Lecture 9—Estimating precipitation over an area

Ok, we ended last time with a statement of how data from a single rain gauge can be reported as either a cumulative mass curve (see) or a hyetograph (see). What I want to talk about today is how we get from these forms of data to what we really wanted, which was estimating the total rainfall on a watershed. To do this we need a few tools:

- We need to know how efficient rain gauges are at collecting all the rain that fell
- We need to know how to handle rain gauges that fail to perform during a rainfall
- We need to be able to incorporate changes to the gauge into the long-term history of the gauge
- We need to know how to incorporate data from several gauges into a unified whole

Last, we're going to talk about how to do an end-run around the entire problem (and incur a bunch of new ones for our trouble).

So, first thing:

RAIN GAUGE EFFICIENCY

Anybody here ever been to NYC on a windy day? You can turn a corner and be nearly blown off your feet. This is because the large buildings in NYC cause very strange wind patterns locally. Specifically, nature doesn't tend to like sharp edges—engineers do. Sharp edges often cause singularities in continuous phenomena, causing instabilities. This is certainly true with wind, where turbulent eddies are often shed off these edges. Take a good look at the rain gauge we have. See where I'm going with this? The sharp edges on the sides of the rain gauge make turbulent eddies that shed over the rain gauge itself. The updrafts retard rain trying to get in the gauge, and this effect is worse for snow (see the graph). A number of attempts have been made to shield the edge, but with limited success.

As a result, a number of *correction factors* have been introduced for rain gauges. Naturally, they're a function of the wind speed. For a standard NWS 8" rain gauge (like ours) the equation is:

$$K_{ru} = 100e^{\left(-4.605+0.062v^{0.58}\right)}$$

and for snow,

$$K_{su} = 100e^{\left(-4.606+0.157v^{1.28}\right)}$$

where *v* is in m/s. *K* is a constant that should be multiplied by the recorded value to get the estimated true rainfall.

MISSING DATA FROM RAIN GAUGES

What to do when furry beasts have ripped out one of your rain gauges? The simplest approach is just to find a gauge that was working, look at a long-term rain recording from that gauge and compare it to the one that's got missing data. You can interpolate from there. HOWEVER, this only works when the rainfall over the watershed tends to be uniform, and the watershed is of fairly low relief (it doesn't work for a thunderstorm in Spearfish, for instance).

In cases where you really need that estimate, here's probably the best way. You need at least three nearby working rain gauges. The formula commonly reported is:

$$P_x = \tfrac{1}{3}\left(\frac{N_X}{N_A}P_A + \frac{N_X}{N_B}P_B + \frac{N_X}{N_C}P_C\right)$$

However, multiple linear regression of the form:

$$P_X = a + b_A P_A + b_B P_B + b_C P_C$$

may be a better choice because it allows for weighting of the stations (to make ones closer count for more, for instance…)

CHECKING FOR CONSISTENCY OF A DATASET

What happens if we decide to: add an Alter shield to a rain gauge? a tree that used to be near our gauge falls over? replace a gauge rippled out by furry beasts? Worry about the growth of a city upwind of our gauge? All of these things may be considered point changes that will affect the long-term trends our gauge records. How to find out if this has been altered? The solution is typically to compare the long-term history of one gauge against an average for many gauges in the area. The two are then plotted on separate axes, and changes in the slope noted (see example). To correct for this regime change, data prior to the change should be altered by multiplying by the ratio of the two slopes (here 0.74/1.19).

ESTIMATING PRECIPITATION OVER A WATERSHED

Ok, so we've got six gauges in a watershed, and a bunch more outside it. How do we take the data from several point sources and amalgamate it into a single whole? Turns out there's several methods. Perhaps the easiest would be simply to average all the gauges in the watershed and be done with it. This is actually done. It tends to be done in areas without too many gauges, and in areas without large variability in precipitation from gauge to gauge.

Let's take a step up in sophistication. Suppose we could figure out, for any point on the watershed, which gauge was closest, and make each gauge responsible only for the area that it's the closest to? How would you do this, anyway? Let's see. Take a watershed with some gauges in it. Start swinging a circle around each gauge until it starts to impinge on another circle. Along that section, we'd draw a line dividing the two arcs. Right? Look familiar? This is EXACTLY how you determine the bisector of a line from way back in geometry. So, we could skip all this circle crap and cut to the chase. Draw a line between any two stations, take the bisector of that line and extend it until it hits someone else's bisector. What you're left with is a very odd shape that outlines the area that's closest to say, station A, in the watershed. These shapes are called *Thiessen polygons* after the proud inventor. To use them, just get the area and multiply by the rainfall at the gauge. Add up all the subarea totals and divide by the total watershed area, *et voilà*, you have the average precipitation over the watershed.

Ok. This seems a *little* hokey. Rain isn't going to fall in weighted polygons, is it? It's going to fall in a continuous smooth surface, like a topo map of rain, right? Hey! What if we make a topo map from the data we've got here? No different from surveying, really. We could just make the topo map, then do the exact same thing we did with the polygons—get the area between each contour and multiply by the average value between two contours (so between the 1mm and 2mm contours you'd multiply by 1.5). These are called *isohyetal diagrams*, and the individual contours are called *isohyets*. There is much bickering about how to contour, but suffice it to say it's done.

AN END-RUN AROUND THE WHOLE MESS

This seems like a lot of work, and it appears that we could get vastly different answers from the same set of data, depending on how we treat the data. Worse, we *still* don't have very good coverage over an area!

One day, *while smoking crack*, meteorologists pondered this problem. "Hey," they thought, "what if we could determine, *in real time*, the amount of rain *at every point*. That would make us look *so cool* on TV!" Thus was born the Doppler weather radar.

During WWII, in the early days of radar, it became apparent that precipitation made it difficult to see incoming aircraft because the radar waves were bouncing off the rain or snow, and not the implements of fiery destruction.

After WWII, military radar stations were turned over for use as thunderstorm monitors, but in 1988 a purpose-built radar network came online. This is the NEXRAD system, also known as WSR-88.

The NEXRAD system uses an empirically derived equation:

$$Z = aR^b$$

Where Z is the strength of the radar return, R is the rainfall rate, and a and b are constants. One common set of coefficients for a and b are 200 and 1.6, respectively, but 228 and 1.5 are often used, too. This causes trouble—take a look. The basic problem is that we have to tune using a set of poorly defined parameters. How to deal with this? That's right, kids—back to the rain gauges. Rain gauges are commonly used to "ground truth" radar data, and a correction added to the radar data (see example). Ever wonder why there was this sudden rush to have every elementary school on the planet get a rain gauge in the early 1990's? This is why.

So. In the US today, rain gauges play as vital a role as they ever did. Rain gauges are the ground truth at a point, and are fleshed out with a more precise algorithm than Thiessen polygons or even contouring—real-time radar data is now used in place of these methods.