On relativistic temperature

David Wallace*

July 1, 2023

Abstract

I revisit the long-running controversy as to the transformation properties of temperature under Lorentz transformations, and argue that, contrary to widespread views in some parts of the literature, 'relativistic thermodynamics' should not be understood as a generalization to relativistic physics of an initially non-relativistic theory but as an application of a general thermodynamic framework, neutral as to its spacetime setting, to the specific case of relativity. More specifically, I observe that the general framework of equilibrium thermodynamics incorporates arbitrary conserved quantities in addition to energy and that when that framework is applied to systems in which momentum is conserved and can be transferred between systems, it gives rise to an unambiguous result as to the thermodynamics of moving (including relativistically moving) systems. This leads to an equally unambiguous prediction as to how thermodynamic temperature transforms under Lorentz transformations (an answer coinciding with Einstein's, and Planck's, original results.) That said, one can identify within this framework other quantities which play a temperature-like role and have different transformation properties; at some level it is a matter of convention and semantics which is 'the' temperature. I conclude with some brief remarks about the situation in general relativity.

1 Introduction

Relativistic thermodynamics is almost as old as relativity itself and yet remains surprisingly controversial. Liu's (1994) history of the subject concludes by describing the theory as 'one of the most recalcitrant in resisting the efforts of relativization'; in recent work Chua (2022) goes further in claiming that relativistic thermodynamics leads to 'a breakdown of the classical non-relativistic concept of temperature'. The issue has acquired a new urgency in the context of recent philosophical criticism of the longstanding claims of analogy between black hole behavior and thermodynamics[1].

*Department of History and Philosophy of Science / Department of Philosophy, University of Pittsburgh; david.wallace@pitt.edu

[1]Notably by Chua (ibid.) in passing, and by Dougherty and Callender (2016) — though see also the response to the latter in Wallace (2018).

It is at first site surprising that any such controversies are compatible with the state of modern thermal physics. There is nothing obviously non-relativistic about modern thermodynamics, or the statistical mechanics that underpins it: to the contrary, it is absolutely routine to apply thermodynamics to systems — the plasma in a fusion reactor, the interior of the Sun, the shock front of a supernova, Big Bang nucleosynthesis — which are not even faintly 'non-relativistic'. At an elementary level, thousands of physics undergraduates calculate the thermodynamic properties of black-body radiation every year without any suggestion that doing so involves a relativistic ingredient not present in similar calculations for the ideal gas; at a more advanced level, one will search the index of monographs on finite-temperature quantum field theory or relativistic astrophysics in vain for any such suggestion.[2]

Things become clearer upon noting the main locus of the controversy: the transformation properties of temperature and other such thermodynamic quantities under Lorentz transformations, something which indeed plays little role[3] in the application of thermodynamics to relativistic systems (where thermal calculations are almost always carried out in a local rest frame). Here one sees the possibility of a controversy that could persist without troubling the physics mainstream. And here one persistently finds the idea that something essentially new is required that cannot simply be found deductively within extant thermodynamics, but which forms part of a generalization or extension of that extant theory. Perhaps one such extension is clearly to be preferred; perhaps it is a matter of convention which is to be used; perhaps the existence of incompatible extensions tells us that thermodynamics itself falls apart in this regime; but in any case, without such an extension thermodynamics is silent on the transformation properties of its quantities.

The main purpose of this paper is to argue that no extension of thermodynamics is required. Rather, extant thermodynamics *already* gives unequivocal answers to these questions, without any need for extension, since:

1. The equation of state for equilibrium thermodynamics expresses the entropy of a system as a function of the energy of that system *and of any other conserved quantities* (and external parameters) for the system.

2. Applied to the specific case of linear momentum as a conserved quantity, thermodynamics unequivocally provides an equation of state for a moving system.

3. In the limit of arbitrarily slow accelerations, the action of transferring linear momentum to or from a system, or equivalently of boosting that system to a different velocity, is reversible, i. e. entropy-conserving.

4. The equation of state, combined with elementary relativistic kinematics, thus uniquely fixes how the thermodynamic temperature of a system

[2] I searched the indexes of Battaner (1996), Kapusta and Gale (2006), and Padmanabhan (2000).

[3] Not no role: observational astrophysics often requires us to consider what a system at thermal equilibrium in one frame looks like in another, a point to which I return in section 7.

(along with its other thermodynamic parameters) changes when that system is reversibly set into motion.

5. Knowing how a system's parameters change when a system is put into uniform motion relative to an inertial frame, in the limit when that is done without disturbing the system's internal degrees of freedom, is equivalent to knowing how they change when the body is left alone but is described from an inertial frame uniformly moving with respect to the first frame.

That said, there is an important aspect of truth in the claims of conventionality. For relativistic systems there are other quantities than thermodynamic temperature that usefully characterize the system — chief among them the 'rest temperature', the temperature that the system has in its rest frame. In that sense there is a conventionality as to what we call 'temperature' for a moving body, closely related to the conventionality as to whether 'mass', for a moving body, means its rest mass or its relativistic (or inertial) mass. But this conventionality does not change the fact that the quantity formally defined as temperature in thermodynamics has unambiguous transformation properties under boosts; more generally, it does not in any way undermine the overall coherence and consistency of thermodynamics, even when applied to relativistically moving bodies.

I should acknowledge right away that little here is really new — certainly nothing calculational is. The literature on relativistic thermodynamics is large and tangled and the boundaries between exegetical novelty, rediscovery and originality are blurry at best. Indeed, from one perspective there is little here not already present in Einstein and Planck's work in the 1900s (Einstein 1907; Planck 1906, 1907), or in Tolman's (1933, 1934) development of that work (for details see Liu, *ibid*). But the very fact that controversy continues a century later shows that perhaps there is still something to be gained from retreading this ground from a more contemporary perspective. (Similarly, for exegetical reasons I think picking through the detailed historical development of the subject risks obscuring more than it reveals and so I have not attempted to do so; in any case I have little to add to the lovely account given by Liu.)

The paper is structured as follows: in section 2 I review the general framework of equilibrium thermodynamics, including the role of conserved quantities and the formal definition of thermodynamic temperature; in section 3 I apply that general framework to linear momentum, and so derive the transformation properties for thermodynamic temperature. In sections 4–5 I apply these results, firstly to the general question of how to construct Carnot cycles between relatively-moving systems and secondly to the apparent paradox that when two bodies of the same rest temperature are moving, each has in its own rest frame a higher temperature than the other. In section 6 I consider an alternative definition of the Carnot cycle, which gives rise to an alternative definition of temperature; in section 7 I consider the relation between the various notions of temperature. In section 8 I briefly consider some related issues in *general* relativity (both for systems in a stationary background spacetime, and for self-gravitating systems in an asymptotically flat spacetime); this section lies a little way outside the main development of ideas in the paper. Section 9 is the conclusion.

Notation: I set both Boltzmann's constant k_B and the speed of light c to 1; I take the signature of the special-relativistic metric to be $(-,+,+,+)$. **Boldface symbols** are 3-vectors; Greek indices range from 0 to 3 and are raised and lowered using the special-relativistic metric; the Einstein summation convention is used. I do not attempt to provide original references for textbook physics (e.g., equilibrium thermodynamics itself).

2 Review of equilibrium thermodynamics

The foundation of equilibrium thermodynamics (called the 'minus first law' by Brown and Uffink (2001)) is that isolated systems evolve towards unique equilibrium states.

But what does 'unique' mean here? If 'isolated' means that energy does not flow into or out of a system during equilibration, then of course different-energy systems will obtain different equilibrium states. But if in addition there are other conserved quantities than energy, and if 'isolated' means that these too cannot be exchanged with the environment, then equilibrium states will be individuated by the values of those other conserved quantities as well as by energy. And if the system's dynamics depends on some externally controlled variable — like the volume, for instance — and if that variable is held fixed during equilibration, then different values of that variable lead to different equilibria.[4]

So a more precise statement is that thermodynamic systems are characterized by some number of conserved quantities, including but not not necessarily limited to energy, and by some (possibly zero) number of external parameters, and that for any given value of the conserved quantities and the external parameters, there is a unique[5] equilibrium state to which an arbitrary state will evolve.

I should pause here to note that this is a somewhat more general conception of thermodynamics than one sometimes sees in the foundational literature, which normally confines itself to systems where the only parameter is volume and the only conserved quantity is energy. But this covers only a small fraction of thermodynamics as it is used in contemporary physics and chemistry: in chemical thermodynamics there are conservation laws tracking the separate conservation of each element; in nuclear chemistry element number is not conserved but quantities like charge and baryon number are; in the thermodynamics of magnetic matter volume is normally fixed and the role of external parameter is played by a magnetic field. If 'thermodynamics' is to be understood broadly

[4]This latter point is recognized in Brown and Uffink's precise statement of the minus first law: "An isolated system in an arbitrary initial state *within a finite fixed volume* will spontaneously attain a unique state of equilibrium" (my emphasis). But volume is not the only possible external parameter, and even for fixed volume there may be conserved quantities other than energy.

[5]Even this, arguably, is insufficiently general, given spontaneous symmetry breaking, but these subtleties are beyond the scope of the present paper.

enough to cover these various applications, it must be understood more abstractly, applicable to whatever conserved quantities and external parameters define the thermodynamic system at hand.

Let us denote any conserved quantities (other than the energy U) by $N_1, \ldots N_K$, and any external parameters by $V_1, \ldots V_M$. The second law of thermodynamics then amounts to the statement that there is a function

$$S(U, N_1, \ldots, N_K, V_1, \ldots, V_M)$$

— the *thermodynamic entropy* — of the conserved quantities and external parameters such that (1) if an isolated system initially at equilibrium is allowed to evolve under externally-induced time dependence of its external parameters and then to return to equilibrium, the value of S will not have decreased; and (2) if two or more systems initially at equilibrium are dynamically coupled so as to be able to exchange energy and other additive conserved quantities, and then the coupling is removed and they are allowed to come to equilibrium, then again the total value of S will not have decreased. (It is common in foundational work to present the second law in more directly operational terms, but in practical applications what matters is the entropy form I state here.)

Assuming (as is standardly done, and as can normally be justified in statistical mechanics) that entropy is a smooth function of the constants and parameters, we can write

$$\mathrm{d}S = \beta \mathrm{d}U + \sum_{i=1}^{K} \theta_i \mathrm{d}N_i + \sum_{i=1}^{M} \alpha_i \mathrm{d}V_i \tag{1}$$

where β, θ_i and α_i are all functions of U, the N_i, and the V_i, given explicitly by

$$\beta = \left(\frac{\partial S}{\partial U}\right)_{N_j, V_j} \quad \theta_i = \left(\frac{\partial S}{\partial N_i}\right)_{U, N_j (i \neq j), V_j} \quad \alpha_i = \left(\frac{\partial S}{\partial V_i}\right)_{U, N_j, V_j (i \neq j)}. \tag{2}$$

As stated this is an entirely formal statement about the space of equilibrium states, but it has an operational interpretation if we take $\mathrm{d}U$, $\mathrm{d}N_i$, and $\mathrm{d}V_i$ to be small but finite changes to the constants and parameters caused by some external intervention (with the system otherwise being kept isolated). In the limit of small changes, $\mathrm{d}S$ becomes the entropy change caused by that intervention once the system returns to equilibrium, and the second law becomes the requirement that that change is nonnegative so long as the system is isolated from its environment.

The *adiabatic assumption*, not always stated explicitly in thermodynamics texts but essential in application, is that if we adjust a system's constants and parameters sufficiently slowly, we can induce it to move through the space of equilibrium states in such a way that entropy increase is arbitrarily small; in other words, that infinitesimal processes are available that arbitrarily approach $\mathrm{d}S = 0$. The limit of such processes are the *reversible processes*; as Norton (2016) has pointed out, they do not strictly represent physically attainable processes but they can be approximated arbitrarily well by such processes, so that it is convenient to proceed under the fiction that they exist rather than carry around epsilons and deltas.

Partly for historical reasons, it is standard to rewrite (1) as

$$\begin{aligned} \mathrm{d}U &= \frac{1}{\beta}\mathrm{d}S + \sum_i \left(-\frac{\theta_i}{\beta}\right)\mathrm{d}N_i - \sum_i \left(\frac{\alpha_i}{\beta}\right)\mathrm{d}V_i \\ &\equiv T\mathrm{d}S + \sum_i \mu_i \mathrm{d}N_i - \sum_i p_i \mathrm{d}V_i. \end{aligned} \quad (3)$$

This expression is sometimes called the *First Law*, and I follow this convention here (without prejudice as to what connection it bears to the historical First Law). $T = 1/\beta$ is the *thermodynamic temperature*; $\mu_i = -\theta_i/\beta$ is a *generalized potential*; $p_i = \alpha_i/\beta$ is a *generalized pressure*.

The canonical example of an external parameter is volume, and the canonical example of a conserved quantity other than energy is particle number; with these choices we have

$$\mathrm{d}U = T\mathrm{d}S + \mu\mathrm{d}N - p\mathrm{d}V \quad (4)$$

and can identify p as the ordinary (mechanical) pressure (and label μ the chemical potential). But the framework is far more general than this and may include more or fewer parameters (the thermodynamics of hot rocks, for instance, normally involves no parameters and no conserved quantities other than energy). To some extent it is up to us which parameters and conserved quantities to include: if we have no interest in any processes which change particle number, for instance, we might drop the $\mu\mathrm{d}N$ term and just treat N as a fixed parameter.

Now suppose that we have two systems, labeled 1 and 2, at temperatures T_1, T_2 with $T_1 > T_2$, and that we perform some 'heat exchange' process that decreases the entropy of the first system by δS_1 and increases that of the second by $+\delta S_2$, while leaving all parameters, and all constants other than energy, unchanged for each system. The two systems can be treated as subsystems of a single larger system, with total entropy equal to the sum of the subsystem entropies, and the *work extracted* $đW$ from the combined system in this process is the net decrease in energy,

$$đW = T_1\delta S_1 - T_2\delta S_2 \quad (5)$$

and the *efficiency* e of the process is the ratio of the work to the energy decrease of the first system,

$$e = \frac{đW}{T_1\delta S_1} = 1 - \frac{T_2}{T_1}\frac{\delta S_2}{\delta S_1}. \quad (6)$$

Since the total entropy increase is nonnegative, $\delta S_2 \geq \delta S_1$, this entails

$$e \leq 1 - \frac{T_2}{T_1} \quad (7)$$

with equality reached in the (idealized) limit where the process is reversible. Reversible processes of this kind are called *Carnot cycles*. Note that this derivation actually continues to make sense if $T_1 < T_2$ but that in this case the efficiency is less than 0, corresponding to the fact that work must be done to induce this transfer.

Carnot cycles are normally realized by some third ('transfer') system that begins at temperature T_1, and go as follows:

1. Entropy δS is reversibly transferred from system 1 to the transfer system, via some change of the transfer system's parameters.

2. The parameter(s) of the transfer system are reversibly adjusted so as to lower the temperature of the transfer system to T_2.

3. Entropy δS is reversibly transferred from the transfer system to system 2, again by some change in the system's parameters.

4. The parameters of the transfer system are reversibly adjusted so as to return it to temperature T_1.

The best-known realization of the Carnot cycle involves using a small box of dilute gas as a transfer system, but this is inessential. In a relativistic context, for instance, it is natural to use a box of black-body radiation, which has no conserved quantities save energy and has volume as its sole external parameter. A empty box initially of negligible volume, placed in thermal contact with another system at temperature T_1 and adiabatically expanded, will reversibly (in the infinite-time limit) fill with black-body radiation at temperature T_1; the same box, adiabatically expanded while thermally isolated, will decrease in temperature until it reaches the desired temperature T_2; the same box, now placed in thermal contact with another system at temperature T_2 and adiabatically contracted to negligible size, will reversibly expel all its energy into that system. In any case, the actual implementation of the Carnot cycle has no significance to the core efficiency results, which rely only on the Second Law and on the existence (in the limit) of reversible entropy-transfer processes which saturate the inequality of the second law.

All this should be familiar, but I want to stress that it is vital to the whole mechanism of Carnot cycles that the informal requirement that the process operates 'in a cycle' is operationally implemented by requiring no changes in other conserved quantities than energy, or in any of the external parameters.

3 Thermodynamics of moving systems

Let us now consider how the above story works if our thermodynamic systems are moving and/or rotating. For simplicity we will assume no external parameters, and consider only conserved quantities, of which there are seven: energy, the three components P_1, P_2, P_3 of linear momentum, and the three components J_1, J_2, J_3 of angular momentum.

We are not used to including momentum (linear or angular) in thermodynamics, and the reason is simple: the theory is most commonly applied to systems confined in some kind of external box, and for those systems momentum (linear or angular) is *not* conserved, but is transferred to the walls of the box in collisional processes. But if the box is moving or spinning, we had better not be so complacent: we will assume that the box itself is part of the system, and include both its energy and its linear and angular momentum in our description.

(A complication immediately arises. Equilibrium is traditionally described as the state a system reaches when all of its macroscopic degrees of freedom are unchanging. But of course a moving body is, well, moving, and precisely because momentum is conserved, that movement does not cease at equilibrium; similarly, in general a rotating body will tumble in space, and that tumbling will not cease as long as angular momentum is conserved.

Nonetheless there clearly is a physically relevant sense of equilibration here: in a cylinder of gas in empty space, tumble and fly though it might, the contents will still reach an appropriately steady state. We can characterize that sense more precisely: suppose (assuming for definiteness classical Lagrangian mechanics) that the system has coordinates $q^1, \ldots q^N$ but that the Lagrangian does not depend on q^1, so that translation in q^1 is a symmetry. (At least locally, any configuration symmetry can be so expressed.) Then the conjugate momentum

$$p_1 = \frac{\partial L}{\partial q^1} \tag{8}$$

is conserved, and there is a self-contained dynamics for the remaining coordinates $q^2, \ldots q^N$, in which p_1 can be treated as a time-independent parameter in the expression for the Hamiltonian. Equilibrium can now be understood with respect to these coordinates. For instance, for a translation-invariant N-particle system we can take the $N-3$ translationally invariant degrees of freedom to collectively reach equilibrium.)

For the most part in this paper, I set aside *angular* momentum and consider only changes of linear momentum. Under this simplification, the First Law becomes

$$\mathrm{d}U = T\mathrm{d}S + \mathbf{v} \cdot \mathrm{d}\mathbf{P} \tag{9}$$

where $\mathbf{v} = (v_1, v_2, v_3)$ is a triple of thermodynamic potentials, respectively conjugate to the momenta P_1, P_2, P_3. (If this notation seems misleading, and at risk of conflating the thermodynamic potential with the velocity... read on.)

The key observation is now that we can act *reversibly* on the system to accelerate or decelerate it. Of course a sufficiently aggressive impulse applied to the system will shake up its internal degrees of freedom, but this will decrease as the gentleness of the impulse increases, and in the limit we can accelerate the system infinitesimally slowly — and hence, thermodynamically reversibly. Under such a transformation the change of energy and momentum will obey

$$\mathrm{d}U = \mathbf{v} \cdot \mathrm{d}\mathbf{P}. \tag{10}$$

But relativistic kinematics also tells us that in this limit, the rest mass of the system will be unchanged. Since we have $U^2 - \mathbf{P} \cdot \mathbf{P} = M^2$, we can differentiate and obtain

$$U\mathrm{d}U = M\mathrm{d}M + \mathbf{P}\cdot\mathrm{d}\mathbf{P} \tag{11}$$

so that in the reversible limit,

$$\mathrm{d}U = \frac{1}{U}\mathbf{P}\cdot\mathrm{d}\mathbf{P}. \tag{12}$$

Comparing this to (10) we can read off

$$\mathbf{v} = \frac{\mathbf{P}}{U}. \tag{13}$$

But of course the right hand side is just the (average) velocity, so in fact the thermodynamic potential for momentum and the velocity can be identified.

Let us now ask how the thermodynamic temperature T changes under a reversible acceleration. We can do so by notion that since asymptotically-slowly-applied accelerations are reversible transformations, entropy is invariant under those transformations — which is to say that entropy transforms as a scalar under Lorentz boosts. Noting that the 4-momentum $\mathcal{P}_\mu = (-U, \mathbf{P})$ transforms as a 1-form, we can read off from the expression

$$\mathrm{d}S = \beta\mathrm{d}U - \beta\mathbf{v}\cdot\mathbf{P} \tag{14}$$

(where as usual $\beta = 1/T$) that the object $(\beta, \beta\mathbf{v})$ transforms as a vector, which we might call the *inverse 4-temperature*. ('4-inverse-temperature' is too ugly, even if more logical.) We can then write the First Law in the manifestly covariant form

$$\mathrm{d}S = -\beta^\mu\mathrm{d}\mathcal{P}_\mu \tag{15}$$

where β^μ is the inverse 4-temperature and $\mathcal{P}_\mu$ is the 4-momentum. (I stress again: this is not a proposed relativistic generalization of the First Law: it is the First Law itself, in this particular physical context.) The physical meaning of inverse 4-temperature is then (minus) the rate of change of entropy with 4-momentum.

We can rewrite the inverse 4-temperature as

$$\beta^\mu = (\beta, \beta\mathbf{v}) = (\beta/\gamma(\mathbf{v}))(\beta\gamma(\mathbf{v}), \beta\mathbf{v}\gamma(\mathbf{v})) \equiv \beta_R v^\mu \tag{16}$$

where $v^\mu = (\gamma(\mathbf{v}), \mathbf{v}\gamma(\mathbf{v}))$ is the 4-velocity and $\beta_R = \beta/\gamma(\mathbf{v})$ is the *inverse rest temperature*. The inverse rest temperature is a scalar and can be defined directly in terms of the inverse 4-temperature via

$$\beta_R^2 = -\beta^\mu\beta_\mu. \tag{17}$$

It has a direct physical interpretation: we have

$$\beta_R\mathrm{d}M = \frac{\beta_R}{2M}\mathrm{d}(-\mathcal{P}^\mu\mathcal{P}_\mu) = -\frac{\beta_R\mathcal{P}^\mu}{M}\mathrm{d}\mathcal{P}_\mu = -\beta^\mu\mathrm{d}\mathcal{P}_\mu = \mathrm{d}S \tag{18}$$

so that inverse rest temperature is the rate of change of entropy with rest mass — or, equivalently, rest temperature $T_R = 1/\beta_R$ is the rate of change of rest mass with entropy. (Note that both rest mass and entropy are scalars, so that this expression is relativistically invariant.) Of course, we can obtain the same result by transforming to the system's rest frame, in which rest mass and energy coincide, as do rest temperature and temperature.

4 Carnot cycles for moving systems: energy transfer at constant momentum

From the transformation rule for inverse 4-temperature, we can immediately say how thermodynamic temperature transforms under Lorentz transformations: it obeys

$$T = T_R/\gamma \tag{19}$$

i. e. reversibly putting a stationary system into motion reduces its thermodynamic temperature. But it is not physically obvious why this should be. To get some insight into this expression — and, more generally, into the form of the Second Law which we have derived — let's consider how a Carnot cycle might explicitly be realized for this system.[6] For simplicity, we will use the black-body-radiation implementation, and will consider running the Carnot cycle between a moving system with velocity v and a stationary system. (Running the cycle between two moving systems can always be implemented using a stationary intermediate, or else just by Lorentz-transforming into a frame where one system is stationary.)

Suppose that the moving system has rest temperature T_R and rest mass M, and the stationary system has temperature T_0. Our Carnot cycle works as follows:

1. Start with a stationary, empty box of radiation with rest mass m. Reversibly boost it so that it is comoving with the moving system.

2. Put the box in thermal contact with the moving system and slowly expand it so that a small amount of energy (at rest temperature T_R) is transferred to the box. This transfer inevitably also transfers some *momentum* to the box.

3. Reversibly accelerate the moving system so that its momentum is back to its original value; reversibly decelerate the box so it is again at rest. At this point, total momentum change (for box or system) is zero, as we require. Since the reversible transformation doesn't change its rest-frame temperature, the box is now at temperature T_R.

4. Reversibly expand or contract the box so that its temperature matches the stationary system.

[6] This notion of a Carnot cycle is explicitly considered by Einstein (1907) and Planck (1906, 1907); see Liu (1994) for the historical details.

5. Put the box in thermal contact with the stationary system and slowly contract it until it is empty.

Evaluating this quantitatively (and working in one spatial dimension for simplicity), we find that at step 1 the 4-momenta of the moving system and the box (written as 4-vectors) are, respectively,

$$
\begin{aligned}
E^{\mu}_{\text{moving}}(1) &= (M\gamma(v), Mv\gamma(v)) \\
E^{\mu}_{\text{box}}(1) &= (m\gamma(v), mv\gamma(v)).
\end{aligned} \tag{20}
$$

At step 2 they transform to

$$
\begin{aligned}
E^{\mu}_{\text{moving}}(2) &= (M\gamma(v) - \delta m\gamma(v), Mv\gamma(v) - \delta mv\gamma(v)) \\
E^{\mu}_{\text{box}}(2) &= (m\gamma(v) + \delta m\gamma(v), mv\gamma(v) + \delta mv\gamma(v)).
\end{aligned} \tag{21}
$$

A reversible transfer of momentum back to the moving system satisfies $\delta E = v\delta P$, so at step 3 we have

$$
E^{\mu}_{\text{moving}}(3) = (M\gamma(v) - \delta m\gamma(v) + \delta mv^2\gamma v, Mv\gamma(v)) \tag{22}
$$

$$
\begin{aligned}
&= (M\gamma(v) - \delta m/\gamma(v), Mv\gamma(v)) \\
E^{\mu}_{\text{box}}(3) &= (m + \delta m, 0).
\end{aligned} \tag{23}
$$

Note that the decrease in energy of the system in this process is

$$
\delta_1 U = \delta m/\gamma(v) \tag{24}
$$

which is less than the energy δm added to the box; this means that the reversible process of transferring momentum back to the system and decelerating the box to rest took net work.

By ordinary thermodynamics, the amount of energy transferred to the stationary system is

$$
\delta_2 U = (T_0/T_R)\delta m \tag{25}
$$

so that the thermodynamic efficiency is

$$
e = 1 - \frac{\delta_1 U}{\delta_2 U} = 1 - \frac{T_R}{\gamma T_0} \tag{26}
$$

and hence the thermodynamic temperature of the moving system is indeed T_R/γ. The decrease in temperature reflects the fact that transferring energy from a moving system without also transferring momentum requires net work.

This Carnot cycle is a direct implementation for our system of the general formula for Carnot cycles in systems with multiple conserved quantities, but it is admittedly difficult to see its physical significance, and there is a clear reason why: a Carnot cycle is a process for transferring *energy* but no other conserved quantity (in this case momentum), and this is not a relativistically invariant distinction: indeed, whether a process *is* a Carnot cycle depends on one's reference frame, since a process that transfers only energy is Lorentz-transformed to a process that transfers both energy and momentum. We can get a more relativistically natural description of the thermodynamics of transfer processes if we consider the more general problem of extracting 4-momentum from the total system via a transfer of 4-momentum from one system to another: if we decrease one system's 4-momentum by $\delta\mathcal{P}^\mu$, transfer some of that 4-momentum to another system, and extract usable 4-momentum (stored in, e.g., the 4-momentum of a single heavy body) $đW^\mu$, then the total entropy change is

$$0 \leq \delta S = \delta_1 S + \delta_2 S = +\beta_1^\mu \delta\mathcal{P}_\mu - \beta_2^\mu (\delta\mathcal{P}_\mu - đW^\mu) \tag{27}$$

or

$$\beta_2^\mu (-đW_\mu) \leq (\beta_2^\mu - \beta_1^\mu)(-\delta\mathcal{P}_\mu). \tag{28}$$

(The minus signs are a consequence of the fact that we are in $(-,+,+,+)$ signature: recall that $\delta\mathcal{P}_\mu = (-\delta U, \delta\mathbf{P})$.)

This provides a bound on the ability of processes (natural or artificial) to extract usable 4-momentum from comoving systems (say, different parts of the accretion disk of a compact object). It might be thought of as a relativistic generalization of the Carnot cycle (and of irreversible heat engines) — so long as it is remembered that the original Carnot cycle remains unambiguously defined, albeit it depends on one's reference-frame.

5 Relativity of temperature

As an application of these results, consider the following apparent paradox, akin to the classic paradoxes of special relativity: two systems A, B with identical rest temperature T are in relative motion, with B having velocity $+v$ relative to A along the x axis. In A's rest frame, B has temperature $T/\gamma(v)$ and so it is entropically favorable for heat to flow from A to B. But in B's rest frame, A likewise has temperature $T/\gamma(v)$ and it is thermodynamically favorable for heat to flow from B to A. At first sight, this is a contradiction.

The resolution of the paradox comes from noting that the statement 'energy is transferred but momentum is not' is not Lorentz-invariant. (The corresponding 4-momentum flow is $(-U, 0)$; obviously, upon Lorentz-transforming this it picks up a nonzero momentum component.) Suppose we begin by working in A's rest frame, and consider a transfer of some small amount of energy ΔU from A to B, without any additional transfer of momentum, so that the 4-momentum transferred $\Delta\mathcal{P}$ (in A's rest frame) has components

$$\Delta\mathcal{P}_\mu = (-\Delta U, 0, 0, 0). \tag{29}$$

Ordinary thermodynamics does indeed tell us that this transfer increases total entropy by

$$\Delta S = \Delta U \left(\frac{\gamma(v)}{T} - \frac{1}{T} \right) = \frac{\Delta U}{T} (\gamma(v) - 1) > 0 \tag{30}$$

since the temperature of B is $T/\gamma(v)$.

Now suppose we instead transfer a quantity of 4-momentum $\mathcal{P}'$ that, in B's rest frame, is purely a transfer of energy ΔU. In A's rest frame this transfer also involves a transfer of momentum: we have

$$\mathcal{P}'_\mu = (-\gamma(v)\Delta U, v\gamma(v)\Delta U). \tag{31}$$

The inverse 4-temperatures of A and B in A's rest frame are, respectively,

$$\beta_A = (1/T, 0, 0, 0) \tag{32}$$

and

$$\beta_B = (\gamma(v)/T, v\gamma(v)/T, 0, 0) \tag{33}$$

and so the change of entropy (continuing to calculate in A's rest frame) is

$$\begin{aligned} \Delta S &= \gamma(v)\Delta U \left(\frac{1}{T} - \frac{\gamma(v)}{T}\right) - v\gamma(v)\Delta U \left(0 - \frac{v\gamma(v)}{T}\right) \\ &= \frac{\gamma(v)\Delta U}{T}(1 - (1 - v^2)\gamma(v)) \\ &= \frac{\Delta U}{T}(\gamma(v) - 1). \end{aligned} \tag{34}$$

In other words, the entropy increase is (a) positive and (b) numerically equal to the previously-calculated entropy increase for a transfer of energy from A to B in A's rest frame. Contrary to intuition, it is actually true that given relatively-moving bodies at the same rest temperature, in each body's rest frame it is entropically favorable for heat to the other body — provided that no momentum is also transferred as part of that flow.

6 Constant-velocity cyclic processes

The relativistic Carnot cycle described above is the flat-footed specification of the Carnot cycle in a context in which momentum is conserved: it describes, as per standard equilibrium thermodynamics, the reversible limit of a process that transfers energy between systems while leaving other conserved quantities constant. There are, however, other reasonable cyclic processes that we could consider. For instance, we might be interested in energy transfers between systems at constant *velocity*. Physically this might describe, e.g., a transfer of usable energy between two spacecraft at different velocities.[7]

We can work out the thermodynamics of this straightforwardly: we have

$$\mathrm{d}U = \mathrm{d}(\gamma(v)M) = \gamma(v)\mathrm{d}M + M\mathrm{d}\gamma(v). \tag{35}$$

Using our existing result $\mathrm{d}M = T_R\mathrm{d}S$, we have

$$\mathrm{d}U = (\gamma(v)T_R)\mathrm{d}S + \frac{U}{\gamma(v)}\mathrm{d}\gamma(v) \tag{36}$$

[7] This notion of a cycle is considered by Ott (1963) and (in correspondence) by the later Einstein; again see Liu (1994) for the historical details.

so that, for energy exchange between two systems at fixed velocities (and hence fixed gamma factors) the maximum efficiency is controlled by an *effective* temperature, $T_v = \gamma T_R$. (Call this the *constant-velocity temperature.*) Unlike the case of constant-momentum energy transfer, constant-velocity energy transfer gets more efficient when a system is moving faster.

We can get further insight into the transformation properties of the constant-velocity temperature by generalizing (35) to

$$\mathrm{d}\mathcal{P}^\mu = \mathrm{d}(v^\mu M) = v^\mu \mathrm{d}M + M\mathrm{d}v^\mu = (v^\mu T_R)\mathrm{d}S + M\mathrm{d}v^\mu \equiv T_v^\mu \mathrm{d}S + M\mathrm{d}v_\mu \quad (37)$$

where now T_v^μ is the *constant-velocity 4-temperature*, defined as the rate of change of 4-momentum with entropy at constant 4-velocity. Manifestly, this transforms as a 4-vector.

It is again helpful to explicitly realize the cycle to transfer energy reversibly between two bodies while keeping velocity constant. Again we work in one spatial dimension, assume a moving system of velocity v, rest temperature T_R and rest mass M, and a stationary system of temperature T_0, and use a box of black-body radiation as our transfer system. A reversible cycle is then:

1. Start with a stationary, empty box of radiation with rest mass m. Reversibly boost it so that it is comoving with the moving system.

2. Put the box in thermal contact with the moving system and slowly expand it so that a small amount of energy (at rest temperature T_R) is transferred to the box. This transfer inevitably also transfers some momentum to the box — but we don't care, since we are not attempting to conserve momentum.

3. Reversibly slow the box until it is at rest (leaving the moving system alone). Since the reversible transformation doesn't change its rest-frame temperature, the box is now at temperature T_R.

4. Reversibly expand or contract the box so that its temperature matches the stationary system.

5. Put the box in thermal contact with the stationary system and slowly contract it until it is empty.

Suppose this process extracts rest mass δm, and hence energy $\delta_1 U = \gamma(v)\delta m$, from the moving system. Once it has come to rest, the box has energy δm: note that this is less than the energy extracted from the moving system, reflecting the fact (known to any driver of a hybrid car) that we can extract work from a moving system by reversibly slowing it down.

By ordinary thermodynamics (as with the constant-momentum cycle), the amount of energy transferred to the stationary system is

$$\delta_2 U = (T_0/T_R)\delta m \quad (38)$$

so that the thermodynamic efficiency is

$$e = 1 - \frac{\delta_1 U}{\delta_2 U} = 1 - \frac{\gamma T_R}{T_0} \tag{39}$$

reproducing our previous result about the effective temperature. The fact that we can extract more energy from a moving system in a constant-velocity cycle simply reflects the fact that usable work can be extracted by slowing the box down.

We can obtain the transformation properties of the effective temperature yet a third way, via section 4's notion of a relativistic heat engine: recall that such engines extract 4-momentum $đW_\mu$ from a 4-momentum flow $\delta\mathcal{P}_\mu$ out of system 1, with the residual 4-momentum flowing into system 2, and that their efficiency is bounded by

$$-\beta_2^\mu đW_\mu \leq (\beta_2^\mu - \beta_1^\mu)(-\delta\mathcal{P}_\mu). \tag{40}$$

Writing the inverse 4-temperature β_i^μ of system i in terms of its rest temperature T_i and 4-velocity v_i^μ, $\beta_i^\mu = v_i^\mu / T_i$, we can rearrange this to

$$-v_2^\mu đW_\mu \leq \left(v_2^\mu - \frac{T_{R2}}{T_{R1}} v_1^\mu \right) (-\delta\mathcal{P}_\mu). \tag{41}$$

A removal of energy from system 1 at constant velocity is characterized by $\delta\mathcal{P}_\mu = v_{1\mu}\delta m$; substituting this in, we obtain

$$-v_2^\mu đW_\mu \leq \left(-v_2^\mu v_{1\mu} - \frac{T_{R2}}{T_{R1}} \right) \delta m. \tag{42}$$

The left hand side is just the energy extracted in the rest frame of the second system, and we have $-v_2^\mu v_{1\mu} = \gamma$, where γ is the gamma-factor of the first system in that rest frame (this is immediate if we perform the calculation in that frame). Since $\gamma\delta m$ is also the energy decrease of the first system in the second system's rest frame, we get

$$-v_2^\mu đW_\mu \leq \left(1 - \frac{T_{R2}}{\gamma T_{R1}} \right) (\gamma\delta m) \tag{43}$$

— in other words, the efficiency of the process is determined by the effective temperature γT_{R1} of the moving system.

These two notions of cyclical processes have been extensively discussed in the literature on relativistic thermodynamics, where the constant-momentum and constant-velocity processes are called the *Einstein-Planck* and *Einstein-Ott* processes, after their respective proponents. To the (imperfect) degree that this literature has reached a consensus, it is that both processes are legitimate extensions of the heat-engine concept to the relativistic regime and that there can be no fact of the matter as to which is correct. (Authors differ on whether this is harmless or problematic: to (Chua 2022), for instance, it reflects that absence of any unequivocal extension of thermodynamics to relativistic physics.)

This is then a good place to reiterate my main point: we do not need to extend thermodynamics to relativity, since its framework is already sufficiently general to incorporate relativity, and indeed makes no assumptions about the spacetime symmetry group of the theory to which it is applied beyond the existence of a time translation symmetry. And in thermodynamics, heat engines are devices which transfer energy but no other conserved quantities, and so it is the constant-momentum process that counts as a heat engine. The constant-velocity process is unambiguously *not* a heat engine in the literal sense of the term (and its reversible limit is unambiguously not a Carnot cycle), simply because it is not a cyclic process in which all conserved quantities except entropy, and all external parameters, are unchanged for each system. This is simply a matter of definition; it is independent of the question of whether this cycle is more physically useful or salient. (Similarly, there are thermodynamic contexts where processes that exchange energy reversibly at constant chemical potential and/or constant pressure are more salient than constant-particle-number and constant-volume processes; for all that, it is the latter and not the former that define Carnot cycles.)

7 So which is the real temperature?

Let's review. We have found no fewer than three quantities that might seem to deserve the name 'temperature'.

1. The *rest temperature* T_R is the rate of change of rest mass with entropy.
2. The *thermodynamic temperature* T is the rate of change of energy with entropy at constant momentum. It is the inverse of the 0-component of the inverse 4-temperature, which is the rate of change of entropy with 4-momentum.
3. The *constant-velocity* temperature T^v is te rate of change of energy with entropy at constant velocity. It is the 0-component of the constant-velocity 4-temperature, which is the rate of change of 4-momentum with entropy at constant velocity.

All three quantities have the same value at zero velocity. The rest temperature is a scalar, and so velocity-invariant; the inverse 4-temperature and the velocity 4-temperature both transform as 4-vectors, so that the thermodynamic temperature is given by $T = T_R/\gamma(\mathbf{v})$ and the constant-velocity temperature is given by $T^v = T_R\gamma(\mathbf{v})$.

But this divergence is a shallow matter of semantics. All three 'temperatures' have been defined from the same starting point: the equilibrium thermodynamics of a system in which energy and momentum are both conserved. If 'temperature' is to have its normal thermodynamic meaning — the rate of change of energy with entropy as all conserved quantities and external parameters are held fixed — then thermodynamic temperature is temperature. On the other hand, the *interesting* part of the thermodynamics entirely concerns the way entropy varies with rest mass, and with that known the full thermodynamics is just a matter of Lorentz transformation. So one could argue that it is the rest temperature that is the genuinely interesting quantity, and we should reserve 'temperature' for that quantity. (This is certainly the practice in astrophysics; likewise in black hole thermodynamics.) But nothing factive hangs on this terminological choice. Meanwhile, the 'constant-velocity temperature' does not really correspond to anything we would normally call temperature in thermodynamics — yet for all that it has a perfectly well-defined operational significance. (Similarly, we could define a 'constant-pressure temperature', if we wished, to parametrize the efficiency of Carnot cycles between systems carried out at fixed pressure rather than fixed volume.)

The whole situation is closely analogous to the status of energy in special relativity. Energy is the generator of time translations; this definition requires only time-translation invariance and does not otherwise constrain the spacetime symmetry group, and it applies just as well in special-relativistic as in non-relativistic physics. Since the notion of time translation is reference-frame dependent (in Galilean-covariant physics, let alone in relativity) it follows that so is energy, and we can determine that the energy of a system in one frame depends on both its energy and its momentum in another frame. We gain deeper insight into relativity once we realize that these transformations can be interpreted as the transformations of a 4-vector, the 4-momentum, of which energy and momentum are coordinates, and that the invariant length of this four-momentum is the energy of a system in the frame at which it is at rest. We might be led by these observations to deemphasize the original concept of energy in favor of the more relativistically natural 4-momentum and rest energy; we might sometimes even shift to using the word 'energy' to describe one or other of these new concepts (as we do, for instance, when we state that $E = mc^2$ — at least if m is likewise the *rest* mass). But the original concept of energy remains, unchanged and none the worse for its encounter with Einstein, and if we decide to shift how we use the word 'energy', that is again only semantics.

(By analogy, 4-momentum is the relativistic generalization of energy, and one understands energy in relativity more deeply once one grasps it, but the original concept of energy as the generator of time translations remains perfectly well defined, albeit reference-frame-dependent, even if one moves from non-relativistic to relativistic physics.)

We can reach the same deflationary conclusion from a statistical-mechanical starting point. Using the canonical distribution to characterize a system, and assuming a quantum microdynamics (though we could as easily use the microcanonical distribution, and/or assume classical mechanics) the density operator for a system with conserved energy and momentum is

$$\widehat{\rho}(\beta,\alpha) = Z(\beta,\alpha)\mathrm{e}^{-\beta\widehat{H}-\alpha\cdot\widehat{P}} \tag{44}$$

where

$$Z(\beta,\alpha) = \mathsf{Tr}\mathrm{e}^{-\beta\widehat{H}-\alpha\cdot\widehat{P}}. \tag{45}$$

Here β and α are Lagrange multipliers and are given, implicitly, as functionals of the expected values $U = \langle \widehat{H} \rangle_{\widehat{\rho}}$ and $\mathbf{P} = \langle \widehat{\mathbf{P}} \rangle_{\widehat{\rho}}$ of the energy and momentum. The entropy S is given by the von Neumann entropy of the density operator:

$$S = -k_B \mathsf{Tr} \widehat{\rho} \ln \widehat{\rho} \tag{46}$$

A Lorentz transformation Λ is represented by some unitary operator $\widehat{U}(\Lambda)$, acting on $\widehat{\rho}$ like

$$\widehat{\rho} \to \widehat{U}(\Lambda) \widehat{\rho} \widehat{U}^{\dagger}(\Lambda) \tag{47}$$

under which the entropy is manifestly invariant. The four operators $\widehat{H}$ and $\widehat{P}_i$ jointly transform like a 1-form: if we write $\widehat{P}_0 = \widehat{H}$ then

$$\widehat{U}(\Lambda) \widehat{P}_\mu \widehat{U}^{\dagger}(\Lambda) = \Lambda^{\nu}_{\mu} \widehat{P}_{\nu}. \tag{48}$$

From this it follows that (β, α_i) transforms like a 4-vector: if we write it that 4-vector as β^μ then

$$\widehat{U}(\Lambda) \widehat{\rho}(\beta^\mu) \widehat{U}^{\dagger}(\Lambda) = \widehat{\rho}(\Lambda^{\mu}_{\nu} \beta^{\nu}). \tag{49}$$

We can immediately recognize β^μ as the inverse 4-temperature; we can also see that the interesting part of the statistical mechanics of any such system involves calculating the density operator for a system in its rest frame, from which its density operator in any other frame can just be obtained by a Lorentz transformation.

The statistical-mechanical description of the system makes manifest something that is reasonably clear even in phenomenological thermodynamics: systems at thermal equilibrium (at least at non-zero temperature) inevitably break Lorentz covariance, in the sense that any such system has a preferred rest frame in which its momentum vanishes. This just follows from the fact that any such system consists of some amount of energetically-excited *stuff* (whether particles, field excitations, or something more exotic) and that stuff will have a rest frame in which its momentum vanishes. We see this in black-body radiation, which (as Chua (2022) observes) is only black-body in one preferred frame; we also see it in finite-temperature relativistic quantum field theory, where the Green functions, and Feynman rules, for the theory explicitly require a choice of rest frame. But this does not mean that the *framework* of equilibrium thermodynamics fails to be Lorentz-covariant, any more than relativistic particle mechanics fails to be Lorentz-covariant because the velocity of a particle picks out a preferred rest frame.

It also makes manifest that, given a system which is at thermal equilibrium with zero momentum with respect to some reference frame, then that same system has a unequivocal and readily-calculated state in another frame moving with respect to that first frame, and that state, in the absence of conserved quantities other than 4-momentum, will be the unique equilibrium state with given 4-momentum. This has some relevance given occasional suggestions (notably by Balescu (1968); see the helpful review in (Chua 2022)) that the statistics of relativistic systems at equilibrium are underdetermined by statistical mechanics when those systems are in overall motion. The cosmic microwave background

(CMB) radiation provides a particularly important example: the Earth is not at rest with respect to the CMB, so what is measured by Earth-orbiting satellites is not black-body radiation but Lorentz-boosted black-body radiation. (This boost is invariably subtracted from the CMB when it is graphically presented.)

8 The view from general relativity

(This section lies slightly outside the main development of ideas in this paper, and may be skipped on first reading.)

If we consider a physical system on a background spacetime with some metric $g_{\mu\nu}$, *in general* there can be no prospect of equilibrium thermodynamics for that system save perhaps in some local sense: thermodynamics requires energy, which requires a notion of time translation invariance, which requires a timelike Killing field (generating a symmetry which can be interpreted as time translation), and most spacetimes lack any such symmetry.

However, something instructive can be said about the thermodynamics of systems that *do* have such a symmetry — such as the spacetime of a stationary black hole, or (in approximation) any more general static arrangement of matter. For the moment let us specifically consider a system in a spherically-symmetric such geometry, such as the Schwarzschild metric

$$\mathrm{d}s^2 = -\left(1 - \frac{2GM}{r}\right)\mathrm{d}t^2 + \left(1 - \frac{2GM}{r}\right)^{-1}\mathrm{d}r^2 + r^2(\mathrm{d}\theta^2 + \cos^2\theta\mathrm{d}\phi^2) \quad (50)$$

that describes spacetime outside a spherically-symmetric mass (or outside the event horizon of an uncharged, nonrotating black hole). Such systems have Killing vectors corresponding to rotations as well as time translations, and hence conserve angular momentum (the Noether charge associated with rotation) as well as energy (which we identify as the Noether charge associated with time translation), but we will simplify just by declaring that we have no interest in modifying the system's angular momentum. We will also assume that our system is *static*: that is, it moves along the Killing vector field lines (we assume it small enough, relative to the scales on which the metric varies, that we may idealize it as pointlike). That will require some external force to maintain it in place (instead of, say, falling into the black hole or other gravitating body) and this means we can treat its spatial location as an external parameter in a thermodynamic description. A stylized description of such a body would be a small box of gas or radiation suspended on a wire from a distant space station; a more realistic case would be a small part of a larger mass of gas or plasma, filling the space (at least outside any central body) and overall static.

As one moves closer into a spherically-symmetric spacetime, one encounters increasing *gravitational time dilation*: the ratio of local time, i. e. the time measured along a Killing vector by the metric, to Killing time (i. e. the parameter that parametrises time translations). Given spherical symmetry, time dilation may be written as a function $\tau(r)$ of some radial coordinate; in the Schwarzschild metric, for instance, we have

$$\tau(r) = \left(1 - \frac{2GM}{r}\right)^{1/2}. \tag{51}$$

We will assume, again as in the Schwarzschild metric, that we have normalized Killing time so that τ is unity at spatial infinity. We may as well use τ itself as our external control parameter, so that the Second Law becomes

$$\mathrm{d}U = T\mathrm{d}S + \alpha \mathrm{d}\tau \tag{52}$$

with α the generalized pressure associated with gravitational time dilation.

Now suppose that we very slowly move the system outwards. In the adiabatic limit this will be reversible, i. e. $\delta S = 0$, and will leave the rest mass of the system invariant, $\delta M = 0$. But it is a standard result of general relativity that the Noether energy and rest mass are related by

$$U = M\tau \tag{53}$$

so that

$$\mathrm{d}U = M\mathrm{d}\tau + \tau \mathrm{d}M. \tag{54}$$

So in our reversible process we have $\mathrm{d}U = M\mathrm{d}\tau$, i. e. $\alpha = M$. Defining the *local temperature* T_R by

$$\mathrm{d}M = T_R \mathrm{d}S \tag{55}$$

and substituting (55) into (52), we obtain $T_R = T/\tau$, i. e. for fixed thermodynamic temperature, T increases as time dilation increases (i. e. as τ decreases). Local temperature, of course, corresponds to temperature as locally measured.

In fact, nothing in our analysis really requires spherical symmetry: the same result holds in a general static spacetime. (It is a standard result of relativistic thermodynamics, originally due to Tolman (1930) and Tolman and Ehrenfest (1930); see, e. g. , (Rovelli and Smerlak 2011) for a modern account and for further references.)

The Hawking radiation of an evaporating black hole provides an illustration of these ideas. If the black hole is placed in a reflecting box and allowed to fill with Hawking radiation, that radiation will be at equilibrium with the black hole itself. As such, its thermodynamic temperature will be constant. The temperature as measured by a local (static) observer will then increase as one approaches the black hole, approaching Planckian temperatures (and thus heralding the breakdown of quantum field theory) at the so-called 'stretched horizon'. (See Wallace (2018) and references therein for details.) If the box is removed, the infalling modes are unexcited (or contain only starlight and the cosmic background radiation) but the analysis remains true for outgoing modes.[8]

[8] 'Outgoing' here means 'originating at asymptotically early times from asymptotically near the event horizon'; these modes have significant amplitude to fall back into the black hole, so that a static observer close to the black hole will see radiation both coming from and falling into the black hole.

So far I have been considering the thermodynamics of a system embedded in a fixed, background spacetime. If we want to do the thermodynamics of a region of spacetime itself (as would be necessary in studying, e. g. , astrophysical self-gravitating objects or in black hole thermodynamics) what matters is the asymptotic symmetries of the system, which are set by its boundary conditions. The normal choice for astrophysical systems is to require the metric to be asymptotically Minkowskian (reflecting the fact that the cosmological constant is negligible on the scales relevant to typical astrophysical systems), in which case its asymptotic symmetries are the Poincaré symmetries, giving rise to conserved angular momentum and 4-momentum. The First Law for such a theory in principle then takes the form

$$\mathrm{d}U = T\mathrm{d}S + \mathbf{v} \cdot \mathrm{d}\mathbf{P} + \boldsymbol{\Omega} \cdot \mathrm{d}\mathbf{J} \tag{56}$$

but, in line with our earlier discussion, in general it makes sense to transform to a rest frame in which $U = M$ and spatial momentum vanishes; we may as well also rotate to a frame in which the angular momentum is in the z direction (and to write $J \equiv J_z$). The First Law then becomes

$$\mathrm{d}M = T\mathrm{d}S + \Omega\mathrm{d}J. \tag{57}$$

If there are other conserved quantities (such as particle number or charge) whose variation is of interest, they too would be included. For black holes, for instance, the no-hair theorem tells us that only charge and the Poincaré invariants are conserved; hence the First Law for black holes takes the familiar form

$$\mathrm{d}M = T\mathrm{d}S + \Omega\mathrm{d}J + \Phi\mathrm{d}Q \tag{58}$$

(see again Wallace (2018) and references therein).

9 Conclusion

Equilibrium thermodynamics is a quite general framework in which entropy is specified as a function of the energy, the other conserved quantities, and the externally-defined constraints, and in which the principle of non-decrease of entropy then sets substantive constraints on the allowable transformations performable on equilibrium systems. There is nothing intrinsically nonrelativistic about this framework: it can be applied without modification to systems in which linear (or angular) momentum is conserved, and thus straightforwardly describes systems in uniform linear motion.

Thermodynamic temperature, by definition, is the inverse of the rate of change of entropy with energy while all other conserved quantities and external constraints are held constant. It is then a direct consequence of the framework of thermodynamics applied to a relativistically-covariant system that if a system at rest with temperature T is slowly accelerated to a speed v then its temperature, in the limit as the rate of acceleration goes to zero, is $T/\gamma(v)$; more generally, $1/T$ can be identified as the 0-component of a 4-vector, and transforms as such.

Notwithstanding this, other quantities — certainly the 'rest temperature', definable equivalently as either the thermodynamic temperature in the rest frame of the system or the inverse of the rate of change of entropy with rest mass; arguably also the 'constant-velocity temperature', definable as the rate of change of energy with entropy at constant velocity — are informative and useful descriptors of a moving system's thermodynamic behavior. In the end it is a matter of convention and semantics what we call 'the' temperature. But that conventional choice is compatible with the unambiguous applicability of thermodynamics to systems in relativistic motion.

Acknowledgements

I am grateful to Eugene Chua and John Norton for helpful conversations, and for thoughtful comments on an earlier version of this manuscript.

References

Balescu, R. (1968). Relativistic statistical thermodynamics. *Physica 40*, 309–338.

Battaner, E. (1996). *Astrophysical Fluid Dynamics*. Cambridge: Cambridge University Press.

Brown, H. R. and J. Uffink (2001). The origins of time-asymmetry in thermodynamics: The minus first law. *Studies in the History and Philosophy of Modern Physics 32*, 525–538.

Chua, E. (2022). T falls apart: On the status of classical temperature in relativity. Forthcoming in *Philosophy of Science*. Preprint at http://philsci-archive.pitt.edu/20744/.

Dougherty, J. and C. Callender (2016). Black hole thermodynamics, more than an analogy? Forthcoming; http://philsci-archive.pitt.edu/13195/.

Einstein, A. (1907). Über das Relativitäts prinzip und die aus demselben gezogenen Folgerungen. *Jahrbuch der Radioactivität und Elektron 4*, 411–462.

Kapusta, J. I. and C. Gale (2006). *Finite-Temperature Field Theory: Principles and Applications* (2nd ed.). Cambridge: Cambridge University Press.

Liu, C. (1994). Is there a relativistic thermodynamics? a case study in the meaning of special relativity. *Studies in the History and Philosophy of Science 25*, 983–1004.

Norton, J. D. (2016). The impossible process: Thermodynamic reversibility. *Studies in History and Philosophy of Modern Physics 55*, 43–61.

Ott, H. (1963). Lorentz-Transformation der Wärme und der Temperatur. *Zeitschrift für Physik 175*, 70–104.

Padmanabhan, T. (2000). *Theoretical Astrophysics, Volume I: Astrophysical Processes*. Cambridge: Cambridge University Press.

Planck, M. (1906). Das Prinzip der relativität und die Grundgleichungen der Mechanik. *Verandlungen der Deutschen Physikalischen Gesellschaft 8*, 136–141.

Planck, M. (1907). Zur Dynamik bewegter systeme. *Sitzungberichte der Preussischen Akademie der Wissenchaften 13*, 542–570.

Rovelli, C. and M. Smerlak (2011). Thermal time and Tolman-Ehrenfest effect: 'temperature as the speed of time'. *Classical and Quantum Gravity 28*, 075007.

Tolman, R. C. (1930). On the weight of heat and thermal equilibrium in general relativity. *Physical Review 35*, 904–924.

Tolman, R. C. (1933). Thermodynamics and relativity. *Science 77*, 291–298.

Tolman, R. C. (1934). *Relativity, Thermodynamics and Cosmology*. Oxford: Oxford University Press.

Tolman, R. C. and P. Ehrenfest (1930). Temperature equilibrium in a static gravitational field. *Physical Review 36*, 1791–1798.

Wallace, D. (2018). The case for black hole thermodynamics, part I: phenomenological thermodynamics. *Studies in the History and Philosophy of Modern Physics 64*, 52–67.