A MOTION COMPENSATION STUDY FOR THE PHARUS PROJECT

by
M.P.G. Otten, M.Sc.
Radar & Communications Division
Physics and Electronics Laboratory
P.O. box 96864
2509 JG The Hague
The Netherlands

SUMMARY

In the PHARUS project, a polarimetric C-band SAR is being developed, which will be preceded by a non-polarimetric test system called PHARS. A motion compensation study is also part of preparatory studies for the final PHARUS design.

A SAR data simulator has been developed as a tool for the investigation of the effects of aircraft motions on the SAR image. From the SAR mapping geometry, a terrain description, the radar parameters, and detailed trajectory and attitude data of a non-manoeuvring aircraft, the simulator generates raw data with a given range resolution. This can be processed, by azimuth compression, into the SAR image.

A secondary purpose of the simulation is to determine the impact of several design parameter choices, and to provide well-defined test-input for SAR processing software that is being developed.

The results of test runs with real flight data have been verified theoretically, and have shown the need for motion compensation. It has also demonstrated a major advantage of simulation, in that it can take many factors into account at the same time, including for instance the SAR processing method, which is hard to do theoretically.

LIST OF SYMBOLS

A	: antenna aperture (area)
a_{LOS}	: line-of-sight acceleration
B_w	: bandwidth
D_{az}	: azimuth shift
$\underline{d}_{la}$	: lever arm distance
H	: altitude
F_n	: receiver noise factor
$f_{s,eff}$	: effective sampling frequency
Δf_d	: utilized Doppler band
κ	: weighting compensation parameter
N_{pr}	: azimuth presumming factor
PCR	: pulse compression ratio:
P_t	: transmitted peak power
R	: slant range
RCS	: radar cross section
ρ_a	: azimuth resolution
r_r	: (slant) range resolution
T_a	: integration (aperture) time
V_{LOS}	: line-of-sight velocity
Φ_{err}	: phase error
λ	: wavelength
θ_g	: depression angle
θ_s	: azimuth beam steering angle

1. INTRODUCTION

The PHARUS project [1] aims at the development of a polarimetric C-band aircraft SAR in the Netherlands, and is carried out by the Physics and Electronics Laboratory TNO (FEL), the National Aerospace Laboratory (NLR), and the Delft University of Technology (TUD). Prior to the realisation phase of the project, i.e. the building of PHARUS and final software development, there is the definition phase, consisting of several preparatory studies, one of which is concerned with motion compensation. FEL and NLR collaborate on this study.

In the definition phase, a test system is also built, called PHARS, so that actual experience with most aspects of SAR can be gained before the final design for PHARUS is made.

The first phase of the motion compensation study was the development of an aircraft SAR simulation program at FEL, called SARGEN. This simulation has recently been taken into use by NLR to carry out the second phase of the study, the motion error analysis.

This paper focuses on phase one of the study, the development of the simulation program SARGEN, and some first results of simulations with 'real' flight data that were made available by NLR.

2. SAR SIMULATION AND MOTION COMPENSATION

Though motion compensation studies can be carried out purely theoretically, a simulation has a number of advantages.

In the first place, practically any investigation that can be carried out theoretically can also be done with a simulation. In fact, the first thing to do would be to validate the simulation with a number of theoretically verifiable test runs. However, the simulation can handle situations which soon become too complicated for theory: all types of errors can be considered, motion can be singled out and combined as desired, and real flight data can be used directly: motion variables will usually not be independent, so that considering all of them 'at the same time' is more realistic than incoherently adding up an error budget. Furthermore, the particular azimuth compression method is taken into account, secondary effects which might otherwise be overlooked will become apparent in a simulated image, and the image quality can be judged in any objective and subjective way.

Apart from a motion compensation study, a simulation can also aid the making of design choices, e.g. by determining ambiguity levels or signal-to-noise ratios, and it can generate known test input data for SAR processing software under development.

3. AIRCRAFT SAR SIMULATION 'SARGEN'

The primary purpose of the simulation is to study azimuth distortion and defocusing caused by aircraft motion. The necessary capabilities and permissible limitations of SARGEN have been established according to this aim. The following simulation model was considered to be adequate:

A situation is simulated where a non-manoeuvring aircraft flies along a trajectory, which is perturbed by motion deviations in all directions and in attitude (roll, pitch, and yaw). The deviations are assumed to be not so large as to cause range walk, but arbitrary otherwise. A certain slant range resolution achieved by pulse compression is assumed, but the pulse compression itself is not simulated. The motion reference point may be displaced with respect to the antenna phase centre. The pointing of the antenna is arbitrary, but fixed during one simulation run. The radar PRF may be fixed or coupled to the aircraft ground velocity. An azimuth presumming factor may also be specified.

The imaged terrain may consist of any collection of point targets in a two dimensional plane, each given by a position, RCS, and phase term[1] (complex reflection coefficient).

Thus the output of SARGEN represents range compressed, and presummed coherent raw data, which then needs azimuth compression to form the SAR image.

Figure 1 shows a block diagram of the complete simulation, i.e. up to the final image, and the organisation of in- and output data.

---

[1] Only relevant in case of closely spaced point targets

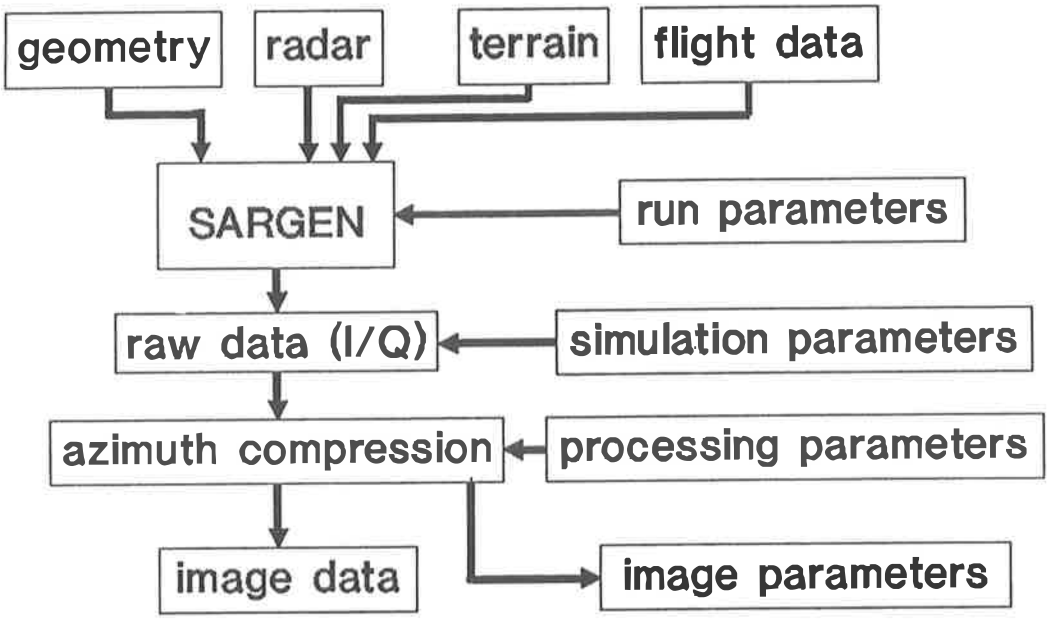

Figure 1. Simulation block diagram

The input parameters are organized in five data sets:

- geometry : parameters which describe fixed distances, like nominal altitude, distance to the imaged swath, displacement between motion reference point and antenna phase centre etc.
- terrain : description of the point target collection
- radar : wavelength, PRF, power, bandwidth, range resolution, antenna dimensions etc.
- flight data: six time sequences: three position- and three attitude variables, plus the time itself (time step is variable).
- run parameters: describes the simulation run itself, i.e. the extent in azimuth and range, and controls the selection of flight variables, and of variable or fixed PRF.

A file of raw data is generated, along with a 'summary' of the simulation parameters, some of which are necessary input for the azimuth compression (velocity, PRF, wavelength, range etc). The processing parameters specify azimuth resolution, weighting parameters, and the image format: size, pixel spacing etc., which are again summarized after processing in the image parameter file, along with the resulting number of lines and pixels.

The simulation has been implemented in the FORTRAN 77 language. Since a simulation such as this can easily become excessively time consuming, a fair amount of effort has been devoted to optimization of the computations. Nevertheless, some limitations remain: study of statistical phenomena (speckle) would require, in this set-up, collections of point targets that are too large to handle, but such investigations were not intended.

Figure 2 shows how, ideally, the simulation would be used in a set-up for motion compensation study, and how, in some cases, this may be approximated by a simpler set-up, which requires only straightforward uncompensated SAR processing. Note however that this alternative set-up presupposes that the motion that is assumed to be compensated is indeed compensated in every respect, which is not always true. For instance: phase errors may be perfectly corrected, but deterioration of the signal-to-noise ratio due to antenna misalignment cannot be undone. Such partial compensation of motion cannot be simulated in the alternative set-up. The alternative approach should therefore only be used, bearing its limitations in mind. Furthermore, since the raw data generation will often be more time consuming then the subsequent SAR processing, it may even be more practical to use the first approach, if an accurate set of flight data is available. One could generate a set of raw data using that set of flight data, after which processing can be done with motion compensation with any desired error imposed on the flight data.

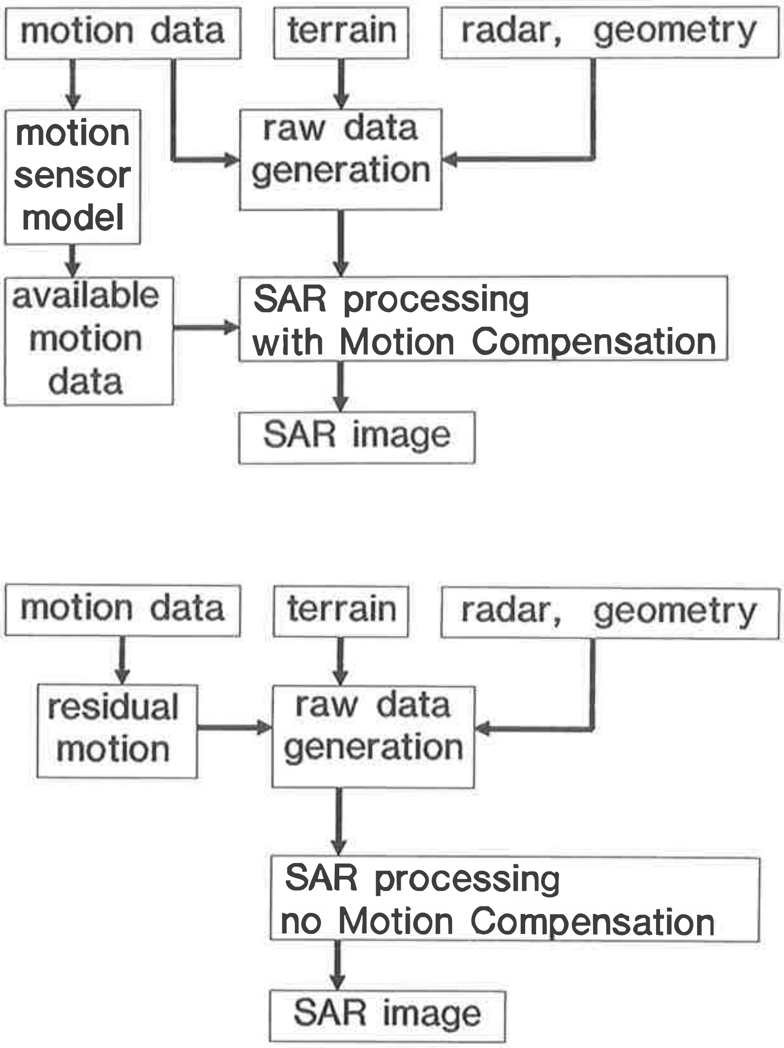

Figure 2. Simulation set-up for motion compensation study
Top : 'ideal' set-up
Bottom: simplified set-up

In principle, autofocus can also be included in a simulation cycle. However, something would have to be done about the statistical properties of the simulated data, to which autofocus algorithms are inherently sensitive. Adding speckle with the proper statistics to the image data, without simulating the underlying physical process, is therefore considered as a future extension.

4. SIMULATION EXAMPLES

Some test simulations were carried out using flight data of an NLR aircraft, a Swearingen Metro II (Fairchild), which is used for Remote Sensing experiments, and which may in future serve as the SAR platform. A fictitious scene was created, described below, containing enough features to reveal any distortion, defocusing, and ambiguities:

Point targets are arranged in cross track lines, at 5 m intervals (less than the range resolution), along track lines, at variable intervals, and one 'diagonal' line. All point targets have RCS=1 m^2, except those constituting the diagonal line, which have RCS=100 m^2, see figure 3. The simulated trajectory is about 600 m at 104 m/s ground speed. The actual flight data are depicted in figure 4 (see figure 5 for roll,pitch, and yaw definitions). Note: the x-position (along track) is the x-position after subtraction of the nominal speed.

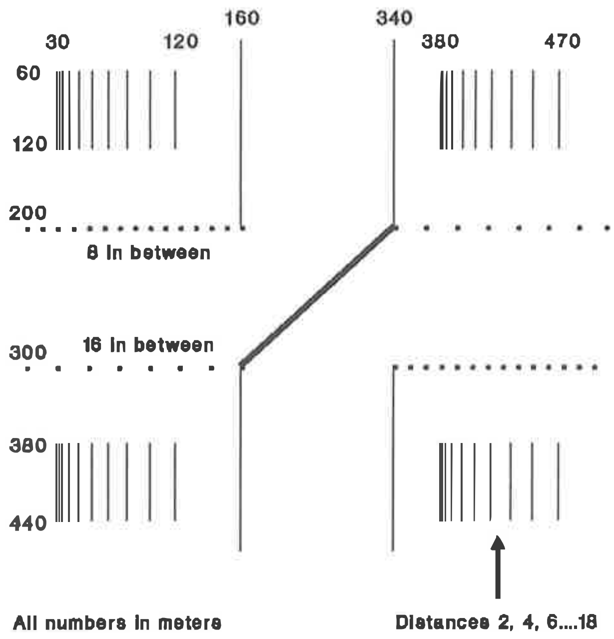

Figure 3. Simulated scene

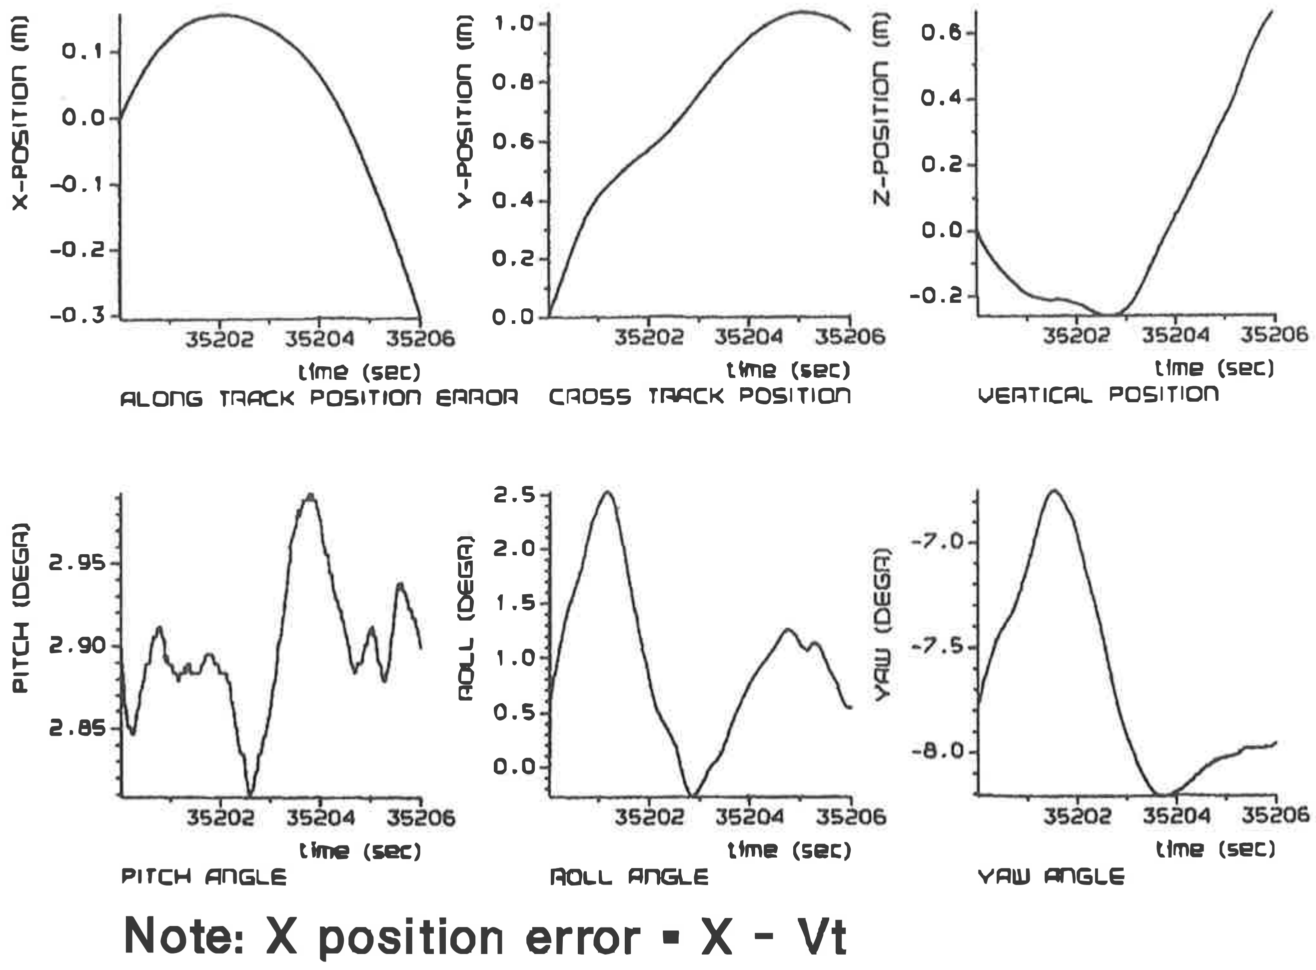

Figure 4. Flight path data over a 6 second interval

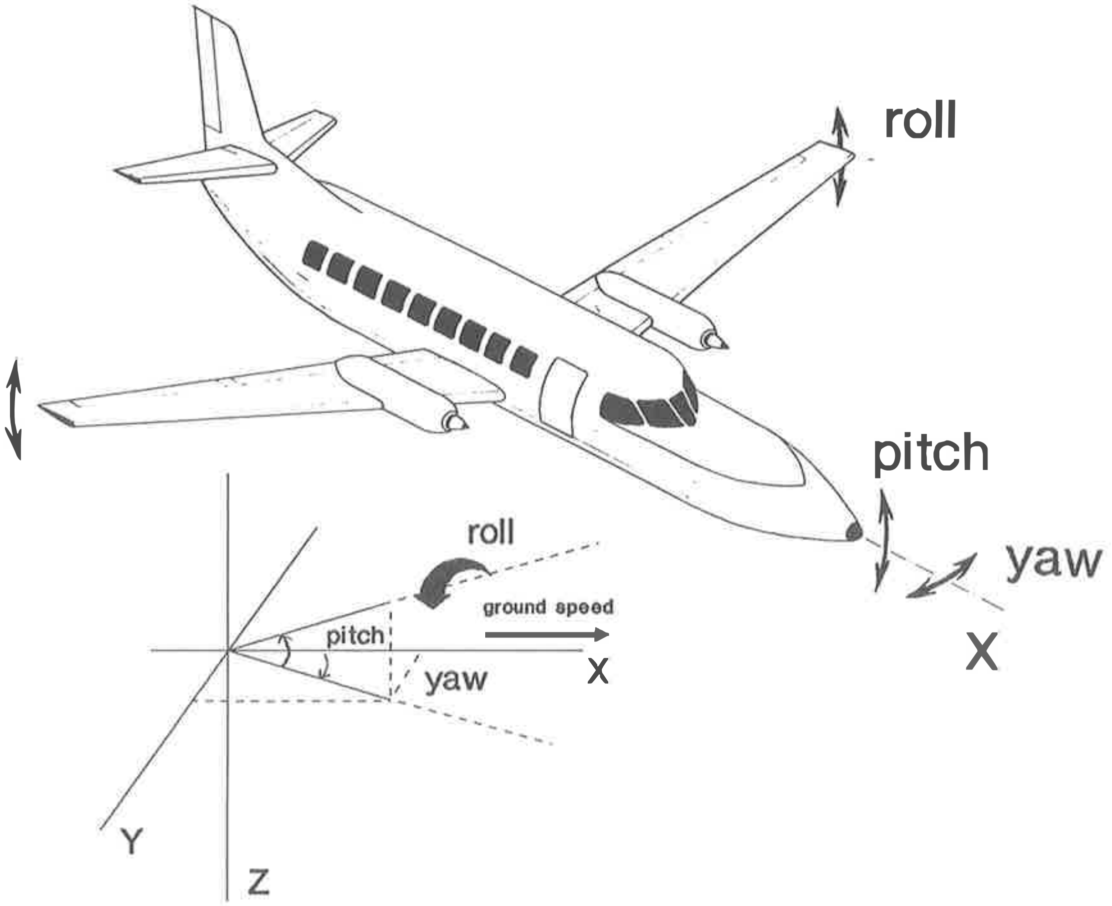

Figure 5. Roll, pitch, and yaw definitions
Top : Roll, pitch, and yaw axes of the airframe
Bottom : Relative orientations with respect to ground speed

Other simulation parameters:

- near slant range — $R = 13$ km
- altitude — $H = 6$ km
- antenna pointing:
  - depression angle: — $\theta_g = 24.4°$
  - azimuth beam steering angle: — $\theta_s = -5.5°$
  - except in the uncorrupted image where — $\theta_s = 0°$
- wavelength: — $\lambda = 0.0566$ m
- PRF : coupled to velocity, — PRF = 3500 Hz, at 100 m/s
- antenna aperture: — $A = 0.35$ m x 0.135 m
  (uniform illumination assumed)
- transmitted peak power: — $P_t = 140$ W
- receiver noise factor: — $F_n = 2$ dB
- bandwidth: — $B_w = 31$ MHz
- pulse compression ratio: — PCR = 400
- range resolution — $\rho_r = 6$ m
- azimuth presumming factor — $N_{pr} = 16$
- lever arm distance — $\underline{d}_{la} = (5.0, 1.0, -0.5)$ m

These parameters - except the lever arm distance, which is not known yet - are generally representative of the actual PHARS parameters.

The azimuth compression is carried out using line-by-line (time domain correlation) and Doppler processing ('deramping'+FFT). In both cases the azimuth resolution ρ_a is 6 m (2 pixels/resolution cell), and raised cosine weighting is applied to achieve -30 dB sidelobes.

Figure 6 shows the uncorrupted SAR image (line-by-line processed). As can be seen in the vertical line structures, the actual resolution is not as good as the theoretical 3-db resolution. This is because a resolution cell only contains 2 pixels, so that a small dip between two point target responses cannot be seen. Around the diagonal, formed by 100 m^2 point targets, a light area is visible: this is the result of the sidelobes (of the synthetic beam) rising above the noise. The maximum signal-to-noise ratio, for the 1 m^2 targets, is found to be about 30 dB, after processing: in the 'uncorrupted' case, the antenna is perfectly aimed, so that the maximum gain is obtained. Since the sidelobes of the synthetic beam are about -30 dB, they do not clearly show, for the 1 m^2 targets, but they do for the 100 m^2 RCS targets, which of course are 20 dB higher.

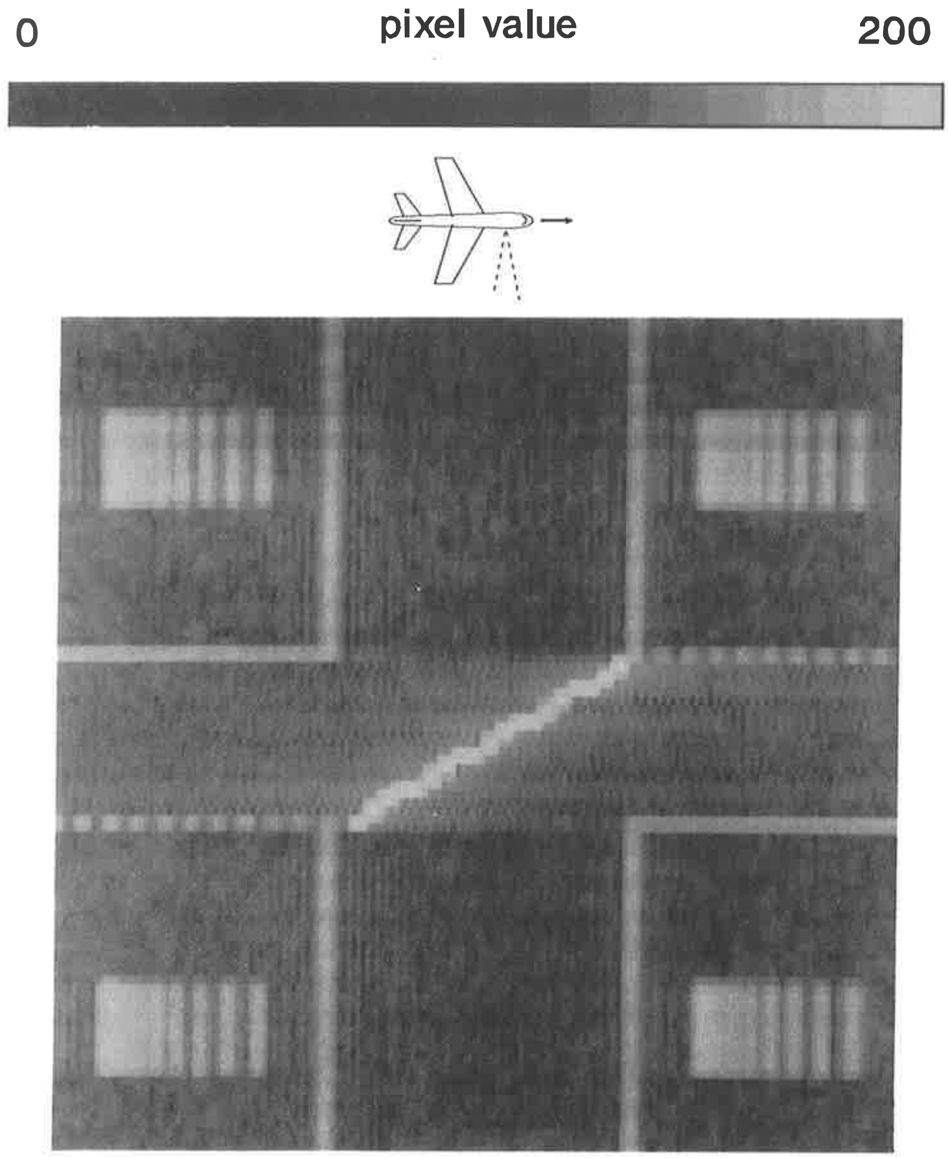

Figure 6. Uncorrupted simulated SAR image

Figure 7 shows the image that results when all motion variables as shown in figure 4 are included, and when the antenna and motion reference point coincide (no lever arm). Since the PRF is coupled to the velocity, the x-motion error is already compensated, and the beam steering angle compensates for part of the average yaw, so that there is little overall gain loss.

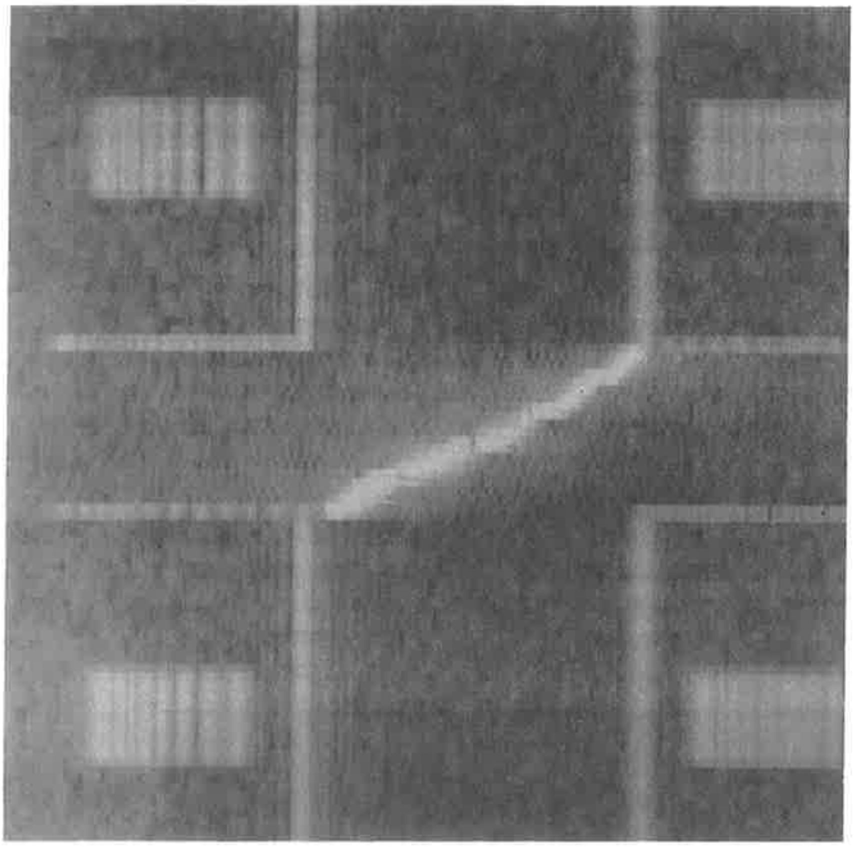

Figure 7. SAR image with motion errors (no 'lever arm' effect)

To get a feel of the degradation that might be expected, some approximations of image shift and defocusing can be made.

To determine azimuth shift, the linear motion components (constant velocity) in the y- and z-direction should be estimated.

The steepest part (beginning of trajectory) of the y-motion corresponds to about 0.45 m/s. The resulting line-of-sight velocity is:

$$V_{LOS} \approx V_y \bullet \cos(\theta_s) = 0.4 \text{ m/s} \tag{1}$$

so that the shift D_{az} at 13 km range is:

$$D_{az} = V_{LOS} \bullet R \;/\; V_x = +50 \text{ m} \tag{2}$$

The maximum z-velocity is 0.3 m/s (second half of trajectory). Along the line-of-sight this becomes:

$$V_{LOS} \approx V_z \bullet \sin(\theta_s) = 0.12 \text{ m/s} \tag{3}$$

yielding an azimuth shift of

$$D_{az} = 15 \text{ m} \tag{4}$$

As for defocusing: A quadratic phase error can be estimated by considering the average cross track acceleration during one aperture time (T_a). T_a is given by:

$$T_a = \kappa \bullet \lambda \bullet R \;/\; (2 \bullet V_x \bullet \rho_a) \approx 0.7 \text{ s.} \tag{5}$$

(κ is 1.1 because of the weighting).

Closer examination of the flight data (not shown here), reveals that the maximum z-acceleration averaged over 0.7 s, is about 0.4 m/s^2, in the middle of the trajectory. The averaged y-acceleration has an extremum of -0.45 m/s, but this occurs at the beginning of the trajectory where the scene happens to be 'empty'. A more typical value is 0.15 m/s^2.

An average cross track acceleration along the line-of-sight of a_{LOS} yields an edge-of-aperture phase error of:

$$\Phi_{err} = (2\pi/\lambda)\ a_{LOS}\ (\tfrac{1}{2}T_a)^2 \qquad (5)$$

A commonly applied limit for this kind of error is $\frac{1}{2}\pi$. The z-acceleration has a line-of-sight component of 0.16 m/s^2, so that $\Phi_{err} = 0.7\pi$. Noticeable defocusing may be expected from this. An y-acceleration of 0.15 m/s^2 leads to $\Phi_{err} = 0.6\pi$. These values that are 'just over the limit' confirm the impression of figure 7: the defocusing is not totally destructive but clearly visible. These are, however, rough approximations. The actual motion contains higher order phase errors, which haven't been considered. Furthermore, when the azimuth compression is done line-by-line, the actual spread of a particular point target response in case of constant cross track acceleration, is determined not only by the quadratic phase error across the aperture but also by the changing orientation (linear phase component) of the aperture itself, since this has the effect of keeping the target in the synthetic beam for a longer or shorter time, depending on the direction of the acceleration. This does not apply to Doppler batch processing, where a batch of pixels is formed from one aperture.

Figure 8 shows the resulting image, when Doppler processing is applied. Reasons for the obvious differences are:

In Doppler processing, a batch of image points is formed from one batch of raw data samples. Therefore, all pixels in one batch will be influenced by the same phase errors across the aperture, and the degradation is practically the same for all pixels belonging to this batch (unless motion errors are extremely large). This causes discontinuities in azimuth from batch to batch (batch seam). It also results in occasional 'double imaging' or disappearance of parts of the scene. Furthermore, the integration length, and thus the location of the batch edges in azimuth, must vary with range. So, points at the same azimuth location, but slightly separated in range, may belong to different batches, and therefore suffer from different degradation effects: this leads to discontinuities in range as well. In general one can say that, for a given trajectory, the response to any particular point target depends on the type of processing, where for Doppler processing, this response is hard to predict. It would be even more difficult to predict, by other means than simulation, how these differences influence the performance of autofocus algorithms.

Adding the lever arm displacement, some more changes appear, see figure 9. In the left section, the yaw rate goes from positive to negative, causing a shift which also goes from positive to negative, so that the left part of the image has shrunk. For example, a yaw rate of -1.4°/s gives rise to a y-velocity of -0.12 m/s, which causes a -14 m shift. Defocusing is hard to predict, since the variations can not be very well approximated by quadratic errors. For instance, the rather sharp change of the roll rate at two time instants, shows up as a 'breaking up' of the upper horizontal line segment in the left section, and of the diagonal in the middle.

As for gain variations: these are small, and cannot be detected visually[2]. There is some gain variation, however, mainly due to yaw, not exceeding 1 dB.

[2] Unfortunately, some apparent large scale intensity fluctuations appear, due to imperfections of the photograph.

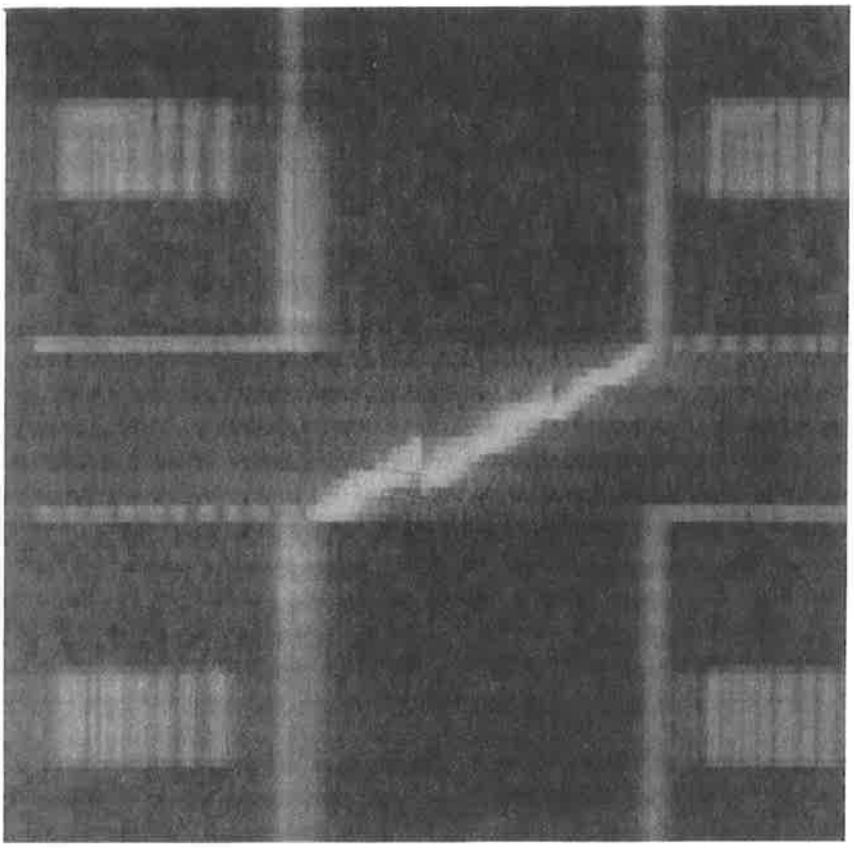

Figure 8. Doppler processed SAR image with motion errors

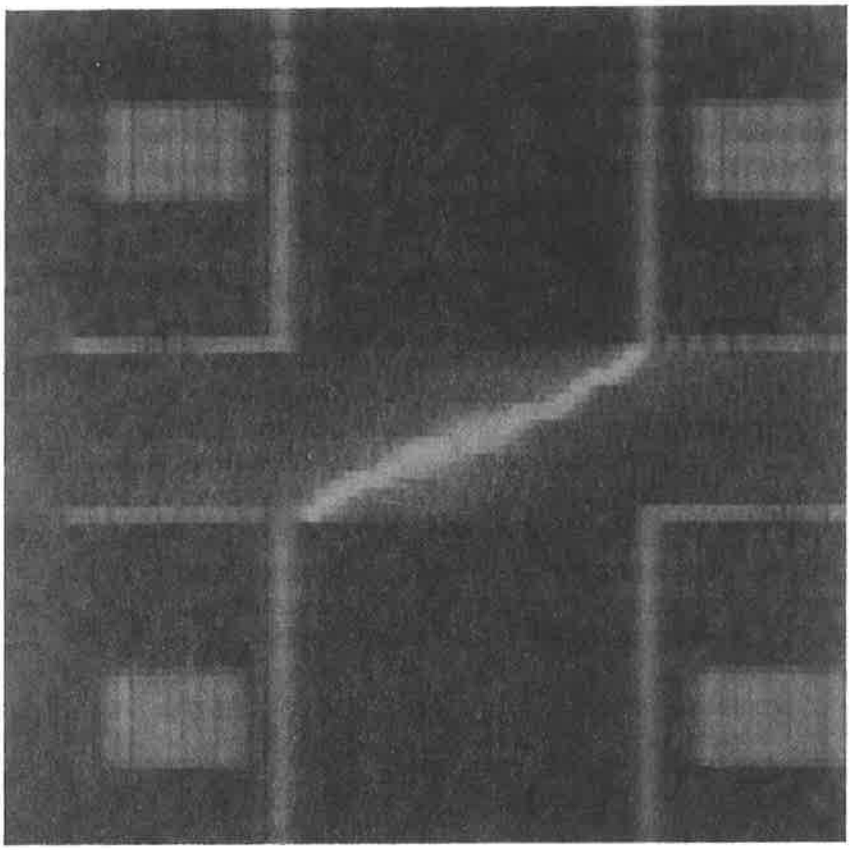

Figure 9. Simulated SAR image with motion errors
including 'lever arm' displacement

The simulation may also be used to determine suitable processing parameters. For instance, when Doppler processing is applied, some frequency margin has to be taken into account to accommodate the ambiguity inherent in a discrete Fourier transform of a not strictly bandlimited signal. This margin can be manipulated by resampling between the focusing and transformation operations[3]. Narrowing the margin makes the computation more efficient, but increases the ambiguities. Figure 10 shows one example of a too narrow frequency margin: the FFT sampling frequency was reduced by a factor of four <u>after</u> the focusing operation. Note that the imaged Doppler frequency range Δf_d is still only a third of the effective sampling frequency $f_{s,eff}$:

$$\Delta f_d = V/\rho_a = 104 \;/\; 6 \approx 17 \text{ Hz} \tag{7}$$

$$f_{s,eff} = 3500 \cdot (104/100)/(16 \cdot 4) \approx 57 \text{ Hz} \tag{8}$$

Though ambiguity levels for such processing schemes could be determined analytically, this is quite difficult, among other things due to the combination of presumming <u>before</u> and averaging <u>after</u> focusing, or 'deramping' of Doppler frequencies. Presumming/averaging can be described as ordinary filter operations, but deramping cannot, since it is not time invariant. In fact, the image clearly shows that the ambiguity varies, due to the varying position in azimuth of the strong targets. In such a case, a simulation provides a quicker way of finding the answers.

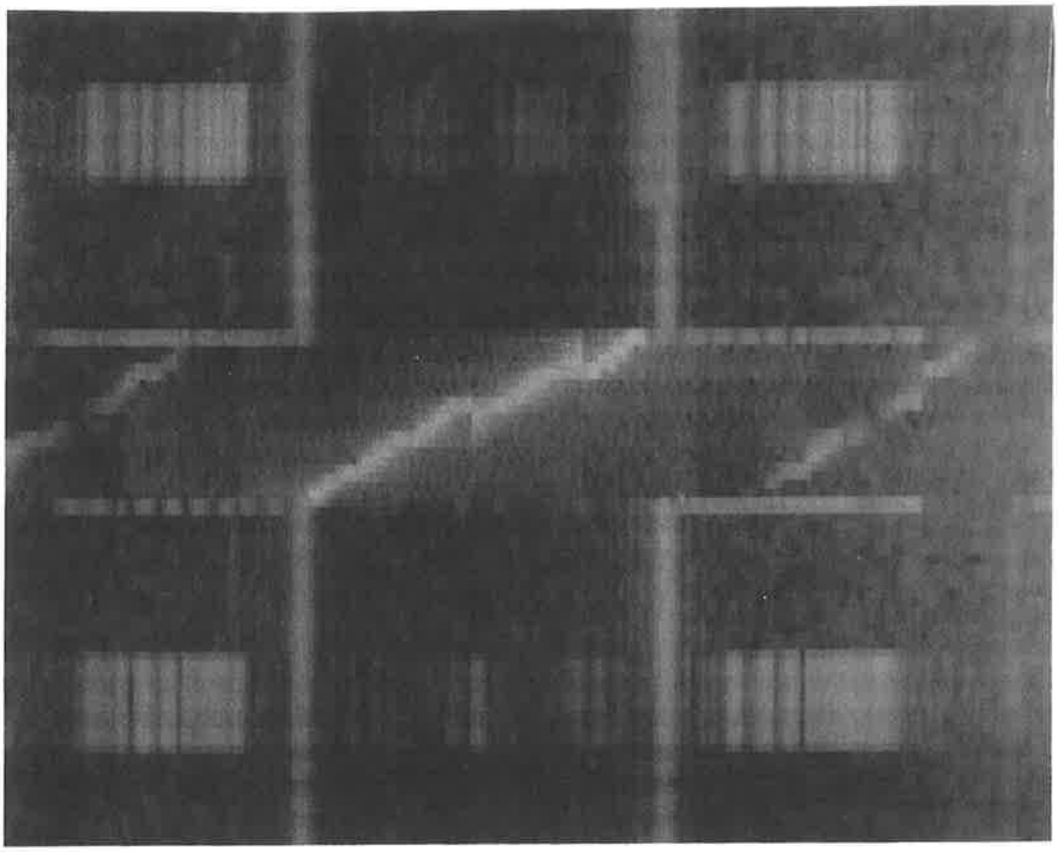

Figure 10. Ambiguities in a Doppler processed SAR image

[3] This yields less degradation than downsampling <u>before</u> focusing

5. CONCLUSIONS

The aircraft SAR simulation is a very useful tool for the PHARUS motion compensation study. Some simulation examples not only show the theoretically verifiable effects, but also some effects, that can be explained qualitatively, but are difficult to quantify theoretically, especially with respect to Doppler processing.

Additional advantages are the possibility to quickly determine the final impact of many design parameters, and the possibility to generate well-defined test data for SAR processing software.

Adding simulated speckle is considered as a future option, when autofocus algorithms are to be investigated using simulated data. Other extensions that are considered are FIR prefiltering, in stead of plain presumming, and in-flight beam scanning.

REFERENCES

[1] P. Hoogeboom
THE PHARUS PROJECT
Paper No.10 of the AGARD Symposium 'High resolution air- and spaceborne radar'
May, 1989, The Hague, The Netherlands.

[2] M.P.G. Otten
Physics and Electronics Laboratory TNO, The Hague
AIRCRAFT SAR SIMULATION 1.0
report FEL-1989-44

DISCUSSION

G.E.Haslam

You have noted that the ambiguity level is dependent upon the processing block size. Is not this phenomenon due to the phase modulation of the spectrum under the presum filter? Please clarify.

Author's Reply

No, the variation in ambiguity is due to down sampling being carried out before FFT-ing and after focusing to the centre of a batch. This down sampling is meant to avoid excessive redundancy, i.e. imaging a much larger Doppler frequency range than the frequency range of interest. Since the focusing is done with the full (original) sampling frequency, pixels in the middle of a batch do not suffer from increased ambiguity by subsequent integration, but pixels near the edge do. So, in fact, the ambiguity level depends on the position of a pixel in a batch. Of course, the effect depends on the reduction of the sampling frequency, which determines the FFT output frequency range.

R.Horn

(1) What kind of method is intended to be used for motion compensation in the PHARUS system?

(2) Is it intended to implement real time motion compensation on board the aircraft?

Author's Reply

(1) The M.C. study is currently in progress: it still has to be decided what method will be required. However, it is very likely that flight path data will be recorded on tape for off-line correction, supplemented by some autofocus procedure. Requirements for the inertial unit still have to be determined, according to the outcome of this study. Also, several autofocus algorithms will be evaluated in this study.

(2) No real-time correction is foreseen in the near future.