



Flood mapping using the Sentinel-1 imagery and the ESA SNAP S1 Toolbox

Document Version 1.0
Date of publication: 5 October 2021

Luigi Selmi

Flood mapping using the Sentinel-1 imagery and the ESA SNAP S1 Toolbox © 2021

by Luigi Selmi is licensed under CC BY-NC 4.0 

Introduction

According to the World Health Organization ([WHO](#)), “between 1998-2017, floods affected more than 2 billion people worldwide”. Even if there has been a decline in the number of deaths due to natural hazards in the last decades, scientists do expect an increase in the number of such events in the future due to climate change [1]. In addition to the death toll, floods cause heavy damage to buildings and infrastructures that amount to huge [economic losses](#). The Copernicus Sentinel-1 mission is a constellation of two satellites that provides Earth observation data in the microwave region of the spectrum that is not affected by the weather conditions, nor it depends on the sunlight. For these reasons it is used successfully to monitor flood events. In order to exploit such advantages, the user of the Sentinel-1 data needs to know the technical characteristics of the products delivered by the ESA, the radar technology and how microwaves interact with different kinds of surfaces. All this information is rarely provided all together in a concise and actionable way. The purpose of this document is to fill the gap between a bare list of instructions without explanations and a theoretical textbook.

We are going to use the imagery from the Synthetic Aperture Radar (SAR) on board the Copernicus Sentinel-1A satellite and the SNAP S1 Toolbox software. Our goal is to map a flood event. We will focus on the flood that occurred in Western Germany in the middle of July 2021, and in particular on the Euskirchen district, in North Rhine-Westphalia, an area on the west side of the Rhine river which was hit hard by the flood.



Flood in Swisttal, NRW, July 15th 2021 (Credit: Bundespolizei Spezialkräfte)

We will use two SAR images. One image acquired on July 15th, during the flood, and another image acquired on July 3rd, before the flood. We will refer to the 1st image as the crisis image and to the 2nd image as the archive image. The archive image is used to distinguish permanent water bodies, such as rivers or lakes, from flooded areas. The idea is quite simple: we will overlay the flooded areas of the crisis image over the archive image to distinguish the flooded areas from the permanent water bodies. We will create a binary image of the flooded areas to be exported in the Google KMZ format, and finally we will import it into Google Earth for 2D and 3D visualization. In order to achieve our goal, we will have to apply some transformations to the SAR data. We will sometimes use the term “image” to mean the product, the file that contains the raster images and the metadata. After a short introduction to the Sentinel-1 system, its data products, the SNAP software, and some information about the interaction between the radar pulses and the Earth surface, we will define a workflow to get from the Sentinel-1 data products to a map of the flooding event, explaining the reasons that motivate the transformation that will be applied at each step of the workflow.

The Copernicus Sentinel-1 Constellation



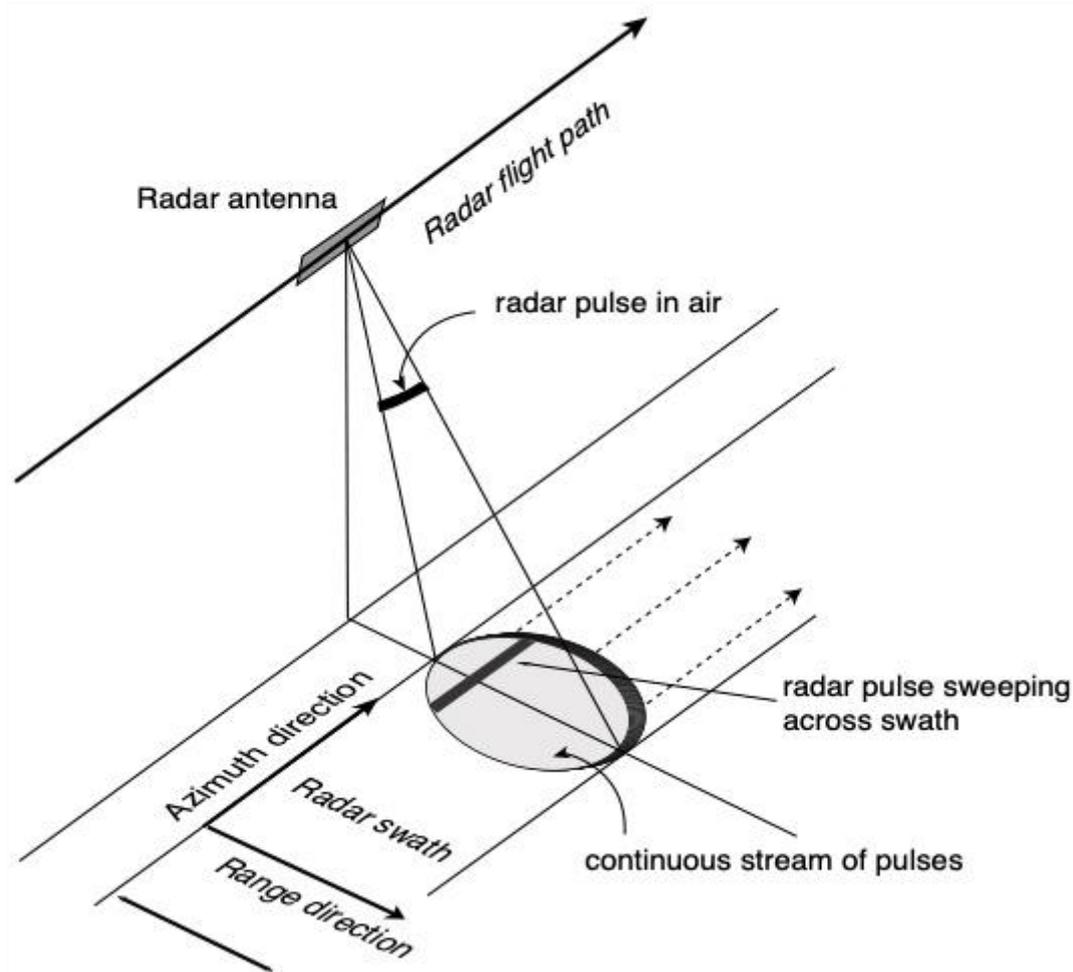
Sentinel-1 (Credit: ESA)

The Copernicus [Sentinel-1](#) constellation is based on two identical satellites for earth observation, launched and operated by the European Space Agency (ESA). Each satellite flies on a Sun synchronous near-polar orbit at 693 km altitude, with a revisit period of 12 days. Since they fly 180° apart from each other, the same area can be revisited every 6 days

at the equator and more frequently at higher latitudes. The primary acquisition mode of the radar is Interferometric Wide (IW) with a 250 km swath and 5 m x 20 m spatial resolution.

The Synthetic Aperture Radar

The Synthetic Aperture Radar on board the Sentinel-1 satellites uses microwaves pulses in the C-band (center frequency 5.404 GHz, or 5.5 cm wavelength) to probe the Earth surface. One advantage of using a radar instead of a passive sensor, is that it does not require the sun to illuminate the target. Another advantage of using microwaves, compared to using the visible or the infrared part of the spectrum, is that the atmosphere is practically transparent in that region of the spectrum, even in the presence of clouds. In other words, contrary to optical passive systems, such as the MRI spectrometer on board Sentinel-2, a radar can work continuously day and night and independently of the weather conditions. The geometry of a radar for earth observation is different from that of a passive system. The antenna looks sideways with respect to the direction of flight and in slant direction with respect to the ground so that two different points on the ground can be distinguished by the difference in time in which the echoes sent back by them are received by the antenna. The SAR antenna on the Sentinel-1 is monostatic, i.e. the transmitter and the receiver are co-located.



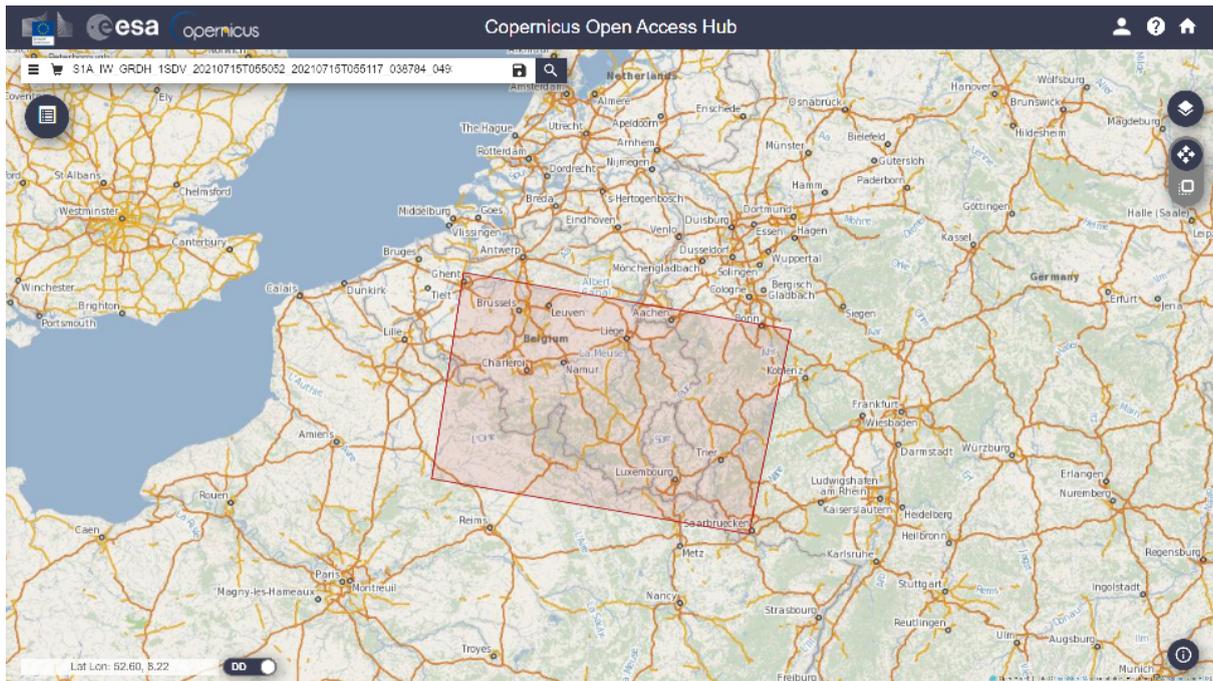
Radar footprint (Credit: NASA - JPL)

The pulses can be transmitted in two polarizations, horizontal (H) and vertical (V). The echoes of the radiation backscattered by the ground can be acquired in co-polarization (HH or VV) and in cross-polarization (HV or VH). Both the amplitude and the phase of the backscattered radiation are recorded by the antenna. The range resolution depends on the duration of the pulse. The shorter the length of the pulse the better the accuracy of the radar to distinguish two scatterers on the ground in the range direction. However, the shorter the pulse length, the lower its energy. Since the power of the backscattered radiation depends on the energy of the transmitted pulse, a tradeoff is established for the pulse length so that the range resolution for the Sentinel-1 can be down to 5 meters in range. The azimuth resolution for a side-looking radar depends on the ratio between the wavelength times the distance from the target and the length of the antenna, therefore to achieve an azimuth resolution of 20 meters in C-band at 700 Km, as the Sentinel-1, the length of the antenna should be around 1750 meters. The Synthetic Aperture Radar achieves the same resolution with an antenna of only 12 meters, using the Doppler shift of the pulse due to the motion of the satellite in the azimuth direction. The imagery collected by the radar is sent to the ground stations for processing and distribution of the products.

The Sentinel-1 Data Products

The ESA delivers different types of [data products](#). Level-0 is the raw SAR product, as it was acquired by the radar. A raw SAR image must be processed in order to provide an interpretable view of the scenery. A [Level-1](#) data product has been pre-processed from the raw data to create images of the amplitude and phase of the backscatter in digital numbers (see [2] and [3] for more information), that is, a Level-1 product is not calibrated and the pixel values must be transformed from digital numbers to backscatter coefficients in order to be compared with other datasets. The pixels in a raster image of a Level-1 product have been projected onto an Earth ellipsoid model, with latitude and longitude, but not projected into a reference system such as WGS84, so a Level-1 image cannot be shown on a map. The Sentinel-1 Level-1 data are available in two separate products: Single Look Complex (SLC), and Ground Range Detected (GRD). The first product contains the information about the amplitude and the phase of the backscatter, while the GRD product contains only information about the amplitude. Level-2 data products have been further processed using inversion algorithms to extract the wind field and the swell over the ocean.

For the task about mapping a flood event over our area of interest, we will download two products from the [Copernicus Open Access Hub](#). After the login one can search for the crisis image. As search criteria, one can select the sensing period (15 July 2021), the Mission (Sentinel-1), and the Product Type.



ESA Copernicus Open Access Hub

We will select the GRD product type because, for flood detection, the information about the phase of the backscattered radiation is not necessary. The size of the GRD product will be $\frac{1}{4}$ the size of the corresponding SLC product. We will also draw a bounding box to select our area of interest, between Euskirchen and the city of Bonn. If any of the two Sentinel-1 satellites was orbiting over the area at the required time, the search will return a list of products that satisfy our requirements. The Sentinel-1A was orbiting over the area of interest, from North to South. The identifier of the product is

S1A_IW_GRDH_1SDV_20210715T055052_20210715T055117_038784_049389_2074

The name of the file follows a [naming convention](#) that provides information about the product. One can see the product details online and also zoom in over the area where the image was acquired before downloading and importing it into SNAP. We can see for example the Relative Orbit Number, or track, that can be used to select the archive product. The same information will be available in SNAP from the product's abstracted metadata. Now we look for the archive image. In order to be compared, the crisis image and the archive image must have been acquired with the radar in the same geometry, that means same track and incidence angle. This requirement can be fulfilled by specifying the same satellite platform S1A, instead of the more generic mission, the same Relative Orbit Number of the crisis product, and leaving the other search parameters but the sensing period, that now is July 3rd 2021. The date of the archive image has been chosen because each satellite, S1A in this case, revisits the same area every 12 days so the first date before July 15th is July 3rd. The identifier of the archive image is

S1A_IW_GRDH_1SDV_20210703T055051_20210703T055116_038609_048E45_3ADC

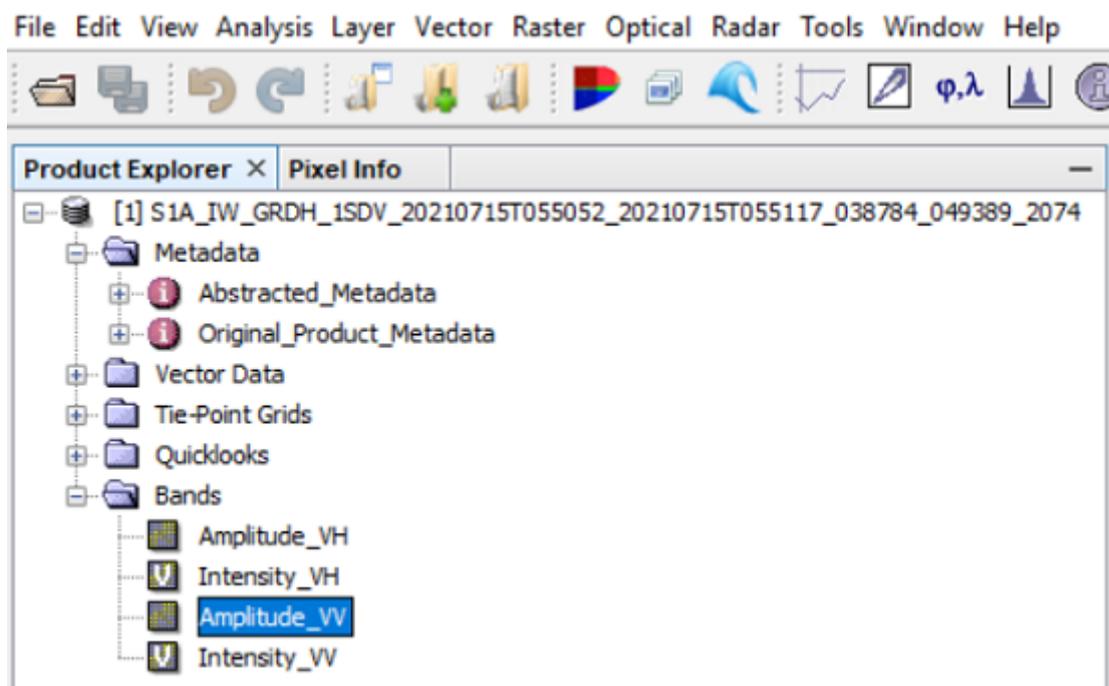
After we have downloaded the two products, we will import them into SNAP. A product is delivered as a zipped folder containing the raster images in TIFF format and other metadata XML files. The raster images are called bands in SNAP.

The SNAP S1Tbx

The SeNtinel Application Platform ([SNAP](#)) is an open source software architecture that is being developed on behalf of the European Space Agency to provide a set of tools for processing and analysing imagery data from the Copernicus Sentinels and other satellites. The S1Tbx is a toolbox dedicated to the Sentinel-1 SAR data. The software is written in Java and can also be used as a [library](#) in external Java applications. A [Python API](#) is also available. In this exercise we use [SNAP version 8.0](#). The software is well documented and supported by a very active [community](#). It can be installed on Linux, Windows and MacOS. A minimum of 8 GB of RAM is required for the proper working of the software.

The Sentinel-1 GRD Product

It is recommended, before starting to apply some transformations to the data, to always have a careful look at the Sentinel-1 product description and its abstracted metadata in order to understand what kind of transformations have been already applied and which one we have to do ourselves. We can open the crisis product by dragging the file into the SNAP's Product Explorer panel.



The crisis products opened in the SNAP Product Explore window

Both the crisis and archive images were acquired in Interferometric Wide ([IW](#)) swath mode and released as a Level 1 Ground Range Detected ([GRD](#)) product. A GRD product has been multilooked and projected to ground range using an Earth ellipsoid model. Multilooking

is a filter that can be used to average the pixel values of an image in range and azimuth. For instance, a multilook of three range looks, computes the mean of three pixels in range and creates a new image with a third of the size in range of the original image. The same can be done in the azimuth dimension. The multilooking operation reduces the size of the image and consequently also its spatial resolution. A multilooked GRD image, acquired in IW mode, has a 20 m x 20 m spatial resolution.

As said, a GRD product is not calibrated. The amplitude of the radar backscatter for each pixel is provided in digital numbers that represent the output of the detector of the radar system after the backscatter has been transformed into an electric signal and amplified. The process to transform the pixel values from digital numbers (DN) to backscatter coefficients is called calibration.

The ground range projection is applied to compute the latitude and longitude coordinates of the pixels. They are computed in the satellite reference system whose origin depends on the orbit of the satellite. If we want to overlay our image on a map we have to project the coordinates onto a standard reference system (datum) such as WGS84, or UTM for the projection on a planar surface. The projection of the data in SNAP is called terrain correction. We will see later that the terrain correction is not just a coordinate transformation.

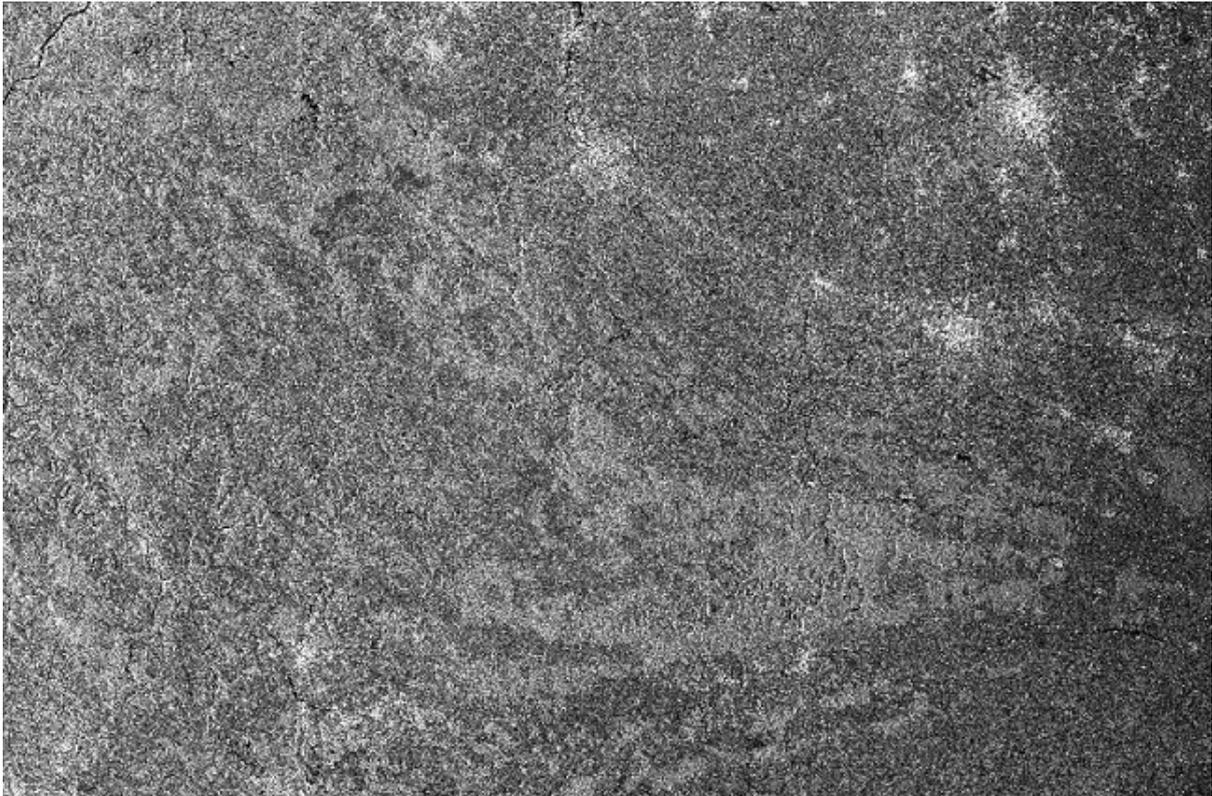
From the description of the Sentinel-1 GRD product we know that in order to use the data in a meaningful way we have to apply a calibration and a coordinate transformation. We can write a draft of the workflow for the processing of our two products as follow

- 1) Calibration
- 2) Terrain Correction

After we have a look at the product description we open one of the bands in the crisis product, to understand what kind of images we have to deal with and what kind of additional operations we have to apply on them to achieve our goal.

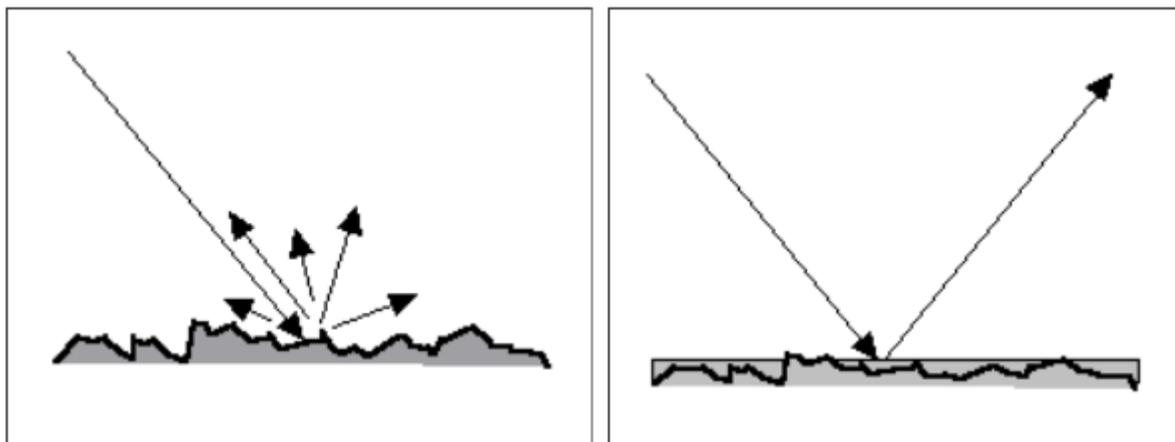
Interpretation of SAR images

If we open the Amplitude_VV band in the crisis product and we zoom in, we may be able to interpret some features like the bright spots that represent urban areas or the black ones that represent water bodies or highways. The origin of the image coordinates system is on the top-left corner, the azimuth direction is from the top to the bottom of the image, the range direction is from the left to the right. The Sentinel-1A satellite was descending to the south, looking at its right side. To put it on a map, it should be rotated about 180 degrees, with the satellite flying southward on the right.



The Amplitude_VV band of the crisis product

On the upper left side we can see the Rhine river and part of the city of Bonn, on the upper right side we can see Brussels and its airport. The basic rule to interpret radar images is that the amount of backscatter, the radiation sent back towards the radar by the illuminated surface, depends on the roughness of the surface. The roughness of a surface is related to the wavelength of the incident radiation. As a rule of thumb, if the wavelength is more than ten times the size of the deviations from a flat surface we can consider the surface to be smooth and its backscatter will be low since most of the incoming radiation will be reflected in a specular direction away from the radar. Other parameters that influence the backscatter are the dielectric constant of the surface, the polarization of the transmitted radiation, and the amount of moisture. The functional relation between all these parameters and the backscatter is quite complex, but the task at hand does not require that level of detail. The area of interest is composed of agricultural areas, forested areas, urban areas and water surfaces. Agricultural and forested areas have a quite strong backscatter and they change partially the polarization of the radiation. Urban areas have a stronger backscatter because of the double reflection (double bounce) on the horizontal and vertical surfaces of the buildings. Flooded areas, like other water bodies such as rivers and lakes, reflect the incident radiation in a specular direction off the radar, so that these areas have a low backscatter and appear black on the image. One other reason for a low backscatter can be a strong absorption of the incident radiation.



Diffuse scattering

Specular reflection

Scattering (Credit: CRISP)

So we know that we have to look at the black spots in the images to detect the flooded areas. Now we need to know what can affect the detection of these areas.

Noise and distortions

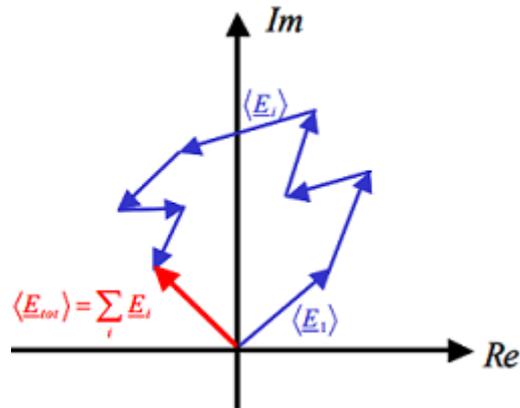
There are three sources of errors that affect all SAR images. One is the thermal noise from the instrument that affects any type of measurement, and the other two are speckle and geometric distortions that are a characteristic of radar imagery.

Thermal Noise

The radar produces noise, like any electronic instrument, that adds a radiometric bias to the backscatter that can affect particularly pixels with low backscatter values [5], such as over water bodies or flooded areas. The Noise-Equivalent Sigma Zero (NESZ) is a parameter that defines the level at which a signal cannot be distinguished from noise. The NESZ for the Sentinel-1 imagery is -22 dB.

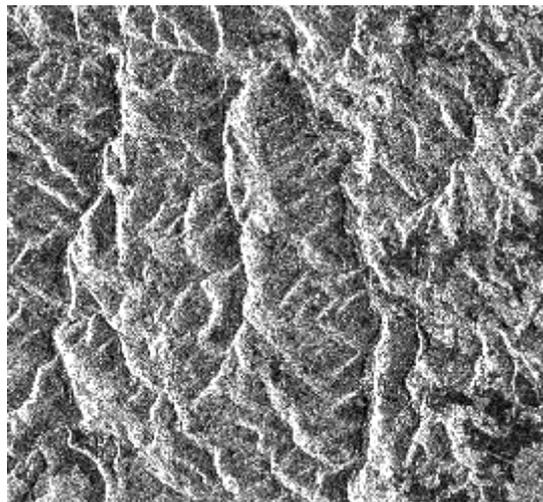
Speckle

The speckle can be seen as a granular salt-and-pepper pattern in the image that does not represent a real difference in the scattering properties of the objects on the surface. The different brightness of the pixels arises from the different distribution of the scatterers within each pixel that result in a different coherent summation of the backscatter. The problem is particularly relevant when the pixel values are used for classification as it may lead to a misclassification of the pixels. The speckle can be reduced by multilooking or by applying some other filter that does not reduce the size of the image.



Geometric and Radiometric Distortions

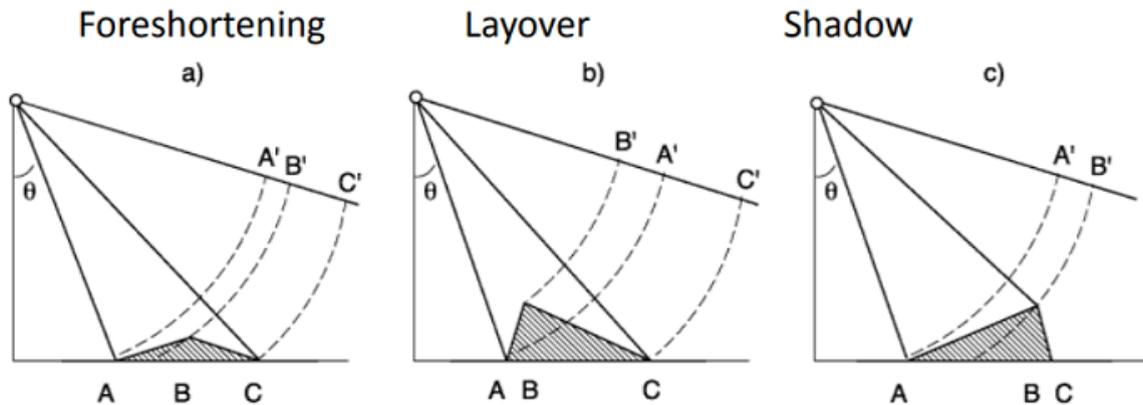
Radar and SAR imagery have distortions over mountainous regions due to the side-looking geometry of the antenna and to the slope of the mountains. Since the backscatter increases for small incidence angles, a mountain with a greater slope, on the side illuminated by the radar, produces a stronger backscatter than one with a smaller slope. The radiometric distortions can be seen in imagery over mountainous regions as narrow bright lines parallel to the flight direction of the satellite.



Mountain area seen by the side-looking radar

Geometric distortions are clearly noticeable in radar imagery over hilly or mountainous areas. The mountains appear to be leaning towards the radar on the left side. There are three types of geometric distortions: foreshortening, layover and radar shadow.

Foreshortening happens when one side of a mountain, illuminated by the radar, appears shorter than the other side. The extreme of foreshortening, called layover, is when the top of a mountain, with a steep slope face on the side of the radar, looks closer to the radar than the base of the mountain. Finally, the side of a mountain opposite to the radar may not be illuminated at all if it is too steep, resulting as black spots in the imagery.



Geometric distortions

The Workflow

We now have enough information about the imagery and the issues that may affect it to decide which operations we should apply in order to create a map of the flood that occurred in our area of interest. We will start by cropping the images to reduce their size. The next step would be to remove the thermal noise. The calibration, which was already included in our workflow, is the mandatory operation that will transform the data from digital numbers to physically meaningful backscatter coefficients. Once we have the backscatter values we may consider to further reduce the speckle. It is up to us to decide whether it might be useful to apply a filter to further reduce the speckle, but not the size of the image, or to reduce both by multilooking again. One other source of errors we did not mention so far is the lack of precise information about the orbit of the satellite available in the product that might affect the precise computation of the coordinates transformation. The next necessary step is the terrain correction, the second operation that we have already included in our workflow. As we said, the first reason for the terrain correction is to project the image onto a geographic reference system otherwise we will not be able to overlay the flooded area on a map. The second reason is to reduce the geometric distortions due to the geometry of the antenna and to the topography. Once we have the image with the proper coordinates and without, or at least reduced, distortions we can create a binary image of the flooded areas to be exported on Google Earth, or any other GIS application. The list of operations that we might want to perform on the products that we have downloaded is as follows

1. Image Subset
2. Thermal Noise Removal
3. Calibration
4. Speckle Filter
5. Multilooking
6. Apply Orbit File
7. Terrain Correction
8. Linear to Logarithmic
9. Thresholding

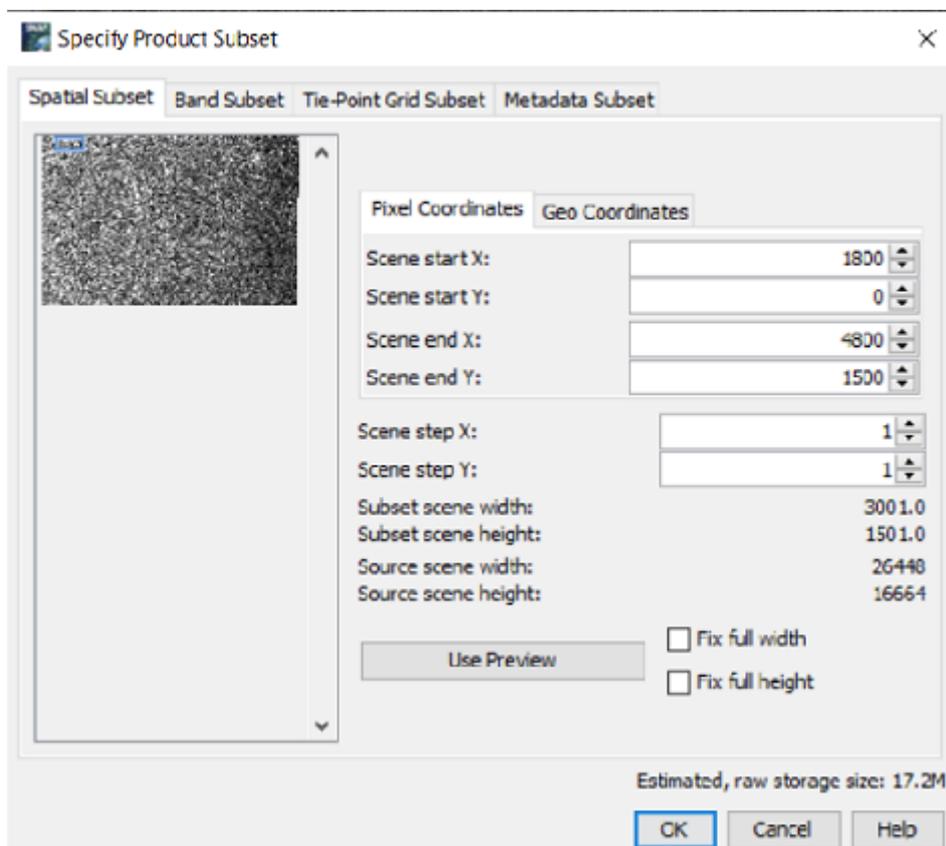
SNAP provides all the algorithms needed to perform the operations in the workflow. We will execute all the steps in the workflow using only the crisis image. Each step will use the output of the previous one as input and will output a new product. After we have completed the workflow for the crisis image, we will define a processing chain to automatically apply the same steps to the archive image.

Image Subset

We want to extract a subset of the images to cover our area of interest. There are three advantages in cropping the images to cover only the area of interest. The first is that in each step of the workflow we will create new products that will be used as input to the next step. Since the size of the GRD products is on average 1GB we may end up using 5GB to process each product. Secondly, the reduced size will speed up all the computations. The third advantage is that by focusing on the area of interest we will have more pixels of the flooded areas and it will be easier for us to find a threshold to create a binary image of those areas. We open the crisis Amplitude_VV image from the SNAP's Product Explorer window. We will use only the vertical co-polarized image since for the task at hand the cross-polarized band does not provide additional information. Once the image is shown we click on

Raster -> Subset

In the popup window we set the 'Scene start X' to 1800, 'Scene start Y' to 0, 'Scene end X' to 4800, and 'Scene end Y' to 1500.



The selected area is visible close to the upper-left corner

The origin of the image is at the top left pixel. The x axis, from left to right, is the range direction and the y axis, from top to bottom, is the azimuth direction. We select the 'Band Subset' tab to deselect all the bands but Amplitude_VV, and then we click on "Ok" to create a new product with a subset of the crisis image. In order to save the new product on the file system, we select it in the 'Product Explorer' then we click on

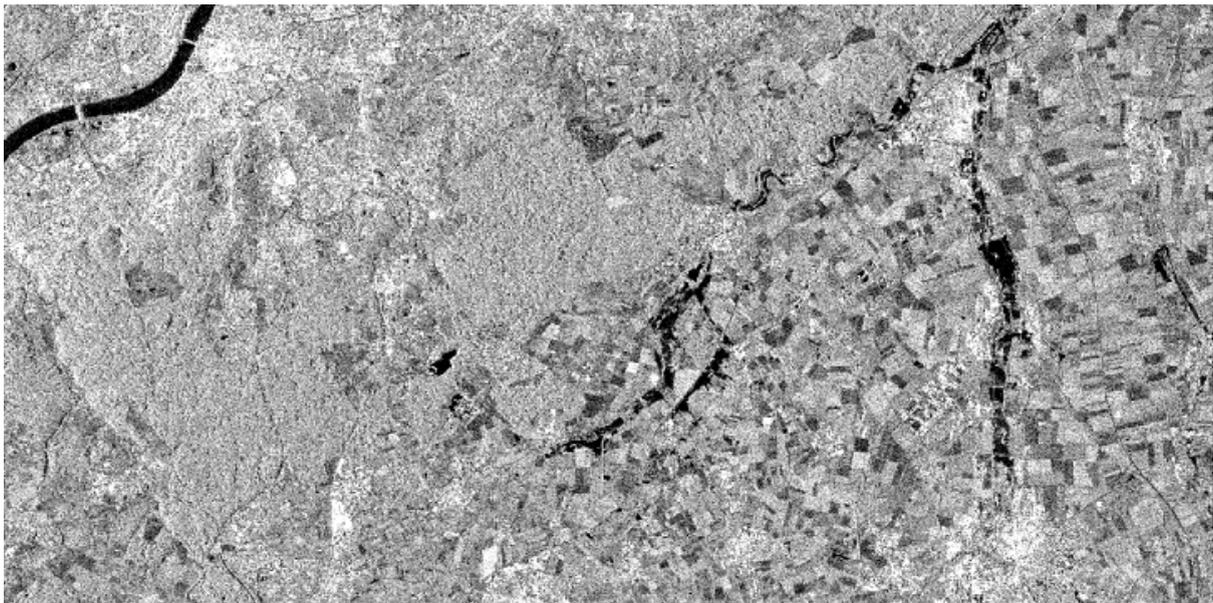
File -> Save Product

to save it in a folder of our choice. It is suggested to create a new folder that will contain the outputs of all the operations that will be performed, e.g. (for the Windows OS)

C:\Users\\flood_mapping

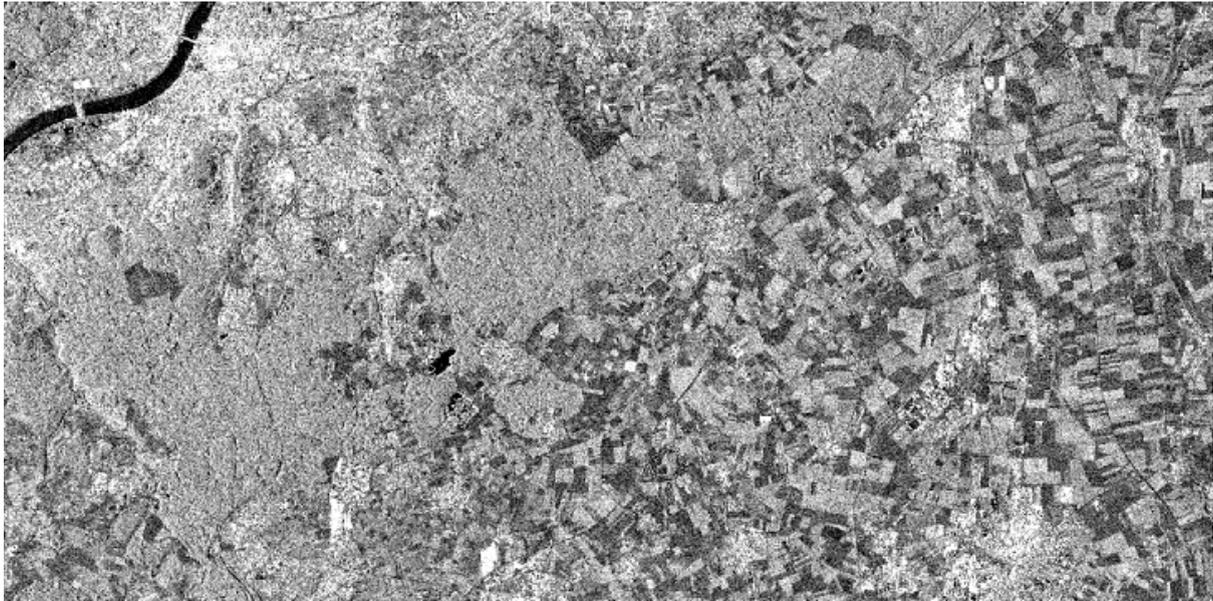
We can shorten a little bit the file name, e.g.

S1A_IW_GRDH_1SDV_20210715_Euskirchen



Amplitude VV image of the flood over the area of interest, July 15th 2021

It is not necessary to subset the archive image now. It will be done automatically when we define the processing chain later. However, the archive image is shown below to get an idea of the difference between the two images, before and during the flood



Archive amplitude VV image of the area of interest, July 3rd 2021

Now we can close the original product since from now on we will use only its subset. The new product contains only the Amplitude_VV band over our area of interest. We select the crisis product, then right-click and 'Close Product'.

Thermal Noise Removal

The product metadata contains a look-up table, in the Original_Product_Metadata section, with the noise measurements that can be subtracted from the intensity values of each pixel using the operator available in SNAP. We click on

Radar -> Radiometric -> S-1 Thermal Noise Removal

In the popup window we add a suffix, e.g. _NR, to the output file name, we select the output folder, and then we click on 'Run'. Now we have a new product

S1A_IW_GRDH_1SDV_20210715_Euskirchen_NR

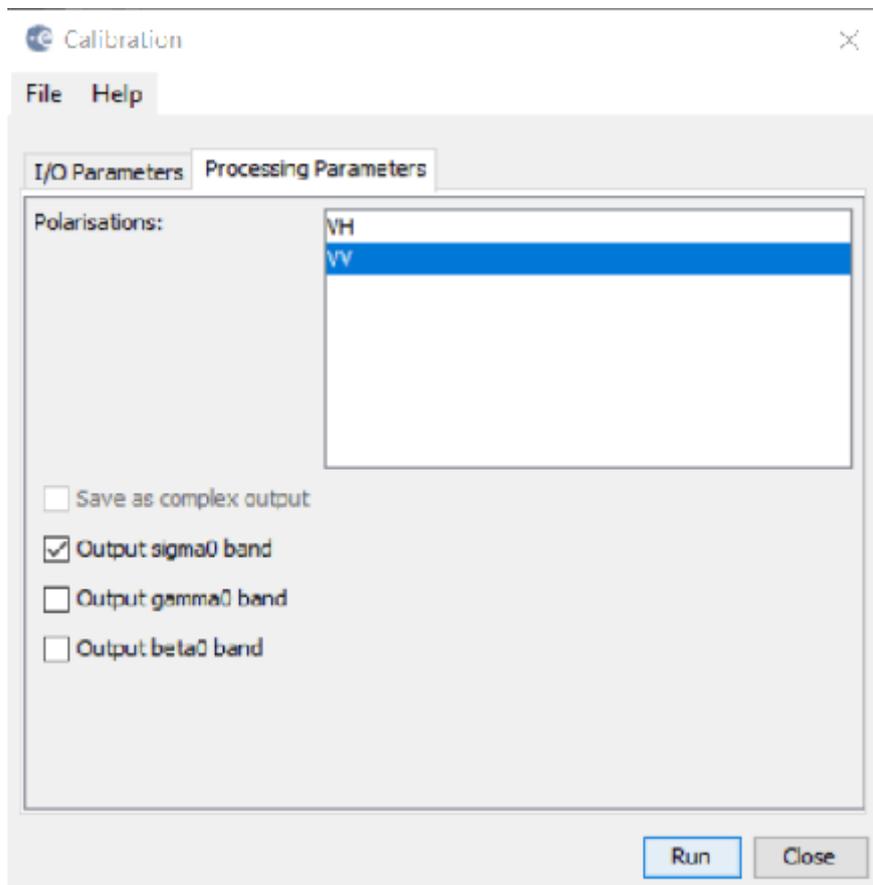
If we open the Bands folder of the output product, we can see that now we have a band called Intensity_VV. This is because the noise is removed from the intensity, that is the squared value of the amplitude.

Calibration

Now we will apply the calibration to transform the pixel values from integer digital numbers (DN) to backscatter coefficients (sigma nought). We click on

Radar - Radiometric -> Calibrate

In the popup window we select the output product of the thermal noise operator. In the 'Processing Parameters' tab we leave the polarization VV, then we select the checkbox 'Output sigma0 band', if it is not already checked, and we click on 'Run'.



Sigma nought set as the output of the calibration

If we open the bands of the calibration output product, we can see that now it is named Sigma0_VV, which stands for sigma nought, or backscatter coefficient. We can read the pixel values from the 'Pixel Info' window that can be opened from the menu bar

View -> Tool Windows -> Pixel Info

| Product Explorer | | Pixel Info X | |
|---------------------|-------------|--------------|--|
| [-] Position | | | |
| Image-X | 787 | pixel | |
| Image-Y | 267 | pixel | |
| Longitude | 7°03'37" E | degree | |
| Latitude | 50°43'44" N | degree | |
| [+] Time | | | |
| [-] Bands | | | |
| Sigma0_VV | 0.12019 | intensity | |
| [+] Tie-Point Grids | | | |
| [+] Flags | | | |

The Pixel Info window

The values that are computed in this step are the backscatter coefficients, also called sigma nought, and represent the ratio between the incident radiation and the backscattered radiation per unit area on a horizontal plane. In other words, the sigma nought does not take into account the topography of the place. The lack of information about the topography may have two consequences. One is that the backscatter is different from the real one if the incidence angle is smaller than it would be on a flat surface. Secondly, also the exact position of the scatterer is different from the real one if the target is in a much higher place because it looks closer to the radar. Since our area of interest does not include high mountains we ignore the radiometric distortions due to the topography and we select the output of the calibration to be the radar cross section (sigma nought). In case there were higher mountains in the area we should have selected the beta nought as output and then applied a terrain flattening operator, to correct the radiometric distortions due to the topography. The name of the new product should be

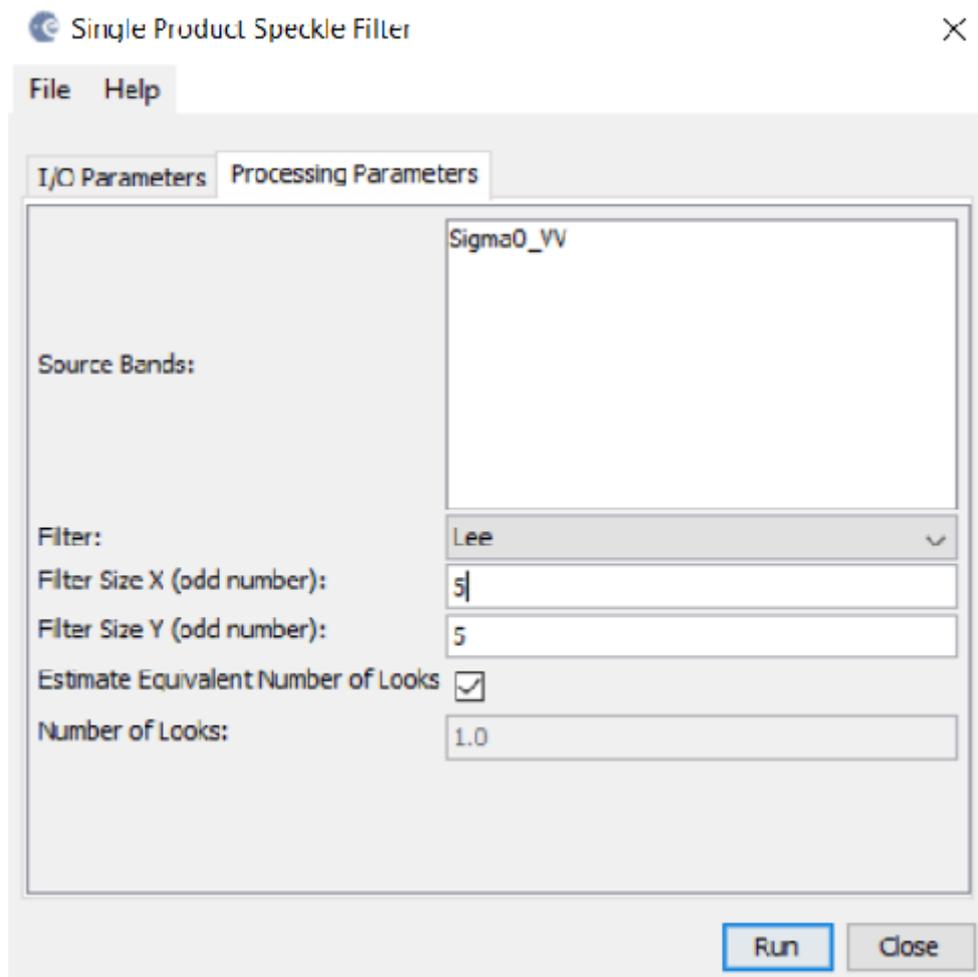
S1A_IW_GRDH_1SDV_20210715_Euskirchen_NR_Cal

Speckle Reduction

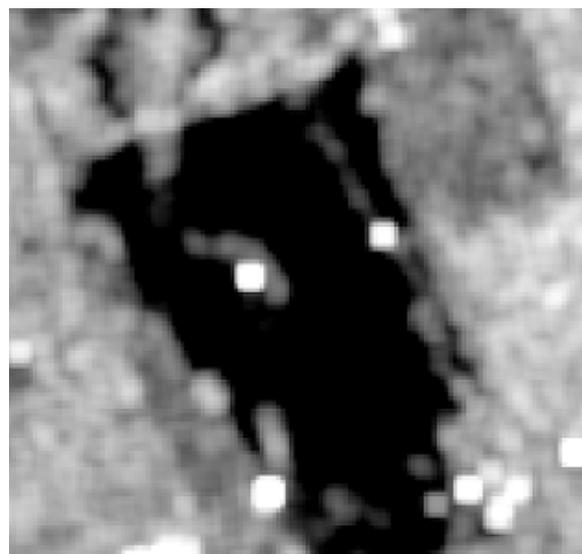
As said, the speckle has been reduced by multilooking in the GRD products, but we might want to further reduce it if we think that it still affects the images we are going to use. It is certainly useful to do an experiment. One reason to do so is that in our case we will need to apply a threshold to separate the flooded areas, with low pixel backscatter values, from the rest and the speckle may lead to a misclassification of some of the pixels. We click on

Radar -> Speckle Filtering -> Single Product Speckle Filter

In the popup window we select the 'Processing Parameter' tab and we set the 'Filter' parameter to 'Lee', leaving both 'Filter Size X' and 'Filter Size Y' to 5.



The filter will compute the average of each pixel using a window of 5x5 neighbor pixels. We can compare the images (better in logarithmic scale) before the application of the speckle filter and after.



If we zoom in over the image on a flooded area, shown on the left, we can see that there is still some speckle, not only over the land but also over the water surface, shown as gray pixels surrounded by dark ones. The image on the right side shows the same area after the application of the speckle filter. The image looks blurred but what is relevant is that many pixels that before could have been misclassified as not flooded areas have been changed and look dark as the surrounding pixels.

After we are done with the speckle filter we should have a new product in the SNAP 'Product Explorer'

S1A_IW_GRDH_1SDV_20210715_Euskirchen_NR_Cal_Spk

Multilooking

This operation can be used to reduce the size of the images by averaging the pixel in azimuth and range direction. It also reduces the speckle, as a byproduct. In our case we will not use it because we have already reduced the speckle in the previous step and applying the multilooking again will further reduce the spatial resolution. We prefer to keep the original spatial resolution of 20x20 meters so we skip this step. In case we are working with a big image and we want to reduce its size, we could proceed as follows: we click on

Radar -> SAR Utilities -> Multilooking

In the popup window we select the source product, e.g. the crisis image, then the 'Processing Parameters' tab to set the 'Number of Range Looks' and the 'Number of Azimuth Looks'. These settings define the number of pixels that will be averaged and mapped to one single pixel in the output multilooked product. If we set those parameters to, e.g. 3, the spatial resolution will become 60x60 m. and the size $\frac{1}{3}$ of the input product in range and azimuth.

Apply Orbit File

In the product's 'Abstracted_ Metadata' we can see a list of 16 'Orbit State Vectors'. Each orbit vector contains two points (x, y, z) used to draw a line that will be used to compute the orbit. The orbit is used for the terrain correction of the image, which is described in the next step, to transform the coordinates from the satellite reference system to a geographic reference system. The orbit available in the product metadata could be enough for the task at hand since we do not make use of the phase of the backscatter like in other SAR applications that requires precise orbit information. Nonetheless, since we will have to reduce the geometric distortions and compare two images acquired on different dates, more precise information about the orbits can improve the results. The precise orbit for each product will be automatically downloaded from the Internet, if available. We decide to use the precise orbit, we click on

Radar -> Apply Orbit File

In the popup window, as a source product we select the output of the speckle filter, we leave all the default parameters and click on 'Run'. The output product should be

S1A_IW_GRDH_1SDV_20210715_Euskirchen_NR_Cal_Spk_Orb

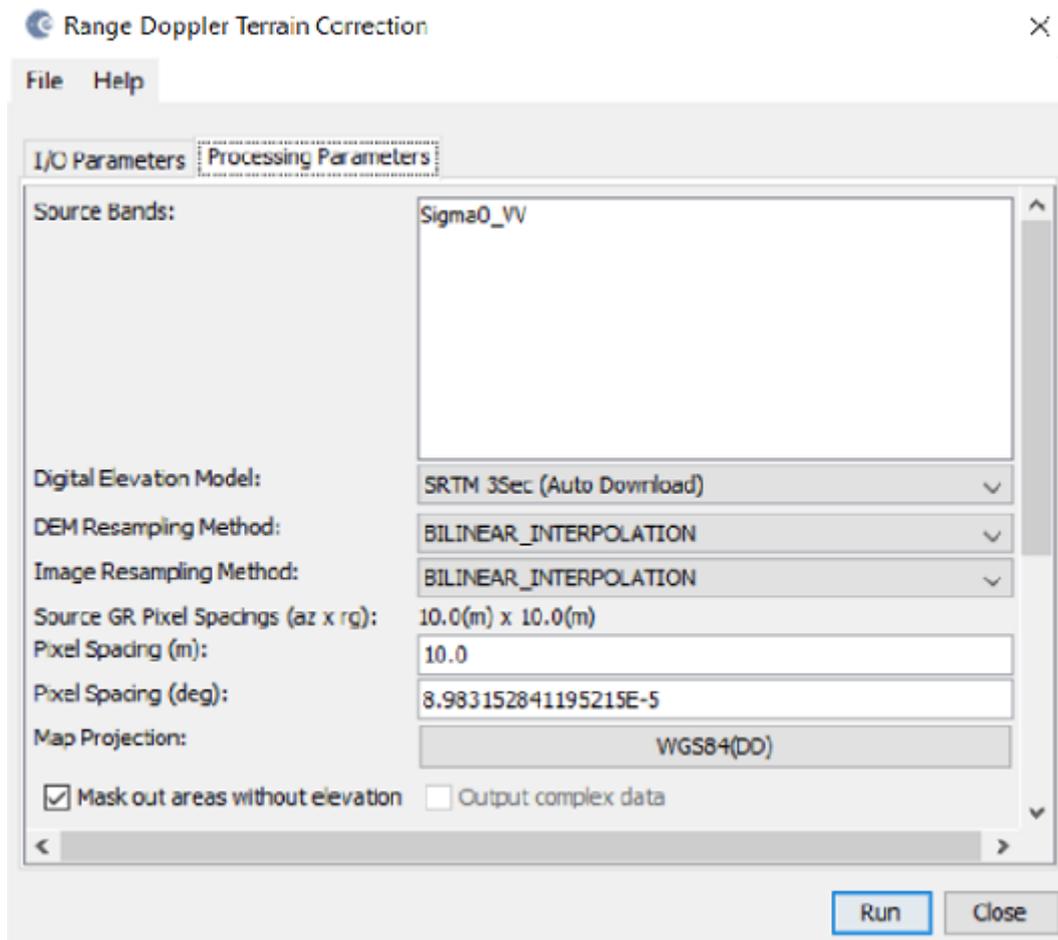
If we look at the abstract metadata of the output product we can see that there are now 22 orbit state vectors. The number of state vectors depends on the size of the subset area. The reason why the precise orbit is not available in the original product is that the precise orbit is made available days-to-weeks after the image acquisition, when more precise orbits are calculated. The precise orbit file may not be available. In this case the operator will fail. Another option is to use the restituted orbit by selecting 'Sentinel Restituted (Auto Download)' for 'Orbit State Vector' in the 'Processing Parameters' tab of the 'Apply Orbit File' popup window. The difference between precise and restituted orbit files is the accuracy, +/- 5 cm for precise orbits, and +/- 10 cm for restituted orbits.

Terrain correction

We apply the terrain correction to project the pixels from the satellite reference system onto a geographic reference system. The origin of the satellite reference system at any given point in time is referred to the orbit of the satellite. Another reason to apply the terrain correction (also referred to as geocoding or orthorectification) is to correct the geometric distortions in the range direction, such as the foreshortening, due to the side-looking geometry of the SAR and to the topography of the terrain. The geocoding depends on the availability of a Digital Elevation Model (DEM). We click on

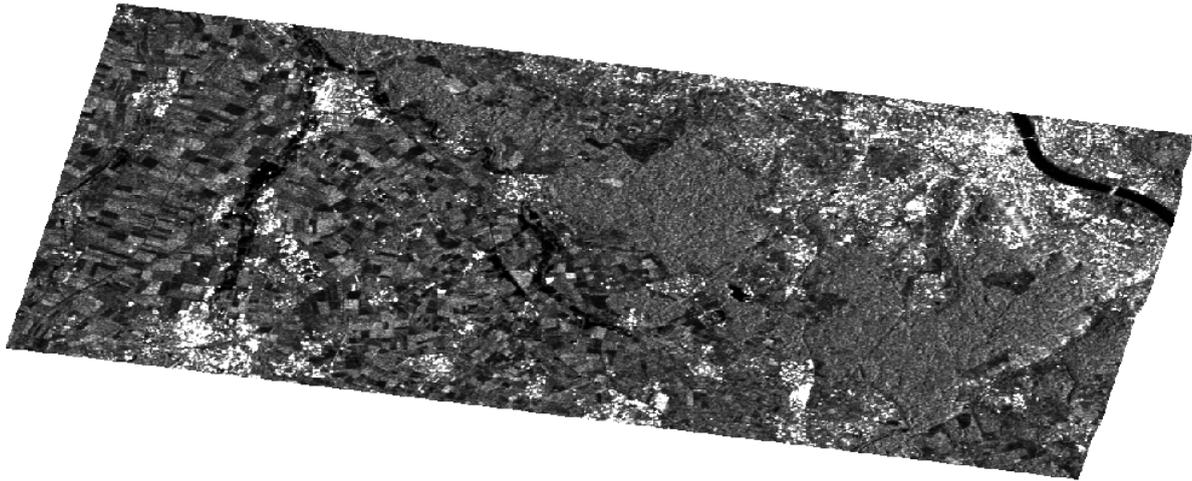
Radar -> Geometric -> Terrain Correction -> Range-Doppler Terrain Correction

In the popup window, we select the precise orbit product as a source product, and we leave all the default parameters.



Range Doppler Terrain Correction window

The pixels will be projected in the WGS84 geographic reference system and the correct distances will be computed using the Shuttle Radar Topography Mission (SRTM) digital elevation model that will be automatically downloaded from the Internet. We could otherwise change the map projection to a 2D coordinate system, e.g. UTM / WGS84 to have x and y coordinates. We click on the 'Run' button, a new product is created with a band Sigma0_VV that shows our area of interest projected onto a geographic reference system.

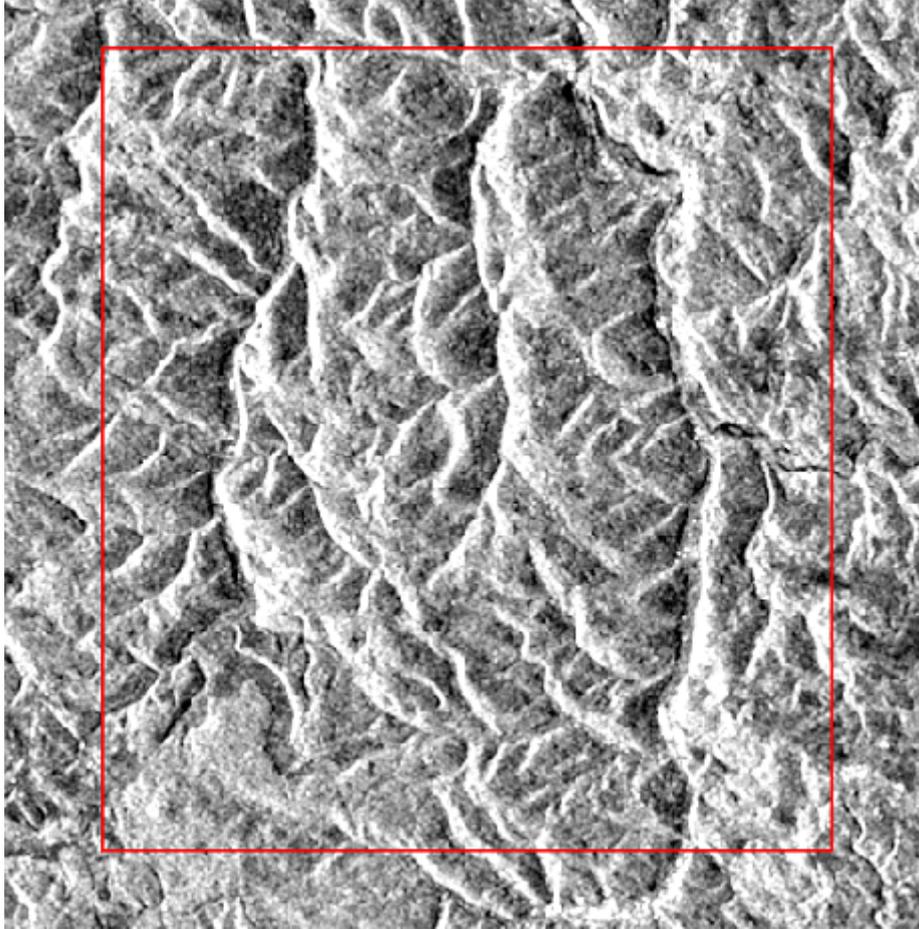


Crisis image Sigma0_VV of the flooded areas after the terrain correction

The terrain corrected product file name should be

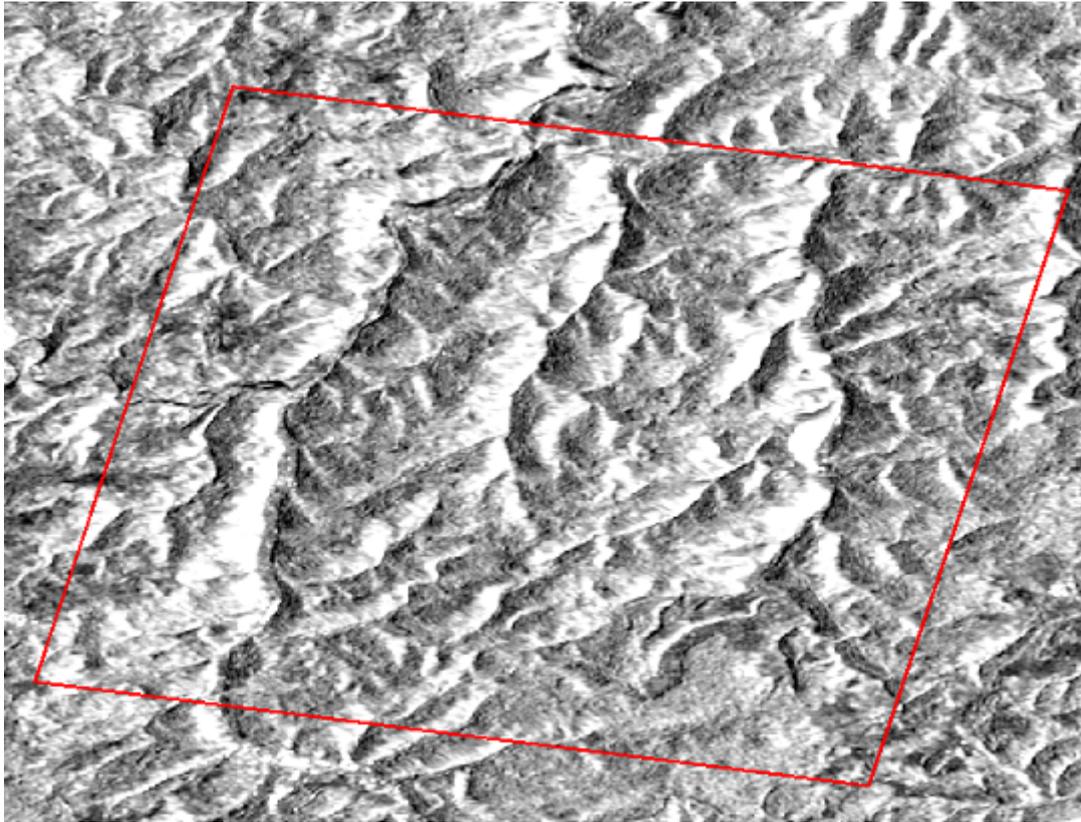
S1A_IW_GRDH_1SDV_20210715_Euskirchen_NR_Cal_Spk_Orb_TC

The subset region that we have selected to be processed is a flat area and the geometric distortions, due to the side looking geometry of the radar and to the topography of the terrain, are not clearly visible. We can see the distortions in another subset area that was not considered for processing. We have already seen this area when we introduced the topic about the geometric distortions, now the same area is marked with a red bounding box. The image before the terrain correction is shown using the satellite reference system whose origin is in the upper-left corner. The mountains look like leaning to the left, towards the radar.



A mountain area before the terrain correction

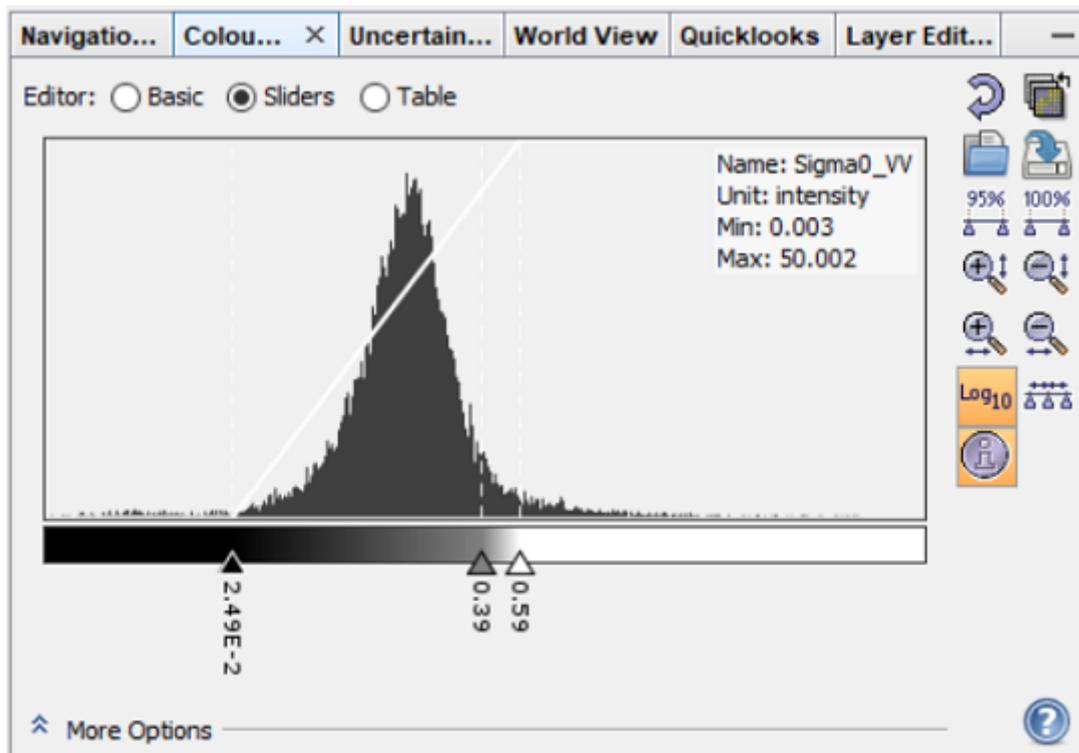
Now we look at the same area after it has been terrain corrected. The area is projected onto a geographic reference system, rotated about 180 degree. The radar was on the right flying to the South and looking at its right side. The leaning effect has been removed.



The mountain area after the terrain correction

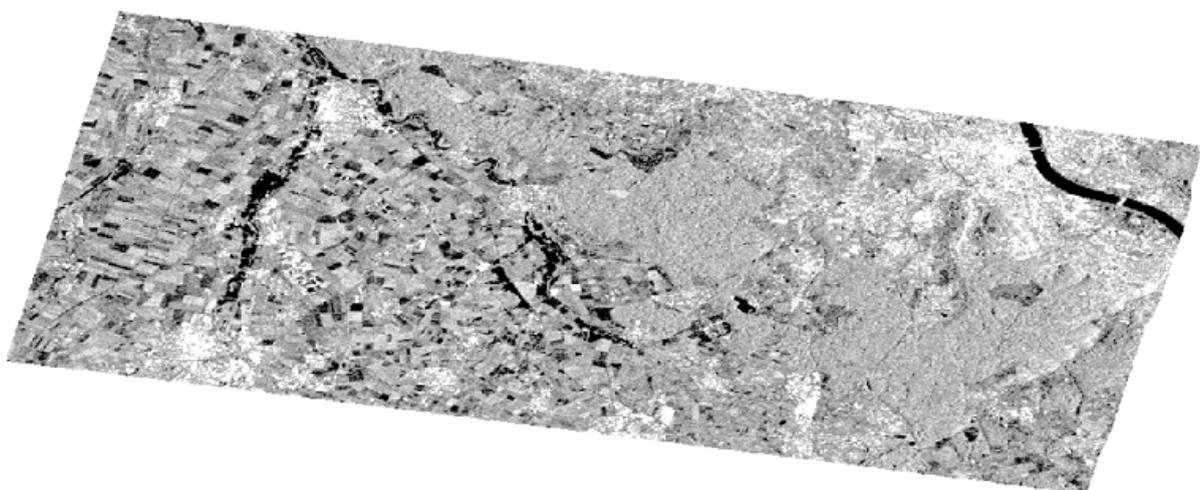
Linear to Logarithmic

We can see that the images are quite dark and that from the histograms, in the 'Colour Manipulation' window, most of the pixels have a small backscatter value and few of them have a high backscatter value. In this case it is better to transform the data from linear to logarithmic values or decibel (dB). We can see that the distribution of the pixel values in decibel is easier to understand and manipulate.



Histogram in log₁₀ of the pixel values of the Sigma0_VV band

We right-click on the band and we select Linear to/From dB. A new virtual band, Sigma0_VV_dB, is created. In the next step we will create a mask, a binary image of the flooded areas, by applying simple arithmetic operations on the pixel values that will be easier to do using the backscatter values in decibel (dB). We can save the new band by right-clicking and then selecting 'Convert Band'.



Terrain corrected Sigma0_VV_db band of the crisis image

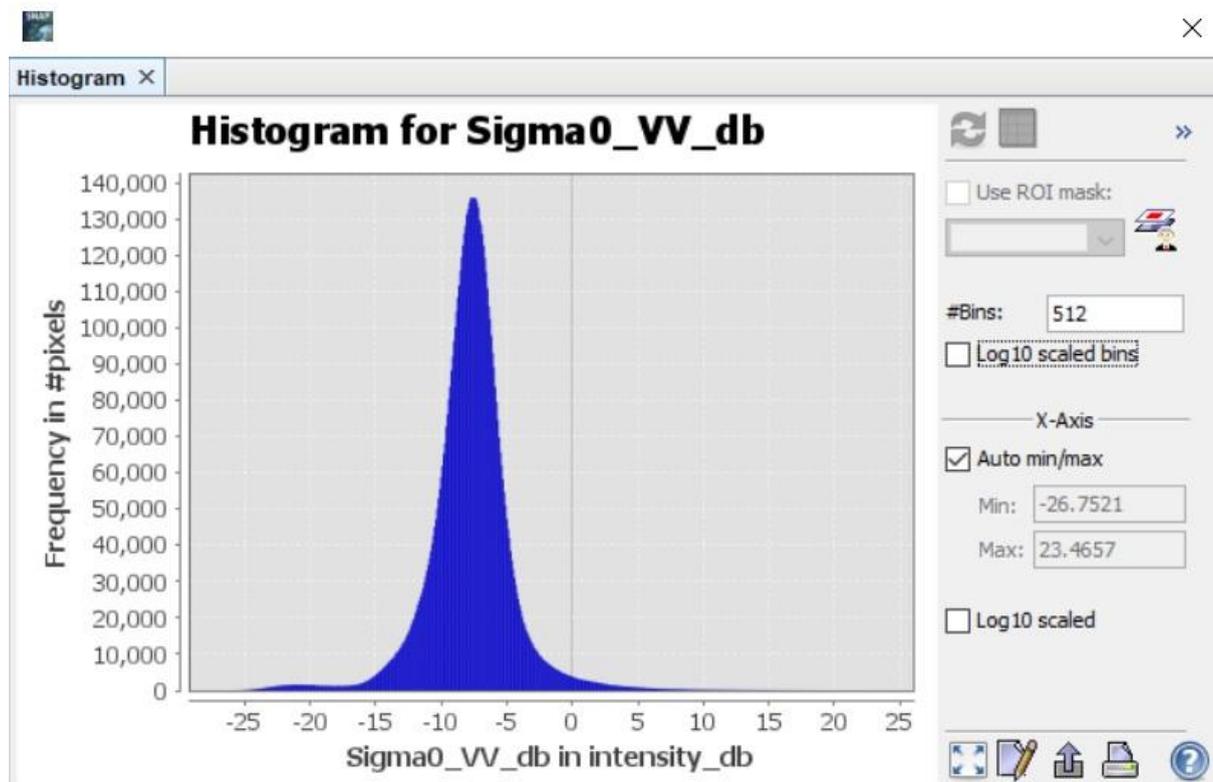
The flooded areas are clearly visible as black spots on the left side of the image and the Rhine river on the right side. This is the last step of our workflow. We will apply the same workflow to the archive image but automatically by defining a processing chain that will contain the same operators.

Thresholding

Now we have the image calibrated, despeckled and terrain corrected. The next step is to separate the flooded areas in the crisis image from the rest. This task is called thresholding in digital image processing. A good reference about the technique is [6]. The objective is to determine a [threshold](#) and assume that the pixels whose values are below the threshold belong to the flooded area, and those that are above that threshold belong to the not flooded area. We will also have to distinguish the flooded areas from the permanent water bodies, i.e. rivers and lakes, using the archive image. This step can be divided into two. First we find a good threshold to separate the flooded areas from the rest, and secondly we create the binary image. The ideal situation to determine a threshold is when the histogram of an image is a bimodal distribution containing two peaks separated by a valley. One mode for the pixels in the flooded areas, with low backscatter values due to the specular reflection on the water surface, and another mode for the pixels of the not flooded areas. To create the histogram of the crisis image, we select the band Sigma0_VV_dB and we click on

Analysis -> Histogram

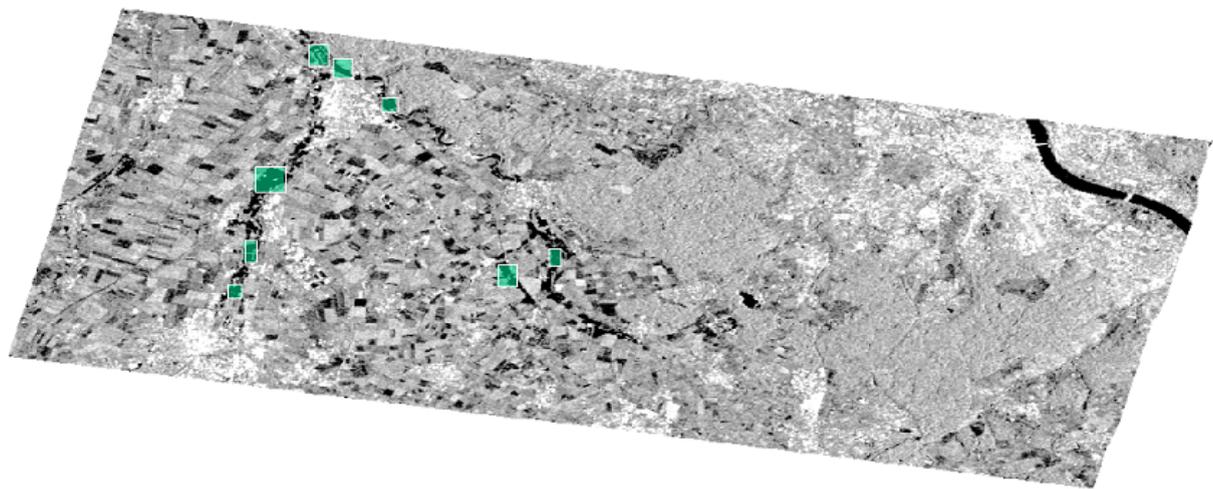
In the popup window we click on 'Refresh View' to load the band and we deselect the 'Log10 scaled bins' checkbox as the data is already in decibel.



In our case, the distribution does not really look like a bimodal one. This happens because the full image still contains large areas that were not flooded so they do not help us to find a threshold to distinguish flooded areas from not flooded ones. We can learn more about the distribution of the pixel values over the flooded areas by focusing our attention to the area around the black spots. SNAP allows us to create geometries, such as rectangles or polygons to select regions of interest in the images. We first create a Vector Data Container that will contain the polygons we will create on the image. We select the terrain corrected product and then

Vector -> New Vector Data Container

In the popup window we write the name of the vector data, e.g. region_of_interest, then we click on 'OK', the polygons will be added to the product's 'Vector Data'. We select the 'Rectangle drawing tool' on the menu bar, we zoom in the Sigma0_VV_dB band and we draw some rectangles over some of the black spots that represents a flooded area

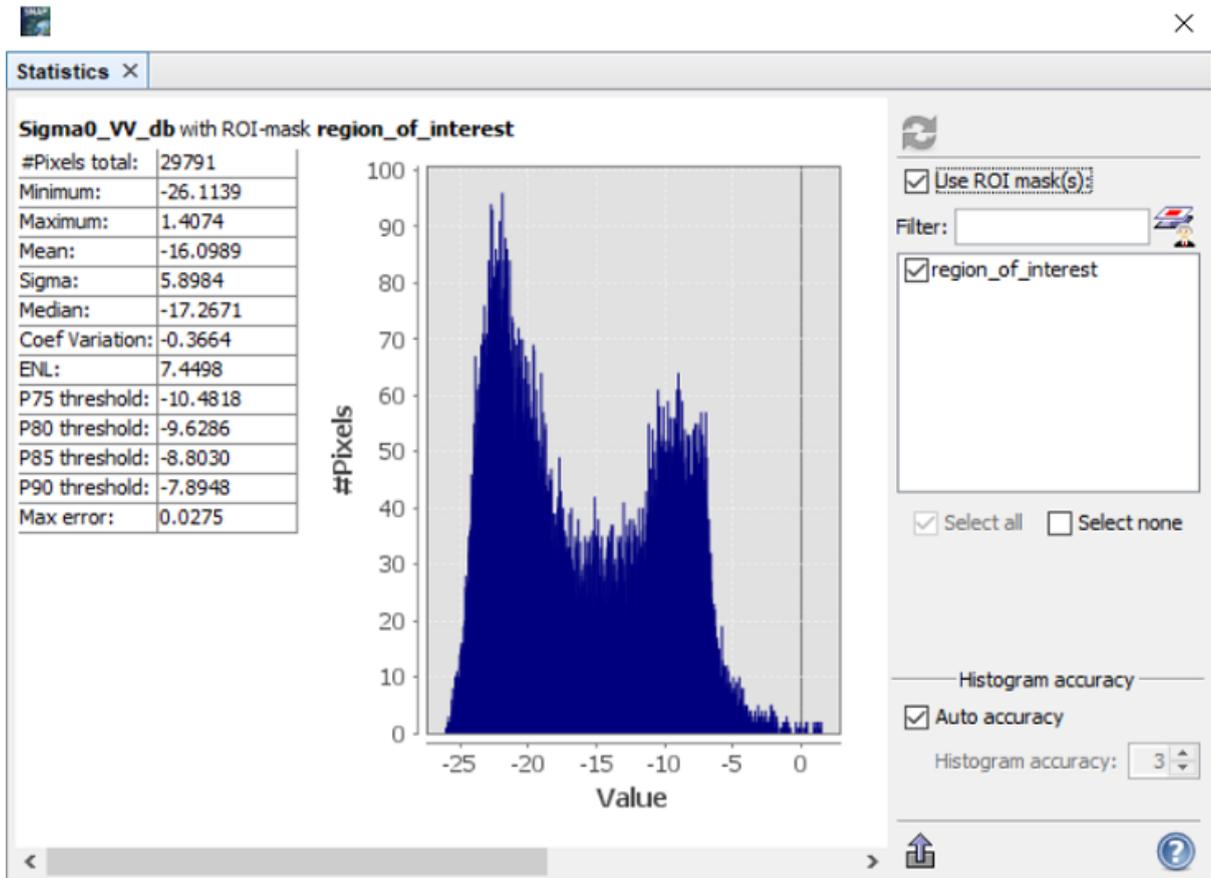


Sample polygons to compute the flooded areas pixels statistics

Our vector data now contains our rectangles and we can look at the statistics and distribution of the pixel values. SNAP provides an additional tool to show the histogram and the statistic of the distribution

Analysis -> Statistics

In the popup window we select the 'Use ROI mask' and the 'region_of_interest' checkboxes and we click on the 'Refresh View' button to load the statistics and the pixel values distribution in our region of interest.



Now we have a bimodal distribution with the left peak at -23 dB and a valley at -15 dB that we can choose as our threshold to create the binary images for the flooded areas and also later for the water bodies. We click on

Raster -> Band Maths

In the popup window we select the terrain corrected product, we write the name of the new band, e.g. surface_water, we deselect the checkbox 'Virtual' so that the binary image will be saved on the file system, and finally we create the expression to split the pixels using the threshold $\text{Sigma0_VV_db} < -15$

Band Maths ×

Target product:
 [2] S1A_IW_GRDH_1SDV_20210715_NR_Cal_Spk_Orb_TC ▼

Name:

Description:

Unit:

Spectral wavelength:

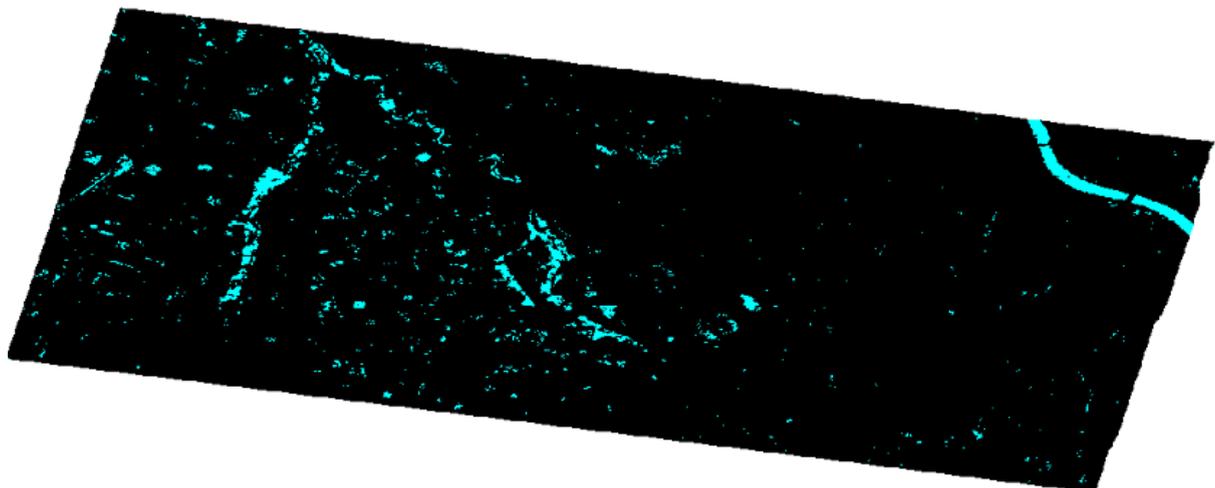
Virtual (save expression only, don't store data)

Replace NaN and infinity results by

Generate associated uncertainty band

Band maths expression:

We click on 'Ok' to close the 'Expression Editor' and the Band Maths window. A new band 'surface_water' with the binary image is created



Flooded areas and water bodies

The pixels colored in cyan, whose value is below the threshold, represent the flooded areas and the water bodies, such as the Rhine river on the right side. The color of the pixels can be set from the 'Table Editor' of the 'Color Manipulation' window

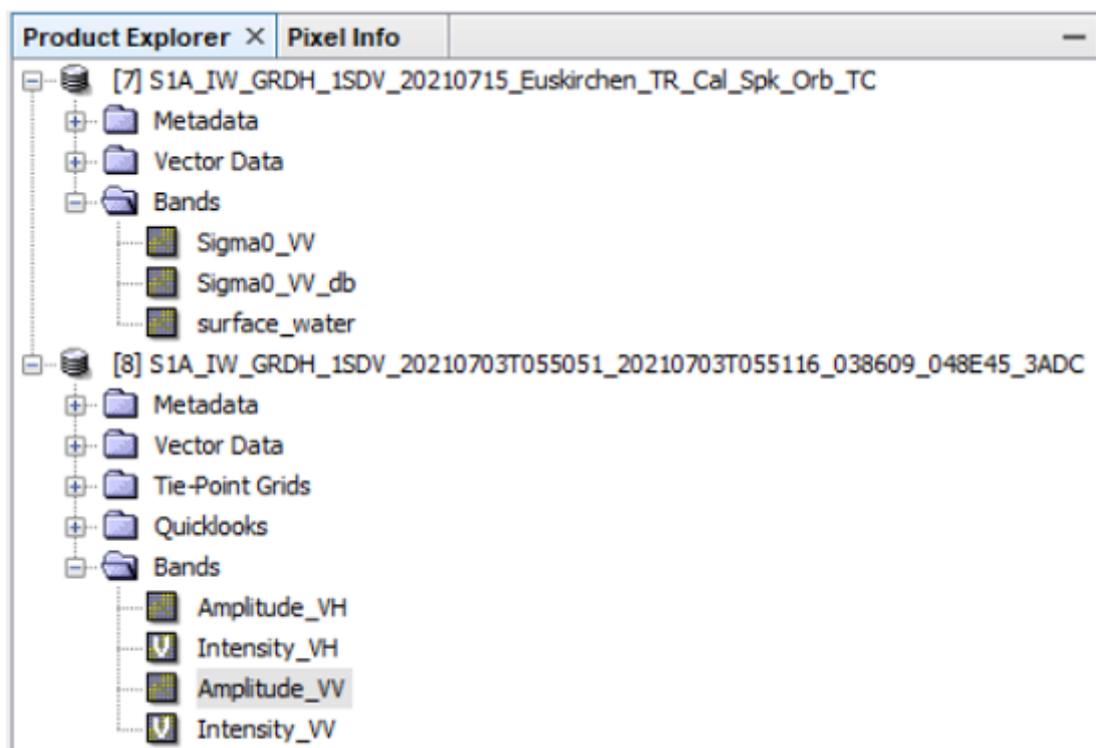
View -> Tool Windows -> Color Manipulation

We can remove the rectangles used to compute the histogram in the binary image from the 'Layer Manager', by deselecting the region_of_interest in the 'Vector Data' section. We could already use this image and overlay it on a map, but our goal is to first remove the water bodies to avoid confusions between the two.

Batch Processing

In our flood mapping task we have processed only one Sentinel-1 product so far: the crisis image acquired on July 15th 2021. The manual work is not only tedious, it also requires some space on the hard disk since at each step a new product is created and saved on disk. The SNAP Graph Builder allows us to design the chain of operators to be applied to the archive product, acquired on July 3rd 2021, without the need to apply each of them manually again.

We can close all the products that we have created so far but the terrain corrected one, and then we import the archive product by dragging the file into the SNAP's Product Explorer window

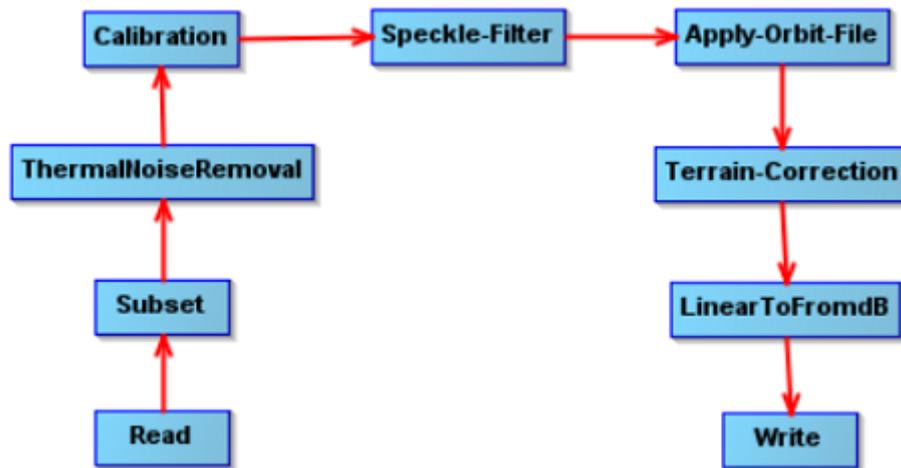


Product Explorer with the archive product imported

To open the Graph Builder we click on

Tools -> GraphBuilder

In the popup window we can add all the operators that we have used to process the crisis product and connect them in a chain like the one shown below



Processing chain

The operators can be added to the graph by right-clicking on the popup window. The GraphBuilder window contains a tab for each operator in the graph to set the parameters as we have done when we were creating the products one by one. Now we set all the operators' parameters in the chain similarly to what we have done for the crisis product.

Read tab

We select the archive product as the source product:

S1A_IW_GRDH_1SDV_20210703T055051_20210703T055116_038609_048E45_3ADC

Subset tab

We select only Amplitude_VV as the source band and we set the following values for the bounding box:

X: 1800 Y: 0

Width: 3000 height: 1500

Thermal Noise Removal tab

We leave all the parameters as they are by default, with the 'Remove Thermal Noise' checkbox selected.

Calibration tab

We leave all the parameters as they are by default, with the 'Output sigma0 band' checkbox selected.

Speckle Filter tab

We set the parameters as we have done for the crisis image:

Filter: Lee

Filter Size X: 5

Filter Size Y: 5

Apply Orbit File tab

We leave all the parameters as they are by default.

Terrain Correction tab

We leave all the parameters as they are by default, with Sigma0_VV as source band.

LinearToFromdB tab

We leave all the parameters as they are by default.

Write tab

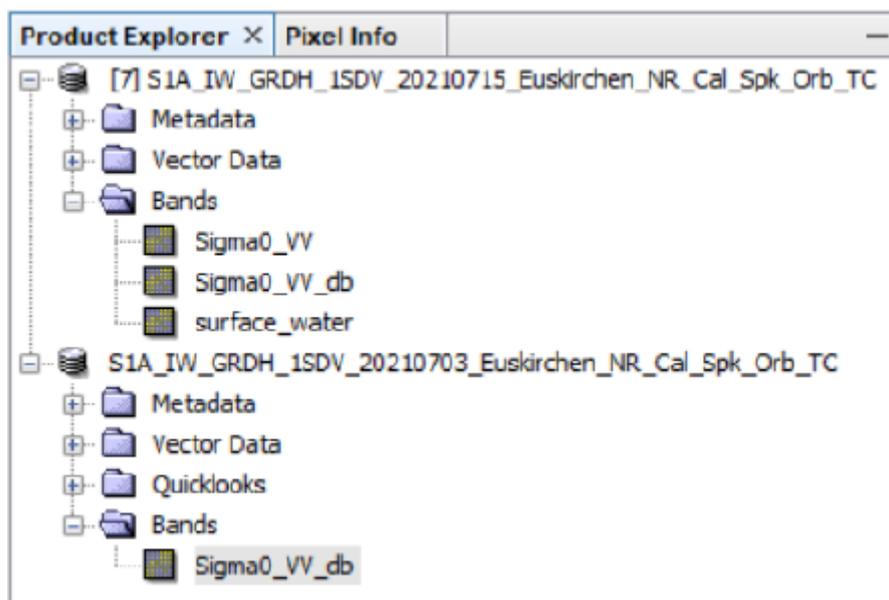
This write operator is directly connected to the LinearToFromdB operator, to not be confused with a second write operator. It will create a new product with the terrain corrected archive image with the pixel values in decibel. We select the output directory, e.g.

C:\Users\\flood_mapping

We can shorten the target product name a little bit, e.g. like this

S1A_IW_GRDH_1SDV_20210703_Euskirchen_NR_Cal_Spk_Orb_TC

Now we are ready to execute our processing chain. We click on 'Run' and a new product is created and shown on the SNAP's 'Product Explorer' window. No intermediate products are saved on disk, only the product that is created by the Write operators. We can close the archive product so we have only the subset products for the crisis and archive images

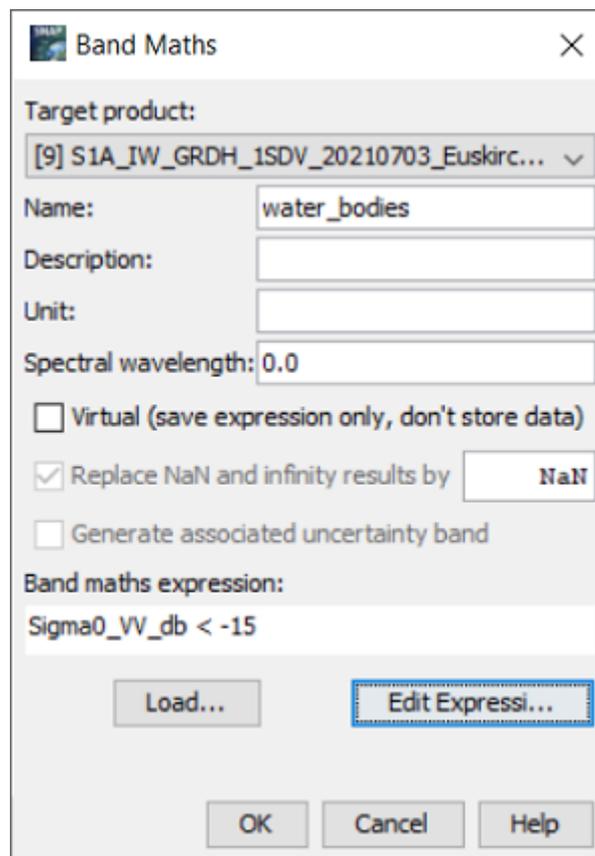


Now we can create the binary image of the water bodies. We select the Sigma0_VV_dB in the subset archive product that we have created using our processing chain, then we click on

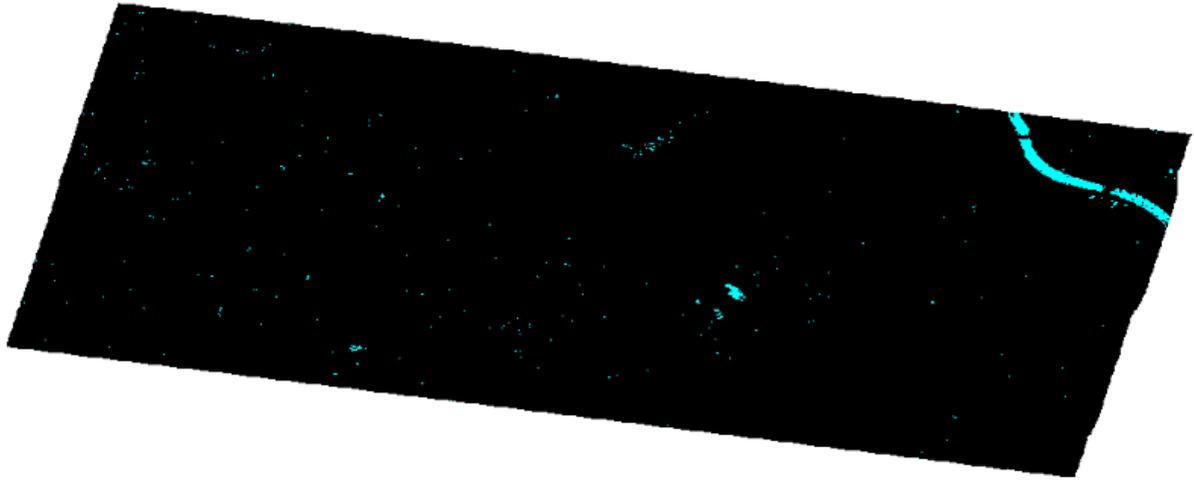
Raster -> Band Maths

In the popup window we check that the subset archive product is selected, we write the name of the new band, e.g. water_bodies, we deselect the checkbox 'Virtual' so that the binary image will be saved on the file system, and finally we create the expression to split the pixels using the threshold the same threshold that we used for the surface water

$\text{Sigma0_VV_db} < -15$



We click on 'Ok' to close the 'Expression Editor' and the Band Maths window. A new band 'water_bodies' with the binary image is created in the subset archive product. We will use this binary image to separate the water bodies (rivers, lakes) from the flooded areas. The pixels color have been changed manually, for presentation.



The binary image with the water bodies

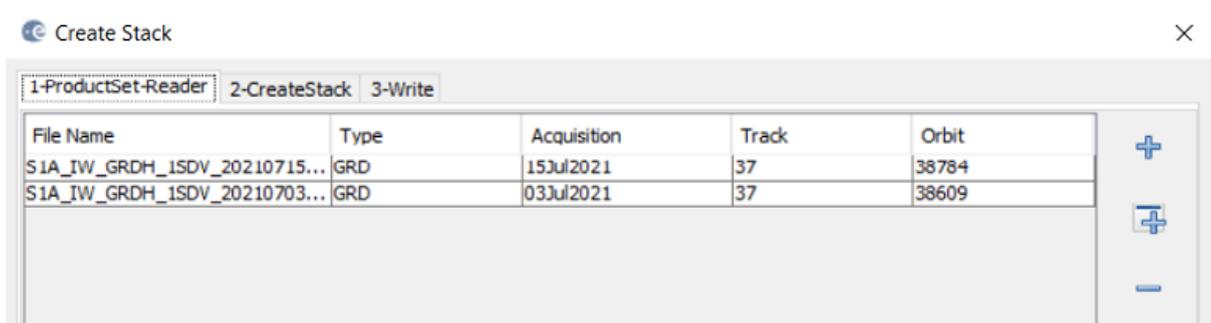
Before moving to the next step we save the subset archive product with the new binary image. We select the archive product and right-click and 'Save Product'

Collocation

At this point we should have two binary images: `surface_water` and `water_bodies`. The first image contains the flooded areas and the water bodies together. In order to separate the flooded areas we have to subtract the `water_bodies` binary image from the `surface_water`. In order to perform such arithmetic operation with two binary images that are in two different products we have to put them together in the same product. In this step we are going to create such a product. The two binary images have been created from the radar images with the satellite in the same geometrical configuration, on the same orbit and observing the surface under the same incidence angle, so we are confident that a pixel in the `surface_water` binary image and a pixel in the `water_bodies` with the same geographical coordinates represent the same area on the Earth surface and so the binary images can be overlaid one on top of the other. To overlay one image over the other we have to create a stack, a single product that will contain both the `surface_water` binary image and the `water_bodies` binary image. We click on

Radar -> Coregistration -> Stack Tools -> Create Stack

In the popup panel we add the two products by clicking on the button 'Add Opened'



In the 'CreateStack' we select 'Orbit' as 'Initial Offset Method' since we have applied the precise orbits. In the 'Write' tab, we remove the date from the output file name since the stack will contain the images taken into different dates, we name the output product, e.g.

S1A_IW_GRDH_1SDV_Euskirchen_Stack

then we click on 'Run'. A new product is created that contains all the bands of the two products from which it has been created.

Flooded Areas

The last step to accomplish our goal is to separate the water bodies from the flooded areas. This is now a simple subtraction between two binary images. We open again the Band Math window

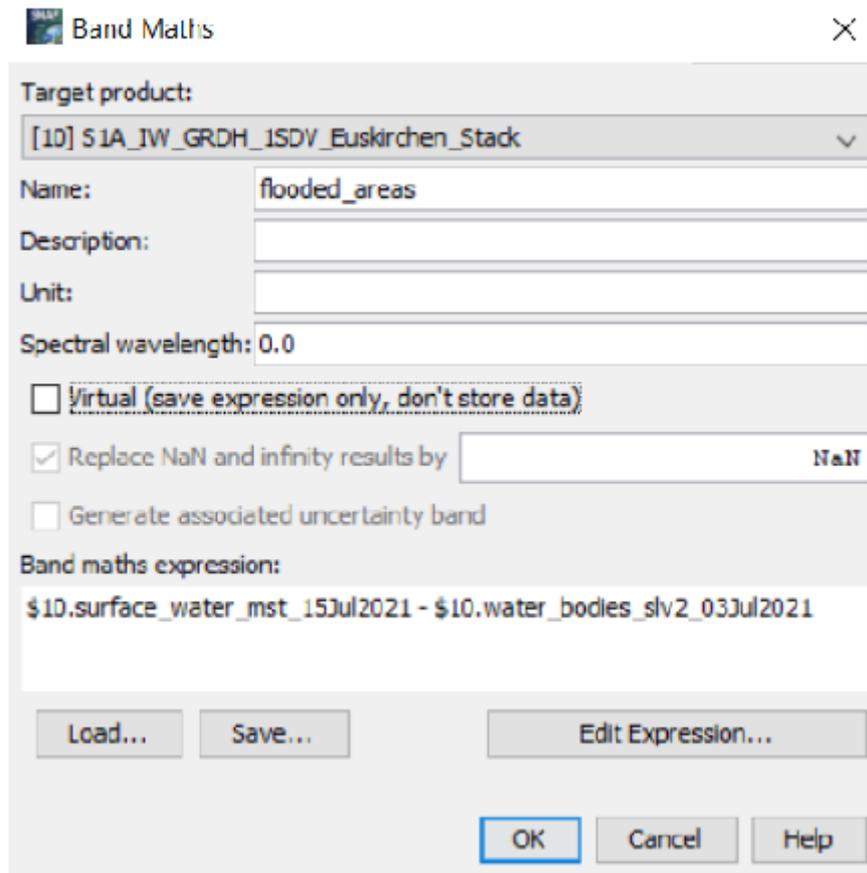
Raster -> Band Maths

In the popup window we select the crisis product, it should be

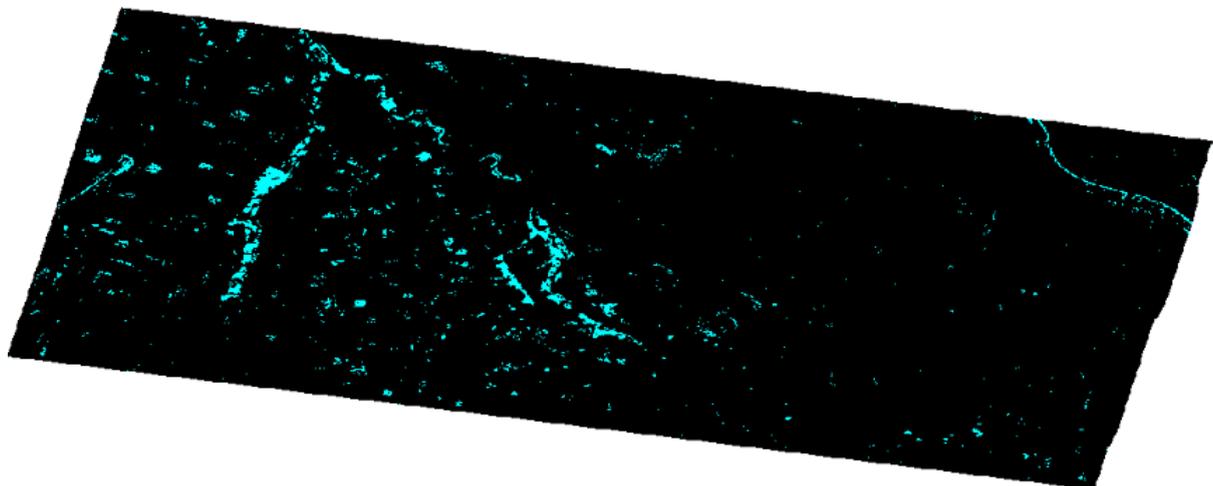
S1A_IW_GRDH_1SDV_Euskirchen_Stack

We write the name of the new band, e.g. flooded_areas, we deselect the checkbox 'Virtual' so that the binary image will be saved on the file system and we click on the 'Edit Expression' button to write the simple expression

surface_water - water_bodies

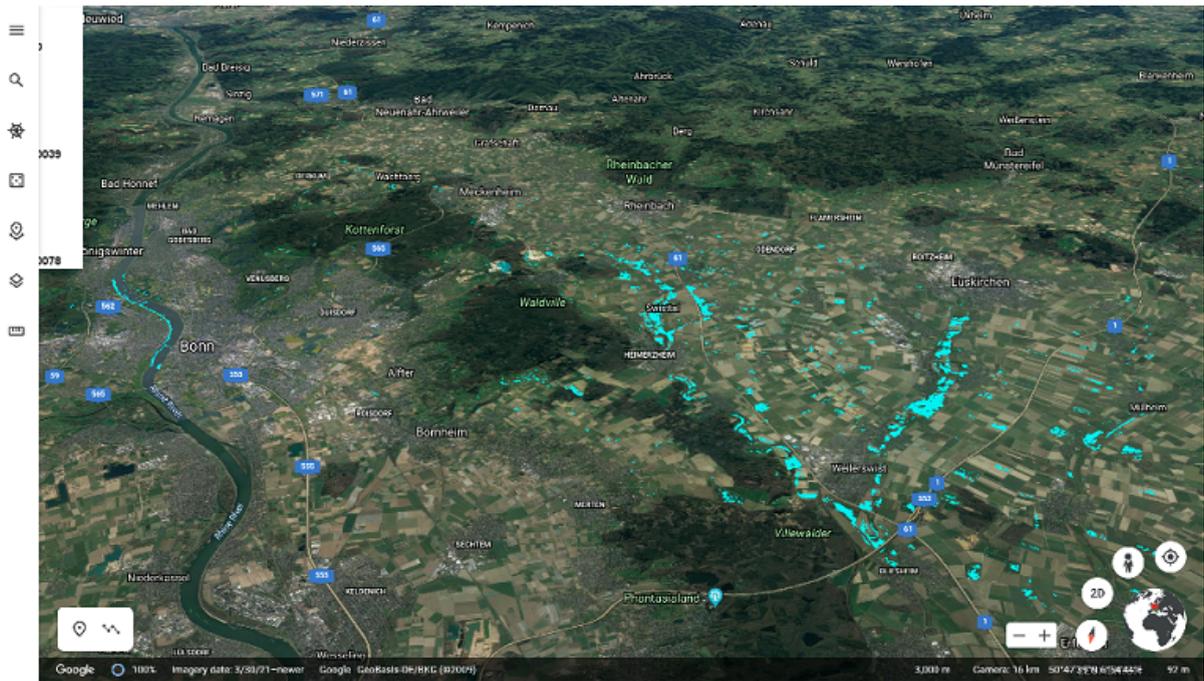


We click on 'Ok' to close the 'Expression Editor' and the Band Maths window. A new band 'flooded_areas' with the binary image is created



Binary image of the flooded areas after the removal of the water bodies

We can project the flooded_areas binary image on Google Earth. We first remove the black background color using the 'Colour Manipulation' tool by selecting 'None' in the Table editor, then we export the binary image by right-clicking on it and selecting 'Export View as Google Earth KMZ'. We save the file that we will upload on Google Earth

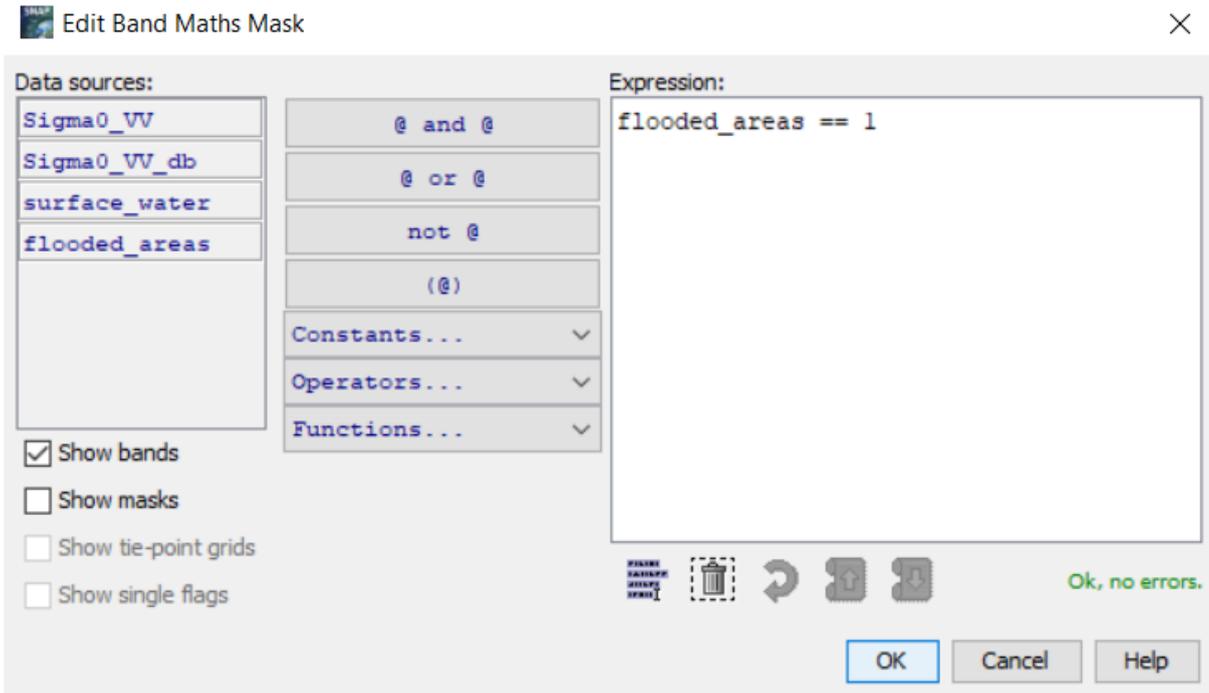


Projection of the flooded areas on Google Earth

Estimation of the flooded area

To compute the area covered by the flood we have to count the number of pixels in the flooded_areas binary image with value 1 and multiply that number for the area covered by a pixel. The area covered by a pixel before the terrain correction is the product of the range spacing by the azimuth spacing, that is 10 m x 10 m, but after the terrain correction the pixel spacing is different because the pixels are geocoded and then [resampled](#). To count the number of pixels of the flooded areas we have to create a mask based on a logical band math expression.

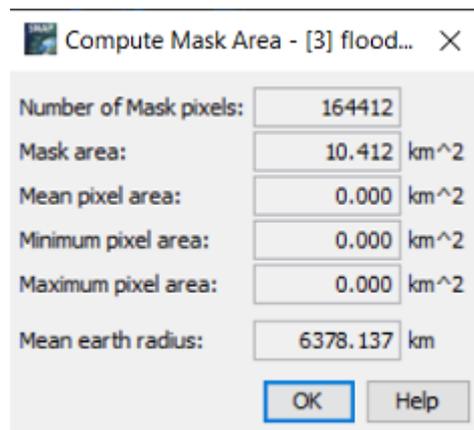
We select the stack product, then from the 'Mask Manager' window we click on the 'f(x)' button. In the 'Logical Band Math Expression' popup window we select the flooded_areas data source and we write the expression as in the picture



Then we click on 'Ok'. We can change the name of the mask, e.g. flood_mask. This mask will select all the pixels whose value is 1 from the flooded_areas binary image. To count the number of pixels in the mask and the area, on the menu bar we click on

Raster -> Masks -> Mask Area

In the 'Compute Mask Area' popup window we select the mask flood_mask



The size of the flooded area is 10.4 Km². The value can be somewhat different whenever we create new binary images.

Conclusion

Synthetic Aperture Radar is a fundamental technology for the monitoring of the Earth surface. Its strengths, compared to the optical sensors, are that it does not depend on the availability of sunlight and can be used even in the presence of clouds. SAR data has been available for some years now, from the US Landsat program and the European Copernicus programme, under an open data licence that allows the commercial reuse of the imagery. Users and companies that could not afford to pay for such imagery are now able to develop products and applications to fully exploit the value of the satellite data.

We have shown how to use the imagery in one important application, flood detection. While the example shown is a simple one it is nonetheless useful and quick to reproduce.

References

1. [Kreienkamp et al. - Rapid attribution of heavy rainfall events leading to the severe flooding in Western Europe during July 2021](#)
2. [Meyer - The SAR Handbook, Chapter 2 - Spaceborne Synthetic Aperture Radar: Principles, Data Access, and Basic Processing Techniques](#)
3. [Moreira et al. - A Tutorial on Synthetic Aperture Radar](#)
4. [Woodhouse - Introduction to Microwave Remote Sensing](#)
5. [ESA - Sentinel 1 - ESA's Radar Observatory Mission for GMES Operational Services, March 2012](#)
6. [Gonzales, Woods - Digital Image Processing, 4th Edition](#)