MP465 – Advanced Electromagnetism

Lectures 15 & 16 (2 April 2020)

B. Plane Waves in Linear Media

1. Plane Waves in a Single Medium

For the rest of this section, we're going to concentrate on *linear* media only, i.e. ones with a given permittivity ϵ and permeability μ. Furthermore, we're going to assume that the medium contains no *free* charges or currents, only bound ones. This means that both ρ and $\vec{J}$ vanish and we have the free sourceless equations

$$\vec{\nabla} \cdot \vec{D} = 0, \qquad \vec{\nabla} \times \vec{E} = -\frac{\partial \vec{B}}{\partial t},$$
$$\vec{\nabla} \cdot \vec{B} = 0, \qquad \vec{\nabla} \times \vec{H} = \frac{\partial \vec{D}}{\partial t}.$$

With $\vec{D} = \epsilon \vec{E}$ and $\vec{H} = \vec{B}/\mu$, these become

$$\vec{\nabla} \cdot \vec{E} = 0, \qquad \vec{\nabla} \times \vec{E} = -\frac{\partial \vec{B}}{\partial t},$$
$$\vec{\nabla} \cdot \vec{B} = 0, \qquad \vec{\nabla} \times \vec{B} = \mu\epsilon \frac{\partial \vec{E}}{\partial t}.$$

If we take the curl of the second of these, we get

$$\begin{aligned} \vec{\nabla} \times \left(\vec{\nabla} \times \vec{E} \right) &= \vec{\nabla} \left(\vec{\nabla} \cdot \vec{E} \right) - \nabla^2 \vec{E} \\ &= -\vec{\nabla} \times \frac{\partial \vec{B}}{\partial t}. \end{aligned}$$

Since $\vec{\nabla} \cdot \vec{E} = 0$ and we can change the order of partial derivatives, this gives

$$-\nabla^2 \vec{E} = -\frac{\partial}{\partial t} \vec{\nabla} \times \vec{B}$$

and then using the last of Maxwell's equations results in

$$\nabla^2 \vec{E} = \mu\epsilon \frac{\partial^2 \vec{E}}{\partial t^2}$$

which we all (hopefully) recognise as a wave eqution! This means that the three components of the electric field all propagate as waves. Furthermore, recall that the constant multiplying the second time derivative is the inverse-square of the speed of propagation, i.e. the field travels through the medium with speed $v = 1/\sqrt{\mu\epsilon}$. Taking the curl of the last equation shows that the magnetic field also satisfies the same equation with the same propagation speed.

Let's look at this speed: we cam rewrite it as

$$\begin{aligned} v &= \frac{1}{\sqrt{\mu\epsilon}} \\ &= \sqrt{\frac{\mu_0\epsilon_0}{\mu\epsilon}}\frac{1}{\sqrt{\mu_0\epsilon_0}} \\ &= \frac{c}{n} \end{aligned}$$

where $c = 1/\sqrt{\mu_0\epsilon_0}$ is the speed of light and the quantity $n = \sqrt{\mu\epsilon/\mu_0\epsilon_0}$ is the *index of refraction* of the medium under consideration. For vacuum, $n = 1$ and for all other media, $n > 1$ and thus electromagnetic (EM) waves – light – travel slower than the speed of light. This is due to the bound charges: as the wave moves through the medium, it interacts with the constituents, and thus sort of "bumps into" all the atoms and molecules, thus slowing it down. (Think about trying to get to the bar to buy a round when the pub's full compared to when it's empty.)

Now, you've all taken MP205 or the equivalent, and so we already know what sort of solutions we get to the wave equation. The particular solution we're interested are plane wave solutions, which we know are characterised by a frequency ω and a wave vector $\vec{k}$. The wave vector points in the direction of propagation and its magnitude is given by the dispersion relation $\omega = |\vec{k}|v$, which for the EM waves we are looking at gives $k = |\vec{k}| = n\omega/c$. (Or alternatively, the wavelength λ is $2\pi/k = 2\pi c/n\omega$, a factor of $1/n$ shorter than the vacuum wavelength.)

The general form for a plane wave solution is $A\cos(\vec{k}\cdot\vec{r} - \omega t + \alpha)$, where A and α are the amplitude and phase of the wave. However, we know that we can write this as $\mathrm{Re}[\tilde{A}e^{i(\vec{k}\cdot\vec{r}-\omega t)}]$ where $\tilde{A} = Ae^{i\alpha}$ is the wave's complex amplitude. Thus, the general form for the solution to the wave equation for the electric field may be written as

$$\vec{E}\,(t, \vec{r}) = \mathrm{Re}\left[\tilde{\vec{E}}_0 e^{i(\vec{k}\cdot\vec{r}-\omega t)}\right]$$

where $\tilde{\vec{E}}_0$ is a constant complex-valued vector. We can also do the same for the magnetic field:

$$\vec{B}\,(t, \vec{r}) \;\; = \;\; \mathrm{Re}\left[\tilde{\vec{B}}_0 e^{i(\vec{k}\cdot\vec{r}-\omega t)}\right]$$

where $\tilde{\vec{B}}_0$ is also a constant complex-valued vector.

Using these forms is incredibly convenient because each time derivative of the complex exponential simply multiplies it by $-i\omega$ and every application of a gradient multiplies it by $i\vec{k}$; by using these tricks with the wave equation, we get $-(\vec{k}\cdot\vec{k})\vec{E} = -\mu\epsilon\omega^2\vec{E}$ (and the same for $\vec{B}$) which just reproduces the dispersion relation. More useful is looking not at the wave equation but at Maxwell's equations. For example, if we use the above form for the electric field, we see

$$\begin{aligned}\frac{\partial \vec{B}}{\partial t} &= -\vec{\nabla}\times\vec{E} \\ &= -\mathrm{Re}\left[i\vec{k}\times\tilde{\vec{E}}_0 e^{i(\vec{k}\cdot\vec{r}-\omega t)}\right]\end{aligned}$$

which says the magnetic field must also be a plane wave with the same wave vector and frequency as the electric field (a fact that was *not* immediate from the wave equation alone). This means that

$$\frac{\partial \vec{B}}{\partial t} \;\; = \;\; \mathrm{Re}\left[-i\omega\times\tilde{\vec{B}}_0 e^{i(\vec{k}\cdot\vec{r}-\omega t)}\right]$$

and thus we see that the amplitudes of the two fields are related by

$$\tilde{\vec{B}}_0 = \frac{\vec{k}}{\omega}\times\tilde{\vec{E}}_0 \quad \Rightarrow \quad \vec{B}(t,\vec{r}) = \frac{\vec{k}}{\omega}\times\vec{E}(t,\vec{r})$$

(since $\vec{k}/\omega$ is real) and thus the magnetic field is perpendicular to both $\vec{k}$ and $\vec{E}$.

If we use this trick with the equation $\vec{\nabla}\cdot\vec{E} = 0$, we see

$$\mathrm{Re}\left[i\vec{k}\cdot\tilde{\vec{E}}_0 e^{i\vec{k}\cdot\vec{r}=\omega t)}\right] = 0 \quad \Rightarrow \quad \vec{k}\cdot\vec{E}(t,\vec{r}) = 0,$$

namely, the wave vector and the electric field are perpendicular. So the upshot is that for EM plane waves, the electric field, the magnetic field and

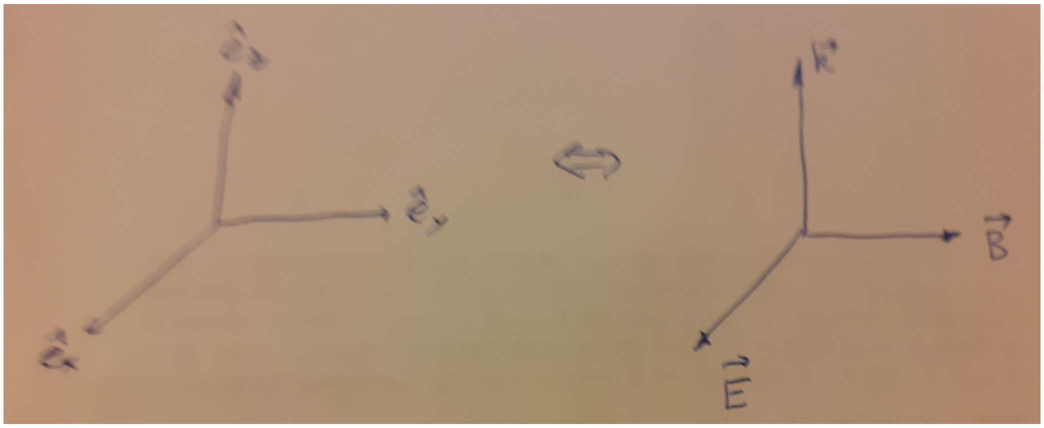

the wave vector form a right-handed triad: $(\vec{E}, \vec{B}, \vec{k})$ have the same relation as $(\hat{e}_x, \hat{e}_y, \hat{e}_z)$, as shown above.

Now, there are two more equations, but they're automatically satisfied: $\vec{\nabla} \cdot \vec{B} = 0$ implies $\vec{k} \cdot \vec{B} = 0$, but we already knew this from $\vec{B} = \vec{k} \times \vec{E}/\omega$. And $\vec{\nabla} \times \vec{B} = \mu\epsilon \partial \vec{E}/\partial t$ gives $\vec{k} \times \vec{B} = -\mu\epsilon\omega\vec{E}$, but notice that

$$\begin{aligned} \vec{k} \times \vec{B} &= \vec{k} \times \left(\frac{\vec{k}}{\omega} \times \vec{E} \right) \\ &= \frac{1}{\omega} \left[\vec{k} \left(\vec{k} \cdot \vec{E} \right) - \vec{E} (\vec{k} \cdot \vec{k}) \right] \\ &= -\frac{k^2}{\omega} \vec{E} \end{aligned}$$

which, since $k^2 = \mu\epsilon\omega^2$, gives full agreement.

So, to summarise, if we have a linear medium with no free sources, then there exist monochromatic plane wave solutions of frequency ω, with the electric field given by

$$\vec{E}\left(t, \vec{r}\right) = \mathrm{Re}\left[\tilde{\vec{E}}_0 e^{i(\vec{k} \cdot \vec{r} - \omega t)} \right]$$

for some complex vector $\tilde{\vec{E}}_0$, where $\vec{k}$ points in the direction of propagation and $k = n\omega/c$, with $n = \sqrt{\mu\epsilon/\mu_0\epsilon_0}$ being the medium's index of refraction. The magnetic field is then

$$\vec{B}\left(t, \vec{r}\right) = \mathrm{Re}\left[\tilde{\vec{B}}_0 e^{i(\vec{k} \cdot \vec{r} - \omega t)} \right]$$

where $\tilde{\vec{B}}_0 = \vec{k} \times \tilde{\vec{E}}_0/\omega$ (and so $\vec{B} = \vec{k} \times \vec{E}/\omega$).

Now, a bit of a comment on this index of refraction thingie: recall that when we indroduced the permeability, it was defined (for nonferromagnetic media) via the magnetic susceptibility by $\mu = \mu_0(1 + \chi_m)$. However, any glance at a list of susceptibilities will quickly show you that they're all generally quite small, with the largest having magnitude of the order of 10^{-4}, Therefore, for almost all materials, $\mu \approx \mu_0$ and we will often (but not always) be able to make this approximation. If we do so, then we see that the index of refraction is just the square root of the material's relative permittivity, $n \approx \sqrt{\epsilon/\epsilon_0}$ or alternatively, $\epsilon \approx n^2\epsilon_0$.

2. Polarisation of EM Plane Waves

Now, you may think we've pretty much figured everything out about EM plane waves in a linear medium, but we haven't. Notice that we still have some freedom to pick some of the quantities involved: the frequency ω, the direction of propagation $\hat{k} = \vec{k}/k$ and the complex electric amplitude vector $\tilde{\vec{E}}_0$. (Note that $k = n\omega/c$ is fixed once we pick ω, provided we know the index of refraction of the medium). Now, we know that $\vec{k}$ and $\tilde{\vec{E}}_0$ are perpendicular, so once the diorection of $\vec{k}$ is picked, then $\vec{k} \cdot \tilde{\vec{E}}_0 = 0$ puts some constraints on the electric amplitude vector.

For example, if we pick a coordinate system such that the positive z-direction is the same as the wave's propagation direction, then $\vec{k} = n\omega\hat{e}_z/c$ and $\tilde{\vec{E}}_0$ must lie purely in the xy-plane, so $\tilde{\vec{E}}_0 = \mathcal{E}_x\hat{e}_x + \mathcal{E}_y\hat{e}_y$, where $\mathcal{E}_x$ and $\mathcal{E}_y$ are complex numbers, so we still have four real numbers that we can pick freely. (Note that once these are picked, the magnetic amplitude vector is

$$\begin{aligned} \tilde{\vec{B}}_0 &= \frac{\vec{k}}{\omega} \times \tilde{\vec{E}}_0 \\ &= \frac{n\hat{e}_z}{c} \times (\mathcal{E}_x\hat{e}_x + \mathcal{E}_y\hat{e}_y) \\ &= -\frac{n\mathcal{E}_y}{c}\hat{e}_x + \frac{n\mathcal{E}_x}{c}\hat{e}_x \end{aligned}$$

which is completely determined.) If we choose to specify both of these complex numbers in polar notation, namely

$$\mathcal{E}_x = |\mathcal{E}_x|\, e^{i\alpha_x}, \qquad \mathcal{E}_y = |\mathcal{E}_y|\, e^{i\alpha_y},$$

then we see the elecric field is

$$\begin{aligned}\vec{E} &= \mathrm{Re}\left[\left(\mathcal{E}_x\hat{e}_x + \mathcal{E}_y\hat{e}_y\right)e^{i(\vec{k}\cdot\vec{r}-\omega t)}\right] \\ &= |\mathcal{E}_x|\cos\left(kz-\omega t+\alpha_x\right)\hat{e}_x + |\mathcal{E}_y|\cos\left(kz-\omega t+\alpha_y\right)\hat{e}_y\end{aligned}$$

Now, the choice of one of the phases is somewhat arbitrary because it can always be eliminated by choosing to start our measurement of time such that one of the components is at the beginning of its cycle. For example, if the time we measure is $t' = t - \alpha_x/\omega$, then the x-component depends on $kz - \omega t'$, so $t' = 0$ corresponds to the choice that this component is a standard cosine $\cos kz$ in space. The argument of the y-component is then $kz - \omega t' - \alpha_x + \alpha_y$, so in reality, the wave depends not on the two phases separately, but on their difference $\delta = \alpha_x - \alpha_y$. Thus, without too much loss of generality, we can take the standard form for our electric field to be

$$\vec{E} = |\mathcal{E}_x|\cos\left(kz-\omega t\right)\hat{e}_x + |\mathcal{E}_y|\cos\left(kz-\omega t-\delta\right)\hat{e}_y$$

and thus to specify any given EM wave, we need to specify the x-amplitude $\mathcal{E}_x$, the y-amplitude $\mathcal{E}_y$ and the relative phase difference δ (taken between $-\pi$ and π). Any such choice is called a *polarisation* of the wave.

To get a rough idea of what such a choice entails, look at the plot on the next page; the two magnitudes specify the amplitude of each of the cosinusoidal oscillations, and the phase difference gives an idea of which component starts its cycle earlier. The plot above is for $\delta > 0$, and we see that the x-component "leads" the y-component: it starts at its peak value of $|\mathcal{E}_x|$, whereas the y-component hasn't yet reached $|\mathcal{E}_y|$ and so is said to "lag" the x-component. For $\delta < 0$, the situations are swapped, and for $\delta = 0$, the two are "in phase". For $\delta = \pm\pi$, one reaches its maximum value when the other reaches its minimum and the components are "completely out of phase".

In fact, the $\delta = 0$ and $\delta = \pi$ cases are special enough to be looked at a bit more: if $\delta = 0$, then (if we take the convention $\alpha_x = 0$ as described above) $E_x = |\mathcal{E}_x|\cos(kz - \omega t)$ and $E_y|\mathcal{E}_y|\cos(kz - \omega t)$, or $E_y = |\mathcal{E}_y/\mathcal{E}_x|E_x$. In other words, the y-component is positively proportional to the x-component. This means if we plot E_y as a function of E_x, we get a line segment with positive slope with endpoints $(|\mathcal{E}_x|, |\mathcal{E}_y|)$ and $(-|\mathcal{E}_x|, -|\mathcal{E}_y|)$, and thus the tip of the vector $\vec{E}$ oscillates back and forth on this line.

Something similar happens if the two components are completely out of phase, $\delta = \pi$. Again, $E_x = |\mathcal{E}_x|\cos(kz - \omega t)$, but now $E_y = |\mathcal{E}_y|\cos(kz -$

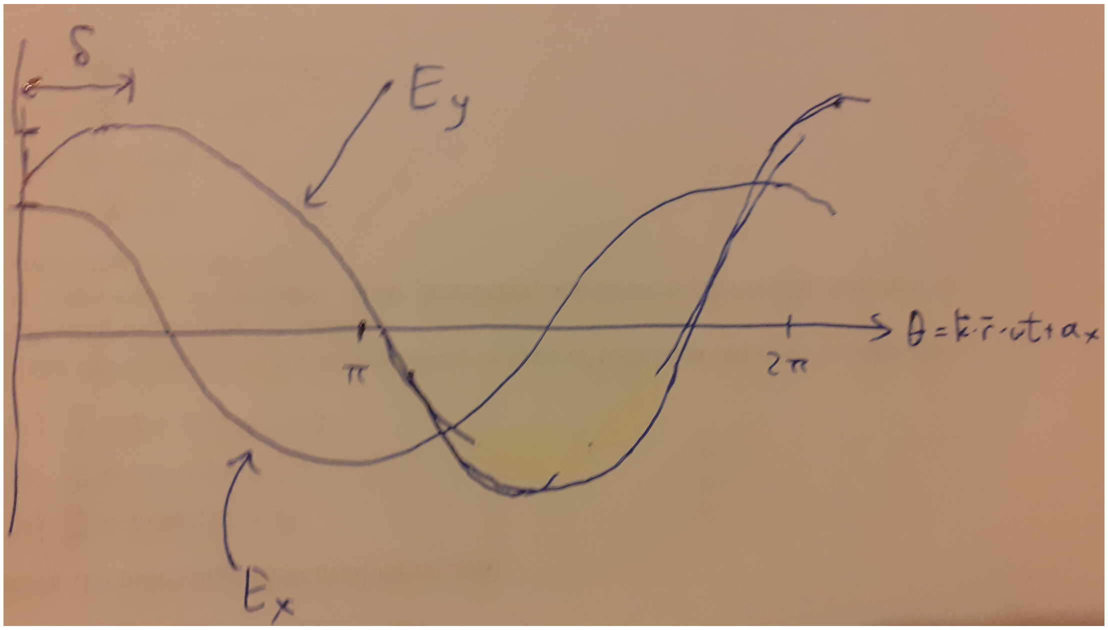

$\omega t - \pi) = -|\mathcal{E}_y|\cos(kz - \omega t)$. This gives $E_y = -|\mathcal{E}_y/\mathcal{E}_x|E_x$, a line segment of negative slope with endpoints $(|\mathcal{E}_x|, -|\mathcal{E}_y|)$ and $(-|\mathcal{E}_x|, |\mathcal{E}_y|)$. But despite the difference in sign, the upshot is that same as for the $\delta = 0$ case: the tip of the vector $\vec{E}$ always stays on a line. Because of this behaviour, we say that any EM wave with either $\delta = 0$ or $\delta = \pi$ is *linearly polarised*.

Another class of polarisation occurs when the two components have the same magnitude ($|\mathcal{E}_x| = |\mathcal{E}_y|$) and the phase difference is either $\pi/2$ or $-\pi/2$. If $\delta = \pi/2$, then, as always, $E_x = |\mathcal{E}_x|\cos(kz - \omega t)$ but $E_y = |\mathcal{E}_y|\cos(kz - \omega t - \pi/2) = |\mathcal{E}_x|\sin(kz - \omega t)$. Or if we let $\theta = kz - \omega t$, $E_x = |\mathcal{E}_x|\cos\theta$ and $E_y = |\mathcal{E}_x|\sin\theta$. Thus, as θ increases (say, if we fix a time and move spatially in the positive z-direction), the tip of $\vec{E}$ traces out a *circle* of radius $|\mathcal{E}_x|$ in an anticlockwise direction. Such an EM wave is said to be *right-circularly polarised*, because the tip of the vector rotates in the same direction as your fingers would wrap if you pointed the thumb of your right hand in the wave's propagation direction.

You can probably guess the next case: if $|\mathcal{E}_x| = |\mathcal{E}_y|$ and $\delta = -\pi/2$, $E_y = |\mathcal{E}_y|\cos(kz - \omega t + \pi/2) = -|\mathcal{E}_x|\sin(kz - \omega t)$, so since $E_x = |\mathcal{E}_x|\cos(kz - \omega t)$, the tip of the electric field traces out a circle of radius $|\mathcal{E}_x|$ in a *clockwise* direction with increasing θ, and so we denote it as a *left-circularly polarised* EM wave.

These are just special cases, of course. In general, the three parameters we need – $|\mathcal{E}_x|$, $|\mathcal{E}_y|$ and δ – could be just about anything. In the next problem set, I'll have you look at the general case and ask you to explain why we describe an EM wave that's not one of the special cases above as *elliptically polarised*. But now, we move onto other important properties of an EM wave...

3. Energy and Momentum of EM Plane Waves

We know an EM field carries both energy and momentum: the expressions giving the energy density and energy current forr a linear medium are, respectively,

$$\begin{aligned} u &= \frac{1}{2}\left(\vec{E}\cdot\vec{D}+\vec{B}\cdot\vec{H}\right) = \frac{\epsilon}{2}|\vec{E}|^2 + \frac{1}{2\mu}|\vec{B}|^2, \\ \vec{S} &= \vec{E}\times\vec{H} = \frac{1}{\mu}\vec{E}\times\vec{B}. \end{aligned}$$

(We also know from Problem Set 3 that $\mu\epsilon\vec{S}$ gives us the momentum density.) These are for a general EM field; what about the specific case of an EM plane wave?

Let's look at the energy density first: we know that, for an EM plane wave, $\vec{B} = \vec{k}\times\vec{E}/\omega$. Recall that $|\vec{a}\times\vec{b}|$ is $|\vec{a}||\vec{b}|$ times the sine of the angle between the two vectors. But since $\vec{k}$ and $\vec{E}$ are perpendicular, $|\vec{B}| = k|\vec{E}|/\omega$. Thus,

$$\begin{aligned} u &= \frac{\epsilon}{2}|\vec{E}|^2 + \frac{k^2}{2\mu\omega^2}|\vec{E}|^2 \\ &= \epsilon|\vec{E}|^2 \end{aligned}$$

since $k/\omega = \sqrt{\mu\epsilon}$. Now, recall that we have the general form

$$\vec{E} = |\mathcal{E}_x|\cos\left(kz-\omega t\right)\hat{e}_x + |\mathcal{E}_y|\cos\left(kz-\omega t-\delta\right)\hat{e}_y$$

and so

$$|\vec{E}|^2 = |\mathcal{E}_x|^2\cos^2\left(kz-\omega t\right) + |\mathcal{E}_y|^2\cos^2\left(kz-\omega t-\delta\right).$$

Thus, the energy density in its full time- and space-dependent form is

$$u(t,x,y,z) = \epsilon\,|\mathcal{E}_x|^2\cos^2\left(kz-\omega t\right) + \epsilon\,|\mathcal{E}_y|^2\cos^2\left(kz-\omega t-\delta\right).$$

This is correct, but in many practical cases, unnecessary. Why? Well, consider the frequencies that might show up in a typical EM wave. For example, radio waves: as anyone with a tuner knows, AM radio frequencies are in the tens or hundreds of kilohertz region and FM radio in the megahertz. And these are on the *low* end of the EM spectrum: visible light has even higher frequencies, and X-rays and gamma rays higher still. So most of the EM radiation you encounter in both everyday life and much of physics tends to oscillate at least 10000 times per second.

This means that, in general, we don't see the oscillatory nature of the wave, but rather its *average* value since it bounces all over the place so quickly. So if we measured, say, u, the actual time the measurement takes could be much longer than the period of oscillation of the wave, and thus the value we'd get would be the *time-averaged* value

$$\langle u \rangle(\vec{r}) \quad = \quad \frac{1}{T} \int_0^T u(t, \vec{r}) \, \mathrm{d}t$$

where T is the period of the oscillation.

So maybe we better look at how we might compute such an average: if we have only a single function, $f(t)$, then we just compute the time-average with the above formula. But note that u and $\vec{S}$ both are made out of the product of two oscillating functions: the forms E^2, B^2 and EB all appear. So suppose we have two functions $f(t)$ and $g(t)$ which oscillate with frequency ω. Then we know there exist complex numbers $\tilde{f}_0$ and $\tilde{g}_0$ such that

$$f(t) = \mathrm{Re}\left[\tilde{f}_0 e^{-i\omega t}\right], \qquad g(t) = \mathrm{Re}\left[\tilde{g}_0 e^{-i\omega t}\right]$$

If $\tilde{f}_0 = f_R + i f_I$ and $\tilde{g}_0 = g_r + i g_I$, it's easy to show that

$$f(t) = f_R \cos\omega t + f_I \sin\omega t, \qquad g(t) = g_R \cos\omega t + g_I \sin\omega t$$

and so

$$f(t)g(t) \quad = \quad = f_R g_R \cos^2\omega t + f_I g_I \sin^2\omega t + (f_R g_I + f_I g_R) \sin\omega t \cos\omega t.$$

The integral from 0 to $T = 2\pi/\omega$ of both $\cos^2\omega t$ and $\sin^2\omega t$ is $T/2$, and of

$\sin \omega t \cos \omega t$ is zero, so we see

$$\begin{aligned} \langle f(t) g(t) \rangle &= = \frac{f_R g_R + f_I g_I}{2} \\ &= \frac{1}{2} \text{Re} \left[(f_R g_R + f_I g_I) + i \left(f_I g_R - f_R g_I \right) \right] \\ &= \frac{1}{2} \text{Re} \left[\tilde{f}_0 \tilde{g}_0^* \right]. \end{aligned}$$

And that's the idea: if we have two functions with the same frequency, then the time average is just half the real part of the complex amplitude of one times the conjugate of the complex amplitude of the other.

So we know that since $u = \epsilon |\vec{E}|^2$, the time-average of this will require us to compute E_x^2. But we know that $E_x = \text{Re}[\mathcal{E}_x e^{i(kz - \omega t)}]$, so $\mathcal{E}_x e^{ikz}$ is what's multiplying $e^{-i\omega t}$, and thus from the above formula,

$$\begin{aligned} \left\langle E_x^2 \right\rangle &= \frac{1}{2} \text{Re} \left[\left(\mathcal{E}_x e^{ikz} \right) \left(\mathcal{E}_x e^{ikz} \right)^* \right] \\ &= \frac{1}{2} |\mathcal{E}_x|^2. \end{aligned}$$

Therefore, since we'll get a similar result for the averge of E_y^2 and $E_z = 0$,

$$\begin{aligned} \langle u \rangle &= \epsilon \left\langle E_x^2 + E_y^2 + E_z^2 \right\rangle \\ &= \frac{\epsilon}{2} \left(|\mathcal{E}_z|^2 + |\mathcal{E}_y|^2 \right) \\ &= \frac{\epsilon}{2} \tilde{\vec{E}}_0 \cdot \tilde{\vec{E}}_0^* \\ &= \frac{\epsilon}{2} |\tilde{\vec{E}}_0|^2 \end{aligned}$$

where we have now extended the idea of the norm of a vector to include complex-valued vectors: $|\vec{a}|^2$ is now defined to be $\vec{a} \cdot \vec{a}^*$, not $\vec{a} \cdot \vec{a}$.

So if we have the complex amplitude vector $\tilde{\vec{E}}_0$ – which we assume we do – we can easily and immediately compute the time-averaged energy density. Notice that it's *constant*; it depends neither on time nor on position, so the average energy of an EM plane wave is distributed uniformly thoughout space.

Now, on to the energy current: the Poynting vector is $\vec{S} = \vec{E} \times \vec{B} / \mu$, so for an EM plane wave, this gives

$$\begin{aligned}\vec{S} &= \frac{1}{\mu}\vec{E}\times\left(\frac{\vec{k}}{\omega}\times\vec{E}\right) \\ &= \frac{1}{\mu\omega}\left[\vec{k}\left(\vec{E}\cdot\vec{E}\right)-\vec{E}\left(\vec{k}\cdot\vec{E}\right)\right] \\ &= \frac{\vec{k}}{\mu\omega}|\vec{E}|^2.\end{aligned}$$

Unsurprisingly, this points (poynts?) in the same direction as the wave is propagating, so energy is being carried along with the wave.

But for the same reasons as discussed above, any measurement of this vector is likely to return an average value, and so what we expect to get is

$$\begin{aligned}\langle\vec{S}\rangle &= \frac{\vec{k}}{\mu\omega}\left\langle|\vec{E}|^2\right\rangle \\ &= \frac{\vec{k}}{2\mu\omega}|\tilde{\vec{E}}_0|^2\end{aligned}$$

for the same reasons as before. Again, the complex amplitude vector contains all the info we need to get the average energy current of an EM plane wave. Also, the average density of momentum is also obtained from this: $\langle\mu\epsilon\vec{S}\rangle = \epsilon\vec{k}|\tilde{\vec{E}}_0|^2/2\omega$.

Now, $\vec{S}$ is the energy current density in much the same way that $\vec{J}$ is the charge current density: if we have a small surface area element $\mathrm{d}\vec{\sigma}$, the $\mathrm{d}P = \vec{S}\cdot\mathrm{d}\vec{\sigma}$ is the little bit of power – the rate of energy change – that flows through it.

This is true in general, but now let's specify to the case of an EM plane wave hitting a surface. If we denote by $\mathrm{d}\bar{P}$ the average power flowing across this surface (we could use $\mathrm{d}\langle P\rangle$ or $\langle\mathrm{d}P\rangle$ to be consistent, but both just look weird), then

$$\begin{aligned}\mathrm{d}\bar{P} &= \langle\vec{S}\rangle\cdot\mathrm{d}\vec{\sigma} \\ &= \frac{|\tilde{\vec{E}}_0|^2}{2\mu\omega}\vec{k}\cdot\mathrm{d}\vec{\sigma} \\ &= \frac{k|\tilde{\vec{E}}_0|^2}{2\mu\omega}\mathrm{d}\sigma\,\cos\theta\end{aligned}$$

where θ is the angle between the wave's propagation vector and the surface's normal. Note that for a fixed area $\mathrm{d}\sigma$, this depends only on the angle θ; thus, it's largest when $\theta = 0$ and smallest when $\theta = \pi/2$.

This explains the temperature variation during the seasons: around June 21$^{\text{st}}$ of each year, the Earth's axis is tilted toward the Sun, and therefore the northern hemisphere receives the light from the Sun more directly (smaller θ, between $-23.4°$ at the Equator to $66.6°$ at the North Pole) whereas the southern hemisphere less so (from $-23.4°$ to $-90°$ at the Antarctic Circle), and so it's warm(ish) in Dublin and cool(ish) in Sydney. Half a year later, around December 21$^{\text{st}}$, the Earth's axis tilts away from the Sun and the situation reverses. You probably all knew that, but now we have the physics and maths to explain it!