Report of WALLABY TWG 4

Tobias Westmeier & Russell Jurek
on behalf of the Working Group

9 November 2010

1 Introduction

This report summarises the activities of WALLABY Technical Working Group 4 on "Source finding and cataloguing" during the first year of the design study. Section 2 describes the data cubes used for source finder testing. In Section 3 we outline the two algorithms tested, namely Duchamp and the Γ test. In Section 4 we present the first results of tests with Duchamp. Section 5 summarises the progress made in the development of a custom source finder for WALLABY. Finally, our main conclusions are presented in Section 6.

2 Data cubes for source-finder testing

As a result of the initial delay in the release of the first ASKAP spectral-line simulation and the limited usefulness of that first simulation for the purpose of source-finder testing, the focus of TWG 4 during the first few months of the design study was on obtaining and creating our own H I data cubes. The parameters of the data cubes selected for source finder testing are summarised in Table 1. The following data cubes were acquired or generated:

2.1 HIPASS data cubes

We extracted two large-scale data cubes from the H I Parkes All-Sky Survey (HIPASS) covering the Virgo cluster and the Magellanic Stream, respectively.

The Virgo cluster data set (Fig. 1b) covers an area of about 24° in size around the Virgo cluster, centred on about $12^{\mathrm{h}}30^{\mathrm{m}}$ in right ascension and +6° in declination. The velocity coverage is 165 to 3 500 $\mathrm{km\,s^{-1}}$, i.e. the entire range of Galactic emission was removed to avoid severe artefacts near diffuse, bright emission in HIPASS. The data set contains a total of about 250 galaxies and a high-velocity cloud. Three small regions have been masked throughout the entire cube because of strong residual emission from bright continuum sources in these regions. The final, masked data cube as well as the list of HIPASS galaxies in this part of the sky can be retrieved from the WALLABY wiki.

The Magellanic Stream data set (Fig. 1a) covers an area of about 35° in size around a section of the Magellanic Stream, centred on about $00^{\mathrm{h}}30^{\mathrm{m}}$ in right ascension and −15° in declination. The velocity coverage is −600 to −100 $\mathrm{km\,s^{-1}}$. The data set contains a large

Parameter	HIPASS	Arecibo	WSRT/ASKAP	Unit
Field size	20×23	15×1	1×1	deg
Redshift range	≤ 0.01	≤ 0.04	$0.02 \ldots 0.04$	
Spectral channels	256	2 202	1 464	
Channel width	62.5	26.1	18.3	kHz
Channel width	13.2	5.5	4.0	$\mathrm{km\,s^{-1}}$
HPBW	864	204	30	arcsec
Pixel size	240	60	10	arcsec
rms noise	11.9	1.4	1.6	mJy

Table 1: Parameters of the data sets used for WALLABY source finder testing. Parameters are provided for the HIPASS data cube of the Virgo cluster, the adjusted Arecibo data cube of the Virgo cluster, and the WSRT and ASKAP model data cubes.

section of the Magellanic Stream and a large number of isolated high-velocity clouds of different sizes. The HVC catalogue published by **Putman et al. (2002)** lists 195 objects in the region and velocity range covered by the data cube, mainly different types of HVCs but also a few galaxies. The final data cube can be retrieved from the WALLABY wiki together with a list of high-velocity clouds in this part of the sky.

2.2 Arecibo data cubes

Thanks to the intermediation of Trish Henning (University of New Mexico), the AGES team ("Arecibo Galaxy Environment Survey") kindly provided us with two data cubes of the Virgo cluster prior to publication. Both cubes cover a narrow (1° wide) strip of several degrees in length near the centre of the cluster. We selected one of the data cubes (Fig. 1c) for source finder testing and slightly adjusted the cube size in all three dimensions. The final data cube contains a large number of galaxies across a velocity range of about 12 000 $\mathrm{km\,s^{-1}}$, dominated by Virgo cluster galaxies at velocities of $\lesssim 3\,000$ $\mathrm{km\,s^{-1}}$.

2.3 WSRT model data cube

This data cube (Fig. 1d) was kindly generated by Paolo Serra (ASTRON) in collaboration with Tom Oosterloo and Gyula Józsa (both ASTRON). About 100 galaxies observed with the WSRT as part of the WHISP project ("Westerbork observations of neutral hydrogen in irregular and spiral galaxies") were artificially shifted to different redshifts and injected into a data cube containing real WSRT noise. The cube was then smoothed to an angular resolution of 30″ as anticipated for WALLABY at that time. The angular dimension of the cube is $1° \times 1°$, and the redshift range is $0.02 \leq z \leq 0.04$ with a spectral channel width of 4 $\mathrm{km\,s^{-1}}$.

The galaxies were spread according to a triangular distribution in redshift (peaking at $z = z_{\mathrm{max}}$) and a flat distribution in right ascension and declination. Input galaxies were picked randomly from the WHISP database which contains spirals but also a large number of irregular galaxies. The resulting data cube is clean in the sense that the deconvolved images of the input galaxies have not been convolved with any dirty beam model.

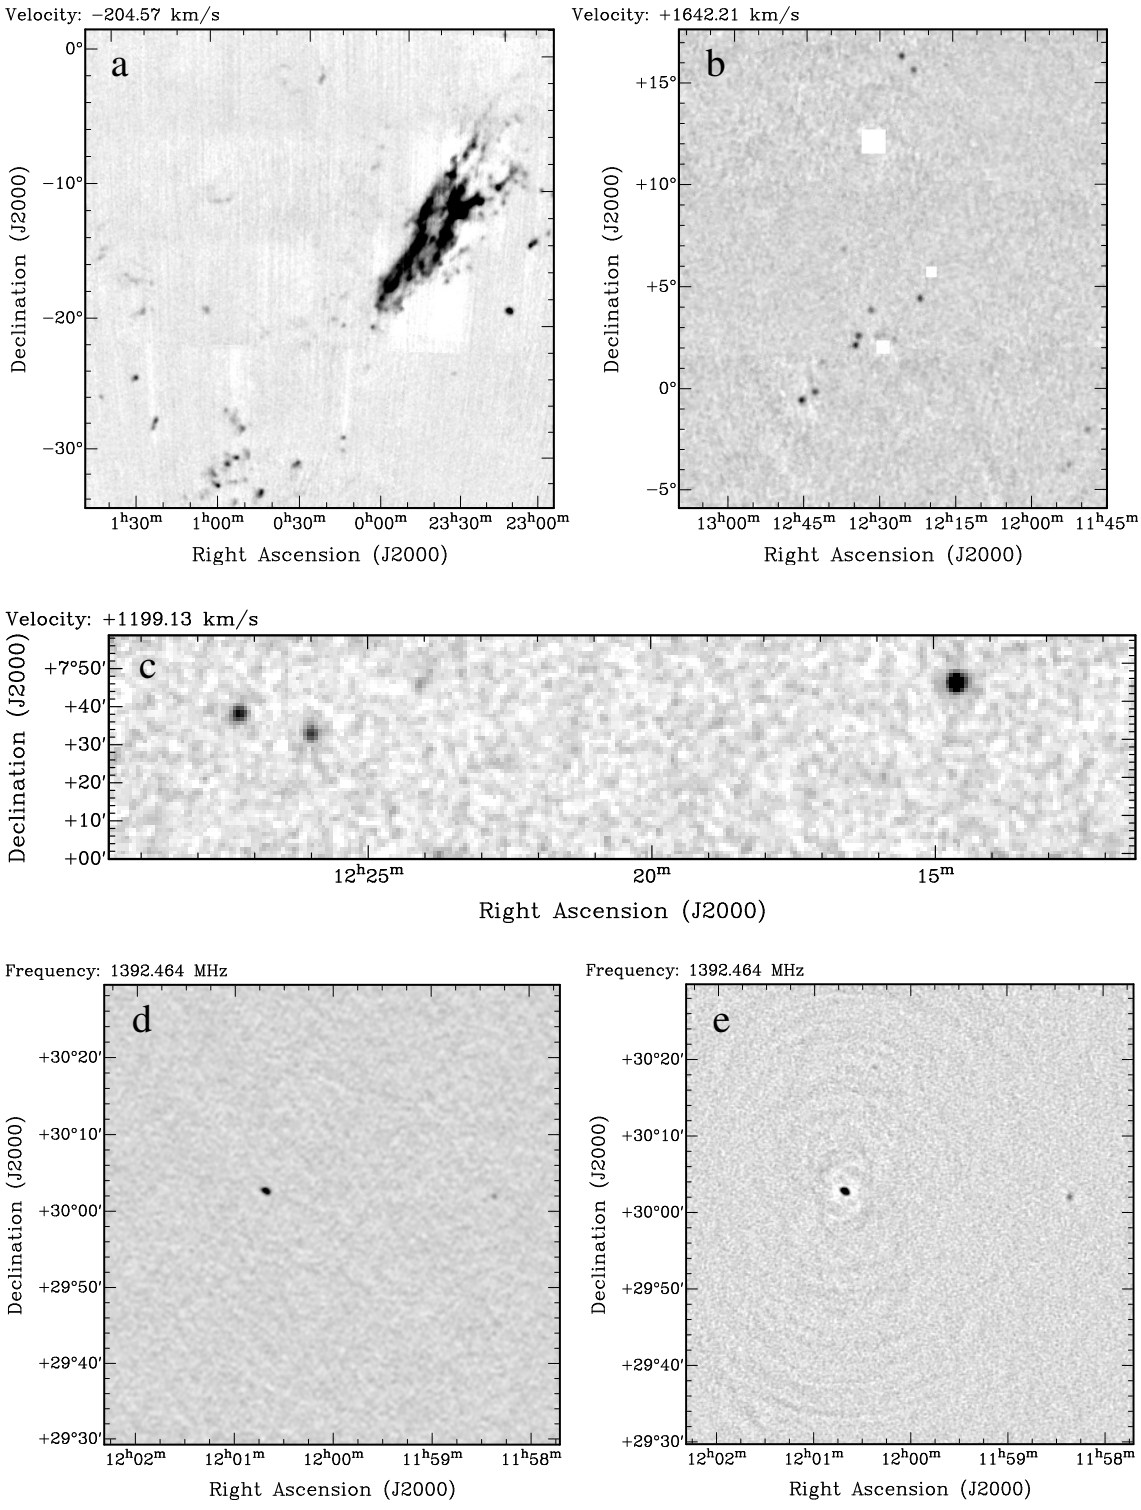

Figure 1: Individual channel maps from our different data cubes for source finder testing. (a) HIPASS data cube of the Magellanic Stream. (b) HIPASS data cube of the Virgo cluster. (c) Arecibo data cube of the Virgo cluster. (d) WSRT model data cube with injected WHISP galaxies. (e) ASKAP model data cube with injected WHISP galaxies.

The density of galaxies in the cube is unrealistically high. This is to ensure that there are enough sources in the cube for a decent statistical analysis of source finder performance. As a result, there may be issues with source confusion in cases where two galaxies partly overlap in phase space.

2.4 ASKAP model data cube

This cube (Fig. 1e) uses the same input galaxy model as the WSRT model data cube provided by Paolo Serra (see Section 2.3 for details). Instead of using WSRT noise, we generated an ASKAP noise cube with the Miriad task UVGEN, Fourier-transformed the visibility data set using INVERT, convolved the WHISP-based source model with the dirty beam provided by INVERT (again assuming a $30''$ beam based on the core array of 30 antennas), and then added the convolved source model to the noise cube produced by INVERT. The resulting data cube has not been deconvolved.

3 Source finding algorithms

3.1 Duchamp

Duchamp is the standard source finder for ASKAP developed my Matthew Whiting (CSIRO). Duchamp makes no assumptions about the morphology of sources. Instead, the software simply scans data cubes for pixels above a specified flux threshold and then merges these detections into objects under various conditions to be specified in the parameter file. Duchamp includes several methods of pre-conditioning and filtering of input data cubes, the most powerful of which is the so-called "à trous" wavelet reconstruction of the cube. By carefully selecting the wavelet scales to be used in the reconstruction one can efficiently filter both high-frequency noise as well as large-scale artefacts. For detailed information about Duchamp please see the **Duchamp User Guide**. The "à trous" wavelet reconstruction of data cubes in Duchamp operates in the following way:

1. The reconstructed image is initialised with 0.
2. The input image is convolved with the wavelet filter function.
3. The difference between input and convolved image is calculated.
4. Everything above a certain threshold is added to the reconstructed image.
5. The filter size is doubled.
6. The procedure is repeated from step 2 until the maximum desired wavelet scale is reached.
7. The final convolved image is added to the reconstructed image.

Duchamp offers wavelet reconstruction in either one, two, or three dimensions, corresponding to reconstruction in the spectral domain, the spatial domain, or the entire three-dimensional phase space, respectively. Different wavelet filter functions can be selected, including a B_3-spline filter, a triangular filter, and a Haar wavelet filter. Another important input parameter is the threshold for wavelet coefficients to be included in the reconstructed image.

3.2 Gamma test

A source finder based on the so-called Γ test has been developed by Benjamin Winkel (University of Bonn) for the Effelsberg Bonn H I Survey (EBHIS). The Γ test is a statistical method for determining the rms noise level, n, of an N-dimensional data set (e.g. a spectrum) with an additional underlying smooth function, $f(x)$:

$$y = f(x) + n. \tag{1}$$

For this purpose, let us define the following two functions:

$$\gamma(q) = \frac{1}{2M} \sum_{i=1}^{M} \left| y_{N(i,q)} - y_i \right|^2, \tag{2}$$

$$\delta(q) = \frac{1}{M} \sum_{i=1}^{M} \left| x_{N(i,q)} - x_i \right|^2. \tag{3}$$

Here, $\gamma(q)$ is essentially the mean difference between the value of a data element, y_i, and that of its q^{th}-nearest neighbour, $y_{N(i,q)}$. Accordingly, $\delta(q)$ defines the mean separation between a data element and its q^{th}-nearest neighbour.

For functions that fulfil the condition defined in Eq. 1 the curve formed by the values of $\gamma(q)$ and $\delta(q)$ for varying q fulfils the linear relation

$$\gamma(q) = A \times \delta(q) + \Gamma \tag{4}$$

where the offset Γ provides an estimate of the square of the rms noise of the data set, $\Gamma = \sigma^2$, and can be determined by a linear fit to the data points. As stated before, the Γ test will provide a good estimate of the rms noise for any underlying function, $f(x)$, as long as that function is smooth. The amplitude of $f(x)$ is not relevant, and a good example for the application of the Γ test is the determination of the noise level in a radio spectrum that is affected by strong baseline ripples. For more information on the Γ test please see the papers by **Evans & Jones 2002** and **Boyce 2003**.

The Γ test can be applied to source finding by assuming that the signal one is looking for does not fulfil the smoothness criterion and therefore leads to a locally increased value of Γ. Tests by **Boyce 2003** have shown that even faint signals extending across not more than a few spectral channels will lead to a significantly increased Γ and a very high signal-to-noise ratio in a spectrum of running-window Γ values over a subset of q values. A simple threshold can then be applied to the Γ spectrum to extract the positions of source candidates. Of course, the Γ test is not restricted to the spectral domain, but can also be applied in the spatial dimension.

4 Results of source finder testing

4.1 Duchamp

We tested Duchamp on several of the data sets described in Section 2. Here, we will summarise the results obtained by applying Duchamp with "à trous" wavelet reconstruction to the WSRT model data cube introduced in Section 2.3.

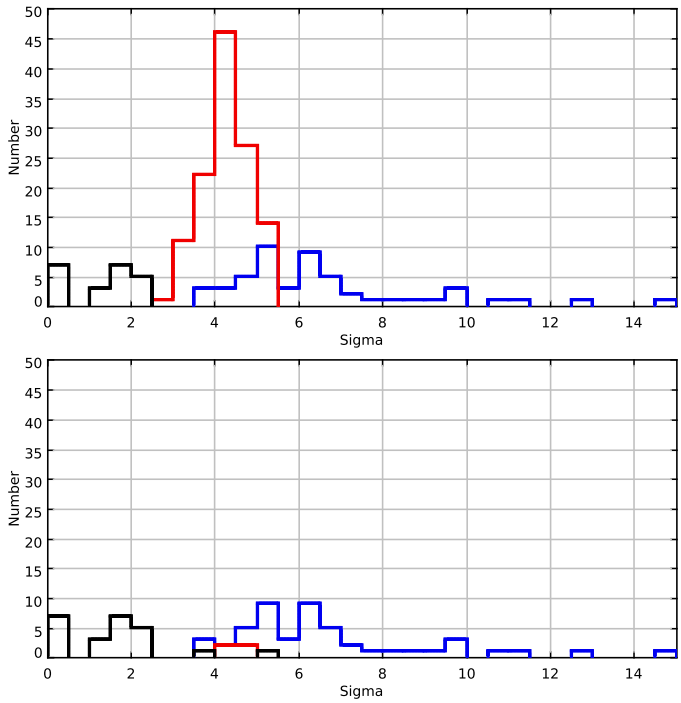

Figure 2: Histogram of sources detected in the WSRT model data cube by Duchamp as a function of peak signal-to-noise ratio, σ, with minimum wavelet scales of 2 (top) and 4 (bottom). Genuine detections are shown in blue, false detections in red, and sources not detected by Duchamp are plotted in black. Exclusion of small wavelet scales is an efficient method of reducing the number of false detections, most of which are due to noise.

4.1.1 Input parameters of Duchamp

First tests on our WSRT model data cube revealed that the most promising way forward is to use one-dimensional wavelet reconstruction in the spectral dimension only. This can be explained by the fact that most galaxies in our test data cube are well-resolved in the spectral domain, but barely resolved in the spatial domain. In addition, we used a B_3-spline filter, a $2\sigma = 3.2$ mJy flux threshold for source finding in the reconstructed data cube, and a 4σ threshold for wavelet coefficients to be included in the reconstructed cube. We ran two different tests with minimum wavelet scales of $s_{\rm min} = 2$ and 4, respectively, for the purpose of filtering out noise peaks. In addition, Duchamp was also run on the original, noise-free input model cube (with very low flux threshold and without wavelet reconstruction) to determine the distribution and parameters of genuine sources in the field for comparison.

4.1.2 Completeness and reliability

Histograms of sources found by Duchamp are shown in Fig. 2 for minimum wavelet scales of 2 and 4. Out of 69 genuine sources in the input model, 22 are not detected by Duchamp (with $s_{\rm min} = 2$; upper panel in Fig. 2). All of these sources have peak flux density levels of $S_{\rm peak} < 2.5\sigma$, implying 100% completeness above 2.5σ. For low peak flux densities,

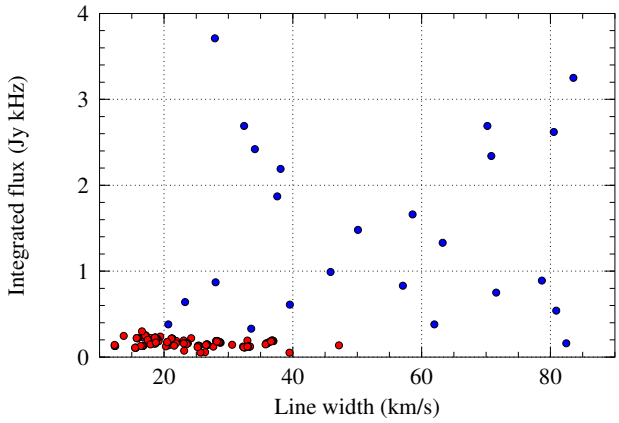

Figure 3: Distribution of genuine (blue) and false (red) detections in integrated flux vs. line width space. Both populations are completely disjunct, allowing a separation of false detections purely based on their parameters.

however, source counts are dominated by false detections with a reliability of only 15% below 5.5σ (but 100% above 5.5σ). The situation is much better for a minimum wavelet scale of $s_{\rm min} = 4$ (lower panel in Fig. 2), in which case the number of false detections is drastically reduced, resulting in a completeness of 96% above 2.5σ and a reliability of 72% below 5σ.

Another method for separating genuine and false detections is through their source parameters. Fig. 3 shows the distribution of false (red) and genuine (blue) detections in integrated flux vs. line width space. Apparently, both populations are completely disjunct, allowing us to discriminate between false and genuine detections based in their distribution in parameter space. A higher-dimensional parameter space might be even more efficient and accurate in separating false detections from genuine ones.

4.1.3 Recovery of position and velocity

Next, we looked into how accurately Duchamp will recover source parameters, such as position, velocity, line width, integrated flux, etc. as compared to the original, noise-free input model. Recovery of source positions is excellent (left-hand panel of Fig. 4) with the exception of cases where Duchamp breaks up a single source into multiple detections. For single detections we determine a standard deviation of $\sigma_{\rm RA} = 3.8''$ and $\sigma_{\rm Dec} = 3.5''$ in right ascension and declination, respectively, which is equal to about 12% of the width of the synthesised beam of $30''$. Recovery of source velocities is very good, too, with errors of no more than about 20 to 30 $\mathrm{km\,s^{-1}}$ even for faint sources (right-hand panel of Fig. 4). For brighter sources of more than about 6σ the standard deviation drops to well below 10 $\mathrm{km\,s^{-1}}$. Again, sources broken up into multiple detections result in significant velocity errors.

4.1.4 Recovery of line width

The left-hand panel of Fig. 5 shows the recovery of line widths (w_{50}) by Duchamp. For bright sources the errors are typically within $\pm 10\%$, but for faint sources below about

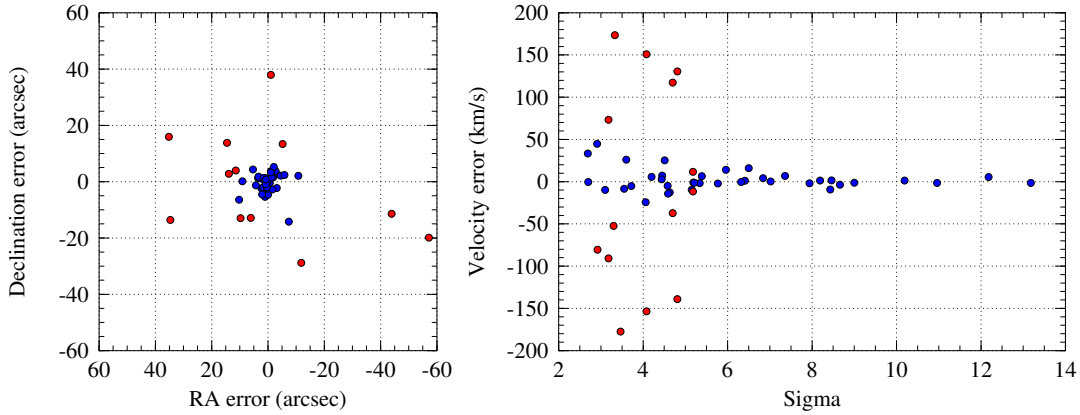

Figure 4: Accuracy of position (left) and velocity (right) of genuine sources as recovered from the WSRT model data cube by Duchamp. Sources broken up into multiple components by Duchamp are plotted in red, single detections in blue.

5σ line widths are systematically underestimated by Duchamp (down to about 60%). Again, the situation is much worse for sources broken up into multiple detections which are drastically underestimated (typically only 10 to 40% of the true line widths).

4.1.5 Recovery of integrated flux

The right-hand panel of Fig. 5 shows the recovered integrated flux of all detected sources as a function of peak signal-to-noise ratio. As evident from the plot, integrated fluxes are systematically underestimated by Duchamp for sources of all fluxes encountered in the WSRT model data cube. For galaxies near the 3σ level only 10 to 20% of the integrated flux is recovered as compared to the original input model, and even for bright sources above 10σ we recover only about 60 to 70% of the integrated flux. This result is alarming for a project like WALLABY that depends on the accurate measurement of fluxes and the H I mass function. The most likely explanation for the discrepancy is that the faint and extended outer H I discs of most galaxies are not picked up by Duchamp and therefore not accounted for in the calculation of integrated fluxes.

4.2 Arecibo Virgo cluster data cube

We also tested Duchamp on one of our single-dish data cubes. For this purpose we selected one of the Arecibo Virgo cluster cubes provided by the AGES team (the one with fewer artefacts and more uniform noise) and slightly cropped the cube in all three dimensions. The final cube covers about $10^\circ \times 1^\circ$ on the sky and the barycentric velocity range of +200 to +12000 $\mathrm{km\,s^{-1}}$.

We then ran Duchamp with "à trous" wavelet reconstruction in the spectral dimension, using the B_3-spline filter, a minimum wavelet scale of 1, and a 4σ threshold for wavelet components to be included in the reconstruction. We ran two tests with different overall flux thresholds of 0.5σ and 1σ, respectively. The detected sources were then inspected by eye to identify whether they could be artefacts (such as noise peaks or RFI) or genuine

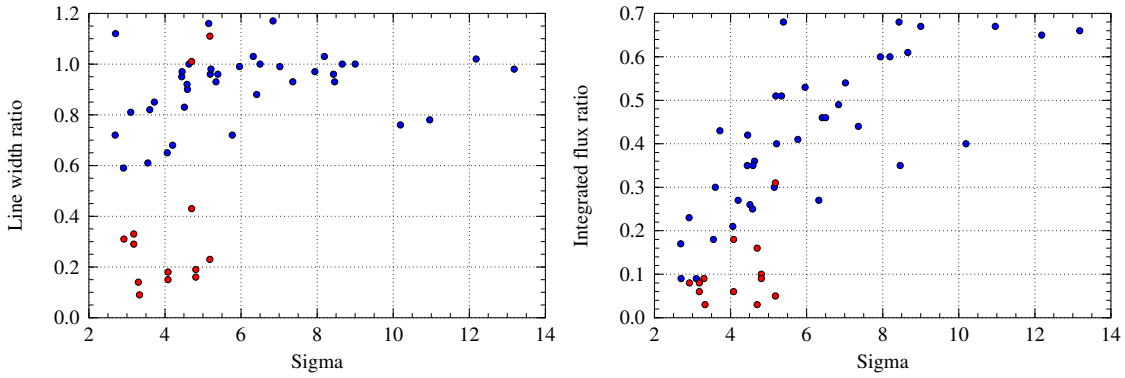

Figure 5: Accuracy of w_{50} line width (left) and integrated flux (right) of genuine sources as recovered from the WSRT model data cube by Duchamp. Sources broken up into multiple components by Duchamp are plotted in red, single detections in blue.

galaxies characterised by a compact morphology and either a Gaussian or a double-horn spectrum.

In the case of a 1σ threshold we detected a total of 107 sources, of which 70 were found to be genuine. This corresponds to an overall reliability of 65.4%. A further 5 detections were presumably due to high-velocity gas around the Milky Way, and 32 detections were found to be artefacts. Most of the artefacts are presumably due to RFI which created a specific "cat paw" pattern on the sky reflecting the layout of the multi-beam receiver. A few more artefacts appear to have been caused by baseline variations at the positions of bright continuum sources.

With a 0.5σ threshold we detected a total of 179 sources, of which 83 were found to be genuine. This results in an overall reliability of only 46.4%. Furthermore, 85 detections were found to be artefacts, and another 6 detections were considered doubtful. A further 5 detections were again caused by foreground high-velocity gas of the Milky Way.

Fig. 6 shows the recovered line width (w_{50}) of all detections plotted versus the barycentric radial velocity. The left-hand and right-hand panels show the situation for a 1σ and 0.5σ peak flux detection threshold, respectively. Most artefacts (red) can be easily identified through their large line widths and clustering around a particular frequency, thus permitting filtering of artefacts to a certain degree on the basis of their characteristics.

In comparison with the tests of Duchamp on interferometer data there are a few significant differences when dealing with single-dish data sets:

- Due to the relatively low angular resolution of the single-dish data most galaxies are spatially unresolved.

- Most of the artefacts in the single-dish data cube are not due to noise, but the result of bandpass instability and RFI.

- The effect of RFI is particularly prominent in the single-dish case because any RFI will directly show up in the imaging data.

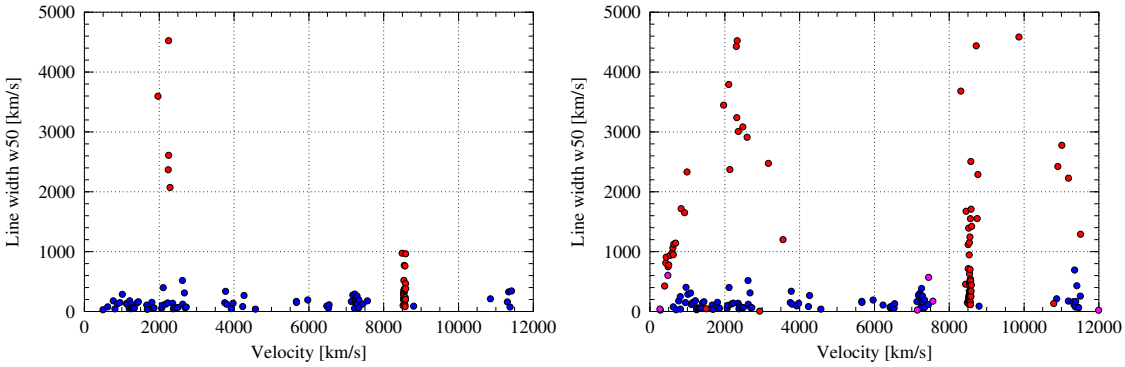

Figure 6: Line width (w_{50}) vs. barycentric radial velocity of the sources detected by running Duchamp with wavelet reconstruction on the Arecibo data cube of the Virgo cluster. The left-hand panel shows the results for a 1σ flux detection threshold and the right-hand panel for 0.5σ. Genuine detections are shown in blue, artefacts in red, and doubtful detections in purple. In the left-hand figure, the artefacts near $v = 2000$ km s^{-1} are due to baseline variations near bright continuum sources, and those near $v = 8500$ km s^{-1} were caused by RFI.

4.3 Gamma test algorithm

Source finder tests with the Γ test algorithm are currently being carried out by Benjamin Winkel (University of Bonn) on some of the data cubes identified in Section 2. These tests are still ongoing, and results as well as a comparison with the performance of Duchamp will be presented later.

5 Development of custom source finder

In addition to testing of existing source finders, we have made significant progress in developing a custom H I source finder for WALLABY in three areas:

1. The development and testing of pre-processing the H I data cubes using three-dimensional iterative median smoothing has been completed.

2. A simulation has been developed to test the optimal way to combine multiple lines of sight to maximise the chance of detecting an H I source.

3. Finally, we have started accumulating a library of wavelets to use for source finding, and are currently developing a wavelet-based source finding tool.

Our progress in each of these areas is described in more detail below.

5.1 Pre-processing of data cubes

First, we present the results of our H I data cube pre-processing development and testing. Pre-processing of H I data cubes improves the performance of the majority of source finders, in particular those that utilise a detection threshold (e.g., Duchamp uses an intensity detection threshold). By improving the signal-to-noise ratio, a given detection threshold

produces more reliable detections. Conversely, a lower detection threshold can be used, which results in more sources being found to the same confidence level. We have developed and tested pre-processing of H I data cubes using three-dimensional iterative median smoothing.

Iterative median smoothing is the process of smoothing data over continually larger scales by taking the median. We are taking advantage of the three-dimensional nature of H I data cubes by smoothing with a three-dimensional kernel. This allows us to smooth on small scales while using many elements to obtain improved statistical power. Using the WSRT model data cube (Section 2.3) we found that two iterations are sufficient. On the first pass we smooth using just the nearest neighbours, then we smooth the resultant data cube using a "sausage" that is five channels long. Using Paolo Serra's WSRT model data cube we find that this implementation of the three-dimensional iterative median smoothing increases the number of sources found by 50% when using Duchamp. The three-dimensional iterative median smoothing also reduces the amount of "fragmentation" of sources into multiple detections and increases the ability of Duchamp to identify the full spatial extent of sources. The statistical output of Duchamp shows that this is the result of the three-dimensional iterative median smoothing improving the signal-to-noise ratio of the WSRT model data cube by decreasing the noise.

Iterative median smoothing is able to improve the signal-to-noise ratio of H I data cubes because it is a non-linear process. Signals in H I data cubes contain edges which non-linear processes preserve. At the simplest level, smoothing while preserving edges results in signal + noise being smoothed within the boundaries of the signal, and pure noise being smoothed outside the regions containing signal. For Gaussian noise (or close to Gaussian noise), taking the median of the noise is virtually equivalent to averaging the noise over these regions, which results in an improved signal-to-noise ratio for the H I data cube.

5.2 Resolution

Next, we discuss our work exploring the effect of the resolution of WALLABY observations on source finding. The high frequency resolution of WALLABY observations improves the science that can be carried out using these observations. The downside to high resolution is that it becomes harder to detect sources in the original data. The higher the resolution of the cubes, the more pixels a signal is dispersed over, which corresponds to a lower signal-to-noise ratio in an individual pixel. This is solved by re-gridding to a coarser resolution. As we are treating H I data cubes as a set of spectra for the purposes of source finding, we are combining multiple lines of sight through a data cube analogous to re-gridding to a coarser resolution. This will improve the signal-to-noise ratio of the spectrum along the principal line of sight, while preserving the full frequency resolution. We have developed a simple simulation to explore the optimal way to combine multiple lines of sight for improving source detection. As expected, the method is most effective for spatially extended sources. The ideal weighting scheme for sources without prior knowledge of their distribution on the sky is a Gaussian weighted profile, as expected. We are currently improving the simulation to be more realistic, explore more weighting schemes, and accommodate sources off-centre from the principal line of sight.

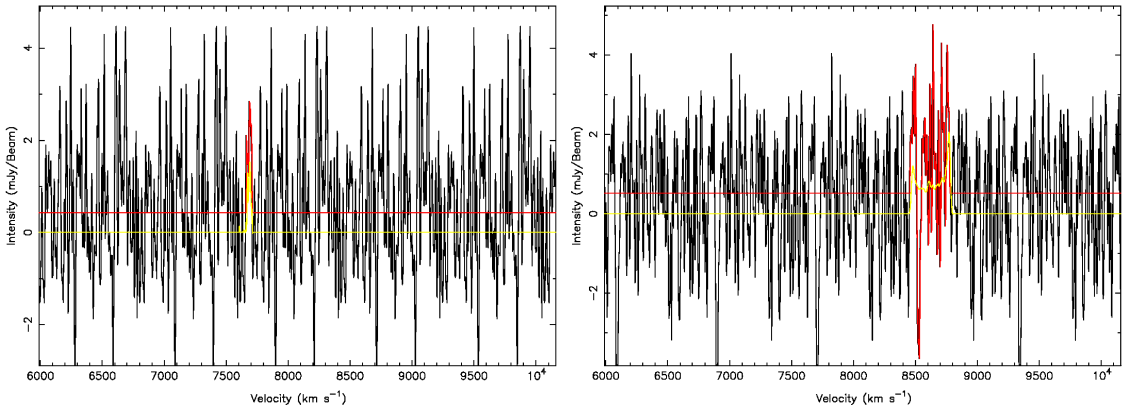

Figure 7: Demonstration of the characterised noise source finder's ability to recover known sources from Paolo Serra's WSRT test cube.

5.3 Wavelet-based source finder

We have begun developing a wavelet-based source finding tool. As a first step, we have started accumulating a library of wavelets that describe the shapes we will search for. These shapes fall into two categories, one describing noise and another one corresponding to galaxy shapes. We have only just begun developing this wavelet-based source finding tool and do not yet have any definitive results to show at this stage. The concept of this source finder is to use a wavelet analysis to decompose a spectrum into shapes that can be separated into noise and sources. The prototype source finder has been successfully applied to the WSRT test cube to demonstrate that with a sufficient library of shapes, wavelet analysis can identify sources. One of the key strengths of a wavelet-based source finder is that it provides an estimate of the shape of the source in the frequency domain. This information is likely to be very beneficial in characterising detected sources, e.g. improving flux estimates.

5.4 Characterised noise source finder

Recently we have explored the idea of characterising the noise in a spectrum and then looking for regions that are different at a statistically significant level. Essentially, we have turned the source finding problem in on itself. A typical source finder tries to find sources based on assumptions and pre-conceptions of what properties they will have that can be selected for. Instead, we are relying on the relatively uniform noise properties and sparseness of ASKAP H I data cubes. This means that for a given spectrum we can use a cumulative frequency distribution to characterise the noise along this line of sight in a statistically meaningful way, with minimal contamination by sources. Using either the Kolmogorov-Smirnov test or Kuiper's test,[1] we can test arbitrary regions of the spectrum on arbitrary scales to see if this region is statistically different from pure noise, indicating the presence of a signal. This source finder is only at a prototype stage, because we have

[1]Kuiper's test is a modified version of the K-S test that, unlike the K-S test, is equally sensitive at the tails of the cumulative frequency distributions as at the centre.

only recently started working on it. The preliminary results however are very promising. We include two plots in Fig. 7 demonstrating the ability of the prototype to recover two sources from the WSRT cube. The black lines are the signal + noise, the yellow line is the signal, and the red line is the reconstructed "signal" recovered by the source finder. The difference in the DC offsets between the signal and reconstruction is an artefact caused by repeating noise in the WSRT test cubes. Note the ability of this prototype to recover weak signals hidden in the noise.

6 Conclusions

- Particularly in the case of faint, edge-on galaxies just above the detection threshold Duchamp tends to break sources up into double (or multiple) components due to their distinct double-horn profile. This problem only occurs for sources with peak flux levels of $S_{\rm peak} \lesssim 5\sigma$ and results in significant errors in both the number of detections as well as the source parameters determined by Duchamp.

  A potential solution to this problem could be to either flag or merge multiple detections in an automated approach based on their parameters and similar positions in phase space. Another option would be to assign to each source a probability for being part of a more extended source, but otherwise leave the source catalogue unchanged. Investigating these options will form an important part of the work in TWG 4 in the second year of the design study.

- Another severe problem encountered during the source finding tests is the systematic underestimation of integrated fluxes by Duchamp. Many spiral galaxies have extended, faint outer gas discs which are often not detected by Duchamp because they fall below the detection threshold. Consequently, up to 90% of the total flux is missing in the integrated flux measurement made by Duchamp for sources near the detection threshold, and even for brighter sources the discrepancy exceeds 30%. This is a particularly severe issue for WALLABY, because the success of the project depends on the accurate measurement of fluxes and H I masses.

  The problem can presumably be solved by changing the way integrated fluxes are being measured. Instead of simply integrating over all pixels above the detection threshold, one could integrate over concentric rings centred on the source position, possibly on a two-dimensional moment map of the source instead of the three-dimensional cube. These options will need to be studied in detail throughout the second year of the design study to develop a satisfactory solution to the missing flux problem.

- Our results have shown that while Duchamp with "à trous" wavelet reconstruction can be successfully employed to detect sources down to below the 3σ level, Duchamp is not particularly well suited to accurately measure source parameters. This outcome suggests that the tasks of source detection and source parametrisation should be handled separately, with a dedicated source parametrisation algorithm running on the final source list produced by the source finder. This strategy would ensure that a more thorough source characterisation algorithm could be employed at selected positions known to contain sources.

In summary, we suggest that the algorithms for source finding and source parametrisation should be developed and optimised separately and in parallel. A third important area that needs to be addressed separately is the preparation and presentation of data products and catalogues.