SAR image post-processing for acquisition trajectory change

Hubert M J Cantalloube

SAR image post-processing for acquisition trajectory change

Hubert M.J. Cantalloube[a]
[a]ONERA Université Paris Saclay, F-91123 Palaiseau, France

Abstract

Locally on a SAR image, the space frequencies can be mapped to the integration time cross radar bandwidth domain (this is the principle of both frequency-domain omega-kappa a.k.a. range-migration and polar formatting types of SAR synthesis algorithms). Thus a SAR image can be post-processed by locally remapping and phasing its space frequencies to refocus the image for example according to an improved trajectory or a terrain elevation change. The main limitations are the non stationarity of the correction in the generic case and the azimuth extend of the focus blur, which constrain respectively the mesh size of the correction and the azimuth window size, result in a computation cost higher that ab nihilo SAR synthesis with the new trajectory/terrain elevation model. This is illustrated with typical airborne SAR image in mountainous area acquired by ONERA airborne radar Sethi.

1 Geometry of the space frequency spectrum of a SAR image

Native SAR images are generally in slant-range coordinates, with the first "range" coordinate X being the distance to a straight line of direction V and origin O, the "nominal trajectory", and the second "azimuth" coordinate Z being the distance measured along the nominal trajectory from the origin point O. In monostatic cases as in **Figure 1**, the nominal trajectory (O, V) is chosen close to the effective acquisition trajectory (typically through a linear regression or the mere straight line between the start and the end of acquisition trajectory) and the range distance is measured radially (so called "zero-Doppler geometry") or along cones of nominal axis with tip angle ("squint of the image") chosen in the middle of the integration interval. Images from ONERA are in squinted geometry, as it is commonly the case with airborne SAR because early designs of SAR system did not have lateral steering capacity for compensating the side wind induced drift angle.

In bistatic case, especially strongly asymmetrical cases such as satellite-to-aircraft, the choice for "nominal trajectory", i.e. the origin axis of the SAR image geometry, is more arbitrary. We can, for example, use a straight line approximation of the receiving aircraft trajectory, but that of the satellite or any intermediate line is also sensible. Hence in the following development, we should keep in mind that the X and Z may be remotely linked to with the effective acquisition geometry, i.e. the effective range to sensor and/or sensor position at (the middle of) integration.

The **figure 2** shows the typical shape of the local spatial frequency spectrum of a squinted geometry SAR image. For illustrating, the image is unweighted, otherwise the edges of the spectrum support would be damped by the weighting window (typically a Hann weighting is used on bandwidth and integration on SAR images because of its 1/n² damping of the far side lobes which is appealing due to the generally very high dynamical range of radar images). The spectrum on figure 2 is also evaluated at a strong point echo in order to limit the effects of the space variability of the spectral support due to the nonlinearity of the effective acquisition trajectory (almost all the energy in the spectrum comes from the point echo, hence its visible edges are the very support at the square centre). The small gradient in phase (colour coded on the figure) is due to the fact the point echo position is not at the exact centre of the middle pixel of the square.

The link between raw signal spectrum and the image frequency is instrumental in frequency-domain SAR image synthesis such as omega-kappa algorithm (a.k.a. range-migration algorithm see [2, 3]). Since the Doppler frequency for a given imaged point is (at least in monostatic case) decreasing with time, signal slow-time frequency (Doppler) can also be mapped to time within integration, which is the principle of the Polar formatting al-

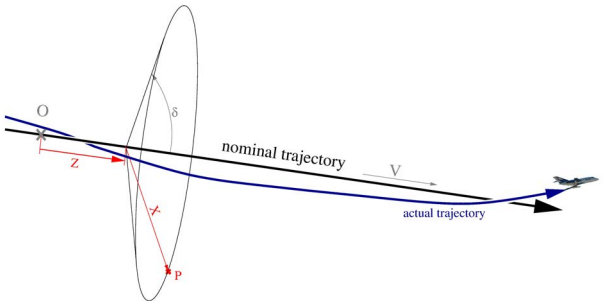

Figure 1 Slant-range SAR image geometry in the typical stripmap monostatic case. A pixel corresponds to the backscatter from a circular region which intersects ground at two positions, side looking actual antenna pattern limits imaging to only one side *P* of the ground-track. Pixel azimuth *Z* is measured along the nominal trajectory from its origin *O*. Range *X* is measured from the nominal trajectory in the squint δ direction.

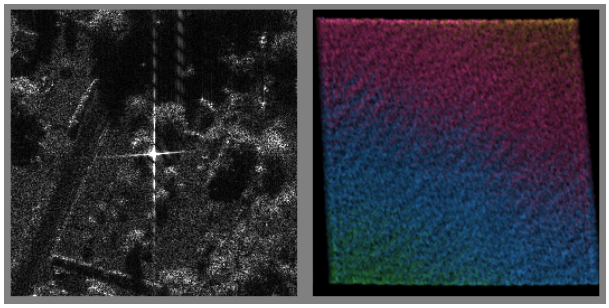

Figure 2 spectrum (right) of a small SAR image square (left). Range axis X and k_X resp. is vertical and azimuth axis Z and k_Z resp. is horizontal. For clarity, image square is centred around a strong point echo and the image is unweighted and squinted.

gorithm (PFA).

Figure 3 shows the basic principle of the omega-kappa frequency domain processing (note this figure is conceptual, and only phase is figured for clarity). If the actual trajectory were a perfect linear acquisition along the nominal trajectory at constant velocity, it is easy to derive the SAR image corresponding to a raw signal spectrum containing only one single non zero frequency (i.e. a given fast-time frequency within the radar band and a given slow-time or Doppler frequency): it is a pure plane wave in the half plane *X>0*,Z corresponding to the image coordinates. This is easy to figure out if one consider first the SAR image corresponding to a single pulse of pure fast-time frequency as it is a spherical wave centred on the antenna position at the pulse transmit/receive (spherical wave spacing is half the pure frequency wavelength because of the two-way propagation of waves).

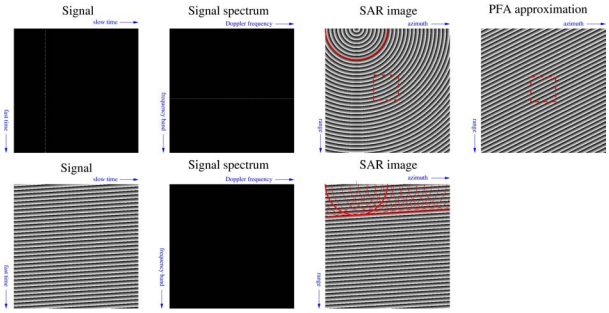

Figure 3 Principle of omega-kappa & PFA SAR synthesis illustrated with a single frequency of the bandwidth from a single pulse of raw signal (top row): SAR image is a spherical wave centred at sensor position, which locally (red square) can be approximated as a plane wave. A pure bandwidth and Doppler frequency in the raw signal (bottom row) results in overlapped spherical waves that accumulate into a conical wave — red circles figure matching wave surface from successive pulses — thus a pure image frequency. The one-to-one correspondence between signal and image spectra frequencies — the Stolt transform — is the base of omega-kappa algorithm.

As a fixed Doppler frequency corresponds to a constant phase shift from pulse to pulse, the spherical waves corresponding to successive pulses combine in conical waves corresponding in the image plane to the expected plane wave. It is possible to imagine following a given constant phase surface (sphere) from pulse to pulse, a cone wave surface emerges as the envelop of the successive spheres. (Unlike what the red overlay on figure 3 suggests, the real wave surface is 1/8 a wavelength nearer because of the contribution of the non tangent part of the spheres — which understandably quickly vanish —, this offset could be formally derived and is known as the principle of stationary phase).

For a small image centred to a given point (the centre of the complete SAR image in the example of figure 3) the spherical waves can be approximated by a plane wave (this is the basis of the polar format algorithm). This establishes the link between a point in a local spectrum in a SAR image and a slow-time position within integration and a radar frequency within the transmitted radar band. This link (of only local relevance) will be used for locally refocus the image.

1.1 Mapping of the integration interval into the space frequency spectrum

In order to derive the space frequencies in the k_X and k_Z coordinates, we need to define local image projection geometry through the vectors derivatives of the 3D position of focus point with respect to the X and Z coordinates. Indeed, without loss of generality, a SAR image (hence a *two* dimensional array) is focused on a given *surface* in the 3D space. This surface is typically the digital terrain model of the area imaged, but not necessarily (think of the deck surface of a moving ship, the wingspan of an aircraft, or the vertical plane along track for nadir imaging SAR). In the "degenerate" case of a monostatic acquisition along perfectly linear trajectory, when the Z axis is the acquisition trajectory axis, the focusing do not depend on the third dimension, but the derivation is still sound for any surface locally mapping to X and Z.

In the monostatic case, the space frequency expressed (dimensionless) in Nyquist frequency units is:

$$k_X = \frac{-2f}{c}\left(\frac{1}{||P-T||}(P-T)\right)\cdot\frac{\partial P}{\partial X}$$

$$k_Z = \frac{-2f}{c}\left(\frac{1}{||P-T||}(P-T)\right)\cdot\frac{\partial P}{\partial Z}$$

where P is the focus point and T the sensor trajectory point within the integration interval where the frequencies k_X and k_Z are evaluated. (f and c are respectively the radar frequency and wave velocity in the media.)

The more general bistatic formula is:

$$k_X = \frac{-f}{c}\left(\frac{1}{||P-T||}(P-T)+\frac{1}{||P-I||}(P-I)\right)\cdot\frac{\partial P}{\partial X}$$

$$k_Z = \frac{-f}{c}\left(\frac{1}{||P-T||}(P-T)+\frac{1}{||P-I||}(P-I)\right)\cdot\frac{\partial P}{\partial Z}$$

With I the corresponding illuminator trajectory point.

Note that the formula is linear in frequency, hence it is only useful to compute the value with f at the middle of the radar bandwidth, the full frequency band is mapped to the segment obtained by scaling (k_X, k_Z) from the origin by $1 - B_{\text{rel}}/2$ and $1 + B_{\text{rel}}/2$.
Note also the modulus of the left hand vector in the bistatic formula (further denominated "line of sight vector") is lower than 2 (the monostatic $T = I$ case excepted of course) this accounts for the "bistatic resolution loss" of $\cos(\beta/2)$ with the bistatic angle β at P.
Notice also the interchangeability of the X and Z variables corroborating the remark in the previous section on the weak link between image coordinates and effective acquisition geometry.

The $\partial P/\partial X$ and $\partial P/\partial Y$ computation, as well as the segment of trajectory (trajectories in bistatic case) included in the integration interval depend of design choices for the SAR processor.
For example, ONERA processor in stripmap mode centres the integration for a generic point P at the trajectory point(s) T (and I) such as the LOS vector makes an angle $\pi/2 - \delta$ with nominal trajectory direction V (δ defines the "squint angle" for the image). Integration interval is defined by an angle α in order to have a constant cross-range resolution (resolution along the tilted horizontal diffraction egret on figure 2 this egret direction is tilted in squinted coordinates, but is perpendicular to the vertical range egret in "zero-Doppler" square coordinates). Namely, in monostatic case, the LOS vector during integration makes an angle less than $\alpha/2$ with the LOS vector at the integration centre. In bistatic case, the angle is enlarged such as the sine of angle between LOS vectors multiplied by the LOS vector modulus equals $2\sin(\alpha/2)$ thus requiring an homogeneous cross-range resolution in the whole image (unlike the range resolution which is downgraded by the varying bistatic angle β in a non-uniform way). Our processor have also a "flash look" mode in which the integration centre T (and event. I) position is fixed for the whole image, integration is defined as in stripmap mode and the "azimuth" Z coordinates is the $\widehat{POV}$ angle (thus in monostatic case, if $O = T$ and V is the sensor flight direction at T, the image is a focused sort of "range-Doppler" image).
The computation of $\partial P/\partial X \text{or} Z$ requires the computation of P Cartesian coordinates and the slope at P of the focus surface for the image (typically the digital terrain model of the area). Our processor provides the Cartesian z-map for the image as an ancillary data, but otherwise it is possible to bracket the Cartesian z coordinate until it hits the terrain model. Derivation of the x, y Cartesian coordinates from the y Cartesian and the X, Z image coordinates is basic geometry in both stripmap and flashlook modes. For a terrain model mapped in Cartesian coordinates $x, y \longrightarrow z$ the terrain slopes $\partial z/\partial x$ and $\partial z/\partial y$ yield the derivatives of the map $x, y \longrightarrow X, Z$ from the horizontal coordinates on the terrain model to image coordinates. The inversion of this 2×2 matrix gives the $\partial P/\partial X \text{or} Z$ wanted. (Similar derivations works for vertical or moving target bound focus planes.)

2 Refocus of a SAR image around a given point

In order to illustrate the process of image refocus, we will take the extreme example of the image of figure 2 synthesised with a linear uniform trajectory with constant attitude. As the acquisition trajectories is neither uniform (aircraft velocity varies $\pm$0.6 m/s) nor linear ($\pm$40 m lateral deviation, altitude is accurate to 1 m but Earth curvature induces a $\pm$9 m vertical deviation from the straight line), the image is strongly and non-uniformly blurred by the motion errors. Non-modelled attitude fluctuations further degrade the image in amplitude and in resolution through parasitic (and adverse) aperture weighting.
Figure 4 shows the full image synthesised with linear constant attitude trajectory, to be compared with the image computed with the trajectory from aircraft navigation unit in **Figure 5** (this is 7 km at the beginning of a 60 km long X-band acquisition with 33 cm resolution on a mountainous area in the South of France). The full scale detail of the mis-focused image at the point reflector of figure 2 shown in **Figure 6** illustrate the azimuth smear caused by the quadratic phase error introduced. The linear phase error results in 63 m error in position for this echo. (note 90% of the quadratic phase originates in *non-modelled radial acceleration* and only 10% is due to along-track velocity error).

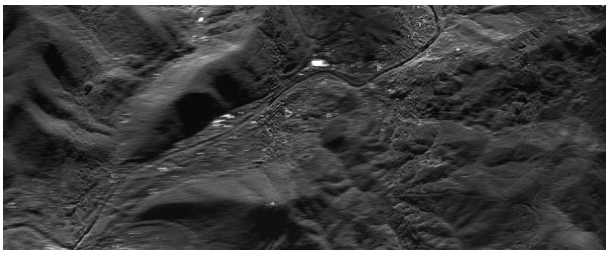

Figure 4 X-band SAR image of mountain area focused with linear uniform attitude-less trajectory. Note the uneven blur, radiometric error and the geometrical distortion.

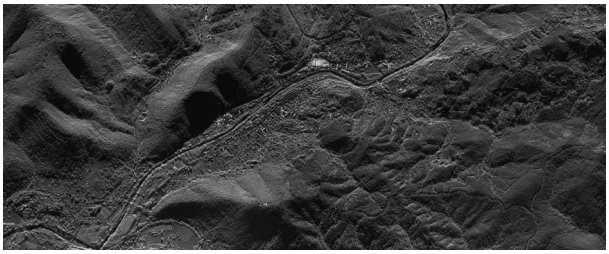

Figure 5 Reference X-band image obtained with the trajectory from aircraft navigation unit.

2.1 Localisation of the effective focus point

For refocusing the image for another trajectory (here the trajectory from aircraft navigation), we need first determine the *effective* position of the middle of the image

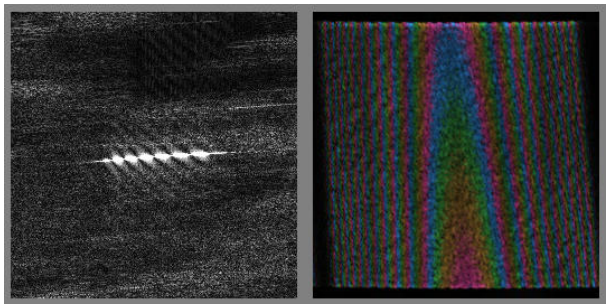

Figure 6 Detail of the mis-focused SAR image (left) and its spectrum (left) at the point echo.

patch with this new trajectory because it differs from the original focus position used for synthesis.
The calculation is the following: First from the X, Y image coordinates, we derive the original 3D Cartesian coordinates $x_{\text{org}}, y_{\text{org}}, z_{\text{org}}$ of the original focus point (as mentioned above, it may include a height search if z-map for image is not available).
Next, we determine the original trajectory points $T_{\text{org},c}$, $T_{\text{org},0}$ and $T_{\text{org},e}$ at integration centre, start and end (event. also illuminator trajectory points if bistatic). Note that in flashlook mode, the fixed integration centre $T_{\text{org},c}$ is already known.
The average vector velocity during integration V_{org} is evaluated from $T_{\text{org},0}$ and $T_{\text{org},e}$. And from $x_{\text{org}}, y_{\text{org}}, z_{\text{org}}$, $T_{\text{org},c}$ and V_{org} we compute the original range R_{org} and radial velocity $V_{\text{rad org}}$ at integration centre.

Since at first order, SAR focusing extracts from the signal energy at the range R_{org} and the radial velocity (hence Doppler frequency) $V_{\text{rad org}}$, we next search on the focusing surface (the terrain model in our case) the Cartesian coordinates $x_{\text{new}}, y_{\text{new}}, z_{\text{new}}$ of the point which have the same range and radial velocity with respect to the new trajectory. In the computation, $T_{\text{new},c}$, $T_{\text{new},0}$ and $T_{\text{new},e}$ are simply the new trajectory points at the corresponding time. Since the location has changed, this search necessarily requires an height search on the terrain model.
This search is much trickier in bistatic case than the optional one for z_{org} above: indeed, we search for a point on a surface (the terrain model) with a given bistatic range and radial velocity with respect to two trajectory points and two velocity vectors. At ONERA, we use an damped Raphson-Newton type search with a gain control law derived from the Armijo approach[1] (the very work of the 60's that was instrumental in the design of the deep learning for neural network).
Important note: the angular aperture from the new focus point with respect to the new trajectory is not necessarily centred and the resulting cross-range resolution may differ from the one required from the initial synthesis.

2.2 Local refocalisation

Once the effective focus point is determined, we need to evaluate for time t in the (original) integration interval, the varied values of the optical path difference between the original focusing $T_{\text{org}} \longleftrightarrow (x_{\text{org}}, y_{\text{org}}, z_{\text{org}})$ (or $I_{\text{org}} \longrightarrow (x_{\text{org}}, y_{\text{org}}, z_{\text{org}}) \longrightarrow T_{\text{org}}$ if bistatic) and the refocused one $T_{\text{nex}} \longleftrightarrow (x_{\text{new}}, y_{\text{new}}, z_{\text{new}})$. Note that the very definition of the effective focus point means that this path difference vanishes at integration centre t_c and has a zero average slope on the interval used for averaging the velocity vectors (in fact, for azimuth weighted images, we use half of the integration interval for averaging).
We also compute the respective centre-frequency positions in the refocused and original spectra by the formulae of section 1.1, with the extra twist that *position in the original spectrum is altered by the slope of the path difference*. Indeed, the original focusing made a wrong assumption on the Doppler value to be compensated by the processing, resulting in a biased t for the centre-frequency position in the original spectrum.
We eventually map the original to the refocused spectrum applying the path difference converted to phase for the frequency at the point in the spectrum.
The mapping is computed from the coordinate in the target refocused spectrum, time t is bracketed in the refocused position table (we compare the direction of k_X, k_Z vectors, because of the linearity in f of section 1.1 formulae). If t is inside integration interval, we compute the frequency (comparing the modulus of the current coordinates with that of $k_X(t), k_Z(t)$ gives frequency relative to the middle frequency). If that frequency is inside the radar bandwidth, we convert the path difference to phase correction at this frequency. We sample the original spectrum at the position tabulated for t scaled by the relative frequency above (for avoiding aliasing in sampling, the original spectrum is oversampled by zero-padding the local image before Fourier transform), apply the phase correction (and optionally an amplitude compensation from antenna pattern difference. This is done in the illustration examples because attitude was fixed in the original trajectory) and write the value to the current target refocused spectrum.
Local refocused image, as in **figure 7**, is just obtained by reverse Fourier transform of the refocused spectrum.
It is the image square of figure 6 refocused on which the middle point echo is perfectly refocused (Ignore the lower phase slope in the spectrum, it just means that -by chance- the echo falls much closer to a pixel centre than on the figure 5. This is irrelevant as the geometrical distortion is *not* compensated here.)

3 Global refocus and rectification of a SAR image

From a local refocus as illustrated in figure 7, to a complete refocused and rectified (i.e. distortion corrected), there are some more technicalities. First, as clearly visible in figure 4 focus blur varies with position in the image (mostly with azimuth). This imply that the refocusing described above is only valid for a small area around the centre point. Second, because the focus blur smears the point response along track as illustrated in figure 6, the image patch input to the refocus must contain the full smear width for

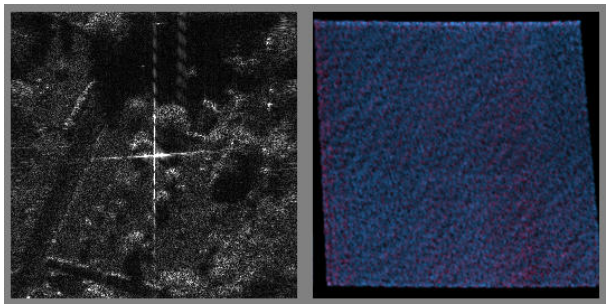

Figure 7 Detail of the locally refocused SAR image (left) and its spectrum (left) at the point echo. Note the smear truncation artefacts at the left and right edges.

the valid part of the refocused patch. This implies that the patch size must include margins (mostly in azimuth) around the valid part eventually used. (The figure 7 is a screen copy of an interactive tool with a fixed 256×256 window, since the point smear is about half of window width, the points on the 1/4 left and right extremities are missing up to half of their Doppler frequencies hence a coarser resolution compared with figure 2. There is also some azimuth wrapping e.g. the tree on the left side.)
The size of the smear margin and the maximum mesh size can be evaluated by sampling across the image the maximum phase slope of the path difference curves (for the smear) and the maximum space variation of the phase values in the path difference curves. In **figure 8**, the margin was 224 pixels (some area have smearing twice bigger than that at our example point echo) and the mesh size 32×256 pixels, refocused patches are smooth tiled with a range and azimuth space domain Hann window.

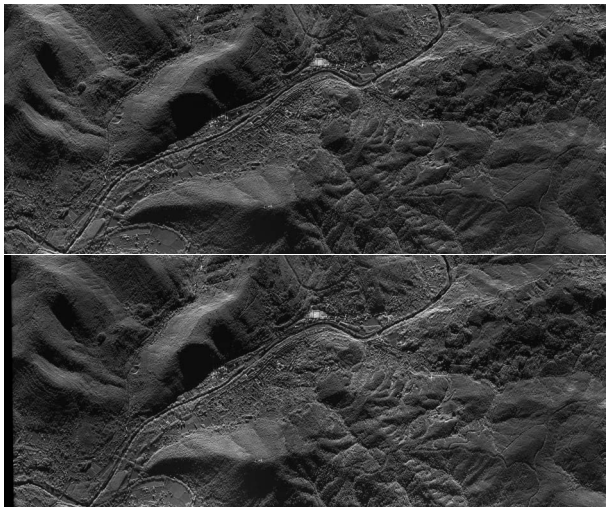

Figure 8 SAR image synthesised with linear constant attitude and refocused on 54×683 overlapping tiles (top) and rectified for geometrical distortion (bottom)

Compared to the SAR image focused directly with the aircraft navigation unit trajectory, the image is well focused but its geometrical distortion is uncorrected (it is a consequence of the fact that the path difference corrected has zero value at origin and zero average slope). The last operation is a coherent remapping of the refocused image to the final geometry. This is done again by smooth tiling the image, 0-pad interpolate the patches and resampling the final position value (through $X_{\text{new}}, Z_{\text{new}} \longrightarrow x_{\text{new}}, y_{\text{new}}, z_{\text{new}} \longrightarrow x_{\text{org}}, y_{\text{org}}, z_{\text{org}} \longrightarrow X_{\text{org}}, Z_{\text{org}}$) in the interpolated refocused patch. The only caution is that, unlike when we map a native SAR image, the patch zero-padding must be made around the refocused $k_X(t_c), k_Z(t_c)$ spectrum centre.

4 conclusion: limitations and usage

Of course, such an extreme refocus as illustrated here requires large margins and thin tiling, thus the computation volume is about ten times bigger than simply redoing the SAR synthesis with the new trajectory, however this may not be the case for smaller correction such as from real-time navigation and post-processed trajectory especially on a limited region of interest. It can alternatively refocus for nearrange bias (radar delay error), target elevation above the terrain model or bistatic clock drift (through range time-proportional correction and radial velocity biasing).
Of course, this method is an asset for refocusing SAR image of which raw signal is not available. For example, if processing is done on the platform and only the synthesised image is transmitted to the ground. It proved also useful for simulating a landscape image from a different viewpoint with a realistic spectrum (in that case, the starting image is of higher resolution, and the final remapped image spectrum is trimmed to the simulated sensor one).

5 adendum: parallel and difference with motion compensation

Prompted by a question from a reviewer, I shall emphasise on the fact that this method is *not* intended to perform motion compensation in the first place. As mentioned in the conclusion above, using this method to focus an image synthesised without motion compensation (using the basic omega-kappa algorithm for linear uniform trajectory illustrated in figure 3) results in a computation some ten times longer than what is needed when using a motion compensated omega-kappa algorithm directly. [4] describes the motion compensated omega-kappa algorithm routinely used at ONERA which produced the image in figure 5. Incidentally, the very same program was used to emulate a *non* motion compensated omega-kappa processor (thus producing image in figure 4) by providing it a linear uniform average trajectory instead of the true non-linear varying velocity trajectory from the aircraft navigation unit. It worth mentioning that early omega-kappa implementations (as described in [3]) was *not* motion compensated, and second order motion compensation was perform by autofocus postprocessing (thus *intentionally discarding the known nonlinearities* of the trajectory!), but modern implementations of omega-kappa as ours can derive the motion compensation from the trajectory.
This method is intended to local or regional image updat-

ing for trajectory, focus altitude or radar parameters *minor* changes. The need that motivated this technique was the real-time processing of SAR images onboard the carrier aircraft. Due to restriction to civilian navigation units, the real-time trajectory is deliberately dithered by the INU firmware, and the accurate trajectory is only available in post-processing some minutes later. The idea was to process the signal during acquisition with the real time trajectory, thus obtaining a SAR image with moderate defocus blur, and later refocus it with the actual accurate trajectory when it becomes available. With a moderate refocus blur, the sub-blocks overlap could be small, and due to the low frequency of the introduced trajectory dither, relatively wide sub-block could be used, resulting in a non dissuasive computation cost.

On the opposite, the extreme case illustrated in this paper, with its wide sub-block (for covering the wide smearing of echoes by the focus blur), short sub-block valid region and large overlap makes it practically irrelevant, it is just a showcase.

Though there is a parallel between motion compensation during omega-kappa SAR synthesis and refocusing image post-processing the main difference is the input tile size and overlap which have an huge impact on performances.

As in post-processing, initial motion compensation consists in adjusting the optical path length during integration (and this adjustment varies from point to point in the image): Locally, zeroth order correction (constant offset) amounts to a range/fast-time translation (a.k.a. ”range migration“), first order correction (slope) amounts to a slow-time/azimuth translation (a.k.a. ”azimuth migration“) and the remaining (quadratic and higher order terms) amounts in azimuth smearing (improperly a.k.a. ”quadratic phase“). In the SAR processor **figure 9** the azimuth smear due to unevenness in trajectory deviation or its projection in the target imaged point is pre-compensated in the signal in each tile sub-bloc. That is why the tile processing is done on a wider buffer (because applying a quadratic phase smears the tile in azimuth). However, the tile is generally much shorter than the integration time (in order to follow unevenness in time of the compensation) and a buffer of twice the tile size (the minimum buffer size for our processor that uses only power of two sizes for Fourier transforms) accommodates for effective Doppler fluctuations of half the Doppler ambiguous frequency of the native signal (this is beyond all practical acquisition cases). Hence the motion compensation does not involves excessively large Fourier transforms. In practice, its cost is roughly the same as the (optional) pulse compression/Doppler decoupling and Doppler trimming first stage of the processor.

On the opposite, in the refocus module of the image post-processor **figure 10** the tile azimuth length should be larger than the defocus smear length + the typical trajectory fluctuation length *at the final pixel size* (e.g. 512 pixels in figure 8) while providing only the refocused image portion along azimuth for the typical trajectory fluctuation length (32 pixels in our example). Furthermore, for smoothing the transition between refocused tiles, the middle valid part of the tile is weighted and the adjacent tile offset is half the valid azimuth length for the tile. Hence, even neglecting the fact that narrow-band compensation only requires 1D Fourier transform in the SAR processor and the further need for zero padding for mitigating the aliasing in the spectrum resampling, the postprocessor yield is at most 1/16 of that of the SAR processor (this is consistent with observed processing times).

However, in the case of the post-correction of a moderate (a few pixels) defocus blur, the azimuth offset between tiles can be close to 50% and the range offset nearly 100% hence the post-processing yield can outperform the abnihilo SAR re-synthesis.

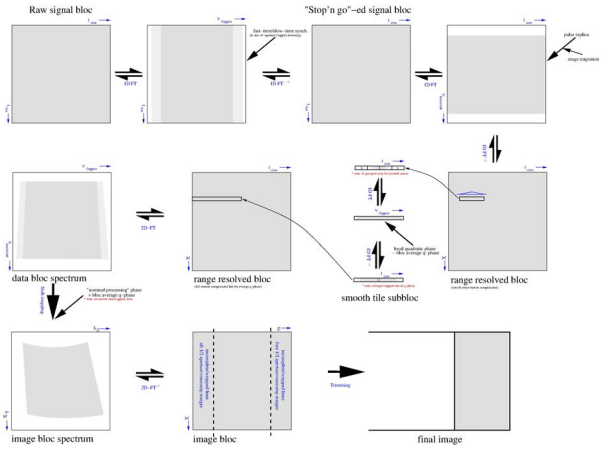

Figure 9 Principle of the motion-compensated omega-kappa algorithm implementation at ONERA.

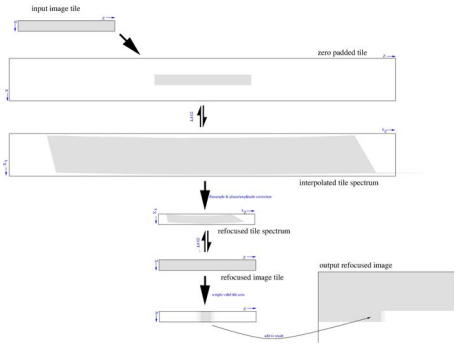

Figure 10 Principle of the refocus module of the image post-processor for comparison.

6 Literature

[1] Armijo L.: Minimization of functions having Lipshitz continuous first partial derivatives, Pacific J. of Math. Vol. 16, No. 1, Jan. 1966, pp. 1-3

[2] Stolt R.H.: Migration by Fourier transform, Geophysics, vol. 43 no. 1, Feb. 1978, pp. 23-48

[3] Soumekh R.: Synthetic Aperture Radar Processing, R.Wiley-interscience, 1999

[4] Cantalloube H. M. J.; Nahum C. E.: Airborne SAR-efficient signal processing for very high resolution, Proceedings of the IEEE, Vol. 101, No. 3, 2013, pp. 784-797