

The Entropy of an Ideal Gas

Randy S

Abstract This article explains how to deduce the entropy of an ideal gas as a function of the total energy E and the total volume V .

Contents

1	The result	2
2	A quick “derivation”	3
3	The key input from quantum physics	4
4	Derivation for $N = 1$	5
5	An explicit set of orthogonal 1-particle states	7
6	Derivation for $N > 1$	9
7	Result for an energy shell	11
8	Is the entropy always finite?	12
9	References	13
10	References in this series	13

1 The result

As in article 66313, Let $\Omega(E, V)$ be the number of mutually orthogonal states whose total energy is $\leq E$ and total volume is $\leq V$. For an **ideal gas**, the function $\Omega(E, V)$ is

$$\Omega(E, V) \propto \frac{\omega^N}{N!} \quad \text{with } \omega = \left(\frac{mE}{\hbar^2} \right)^{D/2} V, \quad (1)$$

where N is the number of molecules,¹ D is the number of dimensions of space, m is the mass of one molecule, \hbar is Planck's constant, and the proportionality factor is dimensionless. Planck's constant enters because the state-counting is ultimately based on quantum physics. The **entropy** is defined by²

$$S(E, V) = k \log \Omega(E, V)$$

where k is **Boltzmann's constant**

$$k \approx 1.38 \times 10^{-23} \text{ J/K}. \quad (2)$$

The definition of entropy includes Boltzmann's constant for historical reasons. Equation (1) says that the entropy of an ideal gas is³

$$S(E, V) = \text{constant} + kN \left(\frac{D}{2} \log E + \log V \right), \quad (3)$$

with a constant term that depends on N, m, \hbar .

This article explains how to derive the result (1).

¹ We could consider a system in which N can vary, but N will be treated as a constant in this article.

² I'm using the notation "log" for the *natural* logarithm.

³ This is the **Sackur-Tetrode** equation, with "constant" as a placeholder for terms that depend only on N and not on E or V .

2 A quick “derivation”

An **ideal gas** is one whose molecules have no internal structure (no rotational or vibrational modes). This is a simplified model of a real gas. This section uses a quick scaling argument to deduce equation (1). The following sections present a more careful derivation.

Here’s the quick scaling argument:

- First, we expect intuitively that the number of possible locations of a single particle is proportional to the volume V , so we expect the number of states of N particles to scale with V like V^N .
- Second, the quantity $\Omega(E, V)$ should be independent of the units we use to express energy and volume. We can form a units-independent combination with the help of two constant quantities: the mass m of one molecule, and Planck’s constant \hbar . (Planck’s constant is involved because a proper derivation uses quantum physics to count the number of orthogonal states.) With these quantities, we can form the combination $\omega = (mE/\hbar^2)^{D/2}V$. This is the unique units-independent combination that is proportional to V .

Combining these two scaling arguments leads to the conclusion that the number of states must scale with ω like ω^N , as shown in equation (1).

3 The key input from quantum physics

Our current understanding of matter is based on quantum physics. A proper introduction to quantum physics would be too much for this article, but we will need one basic concept: some observables are not compatible with each other, meaning that they cannot both be perfectly measured, at least not at the same time. Maybe we shouldn't be too surprised by this, because measurement is a physical process, and we have no good reason to believe that the processes needed for measuring different observables should be perfectly compatible each other. Surprising or not, the existence of mutually incompatibility observables is a fundamental feature of quantum physics.

For the purposes of this article, the most important example is the incompatibility of the observables corresponding to measurements of object's location and momentum along the same direction in space. Those observables can both be measured at the same time, but not with arbitrary precision. Their precision is limited by the inequality⁴

$$\Delta x \Delta p \gtrsim \hbar \quad (4)$$

where Δx and Δp are the resolutions of the location- and position-measurements, respectively (along the given direction in space), and \hbar is Planck's constant. This inequality is a special example of an **uncertainty principle**.⁵ The measurements we normally work with are much coarser than this, so they are consistent with this inequality by a very large margin.⁶

The next section explains how we can use this key input from quantum physics to derive the result (1).

⁴ I'm writing the inequality loosely, using \gtrsim , because I haven't defined things carefully enough for the exact coefficient to be meaningful.

⁵ This name is *not* a good description of what the principle actually says. Names are just names, often relics from a point in history before our understanding of the physics was mature. Don't take them too seriously.

⁶ For a vivid example of this, see <https://physics.stackexchange.com/questions/440399>

4 Derivation for $N = 1$

First consider a single particle ($N = 1$). In the simplest nonrelativistic model of a single quantum particle, we can partition the spectra of the location- and momentum-observables into cells of size Δx and Δp , respectively, along each of D perpendicular directions in D -dimensional space. As long as we choose these cell-sizes to satisfy (4), the theory allows a state to be limited to just one of these cells.⁷

A state limited to one such cell is **orthogonal** to a state limited to another such cell, which means that these two states can be distinguished from each other with certainty in a single measurement event for a suitably chosen observable. Other states can be constructed using quantum theory's **superposition principle**, but we don't need those other states here because they are not orthogonal to the ones we're already counting. Statistical mechanics is based on counting the maximum number of mutually orthogonal states, so we only need to consider states that are limited to individual Δx - Δp cells.

To count the *maximum* number of orthogonal states, we should choose the product $\Delta x \Delta p$ to be as small as the inequality (4) allows. Real molecules have internal structure, and in that case we could have more than one orthogonal state per location-momentum cell. The ideal gas model pretends that the molecules don't have any internal structure, so each minimal Δx - Δp cell contributes only one state when counting the number of orthogonal states.

(Continued on the next page...)

⁷ This isn't quite true: a sharp cutoff at the boundaries of a location-cell is not compatible with a sharp cutoff at the boundaries of a momentum-cell, or conversely, but the counting described here is still valid when the boundaries are fuzzy. Section 5 demonstrates this in detail.

Now, suppose that the particle is confined to a volume V and that its kinetic energy is limited to a maximum value K . The number of location-cells in the D -dimensional “volume” V is

$$\frac{V}{(\Delta x)^D}. \quad (5)$$

The particle’s kinetic-energy observable is $p^2/2m$, where p is its momentum-observable and m is its mass, and momentum has D components in D -dimensional space, so the number of momentum-cells consistent with the maximum kinetic energy K is

$$\frac{(p_{\max})^D}{(\Delta p)^D} = \frac{(2mK)^{D/2}}{(\Delta p)^D}. \quad (6)$$

The location-cell and momentum-cell can be chosen independently (as long as (4) is satisfied), so the total number of orthogonal states consistent with the given V and K is the product of (5) and (6):^{8,9}

$$\Omega(E, V) \Big|_{N=1} \propto \frac{(2mK)^{D/2}}{(\Delta p)^D} \frac{V}{(\Delta x)^D} \propto \frac{(mK)^{D/2} V}{\hbar^D}. \quad (7)$$

For an ideal gas, the total energy E is the same as the total kinetic energy K , so (7) matches the quantity that was denoted ω in equation (1).¹⁰

That’s for a single particle. The next task is to generalize this to N particles.

⁸ I’m ignoring the overall dimensionless numerical coefficients, which don’t affect many calculations in statistical mechanics anyway (article [66313](#)).

⁹ Also remember that this assumes the *ideal gas* model, which means the particle does not have any internal structure. In a more realistic model, we would need to include another factor to account for states that are orthogonal to each other due to different excitations of the internal structure (different rotational or vibrational states, for example) despite being contained in the same location-momentum cell.

¹⁰ Article [73054](#) describes a more general model in which $E \neq K$.

5 An explicit set of orthogonal 1-particle states

This section fills in the details that were promised in footnote 7.

Quantum theory is formulated in terms of operators on a Hilbert space. Elements of the Hilbert space are used to represent states, and linear operators on the Hilbert space are used to represent observables (things that could be measured). For the standard quantum model of a single nonrelativistic particle with no internal structure, elements of the Hilbert space are conveniently represented¹¹ by complex-valued functions $\psi(\mathbf{x})$ of the D spatial coordinates $\mathbf{x} = (x_1, \dots, x_D)$. In this context, $\psi(\mathbf{x})$ is called a **wavefunction**. Each wavefunction must be square-integrable,

$$0 < \int d^D x |\psi(\mathbf{x})|^2 < \infty,$$

and two wavefunctions ψ_1 and ψ_2 are **orthogonal** to each other if¹²

$$\int d^D x \psi_1^*(\mathbf{x})\psi_2(\mathbf{x}) = 0,$$

where the asterisk denotes complex conjugation. The location and momentum observables are represented by linear operators, but for expedience I'll describe them this way instead:

- The particle is localized in a region R if $\psi(\mathbf{x}) = 0$ for all $\mathbf{x} \notin R$.
- The particle has momentum limited to a set S if $\tilde{\psi}(\mathbf{p}) = 0$ for all $\mathbf{p} \notin S$, with

$$\tilde{\psi}(\mathbf{p}) \equiv \int d^D x e^{-i\mathbf{p}\cdot\mathbf{x}/\hbar}\psi(\mathbf{x}).$$

(Continued on the next page...)

¹¹ Notice the two distinct layers of representation: elements of the Hilbert space are used to represent states, and (in this example) functions are used to represent elements of the Hilbert space.

¹² This definition of orthogonality is consistent with the one used in the preceding sections, but I won't explain why here because that would require a review of quantum theory.

Now, partition space into a grid of cells of size Δx in each dimension. For each of these cells R , define a set of wavefunctions

$$\psi_{R,\mathbf{n}}(\mathbf{x}) = \begin{cases} \exp\left(i\frac{2\pi\mathbf{n}\cdot\mathbf{x}}{\Delta x}\right) & \text{if } \mathbf{x} \in R, \\ 0 & \text{otherwise,} \end{cases}$$

where the D components of \mathbf{n} are arbitrary integers. Two such wavefunctions $\psi_{R,\mathbf{n}}$ and $\psi_{R',\mathbf{n}'}$ are obviously orthogonal to each other when $R \neq R'$, because the location-cells don't overlap, and they also turn out to be orthogonal to each other when $\mathbf{n} \neq \mathbf{n}'$. (The fact that the components of \mathbf{n} and \mathbf{n}' are integers is essential for this.) Also, the Fourier transform $\tilde{\psi}(\mathbf{p})$ of $\psi_{R,\mathbf{n}}(\mathbf{x})$ turns out to be *mostly* limited to momenta in a cell of width $\Delta p \sim \hbar/\Delta x$ about $\mathbf{p} = 2\pi\hbar\mathbf{n}/\Delta x$. To see this, use the fact that the Fourier transform of a rectangular-window function $\psi_{R,\mathbf{n}}(\mathbf{x})$ is a sinc function.

Altogether, this shows that if $\Delta x \Delta p \sim \hbar$, then we can indeed construct a set of mutually orthogonal single-particle states with one per Δx - Δp cell, even though the state associated with a given cell is not strictly contained in that cell. This justifies the approach that was used in section 4.

6 Derivation for $N > 1$

To generalize the result (7) to N particles, we need a few more concepts. The most important concept is this one:

- If A and B are two N -particle states, and if a given Δx - Δp cell is occupied in one of them but not in the other one, then those two N -particle states are orthogonal to each other, because measuring the number of particles in that Δx - Δp cell would distinguish between A and B with certainty.

From this, we can already anticipate that the number of orthogonal N -particle states is something like ω^N , the number of ways of assigning N objects to ω cells. To be more careful, we also need these concepts:

- The quantum model of an ideal gas doesn't have any observables tied to individual particles. Instead, it has observables that count the number of particles in a given region and a given range of momenta (subject to the caveat mentioned in footnote 7). When such an observable is measured, it tells us the *number* of particles detected, but it doesn't tell us *which* particles were detected – a concept that simply doesn't apply in this model.¹³ Therefore, we should count the number of ways of assigning the particles to Δx - Δp cells *modulo* permutations of the particles – because in this model, permuting the particles doesn't change the state at all.
- To do the counting correctly for arbitrary N , we would also need to know whether the particles are bosons or fermions: two or more identical bosons can occupy the same Δx - Δp cell, but two or more identical fermions cannot. However, this article assumes that N is small compared to (7), which is true in most practical applications.¹⁴ In this case, the distinction doesn't make a significant difference, because only a negligible fraction of the N -particle states would have cells with ≥ 2 particles even if this is allowed.

¹³ In traditional terminology, the particles are all **identical** or **indistinguishable**. These are just names, not good descriptions (footnote 5).

¹⁴ One exception is the interior of a white dwarf star.

Now we have all the ingredients we need to finish the derivation. As in equation (1), let ω denote the number of orthogonal states when $N = 1$, which is given by (7). If $N \ll \omega$, so that the number of states with more than one particle per cell is negligible (even if it's allowed, as it is for bosons), then the number of orthogonal N -particle states is well-approximated by

$$\frac{\omega^N}{N!}. \quad (8)$$

This doesn't correctly count states that have more than one particle in the same cell, not even if they're bosons, but the number of such states is negligible if $N \ll \omega$. With that qualification, (8) is a good approximation to the number of ways of assigning N particles to ω cells, modulo permutations of the particles.

This completes the derivation of equation (1), including the N -dependence of the coefficient that was not shown there.

7 Result for an energy shell

The preceding sections counted the number of orthogonal states with energy $\leq E$ in a volume $\leq V$, with the result $\Omega(E, V) \propto E^{DN/2} V^N$.

The result is essentially unchanged if we consider states within a narrow range of energies, between E' and E for some very small difference $E - E'$. To see this, use

$$\frac{\Omega(E', V)}{\Omega(E, V)} = \left(\frac{E'}{E} \right)^{DN/2}.$$

This shows that for $N \sim 10^{23}$, the number $\Omega(E', V)$ will be negligible compared to $\Omega(E, V)$ even when the ratio E'/E is only slightly less 1, which means that the interval $E - E'$ can be very narrow.

8 Is the entropy always finite?

Section 5 used 1-particle states that are strictly localized in a bounded region of space (and approximately localized in a bounded part of the momentum domain). Such states exist in strictly-nonrelativistic quantum mechanics, but that kind of model is only an approximation to relativistic quantum field theory. In relativistic quantum field theory, states cannot be strictly localized in any bounded region of space.¹⁵ That's okay, because they can still be localized well enough for all practical purposes, so the derivation shown in the preceding sections is still valid to an excellent approximation.

A more serious issue arises in models that are contrived so that $\Omega(E, V)$ is *infinite*. An example of such a model would be a quantum field theory with a uncountably infinite number of different species, say with a continuum of different masses. Such models do not have any realistic applications, though, and they are generally excluded even from the set of not-necessarily-realistic models that theoretical physicists study for the sake of building general intuition. Theoretical physicists typically consider only models for which $\Omega(E, V)$ is finite, a condition that can be expressed more precisely as a **nuclearity requirement**,¹⁶ so that the system has sensible thermodynamic properties.¹⁷

¹⁵ This is one implication of a general result called the **Reeh-Schlieder theorem**, which is reviewed in Witten (2018).

¹⁶ Section V.5.1 in Haag (1996)

¹⁷ More fundamentally, combining quantum physics and gravity leads to something called the **holographic principle**. This is a model-independent principle, in the sense that any viable combination of quantum physics and gravity should satisfy it. Roughly, this principle implies that the entropy in a bounded region of space is finite, and since it's model-independent, this must be true even in a model that accounts for *everything*, not just in obviously-incomplete models like the one studied in this article.

9 References

Haag, 1996. *Local Quantum Physics: Fields, Particles, Algebras*. Springer

Witten, 2018. “Notes on Some Entanglement Properties of Quantum Field Theory” *Rev. Mod. Phys.* **90**: 45003, <https://arxiv.org/abs/1803.04993>

10 References in this series

Article **66313** (<https://cphysics.org/article/66313>):
“Basic Statistical Mechanics” (version 2022-02-05)

Article **73054** (<https://cphysics.org/article/73054>):
“Why Do Liquids Boil?” (version 2022-02-05)