

SAR Theory/Interpreting Images

(For a more general introduction, see ASF's Frequently Asked Questions about SAR page.)

When the radar signals interact with ground surfaces, they can either be reflected, scattered, absorbed, or transmitted (and refracted). Reflection is often due to a material's high dielectric constant, usually meaning a high water content. Very smooth surfaces also encourage reflection. If the imaged surface is a smooth lake, for example, the incoming radar will be reflected off the lake according to Snell's law - at the same angle as the incidence angle. You've probably seen the reflection of the setting sun in a smooth lake. Such reflections return very little signal strength back to the satellite, resulting in a dark region on an image. Reflections can, however, bounce again off other objects and thereby be redirected back toward the spacecraft, resulting in a stronger return signal. This process is called volumetric scattering (in contrast to surface scattering). It often occurs in vegetation, where (high water content) leaves reflect radar signals which then hit more "wet" leaves and branches, and so on until the signals exit the vegetation. A percentage of the incoming radar is therefore volumetrically scattered back to the spacecraft, giving the vegetation a brighter signature than the smooth lake. Sometimes SAR image interpreters have mistakenly believed that the radar signals were penetrating the vegetation and interacting with the underlying geology. This would only be possible if the vegetation were very dry and therefore did not reflect and, as will be described below, if the vegetation height were small relative to the signal wavelength. Notice that if the vegetation cover grows to a consistent height (for example, if all trees in a forest are relatively the same height), however, the radar return from the tops of the trees will mimic the undulations of the underlying topography (you can still tell it's a hill though it's covered with trees).

When the radar is transmitted through the surface, it will be refracted depending upon the density of the substance according to the index of refraction. The index of refraction equals the velocity of an electromagnetic wave in a vacuum divided by the velocity of an electromagnetic wave in the particular substance and is used to determine how a signal's characteristics will be altered when it passes into a different material. In our case, this means that the radar signals will travel more slowly in the surface material than they did in air, and therefore they will appear to "bend" (refract) toward the surface's normal. The transmitted, refracted signal can then either be absorbed by the material or have another surface interaction (reflection, transmitted, etc.) when it hits a region with different properties. For example radar is often transmitted through cold, dry snow. If the transmitted signal then hits a patch of smooth ice underneath the snow, it will reflect back up through the snow. It often happens that these reflected signals then hit small rocks or denser ice pockets in the overlying snow and bounce yet again, sometimes back toward the spacecraft. So then we could have a combination of transmission, reflection, and scattering - for both surfaces and volumes! You can see how SAR interpretation can become complex.

Surface scattering of incoming radar, determined by surface roughness, is often the main factor influencing radar return. Very smooth surfaces will reflect according to Snell's law, like the reflections off smooth water already discussed. Surface variations on the order magnitude of the radar wavelength will scatter radar as specified by the Rayleigh criterion. The Rayleigh criterion states that $h \cdot \cos(\text{look angle})$, the height of the surface variation in the direction of the incoming radar, must be less than 1/8th the value of the radar's wavelength to be considered "smooth." (The Rayleigh parameter h represents the ground's height variation, while the look angle is measured from nadir to the radar's direction of travel.) Researchers often use values finer than 1/8 (such as 1/25) to distinguish between smooth and rougher surfaces. The rougher the surface, the stronger the radar backscatter return. Therefore the roughness criterians can be used to differentiate between surface geology in unvegetated areas or extent of vegetation in others.

When surface height varies on orders of magnitude greater than the radar signal's wavelength, the geometry between the incoming signal and the surface terrain must also be considered. For one thing, if the radar is coming in perpendicular to the surface, reflections will go right back up to the spacecraft, resulting in a brighter returns. (Just imagine how the signature of a lake would change.) The topography has more complicated effects as well.

Imagine, for example, radar pulses being sent toward a large mountain - not straight down from the spacecraft, but to the side. (Radars do often "look" to one side; if backscatter returns came from both sides, the time of a pulse's return would not uniquely determine the range location it came from. The angle measured from nadir, straight down, to the direction the signal travels is called the look angle.) Just like the setting sun's rays would cast a shadow behind the mountain, so too do the radar signals light up the front side of the mountain while casting a shadow behind it. The radar signals can't get to the mountains back side, therefore little radar can be backscattered from that area. These regions where the radar can't "reach" do look like shadows on the resulting imagery. The shadowing effect increases with greater look angles, just as our shadows lengthen as the sun sets.

The time it takes for a radar signal to travel from the spacecraft to the ground and back again is used to determine how far apart the spacecraft and ground are. In short, the round-trip distance equals the round-trip time multiplied by the speed of electromagnetic waves (round-trip time * $3.0 \cdot 10^8$ m/s = distance). The SAR processor uses this information to generate images: the first radar backscatter returned is assumed to come from the nearest location, and subsequent backscatter returns from adjacent locations (further from the spacecraft). Large topographic variations can confuse this process, however.

When the radar pulse is sent out from the spacecraft, it spreads in a curved fashion, in some aspects like the ripples seen in a calm lake after someone throws a rock in it. As this curve approaches the surface, it might very well hit the top of, say, a pyramid (with slopes steeper than 23 degrees for ERS-1) before it reaches its base. That

is, high objects that are further away in ground range might be hit and therefore return their backscatter before nearer but shorter objects. (Ground range vs. slant range : if you walked to the pyramid from Cairo, you'd reach the base before you climbed to the top; the radar, looking sideways at the pyramid from space, might sense the top before the base.) The SAR processor would therefore think the top of the pyramid was nearer than the bottom, and place the backscatter from the pyramid's top on the image first. This is called image layover.

Another related process contributing to this distortion effect deals with real versus assumed area. The processor is assuming that the land it's imaging is flat (well, relatively - the Earth's surface is assumed to be an ellipse). A steep slope will, however, have much more area than its corresponding base, the area the processor thinks it has. The radar backscatter returns from the many objects along the slope will therefore return at nearly the same time and all be mapped to the hill's base distance in range (distance from satellite - perpendicular to flight direction), resulting in very bright regions on the image. This foreshortening effect, combined with the extreme case of layover mentioned above, causes the near sides of mountains to look very bright and bunched together. Some people say this makes the mountains look like they are "lying down." Decreasing the radar's look angle increases this effect.

(Note that the Alaska SAR Facility's STEP program has made extensive progress in correcting many of the distortions so far described. Some of their software is available on-line, and more is under development.)

Knowledge of these various geometric interactions is used to determine the desired flight path direction over and look angle toward a particular scene of interest. For example, if a fault line of interest runs parallel to the spacecraft flight direction (perpendicular to the look direction), much of its surface area will be aligned to provide a strong backscatter return. Looking down along the fault, however, would imply lower overall topographic variation and the imaged fault line would not be as distinctive.

The radar's polarization also affects how an image will look. The radar signals can be set up to vibrate in either a horizontal (H) or vertical (V) plane or in a circular fashion. Antennas can transmit in one plane and receive in another. Possible variations are HH, VV, HV, and VH for: horizontal transmit, horizontal receive; vertical transmit, vertical receive; and so on. The HV and VH modes are used as particular discriminators. For example, volume scattering is the main interaction which causes the incoming signal to be depolarized (scattered in other directions, at different relative vibrational planes). A H-polarized wave might hit a tree, bounce back and forth among leaves, and be backscattered to the spacecraft in a V-polarized state. Using an HV setup, the antenna could receive these V-polarized returns. Though the overall radar backscatter received would be much lower than, say, an HH setup, the resulting image would have increased variation between regions of volume scattering (higher potential to return V-polarized backscatter - usually vegetation) and regions of surface scattering.

Now that a person has begun using these more science-oriented characteristics to choose the SAR's parameters, some engineering concerns must also be addressed. The spacecraft parameters discussed above, such as the look angle and signal wavelength, as well as the antenna size and other parameters influence the resulting image's resolution.

The range resolution (in the direction perpendicular to the spacecraft's flight direction) is dependent upon the look angle and signal pulse length. The initial theoretical resolution along the slant range is half the pulse length. Multiply the pulse length in time by the speed of electromagnetic radiation (in air, 3.0×10^8 m/s) and divide by 2 to get the slant range resolution (the resolution along the direction in which the radar signals propagate), and divide that value by the cosine of the look angle to obtain the ground range resolution. Note that decreasing the pulse length decreases resolution (makes it better), but that in so doing you also decrease the overall energy in each pulse. Though the signal intensity can be increased, the spacecraft's available power (and heat dissipation capabilities) limit how far you can go.

This resolution is often not acceptable, so signal processing techniques have been developed to improve it. For example, using the resolution theory outlined above, ERS-1's pulse length of 37.1×10^{-6} seconds and its 20.35 degree look angle would result in a ground resolution on the order of several kilometers. It takes some intensive signal processing to obtain the ERS-1 SAR image products' 25 meter resolution. One trick is to have a pulse that changes characteristics over its length. The ERS-1 pulse is called a chirp ; it ramps up in frequency throughout the pulse. This means that, because they contain different frequencies, sections of the pulse can be distinguished from each other and therefore simulate even smaller pulses. The ability to resolve a set of frequencies within the pulse sets a new simulated pulse length, which for ERS-1 results in 25 m rather than several kilometer resolution. The way in which these new "mini-pulses" are resolved, or more accurately the way in which individual chirp returns are discriminated, is outlined in the processing documentation.

Azimuth resolution (resolution in the "along-flight-track" direction) is determined by the footprint or region on the ground which will be hit by one pulse. The resolution is proportional to the wavelength divided by the antenna length. (You might recognize this proportionality from optics, where it is used to determine the diffraction limit and therefore resolution.) The resolution is often assumed to be $.7 \times \text{Range} \times \text{Wavelength} / \text{Antenna Length}$, about several kilometers for ERS-1. Here is where the "synthetic" part of Synthetic Aperture Radar (SAR) comes in to play.

As the spacecraft passes over a target, a ground target appears to be first in front of, then next to (well, relatively), then behind the spacecraft. In the meantime the SAR has sent out many (for ERS-1, about 1000) pulses and therefore recorded a specific target's radar backscatter response about 1000 times. The response to each one of the 1000 pulses was somewhat different, depending on changes in target-sensor geometry and Doppler effects. Just as the Doppler effect changes the sound you hear as a car passes, the motion of the spacecraft relative to the target increases or decreases the signal's frequency. By analyzing the return signals from those 1000 pulses, the target's

Doppler history can be determined. This information allows a target's backscatter to be analyzed as though it had been seen by 1000 different antennas, or correspondingly of a synthesized antenna with length equal to the distance the spacecraft passed through while it was able to get backscatter returns from that target. This large synthesized antenna length decreases (improves) the resolution, and as an effect of the processing methods this resolution holds regardless of the distance between the spacecraft and target. For a further description, see the processing documentation.

So the spacecraft sends out this (sophisticated) pulse and then "listens" to the echo. The pulse is first backscattered by the nearest sensed object, altered in a variety of ways as outlined above. Before that first response is completely recorded, however, pulses are returning from the next locations. They overlap. Then you have pulses being returned from a bit ahead and behind the spacecraft, coming in at different frequencies because of Doppler shifting. Next the spacecraft sends out another pulse, and the whole thing happens again. All of these pulse returns are combining into a big jumble, line after line, so how do we get the pretty pictures? What are the limits? Hold on to your hats; here's where SAR processing enters the scene...