PHYS 2212   Lab Exercise 03: Gauss' Law

**PRELIMINARY MATERIAL TO BE READ BEFORE LAB PERIOD**

Counting Field Lines: Electric Flux

Recall that an electric field (or, for that matter, magnetic field) can be difficult to *visualize*, since the only thing we can see is its *effect* on charges. To aid in the visualization process, we introduce the idea of a "**Field Map**" which represents the field in terms of a set of continuous, directed "lines of force" or "field lines (*FL*)". For electric fields in particular, the basic rules for setting up the correct field map for a given collection of point source charges are as follows:

1. **Field lines always begin on positive charges and end on negative charges. The *number* of field lines so constructed should be proportional to the magnitude of the charge.**

2. **In any particular region of space, the *magnitudes* of the electric field *vectors* are directly proportional to the "*density*" of field *lines* in that particular region.**
    (Note that "field lines" and "field vectors" mean *different* things…)

There are, of course, other rules as well, but we are not interested in them here. Refer back to the relevant section of the textbook if you're interested. For present purposes, we simply note that the two rules above essentially describe a *book-keeping process*, where the counting of field lines conveys useful information, albeit of a *qualitative* nature.

If we wish to be more precise and quantitative, we almost immediately run into some trouble, and the name of that trouble is "*artistic license*": since a field map is mainly meant to convey an idea of the "shape" of the field, there are no hard-and-fast rules about *how many* field lines must be drawn in a particular circumstance. Put in other words, an individual needs to draw only as many field lines as he or she feels necessary in order to get a "feel" for the field.
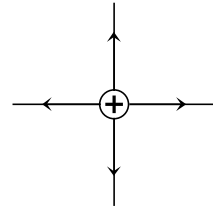

For example, if the field of a +1.0 nC point charge were being drawn, a physics prof. or a math expert might be able to get away with drawing a field map using only four field lines, as shown above.

On the other hand, a less experienced Tech student might need to draw twice as many field lines (below left) before he or she feels "comfortable" in understanding what his or her drawing represents. Of course, it doesn't end there. A really super-duper clueless person (such as a UGA student) might need to draw a *lot* more field lines before reaching any sort of comfort level with his or her field map (below right).
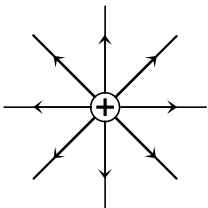
Who's **right**? All of them, of course. As long as each one sticks with their chosen "# Field Lines per unit charge" ratio (i.e. 4 *FL* per nC for the prof., 8 *FL* per nC for the Techie, and fifty-gajillion *FL* per nC for the UGA student), they will all be able to construct valid field maps in subsequent sketches, and be able to draw qualitative—*and even some crude **quantitative**—*information from their sketches.
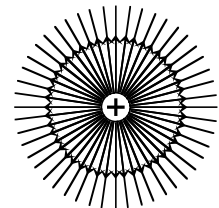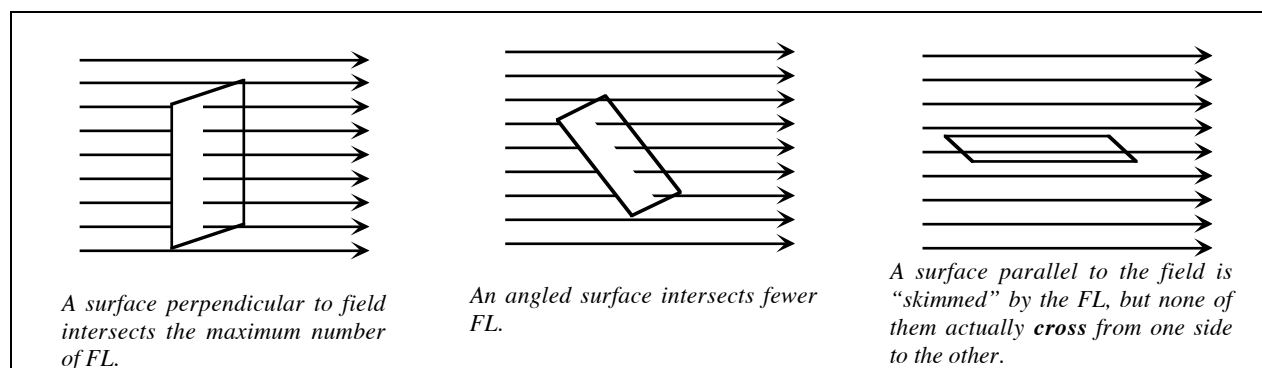

The problem, though, is that these sketches are not cross-compatible. It would be nice to have a way of "counting *FL*" that is not dependent upon the "scale" chosen for our field map. This is particularly important for the next stage of our analysis: drawing accurate *quantitative* information about field magnitude from such a map. To do so, we will need to develop a careful numerical measurement for "Number of *FL*" (which *everyone* can agree upon), and to establish a more rigorous definition of what we mean by "density of *FL*", in Rule 2, above.

We will begin by establishing the following meaning for the phrase "density of field lines":

The density of field lines is measured by counting the **number** of *FL* which **pass through** some surface of total area *A*, provided the area is oriented to be **perpendicular** to the field lines. We would refer to such a surface as a *cross-sectional* area, and denote it suggestively as: $A_\perp$.

It is important at this stage to choose a perpendicular surface for the purpose of counting—after all, if the surface is *parallel* to the field, *FL* do not cross *through* the surface; they merely skim along one side or the other. For a surface of given area *A*, the *maximum* number of *FL* crossings will occur when the surface is oriented to be perpendicular to the field—and that's what we want (at least, for now…).

A surface perpendicular to field intersects the maximum number of FL.

An angled surface intersects fewer FL.

*A surface parallel to the field is "skimmed" by the FL, but none of them actually **cross** from one side to the other.*

Note that it is *also* possible to deal with non-uniform fields, in which the field lines *curve*, simply by allowing the surface in question to curve as well—in such a way that at every point on the surface, *FL* cross at right angles. In

*Surface is constructed in such a way that it is **perpendicular** to every FL which intersects it.*

this situation, the "cross-sectional area" is not that of a flat surface—so we can't simply write it as "*length × width*" or as "$\pi \times (radius)^2$", but we **can** still represent it by the symbol "$A_\perp$".
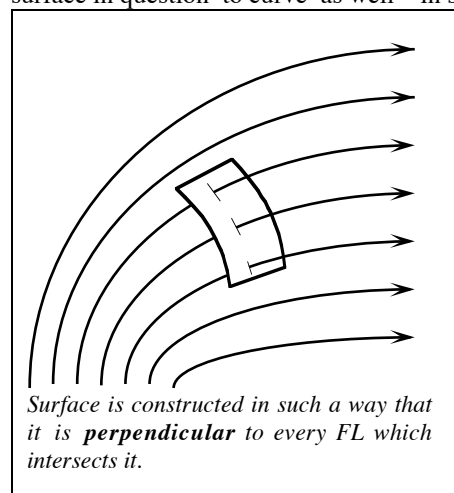
Now, presuming that we can count the number of *FL* which cross the surface, we can express the "density of *FL*" specifically as being the *surface* density of *FL* (i.e. number of *FL per unit area*):

$$\text{Density of } FL = \frac{\#\,\text{FL}}{A_\perp}$$

Of course, this definition is still not "cross-compatible", because the number of *FL* in the numerator will depend upon how many *FL* we've chosen to draw in our field map. But keep in mind that the only reason we're really interested in computing the density of *FL* is to make a calculation of the field magnitude, by invoking Rule 2:

$$|\vec{E}| \propto \text{Density of } FL = \frac{\#\,\text{FL}}{A_\perp}$$

Remember that both $|\vec{E}|$ and $A_\perp$ are mathematically "well-defined quantities—it is the pesky "# *FL*" in the numerator that is a problem. So, let's just re-arrange the expression to put the "known stuff" together on one side:

$$\#\,FL \propto |\vec{E}| \cdot A_\perp$$

The right hand side of this expression is something that *everyone* can agree upon, regardless of how many field lines they have chosen to represent in their diagram. As a result, we will use this product as the basis for establishing a "scale-independent" measurement process for counting field lines that cross a surface. We will refer to this measured quantity as ***electric flux***, or simply flux. Never forget that,—at it's heart—*flux* is describing a book-keeping process, in which we are effectively *counting* the number of *FL* crossing a surface.

Flux: A Formal Definition

Having laid the conceptual groundwork for *why* we introduce the idea of flux, let's look more closely at a precise mathematical definition. Note, too, that the examples and "Tactics Boxes" in section 27.3 of the text are highly relevant (and informative) here, so you may want to devote some time to reviewing *them*, as well.

We start by recognizing that in all the preceding analysis, we were freely choosing our surface as we wished, in order to figure out the qualitative relationships involved. In general, however, we will often face circumstances where we are given a *particular* surface, and are asked to compute an expression for the flux through *that surface* (equivalent to the *number of field lines crossing the surface*). So, our "definition so far" should be viewed only as a "special case":

Flux through a surface perpendicular to the field:
Reviewing what we've done so far, suppose that we are *given* a particular surface such that at *every* point on the surface, the field is directed $\perp$ to the surface. Using the symbol "Φ" to represent flux, we can immediately write:

$$\Phi = |\vec{E}| \cdot A_\perp$$

This expression works because the electric field magnitude is equivalent to the "*FL* surface density" (by Rule 2), and so we are really just writing the mathematical equivalent of:

Total "stuff" = ("stuff" per unit area) x (total area under consideration)

Of course, just as with any other density function, we must be wary of situations where the density is varying over the region being considered. Such cases require us to use integration methods, and such will be the case here. In considering *flux*, this means that if the *field magnitude* varies over the surface, we must proceed by breaking the surface into subregions of area $\delta A_\perp$, computing an expression for the flux $\delta\Phi$ through a single, generic subregion, and then summing over all subregions. In the limit where our sub-regions become infinitesimal, $\delta A_\perp \rightarrow dA_\perp$, and our sum becomes an integral:

$$\Phi = \int d\Phi = \int |\vec{E}| \cdot dA_\perp$$

Admittedly, *writing out* this integral in its full, gory detail, and *evaluating* it can often be a messy process. For now, let's just recognize the *conceptual idea* that a varying electric field magnitude can be handled by breaking the surface into subregions, computing the flux through each subregion, and summing. The success of this process hinges on a crucial property of flux:

- **Flux is an *additive scalar* quantity. That is, Φ is a number which has no vector direction associated with it. Moreover, the total flux through two or more surfaces is the simple mathematical sum of the individual fluxes computed through each of the surfaces separately.**

This result can be justified simply by remembering that flux is effectively measuring a *number of things* (specifically, *number of FL crossing a region)*, and as such, it should combine the way that simple numbers do.

Note, too that the *units* of flux will be the units of (field) x (area) = (N/C)·(m^2) = N·m^2/C. This may seem a bit odd, when we recall that it is supposed to be representing *number of FL*. After all, no sensible person would ever say, "I counted how many field lines crossed this surface, and I found that number to be 4.23 x 10^{-9} N·m^2/C." Sounds goofy, huh? Just remember that technically, "flux" doesn't really *equal* "#FL", but is instead intended to be *proportional to #FL*, using a well-defined and <u>mutually agreed-upon</u> proportionality constant that *everyone* uses.
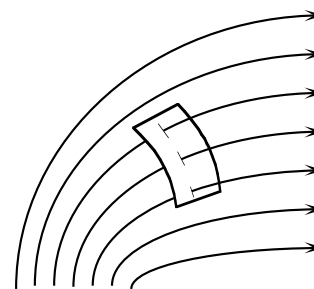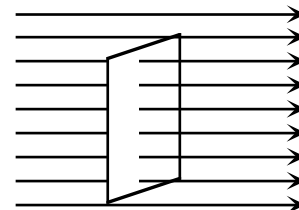
Flux for surfaces which are *not* ⊥ to the field:

Let's make it easy on ourselves, and assume that we are dealing with a *flat* surface. However, let's assume that the surface in question is *not* aligned in a convenient orientation. It should be clear, given our earlier work, that the flux through this surface should be *less than* it would be for the perpendicular surface, simply because the act of "tilting" the surface away from the perpendicular would effectively *reduce* the cross-sectional area of the surface, and decrease the number of *FL* that would cross it's reduced "footprint".
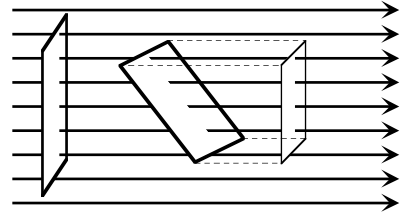
Before we try to analyze how the flux depends on the tilt angle, it will be useful to digress in order to establish a precise description of a surface's *orientation*—that way, we can be more precise about what we *mean* by "tilt angle". The conventional process is <u>not</u> to describe the *plane* of the surface (as we've done so far), but instead to specify the direction of a fixed vector *perpendicular* to the plane of the surface—the so-called "normal vector" for the surface.

*A surface and its normal, in various orientations*
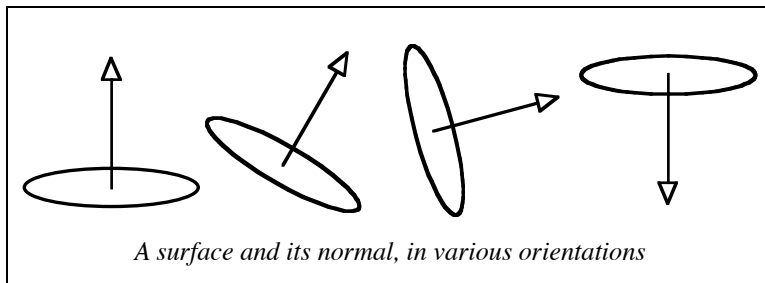
A given surface, no matter what its orientation, will always have a well-defined normal direction (two, actually), and by specifying the direction of that normal, we can unambiguously describe the orientation of the surface in space (at least, to within an overall 180° rotation, as the first and last sketches at left would suggest).

So, for a given surface, we will first establish a normal unit vector $\hat{n}$. [Making it a *unit* vector means that it is dimensionless (i.e. no units) and has a magnitude of exactly "1"—all this vector does is *point* in the direction normal to the surface.] The orientation of the *surface* relative to the field will then be fully specified if we describe the orientation of $\hat{n}$ relative to the field. In particular, let us presume that θ represents the angle between the field direction and the normal unit vector. (Be careful here! Not *all* θ's in *all* problems are actually the "angle between $\vec{E}$ and $\hat{n}$"—sometimes the symbol θ can be used to represent some *other* angle!) When θ = 0°, the normal unit vector is parallel to the field, and the *surface* is perpendicular to the field. This means that the *actual* area A of the surface is also the "cross-sectional" area $A_\perp$, and the problem reduces to the situation we've already worked:

$$\Phi = |\vec{E}| \cdot A_\perp$$

On the other hand, when the normal unit vector is *perpendicular* to the field (i.e. θ = 90°), we will find that the plane of the surface is parallel to the field, and the flux will be zero. This is the case seen before, where the field lines are only "skimming along" the surface, but not crossing from one side to the other:
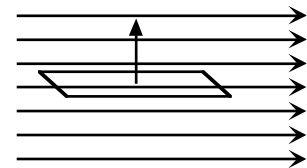
$$\Phi = 0$$
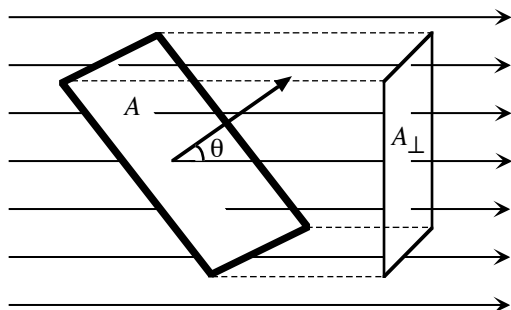
For angles between these two limits, the best way to proceed is to figure out what the "reduced footprint" of the surface would be, by finding its <u>projection</u> onto a plane ⊥ to the field (aka its cross-sectional projection). If A represents the full area, then the cross section will be $A_\perp$, and a little bit of basic geometry tells us that:

$$A_\perp = A \cos\theta$$

(Note how this expression satisfies the θ = 0° and θ = 90° limits…)

4

Thus, for a flat sheet whose normal makes an angle θ with the field direction, we can extend the applicability of our earlier definition by writing:

$$\Phi = |\vec{E}| \cdot A \cos\theta$$

Here, *A* is the actual area of the surface. Once again, we must recognize that if the electric field magnitude varies from point to point on the surface, this simple product must be replaced by an integral, in which the surface is broken into small subregions and the flux through each subregion is computed separately before adding the results to get the net flux.

Interlude—Area as a *vector*:
At this stage of things, it is useful to introduce a new notation that summarizes and simplifies what we've developed so far. This new notation will also make it easier for us to iron out a technique for solving "worst case" flux scenarios (i.e. cases where none of the simplifying assumptions we've use so far are valid).

Note that a flat surface has *two* distinguishing characteristics that we are interested in: its *area A* and its *orientation* (which can be described by specifying $\hat{n}$). We can combine these two properties into a *single* physical quantity, which we will call the **area vector** for the surface:

$$\vec{A} = A\,\hat{n}$$
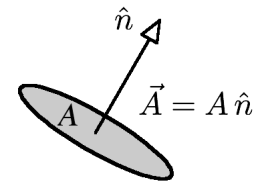
By specifying the magnitude (area) and direction (normal unit vector) of the area vector, we are simply making an "all in one" description of the surface. Of course, as far as $\hat{n}$ is concerned, all we are *really* interested in is the angle that it makes *relative* to the field direction, since our ultimate goal is to write an expression for flux. In terms of the area *vector*, the flux expression becomes:

$$\Phi = |\vec{E}|\,|\vec{A}|\cos\theta$$

We then recognize that this expression for flux can be expressed as a ***dot product*** of the two vectors $\vec{E}$ and $\vec{A}$:

$$\Phi = \vec{E} \cdot \vec{A}$$

Given this definition, we have to confront one "sticky" issue: dot products can yield *negative* values (if the angle between the two vectors is greater than 90°)—what does **negative flux** mean?

Before we try to answer this, let's back up for a bit. Near the middle of Page Four, we noted in passing that specifying the normal unit vector for a surface gave its orientation, <u>to within an overall 180° rotation</u>. (You may have puzzled about what was meant by this statement. You are about to find out…) It should be clear that, given some particular surface, there are always **two** possible choices for the direction of the normal unit vector—one for *each* of the perpendicular directions away from it, as shown at right. Which should we choose? It doesn't matter—**but once we've made the choice, we have to stick with it!** (We will see later that in certain circumstances, there *is* a convention for choosing the direction of $\hat{n}$.)
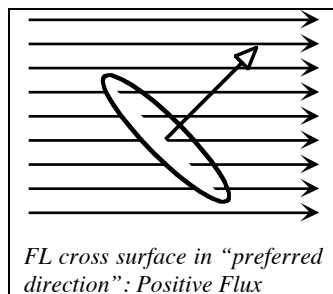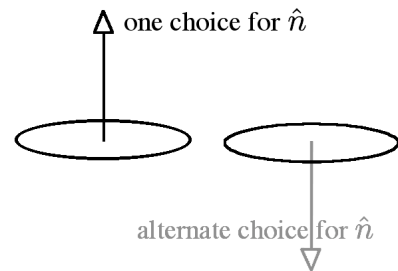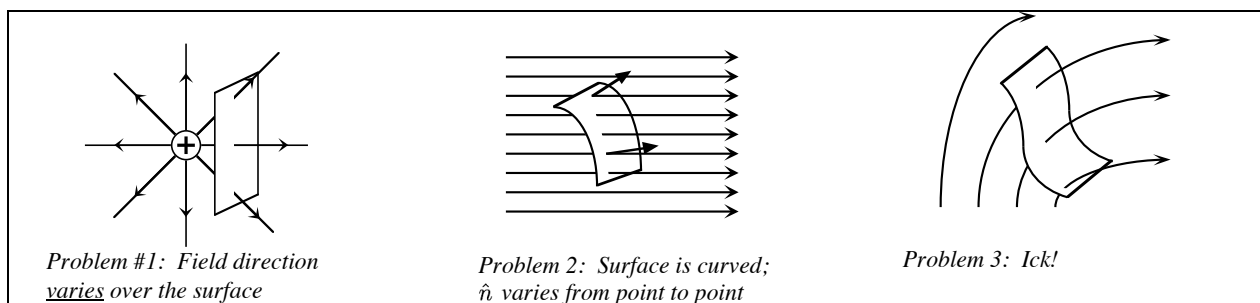
Note that in choosing a normal unit vector, we are establishing a "preferred direction" for crossing the surface: you can either cross "*with $\hat{n}$*" or "*against $\hat{n}$*". With that understanding, we should recognize that in the dot product, positive values occur when $\vec{E}$ crosses the surface in the preferred direction, while negative values occur when $\vec{E}$ crosses the surface in the opposite sense to the preferred direction.

FL cross surface in "preferred direction": Positive Flux

FL cross against the "preferred direction": Negative Flux
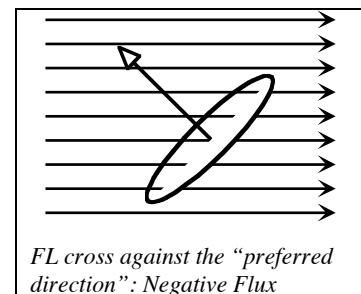
Worst case scenarios—advanced flux calculations:
Finally, we have to recognize that there are still *two* remaining situations that might "gum up" our work:

1. The electric field may vary in *direction* (as well as magnitude) over the surface in question. (An example of this might be a flat sheet in the neighborhood of a single point charge.)

2. The surface may *curve* in such a way that there is *no* single, well-defined "normal direction".

3. Both cases 1 *and* 2 may hold. (OK, that's *three* situations. So, sue me…)

Problem #1: Field direction varies over the surface

Problem 2: Surface is curved; $\hat{n}$ varies from point to point

Problem 3: Ick!

The nifty thing is: we have already worked out the tactic that will resolve all three of these problems! Recall that when the field *magnitude* varied from place to place, we fixed things by breaking the surface into tiny regions δA, on which the field magnitude was *locally* uniform. While we're at it, we might as well *also* require that these tiny subregions of the surface are so small that:

- The field *direction* is locally uniform. That is, the region is so small that the field direction doesn't *noticeably* change as we move around the subregion.

- The surface itself is *locally flat*. That is, we are zooming in so close that we don't *see* the curvature, and the tiny region under consideration has a well-defined *local* normal vector $\hat{n}$ and a local area δA.

Ultimately, we will be integrating, which means that the subregions will be shrunk down to infinitesimal size. When that happens we are able to satisfy *all* "tinyness conditions" at once—infinitesimal subregions are *always* locally flat, and *always* have a locally uniform field magnitude and a locally uniform field direction.

The process of shrinking to infinitesimal subregions, and converting the sum of the fluxes into a *flux integral*, will provide us with the most general possible definition of flux—but keep in mind that the process of evaluating the integral (which may be very easy, or very nasty, depending on circumstances) is still *really* just a mathematical way of **counting "number of field lines"**, in a manner that *everyone* can agree upon.

So, here's how we put it together to get the *most general* expression for flux. At a generic point on the surface, we imagine a tiny subregion which is pointlike, and find an expression for its area dA and the normal unit vector $\hat{n}$ to the surface *at that point*. This allows us to construct the area vector $d\vec{A} = dA\,\hat{n}$. We likewise find an expression for the electric field $\vec{E}$ at that point, and then write the local flux through that small subregion as:

$$d\Phi = \vec{E} \cdot d\vec{A}$$

Then the total flux is obtained by summing over all subregions, which becomes an integral over the surface:

$$\Phi_{tot} = \int d\Phi = \int \vec{E} \cdot d\vec{A}$$

As a final note: such *surface integration* may seem like a daunting task—and in it's most general application, it *is*. However, for *this* class we will most often deal with circumstances where the flux integral simplifies to something "trivially easy" (Sounds *too* good to be true, doesn't it?). At this stage, you would do well to review section 27.3 in more detail—particularly Tactics Box 27.1 (p. 859), which does a good job of itemizing those situations where the flux integral is "trivially easy".