Synthetic Aperture Radar (SAR) Equations in the ASF User Tools

Orion Sky Lawlor
Alaska SAR Facility / Department of Computer Science, University of Alaska at Fairbanks
http://lawlor.cs.uaf.edu/∼olawlor olawlor@acm.org

1 Introduction

There are lots of excellent books on satellite radar processing–in particular, check out the SAR "bible" of Curlander and McDonough [CM91], which is packed full of equations and exciting mathematics. For more pictures and slightly fewer integrals, see Francheschetti and Lanari's recent book [FL99]. Our own Coert Olmsted wrote the relatively readable "Scientific SAR User's Guide" [Olm93] back in 1993. If you're looking for a complete, end-to-end guide to SAR, check these things out.

If you're just looking for the bottom line, or you've stared in slack-jawed horror at the hairy integrals in the above references, you've come to the right place! This is an informal guide to SAR processing in practice, and I will justify every equation.

2 Triangles for SAR

As you can see in this document, there are lots and lots of triangles used in SAR processing.

Everybody knows the Pythagorean Theorem, which relates the lengths of a right[1] triangle's hypotenuse c to the lengths of its other sides a and b, as shown in Figure 1 (A).

$$c^2 = a^2 + b^2$$

This is easy to remember, but it's rare that you have a perfect right triangle. More often, you get a triangle with some arbitrary angle θ not equal to 90 degrees, like in Figure 1 (B). Luckily, there's a slightly modified version of the Pythagorean Theorem that works for all triangles (not just right triangles) called the "Law of Cosines":

$$c^2 = a^2 + b^2 - 2\,a\,b\,\cos\theta$$

[1] Right triangle = 90 degree angle

The triangle we're usually interested in is formed by these three points: the SAR satellite, the center of the Earth, and the "target" point being observed on the Earth's surface. A cross-section of a SAR scene like this is shown in Figure 1 (C). The Law of Cosines as written above thus directly translates to a SAR scene by relating the slant range s to the "earth angle" α.

$$s^2 = e^2 + h^2 - 2eh\cos\alpha$$

$$\alpha = \cos^{-1}\left(\frac{e^2 + h^2 - s^2}{2eh}\right)$$

But we can compute more angles by switching which side we define as c while applying the Law of Cosines–this can give us any of the angles in the triangle. Hence we can relate the Earth radius e to the "look angle" l.

$$e^2 = s^2 + h^2 - 2sh\cos l$$

$$l = \cos^{-1}\left(\frac{s^2 + h^2 - e^2}{2sh}\right)$$

The incidence angle i isn't quite as easy, since it's measured from straight up. If we ignore the variation in "up" due to the ellipsoid and geoid, we can compute the incidence angle as the complement of the angle inside the triangle, like this:

$$h^2 = s^2 + e^2 - 2se\cos(\pi - i)$$

$$i = \pi - \cos^{-1}\left(\frac{s^2 + e^2 - h^2}{2se}\right)$$

Where to summarize, our notation is:

- e Earth radius; distance from center of earth to target, meters.

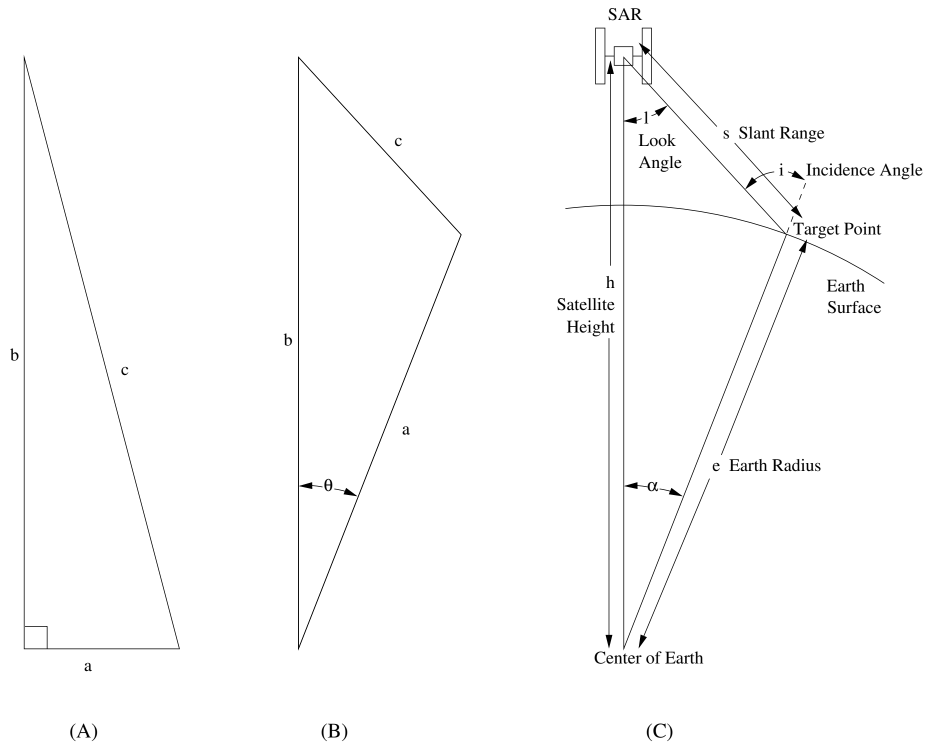

Figure 1: (A) a simple right triangle, to illustrate the Pythagorean Theorem. (B) a non-right triangle, to illustrate the Law of Cosines. (C) application of the law of cosines to a SAR observation.

- s Slant range; distance from satellite to target, meters.
- h height of satellite; distance from satellite to center of Earth, meters. Can be computed from the state vector XYZ position as $h = \sqrt{x^2 + y^2 + z^2}$.
- α ground angle; angle between nadir and target as measured from the center of the Earth, radians.
- i Incidence angle; angle from straight up over to satellite, as measured from the target point, radians.
- l Look angle; angle from straight down over to target, as measured from the satellite, radians.

3 Geocoding

So you've finally processed a SAR image. You want to relate the pixels in the SAR image to some chunk of dirt on the surface of the Earth. Good luck! This is one of the trickiest operations in SAR–even production and commercial software regularly mangles these operations.

3.1 State Vectors

A "state vector" is just the position and velocity of the satellite at a particular time. The trickiest part about state vectors is figuring out the coordinate system they're measured in.

As show in Figure 2, both inertial and fixed-earth coordinate systems are useful in different situations. The ASF tools routines `gei2fixed` and `fixed2gei` convert between these coordinate systems, including both "GHA rotation" and the mysterious "Coriolis term" as described in Appendix A.

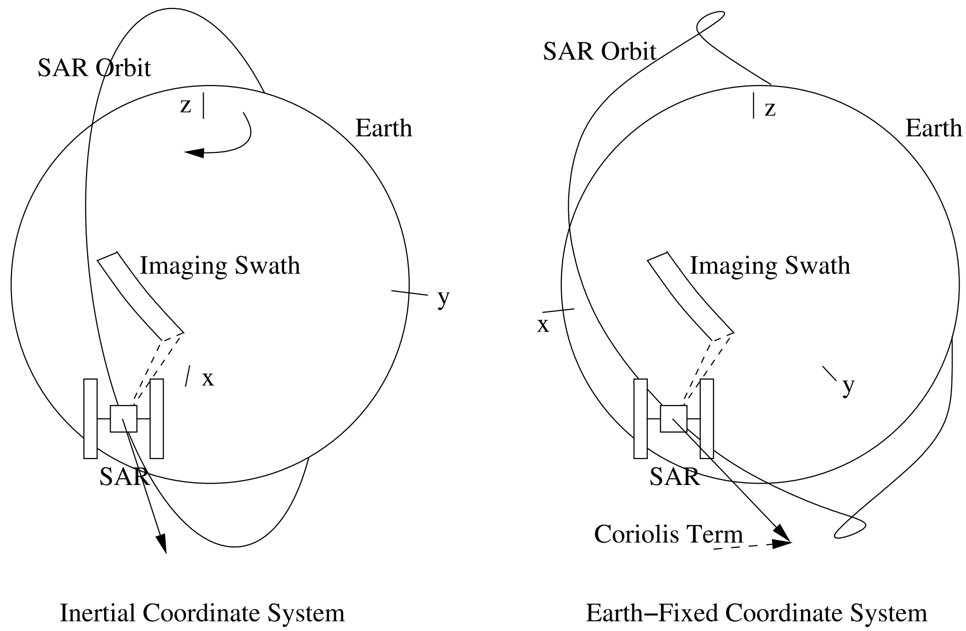

Figure 2: The same satellite pass plotted in both Inertial and Earth-fixed coordinate systems. In an inertial coordinate system, the axes don't move relative to the stars, so orbits are simple; but the Earth does rotate, so ground tracks are weird. In an Earth-fixed coordinate system, the axes are stationary relative to the Earth, so ground tracks are simple; but orbits are weird.

3.2 Doppler to Azimuth Position

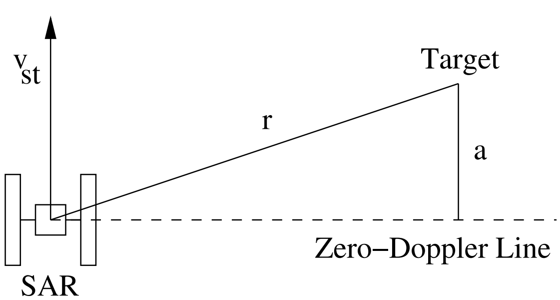

Figure 3: A SAR satellite moving at a velocity of v_{st} relative to the ground sees a target at range r and azimuth a.

There's a nice little scalar equation that relates doppler (e.g., found during SAR processing) to the target's along-track position.

$$f_D = \frac{2\, v_{st}\, a}{\lambda\, r}$$

$$a = \frac{f_D\, \lambda\, r}{2\, v_{st}}$$

- f_D doppler shift of echo radiation, Hz.
- v_{st} magnitude of fixed-earth spacecraft velocity, m/s.
- a target's azimuth location, as measured along the spacecraft's fixed-earth velocity vector, in meters. $a = 0$ lies in the plane perpendicular to the satellite's fixed-earth velocity. Also known as the "along-track distance".
- λ wavelength of radiation, m.
- r slant range to target, m.

Rationale: See Figure 4. $v_{rel} = v_{st}\, a/r$ is the relative velocity of the target along the spacecraft's line of sight. Dividing by $\lambda/2$ turns this closing velocity into a frequency shift.

References: Stare at [CM91] Fig 1.8 to understand the geometry. [CM91] Eq 1.2.4 incorrectly implies that this is an approximation—it's actually exact as defined above (if you're willing to ignore relativity).

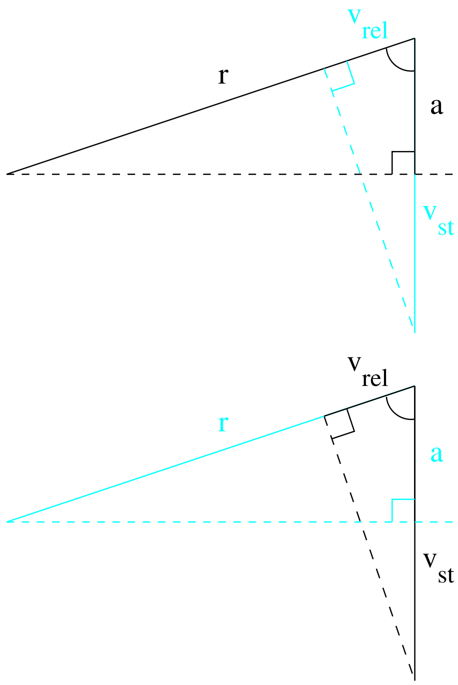

Figure 4: Similar triangles showing that a/r equals v_{rel}/v_{st}. Compare to Figure 3. We've flipped to imagining a stationary satellite and moving target, and project the target's velocity into the satellite range direction.

4 Interferometry

SAR "interferometry" examines the phase relationship between two or more satellite images to determine target motion or deformation, surface topography, and decorrelation. Or, phrased in a more depression fashion, the phase in a given interferogram represents an unknown combination of target motion, surface topography, and random decorrelation phase. Hence often in interferometry, we must make some effort to

4.1 Target Motion

The simplest case for interferometry is when both satellite images are aquired from the same location, and the phase relationship is caused by target motion. The key parameter here is the "path length difference"–the extra total distance travelled by the radiation in the second image. This distance d, in meters, can be computed from the phase p, in radians, via a conversion constant known as the "wavenumber", normally denoted k, in radians per meter. k is related to the wavelength λ, in meters, by $k = 2\pi/\lambda$.

$$d = \frac{p}{k}$$

When the two images use different transmitters, the total path length difference is actually twice the physical motion of the target, because the radiation has to leave and come back along the new displacement. Hence for a two-bounce system, we usually say $d = \frac{p}{2\,k}$, where d now represents the actual *target* motion distance.

You may have noticed that d, the distance travelled, is a scalar, while the actual target motion is usually a vector $\vec{\mathbf{d}}$. Interferometry can only measure displacements along the radar line of sight unit vector $\vec{\mathbf{v}}$;[2] so our scalar distance d is actually a projection of the true target motion $\vec{\mathbf{d}}$ onto the line of sight vector $\vec{\mathbf{v}}$, or $d = \vec{\mathbf{d}} \bullet \vec{\mathbf{v}}$.

Theoretically, given a set of scalar target motions d_i each with a different line of sight vector $\vec{\mathbf{v}}_i$, one can solve the (usually ill-conditioned but overdetermined) linear system $d_i = \vec{\mathbf{d}} \bullet \vec{\mathbf{v}}_i$ for the unknown vector target motion $\vec{\mathbf{d}}$. In practice, this is complicated when the observations are not all simultanious, or when the target motion is not entirely uniform.

4.2 Interferometric Baseline

Figure 5 shows a cross-section of two satellites looking at an object. We call the spatial distance between the two satellites the (spatial) "baseline", which for a given location has only two components–the parallel baseline B_p and the normal baseline B_n. Both are just distances in meters. There is no 3D (for example, B_z) component of the baseline, because while forming an interferogram, we always shift the images along azimuth (out of the page) until they line up. However, sometimes people do speak of a "temporal baseline" as the time, in seconds or days, between the two observations.

[2] Yes, for interferometry there are two different lines of sight; and the phase difference actually only specifies the size of an ellipsoid whose foci lie at the two satellites. For actual interferograms, there isn't much difference between this family of ellipsoids and the family of equal-distance planes perpendicular to the average line of sight.

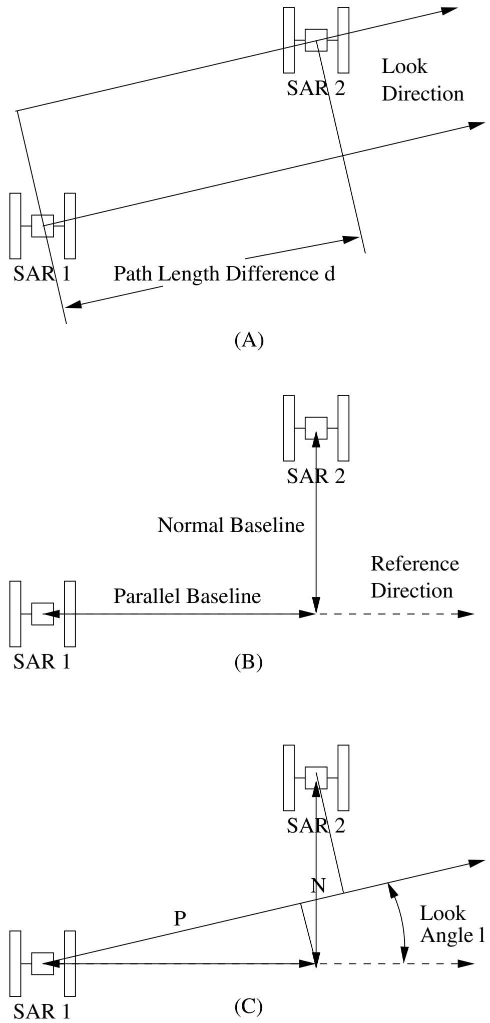

Figure 5: In interferometry, there are two or more satellites looking at the same thing. Because the two satellites are separated in space, there's a look-direction-dependent path length difference between the two satellites (A). The satellite "baseline" is defined as the distance parallel to and normal to some reference direction (B). Given the angle between the look and reference directions, we can compute the path length difference from the baseline (C).

Given the baseline measured from some reference direction, we can compute the path length difference along any other direction using trigonometry. If the angle between the reference and look directions is l radians, then the path length difference d in meters between the satellites is just:

$$d = P + N = B_p \cos(l) + B_n \sin(l)$$

So overall the expected phase difference at a given look angle (for a 2-antenna system) is exactly:

$$p = 2kd = 2k(B_p \cos(l) + B_n \sin(l))$$

Computing the expected phase from a given look angle is sometimes useful to subtract the phase of a known ellipsoid or topography. But more often we want to go the other way–compute the topography or look angle change for a given phase change as measured from an interferogram.

4.3 Phase to Topography

Any change to the look direction l, call it δl, clearly changes the path length d to:

$$d + \delta d = B_p \cos(l + \delta l) + B_n \sin(l + \delta l)$$

For topography, is common to assume δl is small.[3] Thus we can take the differential approximation, which gives:

$$\delta d \approx \delta l(-B_p \sin(l) + B_n \cos(l))$$

We can turn this around and combine with $d = \frac{p}{2k}$ to find the angle change that results from a given phase change:

$$\delta l \approx \frac{\delta p}{2k(-B_p \sin(l) + B_n \cos(l))}$$

Given this topography-driven angle change δl as measured from the phase difference, we can compute the target height above the Earth surface using the Law of Cosines exactly as in Section 2. The situation is shown in Figure 6; the Law of Cosines gives us:

$$(e + t)^2 = s^2 + h^2 - 2sh \cos(l + \delta l)$$

and thus solving for the topographic height t we get

$$t = \sqrt{s^2 + h^2 - 2sh \cos(l + \delta l)} - e$$

Given a correct look angle δl, this is exact but highly nonlinear. It is common to linearly approximate[4] the topographic height from the phase difference δp as

$$t = \delta p \frac{s \sin(i)}{2k(-B_p \sin(l) + B_n \cos(l))}$$

[3] For example, an 1km height difference from a 900km slant range is just under 1/15th of a degree.

[4] I can't seem to re-derive this approximation, but it does appear to work.

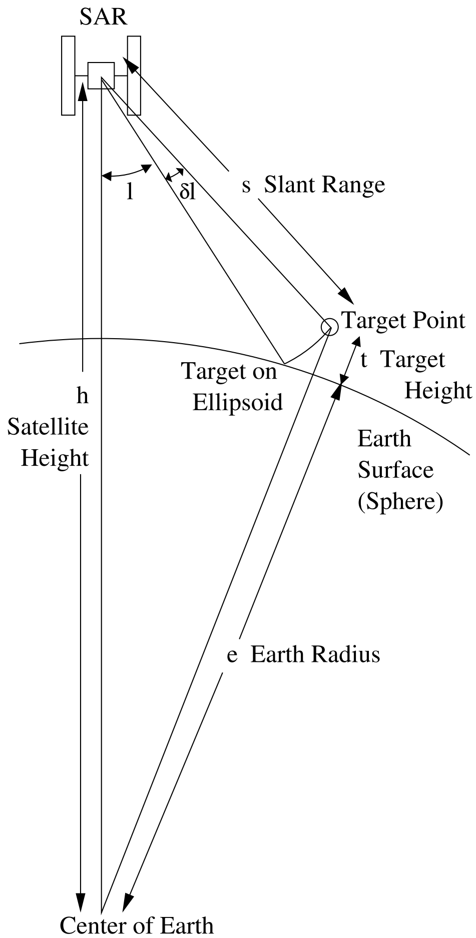

Figure 6: Computing topography height t based on interferometrically-measured look angle difference δl.

- p Phase difference; measured by subtracting the phase of the images to be interfered, radians.

- d Distance target moved between observations, projected into the satellite average line of sight, meters. $d = \frac{p}{2\,k}$

- k Wavenumber; phase change per unit distance, radians per meter. $k = 2\pi/\lambda$, where λ is the radar wavelength in air.

- l Look angle; angle from straight down over to target, as measured from the satellite, radians.

- e Earth radius; distance from center of earth to local ellipsoid, meters.

- s Slant range; distance from satellite to target, meters.

- h Height of satellite; distance from satellite to center of Earth, meters. Can be computed from the state vector XYZ position as $h = \sqrt{x^2 + y^2 + z^2}$.

- t Topography height above a hypothetical spherical Earth of radius e, meters.

- i Incidence angle; angle from straight up over to satellite, as measured from the target point, radians.

- l Look angle; angle from straight down over to target, as measured from the satellite, radians.

- B_p Parallel baseline; distance between satellites measured along the reference direction, meters.

- B_n Normal baseline; distance between satellites measured across the reference direction, meters.

References

[CM91] John C. Curlander and Robert N. McDonough. *Synthetic Aperture Radar Systems and Signal Processing*. John Wiley & Sons, 1991.

[FL99] Giorgio Franceschetti and Riccardo Lanari. *Synthetic Aperture Radar Processing*. CRC Press, 1999.

[Olm93] Coert Olmsted. Scientific sar user's guide. Technical report, Alaska SAR Facility, 1993. http://www.asf.alaska.edu/reference/general/SciSARuserGuide.pdf.

A State Vector Coordinates

In an inertial coordinate system, Newton's laws hold, and the satellite's orbit is just an ellipse–so inertial coordinate systems are used for orbit propagation. However, because the Earth moves beneath the satellite, the ground track is at an angle to the inertial flight path, which significantly complicates geolocations. A typical inertial coordinate system is Geocentric Equatorial Inertial (GEI), where the center of the Earth is the origin, the X axis always points to the spring equinox, and the Z axis points to the north pole.

In fixed-earth coordinates, the Earth is stationary with respect to the coordinate system, which makes it much easier to work out the swath location–so fixed-earth coordinates are often used for geolocations. But because the coordinate system is rotating beneath the satellite, an orbit is quite strange. The standard fixed-Earth coordinate system puts the center of the Earth at the origin again, the X axis at latitude 0 (the equator) and longitude 0 (the Greenwich meridian), the Y axis at latitude 0 longitude 90 degrees east, and the Z axis at the north pole again.

Converting from GEI to fixed-earth accurately is actually pretty tricky. First, you need to figure out the current rotation of the Earth. This is measured by the angle between the Greenwich meridian (the fixed-Earth X axis) and the vernal or spring equinox (the GEI X axis). This is variously called Greenwich (apparent) Sidereal Time (GST) in the International Astronomer's Union (IAU); and somewhat anachronistically the "Greenwich Hour Angle" (GHA) in the ASF tools and CEOS metadata. Since the Earth rotates once a day, the GHA varies between 0 and 360 degrees over the course of a day;[5] so you need accurate timing information to get an accurate GHA.[6] The ASF tools routine `utc2gha` can convert UTC time to a GHA for the vernal equinox.

So, to convert a GEI state vector to fixed-Earth, you find the GHA and then rotate the state vector in the XY plane by that angle, which is easy enough. But you're not done. Velocity in fixed-Earth coordinates isn't just a rotation from the GEI velocity, because we've got to subtract off the velocity of the co-

[5] A sidereal day, relative to the equinoxes.

[6] A one-degree GHA error causes a 100km east-west shift at the equator.

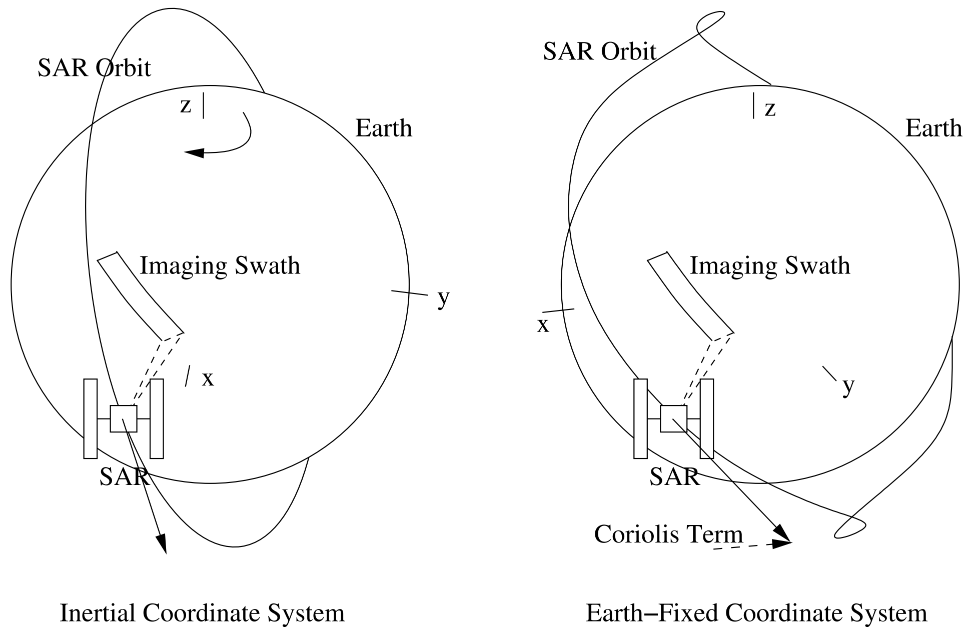

Figure 7: The same satellite pass plotted in both Inertial and Earth-fixed coordinate systems.

ordinate system itself. This "Coriolis term" is often neglected with disastrous results, because the Earth itself rotates at about 1000mph at the equator.

The fixed-Earth Coriolis velocity at a location (x, y, z) (in meters), is $(-\omega y, \omega x, 0)$ (in meters/second). So far, this just says points on the Earth rotate around the north pole, by moving tangent to their XY position—standard cylindrical coordinates stuff. More interesting is the rotation rate ω (in radians/second). Since the Earth revolves one whole rotation per day,

$$\omega = 2\pi/daylength$$

However, since ω is the rotation rate the the Earth relative to the stars,[7] or the "sidereal rotation rate", a day isn't 24 hours long. This is because the Earth actually makes one extra revolution per year; because days are measured relative to the sun, which we also make one complete orbit around. So overall the Earth's sidereal rotation rate is actually $\omega = (2\pi/(24 * 60 * 60)) * (1 + 1.0/365.24218967)$ radians per second, where the funny constant is the length of the "tropical year" (year as measured between equinoxes).

Overall, you can convert a fixed-Earth state vector g into an inertial coordinates state vector i by first rotating both position and velocity by the Greenwich hour angle $theta$, and then adding in the coriolis velocity term:

$$\begin{aligned}
i.x &= g.x \cos\theta - g.y \sin\theta \\
i.y &= g.x \sin\theta + g.y \cos\theta \\
i.z &= g.z \\
i.vx &= g.vx\ \cos\theta - g.vy\ \sin\theta - i.y\,\omega \\
i.vy &= g.vx\ \sin\theta + g.vy\ \cos\theta + i.x\,\omega \\
i.vz &= g.vz
\end{aligned}$$

Note the Z coordinate is unchanged. To convert from inertial back to fixed-Earth, just flip the sign on θ and ω.

[7]GEI is actually in terms of the equinox, not the stars; but if you ignore the 26,000-year precession of the equinoxes, they're the same thing.

B ASF Precision Processor Geolocations

The ASF Precision Processor (PP) produces SAR images in “Ground Range”. This section attempts to explain what the PP means by Ground Range, and how this can be related to classic SAR parameters like observation time and slant range.

Scientists with high accuracy requirements prefer to work with time and slant range, since these are the fundamental axes of a SAR. Luckily Doppler is not a factor for PP geolocations, since after processing to the natural Doppler, ASF PP images are ”deskewed” along azimuth so they geometrically lie at a Doppler of 0 Hz.

B.1 PP y Coordinates

The y coordinate in a PP ground range image is intended to represent along-the-ground azimuth distance, but actually is just a scaled version of time. For example, if the data was processed to a 12.5 meter ground range pixel spacing, the actual distance along the ground may vary due to topographic and ellipsoid effects not modeled by the PP.

We can be more precise by taking ASF CEOS Facility Data Record field 96, AZIMUTH_PIXEL, which we’ll call a (typical value: 12.5 meters per pixel); and field 77, SWATH_VELOCITY, which we’ll call v (typical value: 6626 meters per second).[8] The actual time per azimuth pixel is then just a/v seconds per pixel (typical value: 1.88 milliseconds per pixel).

Thus the Y coordinate is linearly related to observation time. This time, along with the slant range determined in the next section, can be combined with the observation state vectors in the ASF CEOS Platform Position Record to determine the exact imaging geometry.

In practice, ASF L1 images always have the earliest time at the bottom of the image (toward the end of the CEOS disk file); and later times at the top of the image (at the start of the CEOS file). This might be the reverse of what you expect. In the ASF Tools, azPixTime is normally negative for ASF CEOS images, indicating decreasing time with increasing line count.

[8] This is almost equal to the orbital velocity, but scaled down to the Earth height.

B.2 PP x Coordinate

The x coordinate in a PP ground range image actually represents the arc length g in Figure 8. This arc length is the distance from the point directly under the spacecraft (the nadir point) to the point on the Earth’s surface being imaged, measured along an idealized Earth surface, which for the PP is a sphere.

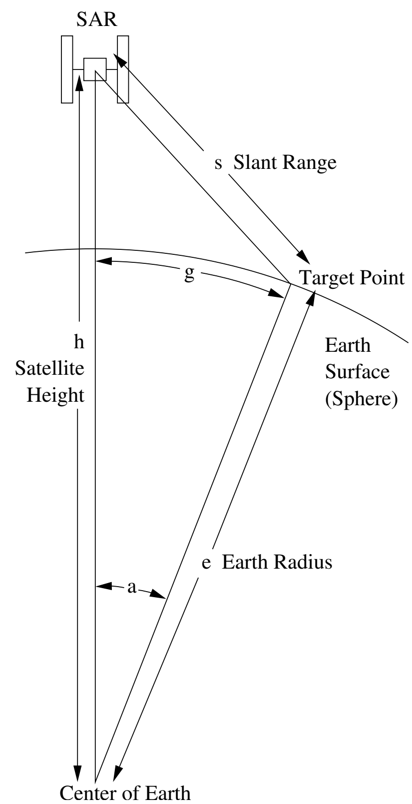

Figure 8: A cross section of the idealized PP observation geometry.

Of course, in reality the Earth is not a sphere. The overall shape of the Earth is ellipsoidal, and even sea level follows local gravitational variations of the geoid, and finally land surfaces have significant topography. Each of these create distortions in the image which the PP is at the moment unable to correct–although our terrain-corrected products take all three factors into account (ellipsoid, geoid, and topography).

The PP computes the ground range arc length g in Figure 8 by multiplying the ground angle a, in radians, by the earth radius e; or $g = ae$ (the usual angle times radius arc length computation). The ground angle a, in turn, can be related to the satellite height h and satellite-ground slant range distance s using the law of cosines as follows:

$$s^2 = e^2 + h^2 - 2eh\cos a$$

And thus we can interchange between PP ground range and slant range as follows:

$$s = \sqrt{e^2 + h^2 - 2eh\cos(g/e)}$$

$$g = e\cos^{-1}\left(\frac{e^2 + h^2 - s^2}{2eh}\right) = g(e, s) \qquad (1)$$

- s Slant range; distance from satellite to target, meters.
- e Earth radius used by PP, meters.
- h Satellite height; distance from satellite to center of Earth, meters. Can be computed from the state vector XYZ position as $h = \sqrt{x^2 + y^2 + z^2}$.
- g PP ground range; arc length along sphere from nadir to target, meters. $g = a\,e$
- a ground angle; angle between nadir and target as measured from the center of the Earth, radians. $a = g/e$.

The slant range determined here can then be used in accurate SAR geolocation.

B.3 PP Earth Radius

Unfortunately, the Earth radius e the PP uses to compute ground range is not the true radius of the Earth at the scene center nor at the nadir; but some other radius.[9] What's worse, the PP Earth radius is not stored directly anywhere in the ASF CEOS.

Currently, the best way to reconstruct the PP's Earth radius e is from the CEOS Facility Data Record. We begin with CEOS field 83, SL_RNG_1ST_PIX, which we'll call s_F; and field 84, SL_RNG_LAST_PIX, s_L. Now, the total ground range distance across the scene should equal CEOS field 97, RANGE_PIXEL, times field 28, ACTUAL_PIXELS (minus one); call this total ground range distance across the whole image Δg. So given a prospective Earth radius e, we can use equation 1 above to compute the ground range distance across the image as $g(e, s_L) - g(e, s_F)$, which in a perfect world would equal the Δg we compute from the CEOS. It usually isn't the same, but we can then adjust our prospective Earth radius e using any standard root-finding algorithm (bisection, secant method, ...) until our $g(e, s_L) - g(e, s_F)$ equals the PP's Δg, which will mean our e matches the earth radius computed by the PP.

So overall we just adjust e until

$$g(e, \texttt{SL\_RNG\_LAST\_PIX}) - g(e, \texttt{SL\_RNG\_1ST\_PIX}) = \texttt{RANGE\_PIXEL} * (\texttt{ACTUAL\_PIXELS} - 1)$$

When using the ASF Tools, the PP earth radius e is computed in exactly this iterative fashion and stored in the .meta file as "earth_radius_pp".

[9] We are still trying to determine exactly how the PP computes e; at the moment it appears to be a geolocation computation based on an incorrect state vector.

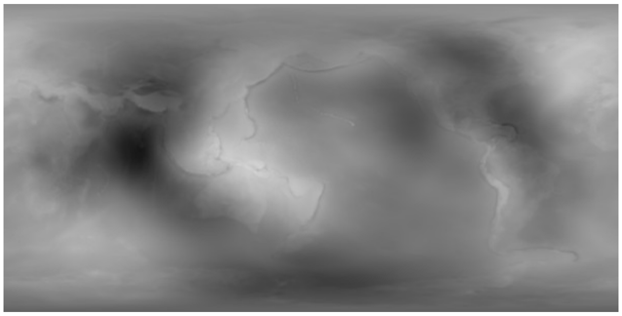

Figure 9: The EGM96 geoid model over the entire Earth. Black represents a geoid 128m below the ellipsoid; white represents a geoid 128m above the ellipsoid.

C Geoid

The trickiest part about measuring stuff on the surface of the earth is the weird terminology. Heck, the field of measuring stuff on the surface of the Earth is called ”Geodesy”, which is weird terminology already!

- Ellipsoid: major and minor radius of earth. Simple mathematical model, easy to convert to and from 3D coordinates. GPS units usually gives you ellipsoid-based heights, and the SAR tools want to process ellipsoid-based heights. Example ellipsoids: Clarke 1866, GRS80, WGS84.

- Geoid: model of the theoretical sea level, which sadly differs from the ellipsoid due to local gravitational variations (e.g., mountain ranges, ocean trenches). The geoid can be up to about 100 meters higher or lower than the ellipsoid, although it's closer to 10 meters higher in Alaska. Page 3 of this presentation shows the difference between geoid and ellipsoid for North America. Among the best current geoid models is EGM96, which can be downloaded as a 1/4” degree raster image.

- Datum: coordinate system used for mapping. Usually much closer to the geoid than to the ellipsoid. Not always mathematically defined; for example, NGVD27 is defined relative to a set of physical sea level guages. Maps and other ”othometric” measurements are always relative to some datum. Example datums: NAD27 (horizontal) and NGVD27 (vertical); NAD83 (horizontal) and NAVD88 (vertical).

If you really want millimeter accuracy across the Earth's surface, you've got to take into account a variety of difficult to model factors: continental drift (centimeters per year, but continuous), earthquakes (up to a few meters after a big one), post-glaciation isostatic ”rebound” (millimeters vertical per year), seasonal and long-term hydrological inflation and deflation (centimeters per year), and so on. Luckily, our current remote sensing images only need to be accurate within about one meter, and so we can ignore all of these effects.

For most purposes, we can actually approximate a vertical datum, like NAVD88, with a geoid, like EGM96. This works because all vertical datums are approximations of the Earth's gravitational surface, and hence only differ by a few meters at most.

The geoid, however, can't be ignored for accurate imaging. For example, EGM96's geoid height near Alaska is as follows:

Geoid Height Meters	Lat,Lon Degrees
15.40	64,-145
13.42	65,-145
10.22	66,-145
14.97	64,-144

So the geoid height in Alaska is nonzero, nonneglible, and non-constant over a frame. Luckily, as you can see the geoid height changes slowly. Hence to first order, you can convert datum-relative heights (like you'd download in a DEM) to ellipsoid-relative heights (like you'd get from a GPS, or use in the tools) by adding in the ellipsoid-relative geoid height. As shown in Figure 10, problem we're fixing here is that the vertical distance between the topography and geoid (the datum-relative height) isn't the same as the vertical distance between the topography and ellipsoid (the ellipsoid-relative height).

C.1 Geoid Tilt

Unfortunately, the mountain ranges that cause geoid undulations don't just affect height–they do this by tilting the local gravitational normal. Think of a

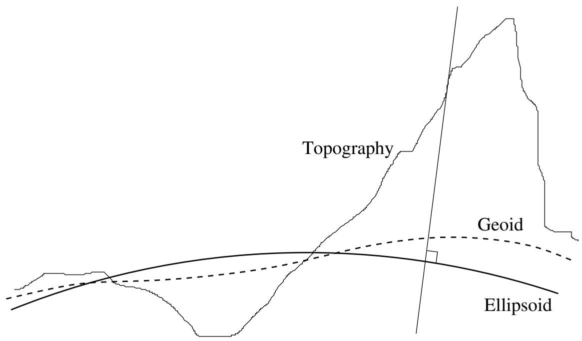

Figure 10: Topography, geoid, ellipsoid, and vertical vector.

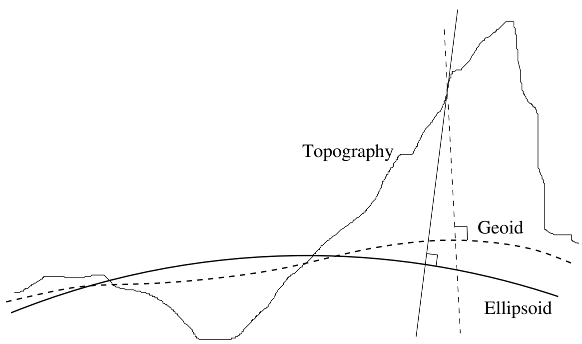

Figure 11: The geoid's "vertical" direction (dashed vertical line, normal to the geoid) isn't the same as the ellipsoid's "vertical" direction (solid vertical line, normal to the ellipsoid). Luckily, this geoid tilt effect is small.

plumb bob hanging next to a mountain range–it'll tilt toward the center of the range. So to be really accurate, as in Figure 11, when changing from ellipsoid to geoid-relative heights there's also a shift in the direction meant by "vertical", which when combined with nonzero topography causes a horizontal position shift.

The maxmium tilt on Earth is in the Andes, where there's about 0.3m horizontal normal tilt per kilometer of elevation (or, equivalently, the ellipsoid and geoid differ an additional 0.3m vertical distance for every kilometer of horizontal motion). For the highest point on Earth, Mt. Everest, at 8.85km high, the horizontal shift at the top due to the geoid would be less than 2.5m.