Class 13 Notes: Interstellar dust: optical properties

We are now entering the second half of the course, where we study the individual components of the ISM. The first component, on which we will focus this week, is dust grains. By dust grains we mean solid particles that are distinguished by having the property that, as solids, they can interact with radiation over a broad wavelength range, rather than at a few resonant frequencies. Thus far our treatment of matter-radiation interaction has focused on line interactions, so we begin this class with a discussion of how dust interacts with radiation. This is a problem in classical electromagnetism – no quantum required for the most part. Once we have this theory in hand, we will use it to explore the thermodynamics and charge state of dust grains.

I. Interaction of radiation with small particles

A. Dielectrics and absorption

We are all familiar with how radiation interacts with large object, where "large" means that the object is much larger than the wavelength of the light with which it is interacting. This is the realm of geometric optics that we learn in high school physics. For the most part, however, interstellar dust is not in this regime, since many dust grains have sizes that range from a few to a few thousand nanometers. Thus we must develop a theory that works for small objects.

Let us start with a review of some of the basic classical theory for the interaction of electromagnetic waves with matter. Suppose we have a constant vacuum electric field $\mathbf{E}$, and we put an electrically-polarisable medium into that field. The medium will polarise in response to the applied electric field, producing a dipole field $\mathbf{P}$. The resulting net field within the medium will be a sum of the two,

$$\mathbf{D} = \mathbf{E} + 4\pi\mathbf{P} \equiv \epsilon\mathbf{E}, \tag{1}$$

where ϵ is the dieletric constant of the medium. A perfect conductor has $\epsilon = 0$, since it excludes electric fields, while empty space has $\epsilon = \epsilon_0 = 1$, since the electric field will just be equal to the applied one [1] We define the related quantity the electric susceptibility χ by

$$\epsilon = 1 + 4\pi\chi, \tag{2}$$

[1] This is another case where we have to be careful about electromagnetic unit systems. As usual in this course, we work in Gaussian units, where the applied electric field $\mathbf{E}$ and the dipole field $\mathbf{P}$ produced by the background medium have the same units, and the net electric field $\mathbf{D}$ is just their sum; the 4π is present in Gaussian units, but would not appear in "rationalised" unit systems like Heaviside-Lorentz or SI units. On the other hand, in the SI system $\mathbf{D}$ and $\mathbf{E}$ have different units, and the relationship between them is $\mathbf{D} = \epsilon_0\mathbf{E} + \mathbf{P}$, i.e., in the SI system there is a distinction between an electric field and a "displacement field", which is obtained by multiplying the electric field by the permittivity of the medium to which is applied. As usual, it is important to understand that the displacement field is not a real thing that is physically distinct from an electric field; you can't do an experiment that can tell the difference between an electric field and

so that empty space has $\chi = 0$.[2]

Now let us apply a time-varying electric field to the medium, for example a propagating electromagnetic wave, i.e., light. For simplicity we will take the light to be monochromatic, and we will orient our coordinate system so that the wave is travelling in the $+x$ direction, so the electric field is $\mathbf{E}e^{i(kx-\omega t)}$. We adopt the usual convention that we use the real part of complex quantities. We can built an arbitrary spectrum of light traveling in this direction out of a combination of such waves with different values of ω. At every point, the medium responds by polarising, such that within it we have a field

$$\mathbf{D} = \epsilon(\omega)\mathbf{E}e^{i(kx-\omega t)}. \tag{3}$$

Note that we have written $\epsilon(\omega)$, i.e., the value of ϵ depends on the frequency with which the electric field is varied. This makes sense: the frequency of variation is not necessarily small compared to the characteristic timescale associated with electrons moving around atoms. Thus the medium does not respond instantaneously, and the response will depend on the properties of the atoms or molecules providing the polarisation, and how their characteristic timescales compare to the frequency of the imposed field. The dependence of $\epsilon(\omega)$ on ω implies that media can be dispersive, meaning that they cause waves of different frequencies to propagate at different speeds through them.

It is also important to note that $\epsilon(\omega)$ need not be real-valued. In our convention of taking the real parts of imaginary-valued quantities, a non-real value of $\epsilon(\omega)$ corresponds to having a medium that shifts the phase of a wave passing through it. However, such phase-shifting necessarily means that the medium changes the intensity of the light passing through it. To see why, recall what we get by plugging our waveform into Maxwell's Laws: a propagating wave must obey a dispersion relation

$$k^2 = \frac{\omega^2}{c^2}\epsilon(\omega)\mu, \tag{4}$$

where μ is the (real) magnetic permeability of the medium. Thus k is a complex number; one can solve for the imaginary part, and in the limit where the real part is much bigger than the imaginary part (i.e. where the attenuation per wavelength is small) we have

$$\mathrm{Im}\, k = \frac{\omega}{2c}\sqrt{\frac{\mu}{\mathrm{Re}\,\epsilon(\omega)}}\mathrm{Im}\,\epsilon(\omega). \tag{5}$$

---

a displacement field. The distinction only exists because SI introduces a new fundamental unit for electric charge or current, rather than defining these quantities in terms of the base mechanical units as is done in the Gaussian system.

[2] Note that there is some ambiguity in notation and terminology here. *Draine* uses α as the symbol for susceptibility, and he refers to it as polarisability. My terminology is consistent with *Jackson*. The distinction between polarisability and susceptibility is very subtle, since both refer to the scaling between an applied electric field and the induced dipole field. Polarisability is usually used to refer to the induced dipole moment of a single molecule, while susceptibility refers to the analogous phenomenon for a macroscopic medium containing multiple molecules. Interstellar dust grains fall somewhere in the grey area between these two cases, hence the ambiguity of terminology.

Except at very low frequencies the magnetic permeability of most materials is $\mu \approx 1$,[3] because magnetisation results from alignment of electron spins and orbits, and the time required for this alignment is of order the precession period of the spin or orbit in the applied magnetic field. Typical values for this time are at least 10^{-10} s, but for optical light typical oscillation periods of the magnetic field are $\sim 10^{-15}$ s. Thus the field oscillates much too rapidly for the atoms to respond, and we can set $\mu(\omega) = 1$ as long as we are considering any wavelength shorter than radio waves. For a medium like the ISM that is not so dispersive as to slow light down much $\mathrm{Re}\,\epsilon(\omega) \approx 1$ as well, so that we simply have

$$\mathrm{Im}\,k \approx \frac{\omega}{2c}\mathrm{Im}\,\epsilon(\omega) = \frac{\pi}{\lambda}\mathrm{Im}\,\epsilon(\omega). \tag{6}$$

The wave varies with position at fixed time as e^{ikx}, so if $\mathrm{Im}\,k > 0$, then the amplitude of the wave varies as $e^{-\mathrm{Im}\,kx}$, i.e., it is exponentially attenuated as it passes through the medium. The power carried in the wave varies as $|\mathbf{E}|^2$, so the attenuation length is $\mathrm{Im}\,k/2$. Thus the optical depth of a path of length L through the medium is

$$\tau = 2\,\mathrm{Im}\,k\,L = \frac{2\pi}{\lambda}L\,\mathrm{Im}\,\epsilon(\omega). \tag{7}$$

B. Scattering and absorption in the long wavelength limit

1. Absorption

Now consider a spherical grain of radius a with complex-valued dielectric function $\epsilon(\omega)$. We must be somewhat careful in how we define this quantity, because we have to choose the scale over which we are averaging. Are we interested in the dielectric function describing how the atoms within a small, uniform, sub-part of a grain behave? Do we mean to consider the response of a medium consisting of many such small particles with vacuum in between them, such as the ISM averaged over scales much larger than a single grain? Or do we mean to average over a region that is much larger than a single grain, but much smaller than the typical inter-grain separation? Any of these choices yield meaningful dielectric functions, but they're not necessarily all the same.

To be clear, in this section we will define $\epsilon(\omega)$ as the dielectric function describing a volume of space containing a single grain, and $\epsilon_g(\omega)$ as the dielectric function of the material of which that grain is made, averaged over a scale much smaller than the grain itself. We defer averaging over scales of many grains for now.

Now consider a plane wave striking a grain. We have already seen that the wave will be attenuated by an amount that depends on the imaginary part of

[3] Again, in Gaussian units. In SI units the equivalent statement would be that $\mu \approx \mu_0$, where μ_0 is the magnetic permeability of free space.

$\epsilon(\omega)$. We define the absorption cross-section of the grain as the optical depth per unit length

$$C_{\rm abs} = \frac{\tau}{L} = \frac{\omega}{c}\,\mathrm{Im}\,\epsilon(\omega) = 4\pi\frac{\omega}{c}\,\mathrm{Im}\,\chi(\omega). \tag{8}$$

2. Scattering

In addition to absorbing radiation, the grains can scatter it. To see how this works, consider that the electric field of the incident plane wave is

$$E_{\rm inc} = E_0 e^{i(kx-\omega t)}. \tag{9}$$

In the limit of a very small grain, $a \ll 1/k$, we can neglect the spatial variation of the electric field over the grain's volume, and approximate the incident electric field as uniform. A uniform electric field applied to a dieletric substance induces a dipole electric field

$$p = \chi(\omega)E_{\rm inc} = \chi(\omega)E_0 e^{-i\omega t}. \tag{10}$$

Since $\mathbf{E}$ varies in time as $e^{-i\omega t}$, p does as well, and a time-varying electric dipole radiates. The time-averaged power per unit solid angle that is radiated at a direction θ relative to the initial plane wave is then given by the dipole radiation formula:

$$\frac{dP}{d\Omega} = \frac{\omega^4}{4\pi c^3}|\chi(\omega)|^2 E_0^2 \sin^2\theta. \tag{11}$$

We define the differential absorption scattering cross section as simply the ratio of this radiated power to the power carried in the incident wave, given by

$$P_{\rm inc} = \frac{c}{4\pi}E_0^2. \tag{12}$$

Thus we have

$$\frac{dC_{\rm sca}}{d\Omega} = \frac{\omega^4}{c^4}|\chi(\omega)|^2\sin^2\theta. \tag{13}$$

The total scattering cross section is simply this integrated over all angles:

$$C_{\rm sca} = \int \frac{dC_{\rm sca}}{d\Omega}\,d\Omega = \frac{8\pi}{3}\frac{\omega^4}{c^4}|\chi(\omega)|^2. \tag{14}$$

Note that, in order to conserve energy, the unscattered portion of the wave that emerges on the far side of the grain must be reduced in amplitude. The extinction cross section is simply the sum of the scattering and absorption cross sections.

This process is known as Rayleigh scattering, and the general result that the scattering cross section varies with wavelength as $(\omega/c)^4 \propto \lambda^{-4}$ is known as Rayleigh's law. It is the reason the sky is blue: the scattering cross section of air molecules (which are much smaller than a wavelength of visible light) varies as λ^{-4}, so the shorter wavelength blue light emitted by the Sun has a much larger scattering cross section than the red light, and it is heavily scattered.

3. Calculating cross sections

We have now figured out how to compute scattering and absorption cross sections given a value for the electric susceptibility $\chi(\omega)$. The next task is to compute its value. Fortunately, this is relatively straightforward problem in electrostatics. We consider a grain of a given shape composed of material with dieletric function $\epsilon_g(\omega)$, we impose an external electric field, and we solve for the resulting electric field. This is just a boundary-value problem, which is trivial to solve numerically, and is solvable analytically for sufficiently simple grain geometries. In particular, if the grains are spheres of volume V, then

$$\chi = \frac{3V}{4\pi}\left(\frac{\epsilon_g - 1}{\epsilon_g + 2}\right). \tag{15}$$

Draine gives more general formulae for ellipsoidal grains, but the scaling is the same as for spheres.

Plugging this into our cross section formulae, we have

$$\begin{aligned} C_{\mathrm{abs}} &= 18\pi \frac{\operatorname{Im}\epsilon_g}{(\operatorname{Re}\epsilon_g + 2)^2 + (\operatorname{Im}\epsilon_g)^2}\frac{V}{\lambda} \\ C_{\mathrm{sca}} &= 24\pi^3\left|\frac{\epsilon_g - 1}{\epsilon_g + 2}\right|^2 \frac{V^2}{\lambda^4}. \end{aligned}$$

It is common to normalize the absorption and scattering coefficients by the geometric size of a grain of equal volume. We define

$$Q_{\mathrm{sca}} = \frac{C_{\mathrm{sca}}}{\pi a_{\mathrm{eff}}^2}, \tag{16}$$

and similarly for absorption, where $a_{\mathrm{eff}} = (3V/4\pi)^{1/3}$ is the radius of a sphere with volume equal to the grain volume. The quantity Q is called the absorption or scattering efficiency, and is dimensionless. Thus for spherical grains we have

$$\begin{aligned} Q_{\mathrm{abs}} &= 12 \frac{\operatorname{Im}\epsilon_g}{(\operatorname{Re}\epsilon_g + 2)^2 + (\operatorname{Im}\epsilon_g)^2}\frac{2\pi a}{\lambda} \\ Q_{\mathrm{sca}} &= \frac{8}{3}\left|\frac{\epsilon_g - 1}{\epsilon_g + 2}\right|^2 \left(\frac{2\pi a}{\lambda}\right)^4. \end{aligned}$$

It is also helpful to consider the limiting behavior of ϵ_g in the long wavelength, low frequency limit. At long wavelengths the real part of the dielectric function ϵ approaches a constant value regardless of whether the medium is an insulator or a conductor. The imaginary part scales linearly with ω for insulators, and scales as $4\pi\sigma_0/\omega$ for conductors, where σ_0 is the conductivity. Formally,

$$\lim_{\omega\to 0}\epsilon_g = \begin{cases} \epsilon_0 + iA\omega, & \text{(insulator)} \\ \epsilon_0 + i(4\pi\sigma_0/\omega), & \text{(conductor)} \end{cases} \tag{17}$$

Plugging these in we find that in either case $C_{\text{abs}} \propto V/\lambda^2$ and $C_{\text{sca}} \propto V^2/\lambda^4$ for large wavelengths.

This tells us three things immediately. First, at long wavelengths, absorption will dominate over scattering as a source of extinction. Second, the total absorption and total extinction will vary with wavelength as λ^{-2} in this limit. Third, the extinction will simply be proportional to the total volume of grain material.

C. Intermediate wavelengths

When we are out of the very long and very short wavelength regimes that we have now solved, life is considerably more complicated, and all the methods available to us are numerical to some degree or another. There are two approaches generally in use: Mie theory and the discrete dipole approximation.

Mie theory is a method for handling scattering by spherical dielectrics of arbitrary size. The basic approach is to decompose the electric field inside and outside the medium into spherical harmonics. The field is then subject to boundary conditions at the surface of the sphere and at infinity, and these give rise to a series of algebraic equations for the coefficients of the spherical harmonics. These can be solved relatively easily on a computer, and, given the coefficients, one can solve for the scattering and absorption cross sections.

The results in detail depend on the dielectric constant of the grain material, but the general outcome is as we might have expected. For $2\pi a/\lambda \ll 1$, we approach the long wavelength regime: extinction is dominated by absorption, and the total extinction cross section varies as V/λ^2. For $2\pi a/\lambda \gg 1$, the extinction cross section approaches twice the geometric cross section – twice because there is always some small-angle scattering. For sufficiently large $2\pi a/\lambda$ the scattering is so forward-peaked that in practice it becomes indistinguishable from the wave simply being transmitted, but formally the light is scattered. Finally, for media of moderate absorptivity, the cross section peaks when $2\pi a/\lambda \sim 1$.

For non-spherical grains even Mie theory is not available, and in general one must resort to a more brute force approach. The general method is called the discrete dipole approximation, and it involves approximating the grain as a series of small dipoles, each of which may represent a single atom or a cluster of atoms – all that is required is that the dipoles themselves be small compared to the wavelength of light in question, so we can use the long wavelength limit for them. Then one numerically solves for their collective response to an incident plane wave.

II. Observational constraints: the grain size distribution

Now that we have some theoretical understanding of how small particles in the ISM interact with light, we are in a position to interpret observations of interstellar dust. Our goal is to use these observations to deduce what interstellar dust must look like. There are four main types of observations of which we can make use, which we will discuss in turn.

A. Extinction

As we discussed back in the first class, one of the first pieces of evidence about the existence of the ISM was the reddening and dimming of distant stars. The observational technique here is fairly simple: find nearby examples of stars of a given spectral type, which are subject to negligible extinction. Then observe more distant stars of the same spectral type and compare the emitted spectra with the reference spectra. The difference is a measure of the wavelength-dependent interstellar extinction along the line of sight to that star. Formally, the extinction in magnitudes is defined as

$$A_\lambda = 2.5 \log_{10}(F_\lambda^0 / F_\lambda), \tag{18}$$

where F_λ is the observed flux and F_λ^0 is the reference flux, corrected to the distance of the observed star. This is related to the optical depth in the usual manner: $A_\lambda = 1.086\tau_\lambda$.

This is often most useful if it is normalised to the amount of hydrogen along the line of sight to the star, and so the quantity typically reported is the extinction per H atom. The H column can be measured if the observations extend into the UV, because then we can measure the column density of atomic H using Lyman α absorption, and the column density of H_2 using Lyman-Werner band absorption around 1000 Å. The H column can also be inferred from 21 cm observations, although in that case one faces the extra complication that the 21 cm beam is always going to be averaging over a much larger area on the sky than the stellar absorption measurement.

The typical result is shown in the slide.

[Slide 1 – typical Milky Way extinction curves]

Observed extinction curves vary from one line of sight to another, and the most common way to characterise this variation is by the dimensionless ratio

$$R_V \equiv \frac{A_V}{A_B - A_V}, \tag{19}$$

where A_V is A_λ integrated over the V band from $500 - 700$ nm, and A_B is A_λ integrated over the B band from $400 - 500$ nm. The quantity $A_B - A_V$ is called the reddening, and is sometimes written $E(B - V)$. This is effectively a measure of the slope of the curve in the $400 - 700$ nm region.

[Slides 2 – extinction curves plus B and V bands]

The different curves shown are typical of lines of sight with the value of R_V indicated by the arrows. Indeed, to good approximation the shape of the extinction curve is fully specified by the value of R_V; lines of sight with the same R_V tend to have very similarly-shaped extinction curves. Typical values of R_V through diffuse gas in the Milky Way have $R_V \approx 3.1$; regions where the ISM is denser tend to have higher R_V values. The lowest observed values of R_V are around 2.1.

The overall shape of the curve is interesting, because it provides a strong constraint on the grain size distribution. Given a particular model for the grain size distribution, one can use the theory of the interaction of radiation with small grains to compute the wavelength-dependent extinction that grain population should produce, and then compare to the observations. While this sort of forward modelling is necessary to test models quantitatively, one can make some general observations about what the grain size distribution must be like simply from observing the form of the extinction curve.

First, note that the extinction generally increases with frequency, so that bluer colours are more extincted than redder ones. This immediately tells us something about the sizes of the dust grains responsible for providing the extinction. An object that is large compared to the wavelength of light with which it is interacting is in the limit of geometric optics. In this limit, the amount of light blocked is obviously not wavelength-dependent; in terms of R_V, in this limit we have $R_V \to \infty$. Recall that the geometric optics limit applies down to a wavelength of order $\lambda \sim 2\pi a$, where a is the size of the object. That there is strong wavelength dependent out to 100 nm implies that the smallest grains must be smaller than $a \sim \lambda/2\pi \approx 15$ nm.

Conversely, the grains cannot all be that small. As we showed in the previous class, in the limit where the size of the scattering object is very small compared to the wavelength of light, the extinction varies with wavelength as $A_\lambda \propto \lambda^4$ – the Rayleigh scattering limit. Integrating this over the B and V bands gives $R_V = 1.2$, lower than any observed R_V value. Thus there must be significant extinction coming from grains that have $a \sim \lambda_V/2\pi \gtrsim 100\ \mu$m.

Second, there is a large bump in the opacity at around 2175 Å, referred to, somewhat unimaginatively, as the 2175 Å bump. The origin of this feature is not completely understood. Its peak is always at the same location, but its width and shape vary from one sightline to another, suggesting that there is no single species responsible for it. The feature may be due to a mix of carbonaceous grains of different sizes, but this remains uncertain.

Third, there are a number of other features toward the infrared end of the spectrum. Three of the more prominent ones are bumps at 3.4, 9.7, and 18 μm. The 9.7 and 18 μm bumps are associated with amorphous silicate grains – basically sand. The 9.7 μm feature comes from stretching of Si-O bonds, and the 18 μm feature from bending of O-Si-O bonds.. The 3.4 μm feature is due to stretching of C-H bonds in carbonaceous grains or hydrocarbon compounds.

B. Polarisation

In addition extincting background starlight, interstellar dust also polarises it. We will not discuss the polarisation mechanism in detail, but a short summary is that, if grains are non-spherical, they are better at scattering or absorbing light that is linearly polarised with the electric field oscillating along the long axis of the grain – in effect they act like little antennae. Stars emit unpolarised light, but as this propagates through interstellar dust, one linear polarisation mode is extincted more than the other, and the light that reaches us is polarised.

If grains were randomly oriented, the net polarisation would be zero. Grains in one orientation would cancel those in other orientations. However, interstellar space is filled with large-scale magnetic fields, and grains of non-zero magnetic susceptibility will preferentially orient themselves with respect to the magnetic field. The exact alignment mechanism, and the degree of alignment, are still not understood, and are hotly debated in the literature. Regardless of the mechanism, somewhat the grains get preferentially lined up in one orientation, producing a net polarisation.

[Slide 3 – polarisation vectors in the Pipe nebula]

The amount of polarisation is dependent on the wavelength of the light, with the maximum usually in the V band. A purely empirical relation, known as the Serkowski law, describes the polarisation degree:

$$p(\lambda) = p_{\max} \exp[-K \ln^2(\lambda/\lambda_{\max})], \tag{20}$$

where $\lambda_{\max} \approx 5500$ Å, $K \approx 1.15$, and $p_{\max}$ along a given sightline varies from 0 to $\sim 0.09E(B-V)/\mathrm{mag}$. The variation in values of $p_{\max}$ probably reflects the degree of uniformity and the orientation of the magnetic field along a given line of sight – lines of sight where the magnetic field varies greatly, or where it is oriented to that the long axes of the grains are close to the line of sight, produce little polarisation.

It is interesting to notice that the total extinction continues to rise into the UV all the way to the point where observations are cut off by hydrogen absorption, whereas the degree of polarisation peaks in V and declines toward the UV. Again, think about the size of the grains responsible for the polarisation compared to the wavelength of light. For wavelengths of light much smaller than grain size, we enter the geometric optics limit, and the degree of extinction is no longer polarisation-dependent: a brick wall does not discriminate between polarisation modes, whereas a field of radio antennae does.

The fall-off of the polarisation degree blueward of V is most easily interpreted as the grains that are responsible for extinction there being either largely spherical or largely non-aligned with the magnetic field. Thus, we conclude that the grains responsible for extinction at $\lambda \lesssim 0.3$ μm, which are $a \approx \lambda/2\pi \lesssim 0.05$ μm in size, are non-spherical or non-aligned. In contrast, the grains responsible for extinction at V band ($\lambda \approx 0.55$ μm), which have $a \approx \lambda/2pi \approx 0.1$ μm, are both non-spherical and aligned.

C. Scattering

A third method of observing interstellar dust is via scattering. When a dust cloud is found near a bright star, we often see the cloud shine in scattered starlight. We can tell that the light is scattered starlight because its spectrum shows the same features as the originating star.

[Slide 4 – the Merope reflection nebula]

A comparison between the intensity of the scattered light we see and the amount of light that we estimate is incident on the cloud lets us measure the albedo of the dust, defined as the ratio of the scattering cross section to the extinction cross section. If we think we know the geometry of the system, we can also estimate the scattering asymmetry parameter $\langle \cos\theta \rangle$, where θ is the scatting angle, as indicated below. The angle brackets indicate the average weighted by the intensity of the scattered light, so that $\langle \cos\theta \rangle = 1$ corresponds to scattering purely in the forward direction, $\langle \cos\theta \rangle = -1$ corresponds to perfect backscattering, and $\langle \cos\theta \rangle = 0$ corresponds to isotropic scattering.

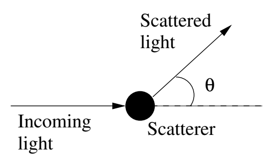

In addition to individual reflection nebulae, we can also study scattering of diffuse galactic light (DGL), contributed by many stars, off interstellar dust. This is much dimmer and hard to observe, but has the advantage that the DGL is relatively isotropic in the plane and has a reasonably well-known dependence on height above the plane, so there is no geometric ambiguity as there is for reflection nebulae.

[Slide 5 – albedo and scattering asymmetry]

We see that albedos in the visible are typically ~ 0.5 μm, falling off slightly toward the UV and rapidly toward the IR. Thus reflection and scattering are about equally important for visible and UV light, while absorption dominates at IR and longer wavelengths. The particles are mildly forward-scattering at optical wavelengths, $\langle \cos\theta \rangle \sim 0.5$. This is inconsistent with the particles being very small, since scattering by small particles is isotropic. Thus the particles that dominate the scattering must have sizes $a \gtrsim \lambda_V/2\pi \approx 0.1$ μm.

D. Thermal emission

Yet another method of observing interstellar grains is by their thermal emission. Since grains are rarely heated to temperatures hot enough to emit optically (and most materials do not survive such temperatures), this emission is generally in the infrared.

[Slide 6 – IR dust emission spectrum]

Most grains are cool enough that they emit primarily in the far IR; this produces the peak at ~ 100 μm, corresponding to grains at temperatures of $\sim 15-20$ K. However, there is also significant emission at significantly shorter wavelengths; producing emission at $\sim 10-30$ μm requires temperatures of at least ~ 50 K. Some of this emission comes from regions near hot stars where the grains are hotter than the average, but there is significant emission at $\sim 10-30$ μm even along lines of sight where there are no young stars. Moreover, emission short of 50 μm accounts for $\sim 1/3$ of the total radiated power, and this cannot be explained by local heating sources.

Instead, the short wavelength emission comes from another source: very small grains, $\lesssim 50$ Å in size, have such small specific heats that they can be heated to $\gtrsim 50$ K by absorption of a single visible photon. This is a stochastic process, since the grain rapidly cools, but such stochastic absorption events can provide a population of warm grains large enough to explain the observed near- to mid-IR emission, provided that there is a significant population of grains at such small sizes. To make the energetics work, these grains must account for $\sim 1/3$ of the total absorption of starlight in the galaxy.

E. Summary of the observational constraints

We have seen that there are a large number of constraints on the sizes of interstellar grains. We can put all of these together, using our theory for how grains of different sizes interact with light of various wavelengths, to obtain an estimated grain size distribution. The process is that one proposes a grain size distribution, then calculates the extinction curve, the polarisation and scattering properties, and the IR emission from this size distribution and the theoretical model, then tries to find the best fit to the data.

In practice, and not surprisingly, this does not generally produce a single unique results. The problem is still somewhat under-constrained. However, all models that do a reasonably good job of matching the data share the feature that the distribution of grain sizes is very broad. The smallest grains must have sizes $\sim$ 10s of Å, in order to explain both the frequency-dependence of the extinction curve in the UV and near-IR emission from warm grains. The largest grains must be at least ~ 0.2 μm in size, in order to explain the bias to forward scattering at optical wavelengths and an extinction curve slope R_V inconsistent with scattering by very small particles.

III. Observational constraints: the total dust mass budget

In addition to using observations to constrain the distribution of grain sizes, we can also use them to constrain the total amount of dust in the ISM. The basic result is known as the Purcell limit, which we will now derive.

A. The Kramers-Kronig relations

The Purcell limit follows from a general result in E&M known as the Kramers-Kronig relations, which put restrictions on the polarisability of physical media. To remind you, we define the dielectric function $\epsilon(\omega)$ of a medium by the statement that, if we impose an external electric field $\mathbf{E}e^{i(kx-\omega t)}$ on that medium, the electric field within it will be

$$\mathbf{D} = \epsilon(\omega)\mathbf{E}e^{i(kx-\omega t)}. \tag{21}$$

The Kramers-Kronig relations follow from a simple statement: the polarisation of the medium, and thus the electric field within it, had better depend only on the current value of the electric field and its value in the past, not on its value in the future. In other words, if we write $\mathbf{D}(x,t)$ as the field inside the medium at position x and time t, then it must be possible to express that field as

$$\mathbf{D}(x,t) = \mathbf{E}(x,t) + \int_0^\infty G(\tau)\mathbf{E}(x,t-\tau)\,d\tau, \tag{22}$$

where $G(\tau)$ is an arbitrary function. The key point here is that the integral only extends to $\tau > 0$, i.e., to the past; the history of the electric field can matter, but its future cannot. Simple substitution shows that $G(\tau)$ is related to our dielectric function $\epsilon(\omega)$ by

$$\epsilon(\omega) = 1 + \int_0^\infty G(\tau)e^{i\omega\tau}\,d\tau. \tag{23}$$

The requirement that $\epsilon(\omega)$ must be of this form implies some very general conditions on it. Physically, we can think of this as follows: the material polarises in response to the varying electric field, but it can only do so based on the field in the past. This limit in time corresponds to a limit in frequency. We will not prove it (see *Jackson*'s E&M text for a proof), but it can be shown that mathematically this implies that

$$\mathrm{Re}\,\chi(\omega') = \frac{2}{\pi}\int_0^\infty \frac{\omega\,\mathrm{Im}\,\chi(\omega)}{\omega^2-\omega'^2}\,d\omega, \tag{24}$$

where $\epsilon(\omega) = 1 + 4\pi\chi(\omega)$, and χ is the electric susceptibility.

This is one of the two Kramers-Kronig relations, and it is a rather surprising result. Remember that the real part of $\epsilon(\omega)$ tells us about the dispersion of a medium, i.e., how waves of different frequencies travel at different speeds within the medium. The imaginary part tells us about the ability of the medium to absorb radiation. The Kramers-Kronig relations show that these two are not independent – if you know one, you must know the other.

B. Implications of the Kramers-Kronig relations

Now let us see what this relation implies fo the ISM. First consider the case $\omega = 0$, corresponding to applying a constant, non-time-varying electric field to our medium. Plugging this into the Kramer's-Kronig relation, we have

$$\mathrm{Re}\,\chi(0) = \frac{2}{\pi}\int_0^\infty \frac{\mathrm{Im}\,\chi(\omega)}{\omega}\,d\omega = \frac{2}{\pi}\int_0^\infty \frac{\mathrm{Im}\,\chi(\lambda)}{\lambda}\,d\lambda, \tag{25}$$

where in the last step we have used the relationship $\mathrm{Im}\,\epsilon = 4\pi\,\mathrm{Im}\,\chi$. Recall that we showed above that the optical depth associated with a path of length L through a medium of dielectric function $\epsilon(\omega)$ is

$$\tau = \frac{2\pi}{\lambda} L \, \mathrm{Im}\, \epsilon(\omega), \tag{26}$$

so substituting this in and recalling that $\mathrm{Im}\, \epsilon = 4\pi \, \mathrm{Im}\, \chi$, we have

$$\mathrm{Re}\, \chi(0) = \frac{1}{4\pi^3} \int_0^\infty \frac{\tau_\lambda}{L} \, d\lambda. \tag{27}$$

Now let us consider a path through the galaxy of hydrogen column density $N_{\mathrm{H}} = n_{\mathrm{H}} L$, where n_{H} is the hydrogen volume density. Thus we have

$$\frac{1}{n_{\mathrm{H}}} \mathrm{Re}\, \chi(0) = \frac{1}{4\pi^3} \int_0^\infty \frac{\tau_\lambda}{N_{\mathrm{H}}} d\lambda = \frac{1}{4\pi^3} \int_0^\infty \frac{1}{1.086} \frac{A_\lambda}{N_{\mathrm{H}}} d\lambda. \tag{28}$$

As we have seen, the quantity in the integral on the right hand side can be determined observationally, so we may consider it known. Let us therefore turn out attention to the left hand side.

Here the trick is that we want to consider not the dielectric response of a single grain or of the material within a grain, as we have done so far, but *the dielectric response of the ISM as a whole*. In other words, we want to average over scales much larger than a single grain, and compute ϵ or χ for this large collection of particles.

The static polarisability of a medium consisting of a bunch of dielectric particles can be computed in a fairly straightforward manner, though we won't go through it in class. It depends on the particles' shape and on their conductivity, and of course on the number density of the particles. The result from electrostatics is that, for spheroidal particles of semi-major axes (a, b, b), the electric susceptibility is

$$\mathrm{Re}\, \chi(0) = n_{\mathrm{gr}} a b^2 F(a/b, \epsilon_g), \tag{29}$$

where n_{gr} is the number density of grains and $F(a/b, \epsilon_g)$ is a dimensionless function that depends on the grains' shape and their internal dielectric constant ϵ_g. Its value is at most unity for spherical grains even if they are perfect conductors, or for elongated grains if they are imperfect conductors; values $\gg 1$ are possibly only if the grains are all elongated needles that conduct perfectly, which is implausible given their likely chemical composition (carbonates, silicates, etc.).

Note, that $\delta \equiv (4/3)\pi n_{\mathrm{gr}} a b^2$ is simply the fraction of the spatial volume that is occupied by grains. Thus we have

$$\mathrm{Re}\, \chi(0) = \frac{3}{4\pi} \delta F(a/b, \epsilon_g), \tag{30}$$

and plugging this in we have

$$\frac{\delta}{n_{\mathrm{H}}} = \frac{1}{3\pi^2 F(a/b, \epsilon_g)} \int_0^\infty \frac{1}{1.086} \frac{A_\lambda}{N_{\mathrm{H}}} d\lambda. \tag{31}$$

We have therefore derived an estimate for the grain volume per unit hydrogen density. Plugging in the observed extinction curve for the Milky Way, we have

$$\frac{\delta}{n_{\rm H}} \gtrsim 3.7 \times 10^{-27} F^{-1} \ {\rm cm}^3/{\rm H}, \tag{32}$$

where we have written the result as a limit because we only have the extinction curve well-measured from roughly $0.1 - 30$ μm, and we have taken $A_\lambda = 0$ outside this range. If we adopt a characteristic grain mass density $\rho_{\rm gr} \approx 3$ g cm^{-3}, typical of rocky materials, we obtain an estimate for the grain mass fraction:

$$\frac{M_{\rm gr}}{M_{\rm H}} \gtrsim 0.0083 \left(\frac{0.8}{F}\right) \left(\frac{\rho_{\rm gr}}{3 \ {\rm g \ cm}^{-3}}\right), \tag{33}$$

where the fiducial value of F comes from our best guess as to grain properties. Thus we obtain a lower limit that the mass of grains is $\sim 1\%$ of the mass of hydrogen in the ISM simply from the extinction curve.