

Development of Low-Irreversibility Engines

Investigators

Christopher F. Edwards, Associate Professor, Mechanical Engineering; Kwee-Yan Teh, Shannon L. Miller, Patrick A. Caton, Hanho Song Graduate Researchers

Introduction

An *engine* is a device that converts some fraction of the energy in a resource into work. The most work that can be developed by a particular engine design is its *reversible* work. The *irreversibility* of an engine is the difference between the reversible work that it could develop and the actual work that it performs; it is the lost work. In this research we investigate the potential to design and implement engines with significantly reduced irreversibility, and thereby, improved efficiency.

The relevance of this work to the objective of GCEP is that significant improvements in efficiency are one of the most effective approaches to reducing greenhouse-gas emissions. Since the approach we take is fundamental and comprehensive, it enables improvements with fuels (such as existing hydrocarbons and possible future hydrogen fuel), and with engines (such as those for transportation, propulsion, and electrical power generation).

Background

Research by Other Groups

While there are many efforts to improve engine efficiency—both in corporate research and in government-sponsored programs—we remain the only effort to our knowledge that is considering the design and implementation limits of engines in a fundamental and comprehensive way. Other efforts remain focused on achieving marginal improvements in the implementation of particular engine designs (e.g., SI or CI piston engines, gas turbine engines), development of advanced engines based on evolved designs (e.g. direct-injection gasoline engines, PEM or SOFC fuel cells), or investigation of the performance capabilities of potential new, but reasonably well-defined engine concepts (e.g., HCCI piston engines, pulse detonation engines).

Our work is fundamentally different from these studies in that it is not predicated on the specification of the means by which the details of the process must occur. In fact the approach is quite the opposite; we compare and contrast the various designs in order to understand what they have in common so as to expose fundamental principles that may be used to develop new, enhanced-efficiency (low irreversibility) designs. Of particular relevance in this regard is our work to understand the key similarities and differences in the thermodynamic limitations of combustion engines and fuel cells.

Previous Results and Conclusions

When this project began it was with a simple observation: That one of the reasons that fuel cell systems can achieve higher efficiencies than combustion-based systems is the destruction of fuel availability during the unrestrained combustion process. This lack of restraint manifests itself in high product gas temperatures with concomitant high entropy such that efficient extraction of the sensible energy resulting from combustion

could not be achieved. In essence the high entropy generation of combustion made the energy less extractable.

During the first year of the project, this hypothesis was investigated in three ways: First, chemical kinetic simulations of combustion processes with simultaneous energy extraction in the form of work were conducted under a number of conditions and using a number of fuels. These simulations showed that under adiabatic conditions the best efficiency was always obtained for processes which did not have simultaneous work extraction but that instead had reaction go to completion before extraction. The apparent reason is that any attempt to extract energy in the form of boundary work results in a reduction of pressure in the product gases. This in turn increases entropy generation more rapidly than the rate of entropy reduction through energy extraction. These results were telling with respect to the extraction hypothesis, but not conclusive since it was still possible that, given the right set of chemical kinetics, thermal state, and rate of extraction, it might still be possible to reduce entropy generation via adiabatic work extraction.

The second approach to investigating this hypothesis was based on the use of HCCI combustion in a piston engine but with very delayed combustion phasing. By delaying the phasing of the combustion process well into the expansion stroke it was possible to set up conditions where energy extraction would occur during combustion, again providing a test of the hypothesis. These results—while only over a limited range of conditions due to extraction profile (slider crank and engine speed) limitations—showed that significant improvements in engine efficiency could be achieved under these conditions (Ref. 1). Further analysis of these experiments showed that the improved efficiency was due to reduced heat transfer losses, not reduction of reaction irreversibility. In effect, these studies showed that the strategy allowed us to reduce the loss of available energy (exergy) by use of simultaneous energy extraction to reduce peak temperatures, but not due to destruction of exergy during chemical reaction.

The third approach to investigating this problem was undertaken and completed in the second year of the project and reported on in our last annual report. This was a dynamical system optimization analysis of the adiabatic reaction problem with arbitrary piston motion (work extraction). This study—based on extensions of Pontryagin's theorem—showed that regardless of the dynamics involved, for an adiabatic system (one which necessarily yields a linear relation between work production and piston motion) the optimal solution must always lie at one of the extremes of the rate of piston motion. In the case of the piston engine problem, this extreme is that in which the rate of motion goes to zero, that is, the no-work-extraction (reaction at constant volume) case.

The importance of this third result cannot be overstated. It confirms the results of the chemical kinetic engine modeling studies initially performed and shows that the basic conclusion—that optimum reaction conditions for an adiabatic system occur without work extraction—is in fact a general result, independent of the dynamics of the system. In effect, it proves that for adiabatic systems our original hypothesis is incorrect. But in addition to this (negative) result, the study provided illumination of the results from the second study: that for non-adiabatic systems—systems which break the linearity condition upon which the proof of the third study rests—there *does* exist an optimal

piston motion (work extraction) profile. These complementary conclusions have important implications for gas turbine and piston engines respectively. For gas turbines—engines which operate with essentially adiabatic work extraction—there is no reason to pursue work extraction during combustion. The optimal configuration—assuming adiabaticity is to be retained—is that which is already employed, reaction going to completion at constant pressure in the combustor before expansion in the turbine. For piston engines however—engines where the peak product temperatures are well in excess of material capabilities, hence heat transfer must be permitted—there exists an optimal reaction-extraction coupling that is not necessarily what is attempted in today's engines. Given the inflexibility of today's engines in terms of extraction (speed can be varied, but the slider-crank profile is ubiquitous) this may or may not pose a significant opportunity for improvement.

At the same time that these studies were being pursued, a fourth effort was also undertaken. That was to understand the unique attributes of fuel cell systems that allow them to achieve highly reversible work extraction, at least at light loads. Although fuel cells have been around for a long time, and their basic thermodynamics and kinetics are well understood, for our purposes, very elementary questions about the details of energy and entropy transfer processes remained unexplained or explained inadequately. For this reason an effort was undertaken to reverse engineer the processes taking place in a fuel cell (in particular a PEM fuel cell) to see if it was possible to accomplish similar processes in an engine that did not use electrical energy extraction but instead used extraction due to gas expansion. In other words, we undertook to understand if it was possible to build a reversible expansion engine.

Efforts in this Past Year

It was against this backdrop that we began our efforts for this past year. Three objectives were selected for investigation based on the previous results. The first was to develop an analysis which allowed us to evaluate the potential for extraction profile shaping (changing piston motion) to optimize engine efficiency in nonadiabatic (piston) engine systems. This study involves simultaneous chemical kinetic and optimization modeling and has the objective of determining to what extent permitting optimized work extraction profiles can improve efficiency. The second was to provide a definitive answer to the question of whether it might be possible to develop a reversible expansion engine and what constraints might be placed on its potential capabilities and configuration. The third objective was to unify our understanding of the characteristics of both combustion (unrestrained reaction) and fuel cell (restrained reaction) engines so as to be able to develop a systematic approach to the design and development of efficient engines. In effect, the objective was to bridge the gap between these two seemingly different engine systems and to understand the relative roles of exergy transfer vs. destruction in these two systems. Finally, a fourth objective was added, but one that operated at only a modest level of investigation. That was to explore what might be accomplished in terms of shaping work addition and removal by using elevated geometric compression ratios in a piston engine combined with dynamic control of valve timing to alter the effective extent of compression and expansion. While this is certainly not the level of flexibility that might be required to implement an optimal expansion process in a nonadiabatic engine, these relatively simple studies—using variable valve timing and

HCCI combustion as tools—can give insight into the tradeoffs that exist when nontraditional work addition and extraction profiles are employed in an actual combustion system (with real losses).

In this report we briefly describe the ongoing (but not yet published) efforts on objectives one and two cited above. Objective four has been completed and published in Publication 2 and will therefore not be covered here except as it informs the studies on the larger discussion of engine efficiency. The majority of this report will focus on objective three: unifying our understanding of engine energy processes so as to be able to obtain a more systematic approach to the design and implementation of chemical (reactive) engines. We believe that we have achieved that unified understanding and that it holds significant promise for allowing us to systematically improve engine efficiency.

Results

Optimization of Nonadiabatic Engines

As stated above, our work during the past year was focused on extending the dynamical system optimization study reported last year to numerical analysis of nonadiabatic systems. The purpose of the analysis is to determine the extent of engine efficiency improvement when the work extraction process is optimized. Our short-term goal is to provide an illustration that would allow us to make a judgment about the significance of the efficiency improvement, and therefore the potential of pursuing this optimization approach for improving engine design. Long-term decisions about this approach will be taken after the initial results are completed.

The model we have developed to explore this question adopts specifications from the single-cylinder engine used in our HCCI research (where our heat transfer observations were made), a simple two-step kinetics model for propane combustion, and an empirical heat transfer correlation model. The problem of maximizing expansion-work output is then cast as an optimal control problem for a dynamical system with volume V , temperature T , and composition (mole numbers) N_i of the gas mixture in the engine as state variables. The rate of volume change, dV/dt is the control variable. We would like to maximize the expansion work (integral of PdV) over one period of the engine cycle, subject to dynamical constraints (differential equations) on temperature due to conservation of energy and on mole numbers based on chemical kinetics.

Solution of this problem is complicated by the fact that there are no effective *global* optimization methods for general nonlinear systems. As such, we start from the conventional slider-crank profile, and seek a *local* optimal solution which maximizes expansion work among feasible work-extraction profiles near the initial point. The optimization process involves three basic steps (Ref. 2):

1. Transcribe the continuous-time optimal control problem into a finite-dimensional constrained nonlinear program (NLP).

The transcription procedure is illustrated in Figure 1. The time domain is discretized into intervals. Within each interval between two time nodes, the state variables are approximated by cubic polynomials that match the states and gradients evaluated at the nodes. The control variable is linearly interpolated.




Figure 1: Transcription of continuous-time ODE system into finite-dimensional constrained NLP.

The polynomial estimates at mid-interval are in turn used to evaluate a mid-interval gradient based on the dynamical system constraints. The difference between this gradient and the vector of mid-interval polynomial derivatives is defined as the defect. It is a function of only the nodal state and control variables. The defect should be zero for the cubic polynomial to be an accurate approximation of the original dynamical system.

The set of nodal state and control variables therefore defines the finite-dimensional approximate model, and the original optimal control problem is transcribed into a nonlinear program subject to the constraint that the defect equals zero at each time interval.

2. Solve the constrained NLP numerically.

This optimization study uses the nonlinear programming solver SNOPT (Ref. 3). SNOPT implements a sequential quadratic programming algorithm to generate a sequence of iterates that converges to a point which satisfies the first-order (local) optimality conditions. At each iteration of the algorithm, SNOPT solves a quadratic program sub-problem which models the original constrained NLP. The quadratic program solution is used as the search direction for the next iterate.

3. Assess the accuracy of the solution. If necessary, repeat steps 1 and 2.

The current status of the effort is that we are in the process of addressing the challenge of transcribing the dynamics for reacting chemical systems. These numerically stiff systems, when well-posed, can usually be solved using ODE solvers that implement adaptive (multi-step, variable-order) algorithms (for instance, `ode15s` in MATLAB). However, such adaptive algorithms introduce inconsistencies in approximating the gradient and Hessian of the transcribed NLP from one iteration to the next and, as such, require special treatment in this application.

Another computational issue pertains to the use of absolute value function $|x|$ to relate engine speed to heat transfer rate. The function is non-smooth and introduces discontinuous derivatives into the transcribed NLP, while SNOPT (and other general-purpose nonlinear programming solvers) are designed to find *locally* optimal solutions, thus requiring functions in a problem to be smooth. Gradient discontinuities may be tolerated only if they occur far from an optimal solution. Again a custom adaptation of

the general solution method is required to perform our analysis. We expect both issues to be resolved within the next few months so that the results of this study might be available in mid summer.

Potential to Develop a Reversible Expansion Engine

The goal of this part of the research has been to explore whether it is possible to design a more efficient, reversible expansion engine. We approach this objective by modeling the expansion engine after a fuel cell since it operates closer to the reversible limit and therefore achieves higher theoretical efficiencies than other reactive engines. In a fuel cell, two half-cell reactions set up an electrical potential difference between the anode and the cathode. Work is extracted by a motor as the electrons move from high potential energy at the anode to low potential energy at the cathode. The analogous expansion engine would similarly have two half-cell reactions, one reacting at a high pressure, the second at a low pressure. Work would be extracted as an intermediate species expands from the high pressure side to the low pressure side of the engine. In both engines, in the limit of a low load, the chemical reactions and work extraction approach reversible operation and therefore maximum efficiency. The ultimate goal of this project is not only to explore the feasibility of building this analogous expansion engine, but to understand in a more general way how the two engines achieve their nearly reversible operation so that this understanding may be applied to other engine designs.

In last year's report, we showed an abstract drawing of the PEM fuel cell next to the analogous expansion engine which is reproduced below in Figures 2 and 3.

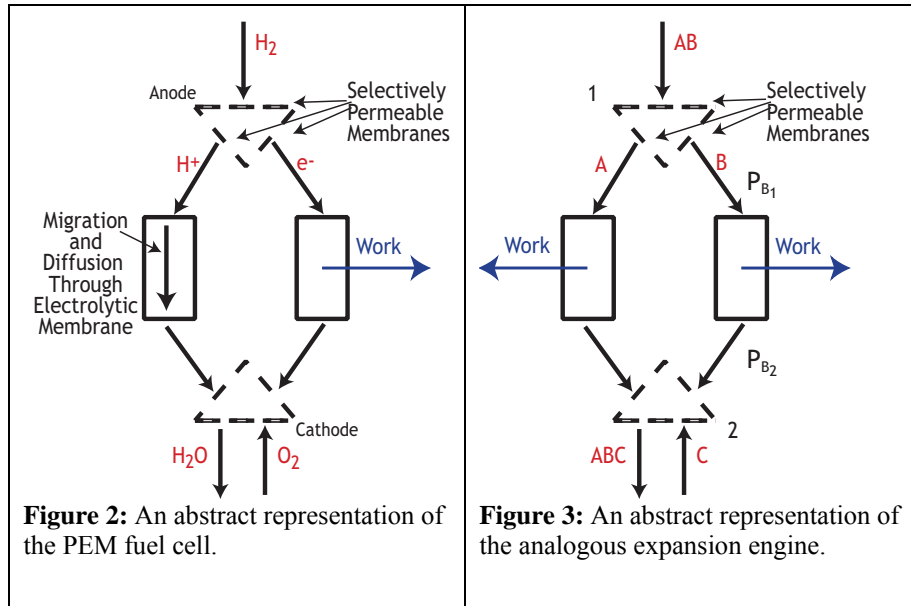


Figure 2: An abstract representation of the PEM fuel cell.

Figure 3: An abstract representation of the analogous expansion engine.

In the expansion engine, reactant AB enters the engine and reacts (with the help of a catalyst) to form species A and B at side 1. Work is extracted as these species expand from side 1 at high pressure to side 2 at low pressure. The extra work extraction due to species A moving from side 1 to side 2 is comparable to the potential work extraction from a proton current in the fuel cell. If we knew how to extract work from a proton

current (imagine a motor with Nafion windings) we could potentially increase the fuel cell's work output.

It is well known that the maximum work output for a fuel cell is

$$W_{out}^{max} = nF(\phi_C - \phi_A) \quad (1)$$

where n is the number of electrons transferred per mole of fuel, F is Faraday's constant and ϕ_C and ϕ_A are the reversible electrical potentials of the cathode and anode respectively. In last year's report (Ref. 1), we showed that the maximum work output for the expansion engine, extracting work solely from the expansion of species B , is

$$W_{out}^{max} = RT(\ln P_1 - \ln P_2) = RT \ln \left(\frac{P_1}{P_2} \right) \quad (2)$$

where R is the universal gas constant, T is temperature, and P_1 and P_2 are the reversible pressures at sides 1 and 2 of the expansion engine. The $\ln P$ term in the expansion engine replaces the ϕ term in the fuel cell. This logarithmic term has consequences for the physical implementation of this expansion engine. The pressures required to balance a typical combustion reaction are much larger than we currently use in expansion engines and would be very difficult to realize in practice. Figure 4 shows the magnitude of $\Delta_r G$ which can be balanced by pressure ratios up to 1000. These magnitudes are much lower than standard combustion reactions indicating that it is unlikely that such energetic mixtures could be used in this engine design.




Figure 4: $-\Delta_r G$ vs. pressure ratio.

While we could potentially start to look for reactants, products and semipermeable membranes that would allow us to build this type of engine, the real benefit of this research has been in understanding the link between fuel cells and other current

expansion engines. Understanding how fuel cells fit into the broader context of reactive engines has allowed us to construct a more systematic and general approach to examining engine design and increasing engine efficiency. This will be evident in the discussion of reversible reaction systems below.

Understanding Reactive Engine Energy Processes

An engine is a device that converts the exergy of a resource into work. As such the goal of our effort in this area is to develop a way to bridge between the abstract notion of resource exergy and the concrete limitations set forth by second-law analysis of specific engine designs. It is, in effect, an attempt to systematically enumerate the possible architectural and implementation choices that can be made in engine design so as to make the problem less one of inspiration and invention and more one of systematic analysis of possibilities and capabilities. We believe that we are well on our way to building that bridge.

The Three Components of Engine Architecture

Confining attention to the space of reactive (chemical) exergy resources, we assert that all engines use three essential processes that occur at least once, but sometimes more than once, in their architecture:

- (1) **Preparation/positioning of the resource:** This involves creation of the correct composition and phase as required by the chemical reaction process as well as positioning of the state of that resource so as to optimize chemical reaction and work extraction.
- (2) **Reaction to an extractable form:** This requires that the chemical resource be activated by some means (thermally or electrically) so as to open pathways for reaction, followed by either conversion to products (chemical species with minimum achievable bond energy but high sensible energy) or to another chemical form that is more easily coupled to extraction than the original resource.
- (3) **Extraction of work from the resource:** This is the essential conversion step wherein the exergy of the resource is reduced and work is done by the engine.

For example, consider the Diesel engine. Two resources are selected for use: a hydrocarbon fuel and atmospheric air. The air resource is *prepared* by compressing it to high pressure and temperature. The high temperature ensures that thermal activation of chemical reaction will occur at the time of fuel injection. The choice of a high compression ratio ensures that the post-reaction resource (high-pressure, high-sensible-energy gases) are well *positioned* for extraction by the expansion process. The fuel resource is also *prepared/positioned* by pressurization to the required injection pressure. *Reaction* to an extractable form occurs following fuel injection. Because of the positioning of the thermal state of the air and the reactivity of the fuel, autoignition occurs (thermal activation) and some part of the total exergy of the resources is transferred into a more extractable form, thermomechanical exergy of products. *Extraction* of work from the resource occurs on the expansion stroke of the engine in the

form of boundary work. Since this takes place in a batch process this work is computed in terms of the integral of PdV for the expansion process.

The gas turbine engine exhibits essentially the same architecture as the Diesel but does so using a steady flow process. Air and a hydrocarbon fuel remain the resources of choice, and both are *prepared* by compression to high pressure. Unlike the Diesel, however, flame holding can be used to provide chemical activation in the gas turbine engine. This relaxes the requirements on fuel properties and post-compression air temperature relative to the Diesel engine, which can be used to advantage in enhanced designs using intercooling and reheating. The essential processes, however, remain the same: *reaction* takes place such that the exergy is transferred from primarily chemical to primarily sensible, and *extraction* occurs by doing boundary work on the rotors of the expansion turbine. Since this takes place in a flowing process, the work output is computed in terms of the integral of VdP for the expansion process.

The PEM fuel cell provides an example that accentuates both the similarities and differences between engines. Here the resources chosen are hydrogen and air. These are *prepared/positioned* by compression (pumping) and humidification. A key difference between the fuel cell and the combustion engines cited above is that the fuel cell does not use a *reaction* stage that transforms the resources into products, but instead it accomplishes a transformation into a more extractable species. In the case of a PEM fuel cell, hydrogen is transformed into protons and electrons. The key here is the formation of free (in the conduction band of platinum) electrons that provide the ability for efficient work extraction when transferred to a motor. Another important difference is that unlike combustion engines that discharge their depleted resource (products) to the environment, this is not possible in the case of electrons. For this reason, a second *reaction* stage is required—one that can provide a sink for electrons from the motor and convert the electrons into a stable form that can be discharged in a charge-neutral fashion to the environment. This electron sink and conversion takes place on the cathode where oxygen, protons, and electrons react to form water. At both electrodes the chemical reactions are activated by a combination of thermal and electrical processes, with the relative contributions depending on the current density demanded by the load and the temperature of the cell. A final observation is that the actual conversion to work does not occur in the fuel cell, but in a motor that is distinct from the cell. As such, a fuel cell is not an engine but a chemical transformation device that is capable of sourcing and sinking a very special chemical species, the electron. (The combination of a fuel cell and an electric motor, however, does constitute an engine.)¹

One of the reasons that fuel cells are considered to be very different from combustion devices is this requirement to have two half-cell reactions in order to be able to create a potential difference. We note here that this is not unique to fuel cells but is inherent in

¹ Another reason for emphasizing that it is the fuel-cell-plus-motor combination that constitutes an engine is to maintain a fair comparison of efficiencies with other devices. While electric motors can achieve efficiencies on the order of 90%, they do incur additional energy loss that is often disregarded in the discussion of fuel cell efficiency. The analogous situation for a piston engine would be to quote indicated efficiencies in place of brake values. In fact the ratio of the two—the mechanical efficiency of the piston engine—is usually of order 90%, essentially the same as that of the electric motor.

any energy system that uses a closed cycle for its working fluid. For example, considering the case of thermal energy conversion in a Rankine cycle, the steam (or other vapor) is supplied to the turbine at high pressure and enthalpy by the boiler system, but is removed by a low pressure (via low temperature) condenser system before being recycled. While it is true that the steam could be discharged to the environment (whereas electrons cannot) it is more advantageous to use a closed cycle where the resource is managed at both the input and output of the extraction device (in this case the steam turbine).² A similar situation exists in chemical energy systems that use carrier species for intermediate reactant transport or conversion. (Black liquor in the pulp/paper cycle and chemical looping for oxygen separation before combustion are two examples that may also be considered.)

Table I: SOFC/Gas Turbine Combined Cycle processes.

Process	Preparation/positioning of the resource	Reaction to an extractable form	Extraction of work from the resource
Reforming	-Mix fuel (natural gas) with steam, preheat -Compress fuel gas -Preheat and compress air	-Thermal activation (nickel catalyst) -Natural gas \rightarrow syngas -Shift Reaction: $CO + H_2O \rightarrow H_2 + CO_2$	
↓			
↓			
Fuel Cell		-Electrical Activation (nickel catalyst) -Anode: $H_2 + O^{2-} \rightarrow H_2O + 2e^-$ -Cathode: $\frac{1}{2}O_2 + 2e^- \rightarrow O^{2-}$	-Electrical work: electrons moving through motor windings
↓			
↓			
Gas Turbine	-Vitiated anode gas mixed with air	-Thermal Activation (flame holding) -Combustion (conversion from chemical to sensible energy)	-Boundary Work: VdP flowing, expansion work

Even more complicated systems such as the SOFC/gas turbine combined cycle using natural gas and air as resources can be broken down into a sequence of these three architectural components: preparation, reaction, and extraction (see Table I). These are the required elements of any energy conversion system. The reason for emphasizing the ubiquitous nature of this simple three-part structure is that once it is accepted, it immediately becomes apparent that the critical component of any energy system is the extraction device. The preparation and reaction stages are absolutely necessary and present significant challenges, but their role is simply that of providing impedance

² Another example of a two-sided thermal energy system, which has been proposed for use with advanced nuclear energy systems (HTGR), is the closed Brayton cycle, with helium as the working fluid.

matching between the exergy resource and the extraction stage. We assert that it is our ability to extract work that sets the key limitations on engine design.

The Limitations Imposed by Extraction

A survey of work-producing devices over the course of history shows a rather surprising result: That we have used only three types of extraction devices in engines at any scale.

- (1) **Impulse machines:** Devices that convert kinetic energy to work such as the water wheel and wind turbine.
- (2) **Expansion machines:** Devices that convert thermal energy to work using either batch or flowing boundary work such as the piston/cylinder and gas turbine.
- (3) **Electrical machines:** Devices that use paired sources and sinks of electrically charged species to do work by virtue of magnetic-field interactions. While this is most commonly done using electrons (because of ease of conduction in metals) any charged species may be used in this fashion. (Recall for example the MHD efforts of the 1970s.)³

This very short list brings three ideas immediately to mind. The first is whether we have systematically explored the possibility to develop an alternative type of extraction device upon which improved engines might be built. (We defer that question for the time being to focus on requirements for optimal design using existing extraction approaches.) The second is that, in addition to work, kinetic energy also provides an entropy-free mechanism for energy transfer. While generation of kinetic energy is often used visibly in certain types of “engines” (e.g., turbojets and rocket motors), it is also used internally as an intermediate carrier within certain energy-transformation devices (e.g., the flywheel of a piston engine or the nozzles of a gas turbine). Finally, since the list is so short, and the requirements for the *positioning* and *reaction* stages are entirely dictated by the extraction device, it may be productive to organize the discussion of efficient engine design by working backward from these few options to the appropriate architecture for the overall machine. That is the approach we have adopted in our work.

Before moving to a discussion of what is possible using these devices, it is important to make one additional observation about the methods we use to develop work: None of the machines that actually produce the work do so with simultaneous chemical reaction. In the case of the gas turbine and piston engines, combustion is completed (or at least it is so intended) before the expansion process begins. Even the fuel cell—often spoken of as

³ The question might be asked is why we do not include potential energy machines in this listing. For example, a pulley with a weight might be considered a machine that uses high-altitude mass as a resource to do work. While we admit that this could be considered as another type of extractor, we would argue that it is of no practical importance to our analysis. A more relevant argument could be made about hydroelectric power systems where potential energy is converted to work. However, since our discussion is confined to the extraction device, this reduces to the problem of an impulse machine since the potential energy of the resource is first converted to kinetic energy during its fall before being extracted by the turbine.

having work extraction during reaction—actually accomplishes the reactions separately from the work extraction. And rocket motors, perhaps the best known case of reaction during expansion (shifting in equilibrium composition with thermal state, if not active burning) are designed with the intent that reaction should not be coupled with extraction. The point of interest here is that all of these systems (except the fuel cell) are cases where the extraction strategy is intended to be adiabatic, and as such our previous proof of optimality without simultaneous reaction holds for these cases. But in counterpoint this raises the question of whether any serious exploration has been devoted to systems in which heat transfer is not only permitted (for example by passive loss) but may be used to advantage in the reaction process. The former situation is one that we have already mentioned in connection with optimization of the extraction profile in the presence of heat loss in a piston engine. The latter is perhaps better viewed as a way to make a reversible chemical transformation than as a way to build a better engine, but as we shall see later, these two objectives will soon merge in our discussion of the possibilities surrounding the development of reversible chemical engines.

Starting at the End: Extraction as the Key to Engine Design

Since design of more efficient electrical and impulse motors is outside of the present scope, the only extractor which needs to be considered further is that based on fluid expansion. There are then three key issues to be considered:

- (1) **Theoretical device performance:** What are the theoretical capabilities of an expansion-based extraction device assuming ideal implementation?
- (2) **Actual device losses:** How well does the device perform relative to its theoretical design specifications? Is it lossy?
- (3) **Positioning the device with respect to environmental interactions:** What is the best way to operate the extraction device with respect to the environment?

The first two of these issues are traditional energy system design questions. How should, or could, it work? How well does it actually work? The third is more subtle and ties back to our discussion of the requirement to have two half-cell reactions in the fuel cell: For a system with one input and one output, there exists a choice of how those connections are made with respect to the environment. In traditional combustion engines, the working fluid is discharged to the environment as “exhaust.” In a fuel cell this is not possible due to the need to maintain overall charge neutrality while generating an electrical potential. In heat engines such as the Rankine and closed Brayton cycles neither connection of the expander is made directly to the environment; the state at both connections is prepared/positioned so as to provide the optimum differential conditions for the expander. The subsystems that set those states instead make the connections to the environment.

This question of whether to discharge a working fluid directly to the environment is critical both for optimization of system performance and for environmental reasons. The environmental aspects are well known: direct health effects, photochemical smog, climate

change. But the system optimization impacts can be more subtle. Recall that exergy is defined as the reversible work that can be extracted from a resource when it is allowed to interact with an environment. The way that this maximum performance is realized is for the resource to be transformed reversibly to the exact conditions of the environment itself; it must be cooled, expanded, and reacted to the concentrations of species present in the environment in order to achieve this maximum work. This implies that any matter that is released in a state that is not in complete equilibrium with the environment still has the potential to do additional work and thereby to improve efficiency. Stated in an inverse manner: Any matter that is discharged with non-zero exergy has the ability to drive environmental change. For these reasons, the ideal situation, both environmentally and for efficiency, is to discharge matter only at the conditions of the environment.

If the choice is made not to discharge matter (as in the case of closed cycles) choices still need to be made with respect to environmental interactions. The reason is that regardless of the overall system design, entropy that enters the system with the energy resource (be it heat or matter) or that is generated within the system (due to device imperfection) must be discharged. This can take place only via heat or matter. So even if matter transfer is proscribed, heat interactions with the environment will still be required and must be managed. This is the case, for example, with conventional steam cycles where the pressure is lowered below atmospheric to improve extraction in the turbine. This is limited, however, by the ability to condense steam and thereby maintain the low-exit-pressure condition. This pressure is ultimately set by the temperature of the entropy sink (via heat) available in the environment.

A final possibility to consider is whether it might be desirable, even if not required, to manage both the inlet and outlet states of the extraction device at some state that is away from that of the environment. Reasons to consider this possibility include material temperature and stress limitations as well as the ability to realize near-adiabatic conditions. Consider, for example, the case of the gas turbine engine. It is well known that the efficiency of this engine improves as the pressure ratio is increased (improving the ability to expand the gas and extract work). Unfortunately the ability to use higher pressure ratios is limited not only by turbomachinery performance but by the turbine inlet temperature. But what if, instead of starting the cycle with an environmental inlet condition—air at 1 bar, 295 K—the air were significantly pre-cooled to say 100 K, perhaps as a consequence of some other process that needed to take place at cryogenic conditions. Now it would be possible to improve the performance of the gas turbine since the complete cycle could be shifted to lower temperatures and therefore use a higher temperature swing while still maintaining the same peak temperature limitation. Whether this benefit outweighs the cost of the work required to obtain the low temperature remains to be seen. In some cases, for example the steam turbine, the benefit clearly outweighs the disadvantages. The point here is that such cases must be considered in a general evaluation of how best to position the extraction.

Returning to the first question—theoretical expansion-device performance—Figure 5 illustrates the situation for ideal gases over the range of specific heat ratios usually encountered in practice. For comparison, Figure 6 shows the temperature achieved by adiabatic combustion of methane with air as a function of equivalence ratio. What these

figures show, is that combustion of a stoichiometric mixture of methane/air at room temperature generates a product temperature that is so high (~2200 K, temperature ratio ~7.5) that it would require a volume expansion ratio in excess of 200:1 to be able to expand the gas back to the reactant temperature. (The ratio of specific heats for the product gases in this case is between 1.25 and 1.4.) Similarly, due to the even weaker dependence of gas temperature on pressure, a pressure ratio well in excess of 1000 (off the graph) would be required to successfully expand this mixture to the environmental temperature (300 K). These values should be contrasted with typical values for piston engines—10 to 20 in volume ratio—and gas turbine engines—up to 50 in pressure ratio.




Figure 5: Temperature ratio as a function of volume or pressure ratio for an isentropic, ideal gas expansion with fixed specific heat ratio k .

One approach to providing a better match might be to use a working fluid with a better specific heat ratio. The line for $k = 5/3$ (1.67) shows the expansion performance when using a monatomic gas like helium as the working fluid. While it is true that the expansion ratios can be significantly reduced (less than 20:1 in volume for a temperature ratio of 7.5), the difficulty here is that you are now committed to operating the system as a heat engine, not an internal combustion engine and ductile materials do not exist which can withstand the 2200 K hot-side temperature at pressure.

In fact it is ductile material working temperatures that impose many limits on engine performance. In addition to the obvious examples like the gas turbine temperature limit cited above, material temperature limitations can also be felt in other ways. One of these is the need for cooling in the piston engine. Given that the intermittent combustion

process reaches temperatures well in excess of material melting points, it is no surprise that active cooling is used to provide thermal management for piston engines. The issue that arises, however, is that under these conditions, the heat loss from the working fluid through the combustion process and expansion stroke is significant and we do a poor job of realizing the ideal process of an adiabatic (and reversible) extraction process. (About 1/3 of the heating value of the fuel is lost to engine coolant.) We note that if peak flame temperatures in the piston engine were reduced significantly, say to 1500 K, then it might be possible to achieve a batch expansion process that was more nearly adiabatic than we have at present.




Figure 6: Temperature for adiabatic combustion (at constant pressure) of methane in air as a function of equivalence ratio. Reactant temperature is 300K.

It is this same issue, excessive heat loss, that sets the upper limit on the compression (and hence expansion) ratio in some Diesel engines. Large heat transfer losses are usually associated with poor surface to volume ratio, but the root cause remains the requirement to cool the load-bearing walls of the cylinder. In this connection it is worth considering whether it may be possible to reduce the peak combustion temperature such that heat loss is less significant. One way to do this is to operate away from stoichiometric equivalence ratios. We note that if a peak temperature of ~ 1500 K were chosen for reasons of adiabaticity, then this would correspond to equivalence ratios in the vicinity of either 0.5 on the lean side, or just over 2.0 on the rich side. The corresponding temperature ratio ~ 5 would require an isentropic expansion ratio in the vicinity of 100:1—still well above typical values. In fact, in order to reduce the volume ratio to ~ 20 ,

the temperature ratio would need to be reduced to about 3:1 (900 K peak) with corresponding equivalence ratios of less than 0.25 or greater than 3.5.

From a traditional engine point of view, these extreme equivalence ratios present a number of problems. One obvious problem is getting reactions to occur at such low temperatures, although catalysis may be used in some situations. Another is low power density. An engine that processes four times the stoichiometric amount of air will produce relatively little output per unit of energy that is cycled internally to support the compression process. (In gas turbine engines this is referred to as the back-work ratio.) If an attempt is made to reduce the amount of work committed to air compression—for example by operating far to the rich side—then it will be necessary to provide some sort of bottoming cycle to the engine in order to consume the fuel.

The point of this discussion should be clear by now: conventional combustion temperatures are poorly suited to gas expansion given the range of volume and pressure ratios over which we can achieve adiabatic, reversible expansion. If power density is not to be sacrificed (by running at very lean or dilute conditions, for example) then some form of either bottoming or topping cycle is needed in order to obtain efficient energy conversion. This is done routinely in electric power generation systems where a steam cycle is used to bottom a gas turbine engine. In this application the steam cycle can be executed in nearly ideal fashion due to the excellent match between the turbine exhaust temperature and the isentropic expansion pressure ratios available using heated steam. For example, efficiencies in excess of 60% have been reported for the GE “H” system combined cycle power plant. Another example is the combination of the solid oxide fuel cell used as a topping cycle for a gas turbine engine. This provides a particularly interesting example in that not only does the SOFC *require* afterburning due to fuel utilization limitations—hence the combustor of the GT cycle is a necessity—but the vitiation of the fuel gas mixture by the fuel cell before being combusted for the gas turbine provides an almost ideal match to its expansion requirements. Combined cycle efficiencies in excess of 70% have been reported for this configuration.

Three approaches may be taken to address this mismatch between our ability to execute adiabatic, reversible expansions and the process of combustion:

- (1) **Compound expansion processes:** We can continue to optimize systems such as the combined cycle power plant (either NGCC or IGCC) so as to maximize extraction despite the mismatch.
- (2) **Improved expansion devices:** We can investigate ways to extend the range of expansion processes to significantly higher volume or pressure ratios while maintaining adiabaticity. This might require materials development or consideration of using sub-ambient temperature conditions (combined with cryogenic air separation for example) to access wider temperature ranges within existing material limitations.
- (3) **Modification of the reaction process:** As illustrated by the SOFC/GT example cited above, it is possible to extract work from the fuel gas stream in a fashion

that has low overall exergy loss. The possibility that we can interact with the fuel either during preprocessing (a partial transformation with work production) or during the final conversion to products in such a way as to improve extractability in the gas expander is worth considering.

Working Backwards: Interacting with the Resource

Once design choices regarding expansion and environmental connections have been made, what remains is to define the best way to position the resource for reaction such that we can maximize the effectiveness of the extraction. For instance, we would like to operate at states that minimize inadvertent exergy loss (e.g. heat loss) but allow for extraction efficiencies near the limits of the device. One additional concern is that in many cases the architecture for the reaction process is itself lossy. This is exemplified by the use of unrestrained combustion reactions, and the associated exergy destruction, that is common in all existing piston and gas turbine engines. We would like to consider whether these architectural irreversibilities can either be reduced (low-irreversibility designs) or eliminated (reversible reaction designs) and whether we can accomplish this while matching results of the reaction process to the expansion process.

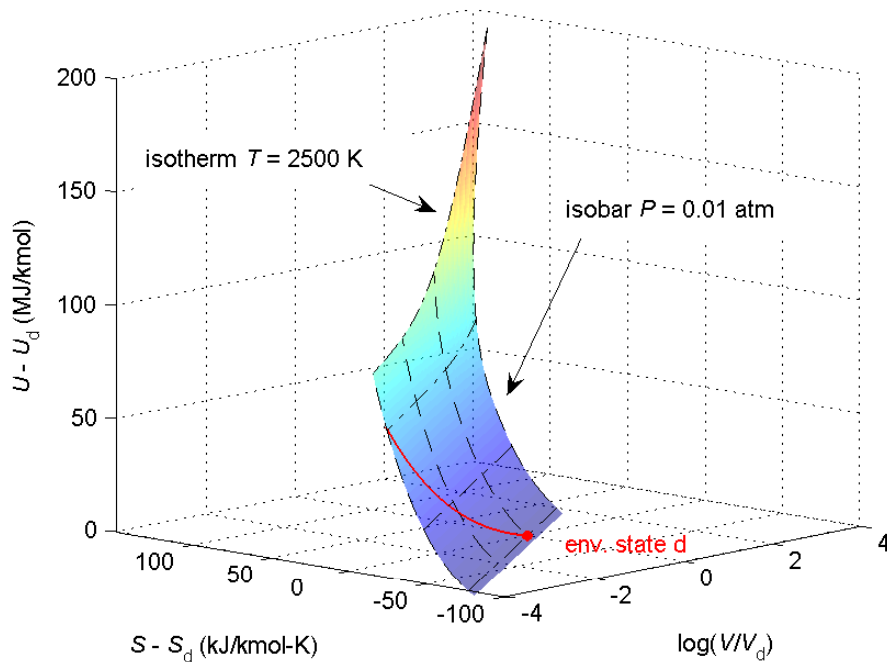


Figure 7: Thermodynamic property surface of water in chemical equilibrium with minor species (e.g. OH).

Figure 7 illustrates the situation, showing a property surface for water in the space of internal energy U , entropy S , and volume V . (This is the surface described by the *fundamental relation* of thermodynamics.) The reason for choosing water as the base species is that it is useful for illustrating both the properties of a typical working fluid (in

steam cycles) as well as for illustrating a typical product of a chemical reaction (stoichiometric hydrogen and oxygen). Please note that the surface depicted is one which shows the properties of water in chemical equilibrium, that is, the effects of dissociation have been included in the construction of the surface.

The importance of this surface is that it shows the range of possible paths by which the exergy of this resource can be coupled, reversibly, to the outside world. For example, the red line terminated by a circular dot, indicates an adiabatic, reversible path that extends from the state of the environment (298 K, 1 atm) up to the maximum pressure considered for the surface (about 10^4 atm). Note that the red curve corresponds to the intersection of the property surface with a plane at constant entropy. If our objective is to use isentropic expansion as a means to extract exergy from this resource, then the goal of the preparation/positioning and reaction stages is to transform our resource in such a way as to make the working fluid land on this isentrope. In the case of a heat-recovery steam generator this is done by providing a well matched heat exchange process and choice of feed pressure to the boiler. In the case of an internal combustion engine, this matching must be provided by the combustion process following reactant positioning.




Figure 8: Thermodynamic property surfaces for (i) stoichiometric mixture of hydrogen and air (in blue); (ii) equilibrium product mixture (in red).

Figures 8 and 9 show the relationship of reactant and product surfaces for the case of stoichiometric hydrogen and air. The product surface, therefore, contains nitrogen as well as water, but dissociation of the water is still permitted in calculation of the surface. The reactants, however, are depicted with a frozen composition, consistent with a traditional combustion process. Although Fig. 8 shows a view that is more nearly normal

to the two surfaces, it is Fig. 9A that shows the essential feature of these surfaces—that they do not intersect. Recall that for a process to be reversible, it must take place in such a way that it corresponds to a continuous path along the property surface. No jumps or discontinuities are permitted. It is not surprising, then, to see that the reactant surface does not connect to the product surface under any conditions—some form of jump (irreversibility) is always required.

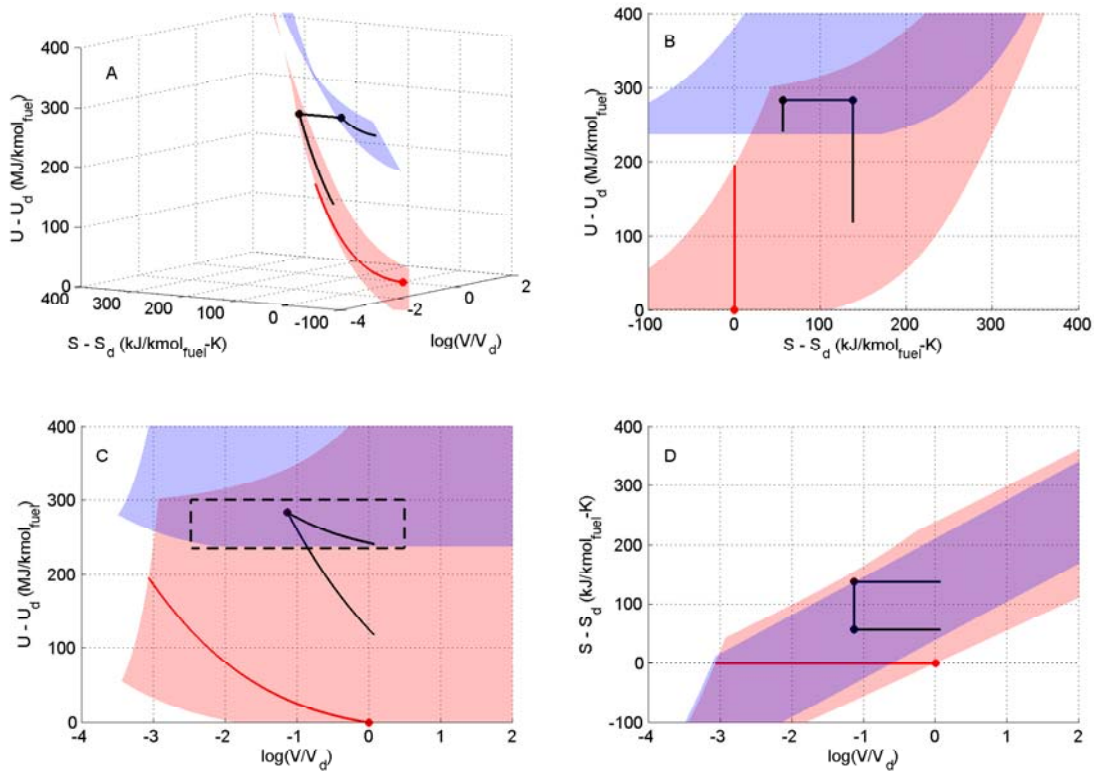


Figure 9: Thermodynamic property surfaces for stoichiometric hydrogen-air and equilibrium product mixtures (same conditions as Fig. 8), in three-dimensional view (subplot A), as well as on the three bounding planes (subplots B-D).

Also shown in the figures is the process path for the ideal, reactive Otto cycle. Best seen in Fig. 9A, this begins with an isentropic compression along the frozen reactant surface to a high internal energy condition, followed by a jump to the product surface at the same volume and energy (but with significantly increased entropy). From this landing point on the product surface, expansion then occurs until the volume returns to its original value. Two observations can immediately be made: The first is that the cycle terminates with a significant amount of internal energy still retained in the product gases. This corresponds to the high exhaust temperatures commonly experienced with SI engines. The second is that the expansion process takes place in a plane that is parallel to, but significantly displaced from, the process connecting to the environmental state. This underscores our previous point that there exists a poor match between stoichiometric

combustion processes and an optimal expansion (at least one that is open to the environment).

Figures 9B-D show projections of these processes onto the three bounding planes. Particularly apparent in Fig. 9B is that the cold reactant mixture exists at a higher entropy than the products (at the conditions for an environmental expansion) even before combustion. Equally important, Fig. 9C shows that there are no conditions within the current bounds of temperature and pressure where a constant U - V combustion process can provide a link to the environmental expansion isentrope. In order to make the reactant surface overlap the isentrope, a sub-ambient-temperature starting condition is needed. (Recall the previous discussion about positioning the expansion process.)

Another way to consider how to position the reaction process is to accept the positioning of the expansion process (let it be determined by the reaction), but to minimize the entropy generation due to combustion. The desirability of this approach stems from the Gouy-Stodola theorem which states that the exergy destroyed in a process is equal to the product of the environment temperature and the entropy generation. As might be surmised from Fig. 9A, the entropy generation is minimum at high internal energies where the two surfaces approach each other. The precise conditions can be established by evaluating the distance between the surface in the entropy direction over a range of U and V conditions. Figure 10 gives this result over the range of conditions in the vicinity of the depicted Otto cycle. (This range of conditions is the dashed rectangle on Fig. 9C.) The figure shows that entropy generation is minimized as you move to the upper right on the U - V plane—to higher reactant internal energies and volumes. This appears to be a consequence of dissociation occurring in the product gases. A final note is that the black line on Fig. 10 shows the path in U - V space along which the compression process occurs. As the reactants are compressed into a small volume (moving to the left) the entropy generation due to combustion is diminished, consistent with the improvement in Otto cycle efficiency with an increase in compression ratio.

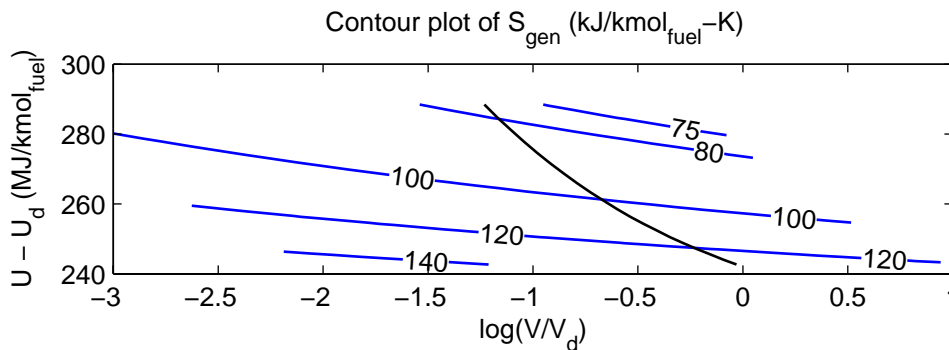


Figure 10: Contour plot of entropy generation during constant U - V combustion.

Fuel Cells and the Possibility of Reversible Reaction

So far our discussion of how to interact with and chemically convert the resource has been in connection with the assumption that the process would proceed from reactants to products in a single-step, inherently irreversible process. But as we have learned from the fuel cell, this does not need to be the case. The fuel cell uses two chemical reactors, each with a restrained, reversible reaction (at least at very low loads) to set up the differential electron source/sink pair needed to operate an electric motor. The key to recognizing what is special here is the restriction to low load. More precisely, it is a restriction that the net current requirement must be very close to (or below) the exchange current of the reaction taking place at each half cell. Upon reflection it can be seen that this reversible limit of operation for the fuel cell corresponds to other reversible limits in thermodynamics: reversible heat transfer and reversible matter transfer. In all cases, the requirement for approaching the reversible limit is that the flux required approach that due to inherent thermal processes (i.e., diffusion). The exchange current density, thermal diffusion of energy (heat), or matter (species diffusion) are all manifestations of the fact that a non-zero temperature state carries with it the potential to support small net exchanges with extremely low entropy generation (relative to the transfer).

Considered in this way, the requirements for a reversible, non-electrochemical reaction system begin to look much like those for the fuel cell: A catalyst must be used to ensure that desired reaction pathways are open over the time scale of interest (and that undesirable pathways are closed). Heat conductors (fins if you like) must be available to interact thermally with the material throughout its bulk. And the ability must exist to do boundary work as the resource flows through the system. (This might be visualized as though the reactor was the inner fluid volume of a gas turbine expander.) For all three interactions (chemical reaction, heat transfer, work transfer) the requirement of reversibility is that the agent enabling the interaction be so finely dispersed as to be within the diffusion length scale of the gas resource. (This is analogous to requiring sufficiently fine dispersion and contact proximity in the vicinity of the triple junction of a fuel cell.) In effect, the outer walls are permitted to move in order to allow energy extraction by work, while the inner surfaces are webbed with heat conductors and catalyst in order to permit a homogenous thermal and chemical state within the system.

We note that the only element missing from this description in order to constitute a fuel cell is the ability to reversibly add or remove species from the mixture. In the case of electrons, this is done with conducting fibers (wires) that are a direct analog of the heat-conducting fins hypothesized above. In the case of protons, it is the nanoscale water channels formed by interaction of sulphonated side chains in Nafion that act as the conductor. In the case of the SOFC, it is the dislocation pathways along which O^{2-} ions flow; and in the case of MCFC, it is the capillary pores in the ceramic matrix that allow the carbonate ions to flow. Since, as we noted above, the half-cells of a fuel cell do not actually do work but instead provide for reversible species transfers of electrons and ions, there is no need in this case to allow moving side walls (to do work). There is also no need to actively manage the thermal state (the heat-conducting elements can be passive).

Consider, for example, the water-gas shift reaction $H_2O + CO \rightleftharpoons H_2 + CO_2$. We know that with available catalysts (or a range of catalysts) this reaction may be maintained in chemical equilibrium over a range of temperatures from about 500-1500 K.

We also know that if the catalyst is formulated correctly, methane and other species formation can be inhibited such that just the four species cited above are present in significant quantities in the system. Figure 11 shows the equilibrium constant K_p for this reaction as a function of temperature.




Figure 11: Equilibrium constant for the water-gas-shift reaction as a function of temperature.

What we propose is to view an input mixture of carbon monoxide and water in equilibrium proportions at some thermal state as the resource supplied to our reversible reactor device. We further propose that by suitable use of heat and boundary work interactions (reversible by virtue of being at a rate close to the thermal diffusion limit) we may coax the mixture of gases to proceed from a relatively small extent of reaction ($K_p \sim 0.5$) near 1400K to a higher extent of reaction ($K_p > 100$) at temperatures near 500 K. This is depicted in Figure 12 which shows a H - P - ξ surface for a mixture with a $C:H:O$ ratio of 1:10:3 (typical of the output of a steam reformer). Since the enthalpy of the mixture is only a function of composition and temperature (for ideal gases) and the composition is only a function of temperature (it is an equimolar reaction), the enthalpy is only a function of the extent of reaction. This allows for a very special result: The shift we desire can be accomplished while producing work by expanding the reacting mixture adiabatically. This will, of course, require that the input mixture be positioned with a sufficiently high pressure to complete the expansion to the required enthalpy change, which is just a resource positioning constraint.




Figure 12: Equilibrium surface of a mixture of CO , H_2O , CO_2 and H_2 in ratio $C:H:O = 1:10:3$. The black line traces an isentropic process stemming from the thermo-mechanical dead state.

The point here is that it is possible to position chemical species in a reversible fashion if the right tools are available: catalysts that reduce activation barriers for certain pathways while raising them for others, the ability to interact through boundary work with the mixture, and the ability to interact reversibly via distributed heat transfer throughout the mixture. So long as transfer rates or reaction velocities are not far in excess of the thermal exchange limits, the process can be made to approach reversibility. Similarly, if higher rates are required, irreversibilities associated with the gradients necessary to raise the rates of transport will be incurred—just as in the case of the fuel cell.

Summary and Conclusions

Through our studies of the irreversibilities and losses of energy systems, from heat engines to fuel cells, to piston and gas turbine engines, a number of parallels have been drawn and seemingly disparate observations unified. These include the following:

- (1) Very few work extraction devices have been developed to date. The vast majority of modern devices use only electrical or boundary work (expansion). Because these are the only devices at hand, these make it relatively easy to organize the study of engine design from the point of view of starting with the requirements of the extractor and working backwards. At the same time serious consideration should be given to whether it is possible to add new options using usually

neglected thermodynamic work modes (e.g., surface tension, a two-dimensional form of boundary work).

- (2) Restricting attention to expansion devices, there exists a particular mismatch between the capabilities of today's expanders and the natural characteristics of combustion (unrestrained) reaction systems. This has been handled traditionally by use of bottoming cycles so as to, in effect, extend the expansion process to more nearly approach the environmental state. Topping cycles should also be seriously considered, particularly those with the potential to break the coupling between power density and optimal expansion. High-temperature fuel cells are one example. Other devices (e.g. thermionic engines) may also be able to provide this service and should be investigated.
- (3) It is often assumed that an expansion device must operate with one of its connections tied to the environment (usually the outlet). That this need not be the case is illustrated both by many existing closed cycles that benefit from managing both device connections and from the fuel cell that requires this for charge conservation. Given the importance of oxygen separations with respect to carbon management, serious consideration should be given to whether cryogenic starting conditions might be used to advantage.
- (4) If improved expansion devices could be constructed—devices capable of higher adiabatic working temperatures or larger expansion ratios—engine efficiency could be improved significantly. For steady-flow expanders, this is likely to be a problem of improving the strength of high-temperature materials. For batch expanders, improvement may come from innovative geometries as well as materials.
- (5) Devices which use inherently irreversible reaction processes (i.e., combustion) might be improved by either of two methods. The first is to reduce the irreversibility due to reaction by positioning the process to take place where the reactant and product entropy surfaces are more closely aligned. This generally requires high temperatures, but evidence exists that changing density may be effective under certain conditions. (Further investigation is needed to explore this.) Using this approach, the losses due to irreversibility can be reduced, but the process will be poorly matched to an expansion to environmental conditions. As such, this approach is likely to require the use of bottoming or topping. The second approach is to tolerate the irreversibility of reaction but to provide a better match to expansion conditions by pre-positioning the reactants before combustion. Again further studies will be needed to establish the possibilities.
- (6) Devices which attempt to execute reactions in a more reversible fashion can be developed (in addition to fuel cells). The key to these devices will be the ability to interact with near-reversible transfers of energy in the forms of heat and boundary work (just as the fuel cell does with ions and electrons). Such devices might be viewed as chemical transformers in much the same way that the half cells of the fuel cell can be viewed in this way. Special materials or devices that

are capable of enabling these reversible transfers and transformations—catalysts, heat-conducting media, suitable reactor-expanders—will be needed to make such an approach possible.

- (7) It appears to be possible to build a reversible expansion engine that would be the analog of the fuel cell, but using expansion work. However, the reaction energy difference that could be utilized is limited due to its logarithmic relationship with pressure. Given that we now understand the role of the half cell in providing separations and transformations, we believe that it will be more productive to pursue the reversible reaction line of attack than to attempt to pursue a practical device that uses isothermal pressure restraint alone.

Future Plans

The immediate plans for the project are to finish up the work on optimal expansion processes for non-adiabatic systems and the analysis of thermodynamics of fuel cells so as to prepare these topics for publication. In parallel with that effort, we will continue to refine our exploration of the design space of engines so as to more closely show the potential design improvements and tradeoffs that are possible. In addition to the steady-flow, reversible-reaction examples that we cited above (e.g., water-gas shift) we have seen evidence in our HCCI work that hydrogen HCCI may be thermodynamically restrained in extent of reaction. We plan to revisit this topic in order to determine if this is so and whether this can be used to advantage for improving the reversibility of piston engine combustion.

Publications

1. Caton, P.A., Song, H.H., Kaahaaina, N.B. and Edwards, C.F., “Strategies for Achieving Residual-Effectuated HCCI Using Variable Valve Actuation,” accepted SAE Journal of Engines, 2005.
2. Caton, P.A., Song, H.H., Kaahaaina, N.B. and Edwards, C.F., “Residual-Effectuated Homogeneous Charge Compression Ignition with Delayed Intake Valve Closing at Elevated Compression Ratio,” accepted International Journal of Engine Research, 2005.

References

1. Edwards, C.F., Teh, K.-Y., Miller, S.L. and Caton, P.A., “Development of Low-Irreversibility Engines,” GCEP Annual Report, 2004.
2. Hargraves, C.R. and Paris, S.W., “Direct Trajectory Optimization Using Nonlinear Programming and Collocation,” *Journal of Guidance, Control and Dynamics*, **10**, 338–342, 1987.
3. Gill, P.E., Murray, W. and Saunders, M.A., “SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization,” *SIAM Journal on Optimization*; **12**, 979–1006, 2001.

Contact

C.F. Edwards: cfe@stanford.edu