Lecture T4 **DRAFT**

Thermodynamic Processes, Heat Engines, and Available Work

Stephen Whitcomb
whitcomb.s@northeastern.edu

This unit asks a single question in three parts: how much work can a process deliver, and what sets the limit? Part I fixes the bookkeeping and derives the reversible-cycle bound. Part II asks why an expanding gas changes temperature, and when it does not. Part III replaces the idealized reservoir with a finite body or a specified environment and defines available work. Throughout we assume T1–T3; enthalpy and the response identities are recalled where used.

Contents

1 Part I. Processes, Cycles, and the Limit on Conversion

1.1. How much work can a temperature difference deliver?

Put a hot brick and a cold brick side by side in an insulated box. Left alone, they end at a common temperature and nothing else happens. Energy has moved, the total energy is unchanged, and no work has been done on anything outside the box. Now imagine that we are allowed to place a machine between them. The machine may do whatever we like internally, provided it returns to its starting condition so that we can run it again. How much work can it lift a weight through before the bricks reach a common temperature?

The question is sharper than it looks. Energy conservation tells us that whatever work comes out was supplied by the bricks, but it does not tell us how much. Conservation alone permits an answer of zero, which is what happens when we let the bricks touch; it also permits an answer equal to the entire energy difference, which never happens. Something other than the first law is selecting a number, and finding that number is the business of this part.

A second setting will occupy Part II. Take a cylinder of compressed gas and let it leak steadily through a valve into a lower-pressure line, with the whole assembly wrapped in insulation. No shaft turns, no weight is lifted, and the pipe walls are rigid. The gas nonetheless changes temperature as it passes the valve, sometimes down and sometimes up, and which way it goes depends on the gas and on the conditions upstream. Here the first law again constrains without selecting: it tells us that a particular combination of energy and pressure is unchanged, and leaves the temperature change to the equation of state.

Before either question can be answered we have to say what we are talking about. Three specifications are needed and none of them is automatic.

What a process specification must contain

(i) **The system and its boundary.** Which matter and which region of space are we tracking, and can matter cross the boundary?

(ii) **The permitted transfers.** Which of energy, volume, and particles may cross, and through what kind of wall or device?

(iii) **The path, not just the endpoints.** Work and heat are not properties of a state. Two processes joining the same pair of equilibrium states can exchange different amounts of each.

The third point is the one that causes the most trouble later, so it is worth saying plainly now. If we are told only that a gas went from (p_1, V_1, T_1) to (p_2, V_2, T_2), we can compute the change in every state function — energy, entropy, enthalpy — but we cannot compute the work or the heat. Those depend on how the change was carried out.

1.1.1 Conventions and a notation record

We write the first law for a closed system as

$$dU = đQ + đW, \tag{1}$$

so that $đW$ is the work done *on* the system and $đQ$ the heat absorbed *by* it. Both are inexact: the bar on the differential is a reminder that $\int đW$ depends on the path, while $\int dU$ does not. When

we discuss engines it is more natural to speak of the work the machine delivers, so we define

$$W_{\text{out}} \equiv -\int đW, \tag{2}$$

and we will always say which of the two we mean. A statement like "the cycle does 1.8 kJ of work" means $W_{\text{out}} = 1.8\,\text{kJ}$ and $W = -1.8\,\text{kJ}$. Sign errors in this subject are almost always convention errors, so the convention is stated once here and not changed.

Symbol	Meaning	SI unit	Convention or held fixed
U	internal energy	J	state function
H	enthalpy, $H = U + pV$	J	state function
S	entropy (physical, not S/k_B)	J K^{-1}	state function
$đQ$	heat absorbed by the system	J	> 0 into the system
$đW$	work done on the system	J	> 0 onto the system
W_{out}	work delivered by the system	J	$W_{\text{out}} = -W$
Q_{h}, Q_{c}	heat drawn from hot / rejected to cold	J	both ≥ 0 as written
T_{h}, T_{c}	reservoir temperatures	K	$T_{\text{h}} > T_{\text{c}}$
C_V, C_p	heat capacities	J K^{-1}	V fixed; p fixed
α	expansivity, $V^{-1}(\partial V/\partial T)_p$	K^{-1}	p, N fixed
κ_T	isothermal compressibility	Pa^{-1}	$-V^{-1}(\partial V/\partial p)_T$
S_{gen}	entropy generated inside the boundary	J K^{-1}	≥ 0, never < 0
η	thermal efficiency	—	$W_{\text{out}}/Q_{\text{h}}$
COP	coefficient of performance	—	benefit over work input

Table 1: Notation record for T4. Entropies are physical entropies with units of J K^{-1}; the dimensionless combination S/k_B and the inverse temperature $\beta = 1/(k_B T)$ do not appear in this unit. Subscripts on partial derivatives always name the variables held fixed.

Composition is held fixed throughout Part I: the systems are closed, with no chemical reaction and no exchange of particles. That restriction is lifted only in the steady-flow discussion of Part III, where matter crosses the boundary by construction.

1.2. Paths, and the entropy changes along them

1.2.1 When may we write $đW = -p\,dV$?

The familiar expression for compression work,

$$đW = -p\,dV \qquad \text{(quasi-static only)}, \tag{3}$$

is not a definition of work. It is a result that requires the system to have a single well-defined pressure p that matches the external pressure at every instant. That is what *quasi-static* means: the process is slow compared with the system's internal relaxation, so the system passes through a continuous sequence of equilibrium states, and the force per unit area on the piston is the equilibrium pressure of the gas.

To see why the restriction matters, consider a gas held behind a piston at 10 bar and release the piston suddenly against a surrounding atmosphere at 1 bar. During the expansion the gas is not uniform: pressure waves cross it, and there is no single number p to insert into (3). What the piston actually feels is the external pressure, so the work done on the gas is

$$W = -\int p_{\text{ext}}\,dV, \tag{4}$$

which here is $-p_{\text{ext}}\,\Delta V$ with $p_{\text{ext}} = 1\,\text{bar}$, and is smaller in magnitude than the quasi-static integral between the same endpoints. Using $\int p\,dV$ for this process would credit the gas with work it never delivered. The same warning applies to any calculation in which a system is released, allowed to relax, or driven faster than it can equilibrate.

Key Point

$\int p\,dV$ evaluates the work only when the system has an equilibrium pressure throughout and the external pressure matches it. For any other path use $\int p_{\text{ext}}\,dV$, or compute W from ΔU and Q.

1.2.2 The four standard paths

Four paths recur constantly, and it pays to have their entropy changes in one place. We take a fixed amount of an ideal gas with constant heat capacities for the explicit expressions, using $pV = nRT$ and $C_p - C_V = nR$; the middle column holds for any closed system.

Path	Held fixed	ΔS (general, closed)	ΔS (ideal gas)
isothermal	T	Q_{rev}/T	$nR\ln(V_2/V_1)$
isochoric	V	$\int C_V\,dT/T$	$C_V\ln(T_2/T_1)$
isobaric	p	$\int C_p\,dT/T$	$C_p\ln(T_2/T_1)$
isentropic	S	0 by definition	pV^γ const, $\gamma = C_p/C_V$

Table 2: Entropy changes along the standard paths. The general column assumes only a closed system of fixed composition traversing equilibrium states; the ideal-gas column additionally assumes $pV = nRT$ with temperature-independent heat capacities. All logarithms take dimensionless arguments.

Two remarks about Table 2. First, the entries are entropy changes between equilibrium states, so they apply to *any* process joining those states, reversible or not; entropy is a state function and does not remember the path. What changes with the path is how that entropy change is split between transfer and generation, a distinction we make precise in Section 1.6. Second, the isothermal entry deserves care: the expression Q_{rev}/T uses the heat that *would* flow along a reversible isothermal path. An irreversible isothermal process between the same states has the same ΔS but exchanges less heat.

1.2.3 Adiabatic is not isentropic

An adiabatic process is one with $đQ = 0$ at every point of the boundary. An isentropic process is one with $dS = 0$. These coincide only when the process is also reversible. The general statement of the second law for a closed system in contact with a boundary at temperature T_b is

$$dS \;\geq\; \frac{đQ}{T_b}, \tag{5}$$

with equality for a reversible process. Setting $đQ = 0$ gives $dS \geq 0$, not $dS = 0$.

Adiabatic and strictly not isentropic

Let an ideal gas expand freely into an evacuated, rigid, insulated chamber, doubling its volume. No work is done on anything ($W = 0$, since the surroundings offer no resistance), and no heat crosses the wall ($Q = 0$), so $\Delta U = 0$ and for an ideal gas $\Delta T = 0$. The entropy change, read from Table 2, is $\Delta S = nR \ln 2 > 0$.
The process is adiabatic and its entropy rose. The rise is generation, not transfer: nothing flowed in. Note also what we did *not* do — we did not evaluate $\int p\,dV$, because the gas had no uniform pressure while it was rushing into the vacuum.

The converse trap is equally common. "Frictionless" does not by itself imply "reversible" either: a piston may be frictionless and still release the gas into a region at a different pressure, generating entropy in the gas rather than at the seal. Reversibility is a statement about the whole process, including whether the system stays arbitrarily close to equilibrium and whether heat crosses any finite temperature difference.

1.3. What the second law forbids

The bound we are after follows from either of two classical prohibitions, which we state in the forms that are easiest to use.

Two statements of the second law

Kelvin–Planck. No process is possible whose sole result is the complete conversion of heat drawn from a single reservoir into work, with the device returned to its initial state.
Clausius. No process is possible whose sole result is the transfer of heat from a colder body to a hotter one.

The phrase "sole result" is doing real work in both. Heat drawn from one reservoir *can* be converted entirely into work — an ideal gas expanding isothermally does exactly that — but only at the cost of leaving the gas expanded. Return the gas to its original volume and the books no longer balance. Likewise heat is moved from cold to hot in every refrigerator, but not as the sole result: work is consumed.

The two statements are equivalent. Suppose a device violated Clausius by moving Q_c from the cold reservoir to the hot one with no other effect. Run an ordinary engine alongside it, arranged to take exactly Q_c from the hot reservoir and return exactly Q_c to the cold one while delivering work. The cold reservoir is then unchanged, and the pair of devices has converted heat from the hot reservoir alone into work, violating Kelvin–Planck. The argument runs in the other direction as well: a Kelvin–Planck violator could drive a refrigerator, moving heat up a temperature gradient with no other effect.

1.4. Carnot's theorem and the thermodynamic temperature scale

1.4.1 The theorem

Consider any device that operates in a cycle while exchanging heat with just two reservoirs, at T_h and T_c, and delivering work. Cyclic operation is essential: at the end of each cycle the device's own state — and therefore its own entropy — is exactly what it was at the start.

Carnot's theorem

Among all devices that operate in cycles between two reservoirs at fixed T_h and T_c, no device has a thermal efficiency greater than that of a reversible device, and all reversible devices operating between the same two reservoirs have the same efficiency. The efficiency is

$$\eta_\mathrm{rev} = 1 - \frac{T_\mathrm{c}}{T_\mathrm{h}}, \tag{6}$$

with T_h and T_c measured on the thermodynamic scale (Section 1.4.2), and with the reservoirs idealized as large enough that their temperatures do not change. Part III removes that last assumption.

The proof most useful for later work is the entropy accounting. Over one cycle, the device returns to its initial state, so $\Delta S_\mathrm{device} = 0$. The hot reservoir loses Q_h at fixed temperature T_h and the cold reservoir gains Q_c at fixed T_c, so

$$\Delta S_\mathrm{tot} = -\frac{Q_\mathrm{h}}{T_\mathrm{h}} + \frac{Q_\mathrm{c}}{T_\mathrm{c}} \;\geq\; 0. \tag{7}$$

Energy conservation over the cycle gives $W_\mathrm{out} = Q_\mathrm{h} - Q_\mathrm{c}$. Eliminating Q_c,

$$\eta \equiv \frac{W_\mathrm{out}}{Q_\mathrm{h}} = 1 - \frac{Q_\mathrm{c}}{Q_\mathrm{h}} \;\leq\; 1 - \frac{T_\mathrm{c}}{T_\mathrm{h}}, \tag{8}$$

with equality exactly when $\Delta S_\mathrm{tot} = 0$, that is, when the device is reversible. A reversible device has $Q_\mathrm{c}/Q_\mathrm{h} = T_\mathrm{c}/T_\mathrm{h}$ regardless of what it is made of, which is the second half of the theorem.

Notice what has and has not been assumed. We assumed cyclic operation, exactly two reservoirs at fixed temperatures, and that the only exchanges with the outside are those two heat flows and the work. We did not assume a working substance, an equation of state, or any particular sequence of steps. We also did not assume anything about the rate: (6) is a bound on a ratio of energies, not a prediction of power, and we return to that point in Section 1.6.1.

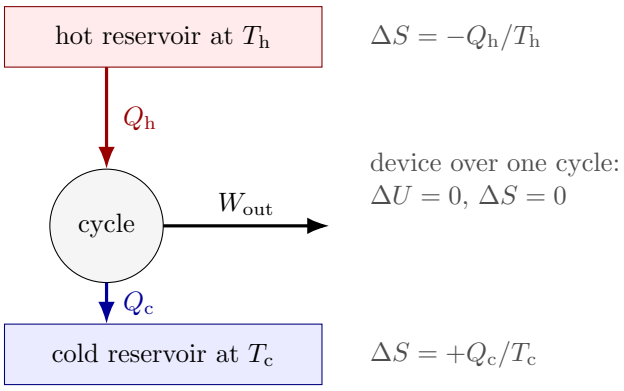

Figure 1: Entropy accounting for a cyclic device between two reservoirs. The device contributes nothing to the entropy balance because it returns to its initial state; the whole of ΔS_tot sits in the reservoirs. Requiring $\Delta S_\mathrm{tot} \geq 0$ with $W_\mathrm{out} = Q_\mathrm{h} - Q_\mathrm{c}$ gives (8). Arrow directions define the signs: Q_h and Q_c are both non-negative as drawn.

1.4.2 Why the ratio of heats defines a temperature

There is a subtlety in (6) that is easy to slide past. Carnot's theorem, proved from Kelvin–Planck alone, says that the efficiency of a reversible two-reservoir cycle depends on the two reservoirs and on nothing else. It therefore has the form

$$\frac{Q_{\mathrm{c}}}{Q_{\mathrm{h}}} = f(\theta_{\mathrm{h}}, \theta_{\mathrm{c}}), \tag{9}$$

where θ is any empirical temperature — a mercury column reading, a resistance — that labels reservoirs consistently. The function f is universal, but it is not yet a ratio of anything.

Now run two reversible cycles in series, the first between θ_1 and θ_2 and the second between θ_2 and θ_3, arranged so that the heat rejected by the first is exactly the heat absorbed by the second. The combination is a reversible cycle between θ_1 and θ_3, so

$$f(\theta_1, \theta_3) = f(\theta_1, \theta_2)\, f(\theta_2, \theta_3). \tag{10}$$

The intermediate label θ_2 has to cancel, which forces $f(\theta_i, \theta_j) = T(\theta_j)/T(\theta_i)$ for some function T of a single variable. That function is the *thermodynamic temperature*, fixed up to one multiplicative constant, and the constant is set by convention through the definition of the kelvin. With this definition, $Q_{\mathrm{c}}/Q_{\mathrm{h}} = T_{\mathrm{c}}/T_{\mathrm{h}}$ for a reversible cycle is not an extra physical assumption; it is how the scale was constructed. What is physical is that the same scale then appears in $dS = đQ_{\mathrm{rev}}/T$ and in the ideal-gas law.

1.4.3 Efficiency, refrigerators, and heat pumps

The same cycle run backwards consumes work and moves heat from cold to hot. The energy and entropy balances are unchanged in form; only the question we ask of them changes, and with it the figure of merit.

$$\begin{aligned} \text{engine:} \quad & \eta = \frac{W_{\mathrm{out}}}{Q_{\mathrm{h}}} \le 1 - \frac{T_{\mathrm{c}}}{T_{\mathrm{h}}}, \\ \text{refrigerator:} \quad & \mathrm{COP_R} = \frac{Q_{\mathrm{c}}}{W_{\mathrm{in}}} \le \frac{T_{\mathrm{c}}}{T_{\mathrm{h}} - T_{\mathrm{c}}}, \\ \text{heat pump:} \quad & \mathrm{COP_{HP}} = \frac{Q_{\mathrm{h}}}{W_{\mathrm{in}}} \le \frac{T_{\mathrm{h}}}{T_{\mathrm{h}} - T_{\mathrm{c}}}. \end{aligned} \tag{11}$$

Each bound in (11) assumes a cyclic device, exactly two reservoirs whose temperatures do not change, and no exchanges other than the two heats and the work. Each is attained only in the reversible limit. Note that $\mathrm{COP_{HP}} = \mathrm{COP_R} + 1$ identically, which simply restates $Q_{\mathrm{h}} = Q_{\mathrm{c}} + W_{\mathrm{in}}$.

The coefficients of performance are not efficiencies and are not bounded by one: delivering Q_{h} to a warm room costs only the work needed to lift Q_{c} through the temperature difference, and that difference can be small. This is the whole argument for heat pumps. A heat pump maintaining a room at 293 K against outdoor air at 273 K has $\mathrm{COP_{HP}} \le 293/20 \approx 14.7$, so in the reversible limit a joule of electricity would deliver nearly fifteen joules of heat indoors, against the one joule an electric heater delivers. Real devices reach a modest fraction of the bound — the compressor is irreversible, and the working fluid must be colder than the outdoor air and hotter than the room for heat to flow at a useful rate at all — but the comparison with resistive heating is still decisive, and it is entirely a consequence of (11).

1.5. One cycle in full

The bound (6) was derived without reference to any particular machine, which makes it powerful and also slightly unsatisfying: it is worth seeing every quantity in it appear in an explicit calculation. We use the Carnot cycle with an ideal gas, which is the simplest reversible cycle that touches only two reservoirs.

The Carnot cycle with one mole of ideal gas

Take $n = 1\,\mathrm{mol}$ of an ideal gas with constant heat capacities, a hot reservoir at $T_\mathrm{h} = 500\,\mathrm{K}$ and a cold reservoir at $T_\mathrm{c} = 300\,\mathrm{K}$. The cycle has four quasi-static legs.

1. Isothermal expansion at T_h, from V_1 to $V_2 = 3V_1$. Since $\Delta U = 0$ for an ideal gas at fixed T, the heat absorbed equals the work delivered:

$$Q_\mathrm{h} = nRT_\mathrm{h}\ln\!\left(\frac{V_2}{V_1}\right) = (8.314\,\mathrm{J\,K^{-1}})(500\,\mathrm{K})\ln 3 = 4567\,\mathrm{J}, \qquad \Delta S = \frac{Q_\mathrm{h}}{T_\mathrm{h}} = 9.13\,\mathrm{J\,K^{-1}}.$$

2. Isentropic expansion from T_h to T_c. No heat, so $W_\mathrm{out} = -\Delta U = C_V(T_\mathrm{h} - T_\mathrm{c})$ and $\Delta S = 0$.

3. Isothermal compression at T_c, back to a volume ratio of $V_4/V_3 = V_1/V_2 = 1/3$. The heat rejected is

$$Q_\mathrm{c} = nRT_\mathrm{c}\ln 3 = 2740\,\mathrm{J}, \qquad \Delta S = -\frac{Q_\mathrm{c}}{T_\mathrm{c}} = -9.13\,\mathrm{J\,K^{-1}}.$$

4. Isentropic compression from T_c back to T_h, with $\Delta U = C_V(T_\mathrm{h} - T_\mathrm{c})$ exactly cancelling leg 2.

The two isentropic legs contribute equal and opposite energy changes and no entropy change, so the net work is $W_\mathrm{out} = Q_\mathrm{h} - Q_\mathrm{c} = 1827\,\mathrm{J}$ and

$$\eta = \frac{1827}{4567} = 0.400 \qquad \text{against} \qquad 1 - \frac{T_\mathrm{c}}{T_\mathrm{h}} = 1 - \frac{300}{500} = 0.400.$$

The volume ratio 3 never appears in the efficiency: it cancels between Q_h and Q_c because the isentropic legs force $V_3/V_4 = V_2/V_1$. That cancellation is why the bound depends on the temperatures alone.

Figure 2 is worth dwelling on, because it makes the efficiency bound visual. On T–S axes the heat exchanged along a reversible leg is the area beneath it, $Q_\mathrm{rev} = \int T\,dS$, and the net work around a closed reversible cycle is the enclosed area. A reversible cycle touching only two reservoirs must be a rectangle, since heat may be exchanged only at T_h or at T_c and the connecting legs must be adiabatic. The rejected heat is the strip below T_c; it cannot be removed without dropping T_c to zero, and that is the geometric content of (6).

Cycles with more than two reservoirs — the Otto, Diesel, and Brayton cycles, in which heat is added over a range of temperatures — are not rectangles, and their efficiencies are lower than a Carnot cycle spanning the same extreme temperatures. Their design is a separate subject, and we leave it aside; what carries over is the accounting.

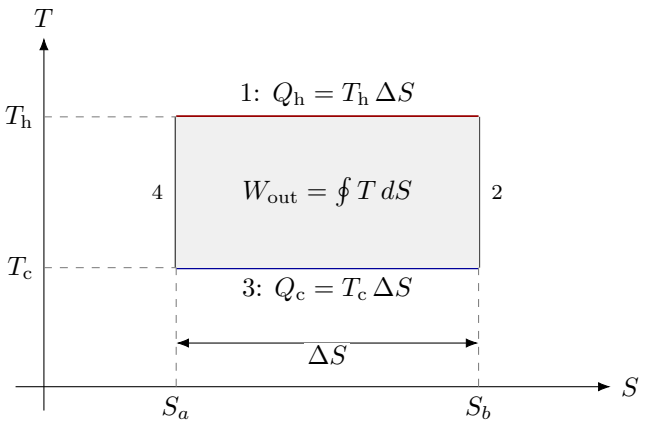

Figure 2: The Carnot cycle on temperature–entropy axes, drawn for the worked example with $T_h = 500\,\mathrm{K}$, $T_c = 300\,\mathrm{K}$, and $\Delta S = 9.13\,\mathrm{J\,K^{-1}}$. Heat absorbed and rejected are the areas under legs 1 and 3; the enclosed area is the net work. On these axes the reversible two-reservoir cycle is a rectangle, and $\eta = 1 - T_c/T_h$ is a statement about the ratio of two rectangles with a common base. Axes are schematic and not to scale.

1.6. Entropy generation and lost work

It is useful to convert the inequality into an equality by naming the deficit. For any process of a closed system whose boundary is at temperature T_b where heat crosses, define the generated entropy by

$$\Delta S = \int \frac{đQ}{T_b} + S_{\mathrm{gen}}, \qquad S_{\mathrm{gen}} \geq 0. \tag{12}$$

The first term is entropy *transferred* with heat; S_{gen} is entropy *produced* inside the boundary by whatever was irreversible there — a finite temperature difference, friction, an unrestrained expansion, mixing. It is zero for a reversible process and positive otherwise, and it is not a state function: it belongs to the process, like Q and W.

Applied to our cyclic engine, (12) gives $S_{\mathrm{gen}} = Q_c/T_c - Q_h/T_h$, and eliminating Q_c in favour of the work,

$$W_{\mathrm{out}} = Q_h\left(1 - \frac{T_c}{T_h}\right) - T_c\,S_{\mathrm{gen}}. \tag{13}$$

The work delivered falls short of the reversible value by exactly $T_c S_{\mathrm{gen}}$. This is our first encounter with a relation of the form "lost work = temperature × entropy generated," and it is worth being careful about which temperature appears. Here it is T_c, the temperature of the reservoir that ultimately receives the rejected heat, and it appears because that is where the generated entropy is finally deposited. In Part III, where the receiving body is an environment at T_0, the same structure appears with T_0 in place of T_c; where the receiving body is finite, the appropriate temperature is the one at which the entropy actually leaves, and we will derive that rather than assume it. The relation in this form goes back to Gouy and to Stodola; Part III gives the attribution along with the general derivation.

Key Point

S_{gen} is defined by (12) relative to a stated boundary. A process that looks reversible from inside one boundary may be generating entropy just outside it; before quoting a value of S_{gen}, say where the boundary is and what temperature it has where heat crosses.

1.6.1 The reversible limit is a zero-power limit

Nothing in this part predicts how long anything takes. The reversible cycle attains (6) because it never lets heat cross a finite temperature difference — but a vanishing temperature difference drives a vanishing heat current, so the reversible engine delivers its maximum efficiency at vanishing power. An equilibrium construction of this kind gives a bound on a ratio of energies. It does not say how fast the process runs, and reading a relaxation time out of it is a category error.

Real design problems therefore trade efficiency against rate, and the analysis of cycles constrained to operate at a finite rate — finite-time thermodynamics — is a distinct subject with its own assumptions about how heat currents depend on temperature differences. We do not develop it here. We flag it because a bound with nothing in it that carries units of time cannot be a statement about how quickly equilibrium is approached.

1.7. What we have, and what is still missing

We can now answer a restricted version of the opening question. If the hot and cold bodies are large enough to count as reservoirs at fixed T_{h} and T_{c}, then any cyclic machine between them delivers at most $(1 - T_{\text{c}}/T_{\text{h}})$ of the heat it draws, and the deficit from irreversibility is $T_{\text{c}}S_{\text{gen}}$. The bound came from cyclic operation plus the second law, with no reference to the working substance.

Three things are still missing, and they set up the rest of the unit.

1. **The gas itself.** We treated the working substance as irrelevant, which is correct for the bound but hides real physics. When a gas expands — freely, or steadily through a valve — its temperature changes in a way that depends entirely on its equation of state, and that dependence is what makes gas liquefaction possible. Part II takes this up.

2. **Finite bodies.** The two bricks of Section 1 are not reservoirs. As the engine runs, the hot one cools and the cold one warms, so T_{h} and T_{c} move and (6) applies only instant by instant. The total work extractable is then a question about the whole trajectory, which Part III answers.

3. **A reference state.** "Useful work" has not been defined. If the machine has to push the atmosphere out of the way, some of the work it does is not available to us. Making that precise requires naming an environment, and it leads to the notion of availability.

The immediate question for Part II is the one the throttle poses: a gas passes through a valve with no work done and no heat exchanged, and its temperature changes. Which property of the gas decides the sign of that change, and why does it reverse at high enough temperature?

2 Part II. Expansion, Throttling, and Gas Cooling

2.1. Three expansions that are not the same process

A gas at high pressure is allowed to expand. Does it cool?

The question as asked has no answer, because "expand" names at least three different processes with three different conserved quantities and three different temperature changes. Keeping them apart is most of the work.

(a) **Quasi-static adiabatic expansion against a piston.** The gas pushes on something and delivers work; no heat crosses the boundary; the process is slow enough to stay in equilibrium. Entropy is constant, and the gas cools because it paid for the work out of its internal energy.

(b) **Free expansion into a vacuum.** Nothing resists, so no work is done; the container is rigid and insulated, so no heat enters. Internal energy is constant.

(c) **Steady flow through a constriction.** Gas is pushed continuously through a valve or a porous plug from a high-pressure line to a low-pressure one, with no shaft and no heat. As we show in Section 2.3, it is the *enthalpy* that is the same on the two sides.

Process (a) is elementary and always cools an ideal gas. Processes (b) and (c) are the interesting ones, because for an ideal gas neither changes the temperature at all. Whatever temperature change they produce is therefore a direct measurement of departures from ideality, which is why they were historically important and why they are still the basis of cryogenic refrigeration.

2.2. Free expansion and the Joule coefficient

Take a rigid insulated vessel divided by a partition, gas on one side and vacuum on the other, and remove the partition. Nothing outside the vessel moves, so $W = 0$; the walls are adiabatic, so $Q = 0$; therefore $\Delta U = 0$. The question is what happens to T when V increases at fixed U.

We want $\mu_{\mathrm{J}} \equiv (\partial T/\partial V)_U$. The triple product rule gives

$$\mu_{\mathrm{J}} = \left(\frac{\partial T}{\partial V}\right)_U = -\frac{(\partial U/\partial V)_T}{(\partial U/\partial T)_V} = -\frac{1}{C_V}\left(\frac{\partial U}{\partial V}\right)_T, \tag{14}$$

so everything hinges on how the internal energy depends on volume at fixed temperature. That derivative follows from $dU = T\,dS - p\,dV$ together with the Maxwell relation $(\partial S/\partial V)_T = (\partial p/\partial T)_V$ established in T3:

$$\left(\frac{\partial U}{\partial V}\right)_T = T\left(\frac{\partial p}{\partial T}\right)_V - p. \tag{15}$$

The combination on the right is sometimes called the internal pressure. It measures the work that must be done against intermolecular attraction to separate the molecules at fixed temperature, and it vanishes identically for an ideal gas, where $p = nRT/V$ makes $T(\partial p/\partial T)_V = p$. Hence

$$\mu_{\mathrm{J}} = -\frac{1}{C_V}\left[T\left(\frac{\partial p}{\partial T}\right)_V - p\right], \qquad \mu_{\mathrm{J}}^{\mathrm{ideal}} = 0. \tag{16}$$

For a van der Waals gas, $p = RT/(v-b) - a/v^2$ per mole gives $(\partial p/\partial T)_v = R/(v-b)$ and therefore an internal pressure a/v^2, so $\mu_{\mathrm{J}} = -a/(C_V v^2) < 0$ whenever the attractive constant a is positive: the gas cools on free expansion, always, at every temperature. There is no inversion here.

That absence is worth noticing, because it tells us that free expansion and throttling are genuinely different questions. Joule's own attempt to measure μ_{J} failed: the effect is small, and the heat capacity of the apparatus swamped the temperature change of the gas. The experiment that worked, and that matters industrially, is the steady-flow one.

2.3. When is throttling isenthalpic?

2.3.1 The plug experiment

Consider a tube containing a porous plug, with gas maintained at pressure p_1 upstream and $p_2 < p_1$ downstream, the whole thing insulated. Rather than think about the messy flow inside the plug, follow a fixed parcel of gas from one side to the other and account for the work done on it by the gas behind and in front.

Let the parcel occupy volume V_1 at (p_1, T_1) before and V_2 at (p_2, T_2) after. The upstream gas acts as a piston pushing the parcel through at constant pressure p_1, doing work $+p_1V_1$ on it; the parcel in turn pushes the downstream gas out of the way at constant pressure p_2, which costs it p_2V_2. With no heat crossing the wall, the first law gives

$$U_2 - U_1 = p_1V_1 - p_2V_2 \qquad \Longrightarrow \qquad U_1 + p_1V_1 = U_2 + p_2V_2, \tag{17}$$

that is, $H_1 = H_2$. The combination $U + pV$ appears because the flow work that pushes fluid across a boundary is inseparable from the energy the fluid carries; this is exactly the circumstance in which enthalpy is the natural variable, and it is the reason enthalpy is the right potential for open devices.

Conditions for an isenthalpic throttle

$H_1 = H_2$ requires all of the following. The flow is **steady**, so no energy accumulates inside the control volume. The device is **adiabatic**, $\dot{Q} = 0$. There is **no shaft work**, $\dot{W}_s = 0$. There is **one inlet and one outlet**, so a single mass flow rate applies. And the **kinetic and potential energy changes are negligible** compared with the enthalpy change, $\frac{1}{2}(\mathrm{v}_2^2 - \mathrm{v}_1^2) \ll h_1 - h_2$ and $g(z_2 - z_1)$ likewise.

The kinetic-energy condition is not a formality. The general steady-flow energy balance per unit mass is

$$q - w_s = \Delta h + \tfrac{1}{2}\Delta \mathrm{v}^2 + g\,\Delta z, \tag{18}$$

where v is the flow speed, and what is conserved when $q = w_s = 0$ is the full stagnation combination $h + \mathrm{v}^2/2 + gz$. A gas emerging from a valve into a wide pipe has small velocity on both sides and the correction is negligible; a gas emerging into a nozzle at high speed does not, and treating a nozzle as isenthalpic is simply wrong. For a typical laboratory throttle the velocity term is smaller than the enthalpy term by several orders of magnitude, but the statement to remember is that it is small, not that it is absent.

Key Point

$h_1 = h_2$ relates the two *end states*. It does not say that enthalpy is constant along the way. Inside the plug the fluid is not in equilibrium: it has no well-defined pressure or temperature, and no thermodynamic path connects the two points. The "isenthalpic curve" drawn on a p–T diagram is the locus of downstream states obtained by varying p_2 with p_1, T_1 fixed, not a trajectory the fluid follows.

Throttling is, of course, thoroughly irreversible. The entropy of the gas rises, $s_2 > s_1$, while no entropy enters: all of it is generated inside. We put a number on the cost of that in Part III.

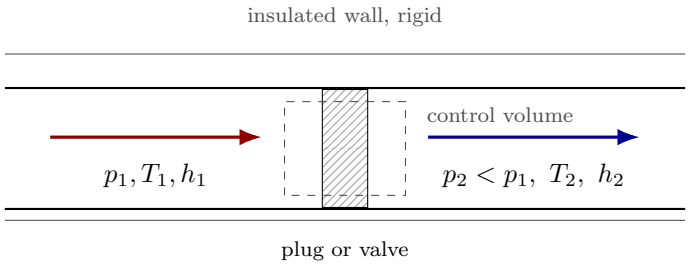

Figure 3: Steady flow through a plug. With no heat crossing the insulated wall, no shaft inside the control volume, steady conditions, and negligible change in flow speed, the steady-flow energy balance (18) reduces to $h_1 = h_2$. The two labelled states are equilibrium states of the fluid in the straight sections of pipe; the fluid inside the plug is not in equilibrium and carries no well-defined p or T.

2.3.2 The Joule–Thomson coefficient

We want the temperature change per unit pressure change along the locus of constant enthalpy. Start from

$$dh = T\,ds + v\,dp, \tag{19}$$

per unit mass or per mole as convenient, and hold T fixed. Using the Maxwell relation $(\partial s/\partial p)_T = -(\partial v/\partial T)_p$, which follows in T3 from the integrability of the Gibbs free energy,

$$\left(\frac{\partial h}{\partial p}\right)_T = T\left(\frac{\partial s}{\partial p}\right)_T + v = v - T\left(\frac{\partial v}{\partial T}\right)_p. \tag{20}$$

Since $(\partial h/\partial T)_p = c_p$ by definition, the triple product rule gives the result we are after.

The Joule–Thomson coefficient

For a single-component fluid of fixed composition in a steady, adiabatic, shaft-free flow with negligible kinetic and potential energy changes,

$$\mu_{\mathrm{JT}} \equiv \left(\frac{\partial T}{\partial p}\right)_h = \frac{1}{c_p}\left[T\left(\frac{\partial v}{\partial T}\right)_p - v\right] = \frac{v}{c_p}\left(T\alpha - 1\right), \tag{21}$$

where $\alpha = v^{-1}(\partial v/\partial T)_p$ is the isobaric expansivity at fixed composition and c_p the heat capacity at constant pressure, both evaluated at the upstream state. Since a throttle always has $dp < 0$, the fluid **cools if and only if $\mu_{\mathrm{JT}} > 0$, that is $T\alpha > 1$**. For an ideal gas $\alpha = 1/T$ exactly, so $\mu_{\mathrm{JT}} = 0$ at every temperature and pressure.

Equation (21) repays a moment's reading. The two terms in the bracket are the two things that happen when a fluid is pushed to lower pressure. Pulling molecules apart against their mutual attraction costs energy and tends to lower the temperature; that is the $T(\partial v/\partial T)_p$ term, which grows with the expansivity. Meanwhile the flow work done on the parcel, $p_1 v_1$ in and $p_2 v_2$ out, need not balance, and the $-v$ term records that. Whichever term wins decides the sign, and which one wins depends on the state.

The ideal-gas result is more than a special case; it is the baseline against which everything else is measured. If μ_{JT} were zero for real gases, there would be no Linde process and no cheap liquid nitrogen.

2.4. What real gases do

2.4.1 Low density: the second virial coefficient

At low enough density we may write the equation of state as $v = RT/p + B(T) + O(p)$, where $B(T)$ is the second virial coefficient with dimensions of molar volume. Substituting into the bracket in (21),

$$c_p\,\mu_{\mathrm{JT}} \;\xrightarrow[p\to 0]{}\; T\frac{dB}{dT} - B \;=\; T^2\,\frac{d}{dT}\!\left(\frac{B}{T}\right). \tag{22}$$

The second form is the useful one. It says that the zero-pressure Joule–Thomson coefficient changes sign exactly where B/T is stationary, so the *maximum inversion temperature* — the temperature above which no amount of throttling from low pressure will cool the gas — is the temperature at which B/T is a maximum. This identification is the standard route to tabulated inversion temperatures, and it is how the values quoted below were obtained.[1]

Notice that B alone does not settle the matter. The Boyle temperature, where $B = 0$, is a different temperature from the inversion temperature, where $(B/T)' = 0$; the two coincide only by accident.

2.4.2 Van der Waals: an inversion curve we can draw

To see the full inversion curve rather than only its zero-pressure end, we need a model with attraction and volume exclusion. The van der Waals equation

$$\left(p + \frac{a}{v^2}\right)(v - b) = RT \tag{23}$$

[1] R. C. Hendricks, I. C. Peller, and A. K. Baron, *Joule–Thomson Inversion Curves and Related Coefficients for Several Simple Fluids*, NASA TN D-6807, Lewis Research Center (July 1972), Appendix B and Table I.

is too crude to be quantitative, as we will confirm, but it produces the right shape and it can be worked through completely. Differentiating (23) at constant p and substituting into (21) gives, after rearrangement,

$$c_p\,\mu_{\mathrm{JT}} = \frac{\lambda v - b}{1-\lambda}, \qquad \lambda(T,v) \equiv \frac{2a(v-b)^2}{RTv^3}. \tag{24}$$

Two checks. As $v \to \infty$ at fixed T we have $\lambda v \to 2a/RT$ and $\lambda \to 0$, so $c_p\mu_{\mathrm{JT}} \to 2a/(RT) - b$, which is exactly (22) with the van der Waals second virial coefficient $B = b - a/(RT)$. And the denominator vanishes when $RTv^3 = 2a(v-b)^2$, which is precisely the condition $(\partial p/\partial v)_T = 0$: the expression blows up on the spinodal locus, where the van der Waals isotherm turns over and the model describes a mechanically unstable state. We return to that pathology below.

Setting the numerator of (24) to zero gives the inversion condition,

$$2a\,(v-b)^2 = R\,T_{\mathrm{inv}}\,b\,v^2 \qquad \Longleftrightarrow \qquad T_{\mathrm{inv}}(v) = \frac{2a}{Rb}\left(1 - \frac{b}{v}\right)^2. \tag{25}$$

Eliminating v between (25) and the equation of state (23) gives the inversion curve in the p–T plane. The algebra is routine and the result is compact in the reduced variables $T_{\mathrm{r}} = T/T_{\mathrm{c}}$ and $p_{\mathrm{r}} = p/p_{\mathrm{c}}$, using the van der Waals critical point $T_{\mathrm{c}} = 8a/(27Rb)$, $p_{\mathrm{c}} = a/(27b^2)$, $v_{\mathrm{c}} = 3b$:

$$p_{\mathrm{r}} = 24\sqrt{3T_{\mathrm{r}}} - 12\,T_{\mathrm{r}} - 27, \qquad \tfrac{3}{4} \le T_{\mathrm{r}} \le \tfrac{27}{4}. \tag{26}$$

The curve (26) is a closed arch, plotted in Figure 4. Three features of it are worth naming.

- It meets $p_{\mathrm{r}} = 0$ at $T_{\mathrm{r}} = 27/4 = 6.75$ and at $T_{\mathrm{r}} = 3/4$. The upper root is the maximum inversion temperature, $T_{\mathrm{inv}}^{\max} = 2a/(Rb) = 6.75\,T_{\mathrm{c}}$; above it the gas warms on throttling from any pressure. The lower root is a low-temperature inversion, reached deep in the liquid-like region where the model is least trustworthy.
- It has a maximum at $(T_{\mathrm{r}}, p_{\mathrm{r}}) = (3, 9)$. Above nine times the critical pressure, this model says throttling never cools, at any temperature.
- Inside the arch $\mu_{\mathrm{JT}} > 0$ and throttling cools; outside it, $\mu_{\mathrm{JT}} < 0$ and throttling warms. The critical point $(T_{\mathrm{r}}, p_{\mathrm{r}}) = (1, 1)$ lies inside.

2.4.3 How good is this, in numbers?

Take nitrogen, and fix the van der Waals constants from the measured critical point rather than from a table of a and b, using $T_{\mathrm{c}} = 126.2\,\mathrm{K}$ and $p_{\mathrm{c}} = 3.4\,\mathrm{MPa}$.[2] Then $b = RT_{\mathrm{c}}/(8p_{\mathrm{c}}) = 3.86 \times 10^{-5}\,\mathrm{m^3\,mol^{-1}}$ and $a = 27R^2T_{\mathrm{c}}^2/(64p_{\mathrm{c}}) = 0.137\,\mathrm{Pa\,m^6\,mol^{-2}}$.

For the coefficient itself at low pressure and room temperature, take $T = 300\,\mathrm{K}$ and a diatomic $c_p = \frac{7}{2}R = 29.1\,\mathrm{J\,mol^{-1}K^{-1}}$:

$$\mu_{\mathrm{JT}} \simeq \frac{1}{c_p}\left(\frac{2a}{RT} - b\right) = \frac{1.10\times10^{-4} - 0.39\times10^{-4}}{29.1}\ \mathrm{K\,Pa^{-1}} = 2.4\times10^{-6}\,\mathrm{K\,Pa^{-1}} \approx 0.24\ \mathrm{K\,bar^{-1}}. \tag{27}$$

[2] Critical constants and all measured or equation-of-state values in this subsection are from NASA TN D-6807, Table I (inversion and Boyle temperatures) and Table II (critical constants and critical Joule–Thomson coefficients).

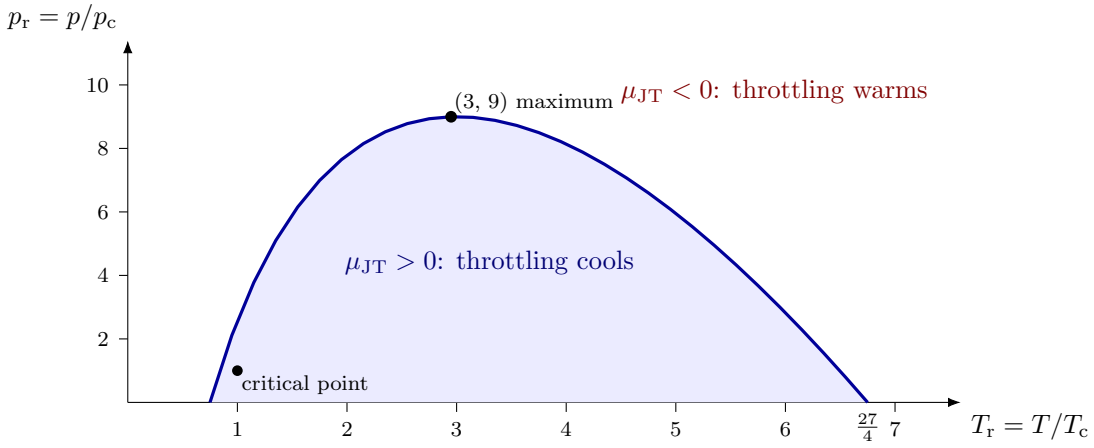

Figure 4: The van der Waals inversion curve (26) in reduced coordinates. Throttling cools only for upstream states inside the shaded arch. The curve is a property of the model, not of a particular gas: every van der Waals fluid has the same reduced inversion curve. Real fluids follow an arch of the same shape but a noticeably smaller one, as Table 3 shows.

Quantity for N_2	van der Waals	From an accurate equation of state	Ratio
maximum inversion temperature	$6.75\,T_c = 852\,\mathrm{K}$	$627\,\mathrm{K}$ $(4.96\,T_c)$	1.36
Boyle temperature	$3.375\,T_c = 426\,\mathrm{K}$	$321\,\mathrm{K}$ $(2.54\,T_c)$	1.33

Table 3: The van der Waals inversion curve has the right shape and the wrong size. Values in the third column are from NASA TN D-6807, Table I, obtained by maximizing B/T for the Bender equation of state; the same report gives 644 K from a second nitrogen equation of state, which indicates the spread between careful determinations. The overestimate of about a third is characteristic: the same report notes that measured reduced inversion temperatures cluster near $5.1\,T_c$ across simple fluids, against the van der Waals prediction of $6.75\,T_c$.

The number is positive, which is the qualitatively important part: nitrogen at room temperature is well inside its inversion arch, and it cools when throttled. Its size should be treated as an estimate of order rather than a measurement. For an independent check at a different state point, the same NASA report gives the Joule–Thomson coefficient of nitrogen at its critical point as $6.2\,\mathrm{K\,MPa^{-1}}$ from an accurate equation of state, against $5.7\,\mathrm{K\,MPa^{-1}}$ measured by Roebuck and Osterberg; the report also observes that isenthalpic measurements are difficult enough that calculations from a good equation of state are often the more reliable of the two.

2.4.4 Where the model fails, and how we can tell

The divergence of (24) at $\lambda = 1$ is instructive rather than embarrassing. It sits on the locus $(\partial p/\partial v)_T = 0$, which is the spinodal; on that locus the compressibility κ_T is infinite, the expansivity α is infinite with it, and the numerator of (21) diverges. In a real fluid c_p diverges there as well, and the ratio remains finite — which is why μ_{JT} is perfectly measurable at the critical point, as the value quoted above shows. The van der Waals expression misses this because it is combined with a constant c_p; the fault lies in that pairing, not in (21).

Key Point

Two statements are often run together and should not be. *Throttling cools gases* is false as stated. *Throttling cools a gas whose upstream state lies inside its inversion curve* is correct, and for helium and hydrogen at room temperature the condition fails.

Helium is the extreme case. Its maximum inversion temperature is a few tens of kelvin, and published determinations disagree substantially: the NASA report collects estimates of 45 K, 54 K, and 47.7 K from different authors, an observed value of 34 K attributed to Crawford, and 55 K from its own second-virial calculation. Whatever the right value in that range, it is far below room temperature, and helium throttled from a room-temperature cylinder warms up. Hydrogen is in the same situation. This is not a detail of cryogenic engineering; it determined the order in which the gases were liquefied historically, and it forces the cascade design described next.

2.5. Turning the effect into liquid

2.5.1 The Linde–Hampson cycle

The Joule–Thomson temperature drop available in one pass is modest — tens of kelvin, not hundreds. The idea that makes liquefaction possible is to use the cold gas that has already been expanded to pre-cool the incoming gas, so that each pass starts colder than the last.

Draw a control volume around the heat exchanger, the valve, and the liquid reservoir, cutting the pipes where the fluid is at ambient temperature T_1. Feed enters at (T_1, p_1) with mass flow $\dot m$; liquid is withdrawn at the low pressure with flow $\dot m_f$; the remaining gas leaves at (T_1, p_2), having been warmed back to ambient by the exchanger. With the whole assembly insulated and no shaft inside, (18) gives

$$\dot m\, h(T_1, p_1) = \dot m_f\, h_f + (\dot m - \dot m_f)\, h(T_1, p_2), \tag{28}$$

and solving for the liquid yield $y = \dot m_f/\dot m$,

$$y = \frac{h(T_1, p_2) - h(T_1, p_1)}{h(T_1, p_2) - h_f} \equiv \frac{\Delta h_T}{h(T_1, p_2) - h_f}. \tag{29}$$

The denominator is positive and roughly fixed by the latent heat. Everything therefore depends on the numerator Δh_T, the *isothermal* enthalpy change on going from high to low pressure at ambient temperature. Writing it as an integral of (20),

$$\Delta h_T = -\int_{p_2}^{p_1} \left(\frac{\partial h}{\partial p}\right)_T dp = \int_{p_2}^{p_1} c_p\, \mu_{\mathrm{JT}}\, dp, \tag{30}$$

we see that liquid is produced only if $\mu_{\mathrm{JT}} > 0$ over the pressure range at ambient temperature. That is the precise sense in which a Linde cycle requires the feed to sit inside the inversion curve, and it is why helium and hydrogen must be pre-cooled — by liquid nitrogen, or by a separate cycle using a different working fluid — before a throttle will do anything useful for them.

Where should the high pressure be set?

Differentiating (29) with respect to p_1 at fixed T_1 and p_2, the yield is stationary when $(\partial h/\partial p)_T = 0$ at the feed state — that is, when the feed sits exactly on the inversion curve. Maximum yield per unit of gas circulated is therefore obtained by compressing to the inversion pressure.

This is not the same as the best design. The compressor work required to reach p_1 grows with p_1, and the cost per unit of liquid is minimized well inside the inversion curve, at pressures substantially below the yield-maximizing value. The NASA report makes the point explicitly, and notes that the early Linde air process ran at an intermediate pressure near 40 atm. The lesson generalizes: a thermodynamic optimum for one figure of merit is not an optimum for another, and the figure of merit has to be named before the word "optimal" means anything.

2.5.2 The alternative: expand against something

A throttle destroys the pressure difference and delivers no work. If instead the gas drives a turbine and the expansion is close to reversible and adiabatic, the relevant coefficient is

$$\left(\frac{\partial T}{\partial p}\right)_s = \frac{T}{c_p}\left(\frac{\partial v}{\partial T}\right)_p = \frac{Tv\alpha}{c_p}, \tag{31}$$

which follows from setting $ds = (c_p/T)dT - (\partial v/\partial T)_p\, dp = 0$. Compare this with (21): the $-v$ term is gone. For any fluid with $\alpha > 0$ — which is nearly all of them, water below 4°C being the standard exception — the right-hand side is positive, so isentropic expansion cools at *every* temperature. There is no inversion curve for an expander, and even helium cools.

The expander also recovers work rather than discarding it, so it is thermodynamically the better device. It is used, in the Claude and Collins cycles among others, wherever it can be. Its limitation is mechanical: liquid droplets damage turbine blades, so the last stage down to the liquid, where the fluid becomes two-phase, is normally still done with a throttle. Real liquefiers therefore use an expander for the bulk of the cooling and a Joule–Thomson valve at the cold end.

2.6. What the pressure drop cost us

We can now say what happens at a valve, but not yet what it costs. Both devices in the last section take a fluid from p_1 to p_2. The expander hands us work; the throttle hands us nothing and raises the entropy of the fluid. The difference between them is the whole of the loss, and to put a number on it we need a reference: some specified environment against which "useful work" is measured.

That is the subject of Part III, which returns to the two bricks of Part I and asks what the maximum extractable work is when the bodies are finite, then defines availability relative to a stated environment and derives the lost-work relation in general form.

3 Part III. Finite Reservoirs, Availability, and Lost Work

3.1. Back to the two bricks

Return to the opening question of Part I. Two finite bodies, a hot one at T_1 and a cold one at $T_2 < T_1$, are otherwise isolated. We may run a reversible engine between them for as long as their temperatures differ. How much work can we extract in total, and what final state are the bodies left in?

Because the bodies are finite, we cannot use $\eta = 1 - T_2/T_1$ with the initial temperatures and multiply by the available heat: the temperatures are the initial ones only at the first instant. As the engine runs, the hot body cools and the cold body warms, the instantaneous Carnot efficiency $1 - T_{\mathrm{c}}(t)/T_{\mathrm{h}}(t)$ shrinks, and the process stops when the two temperatures meet. What we can use is the pair of conservation laws that hold over the whole process, and this is the key idea of the calculation: rather than integrate the changing efficiency, we impose the two global constraints and let them fix the answer.

3.1.1 Two constraints fix the answer

Model each body as having a constant heat capacity, C_1 and C_2, so that a temperature change dT changes its energy by $C_i\,dT$ and its entropy by $C_i\,dT/T$. This is the one substantive assumption, and we flag it: real bodies have temperature-dependent heat capacities, and the closed-form results below are exactly the constant-C case. Let the bodies end at temperatures T_1^f and T_2^f.

The engine operates in a cycle, so over the whole process it returns to its own initial state and contributes nothing to either balance. For *maximum* work the engine must be reversible, which means the total entropy is unchanged:

$$\Delta S_{\mathrm{tot}} = C_1 \ln \frac{T_1^f}{T_1} + C_2 \ln \frac{T_2^f}{T_2} = 0. \tag{32}$$

Energy conservation says the work delivered is whatever energy the two bodies gave up:

$$W_{\max} = C_1\,(T_1 - T_1^f) + C_2\,(T_2 - T_2^f). \tag{33}$$

These are two equations in the two unknown final temperatures, once we add the physical condition that a reversible engine stops driving when the temperatures equalize, $T_1^f = T_2^f \equiv T_f$.

That last condition deserves a word, since it is a conclusion and not a definition. As long as $T_1^f \neq T_2^f$ a reversible engine could still run between them and extract more work, so the work is maximized only when the bodies reach a common temperature. Setting $T_1^f = T_2^f = T_f$ in (32),

$$C_1 \ln T_f + C_2 \ln T_f = C_1 \ln T_1 + C_2 \ln T_2, \tag{34}$$

and solving,

Final temperature and maximum work for finite bodies

For two bodies of constant heat capacities C_1, C_2 coupled through a reversible engine until they reach a common temperature,

$$T_f = \left(T_1^{C_1}\, T_2^{C_2}\right)^{1/(C_1+C_2)}, \tag{35}$$

the *entropy-weighted geometric mean* of the initial temperatures, and the maximum work is

$$W_{\max} = C_1(T_1 - T_f) + C_2(T_2 - T_f). \tag{36}$$

For equal capacities $C_1 = C_2 = C$ this collapses to $T_f = \sqrt{T_1 T_2}$ and

$$W_{\max} = C\left(\sqrt{T_1} - \sqrt{T_2}\right)^2. \tag{37}$$

The results assume constant heat capacities and a reversible engine; the reservoirs are not assumed large.

The final temperature is the geometric mean, not the arithmetic mean. That is the whole content of the calculation, and it is worth seeing why. If we simply let the bodies touch, energy conservation alone gives the arithmetic mean $T_f^{\text{irr}} = (C_1T_1 + C_2T_2)/(C_1 + C_2)$, and no work at all. The reversible engine reaches a *lower* common temperature, because it carries energy out of the bodies as work rather than leaving it all behind as heat, and the geometric mean is always below the arithmetic mean. The gap between the two final temperatures is exactly the room the second law leaves for extracting work.

Two equal bodies of water

Take two 1 kg blocks of water, $C = 4.18\,\text{kJ K}^{-1}$ each, at $T_1 = 360\,\text{K}$ and $T_2 = 280\,\text{K}$.

Reversible engine. $T_f = \sqrt{(360)(280)} = 317.5\,\text{K}$, and

$$W_{\max} = 4.18\left(\sqrt{360} - \sqrt{280}\right)^2 = 21.0\,\text{kJ}.$$

The hot block supplies $C(T_1 - T_f) = 177.7\,\text{kJ}$ of heat and the cold block absorbs $C(T_f - T_2) = 156.7\,\text{kJ}$; the difference, 21.0 kJ, is the work, as energy conservation requires.

Direct thermal contact, for comparison. $T_f^{\text{irr}} = (360 + 280)/2 = 320\,\text{K}$, no work, and

$$S_{\text{gen}} = C\ln\frac{(T_f^{\text{irr}})^2}{T_1T_2} = 4.18\ln\frac{320^2}{(360)(280)} = 0.0658\,\text{kJ K}^{-1}.$$

The overall conversion is $W_{\max}/Q_{\text{in}} = 21.0/177.7 = 11.8\%$, well below the initial-instant Carnot value $1 - 280/360 = 22.2\%$ — because that efficiency applies only at the first instant, and falls throughout as the temperatures close.

3.2. Lost work when the engine is imperfect

Nothing forces the engine to be reversible. Suppose it generates entropy S_{gen} over the process while still bringing the two equal-capacity bodies to a common temperature. Energy conservation gives

the common temperature from the total energy, and the second law fixes how much of the energy difference can leave as work. Carrying the algebra through for equal capacities, the work delivered is

$$W(S_{\text{gen}}) = C(T_1 + T_2) - 2C\sqrt{T_1 T_2}\ \exp\left(\frac{S_{\text{gen}}}{2C}\right), \tag{38}$$

which reduces to (37) at $S_{\text{gen}} = 0$ and decreases monotonically as S_{gen} grows. The work runs out entirely — $W = 0$ — when $S_{\text{gen}} = 2C\ln\!\big[(T_1 + T_2)/(2\sqrt{T_1 T_2})\big]$, which is precisely the S_{gen} of direct contact computed in the example. That is the consistency check: generating the full direct-contact entropy leaves no work, as it must.

Expanding (38) for small S_{gen},

$$W \approx W_{\max} - T_f\, S_{\text{gen}} + O(S_{\text{gen}}^2), \qquad T_f = \sqrt{T_1 T_2}, \tag{39}$$

so the leading cost of a little irreversibility is $T_f S_{\text{gen}}$, with T_f the common final temperature. This is the same "lost work = temperature × entropy generated" structure we met in Part I, and it again picks out a specific temperature: not the environment's, not either body's initial value, but the temperature at which the generated entropy is finally deposited. The lesson is that the temperature in the lost-work relation is always the one where the entropy actually leaves the process, and identifying that temperature is part of the physics, not a convention to be fixed in advance.

3.3. Availability: work relative to an environment

The two-body problem had a natural endpoint — the bodies equalized. Most practical questions are different: we have one system, and around it an effectively infinite environment at fixed temperature T_0 and pressure p_0, and we want the most work extractable as the system comes to equilibrium with that environment. Now the environment is a genuine reservoir, but a new subtlety appears: if the system expands, it has to push the atmosphere out of the way, and that part of its work is not available to us.

3.3.1 Useful work and the dead state

Distinguish the total work the system does on its surroundings from the *useful* work we can actually harvest. If the system volume changes by dV against an atmosphere at p_0, an amount $p_0\, dV$ of the work merely displaces the atmosphere and is unrecoverable. The useful work is

$$đW_{\text{u}} = đW_{\text{by system}} - p_0\, dV. \tag{40}$$

Likewise, heat exchanged with the environment is exchanged at T_0, so it carries entropy $đQ_0/T_0$ that we must account for. The state in which the system has come to complete equilibrium with the environment — same temperature T_0, same pressure p_0, and, if matter is exchanged, same chemical potential μ_0 — is the *dead state*. A system at the dead state can deliver no further useful work, by definition: it is the zero of the availability scale.

The reference environment

Availability is defined only after an environment is specified: a reservoir large enough that exchanging heat, volume, or matter with it does not change its intensive properties T_0, p_0, μ_0. Different choices of environment give different numerical availabilities for the same system. The environment must be named before "available work" has a value.

3.3.2 The availability function

Consider a closed system (no matter exchange) that can exchange heat and work only with the environment, going from some state to the dead state. Over the process, the first law gives $\Delta U = Q_0 + W_{\text{on}}$, where Q_0 is the heat absorbed from the environment and W_{on} the work done on the system. The second law, with all heat crossing at T_0, gives $\Delta S_{\text{tot}} = \Delta S - Q_0/T_0 \geq 0$. Eliminating Q_0 and separating out the atmospheric work $p_0 \Delta V$, the useful work we can extract satisfies

$$W_{\text{u}} \;\leq\; -\big(\Delta U + p_0\,\Delta V - T_0\,\Delta S\big). \tag{41}$$

This motivates the state function

$$A \equiv U + p_0 V - T_0 S, \tag{42}$$

the *availability* (or exergy, when measured from the dead state). Its change sets the ceiling on useful work: for a closed system exchanging heat and work only with an environment at (T_0, p_0),

$$W_{\text{u}} \leq -\Delta A = A_{\text{i}} - A_{\text{f}}, \tag{43}$$

with equality for a reversible process. Note the mixed nature of A: it uses the system's own U, V, S but the *environment's* T_0 and p_0 as coefficients. That is what makes it the right bookkeeping quantity — it charges the system for the atmosphere it displaces and credits it only for entropy the environment will accept.

Key Point

The availability $A = U + p_0 V - T_0 S$ resembles the Gibbs free energy $G = U + pV - TS$ but is not the same object. In G the coefficients are the system's own p and T; in A they are the environment's fixed p_0 and T_0. The two coincide only when the system happens to sit at the environment's temperature and pressure, and it is precisely away from that condition that availability is interesting.

3.3.3 The general lost-work relation

The deficit in (43) has a clean form. Define the lost work as the shortfall from the reversible maximum,

$$W_{\text{lost}} \equiv (-\Delta A) - W_{\text{u}} \geq 0. \tag{44}$$

Tracking the entropy through the derivation above, every bit of this shortfall is generated entropy delivered to the environment at T_0:

Gouy–Stodola relation

For a process exchanging heat with a single environment at temperature T_0, the work lost to irreversibility is

$$W_{\text{lost}} = T_0\, S_{\text{gen}}, \tag{45}$$

where $S_{\text{gen}} \geq 0$ is the total entropy generated. Each unit of entropy generated destroys T_0 units of otherwise-available work. The relation requires a single environment temperature T_0 at which the generated entropy is ultimately rejected; when entropy leaves at some other temperature, that temperature replaces T_0, as the finite-body result (39) showed.

The result is due to Gouy and, for open systems, to Stodola.[3] It converts every entropy-generation estimate into a work penalty in the same units, which is why it is the backbone of exergy analysis: a plant's inefficiencies can be located by finding where S_{gen} is produced and multiplying by T_0.

3.4. Steady-flow availability and open devices

Most real devices are open: fluid streams through them. The steady-flow first law from Part II, per unit mass with no accumulation, is

$$q - w_s = \Delta h + \tfrac{1}{2}\Delta \mathrm{v}^2 + g\,\Delta z, \tag{46}$$

and the corresponding entropy balance for a device exchanging heat only with the environment at T_0 is $\Delta s = q/T_0 + s_{\mathrm{gen}}$, with $s_{\mathrm{gen}} \geq 0$. Eliminating q and neglecting the kinetic and potential terms, the useful (shaft) work per unit mass obeys

$$w_s \leq -\Delta(h - T_0 s) \equiv -\Delta b, \qquad b \equiv h - T_0 s, \tag{47}$$

where b is the *flow availability* or stream exergy (again measured relative to the dead state for an absolute value). The shortfall is once more $w_{s,\mathrm{lost}} = T_0\, s_{\mathrm{gen}}$.

The atmospheric-work term $p_0 \Delta V$ that appeared for the closed system is absent here: in steady flow the device volume is fixed, and the pV work of pushing fluid across the boundaries is already inside $h = u + pv$. This is the same reason enthalpy was the natural variable for the throttle in Part II, and it is worth pausing on: the closed-system availability uses $U + p_0 V$, while the open-system availability uses h, and the difference is exactly the flow work. Getting this right is the difference between a correct and an incorrect exergy balance for a turbine or compressor.

The throttle costed at last

A throttle has $w_s = 0$ and, being adiabatic, $q = 0$; from (47) with $\Delta h = 0$ its entropy generation per unit mass is $s_{\mathrm{gen}} = \Delta s = s_2 - s_1 > 0$, and the flow availability destroyed is $-\Delta b = T_0(s_2 - s_1)$.

For an ideal gas throttled isothermally at the environment temperature from $p_1 = 10\,\mathrm{bar}$ to $p_2 = 1\,\mathrm{bar}$, with $\Delta s = R\ln(p_1/p_2) = R\ln 10$, the destroyed availability is

$$T_0\,\Delta s = RT_0 \ln 10 = (8.314)(300)\ln 10 = 5.74\,\mathrm{kJ\,mol^{-1}}.$$

That is the price of the throttle: the same pressure drop through an ideal reversible expander would have delivered $5.74\,\mathrm{kJ\,mol^{-1}}$ of shaft work at this temperature, and the throttle delivers none. This quantifies the qualitative claim at the end of Part II — the expander is thermodynamically superior precisely by $T_0\, s_{\mathrm{gen}}$ — and it is why liquefiers use an expander wherever liquid droplets do not forbid it.

3.5. What this unit answered, and the question it leaves

We set out to bound the work a process can deliver and to say what sets the limit. The answer came in three layers. Between fixed reservoirs the bound is the Carnot fraction $1 - T_{\mathrm{c}}/T_{\mathrm{h}}$, from cyclic

[3] L. G. Gouy, "Sur l'énergie utilisable," *Journal de Physique Théorique et Appliquée* **8**, 501–518 (1889); A. Stodola's treatment of open systems appeared in his work on steam turbines (1898/1903). Bibliographic details taken from the reference list of Y. Tu and G. Chen, "Rethinking loss of available work and Gouy–Stodola theorem," arXiv:2307.16041 (2023), which also cautions that the appropriate temperature is the one at which entropy is actually rejected, not automatically the ambient T_0 — exactly the qualification attached to (45) and illustrated by the finite-body case.

operation and the second law alone. For the working fluid, whether an expansion cools depends entirely on the equation of state, through $\mu_{\mathrm{JT}} = (v/c_p)(T\alpha - 1)$, and the sign changes across an inversion curve. Between finite bodies, or against a specified environment, the bound is the drop in an availability function, and the shortfall from reversibility is $T_0 S_{\mathrm{gen}}$ — entropy generation, priced in work.

A single quantity runs through all three: entropy generation is what separates the real process from its reversible bound, and $S_{\mathrm{gen}} = 0$ is the signature of the best we can do.

That raises the question this unit cannot answer. Every calculation here ended at a stationary state — the equalized bricks, the dead state, the final temperature of a reversible engine. We found these by setting a first variation of the entropy to zero, or equivalently by imposing conservation and reversibility. But a stationary state need not be a stable one. If we perturb the final state slightly — move a little energy back, let a small volume fluctuation occur — does the system return, or does it run away? The equalized-temperature state feels stable, but we have not shown it, and there are systems whose apparent equilibria are not stable at all.

The tool we used, the vanishing of the first variation, is silent on this. To decide stability we need the *second* variation of the entropy, and the conditions it imposes on the response functions — the heat capacities and compressibilities that appeared throughout this unit. Whether the stationary states we have been computing are stable against all permitted disturbances is the subject of T5.

4 Problems

Problems are numbered consecutively across the unit; solution sketches follow in Appendix A.

Part I: cycles and the reversible bound

Problems are numbered consecutively across the three parts of T4; solution sketches are in Appendix A of the part in which the problem appears.

1. **Two routes, one pair of states.** One mole of ideal gas at $T_1 = 300\,\mathrm{K}$ occupies V_1. It is brought to volume $2V_1$ at the same temperature in two ways: (a) quasi-statically and isothermally in contact with a reservoir at 300 K; (b) by free expansion into an evacuated chamber of equal volume, insulated throughout. For each route compute W, Q, ΔU, ΔS of the gas, ΔS of the reservoir, and S_{gen}. Then state, in one sentence, which of these seven quantities are equal for the two routes and why.

2. **A cycle that is not reversible.** A cyclic device draws $Q_{\mathrm{h}} = 1000\,\mathrm{J}$ from a reservoir at 600 K, rejects heat to a reservoir at 300 K, and delivers $W_{\mathrm{out}} = 350\,\mathrm{J}$ per cycle. (a) Find Q_{c} and S_{gen} per cycle. (b) Find the work that a reversible device would have delivered from the same Q_{h}, and verify (13) numerically. (c) The manufacturer proposes a modification that keeps Q_{h} and Q_{c} fixed but raises W_{out} to 520 J. Show that this is impossible using the first law alone, and identify what it would take to reach 520 J legitimately.

Part II: gas cooling

3. **Inversion from the second virial coefficient.** A gas obeys $v = RT/p + B(T)$ with $B(T) = b_0 - a_0/(RT)$, where a_0 and b_0 are positive constants.

   (a) Show that $c_p \mu_{\mathrm{JT}} = T\,dB/dT - B$ for this equation of state, and evaluate it.

   (b) Find the zero-pressure inversion temperature T_{inv} and the Boyle temperature T_B, and show that $T_{\mathrm{inv}} = 2T_B$ for this model. Confirm that T_{inv} is where B/T is maximal.

   (c) Using nitrogen values $a_0 = 0.137\,\mathrm{Pa\,m^6\,mol^{-2}}$ and $b_0 = 3.86 \times 10^{-5}\,\mathrm{m^3\,mol^{-1}}$, evaluate T_{inv} and compare with the value near 630 K obtained from an accurate equation of state. Which way does the model err, and is that consistent with Table 3?

   (d) Estimate the temperature drop on throttling nitrogen from 60 bar to 1 bar at 300 K, taking $c_p = \tfrac{7}{2}R$. State one reason the estimate is too large and one quantity you would need in order to do better.

Part III: finite reservoirs and available work

4. **Unequal bodies.** A hot body has heat capacity $C_1 = 2C$ at $T_1 = 400\,\mathrm{K}$; a cold body has $C_2 = C$ at $T_2 = 300\,\mathrm{K}$, with $C = 3.0\,\mathrm{kJ\,K^{-1}}$.

   (a) Find the common final temperature T_f under reversible operation and the maximum work, using (35) and (36).

   (b) Find T_f^{irr} for direct thermal contact and the entropy generated. Verify that $T_f < T_f^{\mathrm{irr}}$.

   (c) Verify the small-irreversibility estimate: compute $T_0 S_{\mathrm{gen}}$ for direct contact taking $T_0 = T_f$, and compare with $W_{\max}$. Comment on why the agreement is only approximate here.

5. **An exergy balance for a heat exchanger.** A stream of hot water ($c = 4.18\,\mathrm{kJ\,kg^{-1}K^{-1}}$, $\dot m = 1\,\mathrm{kg\,s^{-1}}$) enters at 350 K and leaves at 320 K, giving up its heat to a cold stream that enters at 290 K and leaves at 315 K, in an insulated counterflow exchanger. The environment is at $T_0 = 290\,\mathrm{K}$.

   (a) Find the cold stream's mass flow rate from an energy balance.

   (b) Find the total entropy generation rate $\dot S_{\mathrm{gen}}$.

   (c) Find the rate of availability destruction $\dot W_{\mathrm{lost}} = T_0 \dot S_{\mathrm{gen}}$, and state what physical feature of the exchanger it is charging for.

A Solution sketches

Part I

Problem 1. Route (a): isothermal and quasi-static, so $\Delta U = 0$, $W_{\mathrm{out}} = \int p\,dV = RT\ln 2 = 1729\,\mathrm{J}$, hence $W = -1729\,\mathrm{J}$ and $Q = +1729\,\mathrm{J}$. The gas gains $\Delta S = R\ln 2 = 5.76\,\mathrm{J\,K^{-1}}$; the reservoir loses exactly $Q/T = 5.76\,\mathrm{J\,K^{-1}}$; $S_{\mathrm{gen}} = 0$.

Route (b): $W = 0$ (nothing resists), $Q = 0$ (insulated), so $\Delta U = 0$ and T is unchanged. The gas entropy change is the same, $\Delta S = R\ln 2 = 5.76\,\mathrm{J\,K^{-1}}$, because entropy is a state function and the endpoints are identical. The reservoir is absent, so its entropy change is zero and $S_{\mathrm{gen}} = 5.76\,\mathrm{J\,K^{-1}}$.

Equal for both routes: ΔU and ΔS of the gas, because both are state functions and the endpoints coincide. Everything else differs, because Q, W, and S_{gen} belong to the path. Note that route (b) never permits $\int p\,dV$; its work is zero, not $RT\ln 2$.

Problem 2. (a) $Q_{\mathrm c} = Q_{\mathrm h} - W_{\mathrm{out}} = 650\,\mathrm{J}$. The device is cyclic, so $S_{\mathrm{gen}} = Q_{\mathrm c}/T_{\mathrm c} - Q_{\mathrm h}/T_{\mathrm h} = 650/300 - 1000/600 = 2.167 - 1.667 = 0.500\,\mathrm{J\,K^{-1}}$.

(b) Reversible operation from the same $Q_{\mathrm h}$ gives $W_{\mathrm{out}}^{\mathrm{rev}} = Q_{\mathrm h}(1 - T_{\mathrm c}/T_{\mathrm h}) = 1000(1 - 1/2) = 500\,\mathrm{J}$. Check (13): $500 - T_{\mathrm c}S_{\mathrm{gen}} = 500 - (300)(0.500) = 350\,\mathrm{J}$, which matches. The efficiency is 0.35 against a bound of 0.50.

(c) With $Q_{\mathrm h}$ and $Q_{\mathrm c}$ both fixed, the first law over a cycle forces $W_{\mathrm{out}} = Q_{\mathrm h} - Q_{\mathrm c} = 350\,\mathrm{J}$ exactly; there is nowhere for the extra 170 J to come from, since $\Delta U = 0$ over a cycle. Reaching 520 J requires reducing $Q_{\mathrm c}$, and $Q_{\mathrm c} \geq Q_{\mathrm h}T_{\mathrm c}/T_{\mathrm h} = 500\,\mathrm{J}$ bounds how far that can go: the largest legitimate work at this $Q_{\mathrm h}$ is 500 J, so 520 J is impossible by the second law as well as unaccounted for by the first. To exceed 500 J one would have to change the reservoir temperatures or draw more heat.

Part II

Problem 3. (a) With $v = RT/p + B$, $(\partial v/\partial T)_p = R/p + dB/dT$, so $T(\partial v/\partial T)_p - v = RT/p + T\,dB/dT - RT/p - B = T\,dB/dT - B$. For $B = b_0 - a_0/(RT)$ we have $dB/dT = a_0/(RT^2)$, hence

$$c_p \mu_{\mathrm{JT}} = \frac{a_0}{RT} - b_0 + \frac{a_0}{RT} = \frac{2a_0}{RT} - b_0.$$

(b) $\mu_{\mathrm{JT}} = 0$ gives $T_{\mathrm{inv}} = 2a_0/(Rb_0)$. The Boyle temperature is where $B = 0$, so $T_B = a_0/(Rb_0) = T_{\mathrm{inv}}/2$. For the second claim, write $B/T = b_0/T - a_0/(RT^2)$; then $d(B/T)/dT = -b_0/T^2 + 2a_0/(RT^3)$, which vanishes at $T = 2a_0/(Rb_0) = T_{\mathrm{inv}}$, and the second derivative is negative there, so it is a maximum, as (22) requires.

(c) $T_{\mathrm{inv}} = 2(0.137)/[(8.314)(3.86\times10^{-5})] = 0.274/3.21\times10^{-4} = 854\,\mathrm{K}$. This exceeds the accurate value near 630 K by about 36%. The model errs high, consistent with Table 3 — unsurprisingly,

since $B = b_0 - a_0/(RT)$ is the van der Waals second virial coefficient and this is the same calculation in different clothing.

(d) At 300 K, $c_p \mu_{\mathrm{JT}} = 2(0.137)/[(8.314)(300)] - 3.86 \times 10^{-5} = 7.1 \times 10^{-5}\,\mathrm{m^3\,mol^{-1}}$, so $\mu_{\mathrm{JT}} = 2.4 \times 10^{-6}\,\mathrm{K\,Pa^{-1}} = 0.24\,\mathrm{K\,bar^{-1}}$ and $\Delta T \approx -(0.24)(59) \approx -14\,\mathrm{K}$.

The estimate is too large for two compounding reasons, either of which suffices. First, μ_{JT} was evaluated in the zero-pressure limit but applied across 59 bar; μ_{JT} decreases as the pressure rises toward the inversion curve, so the average value over the range is smaller than the value used. Second, μ_{JT} is evaluated at the upstream temperature, but the gas cools as it expands, and the correct calculation integrates along the isenthalp rather than multiplying by Δp. To do better one needs $\mu_{\mathrm{JT}}(T, p)$ over the region — equivalently $v(T, p)$ and $c_p(T, p)$ from an accurate equation of state — and one integrates $dT = \mu_{\mathrm{JT}}\, dp$ from the upstream state.

Part III

Problem 4. (a) From (35) with $C_1 = 2C$, $C_2 = C$: $T_f = (400^2 \cdot 300)^{1/3} = (4.8 \times 10^7)^{1/3} = 363.4\,\mathrm{K}$. Then $W_{\max} = 2C(400 - 363.4) + C(300 - 363.4) = C\,(73.15 - 63.42) = 9.73\,C = 29.2\,\mathrm{kJ}$.

(b) Direct contact: $T_f^{\mathrm{irr}} = (2C \cdot 400 + C \cdot 300)/(3C) = 1100/3 = 366.7\,\mathrm{K}$, larger than 363.4 K as expected. The entropy generated is $S_{\mathrm{gen}} = 2C \ln(366.7/400) + C \ln(366.7/300) = C[2(-0.0870) + 0.2007] = 0.0266\,C = 0.0799\,\mathrm{kJ\,K^{-1}}$.

(c) $T_0 S_{\mathrm{gen}} = (363.4)(0.0799) = 29.1\,\mathrm{kJ}$ taking $T_0 = T_f$, which matches $W_{\max} = 29.2\,\mathrm{kJ}$ to within the leading-order approximation. The agreement is not exact even in principle, because (39) is a leading-order expansion in S_{gen}, and here the full irreversibility (direct contact) is not small; the close numerical match reflects the modest temperature ratio, and would degrade for widely separated T_1 and T_2.

Problem 5. (a) Energy balance: $\dot{m}_c c\,(315 - 290) = \dot{m}_h c\,(350 - 320)$, so $\dot{m}_c = (1)(30)/(25) = 1.2\,\mathrm{kg\,s^{-1}}$.

(b) Each stream's entropy rate change uses $\dot{m} c \ln(T_{\mathrm{out}}/T_{\mathrm{in}})$. Hot: $(1)(4.18)\ln(320/350) = -0.3746\,\mathrm{kW\,K^{-1}}$. Cold: $(1.2)(4.18)\ln(315/290) = +0.4148\,\mathrm{kW\,K^{-1}}$. The exchanger is insulated, so no entropy crosses to the environment, and $\dot{S}_{\mathrm{gen}} = -0.3746 + 0.4148 = 0.0402\,\mathrm{kW\,K^{-1}}$.

(c) $\dot{W}_{\mathrm{lost}} = T_0 \dot{S}_{\mathrm{gen}} = (290)(0.0402) = 11.7\,\mathrm{kW}$. This charges for the irreversibility of transferring heat across the finite temperature difference between the two streams: heat falling from around 335 K to around 300 K generates entropy even though no heat is lost to the outside, and that generation is availability that could, in principle, have driven an engine.

B References

- R. C. Hendricks, I. C. Peller, and A. K. Baron, *Joule–Thomson Inversion Curves and Related Coefficients for Several Simple Fluids*, NASA Technical Note D-6807, NASA Lewis Research Center, Cleveland, Ohio (July 1972). Source of the inversion and Boyle temperatures in Table 3, the nitrogen critical constants and critical Joule–Thomson coefficient, the helium estimates, and the Linde operating-point discussion. The measured nitrogen values attributed above to Roebuck and Osterberg are quoted from this report, which cites their 1935 *Physical Review* paper; that paper has not been consulted directly here.

- L. G. Gouy, "Sur l'énergie utilisable," *Journal de Physique Théorique et Appliquée* **8**, 501–518 (1889). Origin of the lost-work relation (45) for closed systems; Stodola's extension to open systems followed around 1898–1903. Bibliographic details for both, and the caution about

which temperature enters the relation, are taken from the reference list and discussion of Y. Tu and G. Chen, "Rethinking loss of available work and Gouy–Stodola theorem," arXiv:2307.16041 (2023). The primary Gouy and Stodola sources have not been consulted directly.

- The availability and steady-flow-exergy framework in Part III is standard and appears in graduate texts on engineering thermodynamics; the presentation here follows the Callen-style development of the maximum-work theorem rather than any single source.