Appendix 1: Tiling the Sky with ALFALFA

Survey Tiles. In order to effectively manage a full (Arecibo) sky survey, that will cover $\sim 12,000^\circ$, it will be necessary to subdivide the sky in sectors, for each of which continuum maps and 3–d spectral data cubes can be processed and archived coherently. We shall refer to each of those sectors as a "tile".

Observations are most likely to be carried out in drift mode, so that gain, beam pattern and ALFA rotation angle remain constant through a given drift and, possibly for a whole declination strip. Calibration and data processing will be greatly simplified. In drift mode, ALFA would be rotated at an angle such that each beam sweeps a track separated from that of the next beam by 108", approximately half a beam width. In order to maintain beam patterns and gain constant for a given declination, it may be advisable to observe at fixed azimuth, e.g. sampling along the local meridian.

A single ALFA strip will cover 7 × 108" = $12'.6$. Nineteen such strips will cover $12'.6 \times 19 = 3^\circ.99 \simeq 4^\circ$ in declination. This is a natural size for a tile, given its commensurability with standard coordinate units. In order for tiles to have a sensible aspect ratio, an RA extent of 20 minutes (tilel size: $\sim 4^\circ$ to 5°) is right. For a declination coverage of 36°, the survey will consist of 648 tiles, each of approximately $20 \cos\delta$ square degrees, where δ is the declination of the tile. Gridding the map to half–beam grid point separation, each tile would consist of 200×133 grid points, albeit keeping data to a tighter RA grid separation may be desirable.

The spectral values will be written in 4 byte real format, so a single, one–polarization N–channel spectrum will be 4N bytes long. Assuming the spectral processor dump rate to be 1 sec (the beam will be oversampled in order to allow better rfi–excision capability), a single ALFA drift strip along the width of a tile will be $7 \times 2 \times 1200 \times 4N$ plus the space allocated to headers. The largest possible number of spectral channels envisaged for ALFALFA is N=8192, which for a bandwidth of 50 MHz (3 levels) would yield ~ 1.3 km s^{-1} channel separation and may make the data useful for a number of galactic studies, in addition to extragalactic ones. Other ALFALFA options are N=2048, 50 MHz bandwidth (9 levels) and N=4096, 100 MHz (3 levels). With N=8192, a single 7–beam strip across a tile will be 0.55 GBytes, plus headers, for the raw data. Other configurational options will be respectively 2 and 4 times smaller in size. In the calculations that follow, we shall assume the largest possible N=8192, which would enable galactic HI studies with the data.

An important consideration in setting the size of the raw data block units is connected to the computational constraints. The data will be bandpass–corrected one ALFA declination, 7–beam strip at a time. If the length of the strip is restricted to the width of a single tile, and assuming that all seven tracks and both polarizations will be simultaneously processed, it is important for expediency that the raw data be loadable in memory all at once. With currently available, inexpensive workstations with 2 GHz of memory, a 0.6 GByte data set has about the right size for efficient processing.

Bandpass correction and first-pass rfi excision will be applied to fully sampled ALFA strip segments of extent comparable with the tile width. After that, strips may be compressed by a factor of about 3 to 6 in the RA dimension, to approximate one–quarter to one half–beam sampling. Sampling somewhat more generously than the Nyquist rate helps with the quality of the gridding process. Polarizations may be added.

After all the declination strips pertaining to a tile will have been observed, bandpass corrected and rfi–excised to first order, a 3–d data cube can be constructed. Assuming, as mentioned above, that the data will have been compressed by a factor 4 in the RA dimension, each 7–beam strip will amount to about 0.045 GBytes, with full spectral resolution. A 19 strip map will amount to 0.85 GBytes. If it will be found that compressing strips in the RA direction before gridding is undesirable, then the data of a full tile may not fit in a 2 GB memory and may need to be broken into quadrants, and separately regridded.