers allow the rejection of entropy with minimal exergy, thus the recuperator could operate over a wider range of pressures without being restricted by the post-compressor temperature. Because this architecture involves external heat transfer, the minimization of s_{out} no longer corresponds to the minimization of $\dot{S}_{gen}$. The attractor definition is then extended to include external destruction, but cannot be used visually here [3].

Figure 7.25a shows that system efficiency increases with the number of intercoolers, but that the optimal pressure ratio does not change with added intercoolers beyond the initial one. While small improvements would continue with an increased number of intercoolers, three intercoolers seems a reasonable choice, especially given a peak pressure of 17.5 bar. Figure 7.25b shows how the exergy distribution changes with added intercoolers for a system at 17.5 bar. As expected, an increased number of intercoolers decreases the compressor irreversibility. Counterbalancing the decrease in compressor irreversibility is a slight increase of recuperator irreversibility. This is reasonable given that increasing the range of the recuperator was one of the primary benefits of the intercooler. These two opposing effects, combined with the fact that neither associated destruction term was particularly large to begin with, result in an

only modest overall gain in efficiency with increased number of intercoolers.

7.3 Conclusions for Use of SOFCs

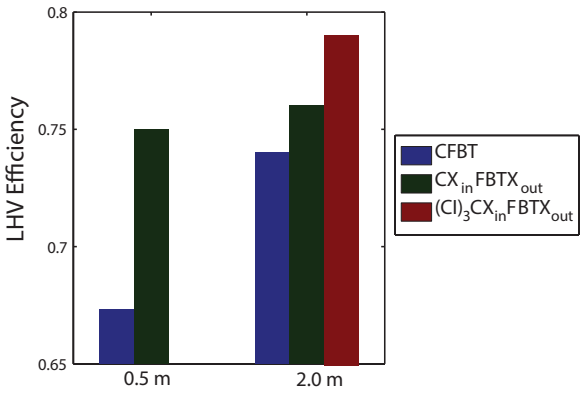

Figure 7.26: Best efficiency scenario for each architecture examined

The SOFC-GT double cycle has been systematically optimized with the attractor method. The attractor diagrams provided the same information as parametric analyses, but also provided understanding of the causal agents in the changes in efficiency.

Fuel utilization, which has a large impact on overall efficiency, has a strong dependence on channel length. Thus, the optimal operating pressure, excess air coefficient, and even architecture, depend upon the available membrane length. Figure 7.26 shows the optimal efficiency for each architecture for the case of a standard fuel cell length of 0.5 m and for the case of effectively unlimited fuel cell length to achieve greater than 99.7% utilization in the fuel cell. (This occurs between 1.0 m and 2.0 m for the range of pressures examined.) For the length-limited case, the best architecture is the $CX_{in}FBTX_{out}$ with a peak LHV efficiency of 75%. For the unlimited length case, the best archiecture is the $(CI)_nCX_{in}HFBTX_{out}$ with a maximum LHV efficiency of 79%.

The efficiencies achieved through the optimization process are significantly higher than for any existing power plant, and provides motivation for more development of SOFC-GT systems. Because the modeled systems do not use HRSGs, the system is also much smaller than the MHI triple cycle and might be built to meet a wide range of desired capacities.

Chapter 8

Synthesis

Engines are composed of sequences of devices that perform energy transfers and transformations with the objective of transforming some fraction of a resource into work. Expanding the range of the types of transfers and transformations leads to the ability to build more complex architectures with the potential to provide improved efficiency. Historically, this has been done largely in an ad hoc manner via incremental improvements to existing systems.

This thesis builds off of recent work by Teh and Ramakrishnan to systematically identify optimally efficient engines for any given set of energy transfers and transformations. These transformations and transfers (degrees of freedom) are performed in devices representative of the current state-of-the-art. The engine architecture is expanded by one degree of freedom at a time, after which the complete system's parameters are recalculated for optimal thermodynamic performance.

Figure 8.1 summarizes the progress made in expanding the scope of optimized architectures. First Teh applied the systematic approach to optimizing the efficiency of work-regenerative piston-cylinder (batch-combustion) engines. Ramakrishnan then translated this research to steady-flow combustion and added internal and external modes of heat transfer. In this dissertation, the steady-flow combustion optimization is completed with the addition of optimal matter transfer. A restrained chemical

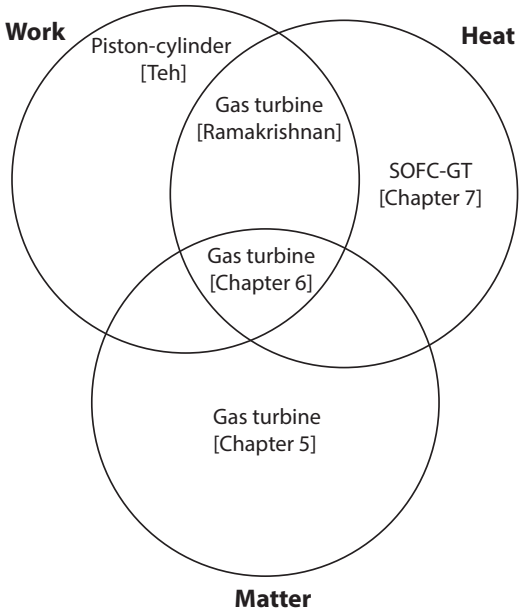

Figure 8.1: Summary of progress in adding degrees of freedom to steady-flow combustion engine optimization.

transformation, in the form of an SOFC, was then introduced to the steady-flow combustion engine and optimized for heat transfer.

Figure 8.2 shows the results of this evolution, starting with a non-optimized simple-cycle gas-turbine engine (labeled 'Simple-Cycle GT). This is typical of most gas-turbine engines today and has a peak efficiency of approximately 47%. The individual points on each curve represent the best efficiency/air-specific work combination for a specific equivalence ratio, starting with 0.2 at the left and increasing in increments of 0.1 moving right. Varying the equivalence ratio is equivalent to varying the peak system pressure. For the simple-cycle, the pressure varies from 117 bar at the leftmost point to 17 bar at the rightmost point.

Ramakrishnan first derived the optimal efficiency possible from a work-regenerative gas-turbine engine (labeled Opt. Work), which resulted in a $CB(TB)_nT$ architecture, in which combustion is first done adiabatically until the maximum temperature is reached, and then continued isothermally through a series of differential stages of

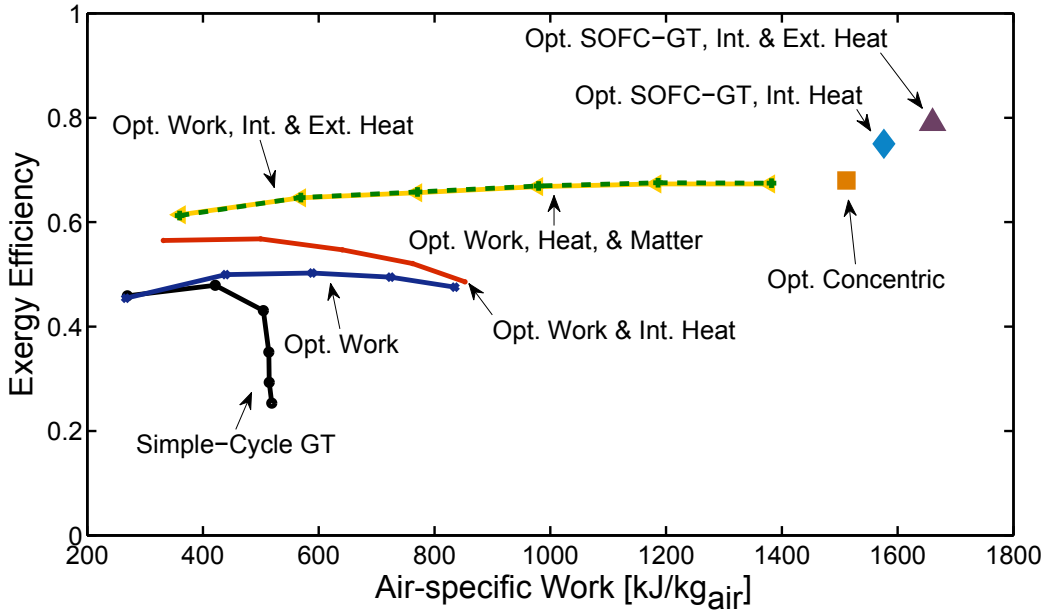

Figure 8.2: Optimal architectures for varying degrees of freedom.
Turbine inlet temperature: 1650 K, turbo-machinery polytropic efficiencies: 0.9,
fuel: 90% CH_4, 5% C_2H_6, 5% N_2

turbines and burners. This already leads to a notable improvement in air-specific work and efficiency over existing architectures.

Additionally optimizing for internal heat transfer results in the red Opt. Work & Int. Heat curve. Internal heat exchange allows for thermal recuperation from the expanded product gas, which otherwise would have been wasted in the exhaust. This brings the peak efficiency up to 57%.

Allowing for external heat transfer in the form of intercooling between compression stages leads to the yellow curve (labeled Opt. Work, Int. & Ext. Heat) and a 10-percentage point jump in peak efficiency without any compromise in air-specific work.

There are only four ways to transfer energy: with matter, with radiation, as heat, or as work. Non-thermal energy (e.g. lasers) is not relevant to combustion engines, which means the only remaining transfer to optimize for is matter. In this study, regenerative and external matter transfers were added as degrees of freedom. Ultimately feed-back and feed-forward matter transfer was not helpful in improving the system efficiency. Water, the only environmental fluid other than air, was introduced at this point. Water provides an alternative to heat transfer in the effort to recuperate thermal energy. The resulting architecture (labeled Opt. Work, Heat & Matter) maintains the same device sequence as Ramakrishnan's optimal heat-and-work sequence, but is optimized by adding approximately 5% of the mass of air additionally as water. The minimal improvement in adding water is a testament to the success of the previous dry work-and-heat architecture in optimizing the transfers of thermal energy. Since water does not significantly affect the chemical reaction, there was no change in the use of chemical exergy.

When water is circulated in a closed, concentric loop with the air, it can expand below the dead state pressure and take advantage of significant mechanical exergy. In the process, it loses the ability to maintain thermal contact in the burner, since closed heat exchangers cannot operate at burner temperatures. This alternative use of water

results in similar efficiencies, also shown in Fig. 8.2 (labeled Opt. Concentric).[1] The air-specific work is higher for the concentric cycle than the injected wet cycle in large part due to the significant fraction of the working fluid existing as water (as opposed to air).

At this point, the architecture is the optimal combination of compressors, burners, turbines, and heat exchangers using chemical fuel, air, and water. It represents an improvement on existing power plants and is entirely composed of existing technologies. In order to take one further step toward improving efficiency, a new chemical reaction device is needed.

In Chap. 7, solid-oxide fuel cells were integrated into the optimal architecture. The fuel cell oxidizes the fuel while directly producing electricity. (As compared to the burner, which produces sensible energy that later is extracted in a turbine.) Using a solid-oxide fuel cell to complete the majority of the chemical reaction significantly reduced the entropy production due to chemical reaction, and therefore increased overall efficiency. A burner is still required after the fuel cell in order to react the incompletely oxidized fuel cell products, but this burner has the benefit of operating at high temperatures, which increases its performance.

Adding the architectures for solid-oxide fuel cell integration to Fig. 8.2 completes this study. The solid-oxide fuel cell, gas turbine engine reaches an efficiency of 75% with internal thermal recuperation. When external heat transfer in the form of compressor intercooling is also allowed, the architecture reaches a remarkable 79% LHV efficiency. Furthermore, the air-specific work is approximately three times that of the Brayton cycle. Compared to modern high-performance power plants that use a combined cycle to recuperate thermal energy, this fuel cell system does not require a heat recovery steam generator (HRSG) and could hypothetically be built on a wider range of scales. This could potentially lead to improved economics over

[1]This architecture, as well as the two fuel cell systems that follow, are only represented by a single point rather than a curve. The integration between the various cycle components is complex enough that the engine only operates at specific equivalence ratios, as opposed to over a wide dynamic range.

current proposals involving a triple cycle of a solid-oxide fuel cell, gas turbine, and steam turbine [139].

Each curve in Fig. 8.2 represents the addition of an additional degree of freedom. At one end of the operating range for each architecture curve, the new degree of freedom is essentially unused and the curve collapses onto the previous curve. The work-optimized curve in blue converges with the black unoptimized curve at low equivalence ratios where the chemical reaction is done entirely in a single combustor. The curve adding in internal heat collapses onto the work-only curve at high equivalence ratios when essentially no thermal recuperation occurs. At low equivalence ratios, the amount of external heat transfer as intercooling is negligible and the yellow and red curves converge.

8.0.1 Optimization Methodologies

Ultimately, the above optimally-efficient engine architectures were derived using a systematic thermodynamic approach developed by Ramakrishnan. A single degree of freedom in the form of a new fluid, device, or transfer mode is added one at a time and the entire system is re-evaluated for optimal parameters. The limit in which each new degree of freedom is "turned off" must result in the preceding optimal architecture. If this were not the case, then every possible permutation of devices would have to be evaluated for each degree of freedom, resulting in combinatorially growth in the number of candidate architectures with the number of degrees of freedom.

However, for the exact reasons that the systematic thermodynamic approach is successful, it is also time intensive. It can only add one degree of complexity at a time, and requires significant expert knowledge of thermodynamics and energy systems.

Based on successes achieved with similar optimization problems, a stochastic numerical approach was also investigated. Unlike the thermodynamic approach, at least hypothetically, many new degrees of freedom can be simultaneously introduced. Relatively little expert knowledge is required, which also has the positive attribute

of avoiding unintentional human bias. Unfortunately, without this knowledge, the problem ends up being far too large for successful computation. The space of possible solutions could not be scanned in a reasonable time frame with a modern computer.

The stochastic approach could potentially be used to solve more targeted problems that incorporate some element of human intuition. For example, larger building blocks for devices could be used. Instead of asking the computer to "discover" a $B(TB)_n$ combustion strategy, a single device could be available that performs isothermal combustion and expansion. Such larger building blocks, combined with upgrades in the algorithmic efficiency could potentially lead to more successful use of the stochastic approach. It is unclear, however, how large the building blocks would have to be in order to develop complex architectures, and whether these would then be so large that the human intuition entirely dictated the methodology, thereby defeating some of the benefits of searching stochastically.

8.0.2 Ramifications for Combined Cycles

The most efficient existing power plants use a simple-cycle gas-turbine coupled to a steam cycle that receives thermal energy from the exhaust of the gas turbine. This type of combined cycle has not yet been studied with the systematic thermodynamic approach. However, the most recent results in Chap. 6 regarding optimal matter transfer give some intuition of how combined cycles fit into the overall optimal architecture landscape.

In Chap. 6, the possibility was introduced of using a separate water stream circulating in parallel with the gas cycle. By closing the water loop, the condensor is set by the saturation pressure of water at ambient temperature. This pressure is significantly lower than ambient pressure, and the water cycle benefits from the extra expansion in the turbine, while only having to re-pressurize a liquid. This in fact is a (concentric) combined cycle. Instead of the traditional strategy of heating the water from the gas turbine exhaust, the water is in thermal contact with the

gas throughout compression and expansion (below the maximum allowable heat exchanger temperature). The water in the concentric cycle therefore gets hotter than water in a standard combined cycle. However, in receiving this thermal energy, the water is removing thermal energy from the gas cycle, which could have been available for expansion.

While this is a trade-off, the thermal energy is extracted, regardless, in one of the two fluid cycles. As such, it seems unlikely that there is a substantial efficiency difference in routing the water with the gas cycle, or simply through an HRSG. The main reason the concentric cycle is more efficient than existing combined cycles is due to the architectural changes that reduce irreversibility during combustion. Thus, it seems reasonable that the concentric cycles serves as an approximate upper bound on the efficiency of a combined gas/water cycle using combustion for chemical conversion.

The decision about which approach to use with water then likely comes down to economics and logistics, not thermodynamics. The concentric cycle is fairly complex, but the mixed air/water version is nearly as efficient and simpler. There is no HRSG, which tends to dominate the physical footprint of combined cycle plants. Thus, the wet cycle could potentially be built to meet a wider range of capacity demands. On the other hand, the optimized wet cycle involves a less-conventional combustion strategy with more devices. Further research into the implementation of the optimized wet cycle would be needed to make an assessment of competitive viability.

8.1 Future Work

Two main lines of research could come out of this current work. The first line is a continuation of the theoretical thermodynamic study of optimal work extraction. On the current trajectory of optimizing steady-flow combustion engines, there are a few areas left to explore. The solid-oxide fuel cell/gas turbine system was not optimized

for matter transfers. This could potentially lead to further gains in efficiency. It would also be a next logical step in the study of matter transfers because the role of water is no longer strictly thermal. Water is a chemical participant in the principal fuel cell reactions (reforming and water-gas shift).

Beyond that, the next step would be to begin to introduce chemical separations. Separations, even through simple flash processes, can allow for more interesting process paths for previously mixed fluids. The ability to condense and recirculate water could lead to benefits – including mixed wet/combined cycles that inject water into air only to separate it out later for expansion. Also significant would be an exploration of carbon dioxide separations and the co-optimization of systems for efficiency and environmental benefit.

The second line of research involves studies of implementation plausibility for the various architectures proposed. Thermodynamics was the only optimization goal for this work, but in order to build systems, they have to be funded, constructed, operated, and maintained. All of these aspects require significant investigation. The choice of the best use of water, in particular, will hinge on these choices since the thermodynamic results yield multiple equivalent possibilities. Some of this work is already underway, particularly with solid-oxide fuel cells. Mitsubishi Heavy Industries' reports indicate that there is a potential commercial pathway toward implementation of high-efficiency power plants. However, to have a positive impact on climate, these systems need to be built soon, and for that to happen, they must be both economic and easily operable.

Hopefully fifty years from now, none of the systems discussed in this dissertation will be newly built because the world will have largely switched to renewable energy resources. However, to get to that point, while minimizing climate risk, fossil fuels today have to be used as efficiently as possible. This thesis provides many promising avenues using current technologies to improve fossil fuel systems in the near future.

Additionally, this work demonstrates the power of applying systematic optimization methodologies to engineering problems. Using this approach to optimize steady-flow combustion engines uncovered initially counter-intuitive conclusions, such as the advantages of compressing air after recuperative heat exchange, or the limited value in adding water to an already thermally-integrated system. Incremental progress can, and certainly has, been made by modifying existing architectures. However, by starting with very simple systems and slowly stepping out complexity using the laws of physics, more dramatic progress is possible.

Bibliography

[1] "World Energy Outlook: Factsheets," Tech. rep., International Energy Agency, 2014.

[2] Teh, K.-Y., *Thermodynamics of Efficient, Simple-Cycle Combustion Engines*, PhD Thesis, Stanford University, Palo Alto, CA, May 2007.

[3] Ramakrishnan, S., *Maximum-Efficiency Architectures for Regenerative Steady-Flow Combustion Engines*, PhD Thesis, Stanford University, Palo Alto, CA, August 2012.

[4] Horlock, J. H., *Advanced Gas Turbine Cycles*, Elsevier Science Ltd., Oxford, UK, 2003.

[5] Bazmi, A. A. and Zahedi, G., "Sustainable energy systems: Role of optimization modeling techniques in power generation and supplyA review," *Renewable and Sustainable Energy Reviews*, Vol. 15, No. 8, 2011, pp. 3480–3500.

[6] Jonsson, M. and Yan, J., "Humidified gas turbinesa review of proposed and implemented cycles," *Energy*, Vol. 30, No. 7, 2005, pp. 1013–1078.

[7] Weinhold, F., "Metric Geometry of Equilibrium Thermodynamics," *Journal of Chemical Physics*, Vol. 63, No. 6, 1975, pp. 2479–2483.

[8] Ruppeiner, G., "Thermodynamics: A Riemannian Geometric Model," *Phys. Rev. A*, Vol. 20, Oct 1979, pp. 1608–1613.

[9] Gilmore, R., "Length and Curvature in the Geometry of Thermodynamics," *Physical Review A*, Vol. 30, No. 4, 1984, pp. 1994–1997.

[10] Callen, H., "A symmetry interpretation of thermodynamics," *Foundations of continuum mechanics, London, Macmillan*, 1974, pp. 51–78.

[11] Johannessen, E. and Kjelstrup, S., "A highway in state space for reactors with minimum entropy production," *Chemical engineering science*, Vol. 60, No. 12, 2005, pp. 3347–3361.

[12] Salamon, P. and Berry, R., "Thermodynamic Length and Dissipated Availability," *Physical Review Letters*, Vol. 51, No. 13, 1983, pp. 1127–1130.

[13] Teh, K.-Y., Miller, S. L., and Edwards, C. F., "Thermodynamic Requirements for Maximum Internal Combustion Engine Cycle Efficiency. Part II: Work Extraction and Reactant Preparation Strategies," *International Journal of Engine Research*, Vol. 9, No. 6, 2008, pp. 467–481.

[14] Teh, K.-Y., Miller, S. L., and Edwards, C. F., "Thermodynamic Requirements for Maximum Internal Combustion Engine Cycle Efficiency. Part I: Optimal Combustion Strategy," *International Journal of Engine Research*, Vol. 9, No. 6, 2008, pp. 449–465.

[15] Ramakrishnan, S., Teh, K.-Y., and Edwards, C. F., "Identification of Optimal Architecture for Efficient Simple-Cycle Gas Turbine Engines," *International Mechanical Engineering Congress and Exposition (IMECE) Conference Proceedings*, Vol. 2009, No. 43796, 2009, pp. 539–548.

[16] Boyd, S. and Vandenberghe, L., *Convex Optimization*, Cambridge University Press, 2004.

[17] Ahadi-Oskui, T., Vigerske, S., Nowak, I., and Tsatsaronis, G., "Optimizing the design of complex energy conversion systems by Branch and Cut," *Computers & chemical engineering*, Vol. 34, No. 8, 2010, pp. 1226–1236.

[18] Mussati, S., Barttfeld, M., Aguirre, P., and Scenna, N., "A disjunctive programming model for superstructure optimization of power and desalting plants," *Desalination*, Vol. 222, No. 1, 2008, pp. 457–465.

[19] Lee, S., Yoon, E. S., and Grossmann, I. E., "Superstructure optimization of chemical process," *SICE 2003 Annual Conference*, Vol. 3, IEEE, 2003, pp. 3171–3176.

[20] Yeomans, H. and Grossmann, I. E., "A systematic modeling framework of superstructure optimization in process synthesis," *Computers & Chemical Engineering*, Vol. 23, No. 6, 1999, pp. 709–731.

[21] Leboreiro, J. and Acevedo, J., "Processes synthesis and design of distillation sequences using modular simulators: a genetic algorithm framework," *Computers & Chemical Engineering*, Vol. 28, No. 8, 2004, pp. 1223–1236.

[22] Shah, P. B. and Kokossis, A. C., "New synthesis framework for the optimization of complex distillation systems," *AIChE Journal*, Vol. 48, No. 3, 2002, pp. 527–550.

[23] Liu, P. and Pistikopoulos, E. N., "Mixed-Integer Optimization for Polygeneration Energy Systems Design," *Optimization in the Energy Industry*, Springer, 2009, pp. 167–191.

[24] Pruitt, K. A., Leyffer, S., Newman, A. M., and Braun, R. J., "A mixed-integer nonlinear program for the optimal design and dispatch of distributed generation systems," *Optimization and Engineering*, Vol. 15, No. 1, 2014, pp. 167–197.

[25] Braun, R. J., Vincent, T. L., Zhu, H., and Kee, R. J., "Analysis, Optimization, and Control of Solid-Oxide Fuel Cell Systems," *Advances in Chemical Engineering*, Vol. 41, 2012, pp. 383.

[26] Bruno, J., Fernandez, F., Castells, F., and Grossmann, I., "A rigorous MINLP model for the optimal synthesis and operation of utility plants," *Chemical Engineering Research and Design*, Vol. 76, No. 3, 1998, pp. 246–258.

[27] Papalexandri, K. P., Pistikopoulos, E. N., and Kalitventzeff, B., "Modelling and optimization aspects in energy management and plant operation with variable energy demands-application to industrial problems," *Computers & chemical engineering*, Vol. 22, No. 9, 1998, pp. 1319–1333.

[28] Jüdes, M., Vigerske, S., and Tsatsaronis, G., "Optimization of the design and partial-load operation of power plants using mixed-integer nonlinear programming," *Optimization in the Energy Industry*, Springer, 2009, pp. 193–220.

[29] Toffolo, A., "A Synthesis/Design Optimization Algorithm for Rankine Cycle Based Energy Systems," *Energy*, 2014.

[30] Grossmann, I. E. and Daichendt, M. M., "New trends in optimization-based approaches to process synthesis," *Computers & chemical engineering*, Vol. 20, No. 6, 1996, pp. 665–683.

[31] Holland, J. H., "Outline for a logical theory of adaptive systems," *Journal of the ACM (JACM)*, Vol. 9, No. 3, 1962, pp. 297–314.

[32] Parsopoulos, K. E., Vrahatis, M. N., et al., "Particle swarm optimization method for constrained optimization problems," *Intelligent Technologies–Theory and Application: New Trends in Intelligent Technologies*, Vol. 76, 2002, pp. 214–220.

[33] Garcia-Nieto, J., Olivera, A. C., and Alba, E., "Optimal cycle program of traffic lights with particle swarm optimization," *Evolutionary Computation, IEEE Transactions on*, Vol. 17, No. 6, 2013, pp. 823–839.

[34] Onwunalu, J. E. and Durlofsky, L. J., "Application of a particle swarm optimization algorithm for determining optimum well location and type," *Computational Geosciences*, Vol. 14, No. 1, 2010, pp. 183–198.

[35] Del Valle, Y., Venayagamoorthy, G. K., Mohagheghi, S., Hernandez, J.-C., and Harley, R. G., "Particle swarm optimization: basic concepts, variants and applications in power systems," *Evolutionary Computation, IEEE Transactions on*, Vol. 12, No. 2, 2008, pp. 171–195.

[36] Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E., "Equation of state calculations by fast computing machines," *The journal of chemical physics*, Vol. 21, No. 6, 1953, pp. 1087–1092.

[37] S. Kirkpatrick, C. G. J. and Vecchi, M., "Optimization by Simulated Annealing," *Science*, Vol. 220, No. 4598, May 1983.

[38] Schwefel, H.-P., *Evolutionsstrategie und numerische Optimierung*, Ph.D. thesis, Technische Universität Berlin, 1975.

[39] Rechenberg, I., "Evolution Strategy: Optimization of Technical systems by means of biological evolution," *Fromman-Holzboog, Stuttgart*, Vol. 104, 1973.

[40] Fogel, D. B., *Evolutionary computation: toward a new philosophy of machine intelligence*, Vol. 1, John Wiley & Sons, 2006.

[41] Koza, J. R., *Automatic synthesis of both the topology and numerical parameters for complex structures using genetic programming*, Springer, 2002.

[42] Toffolo, A., Lazzaretto, A., and Morandin, M., "The HEATSEP Method for the Synthesis of Thermal Systems," *Energy*, Vol. 25, 2010, pp. 976–981.

[43] Reehuis, E., Kruisselbrink, J., Deutz, A., Bäck, T., and Emmerich, M., "Multiobjective optimization of water distribution networks using SMS-EMOA,"

Evolutionary Methods for Design, Optimisation and Control with Application to Industrial Problems (EUROGEN 2011), 2008, pp. 269–279.

[44] Larsen, U., Nguyen, T.-V., Knusen, T., and Haglind, F., "System Analysis and Optimisation of a Kalina Split-cycle for Waste Heat Recovery on Large Marine Diesel Engines," *Energy*, Vol. 64, 2014, pp. 484–494.

[45] Nieman, D. E., Dempsey, A. B., and Reitz, R. D., "Heavy-duty RCCI operation using natural gas and diesel," Tech. rep., SAE Technical Paper, 2012.

[46] Behbahani, S. and de Silva, C. W., "Mechatronic design evolution using bond graphs and hybrid genetic algorithm with genetic programming," *Mechatronics, IEEE/ASME Transactions on*, Vol. 18, No. 1, 2013, pp. 190–199.

[47] Wang, S. and Tai, K., "Graph Representation for Structural Topology Optimization using Genetic Algorithms," *Computers & Structures*, Vol. 82, No. 20, 2004, pp. 1609–1622.

[48] Lohn, J., Hornby, G., Larchev, G., and Kraus, W., "Evolvable Hardware for Space Applications," *Proc. of the AIAA 1st Intelligent Systems Technical Conference*, September 2004.

[49] Daróczy, L., Janiga, G., and Thévenin, D., "Systematic Analysis of the Heat Exchanger Arrangement Problem using Multi-objective Genetic Optimization," *Energy*, 2013.

[50] Emmerich, M., Groß, B., Henrich, F., Roosen, P., and Schütz, M., "Global optimization of chemical engineering plants by means of evolutionary algorithms," *Proc. Aspen World*, 2000.

[51] Manesh, M. K. and Amidpour, M., "Multi-objective thermoeconomic optimization of coupling MSF desalination with PWR nuclear power plant through evolutionary algorithms," *Desalination*, Vol. 249, No. 3, 2009, pp. 1332–1344.

[52] Zebulum, R. S., Pacheco, M. A., and Vellasco, M. M. B., *Evolutionary electronics: automatic design of electronic circuits and systems by genetic algorithms*, Vol. 22, CRC press, 2001.

[53] Emmerich, M., Grötzner, M., Groß, B., and Schütz, M., "Mixed-integer evolution strategy for chemical plant optimization with simulators," *Evolutionary Design and Manufacture*, Springer, 2000, pp. 55–67.

[54] Koch, C., Cziesla, F., and Tsatsaronis, G., "Optimization of combined cycle power plants using evolutionary algorithms," *Chemical Engineering and Processing: Process Intensification*, Vol. 46, No. 11, 2007, pp. 1151 – 1159, Special Issue on Process Optimization and Control in Chemical Engineering and Processing.

[55] Miller, J. F., Job, D., and Vassilev, V. K., *Genetic Programming and Evolvable Machines*, Vol. 1.

[56] McConaghy, T., Palmers, P., Steyaert, M., and Gielen, G. G., "Trustworty Genetic Programming-Based Synthesis of Analog Circuit Topologies Using Hierarchical Domain-Specific Building Blocks," *IEEE Transactions on Evolutionary Computation*, Vol. 15, No. 4, August 2011, pp. 557–570.

[57] Grimbleby, J., "Automatic Analogue Network Synthesis using Genetic Algorithms," *Genetic Algorithms in Engineering Systems: Innovations and Applications*, June 1995.

[58] Beyer, H.-G. and Schwefel, H.-P., "Evolution strategies–A comprehensive introduction," *Natural computing*, Vol. 1, No. 1, 2002, pp. 3–52.

[59] Suzuki, J., "A Markov Chain Analysis on Simple Genetic Algorithms," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 25, No. 4, April 1995.

[60] Rudolph, G., "Convergence analysis of canonical genetic algorithms," *Neural Networks, IEEE Transactions on*, Vol. 5, No. 1, 1994, pp. 96–101.

[61] Rudolph, G., "Convergence of evolutionary algorithms in general search spaces," *In Proceedings of the Third IEEE Conference on Evolutionary Computation*, Citeseer, 1996.

[62] Bäck, T., *Evolutionary algorithms in theory and practice*, Oxford Univ. Press, 1996.

[63] Nix, A. E. and Vose, M. D., "Modeling Genetic Algorithms with Markov Chains," *Annals of Mathematics and Artificial Intelligence*, Vol. 5, 1992, pp. 79–88.

[64] Miranda, J., Polanka, M. D., and Simmons, R. J., "The Use of an Ultra-Compact Combustor as an Inter-Turbine Burner for Improved Engine Performance," *Air Force Institute of Technology, WPAFB, Ohio, Masters Thesis*, 2014.

[65] Korobitsyn, M. A., *New and Advanced Conversion Technologies: Analysis of Cogeneration, Combined and Integrated Cycles*, Ph.D. thesis, Enschede, April 1998.

[66] Chiesa, P. and Macchi, E., "A thermodynamic analysis of different options to break 60% electric efficiency in combined cycle power plants," *Journal Of Engineering For Gas Turbines And Power-Transactions Of The Asme*, Vol. 126, No. 4, Oct 2004, pp. 770–785, 47th International Gas Turbine and Aeroengine Congress and Exhibition, Amsterdam, Netherlands, Jun 2002.

[67] Bolland, O. and Mathieu, P., "Comparison of two CO_2 removal options in combined cycle power plants," *Energy Conversion and Management*, Vol. 39, 1998, pp. 1653 – 1663.

[68] Botero, C., Finkenrath, M., Bartlett, M., Chu, R., Choi, G., and Chinn, D., "Redesign, Optimization, and Economic Evaluation of a Natural Gas Combined Cycle with the Best Integrated Technology CO_2 Capture," *Energy Procedia*, Vol. 1, No. 1, 2009, pp. 3835 – 3842.

[69] Sipöcz, N. and Tobiesen, F. A., "Natural gas combined cycle power plants with CO_2 capture–Opportunities to reduce cost," *International Journal of Greenhouse Gas Control*, Vol. 7, No. 0, 2012, pp. 98 – 106.

[70] Kobayashi, Y., Ando, Y., Kabata, T., Nishiura, M., Tomida, K., and Matake, N., "Extremely High-efficiency Thermal Power System-Solid Oxide Fuel Cell (SOFC) Triple Combined-cycle System," *Mitsubishi Heavy Industries Technical Review*, Vol. 48, No. 3, 2011, pp. 9–15.

[71] Bloom, "ES-5700 Energy Server Data Sheet," Tech. rep., Bloom Energy Corporation, Sunnyvale, CA, 2010.

[72] Roe, P. J., "Fuel cell system and operating method," November 2010.

[73] Zarin Pass, R. and Edwards, C., "Exergy Analysis of a Solid-oxide Fuel Cell, Gas Turbine, Steam Turbine Triple- cycle Power Plant," Vol. 2012.

[74] Zhang, X., Li, J., Li, G., and Feng, Z., "Cycle analysis of an integrated solid oxide fuel cell and recuperative gas turbine with an air reheating system," *Journal of power sources*, Vol. 164, No. 2, 2007, pp. 752–760.

[75] Leah, R. T., "Fuel cell operating cycles and systems," 2003.

[76] Bracco, S., Pierfederici, A., and Trucco, A., "The wet compression technology for gas turbine power plants: Thermodynamic model," *Applied thermal engineering*, Vol. 27, No. 4, 2007, pp. 699–704.

[77] Kavanagh, R. M. and Parks, G. T., "A Systematic Comparison and Multi-Objective Optimization of Humid Power CyclesPart I: Thermodynamics,"

Journal of engineering for gas turbines and power, Vol. 131, No. 4, 2009, pp. 041701.

[78] Hatamiya, S., Araki, H., Katagiri, Y., and Marushima, S., "An experimental and analytical study on the advanced humid air turbine system," *Challenges of power engineering and environment*, Springer, 2007, pp. 290–296.

[79] El-Masri, M., "A modified, high-efficiency, recuperated gas turbine cycle," *Journal of engineering for gas turbines and power*, Vol. 110, No. 2, 1988, pp. 233–242.

[80] Hung, T.-C., Shai, T., and Wang, S., "A review of organic Rankine cycles (ORCs) for the recovery of low-grade waste heat," *Energy*, Vol. 22, No. 7, 1997, pp. 661–667.

[81] Ghazikhani, M., Khazaee, I., and Abdekhodaie, E., "Exergy analysis of gas turbine with air bottoming cycle," *Energy*, Vol. 72, 2014, pp. 599–607.

[82] Chase, D. and Kehoe, P., "GE Combined-Cycle Product Line and Performance," Tech. rep., GE Power Systems, Schenectady, NY, 2000.

[83] "Flex-Plant Solutions: Flexibility for all your generation needs," Tech. rep., Siemens, 2013.

[84] "Gas Turbine Combined Cycle (GTCC) & Integrated Coal Gasification Combined Cycle (IGCC)," Tech. rep., Mitsubishi Heavy Industries, 2015.

[85] Poullikkas, A., "An overview of current and future sustainable gas turbine technologies," *Renewable and Sustainable Energy Reviews*, Vol. 9, No. 5, 2005, pp. 409–443.

[86] Calise, F., d'Accadia, M. D., Palombo, A., and Vanoli, L., "Simulation and exergy analysis of a hybrid Solid Oxide Fuel Cell (SOFC)-Gas Turbine System," *Energy*, Vol. 31, No. 15, 2006, pp. 3278 – 3299, 17th International Conference

on Efficiency, Costs, Optimization, Simulation, and Environmental Impact of Energy on Process Systems.

[87] Stiller, C., *Design, Operation and Control Modelling of SOFC/GT Hybrid Systems*, PhD Thesis, Norwegian University of Science and Technology, Trondheim, Norway, May 2006.

[88] Thorud, B., *Dynamic Modelling and Characterisation of a Solid Oxide Fuel Cell Integrated in a Gas Turbine Cycle*, PhD Thesis, Norwegian University of Science and Technology, Trondheim, Norway, October 2005.

[89] van Gerwen, R. J., *High Temperature and Solid Oxide Fuel Cells*, Chap. 13, 2003.

[90] Hirschenhofer, J., Stauffer, D., Engleman, R., and Klett, M., "Fuel cell handbook," 1998.

[91] Rao, A. and Samuelsen, G., "A thermodynamic analysis of tubular solid oxide fuel cell based hybrid systems," *Journal of Engineering for Gas Turbines and Power*, Vol. 125, No. 1, 2003, pp. 59–66.

[92] Zhang, X., Chan, S., Li, G., Ho, H., Li, J., and Feng, Z., "A review of integration strategies for solid oxide fuel cells," *Journal of Power Sources*, Vol. 195, No. 3, 2010, pp. 685–702.

[93] Arsalis, A., "Thermoeconomic modeling and parametric study of hybrid SOFC-gas turbine-steam turbine power plants ranging from 1.5 to 10 MWe," *Journal of Power Sources*, Vol. 181, No. 2, 2008, pp. 313 – 326.

[94] Bierschenk, D. M., Wilson, J. R., and Barnett, S. A., "High efficiency electrical energy storage using a methane–oxygen solid oxide cell," *Energy & Environmental Science*, Vol. 4, No. 3, 2011, pp. 944–951.

[95] Zhang, X., Su, S., Chen, J., Zhao, Y., and Brandon, N., "A new analytical approach to evaluate and optimize the performance of an irreversible solid oxide fuel cell-gas turbine hybrid system," *international journal of hydrogen energy*, Vol. 36, No. 23, 2011, pp. 15304–15312.

[96] Zhao, Y., Shah, N., and Brandon, N., "Comparison between two optimization strategies for solid oxide fuel cell–gas turbine hybrid cycles," *international journal of hydrogen energy*, Vol. 36, No. 16, 2011, pp. 10235–10246.

[97] Aguiar, P., Adjiman, C., and Brandon, N. P., "Anode-supported intermediate temperature direct internal reforming solid oxide fuel cell. I: model-based steady-state performance," *Journal of Power Sources*, Vol. 138, No. 1, 2004, pp. 120–136.

[98] Achenbach, E. and Riensche, E., "Methane/steam reforming kinetics for solid oxide fuel cells," *Journal of Power Sources*, Vol. 52, No. 2, 1994, pp. 283–288.

[99] Li, Y. and Weng, Y., "Performance study of a solid oxide fuel cell and gas turbine hybrid system designed for methane operating with non-designed fuels," *Journal of Power Sources*, Vol. 196, No. 8, 2011, pp. 3824–3835.

[100] Iora, P., Aguiar, P., Adjiman, C., and Brandon, N., "Comparison of two IT DIR-SOFC models: Impact of variable thermodynamic, physical, and flow properties. Steady-state and dynamic analysis," *Chemical Engineering Science*, Vol. 60, No. 11, 2005, pp. 2963–2975.

[101] Finnerty, C. M., Coe, N. J., Cunningham, R. H., and Ormerod, R. M., "Carbon formation on and deactivation of nickel-based/zirconia anodes in solid oxide fuel cells running on methane," *Catalysis Today*, Vol. 46, No. 2, 1998, pp. 137–145.

[102] Yi, Y., Rao, A. D., Brouwer, J., and Samuelsen, G. S., "Fuel flexibility study of an integrated 25kW SOFC reformer system," *Journal of power sources*, Vol. 144, No. 1, 2005, pp. 67–76.

[103] Klein, J.-M., Hénault, M., Roux, C., Bultel, Y., and Georges, S., "Direct methane solid oxide fuel cell working by gradual internal steam reforming: Analysis of operation," *Journal of Power Sources*, Vol. 193, No. 1, 2009, pp. 331–337.

[104] Stimming, U., Singhal, S., Tagawa, H., and Lehnert, W., editors.

[105] Klein, J.-M., Bultel, Y., Pons, M., and Ozil, P., "Modeling of a solid oxide fuel cell fueled by methane: analysis of carbon deposition," *Journal of Fuel Cell Science and Technology*, Vol. 4, No. 4, 2007, pp. 425–434.

[106] Ahmed, K. and Foger, K., "Kinetics of internal steam reforming of methane on Ni/YSZ-based anodes for solid oxide fuel cells," *Catalysis Today*, Vol. 63, No. 2, 2000, pp. 479–487.

[107] Klein, J., Georges, S., and Bultel, Y., "SOFC fuelled by methane without coking: optimization of electrochemical performance," *Journal of applied electrochemistry*, Vol. 40, No. 5, 2010, pp. 943–954.

[108] Emmerich, M., Grötzner, M., and Schütz, M., "Design of graph-based evolutionary algorithms: A case study for chemical process networks," *Evolutionary Computation*, Vol. 9, No. 3, 2001, pp. 329–354.

[109] Emmerich, M. T., "Optimisation of thermal power plant designs: a graph-based Adaptive search approach," *Adaptive Computing in Design and Manufacture V*, Springer, 2002, pp. 87–98.

[110] Lazzaretto, A. and Toffolo, A., "A Method to Separate the Problem of Heat Transfer Interactions in the Synthesis of Thermal Systems," *Energy*, Vol. 33, 2008, pp. 163–170.

[111] Hillermeier, C., Hüster, S., Märker, W., and Sturm, T. F., "Optimization of Power Plant Design: Stochastic and Adaptive Solution Concepts," *Evolutionary Design and Manufacture*, Springer, 2000, pp. 3–18.

[112] "MATLAB and Statistics Toolbox Release 2013b," .

[113] Goodwin, D. G., "An Open Source, Extensible Software Suite for CVD Process Simulation," *Chemical Vapor Deposition XVI and EUROCVD-14* , Vol. 8, 2003, pp. 155–162.

[114] Maaranen, H., Miettinen, K., and Penttinen, A., "On Initial Population of a Genetic Algorithm for Continuous Optimization Problems," *Journal of Global Optimization*, Vol. 37, 2007, pp. 405–436.

[115] Glover, F., *A Template for Scatter Search and Path Relinking*, Vol. 1363, 1998.

[116] Niederreiter, H., "Random Number Generation and Quasi-Monte Carlo Methods," *SIAM*, 1992.

[117] Chen, Y., Yu, T.-L., Sastry, K., and Goldberg, D. E., "A survey of linkage learning techniques in genetic and evolutionary algorithms," *IlliGAL report*, Vol. 2007014, 2007.

[118] Beasley, D., *Designing a reduced-complexity algorithm for quaternion multiplication*, chap. G1.1.

[119] Schaffer, J. D., *Application of a genetic algorithm to finding high-performance configurations for surface mount device assembly lines*, chap. G1.7.

[120] Bailey, A., Ventresca, M., and Ombuki-Berman, B., "Genetic programming for the automatic inference of graph models for complex networks," *Evolutionary Computation, IEEE Transactions on*, Vol. 18, No. 3, 2014, pp. 405–419.

[121] Mahadevan, P., Krioukov, D., Fall, K., and Vahdat, A., "Systematic topology analysis and generation using degree correlations," *ACM SIGCOMM Computer Communication Review*, Vol. 36, ACM, 2006, pp. 135–146.

[122] Tangmunarunkit, H., Govindan, R., Jamin, S., Shenker, S., and Willinger, W., "Network topology generators: Degree-based vs. structural," *ACM SIGCOMM Computer Communication Review*, Vol. 32, ACM, 2002, pp. 147–159.

[123] "igraph–The network analysis package," .

[124] Paredis, J., "Co-evolutionary constraint satisfaction," *Parallel Problem Solving from NaturePPSN III*, Springer, 1994, pp. 46–55.

[125] Paredis, J., "Coevolutionary computation," *Artificial life*, Vol. 2, No. 4, 1995, pp. 355–375.

[126] Globus, A., Lawton, J., and Wipke, T., "Automatic Molecular Design Using Evolutionary Techniques," *Nanotechnology*, Vol. 10, No. 3, September 1999, pp. 290–299.

[127] Omidvar, M. N., Li, X., Mei, Y., and Yao, X., "Cooperative co-evolution with differential grouping for large scale optimization," *Evolutionary Computation, IEEE Transactions on*, Vol. 18, No. 3, 2014, pp. 378–393.

[128] Stone, S., Pillmore, B., and Cyre, W., "Crossover and Mutation in Genetic Algorithms Using Graph-Encoded Chromosomes," .

[129] Mattingly, J. D., *Elements of Gas Turbine Propulsion*, McGraw-Hill, New York, 1996.

[130] Incropera, F. P., *Introduction to heat transfer*, John Wiley & Sons, 2011.

[131] Min, J. K., Jeong, J. H., Ha, M. Y., and Kim, K. S., "High temperature heat exchanger studies for applications to gas turbines," *Heat and mass transfer*, Vol. 46, No. 2, 2009, pp. 175–186.

[132] Linderoth, S., Bonanos, N., Jensen, K. V., and Bilde-Srensen, J. B., "Effect of NiO-to-Ni Transformation on Conductivity and Structure of Yttria-Stabilized ZrO2," *Journal of the American Ceramic Society*, Vol. 84, No. 11, 2001, pp. 2652–2656.

[133] Smith, G. P., Golden, D. M., Frenklach, M., Moriarty, N. W., Eiteneer, B., Goldenberg, M., Bowman, C. T., Hanson, R. K., Song, S., Gardiner, J., C., W., Lissianski, V. V., and Qin, Z., GRI-Mech Home Page www.me.berkeley.edu/gri_mech, 2006, Last accessed on 01/30/2014.

[134] Afeefy, H., Liebman, J., and Stein, S., *Neutral Thermochemical Data*, Gaithersburg, MD.

[135] Finke, H., Gross, M., Waddington, G., and Huffman, H., "Low-temperature Thermal Data for the Nine Normal Paraffin Hydrocarbons from Octane to Hexadecane," Tech. Rep. 36, Thermodynamics Laboratory, Petroleum Experiment Station, Bureau of Mines, 1954.

[136] Matta, R., Mercer, G., and Tuthill, R., "Power Systems for the 21st Century – 'H' Gas Turbine Combined-Cycles," Tech. rep., Schenectady, NY, 2000.

[137] Ramakrishnan, S. and Edwards, C. F., "Maximum-Efficiency Architectures for Steady-Flow Combustion Engines, III: Heat- and Work-Regenerative Gas Turbine Engines," *Energy* , 2014.

[138] Ramakrishnan, S. and Edwards, C. F., "Maximum-Efficiency Architectures for Steady-Flow Combustion Engines, II: Work-Regenerative Gas Turbine Engines," *Energy* , 2014.

[139] Yamada, K., "Perspectives for decentralized power generation with fuel cells," Keynote address, World Engineers' Convention, Geneva, Switzerland, Sept 2011.

[140] Van Herle, J. and Vasquez, R., "Conductivity of Mn and Ni-doped stabilized zirconia electrolyte," *Journal of the European Ceramic Society*, Vol. 24, No. 6, 2004, pp. 1177–1180.

[141] Contra Costa Generating Station, L., "Application for Certification - Docket# 09-AFC-4," Case document, California Energy Comission, Sacramento, CA, June 2009.

[142] Ramakrishnan, S., Teh, K.-Y., Miller, S. L., and Edwards, C. F., "Optimal Architecture for Efficient Simple-Cycle, Steady-Flow, Combustion Engines," *Journal of Propulsion and Power*, Vol. 27, No. 4, 2011, pp. 873–883.